Strategy Choice In The Infinitely Repeated Prisoners Dilemma

Pedro Dal Bó
Brown University & NBER

Guillaume R. Fréchette*
New York University

October 2011

Abstract

We use a novel experimental design to identify the strategies used by subjects in an infinitely repeated prisoners' dilemma experiment. We ask subjects to design strategies that will play in their place. We find that the strategy elicitation has negligible effects on behavior supporting the validity of this method. We study the strategies chosen by the subjects and find that they include some commonly mentioned strategies, such as tit-for-tat and Grim trigger. However, other strategies which are thought to have some desirable properties, such as win-stay-lose-shift are not prevalent. The results indicate that what strategy is used to support cooperation changes with the parameters of the game. Finally, our results confirm that long run miss-coordination can arise.

* We are grateful to Jim Andreoni, Drew Fudenberg, Asen Ivanov, Thomas Palfrey and seminar participants at UCSD, Princeton, GMU, CMU, UQAM, Tilberg, Oxford, Stanford, Paris 1, Purdue, Washington St. Louis, Bologna, Harvard, Toronto, Columbia , Boston University, HBS, Masstricht University, Bocconi University, CIDE and at the ESA conference in Tucson, the Conference on Social Dilemmas (Rice 2010), and the SAET conference (2011) for very useful comments. We thank Emanuel Vespa for his research assistance, CASSEL (UCLA) and SSEL (Caltech) for the Multistage software as well as Rajeev Advani for its adaptation. Fréchette gratefully acknowledges the support of NSF via grant 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.

The theory of infinitely repeated games has been a very active area of research in recent decades and is central to many applications.[1] The main idea behind this literature is that repeated interaction may allow people to overcome opportunistic behavior. This idea has been supported by a series of experiments.[2] However, less is know about the types of strategies people actually use to overcome opportunistic behavior.

Learning what strategies are actually used is of interest on many levels. First, it can help future theoretical work to identify refinements or conditions which generate those strategies as the ones that would be played. As Rubinstein (1998) writes Folk theorems are statements about payoffs, but the equilibria are not sustained by these payoffs but by strategies. He then states "Understanding the logic of long-term interactions requires, in my opinion, the characterization of the equilibrium strategy scheme. […] The repeated games literature has made little progress toward this target." Second, it can help theorists focus their attention on empirically relevant strategies. For example, an influential literature in biology studies which strategies are likely to survive evolution. Given the complexities of working with the infinite set of all possible strategies they focus on finite subsets of strategies which are chosen to include strategies usually studied by theorists (i.e. Imhof, Fudenberg and Nowak 2007). Identifying in the laboratory strategies that are popular with humans can provide an appealing basis on

---

[1] Within economics, repeated games have been applied to many areas: industrial organization (see Friedman 1971, Green and Porter 1984 and Rotemberg and Saloner 1986), informal contracts (Klein and Leffler 1981), theory of the firm (Baker, Gibbons, and Murphy 2002), public finance (Phelan and Stacchetti 2001) and macroeconomics (Rotemberg and Saloner 1986 and Rotember and Woodford 1990) just to name a few.

[2] Roth and Murnighan (1978) and Murnighan and Roth (1983) were the first papers to induce infinitely repeated games in the lab by considering a random continuation rule. The probability with which the game continues for an additional round induces the discount factor. A large experimental literature now exists on infinitely repeated games. Palfrey and Rosenthal (1994) study an infinitely repeated public good game. Engle-Warnick and Slonim (2004 and 2006) study infinitely repeated trust games. Holt (1985) study a Cournot duopoly which is related to the prisoners' dilemma studied in Feinberg and Husted (1993), Dal Bó (2005), Normann and Wallace (2006), Dal Bó and Fréchette (2011), Blonski et al. (2007) who more specifically study infinitely repeated prisoners' dilemma under perfect monitoring. Schwartz, Young, and Zvinakis (2000) and Dreber, Rand, Fudenberg, and Nowak (2008) study modified prisoners' dilemmas. Aoyagi and Fréchette (2009) and Fudenberg, Rand, and Dreber (2011a) study infinitely repeated prisoners' dilemma under imperfect public monitoring. Duffy and Ochs (2009) and Camera and Casari (2009) and Camera, Casari, and Bigoni (2010) study repeated prisoners' dilemma with random matching. Finally, Cason and Mui (2008) study a collective resistance game, Fudenberg, Rand, and Dreber (2011b) correlate behavior in the dictator game to behavior in an infinitely repeated game, and Cabral, Ozbay, and Schotter (2010) study reciprocity.

[4] Cwojdzinski and Kamecke provide a partial elicitation of infinitely repeated game strategies. Their focus is on finding alternative ways of inducing infinitely repeated games in the laboratory.

which to select strategies to include. Third, it can also help identify in which environments cooperation is more likely to emerge. In other words, the theoretical conditions needed for cooperation may need to be modified once we restrict the analysis to the set of strategies which are actually used. Fourth, identifying the set of strategies used to support cooperation can provide a tighter test of the theory than the previous study of outcomes. It allows us to test whether the strategies used coincide with the ones that theory predicts should be used (i.e. are the strategies used part of a sub-game perfect equilibrium?).

Previous papers have estimated the use of strategies from the observed realization of behavior. There are serious hurdles for identification. First, the set of possible strategies is infinite (uncountable). Second, while a strategy must specify an action after each possible history, for each repeated game we only observe one realized finite history and not what subjects would have done under other histories. Two different approaches have been used to overcome these hurdles. Both methods start by specifying a family of strategies to be considered. They differ in how the best fitting strategies are selected. One approach trades off goodness of fit of a set of strategies versus a cost of adding more strategies (see Engle-Warnick and Slonim 2004 and 2006a, and Camera, Casari, and Bigoni 2010; for a related but Bayesian approach to this see Engle-Warnick, McCausland, and Miller 2004). A second approach uses maximum likelihood estimation to either estimate the prevalence of each strategy in the set under consideration (see, Dal Bó and Fréchette 2011 and Fudenberg, Rand, and Dreber 2010) or by estimating the best fitting strategy while allowing for subject specific heterogeneity in the transitions across states of the strategy (Aoyagi and Fréchette 2009).

In this paper we propose an alternative approach to study strategies: the elicitation of strategies (i.e. a modified strategy method, Selten 1967). [4] We ask subjects to design strategies that will play in their place. A major challenge to the use of the strategy method is that it can affect behavior.[5] We show that not to be the case by combining the strategy method with the "decision by decision" method in such a way that we can compare

[5] Hoffman et al. (1998), Gueth et al. (2001), and Brosig et al. (2003) report evidence that the strategy method may affect behavior. Brandts and Charness (2000) find the opposite.

behavior between both methods but more importantly by comparing behavior with a similar series of experiments without the elicitation of strategies.

We find that the most popular strategies are always defect, tit-for-tat and grim. The prevalence of cooperative strategies is greater in treatments more conducive for cooperation (high probability of continuation and high payoffs from mutual cooperation). More interestingly, the prevalence of tit-for-tat and grim among cooperative strategies depends on the parameters of the games. Grim is more prevalent in treatments with high payoffs from mutual cooperation, while tit-for-tat is more prevalent in treatments with a high probability of continuation.

The prevalence of tit-for-tat stresses the fact that subjects do not only choose strategies that are part of a subgame-perfect equilibrium. In fact, only 54% of the chosen strategies are part of a subgame-perfect equilibrium while 78% is part of Nash equilibrium.

There are previous papers which have also elicited strategies in infinitely repeated and then made them compete in computer tournaments going back to Axelrod (1980b). The focus of this literature was on the relative performance of the different strategies. Since then the literature has moved mostly towards simulations rather than tournaments.[6] In addition to focus on a different population and vary a set of important parameters, our paper shows that the elicitation of strategies does not affect behavior supporting its use.

After the section on experimental design, the results section will cover four broad areas. First, results pertaining to the new method we propose will be presented. This will be followed by an analysis of the strategies used by subjects. Then the limitations imposed by focusing on strategies with memory one will be explored in a set of additional sessions that relaxes that constraint. Finally, a method to recover strategies econometrically will be evaluated.

[6] Examples of recent papers using computer simulations are Nowak and Sigmund (1993) who introduce stochastic strategies, and Nowak, Sigmund, and El-Sedy (1995) who add mutations. Axelrod's (1980a) first competition was a finitely repeated game. Another study that estimates strategies but focuses on the case of finite repetitions is that of Selten, Mitzkewitz, and Uhlirich (1997).

II. Experimental Design

The experimental design is in three phases.[7] In all phases subjects participate in randomly terminated repeated prisoners' dilemma games. A repeated game or supergame is referred to as a match, and is composed of multiple rounds. After each match, subjects are randomly re-matched with a subject.

In Phase 1 subjects simply play the randomly terminated games. In between matches they are reminded of the decisions they took in the last match and of the choices of the person they were matched with. The first match to end after 20 minutes of play marks the end of Phase 1.

In Phase 2, subjects are first asked to specify a *plan of action*, that is a strategy, by answering five questions: "In round 1 select {1, 2}", and the answer to the four questions covering all permutations of "After round 1 if, I last selected [1, 2] and the other selected [1, 2], then select {1, 2}". The choices are presented as drop-down menus and the order in which the 4 questions after round 1 appear is randomized.

After having specified their plan of action, subjects then play the match just as in Phase 1, taking decisions in every round. At this point, the plan of action they specified is irrelevant. After the first match, they are shown what decisions they took in this match, what decisions the person they were matched with took, and what decisions the plan of action they specified would have taken, had it played in their place. They are then asked to specify a plan of action for the coming match. This process (specify a plan; play a match round by round; receive feed back and specify a plan) is repeated for 20 minutes. After 20 minutes of play in Phase 2, the plan of action takes over for the subjects, finish the ongoing match, and play an additional 14 matches; this is Phase 3.

Table 1: Stage Game Payoffs

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

[7] When first reading the instructions, subjects are informed that there are multiple phases, but they are only told about the procedures for Phase 1. Additional instructions are given to them after Phase 1. All instructions and screen-shots are available in the online appendix at https://files.nyu.edu/gf35/public/print/Dal_Bo_2011b_oa.pdf

The stage game is as in Table 1. Each subject is exposed to only one treatment (between-subjects design). The main treatment variables are the payoff from mutual cooperation R and the probability of continuation δ where R takes values 32 or 48 and δ takes values 1/2 or 3/4. One additional treatment is conducted with R = 32 and δ = 9/10. For each of the 5 treatments, 3 sessions are conducted.[8] Payments are based on the sum of the points accumulated in the three phases converted to dollars at the rate of 100 points equal $0.45.

Given those parameters, cooperation can be supported as part of a subgame perfect equilibrium (henceforth SGPE) in all treatments except for the one where δ = 1/2 and R = 32. Furthermore, playing a Grim trigger strategy (cooperating until the other defects and then defect forever) is risk dominant when playing against always defect in both treatments with R = 48 and in the treatment with δ = 9/10.[9]

The design considerations were the following. First, one concern is that asking subjects for a strategy might cue them to something they would not be thinking about absent our intervention. To evaluate to what extent this is a concern we will compare behavior in phase 2 with behavior in Dal Bó and Fréchette (2011), which uses the same parameter values but without the strategy method. Furthermore, the design includes Phase 1 where subjects have time to learn what they want to do in this environment. There is no "cueing" since they are not asked about their strategy at that point.

An additional concern is that subjects may not think in terms of strategies and may not know how to verbalize their plan. To address that concern the design gives feedback about what both subjects did and what the plan of action would have done during phase 2. This gives an opportunity for subjects to realize in what situations the specified plan of action is not doing what they actually want to do. Finally, subjects are incentivized to give the plan of action that makes choices as close as possible to what

[8] Two previous sessions were conducted, however the payments were too low and thus the exchange rate was changed and those 2 sessions are not included in the analysis. Those first two sessions had slightly higher cooperation rates. As a point of comparison for when we describe the results, if those sessions are included the cooperation rate in round 1 of the last match of Phase 1 and Phase 2 would be 0.8 and 0.91 respectively. The comparison can be made with the results presented in Table A4.

[9] The reader interested in the reasons for the choice of parameters is referred to Dal Bó and Fréchette (2011). See also that paper, as well as Blonski et al (2010), for a discussion of the application of the concept of risk dominance to infinitely repeated prisoners' dilemma games.

they want to do since whenever they specify a plan of action, this may be the match where Phase 2 expires and where the plan of action takes over.

A final concern is whether the possible plans of action subjects can specify are sufficient to express the strategies they want. First note that even though simple, the technology at their disposal allows subjects to specify 32 strategies. Also, many of the strategies most often mentioned in the literature can be specified in this way: tit-for-tat, Grim, win-stay lose-shift, and others are available – these (and other) strategies will be defined later. It will also be possible to compare decisions in Phase 2 with the decisions of the plan of action. Any persistent differences would be suggestive that no strategy available was exactly consistent with what the subject wanted to do. Finally, and providing the most direct evidence on this count, we also conducted additional sessions with a slightly different design which allows for additional strategies. These will be described later on in the paper.

III. Choices and the Impact of the Elicitation of Strategies on Behavior

A total of 246 NYU undergraduates participated in these 15 sessions, with an average of 16.73 subjects per session, a maximum of 22 and a minimum of 12. The subjects earned an average of \$24.94, with a maximum of \$42.53 and a minimum of \$12.26. In the treatments with $\delta=1/2$, $\delta=3/4$, and $\delta=9/10$ the average number of rounds per match was 1.92, 3.61, and 8.53 respectively, and the maximum was 9, 22, and 43 respectively. Table A1 in the appendix has more details on each session.

III.a Evolution of choice behavior

We start the presentation of the experimental results by describing behavior in the different treatments. Table 3 presents the cooperation percentages by treatment and by phase (for round 1 and all rounds). In Phase 2, this table reports the subjects' actual decisions, not the choices that the plan of action would have selected. The top panel reports the average for rounds 1 and the bottom panel reports the average for all rounds. It is clear from this table that cooperation rates evolve over time in many treatments. For example, while cooperation decreases with experience under $\delta=1/2$ and R=32, it experiences large increases under $\delta=3/4$ and R=48. As a result, experience magnifies the effects of increases of δ and R on cooperation. This effect of experience across treatments has already been observed in Dal Bó and Fréchette (2011) where strategies are not elicited.

Table 3: Percentage of Cooperation by Treatment

Round 1

First Match of Phase 1				Last Match of Phase 2			
δ \ R	32		48	δ \ R	32		48
½	40.00	<**	71.43	1/2	6.00	<***	58.93
	V		v		^*		^
3/4	36.36	<*	63.04	3/4	22.73	<***	82.61
	V				^***		
9/10	30.00			9/10	62.00		

All Rounds

First Match of Phase 1				Last Match of Phase 2			
δ \ R	32		48	δ \ R	32		48
1/2	30.70	<***	67.69	1/2	5.88	<*	37.50
	V		v		^**		^**
3/4	26.32	<**	54.46	3/4	19.74	<***	86.05
	V				^***		
9/10	21.98			9/10	55.47		

Note: * significance at 10%, ** at 5% and *** at 1%. Statistical significance is assessed by estimating a probit and clustering the variance-covariance at the level of the experimental session.

III.b The impact of eliciting strategies (Comparison with Dal Bó and Fréchette 2011)

The fact that Dal Bó and Fréchette (2011) uses an experimental design which is identical to the one in this paper with the exception that strategies are not elicited allow us to compare cooperation rates in both experiments as test of whether elicitation affects behavior.[10]

Figure 1 shows the average cooperation rate in round 1 across repeated games by treatment in this experiment and in Dal Bó and Fréchette (2011) where strategies are not elicited.[11] For this experiment, the data includes Phase 1 followed by 2. In both cases the figures are for rounds 1 only.[12]

---

[10] Procedures for Phase 1 are almost identical to procedures in Dal Bó and Fréchette (2011). Besides issues of timing (Phase 1 is shorter than the entire experiment in Dal Bó and Fréchette (2011)), the only difference is that subjects were not reminded of the choices they, and the other player, had taken in the previous match between matches in Dal Bó and Fréchette (2011).

[11] Only matches with data for all three sessions are included.

[12] This is done for simplicity.

The main observation is that the evolution of behavior is extremely similar between both experiments for all treatments. However, even though the choices are very similar across experiments, there seems to be slightly more cooperation when strategies are elicited in the δ=1/2 and R=48 sessions and slightly less cooperation in the δ=3/4 and R=48 session.

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

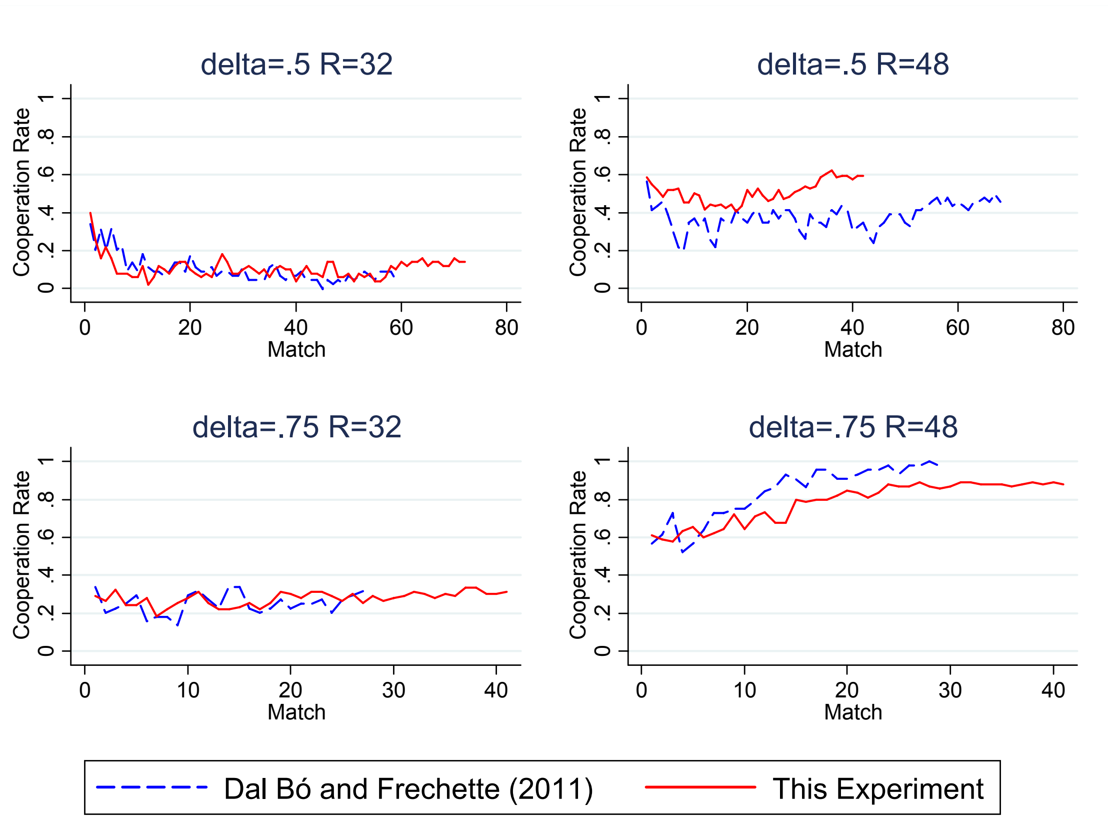

Table 2 shows the cooperation rates for both series of experiments separated by phase (as Dal Bó and Fréchette 2011 does not have phases, matches were assigned to phase 1 and 2 based on whether they started before the mid-point of the session or after the mid-point, the number of phase 1 and 2 matches in the new sessions is comparable to the total number of matches in Dal Bó and Fréchette 2011). By the end of phase 2, taken as a whole, there is only marginal evidence that the new sessions are different from those in Dal Bó and Fréchette (2011) (p-value = 0.09 for round 1 only and > 0.1 for all rounds).[13] The difference in the round 1 only case is driven by the δ=3/4 and R=48

[13] Unless otherwise noted, statistical significance is assessed by estimating a probit and clustering the variance-covariance at the level of the experimental session.

treatment as taken individually it is the only treatment for which there is a statistically significant difference by the end of phase 2 (all other p-values > 0.1). However the difference is already present at the end of phase 1, indicating that the difference is not due to the elicitation method. Moreover, if we consider all matches in phase 2, there are no significant differences with Dal Bó and Fréchette (2011) in any of the treatments (either in round 1 or all rounds). Importantly, the comparative static comparisons across treatments are unaffected by the strategy elicitation.

Table 2: Cooperation Rate by Treatment, Phase and Elicitation of Strategies

		First Rounds Only				**All Rounds**			
		Elicitation of Strategies				Elicitation of Strategies			
		Yes		No*		Yes		No*	
		Phase		Phase°		Phase		Phase°	
δ	R	1	2	1	2	1	2	1	2
Panel A: All Matches									
½	32	0.12	0.09	0.14	0.06	0.09	0.07	0.12	0.08
	48	0.61	0.60	0.37	0.41	0.56	0.52	0.35	0.36
¾	32	0.23	0.24	0.25	0.26	0.19	0.21	0.18	0.23
	48	0.72	0.86	0.75	0.96	0.61	0.80	0.65	0.90
Panel B: Last Match in Each Phase									
½	32	0.04	0.06	0.14	0.02	0.02	0.06	0.10	0.05
	48	0.45	0.59	0.39	0.41	0.31	0.38	0.39	0.41
¾	32	0.23	0.23	0.25	0.30	0.22	0.20	0.16	0.21
	48	0.67	0.83	0.89	0.98	0.57	0.86	0.77	0.96

* Data from Dal Bó and Fréchette 2011.
° Phases 1 and 2 in this case are defined as matches that start before and after the mid-point of the session.

[18] More precisely, Grim is the strategy that starts by cooperating and cooperates unless there has ever been a defection in the past. It is not possible to exactly construct this strategy using the memory one mechanisms available to the subjects. However, it is possible to construct a memory one machine that will be equivalent to Grim in how it plays against any other strategy.

The rate of deviations between plan of actions and actual choices in phase 2 is also relatively low. Only 7% of all choices (across al treatments) do not correspond to what the plan of action would have selected. The median subject makes a choice other than what the plan of action would have selected in only 2% of their choices.

An additional piece of evidence indicating that the strategy method did not affect behavior is the fact that the order of elicitation did not affect choices. The elicitation of strategies consisted of five questions (what you would do in the first round, what you would do if both cooperated, etc). The question about the first round was asked first and the order of the remaining four questions was random. Table OA1 in the online appendix shows what estimates from a probit of order in which the questions were asked on choices by treatment. In no treatment are the estimates for the order jointly significant (at the 10% level).

III.d Changes in Plan of Actions

Finally, we document when subjects changed their plan of action in phase 2. Subjects changed their plan of action in 14% of all matches (past the first match in phase 2). This number varies between 12% and 21% depending on the treatment. It is lowest in the $\delta=1/2$ and R=32 treatment and highest in the $\delta=9/10$ and R=32 treatment. Figure 2 plots the frequency with which subjects changed their plan of action, namely the probability that they changed their plan of action from one match to the next. As can be seen, in every treatment the evolution displays a downward trend, suggesting that subject's behavior is stabilizing.

Figure 2: Frequency of Change in Plan of Action

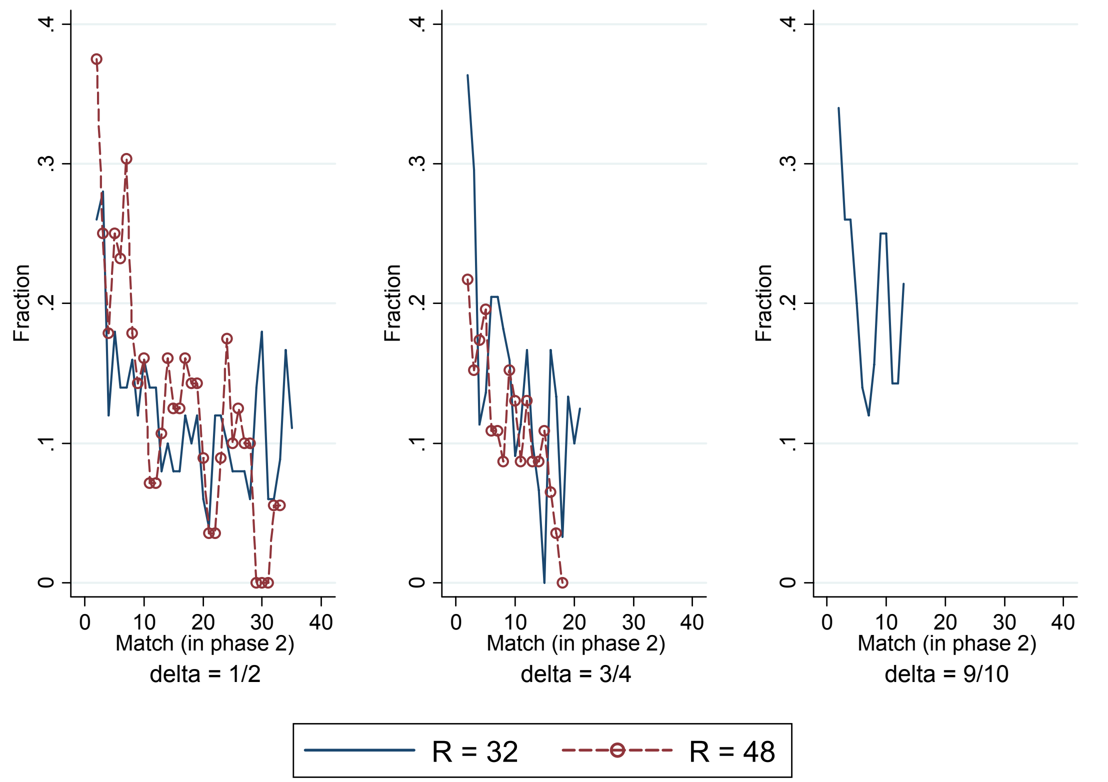

The frequency with which subjects change their plan of action is markedly higher following a match in which their decisions were not the same as what the plan of action they had specified would have taken. This can be seen in Table 4 which breaks down the frequency of changes by whether the choices in the previous match corresponded to the choices the plan of action would have taken.

Table 4: Percentage of Matches Where the Strategy Is Changed

		Decision = Choice of Strategy	
δ	R	Yes	No
½	32	10.08	27.33
	48	11.75	36.25
¾	32	11.02	35.43
	48	9.38	40.32
9/10	32	15.27	45.05
Overall		11.02	35.29

Although Table 4 suggests that subjects adapt their plan of actions to reflect the decisions they make, it could still be possible that some subjects respond to a difference between behavior and plan of action by changing their behavior in future matches. However, if that was the case, choices in this experiment would be systematically different from the Dal Bó and Fréchette (2011) experiment contrary to what it is the case.

IV. Description of strategies

Having shown that the strategy method is not likely to have affected behavior (and hence we believe strategies) we describe now the strategy choices made by the subjects. After first defining some strategies of interest, we describe the strategy choices at the end of phase 2, and then the evolution leading to that choice.

The simplest strategies to consider are always cooperate (AC) and always defect (AD). A strategy already mentioned is tit-for-tat (TFT). TFT starts by cooperating and in subsequent rounds matches what the other subject did in the previous round. The Grim trigger strategy (Grim) also starts by cooperating, and cooperates as long as both players have cooperated last round, and defects otherwise.[18] A strategy that is often discussed in the literature is win-stay lose-shift (WSLS, also know as perfect TFT, or Pavlov), it starts by cooperating, and cooperates whenever both players made the same choice last round, and defects otherwise. It is considered desirable because when it plays itself, it would not defect for ever after a deviation. Finally, another strategy of interest is suspicious tit-for-tat (STFT), it starts by defecting and from then on matches what the other subject did in the previous round.

What can be expected? Table A2 in the appendix shows for each treatment the set of strategies that are part of a Nash equilibrium or a subgame perfect equilibrium. AD is a subgame perfect equilibrium in every treatment and previous experiments on infinitely repeated games have observed situations where some subjects defect. Hence it seems plausible to expect that AD will be selected.

What about strategies to support cooperation? TFT is a likely candidate given that it was the winner in Axelrod's (1982b) tournament. It is also a very intuitive strategy to specify. On the other hand, TFT is not subgame perfect in general. To see this, consider the subgame that follows after player 1 cooperated and player 2 defected. If both players follow TFT after that defection, they will start an infinite sequence of alternating unilateral defections resulting in a total payoff below the payoff from mutual cooperation. This gives both subjects have an incentive to cooperate once to return to full cooperation when TFT tells them to defect if the discount factor is high enough. Thus, maybe TFT will not be popular but instead WSLS, which has performed well in simulations, will be popular. WSLS unlike TFT is a SPGE strategy for a sufficiently high δ. However, WSLS cannot support cooperation for as many values of δ as Grim. The Grim strategy, in that sense, is more robust, and it can support cooperation for much lower values of δ. On the other hand, once it starts defecting, Grim never stops. Thus in that sense it is not forgiving. These different tensions make it unclear which strategy one should expect to see most.

IV.a Final strategy choices

Table 5 shows the distribution of strategies across treatment. Strategies are described by the string of five letters, C for cooperation and D for defection, where the first entry is what it recommends in round 1, the second gives the choice following mutual cooperation, the third column gives the choice after one's own defection if the other cooperated, and so on. Under the label AKA are listed the popular names by which some of these strategies are known. AC stands for Always Cooperate and similarly AD is for Always Defect. Grim, the grim trigger strategy, is the strategy that first cooperates and keeps on cooperating until someone defects. TFT is for Tit-For-Tat, the strategy that first cooperates and then does exactly what the other player did in the previous round (it is sometimes refered to as the matching strategy). WSLS, or Win-Stay-Lose-Shift, which also appears under the name Pavlov and perfect tit-for-tat, starts by cooperating, but then defects if one, and only one, of the two players defects in the previous round, and cooperates otherwise. Finally, STFT is suspicious tit-for-tat which is identical to tit-for-tat except on the first move where it defects.

The most popular strategies across treatments are TFT, Grim and AD. These three strategies on their own correspond to more than two thirds of the data in each treatment and as much as 80% in two treatments. As it can be expected there are large variations in the popularity of AD across treatments. While AD is more prevalent in treatments with low δ and R, TFT and Grim are more prevalent in treatments with large δ and R.

Table 5: Distribution of Elicited Strategies (Last Match)

Strategy*						δ = ½		δ = ¾		δ = 9/10
						R		R		R
					AKA	32	48	32	48	32
C	C	C	C	C	AC		3.57	2.27	8.70	2.00
C	C	C	C	D	AC'			2.27		4.00
C	C	C	D	C				2.27	2.17	2.00
C	C	C	D	D	TFT	6.00	12.50	11.36	32.61	42.00
C	C	D	D	C	WSLS		**1.79**		**2.17**	2.00
C	C	D	D	D	Grim	6.00	**35.71**	**4.55**	**39.13**	**12.00**
C	D	C	D	D		2.00				
C	D	D	D	C			1.79			
C	D	D	D	D			1.79	4.55		
D	C	C	C	D					2.17	
D	C	C	D	C				2.27		
D	C	C	D	D	STFT	12.00	5.36	18.18		10.00
D	C	D	C	C						2.00
D	C	D	C	D	AD'					2.00
D	C	D	D	C				2.27		
D	C	D	D	D	AD'	10.00	**3.57**	**9.09**	**2.17**	**6.00**
D	D	C	D	C					2.17	2.00
D	D	C	D	D		4.00				
D	D	D	C	D	AD'	2.00				
D	D	D	D	C			1.79			
D	D	D	D	D	AD	**58.00**	**32.14**	**40.91**	**8.70**	**14.00**
Cooperative						12.00	53.57	22.73	84.78	64.00
Defecting						86.00	41.07	68.18	13.04	32.00
SGPE						58.00	73.21	54.55	52.17	32.00
SGPE if random						0.03	15.63	9.38	15.63	9.38
Only NE						28.00	12.50	0.00	34.78	44.00
Only NE if random						15.63	18.75	6.25	15.63	9.38
NE						86.00	85.71	54.55	86.96	76.00

Note: AC' (AD') denotes that a strategy will behave as AC (AD) in every history it will reach if choices are perfectly implemented.

Cooperative (Defecting) denotes strategies that are fully cooperative (defecting) with themselves.

Sub-game perfect strategies are denoted in bold, and NE that are not SGPE are underlined.
* The letters in the strategy names denote the recommended action after each possible contingency: initial period, CC, DC, CD and DD, where the second letter designates the other's choice.

If we aggregate strategies depending on whether they lead to cooperation or defection on their path of play we find a similar pattern.[19] The prevalence of cooperative strategies increases with δ and R, while the prevalence of defecting strategies decreases.

Variations in the popularity of specific strategies to support cooperation across treatments are more surprising. Increases in R have greater impact on the prevalence of Grim than of TFT while the opposite is the case for increases in δ. To see this notice that when delta = ½, increasing R doubles the popularity of TFT while the popularity of Grim increases by a factor of about 6. Similarly when delta = ¾, increasing R leads to a less than threefold increase in TFT but a more than eightfold increase in Grim. If, instead, we fix R and look at the impact of increasing delta, we get the opposite; TFT always increases by a greater factor than Grim. This suggests that how subjects choose to support cooperation depends on the particular parameters of the game in a systematic way: as the expected length of the game increases subjects choose to use shorter punishments.

Interestingly, a large proportion of the strategies being chosen are not part of sub-game perfect equilibria (SGPE). In particular, the proportion of strategies that conform to a SGPE when playing against itself reaches a low of 32% under $\delta=9/10$ and $R=32$. The maximum is reached under $\delta=1/2$ and $R=48$ with 73% of strategies being SGPE. The low prevalence of SGPE strategies is related to two observations. First, under $\delta=1/2$ and $R=32$ (that is when cooperation cannot be supported in equilibrium) a significant fraction of subjects (12%) choose defecting strategies that are equivalent in play to AD (DCDDD and DDDCD) but are not SGPE. Second, in the treatments in which cooperation can be supported in equilibrium subjects not only choose SGPE cooperative strategies but also rely heavily on TFT. The famous TFT is not SGPE but can be a Nash equilibrium (see Table A2 in the appendix). In fact, 24% of the strategies chosen are part of Nash equilibria while not being part of a subgame perfect equilibria.

IV.b. Evolution of strategies

Table 6 shows the evolution in the prevalence of the most important strategies (AC, AD, TFT, Grim and STFT) for the first and last repeated game in phase 2, with the

[19] We define a strategy as cooperative (defecting) if it would lead to full cooperation on the path of play against itself (regardless of equilibrium considerations).

exception of WSLS because as can be seen in Table 5 it is not a popular strategy. The evolution of strategies in phase 2 is also reported in Figure A2 (in the appendix). These allow us to study how subjects' choice of strategies evolved. The observed evolution may be due to subjects changing their desired behavior or their better understanding of the functioning of strategies.

Table 6: Evolution of Main Strategies (Percentages for First and Last Match in Phase 2)

	δ = ½						δ = ¾						δ = 9/10	
	R=32			R=48			R=32			R=48			R=32	
	Match		p-v	Match		p-v	Match		p-v	Match		p-v	Match	
Strategy	First	Last		First	Last		First	Last		First	Last		First	Last
AC	0	0	-	9	4	0.05	2	2	-	7	9	0.73	2	2
AD	46	58	0.00	20	32	0.00	39	41	0.70	9	9	-	12	14
TFT	8	6	0.43	21	13	0.24	18	11	0.35	28	33	0.54	28	42
Grim	0	6	0.00	29	36	0.00	2	5	0.39	30	39	0.01	20	12
STFT	16	12	0.35	5	5	-	11	18	0.01	2	0	-	10	10
SGPE	46	58	0.00	54	73	0.00	52	55	0.61	41	52	0.00	42	32
Only NE	38	28	0.07	23	13	0.24	0	0	-	33	35	0.80	28	44

p-v stands for p-value of the test that the percentage of each plan of action is the same between the first and last matc

When it comes to the evolution of plan of actions, there are very few patterns that are true for all treatments. One is that AD is never lower in the last match than it is in the first match. That being said, in some treatments the popularity of AD grows appreciably while in others it does not change. Another pattern that is true in all treatments is that NE plan of actions are at least as frequent in the last match as they are in the first. However, it must be noted that it is not the case that the prevalence of SGPE strategies increases with experience in all treatments. The prevalence of SGPE strategies increases significantly with experience in 3 out of 5 treatments and in one treatment it decreases significantly. In that treatment (delta = 9/10), the decrease can be identified with the decrease in popularity of Grim and with the important increase in popularity of TFT, which is a NE but not subgame perfect. In fact, that is the only treatment in which TFT has a significant change over time, in all other treatments, there is no statistically significant change. On the other hand, the Grim strategy increases significantly in both treatments with R = 48 and in the treatment with delta = ½ and R = 32.

Interestingly the evolution of strategies in some treatments does not have a clear effect on cooperation. For example, under δ=1/2 and R=48 the prevalence of both AD and Grim increase with experience. This suggests that even when subjects gain significant experience they may fail to coordinate on one equilibrium (consistently with Dal Bó and Fréchette 2011).

IV.c. Expected Payoffs

To better understand strategy choice, it is helpful to determine the expected payoffs of different plan of actions given the distribution of plan of actions. Figure 3 presents expected payoffs as a function of the popularity of each plan of action. These are the theoretical expected value given the continuation probability but using the empirical frequency of each plan of action in the last match of phase 2. For each data point, the shape of the marker indicates if the plan of action is a NE, a SGPE, or neither. The figure also indicates the name of the plan of action for plan of actions which are present in at least 10% of the data.

Some noteworthy aspects of the figure are the following. First, the worst performing plan of action is always substantially worst than the best performing one. Second, the most popular plan of actions tend to be close to the best performing plan of actions. In fact, in every treatment, the most popular plan of action is the one with the highest expected payoff (at least weakly).

Figure 3: Expected Payoffs For Each Plan of Action Conditional on the Distribution of Plan of Actions (last round of Phase 2)

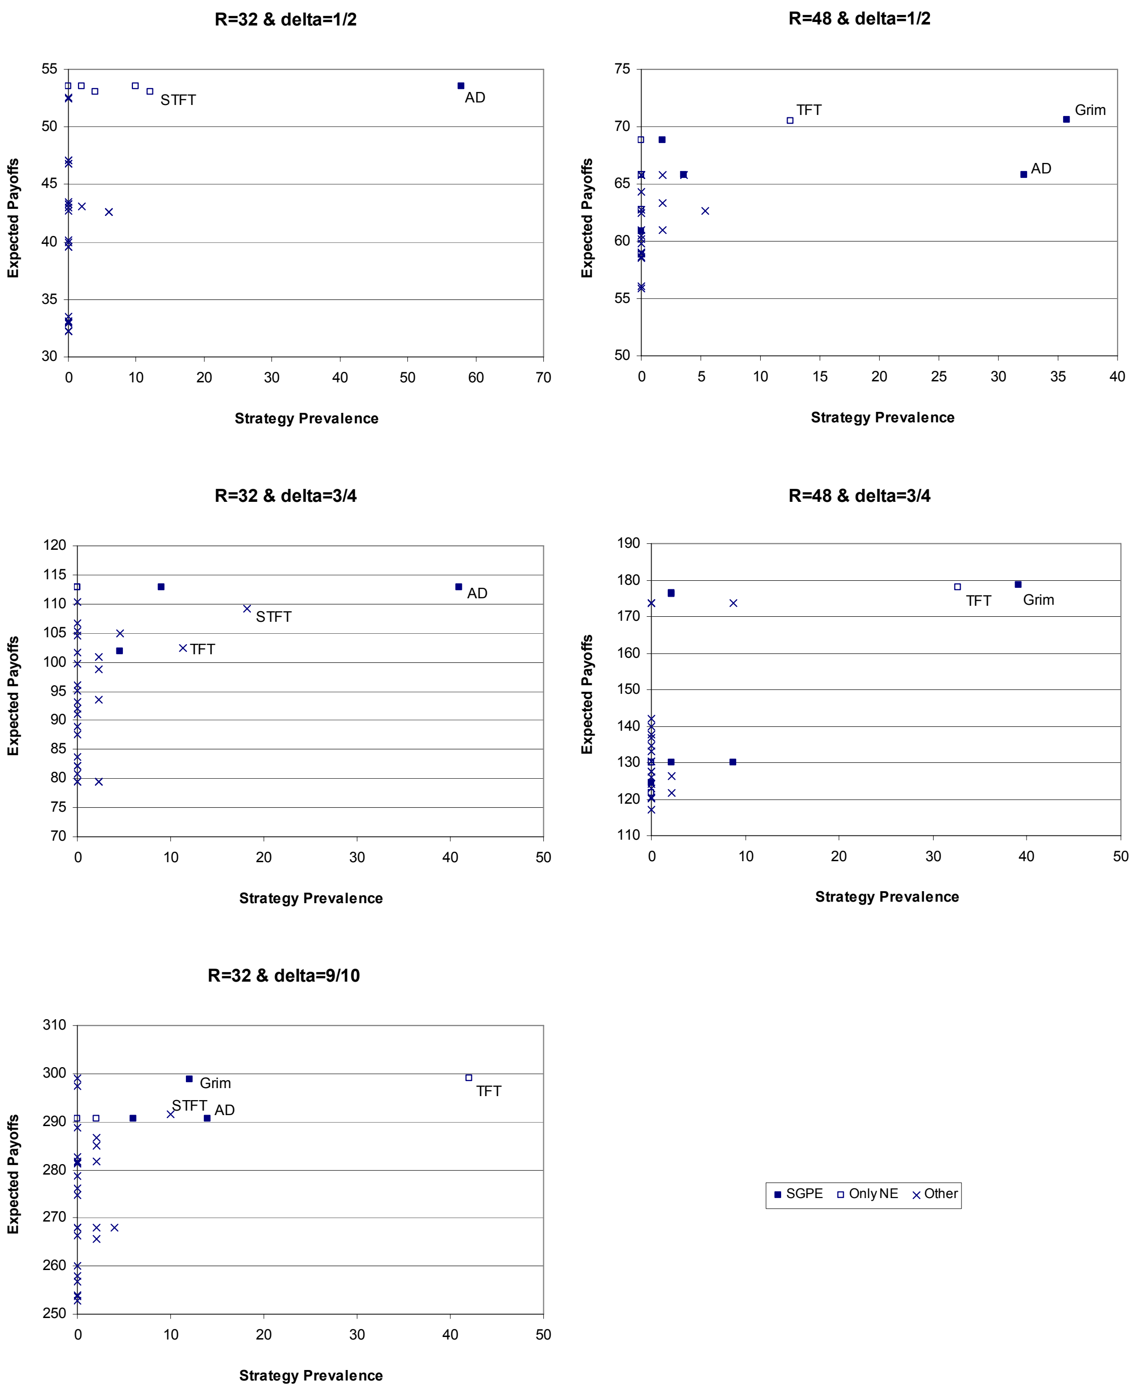

Table 7 provides a calculation of lost payoffs due to suboptimal strategy choice. The first row shows the maximum expected payoff by treatment that can be obtained by one of the memory-1 strategies given the distribution of strategies in the experiment. The second row shows the average expected payoff across strategies (weighted by their prevalence). The difference between the maximum and average expected payoff in one repeated game goes from 1.58 points under δ=1/2 and R=32 to 8.31 points under δ=3/4 and R=48. The monetary value of these losses for the 14 games that subjects expect to play in phase 3 of the experiment goes from 10 cents to 52 cents with an average of 30 cents. The average loss from suboptimal strategy choice corresponds, in average, to less than 4% of the average expected payoffs in phase 3.

Table 7: Expected Losses from Strategy Choice

	δ = ½		δ = ¾		δ = 9/10
	R		R		R
Expected Payoffs	32	48	32	48	32
Maximum (Played)	53.50	70.55	112.91	178.58	299.11
Average (Played)	51.92	67.83	107.64	170.27	293.02
Minimum (Played)	42.64	61.01	79.36	121.78	265.72
Worst Possible	32.27	55.91	79.36	117.15	252.77
Difference (Max – Avg)/Avg	3%	4%	5%	5%	2%
Difference (Max – Min)/Min	26%	16%	42%	47%	13%

The small expected cost of the suboptimal strategy choice are not surprising given the observation from Figure 3 that most of the popular plan of action have high expected value. However, Table 7 also reveals that there are incentives to move away from the least profitable plan of actions, as the plan of action with the lowest expected payoffs that was selected always implied a loss of at least 13%, and in some treatment substantially more, than the best performing plan of action. If one considers the worst possible plan of action (which in most treatments are never selected), then the incentives would appear even startker. Note, however, that Table 7 and Figure 3 suggest that the financial incentives to move away from plan of actions that are NE but not SGPE to others that are

SGPE were often small or non-existent as some NE plan of actions performed very well in expected value terms.[20]

V. Is Memory 1 Enough?

In the experiments presented in the previous sections subjects could only choose strategies that condition behavior on the outcome of the previous period. While we show that this has little effect on behavior, it could still be the case that this restriction in the elicitation of strategies greatly affects the strategies chosen by the subjects.

In this section we present strategy choices when a greater set of possible strategies is given to the subjects. In this menu condition, subjects are offered a menu of strategies. This menu included some of the strategies they could build in the original sessions (AC, AD, Grim, TFT, STFT, and WSLS), but it also included some additional strategies. In particular it allowed for some strategies with "softer" triggers, such as Grim-X which is a Grim trigger strategy that requires X defections before triggering punishment. X is a parameter to be selected by the subject. Similarly there is a TFXT which is similar to TFT but requires X consecutive defect choices before it defects. It also includes a trigger strategy with a finite number of D choices before reverting back to cooperation, and a few other strategies. Subjects always had the opportunity to "build" their plan of action as in the original sessions. The order in which these strategies appeared was randomized. Our goal was to include the most obvious possibilities in terms of strategies in this game. The same R and δ as in the original sessions are used, with the exception of δ=1/2 and R=32 since that treatment is unlikely to generate interesting results given the prevalence of AD.

A total of 182 NYU undergraduates participated in these 12 sessions, with an average of 15.54 subjects per session, a maximum of 20 and a minimum of 12. The subjects earned an average of \$27.51, with a maximum of \$43.13 and a minimum of \$10.55. In the treatments with δ=1/2, δ=3/4, and δ=9/10 the average number of rounds

[20] Note that these estimates of the losses assume that subjects know the distribution of strategies in their treatment. A more appropriate estimate of the monetary costs may have to consider the fact that subjects' may hold different beliefs which cannot be contradicted by the feedback they receive during the experiment. Once the correct beliefs are considered the losses from sub-optimal strategy choice may be even be smaller (see Fudenberg and Levine 1997).

per match was 2.14, 4.16, and 9.42 respectively, and the maximum was 9, 34, and 42 respectively. Table A3 summarizes the treatments that were conducted and basic information about the sessions. Table A4 reports the cooperation rates in these additional sessions as well as in the other sessions for comparisons. The key feature to note is that when compared to the decisions in Dal Bó and Fréchette (2011), the choices in these additional sessions taken jointly are not statistically different at the end of phase 2 (for either all rounds or round 1 only). This provides further evidence that the elicitation of strategies does not affect behavior.

The main result from these additional sessions is that over three quarters of the final strategies in each treatment could be defined using the apparatus of the original sessions. More specifically, 90% of the strategies in the last match for δ=1/2 and R=32; 92% and 77% for δ=3/4 and R=32 and 48 respectively; and 85% for δ=9/10 and R=32. Furthermore, in all but one treatment the most popular strategy that cannot be expressed using the original method ranks as fifth most popular. In the only treatment where it does better, δ=9/10 and R=32, it comes in tied for 3rd, but accounts for only 7% of the strategies. The few non-memory-1 plan of actions that are selected are the following. CD first cooperates for a number of periods determined by the subject, followed by defection until the end. CTFT, a condition form of TFT that accounts for who started to deviate (i.e. it does not punish a deviation which is a response to a prior deviation). GRIMX is a grim trigger that starts after a number X of deviations. RX is a random plan of action which randomizes between cooperation and defection at a certain rate X per round. TFXT is a tit-for-tat that counts the number of deviations of the other and deviates only after X rounds of deviations by the other. Finally, TX is a trigger strategy that punishes for X rounds before reverting to cooperation.

Table 8: Distribution of Elicited Strategies in Additional Sessions (Last Match)

Strategy*					δ = ½	δ = ¾		δ = 9/10
						R		R
				AKA	48	32	48	32
C	C	C	C	AC		0.03	0.11	
C	C	D	D	TFT	<u>0.21</u>	0.21	<u>0.25</u>	<u>0.24</u>
C	D	D	C	WSLS	**0.08**	0.08	**0.11**	0.07
C	D	D	D	Grim	**0.33**	**0.08**	**0.18**	**0.07**
D	C	C	C				0.02	
D	D	D	D					0.02
C	C	C	D				0.02	
C	C	D	D	STFT		0.08	0.02	0.04
C	D	D	D	AD'	**0.02**			
D	C	D	D			0.08		
D	D	D	D	AD	**0.25**	**0.38**	**0.05**	**0.41**
Memory 1 Total					0.90	0.92	0.77	0.85
				CD		0.03	0.02	0.04
				CTFT	0.02		0.09	0.02
				GRIMX	0.04			0.02
				RX	0.02		0.02	0.07
				TFXT			0.07	0.02
				TX	0.02		0.02	

Note: AC' (AD') denotes that a strategy will behave as AC (AD) in every history it will reach if choices are perfectly implemented.

Sub-game perfect strategies are denoted in bold, and only NE are underlined.
* The letters in the strategy names denote the recommended action after each possible contingency: initial period, CC, DC, CD and DD, where the second letter designates the other's choice.

Table 8 reports the percentages of each strategy in the last match of Phase 2. Many of the results from the original session carry over. In particular, AD is still a popular strategy. When it comes to strategies to support cooperation, both TFT and Grim are the most common strategies. It is also still the case that which strategy is used to support cooperation changes with the parameters R and δ. Finally, as before, some non-SGPE are popular, however much of the choices of strategy favor NE strategies.

There are some differences however. Although it is still the case that AD, TFT, and Grim together account for the majority of the data, it is no longer true in each treatment taken individually. In particular, when δ=3/4 and R=48, these three strategies now account for 48% of the strategies. In the other treatments they always account for at least two thirds of the data. The next most popular strategy is WSLS. Hence, in these additional sessions, AD, TFT, Grim, and WSLS taken together account for the majority of the strategy choice by the end of phase 2 in every treatment. This is surprising given that WSLS was almost completely absent in the original sessions, representing at most less than 3% of choices in any given treatment. In these new sessions its popularity goes up to 11% of choices in one treatment. Another change is the popularity of STFT which is substantially decreased in these additional sessions. However, the treatment where it was most popular in the original sessions (δ=3/4 and R=32) is still the treatment where it is most popular in the additional sessions.

Albeit for these small differences, the results clearly indicate that for a majority of subjects in this environment, a plan of action with memory 1 is sufficient to express their strategy. Overall 87% of strategies in the last match have memory 1 or less. In fact, of the subjects that select one of the strategies that were not available in the previous design, and thus could have more than memory 1, 38% design it such that it is in fact a memory 1 or less strategy.

VI. Does econometric estimation recovers the same strategies?

Using data from Phase 2 we can evaluate to what extent strategies can be recovered econometrically from observed behavior. We can compare the estimated prevalence of strategies using the choices from phase 2 to the prevalence of strategies actually chosen by the subjects.

We will study the performance of the estimation procedure proposed in Dal Bó and Fréchette (2011) and also used in Fudenberg, Rand, and Dreber (2010).[21] Let us refer to this approach as the Strategy Population Recovery Method (SPRM). The SPRM is as follows. Denote the choice taken by strategy k in round r by $s_r(s^k)$, the choice taken that round by c_r, and the indicator function taking value 1 when the two are the same and 0 otherwise by y_r (i.e. $y_r=1\{c_r = s_r(s^k)\}$). We model the probability that a choice

[21] Note that the notation used here differs slightly.

corresponds to the one prescribed by a strategy as $\Pr(y_r) = \frac{1}{1+\exp(-1/\gamma)} \equiv \beta$ where γ is a parameter to be estimated. In an environment with two choices, such as this one, one can view the above as resulting from the following: $c_r(s^k) = 1\{s_r'(s^k) + \gamma\varepsilon_r \geq 0\}$ where s_r' (is coded with 1 if the strategy would cooperate and -1 otherwise) and ε is an error term with an extreme value distribution. Using subscripts *imr* to denote subject *i* in match *m* and round *r*, the likelihood that the observed choices were generated by strategy *k* is given by $\Pr_i(s^k) = \prod_M \prod_R \left(\frac{1}{1+\exp(-1/\gamma)}\right)^{y_{imr}} \left(\frac{1}{1+\exp(1/\gamma)}\right)^{1-y_{imr}}$ for a given subject and strategy (where *M* and *R* represent the sets of all matches and rounds). From this we obtain a loglikelihood $\sum_I \ln\left(\sum_K \phi^k \Pr_i(s^k)\right)$ where *K* represents the set of strategies we consider, labeled s^1 to s^K, and ϕ^k are the parameters of interest, namely the proportion of the data which is attributed to strategy s^k. One can think of this in large sample as giving the probability of observing each strategy.

To get a better sense of this approach, consider the case where only one strategy is considered, then the only parameter to be estimated is γ. Suppose the estimate is 0.72, this mean $\beta = 0.8$. In other words, when the strategy suggests cooperation, cooperation is predicted to occur 80% of the time. The quality of the fit can be compared to chance which would predict cooperation with probability 50% (since there are only two choices). Consider now the case of more than one strategy. Imagine the strategies included are: AD, TFT, and Grim, and the estimates of their proportion is one third for each of them, and γ is still 0.72. This implies that in round 1 of a match, the estimated model predicts a 60% cooperation rate. Now suppose we look at a specific subject who first cooperated, but he was matched with someone that defected, then in round 2, the estimated model predicts they will cooperate with a 20% probability. If the person they are matched with cooperated in round 2, the estimated model's prediction for round 3 would now be a 40% chance of cooperation.[22]

This model is estimated on the second half of Phase 2 (pooling together both the original experiments and the additional sessions). The first half of Phase 2 is dropped to limit the importance of matches where behavior is still changing. The variance-covariance matrix is obtained by bootstrap.[23]

The comparison of the estimated prevalence of strategies with the actually chosen strategies allows us to answer several questions. First, if the set of strategies the estimation allow for is the "right" one, does the estimation recover the right proportions. To address this question, the sample is restricted to the cases where the elicited strategies are one of the following: AC, AD, TFT, grim, WSLS, or STFT (the most common strategies). These results are presented in Table 9. The table reports whether the hypothesis that the estimated proportion equal the frequency elicited in the data can be rejected. As can be seen, there are very few cases where the estimated frequency is statistically different from the elicited strategies.

[22] In round one, AD predicts cooperation with probability 0.2, while the other two strategies with probability 0.8. In round 2, all three strategies predict defection with probability 0.8. In round 3, only TFT has cooperation as more likely than defection.

[23] 1000 samples are drawn, an entire session is drawn first, then subjects within that session are drawn (all matches and rounds of each subject).

Table 9: Estimation Performance – Restricted Sample

		AC	TFT	WSLS	Grim	STFT	AD	Gamma
δ = ½, R = 32	Data	0.00	0.06	0.00	0.08	0.17	0.68	
	Estimation	0.00	0.02	0.00	0.00°°°	0.25	0.73	0.26
δ = ½, R = 48	Data	0.03	0.22	0.04	0.38	0.03	0.29	
	Estimation	0.08	0.20	0.01	0.36	0.03	0.32	0.29
δ = 3/4, R = 32	Data	0.03	0.22	0.02	0.07	0.13	0.52	
	Estimation	0.00°°°	0.27	0.00°°°	0.06	0.26	0.41	0.34
δ = 3/4, R = 48	Data	0.14	0.36	0.06	0.35	0.01	0.08	
	Estimation	0.18	0.25	0.00°°°	0.49	0.01	0.08	0.23
δ = 9/10, R = 32	Data	0.02	0.35	0.06	0.13	0.05	0.39	
	Estimation	0.02	0.39	0.01°°	0.16	0.07	0.34	0.29

° Significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.

Clearly the above exercise is made substantially easier by the fact that only a subset of the data is considered. Once the entire data set is taken in consideration, some of the strategies do not correspond to those allowed by the estimator and thus necessarily those choices will be attributed to an incorrect strategy. These results are presented in Table 10. As in the previous table, the frequencies and estimates are reported, but in this case the standard errors are also included.[24]

[24] The appendix includes a table with standard errors for all the estimations preformed.

Table 10: Estimation Performance – Full Sample

		AC	TFT	WSLS	Grim	STFT	AD	Other	Gamma
δ = ½, R = 32	Data	0.00	0.05	0.00	0.06	0.13	0.53	0.22	
	Estimation	0.00	0.04	0.00	0.00°°°	0.28	0.68		0.28
		(0.00)	(0.04)		(0.00)	(0.13)**	(0.14)***		(0.02)***
δ = ½, R = 48	Data	0.03	0.19	0.03	0.33	0.04	0.25	0.13	
	Estimation	0.08	0.08	0.00	0.47	0.02	0.34		0.35
		(0.08)	(0.12)		(0.18)***	(0.05)	(0.13)***		(0.09)***
δ = 3/4, R = 32	Data	0.03	0.18	0.02	0.05	0.11	0.42	0.19	
	Estimation	0.02	0.24	0.00°°°	0.06	0.27	0.40		0.38
		(0.02)	(0.11)**		(0.09)	(0.10)	(0.11)***		(0.05)***
δ = 3/4, R = 48	Data	0.12	0.31	0.05	0.30	0.01	0.06	0.15	
	Estimation	0.22	0.22	0.00°°°	0.45	0.03	0.08		0.27
		(0.14)	(0.26)		(0.28)	(0.04)	(0.05)		(0.05)***
δ = 9/10, R = 32	Data	0.02	0.29	0.05	0.11	0.04	0.32	0.18	
	Estimation	0.00	0.37	0.00°°	0.21	0.09	0.32		0.34
		(0.05)	(0.14)***		(0.13)	(0.07)	(0.12)***		(0.05)***

Bootstrapped standard errors in parenthesis.

° Significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.

* Significantly different from 0 at 10%; ** at 5%; *** at 1%.

The larger the frequency of "Other" strategies, the more difficult it has to be for the estimates to be good. However, in this case the estimates are still fairly close to the elicited frequencies. The estimates picks up the most popular strategies and in most cases the estimates are not statistically different from the elicited strategies. It is also interesting to consider which strategies are statistically different from 0 since when performing such an exercise this would typically be the standard used to determine if a strategy is important or not. In all but one treatment (namely δ = 3/4, R = 48) the two most popular strategies are identified as statistically different from 0. The one exception to this is revealing. In the treatment with δ = 3/4, R = 48 a very high fraction of decisions are to cooperate in every round, consequently it is difficult to distinguish amongst cooperative strategies such as TFT and grim. However, if we test the hypothesis that the fraction of both TFT and grim are equal to 0, we can reject that hypothesis at the 1% level. In fact that hypothesis can be rejected in every treatment where cooperation can be supported in equilibrium.

In the estimation above, although some of the elicited strategies do not correspond to the strategies allowed by the estimation, the most popular ones are all allowed for.

Without prior knowledge, this may be difficult to guarantee. For instance, in Dal Bó and Fréchette (2011), six strategies were considered: AC, AD, grim, TFT, WSLS, and a trigger strategy that starts by cooperating and defects for 2 periods following a defection of the other before returning to cooperation (referred to as T2). Thus, considering this set of possible strategies allow us to study the effect of not having a prevalent strategy in the set of possible ones (STFT) while having one that is not used (T2).

Results for that specification are presented in Table 11. As can be seen by comparing results in Table 11 and 10, in the two treatments where STFT is more important, not including it in the estimation leads to a statistically significant overestimation of the fraction of AD. Having T2 in the set of possible strategies, instead, has no effect on the estimation as it is correctly estimated that no one chooses it.

Table 11: Estimated Performance – Strategy Set from DF 2011

		AC	TFT	WSLS	Grim	STFT	AD	T2	Other	Gamma
δ = ½, R = 32	Data	0.00	0.05	0.00	0.06	0.13	0.53	0.00	0.22	
	Estimation	0.00	0.04	0.00	0.00°°°		0.96°°°	0.00		0.30
δ = ½, R = 48	Data	0.03	0.19	0.03	0.33	0.04	0.25	0.00	0.13	
	Estimation	0.08	0.08	0.00	0.47		0.36	0.00		0.35
δ = 3/4, R = 32	Data	0.03	0.18	0.02	0.05	0.11	0.42	0.00	0.19	
	Estimation	0.02	0.26	0.00°°°	0.06		0.66°°	0.00		0.43
δ = 3/4, R = 48	Data	0.12	0.31	0.05	0.30	0.01	0.06	0.00	0.15	
	Estimation	0.23	0.27	0.00°°°	0.41		0.09	0.00		0.28
δ = 9/10, R = 32	Data	0.02	0.29	0.05	0.11	0.04	0.32	0.00	0.18	
	Estimation	0.00	0.47	0.00°°°	0.18		0.34	0.00		0.36

° significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.

A few more observations are in order. First, the least popular strategy that is identified as statistically significant (Table 10) represents 13% of the elicited strategies. Hence, and to some extent not surprisingly, strategies that are not very popular are difficult to detect. A more interesting observation has to do with the ability to identify strategies in short games. Ex ante it would have seemed reasonable to think that long matches are required to identify strategies, since once needs many transitions to uncover the strategy. With δ = ½, half the games last 1 round, and thus strategy estimation may seem hopeless. However, there is no clear evidence from Table 10 that proportions are estimated any better in the longer games. Given that most strategies used have memory 1, one does not need many interactions to uncover the used strategies.

VII. Conclusions

A growing recent literature studies the strategies used in infinitely repeated games. Several identification hurdles limit the capacity to infer strategies from observed behavior. We overcome these hurdles by asking subjects to design strategies that will play in their place. We find that the strategy elicitation has negligible effects on behavior supporting the validity of this method. We study the strategies chosen by the subjects and find that they include some commonly mentioned strategies, such as tit-for-tat and Grim trigger. However, other strategies which are thought to have some desirable properties, such as win-stay-lose-shift are not prevalent. This last observation is consistent with the Dal Bó and Fréchette (2011) and Fudenberg et al (2010) who, based on the econometric approach described in this paper, find little to no evidence that subjects use WSLS.

As we have shown, in our environment, most subjects seem to rely on strategies that have at most memory 1. We suspect that this result will not hold in general and that in some situations, individuals probably use more complex strategies. Evidence in line with this also comes from Fudenberg et al (2010). They find that subjects mostly use strategies that condition on more than just the last round to support cooperation. However, in most of their treatments where cooperation is an equilibrium, there are no memory 1 strategy that can support cooperation as part of an equilibrium. On the other hand, Aoyagi and Fréchette (2009) who study a different imperfect monitoring technology (and a different stage game), find that in their setting, the evidence is in favor of subjects using a memory 1 strategy. Understanding what features of the environment leads people to use simple or more complex strategy is an interesting question for future work.

In that same vein, we find that the strategies used to support cooperation change with the parameters of the game. One interesting observation is that subjects rely more on TFT as the game becomes longer. This could be rationalized by the subjects thinking of a situation where there is some probability of mistake; as the expected number of rounds increases the expected payoff difference between Grim and TFT becomes larger.

The use of the SPRM to estimate the distribution of strategies results in a very accurate fit. In the full sample, it results estimates of mistakes probabilities between 0..02 and 0.07. Comparing the estimates of the SPRM to the elicited strategies also suggest that it performs reasonably well in terms of identifying the important strategies. However, it can lead to misleading results if one of the key strategies is omitted. It also has difficulty identifying strategies that are present, but only in small amount.

This paper provides a new perspective on the play of the infinitely repeated prisoner's dilemma by allowing us to observe the strategies, as opposed to simply the choices, used by subjects. Much more work remains to be done to develop and broader and deeper understanding of how individuals approach such environments. In particular, what determines the choice of specific strategies in a given environment is an interesting question to be pursued further.

References

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. 1980a. "Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution*, 24(1): 3-25.

Axelrod, Robert. 1980b. "More Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution*, 24(2): 379-403.

Baker, George, Robert Gibbons and Kevin J. Murphy. 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics*, 117(1): 39-84.

Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo. 2010. "Co-operation in Infinitely Repeated Games: Extending Theory and Experimental Evidence." *American Economic Journal : Microeconomics*, forthcoming.

Brandts, J. and Charness, G. 2000. "Hot vs. Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*. 2, 227–238.

Brosig, Jeanette, Weimann, Joachim and Yang, Chun-Lei. 2003. "The Hot Versus Cold Effect in a Simple Bargaining Experiment." *Experimental Economics*, 6:75–90.

Camera, Gabriele and Marco Casari. 2009. "Cooperation among strangers under the shadow of the future." A*merican Economic Review* 99(3), 979–1005.

Camera, Gabriele, Marco Casari, and Maria Bigoni. 2010. "Cooperative Strategies in Groups of Strangers: An Experiment," Purdue University Economics Working Papers 1237.

Cason, Timothy and Vai-Lam Mui. 2008. "Coordinating Collective Resistance through Communication and Repeated Interaction." Purdue University. Mimeo.

Cwojdzinski, Lisa V. and Ulrich Kamecke. Year unknown. "Infinity in the lab: The role of the continuation rule for cooperation in repeated games." Mimeo.

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 and Guillaume R. Fréchette. 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review*, 101(1): 411-429.

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.

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.

Friedman, James W. 1991. "A Non-Cooperative Equilibrium for Supergames." *Review of Economic Studies*, 38(1), 1–12.

Fudenberg, Drew, and David K. Levine. 1997. "Measuring Players' Losses in Experimental Games." *Quarterly Journal of Economics*, 112(2):507-536.

Fudenberg, Drew, David G. Rand, and Anna Dreber. 2010. "Slow to Anger and Fast to Forget: Leniency and Forgiveness in an Uncertain World." American Economic Review (forthcoming).

Green, Edward J. and Robert H. Porter. 1984. "Noncooperative Collusion Under Imperfect Price Information." *Econometrica*, 52(1): 87-100.

Gueth, W., Huck, S., and Rapoport, A. 1998. "The Limitations of the Positional Order Effect: Can it Support Silent Threats and Non-Equilibrium Behavior?" *Journal of Economic Behavior and Organization,* 34, 313–325.

Hoffman, E., McCabe, K.A., and Smith, V.L. 1998. "Behavioral Foundations of Reciprocity: Experimental Economics and Evolutionary Psychology." *Economic Inquiry*, 36, 335–352.

Holt, Charles A. 1985. "An Experimental Test of the Consistent-Conjectures Hypothesis." *American Economic Review*, 75(3): 314–325.

Imhof , Lorens A., Drew Fudenberg, and Martin A. Nowak. 2007. "Tit-for-Tat or Win-Stay, Lose-Shift?" *Journal of Theoretical Biology*, 247: 574-580.

Klein, Benjamin and Keith B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." Journal of Political Economy, 89(4): 615-641.

Murnighan, J. Keith, and Alvin E. Roth. 1983. "Expecting Continued Play in Prisoner's Dilemma Games." *Journal of Conflict Resolution*, 27(2): 279-300.

Normann, Hans-Theo, and Brian Wallace. 2006. "The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment." Royal Holloway. Mimeo.

Palfrey, Thomas R., and Howard Rosenthal. 1994. "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies*, 61(3): 545-565.

Phelan, Christopher and Ennio Stacchetti. 2001. "Sequential Equilibria in a Ramsey Tax Model." *Econometrica*, 69:1491-1518.

Rotemberg, Julio J. and Saloner, G. 1986. "A Supergame-Theoretic Model of Price Wars During Booms." *American Economic Review*, 76: 390–407.

Rotemberg, Julio J. and Woodford, M. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy*, 100: 1153–1207.

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. 1998. *Modeling Bounded Rationality*. Cambridge: MIT Press.

Schwartz, Steven T, Richard A. Young, Kristina Zvinakis. 2000. "Reputation Without Repeated Interaction: A Role for Public Disclosures." *Review of Accounting Studies*, 5(4): 351-375.

Appendix A: Tables

Table A1: Session characteristics

Variable	δ = ½		δ = ¾		δ = 9/10
Payoff from cooperation	**32**	**48**	**32**	**48**	**32**
Number of subjects	18	16	14	18	18
Number of Games	86	63	50	44	25
Phase 1	37	25	16	14	5
Phase 2	35	23	20	16	7
Phase 3	*14*	15	*14*	*14*	*13*
Number of subjects	16	18	16	12	18
Number of Games	82	92	60	54	31
Phase 1	35	46	25	24	8
Phase 2	32	33	21	17	10
Phase 3	15	*13*	*14*	*13*	*13*
Number of subjects	16	22	14	16	14
Number of Games	72	70	43	55	38
Phase 1	24	28	17	23	12
Phase 2	33	28	11	18	13
Phase 3	15	14	15	14	*13*

Note: Italics indicate a phase that started midway through a match.[25]

[25] In phase 3 there should always be 14 or 15 matches (depending on whether the match started in round 1 or midway through a match). The few sessions with 13 complete matches are the results of a parameter in the software that was inadvertently limiting the total number of match/rounds. This was corrected in the additional sessions.

Table A2: Equilibrium Strategies for Memory 1 Strategies

		δ = ½		δ = ¾		δ = 9/10
Strategy	AKA	**32**	**48**	**32**	**48**	**32**
CCCCC	AC					
CCCCD	AC'					
CCCDC			NE		NE	
CCCDD	TFT		NE		NE	NE
CCDCC	AC'					
CCDCD	AC'					
CCDDC	WSLS		SGPE		SGPE	
CCDDD	Grim		SGPE	SGPE	SGPE	SGPE
CDCCC						
CDCCD						
CDCDC						
CDCDD						
CDDCC						
CDDCD						
CDDDC						
CDDDD						
DCCCC						
DCCCD						
DCCDC			NE		NE	
DCCDD	STFT	NE				
DCDCC						
DCDCD	AD'	NE	NE	NE	NE	NE
DCDDC			SGPE		SGPE	
DCDDD	AD'	NE	SGPE	SGPE	SGPE	SGPE
DDCCC						
DDCCD						
DDCDC						
DDCDD		NE	NE			
DDDCC						
DDDCD	AD'	NE	NE	NE	NE	NE
DDDDC						
DDDDD	AD	SGPE	SGPE	SGPE	SGPE	SGPE

Note: The letters in the strategy names denote the recommended action after each possible contingency: initial period, CC, DC, CD and DD. SGPE for strategy "s" denotes that (s,s) is a sub-game perfect equilibrium,NE denotes that the strategy is a NE but not SGPE.

Table A3: Additional Session Characteristics

Variable	δ = ½	δ = ¾		δ = 9/10
Payoff from cooperation	**48**	**32**	**48**	**32**
Number of subjects	18	16	14	20
Number of Games	42	41	44	30
Phase 1	19	13	17	6
Phase 2	8	14	13	9
Phase 3	15	*14*	*14*	15
Number of subjects	18	12	16	12
Number of Games	67	49	41	30
Phase 1	33	19	19	9
Phase 2	20	16	8	7
Phase 3	*14*	*14*	*14*	*14*
Number of subjects	14	14	14	14
Number of Games	63	43	48	31
Phase 1	29	18	22	9
Phase 2	19	11	12	8
Phase 3	15	*14*	*14*	*14*

Note: Italics indicate a phase that started midway through a match.

Table A4: Cooperation Rate by Treatment, Phase and Elicitation of Strategies (in Additional Sessions)

		First Rounds Only						All Rounds					
		Phase 1			Phase 2			Phase 1			Phase 2		
		Elicitation of Strategies			Elicitation of Strategies			Elicitation of Strategies			Elicitation of Strategies		
			Original	Additional		Original	Additional		Original	Additional		Original	Additional
δ	R	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Panel A: All Matches													
½	32	0.14	0.12		0.06	0.09		0.12	0.09		0.08	0.07	
	48	0.37	0.61	0.41	0.41	0.60	0.70	0.35	0.56	0.34	0.36	0.52	0.62
¾	32	0.25	0.23	0.24	0.26	0.24	0.39	0.18	0.19	0.23	0.23	0.21	0.28
	48	0.75	0.72	0.66	0.96	0.86	0.90	0.65	0.61	0.58	0.90	0.80	0.88
9/10	32		0.41	0.32		0.60	0.54		0.44	0.25		0.53	0.37
Panel B: Last Match in Each Phase													
½	32	0.14	0.04		0.02	0.06		0.10	0.02		0.10	0.06	
	48	0.39	0.45	0.44	0.41	0.59	0.62	0.39	0.31	0.30	0.39	0.38	0.62
¾	32	0.25	0.23	0.31	0.30	0.23	0.33	0.16	0.22	0.30	0.16	0.20	0.30
	48	0.89	0.67	0.82	0.98	0.83	0.91	0.77	0.57	0.57	0.77	0.86	0.94
9/10	32		0.58	0.33		0.62	0.52		0.59	0.28		0.55	0.27

Appendix B: Figures

Figure A1: Evolution of Cooperation (first rounds)

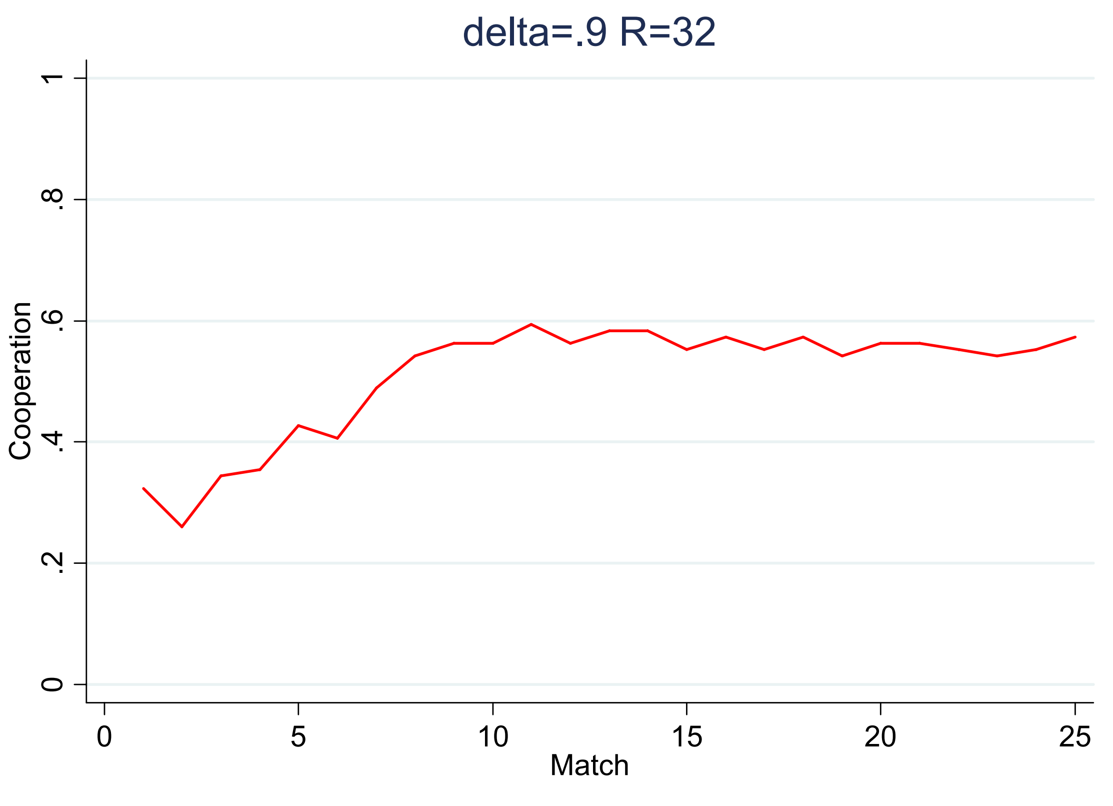

Figure A2: Evolution of Strategies (phase 2)

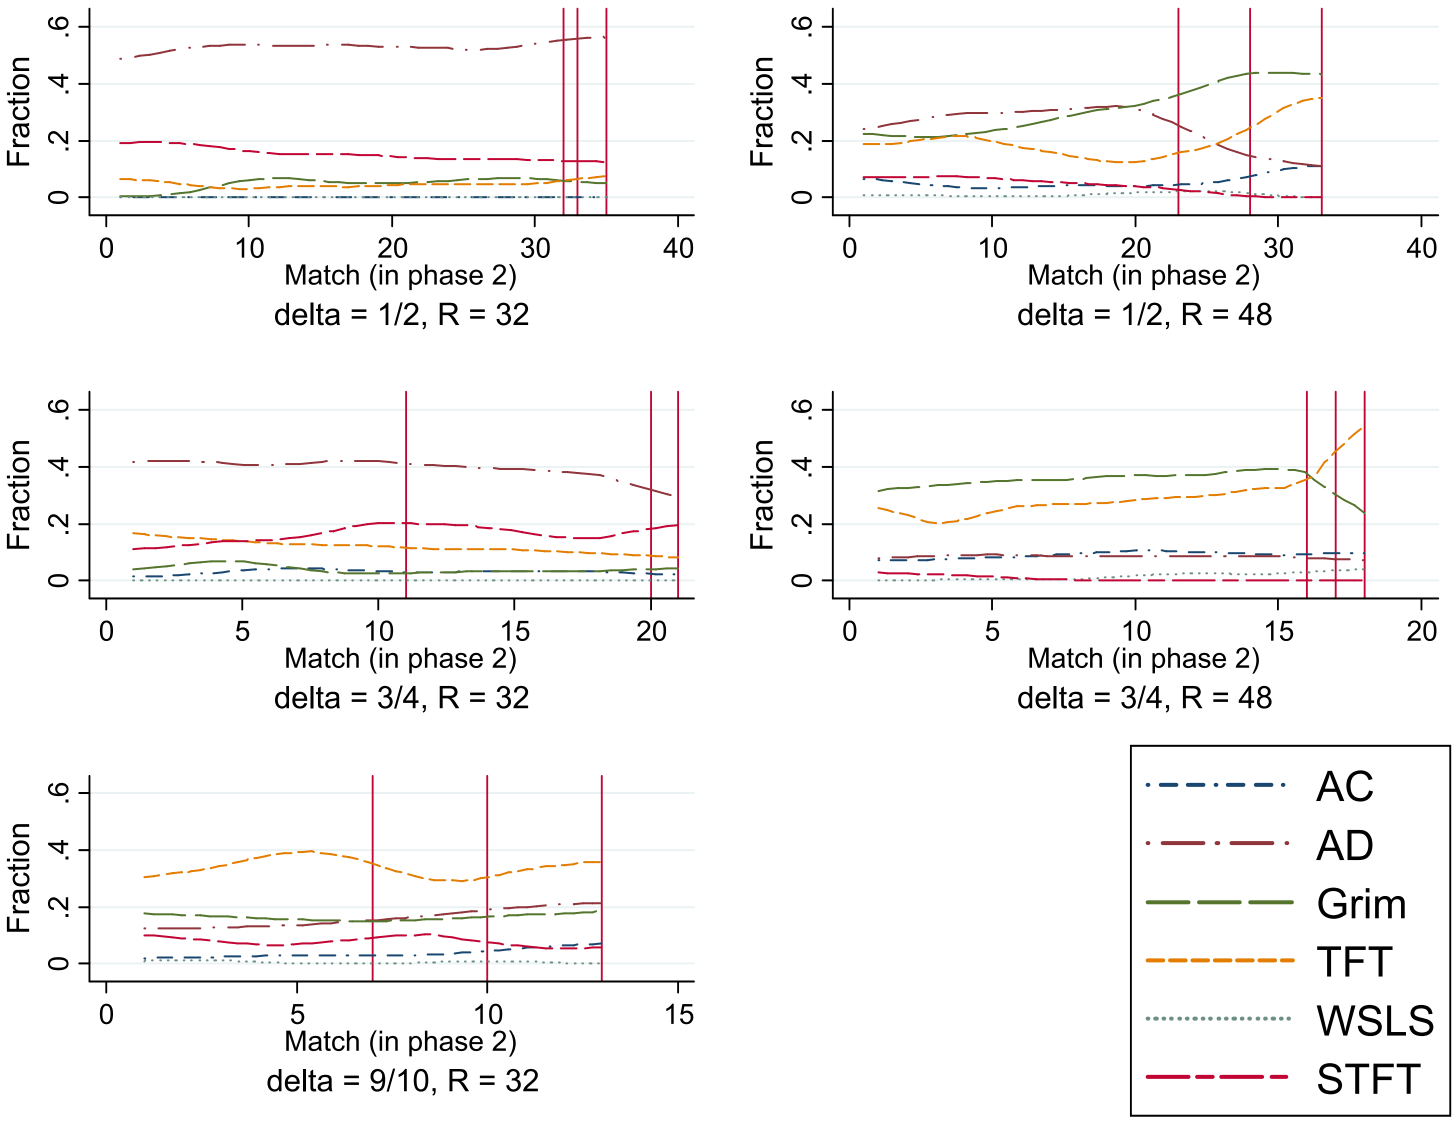

Note: the vertical red lines indicate the end of the phase 2 in each session.

Appendix B: Options in the sessions with a menu of strategies.

(what is in parentheses was not presented to the subjects):

- Select 1 in every round. (AC)
- Select 2 in every round. (AD)
- Select 1 for X rounds, then select 2 until the end. (CD-X)
- Select 1 X% of the time and 2 1-X% of the time. (RANDOM-X)
- In round 1 select 1. After round 1: if both always selected 1 in previous rounds, then select 1 otherwise select 2. (GRIM)
- In round 1 select [1 or 2]. After round 1: if the other selected 1 in the previous round, then select 1. Else if the other selected 2 in the previous round, then select 2. (TFT or STFT)
- In Round 1 select [1 or 2]. After round 1: if both made the same choice (both selected 1 or both selected 2) in the previous round, then select 1. Otherwise select 2. (WSLS or D WSLS)
- In round 1 select 1. After round 1: if the other or me selected 2 X times before then select 2. Otherwise select 1. (GRIM-X)
- In round 1 select 1. After round 1: Select 2 if other selected 2 in all of the previous X [select number] rounds, Otherwise select 1. (TFXT)
- Starts by selecting 1, and keeps selecting 1 until someone selects 2, in that case selects 2 for X periods and then goes back to select 1 until someone selects 2 again, and so on. (T-X)
- In Round 1 select [1 or 2]. After round 1: if the other selected 1 in the previous round then select 1. In round 2 select 1 if other selected 1 in the previous round, and select 2 if other selected 2 in the previous round. If the other selected 2 in the previous round and you selected 1 two rounds prior, then select 2. Else, if the other selected 2 in the previous round and you selected 2 two rounds prior, then select 1. (CTFT or D CTFT).
- Build your own. (This offers the same option as in the memory-1 treatment.)

When [1 or 2] is an option, it was presented as a drop-down menu, and when X needs to be specified, subjects could enter a number in the appropriate box.