SOCIAL LEARNING AND VOLUNTARY COOPERATION AMONG LIKE-MINDED PEOPLE

Simon Gächter
University of Nottingham

Christian Thöni
University of St. Gallen

Abstract
Many people contribute to public goods but stop doing so once they experience free riding. We test the hypothesis that groups whose members know that they are composed only of "like-minded" cooperators are able to maintain a higher cooperation level than the most cooperative, randomly composed groups. Our experiments confirm this hypothesis. We also predict that groups of "like-minded" free riders do not cooperate. Yet, we find a high level of strategic cooperation that eventually collapses. Our results underscore the importance of group composition and social learning by heterogeneously motivated agents to understand the dynamics of cooperation and free riding. (JEL: C91, H41, D23, C72)

1. Introduction

Research in experimental economics and social psychology has repeatedly demonstrated that many people cooperate even in tightly controlled one-shot prisoner's dilemmas and public goods experiments, where the payoff structure entails a dominant strategy to free ride. However, an equally frequent observation is that cooperation declines to rather low levels in repeatedly played cooperation games (Ledyard 1995).

How can we explain these findings, which are puzzling from the viewpoint of rationality and the assumption of selfishness? One explanation is that people are confused to some degree and have to learn to play their dominant strategy, i.e., reduced errors explain the decay (e.g., Palfrey and Prisbrey 1997). Another explanation, which will be central to our paper, focuses less on learning how to play the game but on *social* learning: People differ in their cooperative attitude and

Acknowledgments: This paper is part of a research project "Soziale Interaktionen, Unternehmenskultur and Anreizgestaltung", financed by the Grundlagenforschungsfonds of the University of St. Gallen. We are grateful for very helpful comments by Armin Falk, Ernst Fehr, Steffen Huck, Henrik Orzen, Elke Renner, Martin Sefton, and Jean-Robert Tyran as well as by participants of the Economic Science Association Meeting in Amsterdam and the European Economic Association Meeting in Madrid.

E-mail addresses: Gächter: simon.gaechter@nottingham.ac.uk; Thöni: christian.thoeni@unisg.ch

Journal of the European Economic Association April–May 2005 3(2–3):303–314

learn, during repeat play, about the social behavior of others.[1] The background for this conjecture is the observation from numerous experiments that some people are "conditional cooperators" whereas others are "free riders". The conditional cooperators cooperate if sufficiently many others cooperate as well (e.g., Keser and van Winden 2000; Fischbacher, Gächter, and Fehr 2001; Croson 2002; Falk and Fischbacher 2002; Burlando and Guala 2004; Fischbacher and Gächter 2004). A sizeable minority of people is best described as selfish, because they free ride, whenever this is in their material self-interest. If, as is typical in most experiments, membership in groups is randomly determined, cooperation is very likely to be fragile in repeatedly played experiments, despite the fact that some people are willing to cooperate. The reason is that conditional cooperators, who learn that others take a free ride, are likely to reduce their own contribution, because they do not want to be the "suckers" (see Fischbacher and Gächter 2004).

This analysis suggests that cooperation is bound to be fragile if an agent's social learning about other group members is based on observing their cooperation decisions resulting from a mixture of motivations that are unknown to the agent. By contrast, if conditional cooperators, for instance, would know that the other group members are as well like-minded conditional cooperators, then social learning would be confined to observing cooperative behavior. "Team reasoning" (e.g., Sugden 1993) and subsequent cooperation should be easy if the team players know that they are among like-minded team players. In this case social learning should sustain cooperation and prevent free riding. Likewise, if free-rider types would know that they are among other free riders, free riding should be paramount.

This paper presents experimental evidence on the conjecture that cooperation among like-minded people is substantially different from cooperation in randomly composed groups. To this end, and as we will explain in detail in the next section, we first determine a subject's type and then sort subjects into homogeneous groups of similar types. Subjects are then informed that they will play ten rounds of the public goods game with the same group members who are of their type. We compare the sorted like-minded groups to randomly composed control groups to determine the impact of knowing to be among like-minded group members.[2]

1. In our context, social learning refers to learning about the behavior (or type) of others, whereas in other contexts, e.g., information cascades, social learning means learning from others.

2. Thus, our experiments are related to studies that have also investigated group composition effects. Recent examples comprise Hayashi and Yamagishi (1998), Ehrhart and Keser (1999), Hauk and Nagel (2001), Page, Putterman, and Unel (forthcoming), Coricelli, Fehr, and Fellner (2004), and Riedl and Ule (2004). These studies are mainly interested in endogenous partner selection. By contrast, in our study, group members are exogenously matched. In this respect, our paper is related to Ockenfels and Weimann (1999), Gunnthorsdottir et al. (2001), and Ones and Putterman (2004). However, these papers differ both in their research questions and a number of design details from our study. The paper closest to ours is Burlando and Guala (2004).

To get a further yardstick about the effectiveness of being among like-minded group members, we also conduct "sorted" and "random" experiments in which group members have a punishment option. Punishment is a mechanism that can sustain very high cooperation levels (e.g., Fehr and Gächter 2000; Sefton, Shupp, and Walker 2002; Masclet et al. 2003; Falk, Fehr, and Fischbacher 2004; Carpenter forthcoming; Page, Putterman, and Unel forthcoming).

Consistent with our conjectures, we find that, when among other cooperators, cooperation-minded people can sustain almost efficient cooperation, even in the absence of a punishment option. Contrary to our prediction, we also find that free riders manage to cooperate strategically at nonnegligible levels. Yet, in the absence of a punishment opportunity cooperation among free riders collapses entirely by the final period.

Overall our results suggest that social learning rather than learning about the game drives cooperation and free riding. Because people are heterogeneous with respect to their cooperative attitudes, the exact dynamics entailed by social learning depends on group composition.

2. Procedures

Our goal is to study how like-minded people, i.e., people who know that they share a similar attitude to the cooperation problem, actually cooperate. A suitable instrument to determine like-mindedness should be simple and credibly reveal true preferences. For our purposes, we use a one-shot linear public goods game as the measurement instrument for cooperative attitudes. This game is simple and has the advantage that its payoff structure gives players a dominant strategy to free ride, i.e., to contribute nothing, if they only care about their monetary income.[3] We use the actual contribution level as a measure of the strength of cooperative attitude. The one-shot nature of the game makes a contribution choice unbiased (i.e., nonstrategic) and therefore a credible revelation of true cooperation preferences.[4]

3. To ensure that subjects understand their incentives we administered a set of eight control questions that tested the subjects' understanding of the payoff function. Subjects had to solve all questions successfully before the experiment could start. This was made public knowledge. In our experiments all subjects were able to solve all questions in due time. Before making a binding contribution decision, subjects could also use a "what-if-calculator" that allowed them to calculate their payoffs by inserting combinations of own and others' contributions. Thus, we can safely assume that subjects' contributions in this initial game are well-considered.

4. Ockenfels and Weimann (1999) report an experiment in which they sort people into "cooperative" and "less cooperative" groups. They do this on the basis of observing these people's contribution over ten rounds of a repeated public goods game. They do not get any effect of this sorting on cooperation levels. Since we do get strong effects as we will show below, this suggests that a repeated game and/or still rather heterogeneous "cooperative" and "less cooperative" groups may indeed dilute signals about cooperation preferences.

The details were as follows.[5] The one-shot public goods game was conducted among randomly generated groups of three people. We call this first experiment the *Ranking* experiment. All subjects were endowed with 20 ECU (experimental currency unit). Each subject i decided independently how many ECU (between 0 and 20) to contribute to a linear public good. The contributions of the whole group were summed up and subject i's payoff was

$$\pi_i = 20 - g_i + 0.6 \sum_{j=1}^{3} g_j.$$

This is a standard linear public good where the marginal per capita return is 0.6 (and the social marginal return is 1.8). Therefore, payoffs give money-maximizing subjects a dominant strategy to free ride, i.e., to choose $g_i = 0$, whereas efficiency would require $g_i = 20$, by all i.

Subjects played this *Ranking* experiment just once. They were informed about this both in the instructions and by public announcement. After all participants had chosen their contribution the *Ranking* treatment ended. Subjects received *no* information about the decisions of the other group members and about their earnings at this time.

Subjects then received new instructions for the main experiment, which consisted of a ten-period repeated public goods game with constant group memberships and the same parameters as in the *Ranking* experiment. Subjects were informed about the main experiment only at this stage of the experiment. (Before they started the *Ranking* experiment, they were told that some other part of the experiment would follow, but were not given further details.) This was necessary to ensure that the *Ranking* experiment measures subjects' cooperation preferences as accurately as possible. If subjects knew about the details of the main experiment and that their contribution in the *Ranking* experiment would influence the regrouping procedure then their choice in the *Ranking* experiment could have been strategically biased and would not have been a credible indication of cooperation preferences.[6]

In the instructions as well as in public announcements subjects were informed that groups would be rearranged as follows: All participants in a session were

5. Instructions are available upon request.

6. It is an interesting issue for further research to investigate how knowledge about the sorting procedures would influence contribution choices in the *Ranking* experiment and in the subsequent finitely repeated experiment. Similarly, a further interesting treatment would be to sort people as in our experiment but not to tell them about this sorting mechanism. This would allow teasing apart the effects of being sorted and knowing about it. The results by Gunnthorsdottir et al. (2001) and Burlando and Guala (2004) suggest that even resorting of which subjects are unaware can increase cooperation relative to random matching.

ranked according to their contribution to the project in the *Ranking* experiment. The first group consisted of the three subjects that had chosen the highest contributions in the *Ranking* experiment. The subjects with the fourth-to sixth-highest contribution constituted a second group and so on. The last group consisted of the three subjects who had chosen the smallest contribution in the *Ranking* experiment.

Subjects were then informed about their *new* group members' contributions in the *Ranking* experiment. After the rearrangement of groups and the information about what the new group members contributed in the *Ranking* experiment, the main experiment started. We call this main experiment the *Sorted* experiment.

In order to identify the effect of the sorting mechanism we conducted two kinds of control treatments: the *Ranking-Unsorted* and *Simple* experiments. In the *Ranking-Unsorted* control experiments, subjects played the *Ranking* experiment but this had no effect on the regrouping procedure. The new groups for the main treatment were unsorted, i.e., formed randomly. As in the main treatment subjects received the information about their new group mates' contributions in the *Ranking* experiment. The control experiment *Simple* consisted merely of the ten-period public goods game of randomly composed groups, i.e., there was no *Ranking* experiment. We found no significant differences between the contributions in the *Ranking-Unsorted* and the *Simple* experiments. This also holds for the experiments with punishment described below (both $p > 0.43$, Mann–Whitney tests, independent groups as observations). For the analysis we will therefore pool the data of the *Ranking-Unsorted* and *Simple* treatment and call these observations *Random.* In the following we will refer to the experiments with no punishment as *Sorted N* and *Random N*, respectively.

The experiments in which punishment was available (called *Sorted P* and *Random P*) had exactly the same structure as the *Sorted N* and *Random N* experiments described above. In the *Sorted P* experiments, subjects, after they were sorted into their new groups as a function of their contribution in the *Ranking* experiment, learned both that the public goods game would be played repeatedly with the same new group members and that a punishment option was available at the second stage. The *Ranking* experiment was exactly identical to the previous one (i.e., it involved no punishment). In the public goods game with punishment subjects at the first stage made simultaneous contributions to the public good as in the *Sorted N* and *Random N* experiments, respectively. They were then informed about their group member's individual contributions to the public good and could assign costly punishment points to each group member. One punishment point assigned cost the punishing subject 1 ECU and reduced the punished group member's income by 3 ECU. Each group member could assign up to ten punishment points to each other group member.

In total 231 subjects participated in the ten sessions of our experiments (54 in *Sorted N*; 51 in *Random N;* 72 in *Sorted P* and 54 in *Random P*). The experimental subjects were first-semester undergraduate students from the University of St. Gallen majoring in economics, business, law, or international relations. The experiment was programmed and conducted in z-Tree (Fischbacher 1999). It lasted about 1.5 hours and the subjects earned on average CHF 46 (about €30).

3. Hypothesis and Results

For expositional ease and the data analysis we will divide the newly formed groups in the *Sorted* experiments into three classes, each containing a third of the observations. The third of the groups with the highest average contribution in the *Ranking* treatment is called the class of TOP cooperators. The groups in the middle and lowers third are called MIDDLE and LOW cooperator groups, respectively. In the *Random* experiments, where there is no sorting, we classify *Random* groups ex post according to their average contributions over all ten periods. We classify them into the top, middle, and least cooperative third of groups. They will serve as comparison classes for the *Sorted* experiments.

We are now ready to formulate our hypothesis, which follow from the mounting evidence, mentioned in the introduction, that there is a large degree of heterogeneity in cooperator attitudes. We formulate our hypothesis for the LOW and TOP cooperator groups. For them, being among like-minded group members should matter the most. MIDDLE cooperator groups that consist of people with intermediate degrees of cooperativeness, should behave like the middle third of cooperative groups in *Random N* and *Random P*, respectively.

Hypotheses 1 and 2 formulate our expectations about cooperation levels in the experiments with no punishment (*Sorted N* and *Random N*).

HYPOTHESIS 1. *The average contribution of TOP cooperator groups in the Sorted N experiments is higher than the average contribution of the top cooperative third of groups in the Random N experiments.*

HYPOTHESIS 2. *The average contribution of LOW cooperator groups in the Sorted N experiment is lower than the average contribution of the least cooperative third of groups in the Random N experiments.*

Our next set of hypothesis concerns the experiments with punishment.

HYPOTHESIS 3. *TOP cooperator groups do not need punishment to achieve high cooperation. Therefore, the average contribution of TOP cooperator groups in* *the Sorted P experiments is the same as in the Sorted N experiments. Since we predict no free riding among TOP cooperator groups, they do not punish.*

HYPOTHESIS 4. *The average contribution of LOW cooperator groups is the same as in the Sorted N experiment. Since punishment is costly, LOW cooperator groups do not punish.*

Figure 1 contains our main results. We start with the *Random N* and *Sorted N* experiments. Panel A shows the average contribution (dashed line) in the *Random N* experiments, as well as the mean cooperation levels of the top, middle, and lowest third of these randomly composed groups. The average contribution is relatively stable until Period 8 and then shows the typical endgame effect known from many finitely repeated public goods experiments (e.g., Keser and van Winden 2000). In the final period, contributions are not significantly different between classes (Kruskal–Wallis test, $p = 0.412$).

Panel B depicts the contributions in the *Sorted N* experiments. Period 0 indicates the *Ranking* experiment. We find that contributions in the *Ranking* experiment vary over the whole strategy space. Subjects who were later on sorted into TOP contributor groups contributed on average 18.1 ECU in the *Ranking* experiment; MIDDLE contributors invested 10.1 and LOW contributors 0.8 ECU.

A comparison of average contributions in *Random N* and *Sorted N* shows that, overall, sorting people led to a substantial increase in contributions (see also Burlando and Guala 2004 who report a similar result). In *Random N* average contributions were 9.5 ECU, Whereas in *Sorted N* they amounted to 13.9 ECU. The difference is significant ($p = 0.012$, Mann-Whitney test with group average as observations).

We now test Hypotheses 1 and 2. We find unambiguous support for Hypothesis 1. TOP cooperators, when playing together, contributed significantly more than the most cooperative third of groups in the *Random N* experiments (Mann-Whitney-test, $p = 0.024$). With 4.3 ECU, the difference in cooperation levels is quite substantial (14.1 vs. 18.4 ECU). This comparison is interesting, because the cooperation level by the top third of groups in *Random N* is the upper bound of cooperation that one can expect in randomly composed groups. Thus, when "like-minded" cooperators are sorted together, one can expect a substantially higher and more stable cooperation level than in the best case of randomly composed groups.

Quite to our surprise—and contrary to our Hypothesis 2—we find substantial contributions among the free riders who comprise the LOW contributor group. Although they contributed much less than the groups in the other classes, we find even higher contributions among the sorted LOW contributor groups than among the least cooperative third of groups in the *Random N* experiments. The difference is borderline insignificant (Mann–Whitney test, $p = 0.109$).

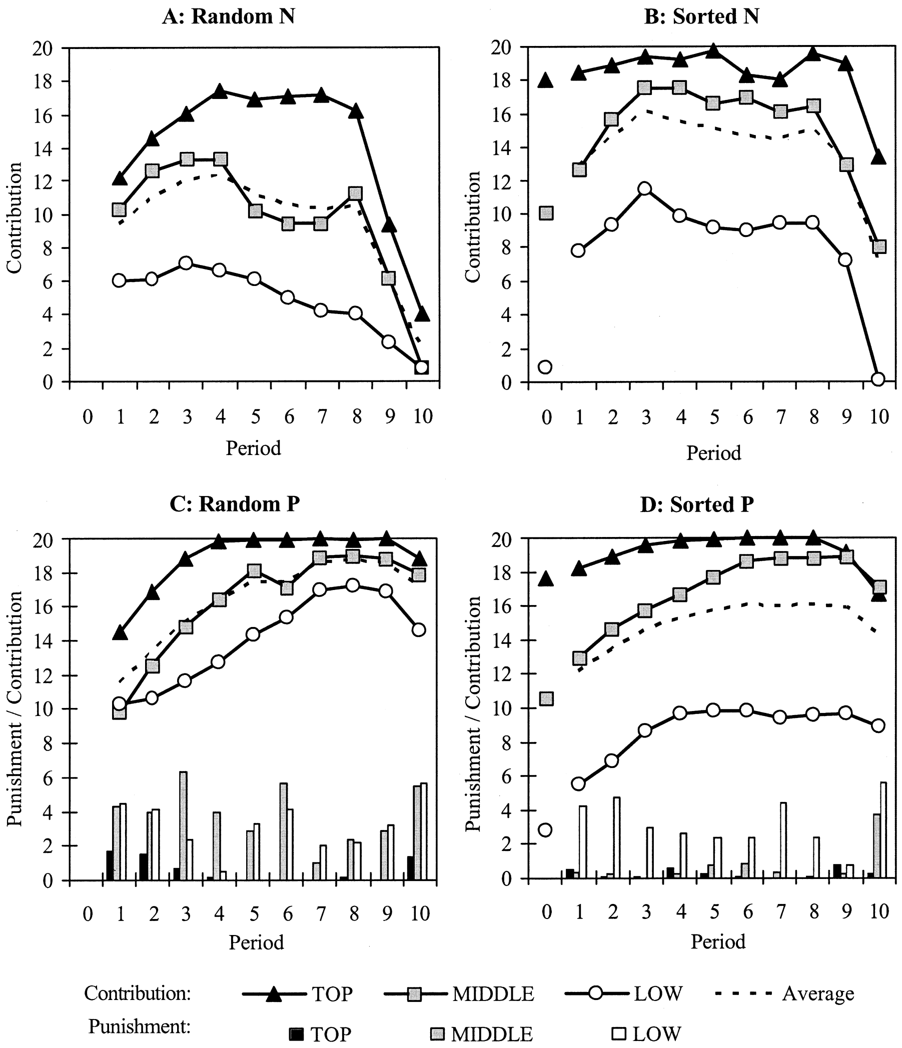

FIGURE 1. Social learning and cooperation in randomly composed groups and groups of like-minded subjects. Lines depict mean contribution levels. Bars denote mean income reduction due to punishment.

We offer two (speculative) explanations for this observation. A first explanation is derived from the bounded rationality of subjects who do not backward induct but know that earning money requires cooperation (Selten and Stoecker 1986). LOW contributors have revealed to each other that they chose the money-maximizing strategy in the *Ranking* experiment. They may therefore believe that there are no cooperators around to free ride on. Thus, they understand that they need to cooperate among themselves if they want to earn money. The second explanation rests on the possibility that LOW contributors actually believe that some other LOW contributors invested nothing in the *Ranking* experiment not because they are free riders, but because they are conditional cooperators with pessimistic beliefs. Then LOW contributors have an incentive to cooperate strategically until the ninth period to induce the conditional cooperators to contribute. They free ride in the final period, when cooperation is not in their rational self-interest anymore. Thus, if for whatever reason LOW contributors believe that some others are conditional cooperators, then rational cooperation is possible even in a finitely repeated cooperation game (Kreps et al. 1982).

There is an endgame effect in all classes, but it is most pronounced among the LOW contributors. In the TOP cooperator groups we find some people who lower their contributions in the final period. Yet, the median contribution is still 20. Overall, we find—in contrast to the *Random N* experiments—that final period contributions differs highly significantly between classes (Kruskal–Wallis test, $p = 0.001$, group averages in Period 10 as independent observations).

We also find that MIDDLE contributor groups contributed substantially more than the middle third of groups in *Random N* (15.0 vs. 9.7 ECU on average). The difference is significant (Mann–Whitney test, $p = 0.022$). A speculative explanation is that many people, who in principle are prepared to cooperate, are hesitant (because they fear free riders) and first want to test waters before they cooperate. Once they are playing together and learn that there are no strong free riders among them because they have been sorted out, MIDDLE cooperators quickly lose their hesitation and cooperate until the final rounds.

Since MIDDLE cooperator groups contributed on average the same as the middle third of groups in the first period of *Random N*, one may argue that the higher level of cooperation among MIDDLE groups in *Sorted N* is even stronger evidence for a like-minded effect than the one observed among TOP cooperator groups. Through the sorting mechanism TOP cooperator groups already started out at a higher cooperation level than the top third of groups in the *Random N* experiments. Our results show that grouping them together allowed the like-minded TOP cooperators to maintain their high cooperation level. However, like-minded MIDDLE cooperator groups when even able to increase their contributions, probably because they expected no strong free riders among them. A similar observation holds for LOW contributor groups who started out at a similar level as the lowest third of groups in *Random N.* Apparently, knowing that they are among like-minded free riders allowed them to maintain a higher level of strategic cooperation than the lowest third of groups in *Random N* who did not have such clear-cut evidence that they are composed of free rider types.

Our next results concern the impact of punishment. We first look at the *Random P* experiment (panel C) and compare it with the *Random N* experiments (panel A). Consistent with previous experiments from finitely repeated public goods experiments with and without punishment we find that average contributions are substantially higher in the presence of the punishment option than in their absence. With punishment, average contributions amounted to 16.4 (ECU; in the absence of punishment contributions were 9.5 ECU $p = 0.000$, Mann–Whitney test). Contributions were also much more homogeneous in the presence of punishment. The bars in panel C denote the average income reduction due to punishment in a particular period. We find that the highest cooperating groups punished only initially and in the final period. Groups in the other classes punished roughly equally and throughout all periods. In the top and middle third of groups punishment was almost exclusively targeted at free riders. In the lowest third of groups there was some punishment of free riders, in particular of strong ones, but also some nonnegligible punishment of cooperators (see also Falk, Fehr, and Fischbacher 2004).

Before we look at our Hypothesis 3 and 4 note that the contributions in the *Ranking* experiment (Period 0) of the *Sorted P* in panel D were almost identical to those of the *Sorted N* experiment in panel B $p = 0.819$, Mann-Whitney test).

Hypothesis 3, which predicts that TOP cooperators' contributions are the same in *Sorted P* than in *Sorted N*, is confirmed (group average contributions are not significantly different; $p = 0.516$, Mann-Whitney test). As predicted, we also find that TOP cooperators did not punish. They punished even less frequently (and less strongly) than the top third of groups in *Random P* ($p = 0.095$, Mann-Whitney test). The presence of the punishment option did not affect TOP contributors because they did not need it to achieve cooperation.

It is also interesting to compare the highest cooperating groups in *Random P* (panel C) and TOP cooperators in *Sorted N.* The cooperation by the highest contributor class in *Random P* indicates the highest cooperation level that one can get in randomly composed groups. The cooperation level by TOP cooperators in *Sorted N* indicates the highest cooperation level that is achievable by like-minded cooperators. We find that TOP cooperators in *Sorted N*, who had no punishment option available, contributed the same as the most successful randomly composed groups who had a punishment option at their disposal. Since TOP cooperator groups in *Sorted N* also had no punishment costs, in terms of efficiency they did even better than the most successfully cooperating groups who could punish misbehavior.

Hypothesis 4, which predicts that (i) LOW cooperators contribute the same in *Sorted N* and *Sorted P* but (ii) do not punish, is partially confirmed. We find support for (i) but have to reject (ii). LOW contributors punished by far the most. Given the observation from the *Random N* experiments that LOW contributors apparently cooperate strategically, it is not too surprising that LOW contributors punished to induce further cooperation. Most punishment was targeted at the free riders, but, as in *Random P*, there was also some punishment of cooperators.

LOW contributors punished about the same as the least cooperative third of groups in *Random P.* Despite this, LOW contributors contributed less in *Sorted P* than the least cooperative third of groups in *Random P* (8.8 ECU vs. 14.1 ECU on average; $p = 0.053$, Mann-Whitney test). The time trend of cooperation was also different. Punishment among like-minded LOW contributors in *Sorted P* only stabilized cooperation. Yet, it strongly increased contributions in the lowest third of groups in *Random P.* Like-minded free riders seem not to be too impressed by punishment inflicted on them by other like-minded free riders.

4. Concluding Remarks

The results of this paper are hard to reconcile with an error-hypothesis but are consistent with social learning by heterogeneous types. The reason is that an error hypothesis would not easily predict that group composition effects matter for cooperation behavior. Yet, this is exactly what we find. Since people are heterogeneous with respect to their attitudes to cooperation, our results suggest that the dynamics of cooperation as produced by social learning will depend very strongly on the extent to which group members are 'like-minded' (see also Ones and Putterman 2004, who aptly talk about an "ecology of collective action"). Our results also confirm that social norms of cooperation are quite easy to sustain in homogeneous groups of people who are aware that others share their attitudes. Like-minded cooperators do not need punishment to uphold cooperation. It is only in heterogeneous groups where punishment is helpful in sustaining cooperation.

We believe that our results are not only of theoretical interest but may also shed light on some management practices that emphasize team spirit. For instance, Ghemawat (1995) describes the group incentive schemes of a large U.S. steel producer. To prevent free riding, management at this company prefers recruiting "farm boys" (p. 697) who share a similar set of values that is thought to hold free riding at bay. The morale is that successful teamwork not only requires the right mix of complementary abilities, but also "team players", who do not shirk their responsibilities even if they safely could.

References

Burlando, Roberto and Francesco Guala (2004). "Heterogeneous Agents in Public Goods Experiments." Working paper, University of Exeter.

Carpenter, Jeffrey (forthcoming). "Punishing Free-Riders: How Group Size Affects Mutual Monitoring and the Provision of Public Goods." *Games and Economic Behavior.*

Coricelli, Giorgio, Dietmar Fehr, and Gerlinde Fellner (2004). "Partner Selection in Public Goods Experiments." *Journal of Conflict Resolution*, 48, 356–378.

Croson, Rachel (2002). "Theories of Altruism and Reciprocity" Evidence from Linear Public Goods Games," Working paper, University of Pennsylvania.

Ehrhart, Karl-Martin and Claudia Keser (1999). "Mobility and Cooperation: On the Run." Working paper 99s-24, CIRANO Montreal.
Falk, Armin, Ernst Fehr, and Urs Fischbacher (2004). "Driving Forces Behind Informal Sanctions." IEW Working Paper No. 59, University of Zurich.
Falk, Armin and Urs Fischbacher (2002). "Crime' in the Lab—Detecting Social Interaction." *European Economic Review*, 46, 859–869.
Fehr, Ernst and Simon Gächter (2000). "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90, 980–994.
Fischbacher, Urs (1999). "z-Tree, Toolbox for Readymade Economic Experiments: Experimenter's Manual." IEW Working Paper No. 21, University of Zurich.
Fischbacher, Urs and Simon Gächter (2004). "Heterogeneous Social Preferences and the Dynamics of Free Riding in Public Goods." Working Paper, University of St. Gallen.
Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001). "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." *Economics Letters*, 71, 397–404.
Ghemawat, Pankaj (1995). "Competitive Advantage and Internal Organization: Nucor Revisited." *Journal of Economics & Management Strategy*, 3, 685–717.
Gunnthorsdottir, Anna, Daniel Houser, Kevin McCabe, and Holly Ameden (2001). "Dispositions, History and Contributions in Public Goods Experiments." Working paper 99-02, University of Arizona.
Hauk, Esther and Rosemarie Nagel (2001). "Choice of Partners in Multiple Two-Person Prisoner's Dilemma Games: An Experimental Study," *Journal of Conflict Resolution*, 45, 777–793.
Hayashi, Nahoko and Toshio Yamagishi (1998.). "Selective Play: Choosing Partners in an Uncertain World." *Personality and Social Psychology Review*, 2, 276–289.
Keser, Claudia and Frans van Winden (2000). "Conditional Cooperation and Voluntary Contributions to Public Goods." *Scandinavian Journal of Economics* 102, 23–29.
Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, 245–252.
Ledyard, John (1995). "Public Goods: A Survey of Experimental Research." In *Handbook of Experimental Economics*, edited by John Kagel and Alvin E. Roth. Princeton University Press.
Masclet, David, Charles Noussair, Steven Tucker, and Marie-Claire Villeval (2003). "Monetary and Non-Monetary Punishment in the Voluntary Contributions Mechanism." *American Economic Review*, 93, 366–380.
Ockenfels, Axel and Joachim Weimann (1999). "Types and Patterns: An Experimental East-West-German Comparison of Cooperation and Solidarity." *Journal of Public Economics*, 71, 275–287.
Ones, Umut and Louis 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.
Page, Talbot, Louis Putterman, and Bulent Unel (forthcoming.) "Voluntary Association in Public Goods Experiments: Reciprocity, Mimiery and Efficiency." *Economic Journal.*
Palfrey, Thomas R. and Jeffrey E. Prisbrey (1997). "Anomalous Behavior in Linear Public Good Experiments: How Much and Why?" *American Economic Review*, 87, 829–846.
Riedl, Arno and Aljaz Ule (2004). "Exclusion, Cooperation, and Social Structure in Network Experiments." Working paper, CREED, University of Amsterdam.
Sefton, Martin, Robert Shupp, and James Walker (2002). "The Effect of Rewards and Sanctions in Provision of Public Goods." CeDEx Working Paper 2002-2, University of Nottingham.
Selten, Reinhard and Rolf Stoecker (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames. A Learning Theory Approach." *Journal of Economic Behavior & Organization*, 7, 47–70.
Sugden, Robert (1993). "Thinking as a Team: Toward an Explanation of Nonselfish Behavior." *Social Philosophy and Policy*, 10, 69–89.