JOURNAL OF PUBLIC ECONOMIC THEORY

Discrimination via Exclusion: An Experiment on Group Identity and Club Goods

SURAJEET CHAKRAVARTY
University of Exeter

MIGUEL A. FONSECA
University of Exeter

Abstract

We study using laboratory experiments the impact on cooperation of allowing individuals to invest in group-specific, excludable public goods. We find that allowing different social groups to voluntarily contribute to such goods increases total contributions. However, a significant proportion of that contribution goes toward the group-specific club good rather than the public good, even when the latter has higher financial returns to cooperation. We find significant evidence of in-group biases, which are manifested by positive in-group reciprocity. That is, club goods allow subjects to display their preferences for interaction with their in-group members, as well as positive in-group reciprocity.

1. Introduction

Individuals often form groups designed to produce club goods. These goods can only be accessed by group members to the detriment of others. Many groups that provide club goods are institutional expressions of various forms of group identities. Membership of those groups is often drawn along ethnic, gender, or religious lines. Examples include some private member clubs, whose externalities are the business networks membership provides, as well as faith schools.[1]

One way to understand the existence of club goods is through a social identity theory framework. This theory argues that membership of social groups shapes individual preferences. This may be because group membership instills particular norms of

[1] For a discussion of the role of private member clubs and their increasing relevance in society, see *The Economist* (2012). Faith schools are learning institutions that are typically operated and/or financed by members of a particular faith and community for the children of that same faith or community. There are 20,000 maintained or government-supported schools in England, of which 7000 are faith schools (Department of Education, UK, 2010).

Surajeet Chakravarty, University of Exeter, Exeter, Devon EX4, UK (s.chakravarty@exeter.ac.uk).
Miguel A. Fonseca, University of Exeter, Exeter, Devon EX4, UK (m.a.fonseca@exeter.ac.uk).

We are very grateful to Tim Miller for his outstanding help in programming the software and running the sessions. Financial support from ESRC grant RES-000-22-394 is gratefully acknowledged. Comments from participants at the 2013 ESA World Meetings are gratefully acknowledged.

Received September 29, 2014; Accepted October 8, 2015.

Journal of Public Economic Theory, 19 (1), 2017, pp. 244–263.

behavior (Akerlof and Kranton 2000), or because it leads individuals to have a concern for the status of the group to which they belong (Shayo 2009). Alternatively, individuals may have greater concerns for the welfare of members of their group than that of outsiders (Tajfel *et al.* 1971; Chen and Li 2009). As societies become more diverse, people often rely on their communities to provide certain club goods in order to maintain their distinct identity. As a result, we observe services such as schools, cultural, and/or leisure activities being provided by different social groups with a view to be used by these specific constituencies alone.

Club goods are typically close substitutes for public goods, and the former's existence can have important implications regarding both the financing and the utility derived from the consumption of the latter. This raises important questions: does the ability of one or more social groups to privately provide group-specific club goods lead to a general decline in public good provision? Are public goods perfect substitutes for club goods or are there externalities from having multiple privately provided public goods whose beneficiaries are different subsets of the population?

In this paper, we study the effect of allowing subjects to contribute toward a group-specific, excludable public good in addition to a pure public good. We wish to understand whether or not total contributions change as a function of the presence of the club good, and to what extent these changes are driven by social identity preferences. Does the membership of a social group affect the way individuals contribute to public goods? Furthermore, do contributions to the club good work as punishments to noncooperative out-group members, or do they function as rewards to in-group cooperators (or both)?

In the first part of our experiment, subjects self-select into one of two social groups depending on their choices in a task that is completely unrelated to the main experiment. In other words, we use a very weak form of group membership, in the tradition of the minimal group paradigm (Tajfel *et al.* 1971). In the second part of the experiment, subjects play a six-player cooperation game in which three subjects belong to one social group and the other three belong to the other group. In every round of our experiment, subjects receive an endowment. They can allocate any part of their endowment in any of three ways. The first is a public account, the public good, whose returns are divided equally between the six players. The second is a group-specific account, the club good, whose returns are divided equally between the three players from their group. Alternatively, they can keep their endowment to themselves.

We implement two main treatments: one where the marginal per capita return (MPCR) of the public good was equal to that of the club good, and another where the MPCR of the club good was twice that of the public good. In the first treatment, the financial returns to cooperating are higher in the public good: although the MPCRs are the same, there are more people contributing to the public good. In the second treatment, the financial returns to cooperation in both club and public good are the same, as the MPCR of the club good is normalized to the number of contributors. We also run two controls: in one, subjects have access to a public good and a club good, but without any group identities; in the other, we induce group identities, but subjects can only contribute to a single pure public good.

We find that total contributions are higher when a club good is available than when subjects can only contribute to the public good. This increase is significant when the MPCR of the club good is normalized to group size. We find strong evidence of in-group effects, which are manifested through positive in-group reciprocity. First, we find positive contribution levels to the club good, even when returns to contributions to the public good financially dominate those to the club good. Second, we find evidence that these positive contributions are significantly correlated with past contributions by other in-group members. In other words, we uncover a hitherto unexplored mechanism of excludable public goods: they are institutions that allow subjects to display their preferences for interaction with their in-group members, and as means to engage in positive in-group reciprocity. Unlike previous evidence from this class of experiments (Blackwell and McKee 2003), we also find evidence that subjects exploit the substitutability between club good and public good contributions. Subjects respond to lower out-group member contributions to the public good by contributing more to their club good. This mechanism is particularly effective when the returns to club good contributions are high.

The remainder of the paper is organized as follows. Section 2 places our paper in the context of the existing literature. Section 3 describes the experimental design and outlines the hypotheses underpinning the experiment. Section 4 presents the results, and Section 5 concludes the paper.

2. Literature Review

There has been much research in experimental economics on privately provided public goods over the past 30 years, but experimental studies on club goods are rare. The focus of most experiments in this area has been on the role of exclusion or ostracism as a sanctioning device. These papers test its effectiveness on mitigating free-riding, rather than on the effect of multiple public goods, some of which may be excludable.

Cinyabuguma, Page, and Putterman (2005) study a 16-player linear public good, in which after each round players can indicate, at a cost of a fraction of their endowment, any number of other subjects to be expelled from the group. Any subject who receives more than 50% of expulsion "votes" is moved to a separate group in the following round, and plays a public good game with members of that group. The public goods game played by expelled subjects has a smaller individual endowment than the main public good game and there is no chance of further expulsion. Maier-Rigaud, Martinsson, and Staffer (2010) report data from a six-player linear public good game. In their experiment, at the end of each round, subjects can vote on one person to be excluded from the group. Any player receiving more than 50% of votes is expelled. Expelled subjects receive a fixed payment for the remainder of the experiment. In both papers, low contributors are more likely to be expelled from the initial group, and contribution levels by nonexpelled subjects are higher than the expelled subjects. Croson *et al.* (2015) study the effect on the behavior of enforcing a rule, stating that the lowest contributor in a group cannot receive any payoff resulting from the public good. They apply this rule to linear public goods, minimum-effort and maximum-effort games. The ability to exclude leads to higher cooperation levels.[2]

Our analysis differs from the previous studies of excludability in public good contribution games in two ways. We introduce excludability by generating two groups and giving subjects the option to contribute either to a nonexcludable public good or to a group-specific good. As such, we are able to study both the potential of exclusion as a response to free riders, as well as the role group-specific club goods play in heterogeneous societies.

[2] Swope (2002) introduces exclusion by enforcing an exogenous rule in which subjects who contribute less than a minimal amount in a given round do not get a return from the public good in that round. Kocher, Sutter, and Waldner (2005) study a public good game where an outsider gets a return from the public good only if he contributes at least as much as a threshold amount. In both papers, having a threshold contribution level leads to significantly higher contributions.

In this sense, our paper is closely related to Blackwell and McKee (2003), Cherry and Dickinson (2008), and Buchan *et al.* (2009, 2011). Blackwell and McKee (2003) study the impact of introducing three group-specific, excludable public goods in a group of 12 people. The authors report on data from four sessions: in each session, 12 subjects were randomly allocated to one of three groups, identified by a color. In each of 10 rounds, participants had to decide how much of their endowment to allocate to a 12-player public good, or to their four-player club good. In all four sessions, the MPCR to the club good was equal to 0.3, but the MPCR to the global public good was different for each session, taking the values of 0.1, 0.15, 0.2, and 0.3. Total contributions to the global public good increase as the global public good MPCR increases. In particular, the authors find that financial motives drive the choice by subjects on which type of good to contribute. As soon as the global good's return to cooperation is as good as that of the club good, subjects contribute more of their endowment to the global good. Furthermore, the authors find no trade-off between club and public goods. Instead, they find contributions decline as experiment progresses, but the ratio of club to global public good contributions stays constant. The authors argue that this is due to the large number of groups within each session.

Cherry and Dickinson (2008) study the effect of allowing subjects to contribute to multiple, nonexcludable, public goods. The authors compare the canonical, one public good case to treatments where subjects may contribute to three public goods. In one treatment, all public goods have the same MPCR; in the other treatment, all three public goods have different MPCRs. They find that total contributions are higher when subjects have multiple public goods at their disposal.

Buchan *et al.* (2009) report on data from a nested public good experiment conducted in multiple locations across the world. Groups of 12 subjects from three different countries could contribute any part of their endowment to one of two public goods: a global public good with an MPCR of 0.25, shared by all 12 subjects; or to a local public good with an MPCR of 0.5, shared only by players in the same country. The authors find that more globalized individuals (or societies) are, the more likely they are to contribute to the global public good. In an extension of this work, Buchan *et al.* (2011) find that individuals who identify more closely with the global identity contribute more toward the global public good, independently of any expectations about others' behavior.

Our experiment complements the findings from Blackwell and McKee (2003) and Buchan *et al.* (2009, 2011) in important ways. We explicitly wish to investigate the role of social identity preferences, while Blackwell and McKee focused on the role of the ratio of club good MPCR to public good MPCR. To this extent, our design makes group choice endogenous, as opposed to random. It also makes group membership more salient by establishing only two groups, as opposed to three groups in Blackwell and McKee's experimental design that could arguably diffuse the out-group bias that is typically prevalent in behavior (e.g., Chen and Li 2009). We are particularly interested in the role in-group preferences play in determining contributions to the local, or club good; and whether these come in addition to, or in place of, contributions to a global or nonexcludable public good. Buchan *et al.* (2009, 2011) study the complementary question, which pertains to the ability of an overarching, global identity to drive contributions to a global good as opposed to a local one.[3]

[3] Our experiment is also related to Chen *et al.* (2014) who look at the impact of multiple identities on cooperation. The authors compare the effect of assimilating identities (university affiliation) versus fragmenting ones (ethnicity).

3. Experimental Design and Hypotheses

In this paper, we induce group identity in participants artificially. Participants choose identity through their selection of their favored artists, following the minimal group paradigm of Tajfel *et al.* (1971). The task of choosing paintings in order to belong to a group is completely unrelated to the main focus of the experiment. Inducing an artificial identity allows us to cleanly study the effect of the identity we induce, rather than other multiple identities a participant may find relevant. We are also able to abstract away from any history of interaction. By combining an artificial identity with strict anonymity, we ensure that this is the only salient factor that influences choices. We can then study the effect of identity while teasing out repeated interaction effects. While studying the effect of particular types of identity such as gender, religion or ethnicity is very important, we feel that working with a generic identity fits the purpose of this study best.

Our experimental design encompasses three stages. Stage 1 assigns participants to two different groups by eliciting their preferences over two artists' paintings. Stage 2 is a problem-solving task designed to reinforce participants' sense of affiliation to their group. Stage 3 is the actual contribution game. We elaborate on each stage below.

Stage 1: *Group Formation and Assignment.* Before each session, participants were shown five pairs of paintings; in each pair, one painting was done by Gustav Klee and the other by Wassily Kandinsky.[4] Each participant stated their preference for one of the paintings in each pair. If participants preferred three or more Klee paintings, they were assigned to the Klee group; otherwise, they were assigned to the Kandinsky group. This meant that our groups were endogenously determined and also we could not guarantee how big each group (Klee or Kandinsky) would be in each session. The variation in group size across sessions was quite small.[5]

Stage 2: *Identity Reinforcement.* After the Klee and Kandinsky groups were formed, subjects participated in the following activity: they were given two paintings and they had to identify their authorship, one of which was painted by Klee and the other by Kandinsky. Participants were allowed to communicate through a chat box for 10 minutes with their fellow group members. Members of a given group (i.e., Klee or Kandinsky) could only see their own fellow group members' comments. Participants received an individual payment for each painting they correctly identified, and no payment if they could not identify a painting's author. This stage was introduced to reinforce the participants' sense of group identity.

Stage 3: *Contribution Game.* Subjects were randomly allocated to groups of six. The composition of the group was either homogeneous (i.e., all six subjects were Klee types or Kandinsky types), or heterogeneous (i.e., three subjects were Klee types and the other three were Kandinsky types). Subjects knew the composition of their own group, but they were not aware of the composition of the other groups in the session.[6]

[4] We chose the artists for no other reason than to be consistent with the existing literature on minimal groups.

[5] Due to this, the number of observations which we have collected does not match in all treatments. This does not affect our statistical analysis.

[6] The experimental protocol we used to elicit groups means that it was possible that in an 18-person session the distribution of artist preferences meant we could not allocate some subjects to groups that

All groups interacted over 20 rounds with fixed matching. Subjects had 20 tokens that they could allocate to a public good account, a group-specific club good account, or to a private account. Payoffs were determined by the following equation, where x_i^p denotes player i's contribution to the public good, and x_i^c denotes player i's contribution to the club good:

$$\pi_i := 20 - x_i^p - x_i^c + a\sum_{j=1}^{6} x_j^p + b\sum_{j=1}^{3} x_j^c, \tag{1}$$

where y is the initial endowment, $0 < a < 1 < 6a$ is the MPCR from the public good, and $0 < b < 1 < 3b$ is MPCR from the club good. There are two distinct subgroups in the six-player group, each with three members.

In all but one treatment, participants could choose between contributing to a nonexcludable public good whose beneficiaries were all six players, or to a club good whose only beneficiaries were their fellow in-group members. Players could also free-ride by not contributing to either good. In other words, our subjects had the opportunity to discriminate against out-group members via exclusion.

Having public and club goods with different group sizes introduces a further complication. If both goods have the same MPCR, then the financial return to cooperation will be higher for the good with the largest number of contributors. Therefore, if both public good and club good have the same MPCR, the public good will be financially more attractive. In this case, subjects face a trade-off between the higher financial reward from contributing to the public good and the higher identity utility from contributing to a club good. The MPCR of the club good, b, is our main treatment variable. In *Club-L*, the MPCR of both public and club goods is equal to 0.4. In *Club-H*, the MPCR of the club good is equal to 0.8, while the MPRC of the public good is equal to 0.4.

We implemented two control treatments. The first, which we denote by *Club-HN*, is identical to *Club-H*, but there is no identity manipulation stage. Subjects were told there were two accounts to which they could allocate their endowment, one of which was available to all six persons in their group and another which was available to themselves and two randomly picked subjects from the six. This treatment is designed to control for the possibility that group size is the main driver for higher contributions in the club good. The second control treatment is denoted by *Ctrl*, where we implemented Klee and Kandinsky identities, and where there were three subjects of each group, but where subjects from either group can only contribute to a single public good shared by all six subjects. At the end of each round, a screen informed subjects of the individual contributions by each of the other five subjects in their group, as well as his identity (Klee or Kandinsky). To mitigate individual reputation effects, the software randomized the display order of individual contributions from round to round, and this was common knowledge. Table 1 outlines the different treatments and number of groups in each treatment.

Each session consisted of 18 participants and lasted between 60 and 90 minutes. At the end of the experiment, subjects were paid individually in cash. Sessions took place in the Fall of 2010. The experiment was conducted using z-Tree (Fischbacher 2007). A total of 144 undergraduate students recruited via ORSEE (Greiner 2004) participated

were either homogeneous or equally divided between Klee and Kandinsky types. When that happened, we allocated those subjects to a companion experiment that shared stages 1 and 2 in its experimental protocol.

Table 1: Experimental design

Treatment	***Ctrl***	***Club-L***	***Club-H***	***Club-HN***
Public Good MPCR (a)	0.4	0.4	0.4	0.4
Club Good MPCR (b)	–	0.4	0.8	0.8
Identity Priming?	Yes	Yes	Yes	No
# of groups	8	5	5	6

in the experiment, who received on average 8.27. The instruction set is available in Appendix A.

We now introduce and discuss our hypotheses.[7] Our hypotheses are predicated on evidence that individuals exhibit a preference for other individuals who share a common social identity in social and economic interactions (Allport 1954). More recently, it has been shown that individuals exhibit higher concerns for the financial welfare of those from their own group in dictator games (Chen and Li 2009).

We start with the effect on the behavior of allowing subjects to contribute to multiple public goods. Cherry and Dickinson (2008) show that increasing the number of public goods available to subjects leads to an increase in total contributions, even if the public goods are perfect substitutes. We therefore hypothesize that adding a (financially dominated) group-specific club good to the canonical public good game will have a similar effect.

HYPOTHESIS 1 (Multiple Public Goods): *Total contributions will be higher in the Club-L treatment than in the Ctrl treatment.*

Before considering the effect social identity may have on contribution levels to club goods, it is important to consider that excludable public goods may be more attractive than pure public goods simply because the number of players contributing to (and benefitting from) them is by definition smaller. As such, the degree of strategic uncertainty inherent to the club good is also smaller. An important literature in experimental economics has already established that cooperation in the public goods game is easier in small groups (Isaac and Walker 1988; Isaac, Walker, and Williams 1994). We therefore hypothesize that if a club good is available and the returns to that club good are normalized to group size, then contributions to the club good will be larger than to the public good. Also, total contributions will be higher than the case where there is no club good.

HYPOTHESIS 2 (Group Size): *In the Club-H treatment, contributions to the club good will be exceeding those to the public good.*

We now turn to the main research question of the paper: the extent to which contributions to the public good can be explained by identity preferences, above and beyond pure group size or multiple goods effects. To achieve this, we take the case where both public good and club goods are available, and their financial returns are normalized to group size. We then compare a treatment where group identities are elicited and made salient to a treatment where they are absent. We hypothesize that club good

[7] We will outline our alternative hypotheses; our null hypothesis is that behavior will be unaffected by the presence of multiple public goods and/or identity, as the dominant strategy Nash equilibrium of all the games under consideration is zero contributions if players are self-interested.

Table 2: Average contributions to public and club good and total earnings

Treatment	***Ctrl***		***Club-L***		***Club-H***		***Club-HN***	
Round	**One**	**All**	**One**	**All**	**One**	**All**	**One**	**All**
PG contributions	8.79	7.65	7.30	5.30	5.67	2.40	8.00	4.32
	(6.04)	(1.18)	(4.40)	(2.16)	(4.10)	(1.04)	(5.41)	(2.64)
CG contributions	–	–	6.20	2.85	7.97	9.09	6.83	6.43
			(3.28)	(0.93)	(4.44)	(3.31)	(5.01)	(2.32)
PG+CG contributions	8.79	7.65	13.50	8.14	13.63	11.49	14.83	10.75
	(6.04)	(1.18)	(5.20)	(2.62)	(5.29)	(2.32)	(5.73)	(2.67)
Total earnings (ECU)	32.31	30.71	31.46	27.98	39.09	36.08	40.77	35.05
	(6.31)	(1.65)	(5.77)	(3.09)	(6.30)	(3.25)	(7.80)	(3.74)
N	48	8	30	5	30	5	36	6

Note: Standard deviations in parentheses. ECU = experimental currency units.

contributions in the identity treatment will be higher than in the treatment where identity is absent.

HYPOTHESIS 3 (Identity): *Club good contributions will be higher in the Club-HN than in Club-HN.*

We complete our set of hypotheses by discussing the trade-off between in-group bias and efficiency. In addition to the opportunity cost of retaining each unit of endowment, which yields a marginal return of unity, subjects in our experiment face an additional dilemma. They can invest their endowment in a group-specific public good, from which they derive some social identity utility, or they can invest in a nonexcludable public good, which is at least as efficient. To understand the nature of this trade-off, we compare behavior in the treatment where both goods have the same MPCR to the treatment where the club good's MPCR is normalized to group size. In the former case, subjects face a salient trade-off between efficiency and in-group love; in the latter, that trade-off is absent. We therefore predict higher club good contributions in the second treatment.

HYPOTHESIS 4 (Identity/Efficiency Trade-Off): *Club good contributions will be higher in the Club-H treatment than in the Club-L treatment.*

4. Results

We begin the analysis of results by looking at the effect of our experimental treatments on contributions to the group-specific good, as well as the to pure public good. Table 2 presents the average contributions in the first round, as well as over the 20 rounds.

Before we proceed, a few remarks on the data analysis are warranted. First, given the relatively small number of independent observations (i.e., groups), we opted to evaluate treatment effects using individual decisions by subjects in a round as the unit of observation. Since the choice variable is bounded below by zero and above by 20, we employ the random-effects Tobit estimator.[8] To account for the fact that we have 20 observations per subject, which may be correlated with each other, we bootstrap standard errors.

[8] Indeed, 26% (30%) of all observations record a contribution level of zero to the club (public) good, and 8% (5%) record a contribution level of 20 to the club (public) good.

Table 3: Random effects Tobit estimates of contribution levels to public and club goods

Dep Var	(1) $x^c_{i,t}$	(2) $x^p_{i,t}$	(3) $x^c_{i,t}+x^p_{i,t}$	(4) $x^c_{i,t}$	(5) $x^p_{i,t}$	(6) $x^c_{i,t}+x^p_{i,t}$	(7) $x^c_{i,t}$	(8) $x^p_{i,t}$	(9) $x^c_{i,t}+x^p_{i,t}$
Club-L	−7.96***	5.33***	−4.22***	−6.72***	3.10***	−5.26***	−3.73***	1.00	-4.01***
	(1.36)	(1.05)	(1.45)	(0.92)	(0.66)	(1.23)	(0.87)	(0.68)	(1.04)
Club-HN	−3.03**	3.41***	−0.04	−6.15***	3.19***	−2.87***	−3.34***	1.16	-1.87*
	(1.44)	(1.03)	(2.02)	(0.82)	(0.69)	(1.04)	(0.96)	(0.79)	(1.07)
Ctrl			−5.45***			−7.40***			
			(1.48)			(1.01)			
OG Public$_{t-1}$							−0.08***	0.09***	0.001
							(0.02)	(0.02)	(0.03)
IG Public$_{t-1}$							0.01	0.04***	0.05***
							(0.01)	(0.01)	(0.02)
IG Club$_{t-1}$							0.26***	−0.11***	0.20***
							(0.03)	(0.02)	(0.03)
Round				−0.29***	−0.47***	−0.60***	−0.27***	−0.40***	-0.51***
				(0.06)	(0.05)	(0.07)	(0.06)	(0.07)	(0.06)
Round× *Club-L*				−0.16**	0.24***	0.08	−0.07	0.20***	0.14*
				(0.07)	(0.06)	(0.09)	(0.07)	(0.07)	(0.07)
Round× *Club-HN*				0.28***	0.03	0.24***	0.20***	0.09	0.21***
				(0.07)	(0.06)	(0.09)	(0.07)	(0.07)	(0.08)
Round× *Ctrl*						0.19**			
						(0.08)			
$(t=1)$				−3.86**	1.46	−2.52			
				(1.73)	(1.12)	(1.99)			
$(t=1)\times$ *Club-L*				5.17**	−1.30	4.49*			
				(2.11)	(1.56)	(2.58)			
$(t=1)\times$ *Club-HN*				4.70**	−0.57	5.54**			
				(2.00)	(1.62)	(2.79)			
$(t=1)\times$ *Ctrl*						0.45			
						(2.56)			
Constant	8.78***	−0.50	12.30***	11.94***	4.17***	18.71***	7.42***	4.22***	13.47***
	(1.29)	(0.80)	(1.35)	(0.64)	(0.59)	(0.85)	(0.97)	(0.84)	(0.97)
N	1920	1920	2880	1920	1920	2880	1824	1824	1824
ρ	0.40	0.29	0.39	0.42	0.34	0.43	0.41	0.32	0.51
	(0.05)	(0.03)	(0.04)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)

Note: Bootstrapped standard deviations in parentheses. ***, **, and *: significance at 1%, 5%, and 10% levels, respectively.

Table 3 presents the results from the estimations. Regression (1) has as its dependent variable the contribution to the group-specific club good by subject i in round t ($x^c_{i,t}$). Its regressors are a set of treatment dummies—the omitted category is the *Club-H* treatment. The intercept coefficient is positive (8.78) and statistically significant: in the *Club-H* treatment subjects allocate a large part of their endowment to the club good. The coefficient on *Club-L* is negative (−7.96) and significant, indicating lower average contributions to the club good in this treatment than in *Club-H*; these contributions are nevertheless positive (*Constant* + *Club-L*=0, $\chi^2(1) = 9.42$, $p < 0.01$). The coefficient on

Club-HN is also negative and significant; contributions to the club good in this treatment are also positive (*Constant* + *Club-HN*=0, $\chi^2(1) = 456.43$, $p < 0.01$).

Regression (2) performs the same analysis using contributions to the public good by subject i in round t, $x^p_{i,t}$, as the dependent variable. The constant term is negative and nonsignificant, driven by the fact that a very large proportion of observations are at the lower censoring limit. The coefficients on *Club-L* and *Club-HN* are both positive and significant, indicating average contributions to the public good in either treatment are higher than in the *Club-H* treatment.

Regression (3) analyzes the treatment effects using the sum of public and club good contributions, $x^c_{i,t}+x^p_{i,t}$, as the dependent variable. Unlike the previous regressions, we now include a dummy variable for the treatment without club goods, *Ctrl*, to allow for the analysis of the effect of the presence of club goods on total contributions. The coefficients on all treatment variables are negative and significant, except the coefficient on *Club-HN*, which is not significant. The coefficient on *Ctrl* is significantly different to that on *Club-L* ($\chi^2(1) = 8.00$, $p < 0.01$). This indicates total contributions are higher when two public goods are available than in the canonical one-public good case.

The results from the summary statistics in Table 2, in addition to the regressions (1)–(3) in Table 3, indicate that total contributions in the *Club-L* treatment are higher than in the *Ctrl* treatment. That is to say, we uncover a multiple public goods effect on contribution akin to that reported by Cherry and Dickinson (2008). This leads to our first finding.

Finding 1: Enabling subjects to contribute to a club good in addition to a public good leads to higher total contribution levels, even when the club good is financially dominated by the public good

Our data also display a group size effect: as shown in Table 2, contributions to the club good are higher than those to the public good in the *Club-H* treatment. This is consistent with the idea that public goods that are provided by smaller groups tend to have higher contribution levels because of lower strategic uncertainty. This is our second finding.

Finding 2: When subjects can contribute to a club good and a public good, and returns to both goods are normalized to group size, contributions to the former significantly exceed those to the latter.

Our data also allow a clear answer to the main research question of the paper. We observe significant differences in behavior between the *Club-H* and *Club-HN* conditions. While regression (3) in Table 3 shows that total contributions are approximately the same in both treatments, regression (1) indicates that average contributions to the club good are higher in the former than in the latter, while the reverse is true for contributions to the public good in regression (2). This is our third finding.

Finding 3: Contributions to the club good can be attributed to social identity preferences.

We can also address the issue of the trade-off between social identity utility and efficiency. We consider two treatments, one (*Club-H*) where the trade-off does not exist, since the MPCR of the public good equals that of the club good, and another where the trade-off is present, since the club good is financially dominated by the public good (*Club-L*). We find evidence for this trade-off in our data: regression (1) shows that club good contributions are significantly lower when the returns to the club good are lower than those to the public good; however, there are still positive contribution levels in the *Club-L* treatment, indicating that financial considerations do not dominate social identity preferences.

Finding 4: We find evidence of a trade-off between social identity utility and efficiency, in that contributions to the club good are affected by how efficient it is relative to the public good.

The experimental research on linear public goods games has consistently found a downward trend in contributions over the course of an experiment (Ledyard 1997). The top panel in Figure 1 outlines average total contributions over the course of the experiment; this negative time trend is also evident in our data. Furthermore, there is also evidence that the effect of identity seems to decline with repeated interaction when employing minimal groups (Eckel and Grossman 2005; Chakravarty and Fonseca 2014). We take these two effects into account in our analysis in two ways. First, we run a second set of regressions where we include a time trend, *Round*, as well as a dummy for round one, which we interact with treatment dummies. The rationale for the latter is a difference in round one contributions between the *Ctrl* treatment and the other three treatments.[9]

Introducing treatment-specific time trends and round one dummies did not change the sign or significance level for the treatment dummies in regression (4) relative to regression (1). We observe a negative time trend in the *Club-H* treatment (−0.29). This trend is significantly more pronounced in the *Club-L* treatment, but almost absent in the *Club-HN* treatment—which is readily observed in the bottom right panel in Figure 1. We also observe strong round one effects in the data: on the one hand, a negative effect in the *Club-H* condition, which is due to an actual increase in club good contributions after round one in that treatment (as per the bottom panels in Figure 1). On the other hand, we observe positive coefficients in the interaction dummy coefficients for the *Club-L* and *Club-HN* treatments, indicating a decline in contributions immediately after round one. In contrast, the results from regression (5) indicate a smaller coefficient on *Club-L* than in regression (2); however, that is accompanied by a weaker time trend relative to the omitted condition. We find no round one effects in the club good contributions. Finally, regression (6) displays qualitatively similar dummy coefficients effects as in regression (3)—although the coefficient on *Club-HN* is now significant. It is worth noting that, although there are no significant differences in round one behavior between *Ctrl* and *Club-H*, we estimate a weaker time trend in the former treatment than in the latter.

Finding 5: There is a trend toward lower contributions as the experiment progresses.

We conclude our analysis of contribution levels by looking at group dynamics. We do so by augmenting regressions (1–3) with treatment-specific time trends, as well as three variables: total contributions by out-group members toward the public good in the previous round, $OG\ Public_{t-1}$; total contributions by other in-group members to the public good in the previous round, $IG\ Public_{t-1}$; and total contributions by other in-group members to the club good in the previous round, $IG\ Club_{t-1}$.[10] These variables are designed to capture the extent to which subjects in our experiment respond to

[9] Using group average contributions as the unit of observation, there are significant differences in average contributions in round one between *Ctrl* and *Club-L* ($z = 2.642$, $p < 0.01$); *Club-H* ($z = 2.642$, $p < 0.01$); and *Club-HN* ($z = 2.843$, $p < 0.01$).

[10] Subjects could not observe contributions to the other group's club good, so we could not generate the variable $OG\ Club_{t-1}$.

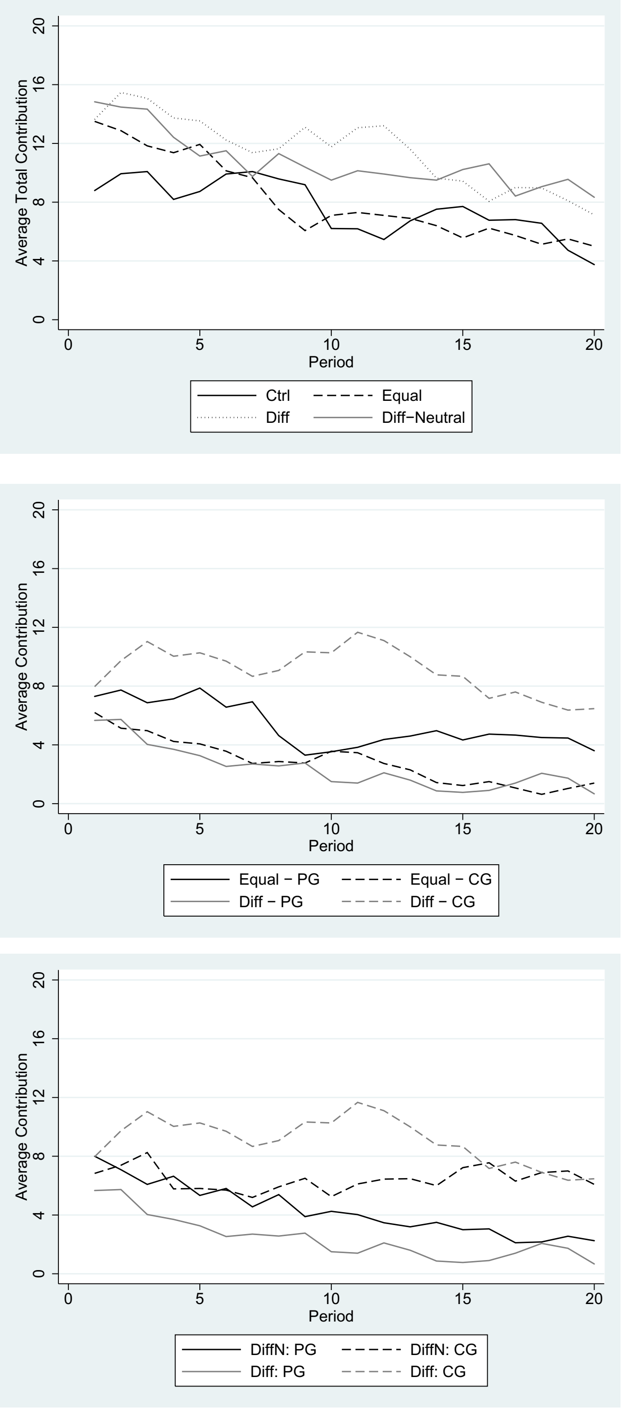

Figure 1: Time series of contributions. Top panel: Average total contributions across all treatments. Bottom left panel: Average club and public good contributions in *Club-L* and Diff treatments. Bottom right panel: Average club and public good contributions in Diff and DiffN treatments.

the end-of-round information about contribution levels. Standard reciprocity models would predict that subjects may wish to reward good behavior in the previous round with higher contributions and bad behavior with low contributions.[11] This reciprocity conjecture predicts a negative correlation between past contributions by other players to the public good and current contributions to the club good. It also predicts a positive correlation between past contributions by other players to the public good and current contributions to the public good. Both correlations should be stronger for in-group members than for out-group members, given that social identity theory predicts subjects that attach a higher concern for the welfare and actions of their fellow in-group members than for out-group members. We do not present results where these three variables are interacted with treatment dummies, as we did not find systematic differences. These results are available in Appendix B.

Regression (7) in Table 3 outlines the results for the augmented regression using $x^c_{i,t}$ as a regress. We find a negative and significant coefficient on $OG\ Public_{t-1}$, a coefficient of zero on $IG\ Public_{t-1}$ and a positive and significant coefficient on $IG\ Club_{t-1}$—this coefficient is significantly larger than that on $OG\ Public_{t-1}$ ($\chi^2(1) = 77.90, p < 0.01$). In other words, a unit increase in total public good contributions by out-group members in $t-1$ leads to a reduction in club good contributions, while the same behavior by in-group members has no effect on club good contributions. However, an increase in club contributions by in-group members in $t-1$ leads to a larger increase in club contributions in t; the last effect dominates.

Regression (8) outlines the results for the regression using $x^p_{i,t}$ as the dependent variable. The coefficient on $OG\ Public_{t-1}$ is positive and significant, as is the coefficient on $IG\ Public_{t-1}$. That is, increases in contribution levels by other players in round $t-1$ lead to higher contribution levels in round t; the effect is nominally stronger for out-group players, but not statistically significantly so ($\chi^2(1) = 2.11, p = 0.147$). We also find a negative and significant coefficient on $IG\ Club_{t-1}$: the more in-group members contribute to the club good in the previous round, the lower contributions are in the present round.

Regression (9) considers the net effect of these two regressions. We find a straightforward substitution effect when considering past behavior by out-group members: higher contributions to the public good by out-group members in $t-1$ lead to higher public good contributions by a given player at the expense of club good contributions. However, higher contributions to the public good by in-group members in the previous round lead to higher total contribution levels in the current round without detriment to club good contributions. The effect of higher club good contributions by in-group members also has a positive net effect, although that comes at the detriment of public good contributions, suggesting that in-group biases play an important role in behavior.

Finding 6: Subjects respond to lower contributions by others to the public good in the previous round by contributing more to the club good in the present round. This effect is stronger with regard to out-group members's behavior.

We conclude our analysis of our data by looking at efficiency. Table 2 outlines average per round payoffs in experimental currency units (ECU) as a function of treatment. The fact that total contributions are higher when both club and public goods are

[11] See Falk and Fischbacher (2006) for a model of reciprocity applied to a public goods game.

available does not necessarily lead to higher efficiency. In the *Club-L* treatment, the small increase in total contributions does not compensate for the fact that some of those contributions are toward a club good whose returns are lower than those to the public good. In this case, the existence of a club good led to a decrease in efficiency, although that difference is not statistically significant ($z = -1.610, p = 0.107$, Mann Whitney [MW] test using mean earnings at group level). However, when both club and public good have equivalent MPCRs, there are noticeable gains in efficiency relative to the treatment with only one public good. This is the case both when group identities are absent (*Ctrl* = *Club-HN*: $z = 2.195, p = 0.028$, MW test), and more so when group identities are induced (*Ctrl* = *Club-H*: $z = 2.489, p = 0.013$, MW test). We find significantly higher efficiency levels in *Club-H* than in *Club-L* ($z = 2.402, p = 0.016$, MW test). We do not find a significant difference between the earnings in the *Club-H* and *Club-HN* treatments ($z = 0.730, p = 0.465$). This constitutes our final finding.

Finding 7: The existence of a group-specific club good leads to efficiency gains only when the financial returns to the club good are equivalent to the returns to the public good.

5. Conclusion

Societies and organizations are often characterized by having individuals belonging to, or identifying with one or more groups. Individuals in such societies participate in nested social dilemmas. They have the opportunity to either contribute to public goods that are enjoyed by every individual in the society, or to club goods that are only enjoyed by their fellow in-group members. Examples of this problem include that of contributions to faith-based schools and or state/public schools. Given that public goods and club goods are often close substitutes that lead us to ask under what conditions will individuals contribute to one type of good and/or to the other, and what are the consequences to overall levels of cooperation, a clear understanding of these questions can help in designing better or more cooperation-inducing institutions.

In this paper, we look at social identity as one possible explanation for individual behavior when facing a choice of contribution to a club or a public good. To this effect, we implemented a nested public good game experiment, in which subjects could either contribute to a pure public good, shared by all subjects, or to a group-specific club good. The groups are endogenously formed during the experiment and this formation is the basis of the social identity of the individuals. There are a number of factors that can determine contributions to club and/or public good. These include financial returns, strategic risk driven by the number of players, the presence of more opportunities to contribute, and finally, bias toward one's own social group.

We observe positive contribution levels to the club good, even when the pure public good is financially more attractive. This is not due to an increase in total cooperation or increased opportunity to cooperate. Subjects actually substitute away from the pure public good into the club good. When the pure public good and the club good have equivalent marginal per capita rates of return, we observe significantly more contributions to the club good than to the public good. The increase in club good contribution away from public goods cannot be explained completely by strategic risk. While individuals also shift their contributions from public to club goods in the treatment without social identity, they do so to a lesser extent.

We also find that while the presence of club goods may lead to increased contributions, it does not necessarily lead to increases in efficiency. Earnings in the treatment where the public good financially dominated the club good were actually lower than those in the treatment which offered only a pure public good. We only observed efficiency gains in the treatments where the MPCR to club good contributions was equivalent to that of the public good. In those treatments, we did observe large efficiency gains.

The main finding of our paper is that we can attribute a significant proportion of those club good contributions to social identity preferences. We also uncovered evidence that the club good also plays a strategic role: subjects respond to low levels of cooperation by out-group members in previous rounds in the experiment by switching contributions away from the public good and into the club good. Our results suggest that club goods, when financially effective, can work not only as an expression of a sense of group identity, but also as a mechanism that enforces cooperation at the population level. Club goods allow subjects to sustain cooperation levels at a much higher level than those achieved in their absence.

Appendix A: Experimental Instructions

A.1. *Club-H* Treatment

A.1.1. Instruction set

Welcome to our experiment. Please remain silent during the course of the experiment. If you have any questions, please raise your hand. You will now take part in a decision-making experiment. The amount you will receive for participating will depend on your decisions and the decisions of other participants. There will be two parts to this experiment. Before each part of the experiment begins, you will receive a set of instructions explaining the details of that particular part.

Once you complete all the decisions in a given part, we will move to the next part of the experiment. You will only receive information about the outcome of your choices at the end of the experiment. To keep track of your choices, we will provide you with a decision form. Your payoff in this experiment will be equal to the sum of payoffs in each of the individual parts. The payoffs throughout the experiment will be denominated in experimental currency units (ECU); 1 ECU is worth 12 pence. Once the experiment ends, your payoff will be calculated and you will receive your payment in cash.

A.1.2. Part 1

In this part, we will show you five pairs of paintings by two artists. For each pair of paintings, you must choose the one you prefer. Once everyone makes their five choices, we will divide participants into two groups according to which artist they preferred.

Once you have been allocated to one of the groups, we will show you a further two paintings. Your task will be to identify which artist painted which painting. You will be allowed to confer with your fellow group members in order to determine the answer to the two questions. To this effect, you will have access to a chat programme, through which you can offer help or get help from your fellow group members.

Messages you post in the chat box will only be visible to members of your own group. You will not be able to see the messages posted by members of the other group and vice-versa. You will be able to communicate with your fellow group members for

10 minutes before submitting your answers. You are free to post as many messages as you like. There are only two restrictions on messages: you may not post messages that identify yourself (e.g., age, gender, location, etc.) and you may not use offensive language. For each correct answer, you will earn 10 ECU. Once everyone submits their answers, the experiment will move to the second part. You will only be informed of your payoff in this part of the experiment at the very end of the session.

A.1.3. Part 2 (only seen after the end of Part 1)

In this part of the experiment, you will be matched with five other participants. You will be interacting with the same five participants until the end of the experiment. There will be 20 rounds in this part of the experiment. At the beginning of each round, each participant will receive 20 ECUs. We will call this your endowment. Your task in each round is to decide how to use your endowment.

You must decide how many ECUs you want to contribute to a public project, which is available to all six participants, how much to contribute to a project that is only available to people of your own group (Klee only if you are a member of the Klee group; Kandinsky only if you are a member of the Kandinsky group) and how many ECU you want to keep for yourself. The consequences of your decision are explained in detail below. Your payoff is given by the following formula:

Your Payoff = 20 ECU + (0.4*Total Contribution to the Public Project − Your Contribution to the Public Project) + (0.8*Total Contribution to the Group Project − Your Contribution to the Group Project).

This formula implies that your payoff in every round is based on three parts:

1. The ECUs you kept for yourself: (20 ECU − Your contribution).
2. The income from the public project, which is 40% of the total contribution from you and from the other five participants.
3. The income from the group project, which is 80% of the total contribution from you and from the other two members of your group (Klee/Kandinsky only).

To fix ideas, let us consider a few numerical examples.

Suppose that each of the six participants contributes 10 ECU to the public project and 0 ECU to their own group's project. This means the total contribution to the public project is 60 ECU and the total contribution to your group project is 0 ECU. In this case, the income from the public project is 0.4*60 = 24 ECU to each of the six participants, while the income from your group's project (only payable to your group's members) is 0 ECU.

Suppose, instead, that each of the six participants contributes 0 ECU to the public project and 10 ECU to their own group's project. This means the total contribution to the public project is 0 ECU and the total contribution to your group project is 30 ECU. In this case, the income from the public project is 0 ECU to each of the six participants, while the income from your group's project (only payable to your group's members) is 0.8*30 = 24 ECU.

Each ECU you keep to yourself raises your payoff by 1 ECU. Each ECU you contribute to the public project raises the total contribution to the project by 1 ECU and causes your income from the public project to rise by 0.4*1 = 0.4 ECU. The income of the other five participants will also rise by 0.4 ECU, so that the total income of the six participants from the public project will go up by 2.4 ECU. Your contribution to the project therefore also raises the income of everyone else.

Each ECU you contribute to the group project raises the total contribution to the group project by 1 ECU and causes your income from the project to rise by 0.8*1 = 0.8 ECU. The income of your other group members will also rise by 0.8 ECU, so that the total income of the three group members from the group project will go up by 2.4 ECU. Your contribution to the group project therefore also raises the income of your fellow group members, but not the income of nongroup members.

In other words, 1 ECU invested in the public project yields a total benefit of 2.4 ECU to the six participants. 1 ECU invested in your group project yields a total benefit of 2.4 ECU to the three members of your group only.

Remember that ECUs earned in one round do NOT carry over to subsequent rounds. You will start every round with the same endowment of 20 ECUs.

Once all participants have made their decisions, you will be informed about your decision, the decision of each participant, the total amount of ECUs contributed to the public and group projects, and your payoff. You will also know whether each person with whom you are playing belongs to either the Kandinsky or the Klee group, but not their exact identity. To this effect, the computer will scramble the order in which the other participants are listed when individual contributions are shown at the end of every round. Once the 20th round is over, the experiment will be over. The computer will select two rounds at random. Your payoff in those two rounds plus the payoff from part 1 will determine your total earnings in the session.

A.2. *Club-HN* Treatment

Welcome to our experiment. Please remain silent during the course of the experiment. If you have any questions, please raise your hand.

You will now take part in a decision-making experiment. The amount you will receive for participating will depend on your decisions and the decisions of other participants.

The payoffs throughout the experiment will be denominated in experimental currency units (ECU); 1 ECU is worth 12 pence. Once the experiment ends, your payoff will be calculated and you will receive your payment in cash.

In this experiment, you will be matched with five other participants. You will be interacting with the same five participants until the end of the experiment.

There will be 20 rounds in the experiment. At the beginning of each round, each participant will receive 20 ECUs. We will call this your endowment. Your task in each round is to decide how to use your endowment.

You will have three alternative ways to use your endowment:

- A public project, which is available to all six participants.
- A group project, which is only available to you and two other participants (who remain the same throughout the 20 rounds).
- Keep it to yourself.

There is no right or wrong way to allocate your endowment. Furthermore, you do not have to allocate money to all three alternatives; it is fine not to allocate any amount to one of the alternatives if you do not wish to do so.

The way you allocate your endowment will have financial consequences to you. We explain them in detail below.

Your payoff is given by the following formula:

Your Payoff = 20 ECU + (0.4*Total Contribution to the public project − Your Contribution to the public project) + (0.8*Total Contribution to the group project − Your Contribution to the group project).

This formula implies that your payoff in every round is based on three parts:

1. The ECUs you kept for yourself: (20 ECU − Your contribution to the public project − Your contribution to the group project).
2. The income from the public project, which is 40% of the total contribution from you and from the other five participants.
3. The income from the group project, which is 80% of the total contribution from you and from the other two members who can also contribute to your group project.

To fix ideas, let us consider a few numerical examples.

Suppose that each of the six participants contributes 10 ECU to the public project and 0 ECU to their own group's project. This means the total contribution to the public project is 60 ECU and the total contribution to your group project is 0 ECU. In this case, the income from the public project is 0.4*60 = 24 ECU to each of the six participants, while the income from your group's project (only payable to your group's members) is 0 ECU.

Suppose, instead, that each of the six participants contributes 0 ECU to the public project and 10 ECU to their own group's project. This means the total contribution to the public project is 0 ECU and the total contribution to your group project is 30 ECU. In this case, the income from the public project is 0 ECU to each of the six participants, while the income from your group's project (only payable to your group's members) is 0.8*30 = 24 ECU.

Each ECU you keep to yourself raises your payoff by 1 ECU.

Each ECU you contribute to the public project raises the total contribution to the project by 1 ECU and causes your income from the public project to rise by 0.4*1 = 0.4 ECU. The income of the other five participants will also rise by 0.4 ECU, so that the total income of the six participants from the public project will go up by 2.4 ECU. Your contribution to the project therefore also raises the income of everyone else.

Each ECU you contribute to the group project raises the total contribution to the group project by 1 ECU and causes your income from the project to rise by 0.8*1 = 0.8 ECU. The income of the two other participants who can contribute to the group project will also rise by 0.8 ECU, so that the total income of the three group members from the group project will go up by 2.4 ECU. Your contribution to the group project therefore also raises the income of those who can contribute to your group project but not the income of the other three participants.

In other words, 1 ECU invested in the public project yields a total benefit of 2.4 ECU to the six participants. 1 ECU invested in your group project yields a total benefit of 2.4 ECU to the three members of your group only.

Remember that ECUs earned in one round do NOT carry over to subsequent rounds. You will start every round with the same endowment of 20 ECUs.

Once all participants have made their decisions, you will be informed about your decision, the decision of each participant, the total amount of ECUs contributed to the public and group projects, and your payoff.

You will also not be able to know the exact identity of the participants you are playing with. To this effect, the computer will scramble the order in which the other participants are listed when individual contributions are shown at the end of every round.

Once the 20th round is over, the experiment will be over. The computer will select two rounds at random, which will determine your total earnings in the session.

Appendix B: Additional Econometric Analysis

Table B1: Tobit estimates of contribution levels to public and club goods

Dep Var	(1) $x^c_{i,t}$	(2) $x^p_{i,t}$	(3) $x^c_{i,t}+x^p_{i,t}$
Club-L	-8.78^{***}	0.13	-9.35^{***}
	(1.71)	(1.47)	(2.39)
Club-HN	-7.94^{***}	2.11	-5.56^{**}
	(1.17)	(1.59)	(2.32)
OG Public$_{t-1}$	-0.17^{***}	0.05	−0.11
	(0.07)	(0.06)	(0.08)
IG Public$_{t-1}$	0.002	0.08^{***}	0.06^{*}
	(0.03)	(0.03)	(0.04)
IG Club$_{t-1}$	0.18^{***}	−0.13	0.09^{*}
	(0.05)	(0.04)	(0.05)
OG Public$_{t-1}$ × *Club-L*	0.09	0.08	0.17^{*}
	(0.08)	(0.07)	(0.09)
OG Public$_{t-1}$ × *Club-HN*	0.12^{*}	−0.001	0.09
	(0.07)	(0.07)	(0.10)
IG Public$_{t-1}$ × *Club-L*	−0.002	−0.06	−0.04
	(0.04)	(0.04)	(0.04)
IG Public$_{t-1}$ × *Club-HN*	0.02	−0.05	−0.01
	(0.04)	(0.03)	(0.05)
IG Club$_{t-1}$ × *Club-L*	0.25^{**}	0.06	0.21^{**}
	(0.11)	(0.07)	(0.09)
IG Club$_{t-1}$ × *Club-HN*	0.08	0.03	0.15^{**}
	(0.05)	(0.05)	()
Round	-0.36^{***}	-0.40^{***}	-0.59^{***}
	(0.06)	(0.06)	(0.08)
Round× *Club-L*	0.10	0.23^{***}	0.29^{***}
	(0.08)	(0.07)	(0.11)
Round× *Club-HN*	0.34^{***}	0.04	0.28^{***}
	(0.08)	(0.08)	(0.10)
Constant	10.69^{***}	4.26^{***}	16.90^{***}
	(1.32)	(1.16)	(1.87)
N	1824	1824	1824
ρ	0.41	0.33	0.51
	(0.03)	(0.03)	(0.21)

Notes: Bootstrapped standard deviations in parentheses.
***, **, and *: significance at 1%, 5%, and 10% levels, respectively.

References

AKERLOF, G. A., and R. E. KRANTON (2000) Economics and identity, *Quarterly Journal of Economics* **115**, 715–753.

ALLPORT, G. (1954) *The Nature of Prejudice*. Cambridge, MA: Addison-Wesley.

BLACKWELL, C., and M. MCKEE (2003) Only for my own neighbourhood? Preferences and voluntary provision of local and global public goods, *Journal of Economic Behavior and Organization* **52**, 115–131.

BUCHAN, N. R., M. BREWER, G. GRIMALDA, R. WILSON, E. FATAS, and M. FODDY (2009) Globalization and human cooperation, *Proceedings of The National Academy of Sciences* **106**, 4138–4142.

BUCHAN, N. R., M. BREWER, G. GRIMALDA, R. WILSON, E. FATAS, and M. FODDY (2011) Global social identity and global cooperation, *Psychological Science* **22**, 821–828.

CHAKRAVARTY, S., and M. A. FONSECA (2014) The effect of social fragmentation on public good provision: An experimental study, *Journal of Behavioral and Experimental Economics* **53**, 1–9.

CHEN, Y., and S. X. LI (2009) Group identity and social preferences, *American Economic Review* **99**, 431–457.

CHEN, Y., S. X. LI, T. X. LIU, and M. SHIH (2014) What hat to wear? Impact of natural identities on coordination and cooperation, *Games and Economic Behavior* **84**, 58–86.

CHERRY, T., and D. L. DICKINSON (2008) Voluntary contributions with multiple public goods. In *Environmental Economics: Experimental Methods*, T. Cherry, S. Kroll, and J. F. Shogren, eds. London: Routledge.

CINYABUGUMA, M., T. PAGE, and L. PUTTERMAN (2005) Cooperation under the threat of expulsion in a public goods experiment, *Journal of Public Economics* **89**, 1421–1435.

CROSON, R., E. FATAS, T. NEUGEBAUER, and A. J. MORALES (2015) Excludability and contribution: A laboratory study in team production, *Journal of Economic Behavior and Organization* **114**, 13–26.

Department of Education (January 19, 2012) Faith schools. Available at http://www.education.gov.uk/schools/leadership/typesofschools/b0066996/faith-schools.

ECKEL, C. C., and P. J. GROSSMAN (2005) Managing diversity by creating team identity, *Journal of Economic Behavior and Organization* **58**, 371–392.

Economist (January 19, 2012) Clubbing together.

FALK, A., and U. FISCHBACHER (2006) A theory of reciprocity, *Games and Economic Behavior* **54**, 293–315.

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

GREINER, B. (2004) The Online Recruitment System ORSEE 2.0: A guide for the organization of experiments in economics. Working Paper 10, Department of Economics, University of Cologne.

ISAAC, R. M., and J. M. WALKER (1988) Group size effects in public goods provision: The voluntary contributions mechanism, *Quarterly Journal of Economics* **103**, 179–199.

ISAAC, R. M., J. M. WALKER and A. WILLIAMS (1994) Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups, *Journal of Public Economics* **54**, 1–36.

KOCHER, M., M. SUTTER, and V. WALDNER (2005) Exclusion from public goods as an incentive system: An experimental examination of different institutions, Mimeo.

LEDYARD, J. O. (1997) Public goods: A survey of experimental research. In *The Handbook of Experimental Economics*, J. H. Hagel and A. E. Roth, eds., 111–194. Princeton, NJ: Princeton University Press.

MAIER-RIGAUD, F. P., P. MARTINSSON, and G. STAFFER (2010) Ostracism and the provision of a public good: Experimental evidence, *Journal of Economic Behavior and Organization* **73**, 387–395.

SHAYO, M. (2009) A model of social identity with an application to political economy: Nation, class and redistribution, *American Political Science Review* **103**, 147–174.

SWOPE, K. J. (2002) An experimental investigation of excludable public goods, *Experimental Economics* **5** 209–222.

TAJFEL, H., M. G. BILLIG, R. P. BUNDY, and C. FLAMENT (1971) Social categorization in intergroup behavior, *European Journal Social Psychology* **1**, 149–178.