

Gutenberg School of Management and Economics
& Research Unit “Interdisciplinary Public Policy”

Discussion Paper Series

Reciprocity towards Groups

David Hugh-Jones and Martin A. Leroch

March 2013

Discussion paper number 1511

Johannes Gutenberg University Mainz
Gutenberg School of Management and Economics
Jakob-Welder-Weg 9
55128 Mainz
Germany
wiwi.uni-mainz.de

Contact details

David Hugh-Jones
Department of Government
University of Essex
Wivenhoe Park, Colchester
Essex, CO4 3SQ

dhughj@essex.ac.uk

Martin A. Leroch
Institute of Political Science
Johannes Gutenberg University Mainz
Jakob-Welder-Weg 12
55128 Mainz

leroch@politik.uni-mainz.de

Reciprocity towards Groups

David Hugh-Jones and Martin A. Leroch*

28 March 2013

People exhibit group reciprocity when they retaliate, not against a person who harmed them, but against another person in that person's group. We tested for group reciprocity in laboratory experiments. Subjects played a Prisoner's Dilemma with partners from different groups. They then allocated money between themselves and other participants. In punishment games, subjects whose partner had defected punished participants from the partner's group more, compared to their punishment of participants from a third group. In dictator-style games, subjects did not exhibit group reciprocity. We examine possible correlates of group reciprocity, including group identification and cooperativeness.

Keywords: reciprocity, groups, conflict, experiment

JEL classification: D74, C92

* David Hugh-Jones: Essex University, dhughj@essex.ac.uk. Martin Leroch: University of Mainz, leroch@politik.uni-mainz.de. This work was supported in part by British Academy Small Research Grant SG101553. We thank Vittoria Levatti, Werner Gueth, Eva Steiger, Johannes Weisser, Ro'i Zultan, David Reinstein, Ryan McKay, Eva van den Broek, Shaun Hargreaves Heap and seminar participants at the Max Planck Institute of Economics, University of Hamburg, University of Warwick and King's College London; and the Max Planck Institute ESI group hiwis and administrative staff: Martin Beck, Nadine Erdmann, Adrian Liebrau, Christian Williges, Christian Streubel, Claudia Zellmann and especially Christoph Goering.

1 Introduction

Human behaviour often follows group boundaries. People kill each other due to their national or ethnic affiliation, beat each other up because they cheer for different football clubs or avoid physical contact with “untouchable” castes. In the sphere of economic life, group-based behaviour may lead to boycotts of products originating from a specific country or specific firms, or to quarrels at the workplace leading to possibly severe inefficiencies.

We argue that reciprocity also follows group boundaries: people sometimes strike back, not against the individual who caused the perceived harm, but against other people in that individual's group. In these cases, those punished may have nothing to do with the initial offence, and the offender himself may be completely unaffected by the retaliation. We call this kind of retaliation *group reciprocity*.

Group reciprocity appears widespread in human conflict, including riots, pogroms and ethnic civil wars. Many such episodes involve calls for revenge against an opposing group for previous harms, whether real or imaginary, recent or ancient.¹ Indeed, grievances over previous conflicts may provide the

¹ For numerous examples, see Horowitz (1985; 2001).

cause or pretext for further outbreaks, leading to cycles of violence.² These cycles hinder development in many of the world's poorest states (World Development Report 2011).

Although group reciprocal *behaviour* has been attested beyond reasonable doubt, we know little about the motivations behind it. Acts of group reciprocity could be undertaken by materially self-interested, rational actors. For example, boycotting a country's exports may be seen as an effective way to change its policies. Striking back against groups may help someone to build a reputation for toughness. Or, not doing so might expose one to punishment from others in one's own group: one eyewitness of the Rwandan genocide stated “ten percent helped; 30 percent were forced to kill; 20 percent killed reluctantly; 40 percent killed enthusiastically” (Mamdani 2001). However, as this quote suggests, it is also possible that some people may group-reciprocate even when there is no strategic incentive to do so. These people will act as if group fairness is a primitive in their utility function: they simply prefer to harm (or help) members of groups who have harmed (helped) them.³

² For the role of grievance in civil wars, see Collier and Hoeffler (2004), Fearon and Laitin (2003), as well as the literature review in Blattman and Miguel (2010). For its role in terrorism see Krueger and Maleckova (2003).

³ We do not propose a particular functional form for group-reciprocal utility in this paper. Moreno (2008) models group fairness on the lines of Rabin (1993).

To investigate whether this is the case, we ran laboratory experiments in which there was no strategic incentive to group-reciprocate. Subjects in the experiments were divided into groups. Each group jointly completed a cooperative task. Then, each subject played a Prisoner's Dilemma (PD) against a partner who was a member of a different group. After play in the PD was revealed, subjects made decisions allocating money between themselves and members of other groups. We compare each subject's allocations to the fellow group members of their partner in the PD, with their allocations to members of a different group with which they had not previously interacted. Using this other group as a control, we can test whether subjects' PD partners' cooperation or defection affected their allocations specifically to the partner's group.

Our main findings are as follows. While the majority of our subjects do not group-reciprocate, and indeed do not discriminate at all between groups, a minority of subjects do group-reciprocate. In particular, we find significant evidence of group reciprocity in punishment decisions, where subjects must pay to harm others. However, we do not observe group reciprocity in dictator games where harming others benefits the subject. Lastly, we examine some possible correlates of group reciprocity. Although group reciprocity is only induced when subjects cooperated and their partner defected, subjects who are more cooperative *in general* are no more likely to group-reciprocate than oth-

ers. Group reciprocity is also not predicted by a psychological measure of group identification.

In the next section we review the relevant literature, and explain our contribution. We then describe our experiments and results.

2 Related Literature

Economists have long been interested in the economic consequences of groups and ethnicity (Becker 1957; Arrow 1972; Akerlof 1976). More recently, economic experimentalists have focused on how group membership affects behaviour, a topic that has long fascinated social psychologists (Sumner 1906; Sherif et al. 1961; Tajfel et al. 1971). Hargreaves Heap and Zizzo (2009) show experimentally that people will put a monetary value on group membership. Further, people discriminate between in- and outgroups. They cooperate more with in-group members (de Cremer and van Vugt 1999; Goette and Huffman 2006, Guala et al. 2009; but cf. Bouckaert and Dhaene 2004); give in-group members more, reward them more for good behaviour, and punish them less (Chen and Li 2009); they also punish offenders more when the victim is an in-group member (Goette and Huffman op. cit., Bernhard et al. 2006). Debate continues over whether this different behaviour is driven by expectations of future reciprocity (Yamagishi and Kiyonari 2000; Yamagishi and Mifune 2008, 2009) or by pure preferences (Guala et al. 2009).

As psychologists have noted (Brewer 1999), ingroup bias *per se* does not imply that people will behave unkindly towards members of an outgroup, only that they will behave more kindly towards members of their ingroup. For example, when asked to allocate a “bad” (exposure to aversive noise) among ingroup and outgroup members, members of a minimal groups did not favour the ingroup (Mummendey et al. 1992). Therefore, ingroup bias alone may not explain the empirical regularity that many real-world conflicts involve appalling intergroup violence (Brewer *op. cit.*).

Dispositions to group reciprocity may be a missing part of the story. Individual biases towards one's own group, which are (relatively) innocuous in ordinary circumstances, can become more dangerous after a perceived threat or insult from another group. Psychologists have investigated indirect group reciprocity (i.e. reciprocity towards groups in retaliation for harms done to other members of one's own group) under the term “vicarious retribution” (Lickel et al. 2006). Stenstrom et al. (2008) show that subjects will verbally express a demand for sanctions in retaliation for a presumed intergroup insult. In Gaertner et al. (2008), subjects allocated unpleasant noise to a group one of whose members had rejected the subject.

Group reciprocity may also help to explain the harmful intergroup behaviour observed in some lab experiments. Abbink and Herrmann (2009) created a ten-round vendetta game between two groups: in each round, subjects could

deduct money from the other group, at a cost to themselves. The low observed rates of this inefficient behaviour (13%) were tripled (40%) by the addition of a symbolic reward. Suggestively, rates of harmful actions were increased by the other group's harmful actions in the previous round, although since the group included the person who chose the previous harmful action, group reciprocity cannot be clearly distinguished here from individual reciprocity. By contrast, Halevy et al. (2008) observe very low rates of intergroup harm when an option to benefit one's own group without harming the other group was available; this was a one-shot game, with less scope for group reciprocity to operate.

Similarly, some results from field data are compatible with group reciprocity. Shayo and Zussman (2010) show that Israeli judges were more likely to find in favour of Israeli plaintiffs against Arab defendants shortly after terrorist incidents. This could be explained by increased solidarity with the ingroup, or by bias against an outgroup perceived as responsible for an attack; a point favouring the former interpretation is that Arab judges were also more likely to find in favour of Arabs against Israelis after an attack. In a military context, Kocher et al. (2008) show that indiscriminate US bombing in Vietnam increased Vietcong control, suggesting that it drove Vietnamese civilians “into the arms of the rebels”.

We see our contribution as follows. We test experimentally, using real money payoffs, whether people will reciprocate the actions of one group member when they make choices affecting other members of that group. Our design rules out strategic motivations, including expectations about others' behaviour towards oneself and reputational considerations. We clearly distinguish group from individual reciprocity, by ensuring that subjects knew they were not harming (helping) the person who had harmed (helped) them, but rather others from that person's group. By having subjects make allocations to multiple out-groups, we also differentiate reciprocity towards a specific group from generalized reciprocity towards third parties (Dufwenberg et al. 2001, Nowak and Roch 2007). Thus, the laboratory setting allows us to identify group reciprocal behaviour more cleanly than would be possible with field data.

Experiments have been run both with lab-created groups (Tajfel et al. 1971; Hargreaves Heap and Zizzo 2009; Guala et al. 2009; Chen and Li 2009; Yamagishi and Kiyonari 2000; Yamagishi and Mifune 2008, 2009; de Cremer and van Vugt 1999) and with naturally occurring groups (Sherif et al. 1961; Fershtman and Gneezy 2001; Bouckaert and Dhaene 2004; Goette and Huffman 2006; Bernhard et al. 2006; Habyarimana et al. 2009). We chose to use lab-created groups in our experiment, for two reasons. First, lab-created groups reduce the likelihood that subjects will bring into the lab set patterns of intergroup behaviour, which might have a strategic or group-reputational com-

ponent. Second, lab groups are relatively easy to create without access to a specific subject pool, and this contributes to our goal of developing a simply reusable experimental paradigm. We also note that the use of lab-created groups arguably provides a stronger test of group reciprocal motivations than using powerful existing identities such as ethnicity. Nevertheless, we believe that future tests of group reciprocity with naturally occurring groups will also be interesting and important.

3 Experimental Design

We ran two experiments. Both experiments shared the following sequence of events. First, subjects were separated into color-coded groups. Second, subjects performed a collective task to increase group identification. Third, subjects were paired with members of other groups, and played a Prisoner's Dilemma (PD). Following Fehr and Fischbacher (2004), we think defection in a PD would be considered a norm violation by many of our subjects.⁴ Fourth, after the PD, subjects completed allocation tasks involving themselves, members of their PD partner's group (but not the PD partners themselves), and members of a third group with whom they had not previously interacted. Defection by a subject's PD partner might affect the subject's behaviour to mem-

⁴ Note, however, that we chose a simultaneous PD game, whereas Fehr and Fischbacher (2004) analysed a sequential PD.

bers of all other groups, not just to members of the partner's group; we use the allocations to the third group to control for this.

Our basic empirical model is

$$(1) \quad \text{giving}_{i,g} = f(\alpha_i + \kappa \text{Self } D_i + \lambda \text{Partner } D_i + \mu (\text{Partner } D_i \times \text{Self } D_i) \\ + X_i\beta + \gamma \text{Pgrp}_g + \delta (\text{Pgrp}_g \times \text{Self } D_i) + \zeta (\text{Pgrp}_g \times \text{Partner } D_i) + \eta \\ (\text{Pgrp}_g \times \text{Partner } D_i \times \text{Self } D_i) + \text{Pgrp}_g \times Z_i\theta + \varepsilon_{i,g})$$

Here, the dependent variable, $\text{giving}_{i,g}$ is the amount given by individual i to a member of group $g \in \{\text{Partner}, \text{Other}\}$. The variables are defined as follows. $\text{Self } D_i$ is 1 if subject i defected in the PD and 0 otherwise. $\text{Partner } D_i$ is 1 if the subject's PD partner defected and 0 otherwise. We interact $\text{Partner } D_i$ with $\text{Self } D_i$ because it is plausible that subjects will react differently if they cooperated, but their partner defected, than if both they and their partner defected. Parameters κ , λ and μ capture the effect of these covariates upon total amounts given. Pgrp_g is 1 if g is the group of the subject's PD partner, and 0 otherwise. Our parameters of interest are ζ and η . These are the effect of the PD partner's defection specifically on giving to the Partner group, above any effect on giving to the Other group. ζ gives the effect for subjects who are themselves PD cooperators and $(\zeta + \eta)$ gives the total effect for subjects who are defectors. Lastly, α_i is a subject specific intercept, X_i is a vector of known covariates which may affect subjects' giving to all groups, and Z_i is a vector of known covariates which may influence subjects' giving to the Partner group alone. Sub-

jects' individual propensities to give, α_i , are highly likely to be correlated with their behaviour in the PD, and therefore with values of *Self D* and *Self D* \times *Partner D*. For this reason, we treat α_i as a fixed effect. If our specification were linear ($f(x)=x$), then we could use first-differencing to eliminate both the fixed effects and all the terms affecting amounts given in general, leaving:

$$(2) \quad \text{giving}_{i,partner} - \text{giving}_{i,other} = \gamma + \delta \text{Self } D_i + \zeta \text{Partner } D_i + \eta (\text{Partner } D_i \times \text{Self } D_i) + Z_i\theta + \varepsilon_{i,partner} - \varepsilon_{i,other}$$

Our data is not suitable for a linear model in either experiment, so that naïve first-differencing will not work, but we use estimation strategies based upon the same idea.

For clarity and ease of comparison, the sequence of treatments and choices in both experiments is illustrated in Table 1 below. The complete description of the experiments follows.

	Subjects	Pre-Group Stage	Group Formation	Post-Group PD	Allocation Task	Post-experiment									
Experiment 1	180 subject in 6 sessions. 3 groups per session. 10 members per group. 1 ECU = 1 EURO.	None.	Random assignment into color-coded groups. Task: group pelmanism game. 5 ECU for winners.	Players endowed with 2 ECU and choose actions A or B. Payoffs in ECU: <table style="margin-left: auto; margin-right: auto;"> <tr> <td></td> <td style="text-align: center;">A</td> <td style="text-align: center;">B</td> </tr> <tr> <td style="text-align: center;">A</td> <td style="text-align: center;">+2,+2</td> <td style="text-align: center;">-2,+5</td> </tr> <tr> <td style="text-align: center;">B</td> <td style="text-align: center;">+5,-2</td> <td style="text-align: center;">-1,-1</td> </tr> </table>		A	B	A	+2,+2	-2,+5	B	+5,-2	-1,-1	10 randomly ordered binary choices allocating money. Players choose A or B. Payoffs shown in Table 2.	Questionnaire. Pelmanism winners revealed. Payment: 2.50 EURO show up fee <i>plus</i> Pelmanism + PD + results of 1 randomly chosen allocation per subject.
	A	B													
A	+2,+2	-2,+5													
B	+5,-2	-1,-1													
Experiment 2	172 subjects in 10 sessions. 6 groups per session. ⁵ 3 members per group. 100 ECU = £1.	Anonymous Prisoner's Dilemma with random other participant. Results revealed at end.	Random assignment into color-coded groups. Task: quiz on animal collective names, with group chat. 1000 ECU for winners. 100 ECU for losers.	Players endowed with 200 ECU and choose actions A or B. Payoffs in ECU: <table style="margin-left: auto; margin-right: auto;"> <tr> <td></td> <td style="text-align: center;">A</td> <td style="text-align: center;">B</td> </tr> <tr> <td style="text-align: center;">A</td> <td style="text-align: center;">+200,+200</td> <td style="text-align: center;">-200,+500</td> </tr> <tr> <td style="text-align: center;">B</td> <td style="text-align: center;">+500,-200</td> <td style="text-align: center;">-100,-100</td> </tr> </table>		A	B	A	+200,+200	-200,+500	B	+500,-200	-100,-100	6 randomly ordered continuous choices allocating money. 3 dictator games with taking frame: other participant has 1000 ECU and subject can take 0-1000 ECU. 3 punishment games: subject has 400 ECU, other participant has 900 ECU. Subject can spend up to 300 ECU to punish. Each ECU spent reduces other's payoff by 3 ECU. Each game played: once with random member of Own group, once with random member of Partner group (not PD partner him/herself), once with random member of (an) Other group. Other participant identified by group and number.	Questionnaire. Initial PD results and quiz winners revealed. Payment: £5 showup fee <i>plus</i> : Sessions 1-5: 1 chosen at random from: initial PD, quiz, second PD or allocation results (from one randomly chosen allocation per subject). Sessions 6-10: at random, <i>either</i> initial PD + quiz <i>or</i> second PD + allocation results.
	A	B													
A	+200,+200	-200,+500													
B	+500,-200	-100,-100													

TABLE 1: EXPERIMENT DESIGN SUMMARY

⁵ One session had 5 groups; 5 groups in total, over 4 sessions, had only 2 members.

3.1 Experiment 1

Experiment 1 was run at the Jena laboratory of the Max Planck Institute for Economics, using the computer software zTree (Fischbacher 2007).⁶ 180 subjects, recruited via ORSEE (Greiner 2004), took part during 6 sessions. Each session lasted about 1 hour. Subjects were paid in Experimental Currency Units (ECU), with 1 ECU = 1 Euro.

In each session, 30 subjects were randomly allocated into three groups named *red*, *green* and *blue*. Each subject was given a player number, running from 1-10 within each group. These numbers, together with the group names, were used to identify subjects to each other throughout the remainder of the experiment. In order to increase group identification and entitativity (the sense that each group forms a single entity), each group played a team game. This was an adapted version of the card game Pelmanism, in which players have to find pairs of identical cards from a face-down deck. In our computer version, the whole group anonymously voted for which card to turnover, and one selected member then observed the votes and chose a card. The group which found the most pairs was awarded 5 ECU. The winning group was not revealed until the end of the experiment.

⁶ Screenshots of the interface are available online in the file screenshots.zip

After the Pelmanism game, subjects were allocated into pairs, each pair containing subjects from two different groups. They were informed of their partner's group and player number. They then played a one-shot PD. Afterwards, subjects were shown their partner's choices and the resulting payoffs for both players. We call each subject's PD partner's group the Partner group; the remaining group is the Other group. For instance, if a red group member played the PD against a green group member, then for her the green group is the Partner group, and the blue group is the Other group.

Choice set	1		2		3		4		5	
Choice	A	B	A	B	A	B	A	B	A	B
ECU allocated to...									ECU allocated to...	
Self	2	1	4	3	1	2	4	5	Target 1	1 2
Target	5	2	3	0	4	1	3	0	Target 2	3 1
	1P-4P: Target is from Partner group; 1O-4O: Target is from Other group.								5P: Target 1 is from Partner group, Target 2 is from Other group. 5O: Target 1 is from Other group, Target 2 is from Partner group.	

TABLE 2: ALLOCATIONS IN EXPERIMENT 1

Lastly, each subject made a set of ten binary choices, allocating money between him- or herself and one or more other subjects (the “target(s)”), always identified by their player number and group. The choices are shown in Table 2. Subjects chose between allocation A and allocation B. There are five pairs of similar choices; we refer to each pair as a “choice set”. In each choice set, subjects faced tradeoffs between maximizing their own income, helping

the target, and increasing or decreasing inequality. Choice sets 1-4 allocate money to the subject him- or herself and to one other target. Choice B always gives less to the target. Choice sets 1 and 2 involve “punishment choices”: the subject can either maximize his own and the other's income, or pay to reduce the other's income. They also have distributional consequences: the income-maximizing choice in set 1 is unequal in the other's favour, compared to the alternative. In set 2, the income-maximizing choice is less unequal than the alternative. Choice sets 3 and 4 are “dictator” choices: the subject can either maximize his own income, or reduce his income but give more to the other. In set 3 the selfish allocation also reduces other-favouring inequality; in set 4 it increases self-favouring inequality. Each of choice sets 1-4 is made once where the target is a member of the Partner group (1P-4P) and once when the target is a member of the Other group (1O-4O). Two allocations, 5P and 5O, were between a Partner group member and an Other group member, with the subject herself not being involved. The choice here was which of the two groups to favour; A is the efficient choice. By comparing subjects' decisions within each pair of choices in a choice set, we can observe the effect of the target's group membership.

Subjects never made allocations involving their previous PD partner – only involving other members of the same group. They could easily observe this, because the PD partner's group and player number, along with the outcome of

the PD, was shown onscreen (in a “history box”) throughout the 10 allocations. The order of the allocations was randomized. At the end of the experiment a single allocation was randomly selected, and the relevant payoffs were implemented.

After the experiment, subjects were given a questionnaire including demographics, measures of group identity, and debriefing questions. The winning group of the Pelmanism game was then revealed. Finally, subjects were called up and paid privately for their winnings from the Pelmanism game, the Prisoner's Dilemma, and the randomly selected allocations (their own and others'). Each subject also received a 2.50 Euro showup fee.

Two points about this design deserve attention. First, allocation into groups and of PD partners within groups was random, and prior to the PD, subjects had no interaction with groups other than their own. Hence, in what follows, we can treat subjects' PD partners' cooperation or defection as an exogenous variable. Second, subjects knew that the allocations were the last part of the experiment. Therefore, in making allocation decisions, they had no strategic reasons to play a particular way (e.g. reputation-building). These comments also apply to Experiment 2. Descriptive statistics for our subjects are shown in Table 3.

Demographics			
<i>N</i>	180		
Male	142 (78.89%)		
Age	Min 19, max 67, mean 24.6, median 24		
Studying:			
... Law	8 (4.44%)	... Natural sciences	38 (21.11%)
... Social sciences	56 (31.11%)	... Other	38 (21.11%)
... Economics	27 (15.00%)	... Not a student	13 (7.22%)
Prisoner's Dilemma choices (self, partner)			
CC	72 (40.00%)	DC	44 (24.44%)
CD	44 (24.44%)	DD	20 (11.11%)
Allocations (those choosing option B)			
Choice set 1P	20 (11.11%)	Choice set 1O	24 (13.33%)
Choice set 2P	10 (5.56%)	Choice set 2O	7 (3.89%)
Choice set 3P	158 (87.78%)	Choice set 3O	154 (85.56%)
Choice set 4P	104 (57.78%)	Choice set 4O	94 (52.22%)
Choice set 5P	71 (39.44%)	Choice set 5O	80 (44.44%)

TABLE 3: DESCRIPTIVE STATISTICS, EXPERIMENT 1

Figure 1 shows the percentage of B choices made in each of the 10 allocations, split by whether the PD partner cooperated or defected. In choice sets 1-4, B is the decision that gives least to the target. In general, PD partner defection increases the proportion of B choices. In choice set 1, partner defection decreases unkind choices (which also reduce the subject's own payoff), but more so towards the Other group than towards the Partner group. Patterns for the other choice sets are less clear.

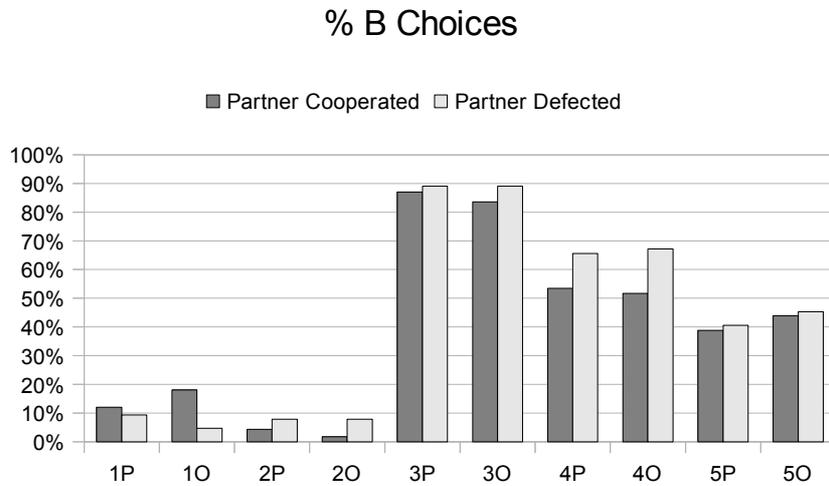


FIGURE 1: ALLOCATION DECISIONS BY PD PARTNER BEHAVIOUR, EXPERIMENT

1

Table 4 tabulates the pairs of choices made by our subjects in each choice set, split up by whether the PD partner cooperated or defected. In order to investigate whether PD partner's defection affected discriminatory behaviour, we examine only those subjects who discriminated either for or against the Partner group. If PD partner behaviour does not affect the level of discrimination, i.e. if $\zeta=\eta=0$ in equation (1), then there should be no association between partner behaviour and the proportion, among those who discriminate in either direction, of those who discriminate against the Partner group. The final row of

Table 4 shows the results of the corresponding Fisher test.⁷ Only the first choice set shows a significant effect: PD partner defection was associated with more discrimination against the Other group. Apparently, subjects who had been defected against were more willing to harm the target when the target was from the Partner group.

Choice set	1		2		3		4		5	
PD Partner	C	D	C	D	C	D	C	D	C	D
Choice P/Choice O										
A/A	94	56	111	56	9	5	47	19	53	28
A/B	8	2	0	3	6	2	7	3	18	10
B/A	1	5	3	3	10	2	9	11	12	7
B/B	13	1	2	2	91	55	53	31	33	19
Fisher test, A/B or B/A vs Partner C or D	0.0350*		0.464		1		0.260		1	

C = PD partner cooperated. D = PD partner defected. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 4: PARTNER AND OTHER GROUP ALLOCATIONS BY PD PARTNER BEHAVIOUR, EXPERIMENT 1

To separate out the effect of partner defection on defectors and cooperators, we run logit regressions. Thus, in equation (1), we take

⁷ We use a Fisher test rather than a chi-squared test because the Fisher test is exact when there are small numbers in each cell, as in this case.

$$f(x) = G_A \text{ if } x > 0; G_B \text{ otherwise}$$

in each of the five choice sets, where G_J is the amount given to the target by choice $J \in \{A, B\}$, and we assume that $\varepsilon_{i,g}$ has a logistic distribution. Fixed effects logits failed to converge in 2 out of 5 regressions, probably because this approach ignores subjects who made the same choice against both Partner and Other group members, leaving insufficient cases for estimation. As a second best, we estimate random effects logits. Table 5 reports selected coefficients. For each choice set, the first column comes from a regression without controls (i.e. with X and Z empty). Controls are not necessary, since the main variable *Partner D* is guaranteed to be exogenous by the experimental design, but appropriate controls will increase the accuracy of our estimates in the nonlinear logit model. The second column adds a vector of controls to both X and Z : dummies for gender, for subjects studying economics, and for group color. The third column adds the controls and also a vector of session dummies.⁸ The coefficients of interest are quite stable across these specifications. *Partner D* significantly predicts less generosity to the partner group in choice set 1 only, although this variable approaches weak significance in choice set 4. The *Partner D* \times *Self D* interaction is significant in choice set 5, but this is not easy to interpret, since the coefficient on *Partner D* is small and insignificant. Thus, group reciprocity was only evident in one out of the five choice sets.

⁸ Regressions for choice set 2 failed to converge and are omitted. It is clear from Table 4 that partner defection would not be significant in choice set 2, under any specification.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Choice set	1	1	1	3	3	3	4	4	4	5	5	5
Pgrp	2.47 (1.64)	-0.67 (9.2)	-2.12 (11.48)	-0.97 (0.8)	-0.98 (2.55)	-0.93 (2.53)	-0.29 (0.53)	-0.78 (1.12)	-0.77 (1.21)	-0.4 (0.41)	-0.49 (0.84)	-0.5 (0.85)
Pgrp × ...												
Partner D	-9.37 (4.2)*	-8.02 (4.05)*	-7.81 (3.88)*	0.67 (1.68)	1.99 (3.04)	1.89 (3)	-1.35 (0.88)	-1.4 (0.93)	-1.63 (1.04)	-0.77 (0.68)	-0.92 (0.72)	-0.96 (0.73)
Self D	6.87 (7.73)	8.12 (9.38)	8.67 (10.17)	0.19 (1.41)	1.78 (2.55)	1.54 (2.51)	0.29 (0.84)	0.41 (0.9)	0.41 (0.97)	0.07 (0.67)	0.11 (0.68)	0.13 (0.7)
Partner D × Self D	0.02 (8.86)	-2.31 (10.65)	-1.87 (10.28)	0.06 (2.74)	1.08 (5.33)	1.32 (5.32)	0.85 (1.47)	0.76 (1.54)	0.92 (1.66)	2.59 (1.2)*	2.79 (1.26)*	2.86 (1.27)*
Log Likelihood	-75.99 (9 df)	-73.32 (19 df)	-72.99 (24 df)	-106.37 (9 df)	-96.7 (19 df)	-96.31 (24 df)	-208.46 (9 df)	-204.23 (19 df)	-196.94 (24 df)	-225.67 (9 df)	-217.99 (19 df)	-210.94 (24 df)
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Session dummies	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes

Dependent variable is choice favouring Partner group (A in choice sets 1-4, B in choice set 5). Controls: gender, group color, economics student, Partner group decision taken before Other group decision. Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 5: RANDOM EFFECTS LOGITS, EXPERIMENT 1

3.2 Experiment 2

Experiment 1 produced evidence for group reciprocity in 1 out of 2 punishment choices, with little evidence for group reciprocity in the dictator choices (sets 3 and 4). To confirm these results, we ran a second experiment. In this design, subjects made continuous rather than binary choices at the final stage. By allowing subjects finer-grained control over their allocations, we hoped to get more accurate estimates of the distribution of group reciprocal preferences in the population. We also changed our group identity manipulation to increase levels of group identification. The new manipulation was a quiz on collective animal names, allowing for computer-based chat among participants of the same group.⁹

Experiment 2 was run at the University of Warwick experimental laboratory, again using zTree. 172 subjects, recruited from the subject pool via ORSEE, participated in 10 sessions. Subjects were paid in ECU, this time with 100 ECU = £1. At the start of the experiment, subjects were informed that the experiment had 4 stages, and that instructions for each stage would be revealed prior to the stage itself. In stage 1, subjects were randomly paired to play an initial PD, the results of which were not revealed until the end of the experi-

⁹ Pretests showed that this manipulation was successful in increasing subjects' levels of group identification.

ment. The aim of this pre-group PD was to observe subjects' propensity to cooperate, before any group assignment took place.

In stage 2, subjects were separated into 6 groups, *red*, *green*, *blue*, *purple*, *orange* and *cyan*.¹⁰ For the remainder of the experiment, all subjects were identified to others by their group name, combined with an individual, within-group number. (In this experiment, subjects were not informed of their own number.) After being assigned to their group, subjects took a collective quiz on animal group names. Subjects could communicate with other group members via computer chat.¹¹ The members of the group with the highest average number of correct answers won 1000 ECU each. Other groups received 100 ECU each. Announcement of the winning group took place after a final questionnaire was answered.

After the quiz, in stage 3, subjects were paired with subjects from other groups, identified by group name and player number, and played a one-shot PD, just as in Experiment 1. Afterwards, as before, subjects were shown their partner's choices and the resulting payoffs.

¹⁰ Most sessions consisted of 18 subjects in 6 groups of 3. However, due to no-shows, some sessions had fewer subjects and one session had only 5 groups. Subjects were not explicitly informed of the groups' sizes.

¹¹ We required subjects not to reveal personal information during the chat. All subjects obeyed this instruction. Chat messages were almost exclusively focused on the quiz answers.

Finally, in stage 4, subjects made 6 allocation decisions allocating money between themselves and another participant (the “target”). Three of these decisions were dictator games, with a “taking” frame: the target was endowed with 1000 ECU, and subjects could take up to 1000 ECU from the target for themselves. This dictator game was played once against a member of the PD partner's group (but not the PD partner himself: subjects knew this because a history of play, including player numbers, was visible on screen), once against a member of a third group subjects had not previously any interacted with, and once against a member of the subject's own group. Again, we refer to these groups as the Partner, Other and Own groups. We do not analyse choices against the Own group further in this paper.¹² The remaining 3 decisions were costly punishment games in which subjects chose the number of ECU to deduct from the other subject's payoff. Subjects chose x between 0 and 300; they received $400-x$ and the target received $900-3x$. Again this game was played once against members of the Partner, Other, and Own groups. The order of allocation tasks was randomized.

After the experiment, subjects answered a questionnaire. Then the winning group of the quiz, the result of the stage 1 PD, and the stage(s) chosen for payment were revealed. Finally, subjects were called up and paid privately, receiving in addition a GBP 5 showup fee. For 5 of the 10 sessions one of the four different stages was selected for payment; in the other 5 sessions, either

¹² We intend to examine these decisions in separate work on in-group favoritism.

stages 1 and 2, or stages 3 and 4 were paid. Under the first payment method (selection of one stage only), if the stage 3 PD is paid, then the subject's stage 4 allocations are not paid. This allows us to examine whether others' intentions alone suffice to induce group reciprocity, or whether those intentions must also have material results. We found few consistent differences between behaviour under the two payment methods, so we pool them in the subsequent analysis.¹³ Table 6 shows descriptive statistics for Experiment 2.

<i>N</i>		172	
Demographics			
Male		78 (45.35%)	
Age		Min 17, max 42, mean 20.8, median 20	
Studying:			
... Law	9 (5.23%)	... Natural sciences	29 (16.86%)
... Social sciences	23 (13.37%)	... Other	79 (45.93%)
... Economics	32 (18.60%)	... Not a student	0 (0.00%)
Prisoner's Dilemma choices (self, partner)			
CC	80 (46.51%)	DC	36 (20.93%)
CD	36 (20.93%)	DD	20 (11.63%)
Allocations (ECU allocated to target, min / median / mean / max)			
Dictator, vs Partner	0 / 10 / 241 / 1000	Dictator, vs Other	0 / 6 / 235 / 1000
Punishment, vs Partner	0 / 900 / 762 / 900	Punishment, vs Other	0 / 900 / 763 / 900

TABLE 6: DESCRIPTIVE STATISTICS, EXPERIMENT 2

Figure 2 shows empirical Cumulative Distribution Functions of the amounts given to the other subject, in the “dictator” or “taking” choice, and in the “pun-

¹³ See the Appendix.

ishment” choice, split up by PD partner behaviour. The first row shows amounts given by all subjects. The second row shows amounts given by subjects who themselves cooperated in the PD. Three facts stand out. First, most subjects gave either 500 or nothing in the taking choice, quite typically for dictator games (Engel 2010; see Bardsley 2008 for a dictator game with a taking frame), and most subjects did not punish (i.e. gave 900) in the punishment choice. Second, subjects became more selfish in the dictator choices if their PD partners defected. Third, in punishment choices, cooperators appear to have punished the Partner group more, but the Other group less, if their PD partners defected.

The data is highly non-normal. In the punishment decisions, more than half of our subjects chose to give the maximum amount of 900 ECU. We therefore analyse the data using a Tobit model: $f(x) = \min\{x, 900\}$ in equation (1). We use the semi-parametric estimator defined by Honore (1992), which removes the subject-specific fixed effects by a process analogous to first-differencing, and which does not rely on assumptions about the functional form of the errors $\varepsilon_{i,g}$. The dictator decisions are more complex: many subjects gave 0 ECU, and many gave 500 ECU, with some donations between 0 and 500 and a few above 500. We analysed these using a fixed effects Tobit with $f(x)=\max\{x,0\}$ in (1)¹⁴.

¹⁴ We also tried a fixed effects ordered logit estimator using the approach of Baetschmann et al. (2011), categorizing the amount given as 0, 1-499, 500, or above 500. The results are qualitatively similar to those reported here, and are available on request.

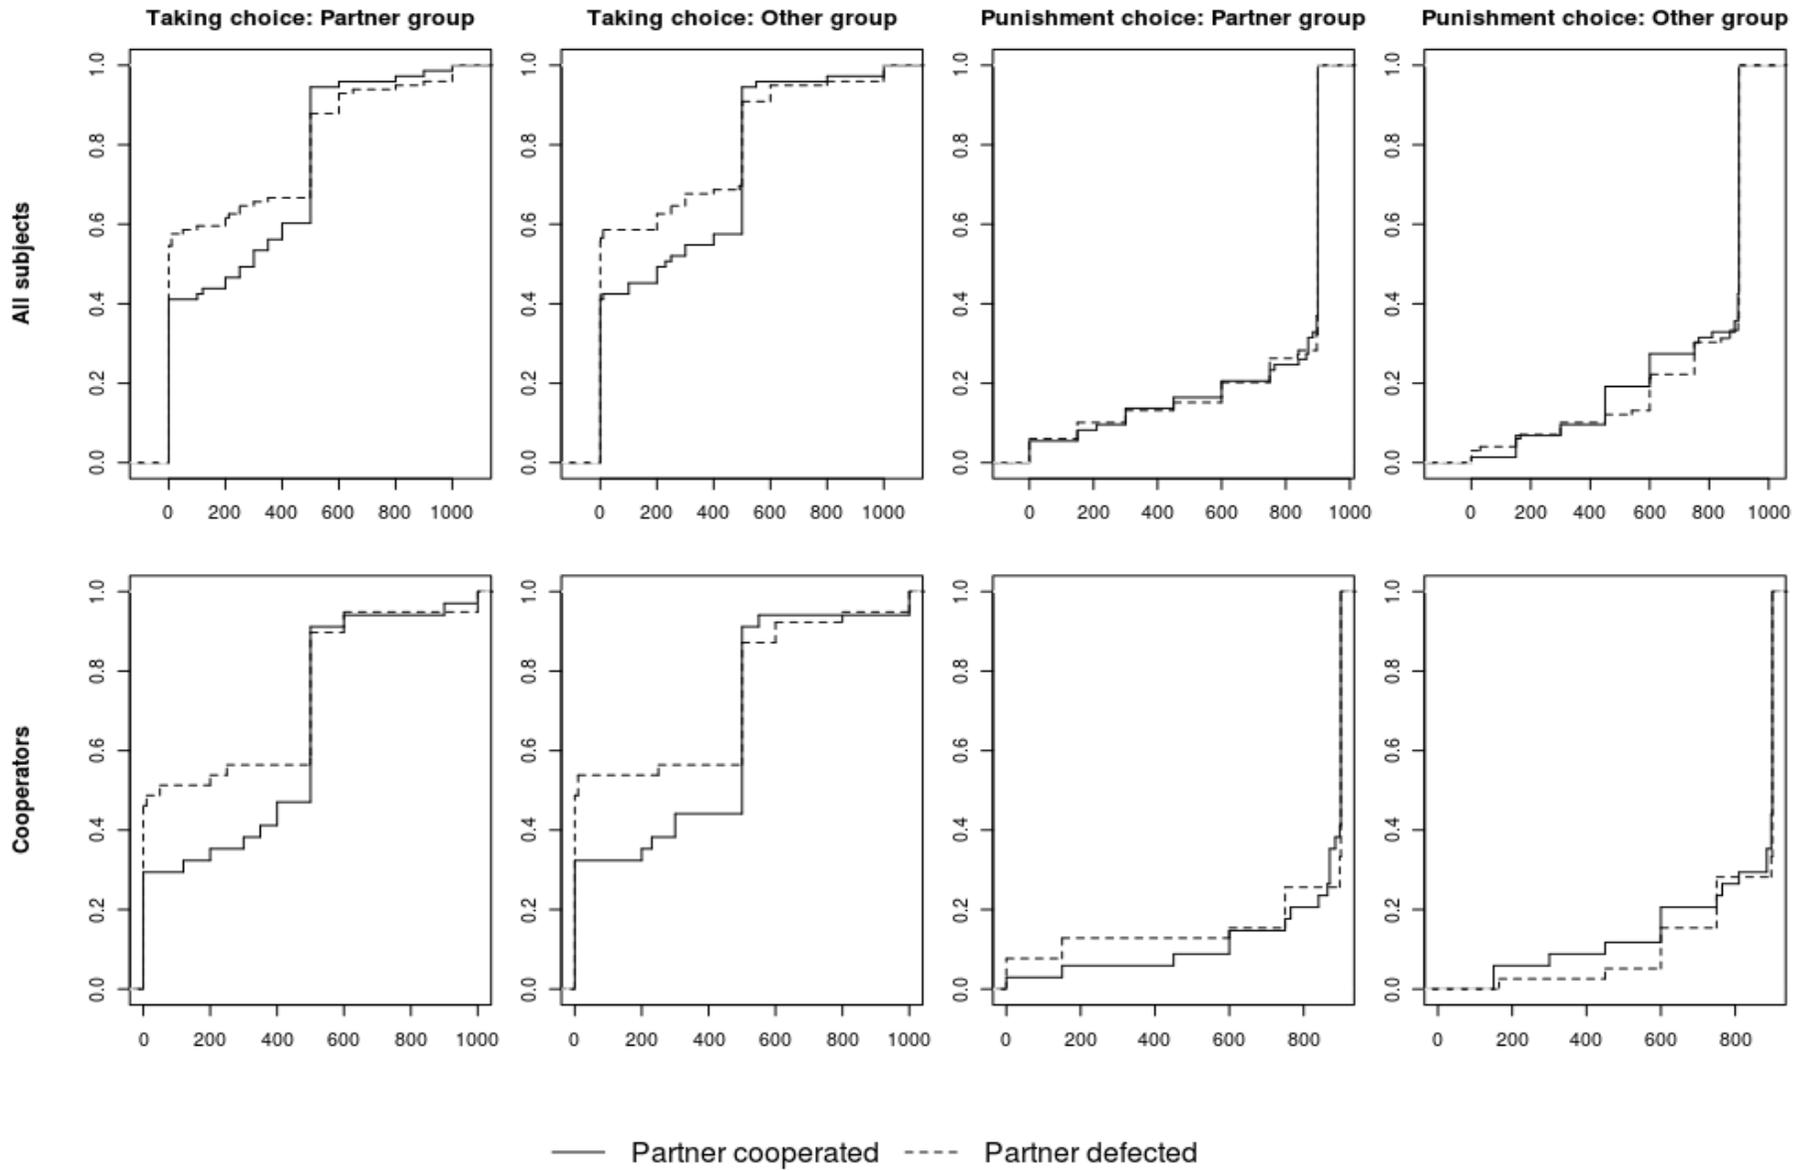


FIGURE 2: EMPIRICAL CDF PLOTS OF AMOUNTS GIVEN, EXPERIMENT 2

	(1)	(2)	(3)	(4)	(5)	(6)
Decision	Punishment	Punishment	Punishment	Dictator	Dictator	Dictator
Pgrp	150 (162.85)	-476.17 (317.49)	-296.56 (206.84)	-33.04 (58.69)	-206.14 (106.46)+	25.4 (141.58)
Pgrp × ...						
Self D	-107.58 (190.17)	-212.23 (181.68)	-91.79 (187.04)	-38.01 (84.74)	11.85 (74.09)	9.35 (79.91)
Partner D	-376.12 (277)	-364.64 (177.36)*	-426.67 (137.63)**	-49.73 (85.42)	26.43 (78)	17.45 (86.36)
Self D × Partner D	443.56 (302.65)	599.8 (216.55)**	391.1 (262.58)	167.11 (125.2)	6.77 (97.39)	30.06 (102.96)
Controls	No	Yes	Yes	No	Yes	Yes
Session dummies	No	No	Yes	No	No	Yes

Controls: Z =(gender, group color, economics student, session % male, Partner group decision taken before Other group decision). Session % male omitted when session dummies included. Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 7: FE TOBIT REGRESSIONS, EXPERIMENT 2

Table 7 shows the results. Columns 1-3 are regressions on the amount given in the punishment decisions. The first column is a simple regression on *Self D*, *Partner D* and their interaction, without controls. The second column adds a vector of controls to Z in equation (1). As before, appropriate controls should increase the efficiency of our estimates. (The Honore estimator does not estim-

ate the fixed effects, nor effects for variables that are constant across a group, so we cannot include controls in X in (1).) Controls include dummies for gender, for subjects who studied economics, and for subjects' color group; in addition we include a measure of the proportion of male subjects in the whole session, since this varied quite widely and we considered that it might affect levels of group reciprocity, and a dummy for the order of the Partner group and Other group decisions. Column 3 adds session fixed effects. The coefficient on PD partner defection is consistently large and negative, and is more precisely estimated when controls are included. The coefficient on the interaction of partner defection and own defection is large and positive; the sum of these two coefficients (i.e. $\zeta + \eta$ in equation (1)) is never significantly different from 0. Columns 4-6 are regressions on the amount given in dictator decisions. The coefficients on both *Partner D* and *Self D* \times *Partner D* are small and not statistically significant. Thus, subjects who cooperated in the PD responded to PD partner defection by giving less to members of the Partner group, but only in punishment decisions, where harming the other participant was costly to the subject.

These results conformed to the basic pattern of Experiment 1: group reciprocity was observed in punishment decisions but not in dictator decisions. In addition, group reciprocity was observed only among those who themselves cooperated in the PD. This finding could be interpreted in two ways. Subjects

may react strongly to betrayal in the PD, but not to mutual defection. Alternatively, subjects who are in general more cooperative may react more to others' defection, whether or not they themselves cooperated in a particular instance (but these subjects are also more likely to cooperate in the PD). To distinguish between these two explanations, Experiment 2 included an initial, anonymous Prisoner's Dilemma without feedback, before subjects were placed into groups. We use subjects' behaviour in this first PD to categorize them as cooperative or non-cooperative "types". Decisions in the first and second PD were correlated: about 50% of cooperators in the first PD defected in the second PD, while about 75% of defectors in the first PD defected again in the second (chi-squared test, $p < 0.01$).

	(1)	(2)	(3)
Decision	Punishment	Punishment	Punishment
Pgrp	161 (186.83)	-521.99 (271.53) +	-441.22 (371.62)
Pgrp × ...			
Self D	-78 (236.31)	-44.88 (183.21)	-32.31 (241.76)
Partner D	-248 (257.7)	-291.5 (208.29)	-195.76 (188.54)
Self D × Partner D	303.85 (308.61)	408.83 (255.08)	162.7 (318.56)
Self D1	-86 (194.64)	132.05 (219.16)	78.63 (194.49)
Self D1 × Self D	-33.76 (258.73)	-237.23 (303.34)	-242.24 (437.06)
Self D1 × Partner D	-427.01 (321.15)	-275.75 (364.69)	-444.21 (366.1)
Self D1 × Self D × Partner D	437.41 (384.82)	324.72 (472.07)	747.73 (592.28)
Summed coefficients:			
Self D + (Self D1 × Self D)	-111.75 (105.34)	-355.36 (207.64) +	-284.26 (229.61)
Partner D + (Self D1 × Partner D)	-675 (191.65) ***	-603.62 (272.32) *	-641.13 (271.71) *
Self D × Partner D + (Self D1 × Self D × Partner D)	741.25 (229.88) **	804.35 (344.13) *	902.42 (352.8) *
Controls	No	Yes	Yes
Session dummies	No	No	Yes

Controls: Z=(gender, group color, economics student, session % male, Partner group decision taken before Other group decision). Session % male omitted when session dummies included. Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 8: FE TOBIT REGRESSIONS, INTERACTIONS WITH INITIAL PD

Table 8 adds *Self DI*, a dummy taking the value 1 if subjects defected in the initial anonymous PD, to the first 3 regressions of Table 7, and interacts it with *Partner D* and *Self D*. If the size of the *Partner D* \times *Self D* interaction in Table 7 is genuinely caused by betrayal in the main PD, then the *Self DI* coefficient should not be significant. On the other hand, if the interaction arises because people with cooperative personalities react more strongly to partner defection, then the interactions with *Self D* should lose significance, and interactions with *Self DI* should be significant instead. In fact, *Self DI* and its interactions are never significant, and, as the “summed coefficient” rows show, the effect of partner defection is consistently larger for subjects who *defected* in stage 1. Thus, behaviour in the initial PD does not help to explain why “betrayal” causes more group reciprocal behaviour than mutual defection. We conclude that only betrayal induces group reciprocity, and that it does so among both more and less cooperative subjects.

Social identity theorists argue that intergroup discrimination is caused by the level of people's identification with their own group – the extent to which they perceive themselves as part of the group and take on its interests as their own (Tajfel et al. 1971; Tajfel 1982). We were interested in whether group identification could help explain group-reciprocal behaviour. In both experiments, the final questionnaire included questions measuring group identification, adapted from Ellemers et al. (1999). We created a variable *Group ID* by sum-

ming answers to these questions.¹⁵ Table 9 reproduces the first 3 columns of Table 7, interacting the other independent variables with *Group ID*. We expected that subjects who showed higher identification with their group would be more likely to group-reciprocate. In fact, while *Group ID* itself predicts less giving to the partner group, its interaction terms are not significant. Indeed, the effect of partner defection is stronger among those with lower *Group ID*. Thus, group identification does not seem to be a useful predictor of group reciprocal behaviour.

¹⁵ Results using an alternative *Group ID* variable, derived from the questionnaire scores by factor analysis, were very similar to those shown here.

	(1)	(2)	(3)
Decision	Punishment	Punishment	Punishment
Pgrp	-68.41 (102.13)	-804.2 (425.85) +	-192.26 (439.18)
Pgrp × ...			
Self D	148.64 (173.06)	171.15 (178.72)	120.69 (250.75)
Partner D	-331.28 (248.02)	-238.14 (329.72)	-257 (240.09)
Self D × Partner D	361.42 (293.75)	288.59 (344.95)	217.99 (369.31)
Group ID	-50.74 (17.61) **	-45.52 (19.36) *	-38.91 (28.34)
Group ID × Self D	30.03 (23.97)	26.56 (21.83)	23.83 (34.51)
Group ID × Partner D	5.27 (41.33)	21.47 (49.08)	12.71 (28.01)
Group ID × Self D × Partner D	14.84 (44.85)	-1.58 (49.61)	1.92 (39.63)
Effect of Partner D at...			
first quartile Group ID	-362.87 (155.12) *	-426.54 (161.31) **	-297.19 (205.57)
third quartile Group ID	-304.95 (432.06)	-178.86 (574.97)	-125.08 (327.21)
Controls	No	Yes	Yes
Session dummies	No	No	Yes

Controls: Z=(gender, group color, economics student, session % male, Partner group decision taken before Other group decision). Session % male omitted when session dummies included. Group ID was centered at its median. Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 9: FE TOBIT REGRESSIONS, INTERACTIONS WITH GROUP IDENTITY,
EXPERIMENT 2

Experiments involving lab-created groups can run the risk of demand effects, in which subjects follow norms implicitly suggested by the experimental design. We took steps to minimize this risk. For example, we randomized the order of allocation decisions, so as to avoid obvious pairings of similar choices made vis-à-vis Partner and Other group members. We also asked subjects about the experiment's purpose in the final questionnaire. Many subjects answered that the experiment was about teamwork; only a few (7 in Experiment 1; 3 in Experiment 2) mentioned e.g. “revenge” or related concepts. Re-running our analyses without these subjects does not affect our results.

3.3 Interpretation

The results above, coming from a new experimental design, need to be interpreted with care. Nevertheless, we here draw some initial conclusions.

Firstly and most importantly: taken together, our experiments provide evidence for group reciprocity. Some subjects responded to their PD partner's defection by reducing their allocations specifically to other members of the partner's group, even at a cost to themselves. Group reciprocity was observed in the absence of strategic incentives, supporting a preferences-based explanation.

Second, group reciprocity did not affect all decisions equally. We saw significant evidence for it only in choices where the Pareto-dominant allocation

would give subjects less than the other participant, but subjects could pay to lower the discrepancy. This is surprising, since we did not observe group reciprocation in the dictator games, where subjects could have punished costlessly (indeed at a profit). One possibility is that group punishment does not obey the law of demand (cf. Anderson and Putterman 2006).

Lastly, not all subjects group-reciprocated. Indeed, the majority of subjects did not discriminate between groups. More cooperative subjects, as measured by an initial anonymous PD, were not more likely to group-reciprocate than less cooperative ones, conditional on their choice to cooperate in the main PD. Also, group reciprocal behaviour was not predicted by psychological measures of group identification, which suggests that a different psychological mechanism is at work than that posited by Social Identity Theory. Other psychological theories not tested here, such as social dominance theory (Sidanius and Pratto 1999), might have more explanatory power.

4 Conclusion

In recent years, evidence from laboratory and field experiments has challenged the motivational assumption of material self-interest, introducing concerns such as fairness and reciprocity. One so far under-emphasized factor, we believe, is that humans, who often act in groups, often also treat other groups as actors. They may therefore assign blame and praise to groups as well as indi-

viduals – with important consequences for human society and politics. At times in history, entire groups have been blamed and scapegoated for their members' real or imagined behaviour. Laboratory experiments could help in understanding the processes behind this phenomenon. In our experiments, subjects responded to one group member's behaviour by treating other members of that group differently.

Discriminatory behaviour may be driven by beliefs rather than by preferences. For example, ethnic discrimination in hiring may come from statistical or screening discrimination (Arrow 1972) rather than from a “taste for discrimination” (Becker 1957). In our experiments, strategic incentives were ruled out: group reciprocity appears to be a matter of pure preference. However, we caution against too narrow an interpretation of this. Group reciprocity may be due to the activation of context-specific norms, rather than to “preferences” thought of as unchanging characteristics of individuals.

The majority of our subjects did not exhibit group reciprocity, and it was not visible in all decisions. To investigate the causes and correlates of group reciprocity, our experimental paradigm needs further refinement. Complementary experiments could be done using homegrown identities such as ethnicity. Lastly, we have still to investigate third-party group reciprocity, in which people retaliate against a group for harm done not to themselves, but to a fellow group member of theirs.

Demonstrating group reciprocity in a reproducible laboratory experiment could open the door to deeper exploration of how it works. We draw an analogy with the minimal group paradigm in psychology. After Sherif's initial research in the 1950s on in-group prejudice, the famous Robber's Cave experiment (Sherif 1961), researchers developed a canonical experimental paradigm to analyse in-group prejudice (Tajfel et al. 1971). By making the behaviour of interest reproducible, the minimal group paradigm catalysed a productive tradition of research on prejudice. We hope that our experiment will similarly encourage further study of this phenomenon under controlled conditions. This could be a valuable complement to field study of the often extreme situations where group reciprocity produces its most visible effects.

Bibliography

- Abbink, K., Herrmann, B., 2009. Pointless Vendettas. Manuscript.
- Akerlof, G., 1976. The economics of caste and of the rat race and other woeful tales. *The Quarterly Journal of Economics* 90, 599-617.
- Anderson, C. M, Putterman, L., 2006. Do non-strategic sanctions obey the law of demand? The demand for punishment in the voluntary contribution mechanism. *Games and Economic Behavior* 54, 1-24.
- Arrow, K. J., 1972. Some mathematical models of race discrimination in the labor market. In: Pascal, A. H. (Ed.), *Racial Discrimination in Economic Life*. Lexington: Lexington Books, 187-203.
- Bardsley, N., 2008. Dictator game giving: altruism or artefact?. *Experimental Economics* 11, 122-133.
- Becker, G. S., 1957. *The Economics of Discrimination*. University of Chicago Press.
- Baetschmann, G., Staub, K.E., Winkelmann, R., 2011. Consistent estimation of the fixed effects ordered logit model. IZA Discussion Paper 5443.
- Bernhard, H., Fischbacher, U., Fehr, E., 2006. Parochial altruism in humans. *Nature* 442, 912-5.
- Blattman, C., Miguel, E., 2010. Civil war. *Journal of Economic Literature* 48, 3-57.
- Bolton, G. E., Ockenfels, A., 2000. ERC: A theory of equity, reciprocity, and competition. *The American Economic Review* 90, 166-193.

- Bornstein, G., 1992. The free-rider problem in intergroup conflicts over step-level and continuous public goods. *Journal of Personality and Social Psychology* 62, 597.
- Bornstein, G., 2003. Intergroup conflict: individual, group, and collective interests. *Personality and Social Psychology Review* 7, 129-145.
- Bouckaert, J., Dhaene, G., 2004. Inter-ethnic trust and reciprocity: results of an experiment with small businessmen. *European Journal of Political Economy* 20, 869-886.
- Brewer, M. B., 1979. Ingroup bias in the minimal intergroup situation: a cognitive-motivational analysis. *Psychological Bulletin* 86, 307-324.
- Brewer, M. B., 1999. The psychology of prejudice: ingroup love or outgroup hate? *Journal of Social Issues* 55, 429-444.
- Chen, Y., Li, S. X., 2009. Group identity and social preferences. *American Economic Review* 99, 431-457.
- Collier, P., Hoeffler, A., 2004. Greed and grievance in civil war. *Oxford Economic Papers* 56, 563 -595.
- De Cremer, D., van Vugt, M., 1999. Social identification effects in social dilemmas: a transformation of motives. *European Journal of Social Psychology* 29, 871-893.
- Dufwenberg, M., Gneezy, U., Güth, W., Van Damme, E., 2001. Direct vs indirect reciprocity: an experiment. *Homo Oeconomicus* 18, 19-30.
- Ellemers, N., Spears, R., Doosje, B., 1997. Sticking together or falling apart: In-group identification as a psychological determinant of group commitment versus individual mobility. *Journal of Personality and Social Psychology* 72, 617.

- Ellemers, N., Kortekaas, P., Ouwerkerk, J. W., 1999. Self-categorisation, commitment to the group and group self-esteem as related but distinct aspects of social identity. *European Journal of Sococial Psychology* 29, 371-389.
- Engel, C., 2010. Dictator games: a meta study. Preprints of the Max Planck Institute for Research on Collective Goods Bonn 7.
- Fearon, J. D., Laitin, D. D., 2003. Ethnicity, insurgency, and civil war. *The American Political Science Review* 97, 75-90.
- Fehr, E., Fischbacher, U., 2004. Third-party punishment and social norms. *Evolution and Human Behavior* 25, 63-87.
- Fehr, E., Gächter, S., 2000. Fairness and retaliation: the Economics of reciprocity." *Journal of Economic Perspectives* 14, 159-181.
- Fershtman, C., Gneezy, U., 2001. Discrimination in a segmented society: an experimental approach. *Quarterly Journal of Economics* 115, 351- 377.
- Fischbacher, U., 2007. z-Tree: turich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171-178.
- Gaertner, L., Iuzzini, J., O'Mara, E. M., 2008. When rejection by one fosters aggression against many: multiple-victim aggression as a consequence of social rejection and perceived groupness. *Journal of Experimental Social Psychology* 44, 958-970.
- Gaertner, L., Schopler, J., 1998. Perceived ingroup entitativity and intergroup bias: an interconnection of self and others. *European Journal of Social Psychology* 28, 963-980.
- Goette, L., Huffman, D., Meier, S., 2006. The impact of group membership on cooperation and norm enforcement: evidence using random assignment to real social groups. *American Economic Review* 96, 212-216.

- Greiner, B., 2004. The online recruitment system ORSEE 2.0 – a guide for the organization of experiments in economics. University of Cologne, Working Paper Series in Economics 10.
- Guala, F., Mittone, L., Ploner, M., 2009. Group membership, team preferences, and expectations. University of Trento, CEEL Working Papers 0906.
- Habyarimana, J., Humphreys, M., Posner, D., 2009. Coethnicity: Diversity and the Dilemmas of Collective Action (Russell Sage Foundation Series on Trust). New York: Russell Sage Foundation Publications.
- Halevy, N., Bornstein, G., Sagiv, L., 2008. 'In-group love' and 'out-group hate' as motives for individual participation in intergroup conflict. *Psychological Science* 19, 405-411.
- Hargreaves Heap, S., Zizzo, D. J., 2009. The value of groups. *American Economic Review* 99, 295-323.
- Honoré, B. E., 1992. Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica: Journal of the Econometric Society*, 533-565.
- Horowitz, D. L., 1985. *Ethnic Groups in Conflict*. Berkeley: University of California Press.
- Horowitz, D. L., 2001. *The Deadly Ethnic Riot*. University of California Press.
- Kocher, M. A., Pepinsky, T. B., Kalyvas, S. N., 2008. Into the arms of the rebels? Aerial bombardment, indiscriminate violence, and territorial control in the Vietnam War. Paper prepared for the ISA's 49th annual convention, San Francisco.
- Krueger, A. B, MaLeckova, J., 2003. Education, poverty and terrorism: is there a causal connection? *The Journal of Economic Perspectives* 17, 119-144.

- Lickel, B., Miller, N., Stenstrom, D. M., Denson, T. F., Schmader, T., 2006. Vicarious retribution: the role of collective blame in intergroup aggression. *Personality and Social Psychology Review* 10, 372.
- Mamdani, M., 2001. *When Victims Become Killers: Colonialism, Nativism, and the Genocide in Rwanda*. Princeton: Princeton University Press.
- Moreno, A. T., 2008. Group fairness and game theory. Universidad de Guanajuato, School of Economics Working Papers.
- Mummendey, A., Otten, S., 1998. Positive–negative asymmetry in social discrimination. *European Review of Social Psychology* 9, 107-143.
- Mummendey, A., Simon, B., Dietze, C., Grünert, M., Haeger, G., Kessler, S., Lettgen, S., Schäferhoff, S., 1992. Categorization is not enough: intergroup discrimination in negative outcome allocation. *Journal of Experimental Social Psychology* 28, 125-144.
- Nowak, M. A., Roch, S., 2007. Upstream reciprocity and the evolution of gratitude. *Proceedings of the Royal Society B: Biological Sciences* 274, 605-610.
- Shayo, M., Zussman, A., 2011. Judicial ingroup bias in the shadow of terrorism. *The Quarterly Journal of Economics* 126, 1447-1484.
- Sherif, M., Harvey, O. J., White, J., Hood, W. R., Sherif, C. W., 1961. *The Robbers Cave experiment: Intergroup Conflict and Cooperation*, Scranton: Wesleyan University Press.
- Sidanius, J., Pratto, F., 1999. *Social Dominance: An Intergroup Theory of Social Hierarchy and Oppression*. Cambridge: Cambridge University Press.
- Stenstrom, D. M., Lickel, B., Denson, T. F., Miller, N., 2008. The roles of in-group identification and outgroup entitativity in intergroup retribution. *Personality and Social Psychology Bulletin* 34, 1570-1582.

Sumner, W. G., 1906. *Folkways: A Study of the Sociological Importance of Usages, Manners, Customs, Mores, and Morals*. Boston: Ginn.

Tajfel, H., 1982. Social psychology of intergroup relations. *Annual Review of Psychology* 33, 1-39.

Tajfel, H., Billig, M. G., Bundy, R. P., Flament, C., 1971. Social categorization and intergroup behavior. *European Journal of Social Psychology* 1, 149-178.

Yamagishi, T., Kiyonari, T., 2000. The group as the container of generalized reciprocity. *Social Psychology Quarterly* 63, 116-132.

Yamagishi, T., Mifune, N., 2008. Does shared group membership promote altruism? *Rationality and Society* 20, 5-30.

Yamagishi, T., Mifune, N., 2009. Social exchange and solidarity: in-group love or out-group hate?. *Evolution and Human Behavior* 30, 229-237.

Appendix: further statistical analyses

The experiment questionnaire asked if subjects had met any of the other participants before. Some subjects (39) answered that they had. These subjects might have had strategic motivations for behaving in a certain way if they expected to meet these subjects again. We are not very concerned about this, because all decisions were anonymous and subjects' roles were assigned ran-

domly. Nevertheless, Table 10 reruns the basic regressions of Table 7, excluding all these subjects. Results are basically unchanged.

Decision	Punishment	Punishment	Punishment	Dictator	Dictator	Dictator
Pgrp	173.5 (186.34)	-677.06 (248.01) **	-368.19 (317.93)	56.67 (40.49)	-184.6 (90.92) *	-16.21 (139.17)
Pgrp × ...						
Self D	-93.19 (210.62)	-83.59 (179.9)	-108.3 (176.24)	-11.28 (81.95)	3.91 (61.36)	-4.5 (62.28)
Partner D	-282.36 (253.45)	-405.56 (176.97) *	-313.14 (142.85) *	-141.67 (79.94) +	-38.59 (64.2)	-58.69 (69.85)
Self D × Partner D	281.85 (285.49)	419.38 (215.59) +	357.27 (232.29)	146.84 (138.3)	51.2 (96.59)	117.16 (112.22)
Controls	No	Yes	Yes	No	Yes	Yes
Session dummies	No	No	Yes	No	No	Yes
N	133	133	133	133	133	133

Controls: Z=(gender, group color, economics student, session % male, Partner group decision taken before Other group decision). Session % male omitted when session dummies included. Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 10: FE TOBITS EXCLUDING SUBJECTS WHO HAD MET OTHER PARTICIPANTS, EXPERIMENT 2

Sessions 6-10 of Experiment 2 used a different payment method, paying either both stages 1 and 2, or both stages 3 and 4 (the PD and the allocations). Table 11 adds a dummy *Payboth* to columns 1, 2, 4 and 5 of Table 7, and interacts it with *Self D*, *Partner D* and their interaction. Results are not very clear. The *Payboth* dummy interacts significantly with *Partner D* in the regression on

punishment choices in column 1, but the coefficient is much smaller and insignificant when controls are added. There is also a significant interaction in the controlled regression on dictator choices (column 4). However, neither the coefficient on Partner D itself, nor the summed coefficient $\text{Partner D} + \text{Payboth} \times \text{Partner D}$, is significantly different from zero. Overall, we think it reasonable to pool the sessions for analysis, as in the main text.

	(1)	(2)	(3)	(4)
Decision	Punishment	Punishment	Dictator	Dictator
Pgrp	183.81 (219.9)	-473.66 (279.66) +	-91.66 (98.92)	-308.2 (181.86) +
Pgrp × ...				
Self D	-12.6 (254.35)	-36.81 (221.53)	47.92 (134.25)	133.93 (141.62)
Partner D	-735.23 (282.64) **	-415.15 (265.28)	-7.08 (147.13)	205.59 (142.73)
Self D × Partner D	687.43 (341) *	416.67 (315.26)	144.33 (194.38)	-150.3 (161.18)
Payboth	-123.81 (229.2)	-61.5 (286.23)	160.14 (106.86)	120.86 (162.02)
Payboth × Self D	-253.09 (278.98)	-154 (348.57)	-216.4 (162.4)	-187.51 (148.29)
Payboth × Partner D	678.23 (329.11) *	77.73 (395.97)	-131.38 (167.77)	-303.39 (139.4) *
Payboth × Self D × Partner D	-324.16 (398.28)	211.31 (472.67)	56.15 (239.47)	224.67 (181.93)
Controls	No	Yes	No	Yes

Controls: Z =(gender, group color, economics student, session % male, Partner group decision taken before Other group decision). Standard errors in parentheses. + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

TABLE 11: FE TOBITS, INTERACTION WITH PAYMENT METHOD, EXPERIMENT 2