

# Reputation and Punishment sustain cooperation in the Optional Public Goods Game

Shirsendu Podder<sup>1</sup>, Simone Righi<sup>2</sup>, and Francesca Pancotto<sup>3</sup>

<sup>1</sup>Department of Computer Science, University College London, United Kingdom

<sup>2</sup>Department of Economics, Ca’Foscari University of Venice, Italy

<sup>3</sup>Department of Communication and Economics, Università degli Studi di Modena e Reggio Emilia

## Abstract

Cooperative behaviour has been extensively studied as a choice between cooperation and defection. However, the possibility to not participate is also frequently available. This type of problem can be studied through the optional public goods game. The introduction of the “Loner” strategy, allows players to withdraw from the game, which leads to a cooperator-defector-loner cycle. While prosocial punishment can help increase cooperation, anti-social punishment – where defectors punish cooperators – causes its downfall in both experimental and theoretical studies.

In this paper, we introduce social norms that allow agents to condition their behaviour to the reputation of their peers. We benchmark this both with respect to the standard optional public goods game and to the variant where all types of punishment are allowed. We find that a social norm imposing a more moderate reputational penalty for opting out than for defecting, increases cooperation. When, besides reputation, punishment is also possible, the two mechanisms work synergically under all social norms that do not assign to loners a strictly worse reputation than to defectors. Under this latter setup, the high levels of cooperation are sustained by conditional strategies, which largely reduce the use of pro-social punishment and almost completely eliminate anti-social punishment.

**Keywords**— Reputation, anti-social punishment, optional public goods game

# 1 Introduction

Explaining the conditions for the evolution of cooperation in groups of unrelated individuals is an important topic for natural and social scientists (Axelrod and Hamilton, 1981; Rainey and Rainey, 2003; Griffin et al., 2004; Nowak, 2006; Perc et al., 2017). Research on the subject has been extensive and has led to the identification of several mechanisms contributing to the success of cooperative behaviour: direct and indirect reciprocity, spatial selection, kin and multi-level selection, and punishment (Nowak, 2006) among others. A large number of situations in which cooperation is difficult to achieve involve the pursuit of collective action to address problems beyond the individual dimension (Smith, 1964). Such problems are frequently studied through variations of the public goods game, where the public goods thrive through the contributions of its participants. While the social optimum is achieved when everyone’s contribution is maximal, individuals always have the temptation to withhold their contributions, free-riding on the efforts of others (Hardin, 1968).

The public goods game has been mainly studied as a choice between contributing (or cooperating) and free-riding (or defecting) (Ledyard, 1995). However, in many situations, the possibility to not participate (or to exit) a situation is also available due, for example, to ostracism (Maier-Rigaud et al. 2010) or to individual choice (Hauert 2002). An individual adopting this strategy chooses to abstain altogether from interactions, giving up not only its costs but also its potential benefits, instead preferring a lower but guaranteed payoff. For this reason, it is frequently called the “Loner” strategy.

This behaviour has been observed in several social animal species. When prides of lions hunt, individuals that actively pursue prey and contribute to their capture can be characterised as cooperators, members participating in the hunt but remaining immobile (or covering only very short distances) can be characterised as defectors, while the members hunting alone can be considered loners (Scheel and Packer, 1991). Similarly, groups of primates frequently display the practice of mutual grooming. In such situations, some individuals both actively receive grooming and groom others (a costly activity in terms of time and effort), while some fail to reciprocate the received benefits. Finally a third category of individuals tend to groom independently, refusing both to help and to accept help from others (Smuts et al., 2008). In humans, loners are those individuals that abstain from participation, from both the construction and from the benefits of the public good, for example refusing to participate in village activities involving conspicuous consumption (Gell, 1986), deciding not to join a business alliance, or quitting a secure job rather than intentionally reducing effort. Finally, at the institutional level, countries can decide to forgo joining geo-political alliances, rather than joining and then not engaging in their activities.

The forces influencing decision making in these situations can be captured in a game theoretical form through the Optional Public Goods Game (OPGG), which introduces the loner strategy into the traditional public goods game. This addition generates significant changes in the game dynamics, undermining the strength of defectors (who cannot exploit loners), while introducing cyclical population shifts. Indeed when most players cooperate, defection is the most effective strategy. When defectors become prevalent, it pays to stay out of the game and to resort to the loner strategy. Finally loner dominance lays the foundation for the return of cooperation (Hauert, 2002; Hauert et al., 2002). Such *rock-paper-scissors* dynamics have been confirmed in experiments (Semmann et al., 2003). Relatedly, the possibility of opting out increases cooperation through enhanced prosociality of those who decide to stay in the game (Orbell and Dawes, 1993), through the exit threat (Nosenzo and Tufano, 2017), and through partner selection (Hauk, 2003).

Two mechanisms have been shown to be particularly important in influencing the ability of a group to engage in effective collective action: punishment and reputation. The former can be described as people willing to sustain an individual cost to punish behaviour that they find inappropriate (Fehr and Gächter, 2002). Experiments (Fehr and Gächter, 2000; Henrich, 2006) have found that pro-social punishment (the punishment of a defector by a cooperator) – whose emergence has been linked to the possibility to abstain from participation (Hauert et al., 2007) – can sustain cooperation in iterated games. More recently however, the positive role of punishment has been criticised, as substantial levels of anti-social punishment (the punishment of cooperators by defectors) have been observed in experiments (Herrmann et al., 2008; Rand et al., 2010; Rand and Nowak, 2011). Although second-order free-riding on anti-social punishment can restore the effectiveness of pro-social punishment in the absence of loners (Szolnoki and Perc, 2017), allowing this type of punishment significantly reduces cooperation, and re-establishes the rock–paper–scissor type of dynamics among strategies (Rand and Nowak, 2011) in the OPGG. Importantly, Rand and Nowak (2011)

show that the strategy of opting out of interactions, combined with a world in which anyone can punish anyone, does not result in any meaningful increase in cooperation. Crucially, it is not the presence of loners that harms cooperation, but the possibility to punish them. Indeed, if loners are shielded from punishment, García and Traulsen (2012) show that cooperators that punish pro-socially prevail, even when anti-social punishment is available. Just like anti-social punishment, anti-social pool rewarding (where agents of some type contribute towards rewarding others of their own type) also destabilises cooperation in fully-mixed populations (dos Santos, 2015) but not structured populations (Szolnoki and Perc, 2015).

Besides punishment, the other key mechanism in sustaining cooperation towards collective action is reputation (Milinski et al., 2002). Adopting a loose definition, reputation emerges when an individual’s actions can be directly or indirectly observed by his peers and used to condition their own behaviour when playing with him. In both well-mixed (Sigmund et al., 2001; Hauert et al., 2004) and structured populations (Brandt et al., 2003), simple reputational systems promote and stabilise cooperation. Reputational systems at different levels of complexity have been studied, ranging from the first-order image scoring norm (Nowak and Sigmund, 1998), to the second-order standing criterion (Leimar and Hammerstein, 2001; Panchanathan and Boyd, 2003) and to the more complex third-order *leading eight* social norms (Ohtsuki and Iwasa, 2004; Ohtsuki et al., 2006). While simpler reputational systems have been observed in animals (Bshary and Grutter, 2006), more elaborate systems are more likely to be the domain of human interactions, due to the complex relationship between actions, reputations and social structures. How reputation evolves when a given behaviour is observed depends on the social norm characterising a population, in other terms their “notion of goodness”, i.e. from the moral value attributed to a given type of action.

Punishment devices and reputation dynamics based on social norms coexist in human social groups (Jordan and Rand, 2020) and therefore their interaction and implication for the success of collective action needs to be studied. In this paper, we extend the theory of OPGG studying the individual and joint impact of punishment and reputation on the sustainability of cooperation in a population where groups of agents engage in repeated interactions. In line with evidence that anti-social punishment exists and heavily influences cooperation (Rand and Nowak, 2011), we allow for all types of punishment to occur, while studying its joint impact with a simple (first-order) social norm that prizes cooperation and penalises the loner and defection choices through bad reputations. By comparing the effect of social norms differing by the relative reputations they assign to defectors and loners, we explore the strategic exchanges between punishment and reputation. Similarly to Panchanathan and Boyd (2004), we consider ephemeral groups of agents that interact repeatedly during each time period and are then reshuffled.

We find that while reputation alone mildly increases cooperation, and punishment alone does not, the two mechanisms synergistically interact leading to high cooperation. Our findings stem from the fact that reputation emerges as a partial substitute of punishment, thus making all types of punishment less necessary. Our results indicate a way through which social groups can sustain cooperation despite the presence of anti-social punishment.

## 2 The Model

A fully-mixed population of  $N$  agents play the OPGG in randomly chosen groups of  $n \geq 4$  as in Hauert (2002). During each interaction, players are given the choice between cooperating (C), defecting (D), or abstaining from the game (L). Cooperators incur a cost  $c$  to participate and contribute towards the public good, while defectors participate but contribute nothing. Loners withdraw from the game to receive the loner’s payoff  $\sigma$  with  $0 < \sigma < (r - 1)c$ . In each game, the sum of the contributions is multiplied by a factor  $r > 1$  and is distributed equally between the group’s participants. If there is only a single participant (or equivalently  $n - 1$  loners in the group), the OPGG is cancelled and all players are awarded  $\sigma$ . At this point, payoffs are awarded and additional games are played among the same agents until, with probability  $1 - \Omega$ , the interactions terminate and the group is dissolved. Note that, unlike Milinski et al. (2002), our setup can explore a form of indirect reciprocity without the need of adding a different game with respect to the OPGG.

Following each round of the game, players have the opportunity to punish any or all of the other  $n - 1$  members of the group based on their played action. Players pay a cost of  $\gamma$  for each player they choose to punish, while they are subjected to a penalty of  $\beta$  for each punishment they receive.

Once the interaction and punishment rounds end, players change strategy based on the distribution of payoffs in the population. Each time-step an agent  $i$  is randomly paired with another agent  $j$  outside (within) its group with probability  $m$  (respectively  $1 - m$ ). Player  $i$  then imitates player  $j$  with probability  $\frac{\exp(u_j)}{\exp(u_i) + \exp(u_j)}$ , where  $u_x$  is the payoff of player  $x$  and the exponential payoff function is used to account for negative utilities. Additionally, we allow mutations to occur in each period with probability  $\epsilon$ .

Extending the OPGG with the classical reputation mechanism of Brandt et al. (2003), we implement a reputational system adopted by the entire population. After a player cooperates, defects, or abstains from the OPGG, he is assigned a reputation according to the population's social norm. A reputation system requires the definition of a social norm, or reputation dynamics (Ohtsuki and Iwasa, 2004), that identifies good and bad behaviour. The simplest possible reputational mechanisms are those that assign reputation on the basis of the mere observation of past actions, such as image scoring (Nowak and Sigmund, 1998). Arguably, most functional human and animal societies associate the highest reputation to cooperative behaviour. For this reason, we focus our attention to the reputation dynamics that prize cooperation above other actions. Like the binary image score, we simplify the reputation values so that players can have only a good (1), intermediate (0) or bad (-1) reputation, always assigning a good reputation to cooperators. We label the *Anti-Defector (AD)* norm assigning reputations of -1 to defectors and 0 to loners. Similarly, the *Anti-Loner (AL)* norm assigns loners a reputation of -1 and defectors a 0. Finally, it is possible to conceptualise norms that do not distinguish between acting as a loner or as a defector. We label these as the *Anti-Neither (AN)* social norm, which assigns 0 to both actions, and the *Anti-Both (AB)* social norm, which assigns -1 to both actions. Similar to punishment, reputations update after every OPGG round. We assume that actions are either observed without error or communicated honestly to everyone so that players' resulting reputations are common knowledge within the population.

To model agents adopting a social norm, we extend the set of behavioural strategies to allow for players to condition their actions on the average reputation of the other players within their group. We condition actions on the average reputation firstly because a player's payoff is dependent on the actions of his group members and secondly because directly observing the actions of others can be difficult in sizeable groups, whereas observing the overall sentiment is arguably easier. Conditioning actions to the average reputation is also a key difference between our setup and Boyd and Richerson (1988) where actions depend on levels of cooperation, and from Takezawa and Price (2010), where they depend on the amounts contributed. Indeed, given that there are three types of actions, each receiving a reputation on the basis of a shared social norm, reputational information can only be used as a signal of the cooperativeness of the environment a player is facing, similar to the case of image scoring. Furthermore, in our model reputation contains both elements of direct and indirect reciprocity as when groups are reshuffled, reputations are carried over. Finally, we significantly extend the strategy set considering strategies whose primary action in a benign environment is not cooperating. The strategy  $X_{Z_C Z_D Z_L}^{k_{\min}, Y}$  defines an agent taking action  $X \in \{C, D, L\}$  if the average reputation of the other members of the group strictly exceeds  $k_{\min} \in [-1, 1)$ , and  $Y \in \{C, D, L\}$  otherwise. Furthermore, the subscripts describe the agent's strategy in terms of its punishment decision taken against cooperators  $Z_C$ , defectors  $Z_D$ , and loners  $Z_L$ , where  $Z_C, Z_D, Z_L \in \{P, N\}$  is the decision to punish or not to punish respectively. In addition to the pure strategies of being an unconditional cooperator, defector and loner, this adds a further 8 strategies:  $C^{-1,D}, C^{0,D}, C^{-1,L}, C^{0,L}, D^{-1,L}, D^{0,L}, L^{-1,D}, L^{0,D}$ , each of which have 8 variants prescribing punishment towards zero, one or more cooperators, defectors and loners. In an effort to minimise the size of the strategy space, and therefore also some of the effects of random drift, we remove counterintuitive strategies, i.e. those that would cooperate only in groups of defectors or loners, defecting or going loner when surrounded by cooperators. For the same reason, strategies of higher complexity, i.e. those implying multiple behavioural shifts and/or more than two possible actions, are excluded from the present analysis and relegated to further studies.

Combining the payoffs of the game and the subsequent punishments, the utility of a player that decides not to opt out after a single round of the OPGG is:

$$u_i = w_i - x_i + r \frac{\left(\sum_j x_j\right)}{n_p} - \beta P - \gamma P'$$

where the first two terms represent the initial endowment ( $w_i$ ) and the initial contribution ( $x_i \in \{0, c\}$ ) of the players. The third term represents the return from the OPGG (where  $n_p$  is the number of agents

that decided to participate in the game, i.e. to not opt-out), and the final two terms represent the losses incurred from punishing ( $\beta$ ) and being punished ( $\gamma$ ). In broad terms,  $P$  represents the number of people in the group who exhibited the actions that player  $i$  punishes, and  $P'$  represents the number of people within the group that punish player  $i$ 's most recent action. For agents that decide to opt-out, the utility reduces to  $u_i = \sigma - \beta P - \gamma P'$ . In summary, if we ignore reputation and punishment, we have 3 strategies, if we only have punishment, we have 24, if we only have reputation, we have 11, and if we have both reputation and punishment, we have 88 strategies.

### 3 Results

Investigating the dynamics of the OPGG (Fig. 1 and Fig. S11 in the Supplementary Materials), in the absence of both punishment and any reputational element, we reproduce the rock-paper-scissor dynamics between the three unconditional strategies as in Hauert (2002) and Hauert et al. (2002). When aggregated over time, this results in agents cooperating or defecting about 20% of the time each, while operating as loners about half of the time. Introducing the full set of punishing strategies as in Rand and Nowak (2011) we successfully reproduce their main result: the proportion of defectors and cooperators further diminishes with respect to the baseline, whilst the proportion of agents acting as loners increases. Having successfully replicated the key results of the extant literature, we investigate two models benchmarked respectively against the traditional OPGG and the OPGG with the full spectrum of punishment. In the first manipulation, we introduce a reputational system to the OPGG, while in the second we add it to the OPGG with punishment. In both cases we report results averaged over the second half of 100 simulations each comprising 200000 time steps.

**Reputation without punishment** We begin by analysing how the introduction of a simple reputational system influences the dynamics of these populations. The AD norm increases cooperation, so that in this circumstance, about 40% of the OPGG actions were cooperative, roughly a 20% increase from the baseline (Fig. 1). Interestingly, the increase in cooperators does not come from the defector ranks, but rather from the loner ranks, which decrease the most with respect to the baseline. The increase in cooperation follows from a change in the patterns of strategy adoption caused by the introduction of reputation. Strategy adoption patterns are visualised, for each social norm studied, in Fig. 2 which also reports the reduced transition matrices (including only the strategies that are - on average - played by more than 10% of the population). Results are computed averaging across time and then across 100 repetitions of each setup. The full population breakdown as well as the complete transition matrices are reported in the left panels of SM Figs. S1 - S4. As noted in SM Figs. S8 and S9, the average population composition is remarkably stable across simulations, ensuring that the results are generalisable beyond a single simulation.

In the case of the AD social norm (that generates the highest levels of cooperation), pure strategies are largely abandoned in favour of conditional strategies. Among the 8 conditional strategies, only two are adopted by a sizeable part of the population:  $C^{0,L}$  and  $D^{0,L}$ . These are strategies that share a common goal: to participate in interactions with a population of agents who are sufficiently cooperative (since they both set the reputational threshold for participation in their group to 0). However, while the former strategy wishes to engage with and enjoy the collaboration of their peers, the latter aims to exploit them. When this is not possible, both strategies act as a loner for a lower but guaranteed payoff. From the transition table (SM Fig. S1, left panel), it emerges that most strategies turn into either  $C^{0,L}$  or  $D^{0,L}$ , and that these two strategies are much more likely to turn into each other than anything else (leftmost panel of Fig. 2). This is a consequence of the interaction of these strategies and the level of cooperation in the population. This effect can be clearly discerned studying the temporal evolution of the proportion of agents playing these two strategies and of the associated actions played (Fig. S14, panel C). While the proportion of  $D^{0,L}$  and  $C^{0,L}$  in the population move systematically in opposite directions, studying the actual actions played we observe oscillations in the proportion of opt-out and cooperative actions, while defection is played by fewer individuals in a stable manner. In other terms the  $C^{0,L}$  strategy uses the reputational information to conditionally cooperate. When in a group there happens to be too many defectors (most of whom are  $D^{0,L}$ s given their relative prevalence) the strategy resorts to opting out, and thus avoiding being exploited. The strategy  $D^{0,L}$  acts similarly, thus avoiding being completely evolved out, but it is unable to spread the defection in the population. Importantly,




Figure 1: **Cooperation is highest when being a defector is strictly less reputable than being a loner, with and without punishment. Punishment and reputation together sustain high levels of cooperation as long as being a loner is not the strictly least reputable action. Without punishment, only the AD social norm sustains cooperation, with all other social norms providing comparable levels of (low) cooperation with the traditional OPGG in the absence of reputation and punishment.** All simulations shown use  $N = 1000$ ,  $n = 5$ ,  $r = 3$ ,  $\sigma = 1$ ,  $m = 0.95$ ,  $\epsilon = 0.1$ ,  $\Omega = 10/11$ . Those including punishment additionally have  $\gamma = 1$ ,  $\beta = 2$ . Results are calculated by averaging across the second half of the simulations, and then across 100 repeated simulations.

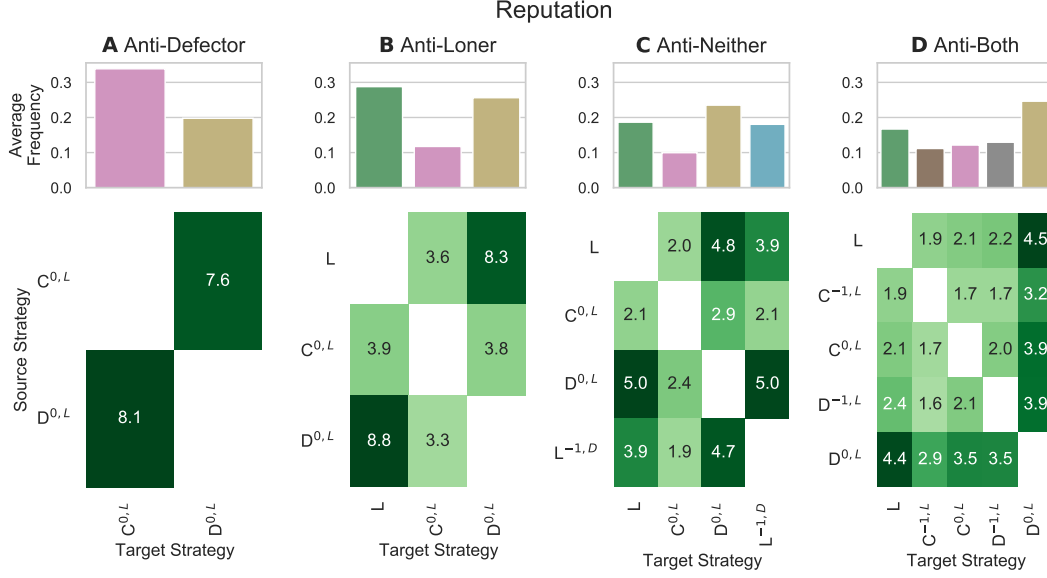


Figure 2: In the absence of punishment, cooperation is sustained only when defection results in a strictly worse reputation than opting out. In this setup, a conditionally cooperating and a conditionally defecting strategy dominate the population.

this dynamic relies on the fact that under the AD social norm, agents playing defection are assigned a strictly worse reputation than loners, thus allowing  $C^{0,L}$  to discriminate between the two behaviours. In other terms, the cyclic behaviour of the baseline model - described by Rand and Nowak (2011) - reemerges, but with the loner strategy being part of two more complex conditional strategies. Importantly, the cycles are - in this setup - quite noisy due to the occasional attractiveness (Fig. S1, left panel) of other strategies, which allows them to be the destination of some transitions and thus to survive albeit in low numbers.

The peculiarity of AD is immediately evident when studying the other social norms. While the anti-defector social norm results in both the greatest population fitness (Fig. S10) and increase in cooperation, all other social norms are unable to sustain levels of cooperation higher than the baseline, resulting in populations largely dominated by loners (Fig. 1) with a cyclical pattern similar to those of Rand and Nowak (2011). The reason for this divergence lays in the reputational damage inflicted by these norms to non-cooperators. Indeed while reputation helps agents condition their behaviour to opponents' types, attributing loners a bad reputation does not inflict on them any direct damage. Since they are already prone to not participating in the game, their payoffs remain independent from any reputational effect. However, attributing loners too low a reputation (in AL and AB) hampers the role of loners in providing a backup option against invading defectors. Furthermore, not attributing defectors a strictly lower reputation (in AN), allows them to more easily exploit cooperators by impeding the latter to discriminate between loners and defectors (the full transition matrices are reported in SM Figs. S1 - S4). In all these cases the effectiveness of  $C^{0,L}$  against defection is more limited hence the strategy has limited evolutionary success (Fig. 2, panels B, C, D and SM Fig. S15). For this reason, in the following we will focus our attention mainly on the AD social norm.

**Reputation with punishment** We now analyse the impact of introducing the AD norm to the OPGG with punishment. Due to the substantial empirical and theoretical evidence for the presence of anti-social punishment in the OPGG, we allow all possible types of punishment to occur. In the presence of the AD norm, the level of cooperation in the population increases significantly with cooperation becoming the most played strategy, at the expense again of loners whose role becomes marginal in this setup. The level of cooperation is not only higher than when solely punishment is used, but is also of the case where only reputation is

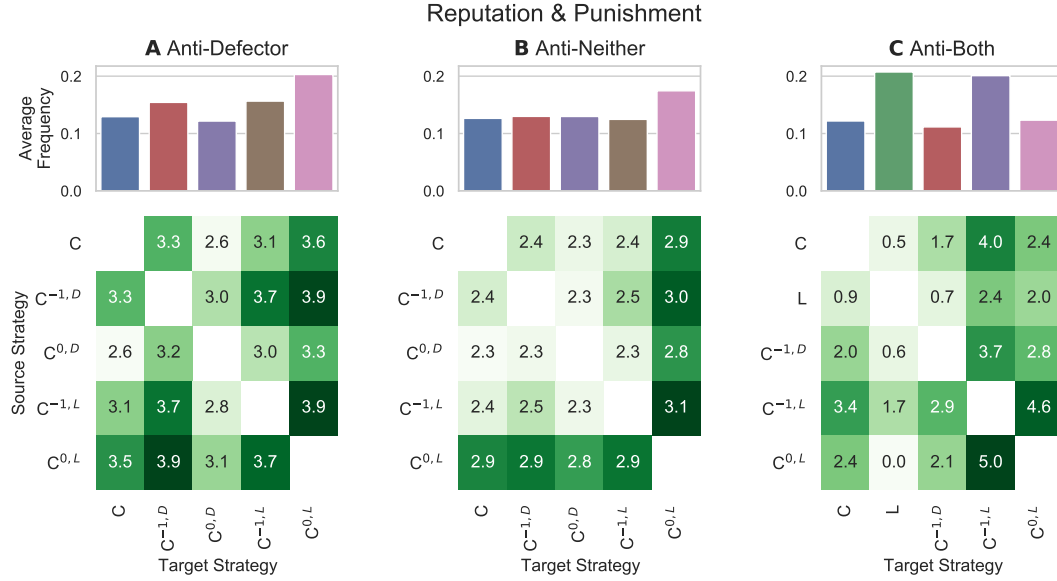


Figure 3: **When punishment is available, high levels of cooperation evolve as long as being a loner is not the strictly least reputable action. Strategies that thrive in these situations conditionally cooperate, becoming indistinguishable between one other in such a cooperative environment.** Under the AB norm, both defecting and exiting are the least reputable actions which creates both the most hostile environments and the least reputable OPGG groups. This results in conditional cooperators being cautious by frequently opting out. The AL social norm is not pictured as the only long-term dominant strategy that exists in significant amounts is the loner strategy. See SM Figs. S1-S4 for the complete transition matrices.



available. Given the success of this cooperation, one could expect punishment and the reputation system to interact in some form, coevolving to produce successful strategies that both punish pro-socially and cooperate conditionally. However, this is largely not the case. Analysing the prevalence of strategies and the transition probabilities between them (Fig. 3 for between strategy transitions, and Fig. 4 for transitions between different punishment patterns for the same basic strategy), we note that – for each of the conditional strategies that become dominant in the population – the variants which employ punishment are quite rare. When punishers are born through mutation, they rapidly turn into their non-punishing counterparts.

While representing consistently low proportions of the population, the types of punishment that show more resilience (being sometimes able to attract non-punishing counterparts), are those that punish defectors and/or loners, i.e. the pro-social types (see SM Figs. S1 - S7 for the complete transition matrices and the prevalence of each strategy within each model and social norm). Thus when both punishment and reputation exist, the latter acts as a substitute for the former. Agents evolve toward strategies that do not punish and away from strategies that do, contributing to the demise of this kind of behaviour, except as a threat against deviators.

Regardless of punishment patterns, the conditional strategies that emerge as prevalent are those whose primary action is to cooperate, which in turn provides the ideal environment for unconditional cooperators to thrive. Indeed, the combination of punishment and reputation ensures that agents that adopt strategies whose primary action is different from cooperation, end up with lower average payoffs given that the established cooperative conditional strategies will be less likely to cooperate with individuals whose primary action is not cooperation. The transition matrix of Fig. 3, shows that the surviving strategies shift between each other without a clear pattern. This is likely due to fluctuations in the group composition when playing the OPGG as these strategies are almost equivalent in largely cooperative environments.

Besides AD, the other social norms excluding AL are also able to increase and stabilise cooperation albeit to a lesser extent (Fig. 1). These social norms produce very similar results to the AD social norm both concerning the emerging conditional strategies (Fig. 3) and the prevalence of pro-social punishment (SM Figs. 12 and 13). The lower the reputation attributed to loners, the lower the frequency of cooperation and population fitness (Fig. S10). Comparing AD and AB against AN social norms (Fig. S12) shows that as the reputability of being a defector increases, the level of pro-social punishment increases to compensate, thus minimising the damage to cooperation and to the fitness of the population. It follows that populations using comparatively more punishment in addition to reputation have lower average fitness.

Finally, for the same reasons as discussed in the reputation only setup, AL fails to stabilise cooperation altogether, resulting in large fluctuations of cooperation levels, high proportions of the loner strategy in the population, and the lowest payoffs among the social norms both with and without punishment.

**Parameter robustness** Our main results discussed so far are robust to an extensive set of parameter variations. These are reported in the Supplementary Materials, Figs. S17 - S24 where the effects of changes to the group synergy factor  $r$ , the loner’s payoff  $\sigma$ , the punishment cost  $\gamma$  and penalty  $\beta$ , probability of repeated interactions  $\Omega$ , OPGG group size  $n$ , the degree of evolutionary group mixing  $m$ , and the rate of mutation  $\epsilon$  are systematically explored (the parameterisations used for each robustness check are reported in SM Table S2). The superiority of combining the reputation and punishment mechanisms remain consistent with our findings, showing higher levels of cooperation when compared to our baseline model without reputation or punishment, and to the baseline model with only punishment. While the model utilising solely reputation does tend to increase cooperation, its behaviour is somewhat more dependent on specific parameter ranges of the degree of mixing  $m$  (SM Fig. S23) and of the stability of groups  $\Omega$  (SM Fig. S21). Cooperation increases when the degree of mixing  $m$  increases, likely due to the pool of available strategies for agents to choose to evolve to. When  $m$  is small, players evolve based on a randomly selected player within their own OPGG groups. The likelihood of a fruitful and cooperative player to exist within this group is lower than when  $m$  is high and players evolve probabilistically based on the payoffs of a randomly selected individual from outside of their group. Instead, considering the likelihood of further interactions  $\Omega$ , in the setup in which both reputation and punishment are active, as long as  $\Omega \neq 0$ , cooperation has a high and similar proportion within the population. This suggests that a combination of both direct (relying on direct experience within the group to condition action) and indirect reciprocity (rely on third party information to condition action) is important in obtaining high levels of cooperation.




Figure 4: **When both punishment and reputation are available, most individuals do not punish, a few punish pro-socially, and almost no one punishes anti-socially.** Nodes represent the punishment variants of each of the 5 most popular strategies within the reputation and punishment model (AD social norm). The values within the nodes are the average proportion of the variant relative to all variants of the strategy. Edges represent the transitions solely between variants, normalised by the total number of transitions between all punishment variants of the strategy. The strategies that do not punish are always the most popular of the variations, those that punish some combination of defectors or loners are common, while those that punish cooperators never survive for long. See SM Fig. S5 for more details.

## 4 Discussion

The OPGG provides a challenging environment for cooperation. The three-way interaction between cooperators, defectors and loners produce natural cyclical dynamics, where cooperative environments favour defection, which in turn makes opting out more advantageous, which in turn leads to the return of cooperation (Hauert, 2002; Hauert et al., 2002). Similar challenges are faced introducing the possibility of agents punishing one another based on their past actions. In this case the cyclical dynamics involve both the presence of pro-social (punishment of defectors and loners) and of anti-social punishment (punishment of cooperators), with each strategy punishing its would-be invader (Rand and Nowak, 2011). In either case, the average cooperation level resulting from the cyclical behaviour is low.

Introducing into the OPGG the AD social norm (assigning cooperators a good reputation and assigning better reputations to agents that opt-out from the game than to those who participate but withhold their contribution), significantly increases cooperation levels unlike the other reputational systems. The effectiveness of the AD norm depends on the fact that it provides a way for agents to condition their behaviour to the propensity of their group to cooperate, while reducing the attractiveness of the loner action, which remains as a backup if the environment becomes too hostile. Accordingly, within the ecosystem of strategies generated by the introduction of this norm, over half of the population end up preferring one of two strategies,  $C^{0,L}$  or  $D^{0,L}$ . Both act as loners if the environment becomes too harsh, but turn to cooperation or defection in a more friendly world. The backup option of abstaining from the OPGG on the one hand prevents defectors from over-exploiting cooperators: if they are present in sufficient numbers, their bad reputation induces  $C^{0,L}$  to become a loner. On the other hand, the same rule supports a conditional defector who is even evolutionarily stronger than the unconditional counterpart, as it viciously exploits reputation and cooperative behaviour while still receiving the loner's payoff when the likelihood of exploiting its group decreases. The evolutionary success and average payoffs of these two strategies are similar, hence agents cycle between them, with their equilibrium favouring conditional cooperators, hence the observed increase in cooperation. It should be noted that in this setup, many other strategies survive - albeit in lower numbers - thus the dominance of the aforementioned strategies is never complete. It is interesting to note that the laxer counterparts of the strategies discussed, namely  $C^{-1,L}$  and  $D^{-1,L}$  are less successful, as they propose cooperation or defection even in groups largely dominated by defectors.

Cooperation turns out to be much more likely when we include the interaction between costly punishment

and a reputational system into the mix, assigning a better reputation to agents that opt-out from the game than to those that participate but withhold their contribution. It should be noted that the emergence of high levels of cooperation in this setup is surprising as the presence of anti-social punishment has been found to reduce the levels of cooperation (Rand and Nowak, 2011 and replicated in Fig. 1) in the OPGG. Given that the success of the combination of punishment and reputation in sustaining cooperation is much higher than the success of either of the two alone, the two mechanisms act in synergy with each other. The ecosystem of strategies that thrive in this environment is made up of strategies whose main (or sole) action is cooperation, who for the most part do not employ the punishment device at all, with a smaller but material proportion punishing pro-socially. For each conditional strategy, there are significant mutual transitions between the non-punishing and the pro-socially punishing variants, indicating a minor but present role for the latter in preserving cooperation. This cooperative environment relies on the regular influx of mutant strategists to mitigate the second-order free-rider problem, moderating the cyclical dynamics it induces. New punishing agents introduced into the population are progressively transformed into the non punishing variants. Their role is confirmed by SM Fig. S24 where in the presence of both reputation and punishment, cooperation levels decrease for lower mutation rates.

All in all, these results point to the fact that in the OPGG, reputational information acts as a cheap substitute to costly punishment. Cooperators are rewarded with a positive reputation, which makes future opponents more likely to cooperate with them. This is in line with the experimental results of Ule et al. (2009), who found that in the presence of reputation, subjects prefer rewarding positive behaviour to punishing those who defect. Since the AD reputational mechanism can be considered as an indirect method of sanctioning anti-social behaviour, our result is also in line with Balafoutas et al. (2014) which experimentally finds that people prefer to punish indirectly rather than directly.

Introducing a reputational system has a differential impact on the different types of punishment. In the absence of reputation, anti-social punishment is able to successfully prevent the invasion of defectors by cooperators – as in Rand and Nowak (2011) – but is unable to do the same with the new conditional strategies under a social norm. Contrastingly, pro-social punishment positively reinforces the reputational information, increasing the level of cooperation in the population. This result is in contrast to Gurerk et al. (2006), that establishes an experimental superiority of punishment institutions over reputational mechanisms for the public goods game and poses doubts on whether this result can be extended to experiments where the OPGG is played.

It should be noted that all strategies with cooperation as the main action are selected in the ecosystem for this setup, and in remarkably similar proportions. This follows from the fact that they always act as cooperators in a strongly cooperative environment, thus becoming indistinguishable from each other (thereby avoiding selection).

Our results qualify the findings of the experimental literature on anti-social punishment (Herrmann et al., 2008; Gächter and Herrmann, 2009, 2011). Such contributions work with designs that rigidly separate individual choices about actions and punishments from reputational concerns. This is particularly true for Rand and Nowak (2011) which was performed on Amazon Mechanical Turk, where professionals play many games. In recent experiments, run in-person and with samples of the general population where individuals knew they were playing with citizens from the same villages or small towns (Pancotto et al., 2021), anti-social punishment was not significant. In these cases, even within an experimental setup, social norms implicitly matter and remove the scope for anti-social punishment.

Reproducing in-silico a simplified version of the combination of social reputation mechanism and punishment systems that characterise complex societies, our contribution shows that these two dimensions interact, shedding light on the surprising success of reputation in a world under the contemporaneous threat of exploitation and of anti-social punishment. Finally, our results contribute to identifying the conditions that allow effective collective action in the presence of the possibility to opt-out of interactions.

## Acknowledgements

Simone Righi gratefully acknowledges funding from “Decentralised Platform Economics: Optimal Incentives and Structure”, Fondi primo insediamento, Ca’Foscari University of Venice. Francesca Pancotto gratefully acknowledges funding from “Data driven methodologies to study social capital and its role for economic

growth”, University of Modena and Reggio Emilia, FAR 2019, CUP E84E19001120005. Shirsendu Podder gratefully acknowledges doctoral funding from the Engineering and Physical Sciences Research Council. The authors acknowledge the use of the UCL Myriad High Performance Computing Facility (Myriad@UCL), and associated support services, in the completion of this work.

## References

- R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981. doi: 10.1126/science.7466396.
- L. Balafoutas, N. Nikiforakis, and B. Rockenbach. Direct and indirect punishment among strangers in the field. *Proceedings of the National Academy of Sciences*, 111(45):15924–15927, 2014. ISSN 0027-8424. doi: 10.1073/pnas.1413170111.
- R. Boyd and P. J. Richerson. *Culture and the Evolutionary Process*. University of Chicago Press, June 1988. ISBN 978-0-226-06933-3.
- H. Brandt, C. Hauert, and K. Sigmund. Punishment and reputation in spatial public goods games. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1519):1099–1104, May 2003. doi: 10.1098/rspb.2003.2336.
- R. Bshary and A. S. Grutter. Image scoring and cooperation in a cleaner fish mutualism. *Nature*, 441(7096): 975–978, 2006.
- M. dos Santos. The evolution of anti-social rewarding and its countermeasures in public goods games. *Proceedings of the Royal Society B: Biological Sciences*, 282(1798):20141994, Jan. 2015. doi: 10.1098/rspb.2014.1994.
- E. Fehr and S. Gächter. Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994, Sept. 2000. doi: 10.1257/aer.90.4.980.
- E. Fehr and S. Gächter. Altruistic punishment in humans. *Nature*, 415(6868):137–140, 2002.
- S. Gächter and B. Herrmann. Reciprocity, culture and human cooperation: Previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, Dec. 2009. doi: 10.1098/rstb.2008.0275.
- S. Gächter and B. Herrmann. The limits of self-governance when cooperators get punished: Experimental evidence from urban and rural russia. *Eur. Econ. Rev*, 55:193–210, 2011. ISSN 0014-2921. doi: 10.1016/j.euroecorev.2010.04.003.
- J. García and A. Traulsen. Leaving the loners alone: Evolution of cooperation in the presence of antisocial punishment. *Journal of Theoretical Biology*, 307:168–173, Aug. 2012. ISSN 0022-5193. doi: 10.1016/j.jtbi.2012.05.011.
- A. Gell. Newcomers to the world of goods: Consumption among the muria gonds. *The social life of things: Commodities in cultural perspective*, pages 110–138, 1986.
- A. S. Griffin, S. A. West, and A. Buckling. Cooperation and competition in pathogenic bacteria. *Nature*, 430(7003):1024–1027, Aug. 2004. doi: 10.1038/nature02744.
- O. Gurerk, B. Irlenbusch, and B. Rockenbach. The competitive advantage of sanctioning institutions. *Science*, 312(5770):108–111, 2006. ISSN 0036-8075. doi: 10.1126/science.1123633.
- G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, 1968. ISSN 0036-8075. doi: 10.1126/science.162.3859.1243.

- C. Hauert. Volunteering as red queen mechanism for cooperation in public goods games. *Science*, 296(5570): 1129–1132, May 2002. ISSN 0036-8075. doi: 10.1126/science.1070582.
- C. Hauert, S. De Monte, J. Hofbauer, and K. Sigmund. Replicator dynamics for optional public good games. *Journal of Theoretical Biology*, 218(2):187–194, 2002.
- C. Hauert, N. Haiden, and K. Sigmund. The dynamics of public goods. *Discrete & Continuous Dynamical Systems - B*, 4(3):575–587, 2004. doi: 10.3934/dcdsb.2004.4.575.
- C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund. Via freedom to coercion: The emergence of costly punishment. *science*, 316(5833):1905–1907, 2007.
- E. Hauk. Multiple prisoner’s dilemma games with(out) an outside option: An experimental study. *Theory and Decision*, 54(3):207–229, 2003. doi: 10.1023/a:1027385819400.
- J. Henrich. Costly punishment across human societies. *Science*, 312(5781):1767–1770, June 2006. ISSN 0036-8075. doi: 10.1126/science.1127333.
- B. Herrmann, C. Thoni, and S. Gächter. Antisocial punishment across societies. *Science*, 319:1362–1367, 2008. doi: 10.1126/science.1153808.
- J. J. Jordan and D. G. Rand. Signaling when no one is watching: A reputation heuristics account of outrage and punishment in one-shot anonymous interactions. *Journal of Personality and Social Psychology*, 118(1):57–88, Jan. 2020. doi: 10.1037/pspi0000186.
- J. O. Ledyard. Public goods: A survey of experimental research. In *Public Goods: A Survey of Experimental Research. S. 111–194 in: J. Kagel/A. Roth (Hrsg.), Handbook of Experimental Economics*. Princeton: Princeton University Press, 1995.
- O. Leimar and P. Hammerstein. Evolution of cooperation through indirect reciprocity. *Proceedings. Biological sciences / The Royal Society*, 268:745–53, May 2001. doi: 10.1098/rspb.2000.1573.
- F. P. Maier-Rigaud, P. Martinsson, and G. Staffiero. Ostracism and the provision of a public good: Experimental evidence. *Journal of Economic Behavior & Organization*, 73(3):387–395, Mar. 2010. doi: 10.1016/j.jebo.2009.11.001.
- M. Milinski, D. Semmann, and H.-J. Krambeck. Reputation helps solve the ‘tragedy of the commons’. *Nature*, 415(6870):424–426, Jan. 2002. doi: 10.1038/415424a.
- D. Nosenzo and F. Tufano. The effect of voluntary participation on cooperation. *Journal of Economic Behavior & Organization*, 142:307–319, Oct. 2017. doi: 10.1016/j.jebo.2017.07.009.
- M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006. ISSN 0036-8075. doi: 10.1126/science.1133755.
- M. A. Nowak and K. Sigmund. The dynamics of indirect reciprocity. *Journal of Theoretical Biology*, 194(4): 561–574, Oct. 1998. doi: 10.1006/jtbi.1998.0775.
- H. Ohtsuki and Y. Iwasa. How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of Theoretical Biology*, 231(1):107–120, Nov. 2004. doi: 10.1016/j.jtbi.2004.06.005.
- H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505, 2006. doi: 10.1038/nature04605.
- J. M. Orbell and R. M. Dawes. Social welfare, cooperators\textquotesingle advantage, and the option of not playing the game. *American Sociological Review*, 58(6):787, Dec. 1993. doi: 10.2307/2095951.
- K. Panchanathan and R. Boyd. A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology*, 224(1):115–126, 2003. ISSN 0022-5193. doi: 10.1016/S0022-5193(03)00154-1.

- K. Panchanathan and R. Boyd. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature*, 432(7016):499–502, 2004.
- F. Pancotto, K. Takács, and S. Righi. Voluntary play increases cooperation in the presence of punishment: A lab in the field experiment. *Mimeo*, 2021.
- M. Perc, J. J. Jordan, D. G. Rand, Z. Wang, S. Boccaletti, and A. Szolnoki. Statistical physics of human cooperation. *Physics Reports*, 687:1–51, May 2017. ISSN 0370-1573. doi: 10.1016/j.physrep.2017.05.004.
- P. B. Rainey and K. Rainey. Evolution of cooperation and conflict in experimental bacterial populations. *Nature*, 425(6953):72–74, Sept. 2003. doi: 10.1038/nature01906.
- D. G. Rand and M. A. Nowak. The evolution of antisocial punishment in optional public goods games. *Nature Communications*, 2(1):434, Aug. 2011. ISSN 2041-1723. doi: 10.1038/ncomms1442.
- D. G. Rand, J. J. Armao IV, M. Nakamaru, and H. Ohtsuki. Anti-social punishment can prevent the co-evolution of punishment and cooperation. *Journal of Theoretical Biology*, 265(4):624–632, Aug. 2010. ISSN 0022-5193. doi: 10.1016/j.jtbi.2010.06.010.
- D. Scheel and C. Packer. Group hunting behaviour of lions: A search for cooperation. *Animal Behaviour*, 41(4):697–709, Apr. 1991. doi: 10.1016/s0003-3472(05)80907-8.
- D. Semmann, H.-J. Krambeck, and M. Milinski. Volunteering leads to rock–paper–scissors dynamics in a public goods game. *Nature*, 425(6956):390–393, 2003.
- K. Sigmund, C. Hauert, and M. A. Nowak. Reward and punishment. *Proceedings of the National Academy of Sciences*, 98(19):10757–10762, 2001.
- J. M. Smith. Group selection and kin selection. *Nature*, 201(4924):1145–1147, Mar. 1964. doi: 10.1038/2011145a0.
- B. B. Smuts, D. L. Cheney, R. M. Seyfarth, and R. W. Wrangham. *Primate Societies*. University of Chicago Press, 2008.
- A. Szolnoki and M. Perc. Antisocial pool rewarding does not deter public cooperation. *Proceedings of the Royal Society B: Biological Sciences*, 282(1816):20151975, Oct. 2015. doi: 10.1098/rspb.2015.1975.
- A. Szolnoki and M. Perc. Second-order free-riding on antisocial punishment restores the effectiveness of prosocial punishment. *Physical Review X*, 7(4):041027, Oct. 2017. doi: 10.1103/physrevx.7.041027.
- M. Takezawa and M. E. Price. Revisiting “The Evolution of Reciprocity in Sizable Groups”: Continuous reciprocity in the repeated n-person prisoner’s dilemma. *Journal of Theoretical Biology*, 264(2):188–196, May 2010. ISSN 0022-5193. doi: 10.1016/j.jtbi.2010.01.028.
- A. Ule, A. Schram, A. Riedl, and T. Cason. Indirect punishment and generosity toward strangers. *Science*, 326:1701–1704, 2009. doi: 10.1126/science.1178883.