

# Infinitely repeated games in the laboratory: four perspectives on discounting and random termination

Guillaume R. Fréchette<sup>1</sup> · Sevgi Yuksel<sup>2</sup>

Received: 26 September 2014/Revised: 9 August 2016/Accepted: 17 August 2016/  
Published online: 19 October 2016  
© Economic Science Association 2016

**Abstract** This paper compares behavior under four different implementations of infinitely repeated games in the laboratory: the standard random termination method [proposed by Roth and Murnighan (J Math Psychol 17:189–198, 1978)] 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 Sabater-Grande and Georgantzis (J Econ Behav Organ 48:37–50, 2002)] or followed by a coordination game [proposed in (Andersson and Wengström in J Econ Behav Organ 81:207–219, 2012; Cooper and Kuhn in Am Econ J Microecon 6:247–278, 2014a)]. 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. Subject behavior in the other two methods display two features: a higher level of stability in cooperation rates and less dependence on past experience. Estimates of the strategies used by subjects reveal that across implementations, even when the discount rate is the same, if interactions are expected to be longer defection increases and the use of the Grim strategy decreases.

**Keywords** Infinitely repeated games · Discounting · Random termination · Prisoner's dilemma

---

**Electronic supplementary material** The online version of this article (doi:[10.1007/s10683-016-9494-z](https://doi.org/10.1007/s10683-016-9494-z)) contains supplementary material, which is available to authorized users.

---

✉ Guillaume R. Fréchette  
frechette@nyu.edu

<sup>1</sup> New York University, New York, NY, USA

<sup>2</sup> University of California, Santa Barbara, USA

**JEL Classification** C9 · C72 · C73 · C91 · C92

## 1 Introduction

The prisoner's dilemma captures in its simplest form a fundamental tension of interest in the social sciences: the conflict between social efficiency and individual optimality. This tension underlies many interesting interactions, economic and otherwise. Examples include, but are not limited to: Cournot competition, the tragedy of the commons, team production with unobservable effort, natural resource extraction, and public good provision. It has long been recognized that repeating such interactions can introduce dynamic incentives, where cooperation can be rewarded and individualistic behavior can be punished in subsequent rounds: thus enabling cooperation to be sustained in equilibrium. Although the theory of infinitely repeated games has been used extensively to explain cooperation in a variety of environments, it has, for the most part, been silent on the issue of equilibrium selection (when players are sufficiently patient, both cooperation and defection are possible equilibrium actions). A growing number of experimental papers have contributed to the literature on infinitely repeated games, not only by testing theoretical predictions, but also by helping to sharpen the predictions of theory when there is multiplicity of equilibria by providing insights into the determinants of cooperation.

The standard method for implementing infinitely repeated games in the laboratory uses random termination [proposed by Roth and Murnighan (1978)] which links the number of expected repetitions of the stage game to the discount factor. 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.<sup>1</sup> 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.<sup>2</sup> Clearly, there is no link between the expected number of rounds played and the discount factor (as in the standard method) 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:

---

<sup>1</sup> In fact, depending on the paper and the application, one (or both) interpretation has been given [see Mailath and Samuelson (2006) Sect. 4].

<sup>2</sup> 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.

the standard random termination method and three other methods that de-couple the expected number of rounds and the discount factor.

There are several reasons for which an experimenter might want to use alternative methods to implement infinitely repeated games in the laboratory. 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 (2015) 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 % of the time. If subjects' learning is influenced by realized outcomes, then it might be difficult for them to learn to cooperate. In some cases, the desire to de-couple the expected number of rounds and the discount factor comes from the opposite need: to reduce the number of rounds in a supergame. Cooper and Kühn (2014a) 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.

Such considerations raise several questions. 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?<sup>3</sup> 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  $\delta$  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  $\delta$ . After the rounds with certainty are played, there is a fixed known probability  $\delta$  that the match continues for an additional round, and payoffs in these rounds are no longer discounted. This procedure was first introduced by Sabater-Grande and Georgantzis (2002) and has since also been used in Cabral et al. (2014) and Vespa (2015).<sup>4</sup>

---

<sup>3</sup> Zwick et al. (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.

<sup>4</sup> 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 (Feinberg and Husted 1993; Feinberg and Snyder 2002), but this changes the environment to a non-stationary one and, thus, for certain games, can introduce different equilibria. In Feinberg and Husted (1993), for example, which studies collusive behavior in duopoly markets, collusion levels are higher in the first part of a supergame (without random termination) relative to the second part (after random termination is introduced).

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 particular strategies in the infinitely repeated game—namely, the Grim trigger strategy and the strategy of always defecting.<sup>5</sup> This method (D+C) was first used in Cooper and Kühn (2014a).<sup>6</sup>

Finally, we consider a new procedure that we refer to as block random termination (BRT). The aim was to develop a procedure whose equivalence to the standard method did not depend on whether or not behavior is different when payoffs are discounted versus when the game is implemented with random termination (as in the D+RT method). Subjects play as in the standard RT, but in blocks of a pre-announced fixed number of rounds. Within a block, subjects receive no feedback about whether or not the match has continued until that round, but 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. To the best of our knowledge, this method is new.<sup>7</sup> Under certain assumptions (to be discussed later), 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. Hence, our main finding is that for the purpose of studying infinitely repeated games and the use of dynamic incentives in the laboratory, the three implementations we studied that involve random termination are viable alternatives.

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.

Our results also indicate that subjects respond to payoff discounting and probabilistic continuation in slightly different ways. For instance, we find strategy choice in an environment where interactions are likely to be short lived to be

<sup>5</sup> The Grim trigger strategy involves first cooperating, followed by cooperation as long as the other player cooperates, but defection forever if either player defects.

<sup>6</sup> Andersson and Wengström (2012) use coordination games with two pareto ranked equilibria that allows agents to support cooperation in the preceding PD, but not necessarily the specific one used in this paper.

<sup>7</sup> It has since been used by Wilson and Wu (2014).

different from one where interactions are long lived, but agents discount future payoffs. Ex-ante, one might have expected that increasing the average number of interactions while keeping  $\delta$  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  $\delta$ , holding payoffs constant, leads to higher cooperation rates. However, estimation of strategies used by subjects in these different environments show 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, and subjects become less likely to support cooperation using a Grim strategy.

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 versus 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 Sect. 3, we discuss the results. We conclude, in Sect. 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).<sup>8</sup>

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

---

<sup>8</sup> 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.

	G	AD
G	$R_{1-\delta}^{\delta\rho}, R_{1-\delta}^{\delta\rho}$	$S\delta^\rho + P_{1-\delta}^{\delta\rho}, T\delta^\rho + P_{1-\delta}^{\delta\rho}$
AD	$T\delta^\rho + P_{1-\delta}^{\delta\rho}, T\delta^\rho + P_{1-\delta}^{\delta\rho}$	$P_{1-\delta}^{\delta\rho}, P_{1-\delta}^{\delta\rho}$

BRT will be  $\rho$  rounds played with certainty with no payoff discounting; the probability that any of these first  $\rho$  rounds is relevant for payments is given by the geometric distribution with parameter  $\delta$ . If the match does not end in the first block of  $\rho$  rounds, then an additional block of  $\rho$  rounds is played, and so on.<sup>9</sup>

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. For a risk-averse agent, a modified condition involving a different critical  $\delta$  would determine subgame perfection.<sup>10</sup> However, Dal Bó and Fréchet (2011) note that this should not have practical relevance given the parameters used in their experiment and 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.

When it comes to the D+C implementation, more assumptions are required for it to be theoretically equivalent to RT. It must be the case that the strategic interaction between any two players in the continuation game starting from round 5 onwards can be summarized by a choice between the Grim and always defect strategies for each agent. Namely, the possible evolution of payoffs from round 5 onwards must be captured by a choice between full cooperation and full defection where miscoordination between the two agents can occur for only one round.<sup>11</sup> This is a potentially restrictive assumption. For example, Dal Bó and Fréchet (2015) provide evidence that strategies beside Grim and AD are often used. On the other hand, a very small set of strategies accounts for the majority of the data. Hence, whether or not this restriction is problematic is unclear.

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.<sup>12</sup> BRT is similar to the strategy method

<sup>9</sup> 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  $\rho \geq 1$ .

<sup>10</sup> Sherstyuk et al. (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.

<sup>11</sup> This assumption is satisfied if subjects are playing strategies which are such that when they play against each other the resulting payoffs are equivalent to the set of payoffs achieved with only Grim and Always defect. For instance, Tit-For-Tat against AD gives the same payoff as Grim against AD.

<sup>12</sup> Brandts and Charness (2000) find no difference between a “hot” and “cold” treatment in two one-shot games.

**Table 1** Stage game in round 1

		(a) Matches 1-12		(b) Matches 13-18		
		C	D	C	D	
C		40, 40	12, 48	C	24, 24	12, 48
D		48, 12	20, 20	D	48, 12	20, 20

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 most implementations of 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 match continues up to a certain round, then only the choices that the subject has made up to that round are relevant for payoffs. In most implementations of the strategy method, when a decision is made for each contingency, these contingencies are mutually exclusive; hence, it is often the case that only one of these choices will be payoff relevant when the uncertainty is resolved. In contrast, with BRT all the choices a subject has made could be used to determine payoffs.

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,  $\delta$  is set to 0.75. In the first part of each session, subjects play 12 matches with the payoff matrix given in Table 1a.

With such a stage game, cooperation can be supported with any discount factor  $\delta$  above 0.29. Moreover, for  $\delta$  greater than  $0.\bar{4}$ , cooperation is risk-dominant in the sense that when focusing only on the strategies always defect and Grim trigger (or Tit-For-Tat) the later risk dominates always defecting. Dal Bó and Fréchet (2011), Blonski et al. (2011) and Fudenberg et al. (2012) [reporting results based on data from Dreber et al. (2008)] find that this criterion correlates with cooperation rates. 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 given in Table 1b. With this stage game,  $\delta$  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  $\delta$  above 0.28.<sup>13</sup> Given the  $\delta$  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, allow us to test whether the comparative static results were the

<sup>13</sup> 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 the same  $\delta$  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.

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 was reversed, 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-matched between matches. Instructions can be found in the online appendix.<sup>14</sup>

The number of periods played with certainty (except in the RT implementation),  $\rho$ , is set to *four*.<sup>15</sup> 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. The stage game in round 4 for the first part of the experiment is given in Table 2a.<sup>16</sup> In D+C there are five rounds; the first four are the same as in D+RT, but the 5th round is given in Table 2b.<sup>17</sup> 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 may 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 Fréchet (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 3 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 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

<sup>14</sup> Available at [http://cess.nyu.edu/frechette/print/Frechette\\_2014b\\_inst.pdf](http://cess.nyu.edu/frechette/print/Frechette_2014b_inst.pdf). The computer interface was implemented using zTree Fischbacher (2007).

<sup>15</sup> As  $\rho$  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.

<sup>16</sup> In part two,  $R$ , the payoff to joint cooperation in round 4, is reduced to 10.1.

<sup>17</sup> In part two, the payoff to  $(G, G)$  is reduced to 30.4.



**Table 2** Stage game in D+C (matches 1–12)

(a) Round 4			(b) Round 5		
	C	D		G	AD
C	16.9, 16.9	5.1, 20.3	G	50.6, 50.6	22.8, 34.2
D	20.3, 5.1	8.4, 8.4	AD	34.2, 22.8	25.3, 25.3

**Table 3** Summary information

	# of subjects	Subjects Per session	Matches Per session	Rounds per match <sup>a</sup>			Subject earnings (\$)		
				Avg	Min	Max	Avg	Min	Max
RT	50	12, 18, 20	18	4.5	1	19	23.9	12.2	32.6
D+RT	48	12, 16, 20	18	7.5	4	22	20.7	14.1	27.2
BRT	42	12, 14, 16	18	6.1	4	20	20.9	13.6	28.7
D+C	52	16, 18, 18	18	5	5	5	20.7	17.3	23.2

<sup>a</sup> Rounds played, though they may not count towards earnings in BRT

(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, for instance, 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 affect 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 (Table 4).

### 3 Results

The first results on aggregate cooperation levels can be found in Table 4.<sup>18</sup> The left panel presents the data focusing only on round one, while the right panel includes all rounds, except for round five 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

<sup>18</sup> 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 Fréchet (2012).

**Table 4** 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

Diff. stands for the difference between part one and two

\*\*\* Significant at the 1 % level (standard errors clustered at session level)

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 one offers a comparison that is easier to interpret. We will analyze the specific behavior within matches later. Table 4 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.

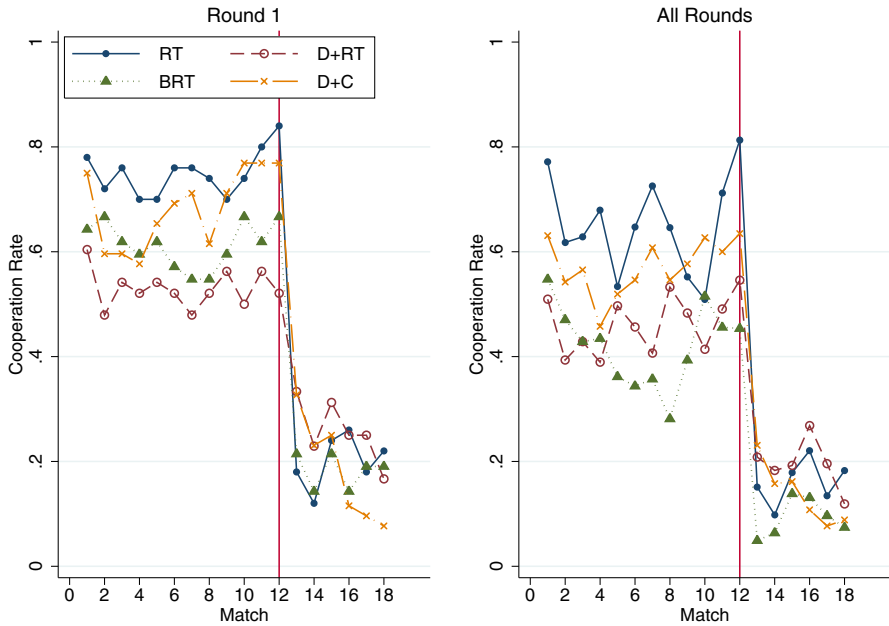
Cooperation rates by matches can be seen in Fig. 1. The figure suggests no clear pattern of changes in cooperation rates over matches. Across treatments, one can see that there are differences in cooperation levels. Furthermore, the changes over matches and across treatments are not necessarily the same for round one and for all rounds.

Looking across treatments, Fig. 1 and the top panel of Table 5 summarize how the cooperation rates compare in the first part of the experiment, where joint cooperation can be supported in equilibrium.<sup>19</sup> When looking at round one, 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 size of the treatment effect—the difference between cooperation rates in parts 1 and 2 - is shown in Table 4. 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 one only or at all rounds).

**Summary of results:** *The comparative static effects are in the same direction for all methods, but there are differences in magnitudes with the standard method (RT) giving the largest treatment effect.*

The results that follow will provide evidence on 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.

<sup>19</sup> The tests in Table 5 include dummy regressors to control for the specific random sequence in a given session.



**Fig. 1** Cooperation rate by match

In the remainder of the paper, we concentrate on the first part of the 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 (Fig. 1) with aggregate cooperation rates dropping below 18 % in all our treatments (Table 4). 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 is observed in the first part is what drives the treatment differences and is critical for understanding the trade-offs associated with the different methods examined.

### 3.1 Matches 1 to 12

To understand the factors that affect the evolution of play over matches, Table 6 reports the marginal effects for the estimates of a probit regression (see Table 7) where cooperation in round one 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);<sup>20</sup> also included are the match number 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

<sup>20</sup> 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.

**Table 5** 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

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

\*\* Significant at the 5 % level (standard errors clustered at session level)

\*\*\* Significant at the 1 % level (standard errors clustered at session level)

round one 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.<sup>21</sup> 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 one of match 1 is included to allow for correlated random effects (reported in Table 8) and is included here for comparability. In the 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 method confirm observations from previous experiments. First, when the opponent in the previous match cooperates, one is more likely to cooperate in the subsequent match.<sup>22</sup> Second, when matches last longer, the subsequent match is more likely to start by cooperation.<sup>23</sup> In the case of D+RT, similar effects are observed; however, as can be seen in Table 6 which reports marginal effects, both channels impact cooperation to a lesser degree. The marginal

<sup>21</sup> By cooperative strategies, in the context of this paper, 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 one are not exogenous of one own's choice and, thus, using only round one avoids issues of endogeneity.

<sup>22</sup> 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.

<sup>23</sup> 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.

**Table 6** Marginal effects (probit) of the factors affecting the evolution of cooperation (matches 2–12)

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

Clustered (session level) standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

effect 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 suggest 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 in cooperation 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 one 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 Fig. 2 shows, the results suggest similar evolutions in that regard across treatments. The figure gives, for each match, the probability of cooperation conditional on the opponent cooperating or defecting in the previous round. In all treatments, over the first few matches, there is a decrease in the probability of cooperation following a defection by the other player.<sup>24</sup> Cooperation rates following a cooperative decision by the other player

<sup>24</sup> In the RT treatment, it starts at 40 % and ends at 18 %. In the other treatments, it starts close to 25 % and decreases to a rate between 7 and 10 %.

**Table 7** Probit estimates of the factors affecting the evolution of cooperation (matches 2–12)

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

Clustered (session level) standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

show an opposite trend, although less pronounced, in all but the D+C treatment.<sup>25</sup> However, the trend for the cooperation rates following cooperation by the 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 % ( $\pm 2$  %) in all treatments in the first match, and it nearly doubles by match 12, reaching 76, 80, 74, and 67 % 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.<sup>26</sup> 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, they represent the last PD game and the choice in the continuation value coordination game.

Figure 3 reveals important differences between the implementation methods. 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 one rate. This is natural in an environment where there is heterogeneity in the subject pool in terms of the strategies used.

<sup>25</sup> The increase in cooperation over the 12 matches (for treatments RT, D+RT, and BRT) is between 7 and 21 % points, depending on the treatments.

<sup>26</sup> 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 Fig. 3 for these treatments.

**Table 8** Random effects probit estimate of the factors affecting the evolution of cooperation (matches 2–12)

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

Clustered (session level) standard errors in parentheses

$\sigma^2$  is the variance of the subject specific random effects.  $\frac{\sigma^2}{\sigma^2+1}$  measures the extent of the individual subject effects and has a minimum value of 0 (no individual subject effects) and a maximum value of 1 (all variance is explained by individual subject effects)

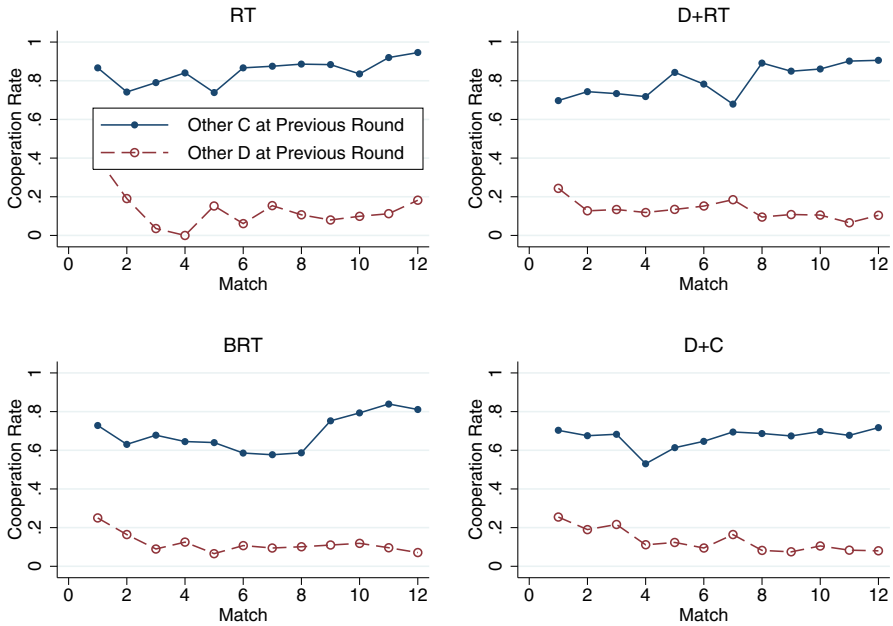
\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

Consider, for example, a population where one half of the population follows Grim while the other half follows Always Defect. If subjects are randomly matched, we would expect cooperation rates to decrease from half to a quarter when we compare the first round to the second round.

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) and suggests different types of strategies to be used in this treatment.<sup>27</sup>

In BRT, cooperation rates in round four are below the round one level, and the round five level is slightly higher. This suggests a small restart effect between blocks, consistent with the results in Table 6. The restart effect between blocks observed in this treatment can potentially pose a problem for analyzing long-term

<sup>27</sup> Our findings in Fig. 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. We postpone further discussion of this to Sect. 3.3.



**Fig. 2** Cooperation rate as a function of the previous choice of the opponent over time

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.<sup>28</sup> In fact, in the last three 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 four than in round one, but the difference is also increasing over time. In addition, it is the only treatment in which cooperative choices in round five are more frequent than in round one. From the graph, we can see a dramatic change in behavior from round four to round five.

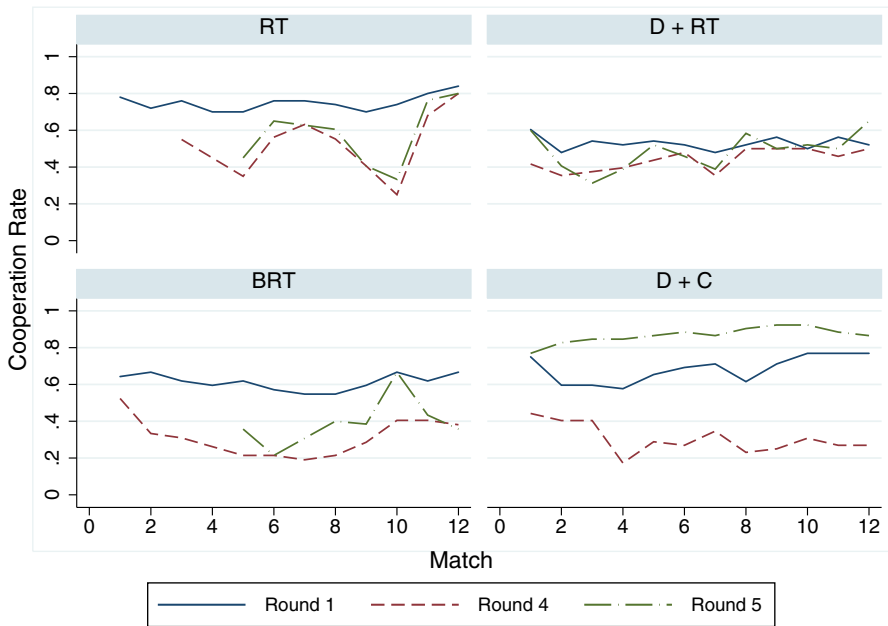
Some of these patterns are better investigated by looking at behavior within a match as shown in Fig. 4. The figure focuses on the second half of the first part of the experiment—namely, matches 7–12.<sup>29</sup> For both RT and BRT, the general pattern seems to be an early decrease in average cooperation followed by a relatively stable period.<sup>30</sup> The D+RT treatment, consistent with earlier observations, presents a smaller decrease in cooperation in the rounds that follow the first

<sup>28</sup> This is done using an ordered probit (cooperation can decrease, stay the same, or increase) clustering by session.

<sup>29</sup> 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.

<sup>30</sup> The drop at the end for RT can be explained by the fact that the sample of matches is changing as we look 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, Fig. 5 presents a similar graph for all matches that lasted at least five rounds, but only looking at the first five rounds. In





**Fig. 3** Cooperation rates in rounds 1, 4, and 5 over time

one. The BRT treatment displays the restart effect mentioned earlier, but if the data is broken down into smaller groups of match (see Fig. 6), 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.

**Summary of results:** *There are variations across methods in how behavior changes with past experiences. D+RT reduces the impact of past experiences and generates the lowest variance in expected payoffs.*

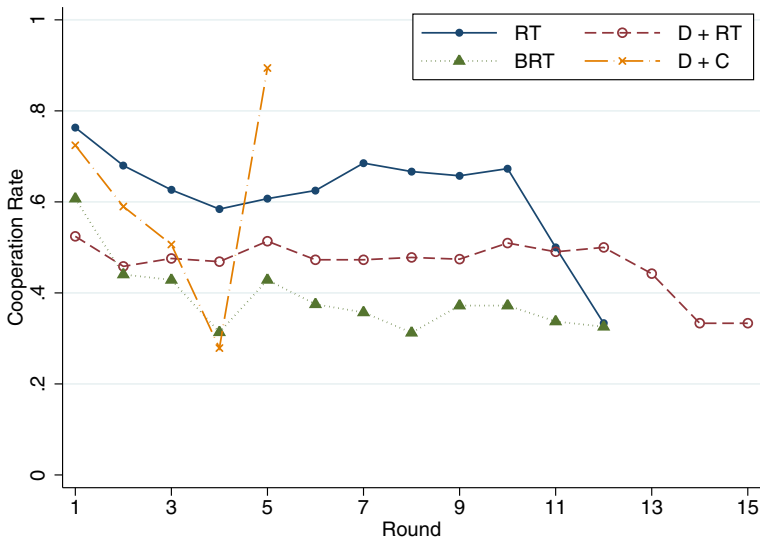
### 3.2 Discounting + coordination

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 9 reports the marginal effects for the estimates of a probit where cooperation in the coordination game (in round five) is regressed on the choice in round one, as well as other controls. Only the choice in round one is included to avoid the endogeneity problems that other rounds would generate. We include subjects' average actions in previous matches to allow strategies in the coordination

Footnote 30 continued

that case, the sample is of the same size for each round of a treatment and as can be seen similar patterns are observed.



**Fig. 4** Cooperation by round, matches 7–12

game to be type-dependent, regardless of the opponent's actions in the previous rounds.<sup>31</sup> 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.<sup>32</sup>

The regression results look significantly different for every block of matches. In particular, in the last block, none of the dummy variables that capture the outcome in round one are significant.<sup>33</sup> 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 one outcome on the cooperation decision in round five. More importantly, the magnitude of the effect is much more pronounced. For instance, when both subjects cooperate in round one, the increase in the probability of cooperation in round five, as opposed to the case where the 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.

<sup>31</sup> We look at actions in 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.

<sup>32</sup> In the case of the last column, including dummy variables for all but one treatment does not qualitatively change the results. Since both dummies are not statistically significant, nor are they jointly different from zero, they are not included. Table 10 reports the probit regression estimates, and Table 11 reports correlated random effects estimates.

<sup>33</sup> Note that for matches 7–9, an F-test shows that the sum of the first three terms is significantly different from 0 ( $p < 0.01$ ). This implies that outcomes of (C, C) and (D, D) in round one generate different levels of cooperation in round five. This suggests that, for a period of time, some subjects employed defection in the last round as a punishment strategy.

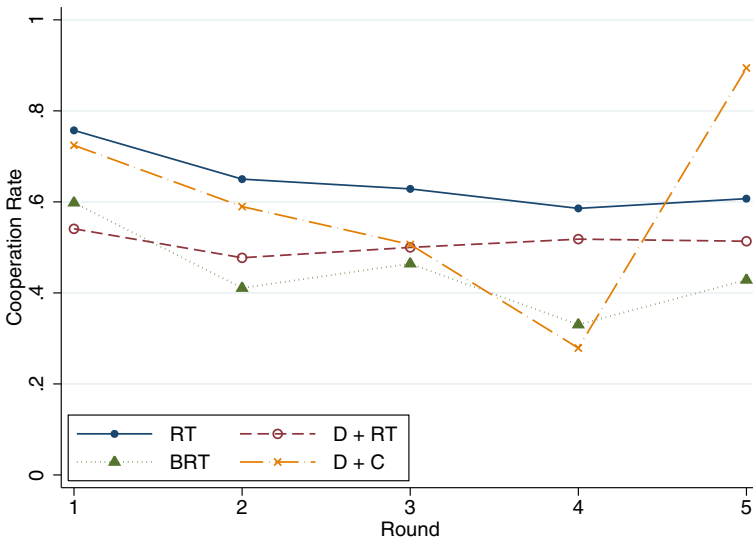


Fig. 5 Cooperation by round, matches 7–12 [only including matches that lasted at least five rounds]

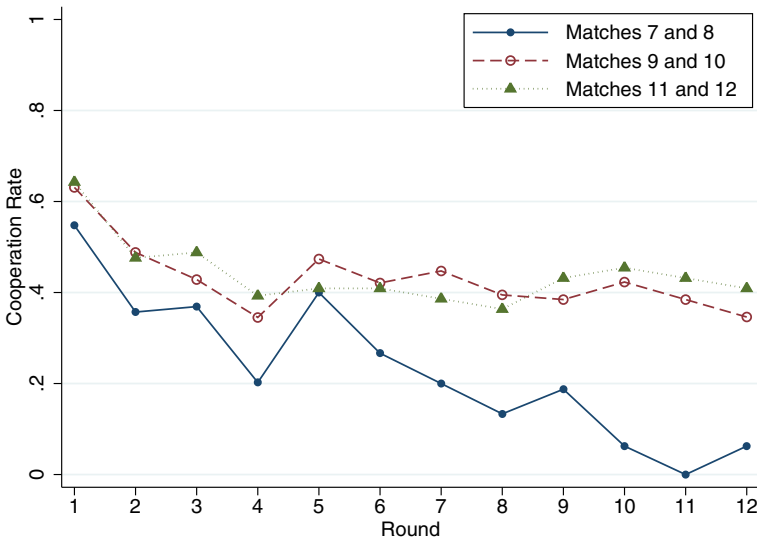


Fig. 6 Cooperation by round in BRT

A less statistical but very telling way of seeing the disconnect between the choice in round five and the choices before that is presented in Fig. 7. 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

five. As can be seen, in D+C, the relation is mostly flat, whereas in all other treatments, there is an important positive relationship.<sup>34</sup>

This is in line with the observation from Fig. 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.<sup>35</sup>

**Summary of results:** *With the D+C method, as subjects gain experience, actions in the coordination game become independent of the evolution of play in the previous rounds. This implies that the coordination game is not used to create dynamic incentives.*

The next section analyzes the strategies that subjects use. Given that behavior in the D+C treatment is substantially different from the other treatments, we focus on the first three methods.

### 3.3 Strategies

This section investigates whether the different methods of implementing infinitely repeated games, although theoretically equivalent, lead to different strategic choices. Our goal is to see if the aggregate differences observed across the different methods can be explained by differences in the strategies used. This can provide insight into the mechanisms driving the behavioral differences across these methods. More broadly, this enables us to study how changing the framing of the environment (from RT to BRT) or the expected interaction length (from RT to D+RT) affects the strategic considerations of the subjects. Understanding how these manipulations affect strategy choice provides insights into behavior that potentially generalize beyond the specifics of this experiment: to other PDs, other social dilemmas, or even to dynamic games more generally.

For this, we employ the strategy estimation procedure introduced in Dal Bó and Fréchet (2011), referred to as strategy frequency estimation method or SFEM, and also used in Fudenberg et al. (2012), Rand et al. (2015), Dal Bó and Fréchet (2015) and Vespa (2015), among others [see Dal Bó and Fréchet (2016) for a summary of findings]. 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.<sup>36</sup>

<sup>34</sup> The figure is almost identical if, instead, the y-axis is computed for round five and all subsequent rounds.

<sup>35</sup> See Embrey et al. (2016) for a review of behavior infinitely repeated PDs.

<sup>36</sup> The interested reader can refer to the online appendix for a more detailed description of the estimation method.

**Table 9** Marginal effects (probit) of the factors affecting cooperation in the coordination game of treatment D+C (round 5)

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

See Table 10

Clustered (session level) standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

**Table 10** Probit estimates of the factors affecting cooperation in the coordination game of treatment D+C (round 5)

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

Clustered (session level) standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

Table 12 reports the estimated frequency of strategies for each of the treatments.<sup>37</sup> In this table, we include only strategies that have a statistically significant positive population share in at least one of our treatments.<sup>38,39</sup>

Results reported in Table 12 stand out, as they indicate that Always Defect, Grim, and versions of Tit-For-Tat account for the most common strategies in all

<sup>37</sup> Standard errors are obtained by bootstrapping. This is done by first drawing sessions and then subjects (both with replacement).

<sup>38</sup> In our estimation, we include the 20 strategies considered in Fudenberg et al. (2012), which cover the commonly considered strategies in repeated prisoner's dilemma experiments. We refer the reader to Table 13 for a description of these strategies. Our estimation results for the entire set of strategies can be found in Table 14.

<sup>39</sup> In Table 14, to ensure that the treatment differences we observe are not driven by differences in the number of observations per match, 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. For example if a match ended in round five 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 this 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 Fréchet (2015) have already pointed out that this method does not perform as well at identifying strategies that are present in small proportions.

of our treatments.<sup>40</sup> 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 unconditional defection to be lowest, at 14 % in RT. In D+RT and BRT, 26 and 25 % of the population always defects. This suggests that subjects use more-cooperative strategies under RT. 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 one, we find it to be 18 % in RT, 46 % in D+RT (different from RT at  $p < 0.1$ ) and 33 % in BRT (not statistically significant from RT).<sup>41,42</sup> This result is in line with our findings in Fig. 4 and Table 5 which report round one cooperation to be highest for RT.

Another feature of our data is the differences in our treatments in terms of the stability of cooperation within a match. In Fig. 4, we see a sharp decline in cooperation from round one to round four with RT. There is a similar decline in BRT. With D+RT, however, the cooperation rate is more stable, with round one's cooperation rate not significantly different from that of round five. This is also seen in Fig. 3. This is a surprising result, as the most-studied and often observed 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. Table 12 reveals that the stability of cooperation rates 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 one, 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 % of the population in D+RT plays strategies that start with defection in round one, but possibly switch back and settle on cooperation, depending on the partner's response.<sup>43</sup> The corresponding share is only 4 % in RT and 8 % in BRT. Additionally, we see that while 32 % of the population is playing Grim-trigger strategies in RT and 21 % in BRT, this share is only 10 % 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. On the other hand, the increase in Always Defect is almost as large as for D+RT. Finally, for Suspicious Tit-For-Tat we see only a very

<sup>40</sup> This is in line with prior research indicating that in perfect monitoring environments, these three strategies can account for the majority of the data. As shown in Fudenberg et al. (2012), when moving to imperfect public monitoring, strategies become more lenient and more forgiving.

<sup>41</sup> 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.

<sup>42</sup> These differences are much less pronounced if we include "suspicious" strategies amongst cooperative strategies. In that case we find 16, 26, and 25 % of defective but not suspicious strategies for RT, D+RT, and BRT respectively.

<sup>43</sup> 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.

**Table 11** Random effects probit estimate of the factors affecting cooperation in the coordination game of treatment D+C (round 5)

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

See Table 10

Clustered (session level) standard errors in parentheses

$\sigma^2$  is the variance of the subject specific random effects

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

modest increase. In contrast to D+RT, these difference are not statistically significant in BRT.

Hence, comparison of the methods in terms of the strategies used provides the most striking contrast when going from RT to D+RT. Since these two treatments are theoretically isomorphic, this amounts to increasing the average number of interactions while keeping the value of cooperation versus defection constant.<sup>44</sup> To summarize our findings, we find this change to generate three key patterns of behavior: (1) the frequency of AD (Always Defect) increases; (2) the popularity of Grim is diminished; and (3) suspicious Tit-For-Tat becomes more popular. These changes are extremely surprising, as it goes against the intuition and folk wisdom about the impact of longer interactions on cooperation in repeated PDs.

In a standard experiment using RT, if  $\delta$  increases, the average number of interactions increases, but the relative value of cooperation also increases.<sup>45</sup> The results presented in this section suggest that increasing  $\delta$  may affect strategy choice via multiple channels, one of which is that it mechanically changes the average

<sup>44</sup> The value of cooperation is often captured by the size of the basin of attraction of AD. Intuitively, if a subject is uncertain about whether he is playing against a subject following the Grim or AD strategy, the set of prior beliefs under which it is optimal to follow the Grim strategy relative to AD increases with  $\delta$ .

<sup>45</sup> Embrey et al. (2016) provides a meta-analysis of experimental research on the finitely repeated PD and analyzes how the interaction length affects cooperations rates.



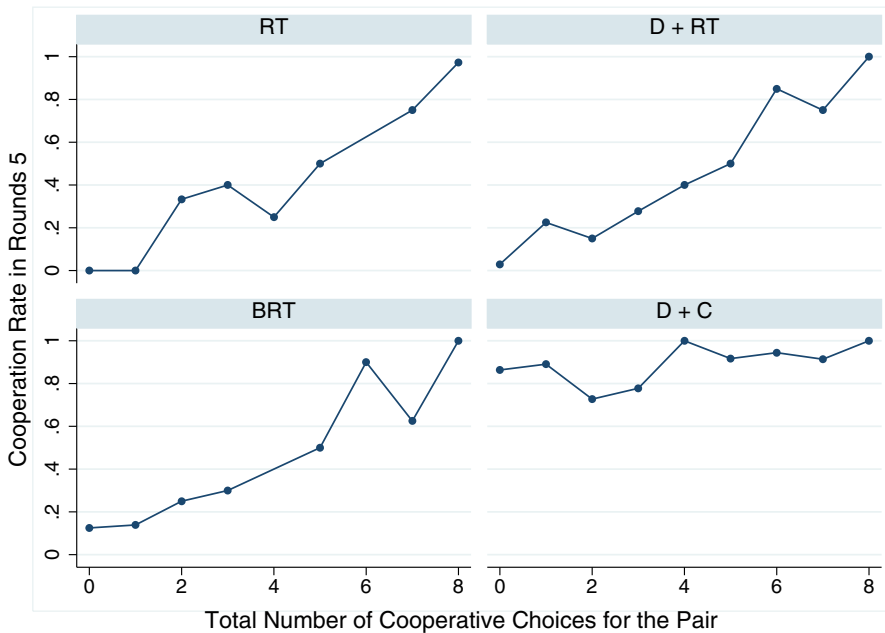


Fig. 7 Cooperation in round 5, matches 7–12

Table 12 Distribution of estimated strategies statistically significant in at least one treatment

	RT	D+RT	BRT
Always defect	0.14 (0.098)	0.26 (0.107)**	0.25 (0.072)***
Grim	0.32 (0.098)***	0.10 (0.061)	0.21 (0.077)***
Tit-For-Tat	0.39 (0.118)***	0.22 (0.095)**	0.33 (0.089)***
2 Tits-For-1 Tat	0.00 (0.055)	0.09 (0.049)*	0.02 (0.068)
2 Tits-For-2 Tats	0.06 (0.044)	0.06 (0.021)***	0.07 (0.043)*
Suspicious Tit-For-Tat	0.02 (0.061)	0.18 (0.057)***	0.05 (0.036)
$\beta$	0.935	0.936	0.901

Bootstrapped standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

number of interactions under most experimental designs. Dal Bó and Fréchette (2015), who directly elicit strategies, find that the impact of increasing  $\delta$  holding stage game parameters constant are for the most part reversed as compared to the three aforementioned ones (and summarized in the previous paragraph). 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 how the discount factor affects strategic choices.

**Table 13** 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 if 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 cooperater	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

**Summary of results:** *Estimated strategies are significantly different in  $D+RT$  relative to RT. This suggests that increasing the interaction length, in the absence of any changes in the discount rate, can affect equilibrium selection.*

## 4 Discussion

We have studied behavior under four different implementations of infinitely repeated games in the laboratory: the standard random termination method (RT) and three other methods that de-couple the expected number of rounds and the discount factor ( $D+RT$ , BRT and  $D+C$ ). Our results can be summarized as follows.

Most importantly, we find that all four methods generate sharp comparative statics in the same direction: cooperation levels drop significantly when parameters of the stage game are changed to make mutual cooperation theoretically unsustainable.

**Table 14** Distribution of estimated strategies

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

Bootstrapped standard errors in parentheses

\*  $p < 0.10$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

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 consider the rest of the game as a finitely repeated game. Cooper and Kühn (2014a, b) show how dynamic incentives can be recovered with this method when subjects are allowed to engage in pre-play communication.<sup>46</sup> Overall, the D+C method induces a markedly different environment that has distinct advantages and disadvantages. It reduces the number of exchanges per match and simplifies the strategy space, but our results suggest that it should not be used without communication when studying infinitely repeated games.

In contrast, all three methods using random termination generate behavior that is consistent with theory. The preferred implementation method among these will depend on the research question at hand and the type of applications that are meant to be modeled in the laboratory. From the perspective of testing the implications of infinitely repeated games in the laboratory, we find that the different methods present various trade-offs.<sup>47</sup>

- The RT method generates the largest comparative statics with respect to parameters of the stage game. It generates fewer rounds per match than D+RT and BRT, hence allows for more matches to be played in one session. This can be important in settings where learning is slow. In addition, it is the most commonly used method making comparisons to other studies more direct. Hence, it provides the default design choice unless the research question presents good reasons to depart from it.
- The D+RT method provides the following benefits relative to RT. It allows the experimenter to vary the average number of interactions independently of the discount factor; and it reduces the impact of past experiences. Hence, this method can be desirable in two ways. First, when longer matches need to be observed; second, when the experimenter wants to reduce variance across sessions. Moreover, if the experiment is meant to represent an environment where payoff discounting is more germane than random termination, using

<sup>46</sup> Cooper and Kühn (2014a) find that free form communication generates cooperation primarily by allowing subjects to communicate explicit punishment threats in response to noncooperative actions. Consistent with this observation, they find that cooperation cannot be sustained when communication is structured and the limited message space does not allow for communication of contingent strategies. Cooper and Kühn (2014b) revisit the impact and nature of communication in their D+C design. They also manipulate the number of rounds played before the coordination game: one or two. As in Cooper and Kühn (2014a), their findings indicate that with free form communication cooperation emerges. Their results are consistent with ours in that, without communication, despite the fact that by the end of the experiment very few subjects in round one choose the high mutual payoff action, the vast majority of subjects do not choose the lowest payoff equilibrium in the coordination game. While lack of variation in round 1 play makes it difficult to look at contingent play in the coordination game, the fact that less than 5 % of subjects play the low payoff action in the coordination game suggests a limited link between choices in the coordination game and history of play in the first round. The same is not true, however, when they allow for free form communication. In that case, there is an important correlation between round one and round two choices, suggesting that subjects create and use dynamic incentives in these treatments.

<sup>47</sup> We are wary of a *cookbook* approach to design and thus provide these with caution. Specific questions each have their own requirements that are difficult to anticipate in the abstract. Thus, these recommendations should be taken with a grain of salt.

D+RT could be important given the observation that this treatment leads subjects to use slightly different strategies.<sup>48</sup>

- The BRT method provides additional benefits of its own. It allows the experimenter to vary the expected number of observations from a match independently of the discount factor or stage game (essentially without changing the expected length of the match). It also reduces the impact of past experiences, but not as much as D+RT. It yields strategic choices closer to RT than D+RT. Hence, it can be a desirable choice when the experimenter needs to observe longer interactions but wants to remain as close as possible to RT.<sup>49</sup>

Hence, the main result of this study is that all four implementations “work” in that they generate comparative static results in the same direction, a direction that is consistent with the theory of infinitely repeated games. However, from the subject’s behavior we can find that only three of these implementations actually generate dynamic incentives. Thus, if the purpose of an experiment is to use an implementation that speaks to the theory of infinitely repeated games, then the D+C method should be avoided.

A more detailed analysis of the subjects’ behavior does reveal some differences in behavior between implementations. These 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 suggest that situations in which agents are very patient, but relationships are likely to terminate for exogenous reasons, may lead to slightly 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.

**Acknowledgments** We would like to thank two anonymous referees, 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 the NSF via Grants SES-0924780, SES-1225779, and SES-1558857.

<sup>48</sup> Of these three methods, this one generates the lowest variance in expected payments. If controlling the budget is a priority, this implementation would be preferable. If this is a concern, however, the method that gives the most control on expected payment is the one presented in Sherstyuk et al. (2013) of using RT and paying only the last round. That method could be adapted to be used in combination with the BRT design.

<sup>49</sup> It should be noted that this method generates a re-start effect (between blocks) that disappears with experience. For this reason, this method is better suited to experimental designs where many matches will be played.

## References

- Andersson, O., & Wengström, E. (2012). Credible communication and cooperation: Experimental evidence from multi-stage games. *Journal of Economic Behavior & Organization*, *81*(1), 207–219.
- Blonski, M., Ockenfels, P., & Spagnolo, G. (2011). Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *The American Economic Journal: Microeconomics*, *3*(3), 164–192.
- Brandts, J., & Charness, G. (2000). Hot versus cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, *2*(3), 227–238.
- Cabral, L. M. B., Ozbay, E., & Schotter, A. (2014). Intrinsic and instrumental reciprocity: An experimental study. *Games and Economic Behavior*, *87*, 100–121.
- Cooper, D. J., & Kühn, K.-U. (2014a). Communication, renegotiation, and the scope for collusion. *American Economic Journal: Microeconomics*, *6*(2), 247–278.
- Cooper, D.J., & Kühn, K.U. (2014b). Cruel to be kind: An experimental study of communication, collusion, and threats to punish. Working paper.
- Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *The American Economic Review*, *95*, 1591–1604.
- Dal Bó, P., & Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review*, *101*(1), 411–429.
- Dal Bó, P., & Fréchette, G. R. (2015). Strategy choice in the infinitely repeated prisoners dilemma. Working paper.
- Dal Bó, P., & Fréchette, G. R. (2016). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, forthcoming.
- Dreber, A., Rand, D. G., Fudenberg, D., & Nowak, M. A. (2008). Winners don't punish. *Nature*, *452*, 348–351.
- Embrey, M.S., Fréchette, G. R., & Yuksel, S. (2016). Cooperation in the finitely repeated prisoner's dilemma. Working paper.
- Engle-Warnick, J., & Slonim, R. L. (2006). Inferring repeated-game strategies from actions: Evidence from repeated trust game experiments. *Economic Theory*, *28*(3), 603–632.
- Feinberg, R., & Husted, T. A. (1993). An experimental test of discount-rate effects on collusive behaviour in duopoly markets. *The Journal of Industrial Economics*, *41*, 153–160.
- Feinberg, R., & Snyder, C. (2002). Collusion with secret price cuts: An experimental investigation. *Economics Bulletin*, *3*(6), 1–11.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for readymade economic experiments. *Experimental Economics*, *10*(2), 171–178.
- Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics*, *15*(3), 485–498.
- Fudenberg, D., Rand, D. G., & Dreber, A. (2012). Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world. *The American Economic Review*, *102*, 720–749.
- Mailath, G. J., & Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. New York: Oxford University Press.
- Rand, D. G., Fudenberg, D., & Dreber, A. (2015). It is the thought that counts: The role of intentions in noisy repeated games. *Journal of Economic Behavior & Organization*, *116*, 481–499.
- Roth, A. E., & Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, *17*, 189–198.
- Sabater-Grande, G., & Georgantzis, N. (2002). Accounting for risk aversion in repeated prisoners' dilemma games: An experimental test. *Journal of Economic Behavior & Organization*, *48*(1), 37–50.
- Schley, D. R., & Kagel, J. H. (2013). How economic rewards affect cooperation reconsidered. *Economics Letters*, *121*(1), 124–127.
- Sherstyuk, K., Tarui, N., & Saijo, T. (2013). Payment schemes in infinite-horizon experimental games. *Experimental Economics*, *16*, 125–153.
- Vespa, E.I. (2015). An experimental investigation of strategies in dynamic games. Working paper.
- Wilson, A., & Wu, H. (2014). At-will Relationships: How an Option to Walk Away Affects Cooperation and Efficiency. Working Paper.
- Zwick, R., Rapoport, A., & Howard, J. C. (1992). Two-person sequential bargaining behavior with exogenous breakdown. *Theory and Decision*, *32*, 241–268.