On the Determinants of Cooperation in Infinitely Repeated Games: A Survey

Pedro Dal Bó
Brown University

Guillaume R. Fréchette*
New York University

March 3, 2016

Abstract

A growing experimental literature studies the determinants of cooperation in infinitely repeated games, tests different predictions of the theory, and suggests an empirical solution to the problem of multiple equilibria. To provide a robust description of the literature's findings, we gather and analyze a metadata set of experiments on infinitely repeated prisoners' dilemma games. The experimental data show that cooperation is affected by infinite repetition and is more likely to arise when it can be supported in equilibrium. However, the fact that cooperation can be supported in equilibrium does not imply that most subjects will cooperate. High cooperation rates will emerge only when the parameters of the repeated game are such that cooperation is very robust to strategic uncertainty. We also review the results regarding the effect of imperfect monitoring, changing partners and personal characteristics on cooperation and the strategies used to support it. *JEL codes: C7, C9.*

*We thank all the authors who made the data we use in this article available, Yunan Ji for her research assistance and Anna Aizer, Masaki Aoyagi, V. Bhaskar, Matthew Embrey, Drew Fudenberg, Ennio Stacchetti, the referees and the editors Steven Durlauf and Janet Currie for useful comments.

1 Introduction

The tension between opportunistic behavior and cooperation is a central feature of human interactions. Understanding whether and how people can overcome the incentives for opportunistic behavior and cooperate is important to economics and other sciences. One of the contributions of game theory to the study of this tension is to recognize the fact that repeated interaction enables credible punishments and rewards that can lead to cooperation.

There have been major advances in the theory of infinitely repeated games, with studies on the determinants of cooperation in increasingly complex environments, from perfect monitoring to imperfect monitoring (both public and private) and others. The applications are many, in economics, other social sciences and the natural sciences. However, the standard theory of infinitely repeated games is sometimes criticized for not providing sharp predictions (when players are sufficiently patient, both cooperation and defection are possible equilibrium actions).[1] Hence, empirical investigations are important, not only to test theoretical predictions, but also to help us sharpen predictions when theory provides multiple ones. Empirical studies of repeated games have been slower to emerge, as many of the important parameters are difficult to observe, such as the discount factor, the probability that an interaction will continue, or the information that agents have about the others' choices. These make experimental investigation a particularly useful method for testing, exploring, and refining these models. Systematically studying how agents behave in such environments helps us learn about the determinants of cooperation.

After a period of initial investigation from the 1970s to 2000, which included only a handful of studies, the last decade has seen a significant increase in the number of laboratory investigations into infinitely repeated interactions: games in which subjects interact for a number of periods without a known termination period, and the experimenter controls the discount factor. This is most typically induced using a random termination period, with all subjects

[1] Tirole [1988] says that "[t]he multiplicity of equilibria is an embarrassment of riches" and Fudenberg and Maskin [1993] state that "[t]he theory of repeated games has been somewhat disappointing. [...] the theory does not make sharp predictions."

knowing the probability with which the game will continue for an additional period–a method first used by Roth and Murnighan [1978].[2]

We review the experimental literature on the determinants of cooperation with a special focus on the prisoner's dilemma. The prisoners' dilemma captures the tension at the heart of social dilemmas in its simplest form, and, as such, it is a useful tool to explore general questions about cooperation. Furthermore, we go beyond a simple review and use metadata assembled from previous experimental studies to revisit the key findings of this literature. For this purpose, we need a common game, and the prisoners' dilemma is the one that has been most studied with random termination in the laboratory.

The use of metadata to revisit the main findings of the literature is useful since the design of each experiment is usually specifically tailored to highlight a certain comparative static or element. Although the knowledge we accumulate this way is often precise, we rarely have a sense of its robustness. How dependent, for instance, is a specific result on the constellation of parameters used in an experiment? On the subject pool? On the details of the procedures and the formulation of the instructions? The hope is that such questions can be resolved over time through replication with different payoffs, instructions, subjects, and procedures used by other researchers. However, it is often difficult to get a good sense of the robust results that come out of a sequence of experiments. With this in mind, we use this metadata to re-examine some of the questions that we and others have studied in specific settings.[3]

We start, in Section 2, by considering environments with perfect monitoring and discuss how the parameters of the game affect cooperation, the evolution of cooperation as subjects gain experience and the choice of strategies. This

[2]Some authors prefer to talk about indefinitely, as opposed to infinitely, repeated games given the method used to implement them. We view the term infinitely repeated games as a reference to the theoretical framework used to analyze the situation of interest and not so much as a description of the implementation. Accordingly, we use the following terminologies interchangeably: randomly terminated games, indefinite games, and infinitely repeated games.

[3]Meta-studies typically gather the results of multiple articles. Although we do that in some subsections, much of the meta-analysis is done by re-analyzing the raw data from multiple studies that we obtained and pooled together (as well as standardized).

is followed by Section 3 on the methods used to induced infinitely repeated games in the laboratory and potential implementation issues. In Section 4, we discuss the impact of personal characteristics on cooperation and the motives behind cooperation. We then study cooperation under imperfect monitoring in Section 5. In Section 6, we present evidence on community enforcement when partners may change from round to round. In section 7, we describe work using stage games other than the prisoners' dilemma and more general dynamic games. The last section explores some promising directions for future research.

2 The Infinitely Repeated Prisoners' Dilemma Game With Perfect Monitoring

In this section, we address the main questions raised by the literature by focusing on the simplest type of infinitely repeated game: infinitely repeated prisoners' dilemma (henceforth PD) games under perfect monitoring and fixed matching. Treatments in these papers differ mainly in the probability of continuation (δ) and the payoff matrix.[4]

A PD payoff matrix consists of four numbers: the temptation payoff (T) which is the payoff for defecting when the other cooperates; the reward from joint cooperation (R): the punishment from mutual defection (P); and the sucker's payoff from cooperation when the other defects (S). For these payoffs to define a PD, it must be that $T > R > P > S$. Furthermore, it is often required that $2R > T + S$ so that alternating between cooperation and defection cannot be more profitable than joint cooperation.

To facilitate comparison across treatments, we work with the normalized matrix in Table 1, which reduces the number of payoff parameters to two: g is the gain from defection when the other player cooperates, and ℓ is the loss from cooperation when the other player defects. This matrix is obtained by applying a monotonic linear transformation to the original matrix, which,

[4]Other differences have to do with factors outside of the standard theory, such as the number of supergames in a session, instructions, and feedback between supergames.

Table 1: Prisoners' Dilemma Row Player's Payoffs

Original

	C	D
C	R	S
D	T	P

Normalized

	C	D
C	$\frac{R-P}{R-P}=1$	$\frac{S-P}{R-P}=-\ell$
D	$\frac{T-P}{R-P}=1+g$	$\frac{P-P}{R-P}=0$

Table 2: Percentage of Cooperation (First Rounds)

	Probability of Continuation		
	0.105	0.5	0.895
Roth and Murnighan (1978)	19	29.75	36.36
Murnighan and Roth (1983)	17.83	37.48	29.07

under general utility assumptions, should not affect behavior.[5]

2.1 Does Repetition Matter?

Roth and Murnighan [1978] and Murnighan and Roth [1983] were the first to experimentally address the issue of whether infinite repetition affects cooperation. They introduced the random termination rule and had subjects play for money against an artificial opponent (a pre-programmed strategy) for one repeated game per treatment.

As Table 2 shows, cooperation tends to increase with the probability of continuation but the effect is not always monotonic in those papers. Moreover, even under high probability of continuation, the subjects are far from making the most of the opportunity to cooperate. This leads Roth [1995] to conclude that "the results remain equivocal."

Feinberg and Husted [1993] combine a random continuation rule (following a fixed number of rounds played for certain) with discounting of payoffs (in every period) in a repeated prisoners' dilemma game. They find that less discounting results in an increase in cooperative outcomes. However, the effect of repetition was not large.

[5]We discuss the empirical validity of this assumption in footnote 10.

Finally, Palfrey and Rosenthal [1994] study contributions to a public good under incomplete information regarding the cost of contribution. They find that contribution rates are greater under a probability 0.9 of continuation than in one-shot games. However, the magnitude of the increase was not large, leading them to conclude that the impact of repetition on cooperation is limited.

A second wave of experimental papers, starting with Dal Bó [2005], provides a more positive answer to the question of whether *the shadow of the future* can increase cooperation and by how much. For example, Dal Bó [2005] finds that when going from one-shot PD games to infinitely repeated ones with $\delta = .75$, cooperation increases fourfold. The main reason for the difference in results is that this second wave of experimental papers allows subjects to participate in several repeated games and gain experience.

We use the data generated by this second wave of experiments to describe the evidence on the effect of repetition on cooperation and the other main results in the literature.[6]

Thus, the results presented here are robust, in the sense that they are not based on one particular paper, but on the data from several experiments conducted by different experimenters in different places with different procedures, instructions and subject pools.

Table 3 displays the variety of treatments used in the second wave of experiments for which we have data (papers are organized by date of publication; then treatments are ordered by their parameters δ, g and ℓ). In this section, we focus on randomly terminated PD experiments with perfect monitoring from 15 papers. None of the treatments combine random termination with other forms of discounting (nor do they vary discounting within a supergame). We also include three articles with one-shot PD games: $\delta = 0$.

[6] We include all the papers in Economics that we are aware of that satisfy the following conditions: i) the stage game is a fixed 2x2 prisoners' dilemma game; ii) there is perfect monitoring; iii) one-shot interaction or repeated interaction through a random continuation rule (and this does not change inside a supergame); iv) pairs are fixed inside a supergame; v) the article was published before 2014 or is one of our own; vi) we have access to the data. We use Internet searches to find articles satisfying these conditions. We also rely on Mengel [2015] to ensure that we include all the relevant one-shot prisoners' dilemma experiments.

Table 3: General Information About Data Sets Used

	Sessions	Subjects	δ	g	ℓ	Supergames
Andreoni and Miller 1993	**1**	**14**	0	1.67	1.33	200
Cooper, DeJong, Forsythe, and Ross 1996	**3**	**33**	0	0.44	0.78	10
Dal Bó 2005	**6**	**276**				
	2	72	0	1.17	0.83	7
	2	102	0	0.83	1.17	9
	1	42	0.75	1.17	0.83	7
	1	60	0.75	0.83	1.17	10
Dreber, Rand, Fudenberg and Nowak 2008	**2**	**50**				
	1	28	0.75	2	2	21
	1	22	0.75	1	1	27
Aoyagi and Fréchette 2009	**4**	**74**				
	2	36	0	0.33	0.11	75
	2	38	0.9	0.33	0.11	10
Duffy and Ochs 2009	**9**	**102**	0.9	1	1	13
Dal Bó, Foster, and Putterman 2010	**28**	**424**	0	1	3	10
Dal Bó and Fréchette 2011	**18**	**266**				
	3	44	0.5	2.57	1.86	71
	3	50	0.5	0.67	0.87	72
	3	46	0.5	0.09	0.57	77
	3	44	0.75	2.57	1.86	33
	3	38	0.75	0.67	0.87	47
	3	44	0.75	0.09	0.57	35
Blonski, Ockenfels, and Spagnolo 2011	**10**	**200**				
	1	20	0.5	2	2	11
	2	40	0.75	2	2	11
	1	20	0.75	1	8	11
	1	20	0.75	1	1	11
	1	20	0.75	0.83	0.5	11
	1	20	0.75	0.75	1.25	11
	1	20	0.75	0.5	3.5	11
	1	20	0.875	2	2	8
	1	20	0.875	0.5	3.5	8
Fudenberg, Rand, and Dreber 2012	**3**	**48**	0.875	0.33	0.33	9
Bruttel and Kamecke 2012	**3**	**36**	0.8	1.17	0.83	20
Sherstyuk, Tari, and Saijo 2013	**4**	**56**	0.75	1	0.25	29
Kagel and Schley 2013	**6**	**114**	0.75	1	0.5	39
Fréchette and Yuksel 2014	**3**	**60**	0.75	0.4	0.4	12
Dal Bó and Fréchette 2015	**41**	**672**				
	3	50	0.5	2.57	1.86	37
	8	140	0.5	0.09	0.57	46
	8	114	0.75	2.57	1.86	25
	10	164	0.75	0.09	0.57	24
	10	168	0.9	2.57	1.86	21
	2	36	0.95	2.57	1.86	7
Total	**141**	**2415**		**Choices:** 157,170		

The information in Table 3 is only for the treatments we use in this section; the papers in which these were originally used might have had many more sessions or treatments, but those did not fit our requirements (deterministic PD with perfect monitoring, fixed and known continuation probability, played in a fixed pair). Also, some of the studies used within-subjects designs; in those, cases we only include the parameters before the first change in treatment. This is to avoid any "contamination" across treatments that we would not account for.

We have data from 15 papers, involving 32 different treatments (combinations of δ, g and ℓ), with 2415 subjects and a total of 157,170 choices.[7] An important difference from the first wave of experiments is that the articles in Table 3 have subjects participating in several supergames under the same combination of parameters, thus allowing them to gain experience. Note that the treatments with the smallest number of supergames have seven supergames.

Remember that the first wave of experiments on infinitely repeated games involved playing a single supergame and only found a modest relation between the probability of continuation and cooperation. We perform a similar exercise using data from the first supergame of each study in Table 3.

We focus on round 1 choices for most of the analysis.[8] While behavior in round 1 provides a somewhat incomplete picture of choices in a supergame, it simplifies the analysis for two reasons. First, supergames will have different numbers of rounds due to the random termination rule, which makes comparisons that use all rounds difficult. Second, behavior in higher rounds is not independent of previous rounds, thus requiring careful consideration of this dependence. Later in the paper, we will analyze choices in the supergame as a whole.

The first two columns in Table 4 present the results of a probit of coop-

[7]Some may worry that 40 percent of the papers considered in Table 3 have one or both of us as a co-author. Relatedly, 50 percent of treatments, 73 percent of subjects and 67 percent of all choices come from an article in which one of us is a co-author. The main results of this section are robust to focusing on data from other authors, as can be seen in the online appendix.

[8]The term round is used to denote a play of the stage game. It is also sometimes referred to as a period. A play of a supergame is often called a match or cycle.

Table 4: The Effect of Repetition on Round 1 Cooperation

	Supergame 1		Supergame 7		Supergame 15	
	Probit	Marginal	Probit	Marginal	Probit	Marginal
	(1)	(2)	(3)	(4)	(5)	(6)
g	-0.202***	-0.0801***	-0.332***	-0.122***	-0.391***	-0.156***
	(0.0377)	(0.0150)	(0.0799)	(0.0311)	(0.117)	(0.0464)
ℓ	-0.142***	-0.0562***	-0.229**	-0.0843**	-0.361*	-0.144*
	(0.0387)	(0.0154)	(0.0961)	(0.0345)	(0.187)	(0.0747)
δ	0.439***	0.174***	1.296***	0.476***	2.186***	0.869***
	(0.135)	(0.0534)	(0.273)	(0.0970)	(0.734)	(0.289)
Constant	0.0895		-0.385*		-0.710	
	(0.118)		(0.210)		(0.496)	
N	2415		2305		1030	

Clustered Standard errors in parentheses

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

eration in round 1 of the first supergame on normalized payoffs and discount factor.[9] Cooperation in the first round of the first supergame is decreasing in both g and ℓ and increasing in δ. However, an increase in the probability of continuation δ has only a small effect, consistent with the evidence from the first wave of experiments.[10]

The second two columns in Table 4 present the results of a probit of cooperation in round 1 of supergame 7 on normalized payoffs and discount factor. We choose the seventh supergame, as this is the highest supergame we can study without losing any treatment (combination of δ, g and ℓ). The last two columns in Table 4 present the results of a probit of cooperation in round 1 of supergame 15 on normalized payoffs and discount factor. We show results from this supergame to confirm that the effect of experience on behavior is not limited to supergame 7. The last four columns in Table 4 show that, after subjects have gained experience playing supergames, the effect of the probability of continuation is much larger than in the first supergame.

Using the results in Table 4, we compute the predicted probability of cooperation for the parameters used by Roth and Murnighan [1978].[11] Table 5 repeats the main result of Roth and Murnighan [1978] and includes the pre-

[9]Standard errors are clustered at the paper level to account for possible effects of instructions, procedures, or software. An additional source of correlation could be session-effects. Given that clustering by paper allows for arbitrary correlations within paper, it also allows for correlations within sessions. For a discussion of session-effects, see Fréchette [2011]. In cases in which there are multiple observations per subject, we will allow for random-effects at the subject level. If, in addition, there are lagged regressors, we will then use a correlated random-effects specification to deal with the initial conditions problem (see Heckman [1981] or Chamberlain [1982] for the static case and Wooldridge [2002] for a clear exposition of the initial conditions problem and methods to address it).

[10]Remember that the normalization of payoffs presented in Table 1 depends on behavior not being affected by linear monotonic transformation of payoffs. To assess the validity of this assumption, we use an analysis similar to the one presented in Table 4, in which we include as controls the multiplicative and additive factors used for the transformation in each treatment. We find that these factors are not statistically significant in the supergames we focus on: 1, 7 and 15. This suggests that the representation of the PD games with only two parameters, g and ℓ, is appropriate. Consistent with this finding, Kagel and Schley [2013] find that cooperation is not affected by expressing payoffs in cents instead of dollars.

[11]The normalized game has $g = \ell = 0.5$.

dicted probabilities from the estimation performed on the data from the new wave of experiments. While the predicted levels based on the new estimates are far from the original cooperation rates in round 1 of the first supergame, the qualitative effect of δ on cooperation is very similar: increasing the continuation probability leads only to a small increase in cooperation.[12] The original experiment reports that increasing δ from 0.105 to 0.895 leads to about a 17-percentage-point increase in cooperation. The new estimates predict about a 13-percentage-point increase in the first supergame. Hence, in both cases, there impact is rather limited.

Table 5: Cooperation in Old and New Experiments (Round 1)

	Continuation Probability		
	0.105	0.5	0.895
Roth and Murnighan	19.00	29.75	36.36
New Estimates (fitted):			
Supergame 1	48.55	55.45	62.19
Supergame 7	29.82	49.30	68.94
Supergame 15	19.57	50.26	80.78

Given the analysis based on the new wave of experiments in Table 4, what would Roth and Murnighan [1978] have found if they had allowed subjects to gain experience? The predicted cooperation rates from the last two rows in Table 5 show that they would have found a much greater effect of repetition on cooperation. The predicted levels show that increasing δ from 0.105 to 0.895 leads to about a 39-percentage-point increase in cooperation in supergame 7 and 61percentage points in supergame 15. This suggests that Roth and Murnighan [1978] would have found a greater impact of repetition on cooperation if their subjects had interacted in several repeated games.

[12] The differences in levels between the observed and predicted behavior may be the result of two important differences in the way the original experiment was conducted relative to those in Table 3. First, subjects played against computerized opponents. Second, subjects were not paid proportionally to their performance. Rather, a prize was given to the "best player." Given those differences, it might not be so surprising that the levels are different.

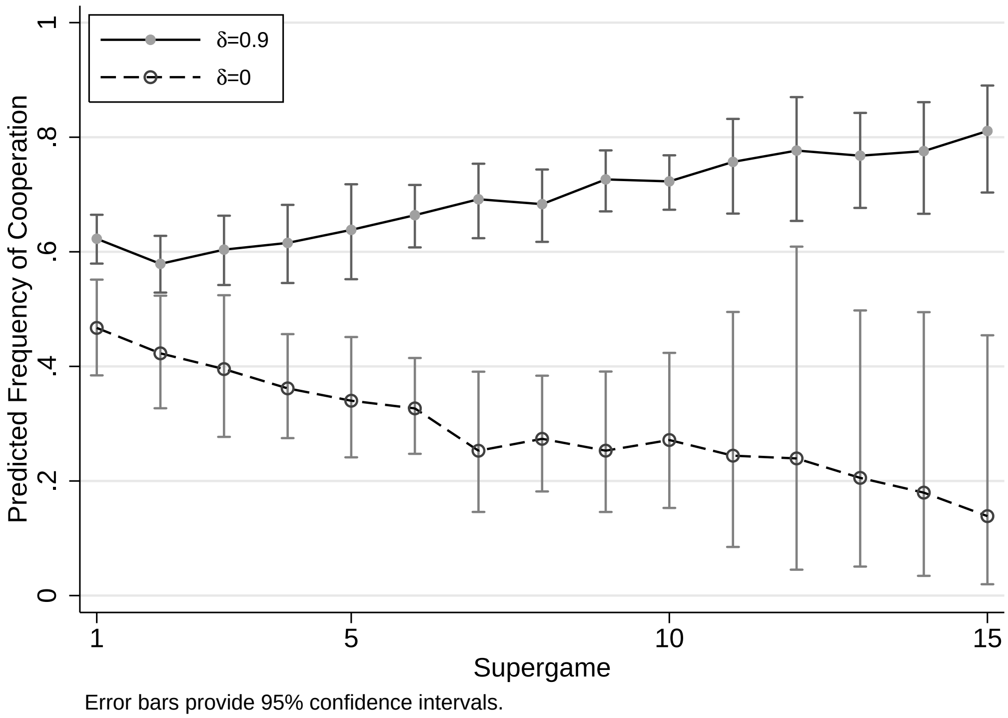

Figure 1: The Impact of δ on Round 1 Cooperation by Supergame:
Predicted Frequency of Cooperation for $g = \ell = \frac{1}{2}$

To further explore how experience affects the impact of repetition on cooperation, we estimate a regression like the ones reported in Table 4 for each supergame. Figure 1 presents the predicted cooperation rates for supergames 1 to 15, again assuming the payoffs of Roth and Murnighan [1978], for a δ of 0 and 0.9, two extreme values for which we have data.[13] By supergame 15, the impact of δ is substantial. In fact, the marginal effect of δ when the payoffs are set at their average values in the sample is 0.87. Hence, having repeated interactions with an uncertain horizon has an important impact on the ability

[13] After supergame 15, the sample size diminishes rapidly. Results beyond supergame 7 should be interpreted with the caveat that the sample starts to change. We note, however, that the upward trend for $\delta = 0.9$ and the downward trend for $\delta = 0$ are present from the very start.

of subjects to support cooperation once they have gained experience.

Two potential concerns with respect to the interpretation of Figure 1 are the focus on round 1 and on cooperation decisions as opposed to cooperative outcomes (joint cooperation). Indeed, in finitely repeated games, round 1 cooperation rates are often much higher than last round cooperation. Some may think that a similar phenomenon (an important decrease in cooperation rates over time) occurs in randomly terminated games. Once one considers rounds beyond the first, the rate of cooperative outcomes is not simply the product of cooperative choices since choices in a pair are no longer independent. Hence, cooperation rates that are relatively high could hypothetically hide much lower cooperative outcomes. Figure 2 presents the predicted rates of joint cooperation in the last round of a supergame, again for the same normalized payoffs ($g = \ell = 0.5$). It shows that the impact of δ on cooperation is not limited to the first round but persists until the end of the supergames.

The analysis presented so far leads us to the following result.

Result 1 *Cooperation is increasing in the probability of future interactions, and this effect increases with experience.*

2.2 The Predictive Power Of Theory

The fact that having randomly terminated repeated interactions affects behavior does not necessary mean that it does as theory predicts. Testing models of infinitely repeated games is complicated by the fact that many outcomes may be consistent with theoretical predictions due to the "folk theorem:" any feasible and individually rational payoff can be supported in a subgame perfect equilibrium if players are sufficiently patient.[14]

[14] On the "folk theorem" under perfect monitoring, see Friedman [1971] and Fudenberg and Maskin [1986]. For a description of the equilibrium payoffs for infinitely repeated PD games for each discount factor, see Stahl [1991]. For a review of the theoretical literature on infinitely repeated games, see Mailath and Samuelson [2006]. Evolutionary stability concepts have been used to reduce the multiplicity of equilibria. For example, Binmore and Samuelson [1992], Fudenberg and Maskin [1990], Fudenberg and Maskin [1993] and Dal Bó and Pujals [2012] find that, under certain conditions, evolutionary stable strategies must be cooperative.

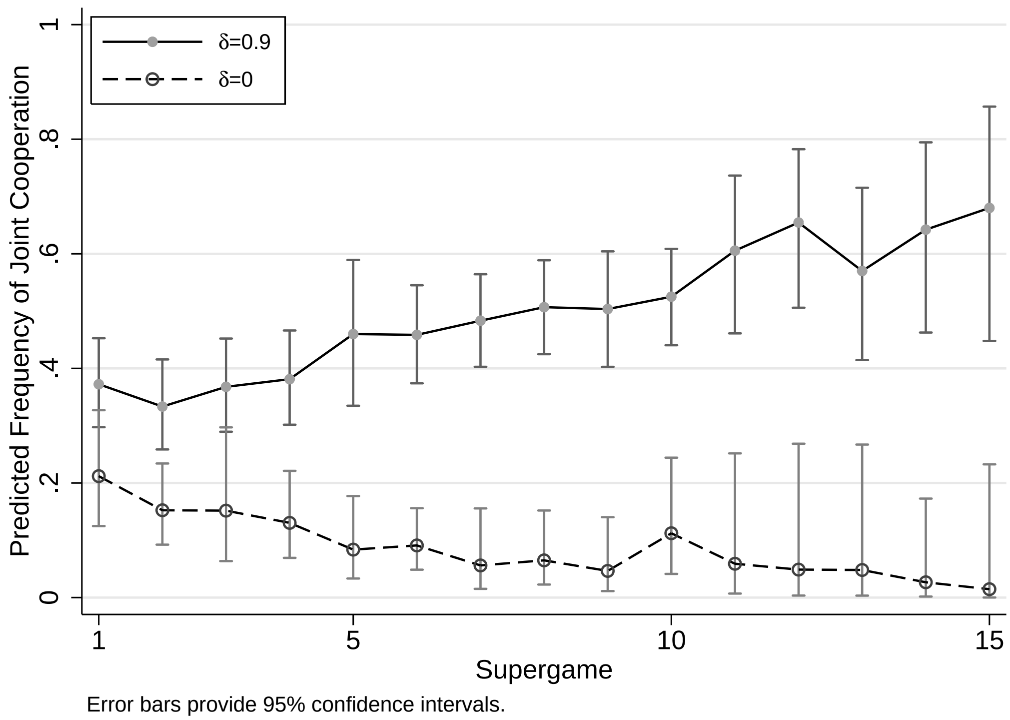

Figure 2: The Impact of δ on Last Round Joint Cooperation by Supergame: Predicted Frequency of Joint Cooperation for $g = \ell = \frac{1}{2}$

However, the standard theory does have testable predictions. For example, cooperation rates should not be higher when cooperation cannot be supported in equilibrium than when it can. Several articles have studied how cooperation depends on whether mutual cooperation can be supported in an equilibrium under the assumption of monetary payoffs and risk neutrality.[15]

For any payoff matrix, we can calculate the minimum δ required to support mutual cooperation in a subgame perfect equilibrium (SPE):

$$\delta^{SPE} = \frac{g}{1+g}. \tag{1}$$

The greater the temptation to defect (g) is, the greater the shadow of the

[15]For now, we exclude cases with equilibria in which alternating cooperation, but not mutual cooperation, can be supported.

future must be for subjects to have an incentive to cooperate.[16]

Roth and Murnighan [1978] and Murnighan and Roth [1983] were the first to show that cooperation levels are greater under treatments for which cooperation can be supported in equilibrium. This issue is revisited in the new wave of articles. In particular, Dal Bó and Fréchette [2011] show that cooperation declines to negligible levels when δ is such that cooperation cannot be supported in equilibrium, while it is higher when cooperation can be supported as part of an SPE.[17] We use the data from the new wave of experiments to study whether cooperation is greater in treatments in which mutual cooperation can be supported in equilibrium.

Table 6 shows cooperation levels depending on whether cooperation can be supported in equilibrium for supergames 1, 7 and 15.[18] The metadata reveal that indeed, cooperation rates are higher when cooperation can be supported as part of an SPE. Cooperation is significantly greater, on average, for treatments in which cooperation can be supported in equilibrium in supergame 1. The difference increases significantly as subjects gain experience: by supergame 7, the difference more than doubles the difference in supergame 1. This is due to a decrease in cooperation when it cannot be supported in equilibrium.

Figure 3 shows the cooperation percentage by treatment in round 1 of supergame 7. For each δ treatment, the figure also shows a vertical line separating the treatments having cooperation as an SPE outcome to the left, and

[16]This critical value δ^{SPE} is calculated as follows. Assume that the other player is following the strategy Grim: cooperate in round 1 and continue cooperating until there is a defection; then, defect forever. The Grim strategy provides the strongest possible punishment for defection. Under the assumption that the other player is following Grim, a player would obtain a payoff of $\frac{1}{1-\delta}$ by also choosing Grim (or any other cooperative strategy), while he would obtain a payoff of $1+g$ by defecting. Therefore, if the other player is following Grim, a player has an incentive to cooperate only if $\delta \geq \frac{g}{1+g}$.

[17]This question also is evaluated by Aoyagi and Fréchette [2009], who report results for $\delta = 0$ and $\delta = 0.9$ and, given their parameters, cooperation can be supported at the higher δ. They find much higher cooperation rates at $\delta = 0.9$.

[18]This table and the next compare cooperation levels in treatments that satisfy some theoretical condition with treatments that do not. In Table 9, we bring this analysis together with the analysis of the effect of δ, g, and ℓ started in Table 4.

Table 6: Equilibrium Condition and Cooperation in Round 1

	Not SPE	SPE	Difference
Supergame 1	34.34	51.23	16.89***
Supergame 7	13.86	48.83	34.97***
Supergame 15	16.67	53.05	36.38***

Significance levels calculated from a Probit regression at the subject level with an RD dummy. S.E. clustered at the paper level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

treatments in which cooperation cannot be supported in equilibrium to the right.[19] As can be seen in Figure 3, while cooperation tends to be higher when it can be supported in equilibrium, there is a lot of variance. Some of the treatments for which cooperation can be supported in equilibrium reach very high levels of cooperation, while others reach only very low levels (compare, for example, the treatment with $\delta = .875$ and $g = \ell = 1/3$ with the treatment with $\delta = .75$ and $g = 1$ and $\ell = 8$). Hence, cooperation being a possible equilibrium action does not necessarily imply that most subjects will coordinate on a cooperative equilibrium. In many of these treatments, a large fraction of subjects chooses to defect instead.

This evidence contradicts the common practice of assuming that players will cooperate whenever cooperation can be supported in equilibrium.[20] Similarly, the experimental literature on coordination games shows that subjects do not necessarily coordinate on Pareto efficient equilibria (see Van Huyck et al. [1990] and Cooper et al. [1992], among others). Infinitely repeated games are, in a way, coordination games for high enough δs. The results show that in infinitely repeated games, subjects do not necessarily coordinate on the best

[19]For example, cooperation can be supported in equilibrium only if $g \leq 1$ when $\delta = 1/2$. Cooperation can be supported in equilibrium for all the available treatments if $\delta = 0.75$ or greater. Cooperation cannot be supported in equilibrium if $\delta = 0$, regardless of g and ℓ.

[20]On this practice, see, among many others, Tirole [1988] (page 253). When discussing tacit collusion Tirole writes: "One natural method (to select from the multiplicity of equilibria) is to assume that the firms coordinate on an equilibrium that yields a Pareto-optimal point in the set of the firms' equilibrium points."

equilibrium, as well.

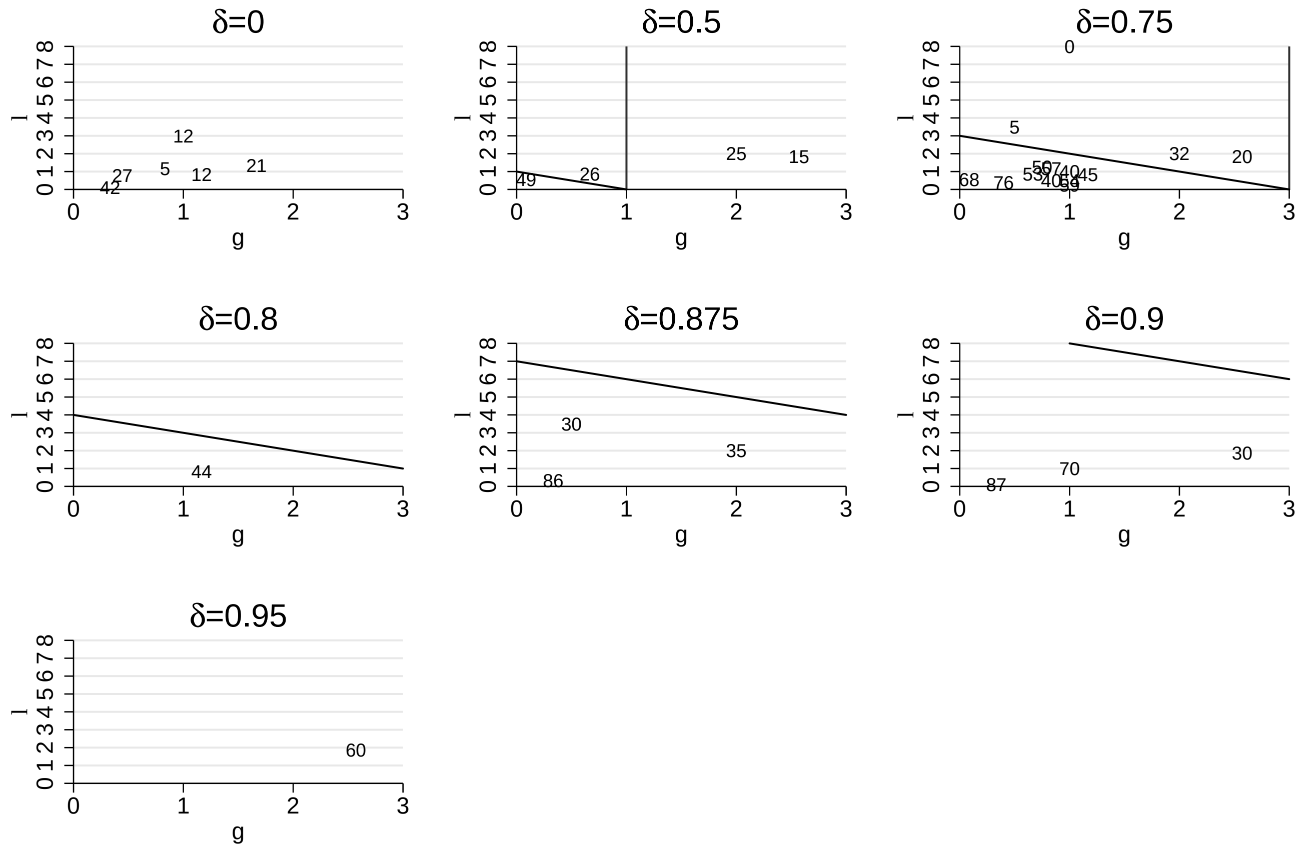

Note: the vertical lines separate the treatments depending on whether cooperation can be supported in an SPE and the diagonal lines depending on whether cooperation is risk−dominant.

Figure 3: Percentage of Cooperation by Treatment (Supergame 7, Round 1)

Result 2 *Cooperation is, on average, greater in treatments in which cooperation can be supported in equilibrium, but there is substantial variation across treatments. The fact that cooperation can be supported in equilibrium does not imply that most subjects will cooperate.*

A second way to study whether the shadow of the future affects behavior as theory predicts is to compare infinitely repeated games with finitely repeated games of the same expected length. Under the assumption of monetary payoffs, cooperation should unravel from the end in finitely repeated games. By comparing behavior in the first round of finitely repeated games and infinitely repeated games of the same expected length, we can test whether cooperation depends on infinite repetition, as the theory predicts, or whether it is simply affected by the length of the games.[21] Engle-Warnick and Slonim [2004] compare behavior in finitely and infinitely repeated trust games and, consistent with theory, find that once subjects have gained significant experience (last ten supergames out of 50) they trust more and are more trustworthy under infinite repetition. Similarly, Dal Bó [2005] finds that cooperation is significantly greater under infinitely repeated PD games than under finitely repeated ones. However, Lugovskyy et al. [2015] compare behavior in finitely and infinitely repeated public good games and do not find that contributions are higher in the latter. They do find, however, different patterns of behavior across the two conditions–in particular, the drop in contributions between the first and last round is greater in the finitely repeated case than in the random termination case. A collapse of cooperative behavior at the end of a supergame seems to be a hallmark of finitely repeated interactions (see Embrey et al. [2016]).[22]

On the whole, the evidence suggests that subjects with sufficient experience behave differently under finitely repeated and infinitely repeated games. However, further work is needed to better understand when and how behavior differs between the two condition.[23]

[21] Roth and Murnighan [1978] also mention the importance of such a comparison, but they do not conduct finitely repeated PD treatments. Instead, they mention prior evidence that cooperation rates are not affected by the horizon for lengths that overlap with the expected duration of their games.

[22] Normann and Wallace [2012] find no difference in average cooperation rates comparing finitely and infinitely repeated PD. The result is not very informative, however, since, as the authors note, subjects participated in only one repeated game and, hence, did not have a chance to learn about the environment. We similarly find that δ has a relatively small impact when focusing on the first supergame. A similar observation can be made about Tan and Wei [2014] who study public goods games.

[23] A third way to study whether the "shadow of the future" affects behavior as the theory predicts is to exploit variations in the set of possible equilibrium outcomes as a function of the parameters of the game. Dal Bó [2005] exploits the fact that a small variation in payoffs

2.3 What Explains Cooperation Levels?

The results from the previous section make it clear that SPE is not enough to understand cooperation levels in infinitely repeated PD games. Then, how does cooperation depend on the parameters of the game? When will the multiplicity of equilibria be resolved in favor of cooperation?

The fact that ℓ affects cooperation, as shown in Table 4, casts doubt that satisfying the SPE condition is enough for cooperation to arise (the SPE condition, $\delta \geq \delta^{SPE}$, depends on δ and g only). Why would ℓ affect cooperation? One reason the literature has considered is that subjects may be uncertain about the choice of the other player. The equilibrium condition (1) is based on the idea that the other player is following the cooperative strategy Grim. However, subjects may not be sure about the strategy of the other player and, hence, will worry about the cost of cooperating when the other defects: ℓ.

To deal with how strategic uncertainty may affect cooperation in repeated games, the experimental literature has borrowed from the literature on coordination games. Harsanyi and Selten [1988]'s concept of risk dominance (RD) has an easy application in symmetric coordination games with two strategies. A strategy is risk dominant if it is a best response to the other player randomizing 50-50 between the two strategies. While there are an infinite number of strategies in infinitely repeated games, one can focus on two strategies: Always Defect (AD) and a cooperative strategy like Grim (G). We say that cooperation is risk dominant if Grim is risk dominant in the game with only Grim and Always Defect (see Blonski and Spagnolo [2015]).

The condition for cooperation to be part of a risk dominant equilibrium of

can result in a large variation in the set of outcomes that can be supported in equilibrium. In that paper, one of the payoff matrices allows for mutual cooperation in equilibrium, while the other payoff matrix allows only for alternating cooperation in equilibrium when $\delta = 0.5$. Consistent with the theoretical predictions, there is more mutual cooperation in the first treatment and some evidence of alternating cooperation in the second treatment. Fréchette and Yuksel [2014] change payoffs in a within-subjects design using $\delta = 0.75$, such that after the payoff change, joint cooperation cannot be supported but alternation still can. They observe an important and rapid decrease in cooperation. There are some signs of alternation, but the cooperation rates are so low that it is difficult to evaluate. More work is needed to verify the robustness of these results.

this simplified game is given by

$$\delta \geq \delta^{RD} = \frac{g+\ell}{1+g+\ell}. \tag{2}$$

Dal Bó and Fréchette [2011] and Blonski et al. [2011] study whether risk dominance or subgame perfection determines choices. They find that risk dominance does explain cooperation, but that it provides only an incomplete picture. The same can be said once we consider the metadata.

Table 7 shows that cooperation rates are higher when cooperation is risk dominant. The difference increases as subjects gain experience: by supergame 7, the difference more than doubles what it is in supergame 1.

Table 7: Risk Dominance and Cooperation in Round 1

	Not RD	RD	Difference
Supergame 1	35.64	54.22	18.57***
Supergame 7	16.10	55.88	39.79***
Supergame 15	20.33	63.06	42.73***

Significance levels calculated from a Probit regression at the subject level with an RD dummy. S.E. clustered at the paper level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

However, there are treatments under which cooperation is risk dominant with low levels of cooperation. Figure 3, which shows the cooperation percentage by treatment in round 1 of supergame 7, also shows diagonal lines separating the treatments having cooperation as risk dominant to the left, and treatments in which cooperation is not risk dominant to the right. Cooperation tends to be higher when it is risk dominant, but there are a lot of differences across treatments (compare, for example, the treatment with $\delta = .875$, $g = \ell = 1/3$ and the treatment with $\delta = 0.875$, $g = 1/2$ and $\ell = 3.5$). Thus, cooperation being risk dominant does not necessarily imply that most subjects will coordinate on a cooperative equilibrium. In some of these treatments, a large fraction of subjects choose to defect instead.

Result 3 *Cooperation is, on average, greater in treatments where cooperation is risk dominant, but there is substantial variation across treatments.*

The fact that cooperation is risk dominant does not imply that a majority of subjects cooperate.

This result is quite surprising, as it indicates that in infinitely repeated games, subjects may not even coordinate on equilibria that are both Pareto efficient and risk dominant.

We say in the previous analysis that the concepts of SPE and RD are not enough to explain cooperation levels. A lot of variation in behavior remains unexplained.

What explains cooperation levels? More precisely, how can the parameters of the game be combined to generate an index that explains cooperation levels? These parameters may not necessarily impact cooperation linearly, as assumed in the analysis presented in Table 4, and they may not discontinuously follow the SPE and RD conditions.

The construction of indices to explain cooperation in the PD has a long history starting with Roth and Murnighan [1978] for infinitely repeated games and with an even longer history in psychology for PD games in general, see for example Rapoport and Chammah [1965]. Besides simply considering subgame perfection as a dichotomous predictor for cooperation, some authors also study whether the magnitude of the incentives to coordinate on cooperation also matters. Roth and Murnighan [1978], for instance, consider the two indices, $\frac{g}{1+g}$ and $\frac{1+g}{1+\ell}$, as "measuring the *difficulty* of achieving a cooperative equilibrium." (p. 196)[24]

In the same spirit, more recently, Blonski et al. [2011] considered a related index, namely $\delta - \delta^{SPE}$, which measures the difference between the continuation probability and the minimum required for cooperation to be an equilibrium outcome. The first panel of Figure 4 shows the relationship between $\delta - \delta^{SPE}$ and cooperation in the first round of the seventh supergame. This index is positively correlated with cooperation, with an apparent increase in

[24]The first index comes from the condition for cooperation to be subgame perfect, while the second comes from the condition for tit-for-tat (TFT) to be a Nash equilibrium. TFT starts by cooperating and then matches what the opponent did in the previous round. Other indices are constructed by only considering payoffs (but not how they relate to δ), as in Murnighan and Roth [1983].

slope when $\delta - \delta^{SPE} > 0$. Table 8 presents the results from a probit analysis of the effect of $\delta - \delta^{SPE}$ on cooperation. While for supergame 7 there is no change in slope when $\delta - \delta^{SPE} = 0$, by supergame 15 the slope is not significantly different from zero when $\delta - \delta^{SPE} < 0$, and it is significantly positive when $\delta - \delta^{SPE} > 0$.

However, a significant amount of variation remains to be explained. In particular, this index does poorly for treatments with a high value of ℓ, as seen in the low levels of cooperation reached in two sessions with high values of ℓ in Blonski et al. [2011]. This leads Blonski et al. [2011] to propose an alternative index, $\delta - \delta^{RD}$, which attempts to capture how risk dominant cooperation is in a given treatment. The second panel in Figure 4 shows the relationship between this index and cooperation in the first round of the seventh supergame.[25] Table 8 shows that the effect of $\delta - \delta^{RD}$ on cooperation is positive, especially when cooperation is risk dominant ($\delta - \delta^{RD} > 0$).

While this index exhibits a high correlation with cooperation, it is not clear that the extent to which cooperation is RD for a particular treatment should be measured in the dimension of the discount factor. Note that the frontier between treatments that have cooperation as RD and treatments that do not is not linear in the δ, g and ℓ space. Hence, ordering the treatments based on the distance to the frontier in the δ dimension may differ from the ordering based on the distance to the frontier in the g or ℓ dimensions.

Dal Bó and Fréchette [2011] address this issue by ordering treatments by how robust they are to strategic uncertainty as captured by the size of the basin of attraction of AD against G: $SizeBAD$. If cooperation can be supported in equilibrium, $SizeBAD$ is defined as the maximum probability of the other player following the Grim strategy such that playing Always Defect is optimal.[26] This index is equal to $\frac{1}{2}$ if $\delta = \delta^{RD}$, and it is decreasing in δ. The

[25] Using data from four studies, Rand and Nowak [2013] present a figure similar to the second panel in Figure 4.

[26] Formally, $SizeBAD$ is defined as follows:

$$SizeBAD = \begin{cases} 1 & \text{if } \delta < \frac{g}{1+g} \\ \frac{(1-\delta)\ell}{(1-(1-\delta)(1+g-\ell))} & \text{otherwise.} \end{cases}$$

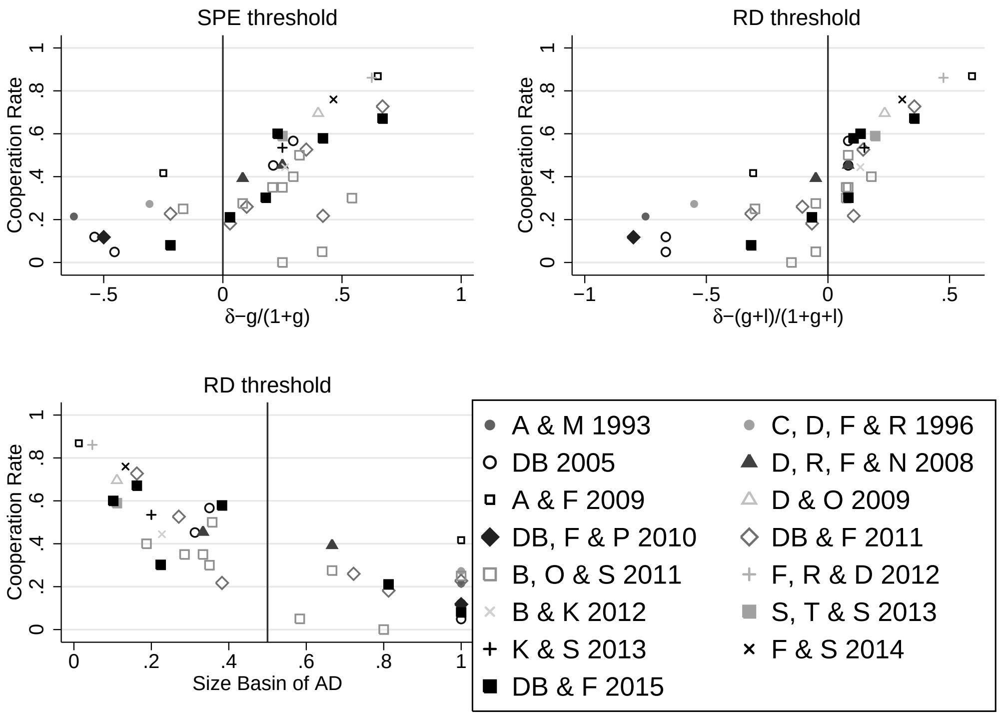

Figure 4: The Relation Between Three indices and Cooperation (Round 1, Supergame 7)

greater $SizeBAD$ is, the weaker cooperation is to strategic uncertainty. For example, if $SizeBAD = 0.8$, a player needs to believe that the other player will cooperate with probability 0.8 or higher for him or her to want to also cooperate. If cooperation cannot be supported in equilibrium, then $SizeBAD = 1$; AD is optimal regardless of beliefs about the other player's behavior.[27]

Dal Bó and Fréchette [2011] suggest that $SizeBAD$ may capture the role of the parameters of the game in determining cooperation. Indeed, they show that, in their data, $SizeBAD$ correlates with round 1 cooperation rates in the last supergame of a session. A similar observation can be made using the metadata. The third panel in Figure 4 shows the relation between $SizeBAD$

[27]Myerson [1991] studies the size of the basin of attraction of Always Defect against a cooperative strategy and shows that it is decreasing in the value of δ.

Table 8: The Impact of the indices on Cooperation (Round 1 - Marginal Effects)

	Supergame					
	7	15	7	15	7	15
	(1)	(2)	(3)	(4)	(5)	(6)
SPE	-0.0986	0.195				
	(0.145)	(0.136)				
$(\delta - \delta^{SPE})$×SPE	0.747***	0.979***				
	(0.0780)	(0.0733)				
$(\delta - \delta^{SPE})$×Not SPE	0.566**	-0.349				
	(0.282)	(0.275)				
RD			0.113**	0.121***	0.225	0.420*
			(0.0451)	(0.0415)	(0.220)	(0.243)
$(\delta - \delta^{RD})$×RD			1.030***	1.677***		
			(0.129)	(0.178)		
$(\delta - \delta^{RD})$×Not RD			0.238***	0.235		
			(0.0574)	(0.273)		
$SizeBAD$×RD					-0.902***	-1.139***
					(0.326)	(0.372)
$SizeBAD$×Not RD					-0.429*	-0.342
					(0.229)	(0.368)
N	2305	1030	2305	1030	2305	1030
Different Slope p-value	0.4622	<0.0001	<0.0001	<0.0001	0.3023	0.1046

Clustered Standard errors in parentheses

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

and cooperation rates in round 1 of the seventh supergame. Cooperation rates are relatively low when cooperation cannot be supported by a risk dominant equilibrium but cooperation rates decrease with $SizeBAD$ when cooperation is risk dominant. As shown in Table 8, by supergame 15, $SizeBAD$ affects cooperation only when cooperation is risk dominant.

Result 4 *Cooperation rates are increasing in how robust cooperation is to strategic uncertainty, especially when cooperation is risk dominant.*

The previous analysis provides a solution to the issue of multiple theoretical predictions. We find that cooperation rates are similar to the ones observed for one-shot PD games when cooperation can be supported in equilibrium but is not risk dominant. Risk dominance does not imply high cooperation rates, either. High cooperation rates are observed only for treatments that are very favorable to cooperation (as captured by a small $SizeBAD$ or large $\delta - \delta^{RD}$).

An interesting question is which of the three indices considered above predicts cooperation best. It is clear that $\delta - \delta^{SPE}$ is inferior to $\delta - \delta^{RD}$ and $SizeBAD$, as $\delta - \delta^{SPE}$ does not depend on ℓ, which has been shown to affect cooperation. Unfortunately, the available data does not allow a robust comparison of $\delta - \delta^{RD}$ and $SizeBAD$ given the high correlation of the two indices in the available data: -0.9172 for supergame 7. This high correlation is due to the choice of treatments from previous papers. Future research should study the explanatory difference between these two indices by appropriately choosing treatments such that the correlation is not as high. Given that the available data do not allow us to clearly compare the performance of $SizeBAD$ and $\delta - \delta^{RD}$, we will focus on $SizeBAD$ in future sections due to its better theoretical properties.

2.4 Learning

The previous section showed that experience can have large effects on cooperation. Are there factors that mediate experiences in systematic ways? Can we pinpoint features that affect how behavior evolves? Two such elements have been previously identified: the length of realized supergames and the choices of others.

Engle-Warnick and Slonim [2006a] note that in infinitely repeated trust games, the length of supergames correlates positively with subsequent choices. Similarly, Dal Bó and Fréchette [2011] show that the length of a supergame affects the likelihood that someone cooperates in the following supergame (even though those are played with different opponents). Dal Bó and Fréchette [2011] also show that a subject who was previously matched with someone who started by cooperating is more likely to cooperate. Both of these observations have been repeated in multiple papers, including Camera and Casari [2009], Sherstyuk et al. [2013], Embrey et al. [2013], Bernard et al. [2014], and Fréchette and Yuksel [2014].

Table 9: Determinants of the Evolution of Behavior (Round 1 Cooperation)

	(1)		(2)	
	Marginal Effects	Standard Errors	Marginal Effects	Standard Errors
g	-0.0215	(0.0143)	-0.0240	(0.0161)
ℓ	-0.0410***	(0.0126)	-0.0379**	(0.0149)
δ	-0.0722	(0.0889)	-0.0567	(0.0950)
SPE	0.0758	(0.0722)	0.0739	(0.0723)
RD	0.398***	(0.0626)	0.391***	(0.0660)
$SizeBAD\times$ RD	-0.969***	(0.177)	-0.958***	(0.182)
Supergame $\times$ Not RD	-0.0008**	(0.0004)	-0.0008**	(0.0004)
Supergame $\times$ RD	0.0038***	(0.0010)	0.0038***	(0.0010)
Length of Previous Supergame$-E$(Length)	0.0057***	(0.0007)	0.0057***	(0.0007)
Other's Coop in Previous Supergame	0.120***	(0.0149)	0.120***	(0.0148)
Turnpike			0.0188	(0.0167)
Complete Stranger			-0.0121	(0.0158)
Coop in Supergame 1	0.292***	(0.0351)	0.292***	(0.0350)
N	45,991		45,991	

Clustered Standard errors in parentheses

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

To confirm this, specification (1) of Table 9 reports the marginal effects of a correlated random effects probit of round 1 choices on: 1) the main factors that were established as important in the previous section, g, ℓ, δ, dummies for subgame perfection and risk dominance, and $SizeBAD$ if cooperation is risk dominant; 2) regressors to account for trends, the supergame number when cooperation is RD and the supergame number when cooperation is not RD; and 3) the regressors of interest, the length in rounds of the previous supergame, and the round 1 choice of the opponent in the previous match. The last regressor gives the mean of the random effects, which is a function of the choice of the subject in round 1 of the first supergame. One can think of this as capturing the "type" of the subject. Specification (2) will be discussed later in this section.

The results clearly indicate that after a longer supergame, subjects are more likely to cooperate and, similarly, more likely to defect following a short supergame.[28] Comparing a situation in which the previous match lasted 15 rounds with one in which it lasted five rounds increases the probability of cooperation by six percentage points. As a reference, in our sample, supergames have between one and 69 rounds, with a mean of 3.12 rounds. Does this mean that subjects do not understand how many rounds to expect for a given δ? Dal Bó [2005] elicited estimates of the expected number of rounds for $\delta = \frac{1}{2}$ and $\frac{3}{4}$ and found that most subjects had, in fact, a good idea of the expected number of rounds. About 71 percent and 56 percent of subjects for $\delta = \frac{1}{2}$ and $\frac{3}{4}$, respectively, knew the exact expected number of rounds, and the average guesses were 1.96 and 3.81, respectively.[29] Hence, this suggests that the effect of realized length on choices comes either from the minority of subjects who do not know the expected length, or from subjects updating their overall evaluation of the value of cooperation through experience. There is an interesting–as yet unexplored–question regarding the way that humans learn in infinitely re-

[28] If instead of the length of the previous supergame, we use the deviation from the expected number, results are qualitatively the same. This is not so surprising given that the specification already controls for δ.

[29] Similarly, Murnighan and Roth [1983] asked subjects about the probability that play would continue for at least two more rounds. While subjects tended to overestimate these probabilities, their beliefs strongly depended on δ.

peated games: Is the impact of the realized length constant throughout or is the impact more important early on?

Similarly clear, and non-negligible, is the effect of the choices of the previous player that one has interacted with. If the other player cooperated in round 1 of the previous match, then a subject is 12 percentage points more likely to cooperate to start his current interaction (which is most probably not with the same person). This process is consistent with subjects having beliefs about the proportion of the population that will cooperate and updating these beliefs as a function of their experiences. As they encounter more subjects who cooperate and update their beliefs upward about the fraction of subjects who cooperate, the value of cooperation increases and, thus, they are more likely to cooperate themselves. Note that this is in line with the previous observations on the role of the basins of attraction. The role of experience and of the basins of attraction suggest that if beliefs, at the beginning of an experience, are close to the dividing line determined by the $SizeBAD$, the specific experiences could lead to very different long-term behavior.

The following result summarizes the above observations.

Result 5 *Cooperation is affected by:*

1. *The realized length of previous supergames.*

2. *The choices of the past subjects with whom one was paired.*

2.5 Strategies

Choices across rounds in infinitely repeated games are likely not independent. In fact, their lack of independence is essential to support cooperation in repeated games: players must condition their behavior on past outcomes. Strategies are contingent plans that describe how players condition their behavior on the past.[30] Cooperation as an equilibrium outcome or as part of a

[30] We do not claim that humans form strategies in their head and then execute them, but simply that we can think of their actions as being the outcome of strategies. Strategies, here, serve as a way to understand the dynamic aspect of how choices interact: both choices that come about and those that are never made given the history of play.

risk dominant outcome are concepts that make sense with respect to strategies. Thus, understanding subjects' choice of strategies is of interest. In addition, the study of the strategies used allows for an additional test of the theory. For instance, if subjects are more likely to cooperate when δ increases, but never *punish* defection, then one would say that behavior is not consistent with the theory of infinitely repeated games, even though the movement of cooperation rates is consistent with the predictions of the theory. Identifying strategies could also help predict behavior better, especially after histories that are rarely reached. Finally, understanding strategy choice may elucidate aspects of human psychology. In other words, characterizing the strategies that subjects use can provide a deeper understanding of behavior.

Two main approaches have been used to study the choice of strategies in infinitely repeated games: eliciting strategies and inferring strategies from choices. The elicitation of strategies consists of asking subjects to submit a strategy: a plan of behavior for each possible contingency. Eliciting strategies has a long and well known history, starting with the work of Robert Axelrod. Axelrod [1980a] reports results for an experiment in which 14 participants (game theorists) submitted strategies for a finitely repeated PD. Axelrod [1980b] considers a randomly repeated PD ($\delta = 0.99654$) and includes 62 entries from game theorists and other participants recruited via announcements in journals for users of small computers. Other papers with tournaments have been written since, but a more common descendent of this method is the computer simulation method popular in evolutionary biology.[31] whether they use tournaments or simulations, these papers search for the best-performing strategy. From the point of view of studying behavior, however, their findings can serve only as suggestive evidence since successful strategies may or may not be popular.

Here, we focus on the strategies that people use rather than strategies that perform well. Even though this question is straightforward, answering it via

[31] A more recent example of a tournament is Selten et al. [1997], who use economics students with some programming experience to program strategies for a finitely repeated asymmetric Cournot duopoly. Examples of computer simulations in evolutionary biology are Nowak and Sigmund [1993] and Nowak et al. [1995].

elicitation is marred by difficulties. First, it is not self-evident that people think in terms of strategies (complete contingency plans). Hence, if we ask them to formulate a strategy at the onset, it might lead them towards behavior that is not what they would have done had we not triggered them to think about strategies. Second, and again because they may not typically think in terms of strategies, people do not necessarily know how to express what it is they are doing as a strategy. Third, the experiment must allow them to specify the strategy that they actually want to use and to provide a way to evaluate whether the important strategies are included. Dal Bó and Fréchette [2015] propose a method to tackle these potential concerns and elicit strategies for multiple parameters.[32]

A second way to study the choice of strategies is to infer strategy choice from the round-by-round choices (cooperate or defect). However, for multiple reasons, recovering strategies from choices in infinitely repeated games is a non-trivial problem, even for games as simple as the PD. First, the space of strategies is infinite (uncountable). Second, observed choices are only a small part of a strategy. Third, coordination creates a lack of variation. For example, if all subjects follow the Grim strategy, then we would observe cooperation in every round but not observe how subjects would react to a defection.

Note that it is not possible to solve these problems with the simple approach of regressing round t play on round $t-1$ outcomes for two reasons. First, that approach rules out strategies that may condition behavior on outcomes from more than a round ago; and, second, it does not fully consider the fact that the behavior of the partner in round $t-1$ depends on one's own behavior at $t-2$.

The various methods that have been proposed to estimate strategies consider the fit of choices and strategies in different ways. While there are many articles proposing different estimation methods (see Engle-Warnick and Slonim [2006b], Engle-Warnick et al. [2007], Aoyagi and Fréchette [2009], Dal Bó and Fréchette [2011], and Camera et al. [2012]), we will focus on the methodology developed by Dal Bó and Fréchette [2011]: the Strategy Frequency Estimation Method

[32]Bruttel and Kamecke [2012] partially elicit strategies so as to compare behavior with other ways of implementing infinitely repeated games in the laboratory.

(SFEM). This method consists of estimating a mixture model of the frequency with which each strategy from a pre-specified set of strategies appears in the pooled data, under the assumption that each subject always uses the same strategies across supergames but may make mistakes. The estimation consists of choosing the frequency of the different strategies and the probability of mistakes that maximize the likelihood of the observed sequences of choices. There are two possible perspectives on how to interpret these frequencies. One can think of each subject as following the same mixed strategy, and the frequencies are the mixing frequencies over the pure strategies.[33] Alternatively, these can be the fractions of subjects following each particular strategy.[34]

There are several reasons for us to focus on this identification method. The main reason is that SFEM has been used in several papers over multiple parameters of the PD, allowing us to compare strategy prevalence across payoff matrices, probability of continuation and monitoring structures. Moreover, we believe that this method has some noteworthy advantages over some of the others methods. A very simple advantage is that SFEM allows researchers to use standard statistical methods and tests. Another advantage is that it uses choices from multiple supergames to identify strategies. To see why this is important, consider the following example. For simplicity, assume that we are concerned about estimating the strategies of a single subject. Suppose that we have data for many supergames in which our subject cooperates in all rounds, but so does his partner. Such data are consistent with many strategies; for instance, always cooperate (AC) and Grim can both generate such sequences. However, imagine that in a few supergames, the subject displays behavior consistent with Grim but not with AC. Then, the SFEM would classify all of them as Grim. Other methods can be equivocal, concluding that much of the

[33]Given that a subject may participate in several supergames, this assumes that the mixing is done before the start of the first supergamge; from then on, the strategy is fixed. Note, however, that simulations show that the SFEM provides the average prevalence of strategies even when subjects randomize among strategies before each supergame–a subject can use different strategies in different supergames.

[34]In either case, the frequencies should be stable (over supergames). Hence, to minimize the chance that this is not the case, the data set used for estimation is typically composed of supergames late in the session.

data are consistent with both AC and Grim.[35]

Beside these advantages, there is also evidence that the SFEM performs well (if the relevant strategies are included in the set of possibilities). First, Fudenberg et al. [2012] tests the procedure on simulated data and find that the method produced results consistent with the data-generating process. Second, as discussed before, Dal Bó and Fréchette [2015] develop a method to elicit strategies while also generating choice-by-choice responses. This allows them to compare the results of the SFEM based on the choices of the subjects versus the strategies that the same subjects selected. From this comparison, they confirm that the estimation results are in line with the strategies that subjects selected across a variety of parameters. In addition, their analysis highlights some of the challenges of recovering strategies. In particular, omitting a strategy that is popular can distort results in important ways, especially if the forgotten strategy is "closer" than others to one that is included. In that case, the "closer" strategy will pick up most of what should have been attributed to the missing strategy. Nonetheless, overall, the SFEM and other approaches to estimation and elicitation reveal some clear regularities in the data, which we describe below.

Table 10 reports the results of five papers that estimate strategies using the

[35]The other methods use to identify strategies vary in several dimensions and have not been used across papers to identify strategies in infinitely repeated prisoners' dilemma games with perfect monitoring and fixed pairs across treatments (the focus of this section). Engle-Warnick and Slonim [2006b] study the number of strategies needed to perfectly predict every choice in infinitely repeated trust games. Their approach uses a penalty function to trade-off how many supergames can be perfectly predicted and how many strategies are required. Camera et al. [2012] study infinitely repeated PDs with random re-matching inside groups (which is covered later in this article). They focus on two-state automata and, for each supergame, find the strategy that best describes a subject's choice given a constant probability of incorrect transitions across the automata states. If the fit is sufficiently accurate (better than chance), the subject is classified as using that strategy. A subject can be classified by more than one strategy. Engle-Warnick et al. [2007] derive a posterior distribution over strategies for each subject using a Bayesian approach. Note that in certain environments, alternative methods may dominate SFEM. Consider, for example, the environment studied by Aoyagi and Fréchette [2009] with a continuous signal. The SFEM is not practical, as it is difficult to know, ex-ante, the reasonable thresholds of the signal for triggering punishments.

Table 10: Frequency (as %) of Key Strategies

	AC	AD	Grim	TFT	STFT	**Total**
Dreber et al 2008§			SFEM Estimation			
$\delta = \frac{3}{4}$, $g = 2$, $\ell = 2$	0	64	7	15	14	**100**
$\delta = \frac{3}{4}$, $g = 1$, $\ell = 1$	0	30	21	40	9	**100**
Dal Bó and Fréchette 2011§			SFEM Estimation			
$\delta = \frac{1}{2}$, $g = 2.57$, $\ell = 1.86$	0	91	0	7	2	**100**
$\delta = \frac{1}{2}$, $g = 0.67$, $\ell = 0.87$	0	76	0	6	8	**90**
$\delta = \frac{1}{2}$, $g = 0.09$, $\ell = 0.57$	1	49	0	24	4	**78**
$\delta = \frac{3}{4}$, $g = 2.57$, $\ell = 1.86$	0	66	0	23	0	**89**
$\delta = \frac{3}{4}$, $g = 0.67$, $\ell = 0.87$	0	11	4	21	8	**44**
$\delta = \frac{3}{4}$, $g = 0.09$, $\ell = 0.57$	2	0	2	55	0	**59**
Fudenberg, Rand, Dreber 2012†			SFEM Estimation			
$\delta = \frac{7}{8}$, $g = 0.33$, $\ell = 0.33$	24	6	12	15	0	**57**
Rand, Fudenberg, and Dreber 2015°			SFEM Estimation			
$\delta = \frac{7}{8}$, $g = 2$, $\ell = 2$	0	18	43	27	5	**93**
Fréchette and Yuksel 2014†			SFEM Estimation			
$\delta = \frac{3}{4}$, $g = 0.4$, $\ell = 0.4$	0	14	32	39	2	**87**
Dal Bó and Fréchette 2015‡			Elicitation			
$\delta = \frac{1}{2}$, $g = 2.57$, $\ell = 1.86$	0	53	6	5	14	**78**
$\delta = \frac{1}{2}$, $g = 0.09$, $\ell = 0.57$	3	25	36	19	3	**86**
$\delta = \frac{3}{4}$, $g = 2.57$, $\ell = 1.86$	2	47	10	10	12	**81**
$\delta = \frac{3}{4}$, $g = 0.09$, $\ell = 0.57$	8	12	35	30	0	**85**
$\delta = \frac{9}{10}$, $g = 2.57$, $\ell = 1.86$	2	14	17	39	7	**79**
$\delta = \frac{95}{100}$, $g = 2.57$, $\ell = 1.86$	0	22	6	25	0	**53**

§ As reported in Fudenberg et al. [2012], out of 11 strategies.
The original estimation in Dal Bó and Fréchette [2011] included fewer strategies.

† Out of 20 strategies.

‡ Out of 32 strategies or more. The numbers reported are from two treatments, one which allows more than 32 strategies.

° Out of 11 strategies.

SFEM for infinitely repeated PDs with perfect monitoring.[36] The table also shows the results from Dal Bó and Fréchette [2015] that elicits strategies. For each treatment of every paper, the estimated percentage of five key strategies is listed. These strategies are represented in Figure 5.[37]

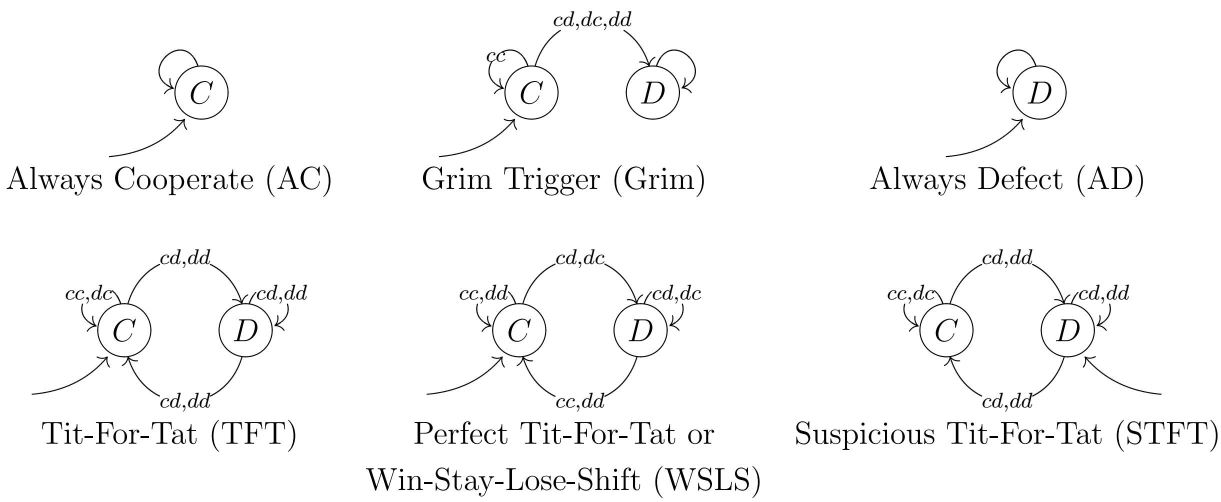

Figure 5: Six Strategies

The most-used strategies include AC and AD, which are self explanatory, and Grim and TFT, which we have already described, but also Suspicious-Tit-For-Tat (STFT). STFT (sometimes referred to as D-TFT) starts by defecting, and then matches the choice of the other player in the previous round. As can

[36] As subjects may be learning about the environment at the beginning of the experimental sessions, these estimates are based on the latter part of the sessions, when behavior is somewhat stabilized. See Dal Bó and Fréchette [2015] for further discussion of this issue.

[37] The representation of strategies as machines works as follows. The machine can be in one of a finite number of states. Transitions between states are governed by conditions. Hence, a machine is defined by the list of its states, the triggering condition for each transition, and its initial state. For example, consider the representation of TFT in Figure 5. The arrow coming from outside with no condition indicates the initial state. The letter inside the circle represents the choice that the strategy indicates in that state. The letters above the arrows indicate the conditions for transitions, which are determined by the pair of actions in the previous round: the first letter for the automaton and the second letter for the opponent's automaton.

be seen, these five simple strategies account for the majority of strategies in all but one treatment (16 out of 17). Furthermore, in most of these treatments, 13 out of 17, these five strategies account for more than three quarters of the strategies.[38]

Given the work of Axelrod, one could have expected subjects to use TFT. On the other hand, TFT is not subgame perfect, and, thus, Grim may be more likely. In fact, both Grim and TFT are present in the data. WSLS (not listed in this table), also known as Perfect TFT and introduced by Fudenberg and Tirole [1991], is a subgame perfect equilibrium for sufficiently patient players and was identified as a successful strategy in simulations by Nowak and Sigmund [1993]. However, it is never a sizeable fraction of the strategies used. AC, which is never an equilibrium, is estimated to have a low prevalence in all treatments. In fact, an even smaller subset of strategies account for most of the behavior: AD, Grim, and TFT. Together, these three strategies cover the majority of strategies in 15 of the 17 treatments and can account for at least 70 percent of them in 11 treatments.[39] Hence, under perfect monitoring, the majority of strategies are simple (two or fewer states), but not necessarily sub-game perfect equilibria (when playing themselves).

Result 6 *Three strategies account for most of the data: AD, Grim, and*

[38]Contrary to this result, in a recent paper, Breitmoser [2015] argues that the behavior of a majority of subjects (after a certain level of experience) in 17 treatments from previous studies is best described by a mixed strategy (after round one). This conclusion is at odds with the basic fact that in our metadata (including the data considered by Breitmoser [2015]), a vast majority of behavior is *perfectly* accounted for by one of the five pure strategies included in Table 10: AD, AC, Grim, TFT, and STFT. For example, for supergames greater than four, those five strategies account for 84 percent of play in supergames lasting more than one round and 79 percent of play in supergames lasting more than two rounds. In fact, in each of the 32 treatments we have, these five strategies always account for the majority. This is true for any restriction of the sort: "supergames with length $>$ some number of rounds." Note that the longest supergame in this sample has 44 rounds, making the likelihood of random choices looking like those from a pure strategy by chance virtually nil. We stress that this exercise does not allow for any mistakes: these strategies fit every choice exactly.

[39]In a recent paper, Romero and Rosokha [2015] propose a new and very flexible method to elicit strategies and they, too, find that these three strategies are the most common.

TFT. Therefore, subjects use punishments to support cooperation, but punishments are not necessarily credible (not SPE).

Note that the popularity of AD, Grim, and TFT seems to extend to more-complex PDs with perfect monitoring. Rand et al. [2015] conduct two treatments in which choices are implemented with error, but the subjects know this, and they are informed not only of the implemented choices, but also of the intended choice of the other player. In Aoyagi et al. [2015], one treatment implements a PD with perfect monitoring where the payoffs are a lottery. AD, Grim, and TFT are estimated to correspond to 64 percet and 28 percent in the two Rand et al. [2015] treatments and 61% in Aoyagi et al. [2015].[40] Although the percentage is low for one of the three, there are many more possible strategies in this case (since behavior can be conditioned on both what the other player intended to do and the outcome). Together, this is suggestive evidence that these three strategies may be important even in more-complex games.

The popularity of simple strategies under perfect monitoring seems to extend to infinitely repeated trust games. Using different estimation methods, Engle-Warnick and Slonim [2006b] find that two-state machines fit more than 80 percent of the observed behavior (also see Engle-Warnick and Slonim [2004]). In particular, they find that trustors tend to use the Grim strategy.

3 Inducing Infinitely Repeated Games in the Laboratory

Before moving on to explore other determinants of cooperation, we take a small methodological detour on the different ways to induce infinitely repeated games in the laboratory. The main method to induce infinitely repeated games in the laboratory is to have a random termination rule: after each round, the game continues with probability δ and ends with probability $1 - \delta$. Subjects are then paid the sum of payoffs from all rounds. Under the assumption of risk neutrality, this method, first introduced by Roth and Murnighan [1978],

[40]Out of 17 strategies for Rand et al. [2015] and 15 for Aoyagi et al. [2015].

induces a preference over outcome that coincides with an infinitely repeated game with discount factor δ.[41]

Several issues may work against the equivalence between infinitely repeated games with exponential discounting and random termination. First, subjects may not understand probabilities and get confused about how the probability of continuation relates to the distribution of supergame lengths. As discussed in Section 2.4, the evidence provided by Murnighan and Roth [1983] and Dal Bó [2005] suggests that subjects, although not perfect, have a good understanding of how δ affects the distribution of supergame lengths.

Second, subjects may think that the experiment cannot last forever, and, therefore, the constant probability of continuation may not be credible (see Selten et al. [1997] for such a critique). For example, if subjects assign probability zero to the event that the experiment will still be going on next year, then, in the last second of the year, they will know that they are participating in the last round, and cooperation should unravel from the end. Logically, this argument hinges on subjects assigning probability zero to a given experimental length, a year in this example. However, there is no reason to believe that in the last second of the year, subjects would still believe that one more round is impossible. Nonetheless, this criticism does lead to an important issue.[42] Many elements outside of the experimenters' control may affect the length of a supergame: the server may crash, or there may be an earthquake. As such, it is likely that subjects do not exactly believe that the probability of continu-

[41] Note that in the theory literature, the model with infinite repetitions and with exponential discounting is sometimes interpreted as modeling a situation with random termination (see, for example, Mailath and Samuelson [2006], section 4.) Which interpretation is more appropriate depends on the application.

[42] From a practical perspective, there are many indications that subjects do not perceive randomly terminated games as they do finitely repeated games. First, there is the evidence from Dal Bó [2005] that cooperation rates are higher in randomly terminated games than in finitely repeated games of the same expected length. Second, in finitely repeated games, Embrey et al. [2016] document in a meta-study that last-round cooperation rates always drop close to zero. There is no similar movement toward no cooperation in randomly terminated games (in cases where cooperation is observed to start with). Third, there is no evidence that cooperation rates drop at the end of sessions implementing randomly terminated games.

ation is the one implemented by the experimenter. Instead, subjects may have a subjective perception of the probability of continuation. Therefore, analyses that rely on the exact measure of the probability of continuation need to be taken with a grain of salt, and more attention should be paid to analyses based on comparisons between treatments (as, for example, increases in δ, which increase subjects' perception of the probability of future interactions). Of course, this is a comment that applies to experimental work in general.

Third, subjects may not be risk neutral. While it is highly plausible that subjects may display preferences for risk that are not neutral, even with low stakes, the evidence provided in Section 4–that risk attitudes do not affect behavior in infinitely repeated games–suggests that deviations from risk neutrality are not likely to be a problem. Sherstyuk et al. [2013] also provide evidence in this direction. They compare behavior in random termination PD games when subjects are paid for all rounds and when they are paid only for the last round. They show that both payment methods induce the same preferences if subjects are risk neutral and different ones if they are not. Consistent with the assumption of risk neutrality, they find that behavior is similar between these two treatments.

Fourth, rematching subjects more than once per session may make it possible to support cooperation even when the parameters are such that, in a single randomly terminated game, cooperation cannot be supported in equilibrium. Such an argument would rely on some form of contagion type equilibria. In section 6, we discuss studies that explore subjects' ability to use such equilibria to support cooperation. More to the point, we use the metadata to determine whether cooperation rates vary systematically across studies as a function of the matching protocol. The three main protocols used are random rematching (random pairing with replacement), perfect stranger (without replacement), and the turnpike protocol.[43] The tradeoff, of course, is that the number of rematches that can be performed is highest for random rematching

[43] The turnpike protocol, first introduced by McKelvey and Palfrey [1992], eliminates the possibility of contagion by matching subjects such that if i was matched with j in a given round, and j was matched with k after that, then i is not matched with k after k is matched with j.

and lowest for the Turnpike design, and, as we have seen, experiencing multiple supergames is important. The regression reported on the last two columns of Table 9 indicates that there are no differences in round 1 choices across these three rematching designs. Note, also, that given the large sample size, there are no power issues.

Other types of designs, besides the Roth and Murnighan [1978] implementation, have also been used to induce infinitely repeated games in the laboratory. For example, Sabater-Grande and Georgantzis [2002], Cabral et al. [2014] and Vespa [2015] use a design with a fixed number of rounds that are played with certainty with payoffs that are exponentially discounted at rate δ; after those rounds, they use a random termination rule with probability of continuation δ. Another design, used by Andersson and Wengström [2012], Cooper and Kühn [2014a] and Cooper and Kühn [2014b], consists of a fixed number of rounds of the PD with discounting, followed by a coordination game. In particular, the coordination game in Cooper and Kühn [2014a] and Cooper and Kühn [2014b] is obtained by considering the expected payoffs in the continuation game of the PD from two relevant infinitely repeated game strategies: Grim and AD.

Fréchette and Yuksel [2014] compare behavior in experiments with random termination rules (RT), experiments with a fixed number of rounds with discounting followed by random termination (D+RT) and fixed rounds with discounting followed by a coordination game (D+C). In addition, Fréchette and Yuksel [2014] introduce a variation on RT in which subjects play in blocks with a fixed number of rounds. Subjects are not informed of whether or not the supergame has ended within a block, and they are paid only for rounds up to termination. A new block of rounds starts only if the supergame has not ended in the previous block. This block random termination design (BRT) allows the researcher to observe behavior in a larger number of rounds than in RT. In addition to the difference in ways to induce infinitely repeated games in the laboratory, Fréchette and Yuksel [2014] consider two different combinations of payoff parameters. In one of these treatments, mutual cooperation can be supported in equilibrium, while in the other treatment it cannot. They find that cooperation is significantly larger in the treatment in which mutual cooperation can be supported in equilibrium under all designs to induce infinitely repeated games. However, the magnitudes of the treatment effects are somewhat different across designs, with the largest effect under RT and the smallest under D+RT.[44]

In conclusion, there are different ways to implement infinitely repeated games in the laboratory, each with its specific advantages and disadvantages. The appropriate method will depend on the application that one has in mind, on the question being asked, and on which design will best allow one to answer that question. More work needs to be done to understand what differences, if any, different implementation methods generate and what those differences teach us about how people perceive and react to dynamic incentives. However, a rough guideline could include the following recommendations. Payment should definitely not be based on a round selected at random. There are theoretical reasons to prefer paying only for the last round, but there is no evidence that, in practice, it produces different results from paying for all rounds. Subjects should not be matched in fixed pairs (across supergames). Also, theory indicates that a turnpike design is the most robust; and intuition suggests that a round-robin matching procedure may be preferable to random re-matching. This being said, our results indicate that round 1 behavior is the same under random re-matching, round-robin re-matching, and a turnpike design (see Table 9). Using payoff discounting followed by random termination or the block random termination method makes sense when observing long interactions is important for the question at hand.[45] Barring such needs and concerns, the standard implementation with random re-matching and payments for all rounds has the advantage of making results more directly comparable to those of other studies.

[44]Interestingly, in the D+C design, behavior in the coordination game is largely independent of cooperation in the last round of the PD. This suggests that subjects do not use the coordination game as a way to provide incentives to behave in the PD rounds. Hence, this design is very useful from the point of view of facilitating the study of communication, which is the topic of papers that have used this method, but, in the eyes of the subjects, it may not be equivalent to an infinitely repeated game.

[45]See, also, the *one-period ahead strategy method* introduced in Vespa [2015], which substantially increases power for the purpose of identifying strategies.

4 The Impact of Personal Characteristics and Motives on Cooperation

As we have seen, there is a significant amount of heterogeneity of behavior even after subjects have gained experience in a particular treatment. While some subjects attempt to establish cooperative relationships, others defect. Could personal characteristics explain this heterogeneity in behavior?

Several articles have explored the connection between risk attitudes and cooperation in infinitely repeated games. Sabater-Grande and Georgantzis [2002] elicit risk preferences by having subjects choose between lotteries and then dividing them in three groups based on their risk aversion. Subjects are matched in each group to play one infinitely repeated game.[46] The authors find that cooperation is negatively related to risk aversion.

Other papers also study the relationship between cooperation and risk aversion, in addition to other personal characteristics. Dreber et al. [2014] does not find a robust relationship between survey responses to question regarding risk attitudes and cooperation. Similarly, Proto et al. [2014] and Davis et al. [2016] do not find a relationship between risk aversion and cooperation. It is not clear why these three papers find no correlation between risk aversion and cooperation, while Sabater-Grande and Georgantzis [2002] does.[47]

Another personal characteristic that has been investigated is whether the subject is an economics major. There is some previous evidence of differences in cooperation rates between economics majors and other students in one-shot prisoners' dilemma games (Frank et al. [1993]). Dal Bó [2005] finds that economics majors tend to cooperate less when cooperation is not SPE, but there

[46] They implement the infinitely repeated game by having 15 rounds with decreasing payoffs (payoffs in round $t + 1$ are 14/15 of the payoffs in round t) and then a probability of continuation of 14/15. Given the payoffs of the PD game, cooperation is both SPE and RD under the assumption of risk neutrality.

[47] One possibility has to do with experience, as subjects participated in only one supergame in Sabater-Grande and Georgantzis [2002], while they participated in several in the other papers. A second possibility has to do with the fact that Sabater-Grande and Georgantzis [2002] is the only of the four papers that groups subjects according to their risk aversion. However, subjects were not aware of being grouped by their risk aversion.

are no significant differences when cooperation can be supported in equilibrium. However, Dreber et al. [2014] and Sherstyuk et al. [2013] do not find a robust relationship between being an economics major and cooperation.

The literature has also not found a robust relationship between gender and cooperation in infinitely repeated games. While some papers find that women tend to cooperate less in some treatments, other papers do not find such a relationship (see Murnighan and Roth [1983], Dreber et al. [2014], Sherstyuk et al. [2013], Davis et al. [2016] and Proto et al. [2014]).

Note, also, that Davis et al. [2016] do not to find a correlation between patience and cooperation.

Proto et al. [2014] study the relationship between cooperation and intelligence (among other personal characteristics). Subjects were separated into two groups based on performance on an IQ test; they then participated in a series of infinitely repeated PD games. Proto et al. [2014] find that, for a high δ, there are no differences in behavior between the groups in the first supergame, but behavior diverges as they gain experience: subjects in the high IQ group learn to cooperate. They do not find such a difference for a treatment with a lower δ. They also do not find that personality traits, as measured by the Big Five (openness, conscientiousness, extraversion, agreeableness and neuroticism), have a systematic effect on cooperation.

Since appealing to social preferences is common in explaining behavior in one-shot games, one would expect that social preferences could also play a role in infinitely repeated games. Dreber et al. [2014] show that there is no correlation between behavior in dictator games and infinitely repeated games in which cooperation can be supported in equilibrium. There is, however, a positive and significant correlation between giving in the dictator game and cooperation in the infinitely repeated prisoners' dilemma game when cooperation cannot be supported in equilibrium. Similarly, Davis et al. [2016] find no correlation between behavior in a trust game and behavior in infinitely repeated games in which cooperation can be supported in equilibrium. In addition, when Dreber et al. [2014] surveyed subjects on their motivations for cooperation, "earning the most points in the long run" was the main motivation for most players. They also show that the type of strategies used by the subjects cannot be better explained by appealing to inequity aversion or altruism. They conclude that social preferences are not important in explaining the heterogeneity of behavior in infinitely repeated games. Rather, behavior is consistent mainly with subjects trying to maximize their monetary payoffs.

Reuben and Suetens [2012] and Cabral et al. [2014] also provide evidence that cooperation in infinitely repeated games is motivated mostly by strategic considerations. By strategic considerations, they mean that subjects cooperate so that their partner will cooperate in the future. In other words, they exhibit forward-looking strategic behavior, as opposed to, for instance, backward-looking reciprocity or simple altruism. Reuben and Suetens [2012] study an infinitely repeated game in which players can submit their actions conditional on whether or not the round is the last one. In addition, one of the players can also condition his or her action on the action of the other player. While the parameters are such that cooperation cannot be supported in equilibrium, they find that cooperation is greater when the round is not the last. That is, if a subject knows the round is the last one, he or she is more likely to defect than if the supergame might continue. This clearly shows that there is a strategic component to cooperation. Cabral et al. [2014] find the same result for infinitely repeated veto games in which cooperation can be supported in equilibrium. When subjects condition on the round being the last one, their choice is more likely to be consistent with a one-shot game than when the supergame might continue. Moreover, Reuben and Suetens [2012] find only a small percentage of cooperation in the last round by subjects informed that the other player cooperates in that round. That is, they find very little intrinsic reciprocity (reciprocating cooperation even when there is no "shadow of the future").

Note that there are as few as a single article studying some of the characteristics discussed in this section. More work needs to be done on some of these characteristics to assess their relationship with cooperation. Nevertheless, at this point, the main findings are as follows.

Result 7 *There is no robust evidence that risk aversion, economic training, altruism, gender, intelligence, patience, or psychological traits have a systematic effect on cooperation in infinitely repeated games in which cooperation* *can be supported in equilibrium. There is evidence consistent with the idea that the main motivation behind cooperation is strategic.*

If most cooperation is motivated by strategic considerations, why do we still observe significant levels of defection in treatments that are very favorable to cooperation? Why do some subjects do not realize that it would be highly profitable to start a cooperative relationship?

It is interesting that altruistic and trusting tendencies (as captured by the dictator and trust games) do not seem to play an important role in infinitely repeated games. Future research should, on the one hand, study why infinite repetition seems to reduce the importance of other-regarding preferences, and, on the other hand, continue to search for personal characteristics that would help us predict who will attempt to establish cooperative relationships. We should also study whether personal characteristics correlate with the strategies used to support cooperation. It is intriguing that we systematically observe a correlation between the choice in the first round of the first supergame and cooperation in later supergames, while few personal characteristics have been found to systematically correlate with behavior.[48] Interestingly, Embrey et al. [2013] find that controlling for whether subjects communicated and agreed to cooperate eliminates the correlation between their first and subsequent choices, suggesting that an important trait might be how subjects react to strategic uncertainty.[49] Hence, a natural question arises: are the subjects who cooperate more likely to coordinate on the Pareto efficient equilibrium in a simple coordination game?

[48]The finding that the first choice a subject makes in a session correlates to the subsequent choices in the session is a side-product of estimating a correlated random effects model for cooperation, as we do in Section 2.4. As reported in Table 9, the round 1 choice correlates with the choices that follow, even after controlling for the structural features of the environment (payoffs and discount), as well as other factors. Previous papers have also made this observation not pooling across different treatments (Dal Bó and Fréchette [2011], Embrey et al. [2013], and Fréchette and Yuksel [2014]).

[49]Embrey et al. [2013] is a PD with imperfect public monitoring where subjects can suggest (and agree to) what they both should play (these agreements are non-binding) before making their choice (which is not directly observable). The paper is described in more detail in the next section.

5 On the importance of information

In order to provide incentives for each agent to behave in a manner that is distinct from behavior in a one-shot game, agents must condition future behavior on past actions in a repeated game. Hence, it is no surprise that the ability of agents to monitor the actions of others affects what can be supported in equilibrium. However, even in the absence of perfect monitoring, agents can, in theory, support equilibria with higher expected payoffs than in a static game.[50] What can be supported will depend on the details of the environment, and an important factor is the type of monitoring: public or private. In other words, do all agents know what they have all observed (besides the knowledge of their own action), or can their signals be different? This section explores the relationship between monitoring and cooperation and is divided in two sections, one for each type of monitoring.

5.1 Imperfect Public Monitoring

This section covers experimental papers, each with a very different form of imperfect public monitoring. The aim is to understand how imperfect public monitoring affects cooperation and the strategies that subjects use to support it.

We start by discussing the evidence regarding cooperation. Figure 6 shows levels of cooperation observed in several articles studying imperfect public monitoring in infinitely repeated games. Cooperation levels from imperfect monitoring treatments are denoted by circles when cooperation can be supported in equilibrium and by empty squares when cooperation cannot be supported in equilibrium. Empty triangles denote cooperation levels from perfect monitoring treatments. As not all papers include treatments with perfect monitoring, we include as a benchmark the levels of cooperation that could be expected under perfect monitoring in the infinitely repeated game based on the analysis from Section 2 (this is done as in Table 5). Similarly, we include as another benchmark the levels of cooperation that could be expected

[50] See Green and Porter [1984], Abreu et al. [1990], and Fudenberg et al. [1994].

from a similar one-shot game. These two benchmarks will help us understand whether subjects are able to reach cooperation levels that are greater than one-shot levels and lower or similar to levels obtained in repeated games with perfect monitoring. In the left panel, this figure presents cooperation levels in round 1 of the seventh supergame as a measure of cooperative intentions. In the right panel, the figure shows levels of joint cooperation in the last round of the seventh supergame as a measure of the capacity to establish long-lasting cooperative relationships (this measure is especially important, as punishments happen in equilibrium under imperfect monitoring). Note that for Aoyagi and Fréchette [2009], Fudenberg et al. [2012]'s payoff matrix 4, and Aoyagi et al. [2015], there is more than one data point because these experiments vary the quality of monitoring for a given payoff matrix.

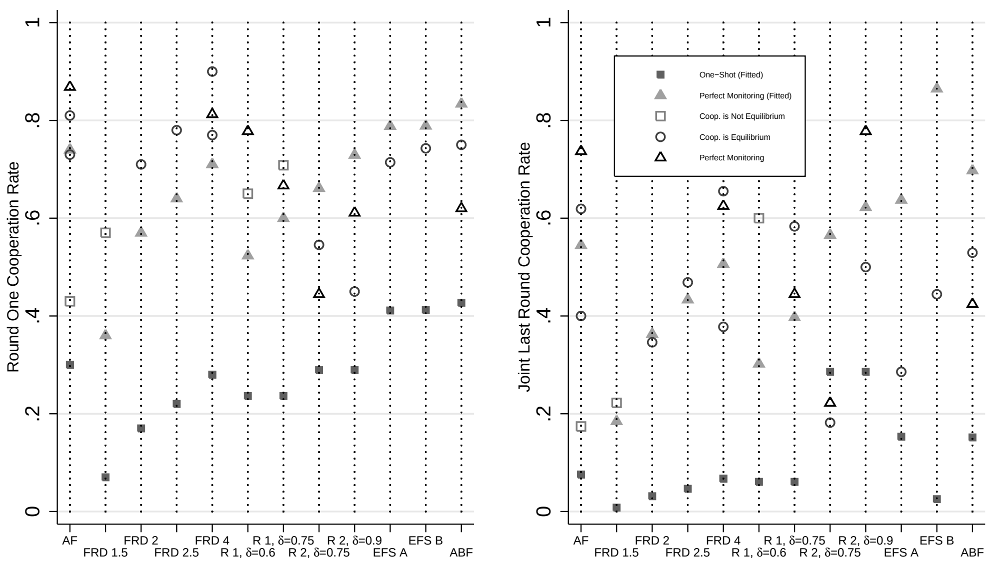

Hollow markers are data. Solid markers are fitted values from the meta−analysis.
For the cases with more than one hollow circle, the highest one always had a more precise monitoring.
AF: Aoyagi and Frechette 2009. n = 166 over four treatments.
FRD: Fudenberg et al. 2012, they differentiate their treatments by (1+g)/l. n = 384 over six treatments.
R: Rojas 2012, parametrizations 1 and 2. n = 158 over eight treatments. The supergames for the two treatments with δ=0.75 in the right panel ended in round one (there is only one session per treatment).
EFS: Embrey et al. 2013, payoff matrix A and B. Includes structured communication. n = 77 over two treatments.
ABF: Aoyagi et al. 2014. Perfect monitoring has random outcomes. n = 134 over two treatments.

Figure 6: Cooperation under Imperfect Public Monitoring (Supergame 7)

Aoyagi and Fréchette [2009] consider a standard PD in which subjects receive only a noisy public signal of each other's action. The signal is one-dimensional and continuous, and, hence, is similar to the price signal in the Green and Porter [1984] environment.[51] However, subjects play the expected value game and thus their signal does not affect realized payoffs.[52] The paper analytically derives the maximal symmetric perfect public equilibrium payoff for a specific distribution of the random component of the public signal, and the authors show that it is decreasing in its variance (the level of noise). However, this result also illustrates that even in the presence of imperfect monitoring, cooperation can be sustained, as the equilibrium payoffs can be above their one-shot level for positive levels of noise. The experiment shows that, indeed, subjects can support cooperation even when monitoring is imperfect. As seen in Figure 6, when it can be supported in equilibrium, round 1 cooperation under imperfect public monitoring in Aoyagi and Fréchette [2009] is similar to cooperation under perfect monitoring. This does not mean that information does not matter: in line with the maximal payoff, behavior is such that, as monitoring becomes noisier, cooperation and average payoffs decrease.

Fudenberg et al. [2012] implement a very different public signal: it is finite and two dimensional, and it affects payoffs. More specifically, after making their choice in the PD, there is some probability ($\frac{1}{8}$ in most treatments) that each player's choice will be altered (to be the other choice). The modified choices determine payoffs, and both players are informed of the implemented (modified) choices.[53] Figure 6 shows the results for each of the four payoff

[51] Green and Porter [1984] study a case win which firms play a Cournot game with shocks to demand that make the actions not directly observable. In their setting, firms observe the product price, which is an imperfect (due to the unobservable random shocks to demand) public (they all observe the same price) signal of the actions.

[52] At the end of each supergame, subjects are informed of their payoffs. If they did not play the expected value of the game, subjects could infer choices from the payoffs and the public signal. Note that, in theory, the way that noise affects payoffs should not matter; only the strategic form of the game is relevant. However, it could be that this does matter for how people behave. The evidence so far is that it does not (see footnote 69 for more details), but in the interest of completeness, we will be precise about how noise determines payoffs in the way each experiment is implemented.

[53] Note that some of the differences in the specific implementation of the public signals across papers may not matter theoretically but could have an effect in practice. For instance,

matrices considered by Fudenberg et al. [2012], who find that subjects can support cooperation despite imperfect monitoring. Moreover, they find that cooperation is greater when it can be supported in equilibrium. For each of the payoff matrices, the figure shows that cooperation under imperfect public monitoring is greater or comparable to what we would expect under perfect monitoring.[54]

Embrey et al. [2013] study imperfect public monitoring in a repeated partnership game. In the stage game, subjects simultaneously choose a costly effort level to contribute to a joint project. The higher the chosen effort levels, the higher the chance that the project will be successful. At the end of the stage, subjects do not observe the effort level chosen by their partner; instead, they observe only whether the project was a failure or a success, with the latter resulting in both subjects receiving a higher payoff. They consider versions of the game with two or three available effort levels. When only two effort levels are available, the game is a PD in expected value (they consider two different payoff matrices: A and B). In this case, the imperfect monitoring environment has a unidimensional, binary–and, thus, very coarse–public signal that has direct payoff consequences. Again, the results indicate that subjects can support cooperation despite imperfect monitoring, as seen in Figure 6. Interestingly, joint cooperation levels in the last round are much lower, consistent with the idea that bad signals trigger punishments.

Using this framework, Embrey et al. [2013] explore the role of renegotiation proofness in equilibrium selection, a selection concept that has drawn a great deal of attention in the theoretical literature.[55] The requirement of

one could imagine that strategies that defect after a "bad" signal and return to cooperation after a "good" signal may be more attractive in a setting such as this one, where it imposes a fee on the party associated with the defection to return on the cooperative path.

[54] In one treatment (labelled FRD 4 in the figure), they also conduct treatments with lower probability that the choice is modified, $\frac{1}{16}$ in one case and perfect monitoring in the other. Although both panels indicate higher frequencies for the case with low noise than no noise, the differences are not statistically significant in either case.

[55] Unlike most papers in this literature, Embrey et al. [2013] include a structured, non-binding communication protocol through which agents can agree to an action combination or to a conditional action plan of the type: do this now; then, if the project is a success, we do this, and if the project is a failure, we do that.

renegotiation proofness reduces how much cooperation can be supported in equilibrium, as it limits the severity of the punishment phases: more-severe punishments are no longer credible, as they would be renegotiated should they be reached.[56] Loosely speaking, the frequent observation that payoffs in prior experiments are below the maximum equilibrium payoffs could be because the harsh punishments required to sustain high levels of cooperation are not credible due to renegotiation concerns. Results from this experiment, however, suggest that subjects do not select the renegotiation proof strategy. By adding a medium effort level to the two-choice stage game, the maximum cooperation that can be supported in a renegotiation-proof equilibrium decreases from the high to the medium effort level. Despite this prediction, subjects rarely select the medium effort level. On the other hand, in line with the renegotiation proofness logic, when the changes in stage games prescribe the use of short (forgiving) punishment phases, among cooperative strategies, forgiving strategies are indeed more common.

Aoyagi et al. [2015] explore yet another implementation of the PD with imperfect public monitoring.[57] The signal an agent receives of the other's choice corresponds to the correct choice with probability 0.9, but with probability 0.1 it is the opposite choice. Payoffs are determined by one's own choice and the signal one receives of the other's choice. Both players are informed of the signal that the other received. Hence, as in Fudenberg et al. [2012], the public signal is finite and two dimensional, and affects payoffs. However, unlike in Fudenberg et al. [2012], for example, if both players cooperate but both are told that the other has defected, they both receive the *sucker's* payoff, while they would receive the *punishment* payoff in Fudenberg et al. [2012]. As seen in Figure 6, round 1 cooperation in supergame 7 is greater under imperfect public monitoring than under perfect monitoring, but this difference is not statistically significant; and, similarly, the rate of joint cooperation in the last round is not statistically different across conditions.

The finding in Aoyagi and Fréchette [2009] that increased noise leads to

[56] They use the renegotiation-proof concept of Pearce [1987].

[57] This paper focuses on private monitoring, which will be described in more detail in the next subsection.

lower payoffs seems specific to their environment. Both Fudenberg et al. [2012] and Aoyagi et al. [2015], who have treatments with perfect and imperfect public monitoring, do not observe lower payoffs under imperfect monitoring. It may be that the relation between monitoring and payoffs depends on the form of the public signal.

Finally, Rojas [2012] uses a clever design to implement imperfect public monitoring.[58] Subjects play a randomly terminated PD in which the specific payoffs can be one of three possibilities corresponding to low, medium, or high demand in a Cournot game.[59] After each round, subjects are informed of their payoffs. However, some payoff numbers are the same for different combinations of the other player's choice and demand. For instance, the payoff to seller one when he selects the *High* output is the same if seller two selects the *Low* output and demand is medium or if seller two selects the *High* output and demand is high. In the perfect monitoring treatment, after choices are made, subjects are informed of the state of demand that prevailed and, hence, they can infer their opponent's choice from their payoff. In the imperfect public monitoring condition, subjects are not informed of the demand state. Given the structure, subjects can reconstruct what must be the prevailing prices (to rationalize profits), and those prices form an imperfect but public signal.[60] The experiment uses a between-subjects design, and four treatments are implemented: one payoff structure with δ of 0.6 or 0.75, and another payoff structure that is more conducive to cooperation, with δ equal to 0.75 or 0.9. The author shows that using Grim and assuming risk neutrality, cooperation can be supported in all four conditions with perfect monitoring, but only with the second payoff specification with imperfect monitoring. Across all

[58] We describe only two of the three treatments implemented here. The third treatment will be mentioned later in the survey.

[59] Unlike most studies, instructions are given *in context*, using terms such as sellers, market, demand, and output.

[60] We note that this mapping is non-trivial and that if subjects do not or cannot perform it, the environment may be perceived as one of private monitoring. Furthermore, although this monitoring structure is public, it has two features that are not standard. First, the public signal is not always common knowledge. Second, the support of the public signal varies with choices.

four treatments, cooperation rates vary between 0.21 and 0.45 under perfect monitoring and between 0.24 and 0.34 under imperfect monitoring. In fact, two treatments find a small increase in cooperation going from perfect to imperfect monitoring. Furthermore, the trend over supergames is either flat or slightly decreasing in three of the four treatments. Overall, the ability of subjects to support cooperation in this environment seems modest even with perfect monitoring, and, thus, the impact of imperfect monitoring seems unclear.

Result 8 *Many subjects cooperate in round 1 even with imperfect public information about their partner's behavior. However, in some of these environments, cooperation is difficult to sustain until the last round.*

It would be informative to study whether certain indices predict cooperation levels, as done for perfect monitoring environments in Section 2.3. Unfortunately, we cannot do that in this survey. First, there are a limited number of treatments with imperfect monitoring, which does not provide sufficient variation. Second, we cannot directly import the concept of the size of the basin of attraction of AD against Grim or TFT from the analysis of behavior under perfect monitoring since, as we will see next, subjects tend to use more-complex strategies under imperfect monitoring.

The literature has not only paid attention to the levels of cooperation reached under imperfect public monitoring, but has also studied the strategies used by subjects in these environments. There is a key difference from the case with perfect monitoring considered so far: under imperfect monitoring, agents must sometimes enter a punishment phase on the equilibrium path, even though they understand that nobody is "cheating"; otherwise, there would be no incentives to cooperate. This difference could be meaningful when it comes to analyzing strategies. In standard perfect monitoring PD games, it is much easier to identify a player's strategy when he or she is matched with other players that use varying strategies or use strategies that have random changes. This allows us to see how the player responds to these different situations. In an experiment, observing a variety of strategies that are somewhat *random* is more likely to happen early on. However, the behavior of interest is what takes place towards the end, when subjects have as much experience as possible. Unfortunately, as behavior stabilizes and subjects coordinate, variability is decreased, making it more difficult to observe the contingent play that would allow us to infer the strategies being used. This is not the case with imperfect monitoring since punishment phases are observed in equilibrium. This makes imperfect monitoring environments potentially more revealing about strategies. However, there is no reason to believe that the strategies used under imperfect monitoring are the same as those used under perfect monitoring. Hence, the study of strategies under imperfect monitoring should be done due to the importance of that environment but not to shed light on the strategies used under perfect monitoring.

Several articles have studied strategies under imperfect public monitoring. First, Aoyagi and Fréchette [2009] estimate that the most commonly used strategy is a two-state threshold strategy which is forgiving (it can come back to cooperation after a defection) and has threshold to move to a punishment phase that is more lenient than the one from the strategy that supports the maximum symmetric equilibrium payoff.

One important contribution of Fudenberg et al. [2012] is to identify that subjects tend to use strategies that are lenient (do not immediately trigger a punishment) or forgiving (come back to cooperation after a punishment phase) in the environment that they study.[61] These include versions of Grim that delay punishments until two or three consecutive bad signals have been observed. They also include variations on TFT that delay triggering a punishment or require more than one good signal to return to cooperation after a punishment has been triggered. This can be seen in Table 11, which summarizes information from perfect monitoring treatments already shown in Table 10 and includes data from studies with imperfect public monitoring.[62] The results

[61] Note that leniency and forgiveness in a finite signal environment must take a very different form from those in a continuous signal environment.

[62] The table separately gives the percentages for AD, Grim, and TFT, which were identified as the most common strategies under perfect monitoring. Note that TFT is a forgiving strategy. "> 2 states (C first)" consists only of lenient and forgiving strategies that start by cooperating. "2 states (C first)" includes some forgiving strategies, such as WSLS, and some that are neither lenient nor forgiving, such as the strategy that starts with cooperation, followed by defection forever.

for Fudenberg et al. [2012] are separated between the one treatment they have with perfect monitoring and the treatments with imperfect public monitoring. One can obtain almost all of the lenient or forgiving strategies that start by cooperating by adding the frequency of TFT and of strategies with more than two states that start by cooperating, abbreviated as "> 2 states (C first)."[63] The striking result from that paper can be observed by noting that TFT and "> 2 states (C first)" go from fitting 47 percent of their data when monitoring is perfect, to 74 percent of the data when monitoring is imperfect. Unlike Aoyagi and Fréchette [2009], Fudenberg et al. [2012]'s results points to many strategies with more than two states.

Looking at Table 11, we see that these papers (Dreber et al. [2008], Dal Bó and Fréchette [2011], Fréchette and Yuksel [2014], Rand et al. [2015], Embrey et al. [2013], and Aoyagi et al. [2015]), together, confirm the finding of Fudenberg et al. [2012]. That is, as monitoring becomes imperfect, a non-negligible fraction of subjects migrate towards forgiving or lenient strategies. Notably, the fraction of strategies accounted for by "> 2 states (C first)," which are all lenient or forgiving, tends to be higher under imperfect public monitoring than under perfect monitoring (51 percent versus 14.6 percent, averaged over studies). Note that, theoretically, there is no need for this change: the Grim strategy can support cooperation in all of these treatments. However, the Grim strategy leads to potentially non-negligible efficiency loss in noisy environments. Hence, the change in strategies could be to mitigate the efficiency cost of non-lenient/non-forgiving strategies in noisy environments. We also note that TFT, which is forgiving, does not seem to increase in popularity under imperfect public monitoring; if anything, it decreases in popularity. Hence, this movement towards leniency and forgiveness also accompanies a movement towards more complexity.

The finding of increased leniency and forgiveness does not seem constrained to discrete signals. Aoyagi and Fréchette [2009] report that as noise increases in their experiment, the median threshold to go from cooperation to defection

[63] This excludes two forgiving strategies, WSLS, and a strategy that punishes a bad signal by a single period of punishment and then returns to cooperation. With the exception of one case, neither of these strategies is popular.

Table 11: Percentage of Strategy Types in Perfect and Imperfect Public Monitoring

Monitoring:	Perfect					Imperfect Public		
	Dreber et al 2008[§]	Dal Bó and Fréchette 2011[§]	Fudenberg et al 2012	Fréchette and Yuksel 2014	Rand et al 2013	Fudenberg et al 2012	Embrey et al 2013[†]	Aoyagi et al 2014
AD	47	40	6	14	18	14	18	19
Grim	14	5	12	34	43	5	6	4
TFT	28	26	15	36	27	7	11	0
AD+Grim+TFT	89	71	33	84	88	26	35	23
Other 2 states (C first)			2	1		2	10	3
> 2 states (C first)	0	24	32	10	7	67	33	53
2 states (D first)	12	4	0	6	5	2	0	0
> 2 states (D first)			8	0		2	14	3
AC	0	1	24	0	0	2	12	20

Results based on SFEM estimation (restricted to studies that included a large set of strategies).

Includes only treatments in which joint cooperation can be supported in equilibrium.

Average over treatments (not weighted) if there are more than one treatment.

Entries left blank were not considered in those specific estimations.

[§] As reported in Fudenberg et al. (2012).

[†] Includes only stage games with two choices.

and from defection to cooperation decreases. In their case, this means that strategies are becoming more lenient (they require a worse signal to trigger a punishment) and more forgiving (they do not require as good a signal to return to cooperation). Therefore, in their setting, leniency and forgiveness are easily changed without changing the complexity of the strategy used.

Result 9 *Imperfect public monitoring results in subjects using more forgiving and lenient strategies relative to perfect monitoring.*

It is worth noting that the different implementations of imperfect public monitoring are not simply the result of the various authors' fancy. Different implementations serve different purposes. To be more specific, the implementation in Aoyagi and Fréchette [2009] allows them to relate payoffs to noise in the public signal, and the one in Embrey et al. [2013] makes it easier to relate choices to renegotiation proofness. On the other hand, if one wants to compare strategies across perfect and imperfect environments, the implementations in Fudenberg et al. [2012] and Embrey et al. [2013] seem preferable since the same set of strategies can be specified. Embrey et al. [2013] also allow an easy comparison with private monitoring, while Fudenberg et al. [2012] has a simpler implementation under perfect monitoring. Finally, Rojas [2012] can turn features on and off such that the environment is *à la* Green and Porter [1984] or Rotemberg and Saloner [1986].

5.2 Imperfect Private Monitoring

Many economic situations involve repeated interactions in which players neither know for sure what the other players did nor know exactly what the other players have observed. A classic example of such imperfect monitoring is found in Stigler [1964], in which he describes the difficulty of colluding in an oligopoly where firms post prices but can offer secret price discounts. Other firms do not observe these discounts. Instead, each firm receives a private signal of the others' behavior through its own sales–sales are an imperfect signal due to random shocks to demand.

Although examples in which private monitoring is present seem ubiquitous, theory has only recently made progress on this front. In fact, Kandori [2002] writes about what players can achieve in repeated games with private monitoring: "This is probably one of the best known long-standing open questions in economic theory." (p. 3) The absence of a commonly observed signal introduces a significant problem: how should agents use histories to coordinate the continuation play to provide the correct incentives, given that they do not know each other's history?

Finding testable implications of the theory of infinitely repeated games can be challenging, but the challenge is even greater with imperfect private monitoring. For instance, establishing whether a specific strategy profile can be supported in equilibrium is often difficult. General results, such as the highest symmetric equilibrium payoff, are still unknown. For these reasons, it is even more important for experiments on this topic to put behavior in relation to something else to make sense of it (cross-treatment comparative statics or a theoretical benchmark).

Without going into detail, we highlight some of the well known results from this literature (see Mailath and Samuelson [2006] for a survey). One important difference between private and public monitoring is that, under some assumptions on the details of the monitoring, no pure strategy profile that supports cooperation can be an equilibrium under imperfect private monitoring.[64] Hence, any equilibrium that supports cooperation will require mixing. Two main approaches are used: the belief-based and the belief-free approaches.[65] Under the belief-based approach, agents mix over simple strategies (for instance, Grim and AD) such that both players make each other indifferent between the two strategies and punishments are credible (see Sekiguchi [1997] and Bhaskar and Van Damme [2002]). Under the belief-free approach, the strategies are constructed such that in every round, subjects are indifferent between cooperation and defection (Piccione [2002] and Ely and Välimäki [2002]).

Not surprisingly, given the challenges highlighted above, evidence from experiments with imperfect private monitoring is scant, with only three papers to date exploring private monitoring with randomly terminated games. The

[64] The typical non-existence of pure strategy equilibria is true under (conditionally) independent signals. Pure strategy equilibria may exist in the case of almost-public monitoring.

[65] See, also, Compte and Postlewaite [2015] for an alternative *behavioral* approach.

first paper that falls in this category is Feinberg and Snyder [2002] which studies a modified PD with three choices (the third action gives both players the *sucker*'s payoff independent of the other's choice). Both sessions are composed of two supergames.[66] In the first supergame, which is played with perfect monitoring, there are eight rounds played for sure, and, after that, each additional round occurs with probability $\frac{2}{3}$. The second supergame has 18 rounds for sure, followed by the same random termination as the first supergame. In that second supergame, imperfect monitoring is introduced by occasionally manipulating the payoff numbers (giving the *sucker*'s payoff to both players), and this is known to subjects. In one treatment (one session), there is ex post revelation of those manipulations, while in the other there is not. In the second supergame, subjects are informed only of their own payoff and they do not know the payoff of the person they are playing with. When subjects are not informed of the payoff manipulations, this introduces private monitoring, since when subjects receive the *sucker*'s payoff, they do not know what their partner did or what signal their partner received. Feinberg and Snyder [2002] report that subjects cooperate 62 percent of the time in the first supergame with no shocks. In the second supergame, with ex-post revelation of payoff manipulations, subjects cooperate in 68 percent of their choices. However, when they are not informed of those manipulations, they cooperate at a rate of only 21 percent. Hence, private monitoring severely hampered the ability of subjects to cooperate in this environment. Note, however, that one difficulty in analyzing these results is that we do not know if cooperation can be supported theoretically. Hence, whether a drop of this magnitude should or should not be expected is unclear.

Matsushima and Toyama [2011] conduct an experiment with an environment closer to that in most of the standard theoretical work on private monitoring. Subjects play a PD and receive an imperfect and private signal of what the other player selected. Their payoffs are determined by the actual choices (in this aspect, it is closer to Aoyagi and Fréchette [2009] than other designs are). In all sessions, $g = \ell = 0.\overline{2}$ and $\delta = 0.967$, making it the treatment with the highest δ conducted so far. The treatment variable is the accuracy of

[66] In total, only 16 subjects participated in this experiment.

the signal; the probability that the signal corresponds to the choice is either 0.9 (high accuracy) or 0.6 (low accuracy). The design is a mix of within- and between-subjects design: two sessions of three supergames with high accuracy, followed by three supergames with low accuracy and two sessions with the reverse order. The authors use two elements to understand the data: the comparison of behavior across the two treatments and the comparison of the data with the implications of a specific equilibrium (symmetric generous tit-for-tat, which is the best symmetric belief-free equilibrium with memory-one). The symmetric generous tit-for-tat (SGTFT) cooperates with some probability in each round. The specific probability is determined by the round (round 1 versus all other rounds) and the signal of the other's choice in the previous round (for rounds other than 1). If subjects select this equilibrium in both treatments, then an implication of the theory is that in the low accuracy treatment, the difference in cooperation rates between the case in which the last signal was *good* versus when it was *bad* should be greater than in the high accuracy treatment. The intuition is that as the monitoring becomes worse, it is more difficult to incentivize agents, and, thus, a more intense reaction is required. The key results are the following: 1) cooperation rates are high in both treatments; 2) cooperation rates are higher when accuracy is higher; 3) cooperation rates are lower than predicted by the SGTFT; and 4) the intensity of the reaction to the last round's signal is *higher* in the high accuracy treatment. What we have established in the previous sections can help us interpret these results. The payoff parameters used in this study are very conducive to cooperation. Using the results of the meta study to predict cooperation rates in a one-shot game with those payoffs over the course of six supergames, we find a predicted average of 43.1 percent.[67] The round 1 cooperation rate in the data for the low accuracy treatment is 43.8 percent. This indirectly indicates that subjects are actually not cooperating more in that treatment than they would in a one-shot game. Relatedly, the intensity of reaction in that treatment is only 0.165, which, although it is not reported, may not be statistically different from zero. On the other hand, in the high accuracy treatment, cooperation in round one averages 78.1 percent and the intensity of reaction is 0.508. This

[67]The paper does not report what happens over time in different supergames.

certainly is in line with subjects supporting cooperation using some strategies that respond to the private signal. It also suggests that the strategies played are different in both conditions, and, thus, interpretations about reciprocity deduced from the comparison to the theoretical implications of SGTFT may be misleading. Furthermore, as we saw in Section 5.1 (and first proposed by Fudenberg et al. [2012]), when moving to imperfect public monitoring, many of the strategies used depend on more than just the signal from the previous round. If this finding extends to private monitoring, an analysis that relies exclusively on memory-one strategies would give the wrong benchmark.

Aoyagi et al. [2015] use a different approach to interpret behavior under imperfect private monitoring. Instead of comparing behavior to a particular equilibrium, they compare behavior across various monitoring environments: perfect, imperfect public, and imperfect private. In order to do this, they use the following design. Under imperfect monitoring, both public and private, subjects choose an action in a PD, and with some probability, the signal of their choice that is communicated to the other player is incorrect. Their payoff is determined by their own choice and the signal they receive of what the other player did. The difference between the public and private treatments is that, under public monitoring, both players know not only their choice and the signal of what the other player chose, as under private monitoring, but also the signal the other player received of their own choice. In the perfect monitoring treatment, subjects know exactly the choice of both players. However, the payoffs are determined by lotteries equivalent to the lottery faced by a subject under imperfect monitoring.[68] Hence, this design keeps the strategic form of the game fixed across treatments–expected payoffs as a function of actions are the same for perfect, imperfect public, and imperfect private monitoring.[69] The parameters are: $g = \ell \approx 0.357$, $\delta = 0.9$, and the probabil-

[68]Under imperfect monitoring, for example, if two players cooperate, they each face a lottery that with some probability they receive the *reward* payoff, unless the signal they receive of the other's action is incorrect (which happens with the complementary probability), in which case they receive the *sucker*'s payoff.

[69] Note that behavior in the perfect monitoring treatment suggests that the outcome of the lottery does not have an important effect on choices, and, if anything, the effect diminishes with experience. Rand et al. [2015] directly studies this question comparing behavior in an

ity that the signal is incorrect is 0.1.[70] Under these parameters, theoretically, (some) cooperation can be supported in all three treatments. The key results with respect to private monitoring are the following. Cooperation can be supported under private monitoring, and this occurs at rates very similar to, albeit slightly lower than, those under perfect and public monitoring. The rates of round 1 cooperation in the last four supergames are 65 percent and 73 percent for perfect and public, and 61% for private. This confirms the result of Matsushima and Toyama [2011] that cooperation can be supported under private monitoring. However, they extend their findings to a different environment and establish that this can be done at rates similar to those under other monitoring structures. In addition, their paper shows that, in line with the intuition from the theory, coordination (of actions) is higher under perfect and public monitoring than it is under private monitoring. Finally, as already mentioned, Fudenberg et al. [2012] establish the popularity of lenient and forgiving strategies under public monitoring, and we confirm that across multiple studies, subjects use more lenient and forgiving strategies to support cooperation under public monitoring than they do under perfect monitoring. A related movement towards more complex strategies is observed here under private monitoring, as well. However, what is new is the observation that under private monitoring, strategies are as lenient as they are under public monitoring, but in terms of forgiveness, they are similar to what is observed under perfect monitoring. In particular, lenient forms of Grim that wait for two or three consecutive bad signals before triggering a punishment phase are significantly more popular under private monitoring than under either perfect or public monitoring.

While more research is needed on imperfect private monitoring, the existing work leads to the following conclusion.

imperfect public monitoring treatment (with the same structure as Fudenberg et al. [2012]) to a similar treatment in which both the implemented choice and the intended choice are revealed to both players (hence, this treatment has perfect monitoring). Although they document some differences between the two treatments, the results, overall, point to most subjects reacting to intentions, not to outcomes.

[70]Depending on the session (there are four sessions per treatment), as few as nine and as many as 22 supergames are completed.

Result 10 *Cooperation is possible under imperfect private monitoring, and subjects seem to move towards adopting more lenient strategies.*

6 Community Enforcement

Previously, we considered situations in which the same set of players interacts repeatedly. However, in many economically relevant situations, the same set of players do not always interact together but switch partners over time. For example, consider how the matching of consumers to firms or the compositions of teams in firms may change from time to time. Moreover, many market and other types of interactions are characterized by free-riding opportunities arising from limited contracting or asymmetric information. Overcoming these opportunities is essential for markets and societies to work efficiently. While this may seem impossible when the changing of partners eliminates the possibility for personal retaliation, a fascinating theoretical literature has shown that cooperation may still be a possible equilibrium outcome (see Kandori [1992], Greif [1993], Ellison [1994] and Okuno-Fujiwara and Postlewaite [1995]).

Opportunistic behavior can be overcome even with changing partners, as groups may follow social norms in which opportunistic behavior is punished by members of the group, even by those who did not suffer the cost of that opportunistic behavior. Does community enforcement lead to cooperation in repeated games with changing partners? A series of recent papers have addressed this question under a variety of environments.

One of the environments studied in the experimental literature on community enforcement is random matching with anonymous players and personal histories (Schwartz et al. [2000], Duffy and Ochs [2009] and Camera and Casari [2009]). Anonymity refers to the fact that players do not learn the identity of their partners. Thus, it is impossible to identify a defector and specifically retaliate against him or her (there is no possibility of establishing a relationship based on personal enforcement, as in the previous sections). Personal histories refer to the fact that subjects observe what their partners do but have no information about what happened in other interactions. Contrary to what one might initially think, Kandori [1992] shows that it may still be possible to support cooperation through community enforcement when the group is small enough (the size of the group that can allow for cooperation in equilibrium depends on the PD payoffs parameters and discount factor). While players neither observe what happened in other pairs nor know the identity of their partners, they can follow a "contagion" strategy that reverts to defection after observing a defection. If players follow this strategy, an initial defection will result in a community-wide collapse of cooperation, and the initial defector will be punished even when no member of the society can identify who the initial defector was.

Duffy and Ochs [2009] study behavior in infinitely repeated games with anonymous random matching and personal histories with groups of six and 14 subjects. Subjects are randomly matched in pairs in each round to play a prisoners' dilemma game with $g = \ell = 1$ and probability of continuation $\delta = 0.9$. In these treatments, cooperation can be supported in equilibrium following Kandori [1992]. They also run sessions in which subjects are matched in fixed pairs for the whole supergame.[71]

Figure 7 shows the evolution of first-round cooperation levels by supergame for both treatments with random matching (groups of six and groups of 14) and for the fixed pairs treatment. In addition, we include as benchmarks the cooperation levels that we would expect for fixed pairs for $\delta = 0.9$ and for one-shot games (this is done as in Table 5 in Section 2).

The evolution of cooperation under fixed pairs is consistent with what one might expected based on the analysis in Section 2. Cooperation increases slightly with experience. For the treatments with random matching, however, cooperation does what one would expect in one-shot games. After even just a little experience, and regardless of the size of the group, random matching leads to levels of cooperation that are quite different from those from repeated interaction under fixed pairs. In conclusion, the results from Duffy and Ochs [2009] suggest that subjects may not use the opportunity to cooperate that arises from the contagion equilibrium identified in Kandori [1992].[72]

[71] This is the treatment that is included in the data described in Table 3.

[72] This result is even more surprising once an element of the experimental design in

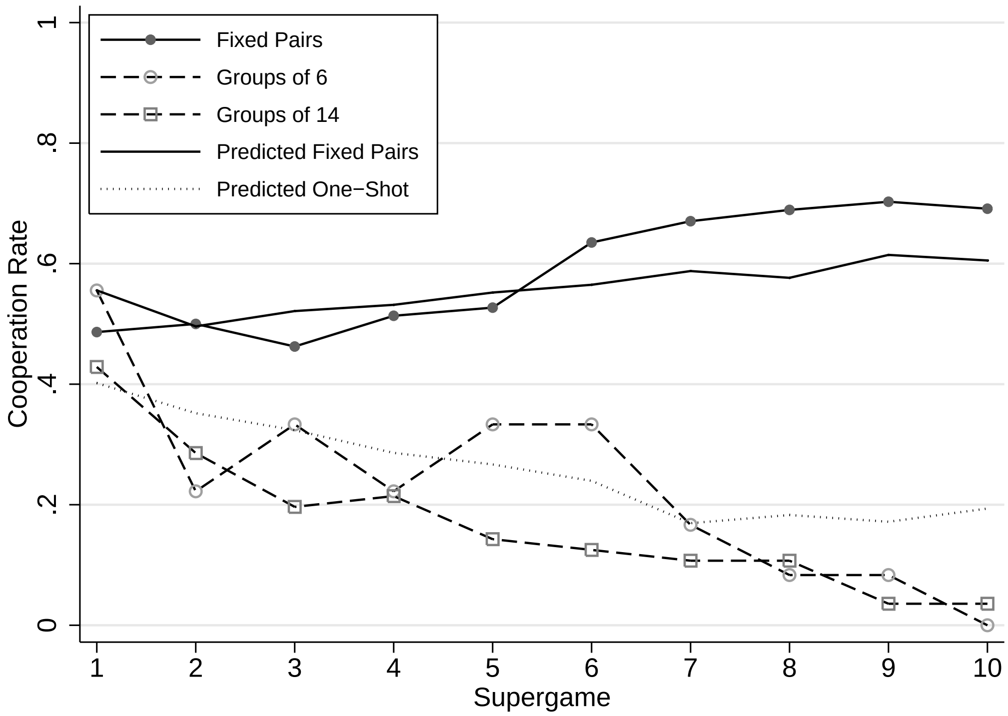

Figure 7: Round 1 Cooperation in Anonymous Random Matching Games with Personal Histories – Data from Duffy and Ochs [2009]

Camera and Casari [2009] study behavior in infinitely repeated games with anonymous random matching and personal histories with groups of four and 14 subjects. In each period, subjects are randomly matched in pairs to play a prisoners' dilemma game with $g = \ell = 1/3$ and probability of continuation $\delta = 0.95$. Figure 8 shows first-round cooperation levels by supergame for both

Duffy and Ochs [2009] is taken into consideration. Each session consisted of six or 14 subjects, depending on the size of the group; in other words, all the subjects in one session play together in all supergames. As such, after the end of a supergame with random matching, a new supergame with random matching would start with the same set of players. Therefore, the likelihood of future interaction is greater than the intended probability of continuation. A similar caveat applies to Schwartz et al. [2000], who study cooperation under anonymous random matching in a modified prisoners' dilemma. They also find very low levels of cooperation when subjects have no information about the past behavior of their partners.

treatments (groups of four and groups of 14). We also include as benchmarks the cooperation levels that we would expect for fixed pairs for $\delta = 0.95$ and for one-shot games (this is done as in Table 5 in Section 2).

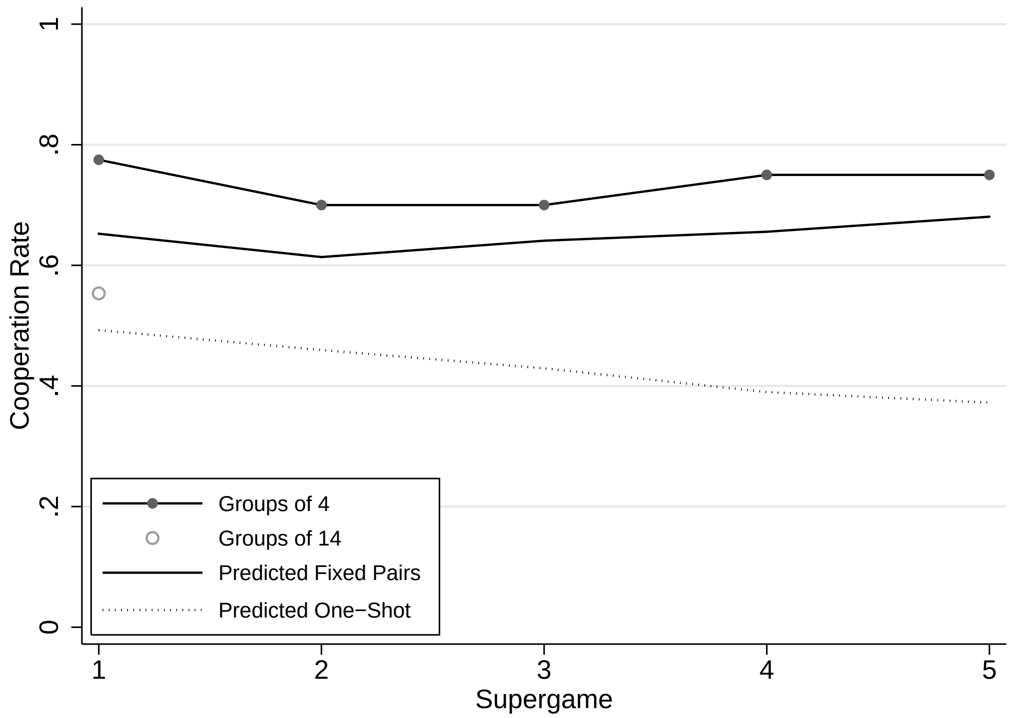

Figure 8: Round 1 Cooperation in Anonymous Random Matching Games with Personal Histories – data from Camera and Casari [2009]

Cooperation is stable across supergames for groups of four and somewhat larger than what could be expected in repeated games with fixed pairs. Data from only one supergame are available for groups of 14. Cooperation levels are smaller than for groups of four and closer to one-shot levels. The result from groups of four leads Camera and Casari [2009] to conclude that cooperation can be sustained even in anonymous settings with personal histories.[73] Note

[73] Bigoni et al. [2013] present data from a similar experiment using university staff instead of students as a sample. They find cooperation rates for staff members that are smaller than for students but still greater than could be expected from one-shot games.

that even though the cooperation rate is high with random matching, the PD payoffs are such that even in a one-shot game, one would expect a relatively high cooperation rate, as shown in Figure 8. Thus, much of the cooperation with random matching may not necessarily pertain to community enforcement. In fact, Camera et al. [2012] estimate a modified version of the SFEM and report that 34 percent of strategies (by far the most popular strategy in their data) correspond to AC. This is much higher than anything reported in Table 10 for perfect monitoring in fixed pairs and is not a strategy that implements community enforcement. They do find, however, that 34 percent of strategies correspond to either Grim or TFT.

In conclusion, while more research needs to be done for different combination of parameters (g, ℓ, δ) and group sizes, the current evidence suggests that cooperation will not arise under anonymous random matching with personal histories unless the group size is small.

Result 11 *Cooperation is unlikely to arise in groups with anonymous random matching unless payoffs are conducive to cooperation and the group is small.*

This result somewhat limits the applicability of contagion equilibrium, as it is unlikely that very small groups will face anonymous interactions.

Can community enforcement lead to more cooperation in less information-starved environments? Several papers have looked at how information can help community enforcement.

Schwartz et al. [2000] study cooperation in an infinitely repeated modified prisoners' dilemma game with random matching and find that providing information about the past behavior of the partners significantly increases cooperation levels. The reason is that subjects are more likely to cooperate if their partner has cooperated before, providing incentives to cooperate.

Duffy and Ochs [2009] also study the impact of information on cooperation in random matching games. They consider two information treatments in addition to the treatment with personal histories discussed above. In one of these treatments, subjects are informed of the average payoff from the current partner's previous match. In another treatment, subjects are informed of the last action chosen by the current partner. They find that cooperation is greater with information, but the difference with the treatment with personal histories is significant only for the average payoff treatment. They conclude that the type of matching is more important than the available information in determining cooperation.

Camera and Casari [2009] also present experimental results from treatments with varying information in addition to their personal histories treatment discussed above. They have two information treatments. Both treatments involve subjects observing all the actions taken by the four members of the group: in one treatment, identities are not revealed, while in the other treatment, the subjects know who did what. We call these two treatments "anonymous public histories" and "perfect information."

Figure 9 shows the evolution of round 1 cooperation for these three treatments. While in the first supergame, there are no significant differences across the three treatments, by the last supergame, cooperation is significantly greater under perfect information than under anonymous private monitoring and anonymous public monitoring. In summary, being able to identify who defected and who cooperated seems more important to support cooperation than simply knowing the distribution of actions without identities.

Duffy et al. [2013] study the impact of providing information on past behavior in random matching infinitely repeated trust games. They find that providing information significantly increases trust and trustworthiness, consistent with the previously discussed evidence from prisoners' dilemma games.

Finally, the treatment with perfect information in Camera and Casari [2009] provides an opportunity to study whether subjects rely on community enforcement or personal enforcement to support cooperation. In that treatment, subjects could support cooperation in equilibrium by relying on personal enforcement alone: since there is no anonymity, subjects could consider the interactions with each other subject as a different repeated game and use some retaliatory strategy, such as Grim, in those particular interactions. Subjects could, instead, rely on some type of community enforcement. For example, subjects could stop cooperating if any member of the group had defected before (all group members are punished). Alternatively, subjects could stop

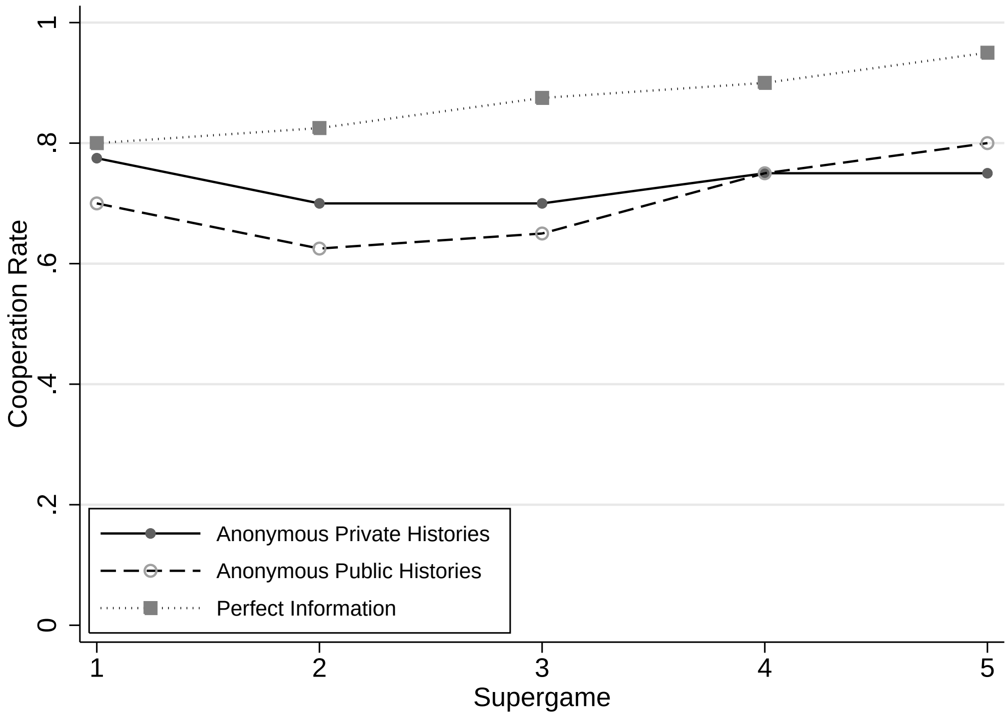

Figure 9: Round 1 Cooperation by Information Treatment – data from Camera and Casari [2009]

cooperating with subjects who had deviated before (only deviators are punished).

To help us understand the type of enforcement used by subjects, Table **??** shows the percentage of cooperation in round 2 as a function of whether the partner in round 2 is the same as in round 1, and as a function of the action in round 1 of the partner in round 2. We focus on round 2 behavior since others' later behavior is likely to be affected by the subject's own behavior, making comparisons difficult to interpret.

When the partner in round 2 is different from the partner in round 1, subjects are more likely to cooperate in round 2 with a partner who has cooperated in round 1 (88.79 percent against 61.90 percent). This suggests that subjects do not consider the relationships in isolation, but use information on

Table 12: Community vs. Personal Enforcement under Perfect Information in Camera and Casari [2009] – Round 2 Cooperation

	Action in round 1 of partner in round 2	
	D	C
Different Round 1 and 2 Partner	61.90	88.79
Same Round 1 and 2 Partner	40.00	95.52

what others have done to others when deciding what to do.[74] Note that the difference in behavior is greater when the partner in round 1 and 2 is the same subject (95.52 percent against 40 percent). Hence, subjects do not rely solely on community enforcement but use personal enforcement, as well. Subjects respond more to a defection that they themselves have suffered than one suffered by another player. This can help explain the greater levels of cooperation observed under perfect information relative to public anonymous monitoring.

Result 12 *In small groups, subjects can use community enforcement to support cooperation, and this is facilitated by providing information about past behavior. In particular, information that allows subjects to punish deviators facilitates cooperation.*

More research is needed to understand the robustness of community enforcement in larger groups and the working of alternative institutions that may provide information. For example, to what degree would it help to have mechanisms that rely on players to voluntarily report their past experience?

[74]Kusakawa et al. [2012] also provide evidence consistent with community enforcement. They consider an environment in which players one and two play a one-shot prisoners' dilemma, and then one of the two players is randomly chosen to play an infinitely repeated PD game with player three. They find that player three is significantly more likely to cooperate with a player who cooperated in the one-shot game. Expecting this, players one and two are more likely to cooperate in the one-shot game when player three observes the behavior from the one-shot game relative to a treatment in which player three does not observe that behavior. This line of research should be continued so that we can better understand the power of community enforcement outside random matching games.

Who would report and how would subjects use these reports?

7 Other stage games and dynamic games

The literature has studied how the "shadow of the future" can limit opportunistic behavior in other environments besides repeated PDs. We have already mentioned in passing some of the other games that have been studied. Closest to the PD, there are papers with a PD structure but with more than two choices. These include Holt [1985], in which subjects play a Cournot duopoly game where they can select one of 18 output quantities, as well as Schwartz et al. [2000] (with random-rematching), Dreber et al. [2008], and Embrey et al. [2013], which use games with a PD structure that offer a third choice. Palfrey and Rosenthal [1994] study the voluntary contribution mechanism. Although these games are all different from the basic PD, many of the results overlap. These papers find some distinct results, including the following: 1) The availability of a costly punishment option increases overall cooperation rates. However, on average, payoffs are not higher in such a situation. Hence, when punishments are available, those who do not use them and cooperate do best. 2) A "somewhat" cooperative option (which is predicted to be selected by renegotiation proofness in this context with imperfect public monitoring) is not popular–subjects in that experiment prefer to either defect or cooperate at the highest possible level.

Other papers study different stage games. Engle-Warnick and Slonim [2004], Engle-Warnick and Slonim [2006a], and Engle-Warnick and Slonim [2006b] study repeated trust games, as discussed in Section 2. Bernard et al. [2014] study an infinitely repeated gift-exchange game. Their parameters are selected such that the returns to sending a gift are modest (as opposed to the typical gift-exchange experiments). With such parameters, it has been shown that gift-exchange is much lower in a typical experiment without random termination. In their setting, gift-exchange can be an equilibrium phenomenon, and, indeed, they observe gift-exchange, despite the lower than usual potential returns to gift giving. Cason and Mui [2014] study a three-players game called the Divide-and-conquer Coordinated Resistance Game. In that game, a *Leader* can attempt a transgression against one or both of the *Responders*, who can either acquiesce or challenge. A successful transgression against a responder takes money away from him. Part of that money is destroyed, the rest goes to the leader if he transgresses against both responders, or is shared (with more going to the leader) if the transgression is towards only one responder. Challenging (by the responders) is costly and is successful only if both responders do it. The experiment explores the effect of repetition and communication on rates of transgression. This setting is quite different from experiments such as the PD because the one-shot game has multiple equilibria (and so does the finitely repeated game). Their findings indicate that repetitions decrease the rate of transgressions, but that this is equally true for finitely repeated and randomly terminated games. Finally, Cabral et al. [2014] study a veto game, as discussed in Section 4.

While there is still much to be understood using the PD, clearly there is enormous scope for work exploring different stage games with random termination to better understand how repetition may affect behavior and to connect this literature with particular applications. For example, the aforementioned paper by Bernard et al. [2014] studies repeated gift-exchange to answer questions related to the labor market.

Another direction of study does not differ in the stage game but in the fact that the stage game changes over time. Infinitely repeated games are only a small subset of dynamic games, and research has started to explore situations in which choices can be conditioned on a state. These fall in two categories: those where the transitions between stage games are determined exogenously and those where they are a function of the agents' choices.

Rojas [2012] experimentally investigates the ability of subjects to support collusion in an environment à la Rotemberg and Saloner [1986]. Every round demand can be low, medium, or high, and subjects know the demand state when they choose their quantity (low or high). The design is such that both main parameter specifications support an equilibrium in which the temptation to deviate is too important, and equilibrium behavior dictates that agents must produce the high quantity when demand is high but can collude for other demand levels. The data indicate that subjects can, indeed, cooperate (collude) in such an environment. In addition, as predicted, subjects are less likely to cooperate (collude) when they learn that demand will be high. When demand returns to a lower level, subjects are more likely to cooperate. This clearly indicates that subjects react in sensible ways to dynamic incentives even in complex settings. Exogenous transitions across states do not destroy the subject's ability to cooperate.

Kloosterman [2015b] also investigates situations in which the transitions between games is exogenous, but his focus is on the implications of beliefs and signals about the future. In a first set of experiments that uses two stage-game PDs, subjects first play one of the two games and then, with a certain probability, they play either the same PD or a different one in all subsequent periods. The design varies which PD a subject first plays and the probability that the PD changes or stays the same for the rest of the supergame. The parameters are such that cooperation in the first choice should depend on the probability of playing, in the future, the PD that is more favorable to cooperation. Indeed, the data support this idea.

The other set of experiments tests how subjects react to imperfect signals about future conditions. In this case, the stage game is an asymmetric partnership game in which the *project* is successful if and only if both subjects exert effort. The asymmetry arises because only one of the two subjects will receive the benefit of the project. Who the beneficiary is in the current period is known, but the signal gives (imperfect) information about the beneficiary in the next round. The experiment varies the precision of the signal and the payoffs to test Kloosterman [2015a]'s surprising prediction that increasing the precision of information about future conditions, in this case, makes cooperation more difficult to support.[75] Much of the data are in line with implications of the theory (reaction to signals, reaction to payoffs, etc.). However, cooperation is higher, not lower, when the signal about future conditions is more

[75]The person who will be the beneficiary in the next period has greater incentives to cooperate today. Therefore, increasing the precision of the signal has two distinct effects on the two players' incentives to cooperate: the expected incentives to cooperate increase for the beneficiary in the next period, while it decreases for the other player. Kloosterman [2015a] shows, that under mild conditions, the negative impact of a better signal overcomes the positive impact, and cooperation is harder to support.

precise. The explanation given for this discrepancy is based on the size of the basin of attractions and, thus, suggests that the earlier results on the role of the basin of attractions in the infinitely repeated PD may extend beyond the simple deterministic and stationary environment.

A recent sequence of papers explores the extent to which behavior in relatively complex dynamic environments, where transitions depend on past choices, is consistent with the prediction of Markov perfect equilibria. Battaglini et al. [2012] explore a repeated legislative bargaining game where, in every round, an endowment can either be invested in a durable public good or consumed (pork). The experiment studies the impact of the voting rule: unanimity, majority, or dictatorship. In theory, a higher voting requirement generates more efficiency. Indeed, the data support this view, and there is more public good as more votes are required for support. Battaglini and Palfrey [2012] also explore a repeated legislative bargaining game without public good. The state variable in this case is the status quo (which takes effect if the proposal on the floor is rejected), which is the last approved proposal. The authors conduct multiple treatment manipulations: concavity of the utility function; the availability of a Condorcet alternative; and the discount factor. The treatment effects are in the direction suggested by the Markov perfect equilibrium. Finally, Battaglini et al. [2015] study a repeated voluntary contribution mechanism in which the public good is durable. They vary whether or not investments are reversible and the number of subjects in the group. The Markov perfect equilibrium has higher levels of public good when investments are irreversible. The data are consistent with that prediction. On the other hand, there is a tendency for over-accumulating the public good early on, a trend that can be corrected later in the reversible case, but not in the irreversible case.

These last three papers show many comparative statics consistent with the predictions of Markov perfect equilibrium. However, in a simple infinitely repeated PD, the unique Markov perfect equilibrium involves defection in every round, which, we have shown, is not what happens in such a game. Hence, a natural question seems to be: when is the Markov assumption a reasonably good predictor of behavior? Or, perhaps: is there evidence that subjects play Markov perfect strategies? To investigate such questions, Vespa [2015] studies a simple version of the dynamic commons problem. In each round, subjects can extract resources from a common pool that is replenished at a given rate. Cooperation involves extracting less than is optimal in a one-shot game to let the resource grow. Vespa [2015] designs an environment that allows him to identify not only the Markov perfect equilibrium payoff, but also the fully efficient symmetric equilibrium payoff.[76] The main finding is that modal behavior is indeed Markov. However, Vespa [2015] does find that the popularity of Markov strategies decreases as the payoff of the efficient outcome increases.

Vespa and Wilson [2015] also explore equilibrium selection in dynamic games by taking incremental deviations from a simple dynamic game in which they add a state to the PD. This simple game allows them to turn certain features of the environment "on" and "off": varying the temptation to defect in one state; changing the transition rule between states; removing specific types of externalities; etc. The paper documents many new results, but the main finding is that, in these simple dynamic games, subjects often condition choices on more than the state using trigger types of strategies. Finally, they propose an index that predicts when outcomes more efficient than the Markov perfect equilibrium are to be expected.

In many situations, agents who do not find a way to cooperate in a relationship will prefer to exit the relationship altogether instead of forever defecting. Wilson and Wu [2014] study how such a possibility affects choices and what the impact of the outside value is on behavior. This offers a new and interesting window on the potential role of contracts and the legal system, as those determine what happens when relations are severed outside of the laboratory. The experiment involves many treatments, all based on the partnership game; in some the outside value is fixed, and in others, the outside value depends on the behavior of the agent while in the relationship. As theory predicts, the

[76] The focus on Markov perfect equilibria in dynamic games is often a matter of convenience, as the environments are too complex to solve for other equilibria. Hence, once must be careful when making the jump from observing comparative statics in line with the Markov prediction to concluding that subjects use Markov strategies. Designing a setting in which we can solve for other equilibria gives more force to the conclusions, either positive or negative.

authors observe much more termination once the outside option exceeds the in-relation minmax. The possibility of termination seems to increase the use of strategies that start with cooperation. Also, punishment is more likely to take the form of termination (rather than simple defection). However, subjects become more lenient: they need more bad signals before triggering a punishment. One rule stands out among the asymmetric termination rules: efficiency is highest when the subject that cooperated the most in the relationship is treated more favorably in case of termination.

Another study that looks into the role of termination is that of Honhon and Hyndman [2015]. Unlike Wilson and Wu [2014], however, terminating a relationship leads the subjects to continue the supergame with someone new. In addition, the paper also investigates the impact of various reputation mechanisms in the case where relations can be terminated. They find that the option to terminate a relationship reduces cooperation but that this can be remedied by certain types of reputation mechanisms. In particular, a reputation mechanism that conveys objective information and carries over across supergames is most effective.

These papers all point in different but extremely interesting directions for future research. Moving toward dynamic games opens the door to many applications of interest. However, the frequent requirement to focus on Markov perfect equilibria means that a better understanding of when such an assumption is reasonable is important. Finding concepts and regularities that connect these more complicated environments back to the simpler case of the PD will probably yield important rewards moving forward.

There are other types of dynamic games that, although distinct in the way they are analyzed, seem to be naturally connected, and we think that future research exploring the similarities, connections, and differences will be useful. In particular, the recent studies of games in continuous time would seem related. Although the theoretical tools to analyze them are different, they, too, generate cooperation in social dilemmas where no such cooperation is predicted in the one-shot game. Two recent experiments on games in continuous time are Friedman and Oprea [2012] and Bigoni et al. [2013]. Both show that subjects can support cooperation under continuous time. However, surprisingly, in continuous time, Bigoni et al. [2013] show that cooperation rates can be lower with a randomly determined horizon as opposed to a deterministic horizon. Understanding the behavioral sources of any differences or similitudes could prove enlightening for theoretical work, through a better understanding of how strategy choices are affected by the details of the strategic environment.

8 Conclusions

The theory of infinitely repeated games has shown that repetition can help people overcome opportunistic behavior and lead to cooperation. While this idea has been applied in many fields of economics and other social sciences, a large experimental literature has only recently considered its validity. In addition to surveying the existing experimental literature, we gather and analyze available experimental data in order to focus on robust results.

We observe that the "shadow of the future" can lead to cooperation and that theory help us understand the conditions under which this will occur. Moreover, the experimental evidence provides a solution to the multiplicity of predictions provided by theory. We find that people do not necessarily coordinate on the best equilibrium. Cooperation is low when it is not robust to strategic uncertainty and increases with robustness to strategic uncertainty. We observe high levels of cooperation only when the parameters are such that attempting to cooperate is not too risky. However, as discussed in Section 2, more work is needed to find the index that best predicts cooperation.

We also see that subjects tend to use simple strategies to punish defection under perfect monitoring. While the fact that subjects use punishments to support cooperation is consistent with theory, one of the most-used strategies (TFT) involves non-credible punishments. Thus, this raises the question of why subjects use it instead of other strategies with better theoretical properties. More work is also needed to understand how the use of strategies changes with the parameters of the game.

We find that the "shadow of the future" can lead to cooperation under imperfect monitoring (both public and private). Much more work is needed to better understand cooperation under imperfect monitoring (especially under private monitoring) and its robustness to the specifics of the environment. In particular, the literature should pay attention to how details of the monitoring technology affect the capacity to coordinate on cooperation and the types of strategies used. For imperfect public monitoring, we see that subjects are more likely to use lenient and forgiving strategies relative to perfect monitoring. More work is needed on the strategies that subjects use under imperfect private monitoring and the connection to those used in theoretical work.

In addition, we see that subjects are able to support cooperation when partners change from period to period. However, so far, this seems true only in very small groups. It seems fair to say that it is much more difficult for subjects to support cooperation when they are randomly rematched than when they play in fixed pairs. In fact, it seems easier for subjects to cooperate even if monitoring is imperfect, but they are in fixed pairs, than if they are randomly rematched. Still more work is needed to understand the strategies subjects use under perfect monitoring with random rematching (to better understand the relative importance of personal and community enforcement) and how other information institutions (such as reputational mechanisms) may help subjects support cooperation.

Communication may be another important determinant of cooperation. As Cooper and Kühn [2014a], Andersson and Wengström [2012], Embrey et al. [2013] and Cooper and Kühn [2014b] discuss, communication could, in principle, have different effects on cooperation.[77] On the one hand, communication before the supergame could help subjects coordinate on cooperation, increasing cooperation rates. On the other hand, communication after a defection could help renegotiation, and the expectation of this could result in lower cooperation rates. Cooper and Kühn [2014a] study the effects of several types of communication on cooperation in a game in which subjects play a social

[77] Others have explored the role of communication. Cason and Mui [2014] studies the effect of communication in repeated divide-and-conquer coordinated resistance game. Bigoni et al. [2014] and Evdokimov and Rahman [2014] study infinitely repeated PDs with imperfect public monitoring in which they allow for delay in choices or feedback, and they explore how the effect of delay is affected by communication. In both papers communication has an important (positive) impact on cooperation.

dilemma first and then a coordination game.[78] Among other results, they find that free-form communication (chats) leads to persistent increases in cooperation, while other forms of communication do not. Although the simplified dynamic game facilitates the study of communication, given the caveats mentioned in Section 3 about this way of inducing infinitely repeated games (finite repetition followed by a coordination game), it seems appropriate to explore the generalizability of the results from Cooper and Kühn [2014a] using other ways to induce infinitely repeated games in the laboratory. More generally, it is important to study the effects of communication for different combinations of parameters in order to replicate the analysis of the determinants of cooperation under perfect monitoring that we offered in Section 2. Will we find high levels of cooperation even when cooperation is not risk dominant? Could it be that once we allow for communication, subjects do coordinate on the Pareto efficient equilibrium? Will subjects use different types of punishments?

While cooperation in repeated games may be affected by elements of the game–such as payoff parameters, probability of continuation, available information, matching or communication–it may also be affected by elements from outside the game. For example, the literature has looked at the role of personal characteristics in explaining differences in behavior but has found no robust relationship. We believe that we should keep looking for personal characteristics that may help explain behavior in repeated games. Another element from outside the game that may affect behavior is history. As we have seen, the experience in a given treatment affects behavior. Similarly, experience in other treatments or other games may affect behavior, as well.[79] We should study whether history is used as an equilibrium selection device in infinitely repeated games. The experiences that subjects had before entering the supergame and the ways in which subjects arrived at a supergame may affect cooperation. Similarly, local norms of cooperation may affect the expectations of behavior that subjects bring to the laboratory. This suggests that a comparative study

[78] The coordination game can be derived from the continuation payoffs that would arise from an infinitely repeated game in which deviations would be punished forever by reversion to the one-shot Nash equilibrium.

[79] Footnote 29 in Dal Bó [2005] presents evidence consistent with this idea.

of behavior across cultures could be of interest.

Given that the formation and termination of groups is usually endogenous, following Bernard et al. [2014] and Wilson and Wu [2014], research could further examine how these realistic features affect cooperation and its determinants.

While, in this article, we have focused on experiments that use the PD as the stage game, we have also mentioned many articles using other stage games (other social dilemmas, trust games, veto games, etc.). To test the robustness of the results, additional experiments with other stage games would be welcome. For example, it would be interesting to know how repetition affects contributions in linear public-good games with a continuum or high number of contributions levels. This could allow for different strategies that may reduce the risk from strategic uncertainty and support cooperation (such as starting small).

Similarly, future articles on dynamic games may shed light on the types of punishments used when actions affect the state. When do subjects coordinate on Markov strategies? When do they rely on punishments that also depend on elements of the history that are not captured by the state? Can the elicitation of strategies help us answer these questions?

Much progress has been made in understanding the determinants of cooperation in infinitely repeated games. Many questions remain to be answered, and we hope that this article provides a clear picture of the state of the literature and suggests productive lines of future research.

References

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58:1041–1063, 1990.

Ola Andersson and Erik Wengström. Credible communication and cooperation: Experimental evidence from multi-stage games. *Journal of Economic Behavior & Organization*, 81(1):207–219, 2012.

Masaki Aoyagi and Guillaume R. Fréchette. Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic Theory*, 144 (3):1135–1165, May 2009.

Masaki Aoyagi, V. Bhaskar, and Guillaume R. Fréchette. The impact of monitoring in infinitely repeated games: Perfect, public, and private. Working Paper, 2015.

Robert Axelrod. Effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(1):3–25, 1980a.

Robert Axelrod. More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(3):379–403, 1980b.

Marco Battaglini and Thomas R. Palfrey. The dynamics of distributive politics. *Economic Theory*, 49(3):739–777, 2012.

Marco Battaglini, Salvatore Nunnari, and Thomas R. Palfrey. Legislative bargaining and the dynamics of public investment. *American Political Science Review*, 106(02):407–429, 2012.

Marco Battaglini, Salvatore Nunnari, and Thomas R. Palfrey. The dynamic free rider problem: A laboratory study. Working Paper, 2015.

Mark Bernard, Jack Fanning, and Sevgi Yuksel. The benefits of being able to fire (and replace) people: An experimental study. Working Paper, 2014.

V Bhaskar and Eric Van Damme. Moral hazard and private monitoring. *Journal of Economic Theory*, 102(1):16–39, 2002.

Maria Bigoni, Gabriele Camera, and Marco Casari. Strategies of cooperation and punishment among students and clerical workers. *Journal of Economic Behavior & Organization*, 94(1):172–182, 2013.

Maria Bigoni, Jan Potters, and Giancarlo Spagnolo. Flexibility, communication and cooperation with imperfect monitoring. 2014.

Kenneth Binmore and Larry Samuelson. Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory*, 57(2):278–305, 1992.

Matthias Blonski and Giancarlo Spagnolo. Prisoners' other dilemma. *International Journal of Game Theory*, 44(1):61–81, 2015.

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

Yves Breitmoser. Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma. *American Economic Review*, 105(5):234–248, September 2015.

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

Luis Cabral, Erkut Y. Ozbay, and Andrew Schotter. Intrinsic and instrumental reciprocity: An experimetnal study. *Games and Economic Behavior*, 87(1): 100–121, September 2014.

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

Gabriele Camera, Marco Casari, and Maria Bigoni. Cooperative strategies in anonymous economies: an experiment. *Games and Economic Behavior*, 75: 570–586, 2012.

Timothy N Cason and Vai-Lam Mui. Coordinating resistance through communication and repeated interaction. *The Economic Journal*, 124(574):F226–F256, 2014.

Gary Chamberlain. Multivariate regression models for panel data. *Journal of Econometrics*, 18:5–46, 1982.

Oliver Compte and Andrew Postlewaite. Plausible cooperation. *Games and Economic Behavior*, 91:45–59, 2015.

David J. Cooper and Kai-Uwe Kühn. Communication, renegotiation, and the scope for collusion. *American Economic Journal: Microeconomics*, 6(2): 247–278, 2014a.

David J. Cooper and Kai-Uwe Kühn. Cruel to be kind: An experimental study of communication, collusion, and threats to punish. Working Paper, 2014b.

Russell Cooper, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–71, May 1992.

Pedro Dal Bó. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economics Review*, 95 (5):1591–1604, December 2005.

Pedro Dal Bó and G.R. Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *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, 2015.

Pedro Dal Bó and Enrique R. Pujals. The evolutionary robustness of forgiveness and cooperation. Working Paper, 2012.

Douglas Davis, Asen Ivanov, and Oleg Korenok. Individual characteristics and behavior in repeated games: An experimental study. *Experimental Economics*, 19(1):67–99, 2016.

Anna Dreber, David G. Rand, Drew Fudenberg, and Martin A. Nowak. Winners don't punish. *Nature*, 452:348–351, 2008.

Anna Dreber, Drew Fudenberg, and David G. Rand. Who cooperates in repeated games: The role of altruism, inequality aversion, and demographics. *Journal of Economic Behavior & Organization*, 98(1):41–55, 2014.

John Duffy and Jack Ochs. Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, 66(2):785–812, 2009.

John Duffy, Huan Xie, and Yong-Ju Lee. Social norms, information, and trust among strangers: theory and evidence. *Economic Theory*, 52(2):669–708, 2013.

Glenn Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61(3):567–588, 1994.

Jeffrey C. Ely and Juuso Välimäki. A robust folk theorem for the prisoner's dilemma. *Journal of Economic Theory*, 102(1):84–105, 2002.

Matthew Embrey, Guillaume R. Fréchette, and Sevgi Yuksel. Cooperation in the finitely repeated prisoner's dilemma. Working Paper, 2016.

Matthew S. Embrey, Guillaume R. Fréchette, and Ennio Stacchetti. An experimental study of imperfect public monitoring: Renegotiation proofness vs efficiency. Working Paper, 2013.

Jim Engle-Warnick and Robert L. Slonim. The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization*, 55(4):553–573, 2004.

Jim Engle-Warnick and Robert L. Slonim. Learning to trust in indefinitely repeated games. *Games and Economic Behavior*, 54(1):95–114, 2006a.

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, 2006b.

Jim Engle-Warnick, William J. McCausland, and John H. Miller. The ghost in the machine: Inferring machine-based strategies from observed behavior. Working Paper, 2007.

Piotr Evdokimov and David Rahman. Cooperative institutions. Working paper, 2014.

Robert M. Feinberg and Thomas A. Husted. An experimental test of discount-rate effects on collusive behaviour in duopoly markets. *Journal of Industrial Economics*, 41(2):153–160, June 1993.

Robert M. Feinberg and Christopher Snyder. Collusion with secret price cuts: An experimental investigation. *Economics Bulletin*, 3(6):1–11, 2002.

Robert H. Frank, Thomas Gilovich, and Dennis T. Regan. Does studying economics inhibit cooperation? *Journal of Economic Perspectives*, 7(2): 159–171, 1993.

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

Guillaume R. Fréchette and Sevgi Yuksel. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. Working Paper, 2014.

Daniel Friedman and Ryan Oprea. A continuous dilemma. *The American Economic Review*, 102(1):337–363, 2012.

James W Friedman. A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38:1–12, 1971. ISSN 00346527.

Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, pages 533–554, 1986.

Drew Fudenberg and Eric Maskin. Evolution and cooperation in noisy repeated games. *American Economic Review*, 80(2):274–79, May 1990.

Drew Fudenberg and Eric Maskin. Evolution and repeated games. Working Paper, 1993.

Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, Cambridge, Massachusetts, 1991.

Drew Fudenberg, David K. Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, pages 997–1039, 1994.

Drew Fudenberg, David G. Rand, and Anna Dreber. Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world. *American Economic Review*, 102(2):720–749, 2012.

Edward J. Green and Robert H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52:87–100, 1984.

Avner Greif. Contract enforceability and economic institutions in early trade: the maghribi traders coalition. *American Economic Review*, 83(3):525548, 1993.

John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA, 1988.

James J. Heckman. *Structural Analysis of Discrete Data with Econometric Applications*, chapter The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process, pages 179–195. MIT Press, Cambridge, MA, 1981.

Charles A. Holt. An experimental test of the consistent-conjecture hypothesis. *American Economic Review*, 75(3):314–325, 1985.

Dorothée Honhon and Kyle Hyndman. Flexibility and reputation in repeated prisoners' dilemma games. Working Paper, 2015.

John H. Kagel and Dan R. Schley. How economic rewards affect cooperation reconsidered. *Economics Letters*, 121(1):124–127, 2013.

M Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59:63–80, 1992.

Michihiro Kandori. Introduction to repeated games with private monitoring. *Journal of Economic Theory*, 102(1):1–15, 2002.

Andrew Kloosterman. Public information in markov games. *Journal of Economic Theory*, 157:28–48, 2015a.

Andrew Kloosterman. An experimental study of public information in markov games. Working Paper, 2015b.

Takao Kusakawa, Kazuhito Ogawa, and Tatsuhiro Shichijo. An experimental investigation of a third-person enforcement in a prisoner's dilemma game. *Economic Letters*, 117(3):704–707, 2012.

Volodymyr Lugovskyy, Daniela Puzzello, Andrea Sorensen, James Walker, and Arlington Williams. An experimental study of finitely and infinitely repeated linear public good games. Working Paper, 2015.

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

Hitoshi Matsushima and Tomohisa Toyama. Monitoring accuracy and retaliation in infinitely repeated games with imperfect private monitoring: Theory and experiments. Working Paper, 2011.

Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica*, pages 803–836, 1992.

Friederike Mengel. Risk and temptation: A meta-study on social dilemma games. Working Paper, 2015.

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

Roger B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, MA, 1991.

H.T. Normann and B. Wallace. The impact of the termination rule on cooperation in a prisoner's dilemma experiment. *International Journal of Game Theory*, 41:707–718, 2012.

Martin A. Nowak and Karl Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, 364(6432): 56–58, 1993.

Martin A. Nowak, Karl Sigmund, and Esam El-Sedy. Automata, repeated games and noise. *Journal of Mathematical Biology*, 33(7):703–722, 1995.

Masahiro Okuno-Fujiwara and Andrew Postlewaite. Social norms and random matching games. *Games and Economic Behavior*, 9:79–109, 1995.

Thomas R. Palfrey and Howard Rosenthal. Repeated play, cooperation and coordination: An experimental study. *The Review of Economic Studies*, 61 (3):545–565, July 1994.

David G. Pearce. Renegotiation-proof equilibria: Collective rationality and intertemporal cooperatio. Working Paper, 1987.

Michele Piccione. The repeated prisoner's dilemma with imperfect private monitoring. *Journal of Economic Theory*, 102(1):70–83, 2002.

Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Higher intelligence groups have higher cooperation rates in the repeated prisoner's dilemma. Working Paper, 2014.

David G. Rand and Martin A. Nowak. Human cooperation. *Trends in Cognitive Sciences*, 17(8):413–425, 2013.

David G. Rand, Drew Fudenberg, and Anna Dreber. It's the thought that counts: The role of intentions in noisy repeated games. *Journal of Economic Behavior & Organization*, 116(1):481–499, 2015.

Anatol Rapoport and Albert M. Chammah. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. The University of Michigan Press, Ann Arbor, 1965.

Ernesto Reuben and Sigrid Suetens. Revisiting strategic versus non-strategic cooperation. *Experimental Economics*, 15(1):24–43, 2012.

Christian Rojas. The role of demand information and monitoring in tacit collusion. *The RAND Journal of Economics*, 43(1):78–109, 2012.

Julian Romero and Yaroslav Rosokha. Constructing strategies in indefinitely repeated prisoner's dilemma. Working Paper, 2015.

Julio J. Rotemberg and Garth Saloner. A supergame-theoretic model of price wars during booms. *American Economic Review*, 76:390–407, 1986.

Alvin E. Roth. *Handbook of Experimental Economics*, chapter Introduction, pages 3–109. Princeton University Press, Princeton, 1995.

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.

Gerardo Sabater-Grande and Nikolaos Georgantzis. Accounting for risk aversion in repeated prisoners' dilemma games: an experimental test. *Journal of Economic Behavior & Organization*, 48(1):37 – 50, 2002.

Steven T. Schwartz, Richard A. Young, and Kristina Zvinakis. Reputation without repeated interaction: A role for public disclosures. *Review of Accounting Studies*, 5(4):351–375, 2000.

Tadashi Sekiguchi. Efficiency in repeated prisoner's dilemma with private monitoring. *Journal of Economic Theory*, 76(2):345–361, 1997.

Reinhard Selten, Michael Mitzkewitz, and Gerald R. Uhlich. Duopoly strategies programmed by experienced players. *Econometrica*, 65(3):517–556, May 1997.

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

Dale O. Stahl. The graph of prisoner's dilemma supergame payoffs as a function of the discount factor. *Games and Economic Behavior*, 3:47–70, 1991.

George J. Stigler. A theory of oligopoly. *The Journal of Political Economy*, 72(1):44–61, 1964.

Lijia Tan and Lijia Wei. Voluntary contribution mechanism played over an infinite horizon. *Pacific Economic Review*, 19:313–331, 2014.

Jean Tirole. *The theory of industrial organization*. MIT press, 1988.

John Van Huyck, Raymond C. Battalio, and Richard Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, pages 234–248, March 1990.

Emanuel Vespa. An experimental investigation of strategies in dynamic games. Working Paper, 2015.

Emanuel Vespa and Alistair J. Wilson. Experimenting with equilibrium selection in dynamic games. Working Paper, 2015.

Alistair J. Wilson and Hong Wu. Dissolution of partnerships in infinitely reapeated games. Working Paper, 2014.

Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press, Cambridge, 2002.