

The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence

Pedro Dal Bó
Brown University

Guillaume R. Fréchette*
New York University

October 2009

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, Gary Charness, Vincent Crawford, Ernesto Dal Bó, John Duffy, Dan Friedman, Alessandro Lizzeri, Muriel Niederle, Alvin Roth, Ennio Stacchetti, Rani Spiegler, Stefano DellaVigna, Thomas Palfrey and seminar participants at Brown, Caltech, HBS, LSE, NYU, Pittsburgh, UC Berkeley, UCL, University of Iowa, University of Maryland, Université Paris 1, University of Calgary, the Conference of the French Economic Association on Behavioral Economics and Experimental Economics (Lyon 2007), the North American Meeting of the Economic Science Association (Tucson 2007), the Allied Social Science Associations Annual Meeting (New Orleans 2008), the Society for Economic Design 5th Conference on Economic Design (Ann Arbor 2008), and the Third World Congress of the Game Theory Society (Evanston 2008) 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 grants SES-0519045, SES-0721111, and SES-0924780 as well as support from the Center for Experimental Social Science (CESS), and the C.V. Starr Center. Dal Bó gratefully acknowledges the support of NSF via grant SES-0720753. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NSF.

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 on 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 will subjects learn that defection is the best individual action?

We present experimental evidence on the evolution of cooperation in infinitely repeated games. For a given continuation probability and payoffs, each subject participated in between 23 and 77 infinitely repeated games. This allows us to study how cooperation evolves as subjects 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 low levels as it has previously been observed in one-shot prisoners' dilemma games (Yoella Bereby Meyer and Alvin E. Roth 2006). This result indicates 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 even after obtaining significant experience. When cooperation can be supported in equilibrium, subjects may fail to make the most of it. 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.

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

Third, we study whether cooperation rises with experience when it is both an equilibrium action and a risk dominant action (as defined later). If we consider together all sessions for which cooperation is risk dominant, we find that cooperation increase in average as subjects gain experience. However, in several of these sessions cooperation decreases with experience and ends far from full cooperation. 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, in some treatments (where the payoff from cooperation and the probability of future interactions are high enough) subjects coordinate and reach a high level of cooperation. In infinitely repeated games, for cooperation to rise to high levels requires more than just being an equilibrium and risk dominant action.

These results show how difficult it is for experienced subjects to sustain high levels of cooperation. They cast doubt on the common assumption that agents 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 and on what happens after subjects have gained significant experience. Early experiments on infinitely repeated games had 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 J. Keith Murnighan 1978; Murnighan and Roth 1983; Charles A. Holt 1985; Robert M. Feinberg and Thomas A. Husted 1993; Thomas R. Palfrey and Howard Rosenthal 1994).² In fact the impact of repetition on rates of cooperation was rather modest, leading Roth to conclude that the results are equivocal (Roth, 1995). However, more recent experiments (Pedro Dal Bó 2005; Masaki Aoyagi and Guillaume Fréchette 2009; and John Duffy and Jack Ochs 2009) yield much more positive results in terms of the ability of subjects to support cooperation in infinitely repeated games. Dal Bó (2005) compares infinitely

² All of these papers used games with a randomly determined length. That is, after each play of the stage game, there would be one more play of the stage game with a fixed probability, or that game ends. That probability is known to all participants. This method for inducing infinitely repeated games in the laboratory was introduced by Roth and Murnighan (1978) and such games are sometimes referred to as indefinitely repeated games or randomly terminated games.

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. Aoyagi and Fréchette (2009) 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 (2009) 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 fixed matching but not under random matching. In a recent paper, Matthias Blonski, Peter Ockenfels and Giancarlo Spagnolo (2007) show that changes in the "sucker payoff" (the payoff from cooperating when the other defects) affect the level of cooperation. As this payoff does not enter the equilibrium conditions for mutual cooperation, their results show these conditions are not enough to understand the determinants of cooperation. In this direction, they provide evidence in favor of risk dominance as an alternative condition.³ 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 reduce the multiplicity of equilibria in infinitely repeated games. Robert Axelrod and William D. Hamilton (1981), Robert Boyd and Jeffrey P. Lorberbaum (1987), Boyd (1989), Yong-Gwan Kim (1994), and Jonathan Bendor and Piotr Swistak (1997) apply the concept of evolutionary stable strategies (ESS) by John M. Smith (1982) to infinitely repeated games with varied results regarding the selection of equilibria. There is also a literature that appeals to bounded rationality. Ariel Rubinstein (1986) and Dilip 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. Ken G. Binmore and Larry Samuelson (1992), David J. Cooper (1996) and Oscar Volij (2002) apply evolutionary refinements to

³ For other papers on infinitely repeated games see Richard D. McKelvey and Palfrey (1995), Jim Engle-Warnick, William J. McCausland and John H. Miller (2004), Engle-Warnick and Robert L. Slonim (2006a and b), Hans-Theo Normann and Brian Wallace (2006), Gabriele Camera and Marco Casari (2009), Dale O. Stahl (2007), Anna Dreber et al. (2008) and Timothy Cason and Vai-Lam Mui (2008).

infinitely repeated games played by finite automata and find that the set of possible payoffs depends crucially on the definition of ESS and the way costs of complexity are modeled (also see Fudenberg and Maskin, 1990 and 1993). In contrast, Volij (2002) shows that always defecting is the unique stochastically stable strategy (Michihiro Kandori et al., 1993, and Peyton H. Young, 1993) in games with finite automata. Philip Johnson, David K. Levine and Wolfgang 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 that exceeds what is required by sub-game perfection (see also Levine and Pesendorfer, 2007). Finally, Blonski and Spagnolo (2001) appeal to the concept of risk dominance as an equilibrium selection criteria in infinitely repeated games.

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 theories. Theories in which subjects always coordinate on defection, even when they are infinitely patient, and theories in which 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.

I. Experimental design

We induce 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 points (one point equals to \$0.006) and where the payoff to cooperation takes one of three possible values: $R=32, 40$ and 48 .

Table 1: Stage Game Payoffs

	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. In each session, a set of subjects participated anonymously through computers in a sequence of infinitely repeated prisoners' dilemma games. Subjects were randomly re-matched with another subject after the end of each repeated game.⁴ In each session subjects participate in as many repeated games as was possible such that the first repeated game to end after 50 minutes of play marks the end of the session. The probability of continuation and the payoff matrix was the same for all repeated games in a session, that is, there was one treatment per session. We conducted three sessions per treatment. The instructions for one of the sessions are in the appendix.

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

II. 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 as in Stahl (1991). Table 2 indicates those treatments under which cooperation can be supported as a sub-game perfect equilibrium action.⁵

Table 2: Cooperation as Equilibrium (SGPE) and Risk Dominant (RD) Action

	R=32	R=40	R=48
$\delta = 1/2$	Neither SGPE or RD	SGPE	SGPE and RD
$\delta = 3/4$	SGPE	SGPE and RD	SGPE and RD

⁴ Random matching allows for a larger number of repeated games in a session than alternative matching protocols like turnpike protocols. While the probability of a pair of subjects interacting together in more than one repeated game is high this is not likely to be a problem for several reasons. First, our results in section III 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, Dal Bó (2005) uses a turnpike protocol with results consistent with other studies that have used random matching protocols. Third, with sessions with similar number of subjects Duffy and Ochs (2009) find that random matching is not enough to develop cooperative strategies across matches.

⁵ More precisely, the critical value of δ over which cooperation can be supported in equilibrium is 0.72 under R=32, 0.4 under R=40, and 0.08 under R=48. While this categorization of treatments is done assuming risk neutrality of the subjects, the results of the paper are robust to considering risk aversion levels of the magnitude typically observed in experiments.

Under $\delta=1/2$ & $R=32$, defection is the only possible equilibrium action and we expect that as subjects gain experience, the levels of cooperation will 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 under monetary payoffs leading to the following question.⁶

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

In the remaining five treatments cooperation can be supported in equilibrium.⁷ However, there is a multiplicity of equilibria under these treatments; defection remains a possible equilibrium action. We study next the issue of equilibrium selection in these treatments. On the one hand, we may expect subjects to learn to coordinate on cooperative equilibria as these equilibria are Pareto efficient. In this case we might expect cooperation to increase with experience and reach levels close to 100%. On the other hand, it may not be realistic to assume that subjects will always learn to coordinate on the Pareto efficient equilibrium. As shown in the literature on coordination games, subjects may fail to coordinate on the Pareto efficient equilibrium when 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).⁸ This leads us to the following question.

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

Finally, we apply the concept of risk-dominance as an equilibrium selection criterion. Risk-dominance was introduced by John C. Harsanyi and Reinhard Selten (1988) and concerns the pairwise comparison between Nash equilibria. In 2×2 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

⁶ See for example the literature on centipede games and the literature on zero-sum games (McKelvey and Palfrey 1992, and James N. Brown and Robert W. Rosenthal 1990).

⁷ 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 payoffs can be supported in equilibrium. In addition, mutual 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 B. Douglas Bernheim and Debraj Ray 1989, and Joseph Farrell and Maskin 1989; also see David Pearce 1992, and references therein, for a review of the different renegotiation proofness concepts).

⁸ This is not to say that subjects never coordinate on the efficient outcome, but rather that they sometimes do not. Gary Charness, Fr chet and Cheng-Zhong Qin (2007) provide an example where subjects often do coordinate on the efficient outcome in a two-stage game (prisoner's dilemma with transfers). Charness (2000) finds that subjects select the efficient outcome in stag hunt games when one-way pre-play communication is allowed.

player. The focus on risk dominance is inspired by the fact that in one-shot games it has shown to be a good predictor of behavior (Russell W. Cooper et al 1990, John B. Van Huyck et al 1990, and Cooper et al 1992).

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. Given the difficulties of applying the concept of risk-dominance to the whole set of possible strategies in infinitely repeated games we follow Roger B. Myerson (1991) in focusing only on a pairwise comparison of a strategy that supports 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.⁹ In 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. We say that cooperation is risk dominant if playing G is the best response to the other player choosing G or AD with equal probabilities. Table 2 shows the treatments under which cooperation is risk-dominant.¹⁰ 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?

III. Main Experimental Results

The 18 experimental sessions were conducted between July 2005 and March 2006. A total of 266 NYU undergraduates participated in the experiment, with an average

⁹ 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.

¹⁰ More precisely, the critical value of δ over which cooperation is risk-dominant is 0.82 under $R=32$, 0.61 under $R=40$, and 0.39 under $R=48$. Note that in the definition we could change the strategy G for any other cooperative strategy that behaves like G when facing AD, for example, tit-for-tat (TFT), and this would not affect the classification of treatments based on risk dominance.

of 14.78 subjects per session, a maximum of 20 and a minimum of 12. 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.

A. General description of behavior

Before answering the three questions raised in the previous section we provide a general description of the observed behavior. The top panel in Table 3 shows cooperation rates by treatment for the first repeated game, on the left for the first round and on the right for all rounds.¹¹ Looking separately at first rounds is important since different repeated games may result in a different number of rounds and the percentage of cooperation may vary across rounds. Cooperation tends to be higher in treatments under which cooperation can be supported as an equilibrium action than when it cannot be supported, but this difference is only significant for first rounds (p-values of 0.023 and 0.151 for first rounds and all rounds respectively). Moreover, it is not the case that an increase in the probability of continuation always results in an increase in cooperation (for instance compare the two treatments with $R=40$) and increases in the payoff from cooperation do not always have a significant effect on cooperation.

Table 3: Percentage of Cooperation by Treatment

First Repeated Game											
First Round						All Rounds					
$\delta \setminus R$	32		40		48	$\delta \setminus R$	32		40		48
$1/2$	34.09	<*	54.00	<	56.52	$1/2$	28.33	<*	39.80	<	41.38
	=		v		^		v***		v*		^
$3/4$	34.09	<	36.84	<*	56.82	$3/4$	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
$1/2$	9.81	<**	18.72	<***	38.97	$1/2$	9.82	<*	17.98	<***	35.29
	^***		^***		^***		^**		^***		^***
$3/4$	25.61	<**	61.10	<*	85.07	$3/4$	20.25	<***	58.71	<	76.42

Note: * significance at 10%, ** at 5% and *** at 1%. See footnote 11 for details about procedures.

¹¹ Throughout the paper (unless specified otherwise) statistical significance is assessed using probit regressions with an indicator variable for one of the two relevant categories. Standard errors are clustered at the level of the session.

The bottom panel of Table 3 shows cooperation rates for all repeated games. Here, in contrast with the first repeated game, increases in the probability of continuation and the payoff of cooperation result in increases in cooperation (the differences are all significant at the 10% level with one exception). Moreover, cooperation is significantly greater in treatments under which cooperation can be supported as an equilibrium action than when it cannot be supported (p-values < 0.001 for both first rounds and all rounds).

These differences between the first and all matches in the comparison of behavior across treatments suggest that experience affects how subjects play in repeated games. The next sections focus on how subjects modify their behavior as they gain experience.

B. 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 4 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.¹² To compare inexperienced versus experienced play we compare behavior in the first 10 interactions with those in interactions 111 to 120.¹³

Cooperation in the treatment where it is not an equilibrium 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 < 0.001). 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 (e.g. 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 (fifth column in Table 4).

¹² 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.

¹³ 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. The results do not hinge on focusing on these interactions.

Table 4: Percentage of Cooperation by Equilibrium Condition and Risk Dominance

Repeated Game Begins in Interaction	First Rounds Cooperation is				All Rounds Cooperation is			
	Not SGPE	SGPE			Not SGPE	SGPE		
		All	Not RD	RD		All	Not RD	RD
1-10	28.57	39.11	31.43	46.53	21.00	34.42	23.56	42.11
11-20	13.04	28.54	20.60	36.26	12.91	27.19	18.10	35.09
21-30	12.23	31.01	14.86	44.34	11.97	33.61	13.48	45.36
31-40	10.61	36.04	14.01	51.83	10.51	38.64	14.63	52.72
41-50	10.20	34.88	14.21	53.99	7.85	34.98	13.81	53.09
51-60	9.75	41.47	18.51	57.47	6.54	39.85	16.32	61.30
61-70	7.14	37.89	17.54	48.98	8.09	40.02	19.21	54.44
71-80	5.65	36.86	20.32	50.00	4.48	39.73	19.10	55.99
81-90	4.72	38.60	20.57	58.42	6.20	44.39	20.75	60.89
91-100	6.11	40.91	22.01	54.88	7.91	47.11	19.28	66.45
101-110	6.64	45.38	17.93	67.62	11.99	46.12	19.50	66.92
111-120	5.50	49.77	22.46	70.61	6.45	55.88	22.60	73.86
121-130	5.77	45.95	21.03	62.05	11.11	43.31	21.99	59.60
131-140	8.33	47.43	30.70	59.49	9.17	42.99	26.23	61.40
141-		46.32	23.86	65.69		47.83	16.57	76.82

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 $\delta=1/2$ & $R=32$. It is clear from this graph that cooperation decreases with experience in all three sessions.

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

The second column in Table 4 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 between 111 to 120 interactions cooperation increased to 50% (p-value of the difference is 0.11). In the sixth column in Table 4 we observe a similar evolution of cooperation from all the rounds in the repeated games (p-value = 0.004). In addition, cooperation rates differ significantly depending on whether cooperation can be supported in equilibrium for the first 10 interactions (p-value = 0.083 for round 1 only and 0.028 for

all rounds) and even more so for the repeated games that start between 111-120 interactions (p -value < 0.001) for both first round and all rounds.

These results support the idea that subjects improve their ability to support cooperation as they gain experience, but only slightly. That they are still very far from coordinating on the efficient outcome is evident by the low levels of cooperation in session where cooperation can be supported in equilibrium. Note that in Figure 1, for sessions where cooperation can be supported in equilibrium, 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.

D. Do subjects learn to cooperate when it is risk-dominant?

In this section we examine whether subgame perfection combined with risk dominance is sufficient for subjects to learn to make the most of the opportunities for cooperation. The third and fifth columns in Table 4 show 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 separated by whether cooperation is risk dominant or not. Cooperation decreases with experience when it is not risk-dominant, but 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 it reached 70.6% (p -value of the difference = 0.016). We observe a similar evolution for all rounds (p -value = 0.001). 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 between 111-120 interactions for both first and all rounds (p -values < 0.05). Nonetheless, the cooperation rate when cooperation is risk-dominant is still far away from full cooperation even when subjects have gained great experience.

We reach an even more nuanced conclusion if we study these treatments at the session level. 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 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 to cooperation are high enough, it is possible for subjects to make the most of the opportunity to cooperate.

IV. Discussion

Our results have a number of interesting implications. First, when playing repeated games, the amount of experience is a critical determinant of outcomes. With experience, subjects reach very low levels of cooperation when it cannot be supported in equilibrium while they may reach very high levels of cooperation when it can be supported in equilibrium. Moreover, the impact on behavior of increasing the continuation probability or the payoff from cooperation is important both in magnitude and statistical significance for experienced subjects, but not for inexperienced ones.

Second, in contrast to the results from one-shot coordination game experiments, in infinitely repeated games subjects might not select an equilibrium that is both payoff dominant and risk dominant. In three of our six treatments, cooperation can be supported as part of a payoff dominant and risk dominant equilibrium, yet it does not always emerge as the selected option.

Overall, this evidence suggests that while being an equilibrium action may be a necessary condition for cooperation to arise with experience, it is not sufficient. Moreover, being risk dominant is not sufficient either. However, subjects do reach high levels of cooperation under very favorable conditions. If subgame perfection and risk-dominance do not fully account for the determinants of end of session behavior, then what other factors matter?

We start by considering here the role of the fundamental parameters of the game in explaining the evolution of cooperation. In the environments considered here, if a subject is contemplating whether to play a defecting strategy (like AD) or a cooperative strategy (like TFT or G), that subject needs to determine which of these two strategies is

the most profitable in expectation based on the parameters of the game and his beliefs about the probability that his partner will play AD or the cooperative strategy. In Figure 1 the horizontal dotted line represents the belief of the cooperative strategy that would leave the subject indifferent between the two strategies for each treatment.¹⁴ If the belief that the other subject will choose the cooperative strategy falls below the dotted line the subject will maximize his expected payoff by playing AD and choosing the cooperative strategy will maximize the expected payoff otherwise. Therefore, it can be expected that the larger the basin of attraction of AD (the larger the set of beliefs that make AD optimal) the less likely a subject will choose to cooperate.¹⁵

To assess whether the size of the basin of attraction of AD can explain subjects behavior as they gain experience, we regress subjects behavior in the first round of the last repeated game on the size of the basin of attraction of AD (see Table 5). We include the size of the basin of attraction and its square and find that the size of the basin of attraction of AD is negatively correlated with the likelihood of cooperation and that its impact diminishes as the basin becomes larger. We also include the difference between the average realized length of the previous repeated games and its expected length. We do this because the realized length of previous repeated game may affect behavior (Engle-Warnick and Slonim 2006a). We find that this variable is highly significant.¹⁶

Table 5: Cooperation in the First Round of the Last Repeated Game (Probit)

Size of Basin Of AD	-7.276***	[2.432]
Size of Basin Square	3.153	[2.061]
Extra Length Of Repeated Games	0.942***	[0.213]
Constant	2.553***	[0.628]
Observations	266	

Clustered standard errors in brackets

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

¹⁴ G is never a best response for $\delta=1/2$ and $R=32$. 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 .

¹⁵ For a related argument see Myerson (1991) section 7.11. Also note that it is not the case that the size of the basins of attraction simply captures how easy or difficult it is to support cooperation in equilibrium in a specific game. For instance, if we consider the difference between the probability of continuation and the minimum probability required to support cooperation we find that this difference does not necessarily order the treatments as the size of the basin of attraction does. For example AD has a larger basin of attraction in the treatment with $\delta=1/2$ & $R=48$ than in the treatment with $\delta=3/4$ & $R=40$, but the difference between the probability of continuation and the minimum required for cooperation is greater in the former.

¹⁶ Adding indicator variables to control for whether cooperation can be supported in equilibrium and whether it is risk dominant does not affect that statistical significance (or sign) of the regressors in Table 5 and neither of these regressors are statistically significant. Note that cooperation is risk dominant whenever the basin of attraction of AD is less than $1/2$.

We conclude from this that the fundamentals of the game significantly affect the evolution of cooperation: games with a large basin of attraction toward cooperation are more likely to generate cooperation than those with large basin of attraction toward defection. In addition, the significant impact of the realized length of the repeated games on behavior suggests that the subjects' own experience in the session is also a key factor.

To further explore the role of previous experience on behavior, we study how cooperation is affected by the behavior of partners and the number of rounds in the previous repeated game. Table 6 presents results from probit estimations in which all matches but the first are considered. In all treatments, a subject who was matched with someone who played cooperate in the first round is more likely to start by cooperating in the following match than someone who was matched with a player who first defected. That behavior, although it could be the result of a host of psychological phenomena, is certainly consistent with updating behavior. That is, if subjects have an estimate of the fraction of the population playing a strategy that starts by cooperating they update their estimate upward after observing a sample consistent with that strategy.

Table 6: Effect of Past Observations on Round 1 Cooperation (Probit)

	$\delta = \frac{1}{2}$			$\delta = \frac{3}{4}$		
	R = 32	R = 40	R = 48	R = 32	R = 40	R = 48
Partner Cooperated in Round 1 of Previous Match	0.424*** [0.11]	0.276*** [0.093]	0.351*** [0.11]	0.260** [0.124]	0.912*** [0.293]	0.581** [0.276]
Number of Rounds In Previous Match	0.018* [0.011]	0.052*** [0.012]	0.055*** [0.015]	0.006*** [0.001]	0.048*** [0.013]	0.002 [0.017]
Subject Cooperated in Round 1 of Match 1	1.113*** [0.167]	0.473*** [0.144]	0.936*** [0.141]	0.169 [0.262]	1.182*** [0.342]	-0.042 [0.206]
Constant	-1.963*** [0.052]	-1.346*** [0.245]	-1.097*** [0.218]	-0.822*** [0.075]	-0.764*** [0.251]	0.626*** [0.228]
Observations	2840	3534	3300	1268	1304	1376

Clustered standard errors in brackets.

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

Subjects are more likely to start by cooperating following a long match than after a short match. This can be seen by the positive and statistically significant coefficient estimate for the length of the previous repeated games consistently with the findings by Engle-Warnick and Slonim (200a) for trust games. This suggests that subjects, while

correct on average (see Dal Bó 2005), may have some uncertainty about the expected length of the game and they update their estimate based on what they observe.

Taken together, these last three observations (the correlation between cooperation and (1) the size of the basin of attraction, (2) the partner's cooperation in the previous repeated game, and (3) the length of the previous repeated game) suggest that subjects' behavior is influenced by their estimate of the value of "investing" in cooperation. The longer the potential interaction is, the more valuable it is. The larger the basin of attraction of cooperation and the more subjects seem to be cooperating, the higher the expected value of starting with cooperation. This also suggests that even in a given environment, details of the history may matter.¹⁷

Points 1 and 2 above suggest that a learning model might well describe some aspects of the behavior we observe in this experiment. But before describing the learning model, we will first study the strategies used by the subjects, which will allow us to simplify the learning. Given that there are an infinite number of strategies but the data is necessarily finite, it is impossible in principle to identify the strategies used by the subjects. We circumvent this problem by restricting our attention to a small set of ex-ante relevant candidates due to their importance in the theoretical literature: Always Defect (AD), Always Cooperate (AC), Grim (G), Tit for Tat (TFT), Win Stay Lose Shift (WSLS) and a trigger strategy with two periods of punishment (T2).¹⁸ The importance of each strategy is estimated by maximum likelihood assuming that subjects have a given probability of choosing one of the six strategies and that they do not change strategies from repeated game to repeated game. We focus on repeated games that started after 110 interactions, that is once behavior is more likely to have stabilized. We assume subjects may make mistakes and choose an action that is not recommended by the strategy. A

¹⁷ As pointed out by a referee, multiple forms of learning can be simultaneously at play: 1) learning to understand the infinitely repeated game which is documented by the finding that subjects' behavior is affected by the length of the games; 2) learning about the properties of the group, which is documented in the finding that the behavior of the partner in the previous match affects the subjects' behavior; 3) Learning about your own risk preferences, which we leave for future research - see Carlos Oyarzun and Rajiv Sarin (2007) for theoretical work on the topic; and 4) boundedly rational learning which is explored in more details later on in the paper where we estimate a belief based learning model.

¹⁸ WSLS is a strategy that starts cooperating and then conditions behavior only on the outcome of the previous round. If either both cooperated or neither cooperate, then WSLS cooperates, otherwise it defects. T2 starts cooperating and a defection by the other triggers two rounds of defection after which the strategy goes back to cooperation. These two strategies are cooperative strategies with punishments of limited length.

detailed description of the estimation procedure is in the online appendix. The estimates of the proportions for each strategy are presented in Table 7 (with the coefficient for T2 being implied by the fact that the proportions must sum to one and gamma captures the amount of noise – as gamma goes to infinity response become purely random).

Table 7 reveals some interesting patterns. First, as expected, cooperative strategies describe the data better in treatments where more cooperative behavior is observed. Second, the cooperative strategy that is most often identified is TFT. While G explains some of the data, its proportion is not statistically significant. Finally, only considering AD and TFT can account for 80% of all the data in matches that start after interaction 110. Moreover, for all treatments it cannot be rejected at the 5% level that subjects only use AD or TFT. At the 10% level, this can only be rejected under two treatments ($\delta=1/2$ & $R=48$ and $\delta=3/4$ & $R=40$). Therefore, in what follows we focus on these two strategies.

Table 7: Estimation of Strategies Used

	$\delta=1/2$			$\delta=3/4$		
	R = 32	R = 40	R = 48	R = 32	R = 40	R = 48
AD	0.920*** (0.085)	0.783*** (0.074)	0.533*** (0.109)	0.648*** (0.119)	0.109 (0.096)	0.000 (0.000)
AC	0.000 (0.000)	0.078 (0.059)	0.072 (0.046)	0.000 (0.000)	0.296** (0.123)	0.079 (0.085)
G	0.000 (0.000)	0.040 (0.040)	0.000 (0.000)	0.000 (0.024)	0.267 (0.202)	0.116 (0.195)
TFT	0.080 (0.085)	0.098 (0.070)	0.376*** (0.112)	0.352*** (0.115)	0.327* (0.186)	0.561*** (0.185)
WSLS	0.000 (0.000)	0.000 (0.007)	0.019 (0.026)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
T2	0.000	0.000	0.000	0.000	0.000	0.244
Gamma	0.362*** (0.098)	0.541 (1.077)	0.428*** (0.061)	0.447*** (0.053)	0.435*** (0.126)	0.287*** (0.061)

Bootstrapped standard errors in parenthesis.

We now study how well a learning model does in fitting the evolution of cooperation in our experiments. Then, we use the estimates of this model to perform simulations showing that the results of this paper would also hold in the very long run.

We simplify the set of possible strategies to only two strategies: AD and TFT. 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). We assume that subjects start with beliefs about the probability their partner uses either AD or TFT and they modify these beliefs based on their observation of their partners' behavior. The updating is allowed to go from Cournot to fictitious play, as in Vincent Crawford (1995) and Yin W. Cheung and Daniel Friedman (1997).¹⁹ We abstract from the complexities of the repeated games by reducing it to the choice in round 1: defect corresponds to AD and cooperate corresponds to TFT. Given beliefs, subjects are modeled as random utility maximizers given the expected return from each choice. We allow for the noise in decision making to decrease with experience. A detailed description of the learning model is in the online appendix.

The estimates are obtained via maximum likelihood estimation for each subject separately.²⁰ 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 the online appendix.

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).²¹ 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

¹⁹ For other relevant models of learning see Roth and Ido Erev (1995) and Colin Camerer and Teck-Hua Ho (1999).

²⁰ 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 (2009) discusses issues and solutions related to pooling data across subjects in estimating learning models and more specifically with respect to hypothesis testing.

²¹ The results from the simulations are robust to varying the number of subjects per simulated session.

line).²² 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 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.

Given that the learning model fits the data well, the simulations to a longer range of repeated games suggest how cooperation would evolve in the long run. The evolution of cooperation in the long run is consistent with what is observed in 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 it 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 it 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 around 40% and the 90% interval sometimes includes full defection

²² In Figure 2 the repeated game axis is displayed in log scale to facilitate the comparison between experimental and simulated data.

but it never includes full cooperation. Thus, in treatments in which cooperation is risk dominant and with a large number of repeated games for subjects to gain experience, full cooperation may fail to arise.

V. Conclusions

The series of experiments presented in this paper shed light on how cooperation evolves as subjects gain experience. The evidence presented suggests that as subjects gain experience they may reassess the gains of attempting to establish a cooperative relationship with a partner and they may modify their behavior accordingly. The subjects' perception of these gains seems to depend both on the given fundamentals of the game affecting incentives (well summarized by the size of the basin of attraction of AD) and also on changing (and to some degree random) elements like the realized length of the previous repeated game and the behavior of their previous partner.

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.²³ 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 equilibrium selection theories that select inefficient outcomes even when players are arbitrarily patient (e.g. Rubinstein 1986 and Volij 2002) and those selecting efficient outcomes whenever they are a possible equilibrium (e.g. Tirole 1988). We hope the evidence provided here will guide future theoretical attempts to study equilibrium selection in infinitely repeated games.

²³ See for example the discussion on Jean Tirole (1988), page 253, regarding equilibrium selection on tacit collusion models: “The multiplicity of equilibria is an embarrassment of riches. We must have a reasonable and systematic theory of how firms coordinate on a particular equilibrium if we want the theory to be predictive and allow comparative statics. One natural method is to assume that firms coordinate on an equilibrium that yields a Pareto-optimal point in the set of the firms' equilibrium profits.” This approach is followed in many papers including Dal Bó (2007).

References

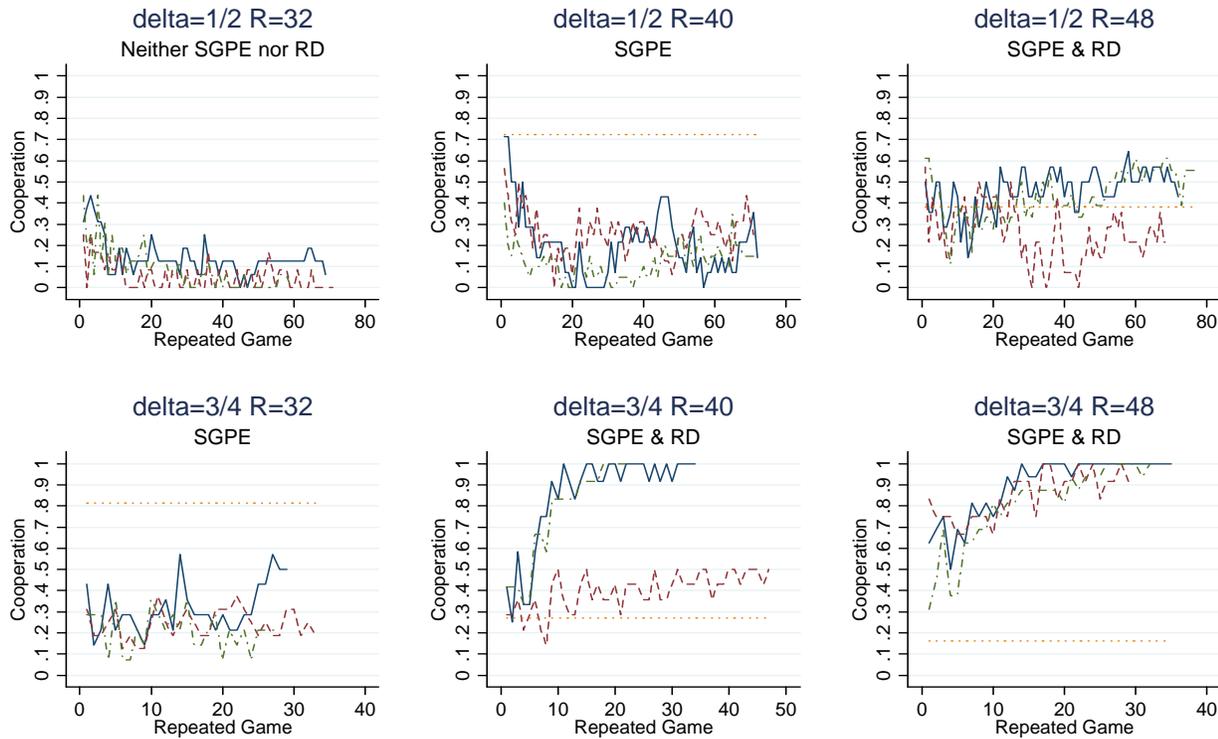
- Abreu, Dilip, and Ariel Rubinstein.** 1988. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata." *Econometrica*, 56(6): 1259-1281.
- Aoyagi, Masaki, and Guillaume R. Fréchette.** 2009. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." *Journal of Economic Theory*, 144(3): 1135-1165.
- Axelrod, Robert, and William D. Hamilton.** 1981. "The Evolution of Cooperation." *Science*, 211(27): 1390-1396.
- Bendor, Jonathan, and Piotr Swistak.** 1997. "The Evolutionary Stability of Cooperation." *American Political Science Review*, 91(2): 290-307.
- Bereby Meyer, Yoella, and Alvin E. Roth.** 2006. "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation." *American Economic Review*, 96(4): 1029-1042.
- Bernheim, B. Douglas, and Debraj Ray.** 1989. "Collective Dynamic Consistency in Repeated Games." *Games and Economic Behavior*, 1(4): 295-326.
- Binmore, Ken G., and Larry Samuelson.** 1992. "Evolutionary Stability in Repeated Games Played by Finite Automata." *Journal of Economic Theory*, 57(2): 278-305.
- Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo.** 2007. "Co-operation in Infinitely Repeated Games: Extending Theory and Experimental Evidence." Mimeo.
- Blonski, Matthias and Giancarlo Spagnolo.** 2001. "Prisoners' Other Dilemma." SSE/EFI Working Paper 437.
- Boyd, Robert.** 1989. "Mistakes Allow Evolutionary Stability in the Repeated Prisoner's Dilemma Game." *Journal of Theoretical Biology*, 136(1): 47-56.
- Boyd, Robert, and Jeffrey P. Lorberbaum.** 1987. "No Pure Strategy is Evolutionary Stable in the Repeated Prisoner's Dilemma Game." *Nature*, 327(6117): 58-59.
- Brown, James N., and Robert W. Rosenthal.** 1990. "Testing the Minimax Hypothesis: A Re-Examination of O'Neill's Game Experiment." *Econometrica*, 58(5): 1065-1081.
- Camera, Gabriele and Marco Casari.** 2009. "Cooperation among strangers under the shadow of the future." *American Economic Review* 99(3), 979-1005.
- Camerer, Colin, and Teck-Hua Ho.** 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica*, 67(4): 827-874.
- Cason, Timothy, and Vai-Lam Mui.** 2008. "Coordinating Collective Resistance through Communication and Repeated Interaction." Purdue University. Mimeo.
- Charness, Gary.** 2000. "Self-serving Cheap Talk and Credibility: A Test of Aumann's Conjecture." *Games and Economic Behavior*, 33(2): 177-194.
- Charness, Gary, Guillaume R. Fréchette and Cheng-Zhong Qin.** 2007. "Endogenous Transfers in the Prisoner's Dilemma Game: An Experimental Test Of Cooperation And Coordination." *Games and Economic Behavior*, 60(2): 287-306.
- Cheung, Yin Wong, and Daniel Friedman.** 1997. "Individual Learning in Normal Form Games: Some Laboratory Results." *Games and Economic Behavior*, 19(1): 46-76.

- Cooper, David J.** 1996. "Supergames Played by Finite Automata with Finite Costs of Complexity in an Evolutionary Setting." *Journal of Economics Theory*, 68(1): 266-275.
- Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross.** 1990. "Selection Criteria in Coordination Games: Some Experimental Results." *American Economic Review*, 80(1): 218-233.
- Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross.** 1996. "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games." *Games and Economic Behavior*, 12(2): 187-218.
- Crawford, Vincent P.** 1995. "Adaptive Dynamics in Coordination Games." *Econometrica*, 63(1): 103-143.
- Dal Bó, Pedro.** 2005. "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games." *American Economic Review*, 95(5): 1591-1604.
- Dal Bó, Pedro.** 2007. "Tacit collusion under interest rate fluctuations." *RAND Journal of Economics*, 38(2): 533-540.
- Dreber, Anna, David G. Rand, Drew Fudenberg, and Martin A. Nowak.** 2008. "Winners don't punish." *Nature*, 452(7185): 348-351.
- Duffy, John, and Jack Ochs.** 2009. "Cooperative Behavior and the Frequency of Social Interaction." *Games and Economic Behavior*, 66(2), 785-812.
- Engle-Warnick, Jim, William J. McCausland, and John H. Miller.** 2004. "The Ghost in the Machine: Inferring Machine-Based Strategies from Observed Behavior." Universite de Montreal. Mimeo.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2004. "The Evolution of Strategies in a Trust Game." *Journal of Economic Behavior and Organization*, 55(4): 553-573.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006a. "Learning to trust in indefinitely repeated games." *Games and Economic Behavior*, 54(1): 95-114.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006b. "Inferring Repeated-Game Strategies From Actions: Evidence From Trust Game Experiments." *Economic Theory*, 54(1): 95-114.
- Farrell, Joseph, and Eric Maskin.** 1989. "Renegotiation in Repeated Games." *Games and Economic Behavior*, 1(4): 327-360.
- Feinberg, Robert M. and Husted, Thomas A.** 1993. "An Experimental Test of Discount-Rate Effects on Collusive Behavior in Duopoly Markets." *Journal of Industrial Economics* 41(2):153-60.
- Fréchette, Guillaume R.** 2009. "Learning in a Multilateral Bargaining Experiment." *Journal of Econometrics* 153(2): 183-195.
- Fudenberg, Drew, and David K. Levine.** 1998. *The Theory of Learning in Games*. Cambridge: MIT Press.
- Fudenberg, Drew, and Eric Maskin.** 1990. "Evolution and Cooperation in Noisy Repeated Games." *American Economic Review*, 80(2): 274-279.

- Fudenberg, Drew, and Eric Maskin.** 1993. "Evolution and Repeated Games." Harvard University. Mimeo.
- Harsanyi, John C. and Reinhard Selten.** 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge: MIT Press.
- Holt, Charles A.** "An Experimental Test of the Consistent-Conjectures Hypothesis." *American Economic Review*, 1985, 75(3): 314–325.
- Kandori, Michihiro.** 1992. "Social Norms and Community Enforcement." *Review of Economic Studies*, 59(1): 63-80.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob.** 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica*, 61(1): 29-56.
- Johnson, Phillip, David K. Levine and Wolfgang Pesendorfer.** 2001. "Evolution and Information in a Gift-Giving Game." *Journal of Economic Theory*, 100(1): 1-21.
- Kim, Yong-Gwan.** 1994. "Evolutionary Stable Strategies in the Repeated Prisoner's Dilemma." *Mathematical Social Science*, 28: 167-197.
- Levine, David K., and Wolfgang Pesendorfer.** 2007. "The Evolution of Cooperation through Imitation." *Games and Economic Behavior*, 58(2): 293-315.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1992. "An Experimental Study of the Centipede Game." *Econometrica*, 60(4): 803-836.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "The holdout game: An experimental study of an infinitely repeated game with two-sided incomplete information." In *Social choice, welfare and ethics: Proceedings of the Eight International Symposium in Economic Theory and Econometrics*, eds. William A. Barnett, Herve Moulin, Maurice Salles, and Norman J. Schofield, 321-352. Cambridge: Cambridge University Press.
- Murnighan, J. Keith, and Alvin E. Roth.** 1983. "Expecting Continued Play in Prisoner's Dilemma Games." *Journal of Conflict Resolution*, 27(2): 279-300.
- Myerson, Roger B.** 1991. *Game Theory: Analysis of Conflict*. Cambridge: Harvard University Press.
- Normann, Hans-Theo, and Brian Wallace.** 2006. "The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment." Royal Holloway. Mimeo.
- Oyarzun, Carlos, and Rajiv Sarin.** 2007. "Learning and Risk Aversion." Texas A&M. Mimeo.
- Palfrey, Thomas R., and Howard Rosenthal.** 1994. "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies*, 61(3): 545-565.
- Pearce, David.** 1992. "Repeated games: cooperation and rationality." In *Advances in Economic Theory: Sixth World Congress*, Vol. 1, ed. Jean-Jacques Laffont, 132-174. Cambridge University Press.
- Roth, Alvin E.** 1995 "Introduction to Experimental Economics." In *The handbook of experimental economics*, John H. Kagel and Alvin E. Roth, eds., 3–109. Princeton: Princeton University Press.

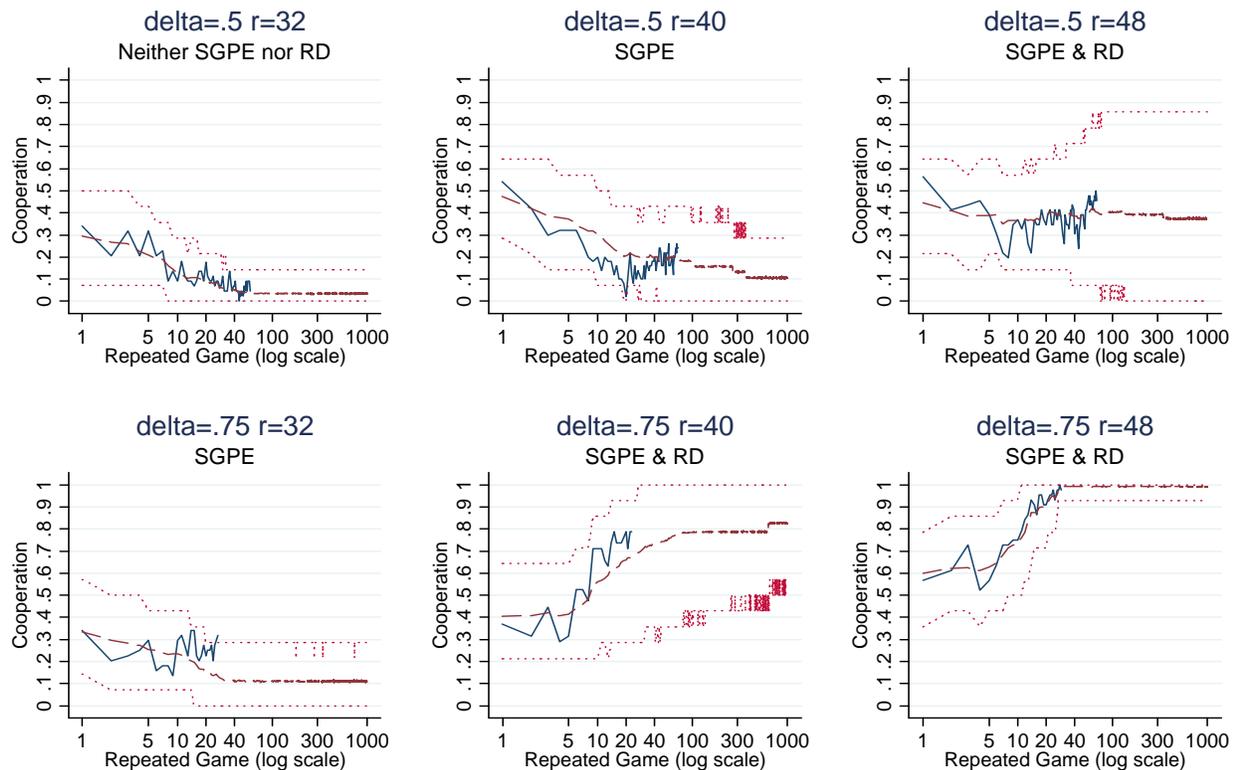
- Roth, Alvin E., and Ido Erev.** 1995. "Learning in Extensive-Form Games: Experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior*, 8(1): 164-212.
- Roth Alvin E., and J. Keith Murnighan.** 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology*, 17(2): 189-197.
- Rubinstein, Ariel.** 1986. "Finite Automata Play the Repeated Prisoner's Dilemma." *Journal of Economic Theory*, 39(1): 83-96.
- Smith, John M.** 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Stahl, Dale O.** 1991. "The Graph of Prisoners' Dilemma Supergame Payoffs as a Function of the Discount Factor." *Games and Economic Behavior*, 3(3): 368-384.
- Stahl, Dale O.** 2007. "An Experimental Test of the Efficacy of Simple Reputation Mechanisms to Solve Social Dilemmas." University of Texas. Mimeo.
- Tirole, Jean.** 1988. *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Van Huyck, John B., Raymond C. Battalio, and Richard Beil.** 1990. "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review*, 80(1): 234-248.
- Van Huyck, John B., Joseph P. Cook, and Richard C. Battalio.** 1997. "Adaptive Behavior and Coordination Failure." *Journal of Economic Behavior and Organization*, 32(4): 483-503.
- Volij, Oscar.** 2002. "In Defense of DEFECT." *Games and Economic Behavior*, 39(2): 309-321.
- Young, Peyton H.** 1993. "The Evolution of Conventions." *Econometrica*, 61(1): 57-84.

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



Note: Solid, dashed and dash-dotted lines denote cooperation rates in sessions 1, 2 and 3 respectively of each treatment. The 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, and dotted lines the 90% interval of simulated data.