Network Architecture, Cooperation and Punishment in Public Good Experiments*

Jeffrey Carpenter† Shachar Kariv‡ Andrew Schotter§
Middlebury College UC Berkeley New York University

May 1, 2012

Abstract

Following Fehr and Gäechter (2000), a large and growing number of experiments show that public goods can be provided at high levels when mutual monitoring and costly punishment are allowed. Nearly all experiments, however, study monitoring and punishment in a *complete network* where all subjects can monitor and punish each other. The architecture of social networks becomes important when subjects

*Some of the results reported here were previously distributed in a paper titled "Network Architecture and Mutual Monitoring in Public Goods Experiments." This research was supported by the Center for Experimental Social Sciences (CESS) at New York University and the UC Berkeley Experimental Social Science Laboratory (Xlab). We are grateful to Jim Andreoni, Boğaçhan Çelen, Andreas Fuster, Tom Palfrey, Matthew Rabin, and Bertil Tungodden for helpful discussions. The paper has benefited from suggestions by the participants of seminars at several universities. Kariv and Carpenter are grateful for the hospitality of the School of Social Science in the Institute for Advanced Studies at Princeton and Institute for the Study of Labor (IZA), respectively. Lastly, while we all share the credit that this research may be due, one of us does not share any blame because he had no idea that the paper had been submitted to this festschrift, commemorating his 65th birthday.

†Department of Economics, Middlebury College, Middlebury, VT 05753, USA (E-mail: jpc@middlebury.edu, URL: http://community.middlebury.edu/~jcarpent).

‡Department of Economics, University of California, Berkeley, 530 Evans Hall # 3880, Berkeley, CA 94720, USA (E-mail: kariv@berkeley.edu, URL: http://emlab.berkeley.edu/~kariv/).

§Department of Economics, New York University, 19 W. 4th Street, New York, NY, 10012, USA (E-mail: abdrew.schotter@nyu.edu, URL: http://homepages.nyu.edu/~as7).

can only monitor and punish the other subjects to whom they are connected by the network. We study several *incomplete networks* and find that they give rise to their own distinctive patterns of behavior. Nevertheless, a number of simple, yet fundamental, properties in graph theory allow us to interpret the variation in the patterns of behavior that arise in the laboratory and to explain the impact of network architecture on the efficiency and dynamics of the experimental outcomes.
JEL Classification Numbers: C91, C92, D62, D63, H41.
Key Words: networks, public goods, monitoring, costly punishment, experiment.

1 Introduction

A perennial question in economics concerns the conditions under which self-interested individuals cooperate to achieve socially efficient outcomes. In a seminal experiment, Fehr and Gächter (2000) show that public goods can be provided at high levels if subjects can monitor the contributions made by other subjects and punish those who are unwilling to contribute. This stands in stark contrast to the experimental results from the familiar public good game in the literature, in which low provision is common (Ledyard, 1995).

A number of experimental papers extend Fehr and Gächter (2000) by making punishment more or less costly to the monitor (Anderson and Putterman, 2005), making punishment only symbolic (Masclet et al., 2003), or by going in the opposite direction and equating punishment with expulsion from the group (Cinyabuguma et al., 2005). In all of these experiments, the authors also find high levels of provision.

A central assumption of nearly all experiments is *full monitoring* – everyone can monitor everybody else. In reality, individuals living in any society are bound together by a *social network*, and often they can only observe the behavior of those who are in their local environment. In public goods experiments, if each subject can observe the actions of only a small number of other subjects, it is not clear that mutual monitoring can give rise to cooperative outcomes. Clearly, *partial monitoring* can be an obstacle to cooperation if, for example, a critical mass of potential punishers is required to deter shirking or punishers are only emboldened to intervene when they know that they are supported by others.

The goal of this paper is to identify the impact of partial monitoring on the effectiveness of mutual monitoring and punishment. We represent the partial monitoring structure by a *graph* that specifies the monitoring technology of the group – that is, who monitors whose actions. Each subject is located at a node of the graph, and subject i can monitor subject j if and only if there is an edge leading from node i to node j. The experiments reported here involve the eight networks [1]-[8] illustrated in Figure 1 below. An arrow pointing from subject i to subject j indicates that i can monitor and punish j.

[Figure 1 here]

The set of networks depicted in Figure 1 has several interesting architectures exemplifying a number of simple yet fundamental concepts in graph theory (defined in the next section) that allow us to interpret variation in the experimental outcomes. Figure 1 also reports the total contribution rates, punishment levels and overall payoff efficiency (net of the costs of being punished and punishing) across networks. From these data we can immediately infer that there are significant differences in subjects' behavior across the different networks.

Our key results are as follows:

- **Cooperation** – Although contributions vary dramatically from network to network, *connected networks* [1]-[4], within which everyone is monitored, demonstrate significantly higher contributions than *disconnected networks* [5]-[8]; however, the complete network [1] does not elicit significantly more contributions than the other connected networks.

- **Punishment** – More punishment is used to maintain or increase contributions in *directed networks* [2] and [5]-[7] where the edges point in only one direction, relative to *undirected networks* [1], [3], [4] and [8]. Our conclusion is that the asymmetry in the relations between any pair of subjects in directed networks gives different monitoring roles to different subjects, which, in turn, increases punishment expenditures.

- **Efficiency** – While there is also considerable variation in net payoffs after subtractions for punishing others and for being punished across networks, the connected networks such as [1] and [4] are the most efficient, whereas the disconnected networks such as [6] and [7] are the least efficient. In addition, adding/removing edges from the graph does not necessarily increase/decease efficiency.

Our experiment demonstrates that network architecture has a significant impact on the behavior of subjects in public goods games and, therefore, on the outcomes achieved in different networks. While the major determinant of contribution levels appears to be whether or not all the subjects are connected, not all connected networks are equally efficient since they elicit different punishment behavior. The complete network [1], for example, yields high contribution levels and high efficiencies both because of and despite the fact that this architecture elicits less punishment than other networks, such as the directed circle [2], where punishment responsibilities are not shared. In short, while the previous literature was correct in pointing out that punishment may increase contributions, it failed to investigate the subtle relationship between network structure and performance. This paper has taken a first step in that direction

Among our other conclusions, the fact that we find it is necessary to take into account the details of the local neighborhood as well as the entire network architecture to explain individual behavior is particularly relevant for future work. The simple summary characteristics of the networks depicted in Figure 1, such as the average distance between subjects, do not fully account for the subtle and complicated behaviors that we observe. To determine the important determinants of individual behavior, it will be necessary to investigate a larger class of networks in the laboratory. This is perhaps one of the most important topics for future research.

The rest of the paper is organized as follows. Section 2 describes the margins along which we extend the previous literature. Section 3 both presents the network concepts that guided our experimental design and the design itself. Section 4 summarizes questions and hypotheses that can be addressed using the experimental data. Section 5 provides the empirical analysis. Section 6 discusses the results and provides some concluding thoughts.

2 Closely related literature

The papers most closely related to ours, Carpenter (2007) and O'Gorman et al. (2009), are those that also allow for costly punishment. Carpenter (2007) compared the complete network [1] to two other connected networks, the directed circle [2] and the undirected circle [3]. He found that contributions in the complete network are as high as in the undirected circle but are significantly higher than in the directed circle; however, the number of other potential punishers, not the network structure, was emphasized.

O'Gorman et al. (2009) compare the complete network [1] to a version of the directed star [5] in which the subjects in the center of the star are changed randomly each round and found that the complete network is less efficient. These results, though based on different designs and reporting some difference in outcomes are important primarily because they suggest that the network architecture affects the provision of the public good. Nevertheless, a more sophisticated and comprehensive analysis is required to detect the underlying properties of networks that facilitate or hinder cooperative outcomes.

Our paper contributes to the large body of experimental work on public goods, but we will not attempt to review this literature. We also contribute to the growing literature on the economics of networks (Jackson, 2008) and to the smaller, more recent, literature on network experiments (Kosfeld, 2004, provides an excellent, if now already somewhat dated, survey).

3 Networks, theory and design

The public good game that we study can be interpreted as follows: In the first stage, subjects simultaneously make voluntary contributions to a public good. The payoff for each subject at the first stage equals her consumption of the public good plus her remaining endowment. At the end of the first stage, subjects are allowed to monitor the contributions of the subjects to whom they are connected by the social network. Thus, we drop the standard assumption that individual contributions are public information and assume that subjects can monitor the contributions of some, but not necessarily all, of the other subjects. In the second stage, subjects are given the opportunity to punish, at some cost, the other subjects to whom they are connected by the network. The terminal payoff for each subject from both stages is given by the *maximum* of either zero or her payoff from the first stage minus the punishment received and the cost of punishing other subjects.

3.1 The networks

We restrict attention to the case of four-person networks, which has several nontrivial architectures. Each network is represented by a graph with four nodes, indexed by $i = A, B, C, D$, where at each node there is a single player. An edge between any two subjects indicates that they are connected and the arrowhead points to the subject whose action 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. We can think of N_i as representing subject i's *neighborhood*. The collection of neighborhoods $\{N_A, N_B, N_C, N_D\}$ completely define a four-person network. The set of networks used in the experimental design is illustrated in Figure 1. Note that edges can be directed – the fact that subject i can monitor subject j ($j \in N_i$) does not necessarily imply that j can monitor i ($i \in N_j$).

Next, we define some key graph-theoretic concepts to which we refer throughout the paper. Our notation and definitions are standard, but to avoid ambiguity we present the concepts in some detail. The first three properties – completeness, connectedness and directedness – are *global* properties of the network architecture whereas degree is a *local* property of the neighborhoods that define the network (the edges into and out of a particular node).

- **Completeness** – A *complete* network is a network in which each pair of nodes is connected by an edge. Otherwise, the network is *incomplete*. Referring to Figure 1, in the complete network [1] every subject directly monitors every other subject. The rest of the networks [2]-[8] we study are incomplete.

- **Connectedness** – A network is *connected* if every pair of nodes i and j is linked by a path and *disconnected* otherwise.[1] Obviously, subjects from disconnected components of a network cannot monitor each other. Referring to Figure 1, networks [1]-[4] are connected whereas networks [5]-[8] are disconnected. Networks [7] and [8] are disconnected, but connectedness is satisfied in a subgraph $\{N_A, N_B, N_C\}$ in which every pair of nodes i and j are connected.

[1] Put precisely, for any pair of players i and j, a *path* from i to j is a sequence $i_1, ..., i_K$ such that $i_1 = i$, $i_K = j$ and there is an edge pointing from i_k to i_{k+1} for $k = 1, ..., K - 1$. Player i is connected to j if there is a path from i to j.

- **Directedness** – A network is *undirected* if the relations between any pairs of nodes is symmetric, so that each edge points in both directions. Otherwise, the network is *directed*. A directed network in which each edge is given a unique direction is called an *oriented* network. In an oriented network, if subject i can monitor subject j ($j \in N_i$) then j cannot monitor i ($i \notin N_j$). In our experimental design the directed networks [2] and [5]-[7] are oriented. Networks [1], [3], [4] and [8] are undirected as all edges are bi-directed and point to both nodes at once.

- **Degree** – Finally, the *degree* of a node is the number of edges that end at that node (a local property). In a directed graph the degree is usually divided into the *out-degree* and the *in-degree*. The out-degree (resp. in-degree) of node i is the number of edges with i as their initial (resp. terminal) node. Clearly, the out-degree of subject i is the number of subjects j that can be monitored by i ($j \in N_i$) and the in-degree of subject i is the number of subjects j that can monitor i ($i \in N_j$).

3.2 The game

The game is formally described using the following notation. Each subject $i = A, B, C, D$ is endowed with y indivisible tokens. At stage one, the subjects simultaneously choose how many tokens $0 \leq g_i \leq y$ to contribute to the provision of the public good. The payoff for each subject i in the first stage can be summarized by

$$\pi_i^1 = y - g_i + \alpha \bar{g}, \tag{1}$$

where

$$\bar{g} = \sum_{j=A,B,C,D} g_j$$

and α is the marginal per capita return (MPCR). Hence, each subject receives the value of the public good ($\alpha \bar{g}$) in addition to the number of tokens retained from her endowment ($y - g_i$). To avoid trivialities, we assume that $0.25 < \alpha < 1$. This condition ensures that contributing is, on one hand, socially efficient, and on the other hand, strictly dominated for any individual subject.

At stage two, after subjects are informed about their individual contributions, they can punish the subjects to whom they are connected in the network at a cost. More precisely, each subject i can punish subject $j \in N_i$ by reducing subject j's payoff from the first stage π_j^1 by p_i^j tokens. The cost of reducing one token from any of the other subjects is $0 < c < 1$ tokens. We also assume that each subject i can spend up to her entire payoff from the first stage π_i^1 toward reducing the payoff π_j^1 of all $j \in N_i$, and that π_i^1 can be reduced at the second stage to zero but that the terminal payoff for the game cannot be negative. The payoff of subject i from both stages of the game can therefore be summarized by

$$\pi_i = \max\left\{0, \pi_i^1 - c\sum_{j \in N_i} p_i^j - \sum_{j: i \in N_j} p_j^i\right\}. \tag{2}$$

By backward induction, it follows that punishment cannot defer free riding in any network architecture. We next briefly illustrate the logic of the backward induction argument and then draw out the important implications of the theory. Since punishing is costly, in theory each subject i will refrain from doing so at the second stage ($p_i^j = 0$ for all i and $j \in N_i$). Because each subject $j \in N_i$ expects that subject i will never punish her, her best response is to contribute nothing in the first stage ($g_j = 0$). Thus, the addition of the second stage has no effect on the outcome of the first stage which is full free-riding, and therefore the prediction of standard theory is that $g_i = 0$ and $p_i^j = 0$ for all i and $j \in N_i$. Note, however, that the aggregate payoff is maximized if each subject i fully cooperates by contributing her endowment ($g_i = y$).

3.3 Experimental design

The experiment was run at the Center for Experimental Social Sciences (CESS) at New York University and at the Experimental Social Science Lab (Xlab) at the University of California, Berkeley. The subjects in this experiment were recruited from all undergraduate classes and had no previous experience in public good or networks experiments.[2] Each experimental session lasted about an hour and a half. A $5 participation fee and subsequent earnings from playing the game were paid in private at the end of the experimental session. The experiments provide us with a rich set of data. Table 1

[2] After subjects read the instructions, the instructions were read aloud by an experimenter. Sample experimental instructions, including the computer program dialog windows are available in an Online Appendix: http://emlab.berkeley.edu/~kariv/CKS_II_A1.pdf. At the end of the instructional period subjects were asked if they had any questions or any difficulties understanding the experiment. No subject reported any difficulty understanding the procedures or using the computer program.

below summarizes the experimental design and the number of observations in each network treatment (the entries have the form a / b where a is the number of subjects and b the number of observations per game).

[Table 1 here]

The endowment y was 25 tokens and the marginal per capita return and the cost of punishing were fixed at $\alpha = 0.4$ and $c = 0.5$, respectively. The network was held constant throughout a given experimental session. In each session, the network nodes were labeled A, B, C, and D. The subject's type (A, B, C, or D) remained constant throughout the session. Each experimental session consisted of 15 independent decision-rounds. To minimize the investment in reputations, each round started with the computer randomly forming four-person networks by selecting one subject of each type. The networks formed in each round depended solely upon chance and were independent of the networks formed in any of the other rounds.

Each round of the experiment consists of two stages, the contribution stage and the punishment stage. The contribution decision was to allocate the endowment between a private good which only benefited the subject and a public good which benefited everyone in the group, according to the payoff function (1). Once all the contributions were recorded, subjects observed the contributions of the subjects to whom they were connected by the network. In addition, all subjects were informed about the sum of the contributions to the public good by all the subjects in their group (including themselves). In the punishment stage, subjects choose if and by how much to reduce the first stage payoff of each of the other subjects with whom they were connected by the network. If they did not wish to reduce the earnings of another subject they had to enter zeros. At the completion of the two stages, the computer informed the subjects of their total payoffs according to the payoff function (2).

4 Research questions

We next identify questions and hypotheses that can be explored using the experimental data. Our experimental design allows us to make direct comparisons to test the importance of each of the four properties of network architecture – completeness, connectedness, directedness, and degree. Recall that completeness, connectedness and directedness are global properties of the network architecture whereas degree is a local property of the neighborhoods that define the network.

- **Completeness** – To test the importance of the network being complete, we compare the complete network [1] to two other networks that are incomplete – the undirected circle [3] and the undirected star [4]. This comparison is salient because, while both the undirected circle and the undirected star are incomplete, they are connected and undirected like the complete network.

- **Connectedness** – To test the importance of the network being connected, we compare the directed circle [2] to line [6] because removing just one edge renders the directed circle disconnected. We also compare the directed circle to disconnected directed circle [7], and the undirected circle [3] to the disconnected undirected circle [8].

- **Directedness** – To examine the impact of the network being directed, we compare the directed and undirected circles [2] and [3], and the disconnected directed and undirected circles [7] and [8]. We also compare the directed star [5] to the undirected star [4]. The only complication here is that the comparison of the directed and undirected stars might be confounded by the fact that directed star is also disconnected while the undirected star is not.

- **Degree** – Finally, we examine local network effects as implied by individual connectivity. In particular, we test whether subjects who are otherwise similarly situated (same in-degree and out-degree) behave differently in different networks. For example, we compare the behavior of all subjects in the directed and undirected circles [2] and [3] to the behavior of subjects A, B and C in the disconnected directed and undirected circles [7] and [8], respectively.

5 Experimental results

In this section we report the results of our experiment. We proceed by systematically evaluating the three global properties of network architecture – completeness, connectedness and directedness. We then detect the effect of differences in network architectures by comparing the behavior of subjects that are symmetrically situated in different networks. Figure 1 above provides a summary of our results. Statistics for three metrics of group performance – average contributions as a fraction of the total endowment, probability of punishing a subject who contributes nothing, and payoffs net of the costs of being punished and punishing – are reported directly next to each network so that one can begin to map outcomes onto architecture. Along the way we will use a combination of nonparametric rank sum tests ($|z|$) and proportions tests ($|r|$). Where appropriate, we will also run parametric regressions that account for individual heterogeneity (using random effects) and learning (using period fixed effects), and add the appropriate controls.

5.1 Completeness

On average, subjects contributed 56 percent of their endowment in our complete network [1]. This behavior is similar to the "stranger" contribution levels found by Fehr and Gächter (2000) and Carpenter (2007) of 58 percent and 61 percent, respectively. While the complete network elicits contributions in line with other experiments, it does not yield the highest average contributions. Indeed, when comparing all the networks, the directed circle [2] generates the highest average contribution of 60 percent of the endowment, a rate that is statistically greater than all the others except the complete network.

[Figure 2 here]

Limiting attention to the undirected connected networks [1], [3] and [4] is most interesting because the only difference between them is whether they are complete. When we consider the summary statistics in Figure 1 and the direct comparisons in Figure 2(a) above, it does not appear that the complete network results in robustly higher contributions. While the mean contribution pooled across rounds is slightly higher in the complete network [1], it is not significantly higher than in the undirected circle [3] and only marginally significantly higher than in the undirected star [4] ($|z| = 0.64$, $p = 0.52$; $|z| = 1.99$, $p = 0.05$, respectively). In fact, when we control for individual heterogeneity and repeated play, we do not find any significant differences in contributions (see column (1) of Table 2 in which the complete network is the omitted network).

[Table 2 here]

Figure 1 above also lists the unconditional probability that a subject who contributes nothing will be punished in each of the networks. In the complete network [1] total free-riders are punished 48 percent of the time; this rate is only slightly higher than in the undirected circle [3] or the undirected star [4], 42 percent and 47 percent, respectively. A proportions test suggests that these three rates are not different ($|r| = 0.47$, $p = 0.64$; $|r| = 0.18$, $p = 0.86$, respectively). However, while the incidence of punishment might not vary between the complete network and undirected circle or undirected star, the severity does. Summarizing punishment is tricky because punishment occurs in relation to contributions, which vary. In Figure 2(b) above we plot the estimated relationship between a subject's contribution and how much she was punished in the case of the complete network, the undirected circle, and the undirected star.

As in Carpenter and Matthews (2009), we utilize a spline specification to allow punishment to diminish more quickly above the implied contribution norm. When we do so, the "knot" that maximizes the regression F statistic, and in this sense fits the data best, occurs when 10 tokens are contributed; that is, the degree to which a subject is punished in relation to how much she contributes increases more rapidly as contributions increase from 0 to 10 tokens than as they increase beyond 10 tokens. As one can see Figure 2(b), punishment levels are much lower in the complete network [1]. This finding is confirmed when we regress positive punishment amounts on network indicators and find that the estimated punishments in the undirected circle [3] and the undirected star [4] are significantly larger than in the complete network (see column (4) of Table 2).[3]

Taken together, relatively high contribution levels and low punishment expenditures make the complete network [1] one of the most efficient architectures. The mean payoff net of the costs of being punished and punishing in the complete network is significantly higher than in the undirected circle [3] ($|z| = 4.13$, $p < 0.01$) and marginally significantly higher than in the undirected star [4] ($|z| = 1.61$, $p = 0.10$). These results are largely replicated when we control for individual random effects and time period fixed effects (see column (5) of Table 2). As one can see in Figure 2(c) above the

[3]Punishment levels might be lower in complete networks because subjects face a coordination problem. Without directly communicating about how much punishment should be levied and how it should be shared, everyone continues to punish, but they are each forced to estimate, on their own, how much to reduce the severity of punishment to account for the actions of the other subjects.

performance of the complete network tends to improve, relative to the undirected circle and the undirected star as the experiment proceeds. In sum, the complete network does not seem to be better at eliciting contributions but, because its punishment levels tend to be lower, it does achieve higher than average efficiency.

5.2 Connectedness

The differences between connected and disconnected networks can be seen by comparing the left column of Figure 1 to the right column. The contribution differences are striking: there is no disconnected network [5]-[8] that yields higher mean pooled contributions than even the lowest-performing connected network, the undirected star [4]. However, to conduct our analysis systematically, remember that we need to compare the directed circle [2] to the line [6] and to the disconnected directed circle [7], and the undirected circle [3] to the disconnected undirected circle [8] directly, which we do in Figure 3(a-c) below. In each comparison, the connected networks yield significantly higher pooled contributions ($|z_{2-6}| = 14.00$, $p < 0.01$; $|z_{2-7}| = 7.31$, $p < 0.01$; $|z_{3-8}| = 4.38$, $p < 0.01$), and these results are mostly robust to an analysis that controls for individual heterogeneity and repeated game effects. If we compare the point estimates (see column (1) of Table 2), we find that the coefficient of the directed circle is higher than the coefficients of the line and the disconnected directed circle ($p < 0.01$). Similarly, the coefficient for the undirected circle is higher than the one for the disconnected undirected circle; however, in this case the difference is not significant at standard levels ($p = 0.23$).

[Figure 3 here]

Although the evidence is mixed, connected networks also appear to elicit at least as much punishment, in terms of both incidence and level. The information in Figure 1 suggests that the probability of punishing a total free-rider is higher in directed circle [2] than in the line [6] ($|r| = 3.49$, $p < 0.01$), higher in the undirected circle [3] than in the disconnected undirected circle [8] ($|r| = 4.38$, $p < 0.01$), and no lower in the directed circle than in the disconnected directed circle [7] ($|r| = 1.10$, $p = 0.27$). The full punishment splines in Figure 4(a-c) below, which are based on an analysis of all the punishment data (including zero punishments), seem to indicate that there is more punishment in connected networks. However, considering only the positive observations, the differences do not appear to be robustly significant (see column (4) of Table 2).

[Figure 4 here]

There also appears to be mixed evidence of an efficiency advantage in connected networks. While the information in Figure 1 indicates that the directed circle [2] yields higher mean payoffs than line [6] and the disconnected directed circle [7] ($|z| = 5.20$, $p < 0.01$; $|z| = 2.49$, $p < 0.01$, respectively), the payoffs in the disconnected undirected circle [8] actually tend to be higher than in the undirected circle [3] ($|z| = 1.81$, $p = 0.07$). Looking at the difference in mean payoff over all rounds of the experiment, as shown in Figure 5(a-c) below, we see little evidence of differences; these results are also seen by comparing the estimates, in which only the payoff difference between the directed circle and the line [6] is significant (see column (5) of Table 2). To summarize, although it is clear that connected networks tend to achieve higher contribution levels, connectedness does not always lead to different punishment levels or robustly higher average efficiency.

[Figure 5 here]

5.3 Directedness

A network is undirected if each edge points in both directions. Otherwise, the network is directed. In studying Figure 1 and Figure 3(d-f) above one can see that the directed circle [2] elicits higher average contributions than the undirected circle [3] ($|z| = 2.20$, $p = 0.03$). However, when one subject is isolated, the disconnected undirected circle [8] does better than the disconnected directed circle [7] ($|z| = 3.76$, $p < 0.01$), and also in the star architecture, the undirected star [4] yields higher contributions than the directed star [5] ($|z| = 11.13$, $p < 0.01$). That said, only the difference between the star networks remains significant ($p < 0.01$) when we include individual random effects and time period fixed effects (see column (1) of Table 2).

Comparing the probabilities of punishing a total free-rider as shown in Figure 1 and the punishment splines shown in Figure 4(d-f) above, we also find that the effect of directedness on punishment is network-dependent. The directed circle [2] yields a higher likelihood than the undirected circle [3] of a free-rider being punished; in fact, the directed circle yields the highest likelihood of *all* the networks of subjects who contributed nothing being punished. Nevertheless, the difference is not quite statistically significant ($|r| = 1.39$, $p = 0.16$) because there are very few observations of total free-riding in these two networks.

Furthermore, there is a large difference between the chance of a total free-rider being punished in the disconnected directed circle [7] and in the disconnected undirected circle [8]. Here the directed network yields both a higher instance of punishment ($|r| = 3.38$, $p < 0.01$) and a significantly higher level of punishment ($p < 0.01$), as can been seen by comparing the point estimates (see column (4) of Table 2). As was the case for contributions, the effect of directedness reverses in this domain as well when we consider the star architecture, as one sees in comparing the undirected star [4] and the directed star [5], but this difference also does not achieve statistical significance ($|r| = 1.52$, $p = 0.13$).

The marginally significant contribution and punishment differences between the directed and undirected networks combine to provide significant payoff differences. Although not obvious from Figures 5(d-f), the directed circle [2] results in higher average payoffs than the undirected circle [3] ($|z| = 2.06$, $p = 0.04$), and the undirected star [4] does better than the than the directed star [5] ($|z| = 7.00$, $p < 0.01$). On the other hand, the disconnected undirected circle [8] yields higher average payoffs than the disconnected directed circle [7] ($|z| = 3.38$, $p < 0.01$). Except for the comparison between the directed and undirected circles, these payoff differences remain significant after controlling for individual heterogeneity and learning (see column (5) of Table 2).

To summarize, more than in the case of both completeness and connectedness, the effect of directedness seems to depend on the structure of the network. When the underlying network is connected, there is some evidence that directedness leads to more cooperation and punishment, and to higher payoffs. When both the directed and undirected networks include a completely disconnected subject we find the opposite of the connected case – directed edges in this structure lead to much more punishment but not more cooperation or higher payoffs. Lastly, when the underlying structure is star-shaped, having a disconnected "prison guard" (the center of the star who can punish the three subjects at the periphery) is particularly bad for contributions and payoffs.

5.4 The impact of global / local network

To this point, we thoroughly studied the global network effect as represented by the completeness, connectedness and directedness of the entire network. We next study the local network effect as represented by the local links a subject has, and discuss both the global and local network effects. Recall that degree is mostly a nodal property, and while our experiment was not designed to systematically add or remove edges (largely because this would confound comparisons of the other properties), we can now use the concept to transition from analyzing the overall performance of a network to analyzing the effect of broader network structures on the behavior of subjects who have the same local neighborhoods (that is, occupy nodes of common out- and in-degrees). The question then is do networks have an impact on the contribution and punishment behavior of subjects inhabiting identical nodes. The answer is that when the underlying structures are significantly different, they clearly do because nodal behavior appears to be affected by significant changes in network architecture.

Since nodes are defined by their out- and in-degrees, we can index them as $N_{n,o,i}$ to indicate a node in network n with o out-degree o and in-degree i. For example, $N_{6,1,1}$ are nodes B and C in the line network [6]. Note that all nodes in the undirected circle [2] also have 1 out-degree and 1 in-degree (they are all defined as $N_{2,1,1}$) so they also share the same same local neighborhood as nodes B and C in the line network [6] although they are in different networks. For notational convenience, we simply leave n unspecified when we consider nodes outside the context of their networks. For example, $N_{n,1,1}$ denotes the generic 1 out-degree and 1 in-degree node. Table 3 catalogues the nodes that exist in our experiments, the networks they are part of, and the subject types that inhabit them.

[Table 3 here]

By design, most nodes exist in several networks; however, some nodes are rare and have only one representative. The most common node is $N_{n,1,1}$ since it occurs in four networks: the undirected circle [3], the undirected star [4], the line [6], and the disconnected directed circle [7]. Other flexible nodes are $N_{n,3,3}$ in the complete networks [1] and undirected star, $N_{n,2,2}$ in the undirected circle [3] and disconnected undirected circle [8], $N_{n,0,1}$ in the directed star [5] and line [6] networks, and the isolated nodes $N_{n,0,0}$ in the disconnected directed and undirected circles. Nodes $N_{n,3,0}$ and $N_{n,1,0}$ exist in only one network and hence will not be discussed. In sum, as catalogued in Table 3 above, five of the seven nodes we study exist in more than one network. Differences in the behavior and outcomes of subjects who occupy the same node in different networks are depicted in Figure 6 below. To test for differences in nodal outcomes by network, we use nonparametric tests and regress contributions, positive instances of punishment, and payoffs on node indicators (see columns (1-3) of Table 4 below).

[Figure 6 here]

5.4.1 $N_{n,3,3}$

Let us first discuss the $N_{n,3,3}$ nodes. Because neither the results of the summary test ($|z| = 0.43$, $p = 0.67$) nor the point estimates different are significant ($p = 0.98$), contributions do not seem to differ for subject who find themselves in nodes $N_{n,3,3}$ in the complete network [1] or the undirected star [4]. However, the subjects who find themselves at the center of the undirected star punish more ($p = 0.05$), and as a result accrue lower payoffs ($p = 0.01$), than the subjects in the complete networks (see columns (1) and (2) of Table 4 below). Clearly, the punishment responsibilities are more salient to the subject at the star's center. In this sense the different network structures change the behavior and outcomes at $N_{1,3,3}$ and $N_{4,3,3}$ significantly.

[Table 4 here]

5.4.2 $N_{n,2,2}$

The $N_{n,2,2}$ node exists in the connected and disconnected undirected circles [3] and [8]. The difference between these networks is that one subject is completely disconnected in disconnected directed circle, while all subjects are connected in the directed circle. Figure 6(b) and the summary tests suggest differences in contributions and punishment between occupiers of the $N_{n,2,2}$ node in these two networks. Subjects in the connected network contribute and punish more ($|z| = 3.88$, $p < 0.01$; $|r| = 1.81$, $p = 0.07$).[4] However, in terms of payoffs, the larger $N_{3,2,2}$ contributions and punishments

[4] While these differences survive the inclusion of period fixed effects, the substantial individual heterogeneity can reduce their significance (see columns (1) and (2) of Table 4 above).

tend to "net out"; there is no evidence of payoff differences between the two $N_{n,2,2}$ nodes ($|z| = 1.11$, $p = 0.27$).

The results thus suggest that having someone completely disconnected from the monitoring network affects the outcomes for the subset of connected subjects. In other words, the connected subjects A, B and C in the disconnected directed circle do not simply ignore subject D. In fact, one could imagine that subject D becomes a scapegoat because she does not have to fear punishment and therefore it would appear to the other subjects that it would be easy for her to free-ride. Ironically, the isolated subjects in the disconnected directed circle actually appear to play with considerable integrity; in fact, they contribute at levels comparable to the three connected members of the network ($|z| = 0.86$, $p = 0.39$).

5.4.3 $N_{n,1,1}$

We see the most variation in behavior among the $N_{n,1,1}$ nodes. The mean fraction contributed by subjects in the directed circle [2] was 0.60 while that same node, in the undirected star [4], the line [6], and the disconnected directed circle [7] produced contribution levels of just 0.49, 0.27, and 0.38 respectively. Using a Kruskal-Wallis test, we find that network structure has a significant impact on nodal performance ($\chi^2 = 177.37$, $p < 0.01$).[5] In short, the $N_{n,1,1}$ nodes elicit high contributions when embedded in connected networks, with the directed circle [2] being the most hospitable network for such nodes and the line being the most inhospitable network for them.

The impact of network structure on nodal performance can also be seen in the punishment behavior of $N_{n,1,1}$ subjects since while the mean punishment given by subjects in the directed circle [2] is 2.67, it is 0.78, 1.10, and 2.26, respectively, for the subjects in the $N_{n,1,1}$ nodes of undirected star [4], the line [6], and the disconnected directed circle [7]. Another Kruskal-Wallis test indicates that these differences are significant ($\chi^2 = 24.33$, $p < 0.01$).[6] Note that while punishments are highest in the directed circle, they are lowest in the undirected star, which is also connected. This is not too surprising since in the star there are three subjects who have the opportunity to punish the

[5] If done pairwise, all the differences are significant when the nonparametric rank sum test is used and most of them survive the addition of individual random effects and time period fixed effects (see column (1) of Table 4).

[6] Here the $N_{2,1,1}$ versus $N_{7,1,1}$ and $N_{4,1,1}$ versus $N_{6,1,1}$ comparisons do not survive when the analysis is done pairwise.

one subject in the center of the star. Again, such a plethora of punishers appears to create a coordination problem: who will be responsible for punishing the subject in the center of the star?

Another factor that might account for the relatively low punishment in undirected star [4] is the fact that the $N_{4,1,1}$ subject have mutual links – not only can they monitor and punish their neighbors, their neighbors can punish them. Occupiers of the $N_{n,1,1}$ node in other networks, however, punish one neighbor and are punished by another. Hence, the subjects in nodes $N_{4,1,1}$ might show more restraint compared to the subjects in other $N_{n,1,1}$ nodes because they are afraid of engaging in punishment feuds (Nikiforakis, 2008). It is also interesting that the mean punishment levels do not differ much between the connected and disconnected directed circles [2] and [7], a result that is not too surprising considering the monitors in the disconnected circle form their own three-person version of the connected circle.

In terms of payoffs, the $N_{n,1,1}$ node clearly does best in connected networks.[7] If one were to look for a common feature to explain this difference, it might be that in the disconnected networks, the line [6] and the disconnected directed circle [7], there exists one subject who is not monitored and cannot be punished. So despite the fact that subjects in these nodes seem to be playing the same monitoring-punishment game locally, globally they realize that there is one subject that has no incentive to contribute. In this sense, the broader network matters.

5.4.4 $N_{n,0,1}$ and $N_{n,0,0}$

The two remaining types of nodes to compare are $N_{n,0,1}$ and $N_{n,0,0}$. Because they all involve subjects who do not punish (out-degree zero), the four instances of these two different types of nodes are combined in Figure 6(d). As to subjects in nodes $N_{n,0,1}$, the results of the nonparametric test suggest that subjects B, C, and D in the directed star [5] contribute significantly more than subject D does in the line [6] ($|z| = 4.18$, $p < 0.01$). This can be accounted for by the higher incidence of punishment in directed star. As to the isolated subjects in nodes $N_{n,0,0}$, perhaps as expected, the behavior of subject D does not appear to depend on being in the disconnected directed circle [7] or the disconnected undirected circle [8].

[7] While all the pairwise comparisons are significant when the nonparametric test is used, only the differences between the connected and disconnected networks survive the analysis (see column (3) of Table 4).

6 Discussion

In the previous section we examined each network property separately in detail. However, it might be helpful to combine the effects of all the properties in one summary analysis. Is it the case that these properties provide the foundation for the successes and failures shown in Figure 1? In Table 5 we regress our three measures of network performance – contributions, punishment, and efficiency – on the three properties of network architecture – completeness, connectedness, and directedness. Among the networks examined here, [1] is the only complete network; all the rest are incomplete. Networks [1]-[4] are connected, while networks [5]-[8] are disconnected. Networks [2], and [5]-[7] are directed, while networks [1], [3], [4], and [8] are undirected. We also add an aggregate measure of degree by calculating the total number of edges for each network.

[Table 5 here]

- **Contributions** – When we control for the other properties we find that connected networks yield significantly more contributions than disconnected networks but that there are no significant differences in the contributions with regard to the other properties. The effect of connectedness is not only statistically significant, five more tokens contributed in the connected networks out of a 25-token endowment is substantial (see column (1) of Table 5).

- **Punishment** – Two properties, completeness and directedness, affect the amount of punishment. Together, these point estimates suggest an interesting interpretation. The fact that the complete network (in which the 'policing' of free-riders is very decentralized) yield lower sanctions and that directed networks (in which it is the responsibility of just one subject to punish a free-rider) yield higher sanctions reinforces the idea that punishers face the coordination problem mentioned above (see column (2) of Table 5).

- **Efficiency** – It might be the case that punishment is too severe in directed networks and that a little ambiguity with regards to punishment enhances payoffs in the complete network. When combined, all four properties significantly affect final net payoffs. Not only do payoffs in the complete network do well while those in directed networks suffer because of differences in the amount of punishment; additional edges also reduce payoffs, perhaps because they tend to be used too often to punish, and connected networks do better (see column (3) of Table 5).

The effect of connected networks on efficiency is particularly interesting because it appears even after controlling for punishment dynamics. In other words, part of the success of connected networks occurs for reasons beyond their ability to distribute punishment. Given the differences between networks with regard to their ability to monitor and punish other subjects, it is important to examine the extent to which network differences in punishment can explain the differences that we see in contributions.

Returning to the contribution regression analysis in Table 2 above, in column (2) we add the lag of received punishment and the lag of contributions. We add the lag of contributions to control for level differences: while free-riders can increase their contributions substantially, high contributors can only increase their contributions slightly, regardless of how much punishment they receive. As expected, punishment is highly significant. For each token of punishment the average subject receives, she increases her contribution by about 0.2 tokens. As important, however, is the fact that the addition of punishment has reduced (compared to those shown in column (1) of Table 2), all the coefficients on the network indictors, some substantially. Clearly, a large part of the variation in contributions previously attributed to the differences in the networks is really due to differences in the amount of punishment generated by the different networks.

In column (3) of Table 2 we examine a robustness check on our punishment analysis. The model in column (2) assumes that subjects have a common response to punishment. In column (3) we add (unreported) interactions that allow the response to punishment to vary by network. In addition to finding that the added interactions do not change the results much, a chi-squared test of the joint significance of the interactions suggest that those interactions add nothing to the results reported in column (2) of Table 2 ($\chi^2 = 2.85$, $p = 0.90$). Thus, assuming a common response to punishment seems reasonable. In short, a major reason for the differences seen in contributions across the networks is that the some networks elicit a lot of punishment and others do not.[8]

[8] We also made sure that the mediation of the network affects on contributions by adding punishment is not due to the fact that the data from all the subjects who cannot

We end our discussion by speculating, based on our results, about the "optimal" number of edges in the standard four-person public good game with punishment. If, in the end, we are mostly concerned about the efficiency with which public goods are provided, then we have to look at payoffs net of the costs of punishing and being punished. If we estimate the effect of the number of edges on net payoffs, we find that all three coefficients of a third-order polynomial specification are significant. The derivative of this function provides us with an estimate of the marginal benefit of adding edges to the monitoring network. The marginal benefit is plotted in Figure 7 below.

[Figure 7 here]

The shape of the plot in Figure 7 is interesting and informative. We see that the marginal benefit actually becomes negative when between five and nine edges are added. In other words, it makes sense to add edges one through four or edges nine through 12, but adding edges to networks that have already between four and eight seems to actually reduce average payoffs. The precision of this estimate nicely summarizes the implications of our results. In three-edge networks like the directed star [5], the line [6], and the disconnected directed circle [7] the problem is straightforward: there is not enough monitoring, and therefore more edges should be added. At the same time, networks with eight edges, like the undirected circle [3], actually have the opposite problem. Here there is more punishment than can be supported by the level of contributions. Ironically, one way to reduce the amount of punishment is to add more edges because this will lower the responsibility of any given monitor, and this ambiguity seems to induce some monitors to reduce punishment to a level that ends up being closer to optimal.

References

[1] Anderson, C. and L. Putterman (2005). "Do Non-Strategic Sanctions Obey the law of Demand? The Demand for Punishment in the Voluntary Contribution Mechanism." *Games and Economic Behavior*, *54(1)*, pp. 1-24.

be punished (in nodes $N_{5,0,1}$, $N_{6,0,1}$, $N_{7,0,0}$, and $N_{8,0,0}$) are dropped in column (2) of Table 2. In another regression, we replaced the missing punishment values with zeros and found that the resulting point estimates are almost identical.

[2] Carpenter, J. (2007). “Punishing Free-Riders: How Group Size Affects Mutual Monitoring and the Provision of Public Goods.” *Games and Economic Behavior*, *60(1)*, pp. 31-51.

[3] Carpenter, J. and P. Matthews (2009). “What Norms Trigger Punishment?” *Experimental Economics*, *12(3)*, pp. 272-288.

[4] Cinyabuguma, M., T. Page and L. Putterman (2005). “Cooperation under the Threat of Expulsion in a Public Goods Experiment.” *Journal of Public Economics*, *89(8)*, pp. 1421-1435.

[5] Fehr, E. and S. Gächter (2000). “Cooperation and Punishment in Public Goods Experiments.” *American Economic Review*, *90(4)*, pp. 980-994.

[6] Jackson M. (2008). *Social and Economic Networks.* Princeton University Press.

[7] Kosfeld, M. (2004). “Economic Networks in the Laboratory: A Survey.” *Review of Network Economics*, *3(1)*, pp. 20-41.

[8] Ledyard, J. (1995). “Public goods: A Survey of Experimental Research.” In: John Kagel and Alvin Roth (Eds.), *The Handbook of Experimental Economics*. Princeton University Press, pp. 111-194.

[9] Masclet, D., C. Noussair, S. Tucker and M-C. Villeval (2003). “Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism.” *American Economic Review*, *93(1)*, pp. 366-380.

[10] Nikiforakis, N. (2008). “Punishment and Counter-Punishment in Public Goods Games: Can we still govern ourselves?” *Journal of Public Economics*, *92(1-2)*, pp. 91-112.

[11] O’Gorman, R., J. Henrich and M. Van Vugt (2009). “Constraining Free Riding in Public Goods Games: Designated Solitary Punishers can Sustain Human Cooperation.” *Proceedings of the Royal Society B*, *276*, pp. 323–329.

Table 1: Experimental design

Network	Nodes	# of obs./subjects
[1]	*A*,*B*,*C*,*D*	240/16
[2]	*A*,*B*,*C*,*D*	240/16
[3]	*A*,*B*,*C*,*D*	300/20
[4]	*A*	135/9
	B,*C*,*D*	405/27
[5]	*B*,*C*,*D*	765/51
	A	255/17
[6]	*A*	240/16
	B,*C*	480/32
	D	240/16
[7]	*A*,*B*,*C*	180/12
	D	60/4
[8]	*A*,*B*,*C*	315/21
	D	105/7
Total		3960/264

Table 2: Comparing outcomes across networks

	(1)	(2)	(3)	(4)	(5)
Dependent variable	Contribution	Contribution	Contribution	Punishment	Payoff
[2] Directed circle	3.974	2.374	2.102	4.471***	-1.621
	(3.396)	(2.072)	(2.095)	(1.574)	(1.382)
[3] Undirected circle	0.301	-0.168	-0.268	2.811*	-3.005**
	(3.195)	(1.946)	(1.967)	(1.557)	(1.311)
[4] Undirected star	-1.673	-0.908	-1.108	3.069**	-0.964
	(2.868)	(1.748)	(1.767)	(1.443)	(1.174)
[5] Directed star	-7.117***	-3.411**	-3.716**	2.607*	-3.307***
	(2.656)	(1.665)	(1.687)	(1.494)	(1.086)
[6] Line	-10.832***	-5.926***	-6.237***	3.984***	-4.049***
	(2.676)	(1.701)	(1.721)	(1.425)	(1.093)
[7] Disconnected directed circle	-4.857	-2.383	-2.724	10.058***	-3.513***
	(3.383)	(2.224)	(2.240)	(1.977)	(1.382)
[8] Disconnected undirected circle	-3.022	-1.385	-1.740	3.187**	-1.264
	(2.987)	(1.926)	(1.943)	(1.572)	(1.225)
Lagged punishment received		0.214***	-0.003		
		(0.03)	(0.25)		
Lagged contribution		0.530***	0.530***		
		(0.03)	(0.03)		
Constant	17.016***	3.408**	3.702**	1.283	32.937***
	(2.423)	(1.546)	(1.570)	(1.254)	(1.047)
Time period fixed effects	Yes	Yes	Yes	Yes	Yes
Punishment×network interactions	No	No	Yes	No	No
Rho	0.68	0.49	0.49	0.41	0.24
Prob > Chi^2	<0.01	<0.01	<0.01	<0.01	<0.01
# of obs.	3960	3080	3080	886	3960

The complete network [1] is omitted. Robust standard errors in parentheses. Tobit regressions adjust for censoring of the dependent variable at zero and one. *, **, *** indicates significance at the 10, 5, and 1 percent levels, respectively.

Table 3: The out-degree and in-degree

Degree	Nodes (subjects)
0-out | 0-in	$N_{7,0,0}$ *(D)*, $N_{8,0,0}$ *(D)*
0-out | 1-in	$N_{5,0,1}$ *(B,C,D)*, $N_{6,0,1}$ *(D)*
1-out | 0-in	$N_{6,1,0}$ *(A)*
1-out | 1-in	$N_{2,1,1}$ *(A,B,C,D)*, $N_{4,1,1}$ *(B,C,D)*, $N_{6,1,1}$ *(B,C)*, $N_{7,1,1}$ *(A,B,C)*
2-out | 2-in	$N_{3,2,2}$ *(A,B,C,D)*, $N_{8,2,2}$ *(A,B,C)*
3-out | 0-in	$N_{5,3,0}$ *(A)*
3-out | 3-in	$N_{1,3,3}$ *(A,B,C,D)*, $N_{4,3,3}$ *(A)*

Table 4: Comparing outcomes across nodes

	(1)	(2)	(3)
Dependent variable	Contribution	Punishment	Payoff
$N_{5,3,0}$	3.883	0.483	-0.668
	(2.801)	(1.685)	(1.241)
$N_{1,3,3}$	13.399***	-2.124	1.988
	(3.365)	(1.815)	(1.367)
$N_{4,3,3}$	13.514***	1.295	-2.021
	(3.941)	(1.939)	(1.611)
$N_{3,2,2}$	13.701***	0.686	-1.017
	(3.184)	(1.743)	(1.297)
$N_{8,2,2}$	10.796***	1.063	0.055
	(3.150)	(1.754)	(1.283)
$N_{2,1,1}$	17.360***	2.345	0.367
	(3.384)	(1.757)	(1.367)
$N_{4,1,1}$	11.131***	0.695	2.043*
	(3.002)	(1.790)	(1.220)
$N_{6,1,1}$	3.749	2.911*	-3.004**
	(2.921)	(1.742)	(1.184)
$N_{7,1,1}$	9.687***	7.942***	-2.438*
	(3.630)	(2.130)	(1.477)
$N_{5,0,1}$	7.070***		-1.536
	(2.668)		(1.094)
$N_{6,0,1}$	2.803		-2.237*
	(3.393)		(1.367)
$N_{7,0,0}$	5.002		1.276
	(5.377)		(2.162)
$N_{8,0,0}$	9.126**		2.728
	(4.288)		(1.752)
Constant	3.611	3.382**	30.947***
	(2.422)	(1.476)	(1.038)
Time period fixed effects	Yes	Yes	Yes
Rho	0.68	0.41	0.23
Prob > Chi^2	<0.01	<0.01	<0.01
# of obs.	3960	886	3960

Node $N_{6,1,0}$ is omitted. Robust standard errors in parentheses. Tobit regressions adjust for censoring of the dependent variable at zero and one. *, **, *** indicates significance at the 10, 5, and 1 percent levels, respectively.

Table 5: The effect of network properties

	(1)	(2)	(3)
Dependent variable	Contribution	Punishment	Payoff
Complete	-8.042	-2.811*	6.050**
	(7.626)	(1.487)	(3.099)
Connected	5.336***	0.565	1.398*
	(2.055)	(0.427)	(0.832)
Directed	2.267	2.143***	-4.145**
	(4.267)	(0.848)	(1.733)
Degree (total)	1.435	0.379	-0.899*
	(1.357)	(0.266)	(0.552)
Rho	0.698	0.288	0.247
Prob > Chi^2	<0.01	<0.1	<0.01
# of obs.	3960	4647	3960

Degree is the total number of edges for each network. Robust standard errors in parentheses. Tobit regressions adjust for censoring of the dependent variable at zero and one. *, **, *** indicates significance at the 10, 5, and 1 percent levels, respectively.

Figure 1: Network architectures and summary statistics

Connected

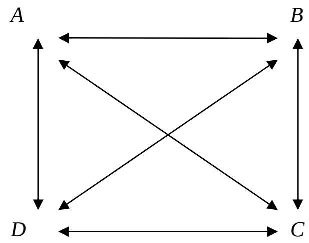

[1] Complete
Fraction(Contributed): 0.56
Pr(punish 0-contributor): 0.48
Mean Payoff: 31.64

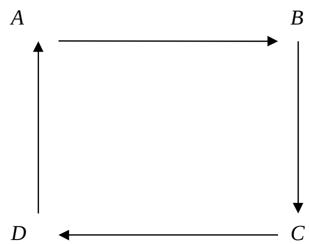

[2] Directed circle
Fraction(Contributed): 0.60
Pr(punish 0-contributor): 0.69
Mean Payoff: 30.02

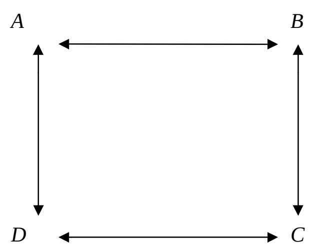

[3] Undirected circle
Fraction(Contributed): 0.55
Pr(punish 0-contributor): 0.42
Mean Payoff: 28.63

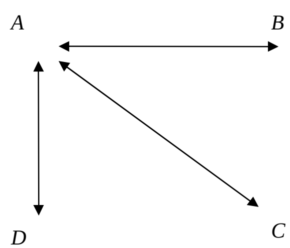

[4] Undirected star
Fraction(Contributed): 0.51
Pr(punish 0-contributor): 0.47
Mean Payoff: 30.68

Disconnected

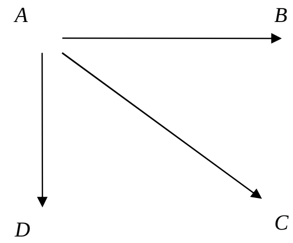

[5] Directed star
Fraction(Contributed): 0.33
Pr(punish 0-contributor): 0.37
Mean Payoff: 28.33

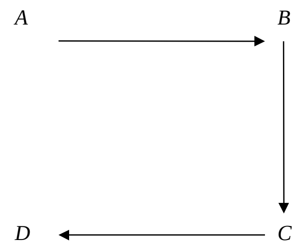

[6] Line
Fraction(Contributed): 0.25
Pr(punish 0-contributor): 0.25
Mean Payoff: 27.59

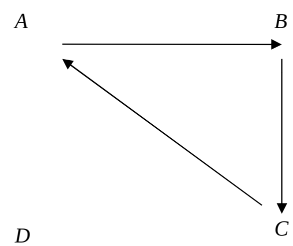

[7] Disconnected directed circle
Fraction(Contributed): 0.38
Pr(punish 0-contributor): 0.51
Mean Payoff: 28.14

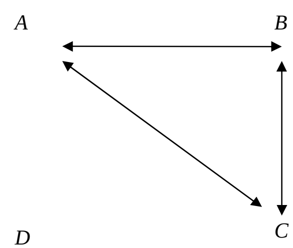

[8] Disconnected Undirected circle
Fraction(Contributed): 0.45
Pr(punish 0-contributor): 0.18
Mean Payoff: 30.37

Figure 2: Comparing complete and incomplete networks

(Figure 2(a) shows mean contributions. Figure 2(b) shows punishment given. Figure 2(c) shows mean payoff.)

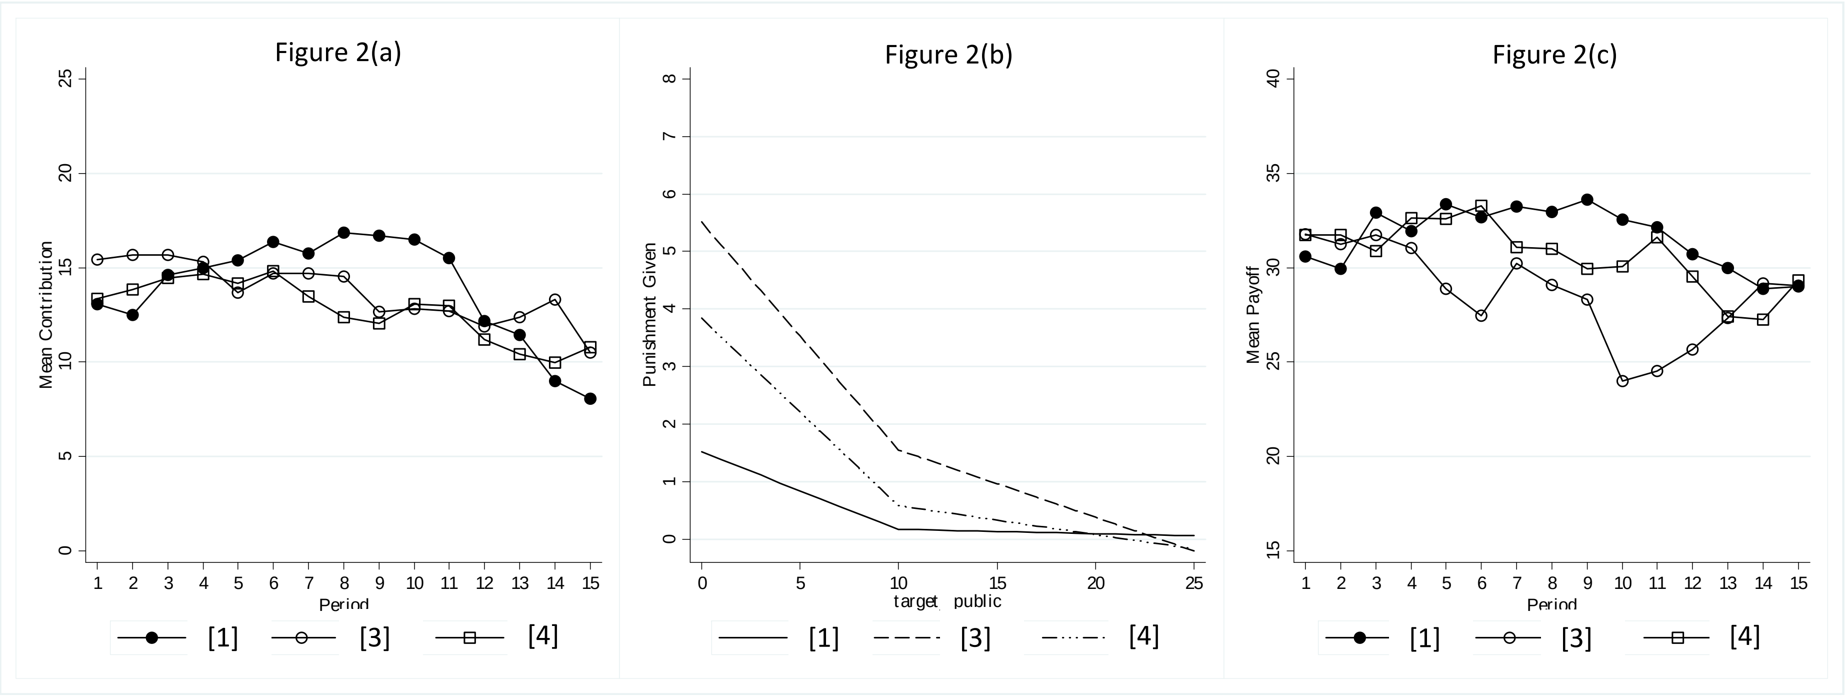

Figure 3: Comparing contributions across networks

(Figure 3(a-c) compare connected and disconnected networks. Figure 3(d-f) compare directed and undirected networks.)

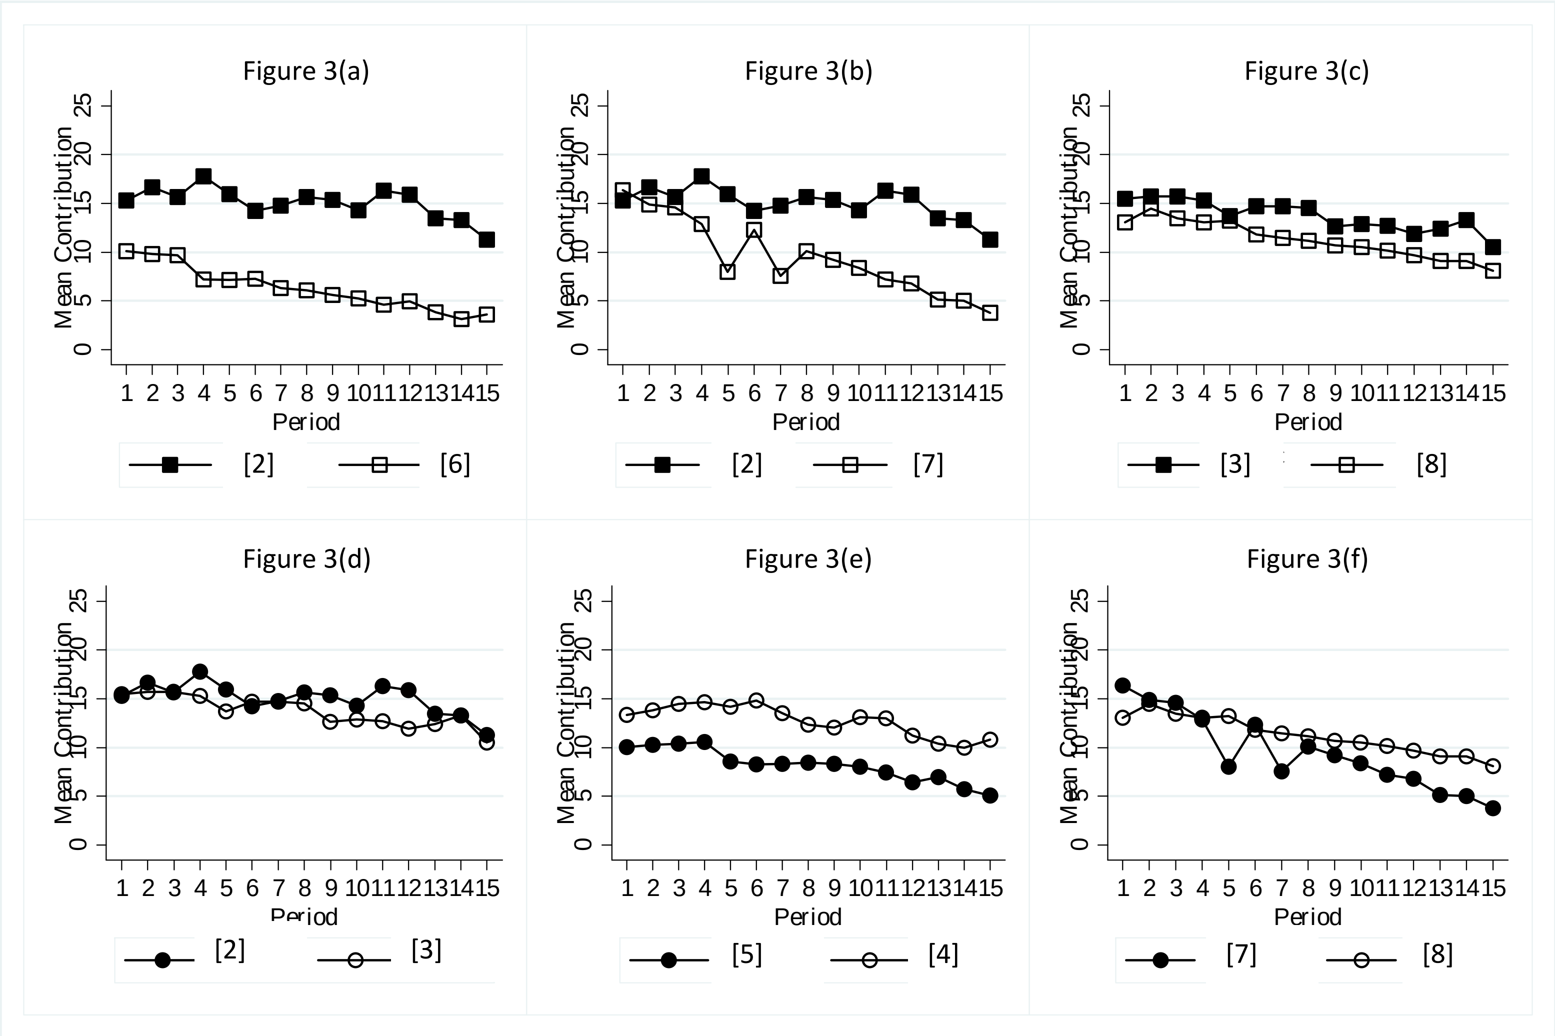

Figure 4: Comparing punishment given – estimated punishment splines

(Figure 4(a-c) compare connected and disconnected networks. Figure 4(d-f) compare directed and undirected networks.)

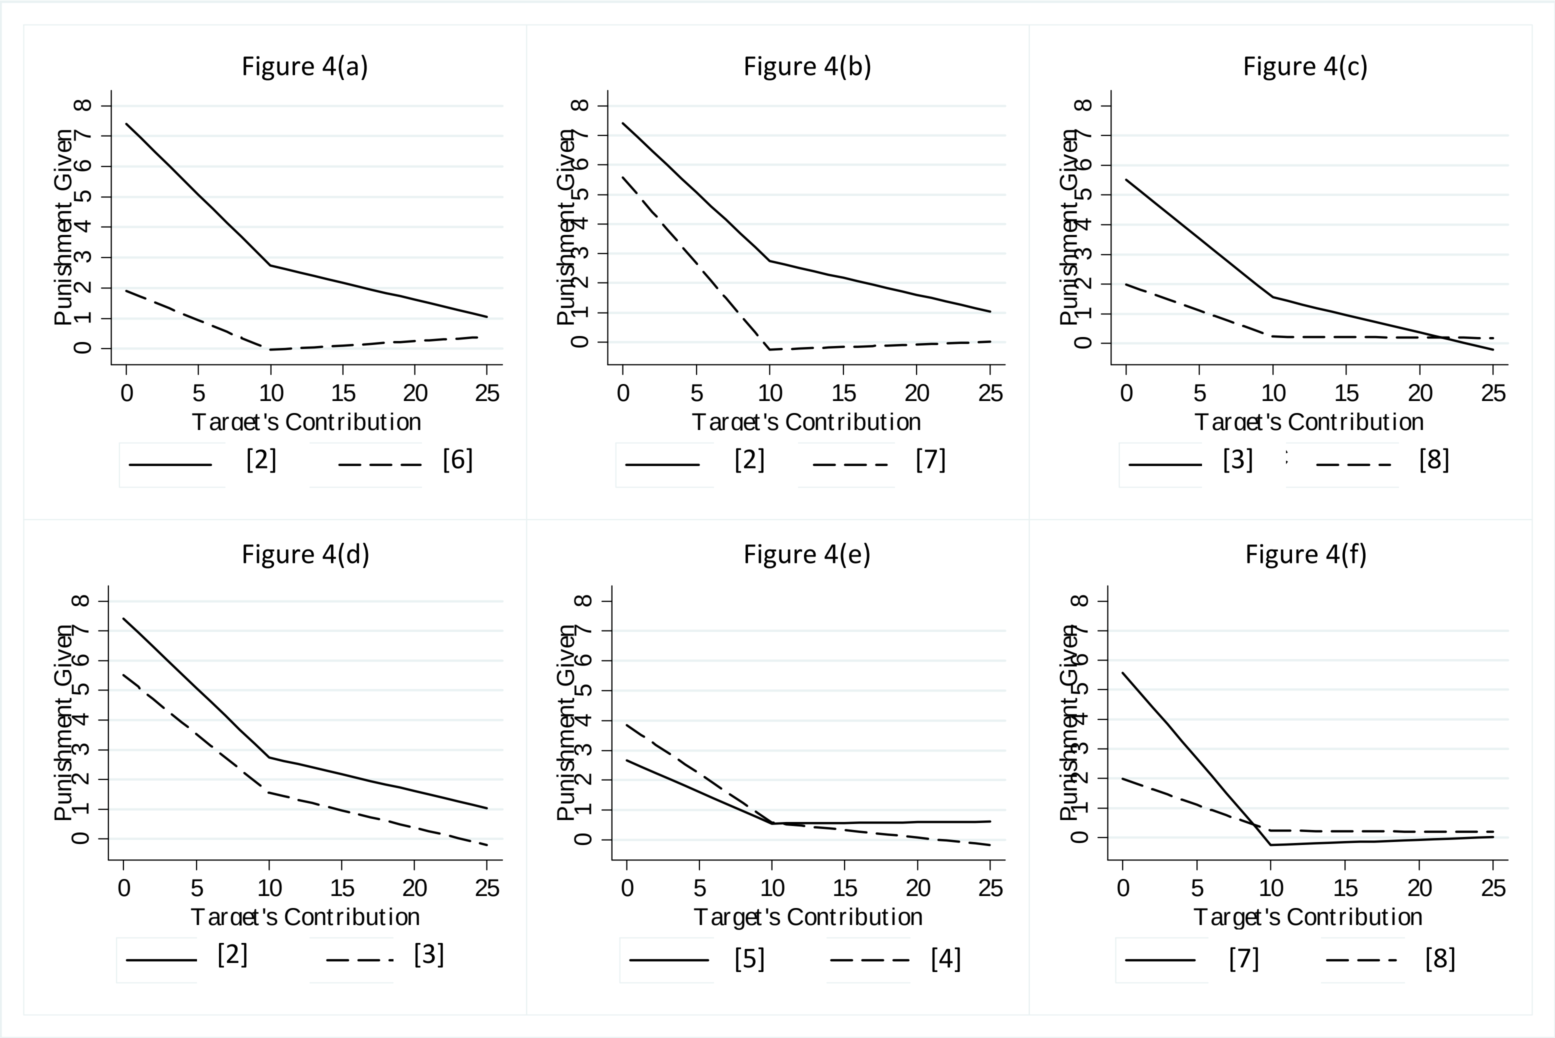

Figure 5: Comparing efficiencies

(Figure 5(a-c) compare connected and disconnected networks. Figure 5(d-f) compare directed and undirected networks.)

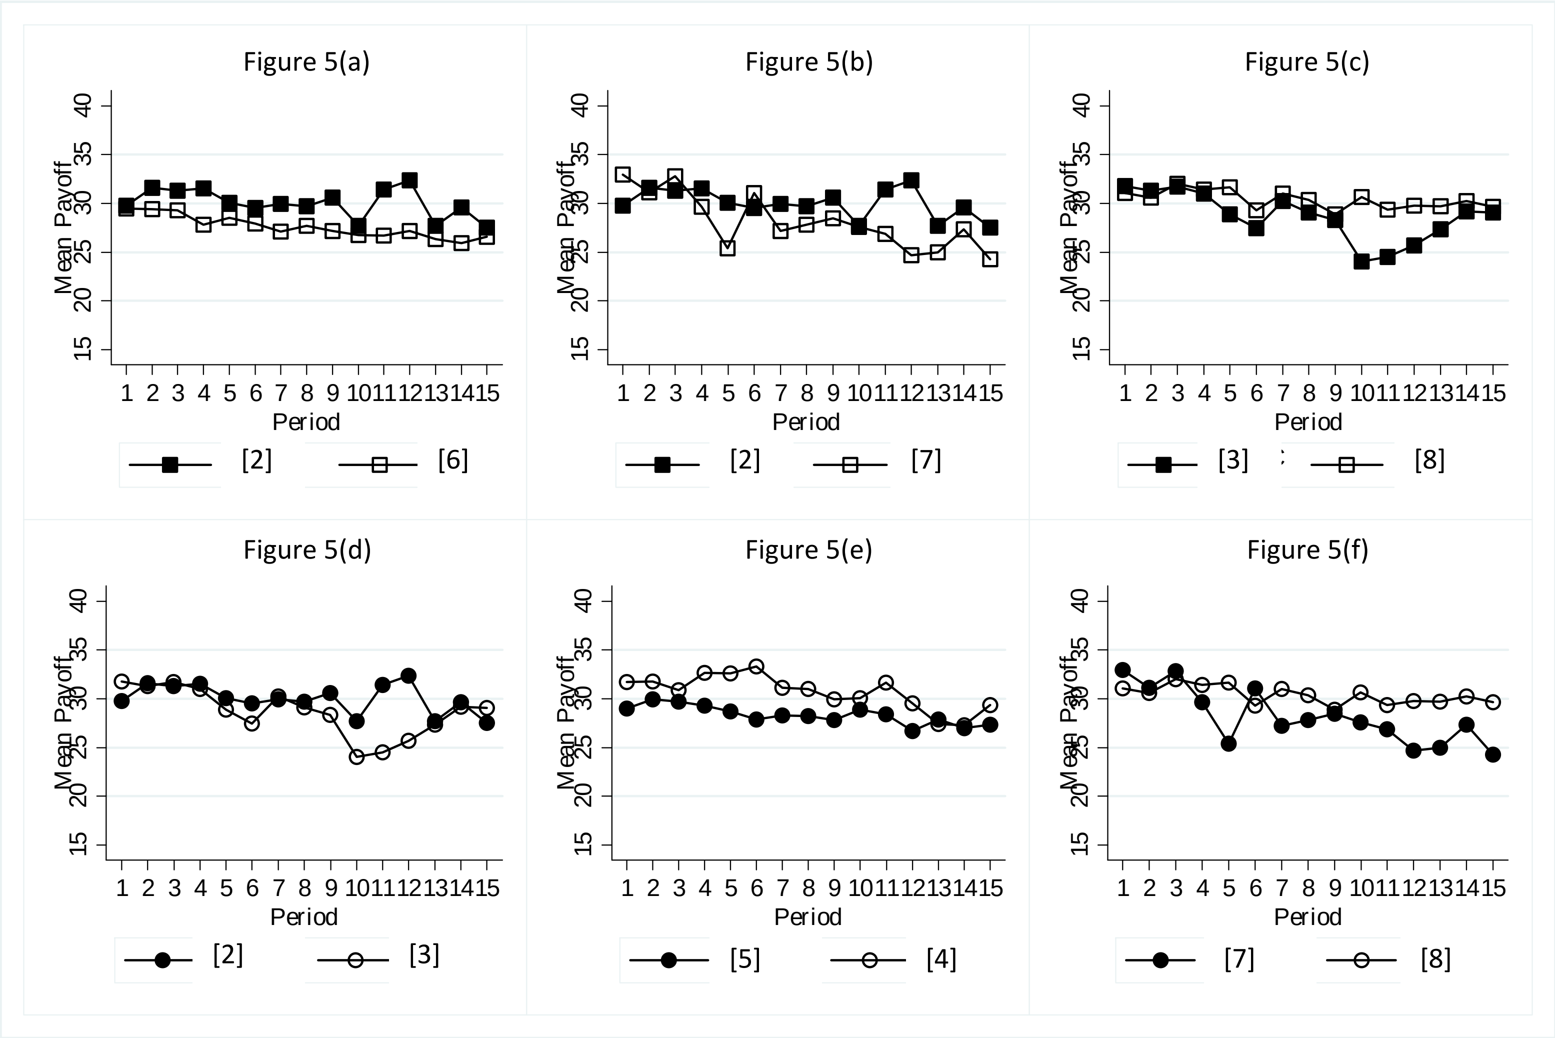

Figure 6: Comparing nodes

(Figure 6(a): $N_{n,3,3}$. Figure 6(b): $N_{n,2,2}$. Figure 6(c): $N_{n,1,1}$. Figure 6(d): $N_{n,0,1}$ and $N_{n,0,0}$.)

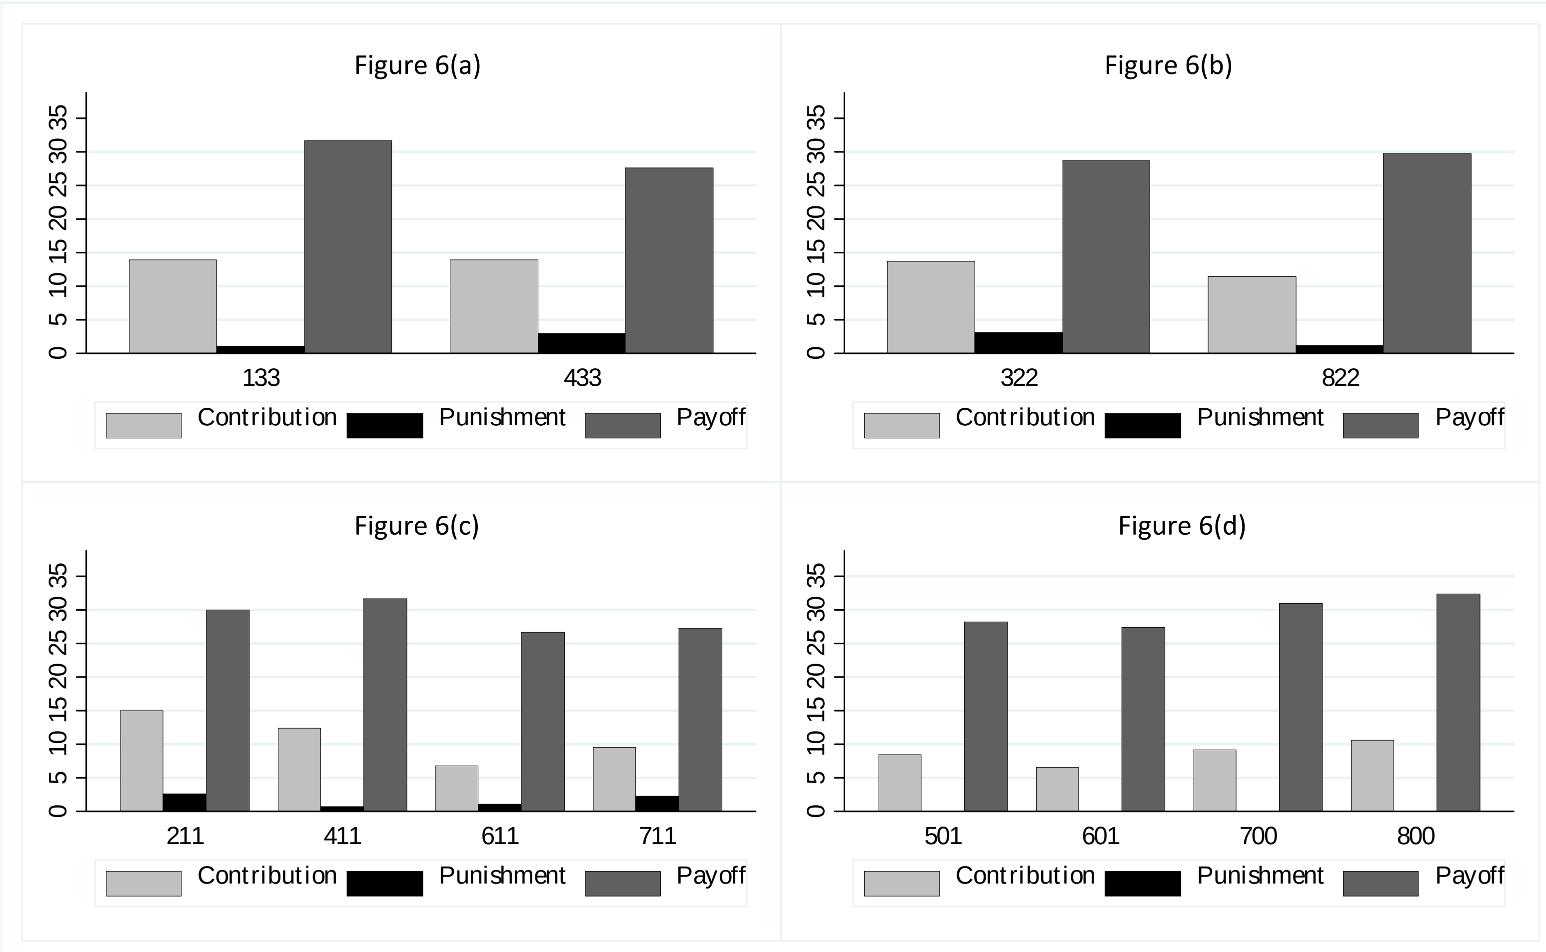

Figure 7: The estimated marginal benefit of adding edges