ELSEVIER

Journal of Public Economics 89 (2005) 1421–1435

JOURNAL OF PUBLIC ECONOMICS

www.elsevier.com/locate/econbase

Cooperation under the threat of expulsion in a public goods experiment

Matthias Cinyabuguma, Talbot Page, Louis Putterman*

Department of Economics, Brown University, Box B, Providence, RI, 02912, United States

Received 28 April 2003; received in revised form 28 November 2003; accepted 28 May 2004
Available online 28 October 2004

Abstract

In a public goods experiment with the opportunity to vote to expel members of a group, we found that contributions rose to nearly 100% of endowments with significantly higher efficiency compared with a no-expulsion baseline. Our findings support the intuition that the threat of expulsion or ostracism is a device that helps some groups to provide public goods.

JEL classification: C91; H41; D23
Keywords: Public goods; Collective action; Punishment

1. Introduction

Using the Voluntary Contributions Mechanism (VCM) to study the free-rider problem, experimentalists have found several stylized facts: Contributions tend to average somewhat over 50% of endowments in the first period of play, but to decay to low levels with repetition (Davis and Holt, 1993). Pre-play face-to-face communication is a powerful tool for fostering cooperation (Isaac and Walker, 1988a; Sally, 1995).[1]

* Corresponding author. Tel.: +1 401 863 3837; fax: +1 401 863 1970.
E-mail addresses: Matthias_Cinyabuguma@Brown.Edu (M. Cinyabuguma), Talbot_Page@Brown.Edu (T. Page), Louis_Putterman@Brown.Edu (L. Putterman).

[1] Recent experiments have extended this finding to audio-visual and on-line communications. See Brosig et al. (2003), Frohlich and Oppenheimer (1998), and Bochet et al. (in press).

doi:10.1016/j.jpubeco.2004.05.011

Cooperation can also be fostered by permitting players to target punishments at one another following their initial contribution decisions (Fehr and Gächter, 2000, 2002; Carpenter and Matthews, 2002; Sefton et al., 2002; Bochet et al., in press; Page et al., in press).

In many real organizations, the exercise of the power to determine who is in and who out of the relevant group provides an additional incentive to contribute to the group effort. Page et al. (in press) found that when subjects helped to determine group composition by ranking one another as prospective partners, higher contributing subjects were able to form groups of their own and to sustain cooperation, while self-interested players had incentives to act like cooperative types to get into groups with higher contributors. In this experiment, both contributions and efficiency were significantly higher in treatments with regrouping than in those in fixed groups.[2]

In Page et al., groups were periodically formed de novo, so that one can think only loosely in terms of "getting into" or "being ejected from" a group. The present study puts into sharper focus one aspect of membership determination, the right to expel members. We model a group that provides its members with a public good through a VCM but where the group has the power to expel members, converting them into non-members excluded from consuming the good. While this remedy is inapplicable to public goods like national defense and clean air, there are nonetheless real-world applications in which a good is public to members but not to non-members of a group. Examples include partnerships, teams, and associations from which members can be expelled for non-fulfillment of their obligations. Commonalities also exist with a broader range of collective action environments, such as ordinary workplaces, and clubs and organizations that provide both public goods and private benefits including sociability. Although our experiment involves a formal sanctioning process, this might also be viewed as an approximate representation of the less formal social sanctions that come into play in some of these settings. Our paper is related to a theoretical paper by Hirshleifer and Rasmusen (1989), which studied a finitely repeated public goods game, where the act of expulsion was costless to the expellers, and found an equilibrium with cooperation. Our paper is also related to a study by Bowles and Gintis (2004) of how a propensity to engage in costly expulsion of non-cooperators in small groups might have emerged in human pre-history through a process of group selection. To our knowledge, however, this is the first experimental study of expulsion as an incentive toward cooperation in a public goods game.

Our paper adds to the emerging literature on mechanisms to sanction free-riding and to a broader literature which has begun to provide a better understanding of behaviors in collective action environments. That emerging literature, including Andreoni (1988), Offerman et al. (1996), Fischbacher et al. (2001), Kurzban and Houser (2001), Ahn et al. (2003), Gunnthorsdottir et al. (2002), and Casari and Plott (2003), makes allowances for

[2] A public goods experiment by Ehrhart and Keser (1999) also had groups form endogenously, but because their subjects were free to join any group unilaterally, cooperative subjects had no protection from free-riders and little cooperation could be sustained. By contrast, Gunnthorsdottir et al. (2002) found that when the experimenters put initial high contributors together, cooperation was sustained more than in randomly formed groups.

the presence of more than one type of agent and for a Bayesian-type solution to finitely repeated games (Kreps et al., 1982). Our findings are difficult to explain unless some agents care about something in addition to monetary payoffs.[3]

The rest of the paper proceeds as follows. In Section 2, we describe the set-up of the experiment. In Section 3, we describe and analyze our experimental results with respect to contribution decisions, expulsion votes, and efficiency (efficiency is linear in net earnings in this experiment). We summarize our findings and provide concluding comments in Section 4. An extended discussion of theoretical issues and of further results is contained in our working paper (Cinyabuguma et al., 2004).

2. Experimental design

In a VCM, subjects are assigned to groups, each subject is provided with an endowment of money or tokens, and each is asked to allocate his or her endowment between a private and a group account. In our baseline treatment, a group of 16 subjects received 10 experimental dollars (hereafter, E$10) at the beginning of each of 15 periods of play.[4] In a given period, a subject i's earnings in experimental dollars were given by

$$\left(10 - C_i\right) + (0.2) \sum_{\text{all } j} C_j \tag{1}$$

where C_i is the amount i assigned (contributed) to the group account and $\sum_{\text{all } j} C_j$ is the sum of contributions to the group account by all members, i included.

In the baseline treatment, group composition remained fixed. Following each contribution decision, subjects were shown on their monitors a list of one another's assignments to the group account in the current period and in the previous two periods, each individual's average assignment so far, and his or her earnings in the current period and so far in the treatment, with each subject identified only by a randomly assigned letter that remained fixed for the treatment, and with subject information displayed in a random order. The task of allocating funds between the two accounts was repeated 15 times, with the terminal period being known from the start to all participants.[5]

The expulsion treatment differed from the baseline treatment as follows. Subjects also began in a group of 16 playing the game just described, but a second stage followed the

[3] Strictly speaking, it is sufficient that enough agents believe that the non-payoff maximizing type may exist, even if none in fact do. However, it is difficult to see how such beliefs could be sustained in a world in which all are payoff maximizers. Many of the papers cited, including Bochet et al. (in press) and Page et al. (in press), provide evidence that the "nonstandard" types do exist and indeed are numerous. See also Andreoni (1995).

[4] One experimental dollar converted to US$0.035 at the end of the session, which included two 15-period repeated game treatments. Average earnings were about US$25 for a 90- to 100-min session (a US$5 minimum was guaranteed each participant, but was never binding).

[5] A decision screen headed "you are now in period __ of 15" made sure that subjects could keep track of the number of periods remaining.

allocation decision in each period but the last. In this stage, subjects were shown the same information about one another's current and past contributions and earnings as in the baseline treatment, but next to the information for each subject was a box that could be checked to register a vote for removing that individual from the original group, called "the Green group". After each such vote, the list was shown again, with the boxes now showing the number of votes cast to expel each individual. If half or more of the current members of the Green group voted to expel a group member, that subject was moved for the remainder of the 15 periods to a secondary group dubbed "the Blue group"(for mnemonic purposes, think of expelled subjects becoming "blue"), where the same allocation game was played, except with endowments of only five experimental dollars each period, and with no voting stage.[6] When the coalition to expel an individual was large enough to be effective, each subject who voted for that person's expulsion was debited 25 experimental cents (approximately 1 real cent). If the required half of current Green group members was not achieved, no money was deducted for the votes cast. Since members saw the number of votes against each individual, votes could have a warning effect when insufficient to expel. As with the baseline treatment, the expulsion treatment continued for 15 periods, with the number of periods being announced at the outset and the current period number being displayed to subjects in each period.

We used two designs of experimental sessions. In the first design, called BE, a baseline treatment was followed by an expulsion treatment. This allowed us to study the treatment effect for the same group of subjects. In the second design, called EE, an expulsion treatment was followed by another expulsion treatment. This allowed us to check for sequencing and learning effects.[7] We refer to the expulsion treatment in the first design as BE1, the first (and only) expulsion treatment in this design. We refer to the first expulsion treatment of the second design as EE1, and the second treatment of this design as EE2.

We conducted four sessions of each design. There were 16 inexperienced participants in each of the eight sessions.[8] In both session designs, subjects learned of the second treatment only after the first treatment ended.[9] Table 1 summarizes the structure of the experiment as a whole.

The expulsion condition was designed so that it would never be rational for a payoff-maximizing agent to vote to expel another group member if there were common knowledge that all subjects cared about their own payoffs only, whereas expulsions might be rational with beliefs that some might have a preference for punishing free riders (compare Kreps et al., 1982).[10] Further discussion of the theoretical considerations behind

[6] That is, for a subject l belonging to the Blue group, period earnings were $(5 - C_l) + (0.2)\sum_{\text{all }k} C_k$ where C_l is what l assigned (contributed) to the group account, and $\sum_{\text{all }k} C_k$ is the sum of contributions to the group account by all members of that group, l included.

[7] We conducted no BB nor EB sessions because our main concern was to contrast B with E behaviors, and the BE order had training advantages.

[8] Subjects were drawn from the general undergraduate population at Brown University. The subjects were seated at computer terminals in a large room and were not permitted to communicate with one another during the experiment. Cash pay-out at the end of the experiment session was in closed envelopes.

[9] At that time, they were told there would be no treatment beyond the second one.

[10] The theoretical prediction for the game played in our experiment differs from that in Hirshleifer and Rasmusen (1989) because in our experiment there is a cost of voting to expel.

Table 1
Summary of designs and treatments

Design	Treatments in design	Description	Total number of sessions in design	Total number of subjects
BE	Baseline	Baseline treatment (15 periods)	4	64
	BE1	Expulsion treatment (15 periods)		
EE	EE1	Expulsion treatment (15 periods)	4	64
	EE2	Expulsion treatment (15 periods)		
Experiment as a whole			8	128

the design are provided in our working paper, which also contains the experiment instructions read by the subjects.

3. Results

Question 1. What was the pattern of contributions in the baseline treatment?

Result 1. There were substantial contributions to the group account in the baseline treatment, but they declined rapidly with repetition.

The middle line of Fig. 1 shows average contributions by period in the baseline treatment. Contributions in period 1 averaged 67% of endowment, but their level dropped to 31% in period 3 and continued to decline, averaging less than 20% throughout the last seven periods and standing at only 8% in period 15.[11] This is much like baseline behavior in other VCM experiments.[12]

Question 2. In expulsion treatments, how often did subjects cast expulsion votes and how many expulsions occurred?

Result 2. In most treatments with expulsion in most sessions, a majority of subjects voted to expel another subject at least once, and in almost all expulsion treatments and sessions, between one and four subjects were expelled from the Green group.

Fig. 2 summarizes for each of the expulsion treatments BE1, EE1 and EE2 the average cumulative number of subjects who had been expelled from the Green group as of the end

[11] The downward trend in contributions is confirmed by an OLS regression in which average contribution is the dependent variable and period number and a constant are the only independent variables. We include one observation per period for each of the 64 BE subjects, for a total of 960 observations. Regressions were estimated using OLS, using fixed effects for individuals, and by the Newey–West adjusted standard error procedure. In all models, the coefficient on the period term is negative and significant at the 0.001 level.

[12] Our experiment is unusual in having a large group size and a low marginal per capita return (MPCR) on cooperation. Isaac and Walker (1988b) conducted VCM experiments in which group size was either 4 or 10 subjects and the marginal per capita return on contributions (MPCR) was either 0.3 or 0.75. The levels and patterns of average contributions in their 10 person groups with MPCR=0.3 are not very different from those in our 16 person baseline groups with MPCR=0.2.

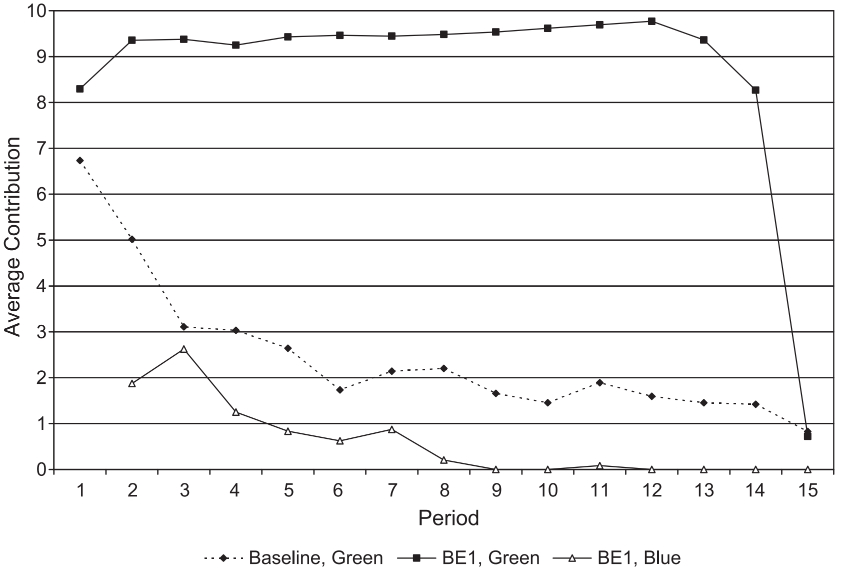

Fig. 1. Average contribution in BE design, by treatment and group.

of each period, and the average number of expulsion votes cast that period. The dashed lines show that an average of 4.5 were expelled by the end of the BE1 treatment, 2.25 by the end of the EE1 treatment, and 1.25 by the end of the EE2 treatment. The solid lines show that the largest number of expulsion votes were cast in each of the expulsion treatments' first period, usually leading to the expulsion of one or more subjects, and that the number of votes cast fluctuated thereafter, with a noticeable uptick in the last periods in

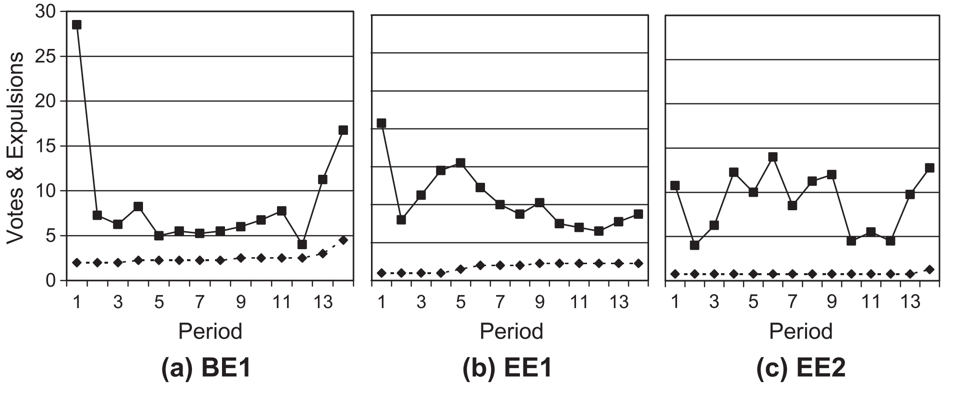

Fig. 2. Votes and expulsions. The top line in (a), (b), and (c) shows the average number of votes to expel in expulsion treatments BE1, EE1, and EE2. The lower line shows the average cumulative number of expulsions for the BE1, EE1, and EE2 treatments.

response to a rise in defections from contributing to the group account.[13] The data show expulsion votes to be well-targeted, with every expelled individual being either the lowest contributor in the group, or a next-lowest-contributor expelled together with a lowest contributor.[14] Nevertheless, a few individuals voted to expel high contributors, and contrary to the conjecture that higher contributors are more likely to vote to expel low ones, we found a negative correlation between an individual's average contribution and the number of votes he cast against low contributors (see Cinyabuguma et al., 2004).

Question 3. What was the effect of the expulsion possibility on cooperation? Was the threat of expulsion sufficient to raise contributions and earnings in the Green group? Overall?

Result 3a. In each of the three expulsion treatments BE1, EE1, and EE2, contributions to the group account were higher than in the baseline treatment even in the first period, indicating an anticipation effect, and they rose, rather than fell, with repetition and evidence of expulsions and expulsion votes.

Consider first the Green (i.e., original) group contributions in the BE1 expulsion treatment, averaged by period in the upper line of Fig. 1. The average proportion of the endowment contributed to the group account began at 83% in BE1 compared with 67% in period 1 of the four baseline treatments, suggesting that subjects adjusted their behaviors in anticipation of the threat of expulsion.[15] The average contribution rose with repetition for the first 12 periods, approaching 98% of endowment in period 12.[16]

[13] The right to vote to expel a member from the Green group was used relatively sparingly. On an average of 73% of the occasions in which a subject could vote, he or she did not vote to expel anyone, while 13.5% of the time, a subject cast only one expulsion vote; 22% of subjects cast no votes to expel during the course of a treatment, while 8% cast only one such vote, 24% cast between two and four expulsion votes in total, and 45% cast five or more expulsion votes in a treatment. (Numbers do not total 100% due to rounding error.) Lest it be supposed that the only real cost of expelling members was the voting charge, note that calculations show that in most of the cases of expulsion in the three expulsion treatments, the person expelled had contributed a positive amount in the period at the end of which he or she was voted out, with the average amount contributed by these individuals being just over E\$6, which, had it continued, would have given E\$6×0.2=E\$1.20 (≈US\$.04) each remaining period to each group member.

[14] Regressions show a significant negative relationship between a subject's contribution to the group account and the number of votes to expel her (Cinyabuguma et al., 2004).

[15] A within-subject comparison of first period contributions for BE subjects shows that of the 42 subjects whose contribution levels differed in the two treatments, 29 contributed more in the expulsion treatment. A Wilcoxon matched pair test shows the first contributions in the expulsion treatment to be higher at the 1% significance level (z=2.66). Session-level comparisons were also done between the average first period contribution in the baseline treatment of the BE sessions and the average first period contribution in the EE1 and EE2 treatments, respectively, using Mann–Whitney tests with one observation for each session, hence four observations for each treatment. Both EE1 and EE2 treatments were found to have higher first period contributions than the baseline treatment of BE, significant at the 5% level in one-tailed tests.

[16] The significance of the upward trend is confirmed by regression estimates. We used the 723 individual-level observations for each individual and each period he or she was in the Green group, up to and including period 13. Making the individual's contribution the dependent variable and the period number the sole explanatory variable, we find a positive coefficient on period that is significant at the 0.001 level regardless of whether the regression is estimated using an OLS model, a model with individual fixed effects, or a model with Newey–West adjusted standard errors.

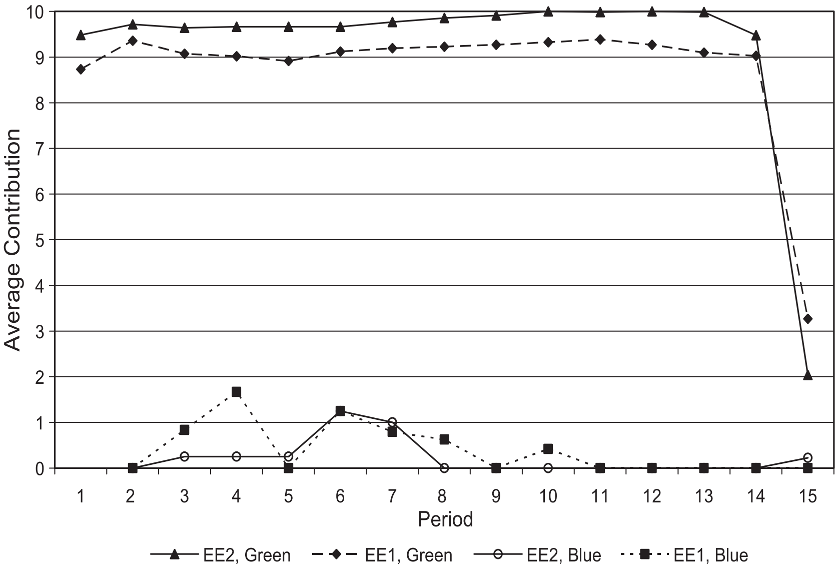

Fig. 3. Average contribution in EE design by treatment and group.

Next, consider the pattern of contributions in the second design, with two expulsion treatments. Fig. 3 shows that the patterns in both expulsion treatments of the EE design were similar to each other and similar to the expulsion treatment of the BE design. In the first expulsion treatment of this design, EE1 (second line from the top), Green group contributions began at 87% of endowment, rose in period 2, and remained between 90% and 94% of endowment through period 14.[17] In the second expulsion treatment, EE2 (the top line), average contributions began still higher, at 94.8%, and rose to the maximum possible for each and every Green group subject in periods ten through thirteen before declining slightly in period 14.[18] Contributions were higher in the second than in the first treatment in every session of the EE design.[19]

Result 3b. Average contributions were significantly higher in expulsion treatments than in the baseline treatment, even when Blue group members are included in the calculation.

[17] Regressions again show a significant upward trend through period 13. In this case, there are 749 individual-level observations for each individual and each period he or she was in the Green group, up to that period. We find a positive coefficient on period that is significant at the 5% level using an OLS model and one significant at the 10% in a model with adjusted standard errors. The corresponding coefficient is positive but not significant in a model with individual fixed effects.

[18] Regressions here too support a significant upward trend, this time even if period 14 is included. There are 857 individual-level observations for each individual and each period he or she was in the Green group up to and including that period. We find positive coefficients on period that are significant at the 1% level using an OLS or an individual fixed effects model, while the coefficient on period is positive and significant at the 5% level using a model with adjusted standard errors.

[19] The same result holds whether looking at Green group members only or at members of both groups.

At the individual level, a within-subject comparison shows that 51 of the 52 BE subjects who were never expelled from the Green group contributed more in the treatment with expulsion than in the baseline treatment, with the remaining subject contributing E$10 every period of both treatments. At the session level, a Mann–Whitney test comparing average contributions in the expulsion and baseline treatments of the four BE sessions, with one observation per session and treatment, finds that average contributions were significantly higher with expulsion, with a p-value of 0.0145 in a one-tailed test, even when contributions of expelled subjects are included. Mann–Whitney tests also show contributions to be higher in each of the two EE expulsion treatments than in the BE baseline treatment, with p=0.0145 even when the contributions of Blue group members are included.[20] (In all three expulsion treatments, the number of subjects sent to the Blue group was never large enough to make contributing to the public good profitable there, as a result of which most Blue group members soon set their contributions to zero—see Figs. 1 and 3).

Result 3c. Although subjects voted out of the Green group earned far less than others, earnings averaged over all subjects (including Blue group members) were nonetheless higher in expulsion than in baseline treatments. Thus, the addition of an expulsion option increased overall efficiency, and not just earnings in the non-expelled group.

Average earnings by period and group are graphed in Figs. 4 and 5. Averaging earnings across the full 15 periods of each treatment but treating the plays of a treatment in different sessions as separate observations, we performed Mann–Whitney tests paralleling those for contributions. The differences in earnings were found to be significant to the same degree as were those in contributions, with one-tailed Mann–Whitney tests yielding p-values of 0.0145 in all three expulsion versus baseline treatment comparisons (each of BE1, EE1, and EE2 versus baseline), even when Green group and Blue group members are included. Average earnings in every one of the 12 group plays of an expulsion treatment (four sessions each for BE1, EE1, and EE2) exceeded average earnings in every one of the four group plays of the baseline treatment.

Result 3d. There was a pronounced decline in contributions when the expulsion threat was removed in period 15 of each expulsion treatment.

Figs. 1 and 3 show sharp last period declines in contributions for Green group members in treatments BE1, EE1, and EE2. This supports the interpretation that high contributions in earlier periods were motivated mainly by fear of expulsion. The sharp decline also contrasts with the much milder end-game declines in cooperation in VCM experiments with communication (Brosig et al.; Bochet et al.) and in VCM experiments in which high contributors come to be grouped together either by the experimenter (Gunnthorsdottir et al., 2002) or by an endogenous grouping process

[20] In fact, average contributions, whether or not the Blue group is included, were higher in each and every one of the 12 plays of an expulsion treatment (that is, the expulsion treatments of each of the four BE sessions, the first expulsion treatment of each of the four EE sessions, and the second expulsion treatment of each of the four EE sessions) than in any of the four plays of the baseline treatment.

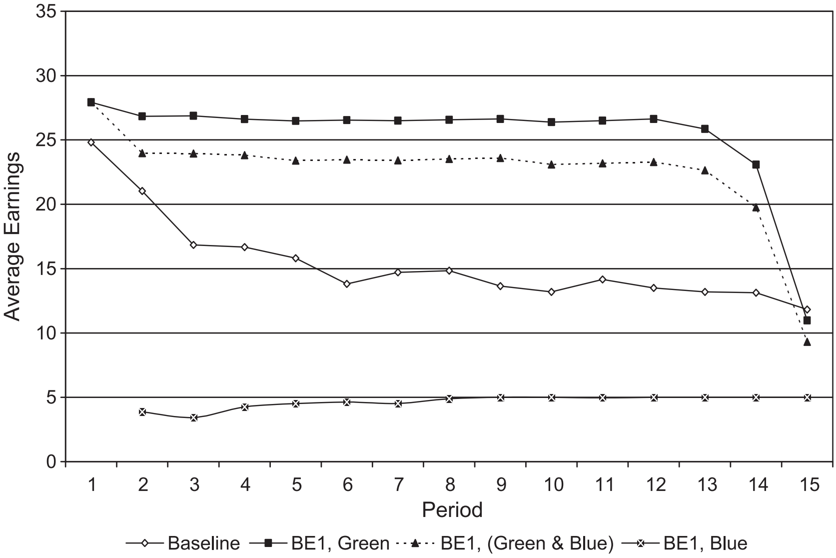

Fig. 4. Average earnings in BE design by treatment and group.

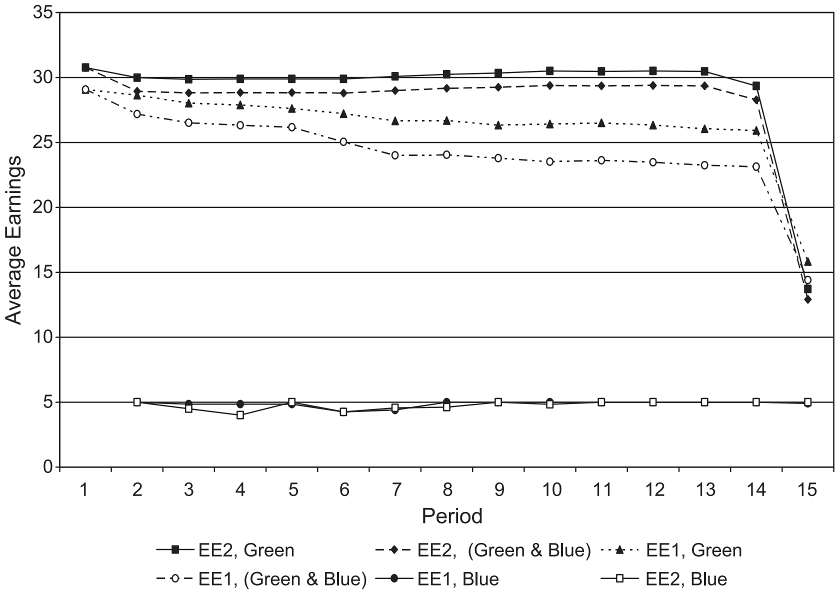

Fig. 5. Average earnings in EE design by treatment and group.

Table 2
Effects of expulsion votes and expulsions on contributions of green group members

Independent variable	Two-limit Tobit regression, dependent variable: $C_{i,t+1}$				
	Periods 2–14			Periods 2–15	
	(1)	(2)	(3)	(4)	(5)
$C_{i,t}$	0.722***	0.523***	0.7489***	1.113***	1.431***
	(0.069)	(0.063)	(0.071)	(0.182)	(0.204)
No. of votes against i in period t	0.408***		0.385***		0.694***
	(0.056)		(0.057)		(0.160)
No. expelled in period t		−0.283**	−0.130	−0.023	0.511*
		(0.123)	(0.122)	(0.361)	(0.311)
No. expelled in period t×no. of votes against i in period t		0.359***	0.228**	0.660**	
		(0.106)	(0.106)	(0.297)	
No. expelled in period t×period 15 dummy				−16.483***	−16.552***
				(0.957)	(0.952)
Constant	4.171***	6.411***	3.953***	4.141	0.592
	(0.923)	(0.883)	(0.934)	(2.636)	(2.812)
No. of observations	2080	2080	2080	2240	2240
Prob>χ^2	0.000	0.000	0.000	0.000	0.000
Pseudo R^2	0.386	0.375	0.387	0.227	0.229

Numbers in parentheses are standard errors.
* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

(Page et al.), which suggests to us that the threat of expulsion not only substitutes for but also reduces the emergence of trust and expectations of cooperation among our subjects.[21]

Question 4. Is there more direct evidence of a signaling effect of expulsion votes?

Result 4. In expulsion treatments, subjects increased their contributions more the more others cast votes to expel them from the Green group, and, especially for the subjects more targeted in the votes, after the expulsion of other group members.

To identify the impacts of expulsions and expulsion votes on the contributions of individual Green group members, we report in Table 2 a series of five 2-limit Tobit regressions in each of which the dependent variable is individual i's contribution to the group account in period t+1 and explanatory variables include i's contribution in

[21] We find it interesting that despite the sharp drop in contributions in period 15 of treatment EE1 the same subjects quickly returned to high contributions in treatment EE2. Indeed, contributions were higher in EE2 than in EE1 although fewer subjects were expelled in the EE2 treatment. Average earnings were also significantly higher in the EE2 than in the EE1 treatment.

period t and a set of individual fixed effects.[22] Columns (1)–(3) report regressions in which the dependent variable is observed in periods 2–14 only, leaving out the pronounced end-game effect, and columns (4) and (5) report regressions for periods 2–15, with an end-game control (as specified below).[23] In regression (1), we include the number of votes cast against i in period t's voting stage, finding it to have a statistically significant and economically important impact on contributions. Clearly, subjects took votes against them as a warning that they could be expelled if they did not raise their contributions.

When instead of the number of votes we used the number of subjects (if any) expelled from the Green group in period t in subject i's session (a specification not shown in the table), its coefficient was insignificant. In regression (2), we include both the number expelled and the product of the number expelled and the number of votes against subject i. The result supports the conjecture that seeing subjects expelled from one's group led i to increase her contribution to the degree that i also received some votes (see the interaction term votes×expelled); non-recipients of expulsion votes on average slightly reduced their contributions, however.[24]

In regression (3), both votes received and the expulsions and votes×expulsion variables are entered simultaneously. The coefficient on number expelled becomes insignificant, while the other two variables remain significant, suggesting that a subject who received expulsion votes increased her contribution, and that she did so all the more the more group members had received enough votes to actually be expelled.

Regressions (4) and (5) include the observations for period 15, in which contributions declined dramatically as the threat of expulsion was removed. Since there were also a number of expulsions of late-game defectors in period 14, this large decline would suggest, spuriously, that expulsions led quite generally to lower contributions, unless controlled by an end-game dummy. The specifications shown thus

[22] The Tobit estimation treats observations as potentially right censored if $C_{i,t+1}$=10 and potentially left censored if $C_{i,t+1}$=0. Data from BE1, EE1, and EE2 treatments are pooled, with included observations being those of all subjects who remained in the Green group for the full 15 periods. Treatment effects (EE1, etc.) cannot be estimated due to the overlap with the set of subject dummy variables. Separate estimates (not shown) without subject fixed effects but with dummy variables for treatment give similar results (with some changes in significance levels) for the other variables, with significant negative coefficients for the EE1 dummy variable in all specifications, and positive coefficients for the EE2 dummy variable which are highly significant in the first three estimates (those ending at period 14) and entirely insignificant in the last two estimates (which include period 15). We also estimated the same five specifications without subject fixed effects, including only the observations from treatments EE1 and EE2, and with a dummy variable for treatment EE2 only. Again, the other coefficients are similar, and for this sample, the EE2 dummy variable has a large (ranging from 0.91 to 2.28) positive coefficient that is significant at the 1% level for every specification.

[23] Contributions of period 1 enter the estimates as a lagged dependent variable.

[24] The typical scenario was that in the first period of play, when expectations of others' choices were still diffuse, contributions varied widely and only the lowest contributors received expulsion votes, with the very lowest being expelled. Seeing that many had contributed less than they had without being expelled, a few of the very highest contributors, who had received no votes, then slightly reduced their contributions, even as most group members raised theirs. Usually, such downward shifts at the top were then reversed as all Green group members shifted towards high contributions, as seen in Figs. 1 and 3.

include an interaction between a period 15 dummy variable and the number expelled.[25] The negative coefficient on the number expelled becomes smaller and still less significant in regression (4), which also includes the term interacting number expelled with votes against i. When this last interaction variable is dropped but the votes variable is entered by itself, in regression (5), both votes against i and the number expelled have significant positive coefficients.

In all, the results in Table 2 support the proposition that expulsions, votes short of expulsion, and the combination of expulsions and votes motivated subjects to raise their contributions to the group account, thus directly contributing to the difference between expulsion treatment and baseline outcomes.[26]

4. Conclusion

In partnerships, communities, and organizations of various types, people often strive to cooperate in the absence of state enforcement. The threat that an uncooperative individual may be ostracized or expelled is one way to foster cooperation. Like a trigger strategy, the expulsion mechanism threatens a dire consequence with irreversibility. In the experiment, much of the effectiveness of the expulsion mechanism can be attributed to the fact that expulsion was used so sparingly, less than three times on average in a treatment of 15 periods. Low contributors were often warned by increasing but still less-than-majority votes before they were expelled, and most of the low contributors responded to the warnings with higher contributions, thus avoiding expulsion. The “dark side” of the mechanism is that by relying so heavily on the implied threat of expulsion, it didn’t allow group members to develop trust in one another’s positive reciprocity, so cooperation collapsed when the threat was removed.

An important theoretical point is that our experiment was designed so that it would never be rational to vote for expulsion if all subjects cared only about maximizing their own payoffs and had common knowledge of this. If each individual cared only for her own material outcome without emotions of annoyance at others’ free riding, and if this were common knowledge, the threat of expulsion would not be credible.[27] Our results thus join others in the recent literature which suggest that human nature includes other types of preferences as well as self-interested preferences, and that this heterogeneity renders credible what otherwise is not a credible threat, making the danger of punishment (in this case, exclusion) a reason to pull one’s weight in a group endeavor.

[25] This specification was found to give clearer results than ones in which an end-game dummy variable appears as a free-standing term.

[26] We took a different approach to testing for the effect of expulsions on contributions in our working paper, estimating Tobit regressions of the change in the average contribution at the group level as a function of the number expelled in the previous period, treatment dummy variables, and a time trend. The coefficient on the number expelled is highly significant and positive.

[27] See Hirshleifer (1987).

Acknowledgement

We wish to thank Sam Bowles, Rob Boyd, Ernst Fehr, and Herb Gintis for suggestions leading to the design of the experiments reported here, and Devin Borland, Chris Horn, George Athanassocoupolos, and Bulent Unel for programming, help in conducting the experiments, and pointers on data analysis. We are grateful to the guest editors and to two anonymous referees for their valuable comments. We would like to thank the MacArthur Foundation Network on Norms and Preferences and the National Science Foundation grant SES-0001769 for support of this research.

References

Ahn, T.K., Ostrom, Elinor, Walker, James, 2003. Heterogeneous preferences and collective action. Public Choice (117), 295–314.

Andreoni, James, 1988. Why free ride? Strategies and learning in public goods experiments. Journal of Public Economics (37), 291–304.

Andreoni, James, 1995. Cooperation in public-goods experiments: kindness or confusion? American Economic Review (85), 891–904.

Bochet, Olivier, Talbot Page, Louis Putterman, in press. Communication and Punishment in Voluntary Contribution Experiments. Journal of Economic Behavior and Organization.

Bowles, Samuel, Gintis, Herbert, 2004. The evolution of strong reciprocity: cooperation in heterogeneous groups. Theoretical Population Biology (65), 17–28.

Brosig, Jeannette, Weimann, Joachim, Ockenfels, Axel, 2003. The effect of communication media on cooperation? German Economic Review 4 (2), 217–241.

Carpenter, Jeffrey, Matthews, Peter, 2002. Social Reciprocity. Middlebury College Department of Economics Working Paper 29.

Casari, Marco, Plott, Charles R., 2003. Decentralized management of common property resources: experiments with a centuries-old institution. Journal of Economic Behavior and Organization 51 (2), 217–247.

Cinyabuguma, Matthias, Page, Talbot, Louis Putterman, 2004. Cooperation Under the Threat of Expulsion in a Public Goods Experiment. Working Paper 2004-05. Department of Economics, Brown University.

Davis, Douglas D., Holt, Charles A., 1993. Experimental Economics. Princeton University Press, Princeton.

Ehrhart, Karl-Martin, Keser, Claudia, 1999. Mobility and Cooperation: On the Run. CIRANO 99s-24, Montreal.

Fehr, Ernst, Gächter, Simon, 2000. Cooperation and punishment. American Economic Review 90, 980–994.

Fehr, Ernst, Gächter, Simon, 2002. Altruistic punishment in humans. Nature 415, 137–140.

Fischbacher, Urs, Gächter, Simon, Fehr, Ernst, 2001. Are people conditionally cooperative? Evidence from a public goods experiment. Economics Letters 71, 397–404.

Frohlich, Norman, Oppenheimer, Joe, 1998. Some consequences of e-mail vs. face-to-face communication in experiment. Journal of Economic Behavior and Organization 35 (3), 389–403.

Gunnthorsdottir, Anna, Daniel Houser, Kevin McCabe, Holly Ameden, 2002. Disposition, History and Contributions in a Public Goods Experiment. Unpublished manuscript, Department of Economics and Economic Science Laboratory, University of Arizona.

Hirshleifer, Jack, 1987. On the emotions as guarantors of threats and promises. In: Dupré, John (Ed.), The Latest on the Best: Essays on Evolution and Optimality. MIT Press, Cambridge.

Hirshleifer, David, Rasmusen, Eric, 1989. Cooperation in a repeated prisoners dilemma with ostracism. Journal of Economic Behavior and Organization 12 (1), 87–106.

Isaac, R. Mark, Walker, James M., 1988a. Communication and free-riding behavior: the voluntary contributions mechanism. Economic Inquiry 26, 585–608.

Isaac, R. Mark, Walker, James, 1988b. Group size effects in public goods provision: the voluntary contributions mechanism. Quarterly Journal of Economics 103 (1), 179–199.

Kreps, David, Milgrom, Paul, Roberts, John, Wilson, Robert, 1982. Rational cooperation in finitely repeated prisoners' dilemma. Journal of Economic Theory 27, 245–252.

Kurzban, Robert, Houser, Daniel, 2001. Individual differences in cooperation in a circular public goods game. European Journal of Personality 15 (S1), S37–S52.

Offerman, Theo, Sonnemans, Joep, Schram, Arthur, 1996. Value orientations, expectations, and voluntary contributions in public goods. Economic Journal 106, 817–845.

Page, Talbot, Louis Putterman, Bulent Unel, in press. Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry, and Efficiency. Economic Journal.

Sally, David, 1995. Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. Rationality and Society 7 (1), 58–92.

Sefton, Martin, Robert Shupp, James Walker, 2002. The Effect of Rewards and Sanctions in Provision of Public Goods. Working Paper, University of Nottingham and Indiana University.