

Indirect Reciprocity and Strategic
Reputation Building in an Experimental
Helping Game

Dirk Engelmann
Urs Fischbacher

CERGE-EI

Charles University
Center for Economic Research and Graduate Education
Academy of Sciences of the Czech Republic
Economics Institute

Indirect Reciprocity and Strategic Reputation Building in an Experimental Helping Game*

Dirk Engelmann[†] and Urs Fischbacher[‡]

July 15, 2003

Abstract

We study indirect reciprocity and strategic reputation building in an experimental helping game. At any time only half of the subjects can build a reputation. This allows us to study both pure indirect reciprocity that is not contaminated by strategic reputation building and the impact of incentives for strategic reputation building on the helping rate. We find that while pure indirect reciprocity appears to be important, the helping choice seems to be influenced at least as much by strategic considerations. Strategic do better than non-strategic players and non-reciprocal do better than reciprocal players, casting doubt on previously proposed evolutionary explanations for indirect reciprocity.

Studujeme nepřímou reciprocitu a strategické budování reputace v experimentální pomocné hře. V libovolném momentě může pouze polovina subjektů budovat reputaci. Tento předpoklad nám umožňuje studovat jak čistou nepřímou reciprocitu, která není kontaminována strategickým budováním reputace, tak účinek stimulů podporujících strategické budování reputace na míru pomoci. Zjišťujeme, že přestože se čistá nepřímá reciprocita ukazuje být důležitá, rozhodnutí pomáhat je ovlivňováno minimálně stejnou měrou strategickým chováním. Strategičtí hráči mají lepší výsledky než nestrategičtí a nereciproční hráči mají lepší výsledky než reciproční, což vzbuzuje pochyby o vysvětlení nepřímé reciprocit navrhované dříve.

JEL Classification: C92

Keywords: indirect reciprocity, reputation, experimental economics

*We thank James Cox, Alan Durell, Ernst Fehr, Simon Gächter, Werner Güth, Steffen Huck, Michael Kosfeld, Manfred Milinski, Wieland Müller, Ronald Oaxaca, Andreas Ortmann, Arno Riedl, Rupert Sausgruber, Dirk Semmann, Georg Weizsäcker, and Viatcheslav Vinogradov for helpful comments and suggestions. Urs Fischbacher acknowledges support from the Swiss National Science Foundation (project number 1214-05100.97), the Network on the Evolution of Preferences and Social Norms of the MacArthur Foundation, and the EU-TMR Research Network ENDEAR (FMRX-CTP98-0238). Dirk Engelmann acknowledges Financial Support by the *Deutsche Forschungsgemeinschaft* (DFG, grant No. EN 459/1-1) and the CERGE-EI foundation (through an ESC-postdoctoral fellowship).

[†]CERGE-EI (Charles University, Prague, and Academy of Sciences of the Czech Republic), P.O.Box 882, Politických veznu 7, CZ-11121 Prague 1, Czech Republic, dirk.engelmann@cerge-ei.cz.

[‡]University of Zurich, Institute for Empirical Research in Economics, Bluemlisalpstr. 10, 8006 Zurich, Switzerland, fiba@iew.unizh.ch.

1 Introduction

Among the recent approaches to conceive a more realistic model of human behavior by extending economic theory by aspects that go beyond narrow self-interest, reciprocity has been prominent, both in theoretical (e.g. Rabin, 1993, Falk and Fischbacher, 1999) and experimental work (e.g. Berg, Dickhaut, and McCabe, 1995, Fehr, Kirchsteiger, and Riedl, 1993). It is still an open question whether reciprocal actions, i.e. rewarding of friendly acts and punishment of hostile acts, result from a motivation for reciprocity (as in Dufwenberg and Kirchsteiger, 2000) or rather from other-regarding preferences (as in the models by Bolton and Ockenfels, 2000, and Fehr and Schmidt, 1999). Direct evidence that a motive for reciprocity plays a role, in particular for negative reciprocal behavior, is provided, for example, by Cox (2002), Blount (1995) and Charness (2002).

The focus of the literature has so far been almost exclusively on direct reciprocity, where a person that is affected by the choice of another person, can directly reward or punish the latter. In contrast, the focus of our experimental study is indirect reciprocity, where friendly or hostile acts of one person towards another are rewarded or punished by a third party.¹ The term “indirect reciprocity” was introduced by Alexander (1987). He argues that individuals’ behavior towards others is not only influenced by their own experience but also by their observations of other’s behavior. According to Alexander, indirect reciprocity works through reputation and status and provides the evolutionary basis for moral systems prescribing cooperation. Indeed, in large societies repeated direct interaction is quite unlikely and hence there is little scope for direct reciprocity. Incentives for adherence to a social norm can thus only be provided by indirect reciprocity based on reputation.

The models by Levine (1998) and by Charness and Rabin (2002) can also be interpreted as models of both direct and indirect reciprocity. In the model by Levine (1998) the intensity of empathy towards another person depends on the altruism of that person. This altruism can, for example, be inferred from the behavior towards a third person, implying indirect reciprocity. In Charness and Rabin (2002) people are assumed to exhibit “concern withdrawal”, i.e. they assign a lower (or even negative) weight to another

¹In an alternative model of indirect reciprocity, a person that has been helped then helps a third party. Boyd and Richerson (1989) study this form of indirect reciprocity in small cyclical networks.

person's payoff in the own utility function if that person has not behaved well, which could, for example, been by the mistreatment of a third person.

Harbaugh (1998) argues (and provides supporting field data) that donations to charity are in part driven by a prestige motive. The benefit of such prestige might come from indirect reciprocity, because donations might be rewarded by outside observers. Indeed, Milinski et al. (2002b) have shown that in an experiment donations to UNICEF are rewarded by other players. This allows for important effects of strategic behavior. If at least some others are indirectly reciprocal, it can pay to strategically build a reputation for being generous. This in turn allows a potential recipient of a donation to increase the revenues by publishing donations. The latter should increase future returns of the donor on his contribution and hence his incentive to give. Apparently charities are aware of the prestige motive (which might be driven by expectations of indirect reciprocity) since it is common practice to announce donors' names and contributions.² The interaction of indirect reciprocity and strategic reputation building can thus have substantial impact on economically relevant interaction.

Looking for explanations for the existence of indirect reciprocity, Nowak and Sigmund (1998a) have conducted simulations of an evolutionary process based on a repeated so-called helping game (a degenerate game where a donor can help a recipient at a cost smaller than the recipient's benefit). The previous behavior of the recipient is reflected in a so-called image score, that the donor can observe. This game is nicely suited to study indirect reciprocity because it precludes (in anonymous sufficiently large groups) any effects of direct reciprocity as opposed to games such as the prisoner's dilemma. They find that maximally discriminating players will eventually take over the population.³ Leimar and Hammerstein (2001), however,

²Andreoni and Petrie (2001) provide further experimental evidence on the prestige motive. They show that subjects, when having the options to contribute both to an anonymous and a broadcast public good, overwhelmingly choose the latter. Beyond indirect reciprocity, reputation building can be crucial for the functioning of markets with repeated one-shot interactions. This is increasingly relevant in markets that are becoming larger and more anonymous and hence less prone to be influenced by direct reciprocity as is exemplified by e-commerce. Bolton et al. (2002) and Keser (2002) provide experimental evidence on the importance of reputation mechanisms in environments with repeated one-shot interactions.

³For a helping game where the image score is based only on the last decision, Nowak and Sigmund (1998b) have shown analytically that discriminators outperform altruists and defectors but that cycles occur due to invasion of discriminator populations by altruists, which in turn yields an advantage to defectors.

show that this result is based on a too restricted initial set of available strategies. They show that subjects who are not indirectly reciprocal but only help in order to keep their own score at a level that induces a high probability of being helped (and hence base their decision only on their own score), could invade and take over a population of image scorers (i.e. players who base their choice only on the recipient's score).⁴ Hence, such strategic reputation building could undermine a cooperative society based on indirect reciprocity.

To assess the interplay of indirect reciprocity and strategic reputation building, and hence the plausibility of the proposed evolutionary explanations for cooperative behavior, it is necessary to clearly distinguish between helping choices based on indirect reciprocity and those based on strategic reputation building. An experimental design that achieves this aim thus has to allow us to identify whether observed helping choices are motivated by strategic reputation building or by indirect reciprocity that is not contaminated by incentives for strategic reputation building by the donor, which we call pure indirect reciprocity.

Our experiment was designed for this purpose. We investigate a helping game where in any period only half of the players have a public image score, the record of their previous behavior. In particular, each subject has a public image score either in the first 40 periods of the experiment or in the last 40 periods. This allows us to identify the motivation behind the helping choices. First, since donors without a public image score interact with recipients with a public image score, we can study pure indirect reciprocity uncontaminated by incentives for strategic reputation building. Second, by comparing the behavior of donors with and without public scores, we can evaluate the relative impact of strategic reputation building on the helping rates. Our design allows us to study the latter issue even within subjects.

Our design is in part based on that of Seinen and Schram (2001) who have conducted an experimental helping game similar to the one studied by Nowak and Sigmund (1998a).⁵ The fundamental difference is our introduction of players with and without score. Seinen and Schram (2001) find that indirect reciprocity is

⁴Nowak and Sigmund (1998a) also study strategies that use both the own and the recipient's score. For a case with perfect information they find that a type who is maximally discriminating with respect to the recipient's score but in addition only helps if the own score is threatening to fall below this threshold, is most frequent. Leimar and Hammerstein (2001), however, argue, using comparable simulations, that this result requires either a substantial influence of genetic drift or a very small cost of helping. Even in the latter case, cooperative image scorers are invaded by sophisticated strategic subjects who build, based on previous experience, a belief whether it pays to have a positive image score.

⁵Wedekind and Milinski (2000) provide the first experimental test of the Nowak and Sigmund model, based on only six

important, since many donors base their helping decision on the image score of the recipient. A substantial part of the donors, however, also base their decision on their own image score, indicating that strategic reputation building is a major force as well. In the main treatment observability of the recipient's score implies that the donor's score will be observable in the future as well. Therefore, the design of Seinen and Schram (2001) does not allow them to study pure indirect reciprocity uncontaminated by strategic concerns. It is also difficult to get a clear idea of the overall impact of strategic reputation building. While they find that a substantial part of subjects base their decisions at least in part on their own score, it is not possible to clearly determine the share of helping choices that are due to strategic reputation building. Hence it is unclear to what extent the invasion predicted by Leimar and Hammerstein (2001) is reflected in the data. These issues were, however, not the focus of their study as it is in ours.

We test three hypotheses. First, indirect reciprocity is important, i.e. the probability that donors help increases in the recipient's image score. Second, subjects strategically build a reputation, i.e. for any given score of the recipient (including an absent score) the average helping rate of donors with a public image score is higher than that of donors without. Third, strategic reputation building weakens the reciprocal relation, i.e. the dependence of the donor's helping rate on the recipient's score is weaker for donors with a public score than for donors without.

We find support for all three hypotheses. There is a clear positive relation between helping rates and recipients' scores for both donors with and without a public score. The latter provides evidence for indirect reciprocity even in the absence of strategic incentives for reputation building, hence for pure indirect reciprocity. Our experiment thus provides, to the best of our knowledge, the first evidence of pure indirect reciprocity in the laboratory.⁶ The average helping rate of donors with a public score is, however, more than twice the average helping rate of donors without. Hence strategic reputation building plays an important role as well. There is also clear evidence on an individual level. 80% of the subjects help clearly more

periods. They find support for indirect reciprocity in the sense that recipients who are helped have higher scores on average than recipients who are not helped. Furthermore, donors who rarely help rather do so when the recipient has a high image score.

⁶We also find substantial helping rates (32%) in interactions where neither the donor nor the recipient has a score. This suggests that motives like altruism or efficiency concerns play an important role as well.

often when they have a score than when they do not, including 25% who never help when they do not have a public score, but often help when they have a public score. The mode of donors' public scores is at 4 helping decisions out of the last 5 decisions. This score is the expected payoff maximizing score in all five sessions. In contrast, if scores are not public, the mode is clearly at 0.

A crucial result is that payoffs relative to group averages are clearly higher for strongly strategic subjects (whose helping rates are generally at least doubled when they have a score) than for the weakly strategic subjects (whose helping rates are not substantially higher when they have a score). They are also higher for the non-reciprocal than for the reciprocal subjects (though not significantly). This would be consistent with the invasion argument by Leimar and Hammerstein (2001) and casts some doubt on the evolutionary explanation for indirect reciprocity suggested by Nowak and Sigmund (1998a).

The paper proceeds as follows. Section 2 presents the helping game and the experimental design. The results are presented in Section 3, followed by a discussion of related literature in Section 4. Section 5 summarizes our results and provides concluding remarks.

2 Experimental Design

We conducted a computerized repeated helping game similar to the game studied by Nowak and Sigmund (1998a) and Seinen and Schram (2001). The experimental software was programmed in z-Tree (Fischbacher, 1999) and the experiments were run in the computer laboratory at the Institute for Empirical Research in Economics of the University of Zurich in Fall 2001. Participants were students from a variety of fields from the University of Zurich and the ETH Zurich and were recruited by telephone. They were randomly assigned to seats in the laboratory. Instructions were provided in written form and participants could read through them at their own pace (see the appendix for an English translation). Donor and recipient roles were labeled A and B in the instructions, but the helping choices were labeled as such, because we considered the game structure so obvious, that the use of the word "help" would not invoke any interpretations that subjects would otherwise not come up with. More neutral terminology would instead appear plainly artificial. At the end of the instructions there were five control questions to check that participants had understood the

key features of the experiment. The experiment started when all participants had answered all the control questions correctly and after an oral summary of the instructions had been given.

There were 16 subjects in each of our five experimental sessions. The helping game was repeated for 80 periods. In each period the subjects were randomly matched (independently between periods) in pairs and the role of donor and recipient were randomly assigned. The donor had the choice whether or not to help the recipient at a cost of 6 “Points”, which yielded a benefit of 15 Points for the recipient. The recipient had no choice to make. Of course, classical game theory based on common knowledge of narrow selfishness and rationality implies no help in the one-shot game and by backward induction for the finitely repeated game. For an infinitely repeated helping game, however, folk-theorem arguments yield many different Nash-equilibria, including one where all players choose the strategy to help if and only if the recipient has a perfect score and on the equilibrium path all players then help.

Each subject had a public score either in the first 40 periods or in the last 40 periods. All subjects were informed about this before the start of the experiment. The common knowledge of this change of roles ensured that subjects were in a symmetric position and hence it precluded that donors with and without score behaved differently because they considered themselves advantaged or disadvantaged.⁷ A public score consisted of the number of times the subject had helped and had not helped in the last 5 times as a donor. In case the subject had so far been in the role of the donor less than 5 times, the score consisted of the total number of help and not help decisions so far. When the recipient had a public score, the donor was informed about this score before making the decision to help. A subject with a public score was also informed about the own score. In case the subject did not have a public score, no score information was displayed (but of course, subjects could easily keep track of their own score). The experimental software did, however, record the hypothetical scores to allow an easy comparison of the choices with and without score. A score that is based on more than the last period allows in principle for punishments, because a player who generally helps can occasionally punish a free-rider without being punished himself if the

⁷While subjects might have considered having a score advantageous or disadvantageous, they knew that they would not be advantaged or disadvantaged over the whole course of the experiment. In a design that simply mixed subjects with and without score, behavioral differences between them might be caused by feelings of being assigned to the inferior role.

indirectly reciprocal players do not demand a perfectly clean record. It is, however, impossible with our information structure to distinguish punishment from occasional defection.

From the second period on, subjects were informed about the outcome of the last period at the same time they were asked to make a decision or were informed that they were a recipient. The upper part of the screen reviewed their role in the preceding period, the donor's decision and the resulting payoff and total payoff so far, as well as the own score if they had a public score in that half of the experiment. A donor was asked for his choice in the lower part of the screen and there he was either informed about the score of the recipient or that the recipient did not have a public score. A recipient was only informed about his role and that he did not have to make a choice. Following period 40, the roles of subjects with and without public score were switched and the scores were set to 0.

At the end of the experiment Points were converted into Swiss Francs at a rate of 1 Point = 0.1 Swiss Franc. Subjects started the experiment with an endowment of 100 Points, corresponding to a show-up fee of 10 Swiss Francs. No additional show-up fee was paid. The sessions took between 64 and 81 minutes and earnings ranged from 6.40 to 55.60 Swiss Francs with an average of 29.36 Swiss Francs (including the 10 Francs initial endowment).

3 Experimental Results

The overall experimental results are displayed in Figures 1 and 2, which show the average helping behavior of donors with and without score for different public scores of the recipients. The latter are represented as the relative number of helping choices, i.e. the number of helping choices in the last five help decisions divided by five, or if the score was based on less than five choices, the number of helping choices divided by the total number of choices (and rounded to multiples of 0.2). In Figure 1 all choices where recipients had a public score based on at least one decision are included, in Figure 2 only the decisions where recipients had a full score, i.e. a score based on five decisions are included. Average helping rates for the individual sessions by score status of donors and recipients are presented in Table 1 (for all scores of the recipients aggregated and only including recipients with a full, or full hypothetical, score). Table 1 shows in particular that helping rates are quite high (32%) even when neither the donor nor the recipient has a score, i.e. in




Figure 1: Donors' average help rate for all recipients with a public score based on at least one decision a situation where indirect reciprocity and strategic reputation building cannot play a role. This suggests that additional motives like altruism or efficiency concerns are important as well.

Figures 1 and 2 provide immediate support for the first hypothesis that donors are indirectly reciprocal and in particular for the existence of pure indirect reciprocity. The helping rate of donors both with and without public score clearly increases with the image score of the recipient, although the relation is monotone only in the case where donors do not have a score and recipients have a full score. A straightforward statistical test confirms the significance of this positive relation. The H_0 hypothesis is that donors do not condition their helping decision on the recipient's score. We conduct a simple binomial test for this hypothesis based on the five sessions as independent observations. To obtain an estimate whether there is a positive relation between the recipient's score and the helping probability we estimate for each session independently a simple linear probability model

$$Help_j = const + \alpha \cdot RsScore_j + \varepsilon_j,$$

with $Help$ a dummy that is 1 if the donor helps the recipient and $RsScore$ the recipient's score and ε_j a normally distributed error term.⁸ Under H_0 in each individual session the probability that the estimated coefficient for the recipient's score is positive is $\frac{1}{2}$ (actually slightly smaller, because of a small positive

⁸The linear probability model may not be the ideal model to analyze the dependence of the help rate on the recipient's




Figure 2: Donors' average help rate for recipients with full public score

		Session	1	2	3	4	5	Total
R with score	D with score		72%	78%	70%	66%	86%	74%
	D without score		45%	42%	22%	27%	49%	37%
R without score	D with score		72%	66%	63%	66%	75%	69%
	D without score		46%	32%	13%	21%	47%	32%

Table 1: Average help rates by score status of donors (D) and recipients (R), recipients with full score only

probability for a flat relation). Hence, since the sessions are independent, the probability for a positive coefficient in all five sessions is (slightly smaller than) $(\frac{1}{2})^5 = \frac{1}{32} < 5\%$. Since we find a positive coefficient in all five sessions we can thus reject the H_0 that there is no positive relation at the 5% level.⁹ This holds independently of whether we include all recipients with a score or only the recipients with a full score and whether we consider donors with or without score (see Table 2).

The affirmative result for the first hypothesis can also be derived from a panel data analysis (with the score. Here, however, we only use it to provide us with an input for the simple non-parametric test whether there is a positive relation between the recipient's score and the help probability. For this test it is also irrelevant if coefficients are not significantly different from 0, as long as they are positive.

⁹The same logic will apply to all our non-parametric tests below. Since all our hypotheses are directed, we can apply one-sided tests throughout.

		Session	1	2	3	4	5
All recipients with score	All donors	.56 (8.1)	.59 (6.4)	.47 (6.1)	.51 (5.2)	.57 (3.9)	
	donors with score	.55 (8.7)	.58 (9.2)	.52 (8.3)	.43 (4.4)	.34 (2.2)	
	donors without score	.53 (7.0)	.75 (10.2)	.36 (4.7)	.64 (8.9)	.66 (4.5)	
All recipients with full score	All donors	.65 (9.0)	.74 (8.8)	.51 (5.3)	.46 (3.9)	.72 (5.7)	
	donors with score	.60 (7.6)	.62 (8.3)	.44 (3.7)	.30 (1.3)	.32 (1.3)	
	donors without score	.69 (10.8)	.93 (13.7)	.49 (6.8)	.69 (10.8)	.87 (7.2)	

Table 2: Coefficients for recipient’s score in linear probability model for helping choice. Adjusted R^2 (in percent) in parantheses.

sessions as independent units of observations). The model is

$$Help_{it} = const + \alpha \cdot RsScore_{it} + \nu_i + \varepsilon_{it},$$

with, as above, $Help$ a dummy that is 1 if the donor helps the recipient and $RsScore$ the recipient’s score, ν_i a session specific effect and ε_{it} a normally distributed error term. We again run the regression separately for donors with and without score, to be able to detect whether a reciprocal relation might be restricted to one group of donors. Table 3 lists the estimates for the coefficients, the standard errors, z-statistics and p-values. No matter whether we include all recipients with a score or restrict the analysis to those with a full score, both for donors with and without scores the coefficient for the recipient’s score is positive and highly significant ($p < 0.1\%$). Obviously, the same results occur if we run the regression jointly for all donors with and without score.

Figures 1 and 2 as well as Table 1 provide clear support for the second hypothesis that donors try to strategically build a reputation. The average helping rate of donors with score is higher than of those

			Coefficient	Std. Error	z	$\Pr > z $
Donors with scores	All R with	α	0.5176	0.0669	7.732	0.000
	score	$const$	0.3671	0.0542	6.774	0.000
	R with	α	0.5075	0.0851	5.965	0.000
	full score	$const$	0.3630	0.0665	5.460	0.000
Donors without scores	All R with	α	0.5558	0.0688	8.081	0.000
	score	$const$	-0.0489	0.0672	-0.728	0.467
	R with	α	0.7190	0.0817	8.803	0.000
	full score	$const$	-0.1702	0.0688	-2.474	0.013

Table 3: Coefficients for the recipient’s score (α) in random effects cross-sectional time-series regression for the help choice

without for any score of the recipient (including absent score, as can be seen from Table 1).¹⁰ The same holds for each individual session. In case only recipients with full score are considered, there is only one tie, in session 1 for a score of 0. Such a dominance of the helping rate of donors with a score over that of donors without a score occurs in one session under the H_0 hypothesis that strategic reputation building is not relevant certainly with probability less than $\frac{1}{2}$. Thus the fact that it holds in all five sessions allows us to reject the H_0 at the 5% level.

This result as well is supported by a panel data analysis. We extend the above model to

$$Help_{it} = const + \alpha \cdot RsScore_{it} + \beta \cdot DWithScore_{it} + \nu_i + \varepsilon_{it},$$

with $DWithScore$ a dummy that takes the value 1 if the donor has a score and 0 otherwise and the other variables as above. Both running the regression for all recipients with a score or only for those with a full score, yields a highly significant ($p < 0.1\%$) coefficient for the dummy, supporting the second hypothesis (see Table 4).

¹⁰If the analysis is restricted to recipients with full score, then the helping rate of both donors with and without score is 0 for recipients with a score 0. There are, however, only 13 interactions with a recipient with a full score of 0. 12 of these are with the same subject and hence all in session 1. Furthermore, since the helping rate for donors without score is already 0 for recipients with a score of 0.2, this tie appears to simply result from censoring.

		Coefficient	Std. Error	z	$\Pr > z $
All R	α	0.5319	0.0485	10.963	0.000
with	β	0.3885	0.0224	17.323	0.000
score	<i>const</i>	-0.0326	0.0461	-0.709	0.478
R with	α	0.6487	0.0590	10.992	0.000
full	β	0.3779	0.0256	14.752	0.000
score	<i>const</i>	-0.1210	0.0474	-2.550	0.011

Table 4: Coefficients for the recipient’s score (α) and for a dummy whether the donor has a score (β) in random effects cross-sectional time-series regression for the help choice

As can be seen in Table 1, in each session for both recipients with and without score the helping rates of donors with score is about twice the helping rate of donors without score. Hence the impact of strategic reputation building is not only statistically significant, but also of substantial magnitude. On the other hand, both for donors with and without score, the average helping rate is only slightly (about 5%) lower if the recipient does not have a score than when he does. Hence recipients without a score are treated roughly as if they had an average score. This indicates that donors do not create any self-serving beliefs that recipients about whom they do not receive information are not helpful, which in turn could justify not helping these recipients.

The importance of strategic reputation building is also very vividly illustrated by Figure 3 which shows the distribution (absolute frequencies on top of bars) of donors’ full (public or private) scores.¹¹ The dark bars show the distribution for donors without public score and the light bars for donors with public scores. In the former case, the mode is with about 40 % at a score of 0, with almost a uniform distribution over the remaining scores. For public scores, in contrast, the mode of the distribution is at a relative score of

¹¹We included all full scores following the donor’s decision, except for scores resulting from a donor’s decision in the last period, because in that case the resulting score could not possibly be relevant for future interaction. The total number of the included scores is 2480, 1227 where the score is public and 1253 where the score is private (the difference is a result of the random allocation of donor and recipient roles, apparently it just happened that players with a private score were chosen slightly more often as donors). Since participants could be a donor (or a recipient) several periods in a row, some of these scores may have never (or several times) been observed.

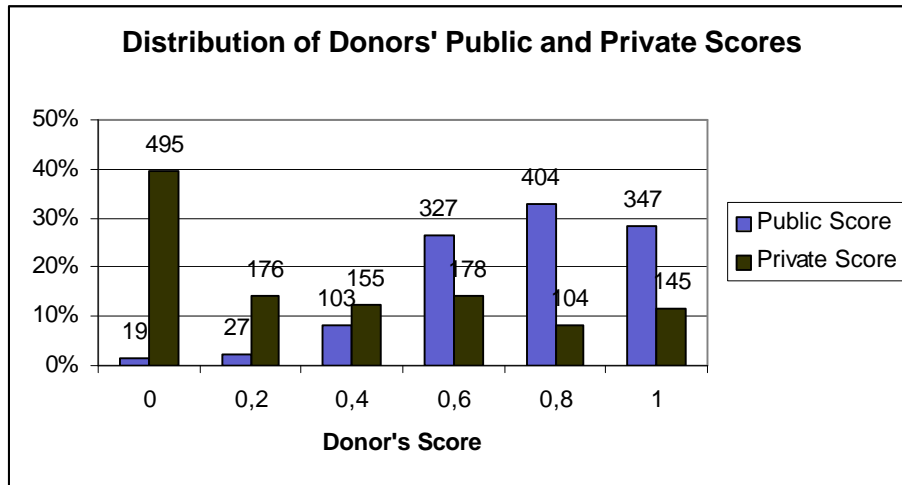


Figure 3: Distribution of public and private (post-decision) donors' scores for all interactions where the donor had a full score (except for following donors' decisions in the last period). Absolute numbers appear on the top of the bars.

0.8 (i.e. 4), with few cases of scores below 0.6 and hardly any below 0.4.¹² Interestingly, in all sessions the score that maximizes expected payoffs for the observed helping rates is 0.8 (see Table 5). This implies that either many subjects manage to strategically optimize their score, or that many subjects use the most prevalent score as a cut-off point. As Table 5 shows, it appears in general optimal to keep the score at or close to 0.8 (this analysis does not take time trends into account, the optimal score is calculated based on average help rates over the whole phase with full scores). The high returns on a score of 0.2 in session 2 and on 0.4 in session 5 can be attributed to low numbers of observations (1 and 3, respectively).

An advantage of our design is that we can also study the importance of strategic reputation building on an individual basis by comparing the helping rates with and without score within subjects. Table 6 shows a classification of subjects. We call a subject strategic if the helping rates are generally higher in the part of the experiment where the subject has a score than in the part where he or she does not. A subject is called strongly strategic if the helping rates with score are in most cases at least twice the helping rate

¹²Of the 19 full scores of 0, 15 come from same subject, the only pure egoist. In all five sessions the mode for private scores is 0. For public scores the mode is 0.8 in three sessions. In one session the mode is 0.6 and in one session it is 1, with 0.8 being the second most frequent score in both cases.

Score	Session						Total
	1	2	3	4	5		
0	0	0	0	0	0	0	0
0.2	-0.6	2.9	0.57	-0.6	-0.6	0.57	0.57
0.4	0.55	0.91	1.11	1.03	2.3	1.03	1.03
0.6	2.1	1.2	1.22	0.87	1.8	1.39	1.39
0.8	2.52	3.01	1.51	2.63	2.85	2.42	2.42
1	2.07	2.48	1.11	0.81	2.62	1.92	1.92

Table 5: Average expected return per period (in Taler) from keeping a certain score, based on average help rates over the whole phase with full scores

without score.¹³ Finally, a pure strategist never helps when he or she does not have a score, but does so several times otherwise.¹⁴ There are several special cases of non-strategic subjects. Simple egoists never help, simple altruists always help and negatively reciprocal altruists always help when the recipient does not have a score or when the recipient’s score is above some cut-off level, but not for a lower score. A subject is classified as reciprocal when there is a clear positive relation between the recipient’s score and the helping rate.¹⁵

As Table 6 shows, the majority of subjects is clearly strategic, there are only 16 (20%) non-strategic types, but even 20 pure strategists and 23 strong strategists. The distribution of types does not depend

¹³For the classification of both strategic and strongly strategic, we allowed deviations from the criteria for one value of the recipient’s score and in that case required it to hold strictly for at least two values of the recipient’s score. In particular we required the criterium to hold for the case where the recipient did not have a score, because the number of observations was much higher than for any single recipient’s score.

¹⁴Note that the pure strategists were not included in the strong strategists.

¹⁵We allowed one exception from the criterium in the sense that for one low score the helping rate was allowed to be higher than for one or several higher scores or for one high score the helping rate was allowed to be lower than for one or several lower scores. In these cases we required at least two either low scores where the helping rate was lower than that for all higher scores or high scores where the helping rate was higher than for each lower score. A flat helping rate in case the donor did not have a score was allowed if the helping rate in case he did have a score showed a clear positive relation. For most subjects, the classification was straightforward, because there was either a clear monotone relation or none at all.

	Pure Strat	Strong Str.	Weak Str.	Non-Str.	Total
Reciprocal	8	12	14	4 (5)	38 (39)
Non-Reciprocal	12	11	7	6 (11)	36 (41)
Simple Egoist				1	1
Simple Altruist				4	4
Negativ Rec. Altr.				1	1
Total	20	23	21	16	80

Table 6: Classification of individual subjects (in absolute numbers). Numbers in parantheses include special types of reciprocal (negatively reciprocal altruist) and non-reciprocal (simple egoist and simple altruist) non-strategic types. These types are listed separately in the third to fifth rows.

substantially on whether the subjects have a score in the first or the second half of the experiment. It also does not differ a lot between sessions, but the concentration of pure and strong strategist is especially high (75%) in session 3. Interestingly, session 3 also has the highest number of reciprocal subjects (62.5%). There is only one simple egoist, and even more surprisingly only one negatively reciprocal altruist, which intuitively appears to be a perfectly reasonable and in particular socially desirable type (helps in general but punishes egoists). Some of the 4 simple altruists might be negatively reciprocal altruists, because they never encountered a recipient with a score below 0.4. Overall, the number of these simple types is much lower than in Seinen and Schram (2001), who find 11% egoists and 36% altruists.¹⁶ Thus we find a higher share of subjects that are reciprocal but also a higher share that are strategic. Interestingly, 40% of the pure strategists and 52% of the strong strategist are also clearly reciprocal. Hence while their primary motive to help appears to be strategic reputation building, they are also concerned with providing incentives for the other subjects and hence instead of just exploiting the cooperative system based on indirect reciprocity, they also stabilize it.¹⁷ The remaining 60% of pure strategists (15% of the total population), however,

¹⁶The numbers are, however, not perfectly comparable, because we use a different classification than Seinen and Schram (2001). In particular, our design allows us to more clearly detect strategic players and hence our classification is finer in that respect.

¹⁷This can be seen as being strategic on a higher level, because due to the matching procedure, donors could profit from inducing others to help, either by later being matched with them again or by indirect effects.

appear to be of the type predicted by Leimar and Hammerstein (2001) to invade the population.

Concerning the third hypothesis that the reciprocal relation is weaker for donors with score than for donors without, Figures 1 and 2 do not provide immediate support. We do, however, find support for this hypothesis if we take a closer look at the data. If we focus on the recipients with a full score (which appears appropriate, because the preceding interactions are subject to more random influence due to initial experimentation of donors), then Table 2 shows that in each session the coefficient of the recipient's score in the linear probability model is higher for donors without score than for donors with score. Hence the relation between the recipient's score and the helping probability is stronger for donors without score than for donors with score in each session, and thus overall is significantly stronger for the donors without score at $p = 5\%$. We can also address this question by a second measure, the adjusted R^2 for each regression. As is shown in Table 2, the adjusted R^2 is higher for the donors without a score than for the donors with a score in each session, and hence overall, the recipient's score is a (significantly) better predictor for the helping probability if the donor does not have a score.¹⁸ If we consider all recipients with a score, however, both these tests fail and hence the overall support for the third hypothesis is weaker than for the first two hypotheses.

These result are again supported by a panel data analysis. We further extend the above model to

$$\begin{aligned} Help_{it} = & \text{const} + \alpha \cdot RsScore_{it} + \beta \cdot DWithScore_{it} \\ & + \gamma \cdot (DWithScore \times RsScore)_{it} + \nu_i + \varepsilon_{it}, \end{aligned}$$

where the interaction term $DWithScore \times RsScore$ captures the additional effect of the recipient's score in case the donor has a score and the other variables are defined as above. As shown in Table 7, the coefficient for the interaction term is significantly ($p < 5\%$) negative if we restrict the analysis to the recipients with full score, supporting the third hypothesis that the impact of the recipient's score is weaker if the donor has a score. If we consider all recipients with a score, however, the coefficient is again negative, but far from significant ($p > 40\%$).

¹⁸The fit of the model is in general not very good because donors' behavior is quite heterogenous. This, however, rather strengthens the result since still a clear effect can be detected.

		Coefficient	Std. Error	z	$\Pr > z $
All	α	0.5641	0.0650	8.681	0.000
R	β	0.4402	0.0771	5.709	0.000
with	γ	-0.0673	0.0963	-0.699	0.485
score	<i>const</i>	-0.0573	0.0553	-1.035	0.300
R	α	0.7614	0.0790	9.636	0.000
with	β	0.5685	0.0926	6.138	0.000
full	γ	-0.2539	0.1186	-2.141	0.032
score	<i>const</i>	-0.2054	0.0616	-3.333	0.001

Table 7: Coefficients for the recipient’s score (α), for a dummy whether the donor has a score (β) and for the interaction term (γ) in random effects cross-sectional time-series regression for the help choice

One could argue that the positive relation between the recipient’s score and the helping probability does not result from indirect reciprocity, but rather from a learning process. Donors might want to find out what is a successful score and may use the observed scores as orientation. Trying to adapt the own score to the observed recipients’ scores would imply to help when one observes a high score and not to help when one observes a low score (though this should strictly be so only in early periods or if subjects are highly myopic, because otherwise the total information one has gathered so far should dominate this period’s recipient’s score). This potential interpretation, however, appears to be valid only for donors with score, because donors without a score do not have an incentive to find out what constitutes a successful score. The support for our third hypothesis, that the relation between recipient’s score and helping probability is stronger for donors without a score, thus contradicts this interpretation and suggests that what we see is indeed indirect reciprocity.

Our affirmative result on the third hypothesis has a potentially important, though at the current state somewhat speculative, implication. If the share of subjects with public score was higher, but the helping rates (conditional on the recipient’s score) for donors with and without public score were not affected, average helping rates would be higher, but the incentives to help in order to build a reputation would be weaker.

Confirming straightforward intuition, strategic reputation building pays, whereas reciprocity does not. Table 8 shows the average payoffs (relative to the average payoffs in the session) of subjects by being reciprocal and strategic, where we summarized the subjects who were classified as pure or strong strategists as strongly strategic and those that were classified as weakly or non-strategic as weakly strategic. Clearly, the strongly strategic outperform the weakly strategic, which does not come as a surprise because being strategic implies, at the same own score, lower costs for helping others. The advantage of the strongly strategic players is, however, remarkably large.¹⁹ More importantly, it pays not to be reciprocal, apparently because being reciprocal distracts from perfectly fine-tuning the own score (or, in case of private scores, is a pure waste).²⁰ This indicates that in an evolutionary game based on this repeated helping game and with the experimentally observed player types, the purely strategic non-reciprocal types would drive out the other types and would eventually undermine the cooperation. Given that the relative payoff of the non-reciprocal strongly strategic players is almost twice that of the reciprocal weakly strategic players, the evolutionary process would be quite fast for any sufficiently payoff-sensitive dynamic.

Figure 4 shows the development of the average helping rates in the first 40 and the second 40 periods. While there is a clear drop in the last two periods in both cases, the helping rate is remarkably stable until the third to last. Since the value of a high score decreases sharply towards the end of the experiment,

¹⁹If we study the data in a more disaggregated way, we find that the payoff for the pure strategic is slightly higher than that for the strongly strategic and the payoff for the weakly strategic is substantially higher than for the non-strategic. Since the numbers of observations is too low for some categories in some sessions to derive meaningful results and since the largest difference is between strongly strategic and weakly strategic we aggregated the data in two categories for the present analysis.

²⁰We see no reason why the individual subjects' payoffs *relative to the session average* should be dependent within one session, hence we use the relative payoffs of individuals as independent observations in the following tests. According to both Mann-Whitney and two-sample t-tests, the differences in relative payoffs between strongly strategic and weakly strategic subjects are significant ($p < 2\%$ for the non-reciprocal, $p < 1\%$ for the reciprocal and $p < 0.1\%$ for the whole sample), but those between reciprocal and non-reciprocal subjects are not ($p > 10\%$ for the strongly strategic, the weakly strategic and the whole sample.) For a stricter test, which does not treat the relative payoffs of individuals as independent, note that the average relative payoffs are larger for the strongly strategic than for the weakly strategic in all 5 sessions (for the non-reciprocal, for the reciprocal as well as for the whole sample) and hence we can reject the hypothesis that the strongly strategic do not do better at $p = 5\%$. The non-reciprocal do better than the reciprocal in only four sessions and hence this test also misses statistical significance.

	Strongly Strategic	Weakly Strategic	Total
Reciprocal	1.14 (20)	0.69 (19)	0.92 (39)
Non-Reciprocal	1.23 (23)	0.87 (18)	1.08 (41)
Total	1.19 (43)	0.78 (37)	

Table 8: Payoffs relative to average session payoff for pure or strongly strategic versus weakly or non-strategic and for reciprocal versus non-reciprocal players, number of players in the respective category in parantheses.

one might have expected helping rates to drop earlier. An analysis of the sources of the end-game effect is remarkably consistent with our classification of subjects into purely strategic, strongly strategic, weakly strategic and non-strategic based on the comparison of behavior with and without public score. A subject who helps primarily in order to strategically build a score would be expected to lower his or her helping rate in the last periods when he or she has a public score. What we observe when we compare individual subjects' helping rates in the last two periods with their overall helping rates when they have a public score is consistent with this expectation. Out of 17 subjects whom we classified as purely strategic and who have been a donor at least once in the last two periods, only 2 increase their helping rate, while 15 lower it.²¹ The corresponding numbers for the strongly strategic players are 4 and 14, for the weakly strategic they are 5 and 8 and for the non-strategic they are 4 (plus 4 with a constant helping rate) and 5. Thus the end-game behavior clearly corresponds to our classification of subjects in terms of strategic behavior. Subjects classified as purely or strongly strategic exhibit a clear drop in helping behavior towards the end while those classified as weakly or non-strategic do not.²² Overall, the helping rate is somewhat lower in the second half of the experiment (51%) than in the first (56%). In the individual sessions, this effect is not consistent. The difference is reversed in one session and virtually zero (0.3%) in another.

²¹Most subjects were a donor only once in the last two periods. For these players, increasing the helping rate means helping this one time and decreasing the helping rate means not helping this one time.

²²Overall helping rates do not differ dramatically between the different categories in the phase when subjects have a score. For purely strategic subjects it is 67%, for strongly strategic 68%, for weakly strategic 84%, and for non-strategic 72%. Therefore, the observation concerning end-game effects is not an artifact of differences in overall helping rates.




Figure 4: Average helping rates over all sessions for the respective 40 periods of the first and the second half of the experiment

Seinen and Schram (2001) analyze the development of different group norms, which are apparent minimal scores that donors demand from a recipient in order to help him or her and which they also try to achieve themselves. They find that over time different groups tend to adopt different norms, since different groups tend towards different helping rates and differences between groups increase over time. We see relatively little evidence for similar tendencies in our data. In both parts of the experiment, helping rates in the first 10 periods differ about as much as in periods 21-30. And they even differ less in the last 10 periods with the exception of a stronger end-game effect in session 3. That effect can be attributed to the high number of strongly strategic players in session 3 and thus it does not indicate the development of a different norm, but rather more violations of the norm.

We also find some support for the alternative notion of indirect reciprocity that players help who have recently been helped. When comparing the average help rates of donors who have been recipients in the last interaction and have been helped with those who have been recipients in the last interaction and have not been helped (see Table 9), we see that for all possible own scores the help rate is higher for the former (one has to consider donors with different scores separately, because otherwise one might incorrectly conclude

own score	0	0.2	0.4	0.6	0.8	1
public score	-	1	0.82	0.91	0.78	0.92
	0	0.64	0.57	0.73	0.55	0.69
private score	0.09	0.32	0.61	0.55	0.78	0.97
	0.03	0.15	0.34	0.44	0.45	0.85

Table 9: Average helping rates for donors who have been recipients in the last interaction and have been helped (top) and not been helped (bottom), by donor’s score and score status.

that being helped increases the probability of helping if in fact donors who have helped more frequently in the past do so in the present but are also helped more frequently). For donors with a (high) public score, it is not entirely clear whether this result really indicates such a form of indirect reciprocity or whether it rather results from subjects realizing that a high score has benefitted them. For donors without public score the effect is also present which provides more conclusive evidence for this form of indirect reciprocity. It is, however, not very consistent in the individual sessions.

4 Related Literature

4.1 Standing Strategy as an Alternative to Image Scoring

Our experiment has pointed out that subjects react to the strategic incentives that are caused by future donors’ indirect reciprocity. This implies that they face a trade-off when they meet a recipient who has a low score. If they are motivated by indirect reciprocity they want to punish this player. Punishing by not giving, however, deteriorates the own score. Hence the own indirect reciprocal motivation may be in conflict with the strategic incentives implied by future donors’ indirect reciprocal actions. Furthermore, as pointed out by Leimar and Hammerstein (2001), a system based on image scoring is susceptible to exploitation by strategic score optimizers. These problems can only be avoided if higher order information is provided, for instance, information about the score of the previous partners of the current recipient.

Sugden (1986) analyzes a model with higher order information and what he calls “standing strategy”. A player is in “good standing” as long as he has helped or has refused to help a recipient not in good standing,

but loses good standing if he does not help a recipient in good standing. Leimar and Hammerstein (2001) show that the standing strategy is evolutionarily stable if the horizon is sufficiently long and the probability of execution errors is small, but considerably larger than the probability of perception errors. They show in simulations that even in a case where these conditions are not all fulfilled, the standing strategy dominates the population and hence while it is not evolutionary stable, it is more robust than an indirectly reciprocal strategy based on a cut-off strategy.

Bolton, Katok, and Ockenfels (2001) conducted an experimental helping game where donors in one treatment were informed only about the current recipient's last choice in the donor role, but in another treatment they were also informed about the recipient's score in that interaction. The underlying model of Bolton et al. (2001) hence corresponds in the first treatment to image scoring models, while in the second treatment it is conceptually closer to that analyzed by Sugden (1986). A crucial difference, however, is that in the situation analyzed by Sugden, there is an incentive both to punish a defector and to help a player who has punished whereas in the underlying model of Bolton et al. there is no incentive to help a punishing player because if one does not help a recipient who has not helped lately, future partners cannot detect whether one has refrained to help an exploiter or a punisher. Hence cooperation based on good standing is more robust than that based on the second order information in the underlying model of Bolton et al.

In Bolton et al. (2001) the presence of second order information clearly increases cooperation, implying that "deeper" information (looking back deeper into the history of play) fosters helping, whereas our results suggest that "more" information (a higher share of players with a public score) might also have adverse effects. The first order information has the strongest influence when the second order information was "help", i.e. for the own probability of being helped it is most important how one last treated another player who has helped. This would be consistent with the conclusion that donors try to follow a standing strategy. If the current recipient's last recipient had helped the last time he was the donor (and hence was unambiguously in good standing), the current recipient's standing can be determined and thus guide the current donor's decision. In contrast, if the current recipient's last recipient had not helped the last time he was the donor, the information is not sufficient to determine this player's standing and hence also the current recipient's standing. Therefore, the information on the current recipient is ambiguous and is thus not expected to have a clear impact on the current donor's decision.

The results of Bolton et al. (2001) agree with ours to the extent that they find that strategic incentives (e.g. costs) matter. In their study this is also shown by a clear end-game effect, similar to our results. Given that in our study the score is based on several periods, one might, however, expect helping rates to drop substantially before the second to last period, which they do not. Bolton et al. (2001) also find that being helped increases the probability of helping in the next period, consistent with the alternative notion of indirect reciprocity suggested by Boyd and Richerson (1989).

Similar to the approach by Bolton et al. (2001), Milinski et al. (2001) compare treatments with first and second order information, but they provide information about all the preceding choices of the recipient and his or her previous recipients, not only about the respective last choices. Furthermore, subjects have the opportunity to track also higher order information, because all decisions are made public and subjects are given names throughout the experiment. While Milinski et al. (2001) find that second order information is taken into account to some extent, they interpret their data to be overall rather in line with an image scoring strategy than with a standing strategy because observed helping rates (for a specific group of players, i.e. those that did not help a recipient who had not helped before and were hence not to blame according to a standing strategy) are not significantly different from expected helping rates under an image scoring strategy (though they were higher) but are significantly lower than expected helping rates under a standing strategy. It appears that subjects took the information that was explicitly provided (first and second order) into account, but were not capable of (or not interested in) keeping track of higher order information.

To us, an effective test of a standing strategy appears difficult to design. An experimental design that provides sufficient information for a standing strategy would either overwhelm subjects with information (if it provides them with the necessary third, fourth and higher order information, as in Milinski et al., 2001) or it would directly suggest to use a standing strategy by providing information about the standing explicitly. Outside the laboratory, people might be able to develop efficient mechanisms to keep track of others' standing, whereas in an experiment, time might be too short to develop such an efficient mechanism. In particular, in non-anonymous interaction outside the laboratory, information about another's past might be retrieved from memory via emotional reactions, which appears to be much harder in an experiment where identification is only via fictitious names or player numbers.

4.2 Effects of Reputation across Games

Milinski, Semmann, and Krambeck (2002a, 2002b) study how an image score acquired in one game (or individual decision) influences the behavior in another. In the experiment by Milinski et al. (2002a) a public goods game is combined with a helping game, where donors are also informed about the recipient's choice in preceding rounds of the public goods game. They find that contributions in the public goods game deteriorate if firstly eight periods of the public goods game are conducted and then eight rounds of the helping game. In contrast, if rounds of the public goods game and the helping game alternate, contributions in the public goods game stay at a substantially and significantly higher level. Furthermore, contributions in subsequent rounds of the public goods game decline quickly if subjects are informed that no further rounds of the helping game follow, but do not if subjects are not given this information. These results confirm that subjects clearly follow the strategic incentives to build a reputation. In the design, however, elements of direct and indirect reciprocity are combined. Rewards and punishment in the helping game of previous behavior in the public goods game are clearly a form of direct reciprocity because the donor was affected by the recipient's behavior in the public goods game. Hence the main result shows that subjects strategically react to the possibility of direct reciprocity, not indirect reciprocity. The experiment provides again, however, clear evidence for indirect reciprocity (and even of a similar impact than that of direct reciprocity), because the decisions in the helping game influence the probability to be helped as much as the decisions in the public goods game.

The experiment by Milinski et al. (2002b) shows that donations made to charity (UNICEF) significantly increase the probability to be helped in a helping game. Hence indirect reciprocity does not only occur within closed groups, but also helping outsiders can improve the own chances of receiving help.

5 Conclusions

We have conducted an experimental helping game where at any time only half of the subjects have a public score and hence a strategic incentive to help. Thus we can study both pure non-strategic indirect reciprocity and the impact of strategic incentives. The interaction of donors with and without public score

is the fundamental difference to the helping experiment by Seinen and Schram (2001). In their experiment, all subjects could build up an image score (or none in the control treatment) and hence it is not possible to clearly distinguish between helping choices that result purely from a motivation for indirect reciprocity and helping choices that are driven by attempts to improve the own score.

We are the first to find clear evidence for indirect reciprocity even in the absence of strategic incentives for reputation building, but we also find very strong effects of strategic reputation building. Specifically, 80% of subjects react to strategic incentives, including more than 50% whose helping rates more than double and 25% who only help when they have an incentive to do so.²³ This is in contrast to the indirect evidence by Seinen and Schram (2001) who classify more than half of the players as non-strategic and in particular find substantially more simple altruists and simple egoists, types that are virtually absent in our experiment.

We also find support for our hypothesis that the subjects who have a public score and hence a strategic incentive to help, exhibit somewhat weaker reciprocal behavior than subjects who do not have a public score. Hence while strategic incentives to build up an image score increase an individual's general propensity to help, it weakens the influence of pure indirect reciprocity. A potential implication is that in a population with a higher share of subjects with a public score, the incentives for strategic reputation building could be weakened.

Concerning the empirical relevance of the invasion predicted by Leimar and Hammerstein (2001), we clearly find strategic non-reciprocal players who also receive higher payoffs than other types. This casts some doubt on the evolutionary explanation for cooperation based on indirect reciprocity suggested by Nowak and Sigmund (1998a) because the types predicted to undermine the cooperation by exploiting the system are clearly present and more successful. Put differently, the argument by Leimar and Hammerstein (2001) that the set of potential types chosen in the simulations by Nowak and Sigmund (1998a) is too restricted is not only valid on theoretical grounds, but is also strongly supported by our experimental data.

²³An alternative explanation for the observed higher helping rates of donors with a public score could be that some donors want to set an example for what they believe to be the right thing to do. Without a public score only the recipient becomes aware of their choice, whereas with a public score their future donors learn their past behavior, which enhances the scope for setting an example.

The exploiting types are out there, so any simulation or evolutionary model that tries to explain some phenomenon in the human society has to take them into account. Therefore, an evolutionary explanation for the presence of indirect reciprocity (that is documented by several experiments, including ours) has to be richer in structure to explain why reciprocal players might survive in the presence of non-reciprocal strategic players. One might argue on the other hand, that the exploiters are a relatively new phenomenon and that we are still in the take-over process. While this view might be correct, we consider it too depressing to share.

References

- [1] Alexander, Richard D., 1987. *The Biology of Moral Systems*. New York: Aldine de Gruyter.
- [2] Andreoni, James and Petrie, Ragan, 2001. "Social Motives to Giving: Can these Explain Fund Raising Institutions?", *Working paper*, University of Wisconsin.
- [3] Berg, Joyce, Dickhaut, John, and McCabe, Kevin, 1995. "Trust, Reciprocity and Social History." *Games and Economic Behavior* 10, 122–142.
- [4] Blount, Sally, 1995. "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences." *Organizational Behavior and Human Decision Processes*, 1995, 63, 131–144.
- [5] Bolton, Gary E., Katok, Elena, and Ockenfels, Axel, 2001. "What's in a Reputation? Indirect Reciprocity in an Image Scoring Game." *Working paper*, Penn State University.
- [6] Bolton, Gary E., Katok, Elena, and Ockenfels, Axel, 2002. "How Effective Are Online Reputation Mechanisms? – An Experimental Investigation." *Working paper*, Penn State University.
- [7] Bolton, Gary E. and Ockenfels, Axel, 2000. "ERC – A Theory of Equity, Reciprocity, and Competition." *American Economic Review*, 90, 166–193.
- [8] Boyd, R. and Richerson, P. J., 1989. "The Evolution of Indirect Reciprocity." *Social Networks* 11, 213–236.

- [9] Charness, Gary, 2002. "Attribution and Reciprocity in an Experimental Labor Market." Forthcoming in *Journal of Labor Economics*.
- [10] Charness, Gary and Rabin, Matthew, 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117, 817–869.
- [11] Cox, James, 2002. "How to Identify Trust and Reciprocity." *Working paper*, University of Arizona.
- [12] Dufwenberg, Martin and Kirchsteiger, Georg, 2000. "A Theory of Sequential Reciprocity." *Working paper*, CentER for Economic Research, Tilburg.
- [13] Falk, Armin and Fischbacher, Urs, 1999. "A theory of Reciprocity." *Working paper* No. 6, Institute for Empirical Research in Economics, University of Zurich.
- [14] Fehr, Ernst, Kirchsteiger, Georg, and Riedl, Arno, 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics*, 108, 437–460.
- [15] Fehr, Ernst and Schmidt, Klaus, 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114, 817–868.
- [16] Fischbacher, Urs, 1999. "Z-Tree: Zurich Toolbox for Readymade Economic Experiments." *Working paper* No. 21, Institute for Empirical Research in Economics, University of Zurich.
- [17] Harbaugh, William T., 1998. "The Prestige Motive for Making Charitable Transfers." *American Economic Review*, 88, 277–282.
- [18] Keser, Claudia, 2002. "Trust and Reputation Building in e-Commerce" *Working paper*, IBM T. J. Watson Research Center.
- [19] Leimar, Olof and Hammerstein, Peter, 2001. "Evolution of Cooperation through Indirect Reciprocity." *Proceedings Royal Society London: Biological Sciences* 268, 745–753.
- [20] Levine, David K., 1998 "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics*, 1(3), 593–622.

- [21] Milinski, Manfred, Semmann, Dirk, and Krambeck, Hans-Jürgen, 2002a. “Reputation Helps Solve the ‘Tragedy of the Commons’.” *Nature* 415, 424–426.
- [22] Milinski, Manfred, Semmann, Dirk, and Krambeck, Hans-Jürgen, 2002b. “Donors to Charity Gain both Indirect Reciprocity and Political Reputation.” *Proceedings Royal Society London: Biological Sciences* 269, 881–883.
- [23] Milinski, Manfred, Semmann, Dirk, Bakker, Theo C. M., and Krambeck, Hans-Jürgen, 2001. “Cooperation through Indirect Reciprocity: Image Scoring or Standing Strategy?” *Proceedings Royal Society London: Biological Sciences* 268, 2495–2501.
- [24] Nowak, Martin A. and Sigmund, Karl, 1998a. “Evolution of Indirect Reciprocity by Image Scoring.” *Nature* 393, 573–577.
- [25] Nowak, Martin A. and Sigmund, Karl, 1998b. “The Dynamics of Indirect Reciprocity.” *Journal of Theoretical Biology* 194, 561–574.
- [26] Rabin, Matthew, 1993. “Incorporating Fairness into Game Theory and Economics.” *American Economic Review*, 83, 1281–1302.
- [27] Seinen, Ingrid and Schram, Arthur, 2001. “Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment.” *Working paper*, CREED, University of Amsterdam.
- [28] Sugden, R., 1986. *The Economics of Rights, Co-operation and Welfare*. Oxford: Basil Blackwell.
- [29] Wedekind, Claus and Milinski, Manfred, 2000. “Cooperation through Image Scoring in Humans.” *Science* 288, 850–852.

Appendix: Instructions

(Original Instructions were in German)

General Instructions

You are taking part in an economic experiment, which is being financed by various research promoting foundations. If you read the following instructions carefully, you can - depending on the decisions you make - influence your own earnings as well as the earnings of the other participants in this experiment. It is, therefore, important that you pay attention to the instructions given below.

The instructions distributed are intended for your personal information only. **Absolutely no communication whatsoever is allowed for the duration of the experiment.** Please address questions you might have to us directly. Violation of this rule leads to the exclusion both from the experiment itself and from all pertaining payments.

The experiment is divided into **periods**. During this experiment we do not deal with francs, but with points. Your income from each period will, therefore, be calculated in points. The total amount of points achieved in the course of the experiment will be converted into francs at the rate of

1 point equals 10 rappen [100 rappen = 1 Swiss Franc].

At the beginning of the experiment you are allotted an endowment of 100 points, thus representing 10 francs.

In each period you form a group with **one** other participant. These groups of two are in each period newly formed at random. It is possible, though not probable, that you will be linked with the same participant in two consecutive periods. You cannot recognize the other participants, and hence do not know whether you have been in a group together with the current other participant before. This guarantees the anonymity of your decision.

Each group consists of one participant with **part A** and one participant with **part B**. Both parts are, in each period, randomly and independently assigned. The probability of being assigned part A for a period is 50 %, irrespective of the part held in the previous period. Therefore, it is possible that you will assume part A or part B in several consecutive periods.

Specific Instructions for the Experiment's Procedure

Decisions to be made by the participants

During each period, in which you assume part A, you determine whether or not you want to help the other participant of your group (who holds part B). If you assume part B no decision is required from you. If you, as the holder of part A, decide to help the other participant of your group, you will be charged with a cost of 6 points, and the other participant of your group is given 15 points. If you decide not to help the other participant of your group, you suffer no cost, and the other participant receives nothing, resulting, for both of you, in the same amount of points as at the beginning of the period.

Participants' information types

The participants differ from each other insofar as other participants are, or are not, informed of the decisions made. Participants, whose decisions are communicated to the other participants, are referred to as **Info types**. The experiment comprises two stages consisting of 40 periods each. At stage one, i.e. during the first 40 periods, one half of the participants are info types. At stage two, the other half of the participants become info types. Thus, you will, like all other participants, be an info type **either** during the **first** 40 periods **or** during the **second** 40 periods. You will always know if you are an info type or not. If during the first 40 periods you were an info type, we will inform you at the end of these 40 periods that for the rest of the experiment you will no longer be an info type and vice-versa. Regardless whether you are an info type or not, you can in each period be matched both with another info type or to a non-info type.

Information on info types

The **last five decisions** made by the info types are being computer-saved, i.e. saved will be the number of times an info type (with part A) granted help and the number of times he denied help. When an info type then assumes part B, this information is given to the other participant of the group (assuming part A). This means that the participant with part A learns how many times the participant with part B granted help during the last five periods and how many times he did not. If at this stage the participant with part B assumed part A in less than five periods, the participant with part A is informed of decisions B made in these periods.

If a participant is not an info type, no information on his decision-making is saved. In particular this means that no one is informed about the decisions made at the stage where one is not an info type. Thus if at stage two of the experiment you are an info type, no information on the decisions you made at stage one will be passed on to another participant.

Participants with part B are given **no** information on participants with part A.

If you are an info type, whose current decisions as participant with part A are passed on to later participants assuming part A, you are, at the beginning of a period, informed of how you decided during the last five periods

with part A (or during less than these five periods if you assumed part A less than five times). This information is submitted to you regardless of which part, A or B, you assume.

Stage two of the experiment

On completion of the 40 periods of stage one and after a short break we will get started with stage two, again consisting of 40 periods. The info types of stage one are no longer info types, and the non-info types of stage one become the info types of stage two. At stage two, all information on the decisions made at stage one is no longer available. This means that the number of periods with part A about which information is released, starts at zero for all participants.

However, the amount of points earned at stage one are carried over to stage two.

Procedure on the Computer

The screen shown to both participants is divided in two sections. The **upper section of the screen** is independent on whether you assume part A or part B.

Information given in the upper section of the screen

Each period reveals, in the **upper section of the screen**, the part you assumed in the previous period as well as the decision the participant with part A made in the last period (see figures 1 and 2 below). Furthermore, you are shown your actual balance of points. As an info type you will also see how many times during the last five periods as A (or during all previous periods as A, if they amount to less than five) you granted help to the participant with part B and how many times you denied it (see the example in figure 1). This is for your information. In the example in Figure 1 you have been the participant with part A during the last period, granting help to the participant with part B. During the last five periods with part A, you granted help twice and denied it three times. The current balance is 121 points. The example in Figure 2 shows the upper section of the screen, if you are not an info type. During the last period you assumed part B and were granted help. Your current balance is 121 points.

Decision-making section for participants A

If you are the participant with part A, you make your decision in the **lower section of the screen**. If the other participant of your group, i.e. the participant with part B, is an info type, you are informed about B's last five decisions (i.e. about the last five periods where he assumed part A). In the event that the other participant of your group, i.e. the participant with part B, is not an info type, you are informed that no information is released to you. The screen below shows that the participant with part B granted help three times and denied it twice during the last five periods where he assumed the part of A.

Below you will see the following question: **“Do you help participant B in this period?”** beside the two fields “Yes” and “No”. Mouse-click one of these fields and activate the “OK” button. **If you choose “Yes” your balance of points will be reduced by 6 points and participant B’s balance will be increased by 15 points. If you choose “No” neither your nor participant B’s balance will be changed.**

Besides, you will learn if you are an info type, which in this example applies. Thus your decision will in future periods, where you assume part B, be revealed to the participant with part A as long as these decisions belong to your last five decisions as the participant with part A and as long as you are at the same stage of the experiment.

Figure 1: Screen for participants with part A

Period 13 of 40
In the last period you were participant A. You granted help. As A during the last 5 periods You granted help twice You denied help three times. Current balance of your points: 121
During this period you are participant A Your participant B during the last 5 periods as A Granted help three times denied help twice. Do you help participant B in this period <input type="radio"/> Yes <input type="radio"/> No Your decisions will be revealed to your future participants A OK

Lower section of the screen for participants with part B

The lower part of the screen only informs you that during this period you are not to make any decision.

Figure 2: Screen for participants with part B

Period
16 of 40
In the last period you were participant B. You were granted help.
Current balance of your points: 121
You are participant B. During this period you make no decision.
continue

Control Questionnaire

Please answer all questions. Wrong answers have no consequences whatsoever! Address any questions to us!

1. Participant A has 121 points, participant B has 112 points. Participant A helps participant B. The balance of points of the participants is:

participant A:

participant B:

2. Participant A has 145 points, participant B has 127 points. Participant A denies participant B help. The balance of points of the participants is:

participant A:

participant B:

3. Suppose you are an info type. During the last five periods you made the following decisions: “help denied”, “help denied”, “help granted”, “help granted”, and “help denied” (in this sequence). You are now again A. In the event that you now help and that in the next period you assume part B: which information on your decisions will be released to participant A?

you granted help times

you denied help times

4. Suppose that during the first stage of the experiment you are an info type. In how many periods, at the most, is the decision you make in period 37 revealed to another participant?

5. Suppose you had the part of B three consecutive times. What is the probability of you again assuming part B during the next period?



CERGE-EI
P.O.BOX 882
Politických vezůů 7
111 21 Prague 1
Czech Republic
<http://www.cerge-ei.cz>