Group size and gossip strategies: an ABM tool for investigating reputation-based cooperation

Francesca Giardini[1], Mario Paolucci[1], Diana Adamatti[2], Rosaria Conte[1]

[1] ISTC-CNR, Roma, Italy,
{francesca.giardini, mario.paolucci, rosaria.conte}@istc.cnr.it
[2] Universidade Federal do Rio Grande - FURG,
dianaada@gmail.com

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. Using a Public Goods Game, we show that gossip, used for assortment under three different strategies, can be effective in large groups, whereas its efficacy is reduced in small groups, with no main effect of the gossiping strategy. We also test four different combinations of gossip and costly punishment, showing that a combination of punishment and reputation-based partner selection leads to higher cooperation rates.
Keywords: evolution of cooperation; reputation; gossip; punishment.

1 Introduction

When cooperators can be easily exploited by those who reap the benefits of cooperation without paying its costs, free-riders will definitely outcompete altruists [18]. How is it possible to reduce the profitability of free-riding in a social dilemma, like a *public goods* game? Models of large-scale human cooperation show that cooperative behavior can be evolutionarily stable if free riders are punished, thus making defection less profitable by means of decreasing the cheaters' payoffs at a cost for the punisher.

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 [7]. According to "strong reciprocity" theory, punishment is a decentralized, spontaneous and effective solution against cheaters, a solution made possible by the presence of strong reciprocators, i.e., individuals who altruistically reward cooperative acts and punish norm violating behaviors [8]. This view has been recently questioned, and there is mounting evidence that the results on strong reciprocity obtained in the lab can be hardly generalized to what happens in the field [13]. Furthermore, punishment inevitably leads to a second-order collective problem because punishing is costly, therefore refusing to punish non-cooperators maximizes individual welfare. A solution to this has been proposed in [20], where the authors use reputation to link reciprocity and collective action, showing that large-scale cooperation can be stabilized. Using a combination of one-shot Public Goods game and a series of mutual aid games they show that the threat of exclusion from indirect reciprocity can sustain collective action. In this setting, social exclusion reduces free-riding and it has no costs for the punishers (what Panchanathan and Boyd call *shunners*), who can withhold help from free-riders without damaging their reputations.

In models of indirect reciprocity supported by reputation, individuals can base their decision to help others on the basis of observation of their past behaviors. Even without direct experience, cooperation can thrive thanks to the exchange of information based on direct observation of others' interaction [19], or on reported experience [1, 23, 19], even when group size increases. If punishing implies paying a cost in order to make the other pay higher costs for his defection, reputation works because information about agents' past actions becomes known to potential partners, and this allows cooperators to avoid ill-reputed individuals. In Axelrod's words[2]: '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, as demonstrated both in lab experiments [22], and in simulation settings, in which reputation has proven to be a cheap and effective means to avoid cheaters and increase cooperators' payoffs [21, 9, 10].

Moreover, prosocial gossip may effectively bypass the second-order free-rider problem, wherein the costs associated with solving one social dilemma might produce a new one [15, 14]). Reputation-based theories of cooperation [1] consider reputation as a by-product of direct observation, thus equivalent to a label or a score, as the well-known *image score* theory developed in [19]. Agents, in that model, can choose whether to help another individual at a certain cost to himself, or to avoid this cost. This decision is based on agents' image scores which are publicly visible. An agent's image score increases when he donates to another agent and decreases in the opposite case, thus working as a reliable indicator of the agents' past behaviors and cooperative attitudes. In such a setting, cooperation can emerge because free-riders have low image scores, therefore they can be easily avoided by cooperators. However, the result depends on the public availability of accurate information:

> 'cooperation based on indirect reciprocity depends crucially on the ability of a player to estimate the image score of the opponent. In the above model, we assume that the image score of each individual is known to every other member of the population'. ([19] p. 575).

Although effective, image score is completely unrealistic, especially if used to account for the evolution of cooperation in human societies. In large groups of unrelated individuals, direct observation is not possible, and records of an individual's past behaviors are usually not freely and publicly available. What is abundant and costless among humans is gossip, i.e., reported information about others' past actions that can be used to avoid free-riders, either by refusing to interact with them, or joining another group in which free-riders are supposedly absent.

2 Gossip as a more realistic information transmission

Unlike previous works [19, 20], in which information about others is publicly available and is used to discriminate between cooperators and non-cooperators in a indirect reciprocity game [19], we design a model in which information is privately transmitted among gossipers. Also, we account for the fact that information is noisy and reputation is sticky, two features that characterize human societies, in which gossip does not necessarily have a positive effect. A third element of novelty of our work is the use of a multi-players game, such as the *public goods game* (*PGG* from now on), in which interactions happen in groups, with a conflict of interest among agents. In such a setting, our work is aimed at investigating whether different reputation-based strategies may have an effect on cooperation rates in mixed populations, and also to compare performances of gossipers, who exchange information about their peers, and punishers, who pay a cost in order to reduce cheaters' payoffs.

We are interested in contrasting benefits and costs of gossip with those of costly punishment, but we also aim at specifying how action strategies may complement reputation spreading. While punishing a free-rider is an action with immediate consequences on the free-rider's payoff, reputation spreading implies an information transmission, but the way in which this information is used is usually not specified. For a complete definition of the gossip and reputation behaviour, we have to define how agents are going to transform this information into action. Here, we propose three reputation-based strategies: gossipers can refuse to contribute to the group (strategy *refuse*, actively look for a better group (strategy *compare*), or apply a more refined form of partner choice. In this latter case, group formation is delegated to a single agent, randomly selected to act as a leader, and then allowed to choose its group mates (strategy *leader*). If the leader is a gossiper, it can use information received about others in order to select the most cooperative partners and avoid the uncooperative ones.

If we define the *object level* as containing all the actions that influence score directly, the *PGG* constitutes the main interaction at the object level for our model. Punishment also happens at the object level, as it is a response strategy influencing scores directly. Distinct from the object level, the *information level* contains observation results, for example compliance information as employed by punisher agents, but also inter-agent information diffusion by gossip. Note that a gossip strategy by itself is incomplete, because it is only specified at the information level; a complete mechanisms need to influence the object level just as punishments does (see Figure 1). Thus, information is applied to the object level by means of costly punishment; by withholding participation to the next game (gossipers under *refuse* strategy); or by weak (gossipers with *compare*) or strong (*leader*) partner selection. These last three mechanisms constitute the application, or *response* as we will call it in the rest of the paper, at the object level of gossip and reputation information.

Building upon results obtained in previous work [3, 11], our agents are cooperative at the object level (with the exception of the free-riders), but can change their behavior on the basis of their peers' actions. Agents start gossiping and

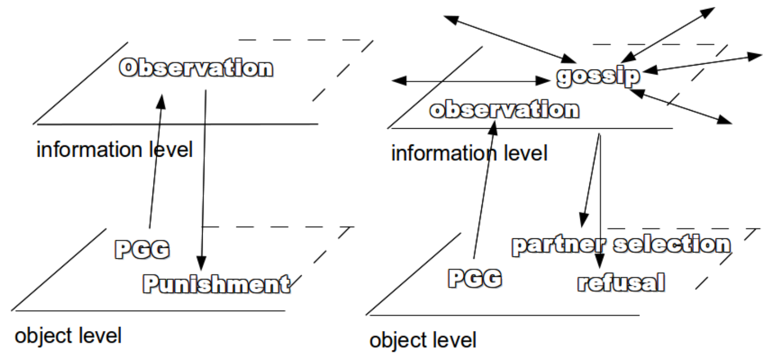

Fig. 1. Actions at different levels. Left: a punishment mechanism is complete as it contains both information specification (observation only) and action specification that effects the object level. Right: gossip and reputation mechanisms dictate how information is diffused and grows, but they don't have an unique response. In this work, we study responses as partner selection and refusal.

punishing, respectively, when the number of free-riders in their group is considered too high, i.e., it exceeds a given threshold (which is set at 0.2 of the group size). The threshold represents an attempt to take into account the fact that free-riders may be difficult to discover and that some missed contributions can go undetected.

It is well known that group size can be a critical factor in models of social interaction. We are interested in understanding the interplay between our strategies and group size, so we report results obtained in small groups (5 agents), medium size group (10 agents), and large groups (25 agents).

To this purpose, we have implemented our model using NetLogo [24]. The implementation is general-purpose and highly customizable, even beyond the purposes of the present paper[3].

The paper is organized as follows: Section 2 explains the simulation model using an ODD protocol, Section 3 presents experiments and results, whereas in Section 4 the results are discussed and some conclusions are sketched.

3 The model

Building up on the simulation framework developed in [11] using NetLogo, we added two new strategies in order to deepen our understanding of the role of reputation mechanisms in supporting cooperation in mixed populations in which different types of agents play a *public goods game* (*PGG*), the classical experimental model used to investigate social dilemmas [17]. In this game, agents in a group decide whether or not to contribute to a public pot at a net personal

[3] The model can be downloaded at http://labss.istc.cnr.it/code/punishment-and-reputation

cost c, in order to create a benefit b (normally with $b > c$). The pot is then divided equally amongst all the participants in the group, without considering their contributions.

In compliance with the ODD protocol [12] for describing ABM models and simulations, the model will be described now in terms of purpose, entities, processes and objectives.

3.1 Purpose

Using a PGG, we investigate how cooperation can be maintained in mixed populations in which there are cooperators, free-riders, punishers and gossipers. While the former two populations play always the same strategy, irrespective of what other agents do, the latter types of agents are reactive. The group's total payoff is maximized when everyone contributes all of their private endowment to the public pool. However, game-theory predicts zero contributions because any rational agent does best contributing zero, regardless of whatever anyone else does.

It is well known that punishment reduces profitability of free-riding and increases cooperation [3, 6], but less is known about the effect of gossiping, which is a kind a informal social control widely used in human societies. This is especially true if we refer to information transmission in a one-to-one way, instead of public information available through direct observation.

In a previous work [11], we have shown the effects of two gossip strategies on the emergence of cooperation: *defect* (a retaliatory strategy that agents played against those with a bad reputation), and *refuse*. Here we consider two alternative strategies, *compare* and *leader*, with the aim of introducing partner choice in the model and testing its performance on cooperation rates. Our hypothesis is that when partner choice is available, a reputation-based strategy for social control should be effective in promoting cooperation and selecting out free-riders. We also predict that there is an interaction effect between group size and gossip strategy, therefore in bigger groups a gossip-based strategy should perform better. For this reason, we test our model in a first experiment for three different group sizes: 5, 10, and 25 agents. To the best of our knowledge, this work is the first attempt to model different gossip strategies and to compare their performance on cooperation rates in a mixed population.

In a second experiment, we also measure different combinations of harsh and mild reactions, with the aim of understanding what happens when punishers and gossipers start defecting in a retaliatory way, and to what extent cooperation is robust to this behavior.

Our simulation allows to:

- explore three different ways of implementing gossip in a PGG, and to test the effectiveness of different gossip-based behaviors on cooperation rates, populations' scores and survival rates;
- identify the most effective combination of costly punishment and gossip in a situation in which Gossipers and Punishers may react to free-riding more or less harshly.

3.2 Entities, state variables and scales

The main entities in our model are agents, either non-reactive or reactive. In the former category we find *cooperators* (C), who always contribute to the common pool, and *free-riders* (FR), who never contribute. In the latter category, *punishers* (P) and *gossipers* (G) start as cooperators, but they change their behavior in response to the percentage of detected free-riders in their group: when the number of known defectors in a group exceeds a threshold set at 20% of group size (following [3]), punishers and gossipers become *active* and apply a counter strategy defined below.

At the onset of the simulation, each agent is endowed with an initial amount of 50 points that can be put in the common pool, or used to punish others; regardless of the strategy, agents are culled from the game when their cumulated payoff goes to zero (*death* of the agent) and they are not replaced. The cost of contributing to the PGG is set to 1, the unit of our utility scale, and the sum of all the contributions is multiplied by a factor set to 3. The public good, i.e., the resulting quantity, is divided evenly among all group members, regardless of individual contributions.

As for counter strategies, punishment works by reducing the payoff of free-riders through the imposition of a cost sustained by the punisher. Punishers pay 1 unit in order to reduce free-riders's payoff by 5. Punishers keeps on punishing until they run out of resources.

Each step, a simple evolutionary algorithm is applied. Agents are ordered by score, and those sitting at the bottom of the ranking are removed and replaced with an identical number of clones generated by a random subset of the surviving ones. The replacement rate is set at 8% of the population, and this kind of algorithm has been already used in the social sciences [3].

For each strategy we calculate the average *score* in time, as the accumulation of points obtained in the PGG, and the population for each strategy, as modified by evolution and death. We also calculate the *cooperation rate* as the ratio of agents who contributed to the last game, a value bounded between 0 and 1. Note that complete cooperation can be reached only if the FR population gets extinct.

Algorithm 1 Description of punishers behaviors

```
While {Number of Timesteps}
    * 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;
```

An important variable is group size, which is known to have an effect on cooperation rates in PGG [3]. We tested our strategies for three different dimensions of groups: small group (5 agents), intermediate group (10 agents) and large group (25 agents).

In the second study, we fine-tuned the reactions, dividing Punishers and Gossipers into two sub-sets [11]:

- Nice Punishers (*Np*) cooperate in the passive state; once active, they punish free-riders at a cost to themselves, but they continue to cooperate in the *PGG*. This behavior will continue until the agent eventually exits the active state when cheaters' rate goes below 0.2 of group size.
- Mean Punishers (*Mp*) contribute in the passive state; once active, they punish free-riders and free ride themselves in the *PGG* until they exit the active state.
- Nice Gossipers (*Ng*) are cooperative in the passive state; once active, they start spreading information about free-riders, and keep on cooperating in the *PGG*.
- Mean Gossipers (*Mg*) contribute in the passive state; once active, they spread information about free-riders, and always defect in the *PGG* until the agent eventually exits the active state.

Algorithm 2 Description of gossipers' behaviors - REFUSE

```
While {Number of Time-steps}
    * Random group formation of the population;
    * Gossipers check others' reputation;
    If number of bad images (known FRs plus anyone who defected and had been marked with a bad image) equals
        or is higher than beta * numagents
          * Reputation diffusion
          * Gossipers refuse the interaction
    * Gossipers take the First stage decision according to their active/passive status;
    * Gather and Distribution of the Public good in each group;
    * Images are updated (bad images added; there is no restoration of bad images if one cooperated)
```

3.3 Process overview and scheduling

The behavior of reactive populations is described in the algorithms 1 to 4. Punishers punish after contributions are made public, whereas gossip reaction works as a proactive strategy that is triggered by agents' reputations.

4 Experiments and Results

We conducted two simulation experiments in which simulations lasted for 200 time steps. In the first set of simulations, we explore a 3x3 set of conditions, with three different group size of 5, 10, and 25 agents. We measured cooperation rates, populations' scores and population size in mixed populations with gossipers playing three different strategies: *compare*, *leader* and *refuse*. Each condition was repeated ten times, for a total of 90 simulation runs. The results are reported in section 4.1.

The second experiment was performed in a 4x2x2 set of conditions, again repeated 10 times each, for a total of 160 simulations. For all plots, points in time are averages over repetitions. Results are reported in 4.2.

4.1 First experiment: testing responses

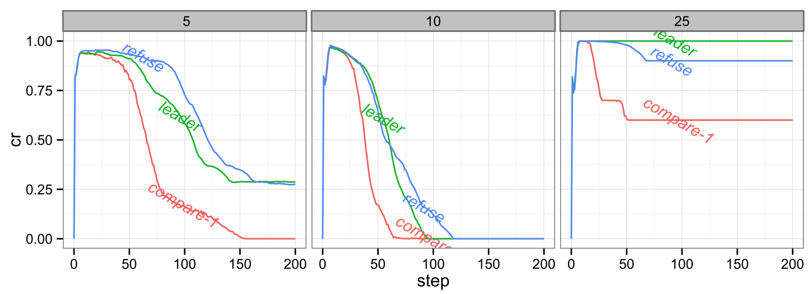

Fig. 2. Cooperation rate (cr) in time by response and group size. In large groups (25 agents), the final cooperation rate is higher, and this is especially true if compared to the intermediary group size (10 agents). Agents in 10-sized groups show the lowest level of cooperation.

First, we wanted to test whether cooperation can evolve and be maintained in a mixed population in which gossip is the only means of social control. Given the strategies employed, cooperation can be generated by any mix of population in a non-active state, except for the *FR*. In Figure 2, the average cooperation rate for all agents is presented. Large groups (25 agents) outperform the others for all responses. In intermediate size groups (10 agents) we observe the lowest cooperation rates. A possible explanation is that with 10 agents in group there are too many agents for direct reciprocity to be effective, but too few for reputation to work. With regard to response strategies, *compare* always leads to lower cooperation levels, while *leader* and *refuse* do not show a consistent

Algorithm 3 Description of gossipers' behaviors - COMPARE

```
While {Number of Time-steps}
    * Random group formation of the population;
    * Gossipers check other agents' reputation;
    If number of bad images (known FRs plus anyone who defected and had been marked with a bad image) equals
          or is higher than beta * numagents
        * Reputation diffusion
        * Gossipers evaluate another group at random, and join it if the percentage of known free-riders there
              is lower than in the current group. Another agent is randomly picked to join the original group of
              the shifting agent.
    * Gossipers take the First stage decision according to their active/passive status;
    * Gather and Distribution of the Public good in each group;
    * First Stage Decisions are made public within the group;
    * Images are updated (bad images added; there is no restoration of bad images if one cooperated).
```

Algorithm 4 Description of gossipers' behaviors - LEADER

While {Number of Time-steps}

- * A given number of agents (depending on the the total population and the group size g) are randomly selected to act as leaders
- **If** Leader is a gossiper
  - * check agents without bad images and tries to build a group with them (if there aren't enough not-bad-image agents he will admits some bad-image to fill up the group)
- **ElseIf** Leader is not a gossiper
  - * gather agents and form a group (G) = (groupsize-1)
- * Gossipers check other agents' reputation;
- **If** number of bad images (known FRs plus anyone who defected and had been marked with a bad image) equals or is higher than beta * numagents
  - * Reputation diffusion (active gossipers inform other gossipers about known cheaters in the group (if gossip=10 they spread info ten times depending on the number of recipients available)
- * Gossipers take the First stage decision according to their active/passive status;
- * Gather and Distribution of the Public good in each group;
- * First Stage Decisions are made public within the group;
- * Images are updated (bad images added; there is no restoration of bad images if one cooperated).

ordering. The *leader* response prevails in larger groups and is the only strategy that supports full cooperation, while *refuse* prevails in smaller groups.

To have a better understanding of the emergence of cooperation, we look into cumulated scores (the results of the *PGG*) for each strategy. These are shown in Figure 3. Free-riders' payoffs are higher in groups of 5 and 10 agents, where, after an initial increase, all other populations reach a maximum score around step 60, followed by a collapse.

To the contrary, in larger (25 agents) groups reputation strategies are effective in keeping free-riders under control. This result supports our view of gossip as a powerful mean to sustain cooperation in large groups. In particular, when gossipers played the *compare* response they achieved the highest payoffs in the population, dominating all other strategies; when they used *leader* or *refuse*, they got the highest payoffs together with cooperators. Punishers, while being an essential ingredient for cheating control, obtain low payoffs in all situations.

Another measure of success is reported in Figure 4 where the final population sizes are shown. For groups of size 5 or 10, the *FR* strategy outperformed other populations, being in several cases the only survivor after 200 steps. Larger groups (25 agents) showed a different pattern: the cooperative strategies, summed together, managed to contain them within the initial simulation steps, and ended out controlling the whole population.

Finally we run a conditional inference tree analysis [16], to isolate, among all the possible variables, the ones that have a major impact on our results. Conditional inference trees are used to estimate a regression between a set of variables, in this case group size and response strategy, ordering them on the basis of the strength of their association to the effect (in this case, cooperation rate).

As shown in Figure 5, in large groups cooperation is higher (leaves 9, 10 and 11). Independent of group size, the group strategy is always divided between *leader* and *refuse* in one branch, and *compare* in the other. *Refuse* is a conservative strategy and *leader* allows for full partner selection, therefore they can promote cooperation and favour gossipers. On the other hand, agents playing the *compare* strategy have only one possibility to change their group, therefore

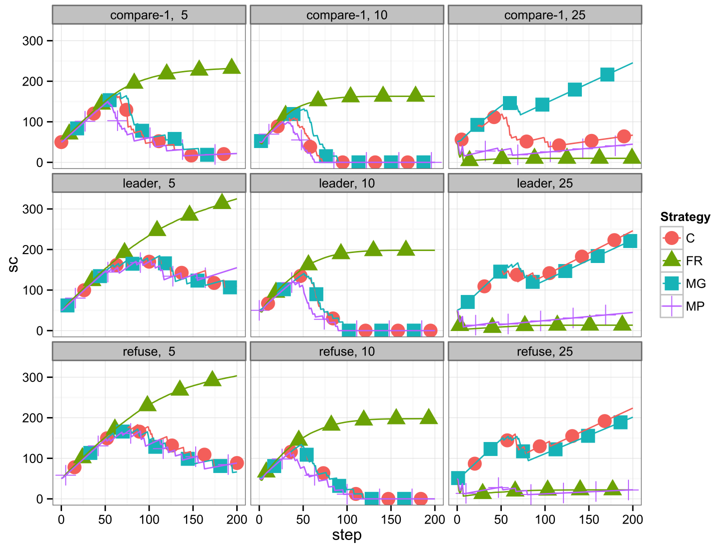

Fig. 3. Average strategy scores in time for each strategy. Rows: gossip response; columns: group size. Free-riders (FR) prevail in group size 5 and 10, while the cooperative strategies fare better in groups of 25 agents. Notice the special case of the *compare-1* gossip strategy separates gossipers (Mg) from punishers (Mp), the latter generally paying the cost of cooperation.

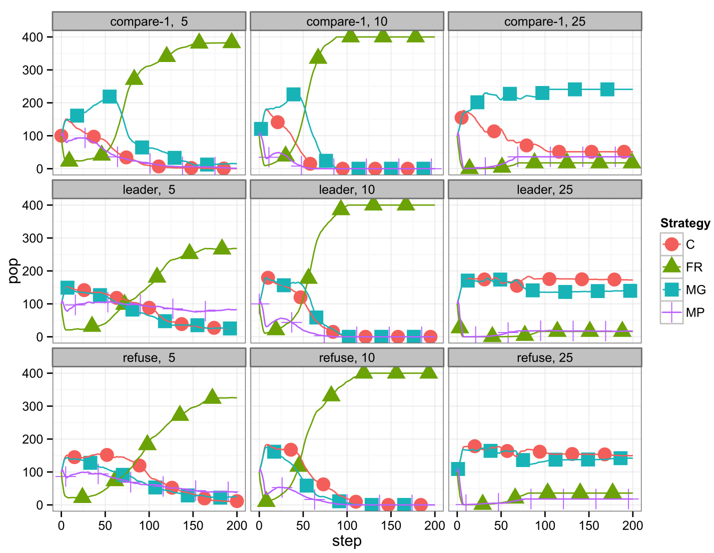

Fig. 4. Number of agents (pop) in time by group size and response. Rows: gossip strategy; columns: group size. The intermediate group level (10 agents per group) shows the higher extinction rate.

they can end up in a group of strangers with a higher number of free-riders than in their original group.

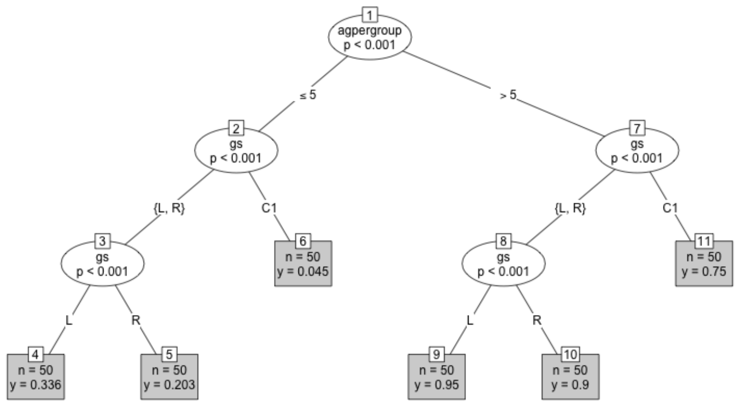

Fig. 5. Conditional inference tree on data for the last 20 steps of the first experiment. The group size (*agpergroup* variable) divides the tree between small groups (≤ 5) and large groups. Subsequently, the gossip response (*gs* variable) divides the tree between *leader* (L) and *refuse* (R) in one branch, and *compare* ($C1$) in the other. The better cooperation rates (Y variables) are in large groups.

4.2 Second experiment: severity of mechanisms

In the second experiment, we wanted to compare two levels of severity for the mechanisms that support cooperation, that is, for punishment and reputation. As with the previous experiment, we record cooperation rates, scores and populations, with the aim to find out the best combination between harsh/mild punishment and harsh/mild gossip.

Results are reported for global cooperation rates in Figure 6. Here, cooperation rates in time for the 4 combinations of Nice and Mean are displayed, in a 3x3 experiments setting with group size of 5, 10 and 25. Gossipers can adopt one of the three responses *compare*, *leader* and *refuse*. Each line shows the global cooperation rate for a single experiment with 100 agents per strategy in combinations of mean and nice variants. We remind that *mean* punishers (Mp) and gossipers (Mg) do not only react to known cheaters in their group by punishing or gossiping, but they also defect, while *nice* ones only apply the relative response.

In this experiment, cooperation was difficult to achieve when gossipers defected and punishers reacted without defecting. On the other hand, we observed very high cooperation rates, up to 1, with the combination of mean punishment

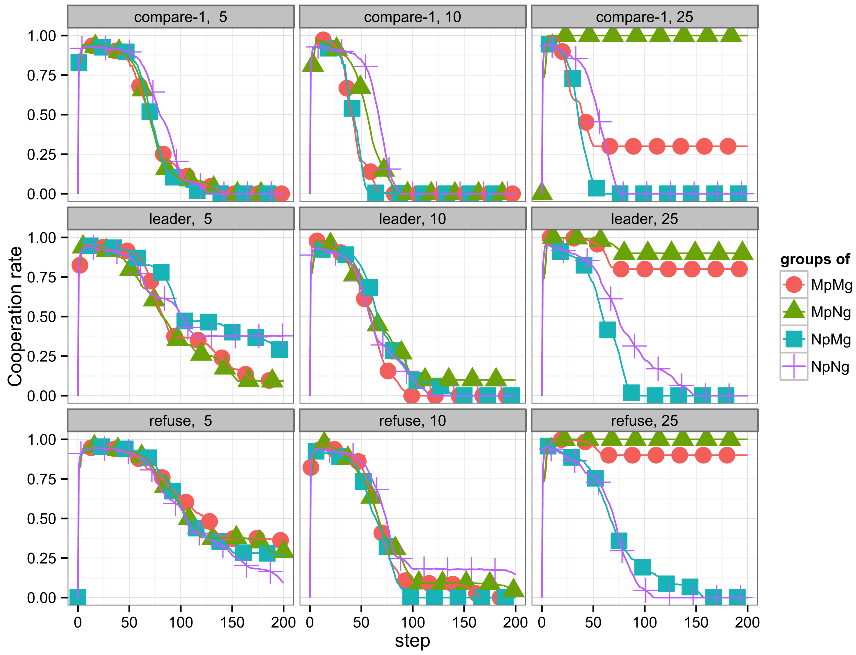

Fig. 6. Cooperation rates for mixed populations in which Punishers and Gossipers strategies are combined. Each line shows the global cooperation rate for a single experiment. The highest cooperation rates are reached in the population in which punishers also defect against known free-riders, whereas Gossipers adopt a milder strategy. This combination outperforms all the other strategies in small and large groups.

and nice gossip in large groups. This suggests that once mean punishers had reduced the number of free-riders, a cooperating gossip (for the *compare* response) or any kind of gossip (for *leader* and *refuse*) was effective in promoting cooperation in large groups.

The conditional inference tree analysis [16], as shown in Figure 7, confirms that in large groups cooperation could reach values close to or equal to 1 when the reactions were in combination, like Mean Punishment-Nice Gossip and gossipers played either *compare* or *refuse* (leaves 10 and 11). This supported our hypothesis about the effectiveness of partner choice in promoting cooperation, especially when cheaters were selected out by punishers.

5 Discussion and future work

Among humans, gossip is a powerful tool for social control, and information transmission might have played a crucial role in the evolution of cooperation, as suggested by Robin Dunbar:

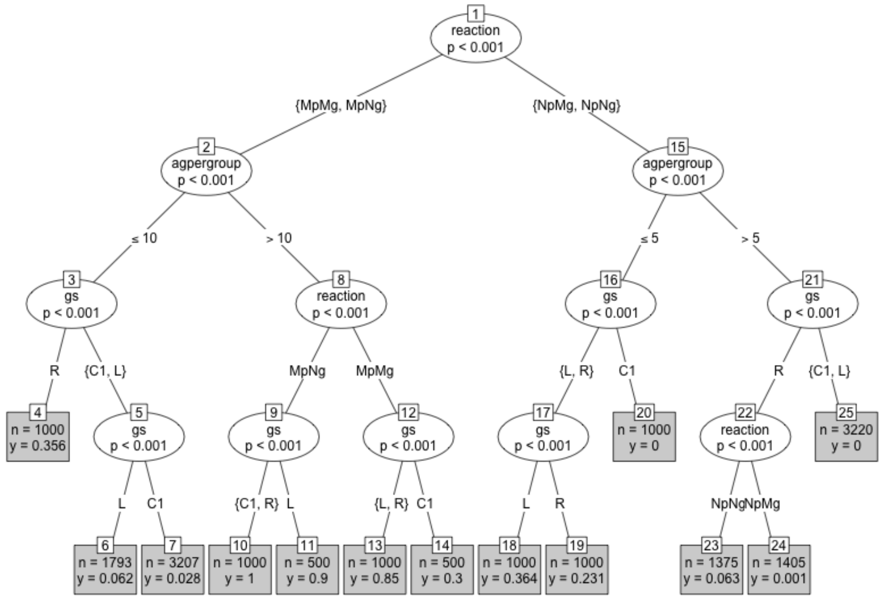

Fig. 7. The conditional inference tree on data for the last 20 steps of the second experiment shows that the main effect on cooperation rates is connected with Mean Punishment (*Mp*) in combination with both Mg and *Ng*. The second level of discrimination is due to group size, and the third mostly to reaction strategies. In general, the *Mp* branch leads to higher levels of cooperation with the exception of small groups for *C1* and *L* reactions.

> 'Lacking language, monkeys and apes are constrained in what they can know ... But language allows us to seek out what has been going on behind our backs. Indeed, we can even be proactive about it and tell our friends and relations what we have seen when we think it might be in their interests to know' ([4], p.103).

In the last decade, research on reputation-supported cooperation has unveiled the importance of getting information about others' past behaviors through direct observation, but less is known about the role played by transmission. In an environment in which free-riders can be easily spotted because of some visible marker, like a score publicly visible, cooperation can easily emerge, but this model cannot be extended to human societies, where free-riders do not show any mark and they also have an incentive in concealing their tendencies and behaviors. In this work, we modeled gossip as information transmission among agents and we linked it with three different ways of using the information received. When we discover that someone is a free-rider we can react in several ways, like refraining from interaction (strategy *refuse* in our model), or avoiding that person and joining another group (strategy *compare*). There are also cases in which humans can actively select their partners, creating a new group in which only reliable cooperators are present (strategy *leader*).

Our data provide some additional insights into the role of gossip spreading on cooperation levels in mixed populations in which gossipers can transmit information and tune their behaviors on the basis of information received from their peers. We show that cooperation rates are higher when agents can compare their present situation and switch to a better one, i.e. they can avoid free-riders, and this solution allows gossipers to get the highest scores in large groups of 25 agents. We also show that the combination of punishment and gossip can lead cooperation to its maximum in large groups, irrespective of the specific gossip strategy. This result is especially interesting because it is in line with ethnographic studies of human societies in which material punishment and gossiping about free-riders usually go hand in hand [5].

Building upon previous work in which we modelled gossip in a more idealtypical way [11], here we made an effort towards a more realistic modeling of gossip spreading, allowing agents to spread gossip at the beginning of each encounter, informing their peers (other gossipers) about the identity of known free-riders in the group. Preventive gossip is a way of warning one's peers against the risks of exploitation, but additional work is needed in order to identify the best conditions for the emergence of gossip. The model that we have developed is just one of several steps required in that direction. It should be supported by experimental and observation data, and, possibly with the help of these data, get refined and replicated.

Acknowledgements

We gratefully acknowledge support from PRISMA project, within the Italian National Program for Research and Innovation (Programma Operativo Nazionale Ricerca e Competitivitá 2007-2013. Settore: Smart Cities and Communities and Social Innovation Asse e Obiettivo: Asse II - Azioni integrate per lo sviluppo sostenibile)

References

1. Richard Alexander. *The Biology of Moral Systems (Foundations of Human Behavior)*. Aldine Transaction, July 1987.
2. Robert Axelrod. *The evolution of cooperation*. Basic Books, 1984.
3. Jeffrey P. Carpenter. Punishing free-riders: How group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior*, 60(1):31–51, July 2007.
4. R. I. M. Dunbar. Gossip in evolutionary perspective. *Review of General Psychology*, 8(2):100–110, 2004.
5. Robert Ellickson. *Order without Law : How Neighbors Settle Disputes*. Harvard University Press, June 2005.

6. Ernst Fehr and Simon Gächter. Cooperation and punishment in public goods experiments. *The American Economic Review*, 90(4):980–994, 2000.
7. Ernst Fehr and Simon Gachter. Altruistic punishment in humans. *Nature*, 415(6868):137–140, January 2002.
8. Ernst Fehr and Frédéric Schneider. Eyes are on us, but nobody cares: are eye cues relevant for strong reciprocity? *Proceedings. Biological sciences / The Royal Society*, 277(1686):1315–1323, May 2010.
9. Francesca Giardini and Rosaria Conte. Gossip for social control in natural and artificial societies. *SIMULATION*, 88(1):18–32, January 2012.
10. Francesca Giardini, Mario Paolucci, and Rosaria Conte. Reputation. In Bruce Edmonds and Ruth Meyer, editors, *Handbook on Simulating Social Complexity*, Understanding Complex Systems, pages 573–577. Springer-Verlag, Heidelberg, 2013.
11. Francesca Giardini, Mario Paolucci, Daniel Villatoro, and Rosaria Conte. Punishment and gossip: sustaining cooperation in a public goods game. In Bogumil Kaminski and Grzegorz Koloch, editors, *Advances in Social Simulation: Proceedings of the 9th Conference of the European Social Simulation Association*. Springer-Verlag, 2013.
12. Volker Grimm, Uta Berger, Finn Bastiansen, et al. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1-2):115–126, September 2006.
13. Francesco Guala. Reciprocity: weak or strong? what punishment experiments do (and do not) demonstrate. Departmental Working Papers 2010-23, Department of Economics, Business and Statistics at Università degli Studi di Milano, July 2010.
14. Garrett Hardin. The Tragedy of the Commons. *Science*, 162(3859), 1968.
15. Douglas D. Heckathorn. Collective action and the Second-Order Free-Rider problem. *Rationality and Society*, 1(1):78–100, July 1989.
16. Torsten Hothorn, Kurt Hornik, and Achim Zeileis. Unbiased Recursive Partitioning. *Journal of Computational and Graphical Statistics*, 15(3):651–674, September 2006.
17. John O. Ledyard. Public goods: A survey of experimental research. In J. H. Kagel and A. E. Roth, editors, *Handbook of Experimental Economics*, pages 111–194. Princeton University Press, 1995.
18. John Maynard-Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1st edition edition, December 1982.
19. Martin A. Nowak and Karl Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393(6685):573–577, June 1998.
20. Karthik Panchanathan and Robert Boyd. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature*, 432(7016):499–502, November 2004.
21. I. Pinyol, M. Paolucci, J. Sabater-Mir, and R. Conte. Beyond accuracy. reputation for partner selection with lies and retaliation. In *Proceedings of the MABS'07. Hawaii, USA.*, volume 5003 of *LNCS*, pages 128–140. Springer, 2007.
22. Bettina Rockenbach and Manfred Milinski. The efficient interaction of indirect reciprocity and costly punishment. *Nature*, 444(7120):718–723, 2006.
23. Claus Wedekind and Manfred Milinski. Cooperation through image scoring in humans. *Science*, 288(5467):850–852, May 2000.
24. U. Wilensky. Netlogo. http://ccl.northwestern.edu/netlogo/., 1999. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.