Exp Econ (2015) 18:15–37
DOI 10.1007/s10683-014-9402-3

ORIGINAL PAPER

Incomplete punishment networks in public goods games: experimental evidence

Andreas Leibbrandt · Abhijit Ramalingam · Lauri Sääksvuori · James M. Walker

Received: 4 April 2013 / Revised: 3 April 2014 / Accepted: 7 April 2014 /
Published online: 26 April 2014
© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract Abundant evidence suggests that high levels of contributions to public goods can be sustained through self-governed monitoring and sanctioning. This experimental study investigates the effectiveness of decentralized sanctioning institutions in alternative punishment networks. Our results show that the structure of punishment network significantly affects allocations to the public good. In addition, we observe that network configurations are more important than punishment capacities for the levels of public good provision, imposed sanctions and economic efficiency. Lastly, we show that targeted revenge is a major driver of anti-social punishment.

Keywords Public goods experiment · Punishment · Cooperation · Networks

JEL Classification C72 · C91 · C92 · D72 · D74

Electronic supplementary material The online version of this article (doi:10.1007/s10683-014-9402-3) contains supplementary material, which is available to authorized users.

A. Leibbrandt (✉)
Department of Economics, Monash University, Clayton, VIC 3800, Australia
e-mail: leibbrandt@gmail.com; andreas.leibbrandt@monash.edu

A. Ramalingam
School of Economics and Centre for Behavioural and Experimental Social Science, University of East Anglia, Norwich, UK

L. Sääksvuori
Department of Economics, University of Hamburg, Hamburg, Germany

J. M. Walker
Department of Economics and Workshop in Political Theory and Policy Analysis, Indiana University, Bloomington, USA

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

1 Introduction

There is widespread evidence that the availability of costly peer sanctioning can have a large positive impact on cooperation in social dilemma settings (e.g., Ostrom 1990; Ostrom et al. 1992; Fehr and Gächter 2000; Walker and Halloran 2004; Sefton et al. 2007). These findings suggest that self-governed monitoring and sanctioning may play an important role in human cooperation and well-functioning of modern societies. However, the prevailing evidence is mainly based on the comparison of two extreme cases; all individuals can punish and be punished by other individuals in a group versus a situation where no one can punish. These criteria are typically not met in the field where various factors such as physical distance, endowments and status, and the social network of actors regularly limit punishment opportunities.

Punishment networks, which define who can punish whom, may play a nontrivial role for inducing more efficient provision of public goods or appropriation from common-pool resources. In particular, it seems plausible that denser punishment networks, where a larger fraction of actors can punish each other, deter actors more effectively from non-cooperative behaviors. This increased deterrence in denser networks may be associated with the threat of being punished by more agents and/or the possibility of larger combined punishment capacity. However, it seems equally plausible that denser punishment networks may deter actors less effectively from non-cooperative behaviors if actors believe that the threat of being punished diminishes as the number of potential targets increases and effective coordination of punishment becomes more difficult. In addition, the increasing number of potential targets and limited individual capacities to sanction may reduce the severity of assigned sanctions. Taken together, there is very little direct evidence on how the network structure and punishment capacity impact public good provision, imposed sanctions and economic efficiency.

In this study, we provide new empirical evidence on the role of punishment networks for facilitating cooperation. We employ a public goods experiment in which we manipulate the structure of punishment networks and punishment capacities. Contribution and punishment decisions are examined across twenty rounds of repeated play in groups of four players who have fixed identifiers. Four networks are examined: a complete punishment network, a 'pairwise' punishment network, an 'untouchable' punishment network and a no-punishment network. In the pairwise network, the group of four is divided into two pairs and punishment can only take place within pairs, although contributions affect the entire group. In the untouchable network, there are three agents that can punish and be punished by each other and one agent who cannot punish or be punished.

By reducing the number of players who can punish a player, the two incomplete networks (pairwise and untouchable) reduce the total capacity of players to impose and receive punishment. For this reason, an additional treatment is conducted in each of the incomplete networks such that punishment capacities were as high as in the complete network. Individual punishment capacities are manipulated in these two networks in order to investigate if observed behavior is driven by the structure of the punishment network or punishment capacity.

These punishment networks were selected for the following reasons. First, arguably, the pairwise networks constitute the most transparent cases to examine

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

issues of targeting sanctions, reputation formation, and limited scope of sanctions. The untouchable networks were selected based on observations from the field where it is common that some agents are temporarily or permanently isolated from others, but cannot be excluded from the benefits of public goods or common-pool resources. Complete and no punishment network conditions are created as benchmarks and to better link our findings to the existing experimental literature. The investigation of punishment behavior in incomplete networks connects our study to numerous examples of common-pool resource management and public good provision settings where the geographical structure and state borders may limit stakeholders' opportunities to sanction each other. At the same time, many of the international agreements designed to protect natural resources and curb environmental deterioration implement governance structures that often allow for accurate monitoring of contributions but limited opportunities to punish detached actors.

A primary finding of this study is that the greater the number of people who can punish and be punished, the greater the contributions to the public good and the greater the amount of punishment used in the group. Further, high contributions are sustained only in the complete and untouchable networks. In addition, the capacity for one individual to punish another plays a less important role on aggregate contribution levels than the network configuration. In particular, higher punishment capacities are unable to stem the observed decline in contributions in the pairwise network, and also play an insignificant role in the untouchable network. Finally, consistent with previous findings, low and high contributors are punished (Hermann et al. 2008), a finding that is consistent with targeted revenge.

This study contributes to the literature testing the effectiveness of various institutional arrangements to overcome the regularly observed sub optimality of voluntary contributions. Among the large body of proposed institutional solutions to the problem of free-riding, opportunities to communicate (Isaac and Walker 1988; Ostrom et al. 1992; Bochet et al. 2006), costly peer punishments (Ostrom et al. 1992, Fehr and Gächter 2000), verbal sanctioning (Masclet et al. 2003), ostracism (Cinyabuguma et al. 2005), combined punishment and reward schemes (Andreoni et al. 2003; Gürerk et al. 2006; Sefton et al. 2007; Leibbrandt and López-Pérez 2014), reputation networks (Milinski and Rockenbach 2006) and leadership structures (Güth et al. 2007) all potentially serve as proximate mechanisms to enhance voluntary cooperation.[1]

[1] Since establishing the seminal finding that costly peer sanctioning can have a large positive impact on cooperation in social dilemmas, numerous additional studies have identified important limitations that may reduce the effectiveness of punishments and hinder the achievement of Pareto improvements through decentralized sanctioning institutions. Among the discussed limitations some particularly notable ones are the threat of counter punishments that may make people less willing to punish free-riders (Denant-Boemont et al. 2007; Nikiforakis 2008) or lead to destructive feuds (Nikiforakis and Engelmann 2011), and anti-socially targeted punishments (Hermann et al. 2008) that may prevent the co-existence of punishments and cooperative strategies (Rand et al. 2010). Likewise, it has been shown that the cost effectiveness of punishments plays an important role when assessing the impact of punishment strategies on cooperation and social efficiency (Egas and Riedl 2008; Nikiforakis and Normann 2008). At the same time, however, it has been shown that various mechanisms allowing participants to effectively coordinate their punishment behavior may enhance the effectiveness of decentralized institutional arrangements (Ertan et al. 2009; Boyd et al. 2010). See Chaudhuri (2011) for a recent article reviewing the experimental literature on sustaining cooperation in social dilemmas.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

In addition, this study connects to an emerging literature examining the role of social and geographic network structures on public good provision when punishment opportunities are absent. Theoretical investigations (Bramoullé and Kranton 2007) and experimental evidence (Yamagishi and Cook 1993; Fatas et al. 2010) point to the fact that contribution levels may differ significantly across networks. Differences in contributions across such networks are explained by conditionally cooperative responses to the restricted spread of information about individual contributions (Fatas et al. 2010).[2]

More closely related to our study are experiments in which punishment opportunities in public goods settings are manipulated (Carpenter 2007a; Kosfeld et al. 2009; O'Gorman et al. 2009; Reuben and Riedl 2009; Nikiforakis et al. 2010; Carpenter et al. 2012; Cox et al. 2013). Reuben and Riedl (2009) study the effectiveness of punishment in privileged groups where some group members generate positive returns from public good contributions. Their findings indicate that punishment is less effective in privileged groups as compared to normal groups. Kosfeld et al. (2009) investigate institution formation in social dilemmas where a subset of players can form a sanctioning institution, while their contributions benefit the outsiders who do not enter the institution. Nikiforakis et al. (2010) vary the effectiveness of punishments across individuals. Their results suggest that institutions with asymmetric sanctioning power can be equally successful in fostering cooperation and efficiency than their symmetric counterparts. Carpenter et al. (2012) manipulate monitoring opportunities and show how properties from graph theory can organize the data patterns that arise in their public goods experiments.

This study differs in several aspects from the previous literature. First, previously unexplored network structures are examined in settings where decision makers receive complete information about individual contributions, sanctions imposed, and sanctions received for all group members. This contrasts with other studies that investigate the joint effect of information dissemination and punishment opportunities in networks where group members do not receive information on individual behavior outside their network (Carpenter 2007a; Carpenter et al. 2012). Second, we use a partner-matching protocol with fixed identifiers. The advantage of fixed identifiers is that this information condition captures the essence of many real networks where individuals have stable positions within a fixed group, not simply a network architecture describing how a random group of individuals occasionally link.[3] Finally, individual punishment endowments and total punishment capacities are controlled for

[2] The influence of exogenous and endogenous networks has been explored in other settings such as coordination games (Keser et al. 1998), stag-hunt games (Charness and Jackson 2007), games of strategic investments (Rosenkranz and Weitzel 2012) and bargaining games (Charness et al. 2007). For a survey, see Kosfeld (2004).

[3] A possible disadvantage is that reputation building is easier in the partner-matching protocol. However, since our primary interest lies in comparing punishment networks and not in disentangling the motivation of individual actors, we believe that the partner-matching protocol is more suited for our purposes. Note, by design, the pairwise network requires that subjects know the decisions with whom they are paired. For experimental control, this implies fixed identifiers be used in the other punishment networks. Clearly our experimental design represents an extreme case in regard to information subjects have on decisions of other group members. Such a design, however, provides a clear benchmark from which the effects of reductions in information could be compared in future studies.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

across groups. Thus, in contrast to many studies, we are able to identify the role of the punishment network and can rule out potential endowment effects.

2 The decision setting

This study includes data from experimental sessions conducted at Indiana University-Bloomington (US) and the University of East Anglia (UK). In each session, 12–20 subjects were recruited from subject databases that included undergraduates from a wide range of disciplines. Via the computer, subjects were privately and anonymously assigned to four-person groups and remained in these groups throughout the 20 rounds in a session. No subject could identify the others in the room that were assigned to their group. Since no information passed across groups, each session involved 3–5 independent groups. At the beginning of each session, subjects privately read a set of instructions, which were then summarized publicly by a member of the research team.[4] Subjects then took a post instruction quiz and were not allowed to continue until all answers were correct. Subjects made all decisions privately.

Stage 1 of each decision round was a linear VCM game. At the beginning of stage 1, each subject was endowed with ten tokens to be allocated between a private account and a group account. For each token placed in his or her private account a subject received 1 token in payment. For each token placed in the group account, each group member received 0.4 tokens in payment. After all subjects had made their decisions in stage 1, they were informed of the aggregate allocations to the group account, and the allocation of each member of their group identified by an anonymous ID letter (A, B, C, or D), which remained the same during all decision rounds.

In stage 2 of each decision round each subject received an additional endowment of six tokens. Subjects were informed that they would make a decision of whether to decrease the earnings of other members in their group by assigning deduction tokens to them.[5] The instructions used neutral language. Each deduction token assigned by a group member to another group member cost the initiator 1 token and decreased the earnings of the recipient by three tokens. Any tokens not used to decrease the earnings of other group members were kept in the subject's private account.

Following stage 2 decisions, each subject received information about the contribution and sanction decisions of every other subject in his/her group.[6] More specifically, each subject reviewed a table which displayed the group account allocation of each subject in their group and the number of deduction tokens each subject assigned to each other subject in the group identified by ID letters. This table also displayed current round and cumulative earnings for each subject. At any point in the experiment subjects could review this same information from the prior round, giving them a complete history of individual decisions from the prior round before

[4] See Sect. C in the Supplementary Material for the instructions. The programs were written using Z-tree (Fischbacher 2007).

[5] This procedure, which parallels that used in Sefton et al (2007), holds constant the resources available for sanctioning across decision rounds and decision making groups.

[6] In the *no-punishment* treatment, subjects received the same information regarding individual group account allocations.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

making their current round decisions. Thus, unlike in many earlier decision settings that have investigated the use of sanctioning mechanisms, it was feasible for subject-specific reputations to develop across rounds.[7] The network treatment conditions are the primary rationale for this particular parameterization.

No sanctions were allowed in the benchmark treatment, the *no-punishment* network. In stage 2, subjects were simply given an additional six tokens, which were placed in their private accounts. Otherwise, this treatment was conducted in the same manner as the treatments that allowed for sanctioning opportunities. As noted in the introduction, there were three treatment conditions that allowed for sanctions: a c*omplete* network*, a pairwise* network, and an *untouchable* network. Experimental conditions varied only in terms of opportunities for sanctioning defined by the network linkages. In the *complete* network condition, subjects had the opportunity to reduce the earnings of all other group members. In the *pairwise* network condition, subjects A and B had the opportunity to reduce the earnings of each other, but not C and D. Likewise, subjects C and D had the opportunity to reduce the earnings of each other, but not A and B. In the *untouchable* network condition, subjects A, B, and C had the opportunity to reduce the earnings of each other, but not subject D. Further, subject D did not have the opportunity to reduce the earnings of any group member. For control purposes, subject D automatically had six tokens allocated to their private account.

Figure 1 illustrates our network treatments. In all network treatments information flow was held the same. Only the punishment opportunities depended on the network. In the figures, an incoming arrow denotes that a player can be punished by the player from whom the arrow originates. An outgoing arrow denotes that a player can punish the receiving group member.

For control purposes, in the initial set of experiments subjects could assign a maximum of two deduction tokens to another group member, reducing that subjects earnings by a maximum of six tokens, regardless of the network structure. Subjects in the *pairwise* network automatically had 4 tokens allocated to their private accounts in stage 2 while subjects A, B and C in the *untouchable* network automatically had 2 tokens allocated to their private accounts in stage 2. Players could use the remaining tokens to sanction players in their network. Thus, in the initial set of experiments, the maximum sanction that a subject could *impose* on another subject was held constant across decision rounds, while the maximum number of punishment tokens a subject could *receive* varied across networks.

An additional set of experiments was conducted in the *pairwise* and *untouchable* networks, where the maximum number of deductions tokens that a subject could *receive* was equal to that of the *complete* network. In the *pairwise-6* treatment each subject could impose up to six punishment tokens on the subject with whom they were paired. In the *untouchable-6* treatment, the three subjects in the punishment network could impose up to three punishment tokens on the other two subjects in their network. Thus, in these treatment conditions, subjects in the networks could have their earnings reduced from punishments by a maximum of 18 tokens, the same as in the *complete* network condition.

[7] Nicklisch and Wolff (2011) and Nikiforakis and Engelmann (2011) study retaliative punishment and allow the development of subject-specific reputations across rounds.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

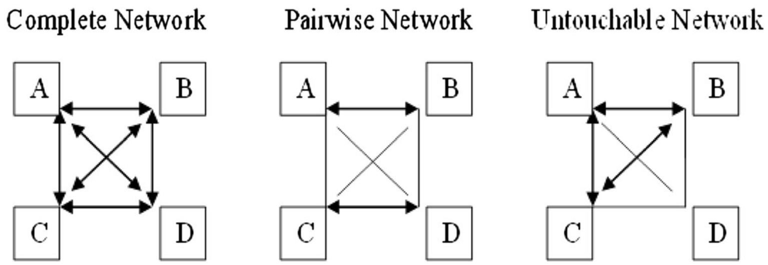

Fig. 1 Punishment networks. In all treatments information flow was held the same, indicated by the *lines* between players. Every player received information about the contribution and punishment decisions of every other player in her group. Only the punishment opportunities depended on the network. An *incoming arrow* denotes that a player can be punished by the player from whom the arrow originates. An *outgoing arrow* denotes that a player can punish the receiving group member

Table 1 presents summary information related to subject groups in each of the conditions. In aggregate, data were collected from 84 four-person groups. In the experiments conducted in the US, the conversion rate of tokens to dollars was 20 to 1. In the U.K., the conversion of tokens to pounds was 30 to 1.[8]

In all treatment conditions, subjects played a finitely repeated game with a known final round. Under the assumption that it is common knowledge that subjects maximize own-earnings, the theoretical prediction is straightforward. The subgame perfect Nash equilibrium for each treatment condition calls for zero allocations to the group account and no-sanctions.[9] As noted earlier, however, experimental studies of the linear VCM game typically find that the level of cooperation observed is not consistent with equilibrium predictions of zero provision of the group good. Moreover, other studies have shown that subjects often pay to sanction other participants when the opportunity is available. However, at the same time subjects react to changes in the price and effectiveness of punishment (Carpenter 2007b), suggesting that players strategically assess the cost and benefits of various sanctioning strategies. At the core of our investigation is the question how the network structure and disposable punishment capacities affect these considerations.

3 Results

As noted in Sect. 2, experimental sessions for the *no-punishment*, *complete*, *pairwise,* and untouchable network conditions were conducted in two locations, the University of East Anglia, UK and Indiana University Bloomington, USA. Recent work suggests that there may be systematic differences in the behavior of subjects in

[8] These differential exchange rates were chosen to create experimental earnings that yielded approximately the same real valued payoffs across locations. Subject's experimental earnings averaged $22 in the US, including a $5 show-up payment, and £15 in the UK, including a £3 show-up payment. Sessions lasted from one to one and one half hours.

[9] In the sanction treatments there are other Nash equilibria, including some that support efficient allocations. However, equilibrium strategies that support efficient allocations rely on non-credible threats to sanction free riders.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 1 Design information for network conditions

Network condition	Number of groups US	Number of groups UK	Total number of independent groups
No-punishment	7	8	15
Complete	9	8	17
Pairwise	6	8	14
Untouchable	8	7	15
Pairwise-6	12	0	12
Untouchable-6	11	0	11

different countries (Hermann et. al. 2008). Controlling for treatment condition, we tested for differences in behavior in the two locations. A detailed analysis is available in Sect. A of the Supplementary Material. In summary, the various tests confirm that there are no statistically significant differences in group allocations and earnings between the two locations. In addition, within each of the three treatments with sanctioning opportunities, the average level of sanctions used by groups is not different between locations. The analysis presented below therefore pools the data from both experimental sites. Results are first presented at the group level, followed by analyses at the individual level. We begin with a graphical presentation and summary statistics which focus on pooled data from the initial set of network conditions and the *pairwise-6* and *untouchable-6* networks.

3.1 Group level results

The discussion of results from the initial treatment conditions focuses on three key outcome variables: (1) tokens allocated to the group account by each four-person group, (2) total tokens used for sanctioning by each four-person group, (3) tokens earned by each group. Figure 2a displays the trajectory, across decision rounds, of mean group allocations, Fig. 2b of sanctions, and Fig. 2c of earnings for the *complete* networks (mean across 17 groups), the *pairwise* networks (14 groups) and the *untouchable* networks (15 groups). Mean group allocations and earnings for the *no-punishment* networks (15 groups) are also presented. To complement the results displayed in Fig. 2a–c, Table 2 presents the means and standard deviations of per-round group allocations, group earnings, and sanctions per group, pooled over decision rounds.

In all treatments, average group account allocations start at around 50 % of the group endowment of 40 tokens. In the *no-punishment* networks, allocations decline over time to levels close to the Nash equilibrium allocation of zero. In the *complete* networks, allocation levels increase slightly and are maintained at around 25 tokens throughout. In the *untouchable* networks group allocations remain steady at around 20 tokens across rounds 1–18. However, allocations are always lower than those in the *complete* networks. Allocation levels in the *pairwise* networks are very similar to those in the *no-punishment* networks, though they are slightly higher after round 5.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

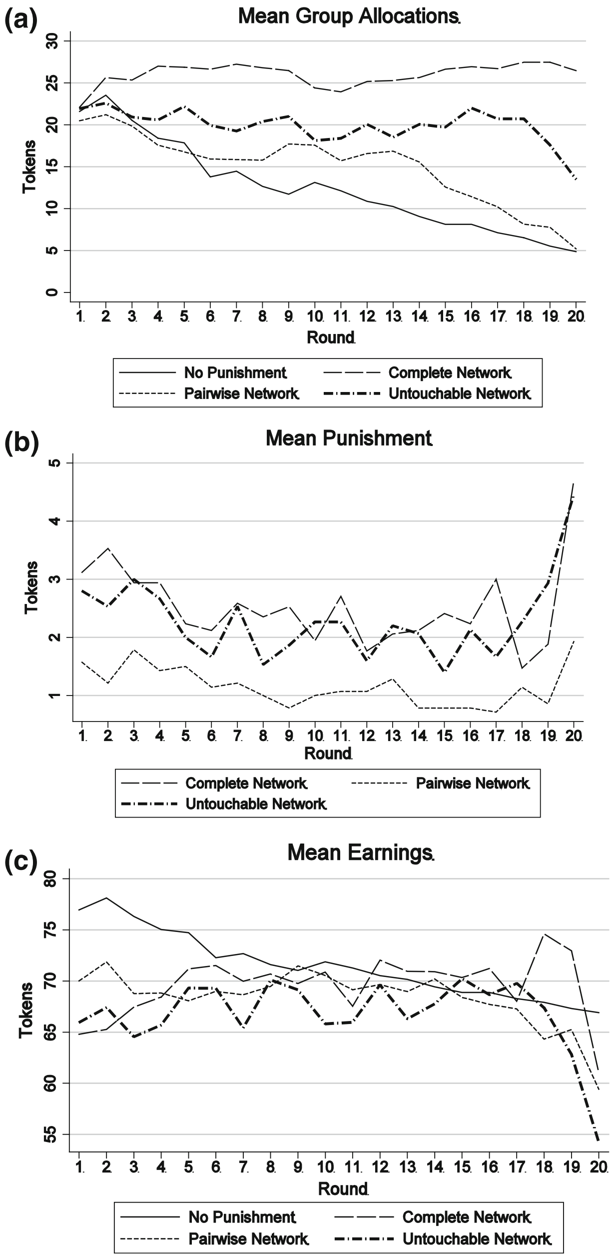

Fig. 2 **a–c** Allocations, sanctions and earnings: initial punishment networks

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 2 Summary statistics: group level data

	Mean allocation (standard deviation)	Mean sanctions (standard deviation)	Mean earnings (standard deviation)
No-punishment	12.52	–	71.512
(15 groups)	(9.134)		(5.481)
Complete	26.017	2.532	69.481
(17 groups)	(11.878)	(1.922)	(13.714)
Pairwise	14.942	1.153	68.351
(14 groups)	(11.406)	(1.0098)	(8.308)
Untouchable	19.92	2.293	66.779
(15 groups)	(9.365)	(1.622)	(9.979)
Pairwise-6	16.867	1.696	67.337
(12 groups)	(7.456)	(1.878)	(7.612)
Untouchable-6	23.691	1.732	71.287
(11 groups)	(11.740)	(1.885)	(9.490)

To complement the graphical presentations, we present evidence below from non-parametric Mann–Whitney tests of differences in behavior across treatments.[10] In our experiment, groups make decisions independently of other groups as they only receive information about themselves. However, a group's decisions are not independent over the 20 rounds in the experiment. Thus the average (allocations, sanctions or earnings) of a group over all 20 rounds serves as an independent observation for these tests. The tests confirm the pattern of results drawn from Fig. 2a–c. Relative to the *no-punishment* networks, group allocations are significantly higher in the *complete* networks ($p = 0.0006$) and the *untouchable* networks ($p = 0.019$), but not in the *pairwise* networks ($p = 0.827$). Further, group allocations in the *complete* networks are significantly higher than in the *pairwise* networks ($p = 0.009$) but not in the *untouchable* networks ($p = 0.117$). The difference between allocations in the *pairwise* and *untouchable* networks is also not statistically significant ($p = 0.097$).

Result 1: The structure of the punishment network significantly affects public good contributions. Incomplete pairwise punishment networks are less effective in increasing public goods contributions.

We next turn to punishment behavior. Recall, in the initial punishment network conditions, subjects were constrained to use no more than two tokens in sanctioning another individual, implying that the number of sanctions that could be imposed varied across network conditions. Yet, as can be seen from Fig. 2b and Table 2, average group sanctions imposed in the *complete* and *untouchable* networks are

[10] In addition to the non- parametric tests, t-tests were also conducted. Unless otherwise noted in the text, the results are robust to both types of tests. In order to further examine differences between punishment networks, OLS, Tobit and panel random effects models were also estimated for group level data. The results are qualitatively similar for all the three models and do not significantly differ from the results obtained by group level parametric and non-parametric tests. This analysis is available in Tables B1 through B5 in the Supplementary Material. We only present the panel estimations for brevity.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

similar in most rounds ($p = 0.850$) and remain steady at around 2.5 tokens per round. In the *pairwise* networks average group sanctions are significantly lower than in the *complete* network ($p = 0.043$) and the *untouchable* network ($p = 0.022$). Thus, the two network structures with greater sanctioning opportunities lead to increased levels of sanctioning in relation to the *pairwise* networks.

Result 2: The structure of the punishment network significantly affects sanctioning levels. Sanctioning levels are lower in incomplete pairwise punishment networks.

While there are significant differences in group allocations across the treatments, after accounting for the costs of sanctioning, there is some evidence that group earnings in the sanctioning networks are marginally lower (or no higher) than in the no-punishment networks. More specifically, earnings in the *no-punishment* networks are higher than those in the other three networks in the first few rounds and in the last round. However, between rounds 5 and 19, there is no systematic difference in earnings across network conditions. The statistical tests confirm that there is no significant difference in earnings between the *no-punishment* networks and the *complete* and *untouchable* networks ($p = 0.533$ and $p = 0.290$ respectively). The non-parametric test suggests a significant difference between the *no-punishment* networks and the *pairwise* networks ($p = 0.049$). This difference, however, is not robust to a standard t test ($p = 0.243$) or to a group-level panel regression where the baseline is the no-punishment treatment ($p = 0.219$ for the *pairwise* treatment dummy).

Result 3: After accounting for the costs of sanctioning, overall earnings across punishment networks and the no-punishment network are similar.

To examine whether results 1–3 are driven by the structure of the punishment networks or differences in absolute punishment capacity, we compare behavior from the *pairwise* networks to the *pairwise-6* networks and the *untouchable* networks to the *untouchable-6* networks. Figure 3a–c display the trajectory of mean group allocations (3a), sanctions (3b) and earnings (3c) for the *pairwise* networks and the *pairwise-6* networks. Table 2 presents mean and standard deviations for both networks. In summary, no statistically significant difference is observed in group allocations, group sanctions, and earnings (allocations, $p = 0.503$; sanction, $p = 0.837$; earnings, $p = 0.471$). In addition, despite the identical group punishment capacity between the *pairwise-6* and *complete* networks, contributions in the *pairwise-6* networks are significantly lower than in the *complete* networks ($p = 0.069$).

Figure 4a–c displays the trajectory of mean group allocations (4a), sanctions (4b) and earnings (4c) for the *untouchable* networks and the *untouchable-6* networks. Table 2 presents mean and standard deviations for both networks. Group allocations start out higher in the *untouchable-6* networks but by round 15, there is no discernible difference in allocations. Interestingly, sanctioning is not higher but slightly lower in the *untouchable-6* networks in all but five rounds. The combination of higher group allocations and lower sanctions across most decision rounds implies that earnings are somewhat higher in the *untouchable-6* networks. However, there are no statistically significant differences between the two untouchable conditions (allocations, $p = 0.452$; sanctions, $p = 0.253$; earnings, $p = 0.312$).

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

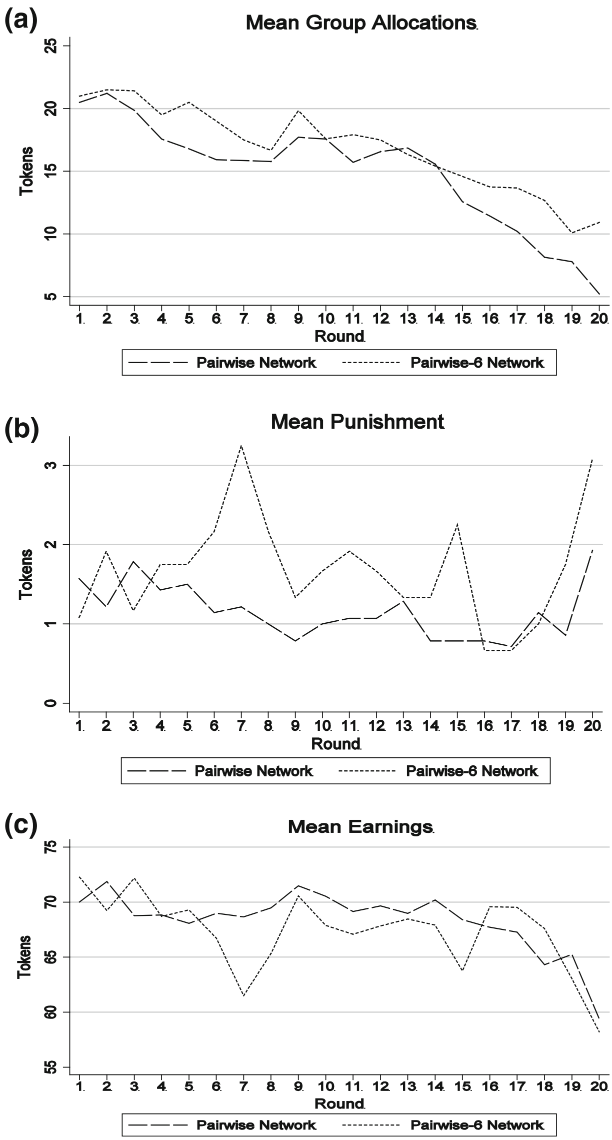

Fig. 3 **a–c** Allocations, sanctions and earnings: pairwise and pairwise-6 networks

Result 4: At the group level, the structure of the punishment network is more important than the absolute punishment capacity in determining group account allocations, sanctions, and earnings.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

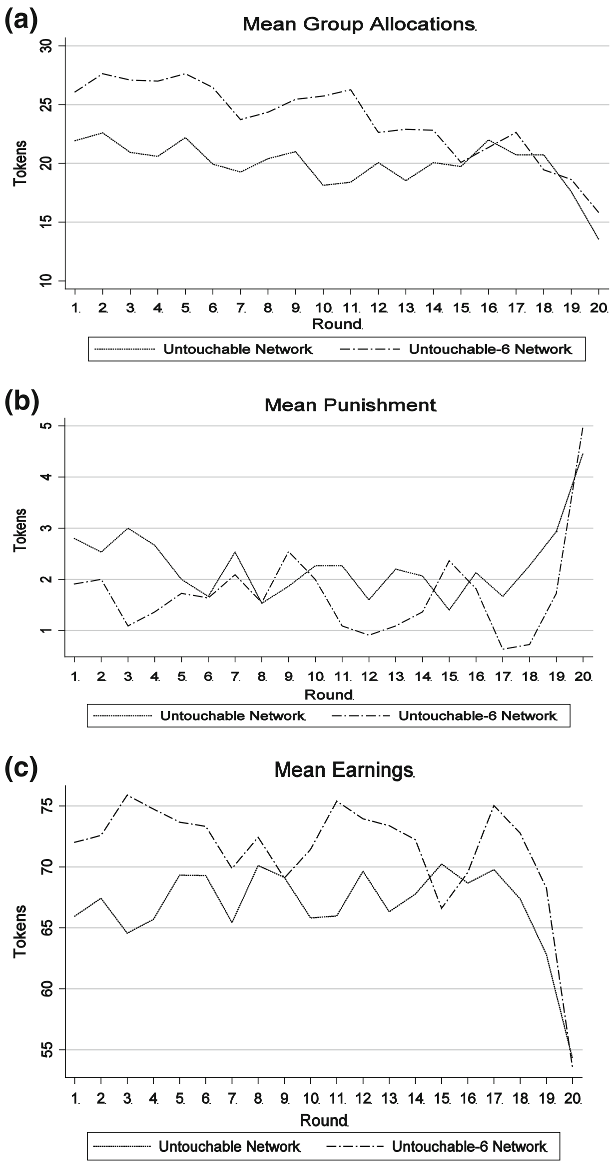

Fig. 4 **a–c** Allocations, sanctions and earnings: untouchable and untouchable-6 networks

3.2 Individual level results in incomplete networks

To complement the group level analysis, we turn to an analysis of decisions of individual group members in the incomplete networks. The nature of individual behavior in repeated public goods settings is often characterized as conditional

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

cooperation. In incomplete networks, the network structure and players' positions in the network are likely to influence how they adjust their behavior to that of the other group members. To better understand the effect of changing network structures on the nature of conditional cooperation, the analyses in the following two sections investigate how the network position in the pairwise and untouchable networks impacts group allocations.[11]

3.2.1 Individual decisions in the pairwise networks

It is an open question whether and to what extent individuals' allocations are influenced by the decisions of subjects that are linked to the punishment network and by the decisions of the other subjects outside the punishment network. More precisely, in the pairwise networks, subject A might be influenced by the allocation of subject B and vice versa (similarly for subjects C and D). However, in our experiment, each individual has information on the decisions of *all* others in his/her group. Thus, it is also possible that, within a group, subject A might be influenced by the decisions of subjects C and D even though he/she cannot be sanctioned by them.

Table 3 presents the results from a regression of individual allocations in a model incorporating the following explanatory variables: lagged allocation of subject i, lagged deviation from the subject with whom subject i is paired in the network, lagged deviation from the mean group allocation of the other pair in the group, lagged sanctions received by i, and round dummy variables. The table reports robust standard errors clustered on independent groups. We estimated an OLS regression, a random effects panel regression and a Tobit regression to account for censoring of the observations. The results are qualitatively the same in all three models. For the sake of brevity, we only report the results from the panel regression. The results indicate that *both* the lagged deviation in allocation from that of one's partner and the lagged deviation from the average allocation of the other pair significantly influence one's allocation decisions ($p < 0.001$ for both coefficients) and that the magnitudes are similar (coefficients for *pairwise* network are −0.273 and −0.256, respectively, and coefficients for the *pairwise-6* network are −0.138 and −0.178, respectively).

Table 3 highlights an additional insight in regard to the effect of received sanctions on allocations to the group account. While the variable *lagged sanction received* is positive, but insignificant, when pooling both pairwise networks, this variable is negative and significant in the *pairwise* networks ($p = 0.014$) and positive and significant in the *pairwise-6* networks ($p = 0.002$). This suggests that in the pairwise network, sanctions have a negative impact on contributions when the punishment capacity is small (for every unit of sanctioning received contributions

[11] An analysis of individual decisions, pooling across all treatments, was also conducted. The findings were consistent with previous studies. Previous round's allocation has a significant positive impact on the current allocation; positive deviations from the average allocation of others in the previous round has a significant negative impact on current allocations; and negative deviations from the average allocation others has a positive impact. This analysis is not included for purposes of brevity, but is available in Table B6 in the Supplementary Material.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 3 Individual allocations in the pairwise and the pairwise-6 networks

	Dependent variable: individual allocations		
	Pairwise network	Pairwise-6 network	Combined pairwise networks
Lagged allocation of *i*	0.944***	0.916***	0.939***
	(0.019)	(0.021)	(0.015)
	[0.000]	[0.000]	[0.000]
Lagged deviation from paired subject in network	−0.273***	−0.138***	−0.198***
	(0.024)	(0.028)	(0.026)
	[0.000]	[0.000]	[0.000]
Lagged deviation from mean allocation of other pair in group	−0.256***	−0.178***	−0.222***
	(0.056)	(0.049)	(0.040)
	[0.000]	[0.000]	[0.000]
Lagged sanctions received	−0.418**	0.295***	0.089
	(0.170)	(0.093)	(0.094)
	[0.014]	[0.002]	[0.343]
Constant	0.629***	0.484	0.436**
	(0.216)	(0.426)	(0.220)
	[0.004]	[0.256]	[0.048]
Observations	1,064	912	1,976
Clusters/groups	14	12	26

Numbers in parentheses are robust standard errors clustered on independent groups. Figures in brackets are *p* values for the two-sided tests of significance

*** Sig. at 1 %, ** sig. at 5 %, * sig. at 10 %

are decreased by 0.418 token); but a positive impact on contributions when the punishment capacity is large (for every unit of sanctioning received contributions are increased by 0.295 tokens).

3.2.2 Individual decisions in the untouchable networks

In the *untouchable* and the *untouchable-6* networks, subjects assigned the positions of A, B or C are allowed to sanction each other. Subjects assigned the position D (the untouchable) face no threat of receiving sanctions. In the analysis below, we investigate the determinants of the allocation decisions of subjects in the A, B, and C positions separately from those in the D position.

Figure 5a and b present the trajectory of mean allocations and earnings by subjects assigned to the A, B, C and D positions across decision rounds. As shown, there is a pronounced decrease in the group account allocations for the subjects in the D position, relative to those in the A, B, and C positions. The mean allocation per round by subjects in the A, B and C positions is 5.89 tokens while the mean per round allocation of subjects in the D position is 3.85 tokens ($n = 26$ groups, $p = 0.017$). Since subjects in the untouchable position also do not spend resources

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

on sanctioning, they earn significantly more than the other group members as seen from the second panel of Fig. 5. The mean per round earnings of subjects in the A, B and C positions is 15.98 tokens while the mean per round earnings of subjects in the D position is 20.75 tokens ($n = 26$ groups, $p < 0.000$). Interestingly, the presence of an untouchable does not appear to have a significant detrimental effect on the willingness to contribute by the other subjects in the same group. There is no significant difference between the mean group account allocation by subjects in the A, B and C positions (5.89 tokens) in comparison to the mean allocation of subjects in the *complete* networks of 6.50 tokens ($n_{complete} = 17$, $n_{untouchables} = 26$, $p = 0.358$).

To examine more closely the factors that influence individual allocations of subjects in the A, B and C positions, Table 4 reports the results from a random effects panel data regression of individual allocations on: the one-period lagged allocation of individual *i*, the one-period lagged deviation of *i*'s allocation from the allocation of D, the one-period lagged deviation of *i*'s allocation from the average allocation of the other members of his punishment network, a one-period lagged variable of sanctions received, and round dummies. The table reports robust standard errors clustered on independent groups.

In summary, allocations of subjects attached to the punishment networks are significantly influenced by their lagged allocations ($p < 0.001$) and the deviation of their lagged allocations from the average allocations of others in the punishment network ($p < 0.001$). In addition, their allocations are also negatively influenced by the deviation of their lagged allocations from the allocation of the untouchable ($p < 0.001$) suggesting that the untouchable can trigger higher contributions of the subjects in the punishment network.[12] Similar to the pairwise networks, punishment capacity appears to determine whether receiving sanctions has a negative (if capacity is small) or positive (if capacity is large) impact on contributions.[13]

Finally, Table 5 presents random effects estimates for the determinants of the allocations of subjects assigned to the untouchable position, D, on the one-period lagged allocation of individual *i,* the one-period lagged deviation of *i*'s allocation from the average allocation of others in the same group, and round dummies. As shown, the allocations of the subjects in the untouchable position are mostly influenced by lagged allocations. The variable, lagged deviation from mean allocations of other subjects in the group, is negative for both untouchable networks and highly significant when pooling data from the *untouchable* and *untouchable-6* networks ($p = 0.009$).

[12] As a test of robustness, the regressions in Table 4 were conducted separately for cases where deviations from D were non-negative versus negative. The results reported in Table 4 are robust to this analysis; although statistically more significant for non-negative deviations. This analysis is available in the Supplementary Material (Tables B7 and B8).

[13] The above analysis highlights the asymmetry in the reactions to sanctions received related to punishment capacities in both incomplete networks. In particular, there is some evidence that sanctions increase future contributions to the public good only when punishment capacities are high. However, the regression estimates indicate that this effect is small; in the untouchable-6 network, the effect is not significant at the 10 % level. This small reaction, combined with the low sanctioning levels observed, leads to the finding that punishment capacities do not significantly affect contributions or efficiency at the aggregate level (Result 3).

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

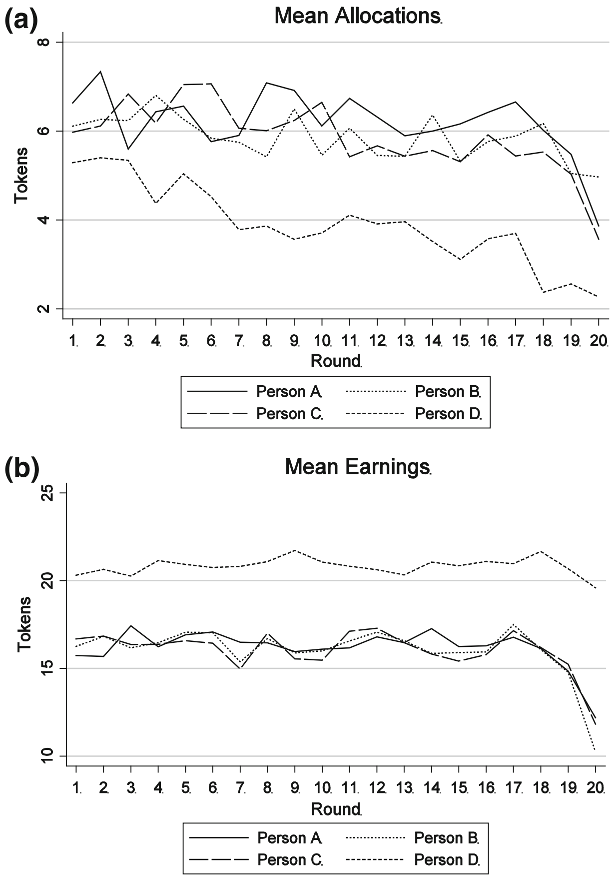

Fig. 5 **a**, **b** Allocations and earnings by network position: combined untouchable networks

Result 5: Subjects condition their contribution on the behavior of subjects in and outside their punishment network.

3.2.3 Patterns of sanctioning behavior

Pooling across treatments and observations within specified intervals, Fig. 6 shows the relationship between average sanctions received by individuals and the deviation of their group allocation from the average allocations of others in the group.[14] Also reported are the number of instances in which sanctions were imposed within each interval. Mean sanctions received are larger when a subject's allocation is below the average allocation of others. Importantly however, there is evidence of 'anti-social' punishment: some subjects are sanctioned even when their allocations are above the mean of others.

[14] Computing the average sanction for each category includes both sanctions imposed and instances in which a sanction was not imposed.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 4 Individual allocations (A, B, C): untouchable and untouchable-6 networks

Dependent variable: individual allocations—persons A, B, C	Untouchable network	Untouchable-6 network	Combined untouchable networks
Lagged allocation of *i*	0.930***	0.948***	0.946***
	(0.031)	(0.036)	(0.021)
	[0.000]	[0.000]	[0.000]
Lagged deviation from allocation of Person D	−0.088***	−0.102**	−0.101***
	(0.021)	(0.049)	(0.022)
	[0.000]	[0.037]	[0.000]
Lagged deviation from mean allocation of others in punishment network	−0.408***	−0.395***	−0.402***
	(0.035)	(0.105)	(0.042)
	[0.000]	[0.000]	[0.000]
Lagged sanctions received	−0.215*	0.113	−0.073
	(0.117)	(0.103)	(0.089)
	[0.066]	[0.271]	[0.415]
Constant	0.298	1.410***	0.712*
	(0.579)	(0.417)	(0.371)
	[0.607]	[0.001]	[0.055]
Observations	855	627	1,482
Clusters/groups	15	11	26

Numbers in parentheses are robust standard errors clustered on independent groups. Figures in brackets are *p* values for the two-sided tests of significance

*** Sig. at 1 %, ** sig. at 5 %, * sig. at 10 %

As discussed above, this study employed a matching protocol with fixed identifiers for each decision maker in a group. An advantage of this protocol is that it captures a critical informational component of some networks. More precisely, unlike previous studies examining sanctioning, this protocol allows for sanctioning imposed on subject *i* by subject *j* to be based directly on lagged sanctions imposed by *i* on *j*. Thus, linkages between sanctions imposed and lagged sanctions received between pairs of subjects within networks (referred to as 'sanctioning pairs') can be examined.

Table 6 presents regressions of individual sanctions imposed on subject *i* by subject *j* as a function of deviations in contributions by *i* from others in the group, one period lagged sanctions imposed by *i* on subject *j*, treatment dummies for the pairwise and untouchable networks[15] and round dummies. Separate regressions are estimated for negative and non-negative deviations. The table reports robust standard errors clustered on independent groups. The results in Table 6 show the usual pattern for sanctioning when deviations are below the average of the others in

[15] The pairwise dummy captures both the *pairwise* and the *pairwise-6* treatments, similarly for the untouchable treatment dummy.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 5 Individual allocations (D): Untouchable and Untouchable-6 networks

Dependent variable: individual allocations person D	Untouchable network	Untouchable-6 network	Combined untouchable networks
Lagged allocation of *i*	0.813***	0.869***	0.831***
	(0.108)	(0.072)	(0.066)
	[0.000]	[0.000]	[0.000]
Lagged deviation from mean allocation of A, B, C in group	−0.133	−0.212*	−0.178***
	(0.086)	(0.121)	(0.068)
	[0.119]	[0.079]	[0.009]
Constant	2.664***	−1.054	1.088
	(0.938)	(1.379)	(0.818)
	[0.004]	[0.445]	[0.183]
Observations	285	209	494
Clusters/groups	15	11	26

Numbers in parentheses are robust standard errors clustered on independent groups. Figures in brackets are *p* values for the two-sided tests of significance

*** Sig. at 1 %, ** sig. at 5 %, * sig. at 10 %

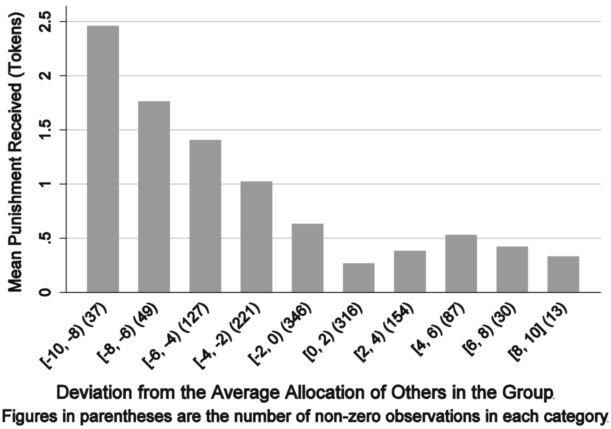

Fig. 6 Mean sanctions received by individuals

the group.[16] Players are punished for low contributions and they receive higher sanctions the lower their contributions are below the average; players receive an additional 0.9 tokens in sanctions for every token they are below the average.

[16] The results are robust to OLS and Tobit specifications.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Table 6 Evidence on targeted revenge in sanctioning pairs

Dependent variable: individual sanctions imposed by *j* on *i*

	Negative deviations	Positive deviations
Absolute value of negative allocation deviations by *i* from average of others in group	0.093***	–
	(0.010)	
	[0.000]	
Absolute value of positive allocation deviations by *i* from average of others in group	–	−0.0009
		(0.005)
		[0.843]
Lagged pairwise sanctions imposed by *i* on *j*	0.182***	0.155***
	(0.045)	(0.034)
	[0.000]	[0.000]
Pairwise	0.133	0.097*
	(0.104)	(0.057)
	[0.202]	[0.088]
Untouchable	0.045	0.060
	(0.067)	(0.048)
	[0.502]	[0.208]
Constant	0.333***	0.375***
	(0.103)	(0.064)
	[0.001]	[0.000]
Observations	3,145	6,165
Clusters [sanctioning pairs]	68 [424]	69 [479]

Numbers in parentheses are robust standard errors clustered on independent groups. Figures in brackets are *p* values for the two-sided tests of significance. Includes round dummies

*** Sig. at 1 %, ** sig at 5 %, * sig at 10 %

We do not find significant evidence showing that (weakly) positive deviations from the group average lead to 'anti-social' punishment.[17] However, there is strong evidence of *targeted revenge*. Players receive sanctions from those they sanctioned in the previous round. Such targeted revenge occurs independently of whether a subject's contribution is greater (positive deviation) or smaller (negative deviation) than the average of other group members.

Result 6: Targeted revenge drives anti-social punishment in our networks.

4 Conclusions

This study contributes to the literature on sanctioning behavior in social dilemma settings by examining the influence of alternative linkages between subjects that

[17] The results are unchanged if a dummy for positive/negative deviations is included instead of the magnitude of such deviations.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

restrict the directional flow of endogenously imposed sanctions, as well as the capacity to sanction at the individual and group level. We find clear evidence that the structure of punishment network affects public good contributions and that the network configuration is more important than the absolute punishment capacity for public good provision, imposed sanctions and economic efficiency. In addition, our experimental design renders it possible to identify targeted revenge as a main driver of anti-social punishment.

The results of this study may have implications for public policy and organizational thinking related to the pervasive conflict of individual interest and collective efficiency. In a world where natural obstacles and manmade institutions limit stakeholder's opportunities to sanction other actors, a proper understanding of underlying group structures and how individual actors connect to each other is crucially important when trying to understand the nature of voluntary cooperation. This study suggests that the nature of incomplete sanction networks may be more important than the group's overall capacity to sanction. This result raises the question of whether and how collective action groups in the field can develop institutions or social norms to overcome such incompleteness.

Acknowledgments Financial support was provided by the National Science Foundation (SES—0849551) and the School of Economics and the Centre for Behavioural and Experimental Social Science, University of East Anglia. The authors would also like to thank Dvin Galstian Pour for programming assistance, as well as David Cooper, Enrique Fatas, Simon Gächter, Elke Renner, Martin Sefton, Sudipta Sarangi, two anonymous referees and seminar participants from Durham, Nottingham, the Experimental Economics Workshop at the Indian School of Business, the 2011 Southern Economic Association Annual Meetings, the 2011 North American Economic Science Association Conference, the 2011 European Economic Science Association Conference and the 5th Maastricht Behavioral and Experimental Economics Symposium for their helpful comments.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

Andreoni, J., Harbaugh, W., & Vesterlund, L. (2003). The carrot or the stick: Rewards, punishments and cooperation. *American Economic Review, 93*(3), 893–902.

Bochet, O., Page, T., & Putterman, L. (2006). Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior & Organization, 60*(1), 11–26.

Boyd, R., Gintis, H., & Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science, 328*(5978), 617–620.

Bramoullé, Y., & Kranton, R. (2007). Public goods in networks. *Journal of Economic Theory, 135*(1), 478–494.

Carpenter, J. (2007a). Punishing free-riders: How group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior, 60*(1), 31–51.

Carpenter, J. (2007b). The demand for punishment. *Journal of Economic Behavior & Organization, 62*(4), 522–542.

Carpenter, J., Kariv, S., & Schotter, A. (2012). Network architecture, cooperation and punishment in public good games. *Review of Economic Design, 95*(1), 1–26.

Charness, G., Corominas-Bosch, M., & Frachette, G. R. (2007). Bargaining and network structure: An experiment. *Journal of Economic Theory, 136*(1), 28–65.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Charness, G., & Jackson, M. (2007). Group play in games and the role of consent in network formation. *Journal of Economic Theory, 136*(1), 417–445.
Chaudhuri, A. (2011). Sustaining cooperation in laboratory public good experiments: a selective survey of the literature. *Experimental Economics, 14*(1), 47–83.
Cinyabuguma, M., Page, T., & Putterman, L. (2005). Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics, 89*(8), 1421–1435.
Cox, J. C., Ostrom, E., Sadiraj, V., & Walker, J. M. (2013). Provision versus appropriation in symmetric and asymmetric social dilemmas. *Southern Economic Journal, 79*(3), 496–512.
Denant-Boemont, L., Masclet, D., & Noussair, C. (2007). Punishment, counterpunishment and sanction enforcement in a social dilemma experiment. *Economic Theory, 33*(1), 145–167.
Egas, M., & Riedl, A. (2008). The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society B—Biological Sciences, 275*(1637), 871–878.
Ertan, A., Page, T., & Putterman, L. (2009). Who to punish? Individual decisions and majority rule in mitigating the free rider problem. *European Economic Review, 53*(5), 495–511.
Fatas, E., Meléndez-Jiménez, M., & Solaz, H. (2010). An experiment analysis of team production in networks. *Experimental Economics, 13*(4), 399–411.
Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review, 90*(4), 980–994.
Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics, 10*(2), 171–178.
Gürerk, Ö., Irlenbusch, B., & Rockenbach, B. (2006). The competitive advantage of sanctioning institutions. *Science, 312*(5770), 108–111.
Güth, W., Levati, M. V., Sutter, M., & van der Heijden, E. (2007). Leading by example with and without exclusion power in voluntary contribution experiments. *Journal of Public Economics, 91*(5–6), 1023–1042.
Hermann, B., Thöni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science, 319*(5868), 1362–1367.
Isaac, M., & Walker, J. (1988). Communication and free-riding behavior: The voluntary contribution mechanism. *Economic Inquiry, 26*(4), 585–608.
Keser, C., Ehrhart, K.-M., & Berninghaus, S. K. (1998). Coordination and local interaction: Experimental evidence. *Economics Letters, 58*(3), 269–275.
Kosfeld, M. (2004). Economics networks in the Laboratory: A survey. *Review of Network Economics, 3*(1), 20–42.
Kosfeld, M., Okada, A., & Riedl, A. (2009). Institution formation in public goods games. *American Economic Review, 99*(4), 1335–1355.
Leibbrandt, A., & López-Pérez, R. (2014). Different carrots and different sticks: Do we reward and punish differently than we approve and disapprove? *Theory and Decision, 76*(1), 95–118.
Masclet, D., Noussair, C., Tucker, S., & Villeval, M.-C. (2003). Monetary and non-monetary punishment in the voluntary contributions mechanism. *American Economic Review, 93*(1), 366–380.
Milinski, M., & Rockenbach, B. (2006). The efficient interaction of indirect reciprocity and costly punishment. *Nature, 444*, 718–723.
Nicklisch, A., & Wolff, I. (2011). Cooperation norms in multiple-stage punishment. *Journal of Public Economic Theory, 13*(5), 791–827.
Nikiforakis, N. (2008). Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics, 92*(1–2), 91–112.
Nikiforakis, N., & Engelmann, D. (2011). Altruistic punishment and the threat of feuds. *Journal of Economic Behavior & Organization, 78*(3), 319–332.
Nikiforakis, N., & Normann, H. (2008). A comparative statics analysis of punishment in public goods experiments. *Experimental Economics, 11*(4), 358–369.
Nikiforakis, N., Normann, H.-T., & Wallace, B. (2010). Asymmetric enforcement of cooperation in a social dilemma. *Southern Economic Journal, 76*(3), 638–659.
O'Gorman, R., Henrich, J., & Van Vugt, M. (2009). Constraining free riding in public goods games: Designated solitary punishers can sustain human cooperation. *Proceedings of the Royal Society, 276*(1655), 323–329.
Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.
Ostrom, E., Walker, J., & Gardner, R. (1992). Covenants with and without a sword: self-governance is possible. *American Political Science Review, 86*(1), 404–417.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.

Rand, D. G., Armao, J. J., Nakamaru, M., & Ohtsuki, H. (2010). Anti-social punishment can prevent the coevolution of punishment and cooperation. *Journal of Theoretical Biology, 265*(4), 624–632.
Reuben, E., & Riedl, A. (2009). Public goods provision and sanctioning in privileged groups. *Journal of Conflict Resolution, 53*(1), 72–93.
Rosenkranz, S., & Weitzel, U. (2012). Network structure and strategic investments: An experimental analysis. *Games and Economic Behaviour, 75*(2), 898–920.
Sefton, M., Shupp, R., & Walker, J. (2007). The effect of rewards and sanctions in provision of public goods. *Economic Inquiry, 45*(4), 671–690.
Walker, J., & Halloran, M. (2004). Rewards and sanctions and the provision of public goods in one-shot settings. *Experimental Economics, 7*(3), 235–247.
Yamagishi, T., & Cook, K. (1993). Generalized exchange and social dilemmas. *Social Psychology Quarterly, 56*(4), 235–248.

Downloaded from https://www.cambridge.org/core. 08 May 2025 at 13:42:09, subject to the Cambridge Core terms of use.