Infinitely Repeated Games in the Laboratory:

Four Perspectives on Discounting and Random Termination*

Guillaume R. Fréchette
New York University

Sevgi Yuksel
New York University

June 5, 2013

Abstract

While infinitely repeated games with payoff discounting are theoretically isomorphic to randomly terminated repeated games without payoff discounting, in practice, they correspond to very different environments. The standard method for implementing infinitely repeated games in the laboratory follows the second interpretation and uses random termination (proposed by Roth and Murnighan [1978]), which links the number of expected repetitions of the stage game to the discount factor. However, we know little about whether or not people treat situations where the future is less valuable than the present in the same way as interactions that might exogenously terminate. This paper compares behavior under four different implementations of infinitely repeated games in the laboratory: the standard random termination method and three other methods that de-couple the expected number of rounds and the discount factor. Two of these methods involve a fixed number of repetitions with payoff discounting, followed by random termination (proposed by Cabral, Ozbay, and Schotter [2011]) or followed by a coordination game (proposed by Cooper and Kühn [2011]). We also propose a new method - block random termination - in which subjects receive feedback about termination in blocks of rounds. We find that behavior is consistent with the presence of dynamic incentives only with methods using random termination, with the standard method generating the highest level of cooperation. The other two methods display two advantages: a higher level of stability in cooperation rates and less dependence on past experience. We also estimate the strategies used by subjects under each method. Those estimates reveal that the average number of interactions, even when the discount rate is the same, affects strategic choices: interactions expected to be longer increase defection and decrease the use of the Grim strategy.

JEL classification: C9; C72; C73; C91; C92

Key words: Infinitely repeated games, discounting, random termination, prisoner's dilemma

* We would like to thank David Cooper, Drew Fudenberg, Ryan Oprea, Andrew Schotter, Ekatarina Sherstyuk, Alistair Wilson, participants at the 2011 ESA Conference and the NYU experimental group for helpful comments, Emanuel Vespa for his help developing the ztree code to conduct the experiment; also CESS and the C.V. Starr Center for financial support. Fréchette gratefully acknowledges the support of NSF via grants SES-0924780 and SES-1225779.

1 Introduction

From a theoretical point of view, whether a game is infinitely repeated and utilities are discounted at a rate δ, or the game ends after every round with probability $(1 - \delta)$ has no impact on how the game is analyzed; the two interpretations are isomorphic. In fact, depending on the paper and the application, one (or both) interpretation has been given (see Mailath and Samuelson [2006] section 4). The study of repeated interactions has a long history in the social sciences. There are clear challenges to studying such games experimentally - namely, one cannot play a game of infinite duration in the laboratory. Hence, in economic experiments, infinitely repeated games are almost exclusively induced using randomly terminated games. This experimental procedure, first introduced by Murnighan and Roth [1983],[1] equates the continuation probability of the repeated game to the discount factor and, thus, gives the experimenter control over the discount factor δ: a crucial parameter from a theoretical perspective. However, implementing infinitely repeated games in this way creates a link between the expected number of rounds played and the discount factor. Clearly, there is no such link when the underlying game that is modeled is infinitely repeated with payoff discounting. Furthermore, in practice, some situations are probably closer to one or the other of these extremes from a descriptive point of view. For instance, some markets have very high turnover (firms exiting frequently), and it probably makes sense to think of those environments as being closer to randomly terminated ones. In other applications, the key agents – such as political parties for instance – are long lived, and it might make more sense to think of them as discounting future payoffs. In this paper, we compare behavior under four different implementations of infinitely repeated games in the laboratory: the standard random termination method and three other methods that de-couple the expected number of rounds and the discount factor.

Although the theory of infinitely repeated games has been very active for a few decades, experimental investigations have only recently become common. To investigate some of the questions that emerge naturally from this literature, it can be important to observe many rounds of a supergame. For example, Vespa [2013] studies a dynamic game in which, given the parameters, the cooperative strategy yields higher payoffs than other strategies only for supergames that last more than seven rounds. However, if the standard random termination method is used, given the discount factor, supergames of this length would be observed only 13 percent of the time. If subjects' learning is influenced by realized outcomes, then it might be difficult for them to learn to cooperate. In

[1] Other methods, used mainly in other social sciences, involve not specifying the number of repetitions or announcing the number of repetitions, but playing for a very long time.

some cases, the desire to de-couple the expected number of rounds and δ comes from the opposite need: to reduce the number of rounds in a supergame. Cooper and Kühn [2011] study communication in an infinitely repeated game. To reduce the difficulty of analyzing messages, they want to reduce the strategy space, which they achieve by limiting the number of rounds per repeated game. They find that allowing for communication has an important impact on behavior.

Such considerations raise several questions. Since theory is, for the most part, silent on the factors that affect cooperation in infinitely repeated games, does varying the number of rounds played for a fixed discount rate change behavior? In a larger context, do agents respond to payoff discounting and probabilistic continuation differently in repeated interactions?[2] For instance, Dal Bó and Fréchette [2012] find that increasing δ has an important impact on the choice of strategies: for example the fraction of subjects who always defect decreases as δ increases. As we will show in this paper, this result is actually the reverse of what we find when we increase the average number of interactions but keep δ constant. From the perspective of testing the implications of infinitely repeated games in the laboratory, do different methods of implementation lead to different conclusions with respect to basic comparative statics of the theory? Finally, from a very practical point of view, if someone has a need to de-couple the discount factor and the number of rounds, what are the impacts of the different implementation methods?

The three variations on the standard randomly terminated (henceforth RT) game we consider are the following. In the RT games, after every round of play, there is a fixed known probability δ that the game continues for an additional round, and a probability $(1 - \delta)$ that the match ends. A match refers to a supergame, and a round is one play of the stage game. One variation involves payoff discounting followed by random termination (D+RT). In this method, a fixed (known) number of rounds are played with certainty, and payoffs in these rounds are discounted at a known rate δ. After the rounds with certainty are played, there is a fixed known probability δ that the match continues for an additional round, and payoffs in these rounds are no longer discounted. This procedure was first introduced by Cabral et al. [2011] and has since been used by Vespa [2013].[3]

Another variation also starts with a fixed number of rounds with payoff discounting, but it is then followed by the coordination game induced by considering only two partic-

[2]Zwick, Rapoport, and Howard [1992] study an infinite horizon game, an alternating bargaining game, with an exogenous termination probability and compare the results to prior experiments using payoff discounting. Results are quite similar even though the experiments use different procedures.

[3]Note that one could also first have a fixed number of rounds without payoff discounting followed by random termination. Such a procedure has been implemented in some experiments, but it changes the environment to a non-stationary one and, thus, for certain games, can introduce different equilibria.

ular strategies in the infinitely repeated game - namely, the Grim trigger strategy and the strategy of always defecting.[4] This method (D+C) was first used by Cooper and Kühn [2011].

Finally, we consider a new procedure that we refer to as block random termination (BRT). Subjects play as in the standard RT, but in blocks of a pre-announced fixed number of rounds. Within a block, subjects get no feedback about whether or not the match has continued until that round, and they make choices that will be payoff-relevant contingent on the match actually having reached that point. Once the end of a block is reached, subjects are told whether the match ended within that block and, if so, in what round; otherwise, they are told that the match has not ended yet, and they start a new block. Subjects are paid for rounds only up to the end of a match, and all decisions for subsequent rounds within that block are void. As in the RT, there is no payoff discounting. Note that this is not the same as the strategy method, subjects make choices conditional on the past history, not for any potential contingency. To the best of our knowledge, this method has not been used before.[5] Under certain assumptions, all three alternative implementations of the infinitely repeated game result in the same theoretical possibilities as random termination.

Our results show that each implementation generates sharp comparative statics: cooperation levels drop significantly when parameters of the stage game are changed to make mutual cooperation theoretically unsustainable. However, analysis of behavior within a match indicates that the cooperation observed with D+C is not supported by dynamic incentives. Under this method, as subjects gain experience, their response to the coordination game becomes independent of the history of play, and subjects' behavior in the first part of the game is similar to behavior observed in a finitely repeated game.

When comparing the other three methods that use random termination, we find the highest levels of cooperation with RT. However, D+RT generates the most stable cooperation rates within a match. Furthermore, we find behavior in D+RT and BRT to be significantly less affected by past experiences within a session. These findings make these methods potentially more desirable than RT when important variations in the realized length of supergames are expected to occur and the samples are small.

In a broader context, our results also indicate that subjects respond to payoff discounting and probabilistic continuation in slightly different ways. For instance, we find

[4]The Grim trigger strategy involves first cooperating, followed by cooperation as long as the other player cooperates, but defection forever if the other defects.

[5]However, Wilson and Wu [2013] use this method after having seen the current paper presented.

strategy choice in an environment where interactions are likely to be short lived to be different from one where interactions are long lived, but agents are impatient. Ex-ante, one might have expected that increasing the average number of interactions while keeping δ constant should increase cooperation. This would be in line with the idea that increasing the number of rounds in a finitely repeated PD increases cooperation, as well as with the observation that increasing δ, holding payoffs constant, leads to higher cooperation rates. Furthermore, if subjects are risk-averse, then moving to a treatment such as the D+RT should make cooperation easier to support (we will come back to this observation in the next section). However, estimation of strategies used by subjects in these different environments shows that, with payoff discounting, subjects are more likely to be suspicious - i.e., reluctant to use strategies that start with cooperation in the first round. In fact, in D+RT the fraction of subjects who always defect increases, even though the expected payoffs from that strategy is lower than in other treatments, and subjects are less likely to support cooperation using a Grim strategy. Revisiting prior experimental evidence with this new found perspective yields corroborative evidence in terms of the strategic impact of having longer interactions (while controlling for the value of cooperation).

It is particularly important to understand the differing effects of these environments on the subjects' strategic considerations and on equilibrium selection as the theory of infinitely repeated games says very little about the factors that affect cooperation. Thus, systematic behavioral differences in repeated interactions with payoff discounting vs. random continuation can have important implications for the application of the theory of infinitely repeated games to these different environments.

The paper is organized as follows. In the next section, we compare the different methods examined theoretically, and we describe the experimental design. In Section 3, we discuss the results. We conclude, in Section 4, with a discussion of the advantages and disadvantages of the different methods. We also discuss the implications of our results beyond implementation of infinitely repeated games in the laboratory.

2 Theoretical Considerations and Design

Denote the stage game payoffs by the following:

	C	D
C	R,R	S,T
D	T,S	P,P

with $T > R > P > S$, which defines a prisoner's dilemma (PD). As is usual for such a game, if

$$\delta \geq \frac{R-T}{P-T} \tag{1}$$

joint cooperation can be supported as part of a subgame perfect equilibrium.

Each of the alternative methods we investigate involves a number of rounds played with certainty, and we denote that number by ρ. Hence, in the case of D+RT, ρ rounds are played with payoff discounting, after which each additional round occurs with probability δ where payoffs are no longer discounted. D+C involves ρ rounds played with payoff discounting, followed by the coordination game below (where G stands for Grim and AD for Always Defect):

	G	AD
G	$R\frac{\delta^\rho}{1-\delta}, R\frac{\delta^\rho}{1-\delta}$	$S\delta^\rho + P\frac{\delta^\rho}{1-\delta}, T\delta^\rho + P\frac{\delta^\rho}{1-\delta}$
AD	$T\delta^\rho + P\frac{\delta^\rho}{1-\delta}, T\delta^\rho + P\frac{\delta^\rho}{1-\delta}$	$P\frac{\delta^\rho}{1-\delta}, P\frac{\delta^\rho}{1-\delta}$

.

BRT will be ρ rounds played with certainty with no payoff discounting; the probability that any of these first ρ rounds is relevant for payments is given by the geometric distribution with parameter δ. If the match does not end in the first block of ρ rounds, then an additional block of ρ rounds is played, and so on. Thus, the probability that the block to be played is the last, given that the previous block was not the last, is given by $\sum_{i=1}^{\rho} (1-\delta)\,\delta^{i-1}$ for $\rho \geq 1$.

How do the different implementation methods affect the condition for cooperation to be part of a subgame perfect equilibrium? In the case of D+RT, if agents are risk-neutral, there is no difference; the condition is the same as in RT. However, as is well known, for a risk-averse agent, a modified condition involving a higher minimal δ can be derived.[6] If agents are risk-averse, then the D+RT method will result in a critical δ between that in condition 1 and the one that would be relevant for RT. However, Dal Bó and Fréchette [2011] note that given the parameters used in their experiment, this should not have practical relevance for the levels of risk aversion typically observed in experiments. This observation also holds true for most experiments conducted in this literature to date, including this one.

[6]Sherstyuk, Tarui, and Saijo [2013] experimentally investigate the effect of paying only in the last round of a match (as opposed to all rounds) which eliminates the need to assume risk neutrality. They find no difference in behavior between the standard payment method and paying only in the last round of a match. Note also that Schley and Kagel [2013] find that behavior is not sensitive to presentation manipulation: i.e.if the payoffs are listed in cents or dollars does not affect cooperation rates.

When it comes to the D+C implementation, more assumptions are required for it to be theoretically equivalent to RT. First, it must be the case that only two strategies are sufficient to capture the evolution of payoffs for both players. Namely, for all strategies used by the subjects, although they do not need to be Grim and AD, it must be that when they play each other, they result in the same payoffs as when Grim and AD play each other (for instance, Tit-For-Tat[7] and AD). Dal Bó and Fréchette [2012] show that strategies beside Grim and AD are used, and thus this assumption might constrain behavior too much. Nonetheless, there is evidence in the Cooper and Kühn data that play in the (period 2) coordination game depends or reacts to what has happened before (in period 1). However, they study a 3 x 3 game and the responsiveness of period 2 to the opponent's choice in period 1 is less than in standard randomly terminated games. What they do see is a tendency to choose a more cooperative action in the coordination game in response to more cooperative choices in period 1.

Finally, BRT is theoretically equivalent to RT. However, one might worry that decisions are made in a different frame of mind, something that some have suggested is a potential problem with the strategy method.[8] BRT is similar to the strategy method in the sense that when subjects make choices, they know that their choices will be payoff-relevant only in some states of the world. However, one should note that with BRT, unlike the strategy method, in every round, a subject considers only one history of play, and the contingency of their choice comes only from random termination. If a round is selected for payment, then the choice that a subject makes in that round is relevant; in the strategy method, a decision is made for each contingency, which is not the case here.

Our experimental design involves a mix of within- and between-subjects design. The implementation method is evaluated across subjects, but the stage game will be varied within-subject. Throughout the experiment, δ is set to 0.75. In the first part of each session, subjects play 12 matches with the following payoff matrix:

	C	D
C	40, 40	12, 48
D	48, 12	20, 20

.

With such a stage game, cooperation can be supported with any discount factor δ above 0.29. Moreover, for δ greater than $0.\overline{4}$, cooperation is risk-dominant in the sense that when

[7] Tit-For-Tat starts by cooperating, then matches what the opponent plays in the previous round from then on.

[8] Brandts and Charness [2000] find no difference between a "hot" and "cold" treatment in two one-shot games.

focusing only on the strategies always defect and Grim trigger (or Tit-For-Tat) the later risk dominates always defecting. Dal Bó and Fréchette [2011], Blonski, Ockenfels, and Spagnolo [2011], and Fudenberg, Rand, and Dreber [2012] (reporting results based on data from Dreber, Rand, Fudenberg, and Nowak [2008]) find that this criterion correlates with cooperation rates.[9] This stage game and discount factor were selected because prior experiments suggests that such parameters will lead to cooperation rates above 0 but below 1, giving us room to observe the different implementation methods having a positive or negative impact on cooperation rates.

In the second part of the experiment, subjects play six matches with the stage game

	C	D
C	24, 24	12, 48
D	48, 12	20, 20

.

With this stage game, δ needs to be above 0.86 for cooperation by both players, (C, C), to be an equilibrium. It is possible in equilibrium for subjects to alternate between (D, C) and (C, D) for δ above 0.28.[10] Given the δ of $\frac{3}{4}$, it is possible for (C, C) to emerge in equilibrium in the first part of the experiment but not in the second. We selected this set of parameters with the idea that it would result in a significant impact across parts 1 and 2, using the RT method, and, thus, allowing us to test whether the comparative static results were the same across all four implementations. Since parameters in which alternation is an equilibrium but (C, C) cannot be supported in equilibrium are few (in fact, we know of only Dal Bó [2005]), we also wanted to add to the body of evidence on this case.

Clearly, cooperation rates could be different if the order were inverted, but since this is not relevant for the questions investigated here, we keep the order constant for simplicity, starting with more repetitions of the case where joint cooperation can emerge since prior evidence suggests that, in general, it is more difficult to generate cooperative behavior. Subjects were informed that the experiment had two parts, and the stage game for the second part was presented to the subjects only after the first part was over. Standard experimental procedures such as neutral language were used. Subjects were randomly re-

[9]Fudenberg et al. [2012], however, do not find that risk dominance correlates with choices in an infinitely repeated game with noise (imperfect public monitoring).

[10]Dal Bó [2005] has a similar treatment (Pd1 with $\delta = 0.5$) where, given the continuation probability, (C, C) cannot be supported in equilibrium, but it is possible to construct an equilibrium in which players alternate between (D, C) and (C, D). He finds alternation between these two outcomes to be slightly higher in this treatment compared to another one with same δ but a different payoff structure, where this cannot be sustained in equilibrium. However, he concludes that there is only weak evidence to suggest that subjects play such an equilibrium.

matched between matches. Instructions can be found in the online appendix.[11]

The number of periods played with certainty (except in the RT implementation), ρ, is set to $four$.[12] This implies that in treatment RT, the expected number of rounds per match is four. In D+RT, there is a minimum of four rounds and the expected number of rounds is seven. In the payoff discounting part of a match, payoffs are discounted by 0.75 every round. Hence, the stage game for Round 4 of the first part of the experiment is

	C	D
C	16.9, 16.9	5.1, 20.3
D	20.3, 5.1	8.4, 8.4

for part 1, while R, the payoff to joint cooperation, is 10.1 for part 2. In D+C there are five rounds; the first four are the same as in D+RT, but the 5^{th} round is

	G	AD
G	50.6, 50.6	22.8, 34.2
AD	34.2, 22.8	25.3, 25.3

for part 1. In part 2, the payoff to (G,G) is reduced to 30.4. In the BRT treatment, the minimum number of rounds is four and the expected number of rounds that will be relevant for payment is four.

When a session of the RT treatment was conducted, the seed for the pseudo-random number generator was picked by the software (based on the internal clock) and saved. Thus, each session of the RT treatment used a different random termination sequence. However, sessions of the other treatments used the same random sequences as the RT treatment. This was to control for the effect that specific experiences in terms of the length of matches have on the evolution of play (the impact of the length of matches on behavior has been documented before in, for example, Dal Bó and Fréchette [2011] for the PD and in Engle-Warnick and Slonim [2006] in the case of the trust game). Three sessions per treatment were conducted at the CESS laboratory at NYU. Subjects were recruited among undergraduate students from multiple majors. Table 1 gives some basic information about the sessions and treatments.

The different methods examined in our experiment imply different expected lengths of interaction between the subjects, which can potentially have an effect on cooperation

[11] Available at https://files.nyu.edu/gf35/public/print/Frechette_2013a_inst.pdf.

[12] As ρ increases, the time in the laboratory required to conduct the alternative implementations becomes longer. Four seemed long enough without making the sessions with alternative methods prohibitively long.

Table 1: Summary Information

	# of subjects	# of sessions	Subjects per session	Matches per session	Rounds per Match* avg	min	max	Subject earnings ($) avg	min	max
RT	50	3	12, 18, 20	18	4.5	1	19	23.9	12.2	32.6
D+RT	48	3	12, 16, 20	18	7.5	4	22	20.7	14.1	27.2
BRT	42	3	12, 14, 16	18	6.1	4	20	20.9	13.6	28.7
D+C	52	3	16, 18, 18	18	5	5	5	20.7	17.3	23.2

* Rounds played, though they may not count towards earnings in BRT.

Table 2: Cooperation Rate

	Round 1				All Rounds			
	(C, C) SPE		not SPE	Diff.	(C, C) SPE		not SPE	Diff.
RT	0.75	$>^{***}$	0.20	0.55	0.65	$>^{***}$	0.15	0.50
D+RT	0.53	$>^{***}$	0.25	0.27	0.47	$>^{***}$	0.18	0.28
BRT	0.61	$>^{***}$	0.18	0.43	0.41	$>^{***}$	0.09	0.33
D+C	0.68	$>^{***}$	0.18	0.50	0.57	$>^{***}$	0.14	0.43

*** Significant at the 1 percent level (standard errors clustered at session level).
Diff. stands for the difference between part 1 and 2.

levels, as previous studies suggest. In infinitely repeated prisoner's dilemma experiments, for parameters in which cooperation can be sustained as a subgame perfect equilibrium, higher discount factors implying longer interactions (due to higher continuation probability in RT) generate higher cooperation levels. However, in these games, the cost of defection is also increasing with the expected length of interaction if some subjects are using Grim strategies. Unlike previous studies, in our experiment, the differences in expected length of interaction across our treatments are compensated for by differences in payoff discounting; thus, they provide no theoretical reason to expect differences in cooperation rates, unless those differences factor in how subjects choose their strategies.

One should also note that observed differences in behavior across our treatments do not, in and of themselves, provide a contradiction to theory. For sufficiently high delta, the theory predicts a multiplicity of equilibria in all treatments. Our findings suggest that differences in these environments, without changing theoretical considerations, can have an effect on the strategies adopted by subjects and, consequently, on equilibrium selection.

3 Results

The first results on aggregate cooperation levels can be found in Table 2.[13] The left panel presents the data focusing only on Round 1, while the right panel includes all rounds, except for Round 5 for D+C. Note that variations across treatments when including all rounds can come about for a variety of reasons. For instance, if cooperation decreases within matches, since D+RT results in longer matches than RT does, it could mechanically result in lower average cooperation over all rounds, even though cooperation rates are the same when looking at the part that overlaps. In that sense, Round 1 offers a comparison that is easier to interpret. We will analyze the specific behavior within matches later.

Figure 1: Cooperation Rate by Match

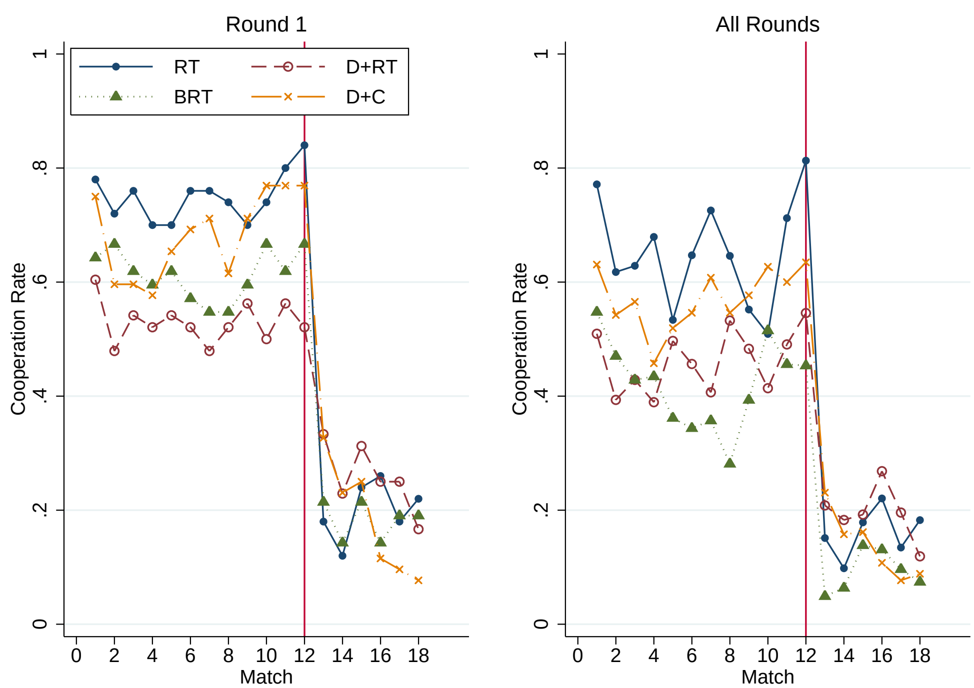

These results clearly show that cooperation is higher when joint cooperation can be supported in equilibrium. More importantly, this result holds true for all four treatments. Cooperation rates by matches can be seen in Figure 1. The figure suggests no clear pattern

[13]Throughout the text, unless noted otherwise, the statistical tests are based on probit estimations allowing for clustering at the session level. For a discussion of potential sources of session-effects, see Fréchette [2011].

Table 3: Comparison of Cooperation Rates Across Treatments

(C, C) SPE

	Round 1				All Rounds			
	RT	D+RT	BRT	D+C	RT	D+RT	BRT	D+C
RT	0.75	$>^{***}$	$>^{***}$	=	0.65	$>^{***}$	$>^{***}$	=
D+RT		0.53	$<^{**}$	$<^{***}$		0.47	=	$<^{**}$
BRT			0.61	=			0.41	$<^{***}$
D+C				0.68				0.57
Not SPE								
RT	0.20	$<^{**}$	=	=	0.15	$<^{**}$	$>^{***}$	=
D+RT		0.25	$>^{**}$	$>^{***}$		0.18	$>^{***}$	=
BRT			0.18	=			0.09	$<^{**}$
D+C				0.18				0.14

** Significant at the 5 percent level (standard errors clustered at session level).

*** Significant at the 1 percent level (standard errors clustered at session level).

The symbol indicates how the cooperation rate of the treatment identified by the row compares (statistically) to the one in the column.

of changes in cooperation rates over matches, but it suggests a few patterns we will explore in more detail. Across treatments, one can see that there are differences in cooperation levels. Furthermore, the changes over matches and across treatments are not necessarily the same for Round 1 and for all rounds.

Looking across treatments, the top panel of Table 3 summarizes how the cooperation rates compare in the first part of the experiment, where joint cooperation can be supported in equilibrium.[14] When looking at Round 1, D+C is in between RT and BRT but is not statistically different from either. All other pairwise comparisons are statistically significant. The standard method has the highest rate and D+RT the lowest. When looking at all rounds, the main change is that the rankings of BRT and D+RT are inverted, with the cooperation rate of D+RT higher in the later case. The bottom panel of Table 3 reports similar information for the second part of the experiment. To summarize, in the second part, the BRT treatment leads to the lowest cooperation rates, while D+RT leads to the highest. The size of the treatment effect - the difference between cooperation rates in parts 1 and 2 - is shown in Table 2. The results can be ordered, with the standard method having the largest treatment effect, followed by D+C, then BRT, and, finally, D+RT with the smallest treatment effect (this order holds, irrespective of looking at Round 1 only or at all rounds).

[14]The tests include dummy regressors to control for the specific random sequence in a given session.

To summarize the results so far, the comparative static effects are in the same direction for all treatments, but there are differences in magnitudes. In particular, cooperation rates in the first part of the experiments, matches 1 to 12, vary across treatments. The results that follow will focus on those matches and provide evidence that clarifies the sources of these differences. First, we will look at factors that affect the evolution of play over matches. Second, we will turn our attention to cooperation within matches, focusing on matches 7 through 12 (after subjects have gained experience). Third, we will explore the strategies used by subjects.

In the remainder of the paper, we concentrate on the first part of our experiment, in which mutual cooperation is possible. In the second part of the experiment, matches 13-18, in line with the theoretical predictions, we observe a sharp decline in cooperation rates (Figure 1) with aggregate cooperation rates dropping below 18 percent in all our treatments (Table 2). Cooperation rates in the second part are so low that there is little to be analyzed in terms of behavior. The variation in cooperation rates that we observe in the first part is what drives our treatment differences and is critical for understanding the trade-offs associated with the different methods we examine.

3.1 Matches 1 to 12

To understand the factors that affect the evolution of play over matches, Table 4 reports estimates of a probit where cooperation in Round 1 of the current match is regressed on observations from the previous match (namely, whether or not the opponent in the previous match first cooperated or not, the length of the previous match, and the length squared);[15] also included are the match and an indicator variable taking value one if the subject cooperated in the first round of the first match and zero otherwise. The dummy variable for whether one's previous opponent cooperated in Round 1 captures an aspect of that opponent's strategy and can be used to update one's beliefs about the probability that other players are using cooperative strategies.[16] Length and length squared can be used to update beliefs about the likely duration of a match (which, in turn, affects the expected value of cooperation). Match is meant to capture, in an economical way, any time trend. The choice in Round 1 of match 1 is included to allow for correlated random effects (reported in Table 8 of the Appendix) and is included here for comparability. In the

[15]Length is redefined to be the number of rounds - 3 in the case of D+RT to make the estimates comparable across treatments. Note, also, that length is the number of rounds used for payments in the case of BRT.

[16]By cooperative strategies, we mean any strategy that starts with cooperation. Note that, although players may update their beliefs about other aspects of the strategy used by others, choices after Round 1 are not exogenous of one own's choice, and, thus, using only Round 1 avoids issues of endogeneity.

Table 4: Probit Estimate of the Factors Affecting the Evolution of Cooperation (Matches 2 to 12)

Dependent Variable: Cooperation in Round 1

	RT	D+RT	BRT	BRT	D+C
Partner cooperated in	0.694***	0.434***	0.184	0.177	0.578***
Round 1 of previous match	(0.216)	(0.164)	(0.152)	(0.155)	(0.153)
Number of rounds	0.214*	0.103***	-0.108	-0.0984**	
in previous match	(0.113)	(0.0316)	(0.0730)	(0.0458)	
Number of rounds	-0.0121	-0.00835*	0.00757	0.00214	
in previous match sq.	(0.00785)	(0.00445)	(0.00639)	(0.00371)	
Two blocks				0.0602	
in previous match				(0.105)	
Three blocks				0.465***	
in previous match				(0.157)	
Match number	0.000287	0.00700	0.0164	0.0152	0.0638***
	(0.0359)	(0.0384)	(0.0238)	(0.0228)	(0.00467)
Subject cooperated in	2.043***	2.423***	0.920**	0.913**	1.541***
Round 1 of match 1	(0.142)	(0.229)	(0.420)	(0.424)	(0.418)
Constant	-1.795***	-1.982***	-0.279	-0.254	-1.450***
	(0.539)	(0.161)	(0.296)	(0.330)	(0.550)
N	550	528	462	462	572

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

case of BRT, we include two specifications, one that additionally accounts for the number of blocks that were played in the previous match. Finally, the specification for D+C drops length and block related regressors since the length is fixed in that treatment.

Results for the standard methods confirm observations from previous experiments. First, when the opponent in the previous match cooperates, one is more likely to cooperate in the subsequent match. See for instance Dal Bó and Fréchette [2011], in which, the authors show that a learning model can account well for the aggregate evolution of that aspect of behavior over matches. Second, when matches last longer, the subsequent match is more likely to start by cooperation. Dal Bó and Fréchette [2011] also observed this, and Engle-Warnick and Slonim [2006] made a similar observation in the context of the trust game. In the case of D+RT, similar effects are observed; however, as can be seen in Table 9 (in the appendix) which reports marginal effects, both channels have a lesser impact on cooperation. The impact of observing someone who first cooperated in the previous match drops from 0.21 to 0.17. The marginal effect of a longer match goes down from 0.06 to 0.04. In the case of BRT, using the same specification, both of these channels lose statistical significance. However, when controls for the number of blocks are added, the results suggests that cooperation rates increase when the previous match has more blocks, but decrease as more rounds are played. In other words, we observe a seesaw pattern of increase for each new block in the previous match and gradual decrease as more rounds occur within the block. Finally, in the case of D+C, we find that cooperation in the first round of the previous match has an impact similar to that found in RT, with a magnitude of 0.21. In addition, in that case, there is a positive trend, with cooperation rates in Round 1 increasing over time.

One question that these results raise is what aspects of learning are affected by the differences across treatments. For instance, does the fact that behavior in BRT reacts less to the observed outcomes mean that subjects do not learn to condition their decisions on what their opponent does? As Figure 2 shows, the results suggest similar evolutions in that regard across treatments. In all treatments, over the first few matches, there is a decrease in the probability of cooperation following a defection by the other player. In the RT treatment, it starts at 40 percent and ends at 18 percent. In the other treatments, it starts close to 25 percent and decreases to a rate between seven and ten percent. Cooperation rates following a cooperative decision by the other player show an opposite trend, although less pronounced, in all but the D+C treatment. The increase in cooperation over the 12 matches is between seven and twenty one percentage points, depending on the treatments. However, the trend of the cooperation rates following cooperation by the

Figure 2: Cooperation Rate as a Function of the Previous Choice of the Opponent Over Time

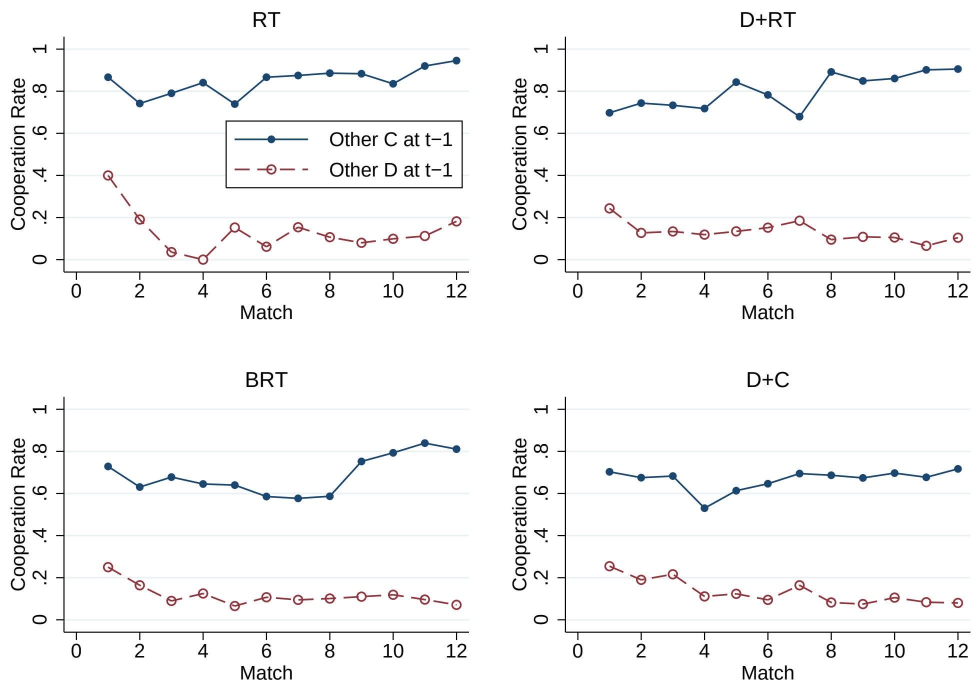

other player is more or less constant in the case of the D+C treatment. As a result of these trends, the difference in the conditional probabilities following a cooperation or defection decision by the other player increases with experience. The difference is 46 percent (plus or minus two percent) in all treatments in the first match, and it nearly doubles by match 12, reaching 76, 80, 74, and 67 percent in the RT, D+RT, BRT, and D+C treatments, respectively.

Figure 3 presents another aspect of the evolution of behavior. For each treatment and each match, the average cooperation rate is shown for rounds 1, 4, and 5.[17] Rounds 4 and 5 are informative because in D+RT, they are the rounds just before and just after the transition to random termination; in BRT, they represent the end of the first block and the start of the second block; and in the D+C treatment, it is the last PD choice and the

[17]There are no sessions in the RT treatment where matches 1 and 2 last at least four rounds. Also, the first match to last at least four rounds in this treatment is the fifth match. Similarly, in the BRT treatment, the first match for which the fifth round within a match is observed is the fifth match. This is the reason why there are missing values in Figure 3 for these treatments.

Figure 3: Cooperation Rates in Rounds 1, 4, and 5 Over Time

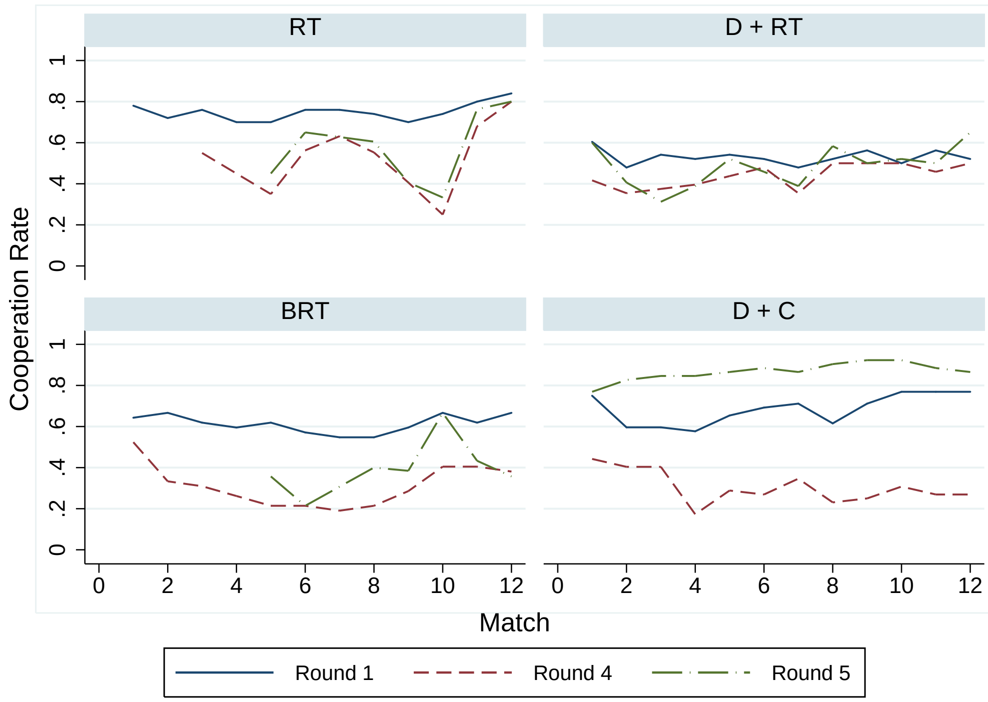

choice in the continuation value coordination game. The first observation is that in RT, cooperation rates in rounds 4 and 5 are very similar, but both tend to be below the Round 1 rate. This is to be expected in an environment without full cooperation.

In D+RT, there is no visible impact of moving from the payoff discounting phase to the random termination phase. However, the fact that cooperation is at the same level in Rounds 1 and later is surprising (since we have found that subjects are more likely to defect after a defection by the other player). The stability of cooperation between Rounds 1 and 5 in D+RT could be consistent with subjects using more-forgiving strategies. However, our findings in Figure 2 show that aggregate response to defection and cooperation exhibit similar differences in D+RT and in RT. The combination of these observations suggests that subjects are more likely to play miscoordinated strategies, such as variations of Tit-For-Tat, rather than Grim in this treatment.

In BRT, cooperation rates in Round 4 are below the Round 1 level, and the Round 5 level is slightly higher. This suggests a slight restart effect between blocks, consistent with the results in Table 4. The restart effect between blocks observed in this treatment can potentially pose a problem for analyzing long-term dynamics within an interaction. However, the restart effect is disappearing over time. If we consider how cooperation between rounds 4 and 5 vary as a function of the match number, we find a statistically significant negative relation.[18] In fact, in the last 3 matches of the first phase, the difference between rounds 4 and 5 is decreasing, and in match 12, there is no restart effect.

D+C is the treatment most different from RT. Not only is cooperation much less likely in Round 4 than in Round 1, but the difference also is increasing over time. In addition, it is the only treatment in which cooperative choices in Round 5 are more frequent than in Round 1. From the graph, we can see a dramatic change in behavior from Round 4 to Round 5.

Figure 4: Cooperation by Round, Matches 7 to 12

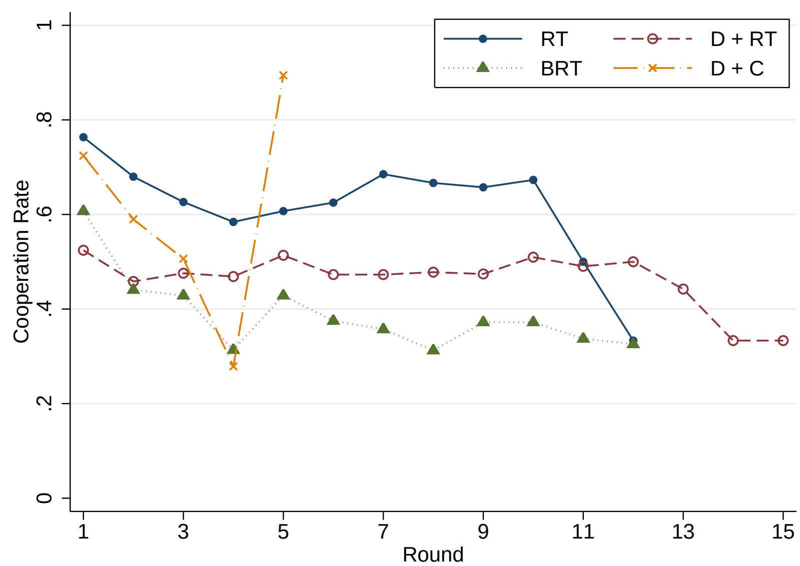

Some of these patterns are better investigated by looking at behavior within a match, and this is what is done in Figure 4. The figure focuses on the second half of the first part of the experiment - namely, matches 7 to 12. We omit the earlier data to look at more-stable and more-experienced behavior, although the picture changes very little when we include all the data. For both RT and BRT, the general pattern seems to be an early decrease in average cooperation followed by a relatively stable period.[19] The D+RT treatment, con-

[18]This is done using an ordered probit (cooperation can decrease, stay the same, or increase) clustering by session.

[19]The drop at the end for RT can be explained by the fact that the sample of matches is changing as we look

sistent with earlier observations, presents a smaller decrease in cooperation in the rounds that follow the first one. The BRT treatment displays the restart effect mentioned earlier, but if the data is broken down into smaller groups of match (see Figure 7 in the Appendix), it is clear that this effect disappears with experience. The most dramatically different behavior is observed in the case of D+C. There is an important decrease in cooperation over the first four rounds, followed by a very high rate of cooperation in the coordination game.

3.2 Discounting + Coordination

Behavior in D+C suggests that subjects do not use dynamic incentives as in the other treatments. To investigate this further, we examine if and how play in the coordination game depends on play in the first four rounds, which, as theory suggests, would be the case if subjects punished deviations from cooperative agreements.

Table 5 reports estimates of a probit where cooperation in the coordination game (in Round 5) is regressed on the choice in Round 1, as well as other controls. Only the choice in Round 1 is included to avoid the endogeneity problems that other rounds would generate. We include subjects' average actions in previous matches (besides the ones under consideration - e.g., in the column for matches 4-6 - this is computed from matches 1-3 and 7-12) to allow strategies in the coordination game to be type-dependent, regardless of the opponent's actions in the previous rounds. The results are reported for various experience levels. In the last column, we also include as a point of comparison the estimates of the same specification for matches 10-12 of our other treatments.[20]

The results for matches 7-9 stand out, as they indicate that play in the coordination game depends to some extent on the Round 1 outcome. The negative impact of individual cooperation in Round 1, combined with the significant positive impact of the interaction term, indicates that a subject is least likely to cooperate in Round 5 if he cooperated in Round 1 when his opponent defected. This suggests that some subjects employ defection in the last round as a punishment strategy if they were a cooperator facing a defector in the first round. Moreover, an F-test shows that the summation of the first three terms is

across rounds. In particular, there are 52 observations in Round 10, but only 12 for Round 11. To eliminate variations due to the fact that the matches that have x-many rounds vary, Figure 6 in the Appendix presents a similar graph for all matches that lasted at least five rounds, but only looking at the first five rounds. In that case, the sample is of the same size for each round of a treatment and as can be seen similar patterns are observed.

[20] In the case of the last column, including dummy variables for all but one treatment does not qualitatively change the results. Furthermore, since both dummies are not statistically significant, nor are they jointly different from zero, they are not included.

Table 10 in the Appendix reports marginal effects, and Table 11 reports correlated random effects estimates.

Table 5: Probit Estimate of the Factors Affecting Cooperation in the Coordination Game of Treatment D+C (Round 5)

	Matches				
	1-3	4-6	7-9	10-12	10-12
	D+C	D+C	D+C	D+C	RT, D+RT, BRT
Partner cooperated	0.491***	0.034	-0.833*	0.444	0.566*
in Round 1	(0.091)	(0.607)	(0.492)	(0.513)	(0.301)
Subject own cooperation	0.349	-0.022	-2.194**	1.211	-0.004
in Round 1	(0.218)	(0.253)	(0.931)	(1.025)	(0.465)
Both cooperated	0.378	0.761	3.509***	0.307	1.118***
in Round 1	(0.318)	(0.710)	(1.286)	(0.738)	(0.366)
Subject own mean cooperation	0.719	-0.945	-0.859	-2.989	1.760**
in rounds < 5 of other matches	(0.612)	(0.735)	(1.405)	(2.355)	(0.694)
Subject own mean cooperation	2.538***	4.337***	4.766***	5.300***	0.302
in Round 5 of other matches	(0.370)	(0.715)	(0.873)	(1.447)	(0.720)
Constant	-2.184***	-2.067***	-0.894	-2.383***	-1.554***
	(0.565)	(0.569)	(0.697)	(0.580)	(0.220)
N	156	156	156	156	242

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

significantly different from 0 ($p < 0.01$). This implies that outcomes of (C, C) and (D, D) in Round 1 generate different levels of cooperation in Round 5.

However, the regression results look significantly different for every block of matches. In particular, in the last block, the impact of the Round 1 outcome on the coordination game decreases dramatically; the magnitude of the first three terms decline and their significance disappear. For the last three matches, the only variable that is predictive of play in the coordination game is the average action in the previous matches. In comparison, the estimation results for our random termination treatments at the same experience level look very different. We see a clear impact of the Round 1 outcome on the cooperation decision in Round 5. More importantly, the magnitude of the effect is much more pronounced. For instance, when both subjects cooperate in Round 1, the increase in the probability of cooperation in Round 5, as opposed to the case where his opponent defects, is 62 percentage points for the cases of RT, D+RT, and BRT taken together, while it is only 3 percentage points in the case of treatment D+C.

A less statistical but very telling way of seeing the disconnect between the choice in Round 5 and the choices before that is presented in Figure 5. On the x-axis is the number of cooperative choices in the first four rounds by either of the players in a pair (hence, the minimum is 0 and the maximum 8 if both players cooperate in all four rounds), and on the y-axis is the probability that a subject cooperates in Round 5. As can be seen, in D+C, the relation is mostly flat, whereas in all other treatments, there is an important positive relationship.[21]

This is in line with the observation from Figure 3 that as subjects gain experience, actions in the coordination game become independent of the evolution of play in the previous rounds. When subjects don't use the coordination game to create dynamic incentives, the first four rounds of the match turn into a finitely repeated game. The rapid decline in cooperation rates observed within a match (for the first four rounds) in this treatment closely resembles behavior reported in finitely repeated PD experiments, providing further evidence that behavior in this treatment does not capture the presence of dynamic incentives.[22]

To summarize our results so far, all four methods find treatment effects in the same direction. However, the magnitude of the effect is slightly different, and so are the coop-

[21]The figure is almost identical if, instead, the y-axis is computed for Round 5 and all subsequent rounds.

[22]See, for instance, Bereby-Meyer and Roth [2006] and Dal Bó [2005]. Note that a decline in the cooperation rate within a match is commonly observed in finitely repeated PDs, but whether or not this unravels over time is a matter of debate. For instance, Andreoni and Miller [1993] and Friedman and Oprea [2012] report that the fall in cooperation happens later, as subjects gain experience.

Figure 5: Cooperation in Round 5, Matches 7 to 12

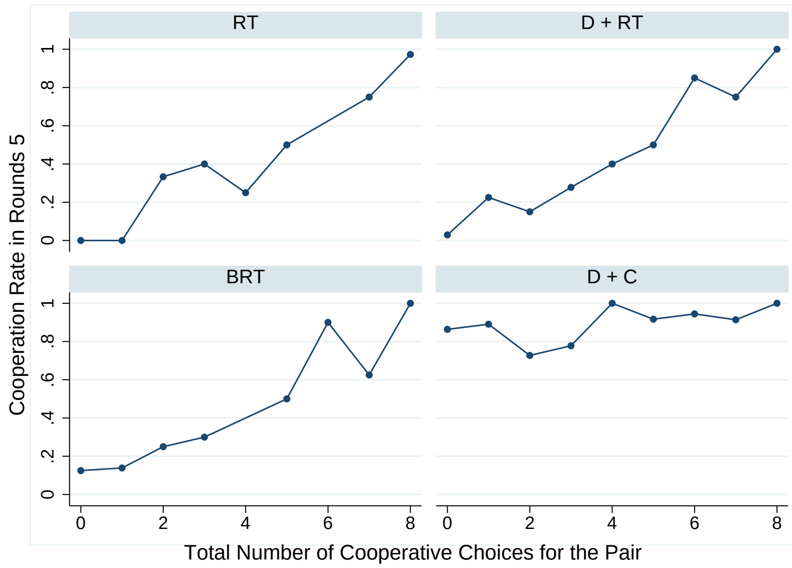

eration rates within a match. There are also variations across treatments in how behavior depends on past experiences over matches. The D+C treatment is substantively different from the other three in that it does not seem to induce dynamic incentives. Next section analyzes the strategies that subjects use.

3.3 Strategies

This section investigates whether the different methods of implementing infinitely repeated games, although theoretically equivalent, lead to different strategic choices. In the remainder of the paper, we will not study behavior in the D+C treatment since it is substantially different from the other treatments. We try to understand if the aggregate differences observed across the different methods can be explained by differences in the strategies used.

For this, we employ the strategy estimation procedure introduced in Dal Bó and Fréchette [2011]), referred to as Strategy Frequency Estimation Method or SFEM, and also used in Fudenberg et al. [2012], Rand, Fudenberg, and Dreber [2013], Dal Bó and Fréchette [2012], and Vespa [2013]. This approach to estimating strategies consists of first computing a vector of the choices that would be prescribed to that subject by each strategy under consideration, given the history of play. The econometric procedure, a mixture model, acts as a signal detection method and estimates via maximum likelihood how close the actual choices are from the prescriptions of each strategy. The key estimates obtained are the proportion in which each strategy is observed in the population sample.

We denote the choice made by subject i in Round r of match m by c_{imr} and the choice that a strategy k indicates to make in Round r of match m for subject i by $s^k_{imr}\left(y_{jm1},...,y_{jm(r-1)};\, s^k_{im1},...,s^k_{im(r-1)}\right)$ if $r > 1$, while the strategy does not depend on previous states or signals in Round 1. Finally, the indicator variable I takes value one if the choice corresponds to the strategy in that round of a given match and zero otherwise: $I^k_{imr} = 1\left\{c_{imr} = s^k_{imr}(\cdot)\right\}$. The probability that a choice corresponds to the one prescribed by a given strategy is modeled as $Pr\left(I^k_{imr}\right) = \frac{1}{1+\exp\left(\frac{-1}{\gamma}\right)} \equiv \beta$, where γ is a parameter to be estimated. This can be motivated from a model in which subjects follow a strategy but make mistakes, as in $c'_{imr} = 1\left\{s'_{imr}(\cdot) + \gamma\varepsilon_{imr} \geq 0\right\}$, where c'_{imr} takes value 0 and 1, s'_{imr} is coded as 1 when the choice should be 1 and -1 when it should be 0 and ε has a logistic distribution. When reporting results, we will report β, as it gives an indication of the quality of the fit, something difficult to read from γ; random choices imply a β of $\frac{1}{2}$, and choices exactly as predicted imply a β of 1. The likelihood that the observed choices for subject i were generated by strategy k are given by

$$prob_i(s^k) = \prod_{M_i}\prod_{R_{im}}\left(\frac{1}{1+\exp(-1/\gamma)}\right)^{I^k_{imr}}\left(\frac{1}{1+\exp(1/\gamma)}\right)^{(1-I^k_{imr})}$$

where M is the set of matches and R the number of rounds in each match. Combining this across subjects and allowing for multiple strategies, each present in a different frequency, ϕ^k, we obtain the following loglikelihood:

$$\sum_I \ln\left(\sum_K \phi^k prob_i(s^k)\right)$$

for the set of strategy K and of subjects I. The parameters of interest ϕ^k give the probability of observing each strategy.

Table 6 reports the estimated frequency of strategies for each of the treatments.[23] In this table, we include only strategies that have a statistically significant positive population share in at least one of our treatments. In our estimation, we include the 20 strategies

[23] Standard errors are obtained by bootstrapping. This is done by first drawing sessions and then subjects (both with replacement).

considered in Fudenberg et al. [2012], which cover the commonly considered strategies in repeated prisoner's dilemma experiments. We refer the reader to Table 12 in the Appendix for a description of these strategies. Our estimation results for the entire set of strategies can be found in Table 13 in the Appendix. In Table 13, to ensure that the treatment differences we observe are not driven by differences in observation length, we also estimate strategies for a subset of observations in D+RT and BRT, focusing only on the rounds in each match that are observed under all methods.[24]

Table 6: Distribution of Estimated Strategies

	RT	**D+RT**	**BRT**
Always Defect	0.14	0.26^{**}	0.25^{***}
	(0.098)	(0.107)	(0.072)
Grim	0.32^{***}	0.10	0.21^{***}
	(0.098)	(0.061)	(0.077)
Tit-For-Tat	0.39^{***}	0.22^{**}	0.33^{***}
	(0.118)	(0.095)	(0.089)
2 Tits-For-2 Tats	0.06	0.06^{***}	0.07^{*}
	(0.044)	(0.021)	(0.043)
Suspicious Tit-For-Tat	0.02	0.18^{***}	0.05
	(0.061)	(0.057)	(0.036)
β	0.935	0.936	0.901

Bootstrapped standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results reported in Table 6 stand out, as they indicate that Always Defect, Grim and versions of Tit-For-Tat account for the most common strategies in all our treatments. Moreover, the results indicate that variations in how subjects behave across the different methods can be linked to differences in the frequency of these strategies. First, we observe

[24]For example if a match ended in Round 5 in RT, we look only at the data from the first five rounds for the equivalent match in D+RT and BRT.
The results are very similar when using all the data versus only the subset. For D+RT, the difference is that Grim and Win-Stay-Lose-Shift are statistically significant using the subset but not all the data, while 2-Tits-For-Tat and 2-Tits-For-2-Tats are statistically significant using all the data but not in the subset. This seems to suggest that identifying strategies with longer memory is more difficult with fewer choices per match. For BRT, Suspicious Tit-For-Tat is not significant using all the data but it is in the subset. On the other hand, 2-Tits-For-2-Tats is not significant in the subset, but significant for all the data. These are relatively small differences considering the number of strategies; and except in the case of Grim, each of these strategies represent less than 10% of the data. Fudenberg et al. [2012] and Dal Bó and Fréchette [2012] have already pointed out that this method does not perform as well at identifying strategies that are present in small proportions.

unconditional defection to be lowest, at 14 percent in RT. In D+RT and BRT, 26 and 25 percent of the population always defects. This suggests that subjects use more-cooperative strategies under RT. However, these differences are not statistically significant.[25] Categorizing strategies as cooperative vs. uncooperative and comparing the relative frequency of these categories provides further evidence in this direction. If we compare the population share of all strategies that start with defection in Round 1, we find it to be 18 percent in RT, 46 percent in D+RT (different from RT at $p < 0.1$) and 33 percent in BRT (not statistically significant from RT). This result is in line with our findings in Figure 4 and Table 3 which report Round 1 cooperation to be highest for RT.

A feature of our data is the differences in our treatments in the stability of cooperation within a match. In Figure 4, we see a sharp decline in cooperation from Round 1 to Round 4 with RT. There is a similar decline in BRT. With D+RT, however, the cooperation rate is more stable, with Round 1's cooperation rate not significantly different from that of Round 5. This is also seen in Figure 3. This is a surprising result, as the most-studied strategies (such as Grim-trigger or Tit-For-Tat) in the literature predict a breakdown of cooperation when a cooperator meets a defector in the first round. Differences in the strategies used across these methods give us an insight into how this result comes about.

As stated above, there is an important increase in the fraction of strategies that with defection when we move from RT to D+RT (18 vs. 46 percent). Looking at the most important changes (ranked by magnitude) we see that: (1) a decrease in Grim ($p < 0.1$); (2) a decrease in Tit-For-Tat ($p > 0.1$); (3) an increase in Suspicious Tit-For-Tat ($p > 0.1$); and (4) an increase in AD ($p > 0.1$). Clearly, few of those differences are statistically significant indicating that identifying the exact channel is difficult. However, the joint hypothesis of a change in Grim and Tit-For-Tat as well as the joint hypothesis of a change in AD and Suspicious Tit-For-Tat are statistically significant ($p < 0.05$ and $p < 0.1$ respectively).

The fact, noted above, that the cooperation rate does not drop as sharply in D+RT can be explained by the decrease in popularity of Grim and the increase in popularity of Suspicious Tit-For-Tat. When a cooperator meets a defector in Round 1, for cooperation to continue in the future, the defector must be playing strategies that potentially switch back to cooperation, and the cooperator must be playing strategies with limited punishment. We see that 20 percent of the population in D+RT plays strategies that start with

[25]To do hypothesis testing between the treatments, we pool data from two treatments and rerun our estimation procedure, allowing for different distribution of strategies in the separate treatments, and use a Wald test. Point estimates for the distribution of strategies following this method are identical to the results we find when the estimation is done separately for each treatment.

defection in Round 1, but possibly switch back and settle on cooperation, depending on the partner's response.[26] The corresponding share is four percent in RT and eight percent in BRT. Additionally, we see that while 32 percent of the population is playing Grim-trigger strategies in RT and 21 percent in BRT, this share is only ten percent in D+RT. Tit-For-Tat type strategies in this treatment imply limited punishment.

Strategy choices in the BRT treatment are fairly close to what is observed in the RT treatment. Although Grim is less popular than under RT, the drop is not as important as under D+RT, and the two are not statistically different. On the other hand, the increase in AD is almost as large as for D+RT, but again it is not statistically different. Finally, unlike under D+RT, Suspicious Tit-For-Tat sees only a very modest increase which is not statistically significant.

4 Discussion

We discuss our results in reverse order, going from the strategies to the aggregate results. The most striking differences in the strategies that subjects use can be found when comparing D+RT to RT. Ex-post, it might not seem too surprising that behavior in BRT is closer to RT than the behavior in D+RT: Since, more rounds per match are experienced in that treatment than in the RT treatment, but fewer rounds are used for payment in BRT than in D+RT (and exactly the same as in RT). In this experiment, going from RT to D+RT affects the number of rounds per match, and, although this should not theoretically affect the strategic environment, our results suggest that it has an effect on the type of strategies people adopt.

Table 7: Change in Popularity of Strategies (Percentage Point Difference)

	Dal Bó and Fréchette 2011		**This Study**
Joint Cooperation Payoff R:	**48 to 32 (decrease of 33%)**		
Continuation Probability δ:	**0.5 to 0.75**	**0.75 to 0.9**	**RT to D+RT**
Always Defect	18	1	12
Grim	-26	-20	-22
Tit-For-Tat	-8	6	-17
Suspicious Tit-For-Tat	8	11	16

[26]The easiest way to see this share of the population is to sum the population share playing strategies that start with defection and then subtract the share playing Always Defect.

When going from RT to D+RT, the average number of interactions increases but the value of cooperation versus defection does not. In a standard experiment using RT, if δ increases the average number of interactions increase but the relative value of cooperation also increases. On the other hand, if the payoff to joint cooperation (R) is decreased, then the relative value of cooperation is decreased but the average number of interactions is not. Hence, a simultaneous increase in δ and decrease in R increases the average number of interactions but can reduce the change in the relative value of cooperation. A prior experiment, Dal Bó and Fréchette [2012], proposes a method to elicit strategies from subjects, and varies both δ and R across treatments. Four of the five treatments in that paper can be used to perform such comparisons: We can compare the treatments $(48, 0.5)$ - where the first number is the payoff in the case of joint cooperation and the second number is δ - to $(32, 0.75)$; and we can compare $(48, 0.75)$ to $(32, 0.9)$.[27] For both comparisons, we find that as the average number of interactions increases: (1) the frequency of AD (Always Defect) increases; (2) the popularity of Grim is diminished; and (3) Suspicious Tit-For-Tat becomes more popular. The only exception is Tit-For-Tat in which case one of the two comparisons goes in the opposite direction. The exact numbers are provided in Table 7.[28] These three patterns are the same as the ones we observe going from RT to D+RT. This is suggestive evidence that for discount factors (and payoffs) such that joint cooperation can be supported in equilibrium, increasing δ may affect strategy choice via multiple channels, one of which is that it changes the average number of interactions. Note also that in the Dal Bó and Fréchette [2012] experiment, if we look at the impact of increasing δ holding R constant, most of the effects on strategies reported in Table 7 are reversed. This highlights the importance of separating the effects of the average number of interactions from those of the relative value of cooperation in order to understand the impact of an increase in δ on strategic choices.

It remains an open question why subjects who know that they will interact for more rounds migrate towards using AD and Suspicious Tit-For-Tat. It does seem sensible, on the one hand, that, when faced with more players using AD, one might want to be more "cautious" and play Suspicious Tit-For-Tat, but this does not answer the question of why the fraction of AD increases in the first place. One might think that the increased leniency and forgiveness of the cooperative strategies used under D+RT make defection more profitable, but that turns out not to be the case given that it is more than compensated

[27]The payoff to defecting when the other cooperates is 50, the payoff to cooperating when the other defects is 12, and the joint defection payoff is 25.

[28]These numbers are obtained using the elicitation methods proposed in that paper combining results from both elicitation methods; however, using the estimation method described in this paper on the choices made in that experiment reveals the same patterns.

by the decrease in payoffs caused by the fact that more subjects defect. In fact the average payoffs of those who defect in every round of a match are 9% lower in the D+RT treatment than in the RT treatment ($p > 0.1$). This is due to the fact that in that treatment one is much more likely to encounter someone who also defects (subjects who always defect face 18% of cooperation under RT but only 10% under D+RT; $p < 0.01$).[29] On the other hand, the decrease in the popularity of Grim seems intuitive: As matches become longer, subjects shy away from strategies that might get them stuck in punishment forever. The fact that similar results are observed in Dal Bó and Fréchette [2012] once we adjust for the attractiveness of cooperation suggests that this phenomenon is robust.

We emphasize, however, that these changes in strategy choice are not important enough to affect the main comparative result, and this suggests that all three methods with random termination can be used to induce infinitely repeated games in the laboratory. All four methods generate sharp comparative statics: Cooperation levels drop significantly when parameters of the stage game are changed to make mutual cooperation theoretically unsustainable. However, analysis of behavior within a match indicates that the cooperation that is observed with the D+C method is not supported by dynamic incentives. With this method, subjects treat the coordination game as independent of the history of play, and they appear to treat the rest of the game as a finitely repeated game. Of course, this does not mean that the D+C method is not useful for other purposes, however it induces a markedly different environment.

The strategic variations across implementation methods discussed above result in small variations in cooperation rates across treatments. In particular, cooperation rates are highest with the RT method. This is extremely surprising, as it goes against the intuition and folk wisdom about the impact of longer interactions on cooperation in finitely repeated PDs (as well as being the opposite of the impact of risk aversion).

Additionally, we find that behavior with this method - discounting followed by random termination, as well as with block random termination - is significantly less sensitive to past experiences within a session. These findings can potentially make these methods, especially discounting followed by random termination, more attractive than RT for experimenters who need to limit the impact of past experiences.

From the perspective of testing the implications of infinitely repeated games in the

[29] When we compare payoffs of those who always defect to those who perfectly follow Grim, we see that average payoffs are much lower for the unconditional defectors in both treatments. It is 43% lower in RT (although not statistically significant), and 45% lower in D+RT ($p < 0.01$). Furthermore, when tested jointly, the difference in average payoffs of AD and Grim is not statistically different when we compare RT and D+RT.

laboratory, we find that the different methods present a trade-off between observing longer games and generating larger comparative statics with respect to parameters of the stage game. All three methods using random termination generate behavior that is consistent with theory. Thus, the *preferred* implementation method will depend on the research question at hand and the type of applications that are meant to be modeled in the laboratory. In particular, it seems that if one needs to observe more rounds, but wants the strategy used to be close to what they are under RT, then the best choice is BRT. On the other hand, if a researcher needs to reduce the expected variation in payment, then D+RT would be the best choice. If experiencing many matches is important (because how to play in a specific game is difficult), then the original method, RT, should be the best option. Finally, if the design is tightly tied to a situation in the field, then the best method would depend on whether typically that situation involves longer or shorter interactions.

In a broader context, our results also suggest that subjects respond to payoff discounting and probabilistic continuation in slightly different ways. To our knowledge, our experiment is the first to report behavioral differences across these environments. Moreover, the impact of longer average interactions while keeping discounting constant are opposite of what we would have expected. Moving outside of the laboratory, these results suggests that situations in which agents are very patient, but relationships are likely to terminate for exogenous reasons, may lead to different strategic choices and, consequently, different dynamics than situations in which agents are less patient, but interactions are less likely to end; even if, from a theoretical perspective, these two environments allow for the same set of equilibrium outcomes.

References

James Andreoni and John H Miller. Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. *The Economic Journal*, 103(418), 1993.

Yoella Bereby-Meyer and Alvin E. Roth. Learning in noisy games: Partial reinforcement and the sustainability of cooperation. *The American Economic Review*, 96(4):1029–1042, September 2006.

M. Blonski, P. Ockenfels, and G. Spagnolo. Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *The American Economic Journal: Microeconomics*, 3(3):164–192, 2011.

Jordi Brandts and Gary Charness. Hot vs. cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, 2(3):227–238, 2000.

Luis M. B. Cabral, Erkut Ozbay, and Andrew Schotter. Intrinsic and instrumental reciprocity: An experimental study. Working Paper, 2011.

David J. Cooper and Kai-Uwe Kühn. Communication, renegotiation, and the scope for collusion. Working Paper, 2011.

Pedro Dal Bó. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *The American Economic Review*, pages 1591–1604, December 2005.

Pedro Dal Bó and Guillaume R. Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review*, 101(1): 411–429, 2011.

Pedro Dal Bó and Guillaume R. Fréchette. Strategy choice in the infinitely repeated prisoners dilemma. Working Paper, 2012.

A. Dreber, D. G. Rand, D. Fudenberg, and M. A. Nowak. Winners don't punish. *Nature*, 452:348–351, 2008.

Jim Engle-Warnick and Robert L. Slonim. Inferring repeated-game strategies from actions: Evidence from repeated trust game experiments. *Economic Theory*, 28(3):603–632, 2006.

Guillaume R. Fréchette. Session-effects in the laboratory. *Experimental Economics*, 15 (3):485–498, 2011.

Daniel Friedman and Ryan Oprea. A continuous dilemma. *The American Economic Review*, 102(1), 2012.

Drew Fudenberg, David G. Rand, and Anna Dreber. Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world. *The American Economic Review*, 102: 720–49, 2 2012.

George J. Mailath and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. New York, Oxford University Press, 2006.

J. Keith Murnighan and Alvin E. Roth. Expecting continued play in prisoner's dilemma games. *Journal of Conflict Resolution*, 27(2):279–300, 1983.

David G. Rand, Drew Fudenberg, and Anna Dreber. It is the thought that counts: The role of intentions in reciprocal altruism. Working Paper, 2013.

Alvin E. Roth and J. Keith Murnighan. Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17:189–198, 1978.

Dan R. Schley and John H. Kagel. How economic rewards affect cooperation reconsidered. Working Paper, 2013.

Katerina Sherstyuk, Nori Tarui, and Tatsuyoshi Saijo. Payment schemes in infinite-horizon experimental games. *Experimental Economics*, 16:125–153, March 2013. Experimental Economics.

Emanuel I. Vespa. Cooperation in dynamic games: An experimental investigation. Working Paper, 2013.

Alistair Wilson and Hong Wu. Dissolution of partnerships in infinitely repeated games: An experimental examination of termination clauses. Manuscript in Preparation, 2013.

Rami Zwick, Amnon Rapoport, and John C. Howard. Two-person sequential bargaining behavior with exogenous breakdown. *Theory and Decision*, 32:241–268, 1992.

Figure 6: Cooperation by Round, Matches 7 to 12 [Only Including Matches That Lasted At Least 5 Rounds]

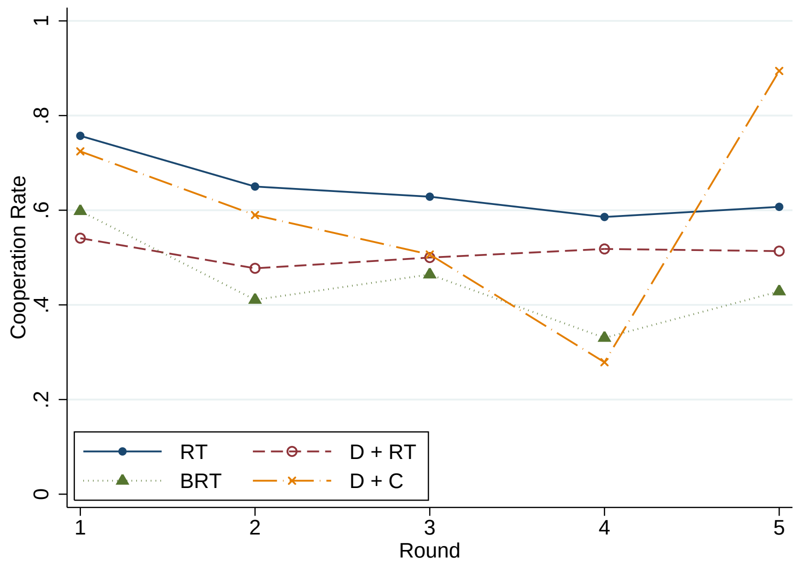

Figure 7: Cooperation by Round in BRT

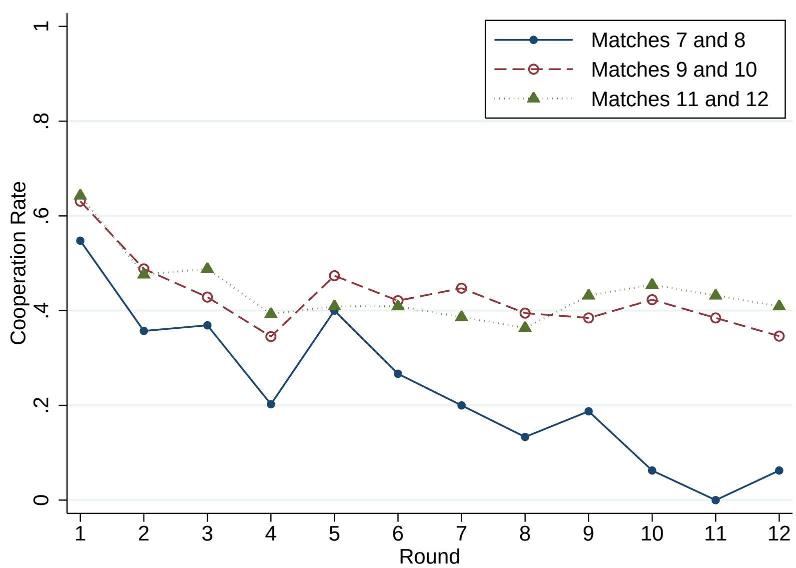

Table 8: Random Effects Probit Estimate of the factors affecting the evolution of cooperation (Matches 2 to 12)

Dependent Variable: Cooperation in Round 1

	RT	D+RT	BRT	BRT	D+C
Partner cooperated in	0.806***	0.434***	0.568***	0.559**	0.604**
Round 1 of previous match	(0.231)	(0.147)	(0.219)	(0.230)	(0.251)
Number of rounds	0.156	-0.017	-0.141	-0.142	
in previous match	(0.099)	(0.061)	(0.120)	(0.122)	
Number of rounds	-0.006	-0.000	0.0098	0.0017	
in previous match sq.	(0.00609)	(0.00662)	(0.01039)	(0.00837)	
Two blocks				0.200	
in previous match				(0.260)	
Three blocks				0.838***	
in previous match				(0.098)	
Match number	0.019	0.017	0.0199	0.0181	0.1118***
	(0.0641)	(0.0656)	(0.0382)	(0.0364)	(0.0171)
Subject cooperated in	3.380***	4.290***	1.606***	1.628***	2.665***
Round 1 of match 1	(1.046)	(0.491)	(0.151)	(0.149)	(0.296)
Constant	-2.502***	-2.932***	-0.928***	-0.878***	-2.084***
	(0.859)	(0.883)	(0.089)	(0.145)	(0.415)
$\frac{\sigma^2}{\sigma^2+1}$	0.53	0.72	0.65	0.65	0.67
N	550	528	462	462	572

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

σ^2 is the variance of the subject specific random effects.[30]

Table 9: Probit Marginal Effects Estimate of the Factors Affecting the Evolution of Cooperation (Matches 2 to 12) See Table 4

	RT	D+RT	BRT	BRT	D+C
Partner cooperated in	0.213**	0.172***	0.070	0.068	0.207***
Round 1 of previous match	(0.086)	(0.065)	(0.058)	(0.067)	(0.153)
Number of rounds	0.059*	0.041***	-0.041	-0.037**	
in previous match	(0.032)	(0.013)	(0.028)	(0.018)	
Number of rounds	-0.00334	-0.00333*	0.00288	0.00082	
in previous match sq.	(0.00208)	(0.00179)	(0.00246)	(0.00142)	
Two blocks				0.02277	
in previous match				(0.03918)	
Three blocks				0.1654***	
in previous match				(0.05183)	
Match number	0.000079	0.00279	0.00624	0.00581	0.02196***
	(0.00986)	(0.0153)	(0.0091)	(0.00874)	(0.00168)
Subject cooperated in	0.673***	0.755***	0.350**	0.347**	0.553***
Round 1 of match 1	(0.044)	(0.028)	(0.153)	(0.154)	(0.142)
N	550	528	462	462	572

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Probit Marginal Effects Estimate of the Factors Affecting Cooperation in the Coordination Game of Treatment D+C (Round 5) See Table 5

	Matches				
	1-3	4-6	7-9	10-12	10-12
	D+C	D+C	D+C	D+C	RT, D + RT, BRT
Partner cooperated	0.114***	0.004	-0.016	0.021	0.218*
in Round 1	(0.041)	(0.073)	(0.012)	(0.025)	(0.116)
Subject own cooperation	0.079	-0.003	-0.048***	0.097	-0.001
in Round 1	(0.066)	(0.029)	(0.014)	(0.080)	(0.178)
Both cooperated	0.078	0.089	0.240**	0.012	0.399***
in Round 1	(0.060)	(0.075)	(0.096)	(0.038)	(0.105)
Subject own mean cooperation	0.154	-0.112*	-0.021	-0.108**	0.674**
in rounds < 5 of other matches	(0.156)	(0.059)	(0.052)	(0.044)	(0.265)
Subject own mean cooperation	0.543***	0.514***	0.114	0.191	0.116
in Round 5 of other matches	(0.111)	(0.111)	(0.091)	(0.143)	(0.273)
N	156	156	156	156	242

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Random Effects Probit Estimate of the Factors Affecting Cooperation in the Coordination Game of Treatment D+C (Round 5) See Table 5

	Matches				
	1-3	4-6	7-9	10-12	10-12
	D+C	D+C	D+C	D+C	RT, D + RT, BRT
Partner cooperated	0.469***	-0.181	-0.833*	0.747	0.981*
in Round 1	(0.108)	(0.475)	(0.492)	(0.841)	(0.530)
Subject own cooperation	0.245	-0.268	-2.194**	1.888	0.209
in Round 1	(0.177)	(0.491)	(0.931)	(1.788)	(0.716)
Both cooperated	0.852***	1.027	3.509***	0.469	1.244***
in Round 1	(0.173)	(0.717)	(1.286)	(0.715)	(0.312)
Subject own mean cooperation	0.920	-0.949	-0.859	-4.000	2.889**
in rounds < 5 of other matches	(0.986)	(0.904)	(1.405)	(4.086)	(1.350)
Subject own mean cooperation	3.384***	5.251***	4.766***	7.780*	0.097
in Round 5 of other matches	(1.040)	(1.301)	(0.873)	(4.310)	(1.015)
Constant	-2.749***	-2.341***	-0.894	-3.789*	-2.396***
	(1.027)	(0.814)	(0.697)	(1.935)	(0.888)
$\frac{\sigma^2}{\sigma^2+1}$	0.45	0.31	0.00	0.54	0.52
N	156	156	156	156	242

Clustered (session level) standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

σ^2 is the variance of the subject specific random effects.

Table 12: Description of Strategies Estimated

Name of Strategy	Description
Always Defect	always play D
Always Cooperate	always play C
Grim	play C until either player plays D, then play D forever
Tit-For-Tat	play C unless partner played D last round
Win-Stay-Lose Shift	play C is both players chose the same move las round, otherwise play D
T2	play C until either player deviates, then play D twice and return to C
Tit-For-2 Tats	play C unless partner played D in both of the last rounds
Tit-For-3 Tats	play C unless partner played D in all of the last 3 rounds
2 Tits-For-1 Tat	play C unless partner played D in either of the last 2 rounds
2 Tits-For-2 Tats	play C unless partner played 2 consecutive Ds in either of the last 3 rounds
Lenient Grim 2	play C until 2 consecutive rounds occur in which either player played D, then play D forever
Lenient Grim 3	play C until 3 consecutive rounds occur in which either player played D, then play D forever
Tit-For-Tat 2	play C of both played C in the last 2 rounds, both played D in the last two rounds, or both played D and C
False cooperator	play C in the first round, then D forever
Suspicious Tit-For-Tat	play D in the first round, then TFT
Suspicious Tit-For-2 tats	play D in the first round, then Tit-For-2 Tats
Suspicious Tit-For-3 tats	play D in the first round, then Tit-For-3 Tats
Suspicious lenient Grim 2	play D in the first round, then Grim 2
Suspicious lenient Grim 3	play D in the first round, then Grim 3
Alternator	play D in the first round, then alternate between C and D

Table 13: Distribution of Estimated Strategies

	RT	**D+RT (all)**	**D+RT (subset)**	**BRT (all)**	**BRT (subset)**
Always Defect	0.14	0.26**	0.29***	0.25***	0.26***
	(0.098)	(0.107)	(0.092)	(0.072)	(0.077)
Always Cooperate	0.00	0.00	0.03	0.00	0.02
	(0.041)	(0.084)	(0.05)	(0.034)	(0.044)
Grim	0.32***	0.10	0.14**	0.21***	0.20***
	(0.098)	(0.061)	(0.073)	(0.077)	(0.077)
Tit-For-Tat	0.39***	0.22**	0.22***	0.33***	0.27***
	(0.118)	(0.095)	(0.084)	(0.089)	(0.095)
Win-Stay-Lose Shift	0.01	0.00	0.04*	0.00	0.00
	(0.05)	(0.061)	(0.022)	(0.041)	(0.049)
T2	0.01	0.01	0.00	0.00	0.00
	(0.054)	(0.048)	(0.075)	(0.014)	(0.014)
Tit-For-2 Tats	0.00	0.00	0.00	0.00	0.00
	(0.037)	(0.037)	(0.009)	(0.026)	(0.039)
Tit-For-3 Tats	0.00	0.05	0.03	0.02	0.02
	(0.05)	(0.042)	(0.02)	(0.028)	(0.032)
2 Tits-For-1 Tat	0.00	0.09*	0.04	0.02	0.02
	(0.055)	(0.049)	(0.043)	(0.068)	(0.037)
2 Tits-For-2 Tats	0.06	0.06***	0.00	0.07*	0.07
	(0.044)	(0.021)	(0.019)	(0.043)	(0.045)
Lenient Grim 2	0.00	0.02	0.03	0.00	0.00
	(0.014)	(0.036)	(0.046)	(0.046)	(0.047)
Lenient Grim 3	0.00	0.00	0.00	0.00	0.00
	(0.022)	(0.003)	(0.021)	(0.043)	(0.024)
Tit-For-Tat 2	0.03	0.00	0.00	0.02	0.07
	(0.039)	(0.024)	(0.01)	(0.05)	(0.055)
False cooperator	0.00	0.00	0.00	0.00	0.00
	(0.012)	(0.007)	(0.002)	(0.03)	(0.031)
Suspicious Tit-For-Tat	0.02	0.18***	0.16***	0.05	0.05*
	(0.061)	(0.057)	(0.049)	(0.036)	(0.03)
Suspicious Tit-For-2 Tats	0.00	0.00	0.00	0.03	0.03
	(0.027)	(0.059)	(0.028)	(0.037)	(0.055)
Suspicious Tit-For-3 Tats	0.00	0.00	0.00	0.00	0.00
	(0.001)	(0.028)	(0.021)	(0.051)	(0.041)
Suspicious lenient grim 2	0.00	0.00	0.00	0.00	0.00
	(0.002)	(0.008)	(0.009)	(0.023)	(0.037)
Suspicious lenient grim 3	0.00	0.02	0.02	0.00	0.00
	(0.004)	(0.039)	(0.038)	(0.004)	(0.011)
Alternator	0.02	0.00	0.00	0.00	0.00
β	0.935	0.936	0.936	0.901	0.898

Bootstrapped standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$