Journal of Public Economics 37 (1988) 291–304. North-Holland

WHY FREE RIDE?

Strategies and Learning in Public Goods Experiments

James ANDREONI*
The University of Wisconsin–Madison, Madison, WI 53706, USA

Received May 1987, revised version received August 1988

Laboratory experiments on free riding have produced mixed results. Free riding is seldom observed with single-shot games; however, it is often approximated in finitely repeated games. There are two prevailing hypotheses for why this is so: strategies and learning. This paper discusses these hypotheses and presents an experiment that examines both.

1. Introduction

The free riding hypothesis has been the subject of laboratory experiments for more than a decade. While the extent of free riding has often varied across experiments, three observations are consistently replicated. First, there is no significant evidence of free riding in single-shot games. Marwell and Ames (1981), for instance, found that subjects generally provide the public good at levels halfway between the Pareto efficient level and the free riding level. Second, when subjects play a repeated game, provision of the public good 'decays' toward the free riding level with each repetition. This decay phenomenon is observed when subjects know the length of the game for sure [Isaac, Walker and Thomas (1984), Isaac and Walker (1988)], and also when they do not [Isaac, McCue and Plott (1985), Kim and Walker (1984)]. Third, free riding is often approximated after subjects play several trials, although exact free riding is seldom realized.

These observations appear to provide mixed support for free riding. It seems clear that the free riding incentives are important – subjects consistently attain outcomes that are closer to the free riding levels than the Pareto efficient levels. On the other hand, the exact predictions of the model are seldom confirmed. The phenomenon of decay is particularly pronounced.

*Thanks to Theodore Bergstrom, Robyn Dawes, Mark Isaac, Gerald Marwell, Michael McKee, Thomas Palfrey, Hal Varian, James M. Walker and some referees for helpful comments. I am especially grateful for the advice and assistance of John H. Miller.

Repetition appears to be necessary for subjects to approach free riding behavior.

Naturally, researchers have looked for explanations of these results. The two hypotheses that are most often proposed are strategies and learning. The learning hypothesis holds that a single shot of the game is not sufficient to allow subjects to learn the incentives. Repeated play allows such learning, and hence learning could explain decay. However, this test of learning is confounded by the fact that repetition allows subjects to signal future moves to each other. This is the basis for the strategies hypothesis. In a repeated game it may be rational for subjects to develop multiperiod strategies that allow for some cooperative behavior, even after the free riding incentives are learned. If this is the case, then these strategies may be responsible for decay.

This paper discusses a laboratory experiment designed to examine the strategies and learning hypotheses directly. Section 2 describes the hypotheses in detail, and indicates how they are tested. The results of the experiment are given in section 3, with a discussion in section 4. The evidence from the experiment suggests, first, that a hypothesis of rational strategic play cannot be supported, and second, that learning may play little or no role in explaining the phenomenon of decay. Moreover, the data are consistent with other predictions based on theories of non-standard behavior, such as altruism, social norms, or bounded rationality.[1] The evidence suggests greater consideration of such non-standard behavior in both theoretical and experimental research.

2. Strategies and learning

2.1. *Theory and evidence*

The experiment reported in this paper is typical of most public goods experiments. It consists of a simple public goods game that is iterated 10 times. Every iteration operates as follows. Five subjects form a group. Each subject in the group is given a budget of 50 'tokens'. The tokens can be redeemed for cash only when they are 'invested' in either a private good (called an 'Individual Exchange') or a public good (called a 'Group Exchange'). A token in the private good earns one cent for the person who invests it. However, earnings from the public good depend on what the group as a whole invents. Each token in the public good earns one half cent for the person investing it, as well as one half cent for each other member of the group. Subjects always move simultaneously, and cannot communicate at any point in the experiment. Subjects are only told the total amount of the public good for their own group. Specific contributions of other individuals,

[1] For examples of such theories see Margolis (1982), Sugden (1984), Frank (1985), Palfrey and Rosenthal (1987), and Andreoni (1987).

and outcomes of other groups, are not known. For ease of reference, the precise details of the experiment are summarized in the appendix.[2]

With the payoffs just described, the equilibrium and efficiency conditions are easily calculated. Investing a token in the public good has a private return of one half cent, while it has a social return of 2.5 cents. Hence, it is Pareto efficient for all subjects to invest all tokens in the public good. On the other hand, since the private return from the private good exceeds the private return from the public good, the rational Nash equilibrium behavior in the single-shot game is to invest zero in the public good, i.e. to free ride. Moreover, the free riding equilibrium is unique. In fact, it is simple to verify that zero investment in the public good is a dominant strategy for each player.

The single-shot equilibrium is easily extended to a finitely repeated game. In the finitely repeated game, each round is an exact replication of the single-shot game. Subjects accumulate earnings each round, but they are not allowed to carry over earnings to succeeding rounds. As shown by Friedman (1986), there is again a unique equilibrium for this game: zero investment in the public good in every round. For both the single-shot and the finitely repeated game the same Nash prediction holds: subjects should invest zero in the public good. This will be called the *free riding* hypothesis.[3]

As already noted, free riding is seldom observed, but instead provision decays with repeated play. The learning hypothesis attempts to explain this by noting that subjects may not immediately understand the incentives of the game, but need repetition to help them learn. Once they recognize the dominant strategy, they will adopt free riding behavior. Since some learn more quickly than others, we should observe, on average, decay toward zero provision. With enough repetition, all subjects will eventually choose their optimal Nash investment. I will call this the *learning* hypothesis.

The second conjecture to explain decay is that rational subjects are playing strategically. This hypothesis is derived from the Kreps et al. (1982) discussion of the Prisoners' Dilemma under incomplete information. The free riding equilibrium rests on an assumption that all subjects believe that all other subjects will be behaving rationally. However, this information may be incomplete. In particular, subject Y may believe that his partners will possibly behave irrationally (perhaps because they have not yet learned the incentives). Then if Y free rides he will educate his partners. As a result, any initial cooperation will unravel to the (less lucrative) free riding equilibrium. Moreover, if Y believes that his partners think *he* does not understand free riding, then by free riding he would reveal himself to be rational. Again, any cooperation would unravel to free riding. Hence, even if all subjects

[2]The instructions provided to the subjects are available from the author on request.

[3]Some have termed this *strong* free riding or *pure* free riding. I will simply call it free riding, since distinctions between strong and weak free riding are not useful here.

understand free riding, they may choose a strategy of investing some in the public good to conceal the fact that they are rational. However, in the known end-period free riding is always optimal. In anticipation of the end-period (using backward induction), subjects are likely to start 'bailing-out.' Hence, it may be an incomplete information Nash equilibrium strategy to cooperate early in the game, but free ride late in the game. Stated differently, decay may be a rational strategy. This will be called the *strategies* hypothesis.[4]

Previous work sheds some light on strategies and learning. For instance, Isaac and Walker (1988, henceforth IW) found conditions under which more than 80 percent of the subjects chose the dominant strategy in the tenth and final round of a game. This suggests that subjects do learn to free ride. IW then repeated their experiment with experienced subjects, i.e. subjects who had participated in a public goods experiment in the past. However, they again observed high levels of provision early and decay with repetition. This suggests that learning alone is not responsible for decay, and provides some support for strategies. The next subsection describes a method of testing both hypotheses directly.

2.2. *Testing strategies and learning*

The experiment reported in this paper is intended to separate learning from strategic play. The design is *subtractive*: subjects participate in a repeated-play environment, but are denied the opportunity to play strategically. Without strategic play, we can isolate the learning hypothesis. Furthermore, by comparing this group to one that *can* play strategically, we can attribute the difference, if any, to strategic play.

Strategies were subtracted by putting subjects in one of two conditions. In the first condition, 20 subjects were randomly assigned (by a computer) to one of 4 groups (containing 5 subjects each). Subjects were told that they would play the game exactly 10 times, but that after each repetition the composition of their group would change in an unpredictable way. In particular, after each decision round, the computer randomly reassigned subjects. Subjects knew they would be reassigned, but were never told which 4 of the remaining 19 subjects were in their group at any time. Thus, no subject could expect to gain by playing strategically.[5] This can be called a

[4]Note that these rational strategies do not include the possibility of punishing strategies like *tit-for-tat*. This is because the game is finite. However, the equilibrium is driven by the *belief* that some subjects might actually adhere to such strategies.

[5]There is, of course, the chance that a subject's actions could eventually feed back to him. However, such feedback will come in a very unpredictable fashion. Moreover, the influence will be mitigated by the actions of all of the other players and their histories. Feedback, therefore, should not figure in the predictions. As will be seen, this is verified by the results.

'repeated single-shot' game. Since subjects in this condition only meet by chance, they will be called *Strangers*.

At the same time as the Strangers were playing, 15 different subjects assembled in an adjacent room. These subjects provided a control group. They played a standard finitely repeated game, again in groups of 5. In each repetition of the game, subjects played with the *same* group of 5 subjects. They knew the composition of their group was fixed, but they did not know which 4 of the remaining 14 subjects composed their group. To contrast with Strangers, call this group the *Partners*. Notice, Partners *can* play strategically. Aside from these two controls on group composition, both Partners and Strangers faced exactly the same game.[6]

The Partners and Strangers conditions test strategies and learning in the following way. Consider strategies first. Suppose a subject is initially investing some positive amount in the public good, but learns in round t that free riding is the single-shot dominant strategy. If she is a Partner – playing strategically – she may continue to contribute to the public good. On the other hand, if she is a Stranger, she has no incentive to continue cooperation – every game for a Stranger is, after all, an end-game. Therefore, under the strategies hypothesis, we expect that giving by Partners will be greater than giving by Strangers, especially early in the game (before the Partners begin to 'bail out'). In the tenth round, however, both Partners and Strangers are playing an end-game, hence both are predicted to free ride.

To isolate the learning hypothesis, the experiment included a 'restart'. The basic experiment just described was performed twice (using a total of 70 subjects). In the second experiment, subjects in both the Partners and the Strangers conditions were unexpectedly told, after their tenth round of play, that they would restart a new set of 10 rounds. Partners would stay in the same group, while Strangers would continue to be randomly reassigned.[7] However, play was suspended after only three additional rounds.[8] If learning is primarily responsible for decay, then both Partners and Strangers should be unaffected by the restart. If either is affected, then this would imply that learning alone cannot explain decay.

Finally, the restart may provide insights into theories of non-standard behavior. Suppose Partners are following a rule-of-thumb for participating in repeated social dilemmas. Then even if they are fully informed and understand free riding, they may deliberately give on the first round. Hence,

[6]The experimental design originally called for 40 subjects – 20 in each condition. However, only 35 subjects agreed to participate, despite attempts to over-book. Hence, 20 were randomly assigned to be Strangers, and 15 to be Partners. This does not affect the result.

[7]In particular, subjects were told that they had finished ahead of schedule, so there was just enough time remaining to complete another set. This was done to make the promise that they would not be restarted a second time appear credible.

[8]Had the budget for subjects been bigger, this would have been unnecessary. Such deceptive practices are, under less restrictive circumstances, not recommended.

Table 1
Average investment in public good per subject.[a]

	Round										
	1	2	3	4	5	6	7	8	9	10	All
Partners	24.1	22.9	21.5	18.8	18.4	16.8	12.8	11.2	13.7	5.8	16.6
Strangers	25.4	26.6	24.3	22.2	23.1	21.9	17.8	19.7	14.0	12.2	20.7
Difference	−1.3	−3.7	−2.8	−3.4	−4.7	−5.1	−5.0	−8.5	−0.3	−6.4	−4.1

[a] $n_P = 30$; $n_S = 40$.

Partners may return to the same point on the restart. For Strangers, on the other hand, the restart represents nothing new. Even if they are playing by a rule-of-thumb, they will be applying a single-shot rule. Hence, Strangers should be unaffected by the restart – even if they are not free riding.

3. Experimental results

The results of the experiment are summarized as a series of six observations. The first three deal primarily with the strategies hypothesis, while the last three deal primarily with the learning hypothesis.

3.1. The strategies hypothesis

Observation 1. Giving by Partners is less than giving by Strangers, in all 10 rounds. Moreover, the difference tends to grow as the last round approaches.

Table 1 reports the average giving per subject in the first 10 rounds. The table includes data on all subjects. Recall that the strategies hypothesis predicts that giving by Partners will be *greater* in all 10 rounds, with the gap *narrowing* as the game progresses. As can be seen, this is exactly the opposite of what occurred, hence is evidence against strategies.

To test whether the difference between Partners and Strangers is significant, assume that the gifts of Partners and Strangers represent random draws from identical probability distributions. A rejection of this hypothesis would indicate that the two populations are significantly different. The test has a chi-squared distribution with one degree of freedom.[9] For the data in table 1, $\chi^2 = 11.3$, which is significant beyond the 0.01 level. Hence, the behavior of Partners and Strangers differs significantly, but in the direction opposite of what the strategies hypothesis predicts.

[9] Pool the data on Partners and Strangers over all 10 rounds. Find the median of the combined sample. If Partners and Strangers are drawn from the same distribution, we should expect half of the observations on Partners, and half on Strangers, to be above the median. If they differ significantly on this respect, we can reject the hypothesis. A discussion of this test can be found in Hogg and Craig (1978, pp. 320–322).

Table 2
Percent of subjects free riding.[a]

	Round										
	1	2	3	4	5	6	7	8	9	10	All
Partners	16.6	13.3	20.0	23.3	33.3	30.0	40.0	40.0	40.0	70.0	34.3
Strangers	15.0	12.5	15.0	15.0	15.0	17.5	22.5	25.0	30.0	42.5	20.5
Difference	1.6	0.8	5.0	8.3	18.3	12.5	17.5	15.0	10.0	27.5	13.8

[a] $n_P = 30$; $n_S = 40$.

Observation 2. The percent of Partners choosing to free ride is greater than the percent of Strangers, in all 10 rounds. The difference is greatest in round 10.

Looking to table 2, it can be seen that Partners free ride more than Strangers in each round, with the overall difference of 34 percent for Partners and 20 percent for Strangers. Again, the difference is the reverse of prediction of the strategies hypothesis, and can be shown to be statistically significant. This is based on a binomial distribution (calling free riding a success and giving a failure). The test for comparing the means of two samples from a binomial distribution has a t-distribution.[10] For proportions in the 'All' column $t = 4.06$, which is significant beyond the 0.01 level.

This difference in free riding is especially strong in round 10. However, recall that in round 10 each group is playing an end-game and each has the same degree of experience. The prediction is identical for both, yet Strangers free ride much less often. Again, this contradicts the strategies hypothesis.

Observation 3. Giving by Partners is least in round 10, but is still above the free riding levels.

Looking again at tables 1 and 2, 70 percent of Partners are free riders in the last round, and average giving is 5.8 tokens (11.6 percent of the Pareto efficient amount). However, only two of the six partners groups reach the free riding equilibrium in round 10. Nine of 30 Partners contribute to the public good, investing an average of 19.4 tokens each (38.8 percent of their tokens). Of that nine, five had been free riders at some point prior to round 10. Hence, many Partners – including those who have shown an understanding of free riding – continue to give even when free riding is a dominant strategy.[11] Again, this opposes the strategies hypothesis.

[10] Let p_1 be the proportion of Partners free riding, and n_1 be the number of observations on Partners. Define p_2 and n_2 similarly for Strangers. Then

$$t = \frac{p_1 - p_2}{\sqrt{p_1(1-p_1)/n_1 + p_2(1-p_2)/n_2}}.$$

[11] Future experiments may want to test whether this end-period effect can be enhanced by using subjects with lots of experience. I thank Mark Isaac and James Walker for this suggestion.

Observations 1, 2 and 3 now provide direct evidence on the strategies hypothesis. When we compare subjects who can play strategically with those who cannot, we expect to observe that strategic play will produce relatively more cooperative outcomes. As was shown, this was not the case. Subjects who simply met like strangers were more cooperative – and significantly so.

3.2. The learning hypothesis

The next three observations focus on the learning hypothesis as an explanation of decay in repeated-play experiments.

Observation 4. Giving by Strangers is greater than giving by Partners in the last round.

As already noted, round 10 is an end-game for both Partners and Strangers. The incentives are thus equivalent for both, and, moreover, they have had the same opportunity for learning. Nevertheless, Strangers give significantly more, and free ride significantly less. Since the situations are exactly the same for both Partners and Strangers, the only way that this observation can be ascribed to learning is if we accept that Partners learn *faster* than Strangers. Most people would, I suspect, be unwilling to accept that such a difference in learning could be so profound. Hence, this observation seems to suggest that learning alone is not responsible for decay.

Observation 5. Strangers appear to be only temporarily affected by the restart.

Observation 6. Partners return to high levels of giving in the restart. Their choices in period 11 largely mirror their choices in round 1. The restart also seems to have a lasting effect.

Table 3 summarizes the average giving in the restart. Average giving by Strangers is 9.9 in round 10, increases to 14.5 in round 11, but is down to 5.3 by round 13. Thus, any effect of the restart appears to be mostly corrected by round 12, and totally eliminated by round 13. This suggests that the Strangers treatment did successfully subtract the strategic play – the restart was treated largely as continuation of the repeated single-shot game.

On the other hand, average giving by Partners in round 11 was only 0.2 below giving in round 1. Of those 15 Partners who restarted, 5 increased their donations, 4 reduced them, and 6 made exactly the same donation as they did in round 1. Furthermore, giving did not fall off rapidly (as it did with the Strangers), but nearly tracked the round 1–3 levels of giving.

Recall that the learning hypothesis predicts that both Partners and Strangers should be unaffected by the restart. This is borne out by Strangers,

Table 3

Average investment in the public good per subject, for subjects who restarted.[a]

	Round												
	1	2	3	4	5	6	7	8	9	10	11	12	13
Partners	19.9	21.3	19.7	18.3	17.5	11.9	10.0	7.0	13.7	5.3	19.7	19.5	16.6
Strangers	22.2	22.9	20.1	21.2	23.0	22.4	15.2	18.8	13.0	9.9	14.5	11.8	5.3
Difference	−2.3	−1.6	−0.4	−2.9	−5.5	−10.5	−5.2	−11.8	0.7	−4.6	5.2	7.7	11.3

[a] $n_P = 15$; $n_S = 20$.

who verify that they are (largely) treating each round as single-shot game. Partners, on the other hand, are strongly affected. In fact, they appear to return to their original round 1 choices. Note that Observation 6 is consistent with the prediction based on non-standard behavior, i.e. that round 1 decisions are part of a reasoned (if non-rational) rule for social dilemmas. This suggests that there may actually be very little learning about free riding, and that subjects mostly understand the incentives from the start – given the opportunity to repeat their moves, they by and large do. Hence, learning, like strategies, is unlikely to provide an explanation of decay.

3.3. Summary

The strategies hypothesis predicts that rational strategic play is responsible for excessive giving in repeated games. However, as Observations 1, 2 and 3 show, subjects who cannot play strategically actually provide more of the public good than subjects who can. This is a contradiction of the strategies hypothesis. Likewise, Observations 5 and 6 indicate that subjects continue to give to the public good, even after revealing an understanding of the free riding incentives. In addition, Observation 4 notes that in round 10 – when all subjects are playing a true end-game, and all have had the same opportunity to learn – Strangers give significantly more to the public good than Partners. It seems unlikely that learning alone can explain this difference. In summary, both the strategies and learning hypotheses are contradicted in the experiment.

4. Discussion and conclusion

This paper contends that neither strategies nor learning can be supported as explanations of decay in public goods experiments. It is natural to ask, therefore, if these hypotheses focus on the right kind of learning. For instance, it is possible that subjects have learned the single-shot dominant strategy, but have not learned the backward induction necessary to understand the equilibrium. Hence, experimenters may turn to a more general learning hypothesis that incorporates learning about both single-shot and repeated-play incentives. This point may also generalize the rational strategies hypothesis. In particular, revealing an understanding of the single-shot equilibria (free riding on round 10) does not necessarily reveal an understanding of the repeated-game structure. This introduces another form of asymmetric information. Therefore, rational players may choose to give on round 11, even after free riding on round 10.

One way to address both of these generalizations of strategies and learning is to use subjects with lots of experience. However, appealing to results from the Prisoners' Dilemma literature, it seems unlikely that even this will change the conclusions. Selten and Stoecker (1986), for instance, examined a finitely repeated Prisoners' Dilemma game where subjects played a series of supergames (of 10 periods), although with a different partner in each restart. The results revealed a restart effect similar to that observed for Partners above and, moreover, the effect did not disappear – even after 25 supergames.

These results suggest that we may need to turn to theories of non-standard behavior. This will require that we look beyond the characteristics of the equilibrium, and examine how subjects make their decisions. For instance, the difference between Partners and Strangers may be due to different decision-making processes, where each process is suitable for the game being played. One example is that subjects may get non-monetary pleasure from cooperative outcomes [Kreps et al. (1982), Stark (1985), and Palfrey and Rosenthal (1987)]. After every round, subjects may be updating prior beliefs on the likelihood of cooperation. One would expect that, for any given outcome, Partners would adjust their priors more rapidly than Strangers. Since Partners and Strangers started period 1 in almost exactly the same place, this could be consistent with what was observed.

Another alternative is that giving is consistent with social norms about participation in social dilemmas. Social norms tend to be enforced by punishing the deviants. Decay may simply represent the groups' struggles to establish a norm. When the game is restarted, it would be natural for Partners to try to establish a norm at the highest level of cooperation achieved in the previous 10 rounds – round 1 giving. Strangers, who have less enforcement power, are likely to ignore the restart.

Further explanation might also be found in what has been called Regret Theory in the economics of uncertainty [Loomis and Sugden (1982)]. For instance, if a subject discovers that she did better than expected, she is 'elated' and will likely choose the same action again. However, if she did worse than expected, she has 'regret'. She is likely to choose more cautiously next time – and probably reduce her donation. The dynamic of this game could result in decay. For example, if the elated subjects repeat their moves, but the regretful subjects cut their gifts in half, then the formerly elated subjects will end up doing worse than expected but the formerly regretful subjects will do better than expected. Hence, giving will ratchet downward after every round. Certainly there are many more possibilities.

In conclusion, the results of this paper fail to confirm that subjects adhere to our standard notions of free riding behavior. This does not suggest foreclosing attempts to confirm rational free riding behavior; however, it does suggest a broader set of alternatives. The study of free riding may benefit from experiments that allow the decision processes of agents to reveal themselves in natural and predictable ways. Focusing on testable alternatives may help us better understand both when and why free riding occurs.

Appendix: Experimental design and procedures

A.1. The single-shot game

Group size. Each group contained five individuals.

Payoff function. Each subject had a budget of 50 tokens. Tokens invested in the Private Exchange earned 1 cent for sure. Tokens invested in the Group Exchange earned 0.5 cent for each subject in the group, regardless of which subject invested it. Earnings in each round were set aside and could not be reinvested in later rounds.

Information. The payoff schedule, group size, total amount of tokens in the group, and the features of the repetition (below) were all common knowledge. After each decision round subjects were told only the total investment in group exchange and their own payoff. They did not know the identities of the other players in their group, the specific contributions of individual group members, or the outcomes of other groups in the experiment. Subjects were instructed not to converse at any time during the experiment.

Subjects. Subjects were recruited from undergraduate Introductory Economics courses at the University of Michigan.

A.2. Finitely repeated play

Repetition. The single-shot game was repeated 10 times for all subjects.

Partners. Fifteen subjects were randomly divided into 3 groups of 5 subjects. These subjects played a finitely repeated game with the other members of their group.

Strangers. Twenty subjects were randomly divided into 4 groups of 5 subjects. After each decision round the computer randomly reassigned the subjects to 4 new groups. Hence, these subjects played a repeated single-shot game.

Replication. Thirty-five subjects were required for each run of the experiment: 15 Partners and 20 Strangers. The basic experiment was run twice, using a total of 70 subjects.

Restart. After the tenth round of the second run, subjects were unexpectedly told that they would restart another set of 10 rounds. The rules of the restart would be identical to the first repeated game. In particular, Partners would be matched with the same group members, and Strangers would continue to be randomly reassigned. After 3 more rounds, however, play was stopped.

A.3. Procedures

Subjects were randomly assigned to either Partners or Strangers conditions. Each subject was given written instructions, which included a schedule of the returns from the group exchange. In addition, the instructions were read aloud to the subjects. Subjects were given a stack of 10 Investment Decision Forms, one for each round, and were asked to record their decisions on the Investment Decision Form for round 1. The forms were then collected by the experimenter and each subject's Subject Number and decision were typed into a computer. The computer calculated the payoff to each subject, and printed an Earnings Report form for every subject. The Earnings Report listed the subject's choice, total group provision of the public good, and the subject's payoff. All information about payoffs, therefore, was strictly private. This was repeated for 10 rounds. For the Strangers, the computer also randomized the groups in each round. For the restart, the subjects were given a new batch of 10 Investment Decision Forms. They were told that they had finished ahead of schedule, so that there was just enough time to play again. This was done to make the promise that the game would not restart a second time appear credible. The standard 10-period experiment was completed in about 50 minutes (including administering instructions). At the end of the experiment, subjects were paid in cash.

A.4. Instructions

The instructions given to both Partners and Strangers were identical, except for one paragraph. The paragraph was under the heading 'Your Group'. Partners received the following paragraph:

> You have been preassigned to a group of 5 participants. The composition of participants in your group will **never** change. During all 10 decision rounds you will be a member of the **same group** of 5 participants.

Strangers received the following paragraph:

> The composition of your group will be changing every decision round. After each decision round you will be **reassigned** to a **new group** of 5 participants. None of the 5 group members will ever have been members of the same group in the past. The chance that any other participant will ever be in a group with you more than one time is very small.

References

Andreoni, James, 1987, Private charity, public goods, and the crowding out hypothesis, University of Wisconsin, Social Systems Research Institute, Working paper 8706.

Frank, Robert H., 1985, The demand for unobservable and other nonpositional goods, American Economic Review 75, 101–116.

Friedman, James W., 1986, Game theory with applications to economics (Oxford University Press, New York).

Hogg, Robert V. and Allen T. Craig, 1978, Introduction to mathematical statistics, 4th ed (Macmillan, New York).

Isaac, R. Mark and James M. Walker, 1988, Group size effects in public goods provision: The voluntary contributions mechanism, Quarterly Journal of Economics 53, 179–200.

Isaac, R. Mark, Kenneth F. McCue and Charles R. Plott, 1985, Public goods provision in an experimental environment, Journal of Public Economics 26, 51–74.

Isaac, R. Mark, James M. Walker and Susan H. Thomas, 1984, Divergent evidence on free riding: An experimental examination of possible explanations, Public Choice 43, 113–149.

Kim, Oliver and Mark Walker, 1984, The free rider problem: Experimental evidence, Public Choice 43, 3–24.

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

Loomis, Graham and Robert Sugden, 1982, Regret theory: An alternative theory of rational choice under uncertainty, Economic Journal 92, 805–824.

Margolis, Howard, 1982, Selfishness, altruism and rationality (Cambridge University Press).

Marwell, Gerald and Ruth E. Ames, 1981, Economists free ride, does anyone else?, Journal of Public Economics 15, 295–310.

Palfrey, Thomas R. and Howard Rosenthal, 1987, Private incentives in social dilemmas, Working paper (Carnegie-Mellon University).

Selten, Reinhard and Rolf Stoecker, 1986, End behavior in sequences of finite Prisoner's Dilemma supergames: A learning theory approach, Journal of Economic Behavior and Organization 7, 47–70.

Stark, Oded, 1985, On private charity and altruism, Public Choice 46, 325–332.

Sugden, Robert, 1984, Reciprocity: The supply of public goods through voluntary contributions, Economic Journal 94, 772–787.