Endogenous Group Formation*

T.K. Ahn† R. Mark Isaac‡ Timothy C. Salmon§
Florida State University Florida State University Florida State University

December 2004

Abstract

While the rules governing the formation of groups engaging in collective action may have significant impact on group size and behavior of members, most experiments on public goods have been conducted with the subjects in fixed groups or of fixed sizes. We study endogenous formation of groups in a public-goods provision game by allowing subjects to change groups under three sets of rules: free entry and exit, restricted entry and free exit, and free entry and restricted exit. We find that the rules governing entry and exit do have a significant impact on individual behavior and group-level outcomes.

JEL Codes: C92, H41, D85
Key Words: Public Goods, Group Formation, Social Networks

1 Introduction

As clearly set-out in Olson (1971), the question of how groups form and engage in collective action is a question of great interest to both economists and political scientists. Olson himself states in the *Logic of Collective Action* that "[T]he movement in and out of the group must no longer be ignored" (p. 36). Looking across different collective action units, it is clear that there are a variety of different rules in use for group formation. Residents in a neighborhood or apartment complex engage in various forms of collective action, from keeping their yards well mowed, participating in a Saturday drive to clean-up their neighborhood park, to organizing neighborhood watch patrols. Most of these activities are on a voluntary basis and whoever participates generates positive externalities to other members of the neighborhood as these activities raise the quality of living and the property values to the entire neighborhood. In most neighborhoods and apartments, entry and exit is free in the sense that current residents can not block the entry and exit of other residents. On the other

*The authors would like to thank the Florida State University and the FSU Foundation as well as John Scholz for research support in funding these experiments. We owe a substantial debt of gratitude to Justin Esarey for programming the experiments. We also thank various participants in seminars or conference presentations at George Mason University, The New and Alternative Directions for Learning Workshop hosted at Carnegie Mellon University, and The Southern Economics Association Conference who provided many helpful comments that improved our experimental design and analysis of data.

†Department of Political Science, Florida State University, Tallahassee, FL, 32306-2230, tahn@fsu.edu. Phone: 850-644-4540 Fax: 850-644-1367.

‡Department of Economics, Florida State University, Tallahassee, FL, 32306-2180, misaac@mailer.fsu.edu. Phone: 850-644-7081 Fax: 850-644-4535.

§Department of Economics, Florida State University, Tallahassee, FL, 32306-2180, tsalmon@fsu.edu. Phone: 850-644-7207 Fax: 850-644-4535.

hand, most of the residential property of Manhattan is owned by Co-ops, where the Co-op boards are allowed to deny entry to potential residents under a variety of different rules[1].

While neighborhoods have more or less fixed sizes, other collective action entities do not. Labor unions, one of the key examples in the *Logic*, may include all the employees of a workplace, or a tiny fraction. The sizes and compositions of individual unions depend in part on their entry or eligibility rules. Early unions were very specific in their eligibility requirements restricting members to specific factories or specific jobs. Some modern unions have attempted to substantially expand their membership beyond these bounds. Further, labor unions have the option of forming larger groups by banding together to pool their resources as evidenced by the AFL-CIO merger in 1955. It is therefore quite an interesting question to ask how different rules for determining group membership might impact not only the effectiveness of the group but also the size of the group.

We can think of real world collective action problems as having two distinct, but dynamically interacting, stages. There is a process of formation of groups that can have a variety of rules based upon the degree of control current members have over the exit of other members or on the entry of new members. This process will determine the size and composition of groups. At another level, the members of existing groups decide on their levels of contribution for the relevant public good, be it cleaner and safer neighborhoods, loyalty to the union to increase bargaining power with employers or more profitable business operation. The degree of success and failure in these collective endeavors in turn affects outsiders' willingness to join a group or the incumbent member's desire to remain in the group, again affecting the size of the group over time. Our goal in this research is to incorporate the dynamic interactions between group formation rules and group behavior building on the rich tradition of experimental research on the provision of public goods.

In most prior experimental literature on public goods, the question of how an existing group performs has been considered a separate issue from how and why groups form in the way that they do and what effect the formation process might have on behavior of the group. Three of the earliest experimental series partially investigate the idea of how group formation dynamics might impact behavior by including group size as an experimental treatment (Marwell and Ames (1979), Isaac and Walker (1988), Kim and Walker (1984)). In these papers, it is the comparative static influences of group size that are of concern, which might be seen as a partial equilibrium approach to studying how behavior changes as group size evolves.

Perhaps the most extensive treatment of the interaction between group size and public good provision has been Isaac, Walker, and Williams (1994). In their paper a standard linear public goods environment is altered both along the dimension of marginal per capita return (MPCR), a measure of the individual self-interested incentives to contribute to provision of the public goods, and of group size (N=4, 10, 40, and 100). The surprising result from this paper was that controlling for the effects of varying MPCR (lower MPCR leads to lower contributions, although with a modified pattern in large groups) large groups were more effective in providing public goods than smaller groups.

Only recently, did experimentalists start investigating endogenous formation of groups in the context of public-goods provision. Gunnthorsdottir, Houser, McCabe, and Ameden (2001), for

[1] There is recent, highly publicized evidence that the entry /exit regulation performed by Co-op boards does have an impact on behavior. A December 16, 2004 *New York Post* story reports a claim that one of the motivating factors behind Mary Tyler Moore's activism against the Co-op board of 927 Fifth Ave. on behalf of two red tailed hawks (Pale Male and Lola) was the fact that the Co-op board had recently denied a potential buyer for her apartment in the building. The man was denied because "He was just what you don't want in a family building." According to the article, a source on the board stated that the event caused Ms. Moore to wage the campaign against the board "as a personal vendetta." While the claim may be unfounded, this does suggest that the board realizes its decisions may have an important impact on the behavior of its residents.

example, match high contributors with other high contributors, low contributors with other low contributors. The result of the regrouping of individuals with cooperative histories was a decrease in the typical decay in contributions to the public good. Moving just a step towards freer group choice, Page, Putterman, and Unel (2002) allow individuals to rank whom they would like to have in their new groups and then a computer matching algorithm matches subjects into new groups. Ranking was a costly exercise for the participants, and the size of all groups was fixed at four by the experimenters. The authors find that this regrouping algorithm by itself or combined with a punishing mechanism, increases efficiency significantly.

Ehrhart and Keser (1999) take the full plunge into allowing groups and group sizes to be formed endogenously. After each period, the nine individuals could move to new groups at a cost leading to groups with 1-9 members. The second stage of each period consisted of each of the members engaging in a standard public goods task with the rest of their current group members. Several parts of their experimental design stand out. First, the incentives of the groups are constructed so that the dominant strategy in the stage game is to contribute zero to the group good and this is independent of group size. Secondly, the MPCR falls as group size increases, so that the addition of free-riders to a group can be harmful to the previous membership. Thirdly, as mentioned, mobility is costly. The authors observe a great deal of mobility, and find some regularities in the data. However, the decreasing MPCR in a group as the number of members increases seems to confound a clear interpretation of the data at least without a complementary control series in which MPCR is constant with group size. Concern, Fehr, and Fellner (2003) allows subjects in a public goods setting to choose partners according to two different mechanisms, bidirectional and unidirectional, and find that contribution levels do vary significantly based on the mechanism used.

We go several steps beyond this existing literature to deal with group formation in a more systematic manner. We treat the *group formation rule* as an experimental treatment variable. Our experimental treatments include group formation under free entry and exit rules similar to Ehrhart and Keser (1999), but we also examine group formation using the rules of "restricted" entry or exit that mimic mechanisms used in the formation of a variety of different types of groups. In naturally occurring situations, there are an almost unlimited number of examples of different mechanisms governing entry and exit. One can join the Sierra Club simply by filling out a card, but joining a country club or a law firm may require the approval of the existing members. In order to get that approval, any one of a countless number of voting mechanisms may be in place: unanimity, supermajority, majority, and so forth. Likewise, one may resign from the Sierra Club at any time, but, on the other hand, a property owner may not be freely allowed to de-annex from a municipality without the approval of the remaining current "members." As a starting point for investigating these issues, we implement different mechanisms in which entry into or exit from a group requires approval of a majority of the incumbent group members. Thus, as in Coricelli, Fehr, and Fellner (2003), we experimentally compare how different group formation rules impact behavior, but we do this in the context in which group sizes are endogenous, not fixed.

Interacting with the rule of group formation is the level of information available to both current and potential group members about the level of contributions of others. When individuals are allowed to choose which group to join and when incumbent members of a group have the power to approve or deny entry and/or exit of members, information becomes an important factor. That is, it matters what level of information individuals have on the past performance of the groups which they might want to join and what level of information voters have on the past behavior of the applicants. It is quite often the case that aggregate information on group decisions is available: How much money did the Friends of the Music School donate to scholarships last year? What are the taxation and expenditure levels of the city that is trying to annex my property? Or, for example, we note that one can access on the web page of the Presbyterian Church, U.S.A. and see the number

Treatment	Group Formation Rules	Sessions	Number of Subjects
Treatment 1	Free Entry, Free Exit	1,4,7,9	48
Treatment 2	Restricted Entry, Free Exit	2,3,5,6	48
Treatment 3	Free Entry, Restricted Exit	8,10,11,12	48

Table 1: Listing of experimental treatments and sessions run.

of members and aggregate contributions of any of their congregations in the country for the past several years. On the other hand, there are cases in which historical individual information may be available. For example, someone proposing to make a large, leadership gift to a philanthropic organization might have an expectation to be told of the specifics of other, similar, large gifts. In addition to information on potential groups, there is the reverse issue of how much do existing groups know about potential new members. Members of some existing groups (for example a law firm or a country club) may know the specific individual history of a prospective member's record in contributing to public goods. On the other hand, a church, scout troop, or political lobbying organization may know nothing about new members. Again, there are many alternative ways in which the levels of information can be experimentally implemented and we have chosen to allow individuals to observe only the aggregate contribution levels of one's own and other groups unless someone applies to enter into or exit from a group under the restricted entry or exit condition. In the restricted entry or exit treatment, the voters can see individual specific history of behavior of the applicants before they vote to approve or deny the request to join or depart. Our experimental design is not intended to mimic the institutional rules of any specific organization but rather to use these mechanisms as an initial foray into such rules to determine if different mechanisms do have much of an impact and if so, in what way.

Section 2 will provide an overview of the design of the experiments. Section 3 will discuss the results and section 4 will conclude.

2 Design of Experiments

The experiment involves three treatments: (1) Free Entry, Free Exit (2) Restricted Entry (with free exit) and (3) Restricted Exit (with free entry). Four experimental sessions were run for each treatment. Twelve subjects participated in each of the twelve sessions for a total of 144 subjects. Table 1 summarizes the overall design of the experiment and the session identification numbers for each treatment. Subjects were recruited mostly from economics courses at a variety of levels at Florida State University. All sessions were conducted in a computer lab using software created with z-Tree (Fischbacher (1999)). In Sessions 1, 2 and 3, subjects were paid a $7 show-up fee and ECUs translated into dollars at a rate of 2 ECUs=$0.01. Average subject earnings were in the range of $12-13 in these sessions. In all other sessions, subjects were paid a $10 show-up fee and ECUs translated into dollars at a rate of 1 ECU = $0.01. Average subject earnings were in the range of $20-25 in these sessions[2]. Sessions lasted on average an hour and a half to two hours.

[2]The change in the show-up fee was an attempt to increase the show-up rate for our subjects. The change in the exchange rate was due to the fact that subject earnings in the first 3 sessions were much lower than predicted and were increased to raise the implied hourly rate. We have tested the effects on behavior from the change and see no statistically significant impact.

Group Size	1	2	3	4	5	6	7	8	9	10	11	12
Individual Optimum	3	3	3	3	3	3	3	3	3	3	3	3
Group Optimum	3	5	6	7	8	9	9	10	11	11	12	13

Table 2: Optimal investment amounts into group account depending on the size of the group.

2.1 The Stage Game

The stage game that each subject faces in each period is a variant of the standard Voluntary Contribution Mechanism (VCM), public goods provision game. In each period, each subject belongs to a group of size $N \in [1, 12]$ and makes a decision on how to divide 15 "tokens" between his own individual account and the group account for the group in which he is a member. The exact group size is endogenously determined within the extremes of one and twelve. Twelve is the maximum group size since twelve subjects participated in each session. In the experiment the decisions are framed as "investment" decisions. In the discussions that follow, we will often refer to the investment in the group account as a *contribution* as this terminology is more natural for researchers, but that was not the language used in the experiments. Let x_i denote the number of tokens individual $i \in \{1, 2, ..., 12\}$ contributes to the group account. Let G_i represent the set of other members in i's group (not including i). The monetary payoff to individual i is

$$\pi_i = 0.5(15 - x_i) + 1.5(x_i + \sum_{j \in G_i} x_j) - \frac{1}{27} x_i^3. \tag{1}$$

Investment to the private account yields .5 Experimental Currency Units (ECUs) to the subject while contributions to the public account generate 1.5 ECUs for that subject and 1.5 ECUs per token for each other member of the group. Individual i also receives 1.5 ECUs for each of the tokens that other group members contribute to the group account. Investment into the individual account is costless but investing a number of tokens x into the group account costs that individual $(1/27) * x^3$ ECUs.

An important feature of this design is that the dominant strategy is not 0 or on the lower boundary of a subject's choices as in many public goods games. The dominant strategy choice for each individual is to invest 3 tokens to the group account and 12 to the individual account. This is independent of the group size and of the number of tokens invested by others. Taking the group choices into account, the Nash equilibrium for the stage game would involve all twelve subjects getting into (or remaining in) the same group and contributing three tokens each to the group account. Each individual earns 59 ECUs per period playing according to the Nash equilibrium. Due to the externality on group account investments, the group optimal contribution level (i.e. the contribution that jointly maximizes payoffs for the entire group) is different and is a function of group size as shown in table 2. The socially efficient or group payoff maximizing scheme involves all 12 subjects getting into (or remaining in) the same group and investing 13 tokens each, returning a payoff of 173.1 ECUs per person per period. The per person payoff in the social optimum is about 290% of that in the Nash equilibrium. One important detail to note is that the equilibrium and efficient solutions differ only in the level of contributions, not in the size of the optimal group.

Another important detail about the payoff schedule induced by this setup is that as group size goes up, investing the optimal level begins to incur a greater risk of losing money. At a contribution level of 8 (socially optimal for $n = 5$) the subject begins to possibly make losses (net earnings -3.46 ECUs) unless his fellow group members contribute positive amounts. As the group size goes up and the socially optimal contribution level increases, the risk of losing money to the subject for

contributing at the socially optimal level increases. In a group of n = 12, contributing 13 tokens yields a personal net payoff of -60.87 ECUs which requires a contribution level of around 3. 64 for each of the other 11 group members in order for the high contributor to break even. So if all group members are contributing at the individual optimum, 3, the high contributor will lose money. Thus contributing at levels past 7 or 8 requires a substantial amount of trust in the willingness of other group members to contribute.

2.2 Experimental Treatments

The experiments begin with a preliminary series consisting of three periods in which subjects make investment choices in groups consisting only of themselves. There are two reasons for having this preliminary series. First, these preliminary periods give the subjects experience with the investment environment without worrying about the group formation or group dynamic issues. Second, it allows us to observe the decisions of the subjects when there are no issues of social efficiency involved in their choices. To make this latter point clear, the choices in the preliminary phase were never observable by other participants and this was communicated very clearly to the subjects. Thus, the choices in the preliminary periods could not be used to signal the degree to which a subject was a high contributor. This completely eliminates any tension between the socially and individually optimal behavior.

After the preliminary phase, the main phase began with one more period in which subjects made an investment choice in a group of only themselves. In all of the subsequent periods (up to a total of 20), a period would begin with the subjects being asked if they wished to switch groups. When asked this question they were able to view a screen showing the average contribution levels and group sizes of the 12 possible groups over the previous 5 periods. This included the number of subjects in each group at the end of the previous period. If a subject indicated that she did wish to change groups, she was presented with a new screen. The new screen was similar to the previous one but showed, along with the aggregate contribution level for each group for the past 5 periods, the number of subjects in each group who had chosen to remain in the group. Subjects were then allowed to choose which group they wished to enter. Subjects were labeled 1-12 and groups were labeled A-L. When the experiment started, subjects were randomly assigned into a letter group to begin in such that each began in a group of 1. Choosing to move to a new group and choosing which group to move to were both costless.

The three experimental treatments differ from one another in terms of whether entry or exit was "free," in the sense that it was automatic as long as a subjects wanted, or "restricted," in the sense that it was subject to rules of joining and departing with a possibility of failure. In the Free Entry/Free Exit treatment, both entry to or exit from a group was unrestricted. In the Restricted Entry Treatment, exit was unrestricted but entry into a new group required approval by the majority of the members of the group to which the applying subject has applied to join. In the Restricted Exit treatment, entry into a group was unrestricted, but exit from a group required approval by the majority of the members of the group from which a subject has applied to depart.

In the Restricted Entry treatment, when a person applies to join a new group, the other subjects who were in the target group at the end of the prior period and chose "no" to the question of if they wished to move to a new group were shown a choice history of each applicant[3]. The choice history consisted of the applicant's actual contributions to the group account and the size of the

[3] Thus those who are in the process of moving do not get to vote on applicants into their current group even if they are rejected by the group they are attempting to move into. Also, if a subject was in say group B at the end of one round, responds "yes" that he wishes to move into a new group and chooses B again, he is automatically accepted back into the group but he is not allowed to vote on new applicants in that same round.

group the applicant was in over the previous five periods. Then they voted "yes" or "no" on the applicant. If more than 50% of the current members of a group voted "yes", to allow entry, the applicant was allowed to enter, otherwise they were rejected and returned to their previous group. A subject did not need the approval from the members of his current group to exit that group.

In the Restricted Exit Treatment, the procedure is quite similar. Subjects are first asked if they wished to move to a new group. If they selected "yes" then they see the same screen as before asking them which group they wished to move to. After they have made their choice, any subject who answered "no" to the question of whether or not they wished to move would see a voting screen with a choice history for all of the other members of their group who are attempting to exit the group. The choice history of the applicants to exit that was shown to the voting members of the group was the same as that in Restricted Entry Treatment. Each of the voting members voted "yes" or "no" on each applicant. If more than 50% of the members voted "yes" on an applicant, she was allowed to exit. Otherwise her application to exit was denied and she was forced to remain in the group. If the departure was approved, the subject would automatically enter whichever group she had previously chosen.

At the conclusion of session 4-12 we included a "bonus" question for our subjects. It was designed to determine if subjects could understand the group optimal contribution calculations. The bonus question asks the subject to identify the group optimal level of contribution when group size is 5. The answer to the question is 8 and subjects were paid an extra $1 if they got within one token above or below. The exact wording of the question was the following:

> This is a bonus question related to the experiment you just completed. If you answer correctly (or within 1 token above or below) you will earn an extra $1.
>
> Assume you are in a group of 5 people (4 plus you). What would be the number of ECUs that each group member would need to invest into the group account to lead to the highest payoff to the entire group?

To help the subjects understand the complexities of the environment, we ran them through an extensive help system explaining all stages of the decision process. We also provided hardcopies of extensive tables summarizing their earnings from any combination of decisions they and their fellow group members might make. During the experiment when a subject was asked to make an investment decision, the software included a "test" button that would allow the subjects to enter a proposed investment level and then see all computations regarding the effect on their payoff and the payoff of others though of course without including the effect from the contributions of others. The instruction script we used can be found in appendix A while screenshots and the payoff tables used are available from the authors upon request[4].

3 Experimental Results

Our analysis will consist of an exploration of the patterns in the data to attempt to identify and explain any systematic behavioral differences that can be attributed to the group formation institutions. While we will not present a full formal model of this environment, the standard game theoretic prediction is readily apparent. Because this is a finitely repeated game, the subgame perfection prediction on the contribution level that the subjects should make per period is 3. In regard to the group formation process, there would be a number of different possible ways to define equilibrium strategies. Any of them must involve subjects using strategies that enable them to

[4] They are also available at http://garnet.acns.fus.eu/~tsalmon.

form a single global group as soon as possible.[5] This prediction does not depend on the institution, though technically we would have to define strategies relating to voting decisions in the second two treatments to have a complete definition of the equilibria. Because having more people in your group is always beneficial, the strategies would be to always allow an entrant and always deny an attempted exit. Another foundation for a baseline hypothesis could be derived from the large volume of prior experimental results on public goods games (Marwell and Ames (1979), Isaac and Walker (1988), Kim and Walker (1984), Isaac, Walker, and Williams (1994) and Andreoni and Miller (1993)). These past studies show general stylized facts that subjects contribute more than the self-interested optimum in the beginning and that the level of contributions drops over time. One might well imagine that similar results would be observed in our experiments. These prior results, however, provide us with no foundation for making predictions regarding how the different group formation institutions might impact the results.

Our hypothesis in beginning this investigation was that we would observe institution specific variation in behavior. We will not present a model that will predict such deviations in large part because our goal was to use this investigation as a means of determining if such deviations do exist in order to inform the possible development of such a model at a later date. We can imagine a number of intuition based hypotheses of the deviations we might see and these will be explained and tested as we present the results below.

The data set created by these experiments is quite rich since there are potentially 4 different types of choices made each subject in each period (i.e. contribution level, move or not, which group to move to and vote yes or no). We will begin presenting our results in a series of graphical and general statistical characterizations of the data in order to give the reader a good idea of the general structure of the data. This will include presenting results on very raw results aggregated across all periods as well as showing several different characterizations of the effects of the institutions on the paths of the experiments. We will then crystallize our results in a series of panel data regressions that can better capture the highly conditional and experience dependent nature of the choices.

3.1 Preliminary Periods Vs. Main Periods

The three preliminary periods of the experiment were run in part to familiarize subjects with the investment decisions before layering on the group choice component. We can also use the subjects' decisions in the preliminary periods as a way of examining the subjects' behavior when social emergence is not an issue. In these preliminary periods any deviations from the dominant strategy of investing 3 tokens in the group account can only be interpreted as mistakes or a lack of understanding of the rules. Figure 1 shows a histogram of these contributions. The very clear result is that 3 is the overwhelmingly modal response with 3.30 being the average[6]. Deviations from this contribution level are minimal and approximately symmetric. Of the 432 choices, 256 (60%) were 3 while 67 (15.5%) were less than 3 and 106 (24.5%) greater than 3. Further, contributions of either 2 or 4 represent very small differences in payoffs to the subject[7] and 336 (77.8%) of the choices are contained in the range of [2,4]. This is strong evidence that the subjects understood the payoffs well despite their complexity and that with no other considerations, they will play the self-interested choice as standard theory would predict. We find no significant difference across

[5] Simple strategies such as "always join group A" or "always join group B" would be examples of strategies that could be part of such an equilibrium but one could of course derive more complex contingent strategies as well.

[6] While we have pooled the choices over all three preliminary periods, one might reasonably ask if there were any observed differences across the three periods perhaps indicating that it took subjects a period or two to learn the optimal contribution level. The average contribution into the group account for the three periods were 3.22, 3.31 and 3.36, indicating that most subjects figured it out from the first choice.

[7] Investing 2 tokens yields 9.20 ECUs, investing 3 yields 9.50 ECUs and investing 4 yields 9.13 ECUs.

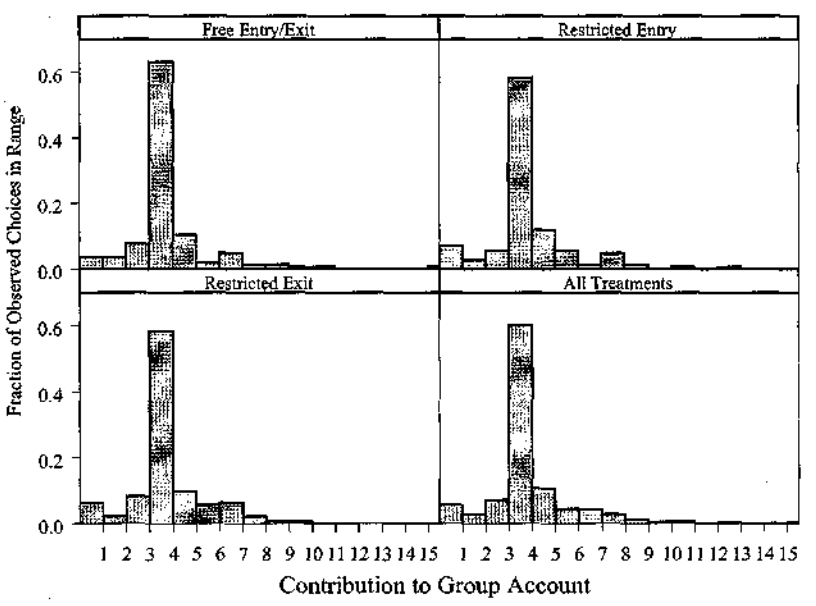

Figure 1: Contributions to group account in preliminary periods.

treatments which is not surprising since the preliminary periods were played in exactly the same manner regardless of treatment. The average contributions are 3.34, 3.33, and 3.24 across the Free Entry/Exit, Restricted Entry and Restricted Exit treatments respectively. None of the paired comparisons are statistically significant in t-tests or in Wilcoxon rank-sum tests. The lack of a difference simply shows that there were no substantial differences in the subject groups across treatments in regard to their ability to understand and find the dominant strategy.

We can then compare the choices in the preliminary periods with those in the main periods. Figure 2 shows the distributions of choices in the main periods for each treatment compared with the combined results from all preliminary periods. In all three treatments we see that the modal choice remains 3 but the percentage of choices less than and equal to 3 goes down as more mass is added to higher contribution choices. Specifically, of the 2,880 choices made in the twenty main periods, 1,071 (37.2%) were 3, 194 (6.74%) were below 3, and 1,615 (56.1%) were above 3.

The overall average contribution to the group account across all sessions in the main periods is 4.53 tokens. When broken down by treatment, the averages are 4.34, 4.85 and 4.41 in the Free Entry/Exit, Restricted Entry and Restricted Exit treatments respectively. The average contribution in the Restricted Entry treatment is significantly higher, in two-sample Wilcoxon rank-sum tests, than that in Free Entry/Exit (p=0.0000) and that in Restricted Exit (p=0.0003), but the contribution levels in these other two treatments are not statistically different. Also notice that in the Restricted Entry treatment, contributions in the range 4 to 7 are about equally frequent and there may be a slight increase in contributions of 7 over 4, while in the other two treatments, the frequency of choices are monotonically decreasing after 3.

Another interesting comparison between the preliminary and main phase involves comparing the contributions in the preliminary phase with those in period 1 of the main phase. This is

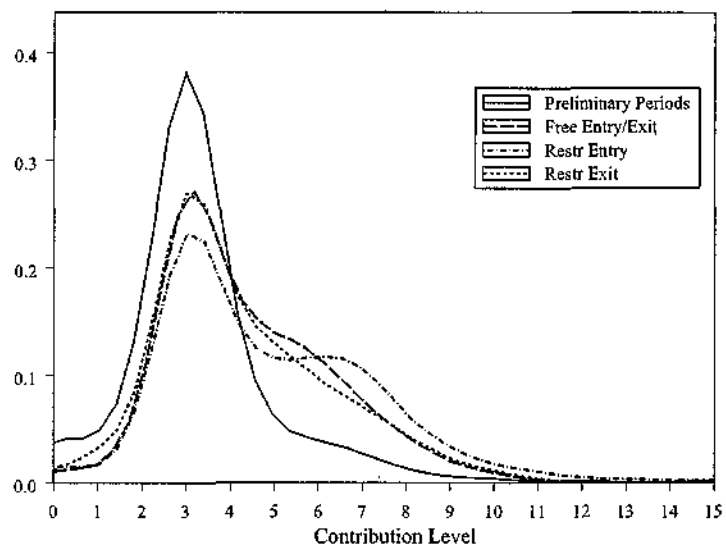

Figure 2: Density graphs of investment choices of all treatments compared to combined data of preliminary periods.

interesting because in both cases subjects are in groups of size 1 but in period 1 of the main phase, subjects can now signal their potential willingness to contribute. A paired Wilcoxon Rank sum test regarding whether or not the difference between the last choice in the preliminary phase and the first choice in the main phase is equal to 0 reveals that there were 80 observations of the same choice, 37 observations in which the subjects chose higher in the main phase and 27 subjects who chose lower in the main phase. This results in a z-score of -1.457 and a *p*—value of 0.15. Thus there is little evidence of substantial amounts of signalling in period 1 of the main phase. Anecdotally it appeared that a few subjects would signal (perhaps unintentionally) in each experiment and this provided a focal point for others to see what groups to join. This seems sensible because if a signal is necessary for coordination then it would be inefficient for all subjects to signal.

We cannot draw many conclusions from this level of analysis regarding the impacts of the institutions but this analysis is enough to conclude that (i) subjects knew the self-interested choice and chose it overwhelmingly when there were no social efficiency concerns, (ii) when social efficiency was added as a possible consideration their contribution choices increased and (iii) at least at a raw level, overall contribution levels were higher in the Restricted Entry treatment than in the other two treatments. The rest of the analysis will continue examining the results to separate out the institutional effects in finer detail.

3.2 Trends Over Time: Group Size, Contribution, and Earnings

The first additional dimension along which to examine the investment choices is how they change over time. Figure 3 shows trends over time of contributions, group size, contributions as a percent of the group optimal levels and earnings. The average group size is computed by summing over subjects and dividing by 12 (the data is the size of the group the subject is in) rather than summing

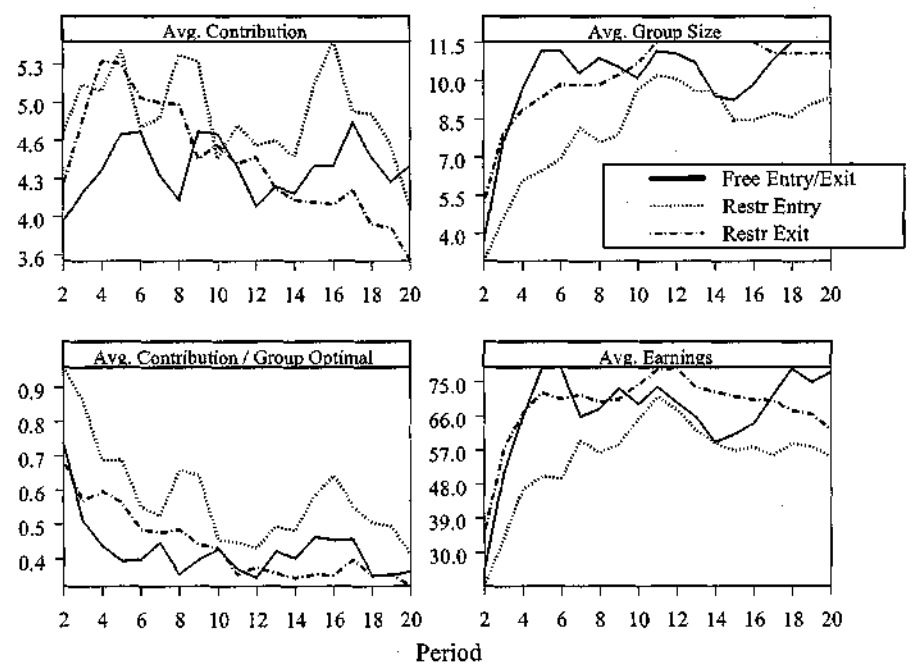

Figure 3: Four figures showing the trends over time of key variables.

over groups with a positive number of members and dividing by the number of such groups. The problem with the latter is that, while perhaps more natural, it distorts what we wish to represent. A division of 12 subjects into two groups of 6 each, and two groups of sizes 1 and 11 would return the same average group size of 6 yet for our purposes these two divisions of subjects to groups, are quite different. By summing over subjects, a division of 6 and 6 returns an average group size 6, while a division of 11 and 1 returns an average group size 10.17. Regardless of how we calculate the group size, though, the results will show that the average group size is the smallest in the Restricted Entry treatment. In fact in all periods, the groups tend to be smallest in the Restricted Entry treatment. Subjects were quite successful in forming large groups in the other two treatments. The average group sizes are 10.10, 8.01, and 10.28, for the Free Entry/Exit, Restricted Entry and Restricted Exit treatments respectively. The difference is highly significant between Free Entry/Exit and Restricted Entry and between Restricted Entry and Restricted Exit but not significant between Free Entry/Exit and Restricted Exit.[8]

Examining the two figures on contribution levels, we see that the raw contributions in the restricted entry treatment are a bit higher than in the other two treatments. Since the size of the groups in that treatment are lower, though, these higher contributions represent a much higher contribution level as a percent of the group optimal levels. The overall averages of contributions as a percent of the group optimal are 0.46, 0.62 and 0.47 respectively[9]. Another interesting pattern

[8]The p-values for Wilcoxon tests for the respective comparisons are $p<0.001$, $p<0.001$ and $p=0.558$.

[9]The p-values of Wilcoxon rank sum tests comparing differences among treatments in contributions as a percent of group optimum are as follows: Free Entry/Exit to Restricted Entry $p < .001$, Restricted Entry to Restricted Exit

	Free Entry/Exit		Restricted Entry		Restricted Exit	
	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	4.116	<0.001	4.9272	<0.001	4.949	<0.001
Time Trend	0.021	0.111	-0.008	0.677	-0.051	0.012

Table 3: Results of simple OLS regressions of average contributions for each treatment regressed on the time period.

	Free Entry/Exit		Restricted Entry		Restricted Exit	
	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	4.116	<0.001	4.9272	<0.001	4.949	<0.001
Time Trend	0.021	0.111	-0.008	0.677	-0.051	0.012

Table 3: Results of simple OLS regressions of average contributions for each, treatment regressed on the time period.

in the contribution data is that in the Restricted Entry and Free Exit/Entry treatments, while the contributions vary substantially over the sessions, the time trend is virtually flat. In the Restricted Exit treatment, however, there is a negative and significant time trend. This can be confirmed by simple regressions of the average contributions as a function of the period and the results can be found in table 3 which shows that the time trend is only significant in the Restricted Exit treatment. Recall that the standard result in prior public goods experiments is a steady decrease in the level of contributions over time. We find this only in the Restricted Exit treatment.

Finally, the average earnings chart shows that the lowest earnings are in the Restricted Entry treatment. This may seem surprising since this treatment had the highest contribution levels, but it is a demonstration of the earnings power of large groups. The overall average per period earnings were 63.97, 52.42 and 65.29 respectively[10].

The trends over time can be summarized as follows: (i) the Restricted Entry treatment features the smallest average group sizes, while the group sizes in the other two treatments are not distinguishable (ii) the contribution level is the highest in the Restricted Entry treatment, while differences between the other treatments are again not significant, (iii) the earnings are the lowest in the Restricted Entry treatment with earnings in the other two treatments about the same, and (iv) a steady decline in contributions to group account is observed only in the Restricted Exit treatment.

3.3 Sample Sessions

To get a better understanding of the dynamic processes involved we will now present pictures of group level data over time. Figures 4-6 show the average investment level, number of members and the group optimal level of investment for that group size for all groups in the respective sessions (1, 2. and 8) that had positive group membership after period 2 of the main phase. The figures include the contributions made by the We have chosen to present data from one "illustrative" session of each treatment. By illustrative we do not mean "average", but rather we specifically chose these sessions because they emphasize the differences between the treatments. Our next section will subject the data to more rigorous statistical tests to determine the degree to which these effects generalize.

Figure 4 displays the path of the sample Free Entry/Exit session. In this session subjects were quickly able to form the global coalition with all subjects in the same group. The initially dominant group, 10, was made focal due to the subject who was initially in that group in period 1 of the main session contributing 7 in that period. In periods 5 through 12 all subjects were in a single group, group 10, but the contribution level was quite low and far below what would have been socially optimal. In period 13 it appears one subject decided to form a new group, group 9, and attempt to

$p = .002$, Free Entry/Exit to Restricted Exit $p = .938$.

[10] The p-values of Wilcoxon rank sum tests comparing differences in treatments in overall means are as follows: Free Entry/Exit to Restricted Entry $p < .001$, Restricted Entry to Restricted Exit $p < .001$, Free Entry/Exit to Restricted Exit $p = .718$.

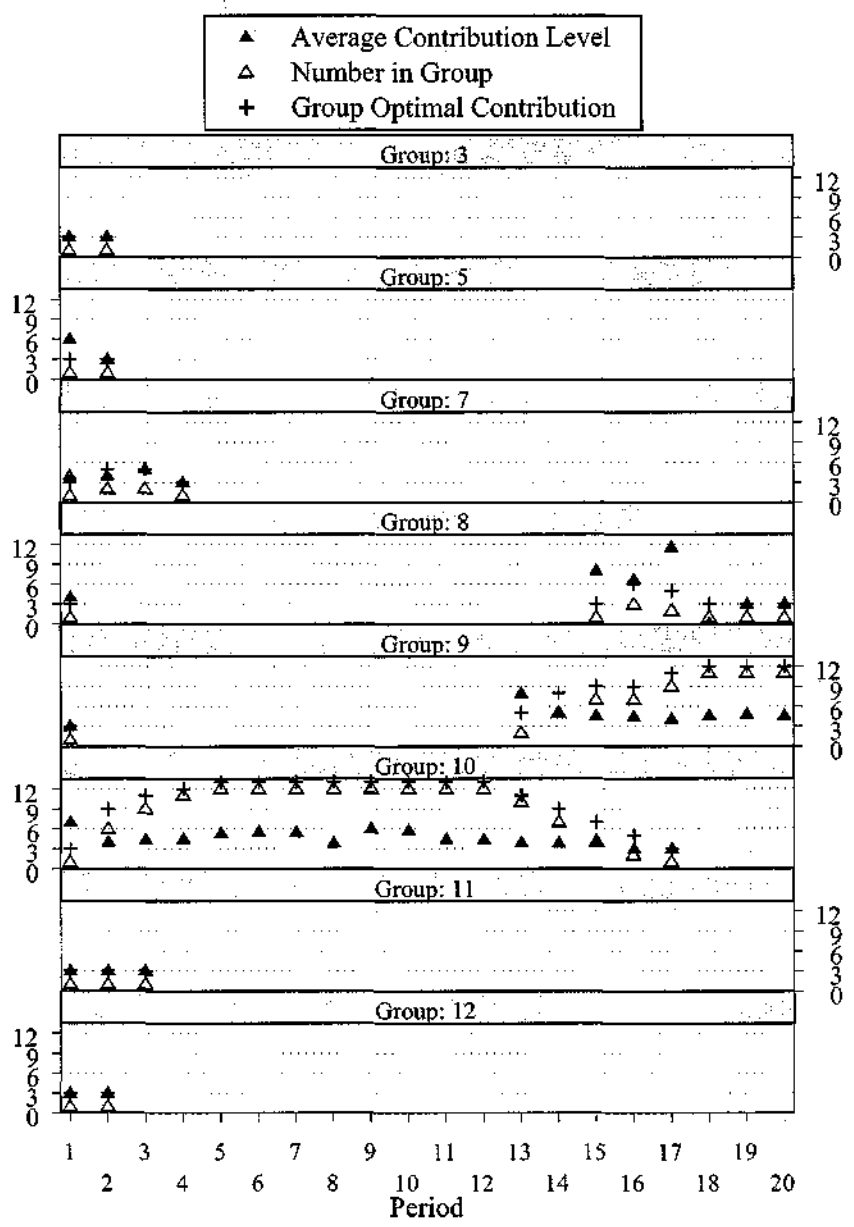

Figure 4: Sample session (session 1) from the Free Exit/Entry treatment showing the time path for all groups with positive membership after period 2.

signal that contributions should be high. Others were happy to join, but the average contribution level quickly fell as they were not as willing to also contribute at the high level signaled by the subject forming the group. Another subject tried again to form a new group in period 15 but had little success in attracting members. By period 18, virtually all subjects had switched from group 10 over to group 9. One might hypothesize that simply the ability to choose who to associate with might lead to. high contribution levels. We find little support for such a hypothesis in these results because the average contribution level is not far above the dominant strategy.

Session 2, as detailed in figure 5, shows a substantial difference in behavior due to the restricted entry mechanism. First, subjects never form into a single group. Groups 1 and 10 retain roughly equal membership throughout the session. Group 10 gets larger faster but the contribution level of the group is low. Group 1 initially had a smaller number of members but the average contribution level was higher and exactly at the group optimal level. Slowly, subjects moved over from group 10 into group 1 and, as they did so, group 1 was not only able to keep its contribution level high, but also to increase the average contribution to keep it approximately at the group optimal level as the group size increased. The reason group 1 did not grow in size faster is not due to a lack of entry attempts but rather to the members of group 1 denying entrance to applicants with low contribution histories. This point will be made clear in the regression analysis in the next subsection regarding the voting data. The intriguing part of the path, though, is that it appears that after entering into group 1, subjects keep their contributions high. One might expect that a subject would increase his contributions to gain entry into the group and then drop back to the dominant strategy level once he has gained membership in the group. The path from this session suggests that the restricted entry mechanism might be effective in teaching people to increase their contributions to gain entrance and then this teaching carries over once someone has entered the group even in the absence of any mechanism to punish "backsliding" in contributions.

The selected sample session from the Restricted Exit treatment is session 8, shown in figure 6. Most subjects initially joined group 3, but a competition between groups 1 and 3 soon developed over which group would be the focal point of the global group. The average contribution level is higher in group 1 than that in group 3 and by period 12 all twelve subjects have joined group 1. The movement of subjects from group 3 to group 1 is somewhat delayed as some of the members of group 3 voted "no" on multiple applications to exit. As subjects joined group 1, the average contribution level went rapidly down to the dominant strategy level of 3. The decline in the level of contributions can be explained by the frustration of subjects who wanted to leave the group, but whose exit was denied. Those frustrated subjects often lowered their contribution levels even below 3. In periods 1-10, when the big group was forming, there were 0 instances of subjects contributing 0 or 1 token and only 2 cases of subjects contributing 2. During periods 11-20 there were 22 attempts to exit the main group with only one being successful. During this time period there were 6 cases of a subject contributing 0, 3 incidents, of contributing 1, and 7 incidents of contributing 2. The single successful exit was by the subject who returned to group 3 (as the lone member) in period 17. The subject had contributed an average of 6.9 tokens per period to the group account between periods 1 and 10 and had attempted to leave the group in periods 11, 14, 15 and 16, unsuccessfully and then successfully in period 17. The subject lowered their contribution to 0 in period 14, 3 in period 15 and back to 0 in period 16. After forming a group of one in period 17, the subject contributes 5 for the remaining periods, perhaps as a way of telling the members of group 1, "this is what you guys should be doing." There were other similar examples of high contributors dropping their contribution levels substantially after their exit was repeatedly denied leading to a reasonable interpretation that the decline in their contributions was due in part to a feeling of frustration or perhaps a strategic move to convince the group to approve their departure.

This result is in-line with observations in other experiments showing that subjects are willing to

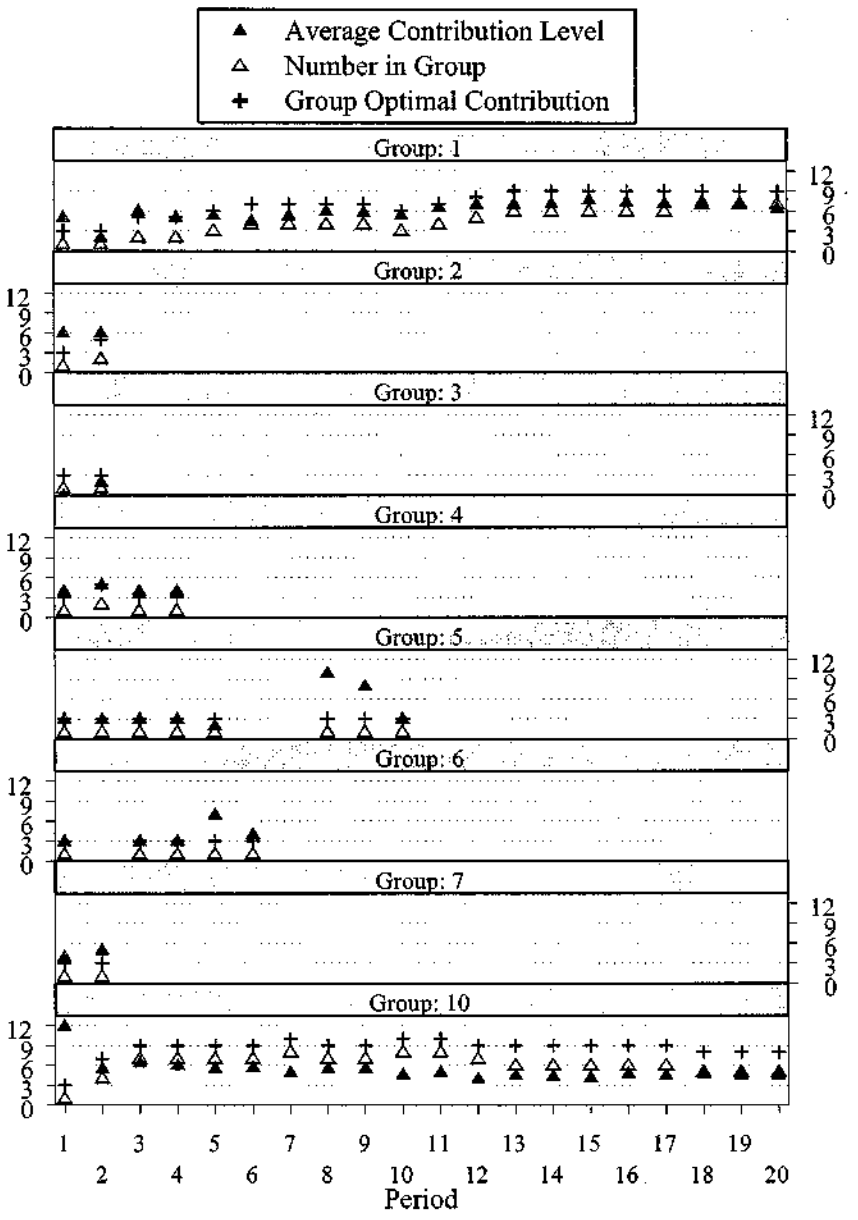

Figure 5: Sample session (session 2) from the Restricted Entry treatment showing the time path for all groups with positive membership after period 2.

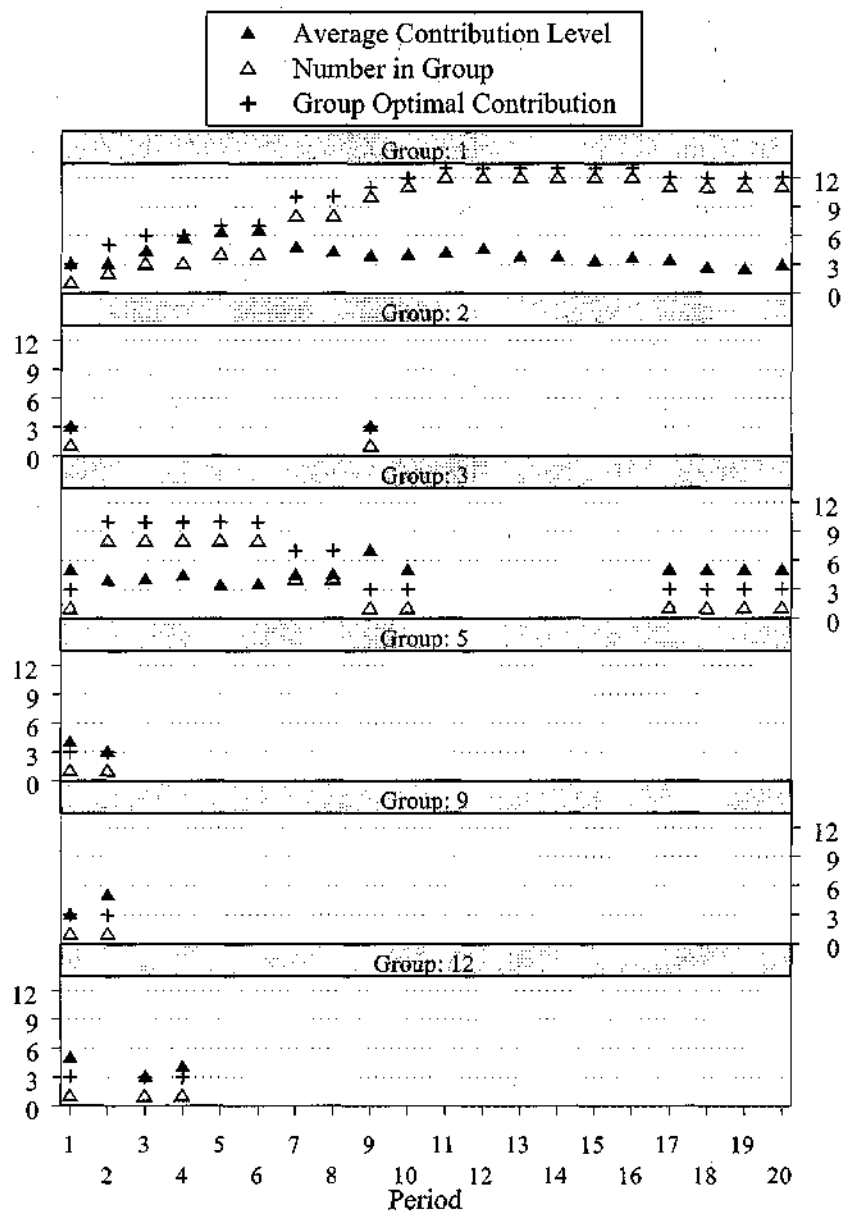

Figure 6: Sample session (session 8) from the Restricted Exit treatment showing the time path for all groups with positive membership after period 2.

Variable	Explanation
Periods in Group	Periods subject has been in their current group
New Group t-2	1 if subject joined current group two periods back, 0 else
New Group t-1	1 if subject joined current group in the prior period, 0 else
New Group t	1 if subject joined current group in current period, 0 else
New Group t+1	1 if subject moves to another group in the next period, 0 else
New Group t+2	1 if subject moves to new group in 2 pds but not in 1 pd, 0 else
Group Optimal	Group optimal contribution given size of current group
Times Move Denied	Times a subject has attempted and failed to join (leave) a group
Failed Move t	1 if subject failed to join (leave) in current period, 0 else
Failed Move t-1	1 if subject failed to join (leave) in previous period, 0 else
Period	Index for the current period (1, 2,..., 20)
Endgame	1 if period=18, 19 or 20, 0 else

Table 4: Variables and explanations for contribution regressions.

engage in costly punishment (see Ostrom, Walker, and Gardner (1992), Fehr and Gächter (2000), Ones and Putterman (2004) and Sefton, Shupp, and Walker (2002)) because contributing less than 3, in addition to decreasing the payoff to the rest of the group, actually decreases the subjects own payoff relative to what they could have achieved by contributing 3. In these other studies, punishment usually takes the form of allowing a subject to pay a certain amount to decrease the payoff of someone else. The form of punishment we observe, decreasing ones contributions even below the individually optimal level, should be quite similar to the form of punishment one might see in more natural social settings, such as in a partnership arrangements when a high contributing partner becomes frustrated with a low contributing partner and decreases his contributions in response. In addition to the punishment interpretation of these results, it is also possible to interpret some of the low contributions as a strategic choice. In the partnership example this would involve a partner decreasing his contributions in order to convince his partners that he is a low contributor so that they would allow him out of the partnership. Such behavior certainly has analogs to behavior observed outside of the laboratory and is a dimension not observable in prior experimental studies. These results suggest that this group formation mechanism should allow for a more natural means of examining both pro- and anti-social behavior than the more artificial punishment schemes previously investigated.

3.4 Individual Level Results

While the figures showing the average contribution levels and group sizes over time are suggestive of a number of possible effects, more careful data analysis is necessary to verify exactly what those effects are. We will therefore present a series of panel data regressions to provide a more rigorous characterization of the factors impacting subjects' decisions on both contributions and votes as well as how a subject's history will impact his probability of success at entering or exiting a group.

Investment Choices Table 5 shows the results from a series of fixed effects estimations of contribution to the group account as a function of several variables. The regression results listed are from a separate regression for each treatment but we have also conducted the fully interacted model with all data points pooled using dummy variables for sessions interacted with all independent variables. The interpretation is the same either way and we chose to present these results as the

	Free Entry/Exit		Restricted Entry		Restricted Exit	
	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	3.033	<0.001	5.030	<0.001	5.988	<0.001
Periods in Group	0.031	0.118	-0.044	0.133	0.027	0.432
New Group t-2	0.491	0.012	0.369	0.171	-0.117	0.600
New Group t-1	0.602	0.002	0.184	0.501	-0.562	0.013
New Group t	0.325	0.126	0.113	0.702	-0.970	<0.001
New Group t+1	-0.265	0.217	-0.052	0.863	-1.072	0.001
New Group t+2	-0.029	0.904	-0.378	0.254	-1.046	0.004
Group Optimal	0.092	0.001	-0.020	0.630	<0.001	1.000
Times Move Denied	-	-	0.460	<0.001	-0.180	0.091
Failed Move t	-	-	-0.473	0.236	-0.574	0.037
Failed Move t-1	-	-	-0.317	0.420	0.037	0.899
Period	0.002	0.877	0.015	0.549	-0.130	<0.001
Endgame	-0.195	0.262	-0.527	0.033	0.088	0.652
Num Obs (Groups)	960(48)		912(48)		912(48)	
$\sigma(\mu)$	1.466		1.536		1.127	
$\sigma(\varepsilon)$	1.482		1.976		1.529	
ρ	0.495		0.377		0.352	

Table 5: Results of three fixed effects regressions (1 for each treatment) with contribution to group account as the dependent variable. ρ is the fraction of the overall variance due to μ_i.

coefficients are easier to interpret[11]. According to the naive self-interest hypothesis, the only thing that should matter to a decision maker is the structure of the incentives leading to a choice of 3 because it is the dominant strategy in each period regardless of group size and the level of others' contribution. Thus theory suggests that we should find a lack of significance in all variables except for the constant which should be 3 with 0 variance. If we believe, however, that at least some subjects take the social efficiency into account or try to send signals to others to contribute more, then several variables might be important. The independent variables used in these regressions are explained in table 4. The specification of each regression is

$$Contribution_{i,t} = \alpha + \beta * X_{i,t} + \mu_i + \varepsilon_{i,t}$$

where i is the index across subjects, t is the index across periods, X is the matrix of regressors, μ_i represents the fixed effect for subject i and $\varepsilon_{i,t}$ is an error term.

In the Free Entry/Exit treatment, four variables have statistically significant coefficients: Constant, New Group $t-2$, New Group $t-1$, and Group Optimal. The constant is highly significant and almost exactly 3, the dominant strategy contribution level. The interpretation of the New Group $t-x$ coefficients is that subjects tend to increase their contribution level by about half a token in the first couple of periods after joining a new group. Subjects in this treatment also seem to respond to the group optimum, but not to a great degree as they raise their contributions by only one tenth of a token per one token increase in the level of the group optimum. The dominant strategy prediction therefore does a reasonable job of explaining the data from this treatment although there were some deviations toward pro-social behavior.

[11] Results of fully interacted model available from authors upon request.

In the Restricted Entry treatment, we find that the constant and two other variables have significant coefficients: Times Move Denied, and Endgame. The coefficient for Times Move Denied indicates that each time a subject's application to a new group is denied, he increases his contribution to the group account by half a token. As explained before, the voters were able to see the contribution history of an applicant for the previous five periods. This result suggests that rejected applicants were able to figure out that the best way to get into a group is to increase their contribution to the group account. One might think that such signaling behavior should lead to the New Group $t + x$ coefficients to also be significant. They, however, would only be significant if most subjects figured out the need to signal prior to being rejected. These results suggest that the signalling strategy is only learned by repeated rejection of a subject's application. Further, since Periods in Group does not have a significant coefficient the indication is that subjects do not appear to drop their contributions over time once they get into a group. This is quite important as it suggests the "teaching" works and subjects do not increase their contributions to get into a group and then immediately or even over time revert to contributing 3. The significance of the Endgame variable indicates that the cooperation does seem to unravel a bit at the end as might be expected since the subjects did know that there would be only 20 periods.

In the Restricted Exit treatment, the constant is 6 and highly significant, but all other significant variables have negative coefficients. This suggests that the base contribution level for this group was quite high, but all of the experience based regressors were pulling the level of cooperation down. In particular, in this Restricted Exit treatment we see that subjects reduce their contribution levels before joining a new group, or perhaps in order to be allowed to exit a current group, as signified by the negative and significant coefficients on the New Group $t + x$ variables. Further, each time a subject's attempt to leave his or her group is voted down by their group members, the subject reduces his or her contribution on average by one half of a token as indicated by the coefficient on Failed Move t of -.574. Since all three coefficients are significant it appears that some subjects learn even without being rejected, perhaps by observing the actions of others, that they have to cease contributing to be let out of the group while others learn that lesson through mainly their own experience in being denied exit. This result was mentioned anecdotally before when describing the subject who reduced their contribution level down to zero, far below the self-interested level, after failed attempts to leave the group. These regression results support both the frustration/punishment story as well as the strategic response story as explanations for high contributors dropping their contributions late in the sessions. We also note that even after taking all of these effects into account, Period still has a negative and significant coefficient indicating that there is a steady decline in contributions over time due to some dynamic not captured in our other variables. As suggested in figure 3 and table 3 and verified in these results, this decrease is only observed in this treatment. The overall picture is that in this Restricted Exit treatment, subjects began being highly cooperative but then learned through interactions in the group formation mechanism to be gradually less so.

Another important detail to both of these last two treatments is that even in the Restricted Entry treatment, the coefficient on the group optimal contribution level is not significant. This implies that it is the interactions governed by the group formation institutions that are driving any cooperative behavior, rather than an independent acknowledgement by most of the subjects of the social benefits of high contributions.

Voting Behavior and Outcomes The next issue we will examine is the determinants of subjects' voting decisions on the entry/exit applications of those attempting to enter/exit into/out of a group. As a reminder, subjects always do better by having additional subjects in their group

Votes					Outcomes of Applications				
	Entry		Exit			Entry		Exit	
YES	372	(62.5%)	147	(26.3%)	**SUCCESS**	70	(59.8%)	17	(20.7%)
NO	223	(37.5%)	411	(73.7%)	**FAILURE**	47	(40.2%)	65	(79.3%)
Total	595	(100%)	558	(100%)	**Total**	117	(100%)	82	(100%)

Table 6: Summary of votes and their outcomes regarding entry/exit attempts in treatments 2 and 3.

Variables	Explanations
Voter-App $t-1$	Voter's Contrib - Applicant's Contrib in pd $t-1$
Voter-App $t-2$	Voter's Contrib - Applicant's Contrib in pd $t-2$
Voter times failed	Number of times voter failed to move
Voter failed $t-1$	1 if voter failed to move in period $t-1$, 0 otherwise
Period	Index of period (1,2,...,20)
Endgame	1 if Period=18, 19, 20, 0 otherwise

Table 7: Explanation of the variables used in the regressions on voting choices.

even if those other subjects are only contributing 3 and they are not harmed, except in a relative sense, even by group members who contribute 0. Therefore, according to a standard naive self-interest model of behavior there is no reason to ever deny someone's entry into a group or allow someone to exit. Any reasonable alternative explanations for denying someone's entry would involve an attempt to teach them to engage in pro-social behavior by contributing more. Explaining why someone would allow the departure of a fellow group member might involve a desire to avoid frustrating them or simply to incur goodwill in hopes that they will eventually return and still be a high contributor.

Table 6 shows the frequencies of YES and NO votes as well as the frequencies of successful and unsuccessful attempts at entry and exit in the respective treatments. The existence of so many NO votes on entry and YES votes on exit indicates that these are definitely more than occasional errors. To uncover the structure we ran a fixed effects logit panel regression of the vote variable (l=yes, 0=no) on a set of independent variables, explained in table 7. To compare the effects of the variables in voting on entry AND exit, we keep the same set of variables for the regression of vote on entry and that of vote on exit. Table 8 shows the regression results.

The key piece of information in determining the likelihood of a voter voting yes to approve an entrant is the difference between the voters contribution level and the applicants. The clear result is that subjects are less likely to vote to admit an applicant into the group if the applicant has contributed less than the voter has. We see the exact opposite in the determination of votes to approve an exit application. In that case the positive coefficient on Voter-App t — 1 indicates that voters are more likely to approve the exit of low contributors. The other significant determinant of a subject's willingness to approve another's departure is whether or not the voter himself has been denied departure. If a voter has just failed to depart himself in the prior period, he is much more likely to vote to allow someone else's exit[12].

We can also examine the outcome of the voting decisions to solidify the results on the determinants of successfully exiting/entering a group. The frequencies of successful and unsuccessful applications for entry and exit are previously shown in figure 6. We ran a logit regression (there

[12]This variable is left out of the regression concerning entry approval votes because there are no observations of someone who was denied in one period and then voted in the next.

	Vote on Entry		Vote on Exit	
	Coeff	**p-value**	**Coeff**	**p-value**
Voter-App $t-1$	-0.257	<0.001	0.276	<0.001
Voter-App $t-2$	-0.253	<0.001	-0.064	0.187
Voter times failed	14.349	0.975	0.104	0.759
Voter failed $t-1$	-	-	1.647	0.017
Period	-0.020	0.553	0.117	0.009
Endgame	-0.319	0.485	-0.319	0.466
Num Obs (Group)	498(34)		361(32)	
Obs per Group	14.6		11.3	
$\ln \mathbf{L}$	-207.329		-145.859	

Table 8: Results of fixed effects logit regressions on votes concerning entry/exit applications.

	Entry		Exit	
	Coeff	**p-value**	**Coeff**	**p-value**
Constant	-0.281	0.857	-3.794	0.128
Applicant Contribution $t-1$	0.416	0.028	-0.210	0.278
Applicant Contribution $t-2$	0.892	<0.001	-0.192	0.303
Num of Voters in Relevant Group	-0.062	0.539	-0.017	0.909
Avg Contrib of Relevant Group $t-1$	-0.874	0.003	0.210	0.056
Avg Contrib of Relevant Group $t-2$	0.208	0.366	0.789	0.092
Period	-0.058	0.286	-0.095	0.318
Num Obs	104		82	
$\ln \mathbf{L}$	-46.725		-29.646	

Table 9: Results of a logit regression with outcome of entry/exit application as dependent variable. "Relevant" group refers to the one the applicant is attempting to enter/exit.

seems little reason to suspect individual specific effects here) with the result of application for entry (exit) as dependent variable, which takes a value of 1 if the application is successful, 0 otherwise. The independent variables should be self-explanatory. Table 9 shows the results[13]. The very clear result is that the prime determinant of whether or not an applicant is admitted into a group is their recent contribution levels. Also notice that the higher the average contribution level of the group being applied to in the previous period, the less likely is the applicant to succeed in their application to enter. The implication is that groups that developed a norm of high contribution are less likely to accept a new member. Examining the data directly reveals this quite clearly as well. There were several subjects who would apply repeatedly to get into a group and be denied until they increased their investment levels. Some were never admitted. In session 2, one subject applied to group 1, the high investment group, 7 times and was never accepted. The subject had invested 3 in every period. Another subject applied 5 times with no success, again always contributing 3, while another applied 5 times and was accepted on the fifth application after 2 periods of increasing his contributions from 3 to 7.

In regard to the outcomes of the exit applications, we see very little that is significant. We see borderline significant results suggesting that groups with higher contribution are more likely to approve an exit. These results in general suffer from a small number of observations because we only observe 17 successful applications to exit a group.

3.5 Responses to Bonus Question

The bonus question was added to the experiments after seeing the results from the first 3 sessions and observing that subjects very clearly were able to identify the dominant strategy. This lead to the question about whether the subjects also clearly understood the group optimal behavior. So in sessions 4 to 12, we added the bonus question asking subjects to tell us what the group optimal contribution level was for a group of size 5.

The correct answer is 8 and subjects were paid for answers in the range 7-9. Figure 7 shows a histogram of the answers the subjects provided to the bonus question split out by treatment. While subjects overwhelmingly figured out the correct dominant strategy in the preliminary periods, this is not the case with the social optimum on this bonus question. In part, this is due to the fact that the social optimum is significantly more difficult to calculate and is indeed a non-linear function of group size. Out of 108 subjects, only 3 subjects answered the question correctly while 29 (26.85%) provided the answer between 7 and 9, obtaining the dollar bonus. Twenty four (22.22%) of the subjects answered 3 indicating that they could not at all distinguish between the individual and social optimum. On the other hand, 16 (14.81%) of the subjects seem to have used the heuristic of "the more the better" providing the maximum answer of 15.

Due to the difficulty of the problem it seems unreasonable to assume that many would get the exact correct answer. The main issue we are interested in is if they can clearly tell that the group optimum would be more than 3. If they can also see that it is less than 15, then that is obviously better. The majority of the subjects did indeed answer between these bounds. Since we are asking this question at the end of the experiment, it is possible that the subject's answer is somehow conditioned on their experience during the experiment. To address this issue we conduct a multinomial logit regression of subjects' answers on a number of independent variables. The results are reported in table 10. In this multinomial logit regression, we divide the answers to three categories: "too low" (0 to 5), "about right" (6 to 10), and "too high" (11 to 15). The comparison

[13] Note that we have exlcuded from these regressions all cases in which a subject has applied to enter a group of size 0 (for the entry regression) and when a subject has applied to get back into the group they were in the previous round. In both cases acceptance to enter is automatic and in the second exit acceptance to exit is automatic as well.

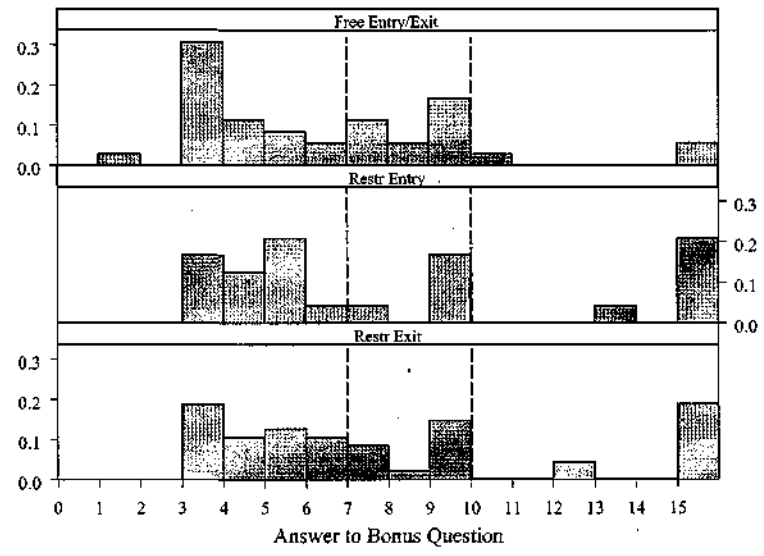

Figure 7: Histogram of answers to bonus question. The region between the dotten lines contains the answers that lead to the $1 payoff. 8 was the correct answer.

group is "two low" and the two panels show the determinants of subjects answers moving from the comparison group to "about right" or "too high." Looking at the first panel, it is apparent that a subject's own contribution level and the average contribution of the members of the groups to which the subject belonged are the most powerful predictor of the likelihood that a subject understands the social externalities of contribution. Higher own contributions, higher group contributions and subjects being in larger groups also tend to make subjects more likely to answer in this middle range. The only variable that seems correlated with answering in the too high range is the subjects own contributions. The treatment dummy variables were only borderline significant in a few cases indicating a lack of a compelling difference between treatments.

4 Conclusion

In this paper we presented an experimental examination of endogenous group formation in a public-goods provision environment. The genesis for the study was a realization that in naturally occurring collective action situations, groups use a variety of entry and exit rules and this should effect the size of groups as well as the level of the collective goods provided. The environmental and institutional features that we utilize to examine this issue are not intended to cover all the complexities of endogenous group formation in naturally occurring situations. Instead, we focused on examining the effects of a few simple mechanisms in an extreme environment that contained large incentives for subjects to just form a single large group and ignore any issues arising from the group formation mechanism. Our results show that even in this extreme environment, group formation mechanisms can have an impact on how groups form as well as on the tendency of individuals to provide the public good.

Bin	Variable	Coeff	p-value
About Right	Constant	-3.738	0.150
6≤Answer≤10	Average Contribution	0.882	0.001
	Average Group Size	-0.947	0.026
	Average Group Contribution	0.211	0.003
	Restricted Entry	-1.372	0.085
	Restricted Exit	-0.213	0.712
Too High	Constant	-4.670	0.113
11≤Answer≤15	Average Contribution	0.636	0.039
	Average Group Size	-0.484	0.283
	Average Group Contribution	0.109	0.158
	Restricted Entry	1.099	0.250
	Restricted Exit	1.499	0.081
	Num Obs	108	
	ln L	-89.010	

Table 10: Logit regression on determinants of success and failure of entry/exit. Reference bin is the "too low" bin of answers in the range 0-5.

Our first clear result is that rules governing entry and exit to/from groups have significant impacts on the size of the groups that form, on both the level and time trend of contributions and on the actual earnings of the subjects. In particular, the mechanism requiring majority approval to allow a new entrant to a group significantly increases the average contribution by allowing high contributors to teach low contributors about pro-social behavior by denying them entrance into the "good" groups until their behavior improves. Although contributions are high in this treatment, group sizes tend to be smaller because subjects were voting to deny entry to low contributors quite often. Due to the structure of the environment that involved substantial returns to adding even low contributing members, this denial of entrants and the resulting smaller groups lead to an overall decline in earnings because a small group of high contributors can not equal the earning power of a large group of low contributors. One might interpret this juxtaposition of results as indicating that subjects were displaying a willingness to give up earnings if it meant they could be in a smaller group of higher contributing members. Thus pro-social behavior in and of itself may have been valued quite highly by some of our subjects.

Our mechanism that placed restrictions on the ability of subjects to exit a group ended up with contribution levels that were quite low and in fact this was the only mechanism that lead to the gradual decline of contributions as observed in most standard public goods experiments. In prior such experiments, subjects were usually assigned to fixed groups for the duration of the experiment while in our treatment subjects could move to a new group but only if the majority of their current group members voted to approve their exit. Applications to exit were rarely approved leading to almost a fixed group treatment. Both situations then lead to a lack of mobility which the evidence suggests leads to subjects who are initially high contributors dropping their contributions over time to the individually optimal level or even below. We were able to show evidence that this occurs in our results both as a form of frustration with other group members and as a strategic means of convincing group members to allow a subject to exit the group. The latter is a particularly intriguing form of behavior not previously experimentally identified that can have substantial implications for how one might want to design group formation mechanisms outside of the laboratory. The essential

lesson seems to be that setting up rules that limit the ability for high contributors to exit leaves them only two options: stay and punish or become such a bad group member that the rest of the group is happy to be rid of them. Neither is a form of behavior that one would want to encourage.

As we mentioned, though, the results from this study are only the beginning of the process to fully understand the types and degree of effects from allowing groups to endogenously form over time. This study establishes an important benchmark result showing that the effects of these forces are so powerful that even in an environment with strong incentives to just form one group and stay in it, i.e. an environment in which the group formation mechanism should have little importance, the mechanism governing that process can still have substantial effects on the outcome. This leaves open a wide range of future possible work involving the investigation of these and other mechanisms in more "reasonable" settings to attempt to get a more accurate view of the magnitude of the effect the group formation mechanism can have in less severe environments. A first extension of this work will therefore involve a similar study in which the only change is to decrease the incentives to form the global group leading to an optimal group size of less than the number of subjects in the experiment. We hypothesize that in such an environment, the process of group formation will become even more important leading to even stronger behavioral effects. This is left to future investigation.

References

Andreoni, J. and J. H. Miller (1993). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *Economic Journal* (103), 570-585.

Coricelli, G., D. Fehr, and G. Fellner (2003). Partner Selection in Public Goods Experiments. Working Paper.

Ehrhart, K.-M. and C. Keser (1999). Mobility and Cooperation: On the Run. Working Paper.

Fehr, E. and S. Gächter (2000). Cooperation and Punishment. *American Economic Review 90*, 980-994.

Fischbacher, U. (1999). z-Tree: Zurich Toolbox For Readymade Economic Experiments. Zurich University.

Gunnthorsdottir, A., D. Houser, K. McCabe, and H. Ameden (2001). Disposition, History and Contributions in Public Goods Experiments. Working Paper Australian Graduate School of Management.

Isaac, R. M. and J. M. Walker (1988). Group Size Hypotheses of Public Goods Provision: An Experimental Examination. *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)*, 1-36.

Kim, O. and M. Walker (1984). The Free Rider Problem: Experimental Evidence. *Public Choice 43*, 3-24.

Marwell, G. and R. E. Ames (1979). Experiments on the Provision of Public Goods. I. Resources, Interest, Group Size, and the Free Rider Problem. *American Journal of Sociology 85(6)*, 1335-60.

Olson, M. (1971). *The Logic of Collective Action; Public Goods and the Theory of Groups.* Cambridge, Mass.: Harvard University Press.

Ones, U. and L. Putterman (2004). The Ecology of Collective Action: A Public Goods and Sanctions Experiment with Controlled Group Formation. Department of Economics Working Paper 2004-01, Brown University.

Ostrom, E., J. Walker, and R. Gardner (1992). Covenants with and without a Sword: Self-Governance Is Possible. *American Political Science Review 86,* 404-417.

Page, T., L. Putterman, and B. Unel (2002). Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency. Working Paper Brown University.

Sefton, M., R. Shupp, and J. M. Walker (2002). The Effect of Reward and Sanctions in Provision of Public Goods. Working Paper.

APPENDIX A: Sample Instructions

Verbal Instruction Script
Base
[ ] — text to be used when entry is restricted
\ \ - text for when exit is restricted

Thank you for participating in today's experiment. I will read through a script to explain to you the nature of today's experiment as well as how to navigate the computer interface you will be working with. I will be using this script to make sure that all sessions of this experiment receive the same information, but please feel free to ask questions as they arise. We ask that you please refrain from talking or looking at the monitors of other participants during the experiment. If you have a question or problem please raise your hand and one of us will come to you.

In the instructions that follow, you will notice that all monetary amounts, earnings, costs, etc. are denominated in Experimental Currency Units (ECUs). At the end of the experiment, your earnings in ECUs will be translated into US$ at the rate of 1 penny per ECU or 100 ECUs= $1. So if you end up with a balance of 5,000 ECUs at the end of the experiment you will be paid $50, 2,500 ECUs $25 and 1,000 ECUs $10. In addition to your earnings from your decisions over the course of the experiment, you will receive your $10 show-up fee regardless of what happens. We will make our payment to you by check at the end of the experiment.

Description of Investment Task

In today's experiment you will be engaged in a task that will be repeated many times. In each period of this experiment you will be a member of a group and will be given 15 tokens that you will be able to invest in two different "accounts": your individual account and the group account. You must invest all 15 tokens each period, but they can be split between the accounts any way you choose.

To assist you in understanding this investment task we have provided you with three tables. Please look first at the one titled "Effect of your Choices on Your Payoff." Since you have 15 tokens to invest, you have 16 different possible ways of allocating them between the two accounts. The first two columns show you what each of those choices are. The third and fourth columns show

you the earnings to you from each allocation. Each token you invest in your individual account earns .5 ECUs while each token invested in the group account yields 1.5 ECUs. Investing in your individual account is costless but investing in the group account is not. For any number of tokens x you invest in the group account you pay a cost of (1/27)*x3. We realize you may not be able to do this computation in your head so the table shows you the cost associated with any choice in the fifth column. Finally the last column shows you your net payoff which is the sum of your earnings from the two accounts less the cost of the investment into the group account.

For example, if you invest 0 tokens in the group account, all would be invested in your individual account and your total earnings would be 7.5 ECUs. If you invested 2 tokens in the group account, you would receive a net total of 9.2 ECUs with a breakdown of 6.5 from the individual account, 3 from the group account and a cost of .3. You should be able to see what the earnings would be from any other investment choice.

You will notice that in the Net Earnings column, your earnings in this example are listed as 9.2+Y. The Y reflects the fact that you receive earnings from the group account from the amounts others invest into it. For every token others invest in the group account, you receive 1.5 ECUs. If you look at the second table, Effect of Choices of Others on Your Payoffs, you will see this effect. Each row represents a different possible total number of tokens other members in your group might invest. For example, if you have invested 2 tokens and the rest of your group invested 0, your earnings are just the 9.2 we calculated above. If they have invested a total of 10 tokens, your earnings rise to 24.2 and if they have invested a total of 100, you earn 159.2.

Note that not all possible amounts of group investments have been listed. If the rest of your group had invested a total of 102, for example, you would earn something in between the amounts listed for 100 and 105. This table should give you a good idea of how your payoffs work.

Just as the choices of others affect your payoffs, your investment choices affect the payoffs of the other members of your group. For every token you invest into the group account, the other members of your group earn 1.5 ECU's each. Please look at the third table, Effects of your Choices on Rest of Group. Here you see for every possible choice you could make both what it generates to you and then what it generates to the rest of your group for the size of any group you might be in. Each group size indicated would include you. So if you were in a group of size 5 (4 people and yourself), and you invested 6 tokens in the group account, you would generate 5.5 ECUs to yourself plus a total of 36 ECUs to the other members of your group for a net total group earnings of 41.5 ECUs. Your actual earnings, of course, would be 5.5 + Y where the Y is determined by the investments of others. The 41.5 is what your choice generates to the rest of the group.

As you can see from the payoff tables, it is possible for you to lose money. If you look back at the second table, Effect of Choices of Others on your Payoffs, you will see that if you select high investments in the group account your earnings will be negative unless the rest of your group invests enough into the group account to make up for the loss. To allow for this possibility, everyone will start with an initial balance of 50 ECUs. As you make investment decisions, this balance will rise as you make money and fall if you have periods with negative earnings. If you lose so much money that your overall balance goes below 0, you will be declared bankrupt. The first time this happens to you, we will re-initialize you, starting you over with a new positive initial balance of 50 ECUs. If your earnings balance goes negative a second time, however, you will be asked to leave the experiment with only your $10.00 show up fee.

Note, these examples were not intended to be taken as suggested investment levels. We chose these examples randomly simply to illustrate how the payoffs work. While we realize these payoff calculations might seem complicated, the tables and the software you will be using should assist in seeing how they work. At this time, are there any questions about how the investment procedure works or how the payoffs are determined?

Begin Software

If you turn to your computer screens now, you will see an investment choice screen similar to what you will see in the experiment. Notice there is a single text box allowing you to enter the number of tokens you wish to invest into the Group Account. Any of your 15 tokens not invested into the Group Account for your group will automatically be invested into your Individual Account. To the right and slightly below the textbox there is a grey button you can press to "test" your choice. If you enter a number of tokens into the box and press this button, at the bottom it will show you the calculations of how many ECUs that choice generates to you and other members of your group. This information is identical to what is summarized in the tables we discussed before. Try a few test amounts now to see how" it works. Note that you can not invest more than 15 tokens per period and you can not invest negative or fractional numbers.

Once you have decided on your investment decision you can press the red button, at the bottom of the screen to accept it. For this practice choice we have placed each of you into a hypothetical group of 5 other members who are just programmed to invest random amounts. In the actual periods of the experiment the other members of your group will consist of other people in this room. Please make an investment decision now so you can see how the program works. This decision will not count towards your actual earnings for the experiment.

Once each of you has made your decision in a period, you will see a screen reporting to you the outcome. The top table reminds you which group you are in, how many people are in it as well as what your choice was. The second row tells you what your payoff was and breaks it down by earnings from your individual account, earnings from your investment to the group account and earnings from the investments of others in your group into the group account.

Below this you see the combined total number of tokens invested by the other members of your group into the group account.

Group Selection

In addition to your investments, you will also choose what group you wish to be a member of \ [subject to certain restrictions] \.

In the main experiment periods, you will make an investment choice in a group consisting only of yourself in period 1. From then on, prior to making an investment choice in a period, you will be asked to choose what group you wish to be a member of. Please press the ok button on your screen now and you will see a screen asking you if you wish to change group. These next few screens all just show you static information and are not affected by any choices you make. They have been designed to show you what the screens would look like in the real experiment.

On your screen you can see the recent history from each of 12 groups labeled A-L. For each group you can see the number of members who were in the group at the end of the prior period. You can also see the total number of tokens that members of that group have invested in the group account as well as the number of members of that group for the previous 5 periods. Note that many groups will have had 0 members over this time period. At the bottom of the screen you are asked to choose if you wish to remain in your current group or if you wish to move to a new group. If you choose to move you will go through the following process: Please check the box indicating you wish to move to a new group and press ok.

The next screen will be very similar to the one you just saw. The only difference is that you will now see each group listed with the number of members who have elected to stay in their group as opposed to the total number who were in the group at the end of the previous period. After observing this information, you can choose which group you wish to move to. You are allowed to choose any group from the list, A-L. This includes groups with no current members. You are also

allowed to choose to move back into the group you were in during any prior period including the most recent one. Those remaining in their group will see a waiting screen at this point until those moving have selected new. groups.

Voting Mechanism

[Your entry to the group is not automatic. After everyone has selected which group they wish to join, the members of the group who chose not to stay in their group from the previous period will vote on whether or not to admit you into their group. After everyone has made a group selection a sample voting screen will appear. Please make a group choice now. The current members of a group will see a choice history for all current applicants to join their group consisting of their investment to the Group Account as well as the size of the group they were in over the prior 5 periods. Those who said "yes" to the question if they wished to move do not get to vote on new members even if they chose to move back into their previous group. They will then vote yes or no on each applicant. If more than 50% of the current members vote YES (to allow entrance), that person will be accepted into the group. If 50% or fewer of the current members vote YES on an applicant, that applicant will return to his or her group from the previous period. After the voting stage has been completed, all subjects will see a new screen. Please note, if you have elected to stay in a group that no one has requested to enter, your voting screen will be blank and you just need to press the ok button to continue.]

Exit **Restriction**

\ Your proposed move to a new group may, however, not be allowed. If you have chosen to move into a new group, the members of the group you are proposing to leave vote on whether or not to approve your departure from the group. After everyone has chosen which group they would like to move to, a screen will appear to all those group members who answered "no" to the question regarding whether they wished to move. If everyone will make a group selection now, a sample screen will appear. Please make a group choice now. The current members of a group will see a choice history for all of the individuals choosing to leave the group consisting of their investment to the Group Account as well as the size of the group they were in over the prior 5 periods. Those who said "yes" to the question if they wished to move do not get to vote on approving departures even if they chose to move back into their previous group. Those that are eligible to vote will then vote yes or no on each applicant. If more than 50% of the current members vote YES (to approve departure) that person will be allowed to exit the group and move into their chosen group. If 50% or fewer of the current members vote YES on an applicant, that applicant will not be allowed to leave the group and will remain in the group. After the voting stage has been completed, all subjects will see a new screen. Please note, if you have elected to stay in a group that no one has requested to exit, your voting screen will be blank and you just need to press the ok button to continue./

The next screen will tell you what group you are now in and the number of other subjects in the group. Once this "group selection" phase is complete, you will then go on to make investment choices just as we discussed before.

Are there any questions about how the group choice procedure will work?

To allow you to familiarize yourself with the investment phase, we will begin with three periods in which everyone will make an investment choice in a group consisting only of themselves. These choices will never be observable by other participants, but they will generate actual earnings. We will then begin the main experiment periods as explained before. In these preliminary periods you will start with a 50 ECU balance and then when we begin the second series, we will restart you with a 50 ECU balance. Your earnings from the preliminary periods will be included in your final

payoff total but you can not draw against them to cover losses in the second series. In period 1 of this series, you will again be making a choice in a group consisting only of yourself and this choice will be observable to other participants in the manner described before. From period 2 of the new series to the end of the experiment, prior to making an investment choice in a period, you will be asked to choose what group you wish to be a member of. There will be a total of 20 periods in this second series.

We ask that you follow the rules of the experiment. Anyone who violates the rules may be asked to leave the experiment with only the $10.00 show-up fee.