

Monitoring and punishment networks in a common-pool resource dilemma: experimental evidence

Ganga Shreedhar, Alessandro Tavoni
and Carmen Marchiori

January 2018

Centre for Climate Change Economics
and Policy Working Paper No. 327
ISSN 2515-5709 (Online)

Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 292
ISSN 2515-5717 (Online)

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its second phase started in 2013 and there are five integrated research themes:

1. Understanding green growth and climate-compatible development
2. Advancing climate finance and investment
3. Evaluating the performance of climate policies
4. Managing climate risks and uncertainties and strengthening climate services
5. Enabling rapid transitions in mitigation and adaptation

More information about the Centre for Climate Change Economics and Policy can be found at: www.cccep.ac.uk.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute. It has six research themes:

1. Sustainable development
2. Finance, investment and insurance
3. Changing behaviours
4. Growth and innovation
5. Policy design and evaluation
6. Governance and legislation

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: www.lse.ac.uk/grantham.

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the authors and do not necessarily represent those of the host institutions or funders.

Monitoring and punishment networks in a common-pool resource dilemma: experimental evidence

Ganga Shreedhar^ψ, Alessandro Tavoni^φ and Carmen Marchiori^φ

January 2018

Abstract

In an experimental study, we explore how imperfect monitoring and punishment network architectures impacts cooperation, punishment and beliefs, in a non-linear common pool resource appropriation dilemma. We find that complete networks (with perfect monitoring and punishment), are the least efficient due to higher punishment, relative to incomplete networks. In addition, high appropriators are sanctioned in all networks, but well-connected and undirected networks elicit higher anti-social punishment. Lastly, although subject's underestimate other's appropriation in all networks, the difference between beliefs and other's appropriation declines with time. This decline occurs faster in complete networks, relative to incomplete but connected networks.

KEYWORDS: Monitoring, Punishment, Networks, Common Pool Resources, Beliefs

JEL CODES: Q57, D83, D62, D70, C91, C92

ACKNOWLEDGEMENTS:

We thank the Grantham Research Institute on Climate Change and the Environment and the ESRC Centre for Climate Change Economics and Policy for funding this research. We also are grateful for comments received at EAERE 2016, Bioecon 18, ICES-ASC 2016, Grantham Research Institute on Climate Change and the Environment Empirical Seminar and the Environment cluster seminar at the Department of Geography and Environment. We extend special thanks to Timo Goeschl and other reviewers.

^ψ Department of Geography and Environment, London School of Economics and Political Science. Address: Houghton Street, WC2A 2AE. Telephone: 04402071075027. E-mail: g.s.shreedhar@lse.ac.uk

^φ Grantham Research Institute for Climate Change and the Environment, London School of Economics and Political Science. Address: Houghton Street, WC2A 2AE. Telephone: 04402071075027. E-mail: a.tavoni@lse.ac.uk

^φ Department of Economics and Management, Università Degli Studi di Brescia. Address: Via San Faustino 74 / b. Telephone: 0302988839-840. E-mail: carmen.marchiori@unibs.it

Monitoring and punishment networks in a common-pool resource dilemma: experimental evidence

I. INTRODUCTION

Peer monitoring and punishment is an important means to mitigate free-riding in social dilemmas. Ostrom et al., (1992, 1994) demonstrate that individuals monitor and punish free-riders even at personal cost, in common-pool resource (CPR) appropriation dilemmas. Fehr and Gächter (2000) find that cooperation is lower without peer punishment in public goods (PG) experiments. But many extant studies implicitly assume perfect monitoring and punishment, i.e., everyone can monitor and punish everyone else.¹

Recent experimental evidence reveals that imperfect peer monitoring and punishment, represented by an underlying network connecting agents, impacts cooperation in voluntary contribution mechanism (VCM) PG experiments (Carpenter et al., 2012, Leibbrandt et al., 2015; Boosey and Issac, 2016). For instance, Carpenter et al. (2012) find the network architecture impacts contributions, punishment and welfare. Leibbrandt et al. (2015) find that contributions to the public good are higher when more punishment opportunities are available, but that punishment is also higher. Boosey and Issac (2016) note that the network structure affects the incidence of anti-social punishment. We extend this literature, by providing new experimental evidence on how the monitoring and punishment network architecture, affects behaviour and beliefs in a CPR appropriation dilemma.

CPR appropriation dilemmas represent strategic situations, where one agent's consumption of resource units, such as fish or water, removes those units from the resource system. Unlike public goods dilemmas, an agent's appropriation inflicts a negative externality on others through the reduction in payoffs, due to the rivalry of resource units (Ostrom et al., 1994).² The baseline appropriation game, pioneered by Ostrom et al., (1992), uses a non-linear payoff function to locate the socially optimal and selfish equilibria in the interior of the strategy space.

¹ See Anderies et al. (2011), Poteete et al. (2010) and Ostrom (2006) for a review of CPR experiments in the lab and field, and Sturm and Weimann (2006) and Noussair and van Soest (2014) for a review CPR experiments alongside other environmental dilemmas. Chaudhuri (2011) and Ledyard (1995) survey PG experiments.

² A PG game models a 'provision' dilemma, where others' contributions generate a positive supply externality because the good is non-excludable and non-rival (Ostrom, 2006). While both games can be strategically equivalent or different frames of the same game (Ledyard 1995), others consider CPR dilemmas as theoretically distinct because of the aspect of rivalry (Apesteguia and Maier-Rigaud, 2006).

This non-linear payoff structure aims to approximate the complexity faced by agents in the field, arising from socio-ecological dynamics (Ibid.).

The role of peer punishment in non-linear CPR appropriation dilemmas is under-examined. Although Ostrom et al. (1992)'s seminal paper is widely cited to demonstrate the effectiveness of peer punishment, the original study reports punishment is erratic and has a limited impact on cooperation (also see Ostrom et al., 1994). In a recent study, Cason and Gangadharan (2015) also observe that peer punishment has a weaker impact on appropriation, resulting in outcomes closer to the Nash equilibrium in non-linear CPR games, compared to linear VCM PG games. They note, "... the effectiveness of peer punishment could vary in different environments; particularly in situations with more complexity that makes it more difficult to identify defectors" (Cason and Gangadharan, 2015, pp. 86). Welfare is also lower, as punishment reduces payoffs.³ More research is required to understand the impact of peer monitoring and punishment, in social dilemmas characterised by complex non-linear payoffs, and resource rivalry.

We address this literature gap, by exploring the role of the network architecture in promoting cooperation in a non-linear CPR appropriation dilemma, using a lab experiment. We exogenously vary the monitoring and punishment network, which in turn jointly determines feedback and punishment opportunities available to agents. Appropriation from the CPR, stated beliefs about other member's appropriation and punishment, is compared over fifteen rounds of play, in groups of four using a stranger-matching protocol.

We select four regular types of networks with non-trivial local and global network properties: the complete, undirected circle, directed circle and line networks. This choice of networks yields two advantages. First, it allows us to examine if outcomes vary as the number of monitoring and punishment opportunities available to agents systematically declines. Second, we can focus on the impact of four graph-theoretical properties of the network architecture, namely completeness, connectedness or directedness, and node degree (network properties are defined in section two). Thus, our contribution is primarily an empirical one: highlighting how beliefs and behaviour vary due to the joint effect of the monitoring and punishment network architecture.

³ On the other hand, recent experiments find behaviour in *non-linear* PG, and CPR games qualitatively similar, as cooperation starts near the socially optimal level and trends towards but remains below the selfish equilibrium (Kingsley, 2015; Kingsley and Liu, 2012; Apesteguia & Maier-Rigaud, 2006).

It is not immediately evident that perfect monitoring and punishment networks necessarily increases cooperation. On the one hand, it allows agents a higher opportunity to punish a larger number of agents to mitigate over-appropriation. On the other hand, agents face a second-order monitoring dilemma, if coordinating costly punishment decisions becomes more difficult in well-connected networks. In this case, agents in imperfect, less-connected networks can focus their punishment on their immediate neighbours, and punishment may be more severe. Besides, agents may find it especially challenging to differentiate between the free-riding and socially efficient appropriation, due to the complexity of the non-linear strategy space. Since perfect monitoring and punishment networks offer greater opportunities for information feedback, it can possibly facilitate a greater (and faster) convergence between beliefs over other's appropriation and their actual behaviour, compared to less-connected networks.⁴

The primary finding from our study is that complete networks – the baseline with perfect monitoring and punishment – elicit the lowest payoffs to agents, and are consequently the least efficient. As average appropriation is not significantly different across networks, this result is primarily driven by heavy pro-social and anti-social punishment received by subjects. Anti-social punishment is higher in undirected and connected networks, but lower and less systematic in directed networks, especially the directed circle. Completeness is significantly associated with higher punishment and beliefs, and lower efficiency. We also find that subject's beliefs underestimated other's appropriation across networks. The difference between beliefs and other's (mean) appropriation declines with time. However, this decline is contingent on the type of network: subjects in the undirected circle and directed circle networks record higher average differences in the last period of play (rounds 11-15), compared to the nodes in the complete network.

Our paper contributes to the experimental literature on how to mitigate the over-appropriation of common-pool resources with peer monitoring. This includes work on peer monitoring and punishment (Cason and Gangadharan, 2015; Kinsley et al., 2015; Janssen et al., 2010; Casari and Plott, 2003; Ostrom et al., 1992), communication with peer monitoring (Ostrom et al., 1992; Cardenas et al., 2004; Cason and Gangadharan 2016), and the relative effectiveness of

⁴ The natural resource governance literature provides empirical support that well-connected social networks elicit both cooperation and over-appropriation. For instance, King (2000) finds that high levels of social interaction and dense social ties amongst local fisherman enabled them to manage unfavourable developments related to the local fishery. Conversely, Bodin and Crona (2008), find that despite dense social ties, fishers in another Kenyan fishing village were reluctant to call out rule-breaking and had few institutions to overcome over-exploitation of local fisheries. One of the possible reasons they outline, is the homogeneity of information and behaviour among key individuals, leading to reduced problem recognition. See Bodin and Crona (2009) for a review of this literature.

rewards over fines (Vyrastekova and Van Soest, 2008). Ostrom et al. (1992), Cason and Gangadharan (2015), Kinsley et al., (2015) are the closest to ours, because they also use a non-linear CPR appropriation game with peer punishment.⁵ While they examine the impact of perfect peer monitoring and punishment, we use perfect monitoring and punishment as the baseline and vary the structure of monitoring and punishment instead. Other key differences relate to the experimental design: we use a strangers-matching protocol, to minimise investment in strategic play and reputation, and to retain our focus on the network architecture.⁶ We also elicit incentivised beliefs in the first stage, as discussed in the following paragraphs.

We also extend the literature on how imperfect peer monitoring and punishment impacts cooperation in linear VCM PG games (Boosey & Isaac, 2016; Leibbrandt et al., 2015; Carpenter et al., 2012).⁷ A consistent finding is that incomplete but connected networks elicit at least as many contributions to the public good as complete networks (Boosey & Isaac, 2016; Leibbrandt et al., 2015; Carpenter et al., 2012). Carpenter et al., (2012) are most closely related to our experiment, as they examine the joint effect of the monitoring and punishment network architecture, using a strangers-matching design. They find that relatively lower punishment in the complete network makes it more efficient than other network structures. On the other hand, Leibbrandt et al., (2015) and Boosey and Issac (2016) vary only the punishment network, and find that subjects use greater punishment when more punishment opportunities are available, but that individual punishment behaviour depends on the type of network. Apart from using a non-linear CPR game, we differ from these studies by eliciting incentivised beliefs about the expected appropriation of other's in the network.⁸

Experimental evidence on beliefs from linear CPR games (Velez et al., 2009) and linear VCM PG games (Croson, 2000; Neugebauer et al., 2009; Fischbacher and Gächter, 2010; Gächter and Renner, 2010; Smith, 2013) offers two consistent findings: first, subjects consistently state

⁵ Related studies on peer monitoring in CPR games includes the following. Casari and Plott (2003) experimentally test a historically-rooted peer punishment institution where punishers receive the 'fines' paid by the punished and find punishment improves cooperation. Vyrastekova and van Soest (2006) use a quadratic CPR game and find punishment is better than rewards at eliciting cooperation. Janssen et al., (2010) incorporate complex temporal socio-ecological dynamics in the CPR game but do not find punishment to be effective in increasing cooperation.

⁶ We also modify the punishment technology, such that punishment points that can be allocated by subjects depend on their earnings in the first stage (following Carpenter et al., 2012; also see Nikiforakis, 2008).

⁷ See Kosfeld (2004) and Choi et al. (2016) for reviews of experiments on networks in the lab.

⁸ Another difference is that we increase the cost of punishment, such that the cost of receiving one deduction point is three tokens. This matches a standard fee-to-fine ratio used in CPR experiments but is higher than Carpenter et al., (2012) who use a fee-to-fine ratio of 0.5:1 (i.e., it costs the punisher 0.5 tokens to assign one deduction point). Leibbrandt et al., (2015) and Boosey and Issac (2016) also use a partners-matching protocol, to allow the formation of agent's reputations. They use a different punishment technology that endows players with additional tokens in stage two to assign punishment. Boosey and Issac (2016) also use a 1:3 fee-to-fine ratio.

‘optimistic’ beliefs by overestimating other’s contributions (or underestimating appropriation), and second, beliefs are positively correlated with contributions. We extend this literature, by eliciting incentivised beliefs across all networks to empirically explore if stated beliefs and beliefs accuracy varies by network architecture.

One means of increasing beliefs accuracy (i.e., reducing the difference between stated beliefs and other’s revealed behaviour) is by incentivising beliefs (Gächter and Renner, 2010). Instead, we explore if feedback about other’s appropriations, which is determined by the network architecture, impacts stated beliefs and belief accuracy. Given the complexity of the non-linear CPR game, uninformed agents may face greater difficulties in thinking of other’s expected appropriations because they lack common knowledge about other’s revealed behaviour in the lab. Conversely, informed agents in better-connected networks can learn more accurately about other’s appropriation in each round.⁹ Greater feedback may lead to a faster convergence between beliefs and behaviour, and therefore higher belief accuracy, in better-connected networks, with a larger share of informed agents.¹⁰ To the best of our knowledge, we are unaware of any other studies which elicit beliefs in either a non-linear CPR dilemma, or alongside the monitoring and punishment network architecture.¹¹

The paper is organized as follows. Section 2 introduces the networks and the structural graph properties; this is followed by a description of the CPR appropriation dilemma and the key hypotheses. Section 3 outlines the experimental design, and section 4 presents the results. We discuss the potential implications for our results and conclude in section 5.

⁹ Notably, previous work has remarked that the greater complexity in non-linear CPR appropriation dilemmas is an important feature impacting appropriation and punishment behaviour (e.g. in Cason and Gangadharan, 2015). Others have noted subjects may experience greater confusion in the CPR appropriation game in the lab (Sturm & Weimann, 2006).

¹⁰ Some intuition can be obtained from models of social learning on networks, although they do not directly model appropriation in CPR dilemmas. For instance, Gale and Kariv (2003) predict that subject’s beliefs and behaviour, ought to converge faster in well-connected networks with more informed subjects (this is experimentally verified in guessing game in Choi et al. (2012)).

¹¹ As the goal of our paper is to examine if the network properties impact reported beliefs in systematic ways, we do not systematically explore if beliefs (and beliefs elicitation) impacts appropriations.

II. CONCEPTUAL FRAMEWORK

Network architecture and properties

Figure 1 (Panel 1A-D), illustrates the four monitoring and punishment networks used in the experiment, and their graph-theoretic properties. A network consists of three agents or appropriators indexed by $i = A, B, C, D$. An edge between two agents indicates that they are connected, and the arrowhead points to the agent whose appropriation can be monitored and punished. For each subject i , N_i denotes the set of subjects $j \neq i$ who can be monitored by i ; in other words, N_i is the monitoring neighbourhood of subject i . An 'undirected' or bi-directional edge between any two nodes implies that the pair of extractors can monitor and punish each other. Conversely, a 'directed' or unidirectional edge denotes that only the extractor to whom the edge points can be observed and punished, but not vice-versa. Thus, if an edge runs from i to j , then i can monitor j ($j \in N_i$), but j cannot monitor i ($i \notin N_j$).

We consider four key properties: completeness, directedness, connectedness, and node degree. The first three are global properties of the network architecture, while the fourth is a local property of the monitoring neighbourhoods that characterise the network. The definition of each network property is briefly summarised below.

1. Completeness [network 0]: each pair of nodes is connected by an undirected edge to represent perfect monitoring, i.e. everyone can observe everyone's appropriations and can choose to punish any extractor in the network. Networks [1-3] are incomplete.
2. Directedness [networks 1-2]: if the edge between pairs of nodes is not bidirectional, the network is a directed network [networks 2 and 3]; otherwise it is undirected [networks 0 and 1].
3. Connectedness [networks 1, 2 and 3]: if every pair of nodes i and j are connected by a path, the network is connected [networks 1-3]; otherwise it is disconnected [network 4]. A 'path' is a sequence of nodes wherein each node in the graph is used only once.
4. Node degree: the number of edges that end at a given node. The out-degree (in-degree) of node i is the number of edges with i as their initial (terminal) node. Specifically, the out-degree of subject i is the number of subjects j that can be monitored and punished by i ($j \in N_i$) and the in-degree of subject i is the number of subjects j that can monitor i ($i \in N_j$). Since nodes are defined by their out- and in-degrees, we can index them as Nn, i, o where n is the number of each network indicated in Figure 1, and i and o represent the number of in- and out-degrees respectively. For instance, the complete

network consists of four identical or symmetric nodes denoted by N033, where each node has three in-degrees and three out-degrees.

The complete network (Figure 1, Panel 1A) represents perfect monitoring and punishment and is the baseline. It is complete, connected, and undirected, with symmetric nodes denoted by N033. The undirected circle network (Panel 1B) is incomplete, connected and undirected, also with symmetric nodes N122. The directed circle (Panel 1C) is a directed network, which is incomplete and connected with symmetric nodes, N211. Finally, the directed line network (Panel 1D) is the only disconnected network, which is incomplete and directed, with asymmetric nodes: N301 (Node A) cannot be monitored or punished (but can extend punishment to Node B); N311 (Nodes B and C) can be monitored and can punish one agent; N310 (Node D), can be monitored and punished, but cannot monitor or punish anyone.

The CPR appropriation dilemma

We modify the baseline CPR appropriation dilemma from Ostrom et al., (1992). In stage one, each agent is endowed with an equal initial amount of effort e . Agents simultaneously choose how to allocate this initial amount of effort between CPR appropriation and an outside-option activity (e.g., agricultural labour), which yields a constant marginal rate of return, w . Specifically, we denote the appropriation effort of agent i by $x_i \in [0, e]$. Hence, the total appropriation effort of all agents within the network is $X = \sum_i x_i$.

The total return on appropriation is given by a non-linear and concave function: $F(X) = aX - bX^2$. This yields a payoff structure where it is initially more profitable to allocate some portion of the endowment to the CPR. However, allocating the entire endowment is counterproductive because $F'(0) > w$ and $F'(ne) < 0$.

Since the CPR is non-excludable and rival in consumption, the return from appropriation to an individual agent is proportional to the ratio between her appropriation effort (x_i), and the total effort of all appropriators (X). The ratio x_i/X is called the individual distribution factor (Apesteguia and Maier-Rigaud, 2006) and captures the negative externality from appropriating rival resource units. The higher x_i is in relation to X , the higher is i 's appropriation from the common pool resource, allowing an appropriator to capture a larger share of returns from the common pool resource ($aX - bX^2$). This setup creates the social dilemma, where each agent has a dominant strategy to appropriate at the inefficient level from the CPR. The first-stage payoff of each agent i is given by:

$$\pi_i^1 = w(e - x_i) + (x_i/X)(aX - bX^2) \quad (1)$$

In the second stage, agents can observe the appropriation of those they are connected to by the network and punish over-appropriation at some cost. Each agent, i can punish j ($j \in N_i$) to reduce j 's payoff from the first stage by p_j^i at a personal unit cost of c_1 , such that $0 < c_1 \leq 1$. Each punishment point received reduces j 's payoff by some factor c_2 ($0 \leq c_2$). The payoff to each agent from stage two is the maximum between 0 and the net return from stage one;

$$\pi_i^2 = \max\left\{0, \pi_i^1 - c_1 \sum_{j \in N_i} p_j^i - c_2 \sum_{i \in N_j} p_i^j\right\} \quad (2)$$

If we assume rational and self-interested agents, by backwards induction, it follows that punishment in stage two cannot deter free-riding in stage one, irrespective of the network architecture. Specifically, as punishment is costly, each subject will issue no punishment in stage two ($p_i^j = 0$, for all i and $j \in N_i$). As subjects expect no punishment in stage two, they will extract at the socially inefficient level in stage one. However, as discussed earlier, evidence from linear VCM PG games suggest that the underlying structural properties of the network affect the provision of the public good. Thus, in the following section, we develop some hypotheses to test the impact of the network architecture for each outcome.

Hypothesis

First, previous results suggest that average contributions are similar in the complete, undirected and directed circle networks (Carpenter et al., 2012; Boosey and Isaac, 2016), but lower in the directed line network, which is a disconnected network (Carpenter et al.; Ibid.). More broadly, they find network disconnectedness decreases contributions to the public good. This yields our first hypothesis:

Hypothesis 1: Average appropriation is similar across the complete, undirected circle and directed circle networks, but higher in the disconnected directed line network.

Second, Carpenter et al., (2012) report that subjects in the complete network tend to use lower punishment, resulting in higher than average net returns compared to many imperfect networks. They also find that punishment is used more heavily in directed networks. Conversely, Leibbrandt et al. (2015) also find that contributions are greater in well-connected punishment networks, but that received punishment is also higher with greater punishment opportunity, leading to similar payoffs. At the individual level, all studies find that free-riders (i.e., those who don't contribute to the public good) and those contribute less than the average contribution

of other's in the group (i.e., low contributors) are punished more heavily. Thus, we also undertake this analysis for each network (for robustness). But given these competing arguments about the effect of the network, we formulate the following null hypotheses for punishment:

Hypothesis 2: Average punishment is similar across the complete, undirected circle, directed circle and directed line networks. Free-riders and low-contributors are punished more heavily in all networks.

Third, we test the null hypothesis that stated beliefs about the expected appropriation of others in the group, do not vary across networks. Agents in the complete network are perfectly informed about other's actual appropriations, and may, therefore, state more accurate beliefs, i.e., the difference between beliefs and other's appropriation may be lower. In particular, less time may be needed for beliefs and behaviour to converge in the complete network. Conversely, poorly informed agents in less-connected networks, lack of knowledge of other's actions, and may need a longer period of learning. The impact of each network type and property on beliefs accuracy is also an open empirical question, and thus we restrict our hypothesis to the effect of network completeness, which is the baseline with perfect monitoring and punishment. This yields our last two hypotheses:

Hypothesis 3: Average beliefs over other's appropriation is similar across the complete, undirected circle, directed circle and directed line networks.

Hypothesis 4: The average difference between beliefs and other's actual appropriation is lower in the complete network, relative to incomplete networks. Furthermore, the convergence of beliefs is also faster in complete networks.

III. EXPERIMENTAL DESIGN

The experiment was held at the London School of Economics and Political Science Behavioural Research Lab (LSE BRL) during November-December 2015. Participation was open only to students previously registered at the LSE BRL and was executed using z-Tree (Fischbacher, 2007). Each of the ten sessions consisted of 16-20 subjects, who were randomly assigned to a computer terminal at the start of each session, and given written instructions.¹²

¹² A concerted effort was made to ensure that subjects understood the instructions, incentives and payoffs, to ensure preferences and beliefs were measured as accurately as possible. Subjects also received payoffs table, with the written instructions. Instructions were first read by the subjects and then read aloud by the experimenter. After this, subjects answered five control questions, which were checked by the experimenter to ensure subjects

We follow Carpenter et al., (2012) to hold one of the four network treatments (i.e., complete, undirected circle, directed circle, or directed line) constant throughout a given session. Subjects were randomly assigned to one of the node labels labelled A, B, C, and D, which was also fixed throughout the session. Each of the 15 rounds started with the computer randomly forming a network group of four, containing one of each type of node label. The network groups formed in each round are independent of the networks formed in any of the other rounds, and we emphasised that the groups of four would be randomly reshuffled in each round. We used this stranger-matching protocol to mitigate investment in reputations from strategic play.¹³

Each round consists of two stages. In stage one, subjects simultaneously choose to ‘place’ tokens from their endowment to the common and private account and enter their beliefs over the appropriation of the other three subjects in the network. The endowment is 25 tokens, and the parameters are $a = 25$, $b = 0.025$, and $w = 5$ (see equation (1)). The Nash equilibrium allocation is sixteen tokens, and Pareto optimal allocation is ten tokens. The ratio of the Nash to Pareto equilibrium strategies is 1.6, and payoffs from the Nash and Pareto equilibrium allocation are 189 and 225 tokens respectively.¹⁴

Along with the appropriation decision, subjects also had to input their beliefs about the expected token placement of the other three group members in the common account, in stage one, in each round. Subjects were incentivised such that they were paid for the accuracy of their estimates, in addition to their earnings from the CPR experiment.¹⁵ After subjects take their decision, everyone is shown their total payoffs from the first stage as shown in equation (1), and by the private and common accounts. This first stage is identical for all subjects in all treatments (including the earnings faced in the CPR game and beliefs).

familiarized themselves with the game, and only then did the experiment proceed. All instructions were framed in a neutral language (e.g. punishment points are referred to as ‘deduction points’). Instructions and screen-shots of each stage of the experiment are given in the Appendices A and B.

¹³ All sessions had 20 subjects, barring two sessions with 16 subjects due to participant absenteeism (one for the undirected circle and directed line respectively). For a typical session of 20 subjects (16): (a) the likelihood in round 1 that a player would meet another player once again during the remaining fourteen rounds was 56% (87.5%) (b) the likelihood that the same group of four players would meet during the remaining fourteen rounds was 2.24% (5.47%). Since the experiment was conducted anonymously, subjects were unable to recognise whether they were matched with a specific player in the past, but only recognised the node label (i.e., A, B, C or D).

¹⁴ The experimental parameters were chosen to ensure a good spread in values between the Nash and Pareto equilibrium allocations as integer values, and to be comparable with the existing CPR literature (Appendix, C, Table C1). The ratio of the Nash to Pareto equilibrium strategies determines the size of the negative externality from the appropriation from the commons account.

¹⁵ Participants were paid a financial incentive for stating correct beliefs, but it was small, to avoid hedging. Subjects received four experimental tokens (\approx GBP 0.5) for each correct guess for each other appropriator, in addition to their experimental earnings from the CPR game.

In stage two, subjects obtained feedback on other's appropriation from the common account and could choose to assign deduction points, based on their network and node position. We used a fee-to-fine ratio of 0.33 in all networks, such that the cost of assigning deduction points is one token ($c_1=1$) and the cost of receiving a deduction point is three tokens ($c_2=3$). If subjects did not wish to reduce the earnings of anyone else, they had to enter zeros. The maximum deduction points subjects could have assigned, was bound by the first stage payoffs. Although negative earnings are possible, this was very uncommon (this occurred in only 7 out of 2880 observations), and was given a value of zero for payment purposes. After both stages concluded, subjects were informed about their total payoffs according to the payoff function, and the cost of deduction points assigned and received, from equation (2). Subjects would then proceed to the next round of the 15 rounds of play, and this number of rounds was common knowledge, alongside the network structure and payoffs.

After the last round was played, we collected data on the subject's gender, and whether they studied economics, followed by a short questionnaire that explores the motivation behind decisions taken. Data on previous experience in experiments was obtained from the LSE BRL. One round was selected at random for payment, and all subjects were paid out for that round (i.e., payoffs from the game and accuracy of beliefs) at the end of the experiment such that 25 Experimental tokens = GBP 1. Each session usually lasted for an hour and 192 subjects who attended earned an average payment of GBP 11.5 each.

IV RESULTS

We report the results in four subsections, corresponding to the each of outcomes, i.e., appropriation, punishment, payoffs and beliefs. In each, we consider if outcomes vary across networks and nodes, and then explore if there are differences by network properties. We use random-effects regression models with robust standard errors clustered at the subject-level, instead of standard statistical tests, due to our experimental design. This allows us to control for potential session-fixed effects using session dummies and learning using round dummies. We also include individual-level controls for sex, experience in lab sessions and economics students.¹⁶ Table 1 reports the summary statistics for the key outcomes by treatment for all rounds.

¹⁶ Of the total sample, 59.35% was female, 16.15% studied economics, and the majority took part in 1-5 experiments (46.35%) and 26% in no experiment previously. Appendix C, Table C2 presents attributes by treatment group, and shows that groups were balanced on these attributes. We replicate the analysis using

Appropriation

We first consider the complete network, the baseline in the literature. Average appropriation pooled across 15 rounds is 61.5% of the endowment, which is close to, albeit marginally higher than, previous studies with perfect monitoring and punishment: for example, 56.6% in Cason and Gangadharan (2015) and 59% in Kingsley (2015).

Figure 2 illustrates that the average appropriation across networks and for each round begins close to the mid-point of the choice interval in all treatments and trends towards the Nash equilibrium. Figure 3, which displays average appropriation by node and sub-period (rounds 1-5, 6-10 and 11-15; with 95% error bars), also echoes this trend. Notably, average appropriation in the complete network (N033) seems to fall marginally in the last period (also see N311). But the figures and summary statistics indicate differences in average appropriation across networks are limited. Results of random-effects models, reported in Table 2, confirm that differences across networks are not statistically significant.

In Table 2, model (1), the coefficients of the network dummies, suggest that appropriation is lower in all networks by around 1-2 tokens compared to the complete network, but this difference is weakly significant only for the undirected circle network (at 10%). In model (2), the coefficient on the Undirected circle becomes insignificant, when we include controls, including beliefs over other's appropriation for that round, and lagged appropriation, beliefs and punishment received (from the previous round). The coefficient of beliefs is negative, but not statistically significant, and that of lagged beliefs is positive and significant at 5%.¹⁷ As expected, the coefficient of lagged appropriation is positive and significant (at 1%), and that of lagged punishment received is negative and significant (at 1%), as expected. Similar results are reported in Table 3, when node dummies are used instead of network dummies. The coefficients on Nodes N122 (in the undirected circle) and N301 (player D, in directed line) are

Random-effects Tobit regressions with qualitatively similar results, which are omitted for brevity, but available on request.

¹⁷ There is a high correlation between the beliefs and lagged beliefs variable. To illustrate, the coefficient on beliefs turns positive but is not statistically significant (both with and without other lagged variables), when we omit when lagged beliefs. However, the inclusion of lagged beliefs as an additional control turns the coefficient on beliefs, negative, as in Table 2, Model (2). We also examine the correlation between beliefs and appropriation, by each network and the results are reported in Table C3. We do not find a clear positive correlation, except in the undirected network (where the coefficient on beliefs is positive and significant at 5%). This departs from the robust correlations between beliefs and contributions documented in linear VCM PG experiments (e.g. in Fischbacher and Gächter, 2010; Smith, 2013). Given the differences between our design, and other studies, it is difficult to conclude why this is precisely the case. One possibility is that the greater complexity of the non-linear payoff space induces greater uncertainty about how subjects expect others to behave. In addition, we average beliefs from guesses for each of the three members (rather than the expected average of the others as in Fischbacher and Gächter (2010)) and include punishment.

negative and weakly significant (Table 3, model (1)), but with the inclusion of the lagged variables, these effects disappear (Table 3, model (2)).¹⁸

Next, we examine the impact of the global network properties of completeness, directedness and connectedness. Recall that any differences in outcomes between the complete and undirected circle yield the effect of completeness. Thus, we restrict the sample in Appendix Table C4 (models (1) and (2)), to the complete and undirected circle network, where the omitted category is the complete network. Controlling for session and round fixed-effects, and individual controls, we find that the dummy on the undirected circle is negative, but not significant, suggesting completeness has no effect on appropriations. Similarly, the effect of directedness is obtained by comparing differences in outcomes between the undirected and directed circle (Table C5, models (1) and (2)). The effect of connectedness by comparing the differences between the directed circle and directed line (Table C6, models (1) and (2)).¹⁹ Again, we find no effect due to both connectedness and directedness. We summarize these findings in our first result:

Result 1: Average appropriation is similar across the complete, undirected circle, directed circle and directed line networks, and does not differ by network completeness, connectedness and directedness or node degree.

Overall, these results support previous findings that average contributions are not statistically different between the complete, undirected circle and directed circle networks. On the other hand, we do not find that cooperation is lower due to directedness, as in Carpenter et al., (2012). As this study differs from previous experiments in numerous ways, a discussion of why the results may differ are somewhat speculative. For instance, we also elicit beliefs, which previous work has found to have either no effect on cooperation (Wilcox and Feltovich, 2000), a negative effect (Croson, 2000) or a positive effect (Gächter and Renner, 2000). How belief elicitation affects appropriation across networks in this setting in an open question. Another possibility for why appropriation is similar across networks, is that free-riding is more

¹⁸ In model (2) in tables 2 and 3, the coefficient on females (relative to males) is negative and weakly significant at around a half-token, controlling for network and other covariates, and session and round dummies.

¹⁹ In addition, we also examine if appropriation varies by nodes labelled A, B, C and D within each network and the results are reported in Tables C7-C10. Appropriation should not vary between nodes in the symmetric networks (i.e., complete, undirected circle and directed circle). This is empirically supported in the complete and directed circle networks. But node A appropriates less than B, C and D in the Undirected circle network (Table C8). This is a puzzling as every subject was randomly allocated, and the label itself should have no systematic effect on behaviour (we do not observe in other groups). Thus, we suspect this may be due to chance. We also find no effect of the node type (i.e. A, B, C, D) on appropriation in the directed line network.

aggressive in non-linear CPR dilemmas, and punishment requires a longer to take effect (Cason and Gangadharan, 2015). We return to this point briefly in the following section.

Punishment

Table 1 shows that received (mean) punishment is the highest in the complete network (11.69 points), especially compared to the directed circle and line networks (3.05 and 3.54 points respectively). From Figure 4, we see that punishment is higher and more erratic in the complete network, in line with the previous CPR literature (the undirected circle network also records higher and somewhat erratic punishment). Figure 5 also shows that nodes in the complete and undirected circle networks receive more punishment on average. Notably, received punishment seems to be marginally declining in the last period (rounds 11-15) in the complete network, directed circle and directed line.

More broadly, this suggests that higher punishment opportunities offered in the complete and undirected circle networks (i.e. the total number of punishment opportunities to each subject and by each round), elicits more positive punishment decisions and at higher levels. Indeed, Figure 6 shows subjects received zero punishment points 54.4% and 56.3% of the times in the complete and undirected circle networks respectively, compared to 85.8% and 80.95% in the undirected circle and line networks (68.1% for the pooled sample). Similarly, received punishment was over 15 points for 25% of the punishment opportunities in the complete network, versus 16.3%, 7% and 9.4% in the undirected circle, directed circle and directed line networks respectively.

Table 2 provides formal support that the complete network elicits higher punishment than the directed circle and line networks. The coefficients in model (3) suggests that controlling for other variables (rounds, sessions and individual controls), received punishment is on average around five to four points lower in the directed circle and line networks (significant at 1%). Table 3, model (3) also provides support that nodes in the directed networks receive less punishment than those in the complete network.

We also consider if punishment severity varies by network, i.e., received punishment divided by the number of punishment opportunities available to all nodes in each round within each network (i.e., 12 opportunities in complete network, 8 in undirected circle network, 4 in the directed circle and 3 in the directed line). From Table 2, Model (4), punishment severity is higher in the undirected circle relative to the complete network (significant at 5%). Similarly,

when we consider punishment severity across nodes, the coefficients on N122 and N311 are positive and significant at 5%.

Network completeness, directedness and connectedness also impact punishment behaviour. In Table C4, we examine the effect of completeness in models (3) and (4), to find that received punishment is lower in the undirected circle, relative to complete networks (at 5%), but punishment severity is not significantly different. Model (3) in Table C5 reveals the effect of directedness, and shows that received punishment is lower in the directed line network relative to the undirected circle network (significant at 1%). Model (4) also shows punishment severity is marginally lower (weakly significant at 10%). Finally, disconnectedness is associated with higher punishment, but this relationship is also not as robust: from Table C6, model (3), the coefficient on punishment is positive and significant at 10%. The result on average punishment by network property is summarized as follows:

Result 2: The complete network elicits the highest levels of received punishment. Completeness is associated with higher punishment, directedness with lower punishment.

Our results broadly correspond to Leibbrandt et al., (2015), that sanctioning is higher in the complete network, versus an incomplete where only two subjects have sanctioning opportunities, but differ from Carpenter et al., (2012) who find that punishment levels are amongst the lowest in the complete network. As previously noted, appropriation is more aggressive, and free-riding takes longer to be identified and punished in non-linear settings. Given this, a more extended period of play may be necessary to realise the long-run benefit of punishment (including the subsequent fall in received punishment and appropriations; Gächter et al., 2008). One possibility is that learning may be faster in the complete network, as it offers higher feedback and equal punishment opportunities. Indeed, we see that received punishment is marginally lower in the last period in Figure 5.²⁰ This suggests that lower punishment could be realised more quickly in the complete network, over a more rounds of play. However, we do not have sufficient data to conclude that is the case, at present.

Now, we proceed to analyse individual level punishment behaviour, to examine if received punishment depends on whether subjects are free-riders or high appropriators. We measure ‘free-riding’ as the absolute difference between the subject’s appropriation and the Pareto

²⁰ A Pearson χ^2 test shows that the difference between average punishment between period two (rounds 6-10) and period three (rounds 11-15) is weakly lower, with $\chi^2 = 44.75$, p-value = 0.083 in the complete network (the corresponding difference is not significant in the undirected circle network, $\chi^2 = 17.19$ and p-value = 0.143).

optimal appropriation of 10 tokens. We use the absolute deviation from the average appropriation of the others in the group, to identify high-appropriators, following Fehr and Gächter (2000). In both cases, an absolute positive difference denotes that the agent has appropriated a higher amount, and is, therefore, a free-rider. Figures 7 and 8 present average punishment received by absolute deviation from the Pareto optimal appropriation and other's (mean) appropriation respectively. Some key facts emerge that are consistent across both figures.

Those who appropriate five tokens and above, from either the Pareto optimal appropriation (the Nash appropriation is 16 tokens) or other's average appropriation, are punished more heavily. As high appropriators are punished even at a personal cost, this is typically called pro-social (or altruistic) punishment in the literature (Fehr and Gächter 2002). We also see anti-social punishment, albeit at lower levels, i.e., punishment targeting those who extract at or around other's (mean) appropriation level (or at the Pareto equilibrium) (Herrmann et al., 2008). This is in line with numerous PG and CPR experiments which report the presence of anti-social punishment, for reasons such as revenge, spite, confusion or to even discourage cooperation (Gintis, 2008; Nikiforakis, 2008; Casari and Plott, 2003; Ostrom et al., 1994).

We also observe that punishment levels are on average higher in the complete and undirected circle network, which is consistent with result two. Notably, the levels of both pro-social and anti-social punishment are higher in the complete network. This finding is supported by the results of the random-effects regression models, in Table 5, which presents the determinants of received punishment, for each network sample. As before, we control for session and round dummies, and individual controls. We also include stage one payoffs as an additional control because the level of received punishment is bounded by earnings from stage one (the results are robust to omitting this variable as well).

First, we consider the complete network. Models (1) and (2) demonstrate that free-riding, measured by positive deviations from the Pareto allocation and Other's (mean) appropriation respectively, receive higher punishment (significant at 1%). In model (1), the coefficient on negative deviations from the Pareto appropriation is positive but insignificant. But in model (2), negative deviations from other's appropriation also receive higher punishment (significant at 1%) – confirming the robust presence of anti-social punishment in the complete network. Models (3) and (4) also confirm the robust presence of pro-social punishment in the undirected circle (in fact, both coefficients are marginally higher). We also see negative deviations from

other's average appropriations, receive higher punishment (significant at 1%). But unlike the complete network, negative deviations from the Pareto, receive lower punishment (significant at 5%).

Model (5), which considers the directed circle network, again reveals that deviations above and below the Pareto optimal allocation receive higher and lower punishment (significant at 1%, and 5% respectively). However, from Model (6), only the coefficient on negative deviations from other's appropriation is positive and significant at 1%; the coefficient on the negative deviation from other's appropriation is positive but insignificant. In the directed line, we also see that pro-social punishment is positive and significant (at 1%). But the coefficient on negative deviations from the Pareto is negative but not significantly different, and negative deviations from other's appropriation is positive and weakly significant at 10% (models (7) and (8) respectively; we also control for the node degree).

More broadly, the analysis so far suggests that different networks elicit specific patterns of punishment behaviour, apart from increasing the overall punishment capacity available to agents. While, pro-social punishment is a robust feature across networks, we find that the complete and undirected circle networks (which are also undirected) elicits systematic anti-social punishment. These findings are summarized in our third result:

Result 3: High appropriators and free-riders are punished in all networks, with pro-social punishment. There is also systematic anti-social punishment, especially in the complete and undirected circle networks. Those who extract less than the Pareto efficient allocation in connected networks, i.e., the undirected and directed circles, receive lower punishment.

Payoffs and efficiency

We then compare the level of efficiency (average payoffs) across networks. Figures 9 and 10 plot the average payoffs from stage two across networks and nodes, and we see that average payoffs start above the Nash equilibrium level, but trends downwards. Table 1 also shows that average payoffs are the lowest in the complete network. This finding is formally supported by the random-effects regressions in Table 2, Model (5), which show that earnings are 17-19 tokens lower in incomplete networks (significant at 1%), controlling for session dummies, round dummies and individual controls. This is also confirmed by regression results across different nodes (Table 3, model (5)). We find that completeness is associated with lower payoffs, but that payoffs are not different due to directedness or connectedness (Tables C4, C5 and C6). As previously discussed, this result is driven by high levels of costly pro-social and

anti-social punishment, and rather limited differences in appropriation across networks. To summarise:

Result 4: The complete network is the least efficient network architecture, and completeness is associated with lower average payoffs.

Beliefs over other's appropriations

We start with discussing if beliefs over other's appropriation from the common account, vary across networks and by network property. Figures 11 and 12 disclose that stated beliefs over other's appropriation start at the Pareto appropriation and trends towards the Nash equilibrium, across networks and nodes. From Table 1, average beliefs tend to be marginally higher in complete networks compared to the undirected circle and directed circle networks.

Table 2, model (6) provides formal support for this when we regress the network dummies on stated beliefs, controlling for session and round dummies, and individual controls. More precisely, stated beliefs is around 2-3 tokens lower in all incomplete networks relative to the complete network (significant at 1%). This finding corresponds with Neugebauer et al. (2009), where subjects report lower beliefs with greater information feedback. In model (7) we add lagged beliefs and other's appropriation in the previous period (following Fischbacher and Gächter, 2010). The coefficient on lagged beliefs is positively and significantly associated with beliefs (at 1%), but that of other's appropriation is positive, but the difference is not statistically significant. The coefficients on the network dummies fall but remain significant. These results are also qualitatively similar to those in models (6) and (7) in Table 3, which assess the impact of node dummies.²¹ Notably, node N301, who receives no feedback on other's appropriation stated the lowest beliefs relative to nodes in the complete network; by a magnitude of around 1.6 tokens (significant at 1%; also see Figure 12). From Tables C4 and C5, directedness is negatively associated with higher beliefs, but the difference is not significant, and disconnectedness is associated with lower beliefs (significant at 5%, by around one token). This yields the following result:

Result 5: Stated beliefs over other's appropriations are the highest in the complete network. Completeness and disconnectedness are associated with higher beliefs. Node N301 states lower beliefs.

²¹ Notably, economics students state higher beliefs (by a magnitude of 1-0.66, significant at 1%) relative to non-economics students in Tables 3 and 4.

Next, we consider the deviation between stated beliefs and other's actual appropriation in more detail. Figures 13 and 14 plot the average deviation between beliefs and other's appropriation, by treatment and node, over time. There is a persistent gap between beliefs and other's mean appropriations, suggesting subjects hold 'optimistic' beliefs. The absolute difference is two tokens on average for the pooled sample. Keeping in mind differences between the current experiment and previous studies, the difference is higher than observed in Velez et al. (2009) who record that individuals expected other members would extract a nearly one unit less than their actual appropriation (non-incentivised total appropriation from others is elicited in a linear CPR game). It is relatively closer to Gächter and Renner (2008) who record that the absolute difference between stated belief and actual average contribution of others is around two tokens (incentivised average contributions of other is elicited in a linear VCM PG game).

Figures 13 and 14 also reveal some differences across nodes and networks. Firstly, the absolute difference between beliefs and other's appropriation systematically falls with time in the complete network and directed line networks. On the other hand, beliefs in the undirected circle and directed circle networks seem to marginally increase in rounds 6-10, and then fall. Lastly, if we consider only node N301 who receives no feedback, the divergence is the highest in the first period, but drops over time.

We examine if differences across nodes are statistically significant by using random-effects models, controlling for sessions fixed-effects and individual controls.²² We add period dummies, where period one (rounds 1-5) is the omitted category. The results are reported in Table 6. Model (1) shows that N301 reports the higher divergence between stated beliefs and appropriation, relative to nodes in the complete network (N033 is the omitted category). The period dummy for rounds 11-15 is negative and significant at 5%, suggesting that controlling for other variables, the deviation declines over time (or equally, belief accuracy increases with time).

To examine if this decline is contingent on the node degree, we interact the node degree dummies with the three-period dummy in model (2). First, the value of the coefficients on N301 and dummy for rounds 11-15, marginally increases and is significant at 1% each. Second,

²² We focus on differences across nodes rather than networks in the text given the unique position of N301. Analogous results examining effects across networks are available in Table C10. The results are qualitatively similar: we find that the coefficients on the undirected circle, and directed circle and line networks are not significant, but the second-period dummy for rounds 11-15 is significant. The interaction terms on the period and undirected and directed network dummies are qualitatively similar.

the interaction term between nodes in the undirected circle and directed circle networks (N122 and N311), and the second category of the period dummy (rounds 11-15) have a positive and significant coefficient (at 5%). This reveals that the nodes in the undirected circle and directed circle networks maintain relatively more optimistic beliefs in the last period i.e., the difference between beliefs and other's appropriation is higher compared to nodes in the complete network in the first period (rounds 1-5), controlling for other variables. This suggests that nodes in incomplete but connected networks, maintain optimistic beliefs for longer, relative to nodes in the complete network.²³ More broadly, this confirms that the network structure impacts learning through time, and this yields our last result:

Result 6: The difference between beliefs and other's actual appropriation is highest for node N301. While the deviation between beliefs and other's (mean) appropriation falls over time, trends vary by the network. Relative to the complete network, the deviation between beliefs and other's (mean) appropriation is higher in the directed circle and directed circle networks, in the last period.

DISCUSSION AND CONCLUSION

This study contributes to the literature on peer monitoring and punishment in social dilemmas, by considering how the network architecture impacts outcomes in a non-linear CPR appropriation game. Broadly, we find evidence that the structure of the network, impacts efficiency, punishment, and stated beliefs over other's appropriation, but that it does not impact appropriation significantly. In particular, our results suggest the network structure impacts both punishment opportunity and the type of punishment used: the excessive use of pro-social and anti-social punishment reduces payoffs in the complete network, making it the most inefficient network. Furthermore, we find that there is faster convergence between stated beliefs and other's appropriation in the complete network, relative to incomplete, but connected networks.

Overall, these findings may hold some potential lessons for policy, keeping in mind the many differences between the field and lab setting. Firstly, it highlights that the structural variation in the distribution of monitoring and punishment opportunities impacts welfare. In fact, the excessive use of anti-social punishment in the complete network provides a cautionary lesson

²³ Neugebauer et al., (2009) do not find that difference between beliefs and other's contributions declines over ten rounds, in groups with and without feedback. We use fifteen rounds to allow subjects greater time to learn the game and modify their behaviour, and find that average differences are more apparent in the last periods during rounds 11-15.

about the cost of exclusively relying on perfect peer monitoring and punishment to enhance cooperation, especially if the primary policy objective is welfare maximisation. Second, it highlights the limited impact of the network on appropriation – suggesting the need for investigation into complementary mechanisms to enhance the effectiveness of peer monitoring, to prevent the destruction of the commons. Some of these may be especially relevant from the network perspective: for instance, the effect of restricting communication opportunities or allowing subjects to ostracise high appropriators, by severing network edges. Third, it is possible that long-run benefits of punishment maybe realised in systematically different ways across network architectures, as they all offer different opportunities for learning and feedback. Some natural extensions to this research are therefore examining how behaviour evolves over a larger number of rounds, as well as allowing for variations in the experimental design to allow for partner-matching, and the systematic examination of beliefs evolve and impact behaviour in more complex, non-linear social dilemmas.

REFERENCES

- Anderies, John M., Marco A. Janssen, François Bousquet, Juan-Camilo Cardenas, Daniel Castillo, Maria-Claudio Lopez, Robert Tobias, Björn Vollan, and Amber Wutich. "The challenge of understanding decisions in experimental studies of common pool resource governance." *Ecological Economics* 70, no. 9 (2011): 1571-1579.
- Apestequia, Jose, and Frank P. Maier-Rigaud. "The role of rivalry: public goods versus common-pool resources." *Journal of Conflict Resolution* 50, no. 5 (2006): 646-663.
- Bodin, Örjan, and Beatrice I. Crona. "Management of natural resources at the community level: exploring the role of social capital and leadership in a rural fishing community." *World development* 36, no. 12 (2008): 2763-2779.
- Bodin, Örjan, and Beatrice I. Crona. "The role of social networks in natural resource governance: What relational patterns make a difference?." *Global environmental change* 19, no. 3 (2009): 366-374.
- Boosey, Luke, and R. Mark Isaac. "Asymmetric network monitoring and punishment in public goods experiments." *Journal of Economic Behavior & Organization* 132 (2016): 26-41.
- Cardenas, Juan-Camilo, T. K. Ahn, and Elinor Ostrom. "Communication and co-operation in a common-pool resource dilemma: a field experiment." In *Advances in Understanding Strategic Behaviour*, pp. 258-286. Palgrave Macmillan UK, 2004.
- Carpenter, Jeffrey, Shachar Kariv, and Andrew Schotter. "Network architecture, cooperation and punishment in public good experiments." *Review of Economic Design* (2012): 1-26.
- Casari, Marco, and Charles R. Plott. "Decentralized management of common property resources: experiments with a centuries-old institution." *Journal of Economic Behavior & Organization* 51, no. 2 (2003): 217-247.
- Cason, Timothy N., and Lata Gangadharan. "Promoting cooperation in nonlinear social dilemmas through peer punishment." *Experimental Economics* 18, no. 1 (2015): 66-88.
- Cason, Timothy N., and Lata Gangadharan. "Swords without covenants do not lead to self-governance." *Journal of Theoretical Politics* 28, no. 1 (2016): 44-73.
- Chaudhuri, Ananish. "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature." *Experimental Economics* 14, no. 1 (2011): 47-83.
- Choi, Syngjoo, Douglas Gale, and Shachar Kariv. "Social learning in networks: a quantal response equilibrium analysis of experimental data." *Review of Economic Design* (2012): 1-23.
- Choi, Syngjoo, Shachar Kariv, and Edoardo Gallo. "Networks in the Laboratory." In *The Oxford Handbook of the Economics of Networks*. : Oxford University Press, 2016-04-14.
- Croson, Rachel TA. "Thinking like a game theorist: factors affecting the frequency of equilibrium play." *Journal of economic behavior & organization* 41, no. 3 (2000): 299-314.
- Fehr, Ernst, and Simon Gächter. "Cooperation and punishment in public goods experiments." *American Economic Review* 90, no. 4 (2000): 980-994.

- Fischbacher, Urs, and Simon Gächter. "Social preferences, beliefs, and the dynamics of free riding in public goods experiments." *The American economic review* 100, no. 1 (2010): 541-556.
- Fischbacher, Urs. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental economics* 10, no. 2 (2007): 171-178.
- Gächter, Simon, and Elke Renner. "The effects of (incentivized) belief elicitation in public goods experiments." *Experimental Economics* 13, no. 3 (2010): 364-377.
- Gächter, Simon, Elke Renner, and Martin Sefton. "The long-run benefits of punishment." *Science* 322, no. 5907 (2008): 1510-1510.
- Gale, Douglas, and Shachar Kariv. "Bayesian learning in social networks." *Games and Economic Behavior* 45, no. 2 (2003): 329-346.
- Gintis, Herbert. "Punishment and cooperation." *Science* 319, no. 5868 (2008): 1345-1346.
- Herrmann, Benedikt, Christian Thöni, and Simon Gächter. "Antisocial punishment across societies." *Science* 319, no. 5868 (2008): 1362-1367.
- Janssen, Marco A., Robert Holahan, Allen Lee, and Elinor Ostrom. "Lab experiments for the study of social-ecological systems." *Science* 328, no. 5978 (2010): 613-617.
- King, A., 2000. *Managing without institutions: The role of communication networks in governing resource access and control* (Doctoral dissertation, University of Warwick).
- Kingsley, David C. "Peer punishment across payoff equivalent public good and common pool resource experiments." *Journal of the Economic Science Association* 1, no. 2 (2015): 197-204.
- Kingsley, David C., and Benyuan Liu. "Cooperation across payoff equivalent public good and common pool resource experiments." *Journal of Behavioral and Experimental Economics* 51 (2014): 79-84.
- Kosfeld, Michael. "Economic Networks in the Laboratory: A Survey." In *Review of Network Economics*. 2004.
- Ledyard, J. "0. 1995. Public goods: A survey of experimental research. JH Kagel, AE Roth, eds., *The Handbook of Experimental Economics*." 111-194.
- Leibbrandt, Andreas, Abhijit Ramalingam, Lauri Sääksvuori, and James M. Walker. "Incomplete punishment networks in public goods games: experimental evidence." *Experimental Economics* 18, no. 1 (2015): 15-37.
- Neugebauer, Tibor, Javier Perote, Ulrich Schmidt, and Malte Loos. "Selfish-biased conditional cooperation: On the decline of contributions in repeated public goods experiments." *Journal of Economic Psychology* 30, no. 1 (2009): 52-60.
- Nikiforakis, Nikos. "Punishment and counter-punishment in public good games: Can we really govern ourselves?." *Journal of Public Economics* 92, no. 1 (2008): 91-112.
- Noussair, Charles N., and Daan P. van Soest. "Economic experiments and environmental policy." *Annu. Rev. Resour. Econ.* 6, no. 1 (2014): 319-337.

- Ostrom, Elinor, James Walker, and Roy Gardner. "Covenants with and without a sword: Self-governance is possible." *American political science Review* 86, no. 2 (1992): 404-417.
- Ostrom, Elinor, Roy Gardner, and James Walker. *Rules, games, and common-pool resources*. University of Michigan Press, 1994.
- Ostrom, Elinor. "The value-added of laboratory experiments for the study of institutions and common-pool resources." *Journal of Economic Behavior & Organization* 61, no. 2 (2006): 149-163.
- Poteete, Amy R., Marco A. Janssen, and Elinor Ostrom. *Working together: collective action, the commons, and multiple methods in practice*. Princeton University Press, 2010.
- Smith, Alexander. "Estimating the causal effect of beliefs on contributions in repeated public good games." *Experimental Economics* 16, no. 3 (2013): 414-425.
- Sturm, Bodo, and Joachim Weimann. "Experiments in environmental economics and some close relatives." *Journal of Economic Surveys* 20, no. 3 (2006): 419-457.
- Van Soest, Daan, and Jana Vyrastekova. "Peer enforcement in CPR experiments: the relative effectiveness of sanctions and transfer rewards, and the role of behavioural types." *Using experimental methods in environmental and resource economics* (2007): 113.
- Velez, Maria Alejandra, John K. Stranlund, and James J. Murphy. "What motivates common pool resource users? Experimental evidence from the field." *Journal of Economic Behavior & Organization* 70, no. 3 (2009): 485-497.
- Vyrastekova, Jana, and Daan Van Soest. "On the (in) effectiveness of rewards in sustaining cooperation." *Experimental Economics* 11, no. 1 (2008): 53-65.
- Wilcox, Nathaniel, and Nick Feltovich. *Thinking Like a Game Theorist: Comment*. Monash University, Department of Economics, 2000. Available at <http://www.uh.edu/~nfelt/papers/belief.pdf>.

TABLES

Table 1: Descriptive statistics by network and node

Network/Node	Sessions	Subjects	Groups	Observations	Appropriation	Punishment	Beliefs	Payoffs
Complete	3	60	225	900	15.37 (5.56)	11.69 (23.78)	13.25 (4.84)	172.36 (51.51)
Undirected circle	2	36	135	540	14.51 (4.99)	6.77 (14.30)	12.69 (3.75)	188.77 (37.21)
Directed circle	2	40	150	600	14.80 (5.29)	2.61 (11.04)	12.26 (3.85)	189.26 (40.34)
Directed line	3	56	210	840	14.81 (6.23)	3.05 (8.99)	12.95 (5.25)	188.71 (48.87)
N311	3	28	210	420	15.01 (6.19)	3.54 (9.92)	13.25 (5.05)	187.94 (48.71)
N301	3	14	210	210	14.37 (5.92)	-	13.37 (5.36)	190.39 (49.52)
N310	3	14	210	210	14.83 (6.61)	2.09 (6.68)	11.92 (5.43)	188.57 (48.73)
Total	10	192	720	2880	14.93 (5.62)	6.61 (17.13)	12.85 (4.61)	183.73 (46.70)

Notes: Mean (standard deviation) is reported for appropriation, received punishment, beliefs and payoffs. Subjects are assigned to one network and to a fixed node type of A, B, C or D, in each session (having 16-20 subjects), and in each of the 15 rounds, they are randomly assigned to a new group.

Table 2: Outcomes across networks

Outcome: Random-effects models:	Appropriation (1)	Appropriation (2)	Received punishment (3)	Punishment severity (4)	Payoffs (5)	Beliefs (6)	Beliefs (7)
Network = 1, Undirected circle	-2.112* (1.196)	-0.937 (0.644)	1.892 (2.325)	2.277** (1.032)	17.377*** (6.147)	-1.809** (0.709)	-0.949*** (0.354)
Network = 2, Directed circle	-1.251 (1.141)	-0.342 (0.606)	-5.229*** (1.860)	0.178 (1.178)	17.014*** (5.854)	-2.740*** (0.811)	-1.338*** (0.397)
Network = 3, Directed line	-2.059 (1.432)	-0.987 (0.820)	-3.761** (1.656)	1.405 (0.860)	19.711*** (6.805)	-3.327*** (0.946)	-1.814*** (0.478)
Beliefs		-0.015 (0.044)					
Appropriation (t-1)		0.454*** (0.039)					
Beliefs (t-1)		0.075** (0.038)					0.528*** (0.033)
Received punishment (t-1)		-0.019*** (0.006)					
Other's (mean) appropriation (t-1)							0.028 (0.023)
Sex	-1.042** (0.518)	-0.520* (0.285)	0.874 (0.724)	0.590 (0.373)	-0.695 (2.258)	-0.408 (0.403)	-0.249 (0.206)
Experience	-0.278 (0.301)	-0.063 (0.174)	-0.236 (0.448)	-0.108 (0.256)	-0.765 (1.499)	-0.405 (0.265)	-0.215* (0.131)
Economics	0.191 (0.649)	-0.099 (0.342)	-1.459** (0.691)	-0.591 (0.401)	6.071** (2.720)	1.403*** (0.417)	0.661*** (0.209)
Constant	13.440*** (1.235)	8.743*** (0.944)	6.114*** (1.886)	1.874** (0.874)	196.595*** (7.334)	11.780*** (0.825)	6.671*** (0.694)
Observations	2,880	2,492	2,670	2,670	2,880	2,880	2,513
Number of subject	192	178	178	178	192	192	192
Session dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Round dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Robust standard errors are clustered at the subject level, and *** p<0.01, ** p<0.05, * p<0.1. The Complete network is the omitted category.

Table 3: Outcomes across nodes

Outcome:	Appropriation	Appropriation	Received punishment	Punishment severity	Payoffs	Beliefs	Beliefs
Random-effects models:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Node degree = 1, N122	-2.124*	-0.952	1.863	2.247**	17.375***	-1.847***	-0.970***
	(1.196)	(0.644)	(2.321)	(1.027)	(6.153)	(0.713)	(0.358)
Node degree = 2, N211	-1.239	-0.344	-5.213***	0.194	16.986***	-2.729***	-1.343***
	(1.139)	(0.605)	(1.858)	(1.177)	(5.857)	(0.812)	(0.400)
Node degree = 3, N311	-1.688	-0.852	-3.344**	1.841**	18.936**	-2.970***	-1.691***
	(1.605)	(0.869)	(1.649)	(0.836)	(7.524)	(1.100)	(0.537)
Node degree = 5, N310	-2.359	-1.292	-4.605***	0.524	19.380**	-4.548***	-2.310***
	(1.793)	(0.932)	(1.694)	(0.916)	(7.664)	(1.071)	(0.639)
Beliefs		-0.016					
		(0.044)					
Appropriation (t-1)		0.454***					
		(0.039)					
Beliefs (t-1)		0.074*					0.525***
		(0.038)					(0.034)
Received punishment (t-1)		-0.019***					
		(0.006)					
Other's (mean) appropriation (t-1)							0.028
							(0.024)
Sex	-1.110**	-0.563*	0.753	0.464	-0.620	-0.538	-0.302
	(0.538)	(0.297)	(0.749)	(0.378)	(2.307)	(0.393)	(0.209)
Experience	-0.279	-0.069	-0.249	-0.122	-0.787	-0.429	-0.226**
	(0.304)	(0.174)	(0.449)	(0.254)	(1.507)	(0.267)	(0.132)
Economics	0.159	-0.112	-1.507**	-0.642	6.122**	1.357***	0.647***
	(0.660)	(0.346)	(0.696)	(0.404)	(2.747)	(0.425)	(0.213)
Node degree = 4, N301	-2.503*				21.553***	-2.866***	-1.627***
	(1.487)				(6.999)	(0.945)	(0.451)
Constant	13.488***	8.815***	6.219***	1.984**	196.577***	11.908***	6.763***
	(1.237)	(0.950)	(1.886)	(0.868)	(7.335)	(0.818)	(0.702)
Observations	2,880	2,492	2,670	2,670	2,880	2,880	2,513
Number of subject	192	178	178	178	192	192	192
Session dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Round dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Robust standard errors are clustered at the subject level, and *** p<0.01, ** p<0.05, * p<0.1. N033 is the omitted category.

Table 4: Received punishment

Networks: Random-effects models:	Complete (1)	Complete (2)	Undirected circle (3)	Undirected circle (4)	Directed circle (5)	Directed circle (6)	Directed line (7)	Directed line (8)
Node degree = 2							-1.414*	-1.327*
							(0.750)	(0.742)
Positive deviation-Pareto	1.286***		1.498***		0.369***		0.348***	
	(0.195)		(0.325)		(0.108)		(0.082)	
Negative deviation-Pareto	0.011		-0.655**		-0.475**		-0.107	
	(0.320)		(0.308)		(0.196)		(0.259)	
Other's (mean) appropriation		0.219		0.025		-0.161	0.086	0.140
		(0.417)		(0.271)		(0.162)	(0.246)	(0.157)
Positive deviation-Other's appropriation		1.406***		2.173***		0.420**		0.381***
		(0.254)		(0.534)		(0.177)		(0.110)
Negative deviation-Other's appropriation		0.839***		0.469***		0.137		0.196*
		(0.227)		(0.179)		(0.096)		(0.118)
Stage 1 payoff	-0.034	-0.108**	0.005	-0.114***	0.020	-0.021	-0.006	-0.027
	(0.023)	(0.048)	(0.022)	(0.040)	(0.020)	(0.022)	(0.027)	(0.022)
Constant	5.447	21.146	2.627	29.161***	-5.585	6.170	0.964	5.473
	(5.418)	(14.222)	(5.706)	(10.220)	(5.875)	(6.832)	(8.775)	(6.014)
Observations	900	900	540	540	600	600	630	630
Number of subject	60	60	36	36	40	40	42	42
Session dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Round dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Outcome variable: Received punishment; Omitted category in models (9) and (10) is Node N311. Robust standard errors clustered at subject level, and *** p<0.01, ** p<0.05, * p<0.1.


Table 5: Difference between beliefs and other's (mean) appropriation across nodes

Outcome: Random-effects models:	Deviation between beliefs and other's (mean) appropriation	
	(1)	(2)
Node degree = 1, N122	0.013 (0.776)	-0.715 (0.920)
Node degree = 2, N211	1.388* (0.824)	0.408 (0.885)
Node degree = 3, N311	1.032 (1.111)	0.678 (1.268)
Node degree = 4, N301	1.137 (0.997)	1.125 (1.636)
Node degree = 5, N310	2.662*** (1.015)	3.202*** (1.175)
Period = 1, 6-10	-0.293 (0.250)	-0.673 (0.435)
Period = 2, 11-15	-0.613** (0.268)	-1.293*** (0.349)
1.degree#1.pd, N122 # Rounds 6-10		0.897 (0.680)
1.degree#2.pd, N122 # Rounds 11-15		1.286** (0.638)
2.degree#1.pd, N211 # Rounds 6-10		1.290** (0.604)
2.degree#2.pd, N211 # Rounds 11-15		1.650** (0.658)
3.degree#1.pd, N311 # Rounds 6-10		0.254 (0.831)
3.degree#2.pd, N311 # Rounds 11-15		0.805 (0.921)
4.degree#1.pd, N310 # Rounds 6-10		-0.155 (1.446)
4.degree#2.pd, N310 # Rounds 11-15		0.193 (1.642)
5.degree#1.pd, N301 # Rounds 6-10		-1.127 (0.946)
5.degree#2.pd, N301 # Rounds 11-15		-0.492 (0.985)
Sex	0.525 (0.414)	0.525 (0.415)
Experience	0.426 (0.278)	0.426 (0.278)
Economics	-1.412*** (0.453)	-1.412*** (0.453)
Constant	1.241 (0.767)	1.595** (0.788)
Observations	2,880	2,880
Number of subject	192	192
Session dummies	Yes	Yes
Individual controls	Yes	Yes

Notes: Robust standard errors are clustered at the subject level, and *** p<0.01, ** p<0.05, * p<0.1. The Complete network and period 1 (rounds 1-5) are the omitted categories.

FIGURES


FIGURE 1: NETWORK TREATMENTS AND PROPERTIES



Panel 1A: Complete network [0]
(Baseline)

Global Properties:
Complete, Connected, Undirected


Node types: Symmetric,
A, B, C, D: N033



Panel 1B: Undirected circle network [1]

Global properties:
Incomplete, Connected, Undirected


Node types: Symmetric,
A, B, C, D: N122



Panel 1C: Directed circle network [2]

Global network properties:
Incomplete Connected, Directed

Node types: Symmetric,
A, B, C, D: N211



Panel 1D: Directed line network [3]

Global network properties:
Incomplete Disconnected, Directed

Node types: Asymmetric,
A: N301
B, C: N311
D: N310

FIGURE 2: APPROPRIATION BY NETWORK

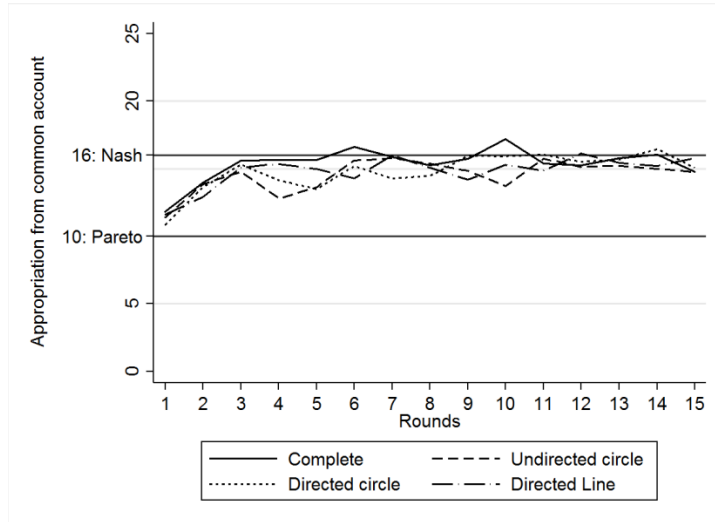


FIGURE 3: APPROPRIATION BY NODE

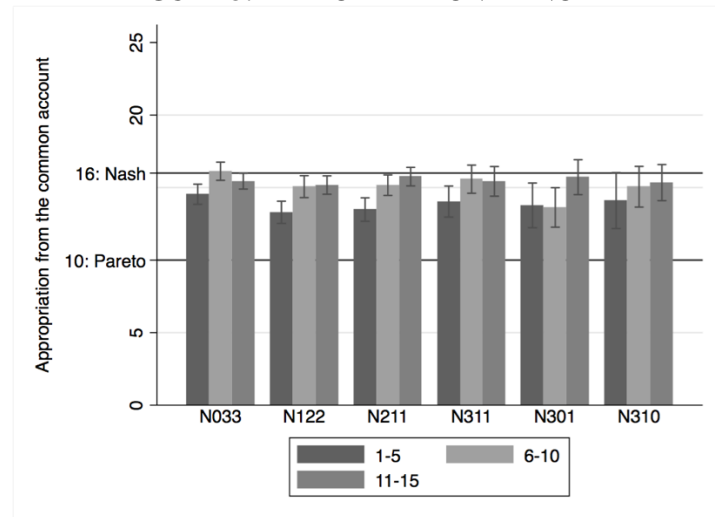


FIGURE 4: RECEIVED PUNISHMENT BY NETWORK




FIGURE 5: RECEIVED PUNISHMENT BY NODE

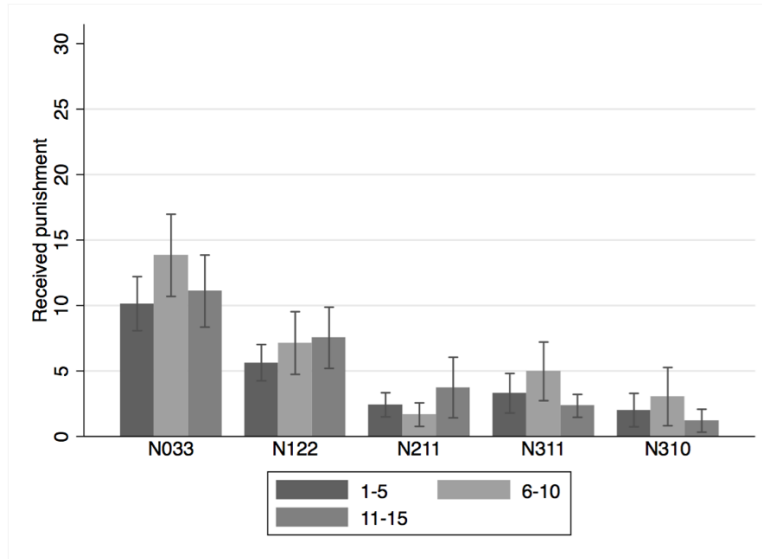


FIGURE 6: RECEIVED PUNISHMENT AS SHARE OF PUNISHMENT OPPORTUNITIES

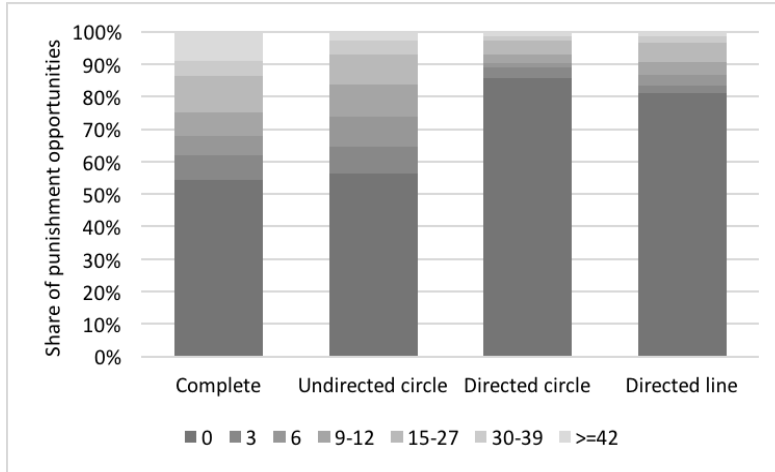


FIGURE 7: RECEIVED PUNISHMENT BY DEVIATION FROM PARETO EQUILIBRIUM AND NETWORK

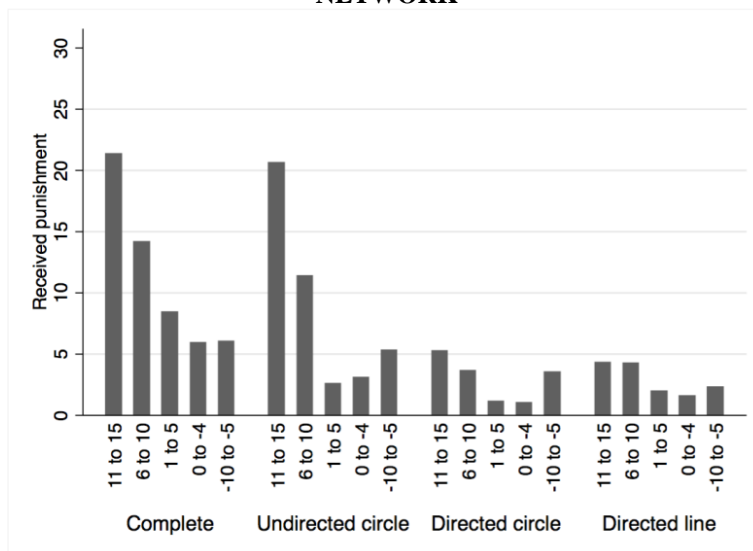


FIGURE 8: RECEIVED PUNISHMENT BY DEVIATION FROM OTHER'S (MEAN) APPROPRIATION AND NETWORK

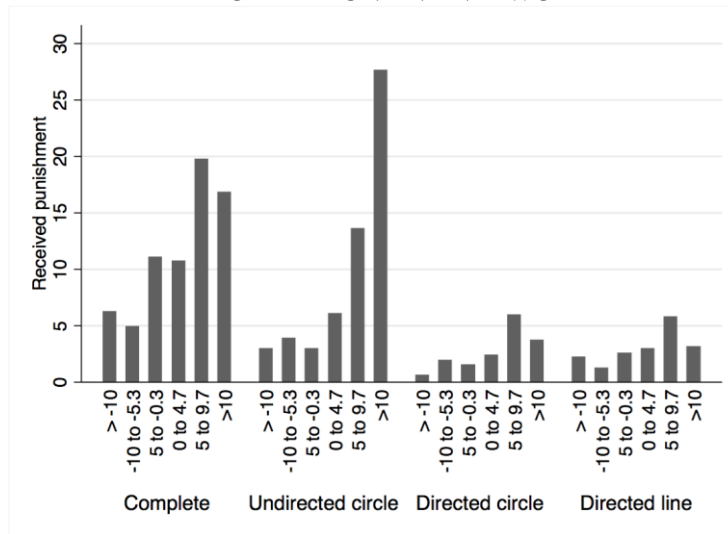


FIGURE 9: PAYOFFS BY NETWORK

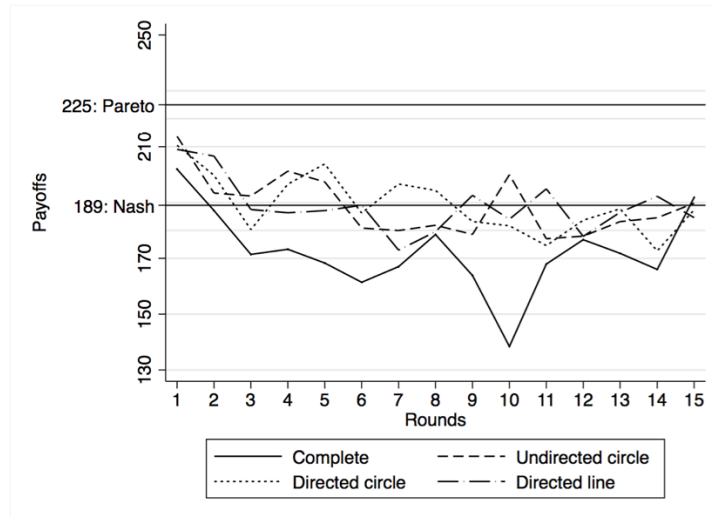


FIGURE 10: PAYOFFS BY NODE




FIGURE 11: BELIEFS ABOUT OTHER'S APPROPRIATION BY NETWORK

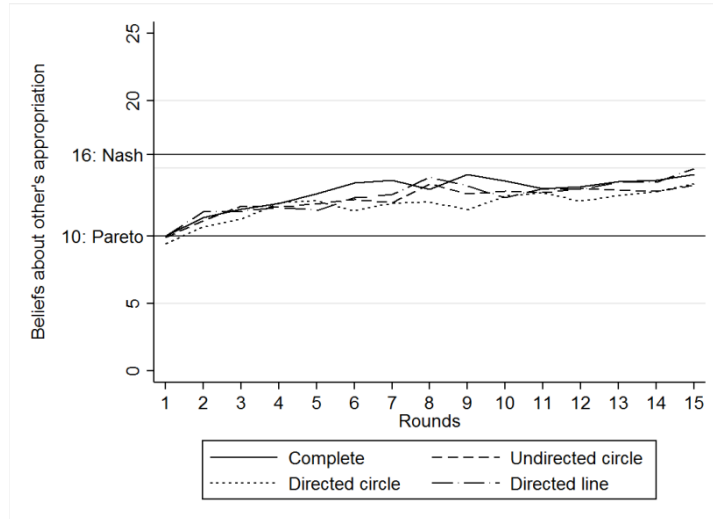


FIGURE 12: BELIEFS OVER OTHER'S APPROPRIATION BY NODE

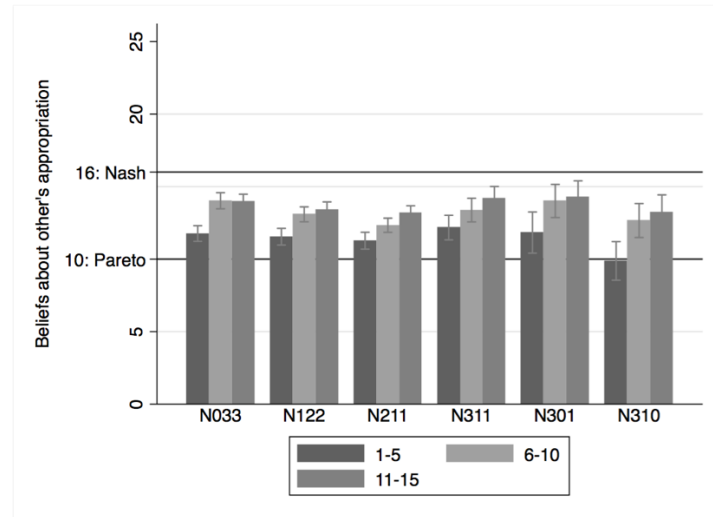


FIGURE 13: DEVIATION BETWEEN OTHER'S (MEAN) APPROPRIATION AND BELIEFS BY NETWORK

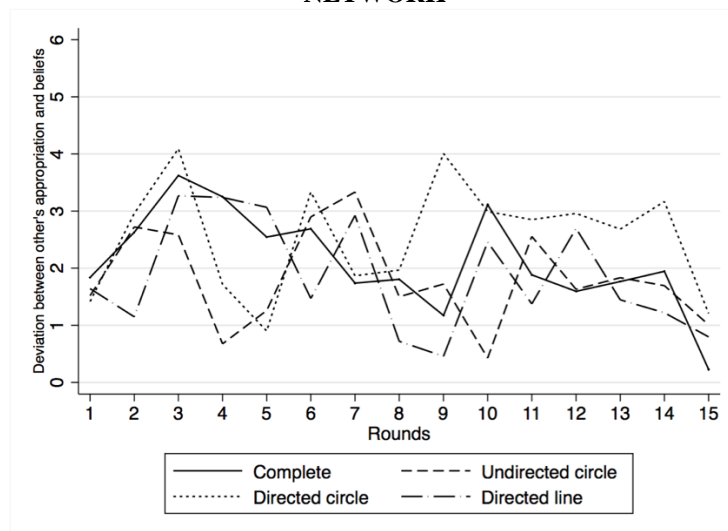


FIGURE 14: DEVIATION BETWEEN OTHER'S (MEAN) APPROPRIATION AND BELIEFS BY NODE

