Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations

Robert Kurzban*†‡ and Daniel Houser†§

*Department of Psychology, University of Pennsylvania, 3720 Walnut Street, Philadelphia, PA 19104; and §Interdisciplinary Center for Economic Science and Department of Economics, George Mason University, 4400 University Drive, MSN 1B2, Fairfax, VA 22030

Unlike other species, humans cooperate in large, distantly related groups, a fact that has long presented a puzzle to biologists. The pathway by which adaptations for large-scale cooperation among nonkin evolved in humans remains a subject of vigorous debate. Results from theoretical analyses and agent-based simulations suggest that evolutionary dynamics need not yield homogeneous populations, but can instead generate a polymorphic population that consists of individuals who vary in their degree of cooperativeness. These results resonate with the recent increasing emphasis on the importance of individual differences in understanding and modeling behavior and dynamics in experimental games and decision problems. Here, we report the results of laboratory experiments that complement both theory and simulation results. We find that our subjects fall into three types, an individual's type is stable, and a group's cooperative outcomes can be remarkably well predicted if one knows its type composition. Reciprocal types, who contribute to the public good as a positive function of their beliefs about others' contributions, constitute the majority (63%) of players; cooperators and free-riders are also present in our subject population. Despite substantial behavioral differences, earnings among types are statistically identical. Our results support the view that our human subject population is in a stable, polymorphic equilibrium of types.

behavioral economics | cooperation | evolution | public goods

The evolution of cooperation in sizable groups of nonkin in humans has been and continues to be the subject of debate (1, 2). Although theoretical models of group selection (3, 4), indirect reciprocity (5–7), and costly signaling (8) have helped to clarify possible evolutionary routes to cooperation, fundamental questions remain about the number and nature of the cognitive mechanisms that underpin human cooperative psychology (9) and whether there are stable individual differences in these mechanisms (10).

Agent-based simulations of the evolutionary dynamics of interacting strategies, which generally embody this assumption, that an agent's type is stable, causing it to use the same strategy until its simulated death, have helped to inform these issues (3). A key result from simulations is that populations, instead of evolving toward agents with homogeneous behavioral strategies, often evolve such that multiple strategies coexist at an equilibrium (11–13). There are many ways that this diversity can be maintained (14), and the details by which this occurs is the subject of intense research activity in various domains (15). Experimental results of the kind reported here can be useful for clarifying whether the assumption of stable types in the realm of human cooperation as assumed by simulations is justified, and, if it is, how best to characterize these types (16, 17).

The laboratory experiment reported here complements simulations by exploring type stability in a potentially fickle human population. In line with types used in simulations (18) and observed in other experimental contexts (19), we consider the hypothesis that people are one of three stable types: (i) coop-

erators, who contribute to generating group benefits at a cost to self, (ii) free-riders, who do not incur these costs, and (iii) reciprocators, who respond to others' behavior by using a conditional strategy. Note that reciprocity is more complex in multiplayer games than in two-player games because the breadth of possible conditional strategies obviously expands (20). However, in our experiment subjects were given information about aggregate behavior only, which simplifies their strategy space.

The dynamics of agent-based simulations are sensitive to the fraction of types in the population and the frequency with which these types interact (13). A goal of the experiments reported here was to investigate whether the type composition of a group has a predictable effect on the dynamics of cooperation in human groups. Public goods games provide a natural environment for this study because the results of previous public goods experiments are consistent with the conjecture that types exist and affect cooperative outcomes.

We use a public goods game with the voluntary contribution mechanism (21) to learn about types and the effect of types on cooperative dynamics. Our experimental subjects are randomly assigned to groups of four people and faced with a decision to divide money (represented by tokens that are provided to them by the experimenter) into private and group accounts. Tokens placed into the private account are kept by the investing individual. Tokens placed in the group account are doubled by the experimenter, and then divided equally among all group members. Hence, each unit invested in the group account increases the aggregate group payoff, but decreases the investing individual's payoff. This generates a social dilemma, where the group optimum is achieved only when each individual makes a decision that is contrary to her or his income-maximizing self-interest (22).

Experimental Design and Procedures

Our experiment included 84 participants in three groups of 24 and one group of 12. We found no statistically significant difference in behavior between the 12- and 24-subject sessions, and therefore data are pooled. Participants were recruited from the undergraduate population by using the recruitment system in place at the University of Arizona's Economic Science Laboratory.

Participants arrived in the laboratory and were seated at computer terminals that were divided by partitions. Except for the instructions, which were printed on paper, the entire experiment was conducted via computer. Once seated, participants were told that they had already earned their show-up payment and that their decisions and earnings would be kept confidential.

PSYCHOLOGY

Subsequently, the instructions for our sequential voluntary contribution mechanism, available on request, were distributed.

The instructions informed participants that they would be assigned to a series of groups, each consisting of four people, and that the members of these groups would be shuffled randomly over an unspecified number of games. Participants were told that they would receive an endowment of 50 tokens per game to divide between two exchanges, and that tokens in the individual exchange earned one cent per token, whereas tokens in the group exchange earned half of one cent for each player.

Each game consisted of all players simultaneously making an initial allocation of tokens between the two accounts. After this initial decision, there were a number of rounds. Each round proceeded as follows. First, one player in each group was provided the current aggregate contribution to the group exchange. Then, that person was given the opportunity to change his or her allocation to the two accounts. The game proceeded round by round until the game ended at a point unknown to the participants. Participants were told that each player would have at least one chance to change his or her contribution decision, and that the game would end at a randomly chosen point. We provided no details on the nature of the randomization process. However, it was emphasized that payoffs in each game would be determined by the final allocation of all group members' tokens to the two exchanges, and that total earnings for the experiment would be determined by summing the earnings from each game.

The rotation of subjects to constitute groups, the order of play within groups, and the length of each game were generated randomly and kept constant for all four sessions. Game lengths were generated by assigning a probability of 0.04 that the game would end after any player's opportunity to change his or her allocation, subject to the constraint that all subjects be allowed to update at least once. We chose this procedure with an eye toward providing enough variation in game lengths to ensure that subjects did not come to expect games to last a specific number of rounds. This technique is important to help make certain that subjects viewed all of their decisions (aside from the first simultaneous contribution) as potentially payoff relevant. Payoff saliency is also an important reason that we chose not to reveal the randomization structure to the subjects: some subjects might mistakenly believe that a small probability of the game ending after any round means that they would always have many opportunities to change their decisions.

Our randomization process generated the following number of opportunities to update contribution decisions (excluding the four initial simultaneous contributions): 16, 7, 23, 32, 32, 34, 4, 17, 31, 8. So, for example, in the first game there was an initial set of simultaneous contributions, and then the game proceeded sequentially until each of the four subjects had had four opportunities to update his or her previous contribution, at which point the game ended and subjects' earnings for that game were calculated.

Participants completed a 10-question quiz that had to be answered correctly before they could proceed. The first game began after everyone had completed the quiz correctly, and subsequent games proceeded automatically after all groups had reached the end of the preceding game.

Participants were paid their experimental earnings privately, ≈$20 on average, and dismissed when the experiment concluded. Subjects were in the laboratory for ≈90 min.

Results

Aggregate Contributions. Each experimental session included at least seven games. Some sessions proceeded slightly faster and included as many as 10 games. Final contributions to the group account displayed the decay typically found in public goods experiments. In particular, average contributions decayed over time from 60% to ≈35% of the subjects' endowment. However,

average contributions mask substantial heterogeneity in behavior among individuals and groups, an issue to which we now turn.

Statistical-Type Classification Algorithm. Our approach to behavioral-type classification is to prespecify a set of behaviors of interest, and then assign one from this set to each subject.[¶] This sort of approach was used, for example, by El-Gamal and Grether (23) in their well known behavioral typing algorithm [see also Houser and Winter (24)]. Although more sophisticated (and cumbersome) procedures are available, the advantage of our classification algorithm is that it provides a simple, fast, and accurate method for inference about individual differences, after which any analysis can be conducted.

The behaviors that interest us are contributing little most of the time (free-riding), contributing a great deal most of the time (cooperating), and contributing an amount roughly equal to the contributions of others (conditional cooperation or reciprocation). Intuitively, our procedure bases inferences about a subject's type on a plot of a subject's contributions against the average contribution to the group account she observed before making his or her own contribution. Contributions by cooperators lie well above the 45° line on this plot. Reciprocators' contributions cluster near the 45° line, whereas free-riders' contributions are small regardless of the average contribution they observe. Fishbacher, Gächter, and Fehr (19), who studied reciprocation by using the strategy method (25) with a one-shot public goods game, used a similar approach to type classification. Indeed, our algorithm can be viewed as a formal extension of their visual-inspection approach.

Although our statistical procedure has the advantage of simplicity, it also risks omitting variables that could be important determinants of behavior. These variables might include, for example, a subject's previous outcomes within the experiment (amounts earned or amounts previous group members have contributed). Moreover, even if all appropriate variables are included, our simple specification might include them in the wrong way (say, by neglecting higher-order terms). Specification errors and omitted variables are a concern for us to the extent that they contaminate our type classifications. In fact, as we discuss below, the decisions subjects made in games used for classification purposes predict well the decisions they made in games played afterward. Hence, the evidence is that our type classifications are valid, and that our model, although simple, nevertheless provides useful and reliable results.

Formally, our statistical procedure is to use a subject's linear conditional-contribution profile (LCP) to classify each subject as one of the three types. We define each subject's LCP as the outcome of an ordinary least-squares regression of his or her contribution decisions on the mean contribution that he/she observed immediately before making a contribution decision. This process enables us to summarize a subject's contribution strategy with two parameters: the intercept and slope of his or her LCP. The intercept provides a measure of how willing a subject is to cooperate even when his or her group counterparts contribute little to the public good. The slope measures a subject's responsiveness, both in direction and magnitude, to others' contributions.

We classify each subject according to his or her estimated LCP, as follows. A subject is classified as a free-rider if and only if the graph of his or her LCP lies everywhere <25. This means that a subject's expected contribution is <50% of the endowment at all levels of others' contributions. Similarly, a subject is classified as a cooperative type if and only if the graph of his or her LCP lies everywhere at or >25. That is, a cooperative type's expected contribution to the public good is more than half of the endowment regardless of others' contributions. We classify a

---

[¶]Houser, Keane, and McCabe (17) provide a type classification algorithm that does not require one to specify the types of interest a priori.

subject as a ==reciprocator if the graph of his or her LCP has a positive slope and lies both above and below the 50% line. Hence, a reciprocator's expected contribution increases as the group average increases.== Subjects with any other LCP are not classified as one of these three types.

Findings. Based on the seven (in-sample) games that each subject plays, our algorithm classifies 17 of our 84 subjects (20%) as free-riders, 11 (13%) as cooperators, and 53 (63%) as reciprocators (three subjects were not statistically classifiable as one of these types).[||] Note that this distribution of types is similar to the distribution reported by Fishbacher, Gächter, and Fehr (19).

Contribution decisions differ substantially among these types. The median per-round contributions of free-riders, reciprocators, and cooperators are statistically significantly different ($P < 0.001$, three-sample Medians test), at 1, 25, and 50 tokens, respectively. Despite this, median earnings per game across types are not statistically different ($P = 0.27$, three-sample Medians test), ranging from 70.0 (conditional cooperators) to 77.5 (free-riders),[**] a narrow band within the possible range from 25 to 125.

To investigate whether individual differences in our experiment are stable, when time allowed we had subjects play up to three additional games, again with randomly reassigned partners. In these out-of-sample games, subjects typed as free-riders continued to contribute less on average than their group counterparts, cooperators more, and reciprocators about the same. ==The median contributions of free-riders, reciprocators, and cooperators in the out-of-sample games are 5, 25, and 50,== respectively, and are statistically significantly different ($P < 0.001$, three-sample Medians test). In addition, we fit a simple censored regression model (contributions on a constant and the average group contribution) to the out-of-sample contribution decisions of each type. Each type's estimated model is statistically significantly different from each other type's ($P < 0.01$ for all three possible pairwise comparisons).[††] Overall then, our results provide evidence that types in our experiment are different from one another and stable over time.

Because we have provided evidence that types are both stable and behaviorally distinct, it is reasonable to distinguish groups by their type compositions. To do this, we assigned to each group a cooperativeness score, with higher scores indicating that the group contained more cooperative types (see Fig. 1). One summary measure of group cooperation is the final amount contributed by a group to the public good. Over the first seven games, there is a positive relationship between this measure and a group's cooperativeness score. Moreover, a simple regression analysis reveals that this relationship is statistically significant, with a one-point increase in the group cooperativeness score associated with a ≈20-token increase in the amount of the final contribution to the public good. This result is not surprising, given that these data were themselves used to determine types. Importantly, however, a similar analysis run on the groups observed in the out-of-sample games reveals a nearly identical relationship, and one cannot reject the hypothesis that a group's cooperativeness score has the same effect on final group contributions in both samples.[‡‡]

Again, because player types are stable and distinct, we investigated whether cooperative dynamics within a newly formed
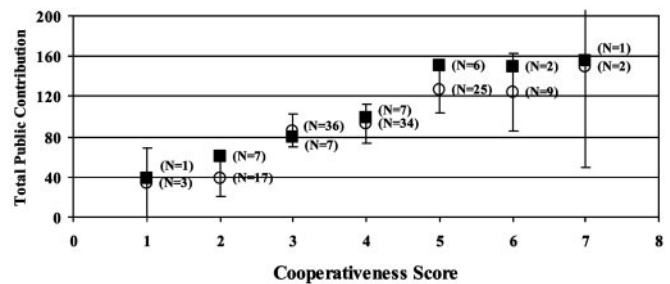

Fig. 1. Displayed are the mean final group contribution for the first seven games (○) and holdout sample (■) and a two-standard error interval above and below that mean. We assign a score of zero to each free-rider in a group and scores of one and two to each conditional cooperator and cooperator, respectively. A group's score is the sum of the scores of its four members. We obtain group scores ranging from one to seven because the random matching process did not generate any group that included only cooperators or only free-riders. The absolute difference between the in-sample means and the holdout sample means is small, with the greatest deviation being 26 tokens. The holdout sample means are generally within, or nearly within, two standard errors of the in-sample means. Of course, not all of these standard errors are tight, because the number of observations for some group scores is quite small (e.g., two and three observations during the first seven games in groups with scores of seven and one, respectively). Shown is evidence that final group contributions are positively related to a group's cooperativeness score and that a group's score predicts its cooperative decisions out of sample.

group can be predicted conditional on knowing that group's type composition. We found that a group's final contribution to the group account can be well predicted by its cooperativeness score (Fig. 1). Moreover, the path leading to this final contribution is also accurately predicted. In particular, the average per-round contribution across out-of-sample groups with a given score generally lies within two standard errors of the estimated mean contribution of in-sample groups with the same score (Fig. 2). However, there are relatively few out-of-sample data points (Fig. 1), and there is significant variation about each mean. In addition, one notices from visual inspection that the paths for different groups are somewhat similar. Consequently, it is not surprising that results from pairwise t tests indicate that the out-of-sample contribution paths for the different groups are largely statistically indistinguishable.[§§]

Note that, as pointed out above, the games lasted different numbers of rounds (between 4 and 34). Clearly, game length also will contribute to the amount of cooperative decay that a group experiences. When the number of observations on a particular group score is large, the game length effect is mitigated by the randomization procedures embedded in our design. For some group scores we have only a small number of observations (see Fig. 1). In these cases, the game length effect will tend to obscure the relationship between the group's cooperativeness score and the final group contribution, and therefore make it more difficult to predict group outcomes based on cooperativeness scores. In this light, our finding that group outcomes are well predicted hints that composition effects are an important source of between-group differences in cooperation.

Discussion

Previous Classification Systems. Our results provide evidence that there are multiple, stable behavioral types that vary with respect to their disposition to cooperate in a group context. Fishbacher, Gächter, and Fehr (19), having found evidence of both free-riding and conditionally cooperative strategies, suggested that groups that

---

[||]The R-squared values for the type regressions ranged from ≈0 to ≈1, with a mean of ≈0.4.

[**]The standard deviations of the earnings distributions for free-riders, cooperators, and reciprocators are 18, 19, and 24, respectively.

[††]Each estimated model was statistically significant ($P < 0.001$) with pseudo R-squared ≈0.02.

[‡‡]Using the in-sample data, we ran a regression of the final contribution amount on an intercept and the group cooperativeness score. The coefficients' point estimates were 11 and 21, respectively, with only the 21 statistically significant ($P < 0.001$). The same regression run on the out-of-sample data generated coefficient estimates of 10 and 25, respectively, with only 25 statistically significant ($P < 0.01$). The null hypothesis that the models' parameters are jointly identical in the two cases cannot be rejected at standard significance levels ($P = 0.59$). The R-squared for both regressions is ≈0.22.

[§§]We ran pairwise, round-by-round t tests to investigate differences in mean contribution paths. Differences were statistically insignificant in all cases except rounds 1–8 between the score-2 and score-5 groups ($P < 0.05$).
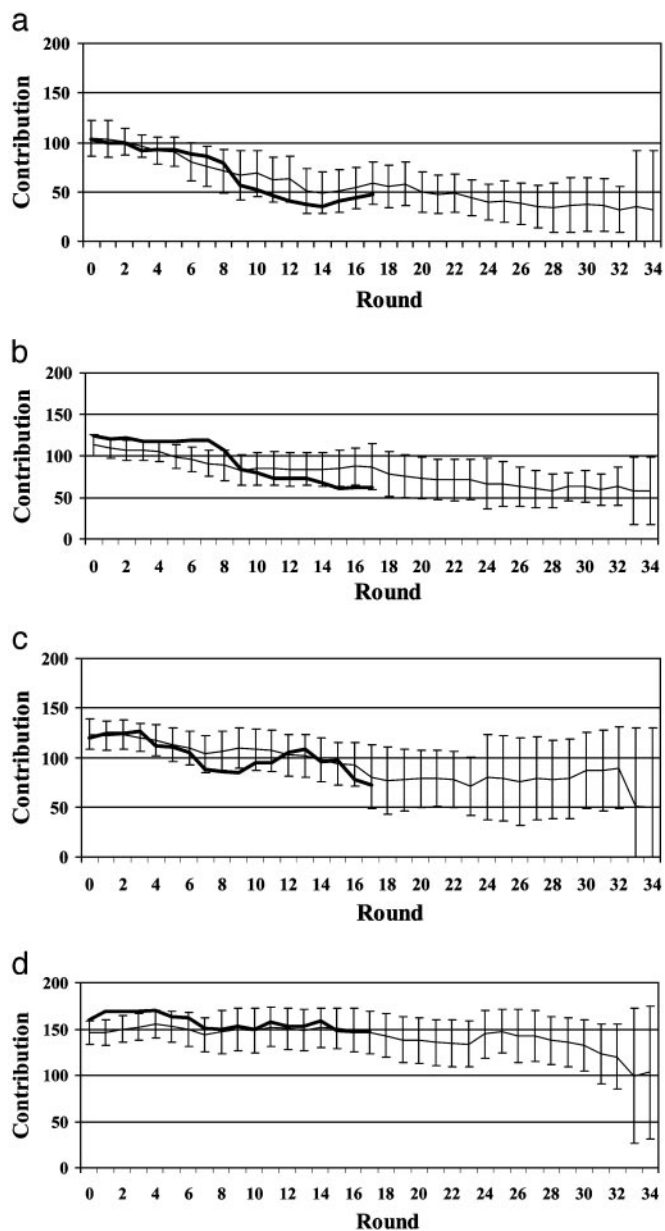
Fig. 2. Mean contribution of groups in the first seven (thin lines) and final three (thick lines) games. (*a*) Groups with score = 2. (*b*) Groups with score = 3. (*c*) Groups with score = 4. (*d*) Groups with score = 5. The horizontal axis indicates the round, where round zero is the initial, simultaneous contribution. The vertical axis indicates the equally weighted (over groups) mean contribution to the public good. We include only groups with cooperativeness scores between two and five because the number of observations on groups with other scores is too low (either one or two) to allow reasonable inferences with respect to dynamics. The pattern of mean contributions over rounds in the first seven games is a prediction of cooperative dynamics in the final three games. The vertical bars extend two standard errors above and below each in-sample mean. For groups with scores of 2, 4, and 5, nearly all of the holdout sample's mean contributions are within two standard errors of the predicted means. These findings suggest that there is a predictable relationship between temporal patterns of group cooperation and the cooperative nature of a group's members.

include both types might be expected to experience cooperative decay and convergence to a noncooperative equilibrium, and then speculated that "the speed of convergence depends on the actual composition of the group." Our results provide direct evidence in support of these and other closely related hypotheses that have been

previously advanced only speculatively in the public goods literature. In addition, because the theoretical efficiency of an institution usually relies on some form of behavioral continuity, evidence that individuals' behavioral rules exhibit temporal stability is valuable to those involved in mechanism design.

Ours is, of course, not the first research to show that there are multiple motivations in experimental games (26) or that individual differences in preferences lead to varied social dynamics (27, 28). Research in this tradition, which has historically used two-option, two-player interactions (29), frequently classified participants into three types, as we do here, although some classification systems include more (30).

It is interesting to note that social psychologists and economists have postulated similar classification systems. The research tradition in social psychology on social value orientation, for example, suggests that people can be classified as competitors (motivated to achieve better payoffs than others), cooperators (motivated to try to increase group welfare), and individualists (motivated to serve their own interests) (31). This set of types nicely parallels work in experimental economics that suggests that people fall into types who are "spiteful" (competitive) (32), "payoff-maximizing" (individualists), and "altruistic" (cooperators) (33).

It is tempting to map these schemes onto our three types, but because of the important differences between two-player and *N*-player interactions (20), two-player, two-option games might or might not be informative of behavior in *N*-player games. For example, although an individual's social value orientation does predict behavior in resource dilemmas (34), it does not do so in public goods games (35).

Direct assessment of types in public goods games is relatively recent, but they have been evident for some time. Early work on public goods environments suggested that some fraction of individuals was "strong free riders" (21), and this proportion changed as a function of the marginal per-capita return. Still earlier, a kind of reciprocal player was identified that used "one-period" contributions or "pulses" to induce reciprocal contributions from others (36). Note that there was some initial skepticism about the importance of types in explaining laboratory data. Pruitt and Kimmel (37), for example, believed that "dispositional qualities" would have "little impact in an impersonal setting as represented by most gaming environments." This view contrasts strongly with a recent comment on the consistency of individual differences in motivations in mixed motive interactions found in experiments; Ketelaar (38) suggested rather that the evidence is that "several different varieties of social motive (and not just one) [are] routinely observed in the adult population."

Currently, individual differences are receiving increasing attention. In addition to Fishbacher *et al.*'s (19) work described above, a valuable contribution closely related to the research reported in this article was done by Casari and Plott (CP) (39). CP model individual differences by assigning individuals parameters of the degree to which they are "spiteful" or "altruistic" in a commons dilemma (which is conceptually similar to a public goods game). While both we and CP use linear parametric models to characterize the nature of other-regarding preferences among our subjects, CP rule out reciprocity, whereas we focus on reciprocal preferences in our effort to develop predictions of group dynamics. Similarly, our sequential design, in contrast to the simultaneous contribution protocol used by CP and the majority of public goods game researchers, allow us to relax CP's assumption that "agents expect the others to act in period *t* as they did in period (*t-1*)." Although this assumption might be correct, our observations imply that expectations about others' behavior might include a dynamic component related to reciprocity.

The Evolution of Cooperative Types and Simulations. We are encouraged by connections between our results, the results of other type classification systems, and the results of evolutionary simulations. Evolutionary game theorists have known for some time

Kurzban and Houser

that populations can achieve stable polymorphic equilibria (40, 41). Lomborg (13), for example, describes evolutionary simulations that lead to stable populations of three types: cooperators, "cautious cooperators," and noncooperators, although the proportions of each varied across simulations. The stability we observe supports the use of types in these simulations and is potentially informative on the key evolutionary issue of whether variation in experimental games might be caused by players using mixed strategies as opposed to the possibility that we are observing a polymorphic population.

Consider also our (unsurprising) result that groups composed of more cooperative types enjoyed greater group cooperation and tended to earn more. For example, three reciprocators when grouped with a cooperator can expect to earn $\approx$40% more than when they are grouped with a free-rider. At the same time, each type's average earnings were statistically identical over the course of our experiment. In a polymorphic equilibrium, payoffs to each type should be equal (14). In this context, the similarity of earnings among our types is intriguing. This result is consistent with the possibility that individual differences reflect genuine strategic differences as opposed to differences in, for example, amounts of confusion between participants (42, 43).

Another promising line of work that could profitably combine both simulation and experimental procedures might be to investigate the properties of our sequential public goods game under a payoff structure that implies interior equilibria. In the laboratory, this structure will help to distinguish motivations for contributions because, for example, positive contributions in our method have ambiguous interpretations. In particular, confusion and prosocial behavior can be conflated (42, 43). (The reason that all participants in our experiment were required to complete perfectly a 10-question quiz is that it helped to ensure that subjects understood the structure of the experimental environment.) In addition to avoiding the confusion confound, public goods games with interior equilibria are able to distinguish rich varieties of behaviors (33) and assess treatment effects in ways that are not possible when the equilibrium is on the boundary (44). At the same time, in their comprehensive

survey, Laury and Holt (44) point out that the main results from the linear voluntary contribution mechanism literature do not change when the equilibrium is moved to any point that lies between the lower boundary and the midpoint of the contribution decision space.

A few additional caveats are in order. Clearly, the notion of type must be understood with respect to a particular class of decisions; if cooperation is very costly, we presume universal defection would obtain (21). Nonetheless, our results add to the growing body of research that suggests that reciprocity is an important motive in group contexts across a range of institutional arrangements (45, 46). Additional work will be required to determine the conditions under which players separate into types, including conditional cooperative ones, and which institutions have the effect of homogenizing play. It seems likely that important variables facilitating reciprocal play are the extent to which gains from trade are available, the amount of repeat interaction, and the availability of information about others' behavior.

In addition, we remain agnostic with respect to the very important issue of the correct ultimate explanation for the existence of cooperative types, and how cooperative strategies are stabilized. Punishment (47–49) and the possibility of non-participation (50, 51) are two elements that look promising as parts of the answer to the question of how adaptations designed for cooperation in groups might have evolved.

One avenue of research, already in progress, is to replicate experiments such as ours to determine whether distributions of types are similar cross-culturally. Such similarities, were they to be found, might help to clarify the origins of heterogeneous preferences. In addition, agent-based simulations, which can be calibrated to yield stable heterogeneity in agents, will be useful for understanding the mix of strategies that can persist in populations. Productive future work, some of which is already underway (52), should be aimed at more closely integrating laboratory and agent-based simulation research.

1. Darwin, C. (1859) *The Origin of Species* (Murray, London).
2. Wilson, E. O. (1975) *Sociobiology: The New Synthesis* (Belknap, Cambridge, MA).
3. Boyd, R., Gintis, H., Bowles, S. & Richerson, P. J. (2003) *Proc. Natl. Acad. Sci. USA* **100,** 3531–3535.
4. Gintis, H. (2000) *J. Theor. Biol.* **206,** 169–179.
5. Leimar, O. & Hammerstein, P. (2001) *Proc. R. Soc. London* **268,** 748–753.
6. Nowak, M. A. & Sigmund, K. (1998) *Nature* **393,** 573–577.
7. Panchanathan, K. & Boyd, R. (2003) *J. Theor. Biol.* **224,** 115–126.
8. Gintis, H., Bowles, S. & Smith, E. A. (2001) *J. Theor. Biol.* **213,** 103–119.
9. Fehr, E. & Fischbacher, U. (2003) *Nature* **425,** 785–791.
10. Van Lange, P. A. M., De Bruin, E. M. N., Otten, W. & Joireman, J. A. (1997) *J. Pers. Soc. Psychol.* **73,** 733–746.
11. Dugatkin, L. A. & Wilson, D. (1991) *Am. Nat.* **138,** 687–701.
12. Aktipis, C. A. (2004) *J. Theor. Biol.* **231,** 249–260.
13. Lomborg, B. (1996) *Am. Soc. Rev.* **61,** 278–307.
14. Maynard Smith, J. (1982) *Evolution and the Theory of Games* (Cambridge Univ. Press, Cambridge, U.K.).
15. Sinervo, R. & Lively, C. M. (1996) *Nature* **288,** 240–243.
16. Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton Univ. Press, Princeton).
17. Houser, D., Keane, M. & McCabe, K. (2004) *Econometrica* **72,** 781–822.
18. Bowles, S. & Gintis, H. (2004) *Theor. Popul. Biol.* **65,** 17–28.
19. Fischbacher, U., Gächter, S. & Fehr, E. (2001) *Econ. Lett.* **71,** 397–404.
20. Boyd, R. & Richerson, P. J. (1988) *J. Theor. Biol.* **132,** 337–356.
21. Isaac, R. M. & Walker, J. M. (1988) *Econ. Inq.* **26,** 585–608.
22. Dawes, R. M. (1980) *Annu. Rev. Psychol.* **31,** 169–193.
23. El-Gamal, M. A. & Grether, D. M. (1995) *J. Am. Stat. Assn.* **90,** 1137–1145.
24. Houser, D. & Winter, J. (2004) *J. Bus. Econ. Stat.* **22,** 64–79.
25. Selton, R. (1967) in *Beitrage Zur Experimentellen Wirt Schaftsforschung*, ed. Sauermann, H. (Mohr Siebeck, Tübingen, Germany), pp. 136–168.
26. Deutsch, M. (1960) *Hum. Rel.* **13,** 123–139.
27. Kelley, H. H. & Thibaut, J. (1978) *Interpersonal Relations: A Theory of Interdependence* (Wiley, New York).
28. MacCrimmon, K. R. & Messick, D. M. (1976) *Behav. Sci.* **21,** 86–100.
29. Messick, D. M. & McClintock, C. G. (1968) *J. Exp. Soc. Psychol.* **4,** 1–25.
30. Au, W. T. & Kwong, J. Y. Y. (2004) in *Contemporary Psychological Research on Social Dilemmas*, eds. Suleiman, R., Budescu, D. V., Fisher, I. & Messick, D. (Cambridge Univ. Press, Cambridge, U.K.), pp. 71–98.
31. Komorita, S. S. & Parks, C. D. (1995) *Annu. Rev. Psychol.* **46,** 183–207.
32. Saijo, T. & Nakamura, H. (1995) *J. Conf. Res.* **39,** 535–560.
33. Cason, T. N., Saijo, T., Yamato, T. & Yokotani, K. (2003) *Games Econ. Behav.* **49,** 81–102.
34. Liebrand, W. B. (1984) *Eur. J. Soc. Psychol.* **14,** 239–264.
35. Parks, C. D. (1994) *Pers. Soc. Psychol. Bull.* **2,** 431–438.
36. Isaac, R. M., McCue, K. & Plott, C. (1985) *Pub. Choice* **62,** 217–236.
37. Pruitt, D. G. & Kimmel, M. J. (1977) *Annu. Rev. Psychol.* **28,** 363–392.
38. Ketelaar, T. (2004) in *Evolutionary Psychology, Public Policy, and Personal Decisions*, eds. Crawford, C. & Salmon, C. (Erlbaum, Mahwah, NJ), pp. 145–168.
39. Casari, M. & Plott, C. (2003) *J. Econ. Behav. Org.* **2,** 217–247.
40. Maynard Smith, J. & Price, G. (1973) *Nature* **146,** 15–18.
41. Peck, J. R. (1993) *J. Theor. Biol.* **162,** 195–228.
42. Andreoni, J. (1995) *Am. Econ. Rev.* **85,** 891–904.
43. Houser, D. & Kurzban, R. (2002) *Am. Econ. Rev.* **92,** 1062–1069.
44. Laury, S. K. & Holt, C. A. in *Handbook of Experimental Economics Results*, eds. Smith, V. & Plott, C. (Elsevier, Amsterdam), in press.
45. Gintis, H., Bowles, S., Boyd, R. & Fehr, E. (2003) *Evol. Hum. Behav.* **24,** 153–172.
46. Kurzban, R., McCabe, K., Smith, V. L. & Wilson, B. J. (2001) *Pers. Soc. Psychol. Bull.* **127,** 187–208.
47. Boyd, R. & Richerson, P. (1992) *Ethics Sociobiol.* **13,** 171–195.
48. Yamagishi, T. (1986) *J. Pers. Soc. Psychol.* **51,** 110–116.
49. Fehr, E. & Gächter, S. (2000) *Am. Econ. Rev.* **90,** 980–994.
50. Hauert, C., De Monte, S., Hofbauer, J. & Sigmund, K. (2002) *Science* **296,** 1129–1132.
51. Semmann, D., Krambeck, H. & Millinski, M. (2003) *Nature* **425,** 390–393.
52. Duffy, J., in *Handbook of Computational Economics*, eds. Judd, K. L. & Tesfatsion, L. (Elsevier, Amsterdam), Vol. 2, in press.

PSYCHOLOGY