

Public Goods Provision and Sanctioning in Privileged Groups*

Ernesto Reuben[†] and Arno Riedl[‡]

Abstract

In public-good provision, privileged groups enjoy the advantage that some of their members find it optimal to supply a positive amount of the public good. However, the inherent asymmetric nature of these groups may make the enforcement of cooperative behavior through informal sanctioning harder to accomplish. In this paper, the authors experimentally investigate public-good provision in normal and privileged groups with and without decentralized punishment. The authors find that compared to normal groups, privileged groups are relatively ineffective in using costly sanctions to increase everyone's contributions. Punishment is less targeted toward strong free riders, and they exhibit a weaker increase in contributions after being punished. Thus, the authors show that privileged groups are not as privileged as they initially seem.

First version: July, 2007 This version: March, 2008

JEL Codes: H41, D01, C92, Z13

Keywords: privileged groups, public goods, punishment, cooperation, collective action

Note: This is the author's version of a work that was accepted for publication in the Journal of Conflict Resolution. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication.

A final version is published in <http://jcr.sagepub.com/cgi/rapidpdf/0022002708322361v1/>

*We thank Enrique Fatas, Rachel Croson, Dan Levin, Jan Potters, Louis Putterman, Frans van Winden, and particularly, Simon Gächter for useful comments and suggestions. Financial support from the Dutch Science Foundation (NWO) through the Evolution & Behavior grant 05112012 and from the EU-Marie Curie RTN ENABLE (MRTM-CT-2003505223) is gratefully acknowledged. All materials for replication are available at <http://jcr.sagepub.com/> and at the authors' personal Web pages.

[†]Northwestern University, e-mail: ereuben@northwestern.edu

[‡]CESifo, IZA, and Maastricht University, e-mail: a.riedl@algec.unimaas.nl

1 Introduction

In public-good provision, the incentives to free ride on the contributions of others generally lead to underprovision of the public good. In his seminal work on collective action, Olson (1965) provides a typology of groups that helps to explain why some groups nevertheless perform better than others in collective-good provision. In particular, he identifies type of group in which at least one member “has an incentive to see that the collective good is provided, even if he has to bear the full burden of providing it himself” (Olson, 1965, p. 50). In such groups, one or more individuals receive a marginal benefit from the public good that exceeds their marginal cost of providing it. Hence, in such groups, at least some level of the public good is voluntarily produced. Correspondingly, Olson also calls them privileged groups.

The theoretical literature on collective action generated important insights about when groups can be regarded as privileged. It has identified important determinants that turn nonprivileged groups into (fully) privileged ones and, hence, circumstances that mitigate underprovision of public goods (see e.g., Sandler, 1992). Considerably less attention has been devoted to the issue of underprovision of public goods in privileged groups, perhaps because it seems less problematic in such groups. A main purpose of this paper is to study precisely this problem and show that privileged groups may not always be as privileged as they seem.

In our view, there are at least three reasons that call for further exploration of voluntary public-goods provision in privileged groups. First, it is clear that many groups confronted with the problem of providing public goods are privileged in the sense of Olson (1965). Contributing to the fight against international terrorism is an obvious example. The United States and the United Kingdom, for instance, can be assumed to contribute because they are likely targets, and hence, their governments perceive the benefits of contributing as being larger than the costs. In contrast, some continental European countries may see themselves as unlikely victims, and hence, perceive the costs as being larger than the benefits. The same could be said of most cases of international security such as containment of “rouge” regimes. Another example is the so-called commons-based peer-production model, in which a large number of anonymous people with varying degrees of interest voluntarily contribute to a common project.¹

Second, even though there are some voluntary contributions in privileged groups, under-provision will occur as long as not all individual incentives are aligned with maximization of social welfare.² That such underprovision is not only a theoretical issue is vividly revealed in national and international political discussions regarding contributions to fighting terrorism and global warming, for instance. As for ordinary, nonprivileged, groups, this asks for studying appropriate mechanisms overcoming the problem of suboptimal collective-good provision.

Third, experimental research has recently identified important informal determinants of voluntary contributions, such as fairness-oriented and reciprocal inclinations. These determinants, however, may work out differently in privileged groups. The asymmetric nature of such groups may cause conflicting perceptions of fairness and reciprocity norms, which in turn may undermine the otherwise very successful enforcement mechanism of decentralized individual punishment (Ostrom et al., 1992; Fehr and Gächter, 2000).

¹One of the best known examples is the Linux operating system. Benkler (2002) discusses additional examples.

²Underprovision is a particular problem for the mostly assumed summation and weakest-links technologies of publicness. For public goods provided with a best-shot technology, privileged groups would, in theory, overcome to a large extent the underprovision problem (see Sandler, 1992).

This paper is the first to experimentally investigate voluntary contributions to a public good in privileged groups vis-à-vis normal, nonprivileged, groups. In particular, we compare behavior of group members with different incentives to contribute in situations with and without the option to costly punish free riders. Our results indicate that without punishment, there is substantial underprovision in privileged groups, albeit contributions are higher than in normal groups. With punishment, privileged groups lose their privileged status completely. Contributions to the public good are not higher than in nonprivileged groups and there is still substantial underprovision. In addition, inequality in privileged groups increases significantly when punishment is possible. Hence, the otherwise powerful institution of decentralized punishment loses some of its bite in privileged groups. The identification of the underlying reasons for this observation is another aim of our paper.

We conducted a linear public-goods experiment (Isaac et al., 1984) in which, in some treatments, subjects are allowed to punish each other after having observed individual contributions to the public good (similar to Fehr and Gächter, 2000). In normal groups, material incentives are such that all subjects have a dominant strategy to free-ride. In privileged groups, there is one subject whose material incentives are such that full contribution is a dominant strategy, whereas other group members' dominant strategy is to free ride completely. For convenience, we refer to subjects with the different material incentives to contribute as high- and low-benefit subjects.³

Note that under the standard assumption of (common knowledge of) narrow selfish preferences, only high-benefit subjects are predicted to contribute (fully) to the public good. In addition, neither high- nor low-benefit subjects are predicted to punish. However, there is plenty of evidence that in normal groups, first, individuals have some (initial) inclination to contribute positive amounts, and second, individual costly punishment is undertaken and effective in decreasing free riding.⁴ These behavioral regularities are often attributed to individuals having a preference for fair outcomes or reciprocal actions. From the outset, however, it is unclear if these type of social preferences can also increase contributions in privileged groups. First, since high-benefit subjects benefit disproportionately from contributing to the public good it is unclear if such behavior should be interpreted as kind towards others or as merely selfish. Since intentions are an important part of reciprocity (Falk et al., 2008) such ambiguity may make low-benefit subjects unwilling to reciprocate high contributions of high-benefit subjects, and high-benefit subjects unwilling to punish low contributions of low-benefit subjects. Second, since contributions of low-benefit subjects can increase income differences with high-benefit subjects, preferences for equality support outcomes in which low-benefit subjects do not contribute to the public good.⁵

³The reasons why individuals may value consumption of the public good differently are numerous. It could be simply a difference in monetary returns. For instance, an individual who owns a large plot of land would benefit much more from a regional irrigation system than an individual who owns only a small plot of land. It may also be that people perceive the importance of the public good differently. For instance, some neighbors more than others might enjoy the existence of a neighborhood swimming pool. Voluntary contributors in commons-based production organizations are another example. They likely receive (perceived or expected) benefits in excess of the costs from contributing to the project. The benefits may range from pure pleasure from contributing to the project to higher likelihood of getting a contract in traditional firms with experience in such commons-based organizations.

⁴Evidence of voluntary contribution without punishment can be found in Ledyard (1995). Examples of studies on decentralized punishment include Ostrom et al. (1992), Fehr and Gächter (2000), Carpenter (2006), and Egas and Riedl (2008).

⁵The same holds for various types of distributional preferences such as a concern for the income of the least

Our results indeed show that differences in punishment as well as in the response to being punished can explain why privileged groups fail to outperform normal groups. The main reason is that the punishment institution produces a much smaller increase in contributions by low-benefit subjects in privileged groups. This in turn can be attributed to relatively unstructured punishment behavior of high-benefit subjects in combination with a weak response in contributions to being punished by low-benefit subjects.

Experimental studies have investigated the role of incentives to cooperate (i.e, the marginal per capita return to contribute) in public-goods settings quite extensively. From these studies, it is well known that changes in the marginal per capita return have significant effects on contribution behavior (see Ledyard, 1995). A few studies have even made cooperation the dominant strategy (Saijo and Nakamura, 1995; Palfrey and Prisbrey, 1996, 1997). Our paper is also related to experiments investigating the effects of heterogeneous endowments in social dilemmas. The early findings of this literature are reviewed in Ostrom et al. (1994) and Ledyard (1995), in which it is reported that inequality leads to lower cooperation levels. However, more recent studies have found mixed results: some report a negative effect of inequality (van Dijk et al., 2002; Cherry et al., 2005), others a positive effect (Chan et al., 1996; Buckley and Croson, 2006), or no effect (Chan et al., 1999; Sadrieh and Verbon, 2006). There are also a few studies using field experiments. Cardenas (2003) finds that real-life income inequalities between Colombian villagers make the use of communication less effective in increasing cooperation inside the lab. Visser and Burns (2006) report that among South African fishermen, lab-induced income inequalities do not mitigate the effectiveness of punishment for promoting cooperation. The work most closely related to this paper is that of Fisher et al. (1995) and Brandts and Schram (2001). In both studies, the effect of unequal incentives to contribute within a group is explored. However, neither of the two investigates privileged groups as we do.

The rest of the paper is organized as follows. In Section 2 we describe the design of the experiment and derive qualitative predictions from theoretical models assuming fairness preferences and reciprocity. In Section 3 we present the results concerning subjects' contribution and punishment behavior. Section 4 closes with a summary of our main results and concludes.

2 Experimental Design

The basic game implemented in the experiment is a linear public-goods game. Each subject takes part in only one of four experimental treatments. The treatments differ in two dimensions: first, the public-goods game is either followed by a decentralized punishment stage a la Fehr and Gächter (2000) (*punishment* condition) or there is no such punishment stage (*baseline* condition); second, subjects in a group differ with respect to the marginal benefit they receive from the public good (*privileged* group condition) or do not differ in this respect (*normal* group condition). The game is played by the same group of three subjects for ten consecutive periods.

The public-goods game consists of a contribution stage in which each subject $i \in \{1, 2, 3\}$ receives an endowment of twenty tokens. Subjects simultaneously decide how many tokens, $c_i \in [0, 20]$, they contribute for the provision of the public good. The marginal monetary benefit for subject i for every token contributed by any group member is given by α_i whereas the marginal monetary benefit from not contributing a token equals 1.

well off (Rawls, 1971). On the other hand, if subjects care for overall efficiency, low contributions by low-benefit subjects are clearly undesirable. For more discussion on distributional preferences see Charness and Rabin (2002) and Engelmann and Strobel (2004).

Hence, from a purely monetary perspective, subject i 's dominant strategy is to contribute nothing if $\alpha_i < 1$. In contrast, if $\alpha_i > 1$, then subject i 's dominant strategy is to contribute the whole endowment. Additionally, if $\sum_i \alpha_i > 1$, then each token contributed increases the sum of earnings in the group, and hence, monetary surplus maximization is attained if everyone contributes all twenty tokens.⁶

In the punishment condition, the contribution stage of the public-goods game is followed by a punishment stage. In the punishment stage, subjects are first informed of the contributions of all other group members. Thereafter, each subject i simultaneously decides how many punishment points, $p_{ij} \in [0, 10]$, to assign to each other subject j . Each punishment point costs the punisher one token and reduces the earnings of the punished subject by three tokens.⁷ After the punishment decision, subjects are informed of the total number of punishment points assigned to them by other members of their group. As in Fehr and Gächter (2000), subjects receive no information about individual punishment behavior of other group members. At the end of a period, earnings of a subject i are given by⁸

$$\pi_i = 20 - c_i + \alpha_i \sum_j c_j - 3 \sum_{j \neq i} p_{ji} - \sum_{j \neq i} p_{ij}.$$

Importantly, in some treatments, groups differ in regard to the values of marginal benefits, α_i , assigned to subjects in a group. In *normal* groups, all subjects have the same $\alpha_i = 0.5$. In *privileged* groups, at the beginning of the experiment, a subject is randomly chosen to receive an $\alpha_j = 1.5$. The two other subjects in a privileged group have the same $\alpha_i = 0.5$. Each subject keeps his or her assigned marginal benefit for all ten periods. For convenience, we refer to subjects with an $\alpha_i = 0.5$ as *low-benefit* subjects and to subjects with an $\alpha_j = 1.5$ as *high-benefit* subjects. Clearly, from a monetary perspective, high-benefit subjects' dominant strategy is to contribute all their endowment. Note that the presence of a high-benefit subject precisely mirrors Olson's (1965) definition of privileged groups as being groups in which at least one member is willing to bear the cost of providing some level of the public good on his own. Note that because of the presence of low-benefit subjects, privileged groups are also prone to underprovision of the public good. The four different treatments are summarized in Table 1.

Normal and privileged groups differ only in the absence or presence of a high-benefit subject instead of a low-benefit subject. Thus, by comparing these group types, we can investigate the effect of the high-benefit subject on overall contribution levels when punishment is or is not available. Furthermore, since low-benefit subjects face the same monetary incentives in both types of groups, any differences in behavior of these subjects across group types must be caused by the presence of the high-benefit subject.

⁶In the experiment, all subjects are informed that they will interact for exactly ten rounds with the same individuals. Furthermore, they know the value of α_i of all group members, but neither during nor after the experiment do they get to know their identities.

⁷In line with other studies (Fehr and Gächter, 2000, e.g.), an upper limit for the amount of punishment points that could be allocated is introduced. The main reason is to avoid giving subjects with high earnings much more punishment power (as punishment is funded out their own earnings). Note also that subjects who are left with less than twenty tokens after the contribution stage can still use all their punishment points but would incur a loss in the respective period. Subjects, however, cannot be punished below zero. See footnote 8.

⁸In case a subject makes losses, these are calculated as follows: $\pi_i = \max[0, 20 - c_i + \alpha_i \sum_j c_j - 3 \sum_{j \neq i} p_{ji}] - \sum_{j \neq i} p_{ij}$.

Table 1: Experimental treatments

	Baseline	Punishment
Normal Groups	Only contribution stage $\alpha_i = 0.5$ for all subjects	Contribution and punishment stages $\alpha_i = 0.5$ for all subjects
Privileged Groups	Only contribution stage $\alpha_j = 1.5$ for one subject, $\alpha_i = 0.5$ for others	Contribution and punishment stages $\alpha_j = 1.5$ for one subject, $\alpha_i = 0.5$ for others

Under the assumption that individuals are selfish money maximizers, nobody is predicted to contribute a positive amount to the public good in normal groups. However, the evidence from previous public-goods experiments suggests some voluntary contribution. Typically, in early periods, subjects contribute substantial positive amounts that decline to low levels with repetition (see, e.g., Ledyard, 1995). In privileged groups in the baseline treatment, low-benefit individuals face the same monetary incentives as subjects in the normal groups. In contrast, high-benefit individuals earn more money the more they contribute to the public good. Therefore, when taking only material incentives into account, it is predicted that public-good provision is higher in privileged groups than in normal groups and that high-benefit subjects contribute all their endowment, whereas low-benefit subjects contribute nothing. Note, however, that the simultaneous presence of high- and low-benefit subjects may interact with nonmonetary motives of voluntary contribution in a nontrivial way. This is shortly discussed below.

Since punishment is costly, and hence, not credible for rational individuals maximizing their own payoff, predictions based on this assumption are the same as without a punishment option. However, there is considerable evidence that punishment is used to enforce high and stable contribution levels (Fehr and Gächter, 2000). Various models of generalized preferences have been proposed to explain contribution and punishment behavior in a range of experiments. In the following we briefly discuss some general predictions regarding punishment and voluntary contributions of such models and relate it to our experimental setup with privileged groups.

For normal groups, the predictions of outcome-based models of inequity aversion are that contribution levels will be relatively low in the baseline treatment and can be high in the punishment treatment (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). For privileged groups, the predictions are different. In the case of no punishment possibilities, any increase of contributions by the high-benefit individual increases his income and does not alter income differences. Hence, high-benefit individuals are predicted to contribute all their endowment, irrespective of their other-regarding preferences. In contrast, any increase in contribution by low-benefit individuals decreases their income and increases income differences. Consequently, low-benefit individuals are predicted not to contribute. Importantly, introducing the punishment opportunity does not change this prediction. The reason is that in the mentioned models of outcome-based inequity aversion the only credible motivation for costly punishment is the reduction of income differences. However, in privileged groups increasing contributions by a low-benefit individual increases inequality. Therefore, high-benefit individuals who are inequity averse have no incentive to

force low-benefit group members to contribute.⁹ Hence, in contrast to normal groups, the opportunity to punish is not predicted to increase contributions in privileged groups.

Predictions of intention-based models of social preferences are more ambiguous. In normal groups, high and low contributions have a straightforward interpretation as being a kind or an unkind action. This provides an incentive to punish low contributors and not to punish high contributors. In privileged groups, however, it is difficult to judge the kindness of actions of high-benefit individuals. In particular, high contributions directly benefit such individuals, and hence, may be interpreted as a selfish act. Consequently, in intention-based models the prediction for privileged groups depends critically on how intentions are modeled. For example, in Falk and Fischbacher (2006) kindness depends on relative income positions. Hence, low-benefit individuals will not regard high contributions by high-benefit individuals as kind, and high-benefit individuals will not consider low contributions by low-benefit individuals as unkind and thus deserving punishment.

In summary, purely outcome-based as well as intention-based models of social preferences suggest that introducing punishment opportunities has a weaker effect on contributions in privileged groups than in normal groups.

3 Results

The computerized experiment was conducted at the CREED laboratory of the University of Amsterdam. The experiment was programmed with z-Tree (Fischbacher, 2007). In total, eighty-one subjects participated in the experiment. Twenty-one (eighteen) participated in the baseline treatment without (with) punishment, and twenty-four (eighteen) in the privileged-group treatment without (with) punishment. This generated between six and eight independent observations per treatment. Each treatment was run in a separate session. A typical experimental session took one hour. On average, subjects were paid out €13.76 (approx. US\$17.50). More detailed experimental procedures and the instructions can be found online in the authors' personal Web pages. The online materials also include the software programs used, detailed descriptive statistics, and all the data required to replicate the results.

3.1 Voluntary Contributions

In this section we first discuss contribution behavior at the group level. Thereafter, we zoom into the behavior of low- and high-benefit subjects.

Figure 1 depicts the groups' mean total contributions during periods in (a) the baseline treatments and (b) the punishment treatments for both group types. In keeping with evidence from comparable public-goods experiments, contributions to the public good are relatively low and decrease over time in the baseline treatment with normal groups. At a higher level, similar dynamics of contributions are observed in privileged groups. On average,

⁹Certainly, as long as earnings end up equal, punishment can be used to enforce other cooperation levels. It would imply that in equilibrium, low-benefit individuals punish high-benefit individuals simply to reduce income differences. However, all such equilibria are Pareto-dominated by the equilibrium in which low-benefit individuals do not contribute.

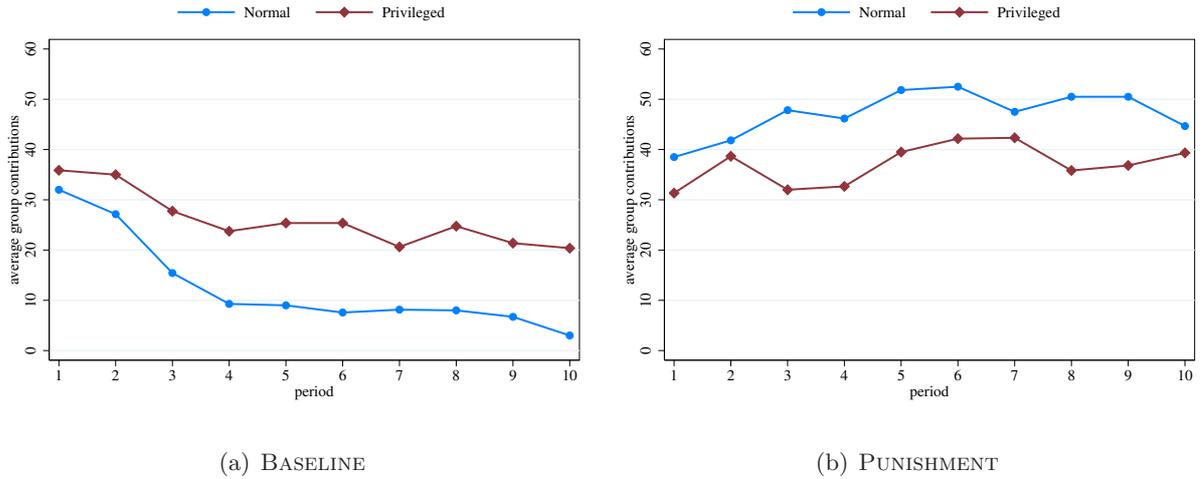


Figure 1: Contributions in the different groups and treatments

normal groups contribute 12.63 tokens, whereas privileged groups contribute 26.04 tokens. A nonparametric Mann-Whitney test clearly rejects equality of distributions ($p = 0.001$).¹⁰

In the baseline treatment, rank order correlations (Spearman’s ρ) of contributions on periods indicate that voluntary contributions significantly decline over time in normal as well as privileged groups ($\rho < -0.462$, $p < 0.001$). However, compared to privileged groups, contributions decline to a lower level in normal groups. This difference may be attributed to convergence toward different equilibrium contribution levels. Indeed, in normal groups, contributions decline to (almost) zero. In privileged groups, last-period contributions are, on average, 20.37 tokens. This is indistinguishable from what is obtained when the high-benefit subject contributes everything and the two low-benefit subjects contribute nothing.

In the punishment treatment, contribution levels are higher than in the baseline treatment, and they do not decline over time for both group types. Note that in contrast to the baseline without punishment, it is no longer the case that contributions are highest in privileged groups. In fact, the average total contribution of normal groups is 47.19 whereas it is only 37.08 in privileged groups. Although a Mann-Whitney test does not reject the hypothesis that average contributions come from the same distribution ($p = 0.120$) it is intriguing that privileged groups lose their privileged status once punishment is possible. Additionally, according to Spearman’s rank-order correlations, contributions stay constant over time in privileged groups ($\rho = 0.149$, $p = 0.129$) but exhibit a significant upward trend in normal groups ($\rho = 0.248$, $p = 0.029$).

The quantitative effect of punishment on contributions also strongly differs between group types. In normal groups, the increase in contributions because of punishment is large (34.56 tokens) and statistically significant ($p < 0.001$). For privileged groups, the increase is also statistically significant ($p = 0.041$) but quantitatively rather

¹⁰Unless otherwise noted, for between-treatment comparisons, we use a Mann-Whitney test, which tests the hypothesis that two independent samples are drawn from the same population. For within-treatment comparisons, we use the Wilcoxon signed-ranks test for matched pairs. If not otherwise indicated, we apply one-sided test statistics and use group averages across all periods as independent observations. All reported p -values are exact values.

small (11.04 tokens). The difference in the effectiveness of punishment is also illustrated by the fact that in normal groups full contribution by all group members is achieved in 35.5 percent of all periods whereas in privileged groups it is observed only in 6.7 percent. Hence, although the opportunity to punish increases contributions in both types of groups the magnitude of this effect is strikingly different between normal and privileged groups. Particularly, it is much stronger in groups in which all subjects benefit equally from the public good. We summarize the findings discussed so far in our first result.

Result 1 VOLUNTARY CONTRIBUTIONS IN NORMAL AND PRIVILEGED GROUPS.

In the absence of punishment, there is underprovision of the public good in normal as well as privileged groups, although voluntary contributions are higher in the latter. If punishment is possible, there is still underprovision but at a lower level, and the difference between group types disappears. Punishment possibilities increase contributions in both group types, but the effect of punishment is much smaller in privileged groups.

Given the diametrically opposed material incentives for high- and low-benefit subjects, it is to be expected that their contribution behavior differs. Figure 2 shows that this is indeed the case. In privileged groups, high-benefit subjects contribute significantly more than low-benefit subjects in treatments with and without punishment. For high- and low-benefit subjects, the average contribution levels are 17.03 and 4.50 tokens, respectively, in the baseline treatment, and 18.43 and 9.32 tokens, respectively, in the punishment treatment. In both cases, equality of distribution is rejected at any conventional level of statistical significance ($p < 0.005$).

Remarkably, high-benefit subjects do not follow their dominant strategy to contribute the full endowment of twenty tokens. This behavior is particularly pronounced and statistically significant in the baseline treatment without punishment ($p = 0.010$). Recall that for high-benefit subjects, contributing the whole endowment is a dominant strategy in absolute as well as relative material terms, and it is also the best for their group's total payoff. Therefore, it is impossible to explain the observed deviations from full contribution on the basis of outcomes only. High-benefit subjects act (at least to some extent) conditionally cooperatively in terms of contributions. It appears that they "punish" low-benefit subjects for low contributions by not contributing all their endowment. To test this conjecture, we regress the contribution of the high-benefit subject on the sum of contributions of the low-benefit subjects in the previous period. In addition, we control for a possible end-game effect by including a dummy variable for period ten.¹¹ The regression results show a positive and statistically significant coefficient for the lagged contributions of low-benefit subjects ($p = 0.033$) and for the last-period dummy ($p < 0.000$). This clearly indicates conditional cooperative behavior of high-benefit subjects up to the last periods. It also explains why their contributions decrease over time until the end-game effect sets in (see Figure 2(a)). In the punishment treatment the coefficient for the lagged contributions of low-benefit subjects is statistically not significantly different from zero ($p = 0.789$). This difference between treatments is consistent with the interpretation that in the presence of the more effective explicit punishment possibility there is no need for high-benefit subjects to resort to lowering

¹¹We use Tobit regressions, censoring contributions at the maximum of twenty tokens. In the baseline treatment, the resulting coefficients are 0.75 (0.35) for the low-benefit subjects' lagged contributions and 89.49 (8.49) for the last-period dummy ($\chi^2 = 118.48$, $p < 0.000$). In the punishment treatment, they are -0.05 (0.20) and 64.57 (12.52), respectively ($\chi^2 = 34.12$, $p < 0.00$). Numbers in parenthesis are robust standard errors (White, 1980).

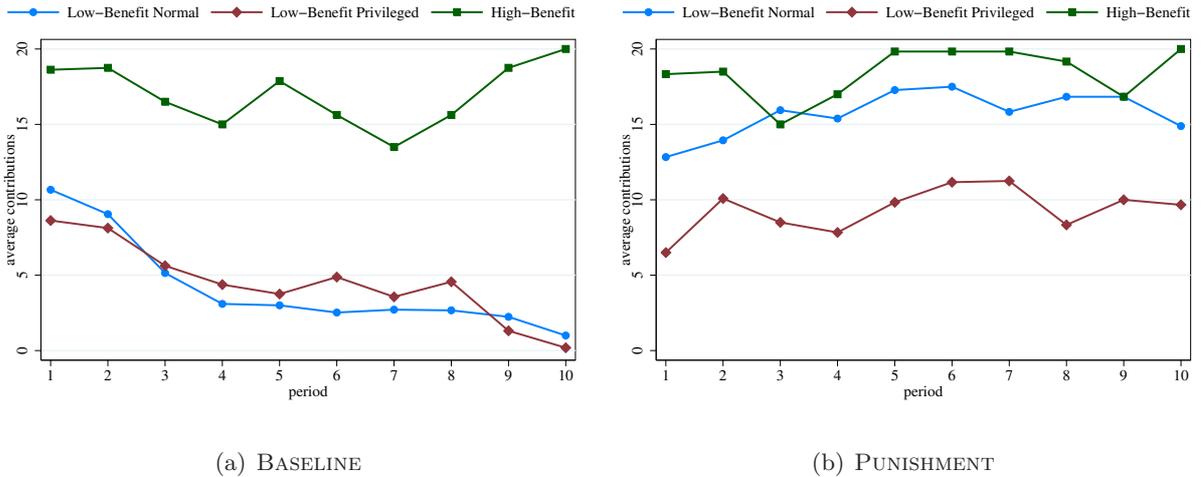


Figure 2: Contributions by low- and high-benefit subjects in each treatment

contributions as a means of disciplining low-benefit subjects.¹² In the next section, we provide evidence showing that high-benefit subjects indeed do punish low-benefit subjects.

In the no-punishment treatment, low-benefit subjects behave very similarly in normal and in privileged groups (see Figure 2(a)). Average contributions in both types of groups are almost equal (4.21 and 4.50 in normal and privileged groups, respectively) and statistically indistinguishable ($p = 0.347$). Contributions of low-benefit subjects also decline significantly over time in both types of groups (Spearman's $\rho = -0.584$ and $\rho = -0.590$, $p < 0.001$). In contrast, in the punishment treatment, contributions of low-benefit subjects differ starkly in normal and in privileged groups (see Figure 2(b)). Average contributions amount to 15.73 tokens in normal groups but are only 9.32 tokens in privileged groups. A Mann-Whitney test clearly rejects equality of distributions ($p = 0.033$). Interestingly, this difference in contributions is present already in the first period. In period one, the average contribution in normal groups is 12.83, which is almost twice as high as in privileged groups, in which it is only 6.50 tokens ($p = 0.004$). Moreover, this difference does not vanish but rather increases over time. Spearman rank-order correlations show that in normal groups, contributions exhibit a statistically significant upward trend ($\rho = 0.247$, $p = 0.028$), whereas no such trend can be detected in privileged groups ($\rho = 0.104$, $p = 0.215$).

Clearly, in privileged groups contributions of low-benefit subjects are significantly higher in the punishment treatment than in the baseline treatment ($p = 0.041$). However, in comparison to the treatment effect for normal groups, the observed increase in contributions is rather meager. For normal groups, the average increase in contributions from the baseline treatment to the punishment treatment is 11.52 tokens, whereas it is only 4.82 tokens for low-benefit subjects in privileged groups ($p = 0.021$). Together, these observations indicate that the effectiveness of punishment concerning low-benefit subjects strongly differs between group types to the disadvantage of privileged groups. We summarize our findings in the following result.

¹²Another possibility for under-provision is spiteful behavior, as reported by Saijo and Nakamura (1995), in groups where $\alpha_i > 1$ for all group members. However, this fails to explain the observed relation with the contribution behavior of low-benefit subjects. Brandts and Schram (2001) also find a small degree of under-provision, which they attribute to errors. However, because of their design, they cannot explore the conditional-cooperation possibility.

Result 2 CONTRIBUTIONS OF HIGH- AND LOW-BENEFIT SUBJECTS.

Without punishment opportunities, although high-benefit subjects contribute significantly more than low-benefit subjects, there is significant underprovision of the public good by high-benefit subjects. With punishment opportunities, contributions of both types of subjects increase. However, contributions of low-benefit subjects in privileged groups reach only 60 percent of the contribution level in normal groups.

3.2 Punishment

The amount of tokens lost because of punishment, per subject and period, is similar in privileged and normal groups (3.72 and 3.23, $p = 0.409$). In privileged groups, low-benefit subjects are punished slightly more than high-benefit subjects (4.03 vs. 3.10). However, this difference is not statistically significant ($p = 0.310$). Over time, punishment decreases significantly in both types of groups (Spearman's $\rho < -0.365$, $p < 0.005$). Allocated punishment points are also similar across types. On average, low-benefit subjects in normal and in privileged groups deal out 0.54 and 0.51 punishment points, respectively. High-benefit subjects in privileged groups assign slightly more punishment points (0.84). The differences are statistically not significant ($p > 0.155$).

In the remainder of this subsection, we discuss possible determinants of punishment and how they differ, first between group types, and second, between high- and low-benefit subjects in privileged groups. Thereafter, we investigate if and how high- and low-benefit subjects react with their public-good contributions to received punishment. This will shed light on the observed differences in contributions of low-benefit subjects in normal and privileged groups.

For the analysis of possible determinants of punishment behavior we apply Tobit regressions. The dependent variable is the amount of punishment points, p_{ij} , subject i allocates to subject j .¹³ As explanatory variables, we use the deviation of j 's contribution from i 's contribution, and the deviation of j 's contribution from the contribution of the third subject in the group k , where we allow for different coefficients for positive and negative deviations. In addition, we control for j 's contribution for cases in which all subjects contributed their whole endowment (this accounts for 21 percent of all periods but only 9 percent of all punishment points) and for a time trend. For privileged groups, we run separate regressions for high- and low-benefit subjects. In the regressions for low-benefit subjects, we included a dummy variable if the target person j is a high-benefit subject and also interact this variable with the period of play.¹⁴ The results are reported in Table 2.

We first discuss behavior of low-benefit subjects in normal and privileged groups (columns two and three in the table). In line with other studies, in both groups, negative deviations from own contributions are strongly

¹³As the amount of punishment points that could be allocated is constrained, we run the Tobit regressions censoring p_{ij} at zero and ten points. All reported standard errors are based on robust estimates, and dependence of observations within groups is accounted for by clustering on each period and group (White, 1980). Note that endogenous censoring could occur for subjects with earnings of less than twenty tokens after the contribution stage, because they might be not willing to punish (strongly), as they would incur losses. However, situations with potential (actual) losses occurred only in 1.94 percent (0.28 percent) of all cases.

¹⁴If we drop these two explanatory variables concerning the high-benefit subject, all reported results stay quantitatively similar and qualitatively the same. The regression coefficients of the other explanatory variables change only marginally, and all reported (non-)significances remain.

Table 2: Punishment assigned to j by i depending on differences in contributions

Dependent variable: p_{ij}	<i>Normal Groups</i>		<i>Privileged Groups</i>	
	All	Low-benefit	High-benefit	
Punishment given by i to j				
Contribution of j	0.121	0.120	−0.086	
c_j	(1.59)	(1.27)	(0.74)	
Positive deviation from c_i	0.148***	−0.105	−0.311	
$\max(c_j - c_i, 0)$	(2.58)	(0.93)	(0.93)	
Negative deviation from c_i	0.312***	0.420***	−0.074	
$\max(c_i - c_j, 0)$	(4.38)	(3.10)	(0.73)	
Positive deviation from c_k	−0.537***	−0.137	−0.166	
$\max(c_j - c_k, 0)$	(3.76)	(1.21)	(1.53)	
Negative deviation from c_k	0.118	0.058	0.132	
$\max(c_k - c_j, 0)$	(1.30)	(0.55)	(0.97)	
Fully contributing group	−4.360***	2.977	−14.606***	
1 if $c_i + c_j + c_k = 60$	(4.74)	(1.17)	(7.11)	
Period	−0.081	−0.019	−0.457**	
t	(1.15)	(0.10)	(2.43)	
j is a high-benefit subject		5.000***		
1 if $\alpha_j = 1.5$		(3.39)		
j is a high-benefit subject \times Period		−0.696***		
t if $\alpha_j = 1.5$		(3.68)		
Constant	−3.253**	−6.292***	2.721	
	(2.17)	(2.60)	(1.22)	
# Observations	360	240	120	
Log-likelihood	−271.96	−191.05	−140.38	
χ^2	124.88	38.65	56.71	

Note: Numbers in parentheses are t -statistics estimated with White’s heteroscedasticity consistent covariance matrix estimator: asterisks indicate significance at the 1% level (***), 5% level (**), and 10% level (*).

and statistically significantly punished. There is no statistically significant difference between normal and privileged groups (Wald test, $p = 0.982$). Positive deviations are punished only in normal groups. Such “perverse” punishment of high contributors occurs at a relatively low level and is also found by other authors (see, e.g., Cinyabuguma et al., 2006; Gächter and Herrmann, 2007; Egas and Riedl, 2008). Probably the most interesting difference between groups is the punishment response toward differences in contributions between the target person j and the third person k . In normal groups, a positive deviation of j compared to k leads to a strong and significant decrease in punishment, whereas in privileged groups, no such effect is detected (Wald test for

the difference between treatments, $p = 0.004$). In privileged groups, the highest contributor is almost always the high-benefit subject. Therefore, the above result indicates that relatively high contributions of such a subject are not rewarded by less punishment. Additionally, low-benefit subjects appear to punish high-benefit subjects simply for being high-benefit subjects, especially at the beginning of the game (see Table 2, column three).

In contrast to low-benefit subjects, punishment behavior of high-benefit subjects is largely independent of relative contributions except that they reward full contributions of low-benefit subjects with a strong and significant decrease in punishment. High-benefit subjects also punish less in later periods as indicated by the statistically significantly negative linear time trend.

Result 3 PUNISHMENT BEHAVIOR IN NORMAL AND PRIVILEGED GROUPS.

In both types of groups, low-benefit subjects' punishment behavior is sensitive to relative contributions. They severely punish deviations from their own contribution. Additionally, in normal groups, low-benefit subjects punish less if the target person contributes more than the third group member. In contrast, high-benefit subjects' punishment behavior is insensitive to the relative contributions of low-benefit subjects.

On average, high-benefit subjects do not punish less than low-benefit subjects. Therefore, it is likely that the unstructured punishment pattern of high-benefit subjects regarding contributions of low-benefit subjects plays a key role in explaining the comparatively low contributions of low-benefit subjects in privileged groups.

Another important aspect for the effectiveness of punishment in increasing contributions is the response of those being punished. Recall that even when punished, the motivation of low-benefit subjects to contribute to the public good might be different between normal and privileged groups. For example, recall that in models of inequity aversion, low-benefit types have little incentive to contribute in privileged groups because any increase in contributions decreases their earnings and increases inequality, especially vis-à-vis the high-benefit subject. To investigate the response to being (not) punished and its potentially differential characteristics across group types, we have estimated the mean change in contributions from period t to period $t + 1$, depending on whether the low-benefit subject was punished in period t or not (controlling for the subject's contribution and the contribution of others in period t). For each group type, we ran two separate ordinary least squares (OLS) regressions, one for periods in which the low-benefit subject was punished and another for periods in which the low-benefit subject was not punished.

Figure 3(a) visualizes the estimation results for normal groups and Figure 3(b) for privileged groups.¹⁵ The lines show the estimated change in contributions as a function of the (low-benefit) subject's contribution in period t for cases in which the subject was punished (brown \diamond) or not punished (blue \bullet) in period t . When not punished, low-benefit subjects have, on average, a tendency to decrease their contributions in both types of groups, particularly if they contributed a large amount. In both normal and privileged groups, subjects who are not punished and

¹⁵Denoting Δc_i as the change in contributions from period t to $t + 1$, c_i as the contribution of subject i in period t and \bar{c}_{-i} as the average contribution of the other two subjects in period t , the resulting regression equations are (t -statistics—in parenthesis—are calculated by clustering on each group and period): (i) for punished subjects in normal groups, $\widehat{\Delta c_i} = 9.39(2.81) - 0.63c_i(4.47) + 0.02\bar{c}_{-i}(0.11)$; (ii) for not punished subjects in normal groups, $\widehat{\Delta c_i} = -4.07(1.76) + 0.042c_i(0.28) + 0.14\bar{c}_{-i}(1.28)$; (iii) for punished subjects in privileged groups, $\widehat{\Delta c_i} = 0.30(0.08) - 0.37c_i(2.04) + 0.36\bar{c}_{-i}(1.18)$; and (iv) for not punished subjects in privileged groups, $\widehat{\Delta c_i} = -1.75(1.28) - 0.23c_i(3.01) + 0.20\bar{c}_{-i}(1.58)$. Figure 3 is drawn with mean values for \bar{c}_{-i} .

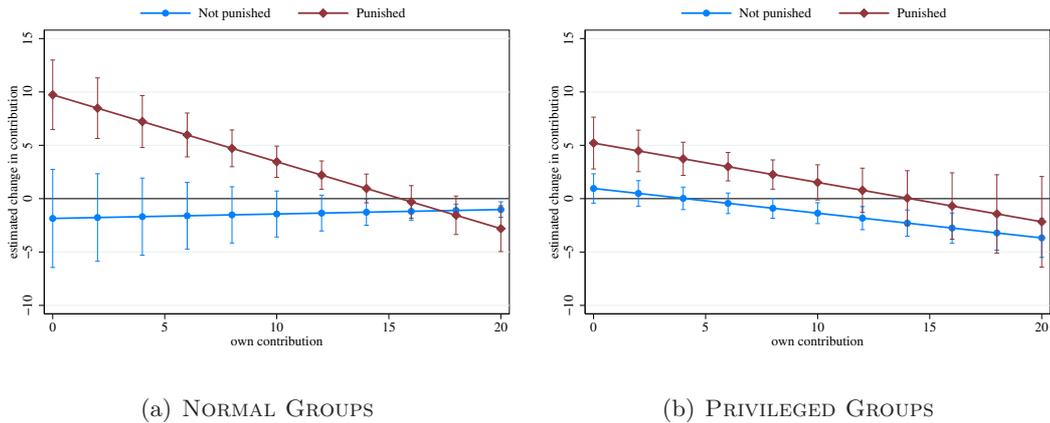


Figure 3: Mean change in contribution by low-benefit subjects depending on previous contributions and punishment

contributed more than 14 tokens significantly decrease their contribution (Wald tests $p < 0.046$), although one can see that the decrease is not large in an absolute sense. If we compare the estimated change in contribution between group types we find it to be statistically different only at very high contribution levels (contributions above 17 tokens, Wald tests $p < 0.033$). In this case, subjects in privileged groups decrease their contribution more if not punished.

When punished, low-benefit subjects respond similarly in both types of groups only if they contributed a large amount of their endowment. For contribution levels of 14 tokens or more, subjects do not respond to punishment by significantly increasing contributions, in normal or in privileged groups (Wald tests, $p > 0.123$).

In stark contrast, low contributors respond to punishment by significantly increasing their contribution, and they do so differently in normal and privileged groups. For instance, after being punished, subjects who free ride completely (zero contribution) increase contributions about twice as much in normal groups as in privileged groups. Wald tests confirm that for contribution levels of 9 tokens or less, punished low-benefit subjects increase their contributions significantly less in privileged groups than in normal groups ($p < 0.047$). Thus, over a large range of contribution levels, punishment produces a much weaker positive response in privileged groups than in normal groups. We summarize these findings in the following result.

Result 4 RESPONSE TO PUNISHMENT IN NORMAL AND PRIVILEGED GROUPS.

In both types of groups, low-benefit subjects who contribute low amounts respond with a significant increase in contributions if punished. Over a large range of low contribution levels, this effect is significantly and substantially weaker in privileged groups.

The fact that high-benefit subjects in privileged groups exhibit a relatively unstructured punishment behavior (Result 3), together with the above finding that strongly free-riding low-benefit subjects exhibit a much weaker response to punishment in privileged groups, can explain why the opportunity to punish has a much smaller effect on the contributions of low-benefit subjects in privileged groups compared to normal groups (Result 2).

3.3 Efficiency and Inequality

The socially beneficial effect of the opportunity to punish free riders is that it increases contributions to the public good, and thereby, also total monetary earnings. The downside is that if costly punishment is actually administered, it destroys resources, and hence, decreases total monetary surplus. In this section we examine how punishment affects earnings of low- and high-benefit subjects in normal and privileged groups.

Taken across all periods, in the baseline treatment with normal groups, average earnings per subject and period are equal to 22.10 tokens, which is statistically not significantly different from the 23.67 when punishment is possible ($p = 0.366$). However, the dynamics of earnings reveal a clear divergence between the baseline and the punishment treatment. In the baseline treatment earnings significantly decrease over time, whereas they show a clear upwards trend in the punishment treatment (Spearman's $\rho = -0.584$, $p < 0.001$ and $\rho = 0.322$, $p = 0.006$). Furthermore, earnings in the second half (periods six through ten) of the public-goods game are significantly larger in the punishment treatment compared to the baseline treatment (21.11 tokens and 24.86 tokens, $p = 0.017$).

In privileged groups, the effect of punishment on earnings is more ambiguous. When taken across all periods, average earnings per subject and period are similar in treatments without and with punishment (33.01 and 33.58 tokens, $p = 0.426$). Earnings also show a decreasing trend in the baseline treatment and an increasing trend in the punishment treatment (Spearman's $\rho = -0.463$, $p < 0.001$ and $\rho = 0.254$, $p = 0.025$). However, unlike in normal groups, in later periods (6-10) the earnings in the punishment treatment are only marginally significantly larger than in the baseline treatment (31.25 tokens without punishment and 36.23 tokens with punishment, $p = 0.091$). The reason for this difference is that compared to normal groups, low-benefit subjects in privileged groups increase their contributions less in the punishment treatment.

Punishment not only affects total earnings but also how these earnings are distributed. A natural and simple measure of inequality is the standard deviation of earnings per period within each group. In normal groups, the average standard deviation across periods is 4.76 in the baseline treatment and 3.43 when punishment is possible. The difference is statistically not significant ($p = 0.183$). In privileged groups, however, punishment significantly increases inequality. Whereas in the baseline treatment, the average standard deviation is 11.56 it increases to 23.69 in the punishment treatment ($p = 0.041$).

The underlying reason for increased inequality in privileged groups is that punishment has a stark differential effect on the earnings of low- and high-benefit subjects. High-benefit subjects gain from the punishment institution: their earnings increase from 42.01 tokens in the baseline treatment to 52.38 tokens in the punishment treatment ($p = 0.071$). In contrast, low-benefit subjects in privileged groups do not gain at all from the punishment institution. In effect, compared to the baseline treatment, their earnings are significantly lower in the punishment treatment (28.51 tokens without punishment and 24.18 tokens with punishment, $p = 0.004$). If we compare earnings of low-benefit subjects in normal and privileged groups, we find that in the treatment without punishment, low-benefit subjects earn significantly more in privileged groups ($p < 0.001$), but this is no longer the case in the treatment with punishment ($p = 0.409$). In other words, low-benefit subjects “exploit” the high-benefit subject when punishment is not possible but are no longer able to do so when punishment is introduced. This is important as it illustrates that privileged groups are not necessarily characterized by the exploitation of the “rich by the poor” as is usually claimed (Olson, 1965; Sandler, 1992). We summarize the findings of this section in our final result.

Result 5 EFFICIENCY AND INEQUALITY.

In normal groups, the opportunity to punish leads to higher overall earnings with no effect on their distribution. In privileged groups, punishment does not significantly increase total earnings but strongly increases inequality at the cost of low-benefit subjects. Consequently, with punishment, low-benefit subjects do not benefit from being part of a privileged group.

4 Discussion and Conclusion

In this paper, we provide evidence that Olson’s (Olson, 1965) privileged groups are not as privileged as conjectured, especially if punishment is possible. In line with standard theoretical predictions, contributions are higher in privileged groups than in normal, nonprivileged groups, and underprovision of the public good is observed in both group types. Surprisingly, in privileged groups, this is not only because of free-riding by low-benefit subjects but also because of nonoptimal contributions of subjects whose dominant strategy is to contribute fully. We observe that high-benefit subjects reciprocate to low contributions of low-benefit subjects by lowering their contributions, despite the fact that contributing fully maximizes absolute and relative individual earnings as well as group earnings. When a punishment option is introduced, privileged groups lose their privileged status completely. Total contributions in normal and privileged groups are virtually identical. In addition, the introduction of punishment significantly increases earnings inequality in privileged groups as high-benefit subjects avoid being exploited by low-benefit subjects. We identify the following behavioral features that account for this relative disadvantage of privileged groups. With a punishment option, low-benefit subjects in privileged groups increase their contributions to a much smaller extent than their counterparts in nonprivileged groups, in which all have the same incentive to free ride. The reason for this is twofold. First, high-benefit subjects exhibit an unstructured punishment pattern, and second, free-riding low-benefit subjects increase their contributions in response to punishment to a much smaller extent in privileged groups than in normal groups.

Formal models of other-regarding preferences capture, to some extent, the possibility that punishment is less effective in privileged groups than in normal groups. However, the theoretical argument supporting this result is not in agreement with our empirical observations. These models predict that low-benefit subjects do not contribute because high-benefit subjects who have a preference for equality are not willing to punish them. In contrast, our results show that high-benefit subjects do punish free-riding low-benefit subjects (albeit in a relatively unstructured way) but that free-riding low-benefit subjects are very reluctant to increase contributions despite the fact that they are punished. This is an interesting finding because it demonstrates that subjects in one role (high-benefit) may be willing to undertake a costly action (punishing), although it increases earnings inequalities, but that subjects in another role (low-benefit) are unwilling to take a costly action (increasing contributions) because it increases inequality. This strongly suggests that an individual’s disutility from punishment depends not only on the received punishment points but also on who is punishing and whether punishment is perceived as fairly administered. To account for such behavior, a more general theoretical model of behavior is called for. Our results clearly indicate that such a model should take into account important heterogeneity parameters and the roles subjects in which are immersed.

Our results also highlight the importance of a common notion of “fair” or “kind” behavior for mechanisms such as decentralized punishment to be effective. In homogeneous groups, contributions to a public good are

unambiguously kind because they increases everyone’s payoff at a cost for the contributing individual. In privileged groups, however, high contributions by high-benefit subjects may be perceived as selfish (because they increase the contributor’s earnings), unfair (because they increase earning inequalities), or kind (because they increase the earnings of others). Low contributions of low-benefit subjects can similarly be interpreted in different ways. The fact that high-benefit subjects punish low-benefit subjects but low-benefit subjects only reluctantly increase contributions is consistent with self-serving interpretations of the “correct” fairness norm (see, e.g. Babcock et al., 1995; Gächter and Riedl, 2005). That diverging perceptions of what constitutes fair behavior can have detrimental effects on cooperation is also suggested by recent research on emotions and social behavior. For instance, Hopfensitz and Reuben (2007) and Reuben and van Winden (2007) find that sanctions enforce cooperative behavior only if people feel they have transgressed a social norm, and hence, deserve punishment (see also Bowles and Gintis, 2005). If low-benefit subjects in privileged groups conceive that they do not deserve sanctioning, a weak response, as observed in our experiment, may be the consequence. In privileged groups, such diverging perceptions of fair behavior are the likely reason for costly punishment being relatively ineffective and wasteful.

In an early study, Johnson and Libecap (1982) argue theoretically that group heterogeneity considerably complicates the private enforcement of cooperative behavior. The results of our paper add to this finding the experimental evidence that the behavioral effectiveness of an institution in overcoming social dilemmas can crucially depend on group characteristics, an observation also in line with results showing that sanctioning institutions can have very different effects on participants behavior with diverging cultural backgrounds (Carpenter and Cardenas, 2007; Gächter et al., 2008). This implies that the design of effective—formal as well as informal—enforcement institutions asks for a proper understanding of the interplay of heterogeneous group characteristics, as in privileged groups, behavioral inclinations and biases, and incentives. Therefore, a promising line for future research is the endogenous formation of groups and institutions in the presence of heterogeneity. For example, if free association to groups is possible, low-benefit subjects might be more reluctant to punish high-benefit subjects if doing so induces their leaving the group. Additionally, if punishment is governed endogenously through voting, low-benefit subjects might respond to punishment more strongly with increasing contributions if punishment from high-benefit subjects can be prohibited.

References

- Babcock, L., Camerer, C. F., and Samuel, S. I. (1995). Biased judgments of fairness in bargaining. *American Economic Review*, 85:1337–1343.
- Benkler, Y. (2002). Coase’s penguin, or, linux and *The Nature of the Firm*. *The Yale Law Journal*, 112:369–446.
- Bolton, G. E. and Ockenfels, A. (2000). A theory of equity, reciprocity, and competition. *American Economic Review*, 90:166–193.
- Bowles, S. and Gintis, H. (2005). Pro-social emotions. In Blume, L. E. and Durlauf, S. N., editors, *The Economy as a Complex Evolving System III: Essays in Honor of Kenneth Arrow*. Oxford University Press, Oxford.
- Brandts, J. and Schram, A. (2001). Cooperation and noise in public goods experiments: applying the

- contribution function approach. *Journal of Public Economics*, 79:399–427.
- Buckley, E. and Croson, R. (2006). Income and wealth heterogeneity in the voluntary provision of linear public goods. *Journal of Public Economics*, 90:935–955.
- Cardenas, J. C. (2003). Real wealth and experimental cooperation: experiments in the field lab. *Journal of Development Economics*, 70:263–289.
- Carpenter, J. and Cardenas, J. C. (2007). An inter-cultural examination of cooperation in the commons. Working paper, Middlebury College.
- Carpenter, J. P. (2006). The demand for punishment. *Journal of Economic Behavior & Organization*, 62:522–542.
- Chan, K. S., Mestelman, S., Moir, R., and Muller, R. A. (1996). The voluntary provision of public goods under varying income distributions. *The Canadian Journal of Economics*, 29:54–59.
- Chan, K. S., Mestelman, S., Moir, R., and Muller, R. A. (1999). Heterogeneity and the voluntary provision of public goods. *Experimental Economics*, 2:5–30.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117:817–869.
- Cherry, T. L., Kroll, S., and Shogren, J. (2005). The impact of endowment heterogeneity and origin on public good contributions: Evidence from the lab. *Journal of Economic Behavior & Organization*, 57:357–365.
- Cinyabuguma, M., Page, T., and Putterman, L. (2006). Can second-order punishment deter perverse punishment? *Experimental Economics*, 9:265–279.
- Egas, M. and Riedl, A. (2008). The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society B – Biological Sciences*, 275(05-065/1):871–878.
- Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency and maximim preferences in simple distribution experiments. *American Economic Review*, 94:857–869.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness: Intentions matter. *Games and Economic Behavior*, 62:287–303.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, 54:293–315.
- Fehr, E. and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90:980–994.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114:817–868.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10:171–178.
- Fisher, J., Isaac, M. R., Schatzberg, J. W., and Walker, J. M. (1995). Heterogeneous demand for public goods: Effects on the voluntary contributions mechanism. *Public Choice*, 85:249–266.
- Gächter, S. and Herrmann, B. (2007). The limits of self-governance in the presence of spite: Experimental evidence from urban and rural russia. Discussion paper no. 2007-11, CeDEx.

- Gächter, S., Herrmann, B., and Thöni, C. (2008). Antisocial punishment across societies. *Science*, 319:1362–1367.
- Gächter, S. and Riedl, A. (2005). Moral property rights in bargaining with infeasible claims. *Management Science*, 51(2):249–263.
- Hopfensitz, A. and Reuben, E. (2007). The importance of emotions for the effectiveness of social punishment. Discussion paper 05-075, Tinbergen Institute.
- Isaac, M. R., Walker, J. M., and Thomas, S. (1984). Divergent evidence on free riding: An experimental examination of possible explanations. *Public Choice*, 43:113–149.
- Johnson, R. N. and Libecap, G. D. (1982). Contracting problems and regulation: The case of fishery. *American Economic Review*, 72(5):1005–1022.
- Ledyard, J. O. (1995). Public goods: A survey of experimental research. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*, pages 111–194. Princeton University Press, New Jersey.
- Olson, M. (1965). *The logic of collective action*. Harvard University Press, Cambridge.
- Ostrom, E., Gardner, R., and Walker, J. (1994). *Rules, Games, and Common-pool Resources*. University of Michigan Press, Ann Arbor.
- Ostrom, E., Walker, J. M., and Gardner, R. (1992). Covenants with and without a sword: Self governance is possible. *American Political Science Review*, 86:404–417.
- Palfrey, T. R. and Prisbrey, J. E. (1996). Altruism, reputation and noise in linear public good experiments. *Journal of Public Economics*, 61:409–427.
- Palfrey, T. R. and Prisbrey, J. E. (1997). Anomalous behavior in public goods experiments: How much and why? *American Economic Review*, 87:829–846.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge.
- Reuben, E. and van Winden, F. (2007). Fairness and shame in the power to take. Discussion paper 05-014, Tinbergen Institute.
- Sadrieh, A. and Verbon, H. (2006). Inequality, cooperation, and growth: An experimental study. *European Economic Review*, 50:1197–1222.
- Saijo, T. and Nakamura, H. (1995). The spite dilemma in voluntary contribution mechanism experiments. *Journal of Conflict Resolution*, 39:535–560.
- Sandler, T. (1992). *Collective Action: Theory and Application*. The University of Michigan Press.
- van Dijk, F., Sonnemans, J., and van Winden, F. (2002). Social ties in a public good experiment. *Journal of Public Economics*, 85:275–299.
- Visser, M. and Burns, J. (2006). Bridging the great divide in south africa: Inequality and punishment in the provision of public goods. Working paper, Göteborg University.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48:817–38.