Fragmented Networks and Sustainable Cooperation in Public Good Experiments

Jakub Caisl*
Charles University

June, 2013

Abstract

Recent experimental studies from a public good game setting found that enforcement can significantly facilitate cooperation (Fehr & Gächter, 2000) and lead to long-term welfare improvement (Gächter et al, 2008). Similar findings usually rest on the assumptions of full information and unrestrained punishment options. In our experiment we relax these assumptions by varying the structure of both information flows and punishment opportunities. We select several non-trivial network designs that differ systematically in their levels of fragmentation, which allows us to identify network fragmentation as an essential property for predictions about cooperative efforts; and assess the relative importance of information and punishment structures for both cooperation and welfare. We find that harsh fragmentation negatively affects public good provision, but mild fragmentation does not, at least in certain contexts. This is essentially because of fragmentation of information - the effect of extending punishment possibilites has only minor, and negative, influence on cooperation.

JEL Classification: D62, D64, D85, C91, C92, H41

Keywords: social networks, public good, punishment, incomplete information, experiment

* Jakub Caisl, Charles University, Prague, Czech Republic; Email: jakubcaisl@gmail.com

1 Introduction

The importance of identifying favourable conditions for cooperation in economics cannot be overstated. It has far-reaching implications for both large and small scale social issues in which foregoing self-interest is necessary to reach optimal collective outcomes – take public goods provision or work-team performance as relevant examples. Traditional economic theory, based on models of purely self-interested agents, offers stark predictions in this respect. Unless collective-interest coincides with self-interest, there is no cooperation and the potential for improving welfare is lost. Experimental research, on the other hand, has shown that individuals do not discard collective goals that easily. They engage in public good provision even if this does not maximize their own profit, albeit such tendency weakens over time (Ledyard, 1994). Furthermore, they are willing to sanction free-riders for their selfish behaviour where appropriate means are available, which greatly facilitates the achievement of optimal group outcomes (Fehr & Gächter, 2000). This can produce large welfare benefits in the long-term perspective, encouraging evolutionary arguments for the importance of norm enforcement in sustaining cooperation (Gächter et al, 2008).

However, such findings almost always rest on the assumption of complete information about the behaviour of others, disregarding the embeddedness of economic relationships in individuals' social networks (Granovetter, 1985). Such networks emerge due to a wide variety of non-economic considerations, such as kinship, ethnicity, religion or political allegiance. They define the general structure of economic relationships within a community, only rarely allowing members to reach everyone else directly. Indeed, usually severe limitations of both information flows and enforcement opportunities occur. Such limitations in turn influence both individual behaviour and collective performance, as documented in the illuminating experimental study authored by Carpenter et al (2012).

As such, networks differ in their social capital, a type of capital defined in Coleman (1988)as: “... a variety of entities, with two elements in common: they all consist of some aspect of social structures and they facilitate certain actions of actors – whether persons or corporate actors – within the social structure.“ The aspect of social structures of particular relevance for us is their fragmentation, a property that measures how easily can individuals reach others within the network (Borgatti, 2006): an individual can either be reached directly, indirectly through a chain of intermediaries, or not at all. Fragmented networks are quite common in everyday life and can emerge due to a wide array of causes. For example, the tendency to create ties with others similar to oneself can lead to ethnic divisions in a society; or specialized departments in a company (McPherson et al, 2001).

Through the lenses of this framework we empirically evaluate two influential theories about social capital of different networks. To the best of our knowledge, no such systematic assessment has been undertaken in previous public good game studies, albeit there is some evidence about the importance of a related measure of network connectedness (Carpenter 2012; Leibbrandt et al 2012).

The first hypothesis states that strong social cohesion is necessary for the emergence and enforcement of norms in a community, which can greatly facilitate public good provision (Coleman, 1988). Highly fragmented groups, in contrast, are unlikely to create and sustain such norms and thus unlikely to reach cooperative outcomes. Anecdotal evidence of highly cohesive, or in Coleman's words "closed", groups with norms promoting non-selfish behaviour is provided to illustrate this point. For example, wholesale diamond market run by a tightly-knit Jewish community in New York can rely on strong norms of trust to deter cheating and thus reduce transaction costs.

Fig. 1 Overview of networks

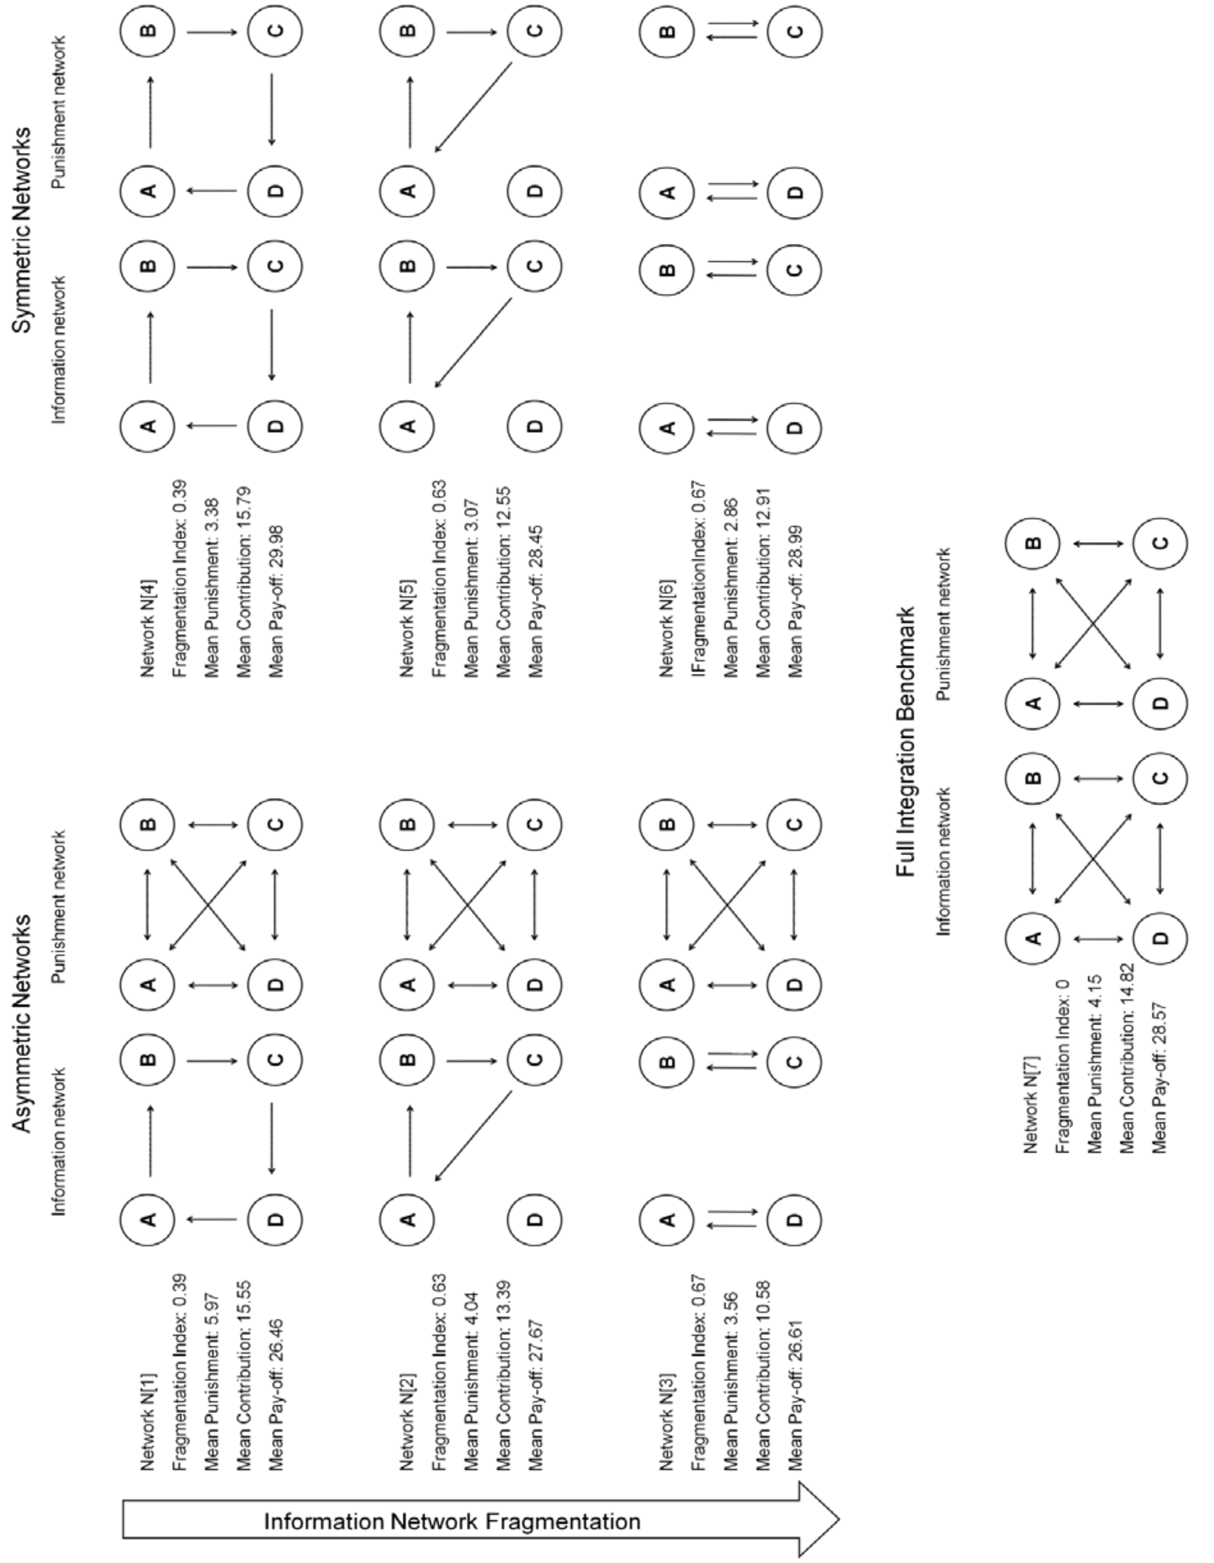

The second hypothesis acknowledges the importance of "network closure" but claims that too cohesive networks can become suffocating and, as a result, a certain amount of fragmentation can prove out to be beneficial (Burt, 2000 & 2001; Granovetter, 1973). Balkundi (2007) supports such claims with a study comparing the efforts of work-teams with different levels of social fragmentation - moderately fragmented teams delivered the best performance, outperforming both the highly cohesive and very loosely-knit teams. While the traditionally invoked argument to support such findings focuses on lack of new ideas in extremely cohesive networks, there is no creative element in our settings. Instead we rely on the previous evidence form Carpenter (2007) that extremely cohesive networks suffer from coordination problems[1] in enforcement of norms that facilitate cooperation.

Our experimental design allows us to clearly distinguish between these two theoretical predictions and assess their validity. To see this, look at Figure 1 above, where the networks used in this experiment are presented. Each network consists of two sub-networks that describe information flows and enforcement structure. A line segment with an arrowhead indicates who can observe (or punish) who – for example a line segment from individual A with an arrowhead towards individual B indicates that A can observe (punish) B. Both networks N[1-3] and N[4-6] differ in their levels of information fragmentation. Intuitively – we call N[1&5] the mildly fragmented scenario since everyone is within the same group; we call the rest of these networks severely fragmented scenarios since in N[2&6] there is one completely isolated individual; and in N[3&6] there are two fully separated sub-groups.

This selection of networks allows us to study in detail how fragmented information affects collective efforts within groups. In particular, we can compare mildly fragmented networks N[1&4] to both fully integrated benchmark N[7] and severely fragmented networks N[2-3] & [5-6]. Thus, we are able to assess the validity of the two contesting hypotheses presented above - whether cohesive networks facilitate public good provision (Coleman, 1988) or whether moderate fragmentation can be beneficial (Burt, 2000).

While such a detailed study of fragmentation is a significant extension of previous experimental findings (see Carpenter et al, 2012; Leibbrandt et al, 2012), the differences are even more important than what unites networks N[1-3] and N[4-6]. Notice that as opposed to N[1-3], networks N[4-6] have unrestricted punishment options since each individual can punish all other network members irrespective of observing his behaviour or not. Therefore each individual has additional punishment opportunities compared to his respective counterpart from N[4-6] (indexed by A, B, C or D), but these opportunities lack precise information about individual behaviour of possible recipients. We study whether individuals are willing to punish others under such uncertain conditions since this can potentially affect group performance, as previously shown in noisy information settings (Ambrus & Greiner, 2012; Grechenig et al, 2010)[2].

The comparison of networks N[1-3] with additional punishment opportunities to networks N[4-6] without them leads to the seminal contribution of this paper: the assessment of the relative importance of information and enforcement structure for public good provision. If networks with unlimited punishment N[1-3] lead to different levels of cooperation or welfare than their corresponding counterparts N[4-6], then enforcement structure clearly co-determines collective outcomes. As such it constitutes a viable target for policy or managerial decisions. Conversely, if no significant difference in between N[1-3] and counterparts N[4-6] can be identified, information

[1] These can be illustrated by the famous case of Kitty Genovese, a girl repeatedly and, in the end, fatally, assaulted in the plain view of a whole apartment building in Queens, New York. Despite the public nature of the act, police was not called till long after the assault as on-lookers assumed that someone else must have already alerted the law enforcement.

[2] Note that our information setting is different from the quoted studies - individuals face no false information and have the option to punish only those they directly observe (with exception of player type D in network N[6]).

structure fully determines cooperation and welfare levels. As such, it should become the central element of theories of fragmentation (Burt, 2001; Coleman, 1988) and the main target of measures to improve group performance, diminishing the importance of enforcement structure in such matters.

To illustrate the importance of such questions we can use a practical example: imagine an anonymous evaluation strategy within a company where employees assess the effort of others in their team. This is broadly consistent with a public goods game, where the team effort is the cooperative outcome of interest and negative evaluation is a means of punishment. The team is separated into two sub-teams that seldom meet and, on the whole, displays a rather poor collective performance. The company is facing a decision on how to modify the evaluation strategy in order to improve the team performance. Currently, the employees are allowed to evaluate only those with whom they are in close contact. The company is thinking about two ways to approach this problem: the evaluation could be extended to colleagues that one worked with, even if only briefly (i.e. the other sub-team); or the two teams could be merged into one to increase information flows. The relative viability of solutions focusing either on enforcement possibilities or on information flows is precisely the issue we try to address in this paper.

The findings of this paper can be summed up in three main points. First, severely fragmented information leads to lower cooperation rates, albeit specific kinds of mild fragmentation do not necessarily have to. Secondly, unlimited punishment opportunities in fragmented networks N[1-3] increase enforcement levels, but this does not lead to corresponding increase in cooperation. Finally, welfare levels are generally lower in networks N[1-3] – the additional punishment does not lead to additional cooperation and thus produces welfare losses. To come back to our main questions: information structure proves to be the main determinant of public good provision, irrespective of enforcement restraints. Welfare-wise, aligning enforcement with information flows becomes important, as without restraints there are unnecessarily high punishment levels that create welfare losses.

The rest of this paper is organised as follows: in section 2 we discuss the differences with related literature, and then move on to formally define networks in section 3. Closely related sections 4 and 5 then proceed to describe the experimental game and procedures. Section 6 presents statistical analysis of data collected and the main findings of this paper. Section 7 elaborates on these findings and concludes.

2 Differences with Related Literature

Our experiment builds on the growing stream of literature focusing on social network analysis[3]. It especially contributes to the debate on whether fragmented networks hinder cooperation in general (Coleman, 1988), or whether certain levels of fragmentation might help avoid excessively cohesive, suffocating social structures (Burt, 2001).

More precisely the research presented here is connected to the recent public good game studies that explore the effect of network architecture on cooperation. The most extensive, and particularly enlightening, among these studies was carried out by Carpenter et al (2012). In a systematic analysis of several monitoring networks they identified two structural properties,

[3] For a broad review of network research in economics see for example Jackson (2010) or Goyal (2009); in sociology see Scott (2000) or Scott & Carringdon (2011); and in experimental science Kosfeld (2004).

connectedness and directedness, as important for cooperation and enforcement - if we simplify a bit connectedness leads to more cooperation and directedness to more enforcement. The results for connectedness are of particular interest, since connected networks usually make it easier to reach other individuals. Note however that connectedness does not capture as much information as fragmentation - it is a binary measure that only states that there are no disconnected individuals in a network but does not go into more detail.

Opting for a narrower approach instead, the work of Leibbrandt et al (2012) is also quite close to ours, especially in terms of networks chosen. They find that while disconnected punishment networks differ significantly in terms of cooperation, they result into similar welfare levels after accounting for expenses on enforcement. Further, more loosely related, experimental research shows how a declining fraction of monitored individuals influences public good provision (Carpenter 2007) or how star information structures can sustain cooperation (O'Gorman et al 2009).

We differ from previous literature in two important respects. Firstly, when we impose limits on information networks, we do so without necessarily restraining enforcement in any way. Thus we can study situations in which individuals do not know precisely how all people around them behave, but can punish every single one of them. This enables us to answer questions central to this paper: do information constraints really matter if punishment options are not limited? Can dense punishment networks sustain cooperation despite the problems posed by severely fragmented information (at least in some scenarios)? Such questions bring us closer to recent studies on imperfect information in public goods game (Grechenig et al, 2010; Ambrus & Greiner, 2012), but such research focuses predominantly on the quality of information provided and does not discuss structural issues.

Secondly, we propose a more sensitive network property - fragmentation - to predict network performance. This measure depends principally on the ease with which an individual can reach other individuals in a given network. In contrast to general network properties studied in Carpenter et al (2012), this is a non-binary measure. Thus, instead of general statements that networks with certain property generate more contributions than those without it, we can identify finer relationships: for example, that higher network fragmentation leads to lower contributions.

3 Definition of a network

In our experiment subjects participate in variations of the public good game that differ systematically in the network structure. Before we proceed to the description of the game itself we need to define the term network and show how networks used vary.

We restrict our attention to four-person networks. In general, network is represented graphically as a collection of nodes, indexed by $i = A,B,C,D$, with a single person at each node. A line segment between two nodes indicates that individuals are conneceted by a tie, with an arrowhead showing its direction. For each individual i, we further define neighborhood N_i as a collection of all individuals $j \neq i$ that i is connected by a tie with. The neighborhood describes the relationships of each subject with other subjects, and consequently the collection of individual neighborhoods $N=\{N_A,N_B,N_C,N_D\}$ completely defines all relationships in a network.

Each of our networks consists of two sub-networks. Firstly, the information network defines whose contributions can an individual observe. For each subject i the neighborhood N_i^I denotes the set of subjects $i \neq j$ who can be observed by i. Secondly, the punishment network defines who can an individual punish. For each subject the neighborhood N_i^P denotes the set of subjects $i \neq j$ who can be punished by i. The network formed by $N^I=\{N^I_A,N^I_B,N^I_C,N^I_D\}$ and $N^P=\{N^P_A,N^P_B,N^P_C,N^P_D\}$ thus completely defines the information flows and punishment opportunities in a network.

The networks used in this experiment are shown in Figure 1 above. The network N[7] is the standard public goods game – we call it full-integration benchmark since both information and punishment networks are complete[4]. The networks N[1-6] are presented in a 3*2 matrix that shows systematic variation in:

1. **Fragmentation of an information network –** To measure the fragmentation of a network, we use the following index

$$F = 1 - \frac{\sum_{i>j} \frac{1}{d_{ij}} + \sum_{j>i} \frac{1}{d_{ij}}}{n(n-1)}$$

,where n stands for number of nodes in a networks and d_{ij} for distance[5] between individuals i and j. By defintion, $F=0$ for benchmark N[7] since it is fully integrated. For both networks N[1-3] and N[4-6] F increases as the number of the network gets higher. Thus, we call the networks N[1&5] mildly fragmented scenarios; and networks N[2-3] & [5-6] the harshly fragmented scenarios.

2. **Symmetricity** – a network is symmetric if the sets N^I and N^P are identical, otherwise it is asymmetric. The networks N[4-6] are symmetric since each subject i can punish other subjects if and only if he/she observes their individual contributions. The term asymmetric is reserved for networks N[1-3], that is for networks that have unlimited punishment opportunities but limited information flows .

This variation in the structure of networks gives us an opportunity to rigorously compare how individual behaviour is affected by increasing levels of information fragmentation (by comparing of N[1-3] and N[4-6] to N[7]); and whether such influence differs for networks where either only information or both information and punishment are fragmented (by comparing N[1] to N[4]; N[2] to N[5]; and N[3] to N[6]).

Besides these general network properties, we also classify individual ties into two types:

1. **Multiplex:** A tie from individual i to individual j is multiplex if j belongs to both N_i^I and N_i^P – individual i both observes and can punish individual j.
2. **Simplex:** If j belongs to only one of the N_i^I and N_i^P sets the tie is simplex - in our case simplex ties mark the cases in which an individual can punish others whose contributions he doesn't directly observe.

[4] A network is complete when each player type i is connected by a tie to all other player types $j\neq i$.

[5] Distance essentially measures the reachability between individuals – if you have to force person A to force some other person B to do something, you don't have so much power over person B as if you could force him/her directly. For a detailed discussion of this index see Borgatti (2006).

4 The game

The subjects are first divided into groups of 4 and then participate in a two-stage public good game whose information and punishment structures vary across networks, as described in the section above. In the first stage of the game each subject receives an initial endowment of 25 tokens. Then subjects are asked to simultaneously divide their 25 tokens between public (g_i) and private accounts (x_i) within their respective groups. The payoff of an individual after 1st stage is computed as follows:

$$\pi_i^1 = x_i + 0.4 \sum_{j=A,B,C,D} g_j$$

Focusing on the right hand side of the above equation, the first term stands for the private account and the second for the public account income. If an individual allocates one token to his private account he simply receives one token for it. If instead he decides to allocate one token to the public account, this token is multiplied by 1.6 and divided equally among all participants in the group, so that each one of them receives 0.4 tokens. This makes it optimal from the group perspective to contribute to the public account (tokens get multiplied) and from the individual perspective to contribute to private account (on keeps a whole token for himself).

In the second stage of the game, each individual i observes his/her income π_i^1 resulting from the first stage and the total amount of tokens contributed to the public account. Furthermore, he/she observes individual contributions of group members j who belong to the information neighbourhood N_i^I and has to estimate individual contributions of others who do not. Considering the observed information, individual i can then decide to punish individually other group members k who belong to the punishment neighbourhood N_i^P, also defined by the network structure. He/she can do so by reducing k's first-stage income π_k^1 by p_i^k tokens. This is costly, since i has to pay 1/3 of a token from his own first-stage income π_i^1 to reduce k's income π_k^1 by 1 token. Player i can't spend more than his/hers first-stage income in this way. In addition, no one's income can be reduced below zero. Thus, the income after the second stage of the game is determined by the following equation:

$$\pi_i = \max\left\{0, \pi_i^1 - \frac{1}{3}\sum_{k\in N_i^P} p_i^k - \sum_{j:i\in N_j^P} p_j^i\right\}$$

At the end of the second stage, after punishment has been assigned, each subject i views the amount of tokens deducted from his first stage pay-off and his final second stage earnings π_i^2. Together these two stages form what we call a period.

5 Experimental procedures

In the section above, we describe the basic parameters and proceedings of the game on one period. Each session consisted of 15 identical and independent periods. At the beginning of each session, subjects were seated in front of the computers and received instructions on how to play the game[6]. While reading the instructions, subjects were allowed to ask clarifying questions. On average it took subjects between 15 and 20 minutes to read the instructions and clarifying questions were seldom – 1 to 2 questions per session.

After reading the instructions, subjects were logged into the session on their PCs. At the beginning of the first period, the computer randomly assigned player types A, B, C, D, each type consisting of one fourth of the whole subject population. The player type stayed constant throughout the whole session and determined to which type of node was an individual assigned. The network structure also stayed constant throughout the whole session. In contrast, the personal composition of groups changed over time as, at the beginning of each period, the computer picked one participant of each player type at random to form a group. As a result, the group composition depended solely on chance and was independent of groups from any other period.

After the game was finished, subjects were asked to fill in a short questionnaire[7] – asking for general information about their background and probing their understanding of the instructions by posing 4 comprehension questions. No significant comprehension problems were reported.

After all subjects filled out the questionnaires, their earnings from the experiment were paid out, privately and in cash. Each subject received a 4 EUR participation fee and, in addition, the tokens earned during the experiment were converted into Euros at a rate of 0.018. That means that each of the 25 tokens which every individual was endowed with at the beginning of each period equalled 0,018 EUR. The average total earnings per subject (including participation fee) was approximately 11,76 EUR. To earn this money individuals had to spend on average one hour and a half in the lab.

We ran the experiment with a total of 256 subjects, spread through 11 sessions[8] at an experimental economics lab at the Vysoká Škola Ekonomická in Prague. The whole experiment was programmed in z-TREE. Subjects were recruited through the Online Recruitment System for Economic Experiments (ORSEE), so the subject pool consisted of university students from different countries (predominantly Czech) and with different area of studies (predominantly economics). Subjects had no previous experience with public good game experiments.

6 Expected behaviour

The current microeconomic theory gives us relatively little guidance for predicting the impact of social networks on individual behaviour. Since individuals are randomly re-matched every round there is no incentive for building reputation and self-interested individuals have no interest in 2nd stage punishment ($p_{ij} = 0$). This is rationally expected by the individuals in the 1st stage of the game who choose to fully free ride as it is their individually most profitable strategy ($g_j = 0$). Nash equilibrium is reached at zero cooperation, irrespective of the social structure. Recently, such theoretical predictions have been transformed by incorporating behavioural norms and beliefs to explain the persistent experimental findings of conditionally cooperative behaviour (Falk & Fischbacher, 2006; Fehr & Schmidt, 1999). However, even such updated theoretical models do not readily incorporate the influence of social networks. As a result, we have to turn to experimental and theoretical research focusing directly on social network analysis.

[6] See online resource 2 for example of instructions
[7] The results presented here survive robustness checks based on the data collected in the questionnaire - see online resource 3 for details.
[8] See online resource 1 for more details on subject distribution

Firstly, drawing heavily on social network analysis literature, we identify information network fragmentation as a key influence on punishment/cooperation. The principal hypothesis, based on the "network closure" argument (Coleman, 1988), states that the more fragmented the information network is, the harder it is to create and enforce norms that would sustain cooperation. Alternatively, focusing on the importance of structural holes and weak ties (Burt, 2001; Balkundi et al, 2007), we could predict that moderately fragmented information does not necessarily lead to decline in punishment and cooperation compared to non-fragmented networks. This approach acknowledges the negative effects of severe information fragmentation but claims that extremely cohesive networks suffer from coordination problems in enforcing of cooperation. Moderate fragmentation can help to coordinate effectively without making the emergence and enforcement of cooperative norms too difficult.

Hypothesis 1: *The more fragmented the information is, the less punishment and cooperation the network generates.*

Secondly, we would like to predict how will additional simplex punishment ties in asymmetric networks N[1-3] affect enforcement (and, consequently, contributions and welfare) . Such ties are likely to generate less punishment than similar multiplex ties in network N[7] due to their uncertain nature – if the sum of contributions of unobserved individuals is higher than group average ($\sum_{j\in N_i^P, j\notin N_i^I} g_j > \left(\sum_{k=A,B,C,D} g_k\right)/4$) one cannot punish a free-rider[9] with a 100% certainty[10]. On the other hand, previous evidence from noisy information settings (Ambrus & Greiner, 2011; Grechenig et al, 2010) suggests that individuals are not afraid to engage in punishment even under such risky circumstances.

Hypothesis 2a: *Individuals will use simplex punishment ties in Networks N[1-3] and will do so even when risk of misdirection cannot be avoided.*

Hypothesis 2b: *However, individuals are aware of risks of misdirection and simplex ties will lead to lower punishment per opportunity than multiplex ties.*

Besides that, simplex ties are likely to create coordination problems in punishment (Carpenter, 2004; Carpenter et al, 2012). In symmetric networks N[4-6] each multiplex tie from individual *i* to individual *j* clearly assigns *i* as the sole person responsible for punishment of *j*. In contrast, asymmetric networks N[1-3] have this responsibility diluted by additional simplex ties aimed at individual *j*. Thus, a multiplex tie from symmetric networks should generate more punishment per opportunity than their counterparts from asymmetric settings.

Hypothesis 3: *Due to coordination problems, the multiplex ties in networks N[4-6] will generate less punishment than their counterparts in networks N[1-3].*

---

[9] We define free-rider as an individual whose contribution is lower than network average. Conversely, we define cooperator as an individual whose contirbution is higher or equal to network average.

[10] Conversely, if sum of contributions of unobserved individuals is lower than group average plus maximum contribution ($\sum_{j\in N_i^P, j\notin N_i^I} g_j \geq 25 + \left(\sum_{k=A,B,C,D} g_k\right)/4$), one cannot punish an above average contributor with a 100% certainty.

Finally, the welfare and cooperation effects of additional simplex punishment ties in asymmetric networks are hard to assess. On one hand, simplex ties in networks N[1-3] are likely to generate some additional punishment of free-riders when compared to networks N[4-6]. On the other hand, simplex ties can also lead to significant amount of misdirected punishment (hypothesis 2a) and to reduced use of multiplex punishment ties (hypothesis 3). Furthermore, the reaction to punishment can differ substantially between symmetric and asymmetric treatments. Given the uncertainty and complexity of such issues, we refrain from proposing a clear-cut prediction of how extension of punishment possibilities affects cooperation and welfare.

7 Results

In this section we present the main results of our experiment. Basic findings are described in Figure 1 above, where the mean individual statistics of group performance – contributions, punishment rates and efficiency - are reported for each network. These convey an initial idea about how individual behaviour differs according to alterations in network structure. In what follows we focus on these differences in more detail by studying:

- how fragmented information affects group outcomes;
- how additional punishment opportunities in asymmetric networks affect enforcement;
- how are the effects of fragmentation on group outcomes mitigated/exacerbated in asymmetric networks.

Throughout this section we use non-parametric Man-Whitney two-sample ranksum tests ($|z|$) and standard one-sample t-tests ($|t|$) for mean comparisons. Where necessary we also run random-effects GLS, Probit or Tobit regressions to account for individual heterogeneity, time effects and additional individual controls.

7.1 Information network fragmentation

We begin our analysis by identifying the systematic effects of information network fragmentation, and other general network properties, on enforcement and cooperation. As discussed before we rank networks according to the values of the fragmentation index: network N[7] is the fully integrated benchmark; networks N [1 & 4] are the mildly fragmented scenarios; and networks N[2-3] & [5-6] are the severely fragmented scenarios.

7.1.1 Punishment and fragmentation

Judging from Figure 2(a-b), fragmentation of information networks is negatively correlated with punishment received *per* individual - the severely fragmented networks clearly generate less punishment than others, as predicted by Hypothesis 1. This is particularly true for symmetric networks N[4-6], where even mildly fragmented network N[4] has lower mean punishment than full-integration benchmark ($|z|>3.79$ for all, $p<0.01$); and punishment rates decline steadily with fragmentation index value (for N[4] to N[5-6] comparisons $|z|>1.95$, $p<=0.05$). For asymmetric networks N[1-3], enforcement rates decrease even more abruptly, particularly when we move from mildly fragmented scenario N[1] to the severely fragmented scenarios N[2-3] (all $|z|>2.32$, $p<0.05$).

However, the punishment rates among these networks are not necessarily lower than benchmark N[7] (in fact, they are almost 2 tokens *per* individual higher in N[1]), indicating that there may be some additional positive effects on punishment rates associated with network symmetricity.

Fig.2 Comparison of mean punishment rates *per* individual
(Fig.2 (a) for asymmetric networks; Fig.2 (b) for symmetric networks)

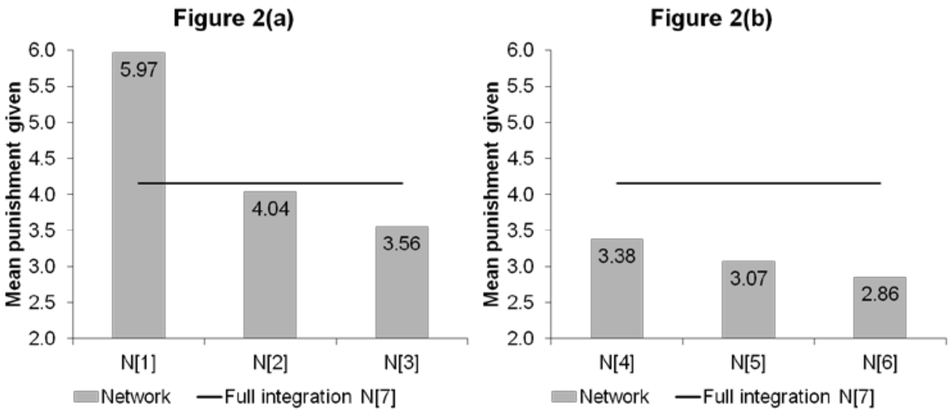

Indeed, by studying Figure 2(a-b) more closely, we see a striking difference between the symmetric and asymmetric networks. In the latter individuals receive on average 1.4 extra tokens of punishment than in the former (|z|=11.32, p<0.01), which broadly confirms that people are not afraid to use the unlimited punishment opportunities as stated in Hypothesis 2a (see section 6.2 for details).

Besides fragmentation and symmetricity, the possible effects of other important properties of information networks need to be assessed. Directedness[11] seems to be of importance, as the higher punishment levels in directed networks N[1-2] & N[4-5] than in the non-directed ones (N[3 & 6]) imply a significant positive influence on enforcement (|z|=1.99, p<0.05). The total number of information ties in a network, on the other hand, does not seem to have any systematic relationship with punishment - for example, the network N[1] with four ties leads to more punishment than both full-integration benchmark with 12 ties and N[2] with three ties. [12]

7.1.2 Cooperation and fragmentation

Severely fragmented information in networks N[2-3]&[5-6] not only lowers punishment but also hinders cooperation. Looking at Figure 3(a-b), the best performer among these networks (N[3]) yields 1.43 tokens less than the mean contributions of 14.82 tokens in benchmark network (|z|=2.26, p<0.05). Other comparisons are even less favourable with the extreme of network N[3] trailing behind by approximately 5 tokens - more than 30% of average amount contributed - behind benchmark (|z|=10.08, p<0.01). This confirms that fragmentation has negative influence on cooperation as predicted in Hypothesis 1. However, there are noteworthy exceptions. Mildly fragmented networks N[1&5] do not generate lower contribution levels than full-integration benchmark, suggesting that strong social cohesion is not always beneficial for cooperation.

[11] In directed networks all ties only have one arrowhead implying their direction, otherwise the network is non-directed. In terms of information N[1] is an example of a directed and N[3] of a non-directed network.
[12] These network properties were chosen according to previous evidence from Carpenter et al (2010) and Carpenter(2004).

Fig.3 Comparison of mean contribution rates
(Fig.3 (a) for asymmetric networks; Fig.3 (b) for symmetric networks)

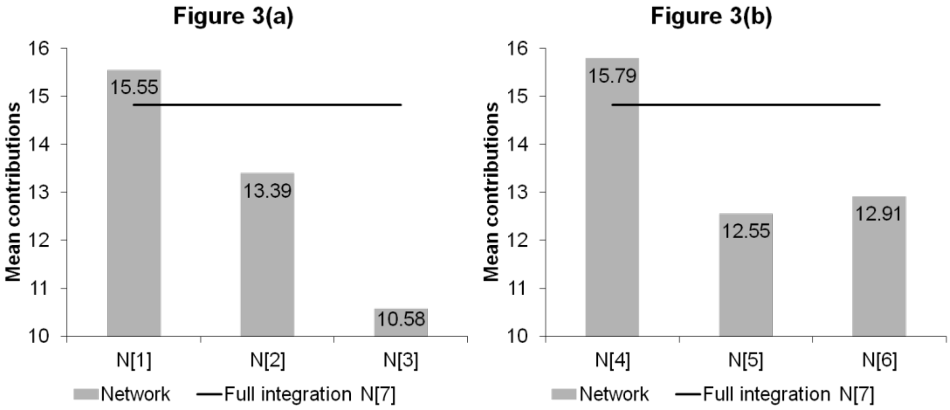

The influence of directedness seems similar as in the case of punishment. The directed networks N[1-2]&N[4-5] outperform the non-directed ones (N[3&6]) in cooperation, with mean contribution of 14.41 tokens for the former and 12.69 tokens for the latter($|z|=6.73$, $p<0.01$). The general result hides some heterogeneity though, with directed network N[5] generating less contributions than non-directed N[6]. This is at least partially caused by the fact that players at isolated node D in N[5] can behave selfishly with no one there to punish them for it.[13]

Finally, symmetricity affects enforcement and cooperation in different ways. While asymmetric networks have much higher punishment rates per individual, no such effect is transmitted onto contributions. Indeed, symmetric networks actually achieve marginally higher contribution rates (by 0.7 tokens per individual; $|z|>2.20$, $p<0.05$) than the asymmetric ones, the difference being particularly dramatic in comparison of networks N[6] and N[3]. This has serious negative implications for welfare levels in asymmetric treatments since punishment is costly.

7.1.3 General network properties and performance

To explore the findings from this section in a more robust way, we present results from regression analysis that controls for general network properties in Table 1. To be more exact, we include a dummy for directed networks, a dummy for asymmetric networks and the values of fragmentation index as our explanatory variables. Column I confirms our previous results for punishment: information fragmentation has a sizeable negative influence on punishment rates, while both directedness and asymmetry have significant positive effects. If we focus on cooperation instead in column II, we see similar results with one exception. The symmetricity dummy becomes insignificant, indicating that additional punishment opportunities in asymmetric networks do not facilitate emergence of norms that would lead to higher public good provision. Consequently, due to higher levels of costly punishment, asymmetric networks suffer from significant welfare losses, as shown in column III.

[13] Players at node D contribute only 7.94 tokens on average, compared to 14.05 tokens contributed by players at other nodes in N[5]. Adding simplex ties makes node D players less selfish, as in N[2] they contribute on average 10.05 tokens. Unfortunately, this disciplining effect is not significant due to low number of observations.

Table 1: The effects of general network properties on individual behaviour			
	(I)	(II)	(III)
Dependent variable	Punishment received	Contribution	Profit
Fragmentation	-2.936***	-5.405***	0.538
	(0.738)	(1.401)	(1.230)
Directed	0.587*	3.539***	0.454
	(0.342)	(0.659)	(0.548)
Asymmetric	1.471***	-0.451	-2.290***
	(0.378)	(0.655)	(0.583)
Constant	4.238***	15.031***	28.294***
	(0.506)	(0.857)	(0.823)
Includes period fixed effects	Yes	Yes	Yes
Observation type	Per individual	Per individual	Per individual
Rho	0.09	0.53	0.16
Prob > Chi^2	<0.01	<0.01	<0.01
# Observations	3840	3840	3840
# Subjects	256	256	256

Note: Random-effects (on individual level) GLS regressions, Tobit to control for upper level of contributions (25) and lower limit of profits (0). Dependent variables are Punishment received (tokens deducted), Contribution (tokens contributed) and Profit (tokens earned). Standard errors clustered at individual level are reported in parentheses.* significant at 10%; ** significant at 5%; *** significant at 1%

7.2 Enforcement and symmetricity

In the previous section we saw that asymmetric networks generate substantially higher punishment than their symmetric counterparts - on average 1.18 extra tokens received per individual (Table 1 Column I). To understand this difference we need to focus on how additional simplex opportunities affect behaviour. To reach this level of detail we need to focus on punishment given *per* opportunity in the rest of the section 7.2[14].

7.2.1 Use of simplex ties

In agreement with our prediction in Hypothesis 2a, individuals[15] are not afraid to use simplex punishment ties in networks N[1-3]. As can be seen from Figure 4(a) they deduct a mean of 1.52 tokens per opportunity in network N[1]; 1.14 tokens in network N[2]; and 0.91 tokens in network N[3], with all these means significantly higher than zero (all $|t| > 11.32$; $p<0.01$). In addition, they are not only willing to use simplex ties often, but also willing to use them in situations where they cannot be absolutely certain whether they are punishing a co-operator or a free-rider[16]. In network N[3], with the lowest simplex punishment mean, over 57.5% of total amount of simplex punishment happens under such uncertain circumstances. In the other two networks, this proportion is even higher – 72.4% in N[1] and 62.7% in N[2].

[14] Throughout the section 7.2 focused on enforcement we exclude players at node D in network N[5], since they cannot punish nor be punished (with exception of Table 3 in section 7.3.1 dealing with punishment distribution).

[15] For relevance of comparison we exclude players at node D in network N[2], since they have no multiplex ties. They punish relatively infrequently and weakly compared to others.

[16] That is when $\left(\sum_{k=A,B,C,D} g_k\right)/4 \leq \sum_{j \in N_i^P, j \notin N_i^I} g_j < 25 + \left(\sum_{k=A,B,C,D} g_k\right)/4$

Fig. 4 Comparison of punishment *per* opportunity against benchmark N[7] (panel (a) – simplex punishment; panel (b) – multiplex punishment)

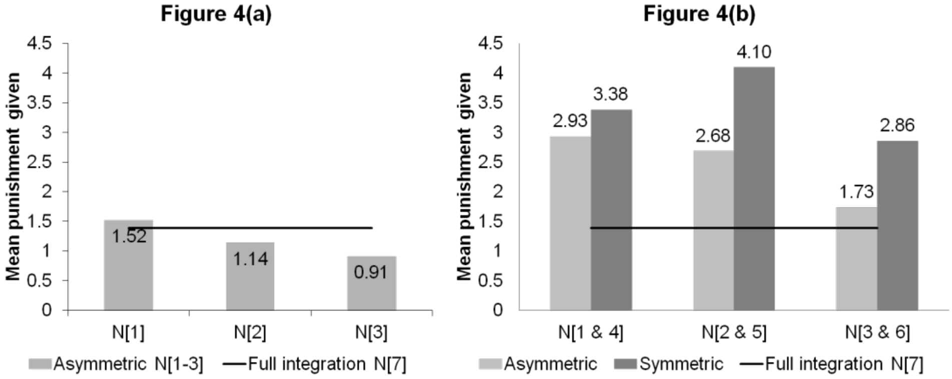

7.2.2 Multiplex versus simplex ties

Even though individuals often punish using simplex ties they still prefer multiplex opportunities with precise information about target's contribution, as can be seen from Figure 4(a-b). To be more exact, in asymmetric networks simplex ties generate approximately half the punishment of multiplex ties (all $|z|>3.17$; $p<0.01$) and the difference becomes even more exacerbated for comparisons with multiplex ties from symmetric networks. This finding confirms the hypothesis 2b that people are aware of the risks of misdirection associated with imprecise information. It also suggests that simplex punishment aggravates punishment coordination issues[17].

Focusing on those issues in more detail, we see from Figure 4(b) that multiplex ties generate more punishment in asymmetric networks than in the full-integration benchmark (significant for N[1-2]; for both $|z|> 2.78$, $p<0.01$). This suggests that changes solely in information networks can alleviate coordination problems and increase *per* opportunity multiplex punishment - people are clearly willing to assume more responsibility for punishment of directly observed individuals when these are not directly observed by anyone else. However, multiplex punishment is still significantly lower in asymmetric than in symmetric networks (all $|z| > 4.24$; $p<0.01$). Thus, changes in information networks are not likely to completely solve coordination issues unless accompanied by corresponding punishment restraints.

To corroborate and extend findings from this section, we use random-effects GLS regressions to compare fragmented networks to omitted full-integration benchmark. In Table 2 column I we see that asymmetric networks N[1-3] have rather similar punishment *per* opportunity rates as the complete network. This masks an important heterogeneity that is explored in column II by including Multiplex dummy equal to 1 for multiplex punishment ties and 0 otherwise. The highly significant coefficient of this dummy shows that individuals deduct on average 1.2 tokens more when using a multiplex tie or, in other words, when punishing an individual whose contribution they can directly observe. Results in column II also reaffirm that fragmented networks in general, and the symmetric ones in particular, reduce coordination problems and increase multiplex punishment - all dummies for networks are positive, significant (except N[2]) and higher for symmetric networks.

[17] When an individual *i* decides whether to punish individual *j*, he takes into consideration whether *j* can be punished by others for two reasons - firstly, he can just selfishly leave the punishment to others; secondly he may consider lowering his punishment if he expects others to punish *j*.

Table 2: Comparing punishment across symmetric and asymmetric networks			
	(I)	(II)	(III)
Dependent variable	Punishment given	Punishment given	Punishment given
Network N[1]	0.605	1.414***	1.030**
	(0.426)	(0.486)	(0.440)
Network N[2]	-0.039	0.872*	0.364
	(0.363)	(0.462)	(0.415)
Network N[3]	-0.200	0.610	0.462
	(0.339)	(0.405)	(0.352)
Network N[4]	1.997***	1.997***	1.847***
	(0.643)	(0.643)	(0.513)
Network N[5]	2.712**	2.712**	2.023**
	(1.066)	(1.066)	(0.845)
Network N[6]	1.473***	1.473***	1.151**
	(0.531)	(0.531)	(0.474)
Multiplex		1.214***	-0.137
		(0.273)	(0.209)
Signal-contribution			-0.167***
			(0.031)
Signal-contribution>0			0.170***
			(0.038)
Sure pun*Signal-cont			-0.261***
			(0.049)
Sure pun*Signal-cont>0			0.303***
			(0.060)
Constant	1.352***	0.137	-0.441
	(0.308)	(0.428)	(0.377)
Includes period fixed effects	Yes	Yes	Yes
Observation type	Per opportunity	Per opportunity	Per opportunity
Rho	0.24	0.24	0.22
Prob > Chi^2	<0.01	<0.01	<0.01
# Observations	8535	8535	8535
# Subjects	249	249	249

Note: Random-effects (on individual level) GLS regressions. Dependent variable is Punishment given (tokens deducted). Standard errors (clustered at individual level) are reported in parentheses.* significant at 10%; ** significant at 5%; *** significant at 1%

Finally, we analyse whether simplex and multiplex punishment has similar motivation. In full information settings it has been shown that whether an individual *i* punishes an individual *j* depends on the difference between their contributions g_j-g_i (e.g. Herrmann et al, 2008). In the case of simplex punishment, where such a comparison of individual contributions is not possible, we have to substitute *i*'s estimate of *j*'s contribution for *j*'s actual contribution. Consequently, the difference that should drive the punishment of individual j becomes $E_i\left(g_j\right)$ - g_i.

Therefore, in column III we add the explanatory variable Signal-Contribution that equals g_j-g_i for multiplex ties and $E_i\left(g_j\right)$ - g_i for simplex ties. In addition, to control separately for punishment of free-riders and co-operators, we add variable Signal-Contribution>0 which equals Signal-Contribution if positive and is 0 otherwise. The coefficients of these variables imply that the difference between own and target's estimated contribution clearly matters for simplex punishment of both below and above average contributors. However, this connection is much weaker than for multiplex punishment, as is apparent from the highly significant coefficients of interaction terms between the Multiplex, Signal-Contribution and Signal Contribution>0 regressors. In other words, the difference between own and target's (estimated) contribution drives both simplex and multiplex punishment, but the sensitivity is much lower for the former. This difference in sensitivity fully explains the lower punishment rates of simplex ties, since the Multiplex dummy becomes insignificant.

7.3 Unrestricted enforcement and collective performance

To recap briefly, people are aware of the risks involved in punishment without precise information and strongly prefer to punish when such information is available. Despite that, they are still willing to use simplex ties for punishment and to do so even under very uncertain circumstances. This results in higher individual enforcement efforts in asymmetric networks than in their symmetric counterparts, which may in turn lead to different contribution rates.

The fact that it does not, deserves an explanation. We provide on by asking two closely related questions. The first is whether punishment in asymmetric networks targets more people who can significantly increase their contributions (free-riders). If not, significant rise in cooperation is unlikely even for higher punishment levels. The second is whether people react to punishment differently across symmetric and asymmetric networks. If reaction to received punishment is lower in the latter type, increased punishment need not result into higher contributions either.

7.3.1 Punishment Distribution

Judging from Figure 5, punishment is not particularly well distributed in asymmetric networks. The blame falls partially on the risky nature of simplex punishment, prone to misdirection – a total bulk of 36.4% tokens deducted is assigned to co-operators (0.4 tokens *per* opportunity), leaving only 64.4% of punishment to target free-riders (0.7 tokens *per* opportunity). Such distributional concerns get particularly grave in network N[1], where 45.6% of simplex punishment hits co-operators.

Fig. 5 – Proportion of punishment against co-operators

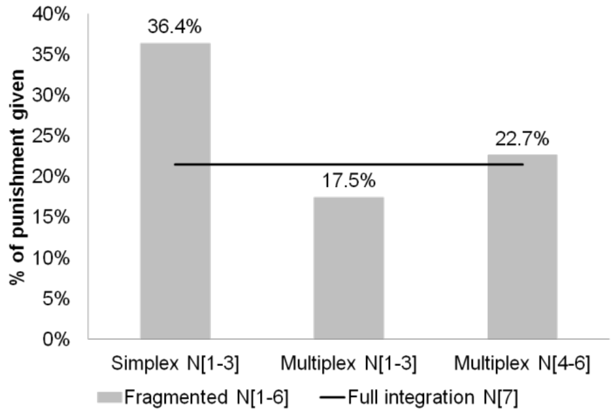

If we focus on multiplex punishment instead, we see that it falls much heavier on free-riders with approximately 80% of it directed towards players with below average contributions. In terms of distribution, symmetric networks N[4-6] gain an advantage since they have higher multiplex punishment *per* opportunity than their asymmetric counterparts (see section 7.2.2). Coming back to Figure 4(b), this affects particularly the networks N[2-3] where multiplex punishment is lower by more than 1 token *per* opportunity than in their symmetric counterparts N[5-6].

These results are further corroborated by results of random-effects GLS regression presented in Table 3. In column I, we limit our analysis to punishment received per individual with contributions equal to or above group average ($g_i \geq \left(\sum_{j=A,B,C,D} g_j\right)/4$). Controlling for individual contributions g_i, we see that network N[1] generates significantly higher levels of punishment directed against co-operators than any other network. This hints at why this network does not generate higher contributions than its symmetric counterpart N[4], in which we observe no such trend.

Table 3: Comparing rates of punishment against free-riders and co-operators

Dependent variable	(I) Punishment received (≥ contribution average)	(II) Punishment received (< contribution average)
Network N[1]	1.704***	2.145**
	(0.569)	(0.864)
Network N[2]	0.354	-2.069**
	(0.510)	(0.822)
Network N[3]	-0.408	-3.165***
	(0.412)	(0.642)
Network N[4]	-0.322	-0.824
	(0.496)	(0.714)
Network N[5]	-0.377	-3.780***
	(0.511)	(1.085)
Network N[6]	-0.247	-4.106***
	(0.475)	(0.897)
Contribution	-0.085***	-0.434***
	(0.025)	(0.038)
Constant	3.539***	10.132***
	(0.517)	(0.815)
Includes period fixed effects	Yes	Yes
Observation type	Per opportunity	Per opportunity
Rho	0.04	0.10
Prob > Chi^2	<0.01	<0.01
# Observations	2053	1787
# Subjects	243	243

Note: Random-effects GLS regressions. Dependent variable is Punishment received (the sum of tokens received by an individual). Column I - individual contributions higher or equal to group average; Column II - individual contributions below group average. Standard errors (clustered at individual level) are reported in parentheses.* significant at 10%; ** significant at 5%; *** significant at 1%

Turning our attention to individuals with below average contributions ($g_i < \left(\sum_{j=A,B,C,D} g_j\right)/4$), we see much more heterogeneity (see column II). Firstly, mildly fragmented networks N[1] & N[5] do not generate lower punishment than full-integration benchmark. In fact, the N[1] dummy is significantly positive, indicating why this network reaches high cooperation rates even despite its high level of punishment against co-operators. Secondly, all severely fragmented networks N[2-3] & N[5-6] lag behind benchmark N[7] in punishment of free-riders, as the negative coefficients of corresponding dummies demonstrate. The effect is relatively less pronounced for the N[2] network, in which the punishment received *per* free-rider is only 2 tokens lower than in benchmark N[7][18].

[18] This is possibly caused by the specific role of player type D in this network. Since player type D is completely unobserved, others can easily assume that he contributes less than others. To the extent that this is true, this can improve punishment distribution. Indeed, type D individuals in this network have lower contributions than others and are targeted more often by simplex punishment.

This possibly makes the network less disadvantaged in cooperation enforcement than other severely fragmented networks.

7.3.2 Reaction to punishment

To see how enforcement distribution affects cooperation, we must look at how subjects change their contributions in response to received punishment. To do so, we employ random effect Tobit regressions whose results are presented in the first two columns of Table 4 below. The dependent variable is the change in contributions from previous round (g_i^t - g_i^{t-1}) and the explanatory variables include network dummies and Lag Punished, which equals 1 if individual was punished in the previous round and 0 otherwise[19]. We also control for lagged contributions and profits, but do not report results of these controls since they are not of direct interest.

Table 4: Comparing contributions and profits across treatments

	(I)	(II)	(III)	(IV)
Dependent variable	ΔContribution (above avg)	ΔContribution (below avg)	Contribution (Period >5)	Profit
N [1]	0.655	0.255	0.700	-2.123**
	(0.866)	(0.835)	(1.339)	(0.950)
N [2]	0.540	-0.916	-2.810**	-0.866
	(0.903)	(0.879)	(1.238)	(0.980)
N [3]	-1.693*	-1.545*	-4.582***	-1.927**
	(0.876)	(0.805)	(1.369)	(0.928)
N [4]	1.317	-0.910	2.640**	1.438
	(0.923)	(0.904)	(1.246)	(1.009)
N [5]	1.308	-1.377	-3.970***	-0.097
	(0.987)	(0.923)	(1.207)	(1.058)
N [6]	0.037	-1.128	-3.842***	0.441
	(0.900)	(0.872)	(1.156)	(0.978)
Lag Punished	-0.677**	1.773***		
	(0.307)	(0.317)		
Constant	-1.633	0.377	18.396***	28.773***
	(1.316)	(1.364)	(1.013)	(0.692)
Includes period fixed effects	Yes	Yes	Yes(Period>5)	Yes
Observation type	Per individual	Per individual	Per individual	Per individual
Rho	0.24	0.23	0.60	0.16
Prob > Chi^2	<0.01	<0.01	<0.01	<0.01
# Observations	1925	1659	2560	3840
# Subjects	241	238	256	256

Note: Random-effects regressions on individual level (Tobit used to control for Contribution, ΔContribution and Profit limits). Dependent variables are Contribution (tokens to public account), ΔContribution (change in contribution from previous round) and Profit (tokens earned). Standard errors reported in parentheses.* significant at 10%; ** significant at 5%; *** significant at 1%

In column I, we limit observations to individuals whose contributions were above the group average in the previous round of the game ($g_i^{t-1} \geq \left(\sum_{j=A,B,C,D} g_j^{t-1}\right)/4$). By doing so, we find that they reduce their contributions significantly after being punished in the previous round. Punishment of co-operators thus weakens cooperation and is wasteful in terms of welfare. When we focus on free-riders ($g_i^{t-1} < \left(\sum_{j=A,B,C,D} g_j^{t-1}\right)/4$) instead in column II, we find quite the opposite reaction. The

[19] We do not use the exact amount of tokens deducted since such modelling is unnecessarily complicated for our purposes.

coefficient for Lag Punished is significant and highly positive – free-riders are quite sensitive to being punished and increase their contributions in response. In turn, such enforcement is likely to be vital to sustaining high contribution levels. The last thing of note is that in network N[3] everyone has marginally stronger tendency to lower contributions over time ($p<0.1$), which helps explain why this network reaches very low cooperation.

7.3.3 Impact on Cooperation

As noted above in section 7.1.2, severely fragmented information leads to sharp decline of public good provision. One of the principal questions of our paper is whether the expansion of punishment opportunities can somehow reverse this trend. Despite the increased levels of enforcement in asymmetric networks, no such reversion occurs due to the serious distributional concerns identified above. This finding holds irrespective of network fragmentation, albeit the causes differ between mildly and harshly fragmented networks:

Both mildly fragmented networks N[1&5] generate a lot of punishment against free-riders, which significantly improves cooperation. The social punishment is particularly high in asymmetric N[1], implying higher contribution levels than in N[5]. On the other hand, N[1] also stimulates the extremely high levels of antisocial punishment, which weakens contributions (see Table 3 for corroboration). These contradictory influences cancel each other out. In the end, both mildly fragmented networks reach similar levels of cooperation as benchmark N[7], as can be seen from Figure 6(a & d). The fact that N[1] doesn't outperform N[5] receives further support in column III of Table 4, which reports results of a Tobit regression comparing contributions over the last ten periods.

Fig. 6 Comparisons of contributions across networks

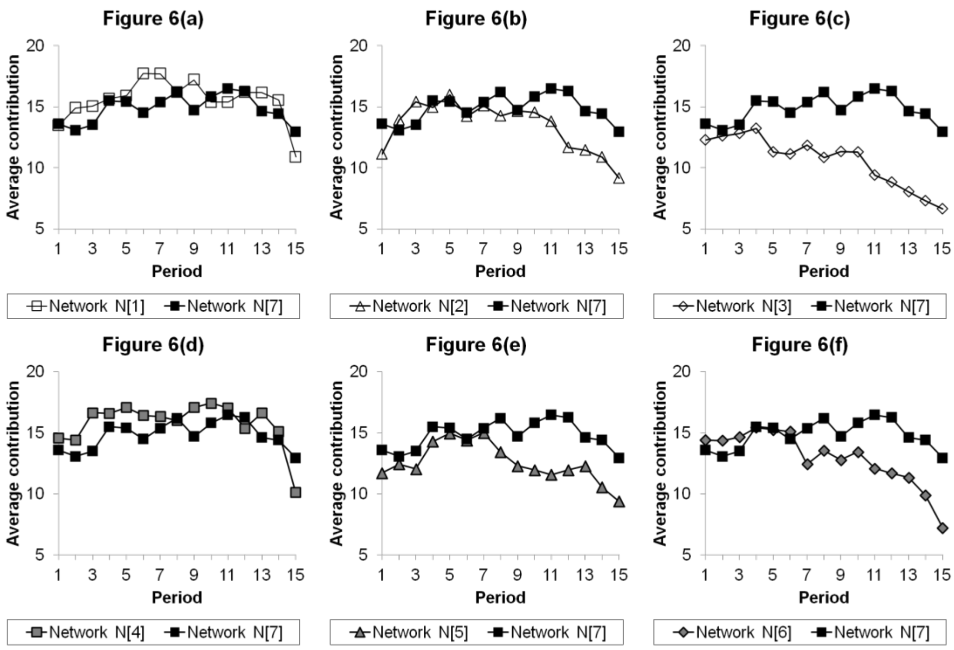

The cases of severely fragmented networks N[2-3] & [5-6] are simpler to explain. All these networks generate less punishment against free-riders than full-integration benchmark (see Table 3 column II for details). The lower levels of social punishment in turn cannot sustain high levels of cooperation – in every single network we can observe sharp decline in contributions over the last ten periods in Figure 6(b-c) & (e-f). This decline is also illustrated by results of a Tobit regression in column III of Table 4, with all coefficients of severely fragmented networks being significant and highly negative. The dummy for network N[2] shows that this network suffers a relatively milder contribution decline, probably due to higher levels of punishment against free riders (see section 7.3.1).

7.3.4 Impact on Welfare

Since higher punishment rates do not elicit significantly higher contributions, asymmetric networks perform poorly in terms of welfare. This is true both for comparisons with their symmetric counterparts and with the full integration benchmark. To see that notice Figure 7 below, which reports the possible welfare gains realized through cooperation[20]. Asymmetric treatments lead to consistently lower welfare gains than other networks, with network N[1] and N[3] reaching respectively only 10% and 11% of maximal gains from cooperation. The effect is less pronounced in network N[2] with 18% of gains realized, probably due to slightly more favourable enforcement distribution in this network (see section 7.3.1).

Fig. 7 Comparison of welfare gains across networks

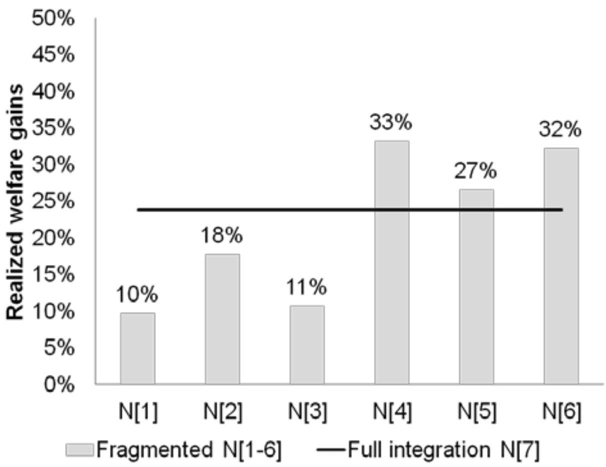

These findings receive support from analysis of individual profit rates. Asymmetric networks never reach higher average profit rates than 27.68 tokens per round, which does not compare favourably to their symmetric counterparts nor to the full-integration benchmark (at least 28.57 tokens; significant for N[1 & 3]: all $|z|>3.26$, $p<0.01$). These differences hold even after accounting for individual heterogeneity and time effects in a Tobit random-effects regression – in column IV of Table 4 networks N[1 & 3] have significantly lower profit rates than benchmark N[7]; and all symmetric networks have consistently higher coefficients than their asymmetric counterparts.

[20] Welfare gains for individual *i* are defined as the additional earnings compared to no-cooperation scenario divided by the additional earnings in a full-contribution scenario ($\frac{\pi_i^2-25}{40-25}$)

8 Concluding remarks

Returning to the principal matter of interest of this paper, the relative importance of information and enforcement structures for collective outcomes, we can distil three important insights from our results.

Firstly, information structure clearly influences enforcement and public good provision. Both punishment and contributions levels decline with fragmented information in symmetric networks N[1-3] and their asymmetric counterparts N[4-6]. This lends some support to the network closure argument that cohesive structures facilitate norm enforcement, which is a form of social capital favourable to cooperation (Coleman, 1988). However, there is an important caveat to this finding – notice that mild fragmentation in directed networks N[1&5] does not significantly lower cooperation rates compared to fully integrated benchmark. The original network closure argument thus needs to be refined in the way proposed by Burt (2000 & 2001): While too loose networks are indeed harmful for public good provision, at least certain types of moderately fragmented social structures (e.g. those leading to directedness) do not necessarily need to be. This is because they can alleviate the problems with coordination of punishment against free-riders in extremely cohesive network N[7] and thus enhance norm emergence and enforcement.

Secondly, extension of punishment opportunities leads to higher punishment even when information stays limited. This happens because people are not afraid to use simplex punishment opportunities, that is opportunities without individual information about target's contribution, in asymmetric networks N[1-3]. While, at first glance, such an increase in enforcement should facilitate emergence of norms promoting cooperation, it creates grave distributional concerns instead. In some cases simplex punishment is poorly aimed (N[1]) and, in others, it creates significant coordination problems that reduce multiplex punishment (N[2-3]). As a result, no additional public good provision is achieved. This broadly ties into recent experimental research on quality of information (Ambrus & Greiner, 2012; Grechenig et al, 2010) that has shown that dubious information quality seriously undermines effectiveness of individual enforcement. We extend the force of this argument by similar results for settings without false or confusing signals, where individuals can just stick to punishing people they directly observe.

Thirdly, comparing the relative importance of information and enforcement structure for cooperation, we observe that information plays a much larger role in determining contributions. The cooperative outcomes are very similar for networks with identical information structure, irrespective of the punishment structure – mildly fragmented information networks N[1&4] display levels of contributions similar to full integration benchmark N[7], while severe information fragmentation in networks N[2-3] & [5-6] prevents high levels of cooperation to be sustained. If anything, overcoming enforcement fragmentation has negative effects on cooperation as in the case of network N[3] compared to N[6]. Thus, we identify fragmented information as a primary cause of low public good provision[21], implying that information structure should play a particularly important role in theories predicting social capital of different networks (Burt 2000, Coleman 1988).

Finally, there are severe welfare losses associated with unrestricted enforcement in asymmetric networks, since higher punishment rates do not improve cooperation. Such negative impacts of additional punishment are mitigated only in network N[2], probably because the

[21] Not to overstate our findings: enforcement structure clearly has its role, since without enforcement cooperation cannot be sustained. What we want to point out here is that once an information structure proves out to be inappropriate for public good provision, no enforcement structure can overcome its problems.

extension of punishment allows disciplining of otherwise unreachable individual at node D. However, even in this case, the comparisons of networks N[2 & 3] or N[2 & 7] show persistent mildly negative welfare effects (albeit these become insignificant in statistically robust regression analysis).

Our results lend support to two types of actions to stimulate collective performance in fragmented social structures: 1) improving information flows within social networks, such as integration of isolated minorities into public debates or promotion of communication across work teams, is likely to improve cooperation; 2) favouring enforcement based on precise information about individual behaviour, such as creation of specialised institutions (courts) for information collection and law enforcement, is likely to increase welfare. Conversely, one should be very careful when trying to solve information deficiencies solely by extending enforcement structure, as this can potentially have grave negative welfare consequences.

Acknowledgements: We are grateful to Michal Bauer, Peter Katuscak, Andrew Schotter, Arno Riedl, Cary Deck, Miroslav Zajicek, Margarida Silva, Ian Levely, Lubos Cingl, Jana Cahlikova and Tomas Miklanek for their comments and suggestions. We would also like to thank participants of GDN 2012 Research Workshop and CERGE-EI brown-bag seminars for further insightful comments. This research was supported by a grant from the CERGE-EI foundation under a program of the Global Development Network. All opinions expressed are those of the author(s) and have not been endorsed by CERGE-EI or the GDN. We also gratefully acknowledge the financial support from Nadace Sophia, as well as the administrative support of Charles University, Faculty of Social Sciences. Finally, the helpfulness of Vysoka Skola Ekonomicka in providing access to its economic laboratory must be mentioned. The first draft of this paper was used as my Diploma thesis at Charles University, see http://ies.fsv.cuni.cz/work/index/show/id/1690/lang/en.

References

Ambrus, A., & Greiner, B. (2012). Imperfect Public Monitoring with Costly Punishment - An Experimental Study. *American Economic Review* , Vol. 102(7), pp 3317-32.

Balkundi, P., Kilduff, M., Barsness, Z. I., & Michael, J. H. (2007). Demographic Antecedents and Performance Consequences of Structural Holes in Work Teams. *Journal of Organizational Behaviour* , Vol. 28, pp 241-260.

Borgatti, S. P. (2006). Identifying Sets of Key Players in a Social Network. *Computational & Mathematical Organization Theory* , Vol. 12, pp 21-34.

Burt, R. S. (2001). Structural Holes versus Network Closure as Social Capital. In R. S. Burt, K. Cook, & N. Lin, *Social Captial: Theory and Research* (pp 31-56). New Jersey: Transaction Publishers.

Burt, R. S. (2000). The Network Structure of Social Capital. *Research in Organizational Behaviour* , Vol. 22, pp 345 - 423.

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

Carpenter, J., Schotter, A., & Kariv, S. (2012). Network Architecture and Mutual Monitoring In Public Goods Experiments. *Review of Economic Design* , Vol. 16(2-3), pp 93-118.

Coleman, J. S. (1988). Social Capital in the Creation of Human Capital. *American Journal of Sociology* , Vol. 94, pp S95-S120.

Falk, A., & Fischbacher, U. (2006). A Theory of Reciprocity. *Games and Economic Behaviour* , Vol. 54, pp 293–31.

Fehr, E., & Gächter, S. (2000). Cooperation and Punishment in Public Goods Experiments. *American Economic Review* , Vol 4., pp 980-994.

Fehr, E., & Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics* , Vol. 114, No. 3, pp 817-868.

Gächter, S., Renner, E., & Sefton, M. (2008). The long-run benefits of punishment. *Science 322* , 1510.

Goyal, S. (2009). *Connections: An Introduction to the Economics of Networks.* Princeton University Press.

Granovetter, M. (1985). Economic Action and Social Structure: The Problem of Embeddedness. *The American Journal of Sociology* , Vol. 91, No. 3, pp 481 - 510.

Granovetter, M. (1973). The Strength of Weak Ties. *The American Journal of Sociology* , Vol. 78, No. 6, pp 1360-1380.

Grechenig, K. R., Nicklisch, A., & Thöni, C. (2010). Punishment Despite Reasonable Doubt. *Journal of Empirical Legal Studies* , Stud. 847.

Herrmann, B., Gächter, S., & Thöni, C. (2008). Antisocial Punishment Across Societies. *Science 319* , pp 1362 - 1367.

Jackson, M. O. (2010). *Social and Economic Networks.* Princeton University Press.

Kosfeld, M. (2004). Economic Networks in the Laboratory. *The Review of Network Economics*, Vol. 3(1) , pp 20 - 42.

Ledyard, J. O. (1994). Public Goods: A Survey of Experimental Research. *Econ WPA - Public Economics* , nr. 9405003.

Leibbrandt, A., Rmalingam, A., Sääksvuori, L., & Walker, J. M. (2012). Broken punishment networks in public goods game: experimental evidence. *Jena economic research papers 004* .

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* , Vol.27, pp 415–44.

O'Gorman, R., Henrich, J., & Van Vugt, M. (2009). Constraining Free Riding in Public Goods Games: Designated Solitary Punishers can Sustain Human Cooperation. *Proceedings of The Royal Society*, Vol. 276(1655), pp 323-329

Scott, J. (2000). *Social Network Analysis: A Handbook (Second Edition).* Sage Publications.

Scott, J., & Carringdon, P. J. (2011). *The SAGE Handbook of Social Network Analysis.* Sage Publications.