

Resistance to extreme strategies, rather than prosocial preferences, can explain human cooperation in public goods games

Rolf Kümmerli^{a,1}, Maxwell N. Burton-Chellew^{b,1}, Adin Ross-Gillespie^{a,c}, and Stuart A. West^{b,2}

^aInstitute of Evolutionary Biology, School of Biological Sciences, University of Edinburgh, Edinburgh EH9 3JT, United Kingdom; ^bDepartment of Zoology, University of Oxford, Oxford OX1 3PS, United Kingdom; and ^cDepartment of Ecology and Evolution, Biophore, University of Lausanne, CH-1015 Lausanne, Switzerland

Edited by Raghavendra Gadagkar, Indian Institute of Science, Bangalore, India, and approved April 15, 2010 (received for review January 21, 2010)

The results of numerous economic games suggest that humans behave more cooperatively than would be expected if they were maximizing selfish interests. It has been argued that this is because individuals gain satisfaction from the success of others, and that such prosocial preferences require a novel evolutionary explanation. However, in previous games, imperfect behavior would automatically lead to an increase in cooperation, making it impossible to decouple any form of mistake or error from prosocial cooperative decisions. Here we empirically test between these alternatives by decoupling imperfect behavior from prosocial preferences in modified versions of the public goods game, in which individuals would maximize their selfish gain by completely (100%) cooperating. We found that, although this led to higher levels of cooperation, it did not lead to full cooperation, and individuals still perceived their group mates as competitors. This is inconsistent with either selfish or prosocial preferences, suggesting that the most parsimonious explanation is imperfect behavior triggered by psychological drives that can prevent both complete defection and complete cooperation. More generally, our results illustrate the caution that must be exercised when interpreting the evolutionary implications of economic experiments, especially the absolute level of cooperation in a particular treatment.

altruism | economic games | inclusive fitness | psychological drives | utility

There is a large empirical literature suggesting that when humans play anonymous one-shot economic games, they cooperate more than would be expected if they were purely self-interested (1–4). A common example comes from the public goods game, in which individuals are arranged in groups of N members, and each is given some quantity of monetary units that they can contribute to a public project. The contributions of all members are summed and multiplied by b (where $b > 1$), then equally divided between the group members. As long as $n > b$, an individual will always gain a higher economic reward in the short term by restricting or withholding their own contributions to the public project, while still benefiting from the contributions made by others. Although this favors zero contribution, individuals still contribute to the public project, and although they contribute less over time, $\approx 10\%$ of individuals continue to contribute in the long run (1).

A highly influential body of research has investigated both the proximate mechanisms that lead to this cooperation and the evolutionary forces that would have selected for them (2, 3, 5–11). From a proximate perspective, it has been argued that the mechanism leading to this cooperation is that individuals value the success of others as well as their own (2). In economics, this is described as utility including prosocial preferences, such as other-regarding preferences or inequality aversion. Considering evolutionary (ultimate) explanations, it has been argued that games with anonymous one-shot encounters cannot be explained by commonly invoked explanations for cooperation between nonrelatives, such as reciprocity and reputation. In response to this, “strong reciprocity”

evolutionary theories have been developed that are suggested to be able to explain the origin of such cooperation (3, 6–10, 12).

However, this body of research is based on the implicit assumption that imperfect behavior by individuals in economic games would not lead to a systematic bias in the level of cooperation. Because the predicted behavior is to contribute nothing to the public good, then any deviations from perfection are automatically perceived as greater-than-expected cooperation (13). Put another way, the experimental design makes it impossible to decouple any form of mistake or error (in the context of the experimental setting) from cooperative decisions. If higher-than-expected levels of cooperation are caused at least partially by imperfect behavior, then this could lead to the development of inappropriate utility functions and an overinterpretation of the evolutionary implications of data from economic games.

Here we empirically test whether the results of previous economic games are due to an intrinsic predisposition toward cooperation, driven by prosocial preferences, or whether they reflect imperfect behavior, which could result from underlying psychological drives. To distinguish between these possibilities, we place individuals in modified versions of the public goods game, in which they would maximize their selfish (strategic) gain by completely (100%) cooperating. This aligns the interest of the individual and others in the group, and so 100% cooperation would be predicted both with and without prosocial preferences. Consequently, if intermediate ($<100\%$) levels of cooperation are observed, then this is inconsistent with either selfish or prosocial preferences, suggesting that the most parsimonious explanation is imperfect behavior triggered by psychological drives. Another way of conceptualizing our experiments is that they provide the appropriate control treatments, which were lacking from previous experiments, in which mistakes do not automatically lead to higher-than-predicted levels of cooperation.

Results

We carried out four independent experiments involving a total of 168 subjects, each comparing behavior in a standard public goods game (control treatment) with that in a modified game, in which competition within groups was either relaxed or completely repressed, such that 100% cooperation was favored. All games were played in groups of four subjects, with four to five such groups playing the game simultaneously, leading to 16 or 20 subjects per session. Subjects received detailed instructions, computer-animated examples, had to complete a comprehension

Author contributions: R.K., M.N.B.-C., A.R.-G., and S.A.W. designed research; R.K., M.N.B.-C., A.R.-G., and S.A.W. performed research; R.K., M.N.B.-C., A.R.-G., and S.A.W. analyzed data; and R.K., M.N.B.-C., A.R.-G., and S.A.W. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹R.K. and M.N.B.-C. contributed equally to this work.

²To whom correspondence should be addressed. E-mail: stuart.west@zoo.ox.ac.uk.

test, and performed a practice round for both treatments, which they played in a consecutive order with reversed treatment orders between sessions. Subjects received 40 credits at the beginning of each experimental round and were asked to decide how many of these credits to contribute to a public project. The games were completely anonymous and involved 6 to 10 rounds of interactions with other group members, over which group composition changed randomly.

Modified Public Goods Game. In our first experiment, we relaxed within-group competition by multiplying all contributions to the public project by a factor greater than the number of players in each group ($b = 5$), so that subjects received 1.25 credits for each credit contributed. The strategy that maximizes economic gains in this case is full cooperation, with 100% contribution to the public project. This game contrasts with the standard public goods game, with $b = 2$, which favors no contribution because subjects only receive 0.5 credits for each credit contributed. As predicted, we found that contributions, and therefore the mean level of cooperation, were higher in the public goods game with $b = 5$ than in the standard public goods game with $b = 2$ [Fig. 1A; linear mixed model (LMM): $t_{63} = 5.70$, $P < 0.0001$]. The standard public goods game showed the usual pattern, with the contributions starting at an intermediate level and then significantly decreasing over the 10 rounds of play [Fig. 1A; generalized linear mixed model (GLMM): $t_{575} = -13.05$, $P < 0.0001$].

However, contrary to expectation, the contributions in the public goods game with $b = 5$ did not significantly increase over the 10 rounds of the game (GLMM: $t_{575} = -1.16$, $P = 0.25$) and remained significantly lower than 100% (Fig. 1A), with only 44% of subjects contributing fully in the last round [95% bootstrap confidence interval (CI): 31–56%]. In addition, we found a symmetry in our results, with the proportion of subjects behaving imperfectly in the public goods game with $b = 5$ (56% withholding from full cooperation) not differing significantly from the proportion of subjects behaving imperfectly in the standard public goods game with $b = 2$ (50% withholding from full defection; Fisher's exact test: $P = 0.60$).

The level of cooperation reflected players' perception of other members of their group. After each game, subjects were asked to score, on a sliding scale, how they perceived their group mates, whereby zero denoted "full collaborators" and 20 denoted "full competitors". Although this score was significantly lower in the public goods game with $b = 5$ (Fig. 2; GLMM: $t_{63} = 11.21$, $P < 0.0001$), it was still significantly greater than zero, with 92% (95% CI: 84–98%) of subjects perceiving their group mates as competitors to some degree rather than full collaborators. Furthermore, there was no evidence for a significant difference in the level of imperfect scoring between treatments (Fisher's exact test: $P = 0.27$), comparing the proportions of subjects scoring their group mates imperfectly in the public goods game with $b = 5$ (not full collaborators = 92%) and the standard public goods game (not full competitors = 84%).

Other Methods for Repressing Competition Within Groups. To test the robustness of our results we carried out three further experiments, in which competition within groups was repressed in different ways, all of which favored 100% cooperation. In all three of these experiments, within-group competition was repressed by introducing group rewards according to the total group contribution (14, 15). Specifically, we calculated the total credits contributed by each group and rewarded each group member depending upon the contribution of his or her group in relation to the other groups in the same session.

In the second experiment, group rewards were added on top of the rewards from the standard public goods game and were chosen in such a way that the return to the individual from the group reward strongly outweighed the cost of contributing to the




Fig. 1. Repression of competition does not lead to complete cooperation in humans. Shown are the mean levels of contribution ($\pm 95\%$ CI) in the different rounds of the four experiments (Exp. 1 to Exp. 4). The open circles show the standard public goods game (control treatment), whereas the filled squares show results from modified games, in which competition among group members was repressed. Although significantly higher levels of contribution were observed in the repression-of-competition treatments, the contributions remained significantly below 100% in all rounds (i.e., none of the confidence intervals includes 100%).

public project. This was reflected by a significant positive correlation between the subjects' rewards and their contribution to the public project (Pearson correlation: $r = 0.41$, $n = 32$, $P = 0.019$), with the optimal strategy to maximize selfish gain in this treatment being again 100% contribution. In the third and fourth experiments, all credits not contributed to the public good were lost, and so the group reward was the only reward a subject could gain, favoring 100% contribution. The third and fourth experiments differed only from one another with respect to whether they featured a comprehension test. We omitted the comprehension test in the third experiment because it could be argued




Fig. 2. Competitors and collaborators. Shown are the extent to which subjects perceived their group mates as competitors or collaborators on a scale from zero (full collaborator) to 20 (full competitor). Data are presented for the two experimental treatments: standard public goods game (control) and the public goods game in which competition among group members was repressed (white bars, experiment 1; dark gray bars, experiment 2; black bars, experiment 3; light gray bars, experiment 4). Even when competition is fully repressed within groups, individuals perceived their group mates not as full collaborators but as competitors to some degree.

to instruct the subjects how to play in this extreme game, biasing their behavior toward 100% contribution. We included the comprehension test in our fourth experiment because if subjects fail to contribute fully even in this biased situation, this would be the strongest support for the idea that humans are not predisposed to cooperate above the predicted level.

The results of the second to fourth experiments supported the conclusions of our first experiment. In all cases, repression of competition led to higher levels of cooperation compared with the standard public goods game (Fig. 1 *B–D*; LMM; second experiment: $t_{31} = 10.36$, $P < 0.0001$; third experiment: $t_{39} = 7.99$, $P < 0.0001$; fourth experiment: $t_{31} = 17.35$, $P < 0.0001$), but the level of cooperation still remained significantly lower than 100% (Fig. 1 *B–D*). In the standard public goods game, the level of cooperation significantly decreased over time (GLMM; second experiment: $t_{159} = -9.83$, $P < 0.0001$; third experiment: $t_{199} = -8.48$, $P < 0.0001$; fourth experiment: $t_{159} = -6.85$, $P < 0.0001$). In contrast, in the repression-of-competition treatments, the level of cooperation significantly increased over time in the second and third experiments (GLMM; second experiment: $t_{159} = 6.55$, $P < 0.0001$; third experiment: $t_{199} = 4.91$, $P < 0.0001$) and showed no significant change in the fourth experiment (GLMM: $t_{159} = 1.47$, $P = 0.14$).

The proportion of subjects contributing fully in the repression-of-competition treatment was significantly lower than 100% and averaged 59% (95% CI: 44–75%), 35% (20–50%), and 72% (56–88%) in the last round of the second, third, and fourth experiments, respectively. In addition, we found a symmetry in our results, with the proportion of subjects behaving imperfectly in the repression-of-competition treatments (withholding from full cooperation) not differing significantly from the proportion of subjects behaving imperfectly in the standard public goods game (withholding from full defection) in the second and third experiments (second experiment: 41% vs. 56%, Fisher's exact test: $P = 0.32$; third experiment: 65% vs. 67%, $P = 1.00$) but was significantly lower in the fourth experiment (28% vs. 78%, $P = 0.0001$).

Although individuals were significantly more likely to perceive group members as collaborators in the repression-of-competition treatments than in the standard public goods games (Fig. 2; GLMM; second experiment: $t_{31} = 13.32$, $P < 0.0001$; third experiment: $t_{39} = 19.10$, $P < 0.0001$; fourth experiment: $t_{31} = 25.94$, $P < 0.0001$), a significantly greater than zero proportion of individuals still perceived group members as competitors to some

degree [second experiment: 94% (95% CI: 84–100%); third experiment: 93% (83–100%); fourth experiment: 66% (50–81%)]. Furthermore, the proportions of subjects scoring their group mates imperfectly did not differ between the treatments with repressed competition and the standard public goods games (second experiment: 94% vs. 78%, $P = 0.15$; third experiment: 93% vs. 75%, $P = 0.07$, marginally significant but in the opposite than predicted direction; fourth experiment: 66% vs. 78%, $P = 0.40$).

Discussion

In our economic games in which competition between group members was reduced or completely removed, individuals would maximize their economic gain by contributing 100% to the public project. Contrary to this expectation, we found that (i) the mean level of contribution was significantly lower than 100% (Fig. 1), with 28–65% of subjects not contributing fully in the last round; and (ii) a large proportion (66–94%) of individuals still perceived their group mates as competitors to some degree (Fig. 2). Furthermore, we found a symmetry with the results of standard public goods games, from both our and previous experiments. Specifically, there was no significant difference in the level of imperfect behavior across our two treatments, examining both the proportion of subjects who contributed an amount that does not maximize their economic gain (three out of four experiments), and the proportion of subjects that incorrectly perceived their group members as competitors or collaborators (seven of eight cases examined). This suggests problems for both the mechanistic and evolutionary implications that have been inferred from previous experiments.

Mechanism and Utility. Our results contradict the argument that the higher-than-expected levels of cooperation are explained by utility functions that involve prosocial preferences. In our experimental treatments in which competition within groups was removed, the interests of the individual, the other members of the group, and the group as a whole are all maximized by complete cooperation. Consequently, if we used the same logic as has previously led to arguments for prosocial preferences to define a utility function that describes the deviations from expected levels of cooperation, it would give a utility function that is negatively influenced by the success of others (an antisocial preference). Clearly a simultaneous positive and negative regard to others is not possible.

Instead, our results suggest that players have a utility that avoids both full defection and full cooperation. This could have resulted from a psychology that avoids extreme behaviors that would be very costly when wrong (bet hedging), avoids irrevocable actions, or that uses a simple rule of thumb that “misfires” and causes errors in the extreme situations of some laboratory experiments (16, 17). The potential explanation that our results just reflect a lack of understanding of the game, or insufficient time to learn, were ruled out by the implementation of comprehension tests, such that individuals could only play once they had demonstrated a full understanding of the game rules. Furthermore, the resistance to full cooperation is demonstrated by our observation that individuals who did not contribute fully usually retained only a small number of credits in all experiments (Fig. 3), which suggests that subjects understood that high contributions were worthwhile but psychological drives prevented full contribution.

Evolutionary Implications. Our results stress the need for caution when interpreting the evolutionary implications of economic games. A general point here is the distinction between what can be inferred from comparisons across experimental treatments, and the absolute behavior in a single treatment (ref. 18, p. 123; ref. 19, p. 369). Qualitative comparisons across treatments have shown that factors such as repeated interactions, punishment,




Fig. 3. Individual contributions to the public project when within-group competition is repressed. Bars show distributions of mean contributions across the four experiments (Exp 1. to Exp. 4). For this analysis, only mean contributions from the last three rounds of each game were considered, to exclude potential noise due to learning effects in early rounds of the game. Most subjects who did not contribute fully tended to retain small numbers of credits, and thus the population of subjects cannot be split into full cooperators (contribute 40) and unconditional defectors (zero contribution).

and competition between groups can be important (1, 9, 14, 15, 20) (Fig. 1). However, note that this shows that humans can respond to such factors, which is different from proving they are important in the real world (21). Comparisons across treatments have also elaborated the role of subconscious cues that could indicate factors such as being watched or group identity, suggesting that these do play a role in natural conditions (22–24).

In contrast, less can be inferred from the absolute level of cooperation in single treatments, such as the occurrence of some cooperation in one-shot interactions of the standard public goods game. One reason is that multiple explanations can always be given as to why a psychological drive to cooperate had arisen. Specifically, any mechanism that provides a direct (to self) or indirect (to others) benefit to cooperation, such as competition between groups, interactions between relatives, and repeated interactions allowing reciprocity or punishment, can potentially account for such drives (25–39). Furthermore, as discussed earlier, there is also the problem that errors can lead to systematic biases, thereby influencing the absolute level of cooperation. Consequently, although our results do not support the strong reciprocity hypothesis, they also do not support alternatives, such as a drive that had arisen in response to a specific factor like reputation (22). Instead, our results challenge the need for an evolutionary explanation that it is specific to one-shot games, and stress that mechanisms hypothesized to explain cooperation in standard public goods games must also explain the lack of cooperation in our experiments.

The above discussion has emphasized the role of underlying psychological drives in explaining the level of cooperation observed in one-shot anonymous games. These explanations, which rely on mechanisms that have been selected outside of the laboratory setting (and could involve genes and/or culture), are the only possible explanation for such cooperation (4, 13, 22, 29, 30, 36, 40, 41). The reason for this is that all evolutionary models of cooperation rely on cooperation providing some direct (to self) or indirect (to others) benefit, because natural selection leads to behaviors that maximize their sum (42). This includes models of strong reciprocity, which—although they are argued to explain cooperation in one-shot encounters—actually rely on population structures leading to interactions between relatives (yielding indirect benefits) and/or group competition (yielding direct ben-

efits to members in cooperative groups) (34–36, 43). This is not to say that factors such as competition between groups or punishment of noncooperators (or any other selective force) do not help explain cooperation in humans, just that if they are important, then this leads to the overall costs and benefits of cooperation differing from that assumed in standard anonymous one-shot public goods games.

Conclusion

We conclude with two general points. First, our results suggest that the higher-than-expected levels of cooperation that have been previously observed in one-shot anonymous public goods games are more easily explained by imperfect behavior than a prosocial preference. This does not mean that there are no situations in which humans show prosocial preferences. However, it does emphasize that if prosocial preferences are to be demonstrated experimentally, appropriate control treatments are required. A key future step is to determine exactly why individuals behave imperfectly. Second, our results illustrate why caution must be exercised when interpreting the evolutionary implications of economic experiments. Although qualitative comparisons across treatments can demonstrate the potential role of different possible explanations for cooperation, numerous explanations can be given for the absolute level of cooperation in a particular treatment.

Methods

Experimental Design. A total of 168 undergraduate students (96 women and 72 men) participated voluntarily in a series of economic games on public goods cooperation. Informed consent was obtained from all subjects. The first experiment with 64 subjects was carried out at the Centre for Experimental Social Sciences at Nuffield College in Oxford. The subjects in this experiment were undergraduate students from different disciplines at the University of Oxford who were recruited using ORSEE (44). The other three experiments with 104 subjects were carried out at the University of Edinburgh with undergraduate biology students. In all experiments, we compared the same control treatment (a standard public goods game) with a modified version of a public goods game, in which competition between group members was relaxed or completely repressed. Both treatments were played in groups of four subjects, with multiple groups making up a session playing the game simultaneously. Subjects played the games on computer terminals, the screens of which were shielded from others to guarantee privacy. The games were completely anonymous (toward other subjects and toward the experimenters) and involved 6 to 10 rounds of interactions with other group members over which group composition changed randomly. The possibility that some subjects encountered each other more than once could not be avoided given the number of subjects per session and the number of rounds, but anonymity ensured that there was no possibility for reciprocity (which would in any event favor greater levels of cooperation and therefore not alter our prediction of 100% cooperation in the repression-of-competition treatments).

The first experiment was played by 64 subjects divided into four sessions, in which each session involved four groups. The first and fourth sessions played the control treatment followed by the repression-of-competition treatment, whereas in the second and third sessions games were played in the reversed order to control for order effects. Both games involved 10 experimental rounds. The second, third, and fourth experiments were played by 32, 40, and 32 subjects, respectively, divided into two sessions with reversed treatment orders between sessions. Sessions involved either four (second and fourth experiments) or five (third experiment) groups, and games involved six experimental rounds.

In all experiments and treatments, subjects were given 40 credits at the beginning of each experimental round and were asked to decide how many of these 40 credits to contribute to a public project. In the control treatment, which remained the same in all four experiments, the contributed credits were doubled and then shared out equally among the four subjects within the group. Consequently, each subject's reward was the sum of retained and shared-out credits, and the strategy that maximizes the selfish gains of individuals would be to contribute none of their credits to the public project. The repression-of-competition treatment of the first experiment differed only from the standard public goods game in the way that contributions to the public project were multiplied by five. This guaranteed that subjects receive 1.25 credits for each credit contributed, such that 100% contribution is the

best strategy no matter what the other subjects in the group do. In this first experiment there were no interactions between groups within a session.

In the second experiment, the repression-of-competition treatment involved a standard public goods game as in the control treatment, but with additional group rewards depending on the summed group contribution to the public project (14, 15). The group reward, which each player within that group received, was zero for the group that contributed least and 60, 130, and 210 for the groups that contributed second-least, second-most, and most, respectively. Consequently, each subject's reward was the sum of retained, shared-out, and group reward credits. The top group prizes largely outweighed the cost of contributing to the public project, such that the strategy that maximized the selfish gains of individuals was to contribute 100% of their credits to the public project.

In the repression-of-competition treatments of the third and fourth experiments, we calculated the total credits contributed by each group and rewarded group members equally depending upon the relative contribution of their group. In the third experiment, we rewarded all individuals of the group with the highest contribution to the public project with 80 credits, followed by 65, 50, 35, and 20 credits to each member of the other groups in decreasing order of total group contribution. Because of reduced session size in the fourth experiment, we adjusted rewards to four instead of five groups: subjects of the same group received 80, 60, 40, and 20 credits depending on whether their group contributed most, second-most, second-least, and least to the public project, respectively. In the event of tied group ranks, the mean number of credits of the tied ranks was calculated and awarded to all subjects of those groups. All credits not contributed to the public project were lost, and so the group reward was the only reward a subject could gain in this game. This game structure completely repressed within-group competition, and therefore the strategy that maximized the selfish gains of individuals was to contribute 100% of their credits to the public project. The repression-of-competition treatments of the third and fourth experiments differed from one another only with respect to whether they featured a comprehension test. We omitted the comprehension test in the third experiment because, under this extreme experimental condition, such a test could be construed to directly instruct the subjects that full contribution is the only option in this game. The comprehension test could therefore perturb the subjects' natural perception of the game and bias contributions toward 100%. We included the comprehension test in our fourth experiment because if subjects failed to contribute fully in this biased situation, this would be the strongest support for the idea that humans are not predisposed to cooperate above the predicted level.

Experimental Procedure. Experiments were carried out in computer labs using the software z-Tree (45), developed at the Institute for Empirical Research in Economics in Zurich. Subjects were randomly allocated a computer terminal on a local network that could be remotely monitored and controlled from the experimenter's terminal. This procedure limited verbal instructions and therefore potential biases to a minimum.

Subjects received instructions for the first treatment via their terminal screens by means of an animated presentation (Microsoft Office PowerPoint 2003). The presentation provided clear information about the game rules, the game sequence, and the rewarding scheme, and included two animated visual examples, which mimicked exactly all steps of the games. After the PowerPoint presentation, subjects were directed to the z-Tree program, whereby they first undertook a comprehension test (except in the third

experiment), in which they had to navigate through two hypothetical rounds where they were asked to calculate the payoff of each subject within the group. Crucially, subjects could only join the game when they had demonstrated full understanding of the game rules and payoff structure. The comprehension test was followed by one practice round and then by the experimental rounds of the game. The practice round allowed subjects to become familiarized with the sequence of the screens and the game in general. On the first screen, subjects were asked to make a decision as to how many of their 40 credits they wanted to contribute to the public project (a rule reminder was added to that screen). This was followed by a screen comparing the contributions of all subjects in that group and by a screen summarizing the subject's reward from this round and the cumulative reward across all rounds. Once all experimental rounds of the first treatment were played, subjects were redirected to the PowerPoint presentation to consult the instructions of the second treatment, and then the sequence of steps repeated as described above. Importantly, subjects did not know beforehand that there was more than one treatment. All credits won during the games were summed and transformed into real monetary rewards.

Once both games were completed, subjects were directed to a questionnaire on which they had to score to what extent they perceived their group mates as competitors or collaborators on a sliding scale from zero (full collaborator) to 20 (full competitor), for both games. They were further asked to provide a unique identification code, which was linked to the subjects' monetary rewards and facilitated anonymous payment via a third party (university and college administrative staff). The maximal possible monetary reward ranged between £27 and £30, depending on experiment. The mean monetary reward across all subjects was £17.0, ranging between £12 and £22.9.

Statistical Analysis. To test whether the level of contribution changed as a response to the rounds of the game, we used GLMM. In each case, "subject ID" was entered as a random factor to account for the fact that repeated measures were taken from the same individual. Exactly the same GLMM procedure was used to test whether the competitor score differed between the control and the repression-of-competition treatments. We further used LMM to test whether the mean contribution across all experimental rounds differed between the two treatments. As with the previous models, we entered subject ID as a random factor. We used bootstrapping (10,000 replicates) with resampling to obtain the 95% CI for the observed proportion of subjects behaving perfectly (i.e., cooperating fully) and the proportion of subjects scoring their group mates imperfectly in the repression-of-competition treatments. Finally, we used Fisher's exact tests to analyze whether frequencies of subjects behaving imperfectly or scoring their group mates imperfectly differ between the control treatment and the repression-of-competition treatment. All statistical computations were carried out with R 2.10.0 (<http://www.r-project.org>).

ACKNOWLEDGMENTS. We thank L. Miller, W. Islam, A. Harris, and M. Bennett from the Nuffield Centre for Experimental Social Sciences for hosting our experiments in Oxford; F. Bayer, M. Cairney, C. El Mouden, A. Gardner, A. Griffin, and P. van den Berg for discussion and their help during the experiments in Edinburgh; and four anonymous referees for their extremely useful comments. This work was funded by the Royal Society, the Leverhulme Trust, the European Research Council, and a Marie-Curie Intra-European fellowship.

- Ledyard JO (1995) *The Handbook of Experimental Economics*, eds Kagel JH, Roth AE (Princeton Univ Press, Princeton).
- Fehr E, Schmidt K (1999) A theory of fairness, competition and cooperation. *Q J Econ* 114:817–868.
- Fehr E, Fischbacher U (2003) The nature of human altruism. *Nature* 425:785–791.
- Burnham TC, Johnson DP (2005) The biological and evolutionary logic of human cooperation. *Anal Kritik* 27:113–135.
- Boyd R, Gintis H, Bowles S, Richerson PJ (2003) The evolution of altruistic punishment. *Proc Natl Acad Sci USA* 100:3531–3535.
- Gintis H (2000) Strong reciprocity and human sociality. *J Theor Biol* 206:169–179.
- Gintis H, Bowles S, Boyd R, Fehr E (2003) Explaining altruistic behavior in humans. *Evol Hum Behav* 24:153–172.
- Gintis H, Bowles S, Boyd R, Fehr E (2005) *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life* (MIT Press, Cambridge, MA).
- Fehr E, Gächter S (2002) Altruistic punishment in humans. *Nature* 415:137–140.
- Fehr E, Fischbacher U, Gächter S (2002) Strong reciprocity, human cooperation and the enforcement of social norms. *Hum Nat* 13:1–25.
- Fehr E, Rockenbach B (2004) Human altruism: Economic, neural, and evolutionary perspectives. *Curr Opin Neurobiol* 14:784–790.
- Bowles S, Gintis H (2004) The evolution of strong reciprocity: Cooperation in heterogeneous populations. *Theor Popul Biol* 65:17–28.
- Hagen EH, Hammerstein P (2006) Game theory and human evolution: A critique of some recent interpretations of experimental games. *Theor Popul Biol* 69:339–348.
- Puurinen M, Mappes T (2009) Between-group competition and human cooperation. *Proc Biol Sci* 276:355–360.
- Burton-Chellew MN, Ross-Gillespie A, West SA (2010) Cooperation in humans: Competition between groups and proximate emotions. *Evol Hum Behav* 31:104–108.
- Haselton MG, Nettle D (2006) The paranoid optimist: an integrative evolutionary model of cognitive biases. *Personality and Social Psychology Review* 10:47–66.
- Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. *Econometrica* XLVII:263–291.
- Frank SA (1998) *Foundations of Social Evolution* (Princeton Univ Press, Princeton).
- West SA (2009) *Sex Allocation* (Princeton Univ Press, Princeton).
- West SA, et al. (2006) Cooperation and the scale of competition in humans. *Curr Biol* 16:1103–1106.
- Levitt SD, List JA (2007) What do laboratory experiments measuring social preferences reveal about the real world? *J Econ Perspect* 21:153–174.
- Haley KJ, Fessler DMT (2005) Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evol Hum Behav* 26:245–256.
- Burnham TC, Hare B (2007) Engineering human cooperation: Does involuntary neural activation increase public goods contributions? *Hum Nat* 18:88–108.
- Batson M, Nettle D, Roberts G (2006) Cues of being watched enhance cooperation in a real-world setting. *Biol Lett* 2:412–414.
- Trivers RL (1971) The evolution of reciprocal altruism. *Q Rev Biol* 46:35–57.
- Alexander RD (1987) *The Biology of Moral Systems* (Aldine de Gruyter, New York).

27. Wedekind C, Milinski M (2000) Cooperation through image scoring in humans. *Science* 288:850–852.
28. Milinski M, Semmann D, Krambeck H-J (2002) Donors to charity gain in both indirect reciprocity and political reputation. *Proc Biol Sci* 269:881–883.
29. Binmore K (2005) *Natural Justice* (Oxford Univ Press, Oxford).
30. Binmore K (2006) Why do people cooperate? *Polit Philos Econ* 5:81–96.
31. Nowak MA, Sigmund K (2005) Evolution of indirect reciprocity. *Nature* 437:1291–1298.
32. Frank SA (1995) Mutual policing and repression of competition in the evolution of cooperative groups. *Nature* 377:520–522.
33. Frank SA (2003) Perspective: Repression of competition and the evolution of cooperation. *Evolution* 57:693–705.
34. Lehmann L, Keller L (2006) The evolution of cooperation and altruism—a general framework and a classification of models. *J Evol Biol* 19:1365–1376.
35. Lehmann L, Rousset F, Roze D, Keller L (2007) Strong reciprocity or strong ferocity? A population genetic view of the evolution of altruistic punishment. *Am Nat* 170:21–36.
36. West SA, Griffin AS, Gardner A (2007) Social semantics: Altruism, cooperation, mutualism, strong reciprocity and group selection. *J Evol Biol* 20:415–432.
37. West SA, Griffin AS, Gardner A (2007) Evolutionary explanations for cooperation. *Curr Biol* 17:R661–R672.
38. Sachs JL, Mueller UG, Wilcox TP, Bull JJ (2004) The evolution of cooperation. *Q Rev Biol* 79:135–160.
39. Bergmüller R, Bshary R, Johnstone RA, Russell AF (2007) Integrating cooperative breeding and cooperation theory. *Behav Processes* 76:61–72.
40. Trivers R (2004) Mutual benefits at all levels of life. *Science* 304:964–965.
41. Trivers R (2006) *Cooperation in Primates and Humans: Mechanisms and Evolution*, eds Kappeler PM, van Schaik CP (Springer-Verlag, Berlin), pp 67–83.
42. Grafen A (2006) Optimization of inclusive fitness. *J Theor Biol* 238:541–563.
43. Gardner A, West SA (2004) Cooperation and punishment, especially in humans. *Am Nat* 164:753–764.
44. 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. University of Cologne, Cologne, Germany.
45. Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. *Exp Econ* 10:171–178.