Published in final edited form as:

Behav Processes. 2013 October ; 99: . doi:10.1016/j.beproc.2013.07.008.

Valuing rewards to others in a prisoner's dilemma game

Vasiliy Safin, Matthew L. Locey, and Howard Rachlin*

Stony Brook University, United States

Abstract

People value rewards to others but discount those rewards based on social distance; rewards to a socially closer person are valued more than identical rewards to a socially more distant person (Jones and Rachlin, 2006). The concept of social discounting can explain cooperation and defection in two-player prisoner's dilemma (PD) games (Axelrod, 1980). The contingencies of a PD game are such that in any single game cooperation is costly to each player herself but beneficial to the other player. From the viewpoint of each player, the costs of cooperation are fully realized, but the benefits of cooperation are discounted by the social distance to the other player. The present experiment measured cooperation and defection in two PD-game conditions with differing reward magnitudes. In one (the 1-2-3-4 condition), the cost of cooperation exceeded its socially discounted benefit, and players were predicted to defect; in the other (the 1-2-9-10 condition), the discounted benefit of cooperation exceeded its cost, and players were predicted to cooperate. Over the course of repeated trials defection increased with the 1-2-3-4 condition but not with the 1-2-9-10 condition. Moreover, participants who rated their partners as closer, relative to random classmates, cooperated at higher rates – consistent with social discounting.

Keywords

Altruism; Cooperation; Patterns of behavior; Prisoner's dilemma; Reciprocation; Social discounting; Tit-for-tat

Altruistic acts have been defined, in economic terms, as "...costly acts that confer economic benefits on other individuals" (Fehr and Fischbacher, 2003). In a prisoner's dilemma (PD) game, cooperation is costly to the cooperating player but also confers benefits on the other player(s). Cooperation in such games therefore fits Fehr and Fischbacher's definition of altruistic behavior. In a two-player PD game, each player gets a moderately high reward if they both cooperate and a moderately low reward if they both defect. However, if one player cooperates and one defects, the cooperator gets a very low reward while the defector gets a very high reward.

For example, in what we call a 1-2-3-4 PD game, mutual cooperation gives each player 3 reward units; mutual defection gives each player 2 reward units; when one player (P) cooperates and the other player (OP) defects, the cooperator gets 1 unit and the defector gets 4 units. Although both players benefit if both cooperate rather than defect (3 units each rather than 2 units each), each player as an individual benefits most by defecting regardless of the other player's choice. Fig. 1a shows these rewards as presented to participants.

Note that in the 1-2-3-4 game the cost to P of cooperating is 1 unit regardless of OP's choice: If OP cooperates, P earns 3 units by cooperating too, and 4 units by defecting; if OP defects, P earns 1 unit by cooperating, and 2 units by defecting too. The cost to P is 1 unit in either case. However, by cooperating, P also ensures that the gain to OP is 2 units (regardless of OP's choice): If P cooperates and OP cooperates, OP earns 3 units (rather than 1, if P had defected); if P cooperates and OP defects OP earns 4 units (rather than 2, if P had defected). The gain to OP is 2 units in either case. Although in absolute terms the benefit of cooperation (2 units to OP) outweighs the cost (1 unit to P), the cost is fully realized whereas the benefit must be discounted by the social distance between the players; the further the social distance, the less the benefit. The typical finding among non-related undergraduates playing repeated 1-2-3-4 PD games, is mutual defection (Silverstein et al., 1998).

Recently Locey et al. (2013) studied behavior in one-shot, 2-player PD games where the cost of cooperation was kept constant but the benefit conferred on the other player was varied. In one condition participants played a single (one-shot) 1-2-3-4 PD game as described above. In the other condition participants played a single 1-2-9-10 PD game, which also holds the cost of cooperation at 1 unit (10 minus 9 or 2 minus 1), but confers a benefit to the other player of 8 units (10 minus 2 or 9 minus 1). The benefit conferred on the other player by cooperating in the 1-2-9-10 PD game was thus 4 times that in the 1-2-3-4 PD game. Therefore players should be more likely to cooperate in the 1-2-9-10 PD game than in the 1-2-3-4 PD game; this was what Locey et al. found.¹

The Locey et al. (2013) experiment was administered as part of an online survey that asked participants to make a single choice between X (choice for defection) and Y (choice for cooperation). The instructions explained that each P was paired with another player (OP) such that P's decision affected not only P's payoff but also OP's. The rewards were hypothetical, but participants were asked to choose as if the rewards were real. Half of the participants played the 1-2-3-4 PD game; the other half played the 1-2-9-10 game. As predicted, the 1-2-3-4 players cooperated significantly less than the 1-2-9-10 players. While these results support the above cost-benefit PD-game analysis, many real-world situations are not one-shot, but rather repeated encounters. The goal of the present study was to test if the results of the Locey et al. (2013) experiment would be found in a face-to-face repeated prisoner's dilemma game with real, rather than hypothetical, rewards.

1. Method

1.1. Participants

Eighty female undergraduate students were recruited through the psychology subject pool at Stony Brook University. Participants were paired and randomly assigned to one of two conditions described below (40 students, 20 pairs per condition).

1.2. Materials

Materials for each participant consisted of a deck of 40 index cards (20 with green circles and 20 with blue circles) that were used to signify cooperation or defection, a computer monitor with instructions, reward schedule, visual representation of points gained so far and (briefly) the number of points gained on the most recent trial. Fig. 1 shows the reward

¹In addition, Locey et al. (2013) found that the percentage of players cooperating in a multi-player PD game varied directly and linearly with the number of players (holding reward to each player constant). In other words, as the number of *other* players benefiting from a player's cooperation increased, so did the players' tendency to cooperate. This finding is further evidence that (socially discounted) reward to others may reinforce cooperation in PD games.

alternatives as presented to the participants. After playing the game, each participant completed a post-experimental questionnaire.

1.3. Procedure

The experiment consisted of two parts – an iterated face-to-face prisoner's dilemma (PD) game, and a questionnaire. The participants were seated on opposite sides of a table, facing each other. Participants were scheduled two at a time. The experimenter assigned each pair to the 1-2-3-4 or 1-2-9-10 condition (in alternating order). The participants were given the following instructions, followed by an opportunity for questions:

You will make a series of choices between two options – a card with a green circle or a card with a blue circle. You choose the card without revealing it to the other player, place it face down on the table and then I will ask you to turn the card over. You will then receive points according to the diagram above. The goal is to get as many points as you can. At the end of the game, you will be able to convert your points into a Starbucks or Barnes and Noble gift card as shown below.

Conversion to Gift Card: 100 – 159 (220 – 399) points = \$5 Gift Card; 160+ (400+) points = \$10 Gift Card. This is not a competition between you and the other player, as both of you will be able to receive a gift card and there is no bonus for having more points than the other player.

Both groups received identical instructions, except the minimum number of points for a gift card was adjusted upwards for the 1-2-9-10 group, as denoted above by the values in parentheses, such that 50% cooperation from both partners resulted in a \$5 gift card for each. The instructions remained on the screen alongside the reward alternatives (see Fig. 1) throughout the experiment. Each possible outcome was then demonstrated with respective payouts. Green signified cooperation; blue signified defection. If both participants played green (both put a green card face down on the table), each received a moderately high reward (3 points for the 1-2-3-4 group, 9 points for the 1-2-9-10 group); if both played blue, each received a moderately low reward (2 points either group); if one participant played green and one played blue, the one who played green received the lowest reward (1 point either group) and the one who played blue received the highest reward (4 points for the 1-2-3-4 group, 10 points for the 1-2-9-10 group).

The participants were then asked to start by choosing a card with a green circle or blue circle, and place it face down on the table. When both cards were face down, participants were asked to turn their card over. The results were then recorded via a keyboard, and the number of points earned flashed on each participant's screen. To slow the pace of the game, every 4 trials were separated by a 15-s pause. After every 16 trials, the participants were asked to shuffle their used cards back into their deck; this was done to discourage estimates about the length of the experiment by means of depletion of the cards in the deck and to replenish the supply of any preferred color. After 40 trials the game ended without forewarning.

After playing the PD game, participants completed a two-page questionnaire. The first page asked the participant to imagine standing on a vast field with the 100 people closest to her, each at a distance inversely proportional to her felt closeness to that person (Rachlin and Jones, 2008). Then they were asked to report the distance to the people ranked 1, 5, 20, and 100 as well the distance to a random psychology student and the person with whom they had just played the PD game.

The second page asked participants to indicate an optimal strategy, if any; the strategy they actually used, if any; and whether they had any prior knowledge of the PD game (see

Appendix A for both questionnaires). Upon completion, the participants were debriefed and given an opportunity for questions.

2. Results

Fig. 2 shows average percent cooperation over the first 8 trials, over all trials, and over the last 8 trials for the two groups. A factorial repeated measures ANOVA with blocks of 8 trials as the within-subject factor and condition as the between-subject factor showed no main effect for block, $F(4, 312) = 1.25, p = .29$, and a marginal effect for condition, $F(1, 78) = 2.80, p = 0.10$. There was, however, a significant interaction between block and condition, $F(4, 312) = 6.30, p = .001$. Considering the last 8 trials only, participants in the 1-2-9-10 condition cooperated significantly more than those in the 1-2-3-4 condition, $t(38) = 3.25, p = .002$. (Unless otherwise indicated, the unit of all statistical analyses was the pair rather than the individual player.)

With regard to change over the course of the experiment, a paired samples t-test comparing the first and last blocks within condition confirmed a significant reduction in cooperation for the 1-2-3-4 condition (53% cooperation in first block and 38% in last block), $t(19) = 3.55, p = .002$. However, for the 1-2-9-10 group, the difference between the means of the first and last blocks (51% and 57%, respectively) was not significant, $t(19) = -1.57, p = .13$.

Eleven of the 80 participants (3 in the 1-2-3-4 group and 8 in the 1-2-9-10 group) reported on the post-experimental questionnaire that they played randomly; 19 others (8 in the 1-2-3-4 group and 11 in the 1-2-9-10 group), failed to report using any particular strategy (e.g., “not sure,” or “tried to mix-it up”). With regard to their actual choices, over the last 20 trials, 44% of the participants (40% in the 1-2-3-4 group and 48% in the 1-2-9-10 group) cooperated about as many times as they defected (9, 10, or 11 times). This indifference may reflect confusion (even though the contingencies were on view throughout the experiment), it may reflect genuine ambivalence between the alternatives, or it may reflect continuous experimentation to test the reactions of the other player.

Cooperation rates were highly correlated within the pairs over the 40 trials ($r = .68, p < .001$), and over the last 8 trials ($r = .54, p = .001$). When partners' cooperation rates are plotted against each other, as in Fig. 3, there is a steeper slope and an appearance of a stronger correlation for pairs in the 1-2-9-10 condition relative to those in the 1-2-3-4 condition ($r = .84$ vs. $r = .55$) over the 40 trials. However, there was no statistically significant difference in correlation between the two groups over the 40 trials.

Turning to the questionnaire, overall, 18% of participants rated their partners as more distant than a random classmate, 29% placed the two at the same distance, and the remaining 53% rated their partners as closer than a random classmate. The ratings differ for the two conditions; in the 1-2-9-10 condition the majority of participants (25 out of 40) ranked their partners as closer than a random classmate, with only 2 participants ranking them oppositely. In the 1-2-3-4 condition, 17 participants ranked their partner as closer than a random classmate, but 12 rated them as more distant. Equidistant ratings were common in both conditions (11 such ratings in 1-2-3-4 and 12 in 1-2-9-10).

As was hypothesized, participants who placed the other player closer than a random student on the post-game questionnaire cooperated at a significantly higher rate than did those who placed her at the same distance or further than a random classmate. This was significant over all 40 trials, $t(77) = 2.06, p = .04$, and marginally significant over the last 8 trials $t(77) = 1.67$, respectively, $p = .10$. There was no significant difference in cooperation between participants who placed their partner further away than a random classmate and those who placed her at the same distance as a random classmate. However, this may have been due to

the low number of participants who placed their partner further away ($n = 14$), as opposed to the number who placed her at the same distance ($n = 23$), or closer ($n = 42$) than a random classmate.

There was no measured relationship between a *partner's* cooperation and relative closeness to the partner. In other words, participants who themselves tended to cooperate also felt closer to their partners than did those who tended to defect, but their partner's cooperation or defection did not appear to be related to this closeness. Moreover, social distance to the partner, as reported in the post-game questionnaire, did not correlate significantly between players, $r = .07$, $p = .65$.

In a two-person PD game, a player may use a tit-for-tat (TFT) strategy. A player (P) using TFT reciprocates the other player's (OP's) choice on the following trial; using TFT, P always cooperates on the trial following OP's cooperation and always defects on the trial following OP's defection. Use of TFT by P has been shown to increase OP's cooperation over repeated trials (Silverstein et al., 1998). We now ask whether participants in this experiment used TFT. We measured P's reciprocation as the difference between two fractions: the probability of P's cooperation following OP's cooperation, minus the probability of P's cooperation following OP's defection. The two fractions are independent; their possible difference ranges from +1 to -1. A difference of +1 indicates strict TFT; a difference of -1 indicates strict tat-for-tit (the inverse of TFT); a difference of 0 indicates no measured reciprocation by P. We call this difference, P's *reciprocation score*.

Reciprocation scores were obtained for all individuals and averaged for each group. The 1-2-3-4 group mean of +0.05 was not significantly different from 0, $t(39) = 1.4$, $p = .16$. However, the mean reciprocation score of +0.11 for the 1-2-9-10 group was significantly higher than 0, $t(37) = 3.42$, $p < .002$ (One pair, with no defections, was excluded from the analysis). An independent samples t -test revealed no significant difference in reciprocation scores between the two groups, $t(76) = 1.4$, $p = .16$. Although the reciprocation score varied and was different from 0 for some players, there was no significant correlation between P's reciprocation score and OP's cooperation over all 40 trials (one's own or partner's) or over the last 20 trials, $r = 0.02$, $p = .90$ and $r = .09$, $p = .43$, respectively.

3. Discussion

Locey et al. (2013) found that participants in a one-shot 1-2-3-4 PD game cooperated significantly less than did participants in a one-shot 1-2-9-10 PD game. However, in the present repeated trials experiment, there was no difference in cooperation between groups in the initial trial, or even over the first 8 trials, between the two groups. But there were several differences between the two situations. The participants in the Locey et al. experiment knew that they would be playing only once, and would not have another chance to choose, whereas those in the present experiment knew that they would be playing repeated trials, and might have been willing to test the contingencies. Also, participants in the Locey et al. experiment played in their dorm rooms or at home, on a computer, without feedback, against an anonymous partner, for a hypothetical reward. On the other hand, participants in the present experiment played in the laboratory, using cards to indicate their choices, with feedback, in the presence of the experimenter, face-to-face with a real partner, for points convertible to a real reward. Any of these several differences, or a combination of them, could have been responsible for the difference in initial-trial results.

By the end of the present experiment, however, participants in the 1-2-9-10 group (where P's cooperation resulted in a greater reward to OP) cooperated at a significantly higher rate than did participants in the 1-2-3-4 condition. Let us look at this result in terms of costs and

benefits of cooperating. The cost of cooperation was equal for the two groups; all Ps received one point less by cooperating than by defecting regardless of OP's choice. However, the benefit of cooperating (in terms of the gain to the other player) was greater for the 1-2-9-10 group (8 points) than for the 1-2-3-4 group (2 points)—again, regardless of OP's choice. However, these benefits must have been discounted by the social distance to OP. We can estimate this discount.

Eq. (1) (Jones and Rachlin, 2006) expresses the discounted value of a reward to another person (v) as a function of the reward's undiscounted value (V), the social distance to the other person (N), and a constant, k :

$$v = \frac{V}{1 + kN} \quad (1)$$

Prior experiments (Locey et al., 2013) estimated kN for an average Stony Brook student relative to a classmate to be about 3.75. This yields a discounted benefit (v) for cooperation of 0.42 units for the 1-2-3-4 group ($V = 2$ units) and 1.68 units for the 1-2-9-10 group ($V = 8$ units). The former is less than the 1-unit cost of cooperation and the latter is greater than the 1-unit cost, predicting mutual defection for the players in the 1-2-3-4 group and mutual cooperation for the players in the 1-2-9-10 group. If PD game choices depended on socially discounted reward, the players in the 1-2-9-10 group would be predicted to cooperate at a higher rate than those in the 1-2-3-4 game – as was found.

In postgame questionnaires, over both groups, players who rated their partners as socially closer than a random classmate (relatively low kN) tended to cooperate more than players who rated their partners as socially more distant than a random classmate (relatively high kN). These results suggest that social discounting was a factor in learning to cooperate or defect in the PD games of this experiment. However, as the social discounting measure always followed the PD game (the order was not counterbalanced), the direction of the effect cannot be determined. Thus, it is possible that participants who cooperated more felt closer to the participant as a result of their own behavior.

Another factor often hypothesized to generate cooperation in PD games is a general tendency to reciprocate in exchanges with another person (Fehr and Fischbacher, 2005). Indeed, when reciprocation by P is explicitly varied, cooperation by OP tends to increase over repeated PD trials, and the increase is proportional to the degree of reciprocation (Baker and Rachlin, 2001). However, there was little evidence that *next-trial* reciprocation was a factor in the present experiments. Although players in the 1-2-9-10 group did reciprocate at a significant rate, this rate was low (+0.11 on a scale in which 0 would indicate no reciprocation and +1 would indicate perfect reciprocation) and it was not significantly different from reciprocation of the 1-2-3-4 group (+0.05). Although a trial-by-trial, tit-for-tat strategy does not appear to be responsible for the participants' behavior, this does not mean that participants did not influence each other. The significant correlation of cooperation between partners in this experiment indicates that there was a strong overall (molar) mutual influence within pairs.

Another possible reason for the difference between the two groups in cooperation is that what counts for cooperation may be ratios rather than differences between rewards. For example, we assumed that the cost of cooperating when the other player cooperated was equal for the two groups ($4 - 3 = 10 - 9$). However, if it were ratios rather than differences that counted, the cost of cooperating would have been greater for the 1-2-3-4 group ($4/3 > 10/9$). Thus, the 1-2-9-10 group may have cooperated more than the 1-2-3-4 group because the effective cost of doing so was less. But such a cost-based model does not explain why

participants should have cooperated at all. With the exception of reciprocation models, nearly all models of PD behavior (e.g., Coombs, 1973) predict only a stronger or weaker tendency to defect.

4. Conclusions

What reinforces cooperation in a PD game? The purpose of the present experiment was to test one potential answer: cooperation by P is reinforced by socially discounted reward to OP. The results constitute some evidence that this was the case. The principal alternative answer is that cooperation is reinforced by the other player's reciprocation, delayed until the following trial (Fehr and Fischbacher, 2005). One answer relies on socially discounted rewards, the other on delay discounted rewards. With respect to delay, Locey and Rachlin (2012) found that when delay of PD-game reciprocation was explicitly varied, cooperation varied inversely with delay. Although the present experiment did find a significant correlation between P's cooperation and stated closeness to OP, and failed to find any evidence of direct reciprocation of cooperation (on the immediately following trial), there was a strong, significant correlation between cooperation by P and OP for both groups – stronger for the group that cooperated more (1-2-9-10). This result suggests that reciprocation may have been a factor but, if it was a factor, it could not have acted by means of a tit-for-tat strategy where P reciprocated OP's choice on the *immediately* following trial. Rather, P may have been sensitive to OP's overall, or molar, rate of cooperation or defection, and reciprocated in kind. Direct sensitivity to molar variables without mediation by immediate contingencies is a pervasive finding in the behavioral literature with human and nonhuman participants, and has been used to account for learning of self-control in situations (such as smoking addiction or alcoholism) where there is no reinforcement of an individual self-controlled act (Rachlin, 2000).

Both social closeness and molar reciprocation may have been simultaneously active in this experiment. Which is the stronger (as well as the degree to which either is innate or learned) are subjects for further research. The question, which is stronger, should be testable since the social discounting explanation of cooperation implies that P's cooperation is essentially independent of OP's behavior during the game (unless N or k changes as a function of OP's behavior), and the molar reciprocation explanation of cooperation implies that P's cooperation is strictly dependent on OP's behavior during the game.

Acknowledgments

This research was supported by Grant R01MH04404916 from The US National Institute of Mental Health.

References

- Axelrod R. Effective choice in the Prisoner's Dilemma. *Journal of Conflict Resolution*. 1980; 24:3–25.
- Baker F, Rachlin H. Probability of reciprocation in repeated prisoner's dilemma games. *Journal of Behavioral Decision Making*. 2001; 14:51–67.
- Coombs CH. A reparameterization of the prisoner's dilemma game. *Behavioral Science*. 1973; 18:424–428.
- Fehr E, Fischbacher U. The nature of human altruism. *Nature*. 2003; 425:785–791. [PubMed: 14574401]
- Fehr, E.; Fischbacher, U. The economics of strong reciprocity. In: Gintis, H.; Bowles, S.; Boyd, R.; Fehr, E., editors. *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press; Cambridge, MA: 2005. p. 151-191.
- Jones BA, Rachlin H. Social discounting. *Psychological Science*. 2006; 17:283–286. [PubMed: 16623683]

- Locey ML, Safin V, Rachlin H. Social discounting and the Prisoner's Dilemma game. *Journal of the Experimental Analysis of Behavior*. 2013; 99:85–97. [PubMed: 23344990]
- Locey ML, Rachlin H. The temporal dynamics of cooperation. *Journal of Behavioral Decision Making*. 2012; 25:257–263. [PubMed: 22754115]
- Rachlin, H. *The Science of Self-control*. Harvard University Press; Cambridge, MA: 2000.
- Rachlin H, Jones BA. Social discounting and delay discounting. *Journal of Behavioral Decision Making*. 2008; 21:29–43.
- Silverstein A, Cross D, Brown J, Rachlin H. Prior experience and patterning in a prisoner's dilemma game. *Journal of Behavioral Decision Making*. 1998; 11:123–138.

Appendix A

The following experiment asks you to imagine that you have made a list of the 100 people closest to you in the world ranging from your dearest friend or relative at position #1 to a mere acquaintance at #100. The person at #1 would be someone you know well and is your closest friend or relative. The person at #100 might be someone you recognize and encounter but perhaps you may not even know their name.

You do not have to physically create the list – just imagine that you have done so.

Now try to imagine yourself standing on a vast field with those 100 people. The actual closeness between you and each other person is proportional to how close you feel to that person. For example, if a given person were 10 feet away from you, then another person to whom you felt twice as close would be 5 feet away from you and one to whom you felt half as close would be 20 feet away. We are going to ask you for distances corresponding to some selected individuals of the 100 on your hypothetical list.

Remember that there are no limits to distance – either close or far; even a billionth of an inch is infinitely divisible and even a million miles can be infinitely exceeded. Therefore, do not say that a person is zero distance away (no matter how close) but instead put that person at a very small fraction of the distance of one who is further away; and do not say that a person is infinitely far away (no matter how far) but instead put that person at a very great distance compared to one who is closer.

Of course there are no right or wrong answers. We just want you to express your closeness to and distance from these other people in terms of actual distance; the closer you feel to a person, the closer you should put them on the field; the further you feel from a person, the further they should be from you on the field. Just judge your own feelings of closeness and distance.

Feel free to use any units you wish (inches, feet, football fields, miles, etc.).

But be sure to indicate what the unit is.

How far away from you on the field is the 1st person on your list?	#	&	units
How far away from you on the field is the 5th person on your list?	#	&	units
How far away from you on the field is the 20th person on your list?	#	&	units
How far away from you on the field is the 100th person on your list?	#	&	units
How far away from you on the field is a randomly selected student from one of your psychology classes?	#	&	units
How far away from you on the field is the person with whom you just played this game (the one sitting across from you at the table)?	#	&	units

Choice behavior questionnaire

1. What do you think the game was about?
2. How did you decide which card to play?
3. How do you think the other player decided which card to play?
4. Do you think there was an optimal strategy? If so, what?
5. Have you ever heard of a prisoner's dilemma game? If so, what do you think it is about?

Fig. 1.

The possible reward outcomes were presented in this format to be visually accessible to the participants. Fig. 1a was presented to participants in the 1-2-3-4 condition and 1b to participants in the 1-2-9-10 condition. The circles were presented in color (green, here as light gray, signified cooperation; blue, here as black, signified defection). Mutual cooperation yielded a higher reward than mutual defection (3 or 9 points each as opposed to 2 points each), but the highest payout on an individual trial could only be gained through defection (4 points for the 1-2-3-4 group; 10 points for the 1-2-9-10 group) when the partner cooperated on that trial. The numbers in the circle indicate the number of points the participant who played that color received when the other player played the adjacent color. Thus when both participants selected the same color, the payout was symmetrical, and when they selected different colors it was asymmetrical.

Fig. 2.

Average cooperation rate between conditions for the first 8 trials, across all 40 trials, and for the last 8 trials. The difference of mean cooperation rates between the two groups was significant for the last 8 trials, $p < .01$. The SEM error bars are calculated for pairs rather than individuals.

Fig. 3.

P's cooperation plotted against OP's cooperation over 40 trials. Each black diamond represents one pair from the 1-2-3-4 group, and the black line represents the best linear fit ($r = 0.55$, $p < .05$). Each gray square represents one pair from the 1-2-9-10 group, and the gray line represents the best linear fit ($r = 0.84$, $p < .01$).