Strategy Choice In The Infinitely Repeated Prisoners' Dilemma

Pedro Dal Bó
Brown University
& NBER

Guillaume R. Fréchette*
New York University

July 2013

Abstract

We use a novel experimental design to identify the subjects' strategies in an infinitely repeated prisoners' dilemma experiment. We ask subjects to design strategies that will play in their place. We find that eliciting strategies has negligible effects on their behavior, supporting the validity of this method. We find the chosen strategies include some common ones such as Tit-For-Tat and Grim trigger. However, other strategies that are considered to have desirable properties, such as Win-Stay-Lose-Shift, are not prevalent. We also find that the strategies used to support cooperation change with the parameters of the game. Finally, our results confirm that long-run misscoordination can arise.

Keywords: infinitely repeated games, prisoner's dilemma, cooperation, strategies, experimental economics.

---

* 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, Boston College, HBS, Maastricht University, Bocconi University, CIDE, UBC, Sciences Po, Carlos III, CEMFI and at the ESA conference in Tucson, the Conference on Social Dilemmas, the SAET conference, and the Conference on New Directions in Applied Microeconomics for very useful comments. We thank Emanuel Vespa and Sevgi Yuksel for their research assistance, CASSEL (UCLA) and SSEL (Caltech) for the Multistage software, as well as Rajeev Advani and Anwar Ruff for its adaptation. Fréchette gratefully acknowledges the support of NSF via grants SES-0924780 and SES-1225779 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.

I. Introduction

The theory of infinitely repeated games has been an 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. While a series of experiments has supported this theory, less is know about the types of strategies people actually use to overcome opportunistic behavior.[2]

Learning which strategies are actually used is of interest on many levels. First, it can help future theoretical work to identify refinements or conditions that lead to these strategies being the ones that are played. Rubinstein (1998, p.142) argues that folk theorems are statements about payoffs and that more attention could be paid to the strategies supporting those payoffs: "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 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 possible strategies, researchers choose to focus on finite subsets of strategies that include those usually studied by theorists (e.g., Imhof, Fudenberg and Nowak 2007). Identifying strategies in the laboratory that are popular

[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 Rotemberg and Woodford 1990) to name just 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) studies a Cournot duopoly that 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, Ockenfels, and Spagnolo (2010) and Fréchette and Yuksel (2013), 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 (2012) study infinitely repeated prisoners' dilemma under imperfect public monitoring while Rand, Fudenberg, and Dreber (2013) explore that environment further by separately identifying the roles of noise and information on behavior. Duffy and Ochs (2009), Camera and Casari (2009) and Camera, Casari, and Bigoni (2012) study repeated prisoners' dilemma with random matching. Finally, Cason and Mui (2008) study a collective resistance game; Fudenberg, Rand, and Dreber (2011) correlate behavior in the dictator game to behavior in an infinitely repeated game; and Cabral, Ozbay, and Schotter (2010) study reciprocity.

with humans can provide an appealing basis for selecting 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 that 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 can. 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 subgame 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 observe only one actual 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, but 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 2012; for a related but Bayesian approach to this, see Engle-Warnick, McCausland, and Miller 2004). A second approach uses maximum likelihood estimation either to estimate the prevalence of each strategy in the set under consideration (see, Dal Bó and Fréchette 2011, Fudenberg, Rand, and Dreber 2012, Camera, Casari, and Bigoni 2012, Vespa 2013, Fréchette and Yuksel 2013, and Rand, Fudenberg, and Dreber 2013) or to estimate 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 studying strategies: the elicitation of strategies (i.e., a modified strategy method (Selten 1967)).[3] 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. We introduce our modified strategy method

[3] Bruttel and Kamecke (2012) provide a partial elicitation of infinitely repeated game strategies. They are not interested in strategies, however; their motivation is to find alternative ways of inducing repeated games in the laboratory.

in ways that limit the chances that the procedure will affect behavior.[4] Our method first allows subjects to play as they would normally to learn what they want to do. We then introduce a hybrid period in which they specify a strategy and receive feedback to help them understand how their strategy would play out, although their choices are still made decision-by-decision. After that, the strategy method is fully implemented. We show that our procedure does not affect behavior by, among other things, 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 of 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 choose only 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% are part of Nash equilibrium.

As far back as Axelrod (1980b), papers have also elicited strategies in infinitely repeated games and then made them compete in computer tournaments. The focus of this literature was on the relative performance of the different strategies. Since then, most of that literature has moved towards simulations rather than tournaments.[5] In addition to focus on a different population and to vary a set of important parameters, our paper shows that the elicitation of strategies has no major distorting impact on choices.

Following the section on experimental design, the results section will cover four broad areas. First, we will present the results pertaining to the new method we propose.

---

[4] Hoffman, McCabe, and Smith (1998), Gueth, Huck, and Rapoport (2001), and Brosig, Weimann, and Yang (2003) report evidence that the strategy method may affect behavior. Brandts and Charness (2000) find the opposite.

[5] 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).

This will be followed by an analysis of the strategies used by subjects. Then, we discuss the limitations imposed by focusing on strategies with memory one and we will study a set of additional sessions that relaxes that constraint. Finally, we will evaluate a method to recover strategies econometrically.

II. Experimental Design

The experimental design is in three phases.[6] In all phases, subjects participate in randomly terminated repeated prisoners' dilemma games. A repeated game or supergame is referred to as a match and consists of multiple rounds. After each match, subjects are randomly rematched with a subject.

In Phase 1 subjects simply play the randomly terminated games. Between matches, they are reminded of the decisions they made 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 four questions after round 1 appear is randomized. After having specified their plan of action, subjects then play the match just as in Phase 1, making decisions in every round. At this point, the plan of action they specified is irrelevant. After the first match, they are shown the decisions they made in this match, the decisions the person they were matched with made, and the decisions that would have been made according to the plan of action they specified, had it been played instead. 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 feedback; and specify a plan) is repeated for 20 minutes. After 20 minutes of play in Phase 2, the plan of action takes over

[6] When first reading the instructions, subjects are informed that there are multiple phases, but they are only told the procedures for Phase 1. They receive additional instructions after Phase 1. All instructions and screen-shots are available in the online appendix at https://files.nyu.edu/gf35/public/print/Dal_Bo_2013b_oa.pdf

for the subjects, finishes the ongoing match, and plays 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

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. Three sessions are conducted per treatment.[7] Payments are based on the sum of the points accumulated in the three phases converted to dollars, with 100 points equaling $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 (cooperate 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.[8]

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 the extent to which 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.

[7] Two previous sessions were conducted for the treatment with R = 48 and δ = 3/4: however, the payments were too low and, thus, the exchange rate was changed, and those two 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.

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

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 subjects an opportunity to determine the situations in which the specified plan of action is not doing what they actually want to do. Finally, subjects are incentivized to choose the plan of action that is closest to what 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 that 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. These strategies include those that are independent of the past, such as Always Defect (AD) and Always Cooperate (AC), and memory 1 strategies for which behavior depends only on the past round, such as Tit-For-Tat (TFT) and Win-Stay-Lose-Shift (WSLS) and others. This method also allows subjects to construct plans of action that will behave exactly as some strategies with memory greater than one. For example, the plan of action that starts cooperating and will keep cooperating if there was mutual cooperation in the previous round functions exactly as the strategy Grim in this setup without mistakes in implementation.[9] In conclusion, many of the strategies most often mentioned in the literature can be specified in this way: TFT, Grim, WSLS, and others are available – these (and other) strategies will be defined later. To address the concern that the set of strategies may be too restrictive, we will compare decisions in Phase 2 with the decisions of the plan of action. Any persistent differences would suggest that no available strategy was entirely 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 that allowed for additional strategies. We will describe these later in the paper.

[9] Since Grim can be implemented by this memory 1 plan of action, we will call both "Grim."

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 δ=1/2, δ=3/4, and δ=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, the 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 δ =1/2 and R=32, it experiences large increases under δ =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), in which 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)

Dal Bó and Fréchette (2011) uses an experimental design that is identical to the one in this paper, with the exception that strategies are not elicited. Therefore we can compare cooperation rates in both experiments to test 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. The main observation is that the evolution of behavior is extremely similar in 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.

[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 made in the previous match between matches in Dal Bó and Fréchette (2011).

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

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

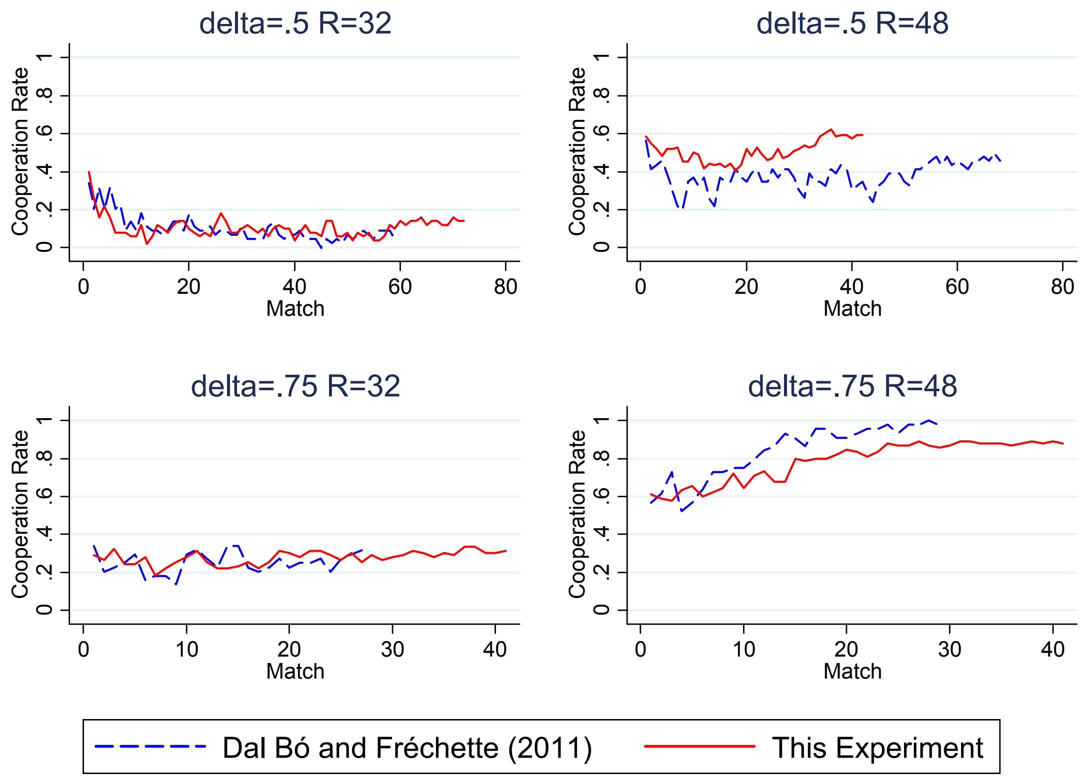

Table 2 shows the cooperation rates for both series of experiments separated by phase (as Dal Bó and Fréchette 2011 does not have phases, matches were assigned to Phase 1 and Phase 2 based on whether they started before or after the mid-point of the session: the number of Phase 1 and 2 matches in the 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).[12] The difference in the round-1-only case is driven by the δ=3/4 and R=48 treatment since, 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

[12] Unless otherwise noted, statistical significance is assessed by estimating a probit and clustering the variance-covariance at the level of the experimental session. For a discussion of potential sources of session-effects, see Fréchette (2011).

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 in 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			
Panel A: All Matches									
		Elicitation of Strategies				Elicitation of Strategies			
		Yes		No*		Yes		No*	
		Phase		Phase°		Phase		Phase°	
δ	R	1	2	1	2	1	2	1	2
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.

The rate of deviations between plans of action and actual choices in Phase 2 is also relatively low. Only 7% of all choices (across all 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 would you do in the first round? What would you 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 estimates from a probit of the 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 Plans of action

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 δ=1/2 and R=32 treatment and highest in the δ=9/10 and R=32 treatment. Figure 2 plots the frequency with which subjects changed their plan of action from one match to the next. As the figure shows, in every treatment, the evolution displays a downward trend, suggesting that subjects' behavior is stabilizing.

Figure 2: Frequency of Change in Plan of Action

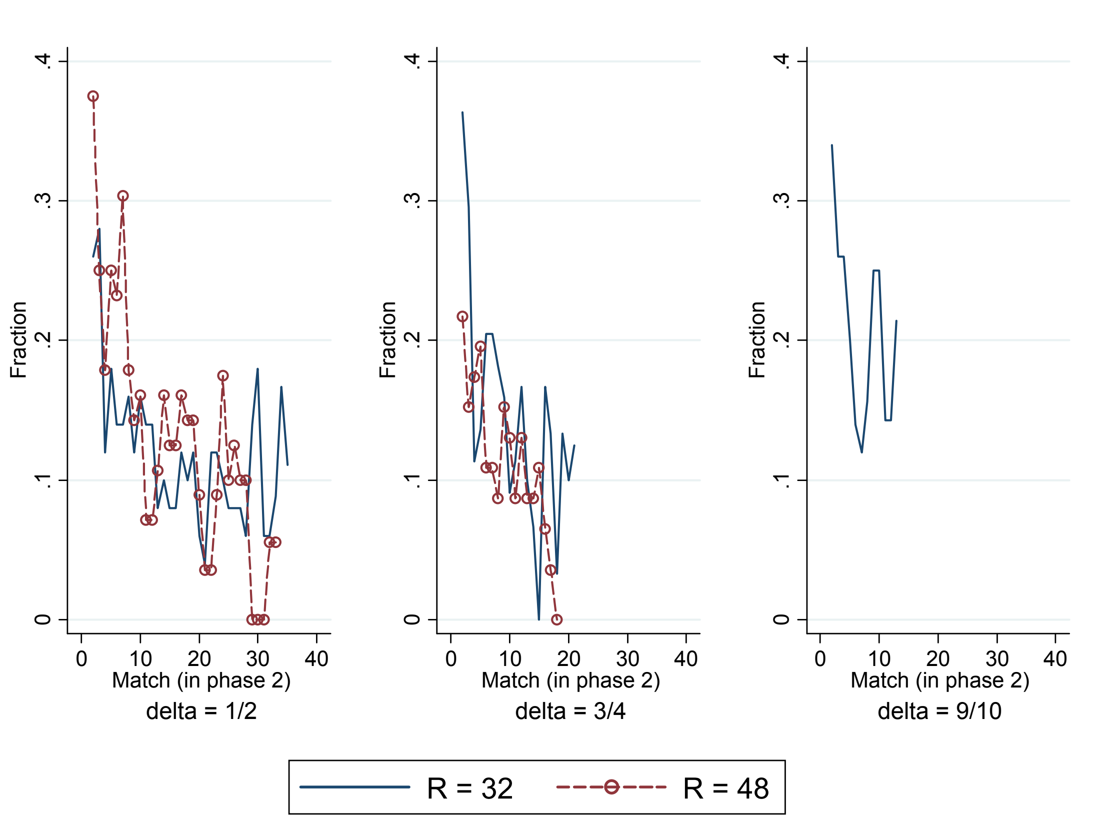

The frequency with which subjects change their plan of action is markedly higher following a match in which their decisions differed from those that their specified plan of action would have made. This can be seen in Table 4, which breaks down the frequency of changes by whether the choices in the previous match corresponded to the choices the plan of action would have made.

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 plans of action 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 had occurred, choices in this experiment would be systematically different from those in the Dal Bó and Fréchette (2011) experiment, and that is not 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 now describe the strategy choices that the subjects made. After first defining some strategies of interest, we describe the strategy choices at the end of Phase 2 and then the evolution leading to those choices.

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, cooperates as long as both players have cooperated in the last round, and defects otherwise.[13] A strategy that is

[13] 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.

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. WSLS is considered desirable because when it plays itself, it would not defect forever after a deviation. Finally, another strategy of interest is suspicious Tit-For-Tat (STFT), which starts by defecting and, from then on, matches what the other subject did in the previous round.

What can be expected? For each treatment, Table A2 in the Appendix shows 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 in which 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. However, 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 lower than the payoff from mutual cooperation. This gives subjects 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, while WSLS, which has performed well in simulations, will. Unlike TFT, WSLF is an 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. However, 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 treatments. Strategies are described by the string of five letters, with C for cooperation and D for defection: 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, such as AC stands for Always Cooperate and AD for Always Defect.

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 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.

If we aggregate strategies according to whether they lead to cooperation or defection on their path of play, we find a similar pattern.[14] 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 true for increases in δ. To see this, note that when delta = ½, increasing R doubles the popularity of TFT, while the popularity of Grim increases by a factor of about six. 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 does. This suggests that how subjects choose to support cooperation depends on the particular parameters of the game in a systematic way — that is, as the expected length of the game increases, subjects choose to use shorter punishments.

[14] 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).

Table 5: Distribution of Elicited Strategies (Last Match)

Strategy*					AKA	δ = ½ R 32	δ = ½ R 48	δ = ¾ R 32	δ = ¾ R 48	δ = 9/10 R 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
Only NE						28.00	12.50	0.00	34.78	44.00
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.
Subgame 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 round, CC, DC, CD and DD, where the second letter designates the other's choice.

Interestingly, a large proportion of the strategies being chosen are not part of subgame perfect equilibria (SGPE). In particular, the proportion of strategies that conform to a SGPE when playing against itself reaches a low of 32% under δ=9/10 and R=32. The maximum is reached under δ=1/2 and R=48, with 73% of strategies being SGPE. The low prevalence of SGPE strategies is related to two observations. First, under δ=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 equilibrium.

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 exception of WSLS because, as Table 5 shows, 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 to 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		p-v
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	0.30
TFT	8	6	0.43	21	13	0.24	18	11	0.35	28	33	0.54	28	42	0.00
Grim	0	6	0.00	29	36	0.00	2	5	0.39	30	39	0.01	20	12	0.00
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	0.00
Only NE	38	28	0.07	23	13	0.24	0	0	-	33	35	0.80	28	44	0.00

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 match.

When it comes to the evolution of plans of action, 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 plans of action 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 three out of five 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. The Grim strategy, however, 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 average cooperation rates. For example, under $\delta=1/2$ 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 (consistent with Dal Bó and Fréchette 2011).

IV.c. Expected Payoffs

To better understand strategy choice, it is helpful to calculate the expected payoffs of different plans of action given their distribution. 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, an SGPE, or neither. The figure also indicates the name of the plan of action for plans of action that 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 worse than the best-performing one. Second, the most popular plans of action tend to be close to the best-performing plans of action. 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 Plans of Action (last round of Phase 2)

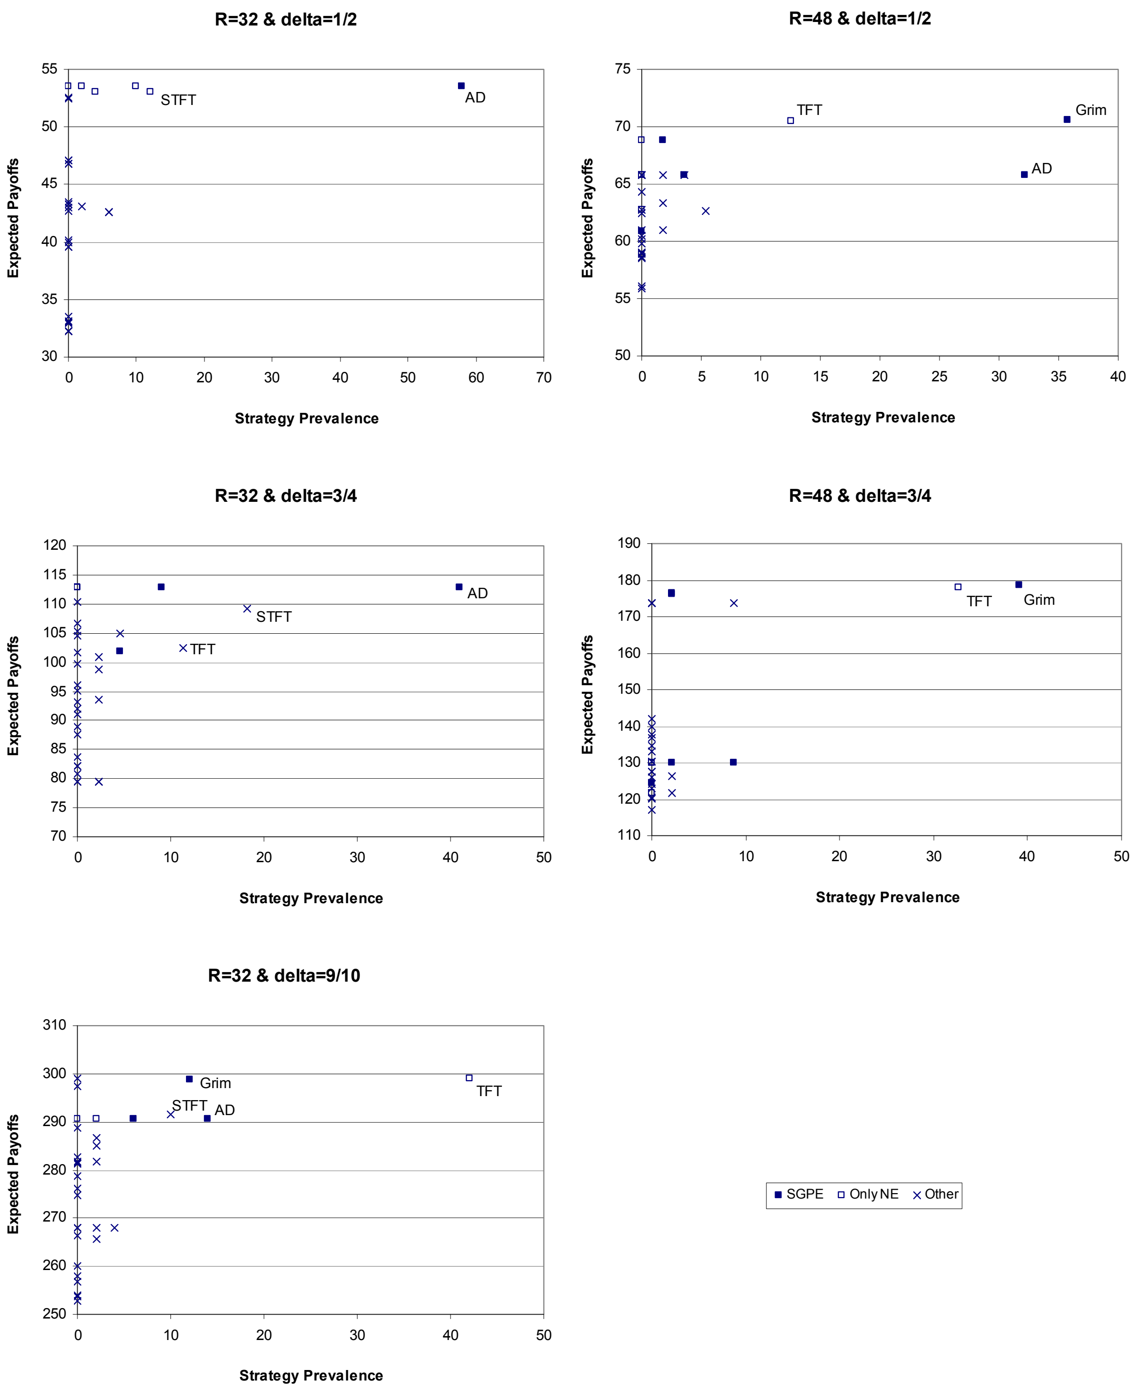

Table 7 provides a calculation of lost payoffs due to suboptimal strategy choice. The first row shows, by treatment, the maximum expected payoff that can be obtained by one of the memory-1 strategies, given the distribution of strategies in the experiment. The second row shows the average expected payoff 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, on 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 plans of action have high expected value. However, Table 7 also reveals that there are incentives to move away from the least profitable plans of action, as the plan with the lowest expected payoffs that was selected always implied a loss of at least 13% — and in some treatments substantially more than the best-performing plan of action. If one considers the worst possible plan of action (which, in most treatments, is never selected), then the incentives would appear even starker. Note, however, that Table 7 and Figure 3 suggest that the financial incentives to move away from plans of action that are NE but not SGPE, to others that are SGPE, were often small or non-existent because some NE plans performed very well in expected-value terms.[15]

V. Is Memory 1 Enough?

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

In this section, we present strategy choices when subjects have a greater set of strategies — a "menu" — from which to choose. 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, with X being a parameter selected by the subject. Similarly, there is 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 defections before reverting back to cooperation, as well as a few other strategies. The complete list of strategies can be found in Appendix C. 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 126 NYU undergraduates participated in these eight sessions, with an average of 16.10 subjects per session, a maximum of 20 and a minimum of 12. [16] The

[15] 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 that 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).

[16] An earlier version of the paper had 12 different sessions with this design. However, a coding error resulted in some of the strategies not being implemented properly. The software was fixed with only minor

subjects earned an average of \$29.37, with a maximum of \$43.45 and a minimum of \$19.74. In the treatments with δ=1/2, δ=3/4, and δ=9/10, the average number of rounds per match was 1.92, 3.75, and 10.93, respectively, and the maximum was 8, 15, and 44, 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; and Table A5 provides an analysis of the statistical significance of the differences. We find that there are no statistically significant differences in cooperation rates between the original sessions and the additional sessions. This suggests that increasing the set of possible strategies that subjects can choose does not affect cooperation rates. However, we find that in two of the treatments, cooperation rates in Phase 2 are significantly different in the additional sessions relative to the sessions without strategy elicitation from Dal Bó and Fréchette (2011). In one of the treatments, the difference appears already before the elicitation of strategies starts (Phase 1), so that the difference in behavior cannot be attributed to elicitation.[17]

Table 8 reports the percentages of each strategy in the last match of Phase 2. While 76% of the subjects chose a strategy from the menu and not from the original mechanism to specify strategies, most of the menu strategies selected are equivalent to those available in the original set (i.e., AD, TFT, Grim, WSLS, among others). The main result from these additional sessions is that 91% of the final strategies could be defined using the apparatus of the original sessions. More specifically, 97% of the strategies in the last match for δ=1/2 and R=32 can be defined using the original set of strategies; 96% and 87% when δ=3/4 and R=32 or 48 respectively; and 85% for δ=9/10 and R=32. Furthermore, in all treatments, the most popular strategy that cannot be expressed using the original method ranks among the least popular strategies. In only one treatment (δ=9/10 and R=32), where it ranks fifth, it is not the least popular. The few selected

modifications to improve the description of the strategies, and one strategy was dropped because the description was too complicated (and was almost never selected). Results are essentially unchanged and available upon requests.

[17] For each treatment, we estimate a probit of cooperation on a dummy variables for the additional sessions, and we compare them with the sessions with only memory one strategies first. We then compare the additional sessions with the sessions from Dal Bó and Fréchette (2011), which do not have strategy elicitation. We consider the first round of the last match in phases 1 and 2. Results are robust to considering all rounds and all matches.

strategies with memory greater than one are the following. GRIM-2 and GRIM-4 are grim trigger strategies that start punishments after two and four deviations respectively. TF2T is a Tit-For-Tat that counts the other's defections and defects only after two rounds of defections by the other. Finally, T-2, T-10 and T-12 are trigger strategies that punish for two, ten and 12 rounds, respectively.

Many of the results from the original session carry over. In particular, AD is still a popular strategy. TFT and Grim are the most common strategies to support cooperation. Moreover, as before, these three strategies account for the majority of strategy choices in all treatments. It is also still the case that the strategies used to support cooperation change with the parameters R and δ. And while the results are less clear than for the original sessions, it is still the case that the prevalence of TFT relative to Grim increases with δ. Finally, as before, some non-SGPE are popular; however many of the choices favor NE strategies.

One change is the popularity of STFT, which is substantially decreased in these additional sessions. Despite this small difference, 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, 91% of strategies in the last match can be implemented using strategies with memory 1 or less.

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

		δ = ½	δ = ¾		δ = 9/10
Strategies Using Original Interface			R		R
	AKA	48	32	48	32
C C C D D	TFT	2.94		13.33	11.76
C C D D C	WSLS				2.94
C C D D D	GRIM	**23.53**	7.14	**3.33**	
C D C D D					2.94
C D D D D		2.94			
D C C C D					2.94
D C D D D	AD'		**7.14**		
D D D D D	AD		**10.71**		**2.94**
Total		29.41	25.00	16.67	23.53
Menu Strategies					
AD		**14.71**	**35.71**	**23.33**	**14.71**
CD-1 (CDDDD)			3.57		
CD-3				3.33	
DWSLS			3.57	**3.33**	
GRIM		**23.53**	**21.43**	**26.67**	**20.59**
GRIM-1 (GRIM)		**2.94**		**3.33**	
GRIM-2				3.33	5.88
GRIM-4			3.57		
STFT					11.76
TFT		23.53	7.14	10.00	11.76
TF1T (TFT)		2.94			
TF2T				3.33	2.94
T1					2.94
T2		**2.94**			2.94
T10					**2.94**
T12				**3.33**	
WSLS				**3.33**	
Total		70.59	75.00	83.33	76.47
Original or Equivalent Strategies		97.06	96.43	86.67	85.29
Cooperative		82.35	39.29	70.00	64.71
Defecting		14.71	53.57	23.33	20.59
SGPE		67.65	75.00	63.33	41.18
Only NE		29.41	0	23.33	23.53
NE		97.06	75.00	86.67	64.71
Grim		50.00	28.57	33.33	20.59
TFT		29.41	7.14	23.33	23.53

Note: Menu strategies in italics denote that they can be generated using the original strategies. 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. Subgame 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 round, CC, DC, CD and DD, where the second letter designates the other's choice.

Using econometric estimation, previous papers have studied the complexity of the strategies used in infinitely repeated games. On the one hand, Fudenberg et al. (2012) find that under imperfect public monitoring many subjects use strategies that condition on more than just the last round to support cooperation. 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. Note, also, that when Fudenberg, Rand, and Dreber (2012) estimate strategies for their one treatment without noise, they find that 31% of strategies use more than two states — lower than for their treatments with noise but higher than that found for the treatments in this paper and in Dal Bó and Fréchette (2011) as well as Fréchette and Yuksel (2013) or Dreber, Fudenberg, and Rand (2008).[18] However, for one of the treatments in Dal Bó and Fréchette (2011) that is not repeated here (namely $\delta = 3/4$, R = 40), Fudenberg, Rand, and Dreber (2012) report that 57% of strategies are estimated to have more than two states. This suggests that in games with perfect monitoring, strategies with more than two states are not the norm, but that for some specific combination of parameters they are more prevalent.

VI. Does econometric estimation recovers the same strategies?

Using data from Phase 2, we can evaluate the extent to which 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 that the subjects actually chose.

We study the performance of the estimation procedure proposed in Dal Bó and Fréchette (2011) and also used in Fudenberg, Rand, and Dreber (2012), Camera, Casari, and Bigoni (2012), Vespa (2013), Fréchette and Yuksel (2013), and Rand, Fudenberg, and Dreber (2013).[19] Let us refer to this approach as the Strategy Frequency Estimation Method (SFEM). The SFEM is as follows. By s^k_{imr}, denote the choice that subject *i* would

[18] See Tables 5 and A3 of Fudenberg et al. (2012). In Dreber et al. (2008), there are 0% of strategies with more than two states in both treatments, while in Dal Bó and Fréchette (2011), there are between 0% and 23% for the four relevant treatments.
[19] Note that the notation used here differs slightly.

make in round r of match m if he followed strategy k. s_{imr}^{k} is a function of past choices by both players and is coded as 1 for cooperate and 0 for defect. The choice that subject i actually made in that round and match is denoted by c_{imr} (also coded 1 for cooperate and 0 for defect), and the indicator function taking value 1 when the two are the same and 0 otherwise is $I_{imr}^{k} = 1\{c_{imr} = s_{imr}^{k}(\cdot)\}$ .We model the likelihood that the observed choices were generated by strategy k as $\Pr_i(s^k) = \prod_{M_i}\prod_{R_{im}}(\beta)^{I_{imr}^{k}}(1-\beta)^{1-I_{imr}^{k}}$ for a given subject and strategy (where M and R represent the sets of matches and rounds) where $\Pr(I_{imr}^{k}) = \frac{1}{1+\exp(-1/\gamma)} \equiv \beta$ and γ is a parameter to be estimated. When β is close to ½, choices are almost random, and when it is close to 1, choices are almost perfectly predicted.

In an environment with two choices, such as this one, one can view the above as resulting from a choice function that implements a strategy with error: $c_{imr} = 1\{s'_{imr}(\cdot) + \gamma\varepsilon_{imr} \geq 0\}$ where s'_{imr} (is coded as 1 if the strategy would cooperate and -1 otherwise) and ε is an error term with an extreme value distribution. The probability (β) of doing the prescribed action is a decreasing function of γ.

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 attributed to strategy s^k. One can think of this in a large sample as giving the probability of observing each strategy.

To get a better sense of this approach, think about the case where only one strategy is considered; then, the only parameter to be estimated is β.[20] Suppose that the estimate is *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 a 50% probability (since there are only two choices). Now, consider the case of more than one strategy. Imagine that the strategies included are: AD, TFT, and Grim, and that the estimates of their proportion is one third

[20] In practice we will estimate γ, but one implies the other.

for each of them, and β is still *0.8*. This implies that in round 1 of a match, the estimated model predicts a 60% cooperation rate. Now, suppose that 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 that he 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.[21]

This model is estimated on Phase 2 data (pooling together both the original experiments and the additional sessions – note that there are no additional sessions for $\delta = 1/2$, R = 32 and, thus, a smaller sample) after dropping the first one quarter of the data. The first quarter of Phase 2 is dropped to limit the importance of matches where behavior is still changing. The variance-covariance matrix is obtained by bootstrap.[22]

The comparison of the estimated prevalence of strategies with the strategies actually chosen allows us to answer several questions. First, if the set of strategies the estimation allows 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 among the following: AC, AD, TFT, grim, WSLS, or STFT (the most common strategies) and the estimated proportion is not 0 (when all six strategies are included).[23] These results are presented in Table 9 (WSLS is not included in Table 9 because its frequency is always estimated to be 0). First, notice the high value of β (correspondingly low values of γ), which indicates a very good fit. Qualitatively, the SFEM mostly performs well at identifying the most important strategies. For instance, in three of the five treatments, it correctly identifies the two most popular strategies and their order of popularity. In one treatment, it correctly ranks all five strategies. In another treatment, it identifies the best three, but not in the correct order. The table also reports whether the hypothesis that the estimated proportions equal the frequency elicited in the data can be

[21] In round 1, 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.
[22] 1000 samples are drawn. A session is selected, then subjects within that session are drawn, finally matches for each subject are randomly chosen
[23] Estimates close to or at the boundary of the parameter space cause difficulties for the estimation of the variance (see, for instance, Andrews 1998). Since this is a challenging issue with no straightforward solution and is not important for the point being made here, we simply re-estimate the model ignoring the strategies that have 0 weight in this case.

rejected.[24] As can be seen, there are four cases (out of 22) for which the estimated and elicited frequencies are statistically different. Note, also, that although they are statistically different, the differences tend to be small in magnitude. In only one of these cases is it a meaningful difference in terms of size or the ranking of strategies (the proportion of TFT in the treatment with $\delta = 3/4$, $R = 32$). When considering the joint hypothesis that all the estimated frequencies are equal to the elicited ones, that treatment is the only one for which the hypothesis of equality can be rejected.[25]

Table 9: Estimation Performance – Restricted Samples

		AC	TFT	Grim	STFT	AD	β	Joint Test§	Number of Obs.
δ = ½, R = 32	Data		0.08		0.19	0.73			1563
	Estimation		0.00		0.24	0.76°	0.98	0.263	
δ = ½, R = 48	Data	0.04	0.22	0.42	0.04	0.29			3288
	Estimation	0.06	0.13	0.48	0.06	0.27	0.96	0.539	
δ = 3/4, R = 32	Data	0.02	0.12	0.13	0.15	0.58			1958
	Estimation	0.02°	0.13	0.10	0.28	0.48	0.94	0.234	
δ = 3/4, R = 48	Data	0.09	0.35	0.42		0.14			2330
	Estimation	0.17	0.13°°°	0.56°		0.15	0.99	0.038	
δ = 9/10, R = 32	Data	0.02	0.49	0.22	0.09	0.18			2204
	Estimation	0.08	0.49	0.12	0.08	0.23	0.96	0.181	

° Significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.
§ Reports the p-value of the joint hypothesis test that each estimated proportion equals the elicited

The exercise of Table 9 is meant mainly to facilitate hypothesis testing between the estimated proportions and the elicited proportions; however, in practice, the sample will not be restricted and certain strategies that subjects use may not be included in the SFEM. Hence, a natural question is whether the results would be qualitatively misleading once all the data are included. 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] To perform these hypothesis tests, we do a bootstrapped Wald test to take into account the variability in both the SFEM estimator and the variance in the elicited strategies. More specifically, when we randomly select the sample on which to perform the estimation, we also select the associated elicited strategies. The bootstrapped Wald statistic is the average of the Wald statistic for each bootstrapped sample testing that the estimator and the elicited frequencies are equal.

[25] To keep Table 9 easy to read, we do not include the standard errors, but these are available in the Appendix and include tests of equality to 0 to show that the estimates have a reasonable amount of precision.

Table 10: Estimation Performance – Full Sample

		AC	TFT	WSLS	Grim	STFT	AD	Other	β	Obs.
δ = ½, R = 32	Data	0.00	0.05	0.00	0.06	0.14	0.53	0.22		
	Estimation	0.00	0.03	0.00	0.02°°	0.23	0.73°°		0.97	2146
		(0.00)	(0.03)		(0.03)	(0.11)**	(0.12)***			
δ = ½, R = 48	Data	0.03	0.19	0.03	0.36	0.03	0.25	0.11		
	Estimation	0.07	0.06°	0.00	0.55	0.05	0.26		0.95	3684
		(0.05)	(0.07)		(0.13)***	(0.05)	(0.09)***			
δ = 3/4, R = 32	Data	0.02	0.10	0.01	0.10	0.12	0.47	0.18		
	Estimation	0.04	0.09	0.00	0.10	0.29°	0.48		0.93	2420
		(0.02)*	(0.05)*		(0.10)	(0.10)***	(0.09)***			
δ = 3/4, R = 48	Data	0.08	0.30	0.03	0.35	0.00	0.12	0.13		
	Estimation	0.18°°	0.11°°	0.00°°°	0.57°°	0.01	0.13		0.97	2772
		(0.07)***	(0.11)		(0.12)***	(0.02)	(0.06)***			
δ = 9/10, R = 32	Data	0.2	0.39	0.01	0.17	0.07	0.14	0.21		
	Estimation	0.08	0.42	0.00°°	0.18	0.11	0.21		0.94	2810
		(0.07)	(0.10)***		(0.11)*	(0.06)**	(0.07)***			

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 pick up the most popular strategies and, in most cases, the estimates are not statistically different from the elicited strategies. In particular, the strategy estimated to be the most popular is always the one that is most popular according to the elicitation. It is also interesting to consider which strategies are statistically different from 0 since 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 = 32), the two most popular strategies are identified as statistically different from 0. What seems to be more difficult is to identify the relative importance of TFT and Grim when cooperation rates are high. This is not particularly surprising since higher rates of cooperation result in fewer observations differentiating between those two strategies. However, in treatments where substantially more subjects choose Grim than TFT, Grim is estimated to be more popular than TFT and vice-versa.

Also note that the strategies that represent less than 10% of the data are rarely identified as statistically significant. Hence, and to some extent not surprisingly, strategies that are present, but 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 one needs many transitions to uncover the strategy. With δ = ½, half the games last only 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. The proximate reason for this in our environment is that most strategies used have memory 1; thus, one does not need many interactions to uncover the strategies being used.

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 two rounds following a defection of the other before returning to cooperation (referred to as T2). Thus, considering this set of possible strategies allows 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).

Table 11: Estimated Performance – Strategy Set from DF 2011

		AC	TFT	WSLS	Grim	STFT	AD	T2	Other	*B*	Number Of Obs.
δ = ½, R = 32	Data	0.00	0.05	0.00	0.06	0.14	0.53	0.00	0.36		2146
	Estimation	0.00	0.04	0.00	0.00°°°		0.96°°°	0.00		0.97	
δ = ½, R = 48	Data	0.03	0.19	0.03	0.36	0.03	0.25	0.00	0.14		3684
	Estimation	0.06	0.06°	0.00°°	0.54		0.31	0.00		0.95	
δ = 3/4, R = 32	Data	0.02	0.10	0.01	0.10	0.12	0.47	0.00	0.30		2420
	Estimation	0.05	0.15	0.00°°°	0.05		0.75°°°	0.00		0.90	
δ = 3/4, R = 48	Data	0.08	0.30	0.03	0.35	0.00	0.12	0.00	0.13		2772
	Estimation	0.21	0.13	0.00	0.51		0.16	0.00		0.97	
δ = 9/10, R = 32	Data	0.02	0.39	0.01	0.17	0.07	0.14	0.00	0.28		2810
	Estimation	0.09	0.60°°°	0.00°°	0.09		0.23°	0.00		0.93	

Results for that specification are presented in Table 11 (table with standard errors in the Appendix). As can be seen by comparing the results in Tables 10 and 11, in the two treatments where STFT is more important, not including it (in addition to other missing strategies) in the estimation leads to a statistically significant overestimation of the fraction of AD. However, 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.

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 such 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 that are thought to have some desirable properties, such as Win-Stay-Lose-Shift, are not prevalent. This last observation is consistent with Dal Bó and Fréchette (2011), Fudenberg, Rand, and Dreber (2012), and Fréchette and Yuksel (2013) who find little to no evidence that subjects use WSLS.

Most subjects seem to rely on strategies that have, at most, memory 1 in this experiment. This result may not hold in general, and in some situations, individuals may use more complex strategies. Understanding the features of the environment that lead people to use simple or more-complex strategies is an interesting question for future work.

We also 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 explained by the subjects thinking of a situation where there is some probability of error; as the expected number of rounds increases, the expected payoff difference between Grim and TFT becomes larger.

The use of the SFEM to estimate the distribution of strategies results in a very accurate fit. In the full sample, the estimated error probability ranges from 0.02 to 0.07 between treatments. Comparing the strategies' frequencies estimated using the SFEM to the elicited frequencies also suggests that SFEM 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 a small amount.

This paper provides a new perspective on the play of the infinitely repeated prisoner's dilemma by allowing us to observe subjects' strategies, as opposed to simply their choices. Much more work remains to be done to develop a 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

Andrews, Donald W.K. 1998. "Hypothesis testing with a restricted parameter space." *Journal of Econometrics*, Volume 84(1): 155-199.

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.

Bruttel, Lisa V. and Ulrich Kamecke. 2012. "Infinity in the lab: How do people play repeated games?" *Theory and Decision* 72(2): 205-219.

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. 2012. "Cooperative Strategies in Anonymous Economies: An Experiment." *Games and Economic Behavior* 75: 570-586.

Cason, Timothy and Vai-Lam Mui. 2008. "Coordinating Collective Resistance through Communication and Repeated Interaction." Purdue University. 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, Drew Fudenberg, and David G. Rand. 2011. "Who Cooperates in Repeated Games?" Working Paper.

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.

Fréchette, Guillaume R. 2011. "Session-Effects in the Laboratory." *Experimental Economics*, 15(3): 485-498.

Fréchette, Guillaume R. And Sevgi Yuksel. 2013. "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination." Working paper.

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. 2012. "Slow to Anger and Fast to Forget: Cooperation in an Uncertain World." *American Economic Review*, 102(2): 720-749.

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.

Rand, David G., Drew Fudenberg and Anna Dreber. 2013. “It's the thought that counts: The role of intentions in reciprocal altruism.” Working paper.

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.

Vespa, Emanuel. 2013. “Cooperation in Dynamic Games: An Experimental Investigation.” Working paper.

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.[26]

[26] 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 round, CC, DC, CD and DD. SGPE for strategy "s" denotes that (s,s) is a subgame 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	12	16	14
Number of Games	77	52	37	30
Phase 1	37	22	10	12
Phase 2	25	15	13	4
Phase 3	15	15	*14*	*14*
Number of subjects	16	16	14	20
Number of Games	82	45	49	25
Phase 1	37	15	21	6
Phase 2	30	16	14	5
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	0.64	0.41	0.60	0.81	0.35	0.56	0.56	0.36	0.52	0.70
¾	32	0.25	0.23	0.27	0.26	0.24	0.35	0.18	0.19	0.21	0.23	0.21	0.21
	48	0.75	0.72	0.61	0.96	0.86	0.78	0.65	0.61	0.49	0.90	0.80	0.70
9/10	32		0.41	0.44		0.60	0.66		0.44	0.37		0.53	0.52
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	0.53	0.41	0.59	0.82	0.39	0.31	0.52	0.39	0.38	0.74
¾	32	0.25	0.23	0.36	0.30	0.23	0.32	0.16	0.22	0.22	0.16	0.20	0.22
	48	0.89	0.67	0.67	0.98	0.83	0.80	0.77	0.57	0.56	0.77	0.86	0.80
9/10	32		0.58	0.56		0.62	0.62		0.59	0.43		0.55	0.61

Table A5: Comparison of Cooperation Rate in Additional Sessions with Original Sessions and Dal Bó and Fréchette (2011) (First Round of Last Match)

	Additional relative to Original				Additional relative to Dal Bó and Fréchette (2011)		
	δ = ½	δ = ¾		δ = 9/10	δ = ½	δ = ¾	
	R	R		R	R	R	
	48	32	48	32	48	32	48
Panel A: Phase 2							
Menu	0.703	0.284	-0.097	-0.006	1.149	0.074	-1.159
	[0.527]	[0.654]	[0.287]	[0.218]	[0.317]***	[0.703]	[0.446]***
Constant	0.226	-0.748	0.939	0.305	-0.22	-0.538	2
	[0.483]	[0.038]***	[0.245]***	[0.213]	[0.237]	[0.260]**	[0.420]***
Observations	90	72	76	84	80	72	74
Panel B: Phase 1							
Menu	0.208	0.382	-0.02	-0.054	0.35	0.308	-0.777
	[0.571]	[0.576]	[0.270]	[0.203]	[0.286]	[0.565]	[0.327]**
Constant	-0.135	-0.748	0.451	0.202	-0.276	-0.674	1.207
	[0.542]	[0.125]***	[0.266]*	[0.202]	[0.224]	[0.057]***	[0.324]***
Observations	90	72	76	84	80	72	74

Note: Results from probit regressions with variable menu equal to one for additional sessions.

Table A6: Estimation Performance – Restricted Samples

(Table 9 with Standard Errors)

		AC	TFT	Grim	STF	AD	γ	β	Number of Obs.
δ = ½, R = 32	Data		0.08		0.19	0.73			
	Estimation		0.00		0.24	0.76	0.26	0.98	1563
			(0.00)			(0.11)***	(0.03)***		
δ = ½, R = 48	Data	0.04	0.22	0.42	0.04	0.29			
	Estimation	0.06	0.13	0.48	0.06	0.27	0.30	0.96	3288
		(0.04)*	(0.10)	(0.14)***		(0.09)***	(0.04)***		
δ = 3/4, R = 32	Data	0.02	0.12	0.13	0.15	0.58			
	Estimation	0.02	0.13°°	0.10	0.28	0.48	0.36	0.94	1958
		(0.02)	(0.06)**	(0.09)		(0.10)***	(0.04)***		
δ = 3/4, R = 48	Data	0.09	0.35	0.42		0.14			
	Estimation	0.17	0.13	0.56		0.15	0.23	0.99	2330
		(0.06)***		(0.13)***		(0.06)	(0.04)***		
δ = 9/10, R = 32	Data	0.02	0.49	0.22	0.09	0.18			
	Estimation	0.08	0.49	0.12	0.08	0.23	0.32	0.96	2204
		(0.07)	(0.09)***	(0.10)		(0.07)***	(0.04)***		

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%.

Table A7: Estimated Performance – Strategy Set from DF 2011

(Table 11 with Standard Errors)

		AC	TFT	WSLS	Grim	STFT	AD	T2	Other	γ	β	Number Of Obs.
δ = ½, R = 32	Data	0.00	0.05	0.00	0.06	0.14	0.53	0.00	0.36			
	Estimation	0.00	0.04	0.00	0.00°°°		0.96°°°	0.00		0.30	0.97	2146
		(0.00)	(0.05)		(0.00)		(0.05)***			(0.04)***		
δ = ½, R = 48	Data	0.03	0.19	0.03	0.36	0.03	0.25	0.00	0.14			
	Estimation	0.06	0.06°	0.00°°	0.54		0.31	0.00		0.35	0.95	3684
		(0.04)	(0.07)	(0.02)	(0.14)***		(0.13)**			(0.05)***		
δ = 3/4, R = 32	Data	0.02	0.10	0.01	0.10	0.12	0.47	0.00	0.30			
	Estimation	0.05	0.15	0.00°°°	0.05		0.75°°°	0.00		0.45	0.90	2420
		(0.04)	(0.07)***	(0.03)	(0.08)		(0.09)***			(0.05)***		
δ = 3/4, R = 48	Data	0.08	0.30	0.03	0.35	0.00	0.12	0.00	0.13			
	Estimation	0.21	0.13	0.00	0.51		0.16	0.00		0.28	0.97	2772
		(0.21)	(0.22)	(0.06)	(0.26)**		(0.20)			(0.18)		
δ = 9/10, R = 32	Data	0.02	0.39	0.01	0.17	0.07	0.14	0.00	0.28			
	Estimation	0.09	0.60°°°	0.00°°	0.09		0.23°	0.00		0.39	0.93	2810
		(0.07)	(0.08)***	(0.00)	(0.07)		(0.06)***			(0.04)***		

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%.

Appendix B: Figures

Figure A1: Evolution of Cooperation (first rounds)

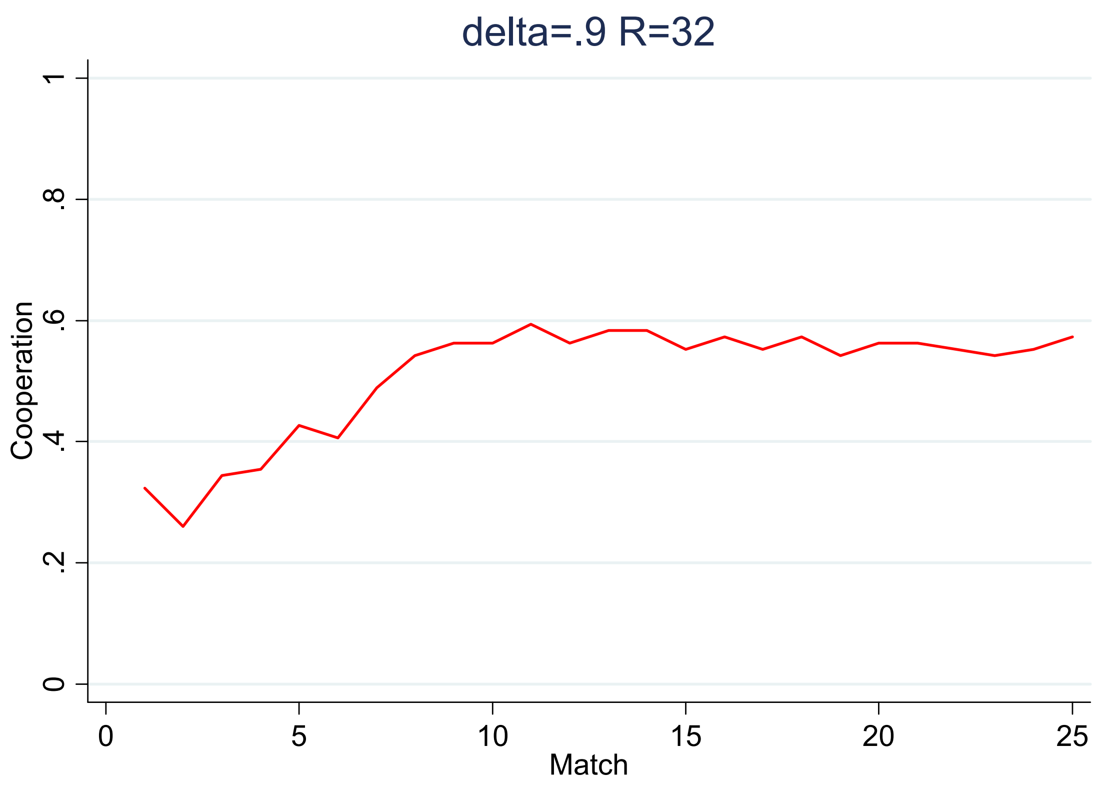

Figure A2: Evolution of Strategies (Phase 2)

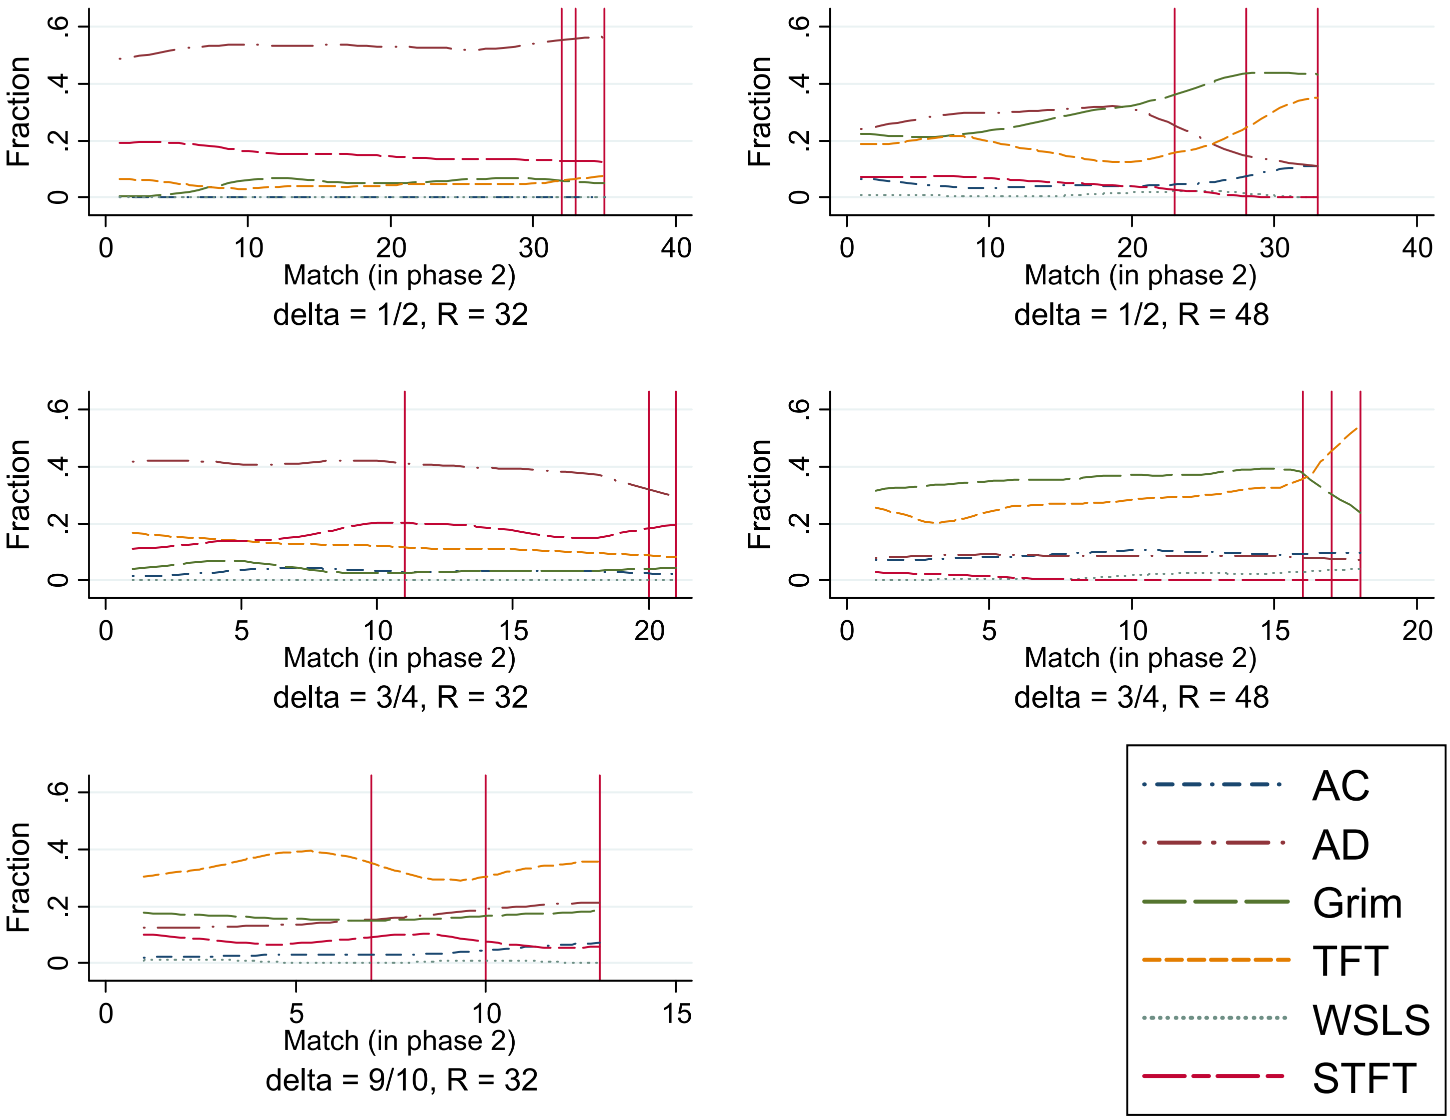

Note: the vertical red lines indicate the end of the Phase 2 in each session.

Appendix C: 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 the 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; 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 in X consecutive rounds either the other or myself selected 2, 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)
- Start by selecting 1 and do so until the other or myself selects 2, in that case select 2 for X rounds. After that go back to the start. (T-X)
- Build your own. (This offers the same option as in the memory-1 treatment.)

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