

Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games

By PEDRO DAL BÓ*

While there is an extensive literature on the theory of infinitely repeated games, empirical evidence on how “the shadow of the future” affects behavior is scarce and inconclusive. I simulate infinitely repeated prisoner’s dilemma games in the lab with a random continuation rule. The experimental design represents an improvement over the existing literature by including sessions with finite repeated games as controls and a large number of players per session (which allows for learning without contagion effects). I find that the shadow of the future matters not only by significantly reducing opportunistic behavior, but also because its impact closely follows theoretical predictions. (JEL C72, C73, C91, C92)

The tension between private incentives that encourage opportunistic behavior and the common good that comes from cooperation is a central feature of human interaction. The main contribution of Game Theory to the study of this tension and its remedies is to recognize that repeated interaction may enable punishment and reward schemes that prevent or limit opportunistic behavior and support cooperation. When there is always a future, as in infinitely repeated games, the credible threat of future retaliation casts “the shadow of the future” in every decision and can overcome opportunistic behavior and support cooperation, thereby solving the tension between private incentives and the common good.

However, the experimental evidence on infinitely repeated games is scarce and in most cases inconclusive or presents methodological problems. In this paper, I report a series of experiments that overcome the shortcomings of previous experiments on infinitely repeated games. I find that the possibility of future interaction modifies players’ behavior, resulting in fewer opportunistic actions and supporting cooperation, closely following theoretical predictions.

Infinitely repeated prisoner’s dilemma games are simulated in the experiment by having a random continuation rule. The experimental design represents an improvement over the existing literature by, among other things, including sessions with finitely repeated games as controls and a large number of players per session (which allows for learning without contagion effects).

I find strong evidence that the higher the probability of continuation, the higher the levels of cooperation. While in the one-shot prisoner’s dilemma games studied here, the cooperation rate is 9 percent, for a probability of continuation of $\frac{3}{4}$, it is 38 percent. The effect of the shadow of the future on the levels of cooperation is greater than previous studies have shown.

But the finding that increases in the probability of continuation result in increases in cooperation is not necessarily evidence in support of the theory of infinitely repeated games. It could be the case that subjects cooperate more the higher the number of rounds in which they will interact; even when there is a final round, punishments are not credible and the future casts no shadow. Then, an increase in the probability of continuation would result in more cooperation, not because of the shadow of the future, but just because of the increase in the expected number of rounds. An innovation of this paper is that I compare the results from infinitely repeated games with the results from finitely repeated games to test whether cooperation depends on the shadow of the future, as theory predicts, or

* Department of Economics, Brown University, Providence, RI 02912 (e-mail: pdalbo@brown.edu). I am grateful to David Levine for invaluable guidance, discussions, and the financial support for this experiment. The support of the people at CASSEL was crucial for this research. I am also grateful to Anna Aizer, Ernesto Dal Bó, Alvin Roth, three anonymous referees, and participants at several seminars for very useful comments.

merely on the length of the games. The lengths of the finitely repeated games were chosen to coincide with the expected lengths of the infinitely repeated ones. I find that the level of cooperation in the final round of the finitely repeated games is similar to the level of cooperation in one-shot games. In addition, these levels of cooperation are lower than those observed in infinitely repeated games, providing evidence that subjects cooperate less when there is no future. This seems to be understood by the subjects at the beginning of the game, resulting in greater levels of cooperation in the first round of infinitely repeated games than in the first round of finitely repeated games of the same expected length. That is, when the expected number of future rounds is the same in both finitely and infinitely repeated games, cooperation is greater under the latter, as theory predicts.

The findings that cooperation increases with the probability of continuation and that infinitely repeated games result in higher levels of cooperation than finitely repeated ones of the same expected length suggests that self-enforcing reward and punishment schemes that limit opportunistic behavior are important in practice as well as in theory.

But the theory of infinitely repeated games provides more than general comparative statics results. It provides precise predictions regarding the set of equilibrium outcomes. I use the fact that small differences in the payoff matrix may result in large differences in the set of equilibrium outcomes to study how closely the behavior of the subjects matches the theoretical predictions. I used two different payoff matrices in the experiment with the peculiarity that, for a probability of continuation of $\frac{1}{2}$, cooperation for both players is an equilibrium in one but not in the other. I find that the percentage of outcomes in which both subjects cooperate is almost 19 percent when it is an equilibrium, while it is less than 3 percent when it is not. These experimental results show that behavior closely, although not perfectly, follows the theoretical predictions that are dependent on the payoff details of the stage game. This provides further support for the theory of infinitely repeated games.

Section I summarizes previous experimental research on the topic. Section II describes the experimental design, and Section III describes

the theoretical predictions that provide the testable hypothesis. Section IV presents the results of the experiment, and the last section concludes.

I. Related Literature

The experimental literature on cooperation generally falls into one of two categories: (a) one-shot and finitely repeated public good and prisoner's dilemma games, and (b) infinitely repeated games. In the first, cooperation can only be considered an anomaly from a theoretical point of view. As such, finding positive levels of cooperation, economists have often concluded that preferences that differ from financial incentives are an important determinant of human behavior (see James Andreoni, 1995; Ernst Fehr and Simon Gächter, 2000; and David K. Levine, 1998, among others). But in most situations, human interaction cannot be represented as a one-shot game nor characterized by a well-defined last stage of interaction. Therefore, it is crucial to study cooperation in situations in which there is always the possibility of future interactions. The possibility of future interactions allows for credible retaliations against opportunistic behavior and casts "the shadow of the future" in every decision. The theory of infinitely repeated games studies cooperation under the shadow of the future and provides a more realistic representation of everyday interactions.¹ Experimental evidence on infinitely repeated games is scarce, however, and in most cases inconclusive or presents some methodological shortcomings.

Alvin E. Roth and J. Keith Murnighan (1978) and Murnighan and Roth (1983) present results of experiments for infinitely repeated prisoner's dilemmas with different continuation probabilities. They find, on average, more cooperation in treatments in which cooperation is an equilibrium outcome than in treatments in which it is not. In addition, Roth and Murnighan (1978) find that the higher the probability of continuation, the greater the number of players who cooperated in the first round of the game (see Table 1). In contrast, Murnighan and Roth

¹ In this interpretation of infinitely repeated games, the discount factor actually incorporates the players' belief regarding the possibility of future interactions.

TABLE 1—PERCENTAGE OF COOPERATION IN THE FIRST ROUND BY PROBABILITY OF CONTINUATION

	Probability of continuation		
	0.105	0.5	0.895
Roth and Murnighan [26]*	19	29.75	36.36
Murnighan and Roth [22] [†]	17.83	37.48	29.07

* Over 121 subjects.

[†] Over 252 subjects.

(1983) present results for experiments with 12 different variations of the prisoner's dilemma in which higher probabilities of continuation did not result in more cooperation in the first round (Table 1).

In addition to presenting contradictory evidence (and offering little hope that opportunistic behavior can be limited by increases on the shadow of the future²), these two papers display methodological problems, some of which are discussed in Roth (1995). In both experiments, subjects played against the experimenter instead of playing against each other. The experimenter followed either the tit-for-tat or the Grim strategy, which can unnecessarily influence the behavior of the subjects based on their observations of past behavior of their partner/experimenter.³ In addition, in both experiments, subjects were not paid proportionately to the "points" they earned during the experiments, so that in fact the subjects were playing a constant-sum game.

Another experiment that employed a random continuation rule is by Robert M. Feinberg and Thomas A. Husted (1993). They combine a fixed continuation probability with different discount factors⁴ to study the effect of repetition on the levels of cooperation in a prisoner's dilemma disguised as a duopoly game. They find that the levels of cooperation increase as the discount factor increases. Nevertheless, that

increase is small and far from the increase needed to fully exploit the possible benefits from cooperation, even when the experiment and instructions were purposely designed to facilitate cooperation. These results are weakened because the payments made to the subjects were quite low and the basic payoffs were not the same in all treatments.⁵

Thomas R. Palfrey and Howard Rosenthal (1994) compare the rate of contribution to a public good under incomplete information regarding the contribution cost when players meet only once with the case when they meet repeatedly with a probability of continuation of 0.9.⁶ They find that repetition leads to more cooperation than one-shot games, but this increase is small (the percentage of contribution goes from 29 to 40). They conclude, "This contrast between our one-shot and repeated play results is not encouraging news for those who might wish to interpret as gospel the oft-spoken suggestion that repeated play with discount rates close to one leads to more cooperative behavior. True enough it does—but not by much."⁷ As the authors suggest later, the power of repeated play may be more evident in a simpler environment.

Finally, three recent papers also present evidence on repeated games with a random continuation rule. Jim Engle-Warnick and Robert L. Slonim (2001) develop a methodology to infer repeated game strategies from observed actions. They use this methodology to study repeated trust games with a probability of continuation of 0.8, and find that trigger strategies are used.⁸ John Duffy and Jack Ochs (2004) find high levels of cooperation (55 percent) in repeated prisoner's dilemma games with a probability of continuation of 0.9, in contrast with low levels of cooperation (6 percent) in treatments in which subjects are randomly rematched after each round. Masaki Aoyagi and Guillaume R. Fréchette (2004) study infinitely repeated prisoner's dilemma games under imperfect public monitoring and find that, as

² Roth (1995) summarizes the results in Roth and Murnighan (1978) by stating: "So the results remain equivocal" (p. 27).

³ The tit-for-tat strategy consists of cooperating in the first round, and thereafter imitating the previous action of the other player. The Grim strategy consists of cooperating until the first defection (other's or own), and defecting thereafter.

⁴ Different discount factors are obtained by different percentages of reduction of the payoffs after every round.

⁵ Another experiment that used a random continuation rule to study repeated oligopoly games is Charles A. Holt (1985). Since this experiment was designed to test for the theory of consistent-conjectures, the results do not provide information regarding cooperation.

⁶ There were at least 20 rounds, after which the probability of continuation was 0.9.

⁷ Palfrey and Rosenthal (1994), p. 548.

⁸ A trigger strategy is a strategy in which a deviation triggers a punishment for a given number of periods.

theory predicts, cooperation decreases with the level of noise of the public signal.

One drawback of the experiments with a random continuation rule is that it is not clear that, as the probability of continuation increases, any increase of cooperation that we may witness can be attributed to an increase in the importance of the future as theory predicts. It could be the case that subjects cooperate more the higher the number of rounds in which they will interact; even when there is a final round, punishments are not credible and the future casts no shadow. Then, an increase in the probability of continuation would result in more cooperation, not because of the shadow of the future, but just because of the increase in the expected number of rounds.⁹ Thus, it is necessary to create controls that will allow one to separate the two effects.

There are also experiments with a finite number of repetitions known to the experimenter but unknown to the subjects. Therefore, in each round, the subjects may assign a positive probability of continuation. For example, Lawrence E. Fouraker and Sidney Siegel (1963) study oligopoly games with an unknown number of rounds and find some cooperation in Cournot duopoly markets, but not in triopoly markets. A more recent paper in this category is Jamie Brown-Kruse et al. (1994), which presents an experiment on repeated price competition in an oligopoly market with fixed capacity constraints. While they observe prices above competitive levels, those prices are far below the monopoly price. In addition, in the treatments in which collusion is more easily supported, the prices are lower. All experiments with a fixed number of rounds unknown to subjects raise the problem that the experimenter cannot control for the players' beliefs with respect to the continuation of the game.¹⁰

Previous experimental results do not provide

much support for the theory of infinitely repeated games, nor for the use of self-enforcing reward and punishment schemes to overcome opportunistic behavior. But given the shortcomings in the design of some experiments (i.e., no real interaction among subjects, random matching protocols that allow for contagion effects,¹¹ final earnings that are not proportional to the payoffs during the game, low earnings, fixed number of rounds unknown to the subjects, and lack of control sessions), and the complicated environment of others (i.e., environments with incomplete information), previous experimental evidence is insufficient to assess the degree to which the theory of infinitely repeated games is supported empirically. This paper presents results from an experiment that was designed to overcome the shortcomings mentioned above and shows not only that the shadow of the future matters, but that its effect is significant and that it closely, but not perfectly, follows the theoretical predictions.

The results presented here for infinitely repeated games contrast sharply with the results for finitely repeated games, presented here and in previous studies, in which large deviations from theoretical predictions are observed.¹² This difference may arise from the fact that people very rarely interact in situations in which there is a well-defined (and commonly known) final stage of interaction. Very early in our lives, we learn that today's behavior may affect the behavior of people we will interact with in the future. Therefore, it is reasonable to expect that subjects will behave in a manner that is consistent with theory in familiar environments, such as the infinitely repeated game, while large deviations from the theoretical predictions are observed in finitely repeated games. In finitely repeated games, if a small fraction of the subjects have trouble adapting their everyday strategies to the experimental environment, this may

⁹ One reason for this is the existence of reputation effects (see David M. Kreps et al., 1992). There is experimental evidence showing that subjects cooperate more in finitely repeated prisoner's dilemma games than in one-shot ones (see Andreoni and John H. Miller, 1993, and Russel Cooper et al., 1996, among others).

¹⁰ This type of design also adds a source of incomplete information, since subjects do not know the beliefs of the other subjects.

¹¹ Note that, given Michihiro Kandori's (1992) contagious equilibrium, random matching is not enough to isolate the different games.

¹² See for example Reinhard Selten and Rolf Stoecker (1986), Andreoni and Miller (1993), Cooper et al. (1996), and Yoella Bereby-Meyer and Roth (2002), all of which present results for ten-round, finitely repeated prisoner's dilemma games. These studies find high levels of cooperation in the first rounds, with cooperation decreasing toward the final rounds of the finitely repeated games.

TABLE 2—STAGE GAME PAYOFFS IN POINTS

PD1				PD2			
Blue player				Blue player			
Red player		C	D	Red player		C	D
Red player	C	65, 65	10, 100	Red player	C	75, 75	10, 100
	D	100, 10	35, 35		D	100, 10	45, 45

result in the well-documented deviations from the perfect information equilibrium predictions.

This idea is supported by analysis of the effect of experience on behavior in the present experiment. In finitely repeated games, subjects start with cooperation rates above 20 percent, but they learn to defect, moving toward equilibrium behavior, as they gain experience. In contrast, subjects learn to cooperate in some of the infinitely repeated game treatments.

II. Experimental Design

This experiment was designed to offer a reliable test of the theory of infinitely repeated games. I use simple stage games: prisoner’s dilemma games. The subjects interacted anonymously with each other through computer terminals. The pairing of subjects was done so that there was no possibility of contagion effects among the different repeated games. I controlled for the subjects’ belief regarding the possibility of future interaction by having a commonly known probability of continuation. The subjects’ final earnings were proportional to the points earned during the experiment, plus a show-up fee. The exchange rate between points and dollars ensured that subjects had significant incentives to try to increase their earnings.

In addition, the experimental design incorporates three important new elements. First, in addition to the random continuation rule sessions, I ran sessions with fixed finite horizon games. The lengths of the fixed finite horizon sessions were chosen to coincide with the expected length of those with a random continuation rule. Therefore, the experimental design allows one to compare the results from infinitely repeated and finitely repeated prisoner’s dilemma games to test whether cooperation depends on “the shadow of the future,” as theory predicts, or merely on the length of the games.

Second, I considered two different prisoner’s dilemma games that resulted in different sets of equilibrium outcomes for some discount factors. In this way, I can study how closely the experimental results follow theoretical predictions that depend on details of the payoff matrices.

Third, a large number of players participated in each session, resulting in a large number of interactions for each treatment and allowing for learning without contagion effects.

Next, I describe the main characteristics of the experiment in greater detail.

Stage Game Payoffs.—I consider two different stage game payoff functions, denoted PD1 and PD2 (see Table 2).¹³

The sets of equilibrium outcomes for the infinitely repeated version of these games are described in the next section.

Public Randomization Device.—According to what is usually assumed in theory, and to allow subjects to coordinate actions and rotate through different outcomes, a public randomization device was provided to the subjects. Every ten seconds a random number between 1 and 1,000 was displayed on a screen at the front of the room. Subjects were told that they could use this number to select one of the actions if they wanted.¹⁴

Subjects’ Total Earnings.—All payoffs in the game were in points. At the end of each session, the points earned by each subject

¹³ For neutrality, in the experiments the actions were called U and D for Red subjects and L and R for Blue subjects. For convenience of the reader, I use here the usual names C and D.

¹⁴ From conversations with subjects and observation during sessions, it is clear that the subjects did not pay attention to the public randomization device.

were converted into dollars at the exchange rate of 200 points = \$1 and paid privately in cash. Subjects were also paid a \$5 show-up fee. In this way, subjects' real earnings in dollars are proportional (up to a constant) to the points obtained during the experiment. In addition, these amounts seem significant enough to influence subjects' behavior. For example, in game PD1, if Blue plays D, the incentive for Red to play D relative to C is 25 points, or 12.5 cents. This amount constitutes a substantial hourly differential, given the number of repetitions in this experiment.

Matching Procedure.—A rotation matching scheme was used to avoid potential interaction and contagion effects between the different repeated games. In each session, subjects were divided into two groups: Red and Blue. In each match, every Red subject was paired with a Blue subject. Subjects were not paired with each other in more than one match. In addition, the pairing was done in such a way that the decisions one subject made in one match could not affect, in any way, the decision of subjects he or she would meet in the future. These features were explained to the subjects.

Given that each subject was matched only once with each subject of the other color, the total number of matches in a session is $N/2$, where N is the number of subjects in a session. Given that there are three treatments per session, in each treatment, there are $N/6$ matches. The large capacity of the experimental lab CASSEL allowed me to run experiments with up to 60 subjects, providing up to ten matches per treatment per subject.

Infinitely Repeated Games.—In half of the sessions, a random continuation rule was used to induce infinitely repeated games. This was done by having one of the subjects—who had been randomly selected as the monitor—publicly roll a four-sided die after each round. The randomization generates an infinitely repeated game, given that there is always the possibility of interacting in future rounds with the same subject.

The probability of continuation δ , of which three different values were considered, is the principal treatment variable in these sessions. One treatment corresponds to the one-shot game (i.e., $\delta = 0$) and the other two are positive

probabilities of continuation (i.e., $\delta = 1/2$ and $3/4$). This treatment variable allows me to control for the subjects' beliefs regarding the probability of continuation.¹⁵ I call these sessions "Dice" sessions.

Finitely Repeated Games.—In the other half of the sessions, subjects played fixed, finite horizon games. I considered three treatments with different length or horizon H : one-shot games ($H = 1$), two-rounds repeated games ($H = 2$) and four-rounds repeated games ($H = 4$). I call these sessions "Finite" sessions. The number of rounds was common knowledge among the subjects. Note that the number of rounds for these treatments corresponds to the expected number of rounds in the random continuation rule treatments.¹⁶ Thus, in the first round, the expected number of rounds to be played is the same in the finitely and infinitely repeated games. This allows one to compare the results of the first round of finitely and infinitely repeated games.¹⁷

¹⁵ It could be argued that the subjects understand that the experiment cannot go on forever and will end at some point. Therefore, the subjects' belief in the possibility of future interactions may depend not only on the roll of the die. The subjects' real discount factor may have two components: one component determined by the roll of the die, and another subjective component which incorporates subjects' belief regarding the experimenter ending the experiment. (Given that subjects were paid at the end of the experiment and that there is a very short span of time between rounds, I disregard the temporal preference component of the discount factor.) It is important to note that if the subjective component is not very sensitive to changes in the random continuation rule, increases in the probability of continuation must result in increases in subjects' expectation of future interaction. Thus, by changing δ , I affect the subjects' belief on the possibility of future interactions. In their experiments, Murnighan and Roth (1983) elicited their subjects' beliefs about continuation probabilities. They found that subjects' estimates that there would be at least two more rounds increased strongly with the probability of continuation.

¹⁶ In infinitely repeated games with a continuation probability of δ , the expected number of rounds is equal to $1/(1 - \delta)$. Therefore, the expected number of rounds in the random continuation session will be 1, 2, and 4 for δ equal to 0, $1/2$, and $3/4$, respectively.

¹⁷ For this comparison to be meaningful, it is crucial that subjects do not expect longer games under Dice than is appropriate. An additional experiment showed that subjects do not overestimate the length of supergames under a random continuation rule. Interestingly, when asked to estimate the length of a supergame under a random continuation rule, 92.7 percent of the subjects provided an estimate equal to or lower than the correct number. These data support the use of

Order of Treatments.—To control for spill-over effects from one treatment to another, two sessions were run for each kind of continuation rule (Dice and Finite) and payoff matrix (PD1 and PD2), changing the order of the treatments. For example, for PD1 and Dice, I run one session with the order ($\delta = 0, \delta = \frac{1}{2}, \delta = \frac{3}{4}$) and another with the inverse order ($\delta = \frac{3}{4}, \delta = \frac{1}{2}, \delta = 0$). I call the first kind of session “Normal” and the second kind “UD” (upside-down).

Sessions.—Given the two-stage games (PD1 and PD2), the different continuation rules (Dice and Finite), the different treatments ($\delta = 0, \frac{1}{2}, \frac{3}{4}$, and $H = 1, 2, 4$), and the change in the order of the treatments (Normal and UD), this experiment consists of eight sessions with three treatments each. Each treatment, or part, consists of one unpaid practice match and $\frac{N}{6}$ paid matches.¹⁸ Each match consists of as many rounds as the continuation rule indicates. Different groups of subjects participated in each session.

III. Theoretical Predictions

If we assume that the payoffs in Table 2 are the actual total payoffs the subjects obtain from the game and that this is common knowledge, that is, if we abstract from problems of interdependent utilities, altruism, taste for cooperation, and reputation effects, the set of subgame perfect equilibria can be calculated using the results in Dale O. Stahl II (1991). The outcomes that can be supported in equilibrium for the different discount factors used in the experiment—and therefore the outcomes that according to theory we should observe—are presented in Table 3.¹⁹

finitely repeated games as controls for the infinitely repeated games. Higher cooperation levels in the latter cannot be explained by subjects misunderstanding the expected number of rounds under a random continuation rule.

¹⁸ The practice match was played after describing the characteristics of the treatment to the subjects. The practice match was a single match that shared all the same characteristics of the treatment, with the only difference that subjects were not paid for their actions. In each session there were three practice matches, one for each treatment.

¹⁹ I focus on equilibrium actions and outcomes instead of equilibrium strategies, since in the experimental data we will not observe the latter. For example, a pair of subjects who cooperate in every round of a repeated game is consistent with both playing the “always cooperate” strategy and the “Grim” strategy.

TABLE 3—EQUILIBRIUM OUTCOMES

	PD1	PD2
0	DD	DD
$\frac{1}{2}$	DD, CD, DC	DD, CC
$\frac{3}{4}$	DD, CD, DC, CC	DD, CD, DC, CC

New equilibria appear as the discount factor increases, allowing the subjects higher levels of cooperation and payoffs, in principle. We can think that some subjects will make the most of this opportunity to cooperate, regardless of the fact that DD remains an equilibrium outcome for high discount factors. In this way, the greater the discount factor, the greater is the percentage of cooperative actions C. For example, if under PD1 all subjects make the most of the opportunity to cooperate, the percentage of cooperative actions goes from 0 percent for $\delta = 0$, to 50 percent for $\delta = \frac{1}{2}$ and 100 percent for $\delta = \frac{3}{4}$. Therefore, I have the following testable hypothesis:

HYPOTHESIS 1: The larger δ , the higher the levels of cooperation.

It is important to note that, for this hypothesis, it is not necessary to assume that the subjects’ only payoffs from the stage game are those in Table 2. With different payoffs, the predictions presented in Table 3 may not be appropriate, but Hypothesis 1 still holds. Dilip Abreu et al. (1990) show that the set of equilibrium payoffs (and consequently the set of outcomes) that can be observed in an infinitely repeated game with a public randomization device (even with imperfect monitoring) cannot decrease when the discount factor increases. Then, for any stage game in which DD is the only Nash equilibrium, increases in the discount factor result in increases in the levels of cooperation if some of the subjects make the most of the opportunity to enforce cooperation.²⁰

In finitely repeated games, the theoretical prediction under perfect information is that no cooperation is possible. Then, under perfect information, we should expect that infinitely repeated prisoner’s dilemma games result in

²⁰ This monotonicity result may not hold without a public randomization device. See George J. Mailath et al. (2002).

higher levels of cooperation than finitely repeated ones. But, as mentioned before, the levels of cooperation in a finitely repeated game may be positive, given reputation effects. To my knowledge, there is no general theoretical result that allows one to compare the set of equilibrium outcomes between finitely and infinitely repeated games for any kind of incomplete information. But it is easy to see that some kinds of incomplete information that allow for positive levels of cooperation in finitely repeated games (i.e., some players always play the Grim strategy) do not hamper cooperation in infinitely repeated games. Thus, under both perfect and incomplete information environments, one can expect more cooperation in infinitely repeated games. In particular, we wish to know whether this is the case in the first round, where the expected number of future rounds of the infinitely and finitely repeated games coincide. Therefore, I have the following testable hypothesis:

HYPOTHESIS 2: *The level of cooperation is higher in the first round of infinitely repeated games ($\delta = 1/2$ and $\delta = 3/4$) than in the first round of finitely repeated games of the same expected length ($H = 2$ and $H = 4$, respectively).*

Finally, from Table 3 we see that the set of equilibrium outcomes is different for PD1 and PD2 when $\delta = 1/2$. Under that discount factor, CC can be observed in equilibrium for PD2 but not for PD1, while CD and DC can be observed in equilibrium for PD1 but not for PD2. Therefore, I have the following testable hypotheses:

HYPOTHESIS 3: *For $\delta = 1/2$, PD2 results in more outcomes CC than PD1.*

HYPOTHESIS 4: *For $\delta = 1/2$, PD1 results in more outcomes CD and DC than PD2.*

The first two hypotheses are quite general in the sense that they do not depend on specific details of the payoff matrices and are robust to perturbations of the stage games. In contrast, the last two hypotheses are quite specific, in the sense that they are closely based on the specified payoff matrices. In this way, the last two hypotheses allow me to test how sensitive behavior is to small payoff differences that re-

sult in large differences in the theoretical predictions.

IV. Experimental Results

The experimental sessions were run between November 2001 and April 2002, with an average length of one hour (excluding the time spent paying subjects). Excluding the subjects selected to be monitors, 390 subjects participated in the experiment, an average of 48.75 subjects per session with a maximum of 60 and a minimum of 30. The subjects were UCLA undergraduates recruited through advertisements on university Web pages and signs posted on campus. Of the subjects, 22.31 percent indicated that they were majoring in economics. The subjects performed a total of 22,482 actions and earned an average of \$18.94, with a maximum of \$25.85 and a minimum of \$12. In the treatments $\delta = 1/2$ and $\delta = 3/4$, the average number of rounds per match was 1.91 and 3.73, respectively. Some descriptive statistics are in Table 4.

Even when subjects participate in a practice match before the paid matches of each treatment, we should expect to see some learning regarding the treatment characteristics and other subjects' behavior during the first matches of each treatment.²¹ Table 5 shows the percentage of cooperation by treatment (row) and by match (column), aggregating over all rounds and all sessions. Match one denotes the first match in the treatment, match two the second match in the treatment, and so on. Each cell shows the percentage of subjects cooperating in that treatment and match, aggregating over all rounds and sessions. As one can see, there is clear learning regarding the difficulties of cooperation in the $\delta = 0$ treatment of the Dice sessions, and in all the treatments of the Finite sessions (that is, in all the treatments with fixed horizons). For example, in the $\delta = 0$ treatment, cooperation decreases from more than 26 percent in the first match to 5 percent in the last match. In contrast, in the $\delta = 1/2$ and $\delta = 3/4$ treatments, there is no such pattern.

To avoid considering actions that are taken with little understanding of the treatment, and

²¹ Recall that every treatment consists of up to ten matches per subject, with each match having as many rounds as the continuation rule of that treatment indicates, e.g., for treatment $H = 2$, there are two rounds per match.

TABLE 4—SESSIONS DESCRIPTIVE DATA

		PD1		PD2	
		Dice	Finite	Dice	Finite
Normal	Date	11/6/01	11/13/01	2/7/02	4/18/02
	Time*	2:30–3:28	4:45–5:31	1:45–2:56	5:15–6:25
	Subjects	42	30	54	48
	Any econ [†]	23.81%	23.33%	12.96%	18.75%
	Ave earning	17.09	13.03	19.91	19.36
	Max earning	19.40	15.23	22.18	21.88
	Min earning	13.48	12.05	15.98	15.48
UD	Date	11/29/01	11/20/01	4/9/02	4/15/02
	Time*	5:10–6:05	5:10–6:05	4:45–5:53	4:45–5:54
	Subjects	42	54	60	60
	Any econ [†]	16.67%	12.96%	31.67%	35%
	Ave earning	14.37	17.77	23.09	22.11
	Max earning	16.23	21.55	25.85	25.10
	Min earning	12.18	12	19.93	17.15

* Starting scheduled time and actual final time.
† Percentage of all economics majors in the session.

TABLE 5—PERCENTAGE OF COOPERATION BY MATCH AND TREATMENT*

		Match									
		1	2	3	4	5	6	7	8	9	10
Dice	$\delta = 0$	26.26	18.18	10.61	11.62	12.63	12.63	5.56	5.26	5.26	5
	$\delta = \frac{1}{2}$	28.36	27.12	34.58	35.53	21.60	19.08	29.84	35.96	28.16	50
	$\delta = \frac{3}{4}$	40.44	28.57	27.78	32.92	46.51	33.09	44.05	53.51	42.26	45.83
Finite	$H = 1$	26.56	18.23	16.67	17.19	11.98	8.02	6.79	10.49	6.14	6.67
	$H = 2$	19.79	15.89	14.84	9.64	11.46	10.80	12.04	10.19	6.58	6.67
	$H = 4$	31.64	30.34	30.47	25.52	25.13	23.77	16.36	19.75	14.91	20.83

* All rounds and sessions.

since most of the learning seems to happen in the first matches, I focus on matches four through ten in most of the analysis of the experimental results.^{22,23}

A. Does Cooperation Increase with the Shadow of the Future?

My first objective is to study how changes in the probability of future interaction affect the levels of cooperation. The experimental results

show the greater the shadow of the future, the higher are the levels of cooperation. Considering the aggregate results for the Dice sessions (matches four through ten, and all rounds), we find that cooperation is just above 9 percent for the one-shot treatment, while it is above 27 percent and 37 percent for $\delta = \frac{1}{2}$ and $\delta = \frac{3}{4}$, respectively. These differences are statistically significant with p -values of less than 0.001.²⁴ Therefore, the experimental results support Hypothesis 1: *the larger δ , the higher the levels of cooperation.*²⁵

²² The elimination of the first three matches does not affect the qualitative results having an effect only on the precise magnitudes.

²³ In addition, for simplicity of presentation, I aggregate the results from Normal and UD sessions for most of the section. The differences between these two types of sessions are discussed at the end of the section.

²⁴ All p -values are calculated taking into account the lack of independence in the decisions of each subject (cluster robust standard errors).

²⁵ Similarly, the percentage of mutually cooperative outcomes (CC) also increases with the discount factor. Considering matches four through ten and all rounds, we find

TABLE 6—PERCENTAGE OF COOPERATION BY ROUND AND TREATMENT*

		Round											
		1	2	3	4	5	6	7	8	9	10	11	12
Dice	$\delta = 0$	9.17											
	$\delta = \frac{1}{2}$	30.93	26.10	19.87	12.50	12.96							
	$\delta = \frac{3}{4}$	46.20	40.76	38.76	34.58	33.04	27.27	24.75	26.28	29.17	26.04	32.29	31.25
Finite	$H = 1$	10.34											
	$H = 2$	13.31	6.90										
	$H = 4$	34.58	21.55	18.97	10.63								

* All sessions, matches four through ten.

The effect of the shadow of the future appears to be large: the percentage of cooperation for $\delta = \frac{3}{4}$ is almost four times greater than for the one-shot treatment. The magnitude of this difference is greater than previously found. For example, in the public good experiments with incomplete information of Palfrey and Rosenthal (1994), the percentage of contributions increases only from 29 percent to 40 percent when the treatment changes from one-shot games to a random continuation rule with $\delta = \frac{1}{10}$. This is also the case if we compare the results of this experiment with the results from Roth and Murnighan (1978) and Murnighan and Roth (1983). In those experiments the percentage of cooperation less than doubles when the probability of continuation increases from $\frac{4}{38}$ to $\frac{34}{38}$. The magnitude of the results presented here supports the idea that infinitely repeated interaction can significantly reduce opportunistic behavior.

B. Infinitely Repeated Games versus Finitely Repeated Games

Our second objective is to compare the levels of cooperation in the Dice and Finite sessions. In Table 6, I present the percentage of cooperation by treatment (row) and by round (column), aggregating over all sessions and matches four through ten. As this table shows, the percentage of cooperation is similar for the one-shot treatments ($\delta = 0$ and $H = 1$) in both types of sessions (p -value = 0.56), showing that there

are no significant differences in the “kind” of people who participated in each session. In the fourth round of the $\delta = \frac{3}{4}$ treatment, the level of cooperation is significantly greater than in the fourth (and last) round of the $H = 4$ treatment (34.58 percent against 10.63 percent, with p -value of less than 0.001). The level of cooperation in the final round of the $H = 4$ treatment is similar to the level of cooperation in one-shot games. Therefore, the absence of a future affects subjects’ behavior in the final round of finitely repeated games: they cooperate less when there is no future. This seems to be understood by the subjects at the beginning of the game, resulting in less cooperation in the first round of a finitely repeated game than in the first round of an infinitely repeated game (34.58 percent against 46.20 percent, with p -value of 0.005). Similar reasoning applies to the comparison of the behavior for $\delta = \frac{1}{2}$ and $H = 2$. That is, when the expected number of future rounds is the same in both finitely and infinitely repeated games, cooperation is greater under the latter, as theory predicts. Therefore, the experimental results support Hypothesis 2: *the level of cooperation is higher in the first round of infinitely repeated games than in the first round of finitely repeated games of the same expected length.*

In addition, note that, for every round, it is the case that the percentage of cooperation in infinitely repeated games ($\delta = \frac{1}{2}$ and $\delta = \frac{3}{4}$) is greater than in finitely repeated games of the same expected length ($H = 2$ and $H = 4$), with p -values of less than 0.01. This difference increases with the round number as finitely repeated games get closer to their final round but infinitely repeated games do not (since there is no final round). Moreover, while cooperation

that mutually cooperative outcomes are just above 1 percent for the one-shot treatment, while they are around 14 percent and 24 percent for $\delta = \frac{1}{2}$ and $\delta = \frac{3}{4}$, respectively.

TABLE 7—DISTRIBUTION OF OUTCOMES BY STAGE GAME AND TREATMENT*

	$\delta = 0$		$\delta = \frac{1}{2}$		$\delta = \frac{3}{4}$	
	PD1	PD2	PD1	PD2	PD1	PD2
CC	2.98	0.27	3.17	18.83	20.68	25.64
CD & DC	20.83	13.98	28.57	25.50	30.34	26.03
DD	76.19	85.75	68.25	55.67	48.98	48.33

* Matches four through ten, and all rounds.

decreases with the round number in finitely repeated games, that is not always the case in infinitely repeated games. In fact, under $\delta = \frac{3}{4}$, cooperation actually increases after the seventh round.²⁶

The difference in behavior in finitely and infinitely repeated games is also clear if we look at individual data instead of aggregated cooperation rates. Unfortunately, as discussed before, looking at individual data does not enable one to identify the strategies used by the subjects. For example, observing a pair of subjects who cooperate in all the rounds of a repeated game is consistent with both of them playing the “always cooperate” strategy or the Grim strategy. Additionally, there is a large number of strategies that can yield cooperation in infinitely repeated games, which can vary in the level of complexity and which cannot be estimated from the observed actions. Notwithstanding this, it is still possible and interesting to compare the proportions of actions that are consistent with the Grim and “always defect” strategies in finitely and infinitely repeated games. In the $\delta = \frac{3}{4}$ treatment, 35.65 percent of the individual actions are consistent with Grim and 39.17 percent are consistent with “always defect.” Instead, in the $H = 4$ treatment, only 21.20 percent of the actions are consistent with Grim and 49.54 percent consistent with “always de-

fect.”²⁷ Therefore, it is clear that a larger proportion of actions is consistent with strategies that support cooperation in infinitely repeated games than in finitely repeated ones.

C. Do Payoff Details Matter?

Our third objective is to compare the outcomes under PD1 and PD2 when $\delta = \frac{1}{2}$ (see Table 7). Remember that CC is not an equilibrium outcome under PD1, but is under PD2. Consistent with this, the percentage of outcomes in which both players cooperate (CC) is significantly lower under the payoff matrix PD1 than under PD2 when $\delta = \frac{1}{2}$ (3.17 percent against 18.83 percent with a p -value of less than 0.001). Note that this is not the case when $\delta = 0$, suggesting that the difference in the percentage of CC when $\delta = \frac{1}{2}$ cannot be attributed to differences in the subjects that participated in the sessions under PD1 and PD2. Thus, the experimental results support Hypothesis 3: *for $\delta = \frac{1}{2}$, the payoffs PD2 result in more outcomes CC than PD1.*

Finally, remember that CD and DC are equilibrium outcomes under PD1 but not under PD2. As theory predicts, the percentage of outcomes in which only one subject cooperates (CD and DC) is greater under PD1 than under PD2 (28.57 percent against 25.50 percent), but this difference is not statistically significant (p -value of 0.505).

For evidence in support of Hypothesis 4, I also look at the levels of alternation between

²⁶ The decrease in cooperation in the first rounds of infinitely repeated games is explained in part by the triggering of punishments by noncooperative actions. The increase in cooperation in later rounds is explained by the end of some of those punishments. The nonmonotonicity of $\delta = \frac{1}{2}$ is apparent when one considers all matches, instead of matches four through ten, since then we have data for up to eight rounds. For all matches with $\delta = \frac{1}{2}$, cooperation goes from 32 percent in round 1 to 17 percent in round 4 and 28 percent in round 8.

²⁷ Given the different theoretical predictions for each prisoner’s dilemma under $\delta = \frac{1}{2}$, it is not sensible to present aggregated information regarding individual actions in this treatment. In addition, given the high proportion of matches that had only one round, the inference of strategies from the observed actions does not provide significantly more information than the aggregate levels of cooperation.

CD and DC in PD1 and PD2.²⁸ For this, I adapt the K index from Anatol Rapoport et al. (1976): $K = (CD) + (DC) - |(CD) - (DC)|$ where (CD) and (DC) denote the percentage of outcomes CD and DC in a given repeated game. For a given repeated game, this index is 100 if CD and DC are the only observed outcomes and occur with the same frequency. Thus, K is an index of alternation in a repeated game. For repeated games with more than one round, and considering only even numbers of rounds, the average K index is 12.86 for PD1 and 7.26 for PD2. This indicates a higher level of alternation under PD1 than under PD2, but this difference is not striking in magnitude. Thus, while there is some weak support of Hypothesis 4, there is a much stronger support of Hypothesis 3.

There are two plausible reasons for the stronger support of Hypothesis 3. First, the lack of communication among subjects may hamper the coordination on alternating asymmetric outcomes, even when there is a public randomization device available. Second, while it can be easily seen that CC can be supported under PD2 but not under PD1, higher levels of strategic sophistication are needed to realize that CD and DC can be equilibrium outcomes in PD1 but not in PD2. While one may not expect subjects to perform the formal calculations, it is reasonable to expect that they will grasp the first feature of the environment better than the second. Subjects' strategic sophistication seems similar to that of game theorists, who also tend to grasp the first feature more easily than the second one!

Let me end the discussion of the results with two comments. First, it is interesting to note that the effect of the shadow of the future increases as subjects gain experience. Subjects clearly learn to cooperate under $\delta = 3/4$, while there is no clear trend under $\delta = 1/2$. In addition, they learn to defect in the finite horizon treatments (one-shot and finitely repeated treatments). This indicates that as subjects gain experience they

learn to defect in all the treatments in which defection is the only equilibrium action under the provided monetary payoffs, while they learn to cooperate in some of the treatments in which cooperation is a possible equilibrium action. Thus, the effect of the shadow of the future increases with experience.²⁹

Second, given that previous studies have found differences in cooperative behavior between economics majors and other students (see Gerald Marwell and Ruth E. Ames, 1981; Robert Frank et al., 1993; Anthony M. Yezer et al., 1996; and David N. Laband and Richard O. Beil, 1999), it is important to note that the support for the hypothesis studied here does not depend on the major of the subjects. For example, for both economics majors and non-economics majors alike, cooperation increases as the probability of future interaction increases, and cooperation is greater in infinitely repeated games than in finitely repeated games. There are differences, however, in behavior across majors. Economics majors tend to cooperate less when cooperation cannot be supported as an equilibrium outcome (one-shot and finitely repeated games), but their behavior is not significantly different from that of other students when cooperation can be supported in equilibrium (infinitely repeated games).

V. Conclusions

This paper presents results from an experiment on infinitely repeated games that overcomes the methodological drawbacks of previous work. The results provide strong support for the theory of infinitely repeated games by showing that the shadow of the future matters, and that it significantly reduces opportunistic behavior,

²⁸ Note that to support Hypothesis 4, it is not necessary to find that subjects actually alternate to cooperate. Given the availability of a public randomization device, the players may follow complicated rules that depend on it and that do not necessarily result in the outcome CD after the outcome DC, nor on the same number of these two outcomes. Regardless of possible random process of alternation, it is interesting to have a measure of the amount of alternation that is observed in the data.

²⁹ Experience also has effects across treatments. The level of cooperation in the $\delta = 1/2$ treatment is lower in the Normal sessions than in the UD sessions (15.3 percent against 37.97 percent, with p -value less than 0.001). That is, in the $\delta = 1/2$ treatment, cooperation was lower when it followed a treatment with low levels of cooperation ($\delta = 0$) than when it followed a treatment with high levels of cooperation ($\delta = 3/4$). These spillover effects can be easily understood if one thinks of infinitely repeated games as coordinations games. The level of cooperation in the previous treatment may signal the type of equilibrium in which the subjects may coordinate in the present treatment. As with coordination games, this result points to the importance of history in equilibrium selection in infinitely repeated games.

closely following the theoretical predictions. Subjects cooperate more the greater the probability of future interaction, cooperate more in infinitely repeated games than in finitely repeated ones of the same expected length, and respond to small payoff changes as predicted by theory. Hence, the infinitely repeated game model not only provides a realistic representation of everyday interaction, but also explains cooperation as equilibrium behavior in a way that is consistent with actual behavior.

In addition, this paper presents suggestive evidence for future research on the effects of learning in infinitely repeated games. Among the treatments in which cooperation was a possible equilibrium action, we found that cooperation increased with experience in some of the treatments but in others did not. Thus, cooperation as an equilibrium does not appear to be a sufficient condition for subjects to learn to cooperate. This different impact of experience in different infinitely repeated games raises the issue of equilibrium selection in infinitely repeated games at both experimental and theoretical levels. To further the understanding of the determinants of cooperation, future studies should seek to uncover the conditions under which players learn to coordinate in equilibria that support cooperation in infinitely repeated games.

REFERENCES

- Abreu, Dilip; Pearce, David and Stacchetti, Ennio.** "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica*, 1990, 58(5), pp. 1041–63.
- Andreoni, James A.** "Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments." *Quarterly Journal of Economics*, 1995, 110(1), pp. 1–21.
- Andreoni, James A. and Miller, John H.** "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence." *Economic Journal*, 1993, 103(418), pp. 570–85.
- Aoyagi, Masaki and Fréchette, Guillaume R.** "Collusion in Repeated Games with Imperfect Public Monitoring." University of California, Los Angeles, Department of Economics Paper: No. 04-22, 2004.
- Bereby-Meyer, Yoella and Roth, Alvin E.** "The Fragility of Cooperation: Learning in Noisy Games." Unpublished Paper, 2002.
- Brown-Kruse, Jamie; Rassenti, Stephen; Reynolds, Stanley S. and Smith, Vernon L.** "Bertrand-Edgeworth Competition in Experimental Markets." *Econometrica*, 1994, 62(2), pp. 343–72.
- Cooper, Russell.** "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games." *Games and Economic Behavior*, 1996, 12(2), pp. 187–218.
- Duffy, John and Ochs, Jack.** "Cooperative Behavior and the Frequency of Social Interaction." University of California, Los Angeles, Department of Economics Working Paper: No. 04-01, 2004.
- Engle-Warnick, Jim and Slonim, Robert L.** "Inferring Repeated-Game Strategies from Actions; Evidence from Trust Game Experiments." Nuffield College, University of Oxford, Economics Papers: No. 2001-W13, 2001.
- Fehr, Ernst and Gächter, Simon.** "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 2000, 90(4), pp. 980–94.
- Feinberg, Robert M. and Husted, Thomas A.** "An Experimental Test of Discount-Rate Effects on Collusive Behavior in Duopoly Markets." *Journal of Industrial Economics*, 1993, 41(2), pp. 153–60.
- Fouraker, Lawrence E. and Siegel, Sidney.** *Bargaining behavior*. New York: McGraw-Hill, 1963.
- Frank, Robert H.; Gilovich, Thomas and Regan, Dennis T.** "Does Studying Economics Inhibit Cooperation?" *Journal of Economic Perspectives*, 1993, 7(2), pp. 159–71.
- Fudenberg, Drew and Levine, David K.** "Measuring Players' Losses in Experimental Games." *Quarterly Journal of Economics*, 1997, 112(2), pp. 507–36.
- Holt, Charles A.** "An Experimental Test of the Consistent-Conjectures Hypothesis." *American Economic Review*, 1985, 75(3), pp. 314–25.
- Kandori, Michihiro.** "Social Norms and Community Enforcement." *Review of Economic Studies*, 1992, 59(1), pp. 63–80.
- Kreps, David M.; Milgrom, Paul; Roberts, John and Wilson, Robert.** "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, 1982, 27(2), pp. 245–52.

- Laband, David N. and Beil, Richard O.** "Are Economists More Selfish Than Other 'Social' Scientists?" *Public Choice*, 1999, 100(1-2), pp. 85-101.
- Levine, David K.** "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics*, 1998, 1(3), pp. 593-622.
- Mailath, George J.; Obara, Ichiro and Sekiguchi, Tadashi.** "The Maximum Efficient Equilibrium Payoff in the Repeated Prisoners' Dilemma." *Games and Economic Behavior*, 2002, 40(1), pp. 99-122.
- Marwell, Gerald and Ames, Ruth E.** "Economists Free Ride, Does Anyone Else? Experiments on the Provision of Public Goods, IV." *Journal of Public Economics*, 1981, 15(3), pp. 295-310.
- Murnighan, J. Keith and Roth, Alvin E.** "Expecting Continued Play in Prisoner's Dilemma Games." *Journal of Conflict Resolution*, 1983, 27(2), pp. 279-300.
- Palfrey, Thomas R. and Rosenthal, Howard.** "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies*, 1994, 61(3), pp. 545-65.
- Rapoport, Anatol; Guyer, Melvin J. and Gordon, David G.** *The 2×2 game*. Ann Arbor: University of Michigan Press, 1976.
- Roth, Alvin E.** "Introduction to Experimental Economics," in John H. Kagel and Alvin E. Roth, eds., *The handbook of experimental economics*. Princeton: Princeton University Press, 1995, pp. 3-109.
- Roth, Alvin E. and Murnighan, J. Keith.** "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology*, 1978, 17(2), pp. 189-98.
- Selten, Reinhard and Stoecker, Rolf.** "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach." *Journal of Economic Behavior and Organization*, 1986, 7(1), pp. 47-70.
- Stahl, Dale O., II.** "The Graph of Prisoner's Dilemma Supergame Payoffs as a Function of the Discount Factor." *Games and Economic Behavior*, 1991, 3(3), pp. 368-84.
- Yezer, Anthony M.; Goldfarb, Robert S. and Poppo, Paul J.** "Does Studying Economics Discourage Cooperation? Watch What We Do, Not What We Say or How We Play." *Journal of Economic Perspectives*, 1996, 10(1), pp. 177-86.