Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 200 (2022) 105416

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Match length realization and cooperation in indefinitely repeated games ☆

Friederike Mengel[a,b,*,1], Ludovica Orlandi[c,2], Simon Weidenholzer[a,3]

[a] *University of Essex, United Kingdom of Great Britain and Northern Ireland*
[b] *Lund University, Sweden*
[c] *Nottingham Trent University, United Kingdom of Great Britain and Northern Ireland*

Received 9 March 2021; final version received 3 December 2021; accepted 23 January 2022
Available online 29 January 2022

Abstract

Experimental studies of infinitely repeated games typically consist of several indefinitely repeated games ("matches") played in sequence with different partners each time, whereby match length, i.e. the number of stages of each game is randomly determined. Using a large meta data set on indefinitely repeated prisoner's dilemma games (Dal Bó and Fréchette, 2018) we demonstrate that the realized length of early matches has a substantial impact on cooperation rates in subsequent matches. We estimate simple learning models displaying the "power law of practice" and show that participants do learn from match length realization. We then study three cases from the literature where realized match length has a strong impact on treatment comparisons, both in terms of the size and the direction of the treatment effect. These results have

☆ We thank Maria Bigoni, Tilman Boergers, Anna Dreber, Matt Embrey, Guillaume Frechette, Dan Friedman, Drew Fudenberg, Ed Hopkins, Tatiana Kornienko, Volodymyr Lugovskyy, Ryan Oprea, Ronald Peeters, Karl Schlag, Giancarlo Spagnolo, Robert Sugden, Emanuel Vespa, Leeat Yariv, Sevgi Yuksel, an anonymous associate editor and two anonymous Reviewers for helpful comments and/or for providing data access. We also thank audiences at the LEG 2019 workshop at Bar Ilan, University of California Davis, the University of East Anglia, the University of Edinburgh, the University of Heidelberg, Oxford University, the SAET 2019 conference, the VIBES seminar series and at the Vienna University of Economics and Business Adminstration for their helpful comments. Financial support by the British Academy (SG162637) is gratefully acknowledged.

* Corresponding author.

E-mail addresses: fr.mengel@gmail.com (F. Mengel), ludovica.orlandi@ntu.ac.uk (L. Orlandi), sweide@essex.ac.uk (S. Weidenholzer).

[1] Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ *and* Department of Economics, Lund University, Tycho Brahes vag 1, Lund (SE).

[2] Department of Economics, Nottingham Trent University, Nottingham NG1 4FQ.

[3] Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ.

https://doi.org/10.1016/j.jet.2022.105416

important implications for our understanding of how people learn in infinitely repeated games as well as for experimental design.

JEL classification: C70; C90

Keywords: Experiments; Indefinitely repeated games; Cooperation; Social dilemmas

1. Introduction

Infinitely repeated games are of enormous importance in many areas of Economics, but also in Politics, Sociology, Biology and many other subjects. The theory of infinitely repeated games has delivered important insights into how repeated interaction changes incentives and how it can enable a wider array of outcomes including, for example, cooperation in social dilemmas. However, it does not always offer sharp predictions. For instance, in the prisoner dilemma both cooperation and defection are equilibrium actions provided that players are sufficiently patient. The multiplicity of outcomes gives an important role to empirical research on infinitely repeated games to narrow down what we can expect empirically in these games. This research has provided key insights on e.g. the determinants of cooperation in social dilemmas, the role of monitoring, or the differences between discrete and continuous time.[4]

Empirical studies of infinitely repeated interactions rely on the equivalent setting of "indefinitely repeated games" using a random continuation probability, as originally proposed by Roth and Murnighan (1978). After every round of play, there is a fixed known probability δ that the game continues for an additional round and a probability $1-\delta$ with which it ends. The length T of the repeated game is hence a random variable with expected value $\mathbb{E}[T] = \frac{1}{1-\delta}$ and standard deviation $\sqrt{\frac{\delta}{(1-\delta)^2}}$. Often researchers observe a subject in more than one indefinitely repeated game. Each of these games is then referred to as a "match" and the researcher consequently observes not just one realization of the random variable T, but a sequence of such "match length realizations". The researcher can make accurate inference from these observations on the underlying infinitely repeated game if either the realized sequences of match length realizations "correctly" represent the infinitely repeated game or if match length realizations are irrelevant for behaviour.[5] Achieving the former can be difficult as there are considerable practical difficulties involved in getting a large enough sample of different sequences of match length realizations.[6] Given that the number of match length realizations is often going to be small in practice, it is crucially important to understand whether match length realization influences behaviour.

In this paper we first demonstrate that the sequence of match length realizations has a substantial, robust and highly statistically significant effect on behavior. Using a large data set from lab experiments studying the infinitely repeated prisoner dilemma compiled by Dal Bó and Fréchette (2018) we show that when participants initially experience relatively long matches subsequent cooperation rates are substantially higher. Specifically, when most matches in the first part of

[4] See Dal Bó and Fréchette (2018) for a review of this extensive literature.

[5] In standard theory only expected match length should matter for behavior. Hence, according to standard theory match length realizations should indeed be irrelevant for behaviour.

[6] We discuss some of these difficulties in detail in Section 4.

an experiment are "long" (above theoretical median length), then cooperation rates are 44% higher in subsequent matches. Intuitively, participants who experience longer matches become more optimistic about the relative benefits of conditionally cooperative strategies and cooperate more. Moreover, by comparing the impact of long matches in the first third to the impact of long matches in the middle third of an experiment on cooperation in the final third, we demonstrate that the effect of early matches is as least as important as the effect of recent ones. This observation is consistent with the "power law of practice" which describes the phenomenon of initially steep- and then flattening out- learning curves. We develop and estimate a simple reinforcement learning model with counterfactuals (Erev and Roth, 1998) that allows us to distinguish learning from match length realizations from learning given the expected match length. We find that indeed learning from match length realization is important.

Our results show that the environment in which early interactions take place matters for subsequent interactions as people learn from match length realizations. While match length realizations and similar stochastic factors are most likely also important factors in the field, they are difficult to isolate from other confounds. The advantage of experimental data is that it allows to separate the effects found in the current paper from alternative explanations. Settings where match length realization is likely to matter include e.g. "cultural differences" between people coming from different work environments (characterized by more or less turnover) or from different social backgrounds. The effects identified in this paper likely apply in many other settings where people learn from stochastic realizations of payoff relevant variables.[7] They also have implications for evidence based policy making. If a policy (e.g. designed to increase cooperative behavior) is evaluated over a certain fixed period, it is possible that the results of the evaluation are affected by early match length realizations even if they are exogenous to the policy evaluated.

Our findings provide valuable insights into how people learn in indefinitely repeated games. As such they can inform the development of new theories of learning in games. They also have important methodological implications for the design of empirical studies. The length of each match is typically drawn at the session level in experiments, meaning that all subjects in a given session experience the same sequences of match length draws. In fact, all papers in the Dal Bó and Fréchette (2018) meta study use this or a very similar design. The number of different sequences of match length realizations for a given treatment ranges between 1 and 10 across the different papers contained in the meta-study. Given our results discussed above we would expect that - with such small numbers of match length realizations - treatment comparisons can be affected. We provide three case studies of papers from the existing literature, which were not part of the Dal Bó and Fréchette (2018) meta study and go beyond the prisoner's dilemma: a continuous time prisoner's dilemma (Bigoni et al., 2015), a public good game (Lugovskyy et al., 2017) and oligopoly games (Embrey et al., 2019). We show that - for each of them - treatment effects differ depending on match length realization. We also run our own experiments and show that in some cases the conclusions drawn from the research might have been different for different match length realizations.

Our paper contributes to a substantial and active literature on indefinitely repeated games, much of it summarized by Dal Bó and Fréchette (2018). Several researchers have documented a

[7] For instance, in common value auctions (as e.g. experimentally studied in Kagel and Levin (1986)) early realizations of common values and of public information could potentially influence subjects' subsequent bidding strategies through learning. Likewise, different realizations of states of the world in experimental financial markets (as in settings similar to Plott and Sunder (1982)) may alter subjects' beliefs on the probability with which states occur and impact trading strategies.

positive effect of the length of the immediately preceding match on cooperation (see e.g. Camera and Casari 2009, Dal Bó and Fréchette 2011, 2018, Fréchette and Yuksel 2017, Bernard et al. 2018).[8] In the context of infinitely repeated trust games, Engle-Warnick and Slonim (2006) find some evidence that there is more trust and trustworthiness in sessions that initially featured long matches as compared to sessions starting out with short ones. As they observe, this gap could have been due to individual subject or session effects since there was already more trust and trustworthiness in the beginning of the initially long sessions. We add to this literature by providing the first comprehensive analysis of the long lasting effects of (the entire sequence of) match length realization on cooperation in infinitely repeated social dilemmas. To the best of our knowledge our paper is also the first to demonstrate that people learn from match length realization using structural estimation of learning models.[9] We also advance the existing literature by discussing in detail the potential implications of these findings for measuring cooperation levels and for the design of empirical studies.

The paper is organized as follows. In Section 2 we demonstrate the main empirical finding of an effect of match length realization on cooperation using the Dal Bó and Fréchette (2018) meta study. Section 3 contains our discussion of the case-studies. We discuss implications of our results in Section 4 and Section 5 concludes. Additional theory, tables, figures and information on our own experiments can be found in an Appendix.

2. Match length realization and cooperation

2.1. The prisoner's dilemma

We consider agents who play a 2×2 indefinitely repeated prisoner's dilemma like the one illustrated in the left panel of Fig. 1. Payoffs satisfy $\mathtt{T} > \mathtt{R} > \mathtt{P} > \mathtt{S}$ and $\mathtt{T} + \mathtt{S} < 2\mathtt{R}$ such that mutual defection is the only Nash equilibrium of the stage game but mutual cooperation maximizes joint payoffs. Following Dal Bó and Fréchette (2018) we can normalize payoffs so that we only have two parameters, see middle panel of Fig. 1. The continuation probability δ indicates the probability with which the game continues for one more round. The number of stages in the indefinitely repeated game is hence a random variable T. It is common in modern experiments to play several such indefinitely repeated games. Usually participants are rematched at the end of one repeated game and play a new game with a new partner. Each such repeated game is often referred to as a "match". Typical experiments differ in the number M of such matches implemented, the expected length of a match (given by $\mathbb{E}[T] = \frac{1}{1-\delta}$) as well as the realized match length. We index the round of play within a match by t and the match by m. T^m is the realized match length of match m, i.e. the number of stages in match m.

A substantial experimental literature has studied how payoff parameters affect cooperation in the prisoner's dilemma.[10] One particularly successful approach, proposed by Blonski and Spagnolo (2015), analyzes a setting where agents can only choose among the repeated game

[8] A similar positive effect is document for the behaviour of the previous opponent, in the sense that subjects are more likely to cooperate when they have been previously matched with somebody starting out with cooperation.

[9] A number of papers estimate learning models in repeated games (see e.g. Erev and Roth (2001), Hanaki et al. (2005), Dal Bó and Fréchette (2011), Ioannou and Romero (2014) or Embrey et al. (2018)), but they do not estimate parameters capturing the extent to which people learn from match length realization.

[10] See e.g. Embrey et al. (2017) and Mengel (2018) for contributions analyzing finitely repeated PD games and Dal Bó and Fréchette (2018) for a survey of the literature on the indefinitely repeated version.

	C	D
C	R	S
D	T	P

PD game.

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

Normalized game.

	GT	AD
GT	$\mathbb{E}[T]$	$-\ell$
AD	$1+g$	0

GT vs. AD game.

Fig. 1. Left: Prisoner's dilemma (PD) game with payoff parameters $T > R > P > S$ and $T + S < 2R$. Middle: Normalized game where joint defection payoff P is subtracted from each cell and all payoffs are divided by $R - P$ (difference between mutual cooperation and defection payoffs). Right: Payoffs in the game induced by Grim-Trigger and Always Defect. GT played against GT yields a payoff of 1 in all $\mathbb{E}[T]$ stages. AD (GT) played against GT (AD) yields once a payoff of $1 + g$ $(-l)$ and zero in the remaining $\mathbb{E}[T] - 1$ stages.

strategies "Grim-Trigger" (GT) and "Always Defect" (AD) and payoffs are given by the expected sum of payoffs of the induced indefinitely repeated game shown in the right panel of Fig. 1 (see also Blonski et al. (2011) and Dal Bó and Fréchette (2011)). Provided "Grim-Trigger" can sustain cooperation in a subgame perfect Nash equilibrium ($\mathbb{E}[T] \geq 1 + g$), the resulting game constitutes a coordination game. The size of the basin of attraction of AD, denoted by $\texttt{SizeBAD}$, is defined as the threshold probability of choosing GT that has to be exceeded to make GT a best response.[11] Formally,

$$\texttt{SizeBAD} = \begin{cases} 1 & \text{if } \mathbb{E}[T] < 1+g \\ \frac{\ell}{\mathbb{E}[T]+\ell-g-1} & \text{otherwise} \end{cases} \tag{1}$$

Note that $\texttt{SizeBAD}$ is decreasing in $\mathbb{E}[T]$ (respectively δ), conveying the intuitive idea that cooperation is easier to sustain under longer expected match durations. Dal Bó and Fréchette (2018) show that $\texttt{SizeBAD}$ indeed predicts cooperation rates very well in a meta-study of indefinitely repeated prisoner's dilemma experiments. They also show that the length of the immediately preceding match has an effect on cooperation rates in the subsequent match.

They suggest that this is either due to a minority of participants who may not understand how match lengths are determined or due to how participants update their overall evaluation of the value of cooperation through experience. They write "there is an interesting - as yet unexplored - question regarding the way that humans learn in infinitely repeated games. Is the impact of the realized length constant throughout or is the impact more important early on?"

As we will see below not only does the length of the immediately preceding match matter, but the entire sequence of match length realizations is important. Further, addressing the question posed by Dal Bó and Fréchette (2018), the impact of realized match length is not constant throughout. Early matches matter at least as much as later matches and sometimes more. We will now demonstrate these patterns empirically (Section 2.2) and then estimate simple learning models to understand how people learn in indefinitely repeated games (Section 2.3).

[11] This corresponds to the probability of AD in the mixed strategy equilibrium of the game in the right panel of Fig. 1.

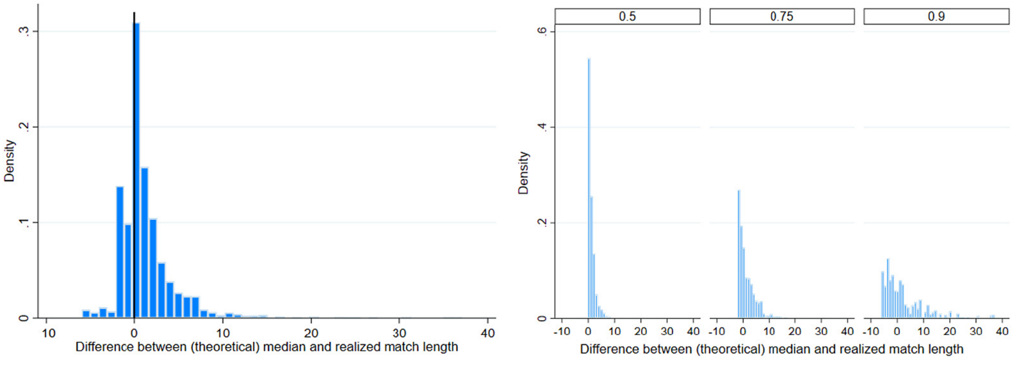

Fig. 2. Distribution of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$.

2.2. *The effect of match length realization*

To study empirically whether there exists a persistent effect of match length in early matches we use the data collected by Dal Bó and Fréchette (2018). They collected data from 141 different sessions of indefinitely repeated prisoner's dilemma experiments with 2415 participants (see Table 3 in Dal Bó and Fréchette (2018)). Some of these sessions are one-shot games ($\delta = 0$), though. In our analysis we will rely on 103 sessions with $\delta > 0$. In all papers contained in their data set the sequences of match length realizations are drawn at session level, i.e. all subjects in a given session faced the same sequence of match length realizations. Fig. 2 shows the distribution of the difference between theoretical median match length and realized match length in the meta-study.[12] The left panel aggregates games with different discount factors. It can be seen that match lengths are, as expected, concentrated around the median with a good amount of variation on both sides. The right panel shows separate graphs for the three most common discount factors $\delta = 0.5$, 0.75 and $\delta = 0.9$. The figure shows that for the longer games with $\delta = 0.75$ and $\delta = 0.9$ most matches are somewhat shorter than what we would expect. However, in all cases, there is a good amount of variation.

We use this variation to study how match length realization in early matches affects subsequent cooperation. We define early matches as the 1st third of matches in a session and create a dummy variable Δ^{1st}_{above} indicating whether more than $\frac{2}{3}$ out of these early matches were (weakly) longer than the theoretical median length.[13] The dummy takes the value 1 in 44% of sessions. Analogously, we can also define dummies Δ^{2nd}_{above} and Δ^{3rd}_{above} which take the value 1 in 42% and 47% of sessions, respectively.

[12] Appendix Figure E.1 shows kernel density estimates.

[13] The reasoning behind these choices is the following. We split matches in three groups (early, middle and late) rather than e.g. two is that it allows us to compare the effect of early (1st third) and middle (2nd third) matches on cooperation in late (3rd third) matches. This allows to address the question whether early experience or recent experience is more important for cooperation. Appendix Table D.1 shows results for alternative splits. The reason we use a dummy is that (i) theoretical medians differ with δ, which means that we cannot just use match length directly, and that (ii) it makes regression results more easily interpretable. Appendix Table D.5 shows results when we use the share of matches above median instead. Last, the reason that we use $\frac{2}{3}$ as a cutoff for the share of long matches is that it produces relatively balanced groups, though some other cutoffs would have produced that too. Appendix Tables D.6-D.7 show the results with alternative cutoffs.

Table 1
Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy Δ^{1st}_{above} indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were longer than the theoretical median match length. Columns (4)-(6): Initial (first stage) cooperation rate in the 3rd third of matches explained by dummies Δ^{1st}_{above} and Δ^{2nd}_{above}. Standard errors clustered at session level. Observations stem from 103 sessions spread across 11 papers.

	Effect of Match Length Realization on subsequent cooperation					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.142***	0.101**	0.226***	0.125**	0.101**	0.207***
	(0.053)	(0.039)	(0.056)	(0.058)	(0.041)	(0.062)
Δ^{2nd}_{above}				0.085	0.069*	0.032
				(0.060)	(0.039)	(0.068)
`SizeBAD`		−0.765***	−0.539***		−0.809***	−0.659***
		(0.069)	(0.094)		(0.068)	(0.149)
`SizeBAD` $\times \Delta^{1st}_{above}$			−0.296***			−0.241**
			(0.094)			(0.104)
`SizeBAD` $\times \Delta^{2nd}_{above}$						0.048
						(0.144)
Constant	0.321***	0.974***	0.747***	0.294***	0.994***	0.844***
	(0.028)	(0.074)	(0.098)	(0.032)	(0.074)	(0.153)
δ f.e.	NO	YES	YES	NO	YES	YES
Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$	-	-	-	0.6903	0.6063	0.0989
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.021	0.219	0.223	0.034	0.251	0.255

Robust standard errors in parentheses $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1 shows the results of regressing cooperation in the final third of matches on Δ^{1st}_{above} as well as `SizeBAD` and an interaction. Early match length has a substantial impact on cooperation in later matches. If at least $\frac{2}{3}$ of these early matches are "long", then cooperation rates are higher for the remainder of the experiment as shown by the positive coefficient on the dummy Δ^{1st}_{above} in column (1). The effect size is substantial, with cooperation rates being 44% higher when initial matches were long as compared to when they were short. As expected, the table also shows a negative impact of `SizeBAD` on cooperation rates (columns (2)-(3)). Interestingly, there is also an interaction effect between `SizeBAD` and Δ^{1st}_{above}. If early matches are long than the detrimental effect of `SizeBAD` is more pronounced. This is intuitive as longer early matches could allow participants to better learn the incentives coming from the game parameters. Conversely, the interaction term also shows that the positive effect of early match length realization is stronger the more favorable the climate is for cooperation. In fact according to Table 1, the effect is positive if and only if `sizeBAD` is smaller than 0.77.

Robustness Appendix D.1 contains tables showing that these results are qualitatively robust to the inclusion of paper fixed effects (Table D.4), to considering different thresholds (Tables D.6-D.8) or to using the share of matches above median instead of a dummy variable (Table D.5). A placebo test shown in Appendix Table D.3 where we regress cooperation in the 1st third of matches on Δ^{3rd}_{above} shows that the results in Table 1 are fundamental and not e.g. driven by correlations of match lengths within sessions or observed or unobserved heterogeneity across papers or treatments, e.g. caused by different ways researchers implement match length draws.[14]

Early vs recent matches Next we ask what is more important for cooperation in the final third of the experiment, early experience, i.e. match length in the 1st third, or recent experience, i.e. match length in the 2nd third of matches? Columns (4)-(6) in Table 1 show the results of regressing cooperation in the final third of matches on both dummies Δ^{1st}_{above} and Δ^{2nd}_{above}. The table shows very clearly that early experience in the 1st third of matches is very important. In all specifications the coefficient on Δ^{1st}_{above} is at least as large as that for Δ^{2nd}_{above} and exhibits a higher level of statistical significance. The interaction effect with `SizeBAD` is also more important for these matches. Early matches seem at least as important as recent matches and potentially, more important.[15]

Experience Does the effect vanish with more experience, i.e. if enough matches are played in the experiment? To answer this question we rerun specification (1) of Table 1 restricting the sample to sessions with (i) at most 12 matches in total, (ii) 12-24 matches, (iii) 24-36 matches etc.[16] Appendix Figure E.2 shows that a positive effect size can be found even in sessions that feature at least 72 matches in a session. The figure also shows a possible downward trend in coefficient sizes as more matches are played, but if at all the trend is slow and suggests that at the very least 80 matches would have to be played in a session for the coefficient to vanish. This can quickly become infeasible especially if the discount factor δ is high. Note also that there is a compositional effect in this analysis as sessions with more matches tend to have smaller δ in the meta-study (t-test, $p < 0.0001$). As with the geometric distribution we would expect more extreme outliers in match length realizations when δ is higher, the compositional effect should artificially exacerbate the effect of experience, i.e. make it seem that with more matches there is less of an effect of match length realization. That we see very little in terms of a downward trend despite this suggests that adding more matches will not easily eliminate the impact of early match length realizations.

To sum up, the results in this section have shown that there can be substantial and non-trivial effects of realized early match length on cooperation rates in the rest of the experiment. Hence, which match length realizations are drawn can potentially affect research results. This is particularly likely if few draws are made (e.g. only one draw per session or treatment). In Section 3 we will study three case studies highlighting this point.

We have also seen that early matches matter at least as much as recent matches. This is in line with a substantial body of evidence on both human and animal learning which shows that learning curves tend to be steeper initially and then flatter. This observation is known as "power

[14] We would not expect realized match length of final matches, which have not yet been played, to affect cooperation in the beginning of the experiment. Hence we would expect zero coefficients on Δ^{3rd}_{above} and the corresponding interaction term. We do indeed find that these coefficients are close to zero and statistically not significant.

[15] Appendix Table D.2 compares the importance of early and recent matches for more different splits. Specifically the table compares the impact of match length realization in the first X-th, second X-th, third X-th...of matches on cooperation in the last X-th of matches, where X ranges from 2,...,10. For all X=2,...,9 the coefficient on the first X-th of matches is larger than that of the (X-1)th Xth of matches.

[16] We choose multiples of 12 to cut the sample as (i) they are close to the 25th, 50th and 75th percentile of match numbers in the overall sample (25th percentile is 11, 50th is 23 and 75th is 34) and (ii) 12 divides by 2, 3 and 4 without remainder allowing us to split the total number of matches in halfs, thirds and quarters as in Appendix Figure E.2.

law of practice" and according to Erev and Roth (1998) dates back to at least Blackburn (1936). In the next subsection we will structurally estimate learning models that do have this property and show that match length realization matters for learning.

2.3. Learning

We study simple learning models which display the "power law of practice". These are straightforward adaptations of previously studied models to an environment where payoffs depend on stochastic realizations of match length. To this end, we consider a set of agents which are recurrently matched to play a series of indefinitely repeated PD games. Following much of the literature (Dal Bó and Fréchette, 2018; Embrey et al., 2017), we restrict attention to the GT and AD strategies.[17] The payoffs of the game induced by these strategies are given in the right panel of Fig. 1, where the expected match length $\mathbb{E}[T]$ is now replaced by the actual match length realization T^m.

2.3.1. Model and identification

In our learning model the choices of agents are determined by propensities which are updated after each match. Propensities can be interpreted as beliefs (as in fictitious play, see e.g. Mookherjee and Sopher 1997) but can also incorporate a much wider set of feelings, such as e.g. familiarity or habituation (as in reinforcement learning, see e.g. Erev and Roth 1998, Boergers and Sarin 1997). Each agent i is endowed with an initial propensity for each strategy, denoted by $\psi_s^{i,0}$ for strategies $s \in \{GT, AD\}$, which may capture pre-game experience, initial inclinations or beliefs. In this section we focus on *reinforcement learning with counterfactuals* (see e.g. Vriend 1997 and Rustichini 1999 and the special cases in Erev and Roth 1998 and Camerer and Ho 1999) where propensities $\psi_s^{i,m}$ for all strategies evolve according to

$$\psi_s^{i,m+1} = \psi_s^{i,m} + \pi(s, s^{-i,m}, T^m).$$

$s^{-i,m}$ denotes the strategy of i's opponent in match m and $\pi(s, s^{-i,m}, T^m)$ gives the payoff earned with strategy s in this case. It remains to specify a choice rule. We assume that i's probability to choose grim trigger in match m is given by the logit choice rule

$$p^{i,m} = \frac{e^{\lambda \psi_{GT}^{i,m}}}{e^{\lambda \psi_{GT}^{i,m}} + e^{\lambda \psi_{AD}^{i,m}}}, \tag{2}$$

where λ is a measure of noise, sometimes also thought of as a measure of rationality of the economic agent. If $\lambda = 0$, then $p^{i,m} = 0.5$, i.e. the agent chooses randomly with uniform probability between GT and AD. By contrast, if $\lambda \to \infty$, then the agent chooses the strategy with the higher propensity with probability 1.[18]

[17] A theoretical justification for why it is sensible to restrict to these strategies is provided in Blonski and Spagnolo (2015) and Blonski et al. (2011).

[18] In environments where payoffs are stable, in the sense that they do not feature an exogenous stochastic element, reinforcement learning with counterfactuals is closely related to (smooth) fictitious play (see e.g. Fudenberg and Kreps 1993 and Fudenberg and Levine 1998) where agents play a (smooth) best response to the belief that future play will follow the past empirical distribution (see e.g. Cheung and Friedman 1997 and Camerer and Ho 1999). The equivalence holds because looking back to previous earnings of strategies is equivalent to forming beliefs based on past behaviour and then computing expected payoffs based on these beliefs. In Appendix A we show the results of simulations where we include also reinforcement learning without counterfactuals and fictitious play and where we also use the linear choice rule instead of the logit choice rule. We find that all models perform very similarly in the simulations.

To estimate this model we first rewrite the updating rule recursively as

$$\psi_s^{i,m+1} = \psi_s^{i,0} + \pi^{i,1}(s, s^{-i,1}, T^1) + \ldots + \pi^{i,m}(s, s^{-i,m}, T^m). \tag{3}$$

If we then denote by $\Delta\Pi^m = \sum_{h=1}^{m} \pi^{i,h}(GT, s^{-i,h}, T^h) - \pi^{i,h}(AD, s^{-i,h}, T^h)$ the cumulative payoff difference between grim trigger and always defect across all past matches, we can re-write (2) as

$$p^{i,m} = \frac{e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \Delta\Pi^m]}}{1 + e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \Delta\Pi^m]}}. \tag{4}$$

Denote by $y_{GT}^{i,m} = 1$ the outcome where player i chooses grim trigger in match m and by $y_{GT}^{i,m} = 0$ where they do not. Given the choice rule above, the likelihood function associated with this binary outcome then takes the following form

$$\begin{aligned} L(\psi_{GT}^{i,0}, \psi_{AD}^{i,0}, \lambda) = \prod_{i=1}^{N} \prod_{m=1}^{M} &\left(\frac{e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \Delta\Pi^m]}}{1 + e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \Delta\Pi^m]}} \right)^{y_{GT}^{i,m}} \\ &\times \left(\frac{1}{1 + e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \Delta\Pi^m]}} \right)^{1 - y_{GT}^{i,m}} \end{aligned} \tag{5}$$

Estimating this model using logistic regression we have that the coefficient for the constant term of the regression is an estimate of $\lambda(\psi_{GT}^{i,0} - \psi_{AD}^{i,0})$ and the coefficient on $\Delta\Pi^m$ is an estimate of λ. Comparing these coefficients hence allows us to identify λ and the difference in initial propensities $(\psi_{GT}^{i,0} - \psi_{AD}^{i,0})$.

Extended model The standard model does not distinguish between learning about the behaviour of opponents - given expected match length - and learning about match length realizations. To allow us to do so, we extend the standard model and consider the following updating rule.

$$\psi_s^{i,m+1} = \psi_s^{i,m} + \alpha\pi^{i,m}(s, s^{-i,m}, T^m) + (1-\alpha)\pi^{i,m}(s, s^{-i,m}, \mathbb{E}[T]). \tag{6}$$

This nests the standard model when $\alpha = 1$. If $\alpha = 0$, then learning about match length realization (T^m) does not matter and all learning relies on expected match length. The higher α the relatively more important is learning about match length realizations. In analogy to above we define $\Delta\Pi_{\mathbb{E}[T]}^m = \sum_{h=1}^{m} \pi^{i,h}(GT, s^{-i,h}, \mathbb{E}[T]) - \pi^{i,h}(AD, s^{-i,h}, \mathbb{E}[T])$ as the cumulated payoff difference between grim trigger and always defect across all past matches conditional on the opponent's choice and assuming that match length is fixed at the expected match length $\mathbb{E}[T]$. The likelihood function for the extended logit choice model can then be written as follows

$$L(\psi_{GT}^{i,0}, \psi_{AD}^{i,0}, \lambda, \alpha) = \prod_{i=1}^{N} \prod_{m=1}^{M} \left(p^{i,m}\right)^{y_{GT}^{i,m}} \left(1 - p^{i,m}\right)^{1 - y_{GT}^{i,m}}$$

with

$$p^{i,m} = \frac{e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \alpha\Delta\Pi^m + (1-\alpha)\Delta\Pi_{\mathbb{E}[T]}^m]}}{1 + e^{\lambda[(\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) + \alpha\Delta\Pi^m + (1-\alpha)\Delta\Pi_{\mathbb{E}[T]}^m]}} \tag{7}$$

The log-likelihood is given by

$$LL(\psi_{GT}^{i,0}, \psi_{AD}^{i,0}, \lambda, \alpha) = \sum_{i=1}^{N}\sum_{m=1}^{M} y_{GT}^{i,m} \ln\left(p^{i,m}\right) + (1 - y_{GT}^{i,m}) \ln\left(1 - p^{i,m}\right)$$

Estimating this model using logistic regression we have that the coefficient for the constant term of the regression is an estimate of $\lambda(\psi_{GT}^{i,0} - \psi_{AD}^{i,0})$. The coefficient on the difference $\Delta\Pi^m$ is an estimate of $\lambda\alpha$ and the coefficient on $\Delta\Pi_{\mathbb{E}[T]}^m$ of $\lambda(1-\alpha)$. Comparing these coefficients allows us to identify λ, α and the difference in initial propensities $(\psi_{GT}^{i,0} - \psi_{AD}^{i,0})$. The parameter we are most interested in here is α, which will tell us to which extent participants learn about match length realization as opposed to using expected match length to make their decisions. Note that our updating rule imposed the constraint that $\alpha \in [0, 1]$. As the unconstrained estimates return values of $\alpha \in [0, 1]$ we report unconstrained estimates below.

Our parameters of interest α, λ and $(\psi_{GT}^{i,0} - \psi_{AD}^{i,0})$ are all estimated at the population level. Wilcox (2006) has shown, though, that in pooled estimations of learning models that contain lagged variables (such as the ones above), the presence of heterogeneity e.g. in the parameter λ can cause biases in the estimated parameters (see also Cabrales and Garcia-Fontes (2000)). Specifically, Wilcox (2006) is concerned with the estimation of a parameter that measures the extent to which learners take into account counterfactuals when updating their propensities. A downward bias in this parameter leads to overestimation of the role reinforcement learning without counterfactuals plays as opposed to models (like belief learning) which take counterfactuals into account. All the learning models we estimate do take counterfactuals fully into account. Hence our estimates are not affected by this particular issue.[19] To study whether similar issue could plague our parameters of interest, most importantly α, we conduct extensive Monte Carlo studies which we report in detail in Appendix B, where we also derive the maximum likelihood estimator of α. Those studies show that heterogeneity in λ can indeed cause downward biases of λ of up to 5% for the most relevant case when λ is assumed to be in the range found in Table 3. Most importantly, they also show that estimates of α are unbiased both when heterogeneity in λ is introduced as well as when heterogeneity in α is assumed in the data generating process. For details see Appendix B.

2.3.2. *Estimation results*

Table 2 shows the parameter estimates for the basic model. Two patterns emerge. First, in terms of the difference between initial propensities we see that participants favor AD over GT $((\psi_{GT}^{i,0} - \psi_{AD}^{i,0}) < 0)$ especially when the horizon is short ($\delta = 0.5$). This is in line with the intuition that a longer horizon makes more cooperative strategies more attractive. Second, there is a relatively large degree of noise with $\lambda \approx 0.027$ on average. Noise decreases somewhat as δ increases. To provide some context to these numbers we note that given the average values of $\Delta\Pi^m$ these estimates together with choice rule (2) imply an $\approx 59\%$ chance of picking the "correct" strategy, i.e. the strategy with the higher value of $\Delta\Pi^m$, on average across all matches m. If we focus on the second half of matches in a session, then this value increases to about 64%, but it is well below 100% in all cases.

Note that the accumulated payoff differences $\Delta\Pi^m$ will tend to be strongly correlated with `SizeBAD` (across all supergames $\rho = -0.7213^{***}$). Hence one might wonder whether control-

[19] We are not interested in comparing learning models with and without counterfactuals. Still, simulations reported in Appendix A show that these types of models perform very similarly with our data.

Table 2
Parameter estimates basic model for all data and separately for $\delta = 0.5$, $\delta = 0.75$ and $\delta > 0.5$. 95% confidence interval in brackets.

	All Data	*Parameter Estimates Basic Model*		
		$\delta = 0.5$	$\delta = 0.75$	$\delta > 0.5$
λ	0.027	0.020	0.034	0.033
	[0.026, 0.028]	[0.020,0.021]	[0.033,0.035]	[0.033, 0.034]
$(\psi^{i,0}_{GT} - \psi^{i,0}_{AD})$	−1.460	−8.700	−0.676	−1.242
	[−1.576, −1.341]	[−9.459, −7.941]	[−0.700, −0.652]	[−1.366, −1.118]
N	37394	16088	18136	21306

Table 3
Parameter estimates extended model for all data and separately for $\delta = 0.5$, $\delta = 0.75$ and $\delta > 0.5$. 95% confidence interval in brackets.

	All Data	*Parameter Estimates Extended Model*		
		$\delta = 0.5$	$\delta = 0.75$	$\delta > 0.5$
α	0.432	0.486	0.589	0.631
	[0.423, 0.440]	[0.484,0.488]	[0.573, 0.605]	[0.605, 0.657]
λ	0.186	0.603	0.172	0.117
	[0.165, 0.207]	[0.533, 0.673]	[0.145,0.200]	[0.096, 0.138]
$(\psi^{i,0}_{GT} - \psi^{i,0}_{AD})$	−0.768	−0.486	−0.191	−0.708
	[−0.913, −0.623]	[−0.578, −0.394]	[−0.384, 0.001]	[−0.989, −0.426]
N	37394	16088	18136	21306

ling for `SizeBAD` might reduce bias in the estimation and/or improve the precision of estimates. Of course doing so makes the interpretation of our parameters of interest in terms of the learning model less straightforward. If we nevertheless control for sizebad we get an estimated value of $\lambda = 0.019$ (confidence interval [0.019, 0.020]) when using all data as in column (1) of Table 2.[20] This corresponds to an about 55% chance of picking the correct strategy across all matches which increases to 60% in the second half of matches.

Table 3 shows the estimates for the extended model. We see again that participants somewhat favor AD over GT in terms of the initial propensities. The model has much less noise than the basic model. We can ask again what our estimates imply in terms of the probability to choose the "correct" strategy - i.e. the strategy with the higher weighted average of $\Delta\Pi^m$ and $\Delta\Pi^m_{\mathbb{E}[T]}$ where the weight is determined by the estimated α. The estimated values of λ imply a probability of choosing this strategy between $\approx 92\%$ when $\delta = 0.5$ and $\approx 76\%$ when $\delta > 0.5$.

Our main parameter of interest is α. In all cases α substantially exceeds zero, showing that participants do learn from match length realizations. This explains why match length realization affects cooperation in subsequent matches. However, it is not the case that only realized match length matters. The estimates clearly suggest that both expected match length and realized match length play a role for participants' learning. Learning from match length realizations is somewhat

[20] When restricting to $\delta = 0.5$ we obtain $\lambda = 0.010$, for $\delta = 0.75$ we get $\lambda = 0.028$ and for $\delta > 0.5$ we obtain $\lambda = 0.028$ in this case.

more important when the horizon is longer ($\delta > 0.5$), which is intuitive as in these cases we can expect more variation in match lengths.[21]

How well do these models predict actual cooperation? Appendix Figure E.3 shows observed cooperation in stage 1 of a match depending on the predicted probability ($p^{i,m}$) of using Grim Trigger according to the basic model (left panel) and the extended model (right panel).[22] The figure shows that there is some prediction error in both models with observed cooperation differing from predicted cooperation by up to twenty percentage points. Prediction errors are generally lower with the extended model.

In sum this section has shown that simple learning models can explain the data reasonably well. Estimates of our extended model clearly show that participants do learn from match length realization. Because participants learn from match length realization, treatment comparisons can be affected by "unusually" long or short match length realizations. And, because learning displays the "power law of practice" early match length realizations will be particularly important.

3. Case studies

We will discuss three applications to illustrate how match length realizations can affect treatment comparisons when indefinitely repeated games are compared with finitely repeated games (subsections 3.1 and 3.2) or when indefinitely repeated games are compared with other indefinitely repeated games (subsection 3.3). The three cases highlighted are not part of the Dal Bó and Fréchette (2018) meta-study data used in Section 2.2 and feature a continuous time prisoner's dilemma (subsection 3.1), a public good game (3.2) and oligopoly games (3.3).[23]

3.1. Cooperation in continuous time

Our first case study is the paper "Time Horizon and Cooperation in Continuous Time" by Bigoni et al. (2015) published in *Econometrica*. Bigoni et al. (2015) compare cooperation rates in a prisoner's dilemma played in deterministic and stochastic continuous time.[24] They consider games of short (20 seconds) and long (60 seconds) expected length, where here we focus on the short games (which is where they find a treatment effect). The deterministic short game lasts 20 seconds. The stochastic short game has a continuation probability of $\delta = \frac{992}{1000}$ and every 0.16 seconds it ends with probability $1 - \delta$. This means that the expected match length in the continuous game is 20 seconds just as in the deterministic game. The expected median length is 13.86 seconds. Bigoni et al. (2015) focus on average cooperation rates in a match. They find that in short games cooperation is higher under deterministic than under the stochastic horizon.

21 We can again control for sizebad. Doing so yields values of $\alpha = 0.437$ and $\lambda = 0.151$ when all data are used. For $\delta = 0.5$ we get $(\alpha, \lambda) = (0.486, 0.434)$, for $\delta = 0.75$ we get $(\alpha, \lambda) = (0.415, 0.151)$ and for $\delta > 0.5$ we get $(\alpha, \lambda) = (0.631, 0.106)$.

22 In the setting with only two strategies GT and AD cooperation in stage 1 uniquely identifies strategy GT. Fudenberg and Karreskog (2020) show that initial play is indeed highly predictive about average cooperation in a match using the same data.

23 Our selection of case studies followed four criteria: (i) the paper should *not* be already included in the meta-study used in Section 2; (ii) it should be on an indefinitely repeated social dilemma; (iii) it has to feature different match length realizations across sessions and (iv) data are publicly available or were made available to us.

24 This important research program combines elements from Dal Bó (2005) studying the role of deterministic vs. stochastic horizon in discrete time and Friedman and Oprea's (2012) study of discrete vs. continuous time under a deterministic horizon.

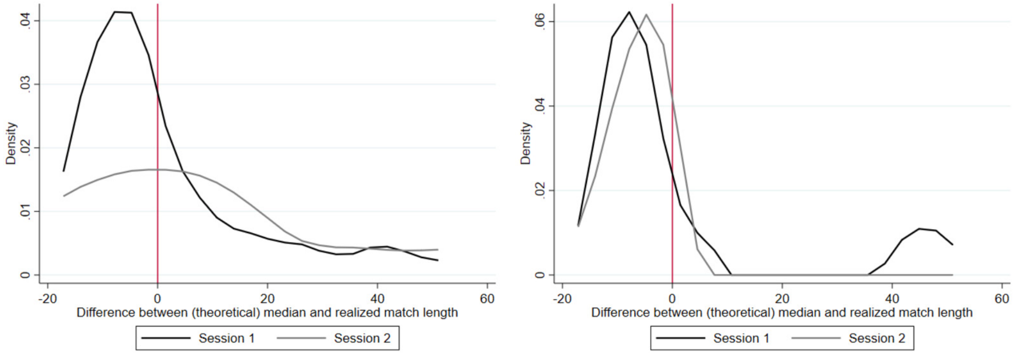

Fig. 3. Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and for the 1st third of matches. 58% of all matches and 81% of matches in the 1st third were shorter than theoretical median match length.

We now study how this result might be affected by match length realizations. Bigoni et al. (2015) conduct two sessions for each treatment condition. In each session there are 24 participants who play 23 matches. Match length is drawn at the session level, i.e. all participants face the same sequence of match lengths. Fig. 3 shows kernel density estimates of the difference between theoretical median match length and realized match length for the two sessions. The left panel shows the entire session and the right panel only the 1st third of the experiment, specifically the first 8 matches (out of 23). It can be seen that in both sessions the vast majority of matches (81%) at the beginning of the experiment (right panel) were shorter than theoretical median length.

To study whether this realization of match lengths could have affected the treatment effect we first replicated Bigoni et al. (2015)'s experiment. We conducted four sessions of the deterministic condition and then two sessions with the same match-length realizations as Bigoni et al. (2015) ("Replication"). Those sessions were conducted as exact replications of their study. See Appendix C for further details. We further conducted two sessions with inverse match length realizations ("Inverse"). For the inverse sessions we determined a sequence of match lengths $(T^m)_{m=1}^{23}$ as follows. For each realized match length T^m in the Replication we compute $\Pr(x \leq T^m)$ and then replace the m-th entry in the sequence by the value T' that satisfies $\Pr(x \leq T') = 1 - \Pr(x \leq T^m)$. Appendix Figure E.4 illustrates how the "inverse" match length sequences are constructed. Last, we conducted 4 sessions where we randomize the sequence of match lengths at the match level ("Match Stoch"). Hence in this treatment we have 96 different realized match length sequences as opposed to just two.

Table 4 gives an overview of the different treatments we conducted as well as the average and median match lengths. The table shows that - compared to the deterministic case - both average and median match length are short in the replication treatments, particularly in the 1st third of the experiment. There the median match length is only 8.48 seconds, much shorter than the 20 seconds in the deterministic case or than the theoretical median of 13.86 seconds. In the inverse condition these match length realizations are naturally longer with the median match length in the 1st third being 21.66, just above the deterministic condition. Last, as expected, when match lengths are drawn at the match level, then, by the law of large numbers, both average and median lengths are close to the theoretical averages and medians.

How does match length realization affect average cooperation rates and the treatment comparison? First, it should be noted that we manage to replicate Bigoni et al. (2015)'s result quite

Table 4
Summary Statistics of the different treatments conducted to replicate Bigoni et al. (2015). Number of Sessions (S) and observations (N) in the different conditions. Average Match Length (Avg ML), average match length in the 1st third of the experiment (Avg ML 1st third), median match length and median match length in the 1st third, average cooperation rate (avg coop) and average initial cooperation rate (avg coop initial). Stars indicate statistical significance (*** 1%, ** 5%, * 10%) of the difference to the deterministic case in random effects OLS regression with standard errors clustered at session level (see Appendix Tables D.9 and D.10).

	S	N	*Avg ML*	*Avg ML 1st third*	*Median ML*	*Median ML 1st third*	*Avg Coop*	*Avg Coop Initial*
Deterministic	4	2208	20	20	20	20	54.04	73.95
Replication	2	1104	22.94	17.6	11.04	8.48	39.58***	51.90***
Inverse	2	1104	17.86	24.0	17.56	21.66	50.90	73.36
Match Stoch	4	2208	19.97	19.57	13.44	13.52	47.40	61.18**

closely. Between the deterministic and replication treatment there is a 14.46 percentage point difference in average cooperation (Table 4) compared to Bigoni et al. (2015)'s 10.9 percentage point difference (Table II in Bigoni et al. (2015)). We fail to replicate the result, though, when we use inverse match lengths. Here the difference in average cooperation rates to the deterministic case is only 3.14 percentage points and not statistically different from zero. With match level draws (Match Stoch) we find a difference to the deterministic case of 6.64 percentage points which is less than half of the effect size than in the replication, but more than twice the effect size of the inverse condition. The difference between the match stochastic condition and the deterministic case is not statistically significant at the 10% level.[25] Having a closer look at the data, we do find, however, that average *initial* cooperation rates (in the first stage of each match) do differ significantly between the Match Stoch and deterministic environments with an effect size of about half of that found in the replication.

Fig. 4 illustrates the different effect sizes that can be obtained for the comparison between the deterministic and stochastic game depending on match length realizations. The largest effect size is obtained in our replication of Bigoni et al. (2015)'s original study where we use the same match length realizations as them. This is true for both average cooperation rates (left panel) and initial cooperation rates (right panel). It can also be seen, though, that the treatment difference for average cooperation rates is not statistically significant for any other match length realization. For initial cooperation rates the treatment difference is smaller in the treatment with match-level randomization and statistically not different from zero for the inverse treatment. We also analyzed two sub-groups from the match-stochastic treatments: those with the smallest and those with the largest share of early matches with above median length. For the latter (MaxMatch) we even find a statistically significant negative treatment effect, specifically higher initial cooperation rates in the stochastic game. This exercise illustrates how treatment comparisons can yield entirely different conclusions depending on match length realization. We should also note that - despite the fact that treatment effects can be strongly impacted by match length realizations - we do not consider this an overall unsuccessful replication. Our exact replication was very successful and in the treatment with "many match length realizations" (MatchStoch) the direction of the effect

[25] In this treatment we use twice the number of observations as in Bigoni et al. (2015). It is still possible that this effect becomes statistically significant with a larger sample size.

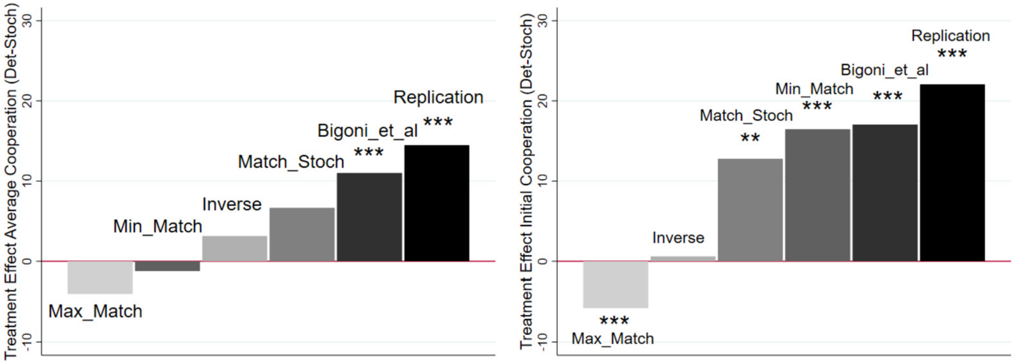

Fig. 4. Different effect sizes obtained for the treatment differences between deterministic and stochastic treatments in average cooperation rates (left panel) and initial cooperation rates (right panel) depending on match length realizations. Bigoni et al.: original effect size in Bigoni et al. (2015); Replication: replication treatment; Inverse: inverse treatment; MatchStoch: treatment with match length realization drawn at match level; MinMatch: only individuals from MatchStoch who had the smallest number of matches 1-7 above median (specifically 1) in this treatment; MaxMatch: only individuals from MatchStoch who had the largest number of matches 1-7 above median (specifically 6) in this treatment.

goes in the same direction as in the original study, even though not always being statistically significant.[26]

We will get back to the question of how to measure the "correct" treatment effect in Section 5. Before we do so, we study two more applications showing how treatment comparisons can be affected by match length realizations.

3.2. Finite and indefinitely repeated linear public good games

Our second case study is the paper "An experimental study of finitely and infinitely repeated linear public goods games" by Lugovskyy et al. (2017) published in *Games and Economic Behavior*. The finitely repeated games they study all have a match length of 5 rounds, while for the indefinitely repeated public good games they draw three sequences of match lengths (using discount factor $\delta = 0.8$). Average match length in the 1st third of sequence 1 is below the mean of 5 used in the finite sessions, specifically 4.4 rounds. By contrast, in sequences 2 and 3 it is above, specifically 6 and 6.6, respectively. Hence initial matches are substantially shorter in sequence 1 compared to the other sequences. Overall, however, the three sequences are very similar with average match length across all 15 matches equaling 5.3, 5.4 and 5.7, respectively. In both the finite and indefinitely repeated sessions participants play 15 matches.

The first hypothesis Lugovskyy et al. (2017) test is that "contributions in repeated games with sequences that have probabilistic end rounds will be greater than or equal to those in repeated games with sequences that have known end rounds". They evaluate this hypothesis by comparing behavior in finite and probabilistic settings for four different pairs of treatments which differ in group size, MPCR and whether participants make a binary contribution choice or not.

[26] Appendix Figure E.5 shows different treatment effects when median cooperation frequencies are compared. Here again, the replication shows the biggest effect and the treatment difference is not statistically significant for the Inverse and MaxMatch condition.

Table 5
Average cooperation rates across all rounds in the finite sessions (column (1)) and across all sessions with a probabilistic ending (column (2)) as in Table 3 in Lugovskyy et al. (2017). We further split the sessions with probabilistic ending in those with initially short matches (S1) and those with initially long matches (S23). Below each cooperation rate we show how the finite setting compares to the rate in question (as in Lugovskyy et al. (2017)). The last column shows the p-value when comparing initially short and long sequences. Following Lugovskyy et al. (2017) standard errors are clustered at the participant level in all regressions.

Decision Setting	*All Rounds Cooperation Rate*				
	Finite	Prob All	Prob S1	Prob S23	S1 vs S23
$N = 4$, MPCR= 0.3	15.0	22.4	10.36	28.41	
		$<^{**}$	$>$	$<^{***}$	$p = 0.000$
$N = 4$, MPCR= 0.6	39.4	44.3	33.52	48.39	
		$<$	$>$	$<$	$p = 0.088$
$N = 2$, MPCR= 0.6	41.1	38.3	31.71	42.64	
		$>$	$>$	$<$	$p = 0.084$
$N = 2$, MPCR= 0.6, Binary	54.5	41.2	36.84	42.65	
		$>^{***}$	$>^{**}$	$>^{**}$	$p = 0.398$

Table 5 shows the results of this analysis. The first two columns ("Finite" and "Prob All") reproduce the analysis in Table 3 in Lugovskyy et al. (2017). The analysis shows that in two of the four treatments cooperation is higher in the finitely repeated game and in the other two it is higher in the indefinitely repeated (probabilistic) game. One each of these comparisons is statistically significant. These and other analysis lead Lugovskyy et al. (2017) to conclude "We do not, however, find consistent evidence that overall cooperation rates are affected by whether the number of decision rounds is finite or determined probabilistically."

When we split out the sessions in those with initially short and those with initially long matches, though, we might have reached a different conclusion. The column "Prob S1" shows cooperation rates as well as comparisons in the session with initially shorter matches. In this case all four comparisons point into the same direction: more cooperation in the finitely repeated game. Only one of the comparisons is statistically significant. It should be noted, however, that the first ($N = 4$, MPCR= 0.3) and third comparisons ($N = 2$, MPCR= 0.6) are both just outside 10 percent statistical significance ($p = 0.150$, $p = 0.102$) in a comparison that is somewhat underpowered.

In the sessions with initially long matches the picture is very different. In this case three out of four comparisons point towards less cooperation in the finitely repeated game. Out of the statistically significant comparisons one each is pointing towards more and one towards less cooperation in the finitely repeated game. Hence while the sessions with initially long matches show more of a similar picture than the overall sample, the sessions with initially short matches behave quite differently and would lead to a different conclusion. It should also be noted that, except for the last comparison ($N = 2$, MPCR= 0.6, Binary), the differences in average cooperation rates across the initially short and long sessions are always statistically significant.[27] Appendix Table D.11 shows that similar conclusions hold when we consider first round cooperation rates only.

[27] This difference is not driven by the shorter matches themselves. If we restrict attention to cooperation rates in the last third of matches only we find a difference (S1-S23) of -18.03, $p = 0.000$ for treatment $N = 4$, MPCR= 0.3, of -17.55, $p = 0.082$ for treatment $N = 4$, MPCR= 0.6, of -7.66, $p = 0.260$ for treatment $N = 2$, MPCR= 0.6 and of $-16, 42$, $p = 0.133$ in treatment $N = 2$, MPCR= 0.6, Binary.

Table 6
Efficiency measure from Embrey et al. (2019). Treatment Effect is the difference between substitutes and complements. The two rightmost columns split out the groups with short initial matches from the rest. As in Embrey et al. (2019)'s main analysis matches 7-10 are considered. Stars are from t-tests with standard errors clustered at group level. *** $p < 0.01$, **, $p < 0.05$, * $p < 0.1$.

with RO	All groups	Short First Matches	Other Groups	Short vs Others
Substitutes	21.0	7.7	28.2	
Complements	26.4	44.2	16.9	
Treatment Effect	− 5.4	−36.4***	11.3**	$p = 0.0030$
without RO	All groups	Short First Matches	Other Groups	Short vs Others
Substitutes	22.7	30.9	18.2	
Complements	9.4	13.8	7.0	
Treatment Effect	13.3**	17.1	11.2	$p = 0.6888$

While on balance it seems to us that Lugovskyy et al. (2017)'s overall conclusion is likely to be robust once "many" match length realizations are considered, the case study shows again how easily different conclusions could have been reached with different match length realizations.

3.3. *Strategy revision opportunities*

Our third case study is the paper "Strategy Revision Opportunities and Collusion" by Embrey et al. (2019) published in *Experimental Economics*. Embrey et al. (2019) explore how the possibility of being able to change a repeated game strategy during the course of play (i.e. to use "behaviour" strategies) affects cooperative behaviour in stylized oligopoly experiments. Their main treatment variations compare games of strategic substitutes and strategic complements with and without revision opportunities (RO). They find that without RO (when strategies have to be encoded upfront) there is more cooperation in games of substitutes than in games of complements. With RO there is more cooperation in games of strategic complements than with substitutes, but the latter difference is not statistically significant. The column "All groups" in Table 6 shows their main treatment effects in terms of efficiency, i.e. in terms of the percentage of the difference between joint profit maximizing payoff and Nash equilibrium payoff realized in the stage game.

Embrey et al. (2019) use a discount factor of 0.875 implying a median match length of 7 stages. There are six different matching groups with different match length realizations. We will group those into two categories: (i) those where first matches are short, specifically where the first three (out of ten) matches all have a below median match length and (ii) all other matches. Table 1 in Embrey et al. (2019) shows that two matching groups (groups 2 and 6) fall into category (i). It should be noted that these matching groups do not have fewer stages overall than others. In fact group 6 has the most stages overall of all groups.[28]

Table 6 shows the average efficiency, defined as the percentage of the difference between joint profit maximizing payoff and Nash equilibrium payoff realized in the stage game. In all cases we follow Embrey et al. (2019) and focus on the average efficiency across matches 7-10.[29] We

[28] Alternatively we could split the sample into two equal sized categories by focusing on the length of the first two matches only, as there are three groups were both the first and second match are below median length. In this case qualitatively the same conclusions do hold.

[29] Appendix Table D.12 shows all matches.

then compare average efficiency in games of strategic substitutes and complements separately for games with and without revision opportunities (RO).

The table shows that treatment effects depend on match length realizations. The overall negative, but statistically insignificant treatment effect (difference between substitutes and complements) with revision opportunities seems driven by the groups with short initial matches, where the effect is almost seven times larger and statistically highly significant. In the other groups the treatment effect reverses sign and is also statistically significant. The difference between groups with short initial matches and other groups is statistically significant at the 1% level. In this case, hence, diametrically opposite conclusions could be reached when matches with short or long initial realizations are studied. This is also illustrated in the left panel Appendix Figure E.6 which shows the treatment effect for all possible selections of two match length realizations.

Without revision opportunities we see a slightly different pattern. There is an overall positive and statistically significant treatment effect. The effect is somewhat larger with short first matches and somewhat smaller in the other groups, but both are positive. The effect is not significant in either of the subgroups, presumably due to lower sample size. Hence in this case the treatment comparison seems robust to match length realizations as is also illustrated in the right panel of Appendix Figure E.6.

Note that when revision opportunities are ruled out then all updating has to take place across (as opposed to within) matches. The analysis in this Section and the fact that match length realization seems to have a stronger impact with revision opportunities could suggest that within match learning might also play an important role for the effect of match length realizations. This could be an additional reason why the effect sizes observed with human players in the Dal Bó and Fréchette (2018) meta study are larger than with computer simulated learners who learn only across matches (see Section 2.3). In the next section we will discuss implications for experimental design in more detail.

4. Discussion

In this section we discuss some implications of our findings as well as potential ways to diagnose and solve potential problems stemming from match length realizations in experimental studies of indefinitely repeated games. Before we go deeper into the different possibilities, it is useful to define the problem. Assume a researcher is interested in measuring how the probability of cooperation depends on the decision environment $\omega \in \Omega$. Here ω can capture things like subject characteristics (age, gender,...), game parameters (sizeBAD, δ,...), decision settings (lab, field,...) or experimental treatments that affect the probability of cooperation $\Pr(C) = f(\omega)$.

Say the experimenter can conduct two treatments ω_0 and ω_1 which cleanly manipulate a dimension of interest in the decision environment and is interested in the effect of these treatments on the probability of cooperation. The experimenter is interested in the treatment effect

$$f(\omega_0) - f(\omega_1).$$

The problem arises because the effect of ω can only be measured together with a match length realization. The probability of cooperation is hence given by $\Pr(C) = f(\omega, T)$, where T is the match length realization drawn from a geometric distribution with mean $\frac{1}{1-\delta}$ and standard deviation $\sqrt{\frac{\delta}{(1-\delta)^2}}$.

Say the experimenter conducts K sessions for each treatment, indexed $k = 1, ..., K$, where each session has N participants, indexed $i = 1, ..., N$ who play each M matches, indexed $m =$

$1, ..., M$. Hence, for each treatment ω_0 and ω_1 the experimenter observes $K \cdot N \cdot M$ individual decisions to cooperate $f(\omega_0, T^m_{0,k})$ and $f(\omega_1, T^m_{1,k})$, which in the usual design depend on K realizations of sequences of match lengths $(T^m_{0,k})^M_{m=1}$ and $(T^m_{1,k})^M_{m=1}$. It is common to focus on average cooperation across the $K \cdot N \cdot M$ observations.[30] This gives the following observed (sample average) treatment effect

$$\frac{1}{K \cdot N \cdot M} \sum_{k=1}^{K} \sum_{i=1}^{N} \sum_{m=1}^{M} \left(f(\omega_0, T^m_{0,k}) - f(\omega_1, T^m_{1,k}) \right), \tag{8}$$

where we have abstracted from individual heterogeneity in the treatment effect as our focus is on match length realization. In general the observed treatment effect will be biased, i.e. will not equal $f(\omega_0) - f(\omega_1)$. A common practice in experimental research is to draw the same sequence of match length realizations for all treatments, i.e. to ensure that $(T^m_{0,k})^M_{m=1} = (T^m_{1,k})^M_{m=1}, \forall k$. Both Lugovskyy et al. (2017) and Embrey et al. (2019) do exactly that, i.e. draw one set of sequences and use it for all treatments. If the influence of match length realization is orthogonal to the treatment, then the resulting treatment effect will *not* be biased. To see this note that in this case the observed treatment effect (8) can be written as follows

$$\begin{aligned}
&\frac{1}{K \cdot N \cdot M} \sum_{k=1}^{K} \sum_{i=1}^{N} \sum_{m=1}^{M} \left(f(\omega_0) + f(T^m_{0,k}) - f(\omega_1) - f(T^m_{1,k}) \right) \\
&= \frac{1}{K \cdot N \cdot M} \sum_{k=1}^{K} \sum_{i=1}^{N} \sum_{m=1}^{M} \left(f(\omega_0) - f(\omega_1) \right) \\
&= f(\omega_0) - f(\omega_1).
\end{aligned}$$

However, in practice, this is often not the case and there will be interactions between the treatment and the effect of match length realization. Both case studies (Lugovskyy et al. (2017) and Embrey et al. (2019)) have illustrated that point. What can researchers do if this orthogonality condition is not given and if - as a consequence - fixing the sequence of match length realizations is not enough?

Diagnosing the problem A first step towards dealing with potential bias in treatment effects induced by match length realizations is to diagnose it. As we have seen in Section 2 there are some regularities in the effect of match length realizations on cooperation that allow identifying and bounding the effect. The most important regularity is that cooperation rates seem to be monotonically increasing in match length, i.e. $f(\omega, T - x) \leq f(\omega, T) \leq f(\omega, T + x), \forall x \in \mathbb{N}$. If monotonicity is indeed given, then it is possible for the researcher to get an idea of how strong the relationship between match length realization and cooperation is in the environment(s) they consider. This can be done by using "very different" match length realizations (across sessions or matching silos) and compare the effect sizes observed for these. Note that in order to do this, it is best if the same sequence of match length realizations is drawn across treatments. More precisely, under monotonicity and equal match length realizations across treatments it is possible to

[30] Sometimes the focus is only on a subset of matches. This could easily be incorporated in the arguments below. It is more common to focus on initial cooperation in each match, but often people also consider average cooperation across the T^m stages of match m. The arguments below apply irrespective of which of these cases is considered.

bound the treatment effect by comparing the average treatment effect under the (initially) shortest realizations (ATE$^-$) and the treatment effect under the (initially) longest realizations (ATE$^+$). As we know the distribution from which match lengths are drawn, we can further compute the probability that match length realizations are even shorter (longer) than the shortest (longest) observed match lengths. Denote this probability by q^- (q^+) and define $q = q^- + q^+$. The researcher can then make statements of the form "With probability $(1 - q)$ the average treatment effect is between ATE$^-$ and ATE$^+$". See Imbens and Angrist (2004) or De Quidt et al. (2018) for a more detailed discussion.[31]

Fixing the problem We now outline some potential ways to "fix" the problem. It is important to note that while none of the proposed solutions is a perfect fix for any possible environment, many of them work in specific environments. Maybe even more importantly many of the solutions discussed can help with diagnosing the problem even if the treatment effect cannot be fully debiased.

The most immediate solution to the problem is probably to simply use as many match length realizations as possible. There are three basic ways to do so: (i) increase K, (ii) increase M or (iii) change the level at which randomization takes place. (i) is a great solution whenever it is feasible. If it is possible to obtain information on the expected size of the effect of match length realization (e.g. from prior literature) then power analysis can be conducted to determine how many sessions K are needed to obtain unbiased treatment effects. (ii) seems less promising. Increasing M works only to the extent that the effect of early match length realization washes out over time. As our analysis in Section 2.3 has shown this does not seem to be the case for match numbers that can reasonably fit in a two-hour experimental session. A possibility that we explored in this paper is (iii) when we randomized match length at the match level in our "MatchStoch treatment". A downside of this approach is that it can induce waiting times as all participants in a matching silo have to wait for the longest match to end before being rematched. This concern is especially important if δ is high and restricts the total number of matches that can be played.

There are some alternatives to the standard random termination method. In a method used e.g. by Sabater-Grande and Georgantzis (2002), Cabral et al. (2014) or Vespa (2019), a fixed (known) number of rounds are played with certainty, and payoffs in these rounds are discounted at a known rate δ. Afterwards, there is a fixed known probability δ that the match continues for an additional round, and payoffs in these rounds are no longer discounted. Andersson and Wengstroem (2012) and Cooper and Kuehn (2014) use a similar method that also starts with a fixed number of rounds with payoff discounting, but is then followed by the coordination game induced by considering only two strategies, "Grim Trigger" and "Always Defect". The first method avoids very short matches, but, as overall match length is still random, does not eliminate the problem. The second method does not have an uncertain match length, but it has the downside that the number of repeated game strategies allowed needs to be restricted ex ante. Further, Fréchette and Yuksel (2017) found that neither of these methods induces behaviour that is consistent with the presence of dynamic incentives.

Fréchette and Yuksel (2017) propose a promising method called block random termination. Participants play as in the standard method, but in blocks of a pre-announced fixed number of rounds. Within a block they receive no feedback about whether or not the match has continued until that round, but they make choices that will be payoff-relevant if it has. Once the end of a

[31] This can be done even if match length realizations differ across treatments as long as the intervals defined by the two sets of match length realizations do have some overlap, but the computation is not as straightforward in this case.

block is reached, subjects are told whether the match ended within that block and, if so, in what round; otherwise, they are told that the match has not ended yet, and they start a new block. Subjects are paid for rounds only up to the end of a match, and all decisions for subsequent rounds within that block are void. With block random termination the length of blocks is a crucial parameter which has to be set carefully.[32] An open question is whether participants learn mainly from the number of stages played or from the number of stages that are payoff relevant.

Other possibilities could include re-sampling approaches by selecting subsets of sessions or matching silos with different match length realizations. Another possibility is to use “inverse” designs (see Section 3.1) more systematically to pair each session with its inverse. Or one could start the experiment with a “training phase”, where participants are given the chance to learn about the distribution of match lengths by observing several realizations.[33] Last, one could use constraints on the realized empirical distribution by e.g. imposing that the mean match length of early matches cannot be more than one standard deviation away from the expected length.[34] In sum, there are many possibilities to deal with the problem of match length realizations. Which one is the most suitable will depend on the specific environment researchers are interested in studying.

5. Conclusions

We have seen that the realized length of early matches in indefinitely repeated games has a substantial impact on cooperation rates in subsequent matches. Using three cases from the literature we also demonstrated a strong impact on treatment comparisons, both in terms of the size and the direction of the treatment effect. Our results have important implications for our understanding of how people learn in infinitely repeated games, for the interpretation of treatment effects when there are stochastic elements, as well as for experimental design.

Theories of learning should take into account how agents learn from sequences of realizations of random variables. One interesting question in this context is which moment of the distribution of match length realizations is most important for learning. In empirical research indefinitely repeated games are implemented using a mean expected match length that derives from the discount factor in the infinitely repeated game considered. Appendix Table D.13 shows, however, that the median match length realization seems a more important determinant for participants’ behaviour than the mean.[35] This raises the question of which sequence of indefinitely repeated games “correctly” represents the infinitely repeated game one ultimately has in mind. This ques-

[32] If blocks are of length one, the method is the same as the standard random termination method. With very long blocks a downside is that the experiment lasts long and fewer matches can be played.

[33] One advantage of such a training phase is that - since there is no strategic interaction - match lengths can be randomized at the individual level for the training phase. This means that many match length realizations can be observed a least for this phase.

[34] Which constraints are effective will depend on the specific treatment comparisons the researcher is interested in. In Section 3.1 we have, for example, seen that cooperation rates are similar for all sequences where mean match length realizations in early matches are within one standard deviation of what we should expect theoretically. However, results in Section 3.2 differ even across sessions where this is the case.

[35] There is research in other contexts suggesting that the median experience might be relevant and that people understand information based on median/rank better especially when there is a lot of skewness in the distribution, which is also the case with match length realizations (Aldrovandi et al., 2015; Wood et al., 2012). One example is consumption of alcohol, where people seem to have a good sense of how their consumption compares to the median, but not to the mean (Wood et al., 2012).

tion has been answered theoretically under standard game theoretic assumptions. But, given how people seem to learn in these games, it might be necessary to rethink this question. For example, in our first case study (subsection 3.1), it is the match stochastic treatment, which, as expected, closely matches the mean match length of the deterministic case. The median match length of the deterministic case is better matched by the inverse treatment, though. Which of those is the more relevant comparison depends on which of these moments is more important for how people learn. If it indeed turns out that median match length is the key statistic determining learning, then future research in both theory and experiments is needed to build and test new models of learning which can accommodate this fact.

For applied work it is important to know that different learning experiences may lead to different behaviors and may be confounded with treatment effects. Not accounting for potential differences in learning experience may lead to falsely claiming effects when there are none or to not finding effects when there are. In Section 4 we have discussed at length which design features empirical studies might use to ensure their conclusions are less vulnerable to match length realization effects. Future research could select the most promising among these and systematically assess how well they work in practice.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2022.105416.

References

Aldrovandi, S., Brown, G., Wood, A., 2015. Social norms and rank-based nudging: changing willingness to pay for healthy food. J. Exp. Psychol., Appl. 21 (3), 242–254.

Andersson, O., Wengstroem, E., 2012. Credible communication and cooperation: experimental evidence from multi-stage games. J. Econ. Behav. Organ. 81 (1), 207–219.

Bernard, M., Fanning, J., Yuksel, S., 2018. Finding cooperators: sorting through repeated interaction. J. Econ. Behav. Organ. 147, 76–94.

Bigoni, M., Casari, M., Skrzypacz, A., Spagnolo, G., 2015. Time horizon and cooperation in continuous time. Econometrica 83, 587–616.

Blackburn, J., 1936. Acquisition of Skill: An Analysis of Learning Curves. IHRB Report 73.

Blonski, M., Ockenfels, P., Spagnolo, G., 2011. Equilibrium selection in the repeated prisoner's dilemma: axiomatic approach and experimental evidence. Am. Econ. J. Microecon. 3 (3), 164–192.

Blonski, M., Spagnolo, G., 2015. Prisoners' other dilemma. Int. J. Game Theory 44, 61–81.

Boergers, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. J. Econ. Theory 77, 1–14.

Cabral, L., Ozbay, E., Schotter, A., 2014. Intrinsic and instrumental reciprocity: an experimental study. Games Econ. Behav. 87, 100–121.

Cabrales, A., Garcia-Fontes, W., 2000. Estimating learning models from experimental data. UPF Economics and Business Working Paper 501.

Camera, G., Casari, M., 2009. Cooperation among strangers under the shadow of the future. Am. Econ. Rev. 99 (3), 979–1005.

Camerer, C., Ho, T., 1999. Experience-weighted attraction learning in normal form games. Econometrica 67 (4), 827–874.

Cheung, Y-W., Friedman, D., 1997. Individual learning in normal form games: some laboratory results. Games Econ. Behav. 19 (1), 46–76.

Cooper, D., Kuehn, K., 2014. Communication, renegotiation, and the scope for collusion. Am. Econ. J. Microecon. 6 (2), 247–278.

Dal Bó, P., 2005. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. Am. Econ. Rev. 95 (5), 1591–1604.

Dal Bó, P., Fréchette, G., 2011. The evolution of cooperation in infinitely repeated games: experimental evidence. Am. Econ. Rev. 101 (1), 411–429.
Dal Bó, P., Fréchette, G., 2018. On the determinants of cooperation in infinitely repeated games: a survey. J. Econ. Lit. 56 (1), 60–114.
De Quidt, J., Haushofer, J., Roth, C., 2018. Measuring and bounding experimenter demand. Am. Econ. Rev. 108 (11), 3266–3302.
Embrey, M., Fréchette, G., Yuksel, S., 2018. Cooperation in the finitely repeated prisoner's dilemma. Q. J. Econ. 301 (355).
Embrey, M., Fréchette, G.R., Yuksel, S., 2017. Cooperation in the finitely repeated prisoner's dilemma. Q. J. Econ. 133 (1), 509–551.
Embrey, M., Mengel, F., Peeters, R., 2019. Strategy revision opportunities and collusion. Exp. Econ. 22 (4), 834–856.
Engle-Warnick, J., Slonim, R.L., 2006. Learning to trust in indefinitely repeated games. Games Econ. Behav. 54 (1), 95–114.
Erev, I., Roth, A., 1998. Predicting how people play in games: reinforcement learning in experimental games with unique mixed strategy equilibria. Am. Econ. Rev. 88 (4), 848–881.
Erev, I., Roth, A., 2001. Simple reinforcement learning models and reciprocation in the prisoner's dilemma game. In: Gigerenzer, G., Selten, R. (Eds.), Bounded Rationality: The Adaptive Toolbox. MIT Press, pp. 215–231.
Fréchette, G., Yuksel, S., 2017. Infinitely repeated games in the laboratory: four perspectives on discounting and random termination. Exp. Econ. 20 (2), 279–308.
Friedman, D., Oprea, R., 2012. A continuous dilemma. Am. Econ. Rev. 102 (1), 337–363.
Fudenberg, D., Karreskog, G., 2020. Learning about Initial Play Determines Average Cooperation in Repeated Games. Mimeo.
Fudenberg, D., Kreps, D.M., 1993. Learning mixed equilibria. Games Econ. Behav. 5 (3), 320–367.
Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games, vol. 2. MIT Press.
Hanaki, N., Sethi, R., Erev, I., Peterhansl, A., 2005. Learning strategies. J. Econ. Behav. Organ. 56 (4), 523–542.
Imbens, G., Angrist, J., 2004. Confidence intervals for partially identified parameters. Econometrica 72 (6), 1845–1857.
Ioannou, C., Romero, J., 2014. Learning with repeated-game strategies. Front. Neurosci.
Kagel, J.H., Levin, D., 1986. The winner's curse and public information in common value auctions. Am. Econ. Rev., 894–920.
Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., Williams, A., 2017. An experimental study of finitely and infinitely repeated linear public goods games. Games Econ. Behav. 102, 286–302.
Mengel, F., 2018. Risk and temptation: a meta-study on prisoner's dilemma games. Econ. J. 128, 3182–3209.
Mookherjee, D., Sopher, B., 1997. Learning and decision costs in experimental constant sum games. Games Econ. Behav. 19, 97–132.
Plott, C.R., Sunder, S., 1982. Efficiency of experimental security markets with insider information: an application of rational-expectations models. J. Polit. Econ. 90 (4), 663–698.
Roth, A.E., Murnighan, J.K., 1978. Equilibrium behavior and repeated play of the prisoner's dilemma. J. Math. Psychol. 17 (2), 189–198.
Rustichini, A., 1999. Optimal properties of stimulus—response learning models. Games Econ. Behav. 29, 1–2 244–273.
Sabater-Grande, G., Georgantzis, N., 2002. Accounting for risk aversion in repeated prisoners' dilemma games: an experimental test. J. Econ. Behav. Organ. 48 (1), 37–50.
Vespa, E., 2019. An experimental investigation of cooperation in the dynamic common pool game. Int. Econ. Rev. 61 (1), 417–440.
Vriend, N.J., 1997. Will reasoning improve learning? Econ. Lett. 55 (1), 9–18.
Wilcox, N., 2006. Theories of learning in games and heterogeneity bias. Econometrica 74 (5), 1271–1292.
Wood, A., Brown, G.D., Maltby, J., 2012. Social norm influences on evaluations of the risks associated with alcohol consumption: applying the rank-based decision by sampling model to health judgments. Alcohol Alcohol. 47 (1), 57–62.