POPULATION LEARNING OF COOPERATIVE BEHAVIOR IN A THREE-PERSON CENTIPEDE GAME

Ryan O. Murphy, Amnon Rapoport and James E. Parco

ABSTRACT

We consider mixed populations ($N = 21$) of genuine (humans) and artificial (robots) agents repeatedly interacting in small groups whose composition is changed randomly from round to round. Our purpose is to study the spread of cooperative or non-cooperative behavior over time in populations playing a 3-person centipede game by manipulating the behavior of the robots (cooperative vs. non-cooperative) and their proportion in the population. Our results convey a positive message: adding a handful of cooperative robots increases the propensity of the genuine subjects to cooperate over time, whereas adding a handful of non-cooperative agents does not decrease this propensity. If there are enough hard-core cooperative subjects in the population, they not only negate the behavior of the non-cooperative robots but also induce other subjects to behave *more* cooperatively.

KEY WORDS • centipede game • cooperation • population dynamics • trust

Introduction

Consider an experimental design in which a finite population (community) of N agents is divided into k groups of n members each who participate in a non-cooperative extensive form game with complete and perfect information. Assume that the stage game is repeated in time, that on each round of play the composition of each of the k groups and the player role of each member of the group in the extensive form game are randomly determined, that the number of rounds is finite and commonly known, and that the outcome information

Rationality and Society Copyright © 2004 Sage Publications. Vol. 16(1): 89–118.
www.sagepublications.com DOI: 10.1177/104346304039876

at the end of each round is limited to the agent's own group. In particular, public information about the decisions or outcomes of all the N agents is not provided, nor is any information given about the previous history of actions of the other members of the group on any particular round. Control over preferences is achieved by paying the agents contingent on their performance.

This experimental design does not allow for reputation-building as information about an agent's past decisions is not available to other group members on following rounds. If the ratio N/n of population size to group size is 'large', or the number of rounds of play is 'small', then repeated encounters are unlikely. Accordingly, the effects of repeated interaction are minimized or completely eliminated (e.g. Crawford 1997). But if the number of rounds is 'large' or the ratio N/n is not too high, repeated encounters are highly likely. In the latter case, the decisions made by an agent on a given round may affect the decisions of the $n - 1$ members of her group on the next round, who, in turn, may affect future decisions of yet other agents in the population. Therefore, the decision of each agent could have far-reaching consequences beyond the present round that may spread in the population precisely because future encounters are likely. Maintaining a long-term horizon and planning ahead, agents may be willing to forego a portion of their payoff on a given round in order to transmit a signal that, despite the anonymity of the repeated interaction, may spread in the population (not unlike rumors or contagious diseases) and thereby change the behavior of its members over time.

Our purpose is to study the spread of either cooperative or non-cooperative behavior in a mixed population of genuine and computerized agents who repeatedly and anonymously interact with one another in small groups whose composition changes randomly from one round of play to another. In particular, we determine if and how the proportion of either 'hard-core' cooperators or 'hard-core' non-cooperators in the population affects population learning under the experimental design described above.[1]

Experimental Procedures for Studying Population Dynamics

There are basically three different experimental procedures that can be used to study how population dynamics might be influenced by the distribution of player types in the population. The first is to classify subjects in terms of some measure of their behavior *prior* to the conduct of the experiment, and then compose stratified populations of players according to their types. This prior sorting may be done with or without the explicit knowledge of the subjects. A second procedure is to sort the subjects into k groups on each round, with or without their knowledge, in terms of some measure of their behavior on previous rounds of the *same game*. For example, Gunnthorsdottir et al. (2001) conducted a public goods experiment with the Voluntary Contribution Mechanism, which is frequently used to examine cooperative behavior and free-riding in social dilemmas (e.g. Davis and Holt 1993; Ledyard 1995). On each round, they sorted their subjects in terms of their contribution decision on previous rounds into either free-riders or cooperators and placed them into different four-person groups without telling them the assignment rule. Their results show that the sorting mechanism that they used helps in keeping subjects with cooperative dispositions together and lead to significant increases, relative to a control group, in cooperators' contributions to the provision of public goods (see also Andreoni 1990; Crosson 1998; and Houser and Kurzban 2002). A third procedure inserts artificial agents – called 'robots' – into the population who are programmed to play in a pre-specified manner. For example, this procedure was used by Calegari et al. (1998) in a study designed to test the competitive equilibrium predictions of a multi-period model of audit pricing and independence. The robots may be 'sophisticated' (flexible or adaptive) in the sense that their decision on any particular round depends on the decisions or outcomes of previous rounds; or they may be 'unsophisticated' (inflexible or static) in that their decisions are fixed across iterations of the stage game.

Each of these three experimental procedures has its shortcomings. Underlying the first procedure is the assumption that the classification of agents into types in terms of some disposition or propensity is task-independent. This assumption may not always be valid as, for example, the same agent may exhibit a high degree of cooperative behavior in one game (e.g. the trust game of McCabe et al. 1996, 1998), but not in another (e.g. the Prisoner's Dilemma game). By sorting subjects into groups in terms of some index of behavior (e.g. cooperativeness) based on previous outcomes of the same game, one introduces the possibility of strategic behavior across rounds. For example, Gunthorsdottir et al. opted not to tell their subjects the assignment rule they used because of their concern that differences in their strategic behavior generated by their knowledge of the sorting rule might confound reciprocity effects. In fact, there is no assurance that their subjects might not have learned this assignment rule. The main problem with the third procedure is that it conceals from the subjects the introduction of robots into the population or the number and the strategies these robots have been programmed to play. Nevertheless, despite these shortcomings, the present study implements the third procedure with a pre-determined number of 'unsophisticated' robots because it is better suited to control the percentage of cooperative or non-cooperative player types in the population.

To study how population dynamics might be influenced by the distribution of player types, we use a version of the centipede game (see, e.g., Rosenthal (1981), McKelvey and Palfrey (1992), and Aumann (1992, 1995, 1998) for theoretical and experimental studies of the 2-person centipede game; see Rapoport et al. (2003) – hereafter RSPN – for a 3-person variant of the centipede game). Described in detail in the next section, the centipede game is ideally suited for our purpose because it allows for intermediate decisions between the two extremes of complete cooperation and complete non-cooperation. Moreover, experimental results that we review below (McKelvey and Palfrey 1992; RSPN 2003) show considerable individual differences in the propensity to cooperate in this game with a distribution of responses that covers the entire range from complete cooperation to complete non-cooperation. Our main hypothesis is that if a handful of artificial agents programmed to cooperate – called 'cooperative robots' – are inserted into a population of genuine subjects playing the centipede game, the propensity to cooperate of the genuine subjects will increase over time. Conversely, the addition of a handful of artificial agents programmed not to cooperate – called 'non-cooperative robots' – will cause this propensity to decrease. These hypotheses are consistent with the findings of the effects of sorting in public good games (e.g. Gunthorsdottir et al. 2001). Although the hypotheses may seem intuitive and straightforward, they do not address the differential effects of cooperative and non-cooperative types on the population dynamics, which are the major issue of the present article.

The plan of the rest of the article is as follows. Section 2 presents and discusses the 3-person centipede game and then reports the findings of RSPN that motivated the present study. Section 3 describes the experimental design and Section 4 reports the experimental results. Briefly, our results show that, as hypothesized, adding cooperative robots increases the disposition of the genuine subjects to play cooperatively relative to a baseline condition that only includes humans. In contrast, non-cooperative robots have no significant effect on the cooperative behavior of our subjects. Section 5 concludes with a summary and discussion of the results.

The Centipede Game

The centipede game is a finite n-person extensive form game with complete and perfect information. Perfect information means that when it is her turn to make a decision, a player knows perfectly all the decisions made by all the players at the previous decision nodes. The assumption of complete information means that everything about the game is commonly known. The two-person centipede game was first introduced by Rosenthal (1981) and later studied theoretically by Aumann (1992, 1995, 1998), Ben-Porath (1997), Binmore (1996), Ponti (2000), Stalnaker (1998), and many others (see Rapoport (2003) for a brief and non-technical review). In a variant of this game presented by Aumann (1992), there are two players called Alice and Ben and a sum of $10.50 lying on the table in front of them. Moving first, Alice can choose one of two alternatives. She can stop the game by taking $10.00 and leaving $0.50 to Ben or she can continue the game by 'passing'. If she decides to 'pass', the amount on the table is increased 10-fold (to $105.00), and it is Ben's turn to play next. He, too, can stop the game by taking $100 and leaving $5 to Alice. If he elects to continue the game by 'passing', the amount is again increased 10-fold (to $1,050.00). The game progresses with player roles being interchanged and payoff increasing 10-fold on each move following each 'passing' decision. The game terminates after three full phases ('innings') of play (i.e. a total of six moves) unless a player elects to end it sooner. In the sixth and final stage, Ben can either choose to end the game and take the $1,000,000, leaving $50,000 to Alice, or he can 'pass'. If he passes, the game terminates with zero payoffs to each of the two players.

Variants of the two-person centipede game have been investigated experimentally by McKelvey and Palfrey (1992), Fey et al. (1996), and Nagel and Tang (1998). In a recent study, RSPN extended the game from $n = 2$ to $n = 3$ players. Figure 1 displays the game tree that was studied by RSPN. The game includes three innings, each

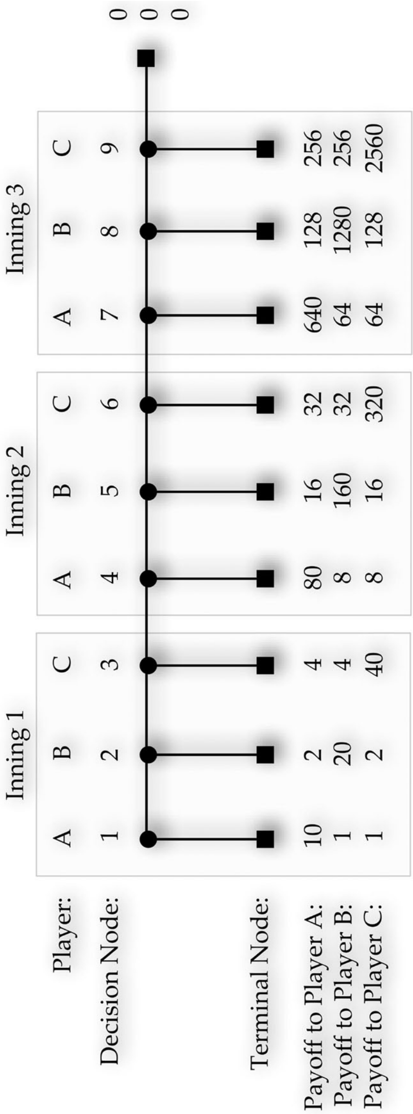

Figure 1. Three-player Centipede Game

consisting of three moves (one per player) for a total of nine moves (decision nodes). It proceeds from left to right with player 1 moving first. At each node, payoffs increase 2-fold, rather than 10-fold as in the game discussed by Aumann, with the player who chooses to stop (moving down) receiving 5/6 and each of the other two players receiving 1/12 of the group payoff. Similar to the game of Aumann, but unlike the two centipede games investigated by McKelvey and Palfrey, a decision to continue (moving right) on the ninth and final decision node yields zero payment to each of the three players.

Analysis

How should rational agents behave if this game is played once? The standard way of answering this question is by backward induction. For the game in Figure 1, the recursive process of backward induction is straightforward. Begin with the ninth and final decision node. At this node, player 3's decision to stop (and receive 2,560) dominates his decision to continue (and receive 0). Therefore, player 3 should stop. Move back to the eighth node, where it is player 2's turn to play. Assuming that player 3 will act rationally and stop at node 9, player 2's decision to stop at node 8 (and receive 1,280) dominates her decision to continue (and receive 256). Hence, a rational player 2, who assumes that player 3 is also rational, should stop at node 8. With the same reasoning applying to all nine stages, a decision to stop at each node is, therefore, prescribed. This is the unique equilibrium solution of the game. Therefore, by the logic of backward induction, an equilibrium is constructed in which player 1 will stop on node 1 with payoffs of 10, 1, and 1 for players 1, 2, and 3, respectively.

Although the reactions of people to the conclusion reached by this process of backward induction vary, most are unwilling to accept it or at least believe that it represents an approach of little practical value. Based on a theoretical analysis of the centipede game, Aumann considers it to be one of the 'disturbing counterintuitive examples to rational interactive decision making' (1992: 219). Based on an experimental investigation of two different 2-person centipede games, McKelvey and Palfrey have reached a similar conclusion.

The backward induction argument is based on the notion of counterfactuals. In applying it, as we did above, it is assumed common knowledge that each player i ($i = 1, 2, 3$) would stop at each of his three decision nodes, even when i knows with certainty that node j ($j > 1$) will never be reached. Much of the theoretical discussion has revolved around the counterfactual nature of this requirement of common knowledge of rationality. In particular, Aumann has drawn a distinction between *substantive* and *material* rationality. Substantive rationality refers to the condition where it is common knowledge that each player i would act rationally at each of her decision nodes v, even if i knows that node v will not be reached. The weaker condition of material rationality stipulates that player i will act rationally only at those of her decision nodes that are actually reached. Aumann (1998) has proved that in the particular case of the centipede game, the somewhat weaker condition of common knowledge of material rationality is sufficient to justify the backward induction process.

Experimental Evidence

However, even this latter condition would seem to be much too strong to be met in practice. The experimental evidence seems overwhelmingly to reject it. McKelvey and Palfrey conducted a 2-person 4-move centipede game experiment in which it was common knowledge that a player would be paired with the same partner only once. Their game also differed from the one in Figure 1 by assigning positive payoffs to the outcome of continuing on the final decision node.[2] The percentages of games that terminated with an exit decision on nodes 1, 2, 3, and 4 were 7.1%, 35.6%, 37.0%, and 15.3%, respectively. (In 4.9% of the cases player 2 continued on the fourth decision node.) In their second experiment that investigated behavior in a 2-person 6-move centipede game with essentially the same design that added two more moves to the 4-move game, the percentages of games that terminated with an exit decision at nodes 1, 2, 3, 4, 5, and 6 were 0.7%, 6.4%, 19.9%, 38.4%, 25.3%, and 7.8%, respectively. (In 1.4% of the cases player 2 continued on the sixth and final decision node.) Thus, only 7.1% of the cases in the 4-move game and 0.7% of the 6-move games supported equilibrium play.

Recently, RSPN conducted two experiments to test the descriptive power of the equilibrium solution in the 3-person 9-move game exhibited in Figure 1. Their experimental design – the same as the one described in Section 1 above – differed from that of McKelvey and Palfrey in several important respects. First, the number of players was increased from $n = 2$ to $n = 3$, thereby rendering the assumption of common knowledge of rationality more difficult to be met in practice. Second, as mentioned earlier, the incentive to continue on the final decision node was removed by assigning zero payments to this decision. Third, the number of iterations of the stage game was increased from 10 to 60, thereby providing more opportunity for learning. Fourth, and most importantly, the population consisted of $N = 15$ subjects (5 groups of 3 subjects each) with group membership and assignment of players to player roles (1, 2, 3) within a group determined randomly on each round. With a relatively small ratio of $N/n = 5$ and a large number of rounds of play (60), future encounters with one or both members of the group were highly likely.[3] In both Experiments 1 and 2 of RSPN, final payment was based on a randomly selected trial method. The subjects were informed that they would receive X% of their individual earnings in three randomly chosen rounds, the same for all N subjects. The value of X was set at 50 in Experiment 1 and at 100 in Experiment 2.

Experiments 1 and 2 of RSPN only differed from each other in the magnitude of the payoffs: US dollars in Experiment 1 (a high-pay condition) vs. US cents in Experiment 2 (a low-pay condition). When the stakes were low, but still of the same order of magnitude as in the 6-move game of McKelvey and Palfrey, the results were very similar, exhibiting the same pattern. Across four different sessions, each including a population of $N = 15$ subjects interacting in this manner for 60 rounds, the percentages of games terminating with a decision to stop on node j were 2.6%, 3.4%, 9.8%, 13.3%, 22.6%, 22.8%, 16.5%, 5.7%, and 3.0% for $j = 1, 2, 3, 4, 5, 6, 7, 8$, and 9, respectively.[4] Altogether, out of 1200 games (4 sessions $\times$ 5 games per round $\times$ 60 rounds per session), equilibrium play was supported in only 2.6% of the time, and in only 3.0% of the cases did the game reach the final decision node. Thus, neither full cooperation (a decision to stop on node $j = 9$) nor full non-cooperation (a decision to stop on node $j = 1$) were supported. These results render the low-pay 3-person centipede game particularly suitable for testing the hypothesis that the *distribution* of stopping decisions across decision nodes would be shifted either to the right or to the left by adding to the population of genuine subjects either cooperative or non-cooperative robots.

The observation motivating the present study was made in the high-pay Experiment 1 of RSPN. Similar to Experiment 2, four separate sessions were conducted each with a population of $N = 15$ subjects. In three of the four sessions, the percentage of subjects stopping on the first node slowly converged to 1 in support of the equilibrium solution. In each of these three sessions, play never proceeded beyond the first inning after approximately 45 rounds of play (see Figure 2 in RSPN). In contrast, the dynamics of the fourth replication exhibited a very different pattern of strong resistance for terminating the game on the first inning. Even on the last few rounds of this session, play moved beyond the third decision node in about 40% of the time. Across all rounds, whereas the final decision node (with a payoff of $2,560.00 for player 3) was reached only on 0, 1, and 2 cases (out of 300) in sessions 1, 2, and 3, respectively, it was reached on 10 cases in session 4. Analyses of individual play, in addition to uncovering considerable individual differences in all four sessions, identified a small subset of hard-core cooperators in session 4. When members of this subset happened to be placed together in the same group by the random assignment procedure, they tended to continue until stopping at one of the three nodes in the third inning. Rather than relying on chance to have sufficiently many subjects of this cooperative type in a population and on the random assignment procedure to place them together in the same group, the present study controls the size of this subset and the strategies played by the members of this subset by introducing preprogrammed robots.

Experimental Design

Subjects

The subjects were undergraduate students who volunteered to participate in a group decision-making experiment for payoff contingent on performance. In addition to receiving a non-salient payment of $5.00, some of the students also received class credit (in courses taught by none of us) for showing up to the experiment on time.[5] The subjects were randomly assigned to 1 of 4 experimental conditions, each including 3 separate sessions, for a total of 12 sessions. The four conditions differed from one another in the presence or absence of robots or the strategies they were programmed to play. The subjects had not previously been instructed on the centipede game or the process of backward induction.

Condition Baseline-0

Condition Baseline-0 included 21 genuine subjects (and zero robots) in each session. It was designed to replicate the low-pay Experiment 2 of RSPN and serve as a control group for the other three conditions.

Condition Coop-6

Condition Coop-6 (for cooperation) included 15 genuine subjects together with 6 cooperative robots. The cooperative robots were programmed to play a simple strategy, unconditional on previous decisions or outcomes, which consisted of continuing with probability 0.95 at each decision node actually reached, and stopping with probability 1.0 at the final decision node.

Condition NonCoop-3

Condition NonCoop-3 (for non-cooperation) included 18 genuine subjects and 3 non-cooperative robots. These three robots were programmed to stop with probability 0.95 at each decision node that was actually reached and stop with probability 1.0 at the final decision node.

Condition NonCoop-6

Condition NonCoop-6 was identical to Condition NonCoop-3 with the only exception that each population included 15 genuine subjects and 6 robots.

Altogether, 63, 45, 54, and 45 genuine subjects participated in Conditions Baseline-0, Coop-6, NonCoop-3, and NonCoop-6, respectively, for a total of 207 subjects.

Sessions lasted approximately 75 minutes. Individual payoffs varied considerably across subjects and conditions from a minimum of \$2.00 (excluding a \$5.00 show-up fee) to a maximum of \$74.00.

Procedure

All 12 sessions were conducted at the Economic Science laboratory (ESL) at the University of Arizona. The procedure was identical to that of RSPN. Upon arrival at the ESL, each participant drew a poker chip from a bag containing 21, 15, 18, or 15 chips for Conditions Baseline-0, Coop-6, NonCoop-3, and NonCoop-6, respectively, to determine their seat assignment. Subjects were seated in individual cubicles, each containing a PC and set of written instructions (see Appendix) that they read at their own pace. When all the subjects completed reading the instructions, questions were answered (there were very few) and the session commenced.

Subjects participated in 90 rounds of the game displayed in Figure 1. On each round, humans and robots were mixed together and then randomly assigned to groups and player role. The same assignment procedure was used in each session. On each round, the 21 subjects were divided into 7 groups of 3 players each. Each group might have included 0, 1, 2, or 3 robots. Each subject viewed the same screen of the game tree in Figure 1, the round number (1 to 90), and her player role (1, 2, or 3). If and when a player was afforded with the opportunity to make a choice, she simply clicked on either the Down or Right branch emanating from the decision node. The robots were programmed to delay their decision by approximately 5–11 seconds, with the delay time decreasing slowly with progressing rounds (this was the median delay time of genuine subjects established in a pre-test) to alleviate suspicion. When selected, the branch changed color (from black to yellow) and a 'commit' button appeared at the bottom of the screen asking the subject to confirm her decision. Subjects could privately change their decision before confirming it as many times as they wanted.

Once the subject confirmed her decision, the screens of the other two group members were immediately updated and the selected branch was identified by changing its color on all three screens. If the decision was to continue ('Right'), the next player in the sequence was prompted to make his choice in exactly the same manner while the two other (inactive) players simply viewed the updated game tree. If the decision was to stop ('Down'), all three subjects were informed that the game played in that round was over. Payoffs were read directly from the screen. Subjects were only informed of the outcome of the game in which they participated; information about the decisions or outcomes of the other population members was withheld. When members of all the seven groups completed viewing their outcome, the session advanced to the next round with subjects being randomly re-assigned to new groups and within group to new player roles.

Final payment was based on the outcomes of 6 (out of 90) rounds. It was commonly known that these rounds had been randomly selected *prior* to the start of the session and written on a sheet of paper that was sealed in an envelope. Subjects were paid their cumulative payoff in these six rounds and dismissed from the ESL one at a time. No post-experimental questionnaire was given, and no formal de-briefing took place.

Results

There are basically three major findings. First, in support of our hypothesis, adding a handful of cooperative robots to the population had the anticipated effect of increasing subjects' propensity to cooperate. Not only do the results reject the null hypothesis that the mean propensity to cooperate in Coop-6 is the same as in Baseline-0, but they also show a steady increase in the propensity to cooperate across rounds in Coop-6, in contrast to a slow decrease in this propensity across rounds in Baseline-0. Second, in contradiction to our hypothesis, adding a handful of non-cooperative robots to the population did not reduce the propensity to cooperate. In fact, the subjects in NonCoop-3 and NonCoop-6 cooperated *more readily* than the subjects in the control condition, although the hypothesis of equal mean propensity in all of these three conditions could not be statistically rejected. Third, our results show considerable individual differences in the propensity to cooperate both between and within conditions. Moreover, if there is a significantly large subset of persistent cooperative subjects, they can successfully negate the effects of the non-cooperative robots. Similarly to the results of Experiment 2 of RSPN, we observe no support for equilibrium play in the baseline condition that included no robots. Evidence supporting these three major conclusions is presented below.

Effects of the Cooperative Robots

Let r_{i1} denote the number of cases that player i reached a decision node in the first inning (nodes 1, 2, or 3). Note that $r_{i1} < 90$. For example, if on a given round player i was assigned to player role 3, and another member of her group (genuine or robot) stopped on either node 1 or 2, then player i never reached a decision node in inning 1 on that round (or, for that matter, in inning 2 or 3). Let

s_{i1} denote the number of cases that player i chose to continue when reaching a decision node in inning 1 ($s_{i1} < r_{i1}$). Finally, define $q_{i1} = s_{i1}/r_{i1}$ as the *propensity* of player i to cooperate on the first inning. q_{i1} is simply the probability of player i continuing on the *first opportunity* she is asked to choose between stopping or continuing irrespective of her player role (1, 2, or 3) in the group. In a similar manner, let q_{i2} denote the propensity of player i to cooperate in the second inning (node 4, 5, or 6), if it is reached. The probability q_{i3} of continuing in the third and final inning is defined in a similar way, except that node $j = 9$ is excluded. (Moving right on the final decision node is an irrational decision rather than an indication of cooperativeness.) Note that, necessarily, $r_{i1} > r_{i2} > r_{i3}$, and that for some players q_{im} ($m = 1, 2, 3$) may not be defined. For example, q_{i3} is not defined for player i if she never reached the third inning on any of the 90 rounds. In this way, each subject i is characterized by a vector of probabilities $\boldsymbol{q}_i = (q_{i1}, q_{i2}, q_{i3})$.

The vectors $\boldsymbol{q}_i$ were used to test the null hypothesis of equality in mean propensities to cooperate in Coop-6 and Baseline-0. To avoid the problem of missing values, subjects for whom one or more of the probabilities q_{im} were not defined were omitted from the analysis. This resulted in removing 2 of the 45 subjects in Coop-6 and 16 of the 63 subjects in Baseline-0. Hotelling T^2 test rejected the null hypothesis ($F_{1,88} = 5.01, p < 0.003$). (Pillai's Trace, Wilk's Lambda, Hotelling's Trace, and Roy's Largest Root tests all yielded the same result.) As listed in Table 1, the mean q_{im} values in Coop-6, computed across all three sessions, were 0.963, 0.774, and 0.379 for $m = 1$, 2, and 3, respectively, with standard errors of 0.016, 0.043, and 0.060. The corresponding values in Baseline-0 were 0.914, 0.580, and 0.362, with standard errors of 0.016, 0.041, and 0.058. Taken together, these results show that the players in Coop-6 had significantly higher propensities to cooperate than the players in the baseline condition that included no robots.[6]

Table 1. Mean Propensity to Cooperate by Inning

Condition	*Inning 1* ($j = 1 - 3$)	*Inning 2* ($j = 4 - 6$)	*Inning 3* ($j = 7 - 8$)
Baseline-0	0.914	0.580	0.362
NonCoop-3	0.954	0.598	0.323
NonCoop-6	0.851	0.651	0.452
Coop-6	0.963	0.774	0.379

In computing the probability of player *i* moving right, we included her decisions in both *uniform* groups that included humans only and *mixed* groups that included both humans and robots. The next analysis only considers the subject's decisions in uniform groups. Restricting the analysis to uniform groups considerably reduces the number of observations. With no robots added, all seven groups in the baseline condition are uniform. However, the expected number of uniform groups on any given round is 2.551 [$7 \times (15/21)^3$], 4.408 [$7 \times (18/21)^3$], and 2.551 [$7 \times (15/21)^3$] in conditions Coop-6, NonCoop-3, and NonCoop-6, respectively.

Table 2 presents the estimated conditional probabilities of stopping computed across subjects participating in uniform groups only. The conditional probabilities are presented by decision node and condition. Each condition includes two rows; the top row displays the frequencies and the bottom row the estimated conditional probabilities. For example, across the three sessions and 90 rounds of play, there were 699 cases in Coop-6 where node $j = 1$ was reached (the expected value is $2.551 \times 3 \times 90 = 688.8$). In 27 of these cases player 1 stopped, resulting in an estimated probability 0.039 of stopping. In 30 of the 672 cases ($699 - 27$) that node $j = 2$ was reached in the same condition, a decision to stop by player 2 was recorded, resulting in an estimated conditional probability 0.045 of stopping. In only 5 of the 64 cases where the ninth decision node was reached in Coop-6, player 3 chose to continue (and obtain zero payoff) rather than stop and take the $25.60. We attribute these very few and obviously irrational decisions (also observed in the other conditions) to either error or a desire to verify the payoff function.

Table 2 shows no support for equilibrium play in the baseline condition. The probability of stopping on the first node is 0.030. Using a similar design with 60 rather than 90 rounds and exactly the same payoffs, RSPN reported a similar probability of stopping on node 1 of 0.026 in their Experiment 2. These two probabilities do not differ significantly from each other. The probability of stopping on the first node in Coop-6 (0.039) is of the same order of magnitude. Comparison of Baseline-0 and Coop-6 shows that the difference between these two conditions is mostly due to the probability of stopping at nodes $j = 3$ to $j = 6$. The probabilities of stopping at nodes 3, 4, 5, or 6 in Baseline-0 are about 4 times, 3 times, 3 times, and 2 times higher than in Coop-6. In the cases where the third inning is reached, the conditional probabilities of stopping at

Table 2. Estimated Conditional Probability of Stopping by Decision Node and Experimental Condition

		Decision Node								
Condition		*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
Baseline-0	*Freq.*	1890	1833	1719	1282	757	317	121	56	29
	Prob.	**0.030**	**0.062**	**0.254**	**0.408**	**0.584**	**0.618**	**0.537**	**0.482**	**0.966**
NonCoop-3	*Freq.*	1197	1159	1097	958	685	388	171	64	25
	Prob.	**0.032**	**0.053**	**0.127**	**0.285**	**0.434**	**0.559**	**0.626**	**0.641**	**0.960**
NonCoop-6	*Freq.*	699	602	533	438	354	253	160	83	51
	Prob.	**0.139**	**0.115**	**0.178**	**0.192**	**0.285**	**0.368**	**0.481**	**0.386**	**0.961**
Coop-6	*Freq.*	699	672	642	602	521	419	290	141	64
	Prob.	**0.039**	**0.045**	**0.062**	**0.135**	**0.196**	**0.308**	**0.514**	**0.546**	**0.922**

Results are based on uniform groups that only include three genuine subjects. The expected number of uniform groups per round is $7 \times (21/21)^3 = 7$, $7 \times (15/21)^3 = 2.551$, $7 \times (18/21)^3 = 4.408$, and $7 \times (15/21)^3 = 2.551$ in conditions Baseline-0, Coop-6, NonCoop-3, and NonCoop-6, respectively. Multiplied by 90, the expected number of such groups per session is 630, 230, 397, and 230.

either node 7 or 8 are about the same in both conditions. For conditions Baseline-0 and Coop-6 separately, we calculated the non-conditional probability of stopping for each of the eight decision nodes. Then, using the individual subject (rather than the individual decision) as the unit of analysis, we tested the null hypothesis of equality of the cumulative distributions of stopping time by the Kolmogorov-Smirnov test. The null hypothesis was soundly rejected ($p < 0.01$).

Next we move from a static to a dynamic analysis of the subjects' decisions using only data from uniform triads of genuine subjects. Across all 21 subjects in Baseline-0, we computed the median node of stopping, one for each round. Then we calculated the running mean of these medians in steps of 10 (the mean of the median node of exit on rounds 1–10, 2–11, . . . , 81–90). The same computation was performed for each of the other three conditions. Figure 2 illustrates the results. The horizontal axis shows the round number from 1 to 90. The numbers on the right-hand vertical axis display the node number j, whereas the numbers on the left-hand vertical axis present the corresponding payoffs to the players who stopped. Note that decision nodes map directly into payoffs for the first player to stop, but that these payoffs increase exponentially in j, not linearly. To illustrate the effect of the subject's decision to stop on her payoff, we have opted to display the running means on a linear payoff scale.

Figure 2 shows that the median node of stopping in Baseline-0 slowly decreased across rounds from about $j = 5$ in the first 10 rounds to a value slightly smaller than 4 on rounds 81–90. As a result, the payoff gained by the subject who stopped was cut, on average, by half, decreasing from about \$1.60 to less than \$0.80. In sharp contrast, the median node of stopping in Coop-6 increased across the first 50 rounds and then stabilized between decision nodes 6 and 7. The corresponding payoff to the subject who stopped increased more than 2-fold from about \$2.30 on rounds 1–10 to about \$5.30 on rounds 81–90. On the last 30 rounds or so, subjects in Coop-6 were earning, on average, 6–7 times more than subjects in Baseline-0. In contrast to the baseline condition, where play moved slowly in the direction of the equilibrium solution, the presence of the cooperative robots resulted in a slow increase in the propensity to cooperate. However, even with 6 cooperative robots added to 15 humans, the effect was not sufficiently strong to induce the

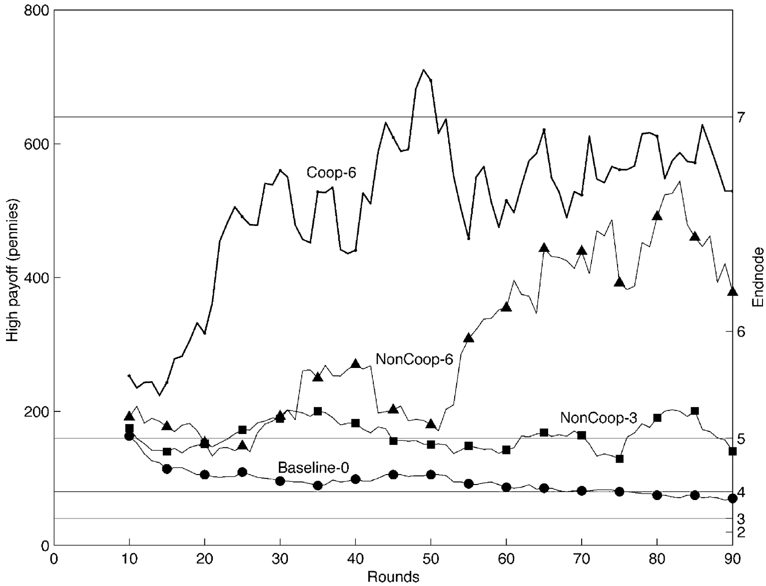

Figure 2. Moving Average of the Median Earnings for Player who Stopped the Game (Genuine Players Only)

humans to cooperate all the way to the final decision node and thereby maximize the group payoff.

Effects of the Non-Cooperative Robots

The vectors of the propensities of cooperation $\boldsymbol{q}_i$ were computed in a similar way for each of the humans in conditions NonCoop-3 and NonCoop-6. They are presented in Table 1. After deleting subjects from whom not all three propensities were defined (8 in NonCoop-3 and 10 in NonCoop-6), the multivariate Hotelling T^2 test was performed to test the null hypotheses of equality in mean propensities to cooperate between any two of the three conditions Baseline-0, NonCoop-3, and NonCoop-6. None of the tests yielded significant results ($p > 0.05$). In particular, we find no evidence that, in comparison to the baseline condition, non-cooperative play by the robots dissuaded cooperative play by the genuine subjects. In fact, subjects in the two non-cooperative conditions tended to cooperate more often than subjects in the baseline condition, although the difference was not significant. The mean q_{im} values in condition NonCoop-3 were 0.954, 0.598, and 0.323 for $m = 1, 2$, and 3, respectively. The corresponding mean propensities in condition NonCoop-6 were 0.851, 0.651, and 0.452.

Shifting the focus to the dynamics of play, Figure 2 shows that the running means of the median stopping node for the two non-cooperative conditions fall between the baseline and the cooperative conditions. In contrast to the baseline condition, there is no evidence for a systematic change in the median propensity to cooperate across rounds in NonCoop-3. The median stopping node in this condition is 5 with a corresponding payoff of \$1.60 to the player who was the first to stop. More surprising is the behavior of the subjects in NonCoop-6 that included twice as many non-cooperative robots as NonCoop-3. Until about round 50, the median stopping node is about the same in both non-cooperative conditions. But after round 50 there is a sharp increase in cooperative behavior in the uniform groups in NonCoop-6. As shown below, this result is mostly due to a small subset of hard-core cooperators who happened to be included in one of the three sessions in this condition. When three members of this subset happened to be assigned to the same group, almost invariably they reached the third inning (68.2% of the time) and often stopped on the final decision node (32.7% of the time).

Individual Differences

Although the subjects in each session were afforded approximately the same number of choices to either stop or continue, there were considerable individual differences in the frequency of stopping decisions. Some subjects were the first to stop on almost every round, demonstrating a strong preference to pre-empt their two group members, whereas others never exercised the stopping option. Table 3 presents the frequency distributions of number of stopping decisions by condition. The individual numbers of stopping decisions across the 90 rounds in the session are grouped in classes of size 10 (except for the first class). For example, 2 subjects in Coop-6 made no stopping decisions, 8 subjects made 1–10 stopping decisions, 9 subjects each had a frequency of stopping decisions between 11 and 20, and so on. Subjects with a relatively large number of stopping decisions did not necessarily stop on the first inning. For example, one of the 45 subjects in Coop-6 stopped a

Table 3. Aggregate Frequency Distributions of the Decisions to Stop (Endnodes 1–8) by Genuine Players across Each Experimental Condition

	Frequency Class										
	0	*1–10*	*11–20*	*21–30*	*31–40*	*41–50*	*51–60*	*61–70*	*71–80*	*> 80*	*Sum*
Baseline-0	3	6	15	9	10	9	9	2	0	0	63
NonCoop-3	6	9	15	1	12	4	6	1	0	0	54
NonCoop-6	3	17	7	9	6	3	0	0	0	0	45
Coop-6	2	8	9	0	9	5	6	5	0	1	45

total of 23 times in the first inning, 45 in the second inning, and 13 in the third inning for a total of 81 stopping decisions. Another subject in the same condition with a total number of 70 stopping decisions never stopped on the first inning. These differences in frequency of stopping decisions resulted in considerable individual differences in payoff.

Individual differences in the 'home grown' propensities to cooperate might have caused considerable differences among the three sessions in the same condition. A sufficiently large number of persistent cooperators might have resisted the effects of the non-cooperative subjects, genuine or not. This, indeed, was the case in one of the sessions in NonCoop-6. To show this, we tallied for each subject the number of her stopping decisions over 90 rounds. Figure 3 exhibits the stopping decisions of all genuine subjects in Sessions 1, 2, and 3 of NonCoop-6. Data for the individual subjects were organized in ascending order. We focus on this condition because of the difference in behavior between Session 2 and the other two sessions.

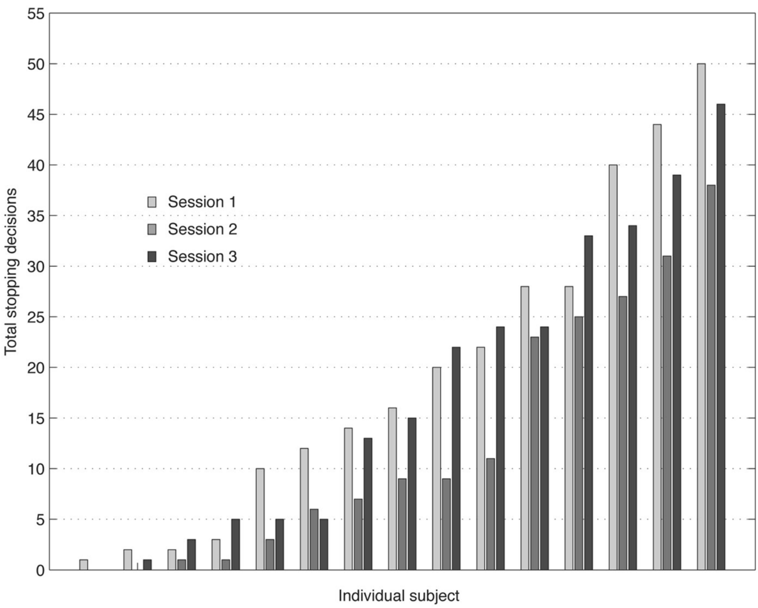

Figure 3 Frequency Distribution of Individual Stopping Decisions in Condition NonCoop-6

Of a total of 296 stopping decisions made by the humans in Session 1, 134 (45%) were made by three subjects only. A total of 186 (62.8%) of all the stopping decisions occurred on the first inning, 87 (29.4%) on the second, and 23 (7.8%) on the third. Stopping decisions in innings 1, 2, and 3 were distributed more or less evenly across the 90 rounds, showing no evidence for learning. Of a total of 272 stopping decisions made by the humans in Session 3, 153 (56.3%) were made by only 4 subjects. The percentages of stopping decisions in the first, second, and third innings were very close to those observed in Session 1: 54.4%, 37.5%, and 8.1%, respectively. Similar to Session 1, stopping decisions in innings 1, 2, and 3 were distributed evenly across the 90 rounds, showing again no evidence for learning.

Although interacting with six non-cooperative robots as in Sessions 1 and 3, the subjects in Session 2 of condition NonCoop-6 behaved quite differently. Altogether, 234 stopping decisions were recorded in this session. Only three subjects accounted for 43.2% of all stopping decisions. In contrast to Sessions 1 and 3, stopping decisions on the first inning were infrequent. Altogether, only 24 (10.3%) stopping decisions occurred in the first inning, and with two exceptions none occurred after rounds 51. Ninety-six stopping decisions (41.0%) were made on the second inning and 114 (48.7%) on the third. For 8 of the 15 subjects both $q_{i1} > 0.95$ and $q_{i2} > 0.95$. In contrast, the number of such persistent cooperators in each of Sessions 1 and 3 was only two. When three of these eight subjects were grouped together by the random assignment procedure, with few exceptions play progressed to the third inning. The cooperative behavior of these subjects affected the behavior of the remaining seven subjects: the number of stopping decisions on the first inning dropped to zero (with two exceptions made by the same subject toward the end of the session), and the number of stopping decisions made on the third inning steadily increased. It is these few hard-core cooperators in Session 2 who are mostly responsible for the differences between the three sessions of condition NonCoop-6 (Figure 3) and the difference between conditions NonCoop-6 and NonCoop-3 (Figure 2).

Discussion and Conclusions

The present study was designed to answer the question to what extent the behavior of a population of agents who interact repeatedly over time can be influenced, if at all, by the distribution of types in the population under the particular experimental design that we have devised. A previous study by RSPN, who explored the emergence of stable patterns of behavior over time using the same 3-person centipede game and a similar experimental design, suggested that populations of genuine subjects may differ from one another in terms of the prior distribution of 'types' of subjects defined in terms of some prior ('home grown') disposition to cooperate. Due to the methodological difficulties of reliably sorting subjects into 'types' either prior to the experiment (e.g. through questionnaires or observed behavior in a previous related game) or sorting them on each round of the game in terms of some index of the subjects' behavior on previous rounds of the same game, we opted to use another design that implants artificial agents ('robots') into populations of genuine agents. The main advantage of this methodology is the control over the proportion of robots of a particular type in the population.

The findings of RSPN articulated the emergence of non-cooperative behavior, in the direction of equilibrium play, when the stakes were unusually high and basically no learning (and no support for equilibrium play) when the stakes were of the same order of magnitude as in previous 2-person centipede experiments conducted by McKelvey and Palfrey. Therefore, when designing the present experiment we anticipated stronger effects of the non-cooperative than the cooperative robots.[7] That was the reason for including only a single condition (Coop-6) with six cooperative robots compared to two conditions, one with three and the other with six non-cooperative robots. The results provided only partial support to our hypothesis. In support of the hypothesis, the addition of cooperative robots did increase the propensity of the genuine subjects to cooperate, namely, to move further right on the game tree in Figure 1 (and increase mean individual payoff). In contrast, the addition of non-cooperative robots in both NonCoop-3 and NonCoop-6 did not change the propensity of the genuine subjects to cooperate in comparison to the baseline condition. Not only did the genuine subjects resist the non-cooperative behavior of the robots, who stopped the game when it was their turn to play with probability 0.95, but when a sufficiently large number of hard-core cooperators happened to be included in the same session, they overturned the effects of the robots and, in fact, led to an increase in the propensity to cooperate of the other subjects. The good news is that subjects seem to ignore non-cooperative behavior of a sizeable proportion of the population that they may attribute to stupidity, greed, or maliciousness and, rather, persist in their attempt to increase, though not necessarily maximize, group payoff.

Caution should be exercised in any attempt to generalize these results beyond the present centipede game, as they may depend on the structure of the game, the ratio n/N of group size to population size, the percentage of artificial agents, and the size of stakes (see Parco et al. 2002). In addition, the present study raises a major methodological issue that warrants brief discussion. As shown by the comparison of the three sessions in NonCoop-6 (Figure 3), it is clear that controlling the percentage of artificial agents in the population is not sufficient by itself for controlling the distribution of 'types', as different subjects may approach the game with different dispositions to cooperate. If it can be established experimentally that the propensity to cooperate generalizes across a class of 'similar' games designed to explore cooperative behavior, then a better methodology would combine prior sorting of subjects into types in terms of their prior disposition to cooperate measured in a different game and the control over the distribution of types in the population by adding pre-programmed artificial agents.

NOTES

We acknowledge support of this research by a grant from the Hong Kong Research Grants Council (Project CA98/99.BM01) to the Hong Kong University of Science and Technology.

1. There is a widely based belief among pious Jews, based on a statement to this effect in the Talmud, that there exist 36 men, called 'Tzadikim' or the Holy Just Men, who somehow sustain the world. These individuals are scattered throughout the world, live in self-imposed confinement, have no acquaintance with one another, are not recognized as holy just men, and are extremely modest and upright. Somehow, through their righteous behavior, interaction with others, and their very being, they avert disasters and sustain the world.
2. The initial payoffs on stage 1 in their game were \$0.40 and \$0.10 for players 1 and 2, respectively. Payoffs increased 2-fold as in Figure 1. If player 2 stopped on the fourth and final stage, the payoffs received by players 1 and 2 were \$0.80 and \$3.20, respectively. If she continued, the respective payoffs were \$6.40 and \$1.60. This payoff scheme, in which continuing on the final decision node maximizes the total group payoff, increases the incentive to continue. Yet, the backward induction applies as in the game in Figure 1.

3. On the average, each player had the same subject as a member of her group during 12.86 rounds, and the same two subjects as members of her group during 0.99 rounds.
4. The 2-person 6-move centipede game of McKelvey and Palfrey and the 3-person 9-move centipede game of RSPN each included three innings. The percentages of games terminating in innings 1, 2, and 3 in the RSPN study were 15.8%, 58.7%, and 25.2%, respectively, compared to 7.1%, 58.3%, and 33.1% in the study of McKelvey and Palfrey.
5. Although, once they arrived at the laboratory, they were told that participating in the experiment was not a condition of receiving class credit. Only those students wishing to participate for financial incentives contingent on performance were asked to stay. Those few students who elected not to stay were given credit, paid $5.00 and dismissed.
6. Players who were excluded from the analysis had, in general, lower propensities to cooperate. For example, a subject who always stopped on the first or second inning never reached the third inning. Therefore, this subject was excluded from the analysis. With only 2 subjects excluded in Coop-6 compared to 23 in Baseline-0, there is a selection effect that biases the results in favor of the null hypothesis. If all the subjects in each of the two conditions were included (e.g. by limiting the comparison to the first two propensities), the difference between the mean propensities would increase.
7. Theoretical results reported by Kandori (1992), who made similar assumptions about repeated interaction in finite communities with imperfect observability, indirectly support this hypothesis. Most relevant to our study is the analysis of a community divided into pairs who play the Prisoner's Dilemma game. Briefly, and similar to our study, Kandori assumes that in each encounter each player observes only the history of action profiles in the stage games that she had played, that the player knows nothing about the identity of players or what has gone in the rest of the community, and that direct communication among the players is non-existent. He shows that under these assumptions a sequential equilibrium strategy consists of a contagious strategy in which defection spreads 'like an epidemic and cooperation in the whole community breaks down' (p. 69).

REFERENCES

Andreoni, J. 1990. 'Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving? *Economic Journal* 100: 464–77.

Aumann, R. J. 1992. 'Irrationality in Game Theory.' In *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, ed. P. Dasgupta, D. Gale, O. Hart, and E. Maskin, pp. 214–27. Cambridge, MA: MIT Press.

Aumann, R. J. 1995. 'Backward Induction and Common Knowledge of Rationality.' *Games and Economic Behavior* 8: 6–19.

Aumann, R. J. 1998. 'On the Centipede Game,' *Games and Economic Behavior* 23: 97–105.

Ben-Porath, E. 1997. 'Rationality, Nash Equilibrium and Backwards Induction in Perfect-information Games.' *Review of Economic Studies* 64: 23–46.

Binmore, K. 1996. 'A Note on Backward Induction.' *Games and Economic Behavior* 17: 135–7.

Calegari, J. M., J. W. Schatzberg and G. R. Sevcik. 1998. 'Experimental Evidence of Differential Auditor Pricing and Reporting Strategies.' *The Accounting Review* 73: 255–75.

Crawford, V. P. 1997. 'Theory and Experiment in the Analysis of Strategic Interaction.' In *Advances in Economics and Econometrics: Theory and Applications*, ed. D. M. Kreps and K. F. Wallis, pp. 206–42. Cambridge: Cambridge University Press.

Croson, R. 1998. 'Theories of Altruism and Reciprocity: Evidence from Linear Public Goods Games.' The Wharton School, University of Pennsylvania, unpublished manuscript.

Davis, D. D. and C. A. Holt. 1993. *Experimental Economics*. Princeton, NJ: Princeton University Press.

Fey, M., R. D. McKelvey and T. R. Palfrey. 1996. 'An Experimental Study of Constant-Sum Centipede Games.' *International Journal of Game Theory* 25: 269–87.

Gunnthorsdottir, A., D. Houser, K. McCabe and H. Ameden. 2001. Disposition, History and Contributions in Public Goods Experiments. Department of Management and Policy, University of Arizona, unpublished manuscript.

Houser, D. and R. Kurzban. 2002. 'Revisiting Kindness and Confusion in Public Good Games.' *American Economic Review* 92: 1062–9.

Kandori, M. 1992. 'Social Norms and Community Enforcement.' *Review of Economic Studies* 59: 63–80.

Ledyard, J. O. 1995. 'Public Goods: A Survey of Experimental Research.' In *The Handbook of Experimental Economics*, eds J. Kagel and A. E. Roth, pp. 111–94. Princeton, NJ: Princeton University Press.

McCabe, K. A., S. J. Rassenti and V. L. Smith. 1996. 'Reciprocity, Trust, and Payoff Privacy in Extensive Form Games.' *Games and Economic Behavior* 24: 10–24.

McCabe, K. A., S. J. Rassenti and V. L. Smith. 1998. 'Game Theory and Reciprocity in Some Extensive Form Experimental Games.' *Proceedings of the National Academy of Sciences* 93: 13421–8.

McKelvey, R. D. and T. R. Palfrey. 1992. 'An Experimental Study of the Centipede Game.' *Econometrica* 60: 803–36.

Nagel, R., and F. F. Tang. 1998. 'Experimental Results on the Centipede Game in Normal Form: An Investigation of Learning.' *Journal of Mathematical Psychology* 42: 356–84.

Parco, J. E., A. Rapoport and W. E. Stein. 2002. 'Effects of Financial Incentives on the Breakdown of Mutual Trust.' *Psychological Science* 13: 292–7.

Ponti, G. 2000. 'Cycles of Learning in the Centipede Game.' *Games and Economic Behavior* 30: 115–41.

Rapoport, A. 2003. 'Centipede Games.' *Encyclopedia of Cognitive Science*. London: Nature Publishing Group.

Rapoport, A., W. E. Stein, J. E. Parco and T. E. Nicholas. 2003. 'Equilibrium Play and Adaptive Learning in a Three-person Centipede Game.' *Games and Economic Behavior* 43: 239–65.

Rosenthal, R. W. 1981. 'Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox.' *Journal of Economic Theory* 25: 92–100.

Stalnaker, R. 1998. 'Belief Revision in Games: Forward and Backward Induction.' *Mathematical Social Sciences* 36: 31–56.

RYAN O. MURPHY was educated in psychology at Wake Forest University (BA), experimental psychology and statistics at Western Washington University (MS), and is currently finishing his PhD at the Eller Business School at the University of Arizona. His research interests are in individual and interactive decision-making, experimental economics, and cognitive psychology.
ADDRESS: University of Arizona, Department of Management and Policy, 405 McClelland Hall, Tucson, AZ 85721, USA [email: rmurphy@eller.arizona.edu].

AMNON RAPOPORT was educated in psychology and philosophy at the Hebrew University of Jerusalem (BA), and quantitative psychology at the University of North Carolina at Chapel Hill (PhD). He has held research, teaching, and administrative positions at the Hebrew University of Jerusalem, Haifa University, and the University of North Carolina at Chapel Hill. He is a Karl Eller Professor of Management and Policy at the University of Arizona. His research interests have been in experimental economics, cognitive and social psychology, political science, management, and marketing. His writings include six books and about 170 articles on individual and interactive decision-making.

ADDRESS: University of Arizona, Department of Management and Policy, 405 McClelland Hall, Tucson, AZ 85721, USA [email: amnon@u.arizona.edu].

JAMES E. PARCO was educated in economics at the United States Air Force Academy (BS), earned an MBA at William and Mary, and completed his PhD at the Eller Business School at the University of Arizona. Parco is actively serving his country as a Major in the US Air Force and is presently stationed at the US Embassy in Tel Aviv, Israel. He is scheduled to join the faculty at the Air Force Academy in the Fall of 2003. His research interests are in bargaining, individual and interactive decision-making, and experimental economics.

ADDRESS: United States Air Force Academy, Department of Management, Colorado Springs, CO 80840, USA [email: james.parco@usafa.af.mil].

Appendix

Three-player Group Decision-making Experiment

Instructions

This experiment has been designed in order to study decision-making behavior in small groups. The instructions for this experiment are simple. If you follow them carefully and make good decisions, you may earn a considerable amount of money. The participants may earn different amounts of money in this experiment because each participant's earnings are based partly on his/her decisions and partly on the decisions of the other group members. *The money you earn will be paid to you, in cash, at the end of the experiment*. Therefore, it is important that you do your best. A research foundation has contributed the money to study group decision-making behavior.

In the event that you have any questions after reading these instructions, please raise your hand and the supervisor will come to answer them.

Description of the Task

During each trial of this experiment, you will be participating in a game that requires three players to take turns in sequence. When it is his/her turn to play, each participant (player) has to choose between two decisions:

- Move right
- Move down

If any player chooses to move down, then the game will end.

If a player chooses to move right, then the next player in the sequence will be faced with the same choice – move right or move down. If he/she is the last in the sequence, the game will end no matter what decision he/she makes.

There are a large number of players in the room who take part in this experiment. At the beginning of each trial, the computer will divide all the players in the room into separate groups of three players each. Thus, on each trial, the computer will randomly match you with two other players and randomly assign each group member the roles of Player 1, Player 2, and Player 3. Player 1 will be asked to make the first decision in every trial (followed in turn by Player 2 and Player 3). Each player will have a maximum of three decisions to make during each trial. A trial will end when either one of the three group members moves down or Player 3 moves right on his/her third turn.

Below is a pictorial representation of the game. The circles with the player numbers identify which player makes a decision (either down or right), given that the game has progressed to that circle. The arrows pointing right and down represent the two decisions. The terminal rectangles contain the payoff information. All the payoffs are in pennies (100 pennies = $1.00 US). All trials end at one of these terminal rectangles. The top number in each rectangle identifies the payoff for Player 1. The middle number is the payoff for Player 2. And the bottom number is the payoff for Player 3.

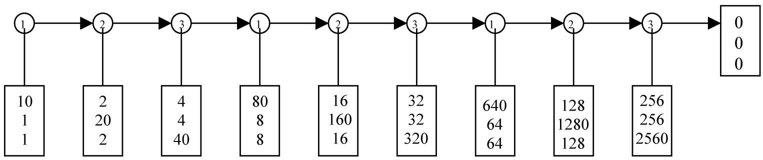

The diagram describes the choice that each player has to make and the resulting payoffs for all three players.

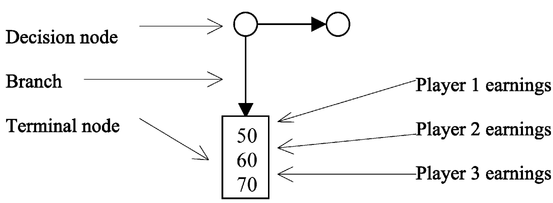

When a trial starts, each member of your group will be presented with a similar picture of the game on his/her computer monitor. Each trial will start with Player 1 at the farthest left decision node. When it is your turn, the computer will ask you to make your decision, and will identify the two branches available for selection. To make your decision, simply move the mouse pointer to one of the two branches connecting to either (1) the decision node of the next player (move right); or (2) the terminal node (move down). When you click on the branch, the computer will highlight it.

You will then be able to choose either branch until you confirm your decision. Once you are satisfied with your decision, click on the 'Commit' button at the bottom of the screen to confirm your choice.

Please take time now to study the game and its possible payoffs. You will observe that when a player moves down, his payoff exceeds the payoffs of each of the other two members 10-fold. You will also observe that the payoffs tend to increase in magnitude as the game moves to the right.

Procedure

You will participate in 90 trials, all having the same payoffs. Because communication with the other group members is only conducted via the computer and the assignment of group members is made randomly, you will not know the identity of the other two players in your group, nor will either of them know your identity. Any other form of communication during the entire experiment is strictly forbidden.

Each trial follows the same sequence. First, the computer will randomly match you with two other players and randomly assign the roles of Player 1, Player 2, and Player 3. Thus, you will be playing against different players on different trials. Next, Player 1 will be asked to make a decision. Each player will then move in the pre-determined sequence until the trial ends. When it ends, the computer will determine your payoff for the trial. During the experiment, you will be able to review all of your past decisions and payoff history. To view past results, simply click on the 'Previous Trial Results'. After you complete reviewing your past decisions and results of the trial, you will move to the next trial, where you will be randomly re-matched with two other players.

Payment at the End of the Session

At the end of the session, the computer will randomly select 6 out of the 90 trials and compute your payoff as the cumulative sum of these 6 trials.

Please look up to indicate to the supervisor that you have completed reading the instructions. The supervisor will start the experiment in just a few minutes.