



**Rational Cooperation in the Finitely Repeated Prisoner's Dilemma:
Experimental Evidence**

James Andreoni, John H. Miller

The Economic Journal, Volume 103, Issue 418 (May, 1993), 570-585.

Stable URL:

<http://links.jstor.org/sici?sici=0013-0133%28199305%29103%3A418%3C570%3ARCITFR%3E2.0.CO%3B2-O>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Economic Journal is published by Royal Economic Society. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/res.html>.

The Economic Journal
©1993 Royal Economic Society

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2002 JSTOR

RATIONAL COOPERATION IN THE FINITELY REPEATED PRISONER'S DILEMMA: EXPERIMENTAL EVIDENCE*

James Andreoni and John H. Miller

In the finitely repeated prisoner's dilemma, it is well known that defection in every game is the unique dominant-strategy Nash equilibrium. This follows from the familiar backward-induction arguments. Kreps *et al.* (1982), however, show that if there is incomplete information about the types of players then cooperation early in the game can be consistent with rational behaviour.¹ Suppose, that both players believe that there is a small chance that their opponent may be altruistic. For instance, the opponent may get extra pleasure from mutual cooperation or may even adopt a tit-for-tat strategy. Then it could be in each player's best interest to pretend, at least for some time, to be an altruistic player in order to build a reputation for cooperation, until the game eventually unravels to mutual defection.

The sequential equilibrium reputation hypothesis has become influential in many literatures. It has become important to know whether this hypothesis has good predictive power, and whether individuals will rationally build reputations. It is also of interest to know whether some fraction of the population actually have altruistic motives. In literatures on social dilemmas there has been extensive discussion about whether altruistic concerns, like gaining extra pleasure from mutual cooperation, are necessary for the characterisation of preferences.²

There is some evidence from experiments on both reputation building and altruism. Camerer and Weigelt (1988) consider an eight-period game of loan contracts, and find that the behaviour of subjects largely meets the sequential equilibrium prediction, although lenders are slightly more optimistic about the probability of repayment than the experimental controls merit. Camerer and Weigelt refer to this optimism as 'homemade priors' that subjects bring to the experiment from outside and use to supplement the priors controlled for by the experimenter. Adjusting for these homemade priors, Camerer and Weigelt find a close match with the theory.³ Selten and Stoecker (1986) also find that, with sufficient experience, subjects appear 'to learn the sequential equilibrium,

* We are grateful to Robyn Dawes, Paul Milgrom, John Carter, and two anonymous referees for helpful comments, and to Dan Schneidewend and Soren Hauge for expert programming and research assistance. Andreoni also thanks the National Science Foundation, grant SES 8821204, for financial support. Errors are the responsibility of the authors.

¹ See also Kreps and Wilson (1982), and Milgrom and Roberts (1982).

² See, e.g. Palfrey and Rosenthal (1988), Andreoni (1989, 1990), and Cooper *et al.* (1990) for a discussion of this.

³ See Camerer and Weigelt (1988) for a more complete discussion of other experiments that pertain to reputation building.

although subjects are more cooperative than predicted. McKelvey and Palfrey (1992) find significant evidence of reputation building in centipede-game experiments, but also find important effects of altruism. Further evidence on altruism comes from the literatures on prisoner's dilemma and public goods, where experiments indicate that subjects may have their own altruistic preferences that interfere with the incentives of the experiment. Such effects have been identified by Palfrey and Rosenthal (1988), among others.⁴

This paper examines cooperation in the finitely repeated prisoner's dilemma by directly testing the model posed by Kreps *et al.* (1982). We consider a series of finitely repeated prisoner's dilemma games in which we manipulate subjects' beliefs about their opponent's type. By raising the probability that a player will have an altruistic opponent, we increase the benefits to reputation building. Of course, if subjects really do have altruistic preferences, then we cannot completely control the homemade priors that subjects may bring to the experiment from outside. Therefore, we also include a control group that plays repeated single-shots of prisoner dilemma, and thus has no opportunity to build reputations. By comparing this group to the others, we are able to measure the effect of reputation building over 'homemade altruism', that is, people's natural tendency to cooperate.

The result is that the sequential equilibrium reputation model appears to be a good predictive model of cooperative behaviour in the finitely repeated prisoner's dilemma. Subjects seem to undertake significant efforts to build reputations for altruism. However, we also find that those reputations are well-deserved. In the group that cannot build reputations, we find a consistent pattern of cooperation that does not deteriorate, even after 200 single shots of the prisoner's dilemma. Hence, there clearly appears to be a significant number of 'altruistic types' in the population. This finding is consistent with evidence from other social sciences. For instance, in a detailed analysis of prisoner's dilemma experiments, psychologists Kelley and Stahelski (1970) conclude, 'There are two stable types of individuals which may be described approximately as cooperative and competitive personalities' (p. 66).⁵ These cooperative or altruistic players appear to form the basis for reputation building. While in theory all that is required for cooperation is sufficient beliefs that altruists exist, in practice such beliefs appear to be consistent with actual tastes for cooperation.

I. THEORETICAL BACKGROUND

Fig. 1 shows the prisoner's dilemma payoff matrix used in the experiments reported here. Kreps *et al.* (1982) describe an equilibrium in the finitely

⁴ See Roth and Murnighan (1978) and Roth (1988) for reviews of prisoner's dilemma experiments. For a historical review of prisoner's dilemma experiments in psychology and sociology, see Rapoport and Chamah (1965). See Dawes and Thaler (1988) for a recent review and discussion of cooperation in providing public goods.

⁵ For reviews and discussions of the psychology and sociology literatures on cooperation and altruistic behaviour, see Dawes (1980) and Piliavin and Charng (1990).

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	7, 7	0, 12
	Defect	12, 0	4, 4

Fig. 1. The prisoner's dilemma.

repeated prisoner's dilemma in which two rational players both believe that there is a small probability, δ , that the other is 'irrational'. They give two examples of irrationality. First, the opponent may be playing a tit-for-tat strategy, which begins by cooperating, and then plays whatever its opponent played on the last round. Second, players could believe the opponent may get extra utility from mutual cooperation, such that cooperation is the best response to cooperation. In each case, a sufficiently high δ can lead each player to adopt a strategy of the sort 'cooperate until round T , or until my opponent defects, and defect thereafter'. Higher values of δ will tend to increase the amount of cooperation.

A strict interpretation of the Kreps *et al.* theory is that no 'irrational' or 'altruistic' types need to exist, but only that there are sufficient beliefs that such types exist. This has been called the rationality hypothesis. It has been noted that this strict interpretation of the model requires players' beliefs about types to differ from the actual distribution of types (Samuelson, 1987). Hence, the model would be more natural if some players actually were altruistic. This alternative has become known as the altruism or 'warm-glow' hypothesis, and has been suggested by many researchers.⁶ All of the alternative models of altruism can be viewed as one of, or some combination of, three similar models. Each model includes a single altruism parameter, α , but makes different assumptions about its use. The three models are: (i) *Pure Altruism*. Let p_i be the payoff of person i . Then under this model, the utility of player i is $U_i = p_i + \alpha p_j$, $0 \leq \alpha \leq 1$. Hence, these players care directly about the payoff of the other player. (ii) *Duty*. Utility $U_i = p_i + \alpha$, where $\alpha \geq 0$ whenever i cooperates and is zero otherwise. Here, i feels an obligation to cooperate. (iii) *Reciprocal Altruism*. Utility $U_i = p_i + \alpha$, where $\alpha \geq 0$ whenever both i and i 's opponent cooperate and is zero otherwise. Here there is special pleasure in successful cooperation. This is the model suggested by Kreps *et al.* (1982).

In general, (i) and (ii) can support three equilibrium strategies in the single shot game, depending on the payoff parameters in the prisoner's dilemma. The

⁶ In addition to Kreps *et al.* (1982), see Palfrey and Rosenthal (1988), Dawes (1980), Stark (1985), Camerer (1988), and Cooper *et al.* (1990), among others.

first two are dominant strategies for either defection (if α is sufficiently small) or cooperation (if α is sufficiently large). The third equilibrium strategy is a matching strategy, where players find cooperation a best response to cooperation, and defection a best response to defection. Hence, if the chance is sufficiently high that one's opponent will cooperate, then the optimal strategy for an altruist is to cooperate. For the parameter values chosen for our experiment, it is easy to show that only the first two strategies are possible equilibrium strategies in the one shot game.⁷ This means that if (i) or (ii) is the correct model, then only a dominant strategy of cooperation or a dominant strategy of defection are possible. Model (iii) differs from these in that the matching can be an equilibrium strategy if α is greater than the temptation payoff (12 here) minus the cooperative payoff (7 here). If α is less than this amount, then there is a dominant strategy for defection. Also, there can be no dominant strategy of cooperation under reciprocal altruism, no matter how big α . Other theories of altruism can be generated with combinations of the above, so these models need not be mutually exclusive. For the parameters chosen, however, it is clear that extra utility from mutual cooperation is essential for cooperation to emerge. For this reason, and to maintain consistency with *Kreps et al.* (1982), we will focus further discussion on the reciprocal altruism model of warm-glow alone.

If people are altruistic, then one implication is that cooperation can be maintained in single-shot plays of the prisoner's dilemma, even without the possibility of reputation building. When we consider the finitely repeated prisoner's dilemma with altruism, we also get a sequential equilibrium prediction that repeated play should increase cooperation. Unless there is common knowledge that everyone in the population has an α parameter that supports the cooperative equilibrium, then it may pay all subjects to build a reputation for being altruistic, but to defect late in the game. In this way, the sequential equilibrium predictions are the same for both altruistic and non-altruistic populations. However, we should not expect all patterns of play to be independent of the degree of altruism. According to the rationality hypothesis as it becomes increasingly clear that the population is all 'rational', cooperation should become increasingly difficult to maintain. Hence, in a series of finitely repeated games, against a variety of opponents, defection should tend to occur earlier and earlier in each repeated game (Selton and Stoecker, 1986). On the other hand, if people are altruistic, then as the true proportion of altruists becomes known one could observe cooperation extending until later and later, since altruistic subjects will become more confident that cooperation will be reciprocated.

The next section will outline the design of the experiment used to discuss reputation building, and will specify the hypotheses we will examine.

⁷ Consider duty. For cooperation to be a best response to cooperation, it must be that $\alpha \geq 5$ (i.e. $12 - 7 = 5$). For defection to be a best response to defection, it must be that $\alpha \leq 4$ (i.e. $4 - 0$). These conditions cannot be met simultaneously, eliminating the mixed strategy. Similar results hold for pure altruism.

II. EXPERIMENTAL DESIGN

This experiment was run with four conditions, each requiring 14 subjects.⁸ Subjects interacted over a computer network. Each subject participated in only one session, and each session ran only one condition. Subjects were randomly assigned to computer terminals at the start of the experiment. They were given written instructions, which were read aloud at the start of the session. Each session lasted less than 90 minutes, and subjects earned an average of \$11.65. All subjects were paid privately, in cash, at the end of the experiment. Subjects were recruited from introductory microeconomics courses at the University of Wisconsin. A copy of the subjects' instructions is included as an appendix to this paper.

The four conditions are the following:

1. *Partners.* The computer randomly paired the 14 subjects, and each subject played a 10-period repeated prisoner's dilemma with their partner. All pairings were anonymous. Subjects then received a summary of their earnings from the ten-period game. They were then randomly rematched with another partner for another 10-period game. This was done for a total of 20 10-period games, that is, for a total of 200 rounds of the prisoner's dilemma.

In the partners condition, subjects play a series of finitely repeated games, each time with a new partner. In every ten-period game, therefore, subjects can gain from reputation building. However, since there are only 14 subjects, as the experiment progresses subjects should get greater and greater knowledge about the true distribution of types. Under a hypothesis of no altruism, one would predict that with each 10-period game cooperation should become harder to sustain, especially near the end of the experiment (Selten and Stoecker, 1986). However, if there really are altruistic types, then a set pattern of cooperation can be sustained throughout the experiment. In fact, under altruism, cooperation could increase toward the end of the experiment as people become more certain that altruism will be reciprocated.

2. *Strangers.* The computer randomly paired the 14 subjects for every iteration of the prisoner's dilemma, for a total of 200 iterations. That is, each subject had a new partner every iteration. To make sure that there were no presentation differences between the strangers and the partners, subjects were also given summaries of their performance every 10 rounds, as was done with the partners.

For strangers, there is no incentive for any one subject to build a reputation. Under a perfect rationality assumption,⁹ there should be no cooperation in this group, especially by the end of the experiment. Note, however, that the number of subjects is small relative to the number of rounds so it is not inconceivable that the group as a whole could build a reputation (Kandori, 1992). We examine this hypothesis below.

3. *Computer50.* The instructions for this group were identical to the

⁸ Unfortunately, only 12 subjects could be recruited for condition 4, computero.

instructions for the partners, except that this group was given a 50% chance of meeting a computer partner in any 10-period game, rather than another subject in the room. They were told the computer would play the tit-for-tat strategy (called 'copy cat' in the instructions).

This conditions takes the Kreps *et al.* (1982) hypothesis literally. Relative to the partner's group, subjects in this condition should have greater confidence that they may be playing an altruistic opponent, and under the sequential equilibrium reputation hypothesis should be more cooperative than the partners. If they are not more cooperative, this would contradict a model of sequential equilibrium.

4. *Computero*. This condition is equivalent to the computer50 condition, except that subjects were told that the chance of playing the tit-for-tat partner was 1/1000, i.e. 0.1%.

If more cooperation is observed in the computer50 condition than in the partners condition, it could be that common knowledge of the tit-for-tat strategy, rather than changes in the probability of playing an altruist, could be influencing play. For this reason we also ran the computero condition, which is equivalent to computer50 except that subjects were told that the chance was approximately zero that they would actually play the computer. In this case, the tit-for-tat strategy was common knowledge, but the probability of playing an altruist was not (directly) increased. If common knowledge of the tit-for-tat strategy alone is sufficient to encourage altruism, then computeros should be more cooperative than partners.

During each iteration of the game subjects in each condition were told their last round decision, the decision of their opponent, and their earnings. They were also told how many rounds remained with their current partner (zero for strangers), and how many rounds remained in the experiment in total. All subjects were given a recap of their earnings every 10 periods, and all terminals beeped at the recap. In addition, every time the computer randomly rematched the subjects, the words 'New Partner' flashed on the computer screen. Subjects were also given an option of reviewing all of their previous periods of play at any time by hitting a single key on the keyboard, and then paging up and down.

Subjects in all four conditions also participated in an unanticipated 'restart'. After all 200 rounds of the main experiment were complete, subjects were told that they would play an additional 10-period game. For these ten periods the subjects were matched with a computer player rather than another person, and the computer played the tit-for-tat strategy for sure. This was announced verbally to all subjects, and a description of the tit-for-tat strategy was written on a chalkboard for all subjects to see.

The purpose of the restart is to gauge the strategic sophistication of the subjects, independent of any altruism. The optimal strategy is to cooperate for 9 periods, and defect on the tenth. One might suspect that experience with the repeated game, and especially experience with tit-for-tat players, might increase a subject's ability to choose the optimal strategy when faced with the sure prospect of playing a computer using tit-for-tat. Hence, if subjects learn

the sequential equilibrium, one might expect that partners, computer50s, and computer0s will perform better in the restart than strangers, and that perhaps even that computer50s should perform better than the rest.

III. RESULTS

In this section we present evidence that the behaviour of subjects is consistent with the predictions of the sequential equilibrium reputation model. However, the data also suggests that, in addition to holding beliefs that a fraction of the population may be altruistic, a significant share of the subjects actually appear to be altruistic.

The Sequential Equilibrium Hypothesis

Fig. 2 illustrates the average percent cooperation across all ten rounds of the repeated game, for all 20 10-period games. Comparing the partners, strangers and the computer50s, we see patterns that are consistent with the predictions of the sequential equilibrium reputation model. First, partners are more

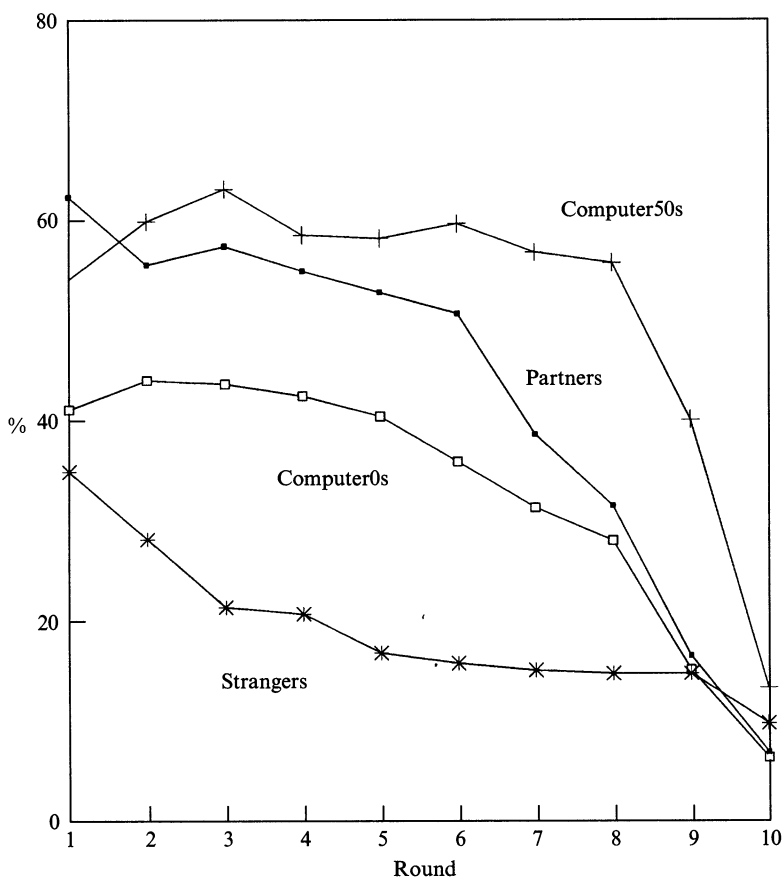


Fig. 2. Percent cooperation by round. Averaged over all 20 10-period games.

cooperative than strangers. Using a Mann–Whitney nonparametric test, we see that subjects in the partners condition cooperate significantly more than strangers, with $z = 3.079$, which is significant at the $\alpha \leq 0.001$ level.⁹ Likewise, computer50s are also significantly more cooperative than strangers ($z = 3.359$, $\alpha \leq 0.001$). Second, for both computer50s and partners, cooperation is highest in the early rounds and declines near the endgame. Splitting each condition into two groups, rounds 1–5 and rounds 6–10, we can compare the behaviour of subjects in the two halves of each 10-round game. Both partners and computer50s are significantly more cooperative in the first five rounds of each game ($z = 3.217$, $\alpha \leq 0.001$ for partners, and $z = 1.723$, $\alpha \leq 0.04$ for computer50s). However, for strangers, there is no significant difference between the first five and last five rounds of each repeated game ($z = 1.36$). All of these results are consistent with the reputation building hypothesis. Third, the computer50s are significantly more cooperative than partners. The significance of this difference shows up entirely in the second half of the 10-round game. Looking at only rounds 1–5, the levels of cooperation are roughly the same for partners and computer50s ($z = 0.850$). However over the final 5 rounds of each 10-round game the computer50s are significantly more cooperative than partners ($z = 2.137$, $\alpha \leq 0.01$). This implies that the main difference between partners and computer50s is that subjects in the computer50 condition simply wait until later in the game to defect. Again, this is fully consistent with the predictions of the sequential equilibrium hypothesis.

Next we examine the computero condition. First we can observe that the behaviour of the computeros is not significantly different from partners¹⁰ ($z = -1.28$). However, like the partners, computeros are significantly more cooperative than strangers ($z = 1.93$, $\alpha < 0.03$), and significantly less cooperative than computer50s ($z = -2.11$, $\alpha < 0.02$). Hence, simply making the tit-for-tat strategy common knowledge does not by itself appear to have any significant impact on cooperation.¹¹

Finally, it is interesting to note that the average level of cooperation in the end-game, round 10, is virtually identical for partners, strangers, and computeros, and is only slightly higher for computer50s. This is also consistent with the sequential equilibrium hypothesis.

⁹ The Mann–Whitney test statistic is approximately normal. This test will also be used in all subsequent test statistics reported. To calculate the statistic, begin by finding the average percentage cooperation of each subject in the two samples to be compared. Pool the samples and rank them. The statistic then looks for significant differences in the rank sums across conditions. This non-parametric test is superior to tests based on means for samples of this size because it is not easily influenced by the actions of a small number (i.e. one or two) of subjects.

¹⁰ The result is similar if we consider rounds 1–5 and rounds 6–10 separately, with $z = -1.598$ and $z = -1.031$ respectively.

¹¹ One curious observation from Fig. 2 is that for both computer50s and computeros cooperation actually peaks in the second or third round, rather than the first. This appears to be due to a small number of ‘testers’, who, in early rounds of the experiment, began every 10-period game by defecting in order to ‘test’ if they were playing a tit-for-tat opponent. However, subjects quickly learned the futility of this, and over the last half of the experiment there was very little behaviour that could be seen as testing. These testers may explain why the mean level of cooperation among the computeros is actually below that of the partners. In particular, testers in early games may have reduced the ‘homemade priors’ on altruism, diminishing the expected benefits to reputation building.

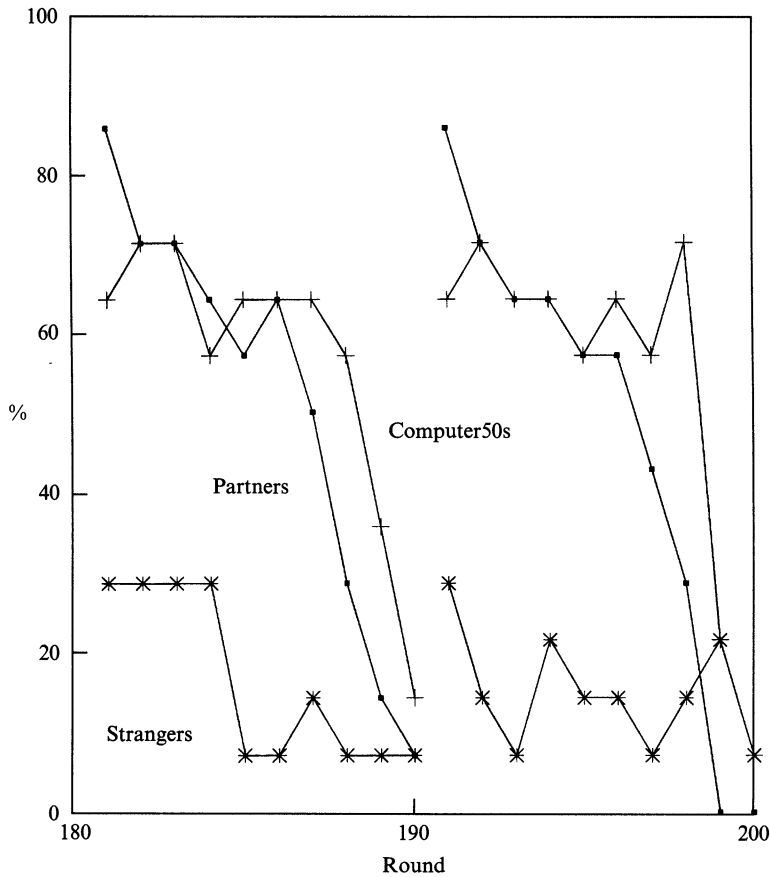


Fig. 3. % cooperation by round: Final 20 rounds of play.

The typical patterns of play in this experiment can be seen most clearly in Fig. 3, where we show the final two 10-period games (rounds 181 to 200) for the partners, strangers, and computer50s.¹² As predicted, cooperation by the partners peaks in round one, at 86%, and stays above 50% for 4–6 rounds before falling to zero. For computer50s, cooperation is level at about 60–70% for the first 8 periods, until it falls to about 7–14%. Strangers, on the other hand, vary cooperation frequently over the ten-period set, with cooperation between 7 and 28%. Again, these patterns are consistent with the sequential equilibrium reputation hypothesis.

The Altruism Hypothesis

In the last subsection we saw evidence that subjects were willing to build a reputation for altruism. This leaves the more subtle question of whether some

¹² For ease of presentation, we did not include the computeros in the figure. However, their relative position is like that in Fig. 2; they are significantly more cooperative than strangers, but not as cooperative as the other two.

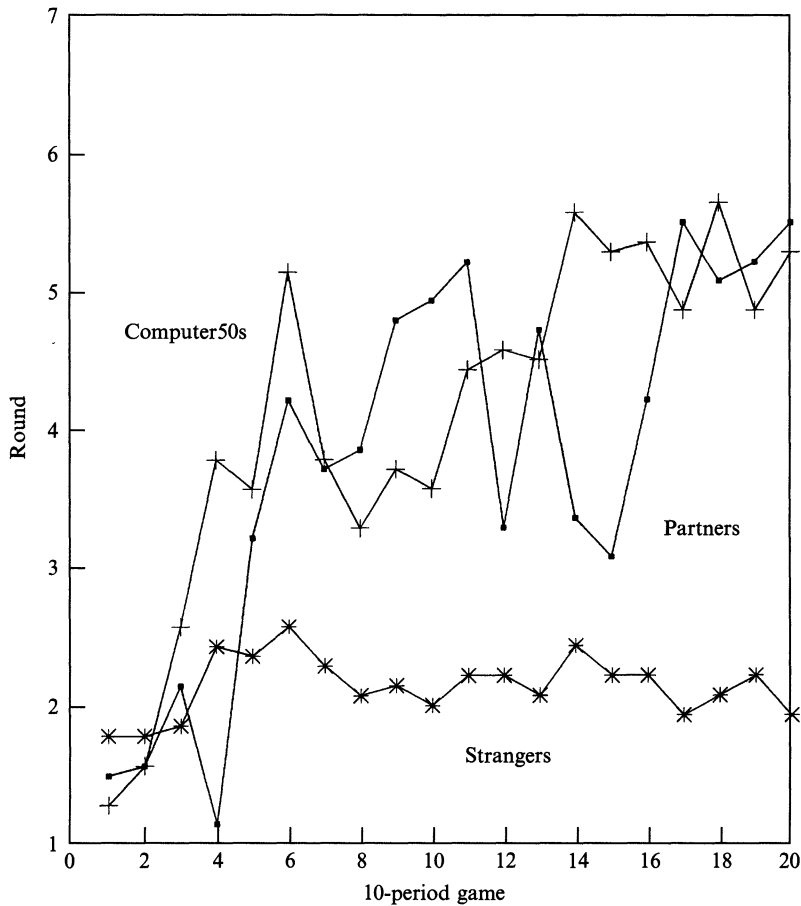


Fig. 4. Mean time until first defection.

subjects actually are altruists. Evidence on this can be found in Fig. 4. This shows the mean round of first defection for partners, computer50s and strangers.¹³ Contrary to the rationality hypothesis, partners and computer50s waited longer until their first defection as the experiment progressed, even in the final games. This is consistent with a hypothesis of altruism in which subjects continue to update their priors on the degree of altruism in the population throughout the experiment. Looking at the strangers, after a brief initial increase in the percent of cooperation, the mean time until first defect remained remarkably stable over the course of the experiment. This again is

¹³ Results similar to those in Fig. 4 obtain if the median is used. In calculating the means and medians, subjects who played all-cooperate were assumed to defect on round 11. However, there were very few such subjects. The computeros are not presented, but they were again more cooperative than strangers, but less than partners. One can also note in Fig. 4 that there is more variance among partners and computer50s than among strangers. This, probably, reflects floor effects among strangers, who are mostly choosing to defect, rather than reflecting any significant behavioural differences.

consistent with the view that, after learning the distribution of altruism in the population, a stable set of cooperative players developed.

Looking more closely at the strangers we can see additional evidence of altruism. If there is no real altruism, strangers should reach mutual defection at some point in the game and remain there. If there is a possibility for group reputations, then this too should be diminished as the end of the experiment approaches. As can be seen in Fig. 3, cooperation did not deteriorate at the end of the game. In fact, a detailed look at the data shows that the general pattern of cooperation among strangers illustrated in Fig. 2 is representative of the level of cooperation throughout the entire experiment.

The fact that strangers develop a stable pattern of cooperation suggests that perhaps they may be playing a Nash equilibrium game where subjects have incomplete information about the altruism of their opponents. We can examine this hypothesis by considering the model of reciprocal altruism discussed earlier. If this were the true model, then individuals would know their own altruism parameter α , but not their opponent's. All subjects would have prior beliefs about the distribution of α 's and would be playing a Nash equilibrium game of imperfect information. Assuming that all subjects have common priors about the distribution of types, then we can solve for a critical value of α , α^* , such that all subjects with $\alpha > \alpha^*$ will always cooperate and all subjects with $\alpha < \alpha^*$ will always defect. Those with α equal to α^* will, in equilibrium, choose a mixed strategy.¹⁴ Given the payoff parameters specified in Fig. 1, those with $\alpha = \alpha^*$ will cooperate with probability $p^* = 4/(\alpha^* - 1)$, where α^* is some positive number greater than 5 (since this is the difference between the temptation payoff 12 and the cooperation payoff 7).

The above equilibrium indicates that we should observe three types of subjects: *cooperators* who only cooperate, *defectors* who only defect, and *mixers* who cooperate with probability p^* . Notice that this imperfect-information equilibrium imposes a certain amount of symmetry on the outcome of the game. Suppose a mixer observes cooperation with probability $p_0 > p^*$. Then a mixer should update his beliefs about the distribution of types and become a cooperator. Likewise, if the observed p is less than p^* the subject should become a defector. Let p_m be the probability of cooperation by a mixer, and let p_n be the probability of cooperation by a non-mixer, that is, cooperators and defectors combined. Define π as the proportion of mixers in the population. Then the probability of cooperation that a mixer actually observes is $p_0 = \pi p_m + (1 - \pi) p_n$. Since in equilibrium $p_m = p^*$ and $p_0 = p^*$, it follows that $p_n = p^*$ in equilibrium as well. This implication of the imperfect-information equilibrium will serve the basis of our test of the model.

To examine this incomplete information equilibrium, we begin by examining the strangers condition for the last half of the experiment. By this point subjects have experience in 100 games, and should have a well developed sense of the

¹⁴ Solving for the equilibrium can be sketched as follows. Let $f(\alpha)$, $0 \leq \alpha < \infty$ be the distribution of types. Then let α^* be the critical level of α and let p^* be the equilibrium probability of cooperation. Then the equilibrium can be solved from the equations $p^*(7 + \alpha^*) + (1 - p^*)0 = p^*12 + (1 - p^*)4$, and $p^* = \int_{\alpha^*}^{\infty} f(\alpha) d\alpha$.

probability of cooperation in the population. Then define defectors as those subjects whose behaviour is not significantly different from a strategy of total defection. At the 99% confidence interval, this requires defecting at least 93.7% of the time. This method identifies five subjects who can be classified as defectors, with defection rates of ranging from 97 to 100%, and an average rate of 98.8% defection. Defining cooperators as those whose behaviour does not significantly differ from total cooperation, we can identify one subject who cooperated 94% at the time. Another seven subjects fell in between, and can readily be classified as mixers. These subjects cooperated from 10 to 32% of the time, and with average rate of 20% cooperation. There was one subject, subject 7, who displayed an unusual pattern of trying to use the 10-round summary as a coordination device, and hence often cooperated significantly in early rounds in each 10-round set. This subject had an overall level of cooperation of 42%, well above the other mixers. Hence, it is unclear whether subject 7 should be classified as a mixer or a cooperator. As a result we will present the data with both classifications of this subject, beginning with the classification as a cooperator.

With this classification, we find that the average probability of cooperation by mixers is $p_m = 0.2000$, and the average cooperation that the mixers observed¹⁵ is $p_0 = 0.1958$. The combined cooperation of the cooperators and defectors is $p_n = 0.2028$, while the overall probability of cooperation is 0.2015. These numbers are all strikingly similar, and are consistent with the imperfect-information equilibrium explanation. Reclassifying subject 7 as a mixer, we find $p_m = 0.2275$, $p_0 = 0.1988$, and $p_n = 0.1667$. Again, these numbers are all close in value, and not significantly different. Similar results hold up for the total experiment in general.¹⁶ With the original classification of subjects, we find $p_m = 0.1842$, $p_0 = 0.1808$ and $p_n = 0.2007$, with overall cooperation of 0.1917. By classifying subject 7 as a mixer, we find $p_m = 0.2081$, $p_0 = 0.1870$, and $p_n = 0.1717$. Again, as predicted by the imperfect-information equilibrium, p_m and p_n are very similar.

This suggests that the behaviour in the stranger condition is consistent with an imperfect-information equilibrium in which individuals share a common prior on the probability of experiencing cooperation, p^* , of about 0.20. Two previous studies have also estimated subjects' subjective priors on cooperation. Camerer and Weigelt (1988) estimated 'homemade priors' of 0.17 that an opponent would play cooperatively, and McKelvey and Palfrey (1992) estimated the proportion of altruists to be 0.5 and 0.10. The similarity of these estimates to our own is a pleasant surprise.

The Restart

After completing the main experiment, all subjects were told that they would play against the tit-for-tat strategy for a 10-period repeated game. The striking

¹⁵ The value p_0 is determined by finding the actual level of cooperation observed by each mixer.

¹⁶ For the entire experiment, mixers range from 7.5 to 24.5% cooperation, defectors range from 95.5 to 100% defection, and two cooperators have 97 and 62.5% cooperation.

result from the restart is that all conditions perform almost identically. This is true in the individual data as well as the aggregate data. In the partners group, 8 of 14 subjects chose to the optimal strategy of cooperation until the final round. In both the strangers and the computer50s, the number was 7 of 14 subjects, while the computereros had 8 subjects choose the optimal strategy.¹⁷ The restart shows that the level of sophistication of subjects in all conditions was about the same. For instance, strangers, who had no experience with the finitely repeated play, were just as successful at exploiting the computer strategy as were the computer50s, who had experience playing tit-for-tat opponents. Hence, while subjects generally exhibit behaviour consistent with the sequential equilibrium prediction in the main experiment, the restart shows that they do not uniformly demonstrate the strategic sophistication that we ascribe to sequential equilibrium players in theory.

IV. CONCLUSION

This paper presented experiments designed to examine the sequential equilibrium reputation hypothesis in the finitely repeated prisoner's dilemma. Our results support the sequential equilibrium prediction. Subjects in a finitely repeated prisoner's dilemma were significantly more cooperative than subjects in a repeated single-shot game. Moreover, by increasing subjects' beliefs about the probability that their opponent is altruistic, we can further increase reputation building.

Several findings in the experiment suggest that, rather than simply believing that some subjects may be altruistic, many subjects actually are altruistic. Play in the repeated single-shot game is consistent with a model of warm-glow in which people get additional utility from mutual cooperation, and our results suggest that there is a stable fraction of such altruists in the population. The evolution of play in the repeated games is also consistent with the altruism hypothesis. Rather than defecting earlier in each of the series of repeated games, subjects continue to increase their waiting time until their first defection, even as the experiment nears the end.

In summary, subjects appear very willing to build reputations for altruisms. However, it seems important to the observed play of the game that some subjects actually are altruists. In contrast to the strict, purely rational, version of the reputation building hypothesis, there may be no real difference in the beliefs that an opponent is an altruist and the actual chance it is so.

University of Wisconsin

Santa Fe Institute and Carnegie Mellon University

Date of receipt of final typescript: July 1992

¹⁷ There were also other similarities across conditions. Every condition had one or two 'alternators' who began with defection and alternated getting the temptation and the sucker payoff, until round 10 when they took the mutual defection payoff. Each condition had one subject playing all-cooperate, except the computereros who had the only subject who played all-defect. All conditions, except the strangers, had one subject who cooperated until round 8, and defected for rounds 9 and 10.

APPENDIX

Subjects' Instructions for the Computer50 Condition

THE UNIVERSITY OF WISCONSIN
 Department of Economics

Subjects' Instructions

WELCOME

This experiment is a study of economic decision making. The instructions are simple. If you follow them carefully and make good decisions you may earn a considerable amount of money.

The money you earn will be paid to you, in cash, at the end of the experiment. A research foundation has provided the funds for this study.

The One-Round Decision

In this experiment you will be paired with one other player. You will be paired with this player through a computer network – at no time will your true identity be revealed to the other participants. The other player, like yourself, was recruited from an economics course at the UW.

Both you and the other player will have two possible choices. You can choose *LEFT* or you can choose *RIGHT*. If you both choose *LEFT* you will both get a payoff of 7 cents. If you both choose *RIGHT* you will both get a payoff of 4 cents. If you choose *RIGHT* but the other player chooses *LEFT*, you will get a payoff of 12 cents, but the other player will receive 0 cents. Likewise, if you choose *LEFT* but the other player chooses *RIGHT*, then you receive 0 cents and the other player receives 12 cents. These payoffs are summarized in the table below. The **bold** number in the top portion of each box is the payment received by you, the number in the bottom is the payment received by the other player:

		Payoff From Your Move	
		LEFT	RIGHT
Other Player's Payoff	LEFT	7 7	12 0
	RIGHT	0 12	4 4

When choosing your move, you *will not* know the choice of the other player. You must make your choice *without* knowing what the other player will choose. After all players in the experiment have made their choices, the computer will report to you the move chosen by the other player and your payoff from *this round* of play.

Sets of Rounds

You will play the one-round game just described in 10 one-round sets. That is, each set of play will consist of 10 one-round games. To begin a set of rounds, the computer

will randomly match you with another player in the room. *You will then play the one-round game just described with the same player for a total of 10 rounds.* That is, all 10 rounds in the set will be played with the same other player.

After the 10th round a new set will begin. The computer will randomly reassign you to play with another player: *every 10 rounds you will be randomly reassigned to a new subject.* You will **never** be assigned to play with the same person for more than 10 rounds.

Reminder: *During each 10-round set, you will be playing each one-round game with the same other player for all 10 rounds.*

At the end of each round, the computer will tell you your move in the last round, the other player's move, and your earnings from that round. At the end of each set, the computer will tell you your total earnings for the entire 10-round set.

We will play this game for a total of 20 sets of 10 rounds each. That is, there will be 20 sets, and each will have 10 decision rounds. Thus, during the course of the experiment you will play a total of 200 one-round games.

Computer Players

At the beginning of every set there is a chance that you will be randomly paired with a computer player, rather than a fellow participant in the experiment. For every 10-round set, the chance that you will be paired with a computer player is $1/2$. That is, there is a 50% chance that you will be assigned the computer player. If you are not paired with the computer, you will be matched with another person in the experiment.

Computer Moves

The computer player is always programmed to use a very simple 'copy cat' rule. The computer will start every 10-round set by choosing *LEFT*. After that the computer will make the same choice that you made on the previous round. For example, if you choose *LEFT* on round 1, the computer will choose *LEFT* on round 2. If you choose *RIGHT* on round 2, the computer will choose *RIGHT* on round 3. And so on.

Confidentiality

Your identity in the experiment will not be made known to any other participant at any time in the experiment. *Your decisions and payoffs are confidential.*

Do not discuss your choices or payoffs with any other player!

Thank you and Good Luck!

REFERENCES

- Andreoni, James (1989). 'Giving with impure altruism: applications to charity and Ricardian equivalence.' *Journal of Political Economy*, vol. 97, pp. 1447-58.
- Andreoni, James (1990). 'Impure altruism and donations to public goods: a theory of warm glow giving.' *ECONOMIC JOURNAL*, vol. 100, pp. 464-77.
- Camerer, Colin (1988). 'Gifts as economic signals and social symbols.' *American Journal of Sociology*, vol. 94, pp. S180-214.
- Camerer, Colin and Weigelt, Keith (1988). 'Experimental tests of the sequential equilibrium reputation model.' *Econometrica*, vol. 56, pp. 1-36.
- Cooper, Russell, DeJong Douglas V., Forsythe, Robert and Ross, Thomas W. (1990). 'Cooperation without reputation.' Working paper, University of Iowa.
- Dawes, Robyn M. (1980). 'Social dilemmas.' *Annual Review of Psychology*, vol. 31, pp. 169-93.
- Dawes, Robyn M. and Thaler, Richard (1988). 'Anomalies: cooperation.' *Journal of Economic Perspectives*, vol. 2, pp. 187-98.

- Kandori, Michihiro (1992). 'Social norms and community enforcement.' *Review of Economic Studies*, vol. 59, pp. 61-80.
- Kelley, Harold H. and Stahelski, Anthony J. (1970). 'Social interaction basis of cooperators' and competitors' beliefs about others.' *Journal of Personality and Social Psychology*, vol. 16, pp. 66-91.
- Kreps, David M., Milgrom, Paul, Roberts, John and Wilson, Robert (1982). 'Rational cooperation in the finitely repeated prisoners' dilemma.' *Journal of Economic Theory*, vol. 27, pp. 245-52.
- Kreps, David M. and Wilson, Robert (1982). 'Reputation and imperfect information.' *Journal of Economic Theory*, vol. 27, pp. 253-79.
- McKelvey, Richard D. and Palfrey, Thomas R. (1992). 'An experimental study of the centipede game.' *Econometrica*, vol. 60, pp. 803-36.
- Milgrom, Paul and Roberts, John (1982). 'Predation, reputation and entry deterrence.' *Journal of Economic Theory*, vol. 27, pp. 280-312.
- Palfrey, Thomas R. and Rosenthal, Howard (1988). 'Private incentives in social dilemmas: the effects of incomplete information and altruism.' *Journal of Public Economics*, vol. 35, pp. 309-32.
- Piliavin, Jane Allyn and Charng, Hong-Wen (1990). 'Altruism: a review of recent theory and research.' *Annual Review of Sociology*, vol. 16, pp. 27-65.
- Rapaport, Anatol and Chamma, Albert M. (1965). *Prisoner's Dilemma*. Ann Arbor: University of Michigan Press.
- Roth, Alvin E. (1988). 'Laboratory experimentation in economics: a methodological overview.' *ECONOMIC JOURNAL*, vol. 98, pp. 974-1031.
- Roth, Alvin E. and Murnighan, J. Keith (1978). 'Equilibrium behaviour and repeated play of the prisoner's dilemma.' *Journal of Mathematical Psychology*, vol. 17, pp. 189-98.
- Samuelson, Larry (1987). 'A note on uncertainty and cooperation in a finitely repeated prisoner's dilemma.' *International Journal of Game Theory*, vol. 16, pp. 187-95.
- Selten, Reinhard and Stoecker, Rolf (1986). 'End behaviour in sequences of finite prisoner's dilemma supergames: a learning theory approach.' *Journal of Economic Behaviour and Organization*, vol. 7, pp. 47-70.
- Stark, Oded (1985). 'On private charity and altruism.' *Public Choice*, vol. 46, pp. 325-32.