

The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence

Pedro Dal Bó
Brown University

Guillaume R. Fréchette*
New York University

April 2007

Abstract

A usual criticism of the theory of infinitely repeated games is that it does not provide sharp predictions since there may be a multiplicity of equilibria. To address this issue we present experimental evidence on the evolution of cooperation in infinitely repeated prisoners' dilemma games as subjects gain experience. We find that cooperation decreases with experience when it cannot be supported as an equilibrium outcome. More interestingly, the converse is not necessarily true: cooperation does not always increase with experience when it can be supported as an equilibrium outcome. Nor is a more stringent condition, risk dominance, sufficient for cooperation to arise. However, subjects do learn to cooperate when the payoff to cooperation and the importance of the future is high enough. These results have important implications for the theory of infinitely repeated games. While we show that cooperation may prevail in infinitely repeated games, the conditions under which this occurs are more stringent than the sub-game perfect conditions usually considered.

* We are grateful to Anna Aizer, Ernesto Dal Bó, John Duffy, Dan Friedman, Alessandro Lizzeri, Muriel Niederle, Alvin Roth, Ennio Stacchetti and seminar participants at HBS and NYU for very useful comments. We thank Mark Dean for his research assistance, CASSEL (UCLA) and SSEL (Caltech) for the Multistage software as well as Rajeev Advani and Samuel Mati for its adaptation. Fréchette gratefully acknowledges the support of NSF via grant SES-0519045 as well as support from the Center for Experimental Social Science, and the C.V. Starr Center.

Introduction

A usual criticism of the theory of infinitely repeated games is that it does not provide sharp predictions since there may be a multiplicity of equilibria.¹ For example, in infinitely repeated prisoners' dilemma games with patient agents, both cooperate and defect may be played in equilibrium.

Even though the theory of infinitely repeated games has been used to explain cooperation in a variety of environments, no definitive solution has been provided to the problem of equilibrium selection: when both cooperation and defection are possible equilibrium outcomes, which one should we expect to prevail? Previous experimental evidence has shown that subjects often fail to coordinate in a specific equilibrium when they play a small number of infinitely repeated games: some subjects attempt to establish cooperative agreements while others defect. But how would behavior evolve as subjects learn from previous repeated games? Would cooperation prevail when it can be supported in equilibrium? Or are subjects condemned to learn that defection is the best individual action?

We present evidence on the evolution of cooperation in infinitely repeated games based on a series of experiments. For a given continuation probability and cooperation payoff, each subject participated in between 23 and 77 infinitely repeated games. This allows us to study how cooperation evolves as subject gain experience. First, we find that in treatments in which cooperation cannot be supported in equilibrium, the level of cooperation decreases with experience and converges to levels comparable to those observed in one-shot prisoners' dilemma games. This supports the idea that being a possible equilibrium action is a necessary condition for cooperation to arise with experience.

Second, we find that in treatments in which cooperation can be supported in equilibrium, the level of cooperation does not necessarily increase and may remain at low levels. When cooperation can be supported in a subgame perfect equilibrium, subjects may fail to make the most of it: being a possible equilibrium outcome is not enough for

¹ Fudenberg and Maskin (1993), for example, state that "The theory of repeated games has been somewhat disappointing. ...the theory does not make sharp predictions."

cooperation to rise with experience. Together, this evidence suggests that while being an equilibrium action may be a necessary condition for cooperation to arise with experience, it is not sufficient.

Third, we find that cooperation being both an equilibrium action and a risk dominant action (as defined later) is not sufficient either for cooperation to rise with experience. Risk dominance has been used as a selection criterion in the study of coordination games. While the experimental evidence on one-shot coordination games suggests that actions that are both Pareto efficient and risk-dominant are usually selected, our evidence suggests that those conditions are not sufficient in infinitely repeated games. However, we do observe that it is possible for subjects to learn to cooperate and reach high level of cooperation if the payoff from cooperation and the probability of future interactions are high enough. In infinitely repeated games, for cooperation to rise to high levels more is needed than just being an equilibrium and risk dominant action.

These results show how difficult it is for cooperation to arise even for experienced subjects. These results cast doubt on the common assumption that subjects will make the most of the opportunity to cooperate whenever it is possible to do so in equilibrium.

While there is a previous experimental literature in infinitely repeated games, this literature has not focused on the evolution of cooperation. Previous experimental evidence on infinitely repeated games has shown that cooperation is greater when it can be supported in equilibrium but that subjects fail to make the most of the opportunity to cooperate (see Roth and Murnighan, 1978, Murnighan and Roth, 1983, Palfrey and Rosenthal, 1994, Aoyagi and Fréchette, 2003, and Dal Bó, 2005). In addition, Dal Bó (2005) compares infinitely repeated and finitely repeated prisoner's dilemma games of the same expected length and finds that cooperation is larger in the former as theory predicts.² Aoyagi and Fréchette (2003) show that in infinitely repeated prisoner's dilemma games with imperfect public monitoring the level of cooperation increases with the quality of the public signal. Duffy and Ochs (2003) compare the levels of cooperation in random matching and fixed matching infinitely repeated games with high continuation probability. They find that cooperation increases as subjects gain more experience under

² Engle-Warnick and Slonim (2004) compares finitely and infinitely repeated trust games finding similar results. For more evidence on repeated trust games see Engle-Warnick and Slonim (2006).

fixed matching but not under random matching. Our experimental design differs from the previous literature in that, for several combinations of continuation probabilities and payoffs to cooperation, we allow subjects to participate in a large number of repeated games. In this way we can study how cooperation evolves under different treatments as subjects gain experience.

There have been theoretical attempts to solve the problem of multiplicity of equilibria in infinitely repeated games. Axelrod and Hamilton (1981), Boyd and Lorberbaum (1987), Boyd (1989), Kim (1994), and Bendor and Swistak (1997) apply the concept of evolutionary stable strategies (ESS) by Maynard Smith's (1982) -- or modifications -- to infinitely repeated games with varied implications regarding the selection of equilibria. There is also a literature that appeals to bounded rationality, in the form of finite automata. Rubinstein (1986) and Abreu and Rubinstein (1988) look at the set of equilibrium payoffs in repeated games played by infinitely patient finite automata with lexicographic costs of complexity and find that whether efficiency can be achieved depends on the particular equilibrium concept. Binmore and Samuelson (1992), Cooper (1996) and Volij (2002) apply evolutionary refinements to infinitely repeated games played by finite automata and find that the set of possible payoffs depends radically on the definition of ESS and the way costs of complexity are modeled (also see Fudenberg and Maskin, 1990 and 1993). Blonski and Spagnolo (2001) appeal to the concept of risk dominance as an equilibrium selection criteria in infinitely repeated games. In contrast, Volij (2002) shows that always defect is the unique stochastically stable strategy (Kandori et al., 1993, and Young, 1993) in games with finite automata. Johnson, Levine and Pesendorfer (2001) study stochastically stable strategies in a random matching gift-giving repeated game with local information systems. They find that cooperation (gift-giving) is stochastically stable only if the payoff from cooperation is above a critical value which exceeds what is required by sub-game perfection (see also Levine and Pesendorfer 2007).

This variety of theoretical results underscores the need for empirical data to solve the issue of multiplicity of equilibria in infinitely repeated games. The experimental results we present can inform future theories. Theories claiming that subjects will coordinate on defection even when they are infinitely patient and theories claiming that

they will always coordinate on cooperation are not supported by the data. However, we find empirical support for theories predicting cooperation under sufficiently favorable conditions.

2. Experimental design

This experiment consists of 18 sessions. In each session, a set of subjects participated anonymously through computers in a sequence of infinitely repeated prisoners' dilemma games. We induce an infinitely repeated game in the lab by having a random continuation rule: after each round the computer decided whether to finish the repeated game or have an additional round depending on a random number. We consider two probabilities of continuation: $\delta=1/2$ and $\delta=3/4$. The stage game is the simple prisoners' dilemma game in Table 1 where the payoffs are denoted in cents and where the payoff to cooperation takes one of three possible values: $R=32, 40$ and 48 .

Table 1: Stage Game Payoffs (in cents)

	C	D
C	R, R	12, 50
D	50, 12	25, 25

Therefore we have two main treatment variables, the probability of continuation and the payoff from cooperation, resulting in a total of six treatments. Each session consisted of a sequence of infinitely repeated games for one treatment only (between subjects design) and we run three sessions per treatment. In each session subjects participate in as many repeated games as was possible in one hour of play (the first repeated game to end after one hour of play marks the end of the session). Subjects were randomly re-matched with another subject after the end of a repeated game.³ Different

³ Random matching allows for a larger number of repeated games in a session than alternative matching protocols like complete strangers or turnpike protocols. The possibility of a pair of subjects interacting together in more than one repeated game is not likely to be a problem for several reasons. First, our results in section 4 suggest that the matching protocol does not introduce additional repeated games effects –for example, cooperation reaches one-shot levels when it cannot be supported in equilibrium. Second, Duffy and Ochs (2003) found that cooperative norms do not emerge under random matching. Third, Dal Bó

groups of subjects participated in each session. The instructions for one of the sessions are in the online appendix.⁴

The treatments and results are organized around three questions that derive from the theoretical background described in the next section.

3. Theoretical Background

If we assume that the payoffs in Table 1 are the actual total payoffs the subjects obtain from the stage game and that this is common knowledge, the set of subgame perfect equilibria can be calculated using the results from Stahl (1991). Table 2 indicates those treatments under which cooperation can be supported as a sub-game perfect equilibrium action.⁵

Table 2: Cooperation in Equilibrium

	R=32	R=40	R=48
$\delta = 1/2$	NO	YES	YES
$\delta = 3/4$	YES	YES	YES

Under $\delta = 1/2$ & $R = 32$, only defection is a possible equilibrium action and we expect that as subjects gain experience the levels of cooperation decrease to one-shot levels. However past experimental evidence indicates that there are games in which observed behavior does not converge to the unique equilibrium leading to the following question.

QUESTION 1: Do subjects learn to defect when it is the only equilibrium action?

Under all treatments except $\delta = 1/2$ & $R = 32$ cooperation can be supported in equilibrium⁶ and we may expect that cooperation increases with experience and reaches

(2005) uses a turnpike protocol with results consistent with other studies that have used random matching protocols.

⁴ Available at http://homepages.nyu.edu/~gf35/print/df_online_appendix.pdf.

⁵ More precisely, the critical value of δ over which cooperation can be supported in equilibrium is 0.72 under $R = 23$, 0.4 under $R = 40$, and 0.8 under $R = 48$.

⁶ In fact, it can be shown following Stahl (1991) that for all the treatments, but $\delta = 1/2$ and $R = 32$, the whole set of feasible and individually rational payoff can be supported in equilibrium. In addition, mutual

levels close to 100%. However, there is a multiplicity of equilibria under the treatments for which cooperation can be supported. While we may assume that subjects will learn to coordinate on the Pareto efficient equilibrium, this may not be the case. As shown by an extensive literature on coordination games, subjects may fail to coordinate on the Pareto efficient equilibrium if the costs from not coordinating are too high for the subject playing the Pareto efficient action (see Cooper et al., 1990, and Van Huyck et al., 1990).⁷ These alternative hypotheses lead us to the following question.

QUESTION 2: Do subjects learn to cooperate when it is an equilibrium action?

Previous literature has studied the concept of risk-dominance as an alternative equilibrium selection criterion. Risk-dominance was introduced by Harsanyi and Selten (1988) and concerns the pairwise comparison between Nash equilibria. In 2x2 coordination games an equilibrium is risk-dominant if its equilibrium strategy is a best response to a mixture that assigns probability of one-half to each strategy by the other player. While risk-dominance is easy to define and use in 2x2 games, it presents complications in general simultaneous-moves games. Its application to infinitely repeated games also faces the problem that two or more strategies can be identical to each other on the path of the game making it impossible to rank them (for example, consider different trigger strategies). Given the difficulties applying the concept of risk-dominance to the whole set of possible strategies in infinitely repeated games we focus only on a pairwise comparison of all strategies that support cooperation against the ultimate defection strategy: “always defect” (AD). Blonksi and Spagnolo (2001) show that the strategy “grim” (G) risk dominates AD if there is any cooperative strategy that dominates AD.⁸ In

cooperation can be supported in these five treatments with continuation payoffs in the efficient frontier of the set of feasible and individually rational payoffs. Therefore, mutual cooperation is renegotiation proof under most renegotiation proofness concepts in the five treatment in which it can be supported in equilibrium (see Bernheim and Ray, 1989, and Farrell and Maskin, 1989; also see Pearce, 1992, and references therein, for a review of the different renegotiation proofness concepts).

⁷ This is not to say that subjects never coordinate on the efficient outcome, but rather that they sometimes do not. Charness, Fréchet and Qin (2006) provide an example where subjects often do coordinate on the efficient outcome in a game where the possibility of contingent rewards transform (in equilibrium) the second stage game to a coordination game.

⁸ The grim strategy is the strategy that starts by cooperating and continues to do so as long as the other player cooperates, but defects forever following a defection by the other player.

other words, G is the “less risky” of the cooperative strategies when matched with someone playing AD, and we only need to focus on the comparison between G and AD.

Table 3 shows the treatments under which G is risk-dominant against AD.⁹

Table 3: Risk-Dominant Cooperation

	R=32	R=40	R=48
$\delta=1/2$	NO	NO	YES
$\delta=3/4$	NO	YES	YES

If subjects learn to cooperate when cooperation is an efficient equilibrium action and is also risk-dominant (as defined above and argued by Blonksi and Spagnolo, 2001), we should observe that cooperation increases with experience and reaches levels close to 100% under the following three treatments: $\delta=1/2$ & R=48, $\delta=3/4$ & R=40, and $\delta=3/4$ & R=48. This reasoning leads us to the following question.

QUESTION 3: Do subjects learn to cooperate when it is risk-dominant?

4. Experimental Results

We conducted 18 experimental sessions between July 2005 and March 2006. A total of 266 subjects participated in the experiment, with an average of 14.78 subjects per session, a maximum of 20 and a minimum of 12. The subjects were NYU undergraduates recruited through email solicitation at the beginning of the semester. The subjects earned an average of \$25.95, with a maximum of \$42.93 and a minimum of \$16.29. In the treatments with $\delta=1/2$ and $\delta=3/4$ the average number of rounds per match was 1.96 and 4.42 respectively, and the maximum was 9 and 23 respectively. Some descriptive statistics of the sessions are presented in Table 4.

⁹ More precisely, the critical value of δ over which G is risk-dominant against AD is 0.82 under R=32, 0.61 under R=40, and 0.39 under R=48.

Table 4: Session characteristics

Sessions	Payoff from cooperation	Delta 0.5			Delta 0.75			
		32	40	48	32	40	48	
1-6	Number of subjects	16	14	14	Number of subjects	14	12	16
	Number of Games	69	72	72	Number of Games	29	34	35
	Number of rounds	2.01	2.08	2.01	Number of rounds	5.55	3.94	4.31
	Average Payoff	24.22	29.17	30.44	Average Payoff	28.39	31.35	42.93
	Maximum Payoff	22.29	24.90	28.05	Maximum Payoff	26.26	29.78	39.99
	Minimum Payoff	18.79	22.10	25.20	Minimum Payoff	24.60	27.65	35.23
7-12	Number of subjects	12	16	14	Number of subjects	16	14	12
	Number of Games	71	71	68	Number of Games	33	47	29
	Number of rounds	1.89	2.03	2.00	Number of rounds	4.58	3.15	5.14
	Average Payoff	21.45	26.26	26.62	Average Payoff	26.47	29.29	39.56
	Maximum Payoff	20.53	23.71	23.82	Maximum Payoff	24.14	27.41	36.63
	Minimum Payoff	17.75	21.61	21.43	Minimum Payoff	21.02	25.55	33.32
13-18	Number of subjects	16	20	18	Number of subjects	14	12	16
	Number of Games	59	72	77	Number of Games	27	23	32
	Number of rounds	2.05	1.74	1.84	Number of rounds	4.70	5.43	4.06
	Average Payoff	19.50	22.05	29.36	Average Payoff	21.59	27.58	34.71
	Maximum Payoff	18.35	20.15	27.39	Maximum Payoff	20.51	24.97	30.98
	Minimum Payoff	16.29	18.16	25.01	Minimum Payoff	19.57	23.91	25.76

4.1 General description of behavior.

Before answering the questions from the previous sections we provide a general description of the observed behavior. The first panel in Table 5 shows cooperation rates by treatment for the first repeated game, on the left for the first round and on the right for all rounds. Looking separately at first rounds is of importance since different repeated games may result in a different number of rounds and the percentage of cooperation may vary across rounds. Cooperation is significantly larger under $\delta=3/4$ & $R=48$ than under $\delta=1/2$ & $R=32$ in the first repeated game (p-value of 0.032 and 0.002 for first round and all rounds respectively).¹⁰ In addition, cooperation tends to be higher in treatments under which cooperation can be supported as an equilibrium action than when it cannot be

¹⁰ The statistical tests are t-tests with the variance-covariance modified by clustering at the subject level. This is true of every test reported unless noted otherwise. The qualitative results are largely robust to clustering at the session level or using the nonparametric Wilcoxon test with session averages as observations. The results on the evolution of cooperation that we present in the following sections are also largely robust to considering linear trends. Results from using these alternative statistical tests are provided in the online appendix at http://homepages.nyu.edu/~gf35/print/df_online_appendix.pdf.

supported, but this difference is only significant for first rounds (p-values of 0.087 and 0.167 for first rounds and all rounds respectively). However, it is not the case that an increase in the probability of continuation always results in an increase in cooperation (compare the two treatments with R=40) and increases in the payoff from cooperation has no significant effects on cooperation.

Table 5: Percentage of Cooperation by Treatment

First Repeated Game											
First Round						All Rounds					
$\delta \setminus R$	32		40		48	$\delta \setminus R$	32		40		48
0.5	34.09	<*	54.00	<	56.52	0.5	28.33	<	39.80	<	41.38
	=		v		^		v		v*		^
0.75	34.09	<	36.84	<*	56.82	0.75	21.76	<	26.36	<***	56.10

All Repeated Games											
First Round						All Rounds					
$\delta \setminus R$	32		40		48	$\delta \setminus R$	32		40		48
0.5	9.81	<***	18.72	<***	38.97	0.5	9.82	<***	17.98	<***	35.29
	^***		^***		^***		^***		^***		^***
0.75	25.61	<***	61.10	<***	85.07	0.75	20.25	<***	58.71	<***	76.42

Note: * significance at 10%, ** at 5% and * at 1%.

The second panel of Table 5 shows cooperation rates for all repeated games. Here, increases in the probability of continuation or the payoff of cooperation result in a significant increase in cooperation (p-value of less than 0.001 and 0.001 for first round and all rounds respectively). In addition, cooperation is significantly greater in treatments under which cooperation can be supported as an equilibrium action than when it cannot be supported (p-values of less than 0.001 and 0.001 for first rounds and all rounds respectively).

These differences between behavior in the first match and all matches suggest that experience affects how subjects play in repeated games in an important way. The next sections focus on how subjects modify their behavior as they gain experience.

4.2 Do subjects learn to defect when it is the only equilibrium action?

To answer this question we study the evolution of cooperation under $\delta=1/2$ & $R=32$, the treatment in which cooperation cannot be supported in equilibrium. The first

column of Table 6 shows the percentage of subjects that choose to cooperate in the first round of each repeated game in this treatment with the repeated games aggregated according to the interaction in which they started.¹¹ To compare inexperienced versus experienced play we compare behavior in the first 10 interactions with those in interactions 111 to 120.¹²

Table 6: Percentage of Cooperation by Equilibrium Conditions (First Rounds)

Repeated Game Begins In Interaction	First Round Cooperation is		All Rounds Cooperation is	
	Not SGPE	SGPE	Not SGPE	SGPE
1-10	28.57	39.11	21.00	34.42
11-20	13.04	28.54	12.91	27.19
21-30	12.23	31.01	11.97	33.61
31-40	10.61	36.04	10.51	38.64
41-50	10.20	34.88	7.85	34.98
51-60	9.75	41.47	6.54	39.85
61-70	7.14	37.89	8.09	40.02
71-80	5.65	36.86	4.48	39.73
81-90	4.72	38.60	6.20	44.39
91-100	6.11	40.91	7.91	47.11
101-110	6.64	45.38	11.99	46.12
111-120	5.50	49.77	6.45	55.88
121-130	5.77	45.95	11.11	43.31
131-140	8.33	47.43	9.17	42.99
141-		46.32		47.83

Cooperation was 29% in the first round of the repeated games that begin within the first 10 interactions, dropping to 5.5% in the repeated games that begin within interactions 111 to 120 (this difference is significant with p-value below 0.01). For any repeated game that starts after 50 interactions cooperation is always below 10%. These levels are similar to the levels observed in one-shot prisoners' dilemmas (for example see Cooper et al, 1996, Dal Bó, 2005 and Bereby-Meyer and Roth, 2006). The evolution of cooperation is similar if we aggregate the data from all rounds (see column 3 in Table 6).

¹¹ We use the word interaction to number each decision stage regardless of the repeated game. For example, if the first repeated game lasted for 5 rounds, the first round of the second repeated game is the sixth interaction. We use the word round to number decision stages inside a repeated game.

¹² We do have data on repeated games that started even later, but because there are slight variations in total number of interactions and length of particular repeated games across sessions, the sample size is stable only up to interactions 111-120.

We also reach similar conclusions if we present the data in groups of 10 repeated games instead of interactions.

From the aggregated data in this treatment it is clear that subjects learn to defect and cooperation reaches negligible levels when cooperation cannot be supported in equilibrium. We reach a similar conclusion when we study the evolution of cooperation in each session under this treatment. Figure 1 displays the proportion of cooperation in the first round of each repeated game by session and treatment. The first graph in Figure 1 displays the evolution of cooperation for the three sessions with $\delta=1/2$ & $R=32$: session 1 as a solid line, 2 as dashed line and 3 as a dash-dotted line. It is clear from this graph that cooperation decreases with experience in all three sessions.

4.3 Do subjects learn to cooperate when it is an equilibrium action?

The second column in Table 6 shows the percentage of subjects that choose to cooperate in the first round of the repeated games under which cooperation can be supported in sub-game perfect equilibrium. Initially, cooperation was 39%, but in repeated games after 111 to 120 rounds of experience cooperation increased to 50% (p-value of the difference less than 0.01). We observe a similar evolution of cooperation from all the rounds in the repeated games (see the fourth column in Table 6). In addition, there are important differences with respect to the treatment in which cooperation is not a possible equilibrium outcome. The difference in cooperation rates across the equilibrium and non equilibrium situation is statistically significant for the first 10 interactions (p-value < 0.05) and even more so for the repeated games that start after 111-120 interactions (p-value < 0.01) for both first round and all rounds.

These results support the idea that subjects improve their ability to make the most of the opportunity to cooperate as they gain experience, but only slightly. They are still very far from all coordinating on the efficient outcome.

The study of each session separately shows that cooperation being a possible equilibrium outcome does not necessarily lead to high levels of cooperation as subjects gain experience. The last five graphs in Figure 1 displays the evolution of cooperation by session and treatment for the treatments in which cooperation can be supported. Note that cooperation is lower in the last repeated game than in the first repeated game in 8 of these

sessions and it is higher in 7. While there is large variation in the evolution of cooperation across these treatments, it is clear that cooperation being a possible equilibrium outcome does not necessarily lead to increasing levels of cooperation as subjects gain experience.

4.4 Do subjects learn to cooperate when it is risk-dominant to do so?

In this section we examine whether subgame perfection combined with risk dominance may be sufficient for subjects to learn to make the most of the opportunities for cooperation. Table 7 shows the percentage of subjects that choose to cooperate in the first round and all rounds of the repeated games for treatments under which cooperation is an equilibrium, separately by whether cooperation is risk dominant ($\delta=1/2$ & $R=48$, $\delta=3/4$ & $R=40$, and $\delta=3/4$ & $R=48$) or not ($\delta=1/2$ & $R=40$, and $\delta=3/4$ & $R=32$).

Table 7: Percentage of Cooperation by Risk Dominance (First Rounds)*

Repeated Game Begins in Interaction	First Round Cooperation is		All Rounds Cooperation is	
	Not RD	RD	Not RD	RD
1-10	31.43	46.53	23.56	42.11
11-20	20.60	36.26	18.10	35.09
21-30	14.86	44.34	13.48	45.36
31-40	14.01	51.83	14.63	52.72
41-50	14.21	53.99	13.81	53.09
51-60	18.51	57.47	16.32	61.30
61-70	17.54	48.98	19.21	54.44
71-80	20.32	50.00	19.10	55.99
81-90	20.57	58.42	20.75	60.89
91-100	22.01	54.88	19.28	66.45
101-110	17.93	67.62	19.50	66.92
111-120	22.46	70.61	22.60	73.86
121-130	21.03	62.05	21.99	59.60
131-140	30.70	59.49	26.23	61.40
141-	23.86	65.69	16.57	76.82

*Only considers treatments in which cooperation is a SGPE action.

Table 7 shows that while cooperation decreases with experience when it is not risk-dominant, it increases with experience when it is risk dominant. While in the first rounds of the early repeated games in the risk-dominant treatments cooperation was 46.5%, in later repeated games (after 110 interactions) it reached 70.6% (p-value less

than 0.01). We observe a similar evolution for all rounds. These aggregated data suggests that being a possible equilibrium outcome and risk-dominant does lead to increasing levels of cooperation as subjects gain experience. The difference in cooperation rates across the risk-dominant and non-risk-dominant case is statistically significant both at the beginning and for the repeated games that start after 111-120 interactions for both first and all rounds (p -values < 0.01). Nonetheless, the cooperation rate when cooperation is risk-dominant and after subjects have gained much experience is still far away from full cooperation.

We reach an even more nuanced conclusion if we study these treatments by session. The graphs in Figure 1 for $\delta=1/2$ & $R=48$, $\delta=3/4$ & $R=40$, and $\delta=3/4$ & $R=48$ display the evolution of cooperation for the sessions under the treatments in which cooperation can be supported and is risk-dominant. Cooperation is lower in the last repeated game than in the first repeated game in 3 sessions and higher in 6. While there is large variation in the evolution of cooperation, it is clear that cooperation being risk-dominant does not necessarily lead to increasing levels of cooperation as subjects gain experience. However, all session in the treatment with $\delta=3/4$ & $R=48$ reach high levels of cooperation. This suggests that if both the probability of continuation and the payoff of cooperation are high enough it is possible for subjects to make the most of the opportunity to cooperate.

4.5 Strategies

It is of interest to study the strategies used by subjects. Unfortunately, looking at the individual data does not allow one to identify the strategies. For example, observing a pair of subjects that cooperate in every round is consistent with a number of cooperative strategies like always cooperate (AC) and any kind of trigger strategy including grim (G). However, it is still possible and interesting to study what proportion of the observed behavior is consistent with some theoretically important strategies and how their prevalence changes with the treatment variables.

Table 8 indicates the fraction of behavior consistent with the strategies always defect (AD), always cooperate (AC), grim trigger strategy (G), and tit-for-tat (TFT). Note

that AD and AC are mutually exclusive but AC, G and TFT are all consistent with a subject that cooperates in every round if the other player always cooperates as well.¹³

Table 8: Categorization of Observed Behavior

All Repeated Games						
	Treatment	Always D	Always C	Grim	Tit-for-tat	AD + Grim
$\delta = 0.5$	R = 32	86.37	6.31	7.91	7.98	94.28
	R = 40	74.30	12.50	15.37	15.46	89.68
	R = 48	56.28	31.47	33.74	33.65	90.02
$\delta = 0.75$	R = 32	61.20	11.74	15.24	16.01	76.45
	R = 40	26.90	47.69	49.85	50.89	76.75
	R = 48	8.59	68.80	74.51	75.00	83.10
	Average	59.93	24.87	27.61	27.85	87.54
Repeated Games that Start In the First 10 Interactions						
$\delta = 0.5$	R = 32	62.95	14.73	21.43	21.43	84.38
	R = 40	60.82	19.88	27.19	27.19	88.01
	R = 48	49.16	35.29	35.29	34.03	84.45
$\delta = 0.75$	R = 32	65.91	13.64	18.94	18.94	84.85
	R = 40	42.24	20.69	21.55	20.69	63.79
	R = 48	24.26	30.88	38.24	41.18	62.50
	Average	53.45	22.64	27.53	27.53	80.98
Repeated Games that Start In Interactions 110-120						
$\delta = 0.5$	R = 32	90.50	1.50	4.50	5.00	95.00
	R = 40	72.38	16.67	19.52	18.57	91.90
	R = 48	49.51	36.89	37.86	38.84	87.38
$\delta = 0.75$	R = 32	64.63	14.02	17.07	17.68	81.71
	R = 40	23.08	54.81	58.65	59.62	81.73
	R = 48	0.00	92.22	93.33	95.00	93.33
	Average	53.1	33.83	36.18	36.75	89.28

The percentage of observed behavior that can be explained by AD is decreasing in the probability of future interaction and the payoff to cooperation (from 86% of the observations under $\delta=1/2$ & $R=32$ to 9% under $\delta=3/4$ & $R=48$). In contrast, the percentage that can be explained by the cooperative strategies (AC, G, and TFT) is

¹³ If two subjects (1 pair) played one repeated game and one subject always cooperated while the other always defected, the fraction of behavior consistent with AD would be 0.5, the fraction consistent with AC would also be 0.5 and the fraction consistent with G and TFT would be 0. If on the other hand both subjects always cooperated, then the fraction of AD would be 0, and for all others (AC, G, TFT) it would be 1.

increasing in the probability of future interaction and the payoff to cooperation. For example, the percentage of observations that can be explained by G goes from 8% under $\delta=1/2$ & $R=32$ to 75% under $\delta=3/4$ & $R=48$. These differences increase with experience as the comparison of the second two panels of Table 8 show. For example, in the repeated games starting in interactions 110 to 120 AD and G explain respectively 90% and 5% of the observations under $\delta=1/2$ & $R=32$ and 0% and 93% under $\delta=3/4$ & $R=48$.

Finally, note that focusing on AD and G allows us to explain an important part of the data (88%) and this percentage increases with experience.

5. Cooperation in the long run

The analysis of the evolution of cooperation as subjects gain experience thus far is limited by the number of repeated games that a subject may play in a given session (in this case the subjects played between 23 and 77 repeated games). In this section we study how behavior would evolve under an even greater number of repeated games by estimating a model of learning based on our experimental data and carrying out simulations that allow us to study the evolution of cooperation over a longer period than what we are able to observe in the laboratory.

To motivate our learning model we start by focusing on AD and G strategies and calculating the basins of attraction of each strategy.¹⁴ Imagine that someone is considering whether to play AD or G. This subject needs to determine which of these two strategies is the most profitable in expectation based on his beliefs about the probability that his partner will play AD or G. In Figure 1 the horizontal dotted line represents the belief that would leave the subject indifferent between the two strategies for each treatment.¹⁵ If the belief falls below the dotted line the subject will maximize his expected payoff by playing AD and if the belief is above the line, choosing G will maximize the expected payoff. These lines denote the limits of the basin of attraction of

¹⁴ As mentioned earlier (see Table 8), almost 90% of the entire data set is consistent with the AD and G strategies. Thus focusing on those will substantially simplify the analysis without being too far from the actual data. These 2 particular strategies also have the advantage of being very familiar to theorists.

¹⁵ There is no indifference line for $\delta=1/2$ and $R=32$ since G is never a best response in this treatment. For the other treatments it is 0.72, and 0.38 for δ equal to 0.5 in increasing order of R and 0.81, 0.27, and 0.16 for δ equal to 0.75 again in the order of increasing R.

the two strategies. As can be seen in Figure 1, when cooperation starts below the limit of the basins of attraction it is more likely that cooperation will diminish with experience while the opposite happens when cooperation starts above the limit.

The fact that basins of attractions seem to explain how behavior separates across treatments is very suggestive of an adaptive model of behavior. It may be that after learning the rules of the game, subjects form beliefs about the fraction of subjects that will play G and the fraction that will play AD. Given these beliefs they choose the strategy that would maximize their expected payoffs, and play according to that strategy. Upon observing the behavior of their partners, subjects update their beliefs and this determines the strategy used in the repeated games that follow. If this explanation is correct, we should observe a positive correlation between a subject's first round choice and the first round choice of his partner in the previous repeated game. Table 9 shows the results from a random effects Probit estimation of the first round decision as a function of the first round choice of the previous repeated game partner. We control for the individual choices in round 1 of the first repeated game to allow the mean of the random effects to be different for subjects who cooperated and those who did not in the first decision of the entire experiment. The coefficient ρ gives the fraction of the variance attributable to the subject specific component.

**Table 9: Effect of Past Observations on Round 1 Cooperation
(Correlated Random Effects Probit)**

	$\delta = 0.5$			$\delta = 0.75$		
	R = 32	R = 40	R = 48	R = 32	R = 40	R = 48
Opponent Cooperated in Round 1 of Last Repeated Game	0.426*** (0.117)	0.349*** (0.070)	0.590*** (0.064)	0.391*** (0.105)	0.920*** (0.117)	0.857*** (0.130)
Cooperated in Round 1 Of First Repeated Game	1.149*** (0.218)	0.721*** (0.245)	1.564*** (0.347)	0.497 (0.375)	1.725*** (0.498)	0.495 (0.346)
Constant	-2.167*** (0.144)	-1.687*** (0.186)	-1.719*** (0.184)	-1.277*** (0.232)	-0.653** (0.306)	0.568** (0.261)
P	0.271 [§] (0.063)	0.398 [§] (0.057)	0.688 [§] (0.053)	0.553 [§] (0.069)	0.657 [§] (0.066)	0.502 [§] (0.079)
Observations	2840	3534	3300	1268	1304	1376
Number of Subjects	44	50	46	44	38	44

Standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

[§] significant at the 1% level using a likelihood ratio test

The variable of interest, partner's choice in round 1 of the last repeated game, is statistically significant and positive in every treatment. This suggests that, when a subject observes his partner cooperating, he is more likely to start by cooperating himself in the following repeated game.

5.1 A Learning Model

Given the evidence that subjects modify their behavior as a function of the observed behavior of their partners from previous repeated games we study in this section a model in which subjects learn from the behavior of their partners about the distributions of strategies in the population.

We model the way subjects update their beliefs about the probability of facing different strategies using a belief based learning model (see Fudenberg and Levine, 1998). Estimating a learning model from repetitions of randomly terminated repeated games can be difficult as the set of possible strategies within a repeated game can be extremely large. However, since more than three quarters of the data can be explained by two very simple strategies (AD and G), as seen in section 4.5, we will abstract from the complexities of the repeated game by reducing it to the choice in round 1: defect corresponds to AD and cooperate corresponds to G. The expected return from each choice is given by the theoretical values. Using the theoretical values has the advantage that it is not affected by the small sample variations in the sample, but implies that subjects think everyone else is using either AD or G perfectly.

We will estimate a model which allows for behavior ranging from Cournot to fictitious play, as in Crawford (1995) and Cheung and Friedman (1997). The learning model is as follows. Subjects in the first repeated game have beliefs about the probability their partner uses either AD or G. These beliefs are tracked by two variables: β_{it}^{AD} and β_{it}^G such that the belief by subject j at time t that his partner will play AD is $\beta_{it}^{AD} / (\beta_{it}^{AD} + \beta_{it}^G)$. In the first repeated game each subject has a given β_{i1}^{AD} and β_{i1}^G . After the first repeated game is played they update their beliefs as follows $\beta_{it+1}^k = \theta_i \beta_{it}^k + 1(a_j^k)$ where θ_i discounts past beliefs ($\theta_i = 0$ gives Cournot dynamics and $\theta_i = 1$ is fictitious play), k is the action and $1(a_j^k)$ is an indicator function that takes value 1 if subject j (with

whom i is paired) took the action k (G or AD). Given those beliefs, subject i is modeled as a random utility maximizer where each choice yields

$$U_{it}^a = \frac{\beta_{it}^{AD}}{\beta_{it}^{AD} + \beta_{it}^G} u^a(a_j^{AD}) + \frac{\beta_{it}^G}{\beta_{it}^{AD} + \beta_{it}^G} u^a(a_j^G) + \lambda_{it} \varepsilon_{it}^a \text{ where } u^a(a_j^k) \text{ is the average payoff}$$

from taking action a when subject i is paired with j and takes action k . The parameter λ_{it} is a scaling parameter that measures how well the subject best-responds to his beliefs where $\lambda_{it} = \lambda^F_i + \phi_i^{t-1} \lambda^V_i$ with $\phi_i \in [0,1]$, and with λ^F_i and λ^V_i being positive and representing the fixed and variable parts of the scaling parameter. That is, we allow for the noise in decision making to decrease with experience. Finally, ε_{it}^a is an idiosyncratic error term assumed to have a type I extreme value function. Given the distributional assumption on the error terms, this gives rise to the usual logistic form for the

probabilities:
$$p_{it}^a = \frac{\exp\left(\frac{1}{\lambda_{it}} U_{it}^a\right)}{\exp\left(\frac{1}{\lambda_{it}} U_{it}^{AD}\right) + \exp\left(\frac{1}{\lambda_{it}} U_{it}^G\right)}$$
. Thus the parameters to be estimated

are β_{it}^{AD} , β_{it}^G , θ_i , λ^F_i , λ^V_i and ϕ_i .

The estimates are obtained via maximum likelihood estimation for each subject separately.¹⁶ We have between 23 and 77 round 1 observations per subject. Subjects whose round 1 action is always the same are dropped from the estimation sample –this represents 19.55% of the data. Summary statistics of the estimates are presented in Table A1 in the appendix.

Note that we are not interested in the specific estimates of the learning model but that our focus is on the out of sample predictions we present in section 5.3.¹⁷ However, there are several characteristics of the estimated model that are of interest. First, many subjects initially believe one strategy or the other is likely, but not both. Second, there is

¹⁶ An alternative would be to pool the data. However, for the purpose of this paper and given the number of observations per subjects, obtaining subject specific estimates seem reasonable. Fréchet (2005) discusses issues and solutions related to pooling data across subjects in estimating learning models and more specifically with respect to hypothesis testing.

¹⁷ We have estimated simpler versions of the learning model and while the fit with the experimental data was decreased, the out of sample predictions were similar. This suggests that our prediction on the long run evolution of cooperation is robust to changes in the learning model and we expect that alternative models will generate similar qualitative results (see for instance Roth and Erev, 1995, and Camerer and Ho, 1999).

a lot of variability in the discounting of past evidence. In some treatments many subjects look like they exhibit the Cournot dynamics while in others more subjects look like fictitious players. As δ and R increase subjects tend to pay more attention to the distant past, while for low δ and low R subjects tend to look only at the immediate past. Finally, note that the effect of experience in reducing noise tends to increase the larger δ and R are.

Correlations among the parameters also yield some additional insights. The fact that θ_i is negatively correlated with the strength of beliefs in the first repeated game ($\beta_{i1}^{AD} + \beta_{i1}^G$) suggests that subjects who pay less attention to the far past have stronger initial beliefs. The positive correlation between θ_i and ϕ_i suggest that subjects who pay more attention to the far past have flatter scaling parameters across time. The negative correlation between ϕ_i and the two terms of the scaling parameter (λ^F_i and λ^V_i) suggests that subjects with a greater initial scaling parameter have also a steeper scaling parameter across time.

Using these estimates we perform simulations to assess how well the learning model fits the data obtained in the experimental sessions. These simulations consist of 1000 sessions by treatment using the learning model previously estimated and adding the subjects that always played the same action and assuming that they would do so irrespective of the choices of the subjects they are paired with. The session size is taken to be 14 (which is the closest to the mean session size).¹⁸ The composition of each session is obtained by randomly drawing (with replacement) 14 subjects (and their estimated parameters) from the pool of subjects that participated in the corresponding treatment. Figure 2 displays the average simulated evolution of cooperation across repeated games by treatment (dashed line), in addition to the observed evolution (solid line). The vertical lines denote the end of each experimental session conducted in the lab. The dotted lines denote the upper and lower bounds to the interval that includes 90% of the 1000 simulated sessions.

The simulations based on the estimated learning model track well the evolution of cooperation observed in the data. First, note that for every treatment in which cooperation

¹⁸ The results from the simulations are robust to the number of subjects per simulated session.

is lower (greater) in the last repeated game than in the first repeated game the same is true for the simulations. Second, the experimental data are largely within the 90% interval generated by the simulations. Finally, for the range of repeated games for which we have experimental data from all three sessions, the average level of cooperation in the simulations while obviously less noisy is generally similar to the observed levels with differences of 5% on average, but that never exceed 19%.

5.3 Simulation of the evolution of cooperation in the long run

Given that the learning model fits the data well, we now extend the simulations to a longer range (1000 repeated games) to study how cooperation evolves in the long run. Figure 3 shows the evolution of the average level of cooperation across the simulations and the 90% interval.

The results are consistent with those of the experimental sessions (most of the convergence in behavior happens in the first 100 repeated games). In the treatment in which cooperation cannot be supported in equilibrium the simulated levels of cooperation converge to one shot levels (less than 5%). In addition, the 90% interval includes full defection from very early repeated games and never includes full cooperation. In treatments in which cooperation can be supported in equilibrium but G is not risk dominant, cooperation decreases with experience, converging to levels close to those observed in one shot games. In this case as well, the 90% interval includes full defection from very early repeated games and excludes full cooperation. In fact no simulated session under these treatments achieved full cooperation in any repeated game.

In contrast, for the treatments in which cooperation can be supported in equilibrium and G is risk dominant cooperation may reach much higher levels. For two of the treatments in this group, $\delta=3/4$ & $R=40$ and $\delta=3/4$ & $R=48$, cooperation does reach high levels after subjects have gained experience and the 90% interval includes full cooperation after 30 repeated games. In the case of $\delta=3/4$ & $R=48$ the mean level of cooperation is practically 100%. However, in the remaining treatment, $\delta=1/2$ & $R=48$, cooperation remains between 40% and 50% and the 90% interval sometimes includes full defection but it never includes full cooperation. Thus, in treatments in which G is risk

dominant and with a large number of repeated games for subjects to gain experience, full cooperation may fail to arise.

Figure 4 shows the simulated distribution of cooperation by treatment. The figure shows for each treatment and repeated games 1 and 1000 the proportion of all the simulated sessions that have a given number of subjects (out of 14) choosing to cooperate in the first round. Figure 4 also shows the limit of the basins of attraction of G and AD for each treatment: for example under $\delta=3/4$ & $R=48$ if the subject expects that 3 or more of the subjects in the session play G then G is the best response and AD is the best response otherwise.

The distributions of the number of cooperative actions are unimodal for every treatment in repeated game 1 and this mode is always located in the interior (from 4 cooperative actions in $\delta=1/2$ & $R=32$ and $\delta=3/4$ & $R=32$ to 8 in $\delta=3/4$ & $R=48$). The distribution of cooperative actions is quite different in repeated game 1000. In the treatment in which cooperation cannot be supported in equilibrium ($\delta=1/2$ & $R=32$) the mode and median is full defection. In the treatments under which cooperation can be supported in equilibrium but is not risk dominant, the mode decreases with experience and converges to 1 cooperative action over 14 in a session. In the treatments in which cooperation is both an equilibrium action and risk dominant, the results are diverse. In the case of $\delta=1/2$ & $R=48$ the distribution in repeated game 1000 is bimodal, with modes in 1 and 11 cooperative actions over 14. This bifurcation in the evolution of cooperation resembles the continental divide results from the coordination games literature (see Van Huyck et al. 1997). For the other two treatments in this group ($\delta=3/4$ & $R=40$ and $\delta=3/4$ & $R=48$) the distribution moves to higher levels of cooperation with experience with an extreme result for $\delta=3/4$ & $R=48$ where the mode and median converge to full cooperation.

5. Conclusions

The series of experiments presented in this paper shed light on how cooperation evolves as subjects gain experience. We vary both the probability of continuation and the payoff from cooperation. We find that cooperation may not prevail even when it is a

possible equilibrium action. This provides a word of caution against the extended practice in applications of the theory of infinitely repeated games of assuming that subjects will cooperate whenever it is an equilibrium action. Moreover, cooperation may not prevail even under more stringent conditions (risk dominance) indicating how difficult it is for cooperation to prevail in repeated games. However, cooperation does prevail under some treatments –namely, when the probability of continuation and the payoff from cooperation are high enough. This evidence contradicts some equilibrium selection theories that select inefficient outcomes even when players are arbitrarily patient. We hope the evidence provided here will guide future theoretical attempts to study equilibrium selection in infinitely repeated games.

Appendix 1

Table A1: Summary Statistics of Learning Model Estimates

Treatment	P(G)	All P(G)*	Mean		λ^F	λ^V
			θ	Φ		
$\delta=1/2$ R=32	0.79	0.61	0.21	0.23	4.49E+13	1.33E+299
$\delta=1/2$ R=40	0.53	0.48	0.58	0.15	3.37	2.93E+304
$\delta=1/2$ R=48	0.46	0.40	0.82	0.10	1.94	2.21E+301
$\delta=3/4$ R=32	0.39	0.31	0.32	0.31	6.14E+17	7.46E+287
$\delta=3/4$ R=40	0.32	0.34	0.55	0.28	1.16	5.86E+280
$\delta=3/4$ R=48	0.25	0.47	0.32	0.06	3.10	3.4E+297

Treatment	P(G)	All P(G)*	Median		λ^F	λ^V
			θ	Φ		
$\delta=1/2$ R=32	1.00	1.00	0.00	0.00	2.89	3.217E+36
$\delta=1/2$ R=40	0.71	0.53	0.64	0.00	1.85	2.973E+17
$\delta=1/2$ R=48	0.33	0.17	0.98	0.00	0.26	5.04E+102
$\delta=3/4$ R=32	0.00	0.00	0.00	0.00	3.33	2258.90
$\delta=3/4$ R=40	0.00	0.00	0.54	0.00	0.23	10.08
$\delta=3/4$ R=48	0.00	0.26	0.26	0.00	0.28	3.058E+71

*Assigns P(G)=1 (0) to subject that always choose to cooperate (defect) in first round

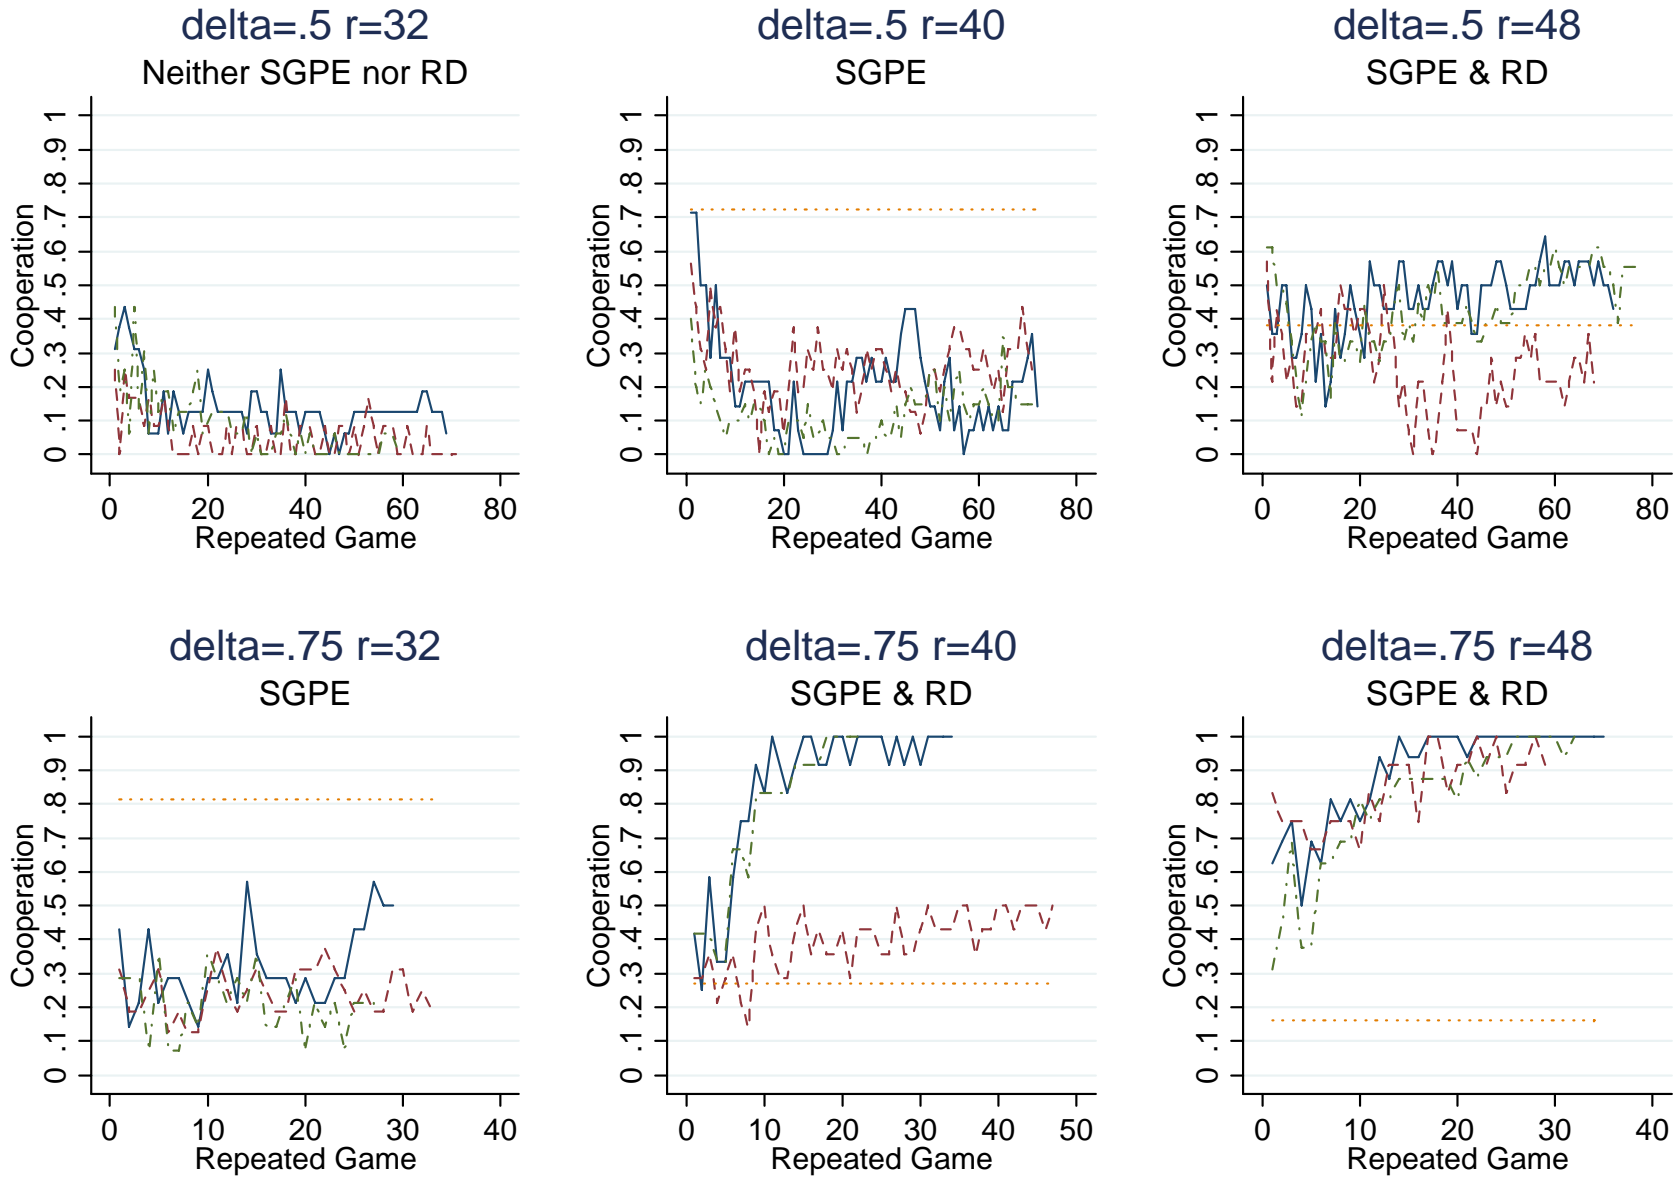
References

- Abreu D. and A. Rubinstein (1988). "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica*, 56(6), 1259-81.
- Aoyagi, M., and G. Fréchette (2003). "Collusion in Repeated Games with Imperfect Public Monitoring," HBS, mimeo.
- Axelrod, R. and W.D. Hamilton (1981). "The Evolution of Cooperation," *Science*, 211(27), 1390-6.
- Bendor, J. and P. Swistak (1997). "The Evolutionary Stability of Cooperation," *American Political Science Review* 91(2), 290-307.
- Bereby Meyer, Y. and A.E. Roth (2006). "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation," *American Economic Review*, 96(4), 1029-42.
- Bernheim, B.D. and D. Ray (1989). "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior* 1(4), 295-326.
- Binmore, K.G. and L. Samuelson (1992). "Evolutionary Stability in Repeated Games Played by Finite Automata," *Journal of Economic Theory*, 57, 278-305.
- Blonski, M. and G. Spagnolo (2001). "Prisoners' Other Dilemma," SSE/EFI Working Paper 437.
- Boyd, R. (1989). "Mistakes Allow Evolutionary Stability in the Repeated Prisoner's Dilemma Game," *Journal of Theoretical Biology*, 136, 47-56.
- Boyd, R. and J.P. Lorberbaum (1987). "No Pure Strategy is Evolutionary Stable in the Repeated Prisoner's Dilemma Game," *Nature*, 327, 58-59.
- Camerer, C. and T.H. Ho (1999). "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica* 67(4), 827-74.
- Charness G., G. R. Fréchette and C.-Z. Qin (2006) "Endogenous Transfers in the Prisoner's Dilemma Game: An Experimental Test Of Cooperation And Coordination," *Games and Economic Behavior*, forthcoming.
- Cheung, Y.W. and D. Friedman (1997). "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46-76.
- Cooper, D.J. (1996). "Supergames Played by Finite Automata with Finite Costs of Complexity in an Evolutionary Setting," *Journal of Economic Theory*, 68, 266-75.
- Cooper, R.W., D.V. DeJong, R. Forsythe, and T.W. Ross (1990). "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80(1), 218-33.
- Cooper, R.W., D.V. DeJong, R. Forsythe, and T.W. Ross (1996). "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games," *Games and Economic Behavior*, 12(2), 187-218.
- Crawford, V.P. (1995). "Adaptative Dynamics in Coordination Games," *Econometrica*, 63(1), 103-43.

- Dal Bó, P. (2005). "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games," *American Economic Review*, 95(5), 1591-604.
- Duffy, J. and J. Ochs (2003). "Cooperative Behavior and the Frequency of Social Interaction," University of Pittsburgh, mimeo.
- Engle-Warnick, J. and Robert S. (2004). "The Evolution of Strategies in a Trust Game," *Journal of Economic Behavior and Organization*, 55, 553-573.
- Engle-Warnick, J. and R.L. Slonim (2006). "Inferring Repeated-Game Strategies From Actions: Evidence From Trust Game Experiments," *Games and Economic Behavior*, 54 (1), 95-114.
- Farrell, J. and E. Maskin (1989). "Renegotiation in Repeated Games," *Games and Economic Behavior* 1(4), 327-360.
- Fréchette, G. R., (2005). "Learning in a Multilateral Bargaining Game," mimeo.
- Fudenberg, D. and D.K. Levine (1998). *The Theory of Learning in Games*, Cambridge: MIT Press.
- Fudenberg, D. and E. Maskin (1990). "Evolution and Cooperation in Noisy Repeated Games," *American Economic Review*, 80(2), 274-279.
- Fudenberg, D. and E. Maskin (1993). "Evolution and Repeated Games," Harvard University, mimeo.
- Harsanyi, J.C. and R. Selten (1988). *A General Theory of Equilibrium Selection in Games*, Cambridge: MIT Press.
- Kandori, M. (1992). "Social Norms and Community Enforcement," *Review of Economic Studies*, 59, 63-80.
- Kandori, M., G.J. Mailath, and R. Rob (1993). "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61(1), 29-56.
- Johnson, P., D.K. Levine and W. Pesendorfer (2001). "Evolution and Information in a Gift-Giving Game," *Journal of Economic Theory* 100 (1), 1-21.
- Kim, Y. (1994). "Evolutionary Stable Strategies in the Repeated Prisoner's Dilemma," *Mathematical Social Science*, 28, 167-97.
- Levine, D.K. and W. Pesendorfer (2007). "The Evolution of Cooperation through Imitation," *Games and Economic Behavior*, 58 (2), 293-315.
- Murnighan, J.K. and A.E. Roth (1983). "Expecting Continued Play in Prisoner's Dilemma Games," *Journal of Conflict Resolution*, 27(2), 279-300.
- Palfrey, T.R. and H. Rosenthal (1994). "Repeated Play, Cooperation and Coordination: An Experimental Study," *Review of Economic Studies*, 61(3), 545-65.
- Pearce, D. (1992). "Repeated games: cooperation and rationality," in J.J. Laffont, ed., *Advances in Economic Theory: Sixth World Congress*, Vol. I. Cambridge University Press.


- Roth A.E. and I. Erev (1995). "Learning in Extensive-Form Games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, 8, 164-212.
- Roth A.E. and J.K. Murnighan (1978). "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma," *Journal of Mathematical Psychology*, 17, 189-97.
- Rubinstein, A. (1986). "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83-96.
- Stahl, D.O. (1991). "The Graph of Prisoners' Dilemma Repeated Game Payoff as a Function of the Discount Factor," *Games and Economic Behavior*, 3, 368-84.
- Van Huyck, J.B., R.C. Battalio, and R. Beil (1990). "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80(1), 234-48.
- Van Huyck, J.B., J.P. Cook, and R.C. Battalio (1997). "Adaptative Behavior and Coordination Failure," *Journal of Economic Behavior and Organization*, 32, 483-503.
- Volij, O. (2002). "In Defense of DEFECT," *Games and Economic Behavior*, 39, 309-21.
- Young, P.H. (1993). "The Evolution of Conventions," *Econometrica*, 61(1), 57-84.

Figure 1: Evolution of Cooperation by Treatment and Session (first rounds)



Note: horizontal dotted lines denote the limit of the basis of attraction of Always Defect versus Grim.

Figure 2: Simulated Evolution of Cooperation Implied by the Learning Estimates



Note: solid lines represent experimental data, dashed lines the average simulated data, dotted lines the 90% interval of simulated data, and vertical lines denote the end of an experimental session.

Figure 3: Simulated Evolution of Cooperation Implied by the Learning Estimates









Figure 4: Distribution of Outcomes in Simulated Sessions

Repeated Game 1 (dashed) vs. Repeated Game 1000 (solid). Vertical lines denotes limit of basins of attraction.

