Journal of Behavioral and Experimental Economics 94 (2021) 101742

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Economics

journal homepage: www.elsevier.com/locate/jbee

Conspiracy against the public - An experiment on collusion

Åshild A. Johnsen [1], Ola Kvaløy [2,*]

[1] Frisch Centre, Oslo
[2] University of Stavanger Business School

ARTICLE INFO

Keywords:
Infinitely repeated prisoner's dilemma game
Negative externality
Cooperation
Collusion
Experiment

ABSTRACT

We study to what extent collusive behavior is affected by the awareness of negative externalities. Theories of outcome-based social preferences suggest that negative externalities make collusion harder to sustain than predicted by standard economic theory, while sociological theories of social ties and intergroup comparisons suggest that bilateral cooperation can be strengthened if there exist outsiders that gain from cooperative break down. We investigate this in a laboratory experiment. Subjects play the infinitely repeated prisoner's dilemma with and without a negative externality. The externality is implemented by letting subjects make a positive contribution to a public good if they choose to deviate from cooperation between the two, i.e. cooperation is collusive since the gains are at the expense of the public. We find that this negative externality tends to increase collusive behavior. Initially, the level of cooperation is lower, but as subjects gain experience and observe that their partners choose to cooperate despite the negative externality, they cooperate as least as much as in the baseline treatment.

1. Introduction

There are two main motives for cooperative behavior in repeated game prisoner's dilemma. First, it may be profitable: Cooperation builds reputation and can thus be sustained as an equilibrium even if parties are selfish and fully rational (see Dal Bó and Fréchette (2018) for a survey). Second, it may satisfy social preferences (e.g. Andreoni and Miller, 1993). People sometimes reciprocate cooperative behavior even if it has material costs.

However, the effect of social preferences on cooperative behavior is not so straightforward in many real-world social dilemmas. The reason is that cooperation sometimes yields negative externalities. Collusion is an important example. Firms can cooperate on price increases or quantity restrictions in order to achieve higher profits, but this comes at the public's expense.

One question is how these negative externalities imposed on the public affect the behavior of colluding parties. Prominent theories of social preferences do not give a clear answer. Theories of outcome-based social preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) suggest that collusion potentially may be harder to sustain. If people put positive weight on the utility of third parties, breaking the collusive partnership becomes less costly. Theories of intention-based reciprocity, however (Rabin, 1999; Dufwenberg and Kirschsteiger, 2004), suggest that negative externalities may not reduce cooperation even if people care about third parties, since they may also care about the intentions behind the actions. A negative outcome for a third party may potentially even give positive signals to a collusive partner.

Sociological theories explicitly open for this conjecture. Theories of social ties, in-group and out-group trust and intergroup comparisons suggest that bilateral cooperation may be *strengthened* if there exist outsiders that gain if their cooperation breaks down (see e.g. Turner, 1975; Brewer and Kramer, 1986; Tajfel, 1982; Putnam, 2007).

Financial support from the UiS Business School and the Norwegian Research Council is gratefully acknowledged. For constructive comments, thanks to Björn Bartling, Alexander Cappelen, Shachar Kariv, David Laibson, John List, Mari Rege, Anja Schöttner, Bertil Tungodden, Marie Claire Villeval, Roberto Weber, and participants at PhD courses/ workshops at the Choicelab, the Nordic Conference on Behavioral and Experimental Economics, the Zurich Workshop on Economics, and the Stavanger Workshop on Behavioral Economics"People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices."(Adam Smith, (1776), pp. 111)

* Corresponding author.
E-mail addresses: a.a.johnsen@frisch.uio.no (Å.A. Johnsen), ola.kvaloy@uis.no (O. Kvaløy).

https://doi.org/10.1016/j.socec.2021.101742
Received 8 July 2020; Received in revised form 12 March 2021; Accepted 28 June 2021
Available online 10 July 2021

Anecdotal evidence from court cases against illegal cartels supports this conjecture: the secret "conspiracy" against the consumers creates bonds between the colluding parties that makes the cartel agreement more robust.[1] A recent paper by Malmendier and Schmidt (2017) formalizes a similar idea: If parties cooperate, despite the fact that their cooperation yields negative externalities, then this may increase the weight that the cooperating parties attach to each other's welfare. They find support for their theory in a gift exchange experiment with negative externalities.

In this paper, we investigate – by use of a controlled lab experiment - the effect of negative externalities on cooperative behavior in a repeated prisoner's dilemma (PD). In the baseline treatment (taken from Dal Bó and Fréchette, 2011), pairs of subjects simultaneously choose between cooperation and defection in an infinitely repeated PD, i.e. a game that each period ends with a probability less than one. In the "public good treatment", subjects play the same repeated PD, with the difference that they contribute to a public good when they defect from the bilateral cooperation with their partner. That is, if they engage in bilateral cooperation they do not contribute to the public good, and hence the gains from cooperation are at the expense of the public. In the experiment, the public good is represented by a student organization that provides services with public good properties to the students at the university.

Our paper relates most closely to Malmendier and Schmidt (2017), who investigate the negative externalities of gift giving, e.g. business gifts, where the giver hopes to get favorable treatment from the receiver. They find that the gift triggers an obligation to repay the gift, even if the gift is given with the intention to affect the decision of the recipient at the expense of a third party. In fact, the gift has a stronger effect when it imposes a negative externality on a third party. They denote this obligation to repay a "bond" between gift giver and recipient.

In line with Malmendier and Schmidt (2017), our conjecture is that the negative externality imposed by cooperation, increases the level of cooperation. To some extent, this is what we find: Cooperation tends to increase more in the public good treatment. Initially, the level of cooperation is lower, but as subjects gain experience and observe that their partners choose to cooperate despite the negative externality, they cooperate as least as much as in the baseline treatment.

These results fit well with the mechanism proposed by Malmendier and Schmidt: If subjects in the public good treatment think that their partner might care about the public good, they will initially expect less cooperation compared to subjects in the baseline treatment. Hence, when a subject in the public good treatment experiences cooperation, given the lower expectations, cooperation is seen as a more favorable act than in the baseline, that triggers a positive response. This is modelled as attaching more weigh to the welfare of the colluding partner compared to subjects in the baseline treatment. In our setting, the "bond" is reflecting the pair's ability and willingness to form a collusive/ cooperative unit.

Our paper contributes to a small but growing literature investigating prosocial behavior in situations where the behavior comes at the expense of third parties. As noted by Malmendier and Schmidt, such negative externalities have largely been ignored in the theoretical and experimental literature. In a related experiment, Pan and Xiao (2016) study to what extent gift giving triggers the obligation of the receiver to favor the gift giver. They show that the pure gift effect is present even when it leads to a less efficient outcome, i.e. when it comes at the expense of a third party. Currie, Lin, and Meng (2013) study how gift-giving affects third parties in Chinese hospitals. A pair of trained actors visits physicians and plays the role of patients. If the first patient gives a small gift to the physician, s/he receives better service and is less likely to be prescribed unnecessary and costly medication. If the first "patient" introduces the second "patient" as a friend, this patient also receives better service.

An interesting recent paper by Bland and Nikiforakis (2015) studies how third-party externalities affect equilibrium selection in one-shot coordination games. They find that subjects are willing to incur a cost to try to avoid imposing large negative externalities on third parties. However, they ignore the negative externalities if the incentives are sufficiently at odds with third party interests. Like us, they demonstrate that third party externalities affect equilibrium selection, but in contrast to us, they do not find that third party externalities lead to stronger coordination on self-interested outcomes. The plausible explanation is that we study repeated interaction, in which cooperative bonds can develop over time. In fact, consistent with Bland and Nikiforakis, we also find that subjects put positive weight on the welfare of third parties in the first round of the repeated game.

Another strand of literature that (implicitly) studies the effect of negative externalities on prosocial behavior is the experimental papers on bribery games, which are similar to repeated gift exchange games. In Abbink, Irlenbusch, and Renner (2002) and Abbink (2004), the gains from corruption are at the expense of the public. The reciprocal relationship between the briber and the bribee is undesirable with regard to social welfare and is subject to punishment when discovered. They find that bribery relationships do develop over time despite negative externalities, but they do not find that the public aspect leads to *more* corruption. In a similar bribery game, Barr and Serra (2009) finds that bribers are *less likely* to offer bribes when the negative externalities imposed on the public are high, and main difference from Abbink et al (2002) and is that the experiment is more clearly framed as corruption, such that the negative externality becomes more apparent.

In contrast to these papers, our negative externality is an absence of a positive contribution to a public good, rather than a direct reduction in a third party's payoff. Moreover, in our paper, it affects an external organization, not an individual who is taking part in the experiment. Our negative externality is thus more abstract and subtle. Taken together, these papers show that the severity and framing of the negative externality matters. It also indicates that in standard collusion games, where the negative effect on social surplus may feel less severe and concrete than in bribery games, we are more likely to observe that the externality tightens the bond between the collusive parties.

By being able to cooperate to a larger extent, the subjects in the public good treatment earn more than the baseline subjects, which aligns well with the summary of what we have learnt about cooperation in repeated prisoner's dilemma games, namely that subjects are mainly tying to maximize their monetary payoffs (Dal Bó and Fréchette, 2018). Gneezy et al. (2019) show that the mechanism through which bribes "work" is mostly greed and not reciprocity. In our experiment we cannot disentangle greed from reciprocity, but our main contribution is to show that the presence of a third party can facilitate cooperation.

On a more general level, our paper is related to the experimental literature on collusion. For recent surveys, see Potters and Suetens (2013) and Armstrong and Huck (2014). The public aspect we introduce is usually present only implicitly in this literature. Subjects in the room participate in a market, and prices increase when they collude. Our paper also illuminates the experimental literature on cooperation in infinitely repeated games, see Dal Bó and Fréchette (2018) for a recent survey. A general finding is that higher defection payoff reduces the level of cooperative behavior. Our experimental results indicate that this does not necessarily hold in situations where the defection payoff is somewhat dubious. An uncertain change in defection pay-off – which is due to uncertainty about opponent's social preferences - might change the *expectation* for cooperative behavior, which in the end is not fulfilled by the cooperative parties. If it is interpreted as a potential increase in defection payoff, but turns out not reduce cooperation, it may in fact

[1] See e.g **Hammond (2005)** on the case against the Lysine cartel, where the world's five largest producers of lysine, an animal feed additive, succeeded in doubling the world **price** of lysine for several years. This cartel overcharged consumers and customers by an estimated US$ 140 million. The cartel was prosecuted and charged after a member was caught on tape saying that "Our competitors are our friends. Our customers are the enemy."

Table 1
The game matrix with players' pay-off in both treatments.

	Player 2		
		Collude	Defect
Player 1	Collude	40, 40	12, 50
	Defect	50, 12	25, 25

Table 2
Payoff to the public good (student organization StOr) and the players in both treatments.

	Collude-Collude	Collude-Defect	Defect-Collude	Defect-Defect
Player 1	40	12	50	25
Player 2	40	50	12	25
StOr, baseline treatment	0	0	0	0
StOr, PG treatment	0	25	25	50

lead to more robust cooperative relationships for those who choose to cooperate.

The rest of the paper is organized as follows: Section 2 presents the experimental design, while Section 3 presents the behavioral predictions. The results are presented in Section 4, while Section 5 concludes.

2. Experimental design and procedure

The experiment uses a between-subject design and consists of two treatments. The baseline is a standard infinitely repeated prisoner's dilemma game where the gains from cooperation are not at the expense of the public. The second treatment is based on the same infinitely repeated prisoner's dilemma as the baseline, but now the gains from cooperation are at the expense of the public, hence it can be referred to as collusion. The public is represented by a student organization, *StOr*.

The baseline is a replication of a treatment in Dal Bó and Fréchette (2011), who study cooperation in infinitely repeated games. The subjects are divided into pairs, denoted "matches". Each match consists of an ex-ante unknown number of rounds. Infinity is simulated using what is known as a random continuation rule. The random continuation rule assigns a fixed probability of continuation, and in this experiment it is equal to 3/4. In each match, all the subjects in a room are divided into random pairs. They play the first round and then a lottery decides if there will be another round. There is a 75% chance that this round will be followed by another round. This means that in expectation each match lasts for 4 rounds. When the match ends, the subjects are all randomly re-matched, and the same procedure is repeated until the experiment is over, which is after an hour. The shortest session consists of 31 matches, the longest session consists of 42 matches. Before the subjects leave the room, they also make one more decision (more below) and fill in a questionnaire.

The subjects are paid for the outcome from every decision made (see e.g. Sherstyuk, Tarui, and Saijo, 2013). The payoff when both subjects collude is equal to 40 experimental units (ECU) for both subjects. Temptation payoff is 50, sucker's payoff is 12, and if they both defect, they both get 25 each. The exchange rate is 10 ECU for 1 NOK (or ca. 75 ECU per 1 EUR).

The public good treatment has the exact same payoffs as the baseline Table 1. The only difference is that now the subject contributes 25, via the experimenter, to a public good if he or she chooses defect. That is, if a subject defect, the experimenter pays 25 to a public good (see below). It is not withdrawn from the subject's payoff. The treatments are identical in all other regards. When both subjects choose defect, they both contribute 25 to the public good, 50 in total. When one of the subjects chooses defect, and one subject chooses collusion, the defecting subject contributes 25 to the public good. Finally, when both subjects choose to collude, the contribution to the public good is zero. The payoffs are shown in Table 2:

The public good in this experiment is provided by the student organization StOr at the home university of the subjects, the University of Stavanger. StOr is a non-party affiliated interest organization, where all students at the UiS are members. The organization is responsible for life on campus, student welfare, student elections, student organizations, international students, exchange programs, legal issues regarding exams, syllabus and so on. In sum, it provides services that have public good properties. The services provided by this organization allow multiple agents to consume most of it at the same time (non-rival), and it is not possible to exclude subjects who did not contribute to the good from consuming it (non-excludable). The contribution to the public good is a fixed amount, and since the organization already exists, there is no provision point that needs to be reached. When the subjects contribute to the public good, it will translate into a very small increase in the

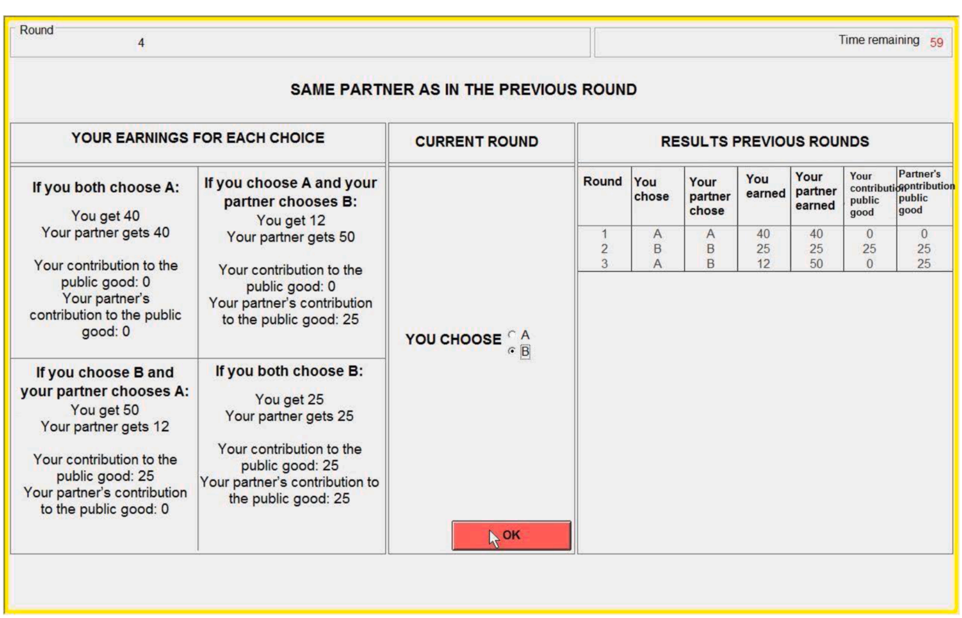

Figure 1. Screen-shot public good treatment, translated version.

Table 3
Payoff matrix.

	C	D
C	R	S
D	T	P

provision of the public good. This is meant to capture the fact that when firms refrain from colluding, this increases efficiency in the economy, from which both consumer and firms benefit (although to a smaller extent than the consumers). The instructions and the design were presented in a neutral language (*A* is collusion, *B* is defection/ not colluding), and the subjects were provided with an overview of all the information within each match, but not between matches. Figure 1 is a translated screen-shot from the public good treatment, see the Appendix for instructions for both treatments.

A total of 120 students at the University of Stavanger (Norway) participated in the experiment, with 20 students in each session. The subject sample in each treatment is similar to that of recent papers investigating cooperation using infinitely repeated prisoner's dilemma games (see e.g. Duffy and Ochs, 2009; Fréchette and Yuksel, 2013; Fudenberg, Rand, and Dreber, 2012; Kagel and Schley, 2013; Sherstyuk, Tarui, and Saijo, 2013). The subjects earned an average of EUR 55. The subjects were rematched on average 34 times. As the shortest session lasted 31 matches, the analysis is based on the first 31 matches for every session. The average number of rounds per match was 4.1, and the maximum number of rounds was 24. All instructions were given both written and verbally. The experiment was conducted and programmed with the software z-Tree (Fischbacher, 2007).

After the prisoner's dilemma game and before the questionnaire, the subjects were given a final task. Subjects in both treatments were asked if they preferred either *ECU 40 for themselves*, or *ECU 25 for themselves and ECU 25 to StOr* (incentivized). This decision involves the same payoffs as the prisoner's dilemma, but the difference is that it does not involve any interaction with other subjects.

3. Behavioral Predictions

In order to explain favoritism in gift exchange, Malmendier and Schmidt (2017) proposes a model of social preferences with endogenous references groups. We will here briefly present their novel behavioral assumption, leading to our behavioral prediction.

Consider a two player game of perfect information, where player i chooses strategy s_i, and s denote a pure strategy profile for both players. The utility of player i is given by:

$$U^i = m^i(s) + \alpha_i^j(s\sigma)\cdot m^j(s),$$

where $m^i(s)$ is player i's material payoff, and $\alpha_i^j(s\sigma)$ is the weight player i puts on the payoff of player j. This weight depends on the expected strategy profile, σ. The profile σ can consist of mixed strategies, and is expected to be played because of e.g. past experience, social norms, or simply because σ is an equilibrium that the players expect to be played. Formally, Malmendier and Schmidt makes the following assumption:

If player j chooses a strategy s_j that increases (decreases) player i's payoff that i would have received if j had chosen the expected strategy σ_i, then the weight of player j's payoff in player i's utility increases (decreases) compared to the weight if j had chosen σ_j. That is: $m^i(s_j, \sigma_i) \geq m^i(\sigma_j, \sigma_i) \rightarrow \alpha_i^j(s_j, \sigma_i\sigma) \geq \alpha_i^j(\sigma\sigma)$ and equivalently $m^i(s_j, \sigma_i) \leq m^i(\sigma_j, \sigma_i) \rightarrow \alpha_i^j(s_j, \sigma_i\sigma) \leq \alpha_i^j(\sigma\sigma)$.

Consider now the payoff matrix for player i (the row player) in the prisoner's dilemma game (where C denotes cooperation and D defection) in Table 3:

The PD payoff matrix consists of the reward from cooperation (R), the temptation payoff (T), which is the payoff to defecting when the other cooperates, the punishment from mutual defection (P), and the sucker's payoff from cooperation when the other defects (S). For these payoffs to define a PD game, it must be that T > R > P > S. Given the utility function above, the players do not actually play the same stage game every period. The reason is that if a player identifies a different strategy than expected, then the payoffs in the game change.

Consider now our experiment and assume that at least some players care about the public good. In the first period, T and P are then higher in the public good treatment than in the baseline. We will thus expect to

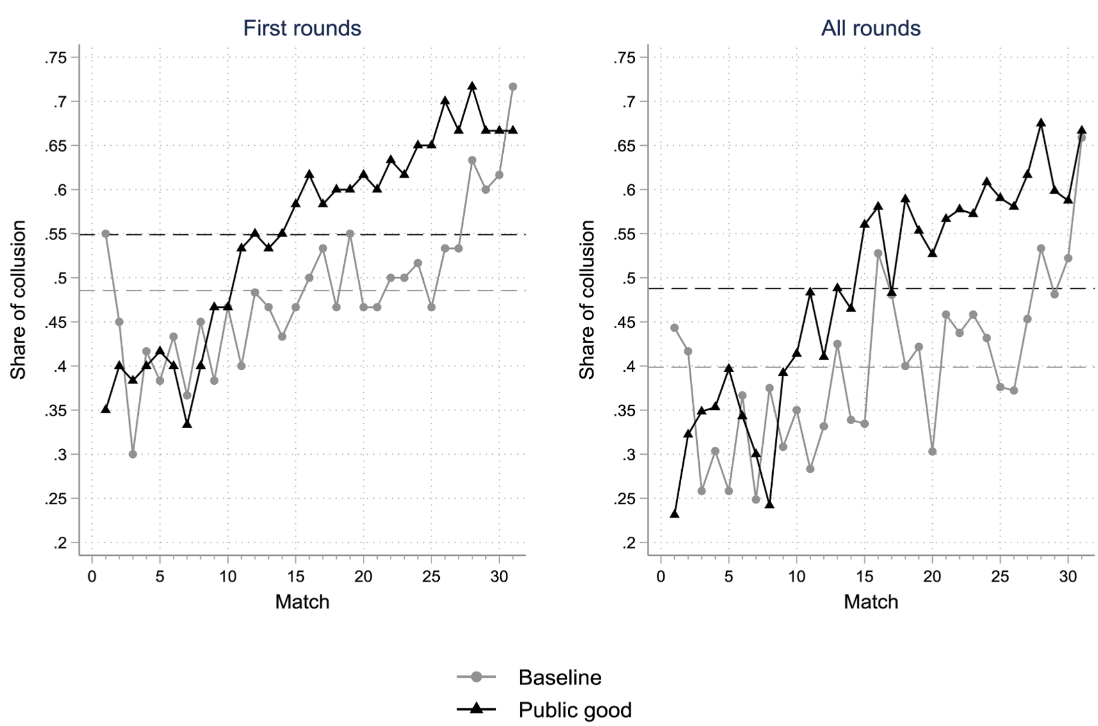

Figure 2. Evolution of collusion.

Table 4
Individual Collusion rates.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	R1_M1	R1_M1	M1	M1	R1	R1	All	All
PG treatment	-0.20**	-0.19**	-0.21***	-0.21***	0.06	0.08	0.09*	0.10*
	(0.09)	(0.09)	(0.07)	(0.06)	(0.07)	(0.07)	(0.05)	(0.05)
Age		0.10**		0.13***		0.04		0.04
		(0.04)		(0.03)		(0.04)		(0.03)
Female		-0.01		-0.02		-0.11		-0.07
		(0.10)		(0.07)		(0.07)		(0.06)
A&E		-0.20		-0.11		-0.04		-0.03
		(0.16)		(0.12)		(0.12)		(0.10)
SS		-0.17*		-0.09		-0.06		-0.05
		(0.10)		(0.07)		(0.08)		(0.06)
Member PG		0.21		0.11		0.01		0.01
		(0.19)		(0.09)		(0.11)		(0.08)
Observations	120	120	120	120	120	120	120	120
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Mean baseline	0.55	0.55	0.44	0.44	0.49	0.49	0.40	0.40
SD baseline	0.50	0.50	0.43	0.43	0.34	0.34	0.26	0.26
P-value WCB	0.19	0.15	0.04	0.04	0.57	0.45	0.50	0.44
P-value RI	0.20	0.20	0.10	0.10	0.60	0.50	0.50	0.50
P-value M-W	0.03		0.02		0.35		0.10	

Note: The independent variable in column (1) and (2) is equal to 1 if individuals colluded, 0 otherwise. The independent variable in column (3) and (4) is equal to the individual average collusion rate in the first match. The independent variable in column (5) and (6) is equal to the individual average collusion rate in the first rounds. Individual average collusion rates in all rounds of 31 matches in column (7) and (8). Controls: Age is standardized, with mean equal to zero, and standard deviation equal to one. A&E is equal to 1 for students at the faculty for arts and education and 0 otherwise, SS is equal to 1 for students at the faculty for social sciences, and students at the faculty of science and technology constitutes the reference category. Linear regression models, robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. P-values: *ri* denotes randomization inference, M-W denotes Mann-Whitney U test, and WCB denotes Wild Cluster Bootstrap procedure bootstrap clustering by session, Webb weights (MacKinnon et al. 2019).

Table A1
Balance.

VARIABLES	(1) Age	(2) Female	(3) A&E	(4) SS	(5) S&T	(6) Member PG	(7) IMP PG
PG	0.03	0.17*	0.02	0.00	-0.02	-0.03	0.32
	(0.18)	(0.09)	(0.05)	(0.08)	(0.09)	(0.05)	(0.38)
Observations	120	120	120	120	120	120	120
Mean baseline	-0.01	0.50	0.08	0.28	0.63	0.10	6.62
SD baseline	0.88	0.50	0.28	0.45	0.49	0.30	2.09
P-value	1.00	0.70	1.00	1.00	1.00	0.70	0.30

Note: Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0. Age is standardized, with mean equal to zero, and standard deviation equal to one. Field of study: arts and education (A&E), the faculty of social sciences (SS), the faculty of science and technology (S&T). Member PG: "Have you previously been/ are you currently a member of StOr?", 1 yes, 0 no. Imp PG: *How important is the work done by StOr, in your opinion?*", scale 1-10. P-value from randomization inference in final row, same procedure as the main analysis. F-test of joint orthogonality tests whether the observable characteristics are jointly unrelated to treatment status: Prob > F = 0.83.

see more cooperation in baseline. Assume now that a player *i* observes player *j* playing C in the public good treatment, despite expected play of D. She may then attach more weight α_i^j to player *j*'s welfare, and hence the R/T ratio increases, which can lead to more cooperative play in the public good treatment than in the baseline.

There are of course numerous equilibria in super games like this, and it is straightforward to construct equilibrium strategies where the players after *n* periods cooperate more in the public good game. This is also our behavioral prediction: We will initially see less, but then more cooperation (denoted collusion) when cooperation is at the expense of the public good.

4. Experimental results

Table A1 in the appendix provides a summary of the subjects randomly placed into each treatment and test for balance in predetermined variables across treatments for age, field of study, gender, previous membership in the organization, as well as attitudes towards the organization representing the public good.[2] Considering an F-test for the joint significance, we see that the data are balanced across these characteristics.[3] On a scale from 1 to 10, the subjects on average rate the importance of the student organization's work to 6.8/10, which supports that they value the services of the organization. For more details on each session, see Table A2 in the appendix.

Figure 2 below illustrates how collusion evolves in the baseline and the public good treatment. The left panel illustrates the collusion rate of the first round of each match for both treatments. The right panel illustrates the collusion rate of each match – averaged across all rounds

[2] The survey data is collected after the subjects have finished the experiment, hence, we only include previous membership and not attitudes towards the public good organization as control variables in the main regressions, as the latter could be influenced by the subjects' experiences in the experiment.

[3] Results from robust linear regression indicates imbalance with respect to gender (significant at the ten percent level), but randomization inference does not. Including gender as a control does not affect the results.

Table A2
Session characteristics.

	Session	Subjects	Games	Average no. of rounds	Average Payoff
DBF	1	12	34	3.9	31.4
	2	14	47	3.2	29.2
	3	12	23	5.4	27.6
Baseline	2	20	32	4.6	82.8
	3	20	32	3.6	56.3
	4	20	31	4.1	64.2
Public good	1	20	43	4.3	101.6
	5	20	33	3.5	57.7
	6	20	32	4.6	86.0

Note: The payoffs for Dal Bó and Fréchette (2011) are in 2011 US dollars, while the payoffs in this experiment are in 2013 US dollars, exchange rate May 24th: 5.92 NOK per US dollar. Subjects are on average 24 years old, 58.3 percent of the sample are female, 63 percent are from the faculty of science and technology, 28 percent are from the faculty of social sciences, and 8 percent are from the faculty of arts and education.

dashed line marks the average collusion rate in the baseline (0.48 considering only first rounds, 0.40 for all rounds), while the black dashed line marks average collusion rates for the public good treatment (0.55 considering only first rounds, 0.49 all rounds).

Are these overall differences statistically significant? In the following section, we investigate these results using both regression analysis and randomization inference. Table 4 presents the results from robust linear regression models with collusion as the dependent variable, and the analysis is at the individual level. We show the results from randomization inference in the final row of Table 4. Using Fisherian randomization inference, we consider what would have happened under all possible assignments of the six sessions to treatment.[5] Consider the null hypothesis that the public good treatment has no effect. If the null hypothesis is true, we know exactly what the outcome would be in alternative allocations of treatment. This allows for the calculation of exact p-values. We re-randomize sessions to all 20 possible combinations of treatment and baseline and investigate whether our actual experimental results are unusual or not. *P*-values are given by the share of times we observe a difference in means between the treatment and baseline that is

Table A3
Generalized linear models

VARIABLES	(1) R1_M1	(2) R1_M1	(3) M1	(4) M1	(5) R1	(6) R1	(7) All	(8) All
PG treatment	-0.82**	-0.83**	-0.65***	-0.66***	0.12	0.16	0.20*	0.23*
	(0.38)	(0.39)	(0.20)	(0.19)	(0.13)	(0.13)	(0.12)	(0.12)
Age		0.43**		0.32***		0.07		0.08
		(0.20)		(0.06)		(0.06)		(0.06)
Female		-0.05		-0.07		-0.20		-0.17
		(0.41)		(0.19)		(0.13)		(0.12)
A&E		-0.87		-0.23		-0.07		-0.06
		(0.73)		(0.34)		(0.22)		(0.22)
SS		-0.79*		-0.28		-0.11		-0.11
		(0.46)		(0.21)		(0.15)		(0.14)
Member PG		0.93		0.33		0.02		0.03
		(0.85)		(0.24)		(0.22)		(0.18)
Observations	120	120	120	120	120	120	120	120
Controls	No	Yes	No	Yes	No	No	No	No
Mean baseline	0.55	0.55	0.44	0.44	0.49	0.49	0.40	0.40
SD baseline	0.50	0.50	0.43	0.43	0.34	0.34	0.26	0.26
P-value WCB	0.08	0.06	0.00	0.00	0.52	0.39	0.44	0.37

Note: The independent variable in column (1) and (2) is equal to 1 if individuals colluded, 0 otherwise. The independent variable in column (3) and (4) is equal to the individual average collusion rate in the first match. The independent variable in column (5) and (6) is equal to the individual average collusion rate in the first rounds. Individual average collusion rates in all rounds of 31 matches in column (7) and (8). Controls: Age is standardized, with mean equal to zero, and standard deviation equal to one. A&E is equal to 1 for students at the faculty for arts and education and 0 otherwise, SS is equal to 1 for students at the faculty for social sciences, and students at the faculty of science and technology constitutes the reference category. Column (1) and (2): Logistic regression, reporting estimated coefficients. Column (3)-(8): Poisson models. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Poisson regression with the Huber/White/Sandwich linearized estimator of variance is a permissible alternative to log linear regression (Wooldridge, chapter 18; Cameron and Trivedi, chapter 17.3.2). P-values W-B denotes Wild Cluster Bootstrap clustering by session (Rao score/Lagrange multipler test (scoretest)).

within each match - for both treatments.[4]

In the first round of the first match in the left panel, we see that 55 percent of the individuals in the baseline choose to collude, compared to 35 percent in the public good treatment. Looking at the collusion rate in the first match overall in the right panel, we see that 44 percent in the baseline choose to collude, compared to 23 percent in the public good treatment. Hence, collusion rates seem to be lower in the first match in the public good treatment compared to the baseline. However, we can see that this changes as the subjects gain experience. We see in both panels that collusion rates increase for both the baseline and the public good as the subjects gain experience. Furthermore, the collusion rates seem to be increasing more in the public good treatment. The gray

equal to or bigger than (in absolute value) the realized outcome in a two-sided test.[6]

First, we investigate whether the first round of the first match differs between treatments in (1). The results from the robust linear regression model indicates that it does – the share cooperating in the public good treatment is 20 percentage points lower than in the baseline, and it is significant at the five percent level. In the lower panel of Table 4 we show the results from randomization inference by including the p-values from a two-sided test under the sharp null hypothesis of no difference

[4] The collusion rate for each match is equal to the average collusion rate over all rounds within each match. We only use data where we have observations for all individuals/sessions (31 matches).

[5] Randomization inference provides an excellent test for analysing data from randomized experiments, especially in samples with a small number of observations and with clustered randomization, as in this experiment. We perform randomization inference using the *ritest* developed by Simon Heß (2017).

[6] For more on randomization inference, see e.g. Young (2019), Imbens and Rubin (2015), and Athey and Imbens (2016, 2017).

Instructions Public Good Treatment

Welcome!

- In this experiment you will be asked to make some decisions.
- You will get the opportunity to earn money which will be paid to you in cash and anonymously when the experiment is over.
- In this experiment we use what we denote experimental currency units, ECU. By the end of the experiment your total earnings will be converted into Norwegian Kroners according to the following rate: 10 ECU = 1 NOK.
- Your earnings depend partly on your decisions, partly on others' decisions, and partly on chance.
- We will now go through the instructions in detail. You will be given sufficient time to read the instructions. The experiment will last for about one hour.
- If you have any questions regarding the instructions, please raise your hand and we will come to you.
- The experiment will be conducted using computers, and talking or communicating with others during the experiment is not allowed.
- Please turn your cellular phone off and put it away.

Instructions

1. All participants will be randomly matched into pairs several times during this experiment, and each time you and your partner will be asked to make some decisions. Each matched pair plays a sequence of rounds.
2. The number of rounds in each match will vary from match to match. The number of rounds in each match is determined by a lottery. After one round has been played, there is a 75 percent chance that there will be another round. Another way of saying this is that there will be another round in 3 out of 4 times. This means that when you have finished playing the first round, there is a 75 percent chance that there will be a second round. In other words - when round 2 is finished, the probability of a third round is still equal to 75 percent.
3. When the outcome of the lottery is that there will not be another round, all participants are randomly re-matched again. Your earnings from the previous match will be set aside on your personal account. The number of rounds which you and your new partner will meet will be decided by the same lottery as described in 2.
4. In the table below, you can see what you earn, what your partner earns, and what is contributed to the public good for all the four possible choice sets:

If you both choose A: You get 40 Your partner gets 40 Your contribution to the public good: 0 Your partner's contribution to the public good: 0	**If you choose A and your partner chooses B:** You get 12 Your partner gets 50 Your contribution to the public good: 0 Your partner's contribution to the public good: 25
If you choose B and your partner chooses A: You get 50 Your partner gets 12 Your contribution to the public good: 25 Your partner's contribution to the public good: 0	**If you both choose B:** You get 25 Your partner gets 25 Your contribution to the public good: 25 Your partner's contribution to the public good: 25

In each round you and your partner choose between choice **A** and choice **B**. You and your partner choose simultaneously, and you will get to know your partner's decision after you have made a decision. Your earnings depend on your decision, but they also depend on what your partner's decision is.

As you saw in the table above, you have the opportunity to contribute to a public good in this experiment. The public good is the student organization ***StOr*** here at the University of Stavanger, which works to promote the students' interests (see introduction

Figure A1. Instructions public good treatment

between the public good treatment and the control. A coefficient equal to or larger than (absolute value) -0.20 occurs in 4 out of the 20 re-randomizations, producing a p-value of 0.20. Hence, randomization inference does not provide support for the results from regression (1), and we conclude that we do not have strong evidence in favor of for any treatment effect for the first round of the first match. Adding individual control variables in (2) do not change this result.

Column (3) presents the results when the dependent variable is equal to the individual's overall collusion rate in the first match.[7] The coefficient is similar to that in (1), and significant at the one percent level. The difference (treatment effect) between treatments is equal to about 0.6 standard deviations (0.21/0.38). This result is supported by

[7] The individual's collusion rate is calculated in two steps: First, we calculate each individual's rate of collusion within each match, and second, we average over the average collusion rate for each individual.

below). When you choose **B**, you contribute to the public good, and by "public good" we always mean ***StOr***. When you choose **A**, you do not contribute to the public good. The same goes for your partner - when your partner chooses **B**, he/ she contributes to the public good. When your partner chooses **A**, he/ she does not contribute to the public good.

Summary

The number of rounds in each match is decided by a lottery. After each round there is a 75 percent chance that there will be another round. When there is no new round, all participants are re-matched. Below you can see what the screen will look like when you are making a decision. The left part of the screen shows what you earn for each choice set. In the middle you see where you make your decision. In the right part of the screen you will see the results from the previous rounds with the current partner. Take your time - feel free to take 30 seconds to make your decision and remember to press the OK-button when you have made your decision.

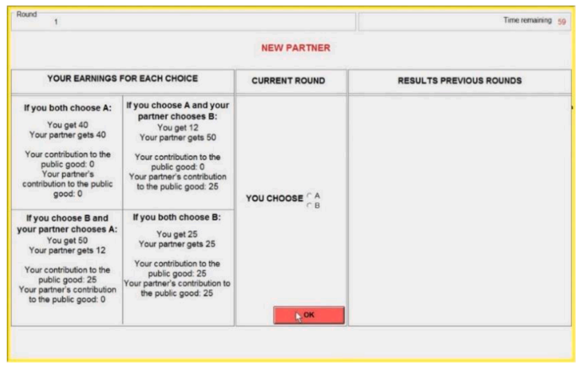

The lotteries are all drawn by the computer, and they are always randomly drawn.

Please follow the messages which appear on the screen. In the end you will be asked to fill out a short questionnaire, and you will be informed about your total earnings converted into NOK. On the pc cabinet you can see a white sticker with the logo of the university, and a number, for instance D10136. Please write down this number and your total income on the receipt when the experiment is over. When we tell you that the experiment is over, you can leave the room with the receipt in hand. Please bring this to the EAL building, office H-161, to collect your earnings.

About StOr

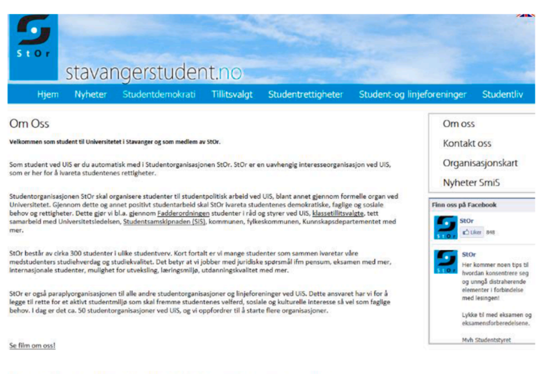

Figure A1. (*continued*).

randomization inference: An equally large or larger coefficient occurs in 2 out of the 20 re-randomizations – when the allocation of sessions is equal to the actual outcome, and when the baseline sessions are rerandomized as public good sessions (and baseline as public good sessions) - making it marginally significant at the ten percent level. Adding individual control variables in (4) do not change this result. The Mann-Whitney U test, logistic regression models for the binary outcome in (1) and (2), and Wild Cluster Bootstrap procedure (Roodman et al., 2019) clustering by session for both the linear and logistic regression models all produce consistent results with *finding 1*, see the bottom rows of Table 4 and Table A3 in the appendix.

Finding 1: Collusion is lower in the first match in the public good treatment.

In column (5), the dependent variable is equal to the average individual's collusion rate over all first rounds. The public good treatment coefficient is positive, but not significant. The effect size is small and equal to 0.17 standard deviations. Randomization inference also shows that the observed outcome occurs frequently; hence, our results do not imply that there is any difference between treatments when we only look at first rounds, and these results do not change when we add individual controls in (6).

In (7), we average over all matches and all rounds. The results from the robust linear model indicate that collusion is higher in the public good treatment, but it is only marginally significant at the ten percent level (supported by Mann-Whitney U-test, see Table 4, and Poisson model with robust SEs, see Table A3). The group means differ by about 0.30 standard deviations. However, the results from randomization inference show that this result is not very unusual (p-value 0.50). These results do not imply that there is any difference between treatments. This result also holds when controlling for individual characteristics in (8), and Wild Cluster Bootstrap procedure clustering by session for both the Logit and Poisson regression models all produce consistent results with finding 2, see Table 4 and Table A3 in the appendix.

Finding 2: As subjects gain experience, they cooperate as least as much as in the baseline treatment.

Comparing payoffs between treatments, we find that the subjects in

Instructions Baseline

Welcome!

- In this experiment you will be asked to make some decisions.
- You will get the opportunity to earn money which will be paid to you in cash and anonymously when the experiment is over.
- In this experiment we use what we denote experimental currency units, ECU. By the end of the experiment your total earnings will be converted into Norwegian Kroners according to the following rate: 10 ECU = 1 NOK.
- Your earnings depend partly on your decisions, partly on others' decisions, and partly on chance.
- We will now go through the instructions in detail. You will be given sufficient time to read the instructions. The experiment will last for about one hour.
- If you have any questions regarding the instructions, please raise your hand and we will come to you.
- The experiment will be conducted using computers, and talking or communicating with others during the experiment is not allowed.
- Please turn your cellular phone off and put it away.

Instructions

1. All participants will be randomly matched into pairs several times during this experiment, and each time you and your partner will be asked to make some decisions. Each matched pair plays a sequence of rounds.
2. The number of rounds in each match will vary from match to match. The number of rounds in each match is determined by a lottery. After one round has been played, there is a 75 percent chance that there will be another round. Another way of saying this is that there will be another round in 3 out of 4 times. This means that when you have finished playing the first round, there is a 75 percent chance that there will be a second round. In other words, when round 2 is finished, the probability of a third round is still equal to 75 percent.
3. When the outcome of the lottery is that there will not be another round, all participants are randomly re-matched again. Your earnings from the previous match will be set aside on your personal account. The number of rounds that you and your new partner will meet will be decided by the same lottery as described in 2.
4. In the table below, you can see what you earn and what your partner earns for all four possible choice sets:

If you both choose A: You get 40 Your partner gets 40	**If you choose A and your partner chooses B:** You get 12 Your partner gets 50
If you choose B and your partner chooses A: You get 50 Your partner gets 12	**If you both choose B:** You get 25 Your partner gets 25

In each round you and your partner choose between choice **A** and choice **B**. You and your partner choose simultaneously, and you will get to know your partner's decision after you have made a decision. Your earnings depend on your decision, but they also depend on what your partner's decision is.

Summary

Figure A2. Instructions baseline

the public good treatment receive significantly higher payoffs compared to the baseline. On average, individuals in the baseline treatment earn USD 63.1 based on the 31 first matches, compared to USD 71.2 in the public good treatment.[8]

[8] Two-sample Wilcoxon rank-sum (Mann-Whitney) test, Prob > |z| = 0.0003. Average payoffs based on all matches: USD 67.7 in baseline and 81.8 in the public good treatment. (Prob > |z| = 0.0002). Average payoff based on the first 31 matches: USD 63.1 versus 71.2, (Prob > |z| = 0.0003).

The number of rounds in each match is decided by a lottery. After each round there is a 75 percent chance that there will be another round. When there is no new round, all participants are re-matched. Below you can see what the screen will look like when you are making a decision. The left part of the screen shows what you earn for each choice set. In the middle you see where you make your decision. In the right part of the screen you will see the results from the previous rounds with the current partner. Take your time - feel free to take 30 seconds to make your decision and remember to press the OK-button when you have made your decision.

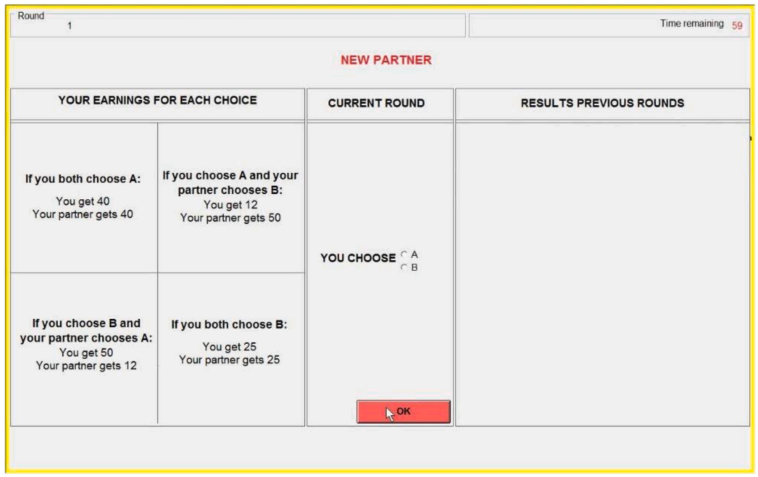

The lotteries are all drawn by the computer, and they are always randomly drawn.

Please follow the messages which appear on the screen. In the end you will be asked to fill out a short questionnaire, and you will be informed about your total earnings converted into NOK. On the pc cabinet you can see a white sticker with the logo of the university, and a number, for instance D10136. Please write down this number and your total income on the receipt when the experiment is over. When we tell you that the experiment is over, you can leave the room with the receipt in hand. Please bring this to the EAL building, office H-161, to collect your earnings.

Figure A2. (*continued*).

5. Conclusion

This paper studies to what extent collusive behavior is affected by the awareness of negative externalities. Theories of outcome-based social preferences suggest that negative externalities make collusion harder to sustain than predicted by standard economic theory, since collusion reduces social surplus. On the other hand, sociological theories of social ties and intergroup comparison suggest that bilateral cooperation can be strengthened if there exist outsiders that gain from cooperative break down.

Our experimental results give support for the latter: Cooperation tends to be strengthened when it comes at a third party's expense. Initially, the level of cooperation is lower, but as subjects gain experience, they cooperate as least as much when cooperation is at the expense of the public. Malmendier and Schmidt's (2017) model of social preferences with endogenous reference groups provides a potential explanation for our experimental results: Negative externalities may lower expectations for collusion. If they still collude, despite negative externalities, then this may increase the weight that the colluding parties attach to each other's welfare and create stronger ties. Our finding is consistent with the experimental findings of Malmendier and Schmidt, and a few other gift exchange papers with third party externalities. In contrast to this literature, we study a *repeated simulations move game* – i.e. the prisoner's dilemma - with third party negative externalities. This resembles well the strategic situation facing two colluding firms that tacitly decide to keep prices above marginal costs, and thereby capture surplus from consumers.

We believe we will see a growing literature examining the effect of negative externalities on prosocial behavior and equilibrium selection, both in social dilemmas and coordination games with strategic uncertainty. The importance of third-party externalities and the inconclusive predictions from the prevailing theories of social preferences suggest that investigating these questions is important. As Bland and Nikiforakis (2015) notes, the fact that third-party externalities were found to have an effect in their coordination games "suggests that it is interesting to explore in future research how they affect tacit coordination in different classes of coordination games, e.g., when decision-makers' incentives are not aligned". Our paper contributes to this research agenda.

Appendix

Questionnaire

In order to learn more about the individuals' social preferences in absence of coordination issues, we asked the subjects whether they preferred either ECU 40 for themselves, or ECU 25 for themselves and ECU 25 to the public good (incentivized) after the experiment had ended. If subjects felt that they had contributed enough to the public good in the public good treatment already,[9] contributions to StOr should be lower in the public good treatment compared to the baseline. The contributions were high and did not differ across treatments; 68% in the baseline chose to contribute to the public good, compared to 65% in the public good treatment (Two-sided Mann-Whitney test: p-value 0.70).

In the questionnaire following the experiment the subjects were asked to rate the importance of the student organization's work. On a scale from 1 to 10 (highest), the subjects on average report 6.8 (std.err. 0.19), which supports that they value the services of the organization, and there is no difference across the treatments (Two-sided Mann-Whitney test: p-value 0.43). We also asked whether they had been

[9] Conditional on that subjects care equally much about the public good in both treatments, which we show in the next paragraph that they seem to do.

active members of the organization and there was no treatment difference.[10]

Ideally, we should have elicited beliefs about cooperation at the start of each match. However, we only asked the subjects about their expectations after the experiment. Differences in beliefs at the start of the experiment and potential changes throughout the experiment are not captured by these questions, and furthermore, the answers are likely to be influenced by experiences throughout the experiment. Due to these issues we have not made use of these questions in the paper.

First, the subjects were asked what option best described how they reasoned when they were matched with a new partner the first round. In both treatments around 60 percent of the subjects reported that they expected that their partner would cooperate (58 percent in baseline versus 62 percent public good, p-value 0.71). Around 75 percent in both treatments expected their partner to cooperate in the second round, if partner also chose to cooperate in round 1 (73 percent in baseline versus 77 percent in public good, p-value 0.67). If their partner did not cooperate in the first round, 3 percent of subjects in baseline expected cooperation, while 12 percent of the subjects in the public good treatment expected cooperation in round 2 (p-value 0.08).

Second, the subjects in the public good treatment were asked whether their choices and their partners choices were affected by the fact that not colluding also involved contributing to a public good. Their beliefs about their partners (in parenthesis) were almost identical to their own responses. 7 (10) percent reported increased motivation to collude, 57 (52) percent reported that it did not influence them, and 37 (38) percent reported increased motivation to not collude (Wilcoxon signed-rank test: p-value 0.97).

Randomization

Session statistics

The average number of rounds per match is 4.1, and the maximum number of rounds is 24 (Dal Bo´ and Fre´chette (2011): average 4.4 and maximum 24):

Robustness checks

Instructions

Figure A1: Instructions baseline

Figure A2: Instructions public good treatment

References

Andreoni, James, & Miller, John H. (1993). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *Economic Journal, 103*, 570–585.

Abbink, Klaus (2004). "Staff rotation as an anti-corruption policy: an experimental study", *European Journal of Political Economy*20 4, pp. 887–906.

Abbink, Klaus, Irlenbusch, Bernd, & Renner, Elke (2002). "An experimental bribery game". *Journal of Law, Economics, and Organization, 18*(2), 428–454.

Armstrong, Mark, & Huck, Steffen (2014). Behavioral Economics and Antitrust. *The Oxford Handbook of International Antitrust Economics, 1*, 205–228.

Athey, Susan and Guido Imbens (2016). "The econometrics of randomized experiments", *arXiv preprint arXiv:1607.00698*. In preparation for *Handbook of Development Economics*.

Athey, Susan, & Imbens, Guido. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives, 31*(2), 3–32.

Barr, Abigail, & Serra, Danila (2009). The effects of externalities and framing on bribery in a petty corruption experiment. *Experimental Economics, 12*(4), 488–503.

Bland, James, & Nikiforakis, Nikos (2015). Coordination with third-party externalities. *European Economic Review, 80*, 1–15.

Brewer, Marilynn B., & Kramer, Roderick M. (1986). Choice behavior in social dilemmas: Effects of social identity, group size, and decision framing. *Journal of Personality and Social Psychology, 50*(3), 543.

Bolton, Gary E., & Ockenfels, Axel (2000). ERC: a theory of equity, reciprocity and competition. *American Economic Review, 90*, 166–193.

Currie, Janet, Lin, Wanchuan, & Meng, Juanjuan (2013). Social networks and externalities from gift exchange: Evidence from a field experiment. *Journal of public economics, 107*, 19–30.

Dal Bo´, Pedro, & Fre´chette, Guillaume R (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review, 101* (1), 411–429.

Dal Bó, Pedro, & Fréchette, Guillaume R. (2018). On the Determinants of Cooperation in Infinitely Repeated Games: A Survey. *Journal of Economic Literature, 56*(1), 60–114.

Duffy, John, & Ochs, Jack (2009). Cooperative behavior and the frequency of social interaction. In: *Games and Economic Behavior, 66*(2), 785–812.

Dufwenberg, Martin, & Kirchsteiger, Georg (2004). A theory of sequential reciprocity. *Games and Economic Behavior, 47*(2), 268–298.

Fehr, Ernst, & Schmidt, Klaus M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*, 817–868.

Fischbacher, Urs (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics, 10*(2), 171–178.

Fre´chette, Guillaume R, & Yuksel, Sevgi (2013). Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics, 20*, 279–308.

Fudenberg, Drew, Rand, David G, & Dreber, Anna (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review, 102*(2), 720–749.

Gneezy, Uri, Saccardo, Silvia, & Van Veldhuizen, Roel (2019). Bribery: Behavioral drivers of distorted decisions. *Journal of the European Economic Association, 17*(3), 917–946.

Hammond, Scott D (2005). *Caught in the act: Inside an international cartel*. US Department of Justice.

Heß, Simon (2017). Randomization inference with Stata: A guide and software. *Stata Journal, 17*(3), 630–651.

Imbens, Guido W., & Rubin, Donald B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Kagel, John H, & Schley, Dan R (2013). How economic rewards affect cooperation reconsidered. *Economics Letters, 121*(1), 124–127.

Malmendier, Ulrike, & Schmidt, Klaus (2017). You owe me. *American Economic Review, 107*, 493–526.

Pan, Xiaofei, & Xiao, Erte (2016). It's not just the thought that counts: An experimental study on the hidden cost of giving. *Journal of Public Economics, 138*, 22–31.

Potters, Jan, & Suetens, Sigrid (2013). Oligopoly experiments in the current millennium. In: *Journal of Economic Surveys, 27*(3), 439–460.

Putnam, Robert D. (2007). "E pluribus unum: Diversity and community in the twenty-first century". *Scandinavian Political Studies, 30*, 137–174.

Rabin, Matthew (1999). Incorporating fairness into game theory and economics. *American Economic Review, 83*(5), 1281–1302.

Roodman, D., Nielsen, M.Ø., MacKinnon, J. G., & Webb, M. D. (2019). Fast and wild: Bootstrap inference in Stata using boottest. *The Stata Journal, 19*(1), 4–60.

Sherstyuk, Katerina, Tarui, Nori, & Saijo, Tatsuyoshi (2013). Payment schemes in infinite-horizon experimental games. *Experimental Economics, 16*(1), 125–153.

Smith, Adam (1776). An Inquiry into the Nature and Causes of the Wealth of Nations. Ed. by Jim Manis. PSU-Hazleton, Hazleton, PA 18202: The Electronic Classics Series.

Tajfel, Henri (1982). "Social psychology of intergroup relations". *Annual review of psychology, 33*(1), 1–39.

Turner, John C. (1975). Social comparison and social identity: Some prospects for intergroup Behavior. *European journal of social psychology, 5*(1), 1–34.

Young, Alwyn. (2019). Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results. *The Quarterly Journal of Economics, 134*(2), 557–598.

[10] 10 percent of the control group and 7 percent of the public good treatment reported yes, p-value 0.51.