History matters: Match length realizations and cooperation in indefinitely repeated games*

FRIEDERIKE MENGEL [†]
University of Essex
and Lund University

LUDOVICA ORLANDI[‡]
University of Essex

SIMON WEIDENHOLZER[§]
University of Essex

August 15, 2020

Abstract

Experimental studies of infinitely repeated games typically consist of several indefinitely repeated games ("matches") played in sequence with different partners each time, whereby the 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 & Fréchette, 2018) we demonstrate that the realized match length of early matches has a substantial impact on cooperation rates in subsequent matches. We show that the effect is in line with a class of learning models displaying the "power law of practice". 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 important implications for our understanding of how people learn in infinitely repeated games as well as for experimental design.

Keywords: Experiments, Indefinitely Repeated Games, Cooperation, Social Dilemmas .

JEL Classification Numbers: C70, C90. .

*We thank Maria Bigoni, Anna Dreber, Matt Embrey, Guillaume Frechette, Dan Friedman, Drew Fudenberg, Ed Hopkins, Volodymyr Lugovskyy, Ryan Oprea, Ronald Peeters, Giancarlo Spagnolo and Robert Sugden for helpful comments and for providing data access. We also thank audiences at the LEG 2019 workshop at Bar Ilan, the University of East Anglia, the University of Heidelberg, the SAET 2019 conference, and at the Vienna University of Economics and Business Adminstration for their helpful comments. Financial support by the British Academy (SG162637) is gratefully acknowledged.

[†]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ *and* Department of Economics, Lund University, Tycho Brahes vag 1, Lund (SE). *e-mail*: fr.mengel@gmail.com.

[‡]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ. *e-mail*: lorlan@essex.ac.uk.

[§]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ. *e-mail*: sweide@essex.ac.uk.

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 key insights on how repeated interaction changes incentives and can sometimes enable a wide 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 lab experimental 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.[1]

Infinitely repeated games are implemented in the lab as "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 the match ends. A match refers to a supergame (a "repeated game"), and a round (or stage) is one play of the stage game. To allow for learning, experiments typically feature repetitions. Participants play several matches and are rematched in between. The length of each match (in terms of number of rounds) is determined randomly, giving rise to a sequence of match length realizations. For a given termination probability, δ, participants may thus experience different realizations of match length and consequently different sequences of match lengths. This design is unproblematic if either the realized sequences of match length realizations "correctly" represent the infinitely repeated game or if match length realization is irrelevant for behaviour.[2] 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.[3] If the number of match length realizations is "small", it is difficult to know whether the realized distribution is "representative enough" for the infinitely repeated game under consideration, as this will depend on how people learn and which moments of the distribution are the most critical for their learning. For this reason it is 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 behaviour. Using a data set on the infinitely repeated prisoner dilemma compiled by Dal Bó and Fréchette (2018) we show that when subjects initially experienced relatively long matches subsequent cooperation rates are substantially higher. Specifically, when most matches in the first part of an experiment were "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.

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

[2]In standard theory only expected match length (captured by δ) should matter for behavior. Hence, according to standard theory match length realization should indeed be irrelevant for behaviour.

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

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 one 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 show that fairly simple reinforcement learning models in the spirit of Erev and Roth (1998) are consistent with these findings.

Our results hence show that the environment in which early interactions take place matters for subsequent interactions as people learn from match length realizations. As such our results speak to our understanding of what is often termed "cultural differences". They can help explain why people growing up in different social backgrounds (characterized by more or less stable interactions) or people coming from different work environments (characterized by more or less turnover) might show different behaviour in the exact same situation. Our results also show that we would not expect these groups to behave differently in *all* situation. If the environment is very unfavorable to cooperation we would not expect a behaviour difference between these groups based on our results in this paper. The results 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 do not only yield insights into how people learn in indefinitely repeated games, they also have important methodological implications for experimental design. The length of each match is typically drawn at the session level, 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: a continuous time prisoner's dilemma (Bigoni, Casari, Skrzypacz, & Spagnolo, 2015), a public good game (Lugovskyy, Puzzello, Sorensen, Walker, & Williams, 2017) and oligopoly games (Embrey, Mengel, & Peeters, 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 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, Fanning, and Yuksel 2018).[4] 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

[4]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.

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 discuss in detail the potential implications of these findings for measuring cooperation levels and for experimental design.

The paper is organized as follows. In Section 2 we demonstrate the main empirical finding of an effect of match length realization on cooperation. We also introduce the learning models and show that it can be explained by learning. Section 3 contains our discussion of the case-studies and Section 4 concludes. Additional theory, tables, figures and information on our own experiments can be found in an Appendix.

2 The Effect of Match Length Realizations

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 Figure 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 figure 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.

	C	D
C	R	S
D	T	P

PD game.

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

Normalized game.

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

GT vs. AD game.

Figure 1: Left: Prisoner's dilemma (PD) game with payoff parameters $\mathtt{T} > \mathtt{R} > \mathtt{P} > \mathtt{S}$ and $\mathtt{T} + \mathtt{S} < 2\mathtt{R}$. Middle: Normalized game where joint defection payoff $\mathtt{P}$ is subtracted from each cell and all payoffs are divided by $\mathtt{R} - \mathtt{P}$ (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.

A substantial experimental literature has studied how payoff parameters affect cooperation in the prisoner's dilemma.[5] One particularly successful approach, proposed by Blonski and Spagnolo (2015), analyzes a setting where agents can only choose among the repeated game 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 Figure 1 (see also Blonski, Ockenfels, and Spagnolo (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 SizeBAD, is defined as the threshold probability of choosing GT that has to be exceeded to make GT a best response.[6] 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 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 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?"

The remainder of this section is dedicated to answering this question. Doing so will not only yield valuable insights into how people learn, but it will also make an important methodological point regarding how indefinitely repeated games should be implemented in the lab. 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 show that they are in line with several simple learning models (Section 2.3).

2.2 Empirical Analysis

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 use data from 141 different sessions of indefinitely repeated prisoner's dilemma experiments with 2415 participants (see Table 3 in Dal Bó and Fréchette (2018)). In all papers contained in their data set the sequences of match length realizations is drawn at session level, i.e. all subjects in a given session faced the same sequence of match length realizations. Figure 2 shows the distribution of the difference

[5] See e.g. Embrey, Fréchette, and Yuksel (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.

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

between theoretical median match length and realized match length in the meta-study.[7] 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.

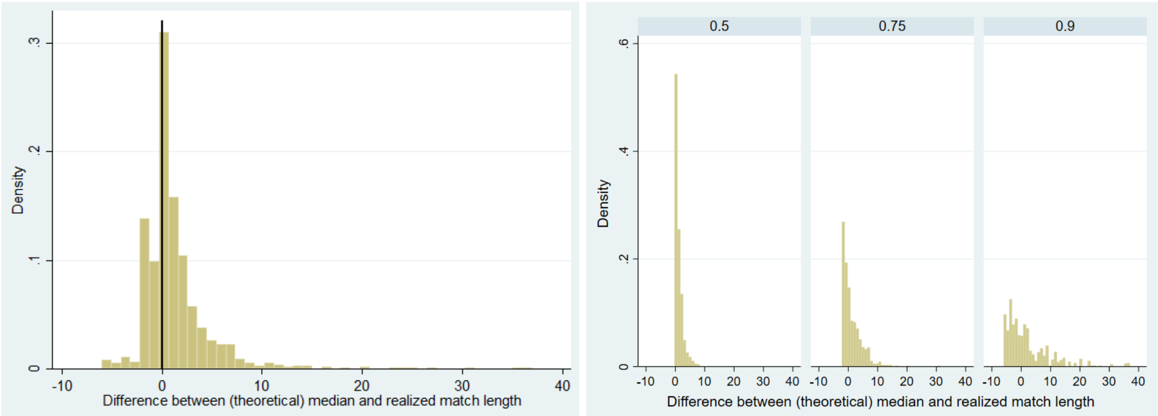

Figure 2: Distribution of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$.

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.[8] 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.

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

[7] Appendix Figure D.1 shows kernel density estimates.

[8] 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 C.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 C.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 C.6-C.7 show the results with alternative cutoffs.

	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<0.01, ** p<0.05, * p<0.1

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 141 sessions spread across 15 papers.

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 C.1 contains tables showing that these results are qualitatively robust to the inclusion of paper fixed effects (Table C.4), to considering different thresholds (Tables C.6-C.8) or to using the share of matches above median instead of a dummy variable (Table C.5). A placebo test shown in Appendix Table C.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.[9]

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

[9]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.

these matches. Early matches seem at least as important as recent matches and potentially, more important.[10]

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.[11] Appendix Figure D.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 law of practice" and according to Erev and Roth (1998) dates back to at least Blackburn (1936). In the next subsection we will outline a class of learning models that implies this property and show that the results above can be explained by such a learning model.

2.3 Learning

To study how match lengths can affect learning we discuss several 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 & Fréchette,

[10] Appendix Table C.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.

[11] 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 D.2.

2018; Embrey et al., 2017), we restrict attention to the GT and AD strategies.[12] The payoffs of the game induced by these strategies are given in the right panel of figure 1, where the expected match length $\mathbb{E}[T]$ is now replaced by the actual match length realization T^m. In the learning models 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 prepositions as well as initial beliefs. The models we study differ only in how these propensities are updated after each match.

Under *reinforcement learning without counterfactuals* (see e.g. Erev and Roth (1998) and Roth and Erev (1995)) agents increase the propensity of the strategy chosen by the payoff received, i.e.

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

where $s^{-i,m}$ denotes the strategy of i's opponent in match m, $\mathbb{1}(s^{i,m} = s)$ indicates whether agent i uses strategy s in match m or not, and $\pi(s, s^{-i,m}, T^m)$ gives the payoff earned with strategy s in this case.[13] Note that the propensity for strategies not chosen does not change. This can be seen as a pretty simplistic way of learning which is solely driven by ones' own experience. Match length realizations influence learning through their role in determining payoffs in the game. Note that this form of reinforcement learning features what Erev and Roth (1998) call "force of habit" where frequently chosen actions are reinforced more frequently. This is not the case under *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 also strategies that have not been played are reinforced and propensities for all strategies evolve according to

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

In environments were payoffs are stable, in the sense that they do not feature an exogenous stochastic element, reinforcement learning with counterfactuals is very 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 order to specify (smooth) *fictitious play* in the present environment we need to specify beliefs about play of the others as well as beliefs about match length realizations. Here we show a version where agent i simply uses the average previous match length realization, given by $\bar{T}^m = \frac{1}{m}\sum_{k=1}^{m} T^k$, and the share of her opponents previously choosing grim trigger, given by

[12] 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).

[13] While this cumulative formulation of reinforcement learning is more common in the learning literature, some authors have used average payoffs instead (Bigoni et al., 2015). We chose the former because it is more common. With the linear choice rule (see below) both choices will produce similar results.

$\sigma^{-i,m} = \frac{1}{m}\sum_{k=1}^{m}\mathbb{1}(s^{-i,k} = GT)$, where $\mathbb{1}(s^{-i,k} = GT)$ equals one if agent i's opponent had chosen GT in period k and is zero otherwise. We adopt a propensity based formulation of fictitious play as in Camerer and Ho (1999) and Hopkins (2002) where propensities are simply given by the expected payoffs $\hat{\pi}$ given these beliefs, i.e.

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

It remains to specify how these propensities translate into choices. For all three models we use Luce's linear probability choice rule where choice probabilities are linearly proportional to propensities.[14] More formally agent i's probability to choose GT in match m is given by

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

and AD is chosen with the remaining probability, i.e. $p_{AD}^{i,m} = 1 - p_{GT}^{i,m}$.

To quantify the effect of match length realizations on cooperation rates in these models, we have run several simulations for each of these three learning models, using key game parameters from the meta dataset of Dal Bó and Fréchette (2018)'s data. The payoff parameters of the underlying PD games were chosen to correspond to the 25th, 50th and 75th percentile of the distribution of `SizeBAD`. In each simulation run, 16 agents were matched against each other in a round robin tournament and match length realizations were drawn at the session level using the discount factor $\delta = 0.75$, thus targeting the median group size and the most frequent discount factor in Dal Bó and Fréchette (2018). Our exercise contains 4000 simulated experiments played over 15 matches for each of the three payoff configuration leading to 960,000 indefinitely repeated PD games for each of the three learning models. More details can be found in Appendix A.2.

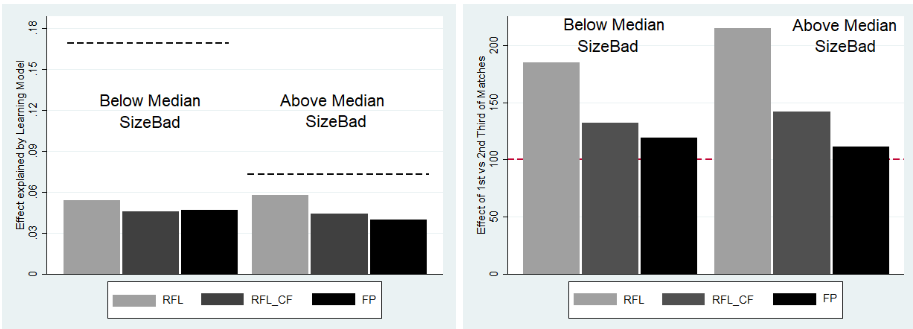

Figure 3: Left Panel: Effect Size for Δ_{above}^{1st} in simulated data for the three models (bars) compared to effect size in Dal Bó and Fréchette (2018) data (dashed line) in two samples: below median values of `SizeBAD` (left three bars and thick dashed line) and above median values of `SizeBAD` (right three bars and thick dashed line). Right Panel: Relative Effect size of Δ_{above}^{1st} and Δ_{above}^{2nd} in simulated data for the three models (bars) in two samples: below median values of `SizeBAD` (left three bars) and above median values of `SizeBAD` (right three bars). The dashed line indicates $\Delta_{above}^{1st} = \Delta_{above}^{2nd}$.

[14]This is a special case of the power probability form and is the most frequently used in the reinforcement learning literature. Under fictitious play the exponential or logit choice rule is more commonly used. Its main advantage lies in being able to be directly deduced from the maximization of an underlying perturbed payoff function (see e.g. Fudenberg and Levine (1998)).

For each of these three datasets we run regressions identical to the one shown in column (1) of Table 1, thus measuring the impact of long match length realizations in the first third of matches on initial cooperation rates in the remainder of matches. Standard errors are clustered at the run level. We then compare the effect size obtained in these regressions to the empirically observed one. As we have seen that there is an interaction effect between `SizeBAD` and the effect of early match length realization (Table 1) we also split the sample in below and above median `SizeBAD`.

The left panel of Figure 3 shows the results. Reinforcement learning (without counterfactuals) shows a slightly larger effect compared to the other models.[15] For all models the effect size does not differ across the two subsamples of low and high `SizeBAD` ($p > 0.1$). This differs from the human data, where for small values of SizeBAD early match length realization has a much larger effect size, though the difference is just outside of conventional levels of statistical significance ($p = 0.1753$). In this sample the effect size obtained purely through the learning models is about a third of the overall effect size. In the sample where defection is relatively attractive (high SizeBAD), the simulated effect size is about 80% of the empirical effect size with human players.

The right panel of Figure 3 shows the ratio between Δ^{1st}_{above} and Δ^{2nd}_{above} for our three learning models. It can be seen that this ratio is always above 100%, i.e. $\Delta^{1st}_{above} > \Delta^{2nd}_{above}$. This illustrates the consequences of the "power law of practice", which implies that early realizations of match lengths may be more important for cooperation rates than later experience.[16] To understand why this is the case,[17] note that in all models the probability of choosing GT and AD are determined by the ratio of their propensities. In the case of the reinforcement learning models these propensities are generally increasing over time and consequently higher in later matches. Thus, the impact of a long match and of the associated high payoffs on the choice probabilities will be larger early on when propensities are small than later when they are large (keeping the ratio of propensities and the play of the opponent constant). Thus, whatever happens early on shifts the pattern of play in a more pronounced way than later experiences.[18]

We have seen that all three learning models can explain the direction of the empirically observed effect of match length realization. However, we have also seen that effect sizes are larger empirically than what we would expect from the models. This is particularly the case when sizeBAD is small, i.e. when cooperation is relatively attractive. One possible explanation why effect sizes are larger with human players is that they may stop learning after some rounds, while the simulated learning models keep learning. Hence for the simulated learners the effect of initially long matches is (partially) corrected when later matches are shorter.

[15]Given the substantial sample size in the simulations all the effects obtained are highly statistically significant ($p < 0.0001$).

[16]This figure does not show effect sizes from human players as they are much larger than the scale of the figure if sizeBAD is small (below median). In this case the estimated effect size of Δ^{1st}_{above} is 0.169^{***} and for Δ^{2nd}_{above} it is 0.001. When sizeBAD is above median these effect sizes are much closer.

[17]We explore these arguments more formally in the appendix.

[18]Appendix Tables C.9 and C.10 contain the corresponding regression tables and show that indeed early matches are more important than recent matches in these models. There is no statistically significant difference, though, in the importance of early and more recent matches in the fictitious play model (Table C.11).

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).[19]

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.[20] 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.

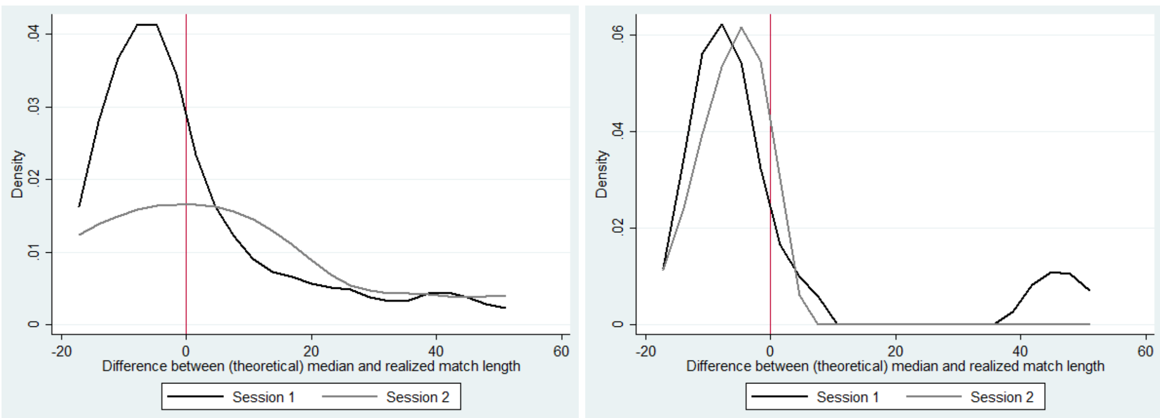

Figure 4: Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and for the 1st third of matches. 58% of all matches and 81% of matches in the 1st third were shorter than theoretical median match length.

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. Figure 4 shows kernel density estimates of the difference

[19]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.

[20]This important research programm 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.

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.

	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**

Table 2: 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 C.12 and C.13).

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 B 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 D.3 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 2 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 closely. Between the deterministic and replication treatment there is a 14.46 percentage point difference in average cooperation (Table 2) 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.[21] 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.

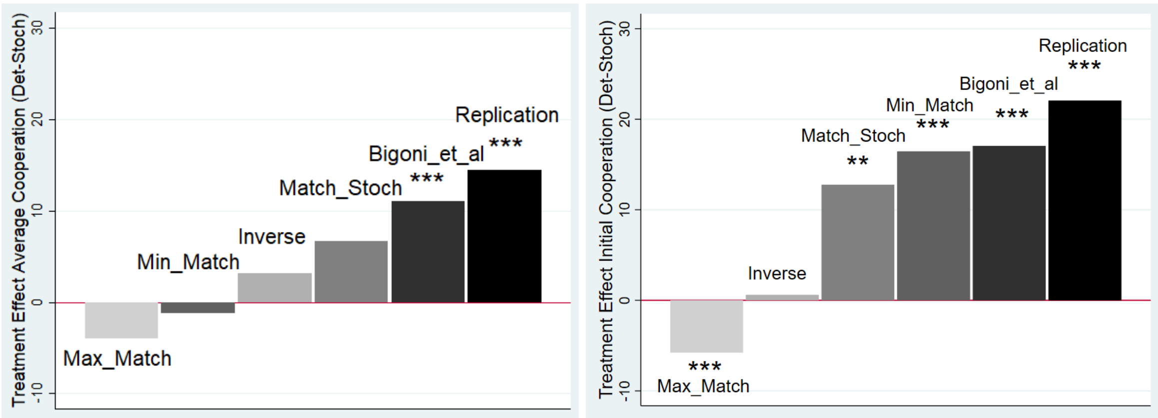

Figure 5: Different effect sizes obtained for the treatment differences between deterministic and stochastic treatments in average cooperation rates (left panel) and initial cooperation rates (right panel) depending on match length realizations. Bigoni et al: original effect size in Bigoni et al. (2015); 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;

Figure 5 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.

[21] 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.

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 in the treatment with "many match length realizations" (MatchStoch) the direction of the effect always goes in the same direction as in the original study.[22]

We will get back to the question of how to measure the "correct" treatment effect in Section 4. 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 equalling 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 hypotheses 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.

Table 3 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

[22] Appendix Figure D.4 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.

	All Rounds Cooperation Rate				
Decision Setting	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 3: 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.

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.[23] Appendix Table C.14 shows that similar conclusions hold when we consider first round cooperation rates only.

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

[23]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.

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 4 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.

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	$p = 0.6888$

Table 4: 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.$^{***}p < 0.01$, $^{**}, p < 0.05$, $^{*}p < 0.1$.

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.[24]

Table 4 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.[25] We 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 D.5 which shows the treatment effect for all possible selections of two match

[24] 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.

[25] Appendix Table C.15 shows all matches

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 D.5.

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 learning and experimental design in more detail.

4 Discussion and Conclusions

We have seen that the realized match length of early matches in indefinitely repeated games has a substantial impact on cooperation rates in subsequent matches. We also studied 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. As discussed in the Introduction our results have implications for policy evaluation and can shed some light on what is often termed "cultural differences". They can help explain why people growing up in different social backgrounds (characterized by more or less stable interactions) or people coming from different work environments (characterized by more or less turnover) might show different behaviour in the exact same situation. Our results also have important implications for our understanding of how people learn in infinitely repeated games as well as for experimental design. We now discuss these in turn.

Our results have shown that people do learn and what they learn is affected by match length realization. Models displaying the "power law of practice" are able to explain the key patterns in our results quite well and can explain part of the empirical effect of match length realization. One interesting question is which moment of the distribution of match length realizations is most important for learning. 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 C.16 shows, however, that the median match length realization seems a more important determinant for participants' behaviour than the mean.[26] This raises the question of which sequence of indefinitely repeated games "correctly" represents the infinitely repeated game one ultimately has in mind. This question has been answered theo-

[26]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, Brown, & Wood, 2015; Wood, Brown, & Maltby, 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).

retically 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.

We now discuss the consequences for experimental design when treatment comparisons can be affected by match length realizations. First note that drawing the same sequence of match length realizations for all treatments does not solve the issue. 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. As both of these case studies have illustrated this does not solve the issue, as there can be interactions of the treatment effect with match length realization.

The most obvious solution to the problem is to simply use as many match length realizations as possible, e.g. by randomizing match length at the match level, as we did in our "Match-Stoch treatment". The main downside of this approach is that it can induce waiting times as all participants in a session or matching silo have to wait for the longest match to end before being rematched. This also restricts the total number of matches that can be played overall. As the overall effect on length is increasing in the number of matched subjects, the effect on expected length can be potentially mitigated by matching silos at either the cost of shorter matches or possibly allowing for rematching of the same individuals. Instead of using many match length realization one could try and use many matches and hope 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.

What could be other possible solutions? Fréchette and Yuksel (2017) studied three alternatives to the standard random termination method used to examine infinitely repeated games in the laboratory. In a method used e.g. by Sabater-Grande and Georgantzis (2002), Cabral, Ozbay, and Schotter (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. After the rounds with certainty are played, there is a fixed known probability δ that the match continues for an additional round, and payoffs in these rounds are no longer discounted. 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 particular strategies in the infinitely repeated game - namely, the Grim trigger strategy and the strategy of always defecting. Fréchette and Yuksel (2017) also propose their own method - block random termination - in which subjects receive feedback about termination in blocks of rounds. Specifically, the play as in the standard method, but in blocks of a pre-announced fixed number of rounds. Within a block, subjects receive no feedback about whether or not the match has continued until that round, but they make choices that will be payoff-relevant, contingent on the match actually having reached that point. Once the end of a block is reached, subjects are told whether the match ended within that block and, if so, in what round; otherwise, they are told that the match has not ended yet, and they start a new block. Subjects are paid for rounds only up to the end of a match, and all decisions for subsequent rounds within that block are void.

Can these methods solve the problem of match length realizations? The first method could possibly mitigate the impact of match length realizations, as at least very short matches are avoided. On the other hand by systematically eliminating short matches, this could distort the results and in a social dilemma e.g. lead towards higher cooperation. In addition, as overall match length is still random, this method 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. This can be undesirable in some contexts. Further, Fréchette and Yuksel (2017) found that neither of these methods induces behaviour that is consistent with the presence of dynamic incentives. Last, with block random termination the length of blocks is crucial. If it is very short, then effect of match length realization will at best be mitigated.[27] With long blocks, there is the downside that the experiment lasts "unnecessarily" long and as a consequence only fewer matches can be played. An open question is whether the fact that many stages are played affects behaviour even if participants learn afterwards that some of them were payoff-irrelevant. Similar it could be that participants' behaviour reacts only to information on how many stages were payoff relevant as opposed to how many stages were played.[28] One possibility to address the latter concern might be to use the block random method and give feedback about match length realization only at the very end of the experiment. This, however, would make learning much harder for participants.

How should one then design experiments to get the "correct" treatment effect? By running sessions with "many" match length realizations - as in our match stochastic treatment - we should be able to adequately capture the average effect for the stochastic game. Having many match length realizations also allows to study how sensitive the specific results obtained are to match length realization. In many cases results may be robust to match length realization. If it is not possible to run "many" different sequences, the solution we recommend at the moment is to use "very different" match length realizations (across sessions or matching silos) and attempt to bound the treatment effect by comparing the (initially) shortest with the (initially) longest realizations. This still, of course, requires a good enough sample size for each match length realization. For future research, studying the other possibilities mentioned and finding out which method works best seems a worthwhile programme to us.

References

Aldrovandi, S., Brown, G., & Wood, A. (2015). Social norms and rank-based nudging: Changing willingness to pay for healthy food. *Journal of Experimental Psychology: Applied*, *21*(3), 242-254.

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

[27] If blocks are of length one the method is the same as the standard method

[28] Some indication that this might be the case is given by Fréchette and Yuksel (2017) themselves. They find that, even when the discount rate is the same, if interactions are expected to be longer defection increases and the use of the Grim strategy decreases.

Bernard, M., Fanning, J., & Yuksel, S. (2018). Finding cooperators: Sorting through repeated interaction. *Journal of Economic Behavior and Organization*, *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. *American Economic Journal: Microeconomics*, *3(3)*, 164-192.
Blonski, M., & Spagnolo, G. (2015). Prisoners' other dilemma. *International Journal of Game Theory*, *44*, 61-81.
Boergers, T., & Sarin, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, *77*, 1-14.
Cabral, L., Ozbay, E., & Schotter, A. (2014). Intrinsic and instrumental reciprocity:an experimental study. *Games and Economic Behavior*, *87*, 100-121.
Camera, G., & Casari, M. (2009). Cooperation among strangers under the shadow of the future. *American Economic Review*, *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 and Economic Behavior*, *19*(1), 46–76.
Cooper, D., & Kuehn, K. (2014). Communication, renegotiation, and the scope for collusion. *American Economic Journal: Microeconomics*, *6*(2), 247-278.
Dal Bó, P. (2005). Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economic Review*, *95*(5), 1591–1604.
Dal Bó, P., & Fréchette, G. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, *101(1)*, 411-429.
Dal Bó, P., & Fréchette, G. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, *56*(1), 60-114.
Embrey, M., Fréchette, G. R., & Yuksel, S. (2017). Cooperation in the finitely repeated prisoner's dilemma. *The Quarterly Journal of Economics*, *133*(1), 509–551.
Embrey, M., Mengel, F., & Peeters, R. (2019). Strategy revision opportunities and collusion. *Experimental Economics*, *22 (4)*, 834-856.
Engle-Warnick, J., & Slonim, R. L. (2006). Learning to trust in indefinitely repeated games. *Games and Economic Behavior*, *54*(1), 95–114.
Erev, I., & Roth, A. (1998). Predicting how people play in games: Reinforcement learning in experimental games with unique mixed strategy equiloibria. *The American Economic Review*, *88*(4), 848-881.
Fréchette, G., & Yuksel, S. (2017). Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics*, *20*(2), 279-308.
Friedman, D., & Oprea, R. (2012). A continuous dilemma. *American Economic Review*, *102*(1), 337–63.
Fudenberg, D., & Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, *5*(3), 320–367.
Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games* (Vol. 2). MIT press.
Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica*, *70*(6), 2141-2166.
Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., & Williams, A. (2017). An experimental study of finitely and infinitely repeated linear public goods games. *Games and Economic Behavior*, *102*, 286-302.
Mengel, F. (2018). Risk and temptation: A meta-study on prisoner's dilemma games. *Economic Journal*, *128*, 3182-3209.
Mookherjee, D., & Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior*, *19*, 97-132.
Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, *8*(1), 164–212.
Roth, A. E., & Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Pyschology*, *17*(2), 189-198.
Rustichini, A. (1999). Optimal properties of stimulus—response learning models. *Games and Economic Behavior*, *29*(1-2), 244–273.

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

Vespa, E. (2019). An experimental investigation of cooperation in the dynamic common pool game. *International Economic Review*, *61*(1), 417-440.

Vriend, N. J. (1997). Will reasoning improve learning? *Economics Letters*, *55*(1), 9–18.

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 and Alcoholism*, *47*(1), 57-62.

Online Appendix for "History Matters: Match Length Realizations and Cooperation in Indefinitely Repeated Games"

F. Mengel, L. Orlandi, S. Weidenholzer

Contents

A Additional details learning models

In this Appendix we provide additional results for the learning models discussed in Section 2.3 as well as more details on the simulations discussed in that Section.

A.1 Comparative Statics of Choice Probabilities

In a first step consider the reinforcement model without updating of counterfactuals. Assume that agent i chooses GT in match m and consequently earns a payoff of $\pi^{i,m}(GT, s^{-i,m}, T^m)$. Then the change in choice probabilities from match m to match $m+1$ is given by:

$$\Delta p_{GT}^{i,m} = p_{GT}^{i,m+1} - p_{GT}^{i,m} = \frac{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m} + \pi(GT, s^{-i,m}, T^m)} - \frac{\psi_{GT}^{i,m}}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m}}$$

We are interested in comparing the impact in early matches to later matches, keeping the environment constant. Thus, we keep the ratio of propensities fixed by setting $\psi_{AD}^{i,m} = \lambda \psi_{GT}^{i,m}$ where $\lambda > 0$ is a constant. Thus,

$$\Delta p_{GT}^{i,m} = \frac{\lambda \pi(GT, s^{-i,m}, T^m) \psi_{GT}^{i,m}}{\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m}\right)\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)\right)}$$

Generally, propensities will be larger in later matches. Thus, in order to compare the impact of earning some payoff early to earning the same payoff later on (for a fixed ratio of propensities), we consider the change in $\Delta p_{GT}^{i,m}$ as the level of propensities $\psi_{GT}^{i,m}$ varies:

$$\frac{\partial \Delta p_{GT}^{i,m}}{\partial \psi_{GT}^{i,m}} = -\frac{\lambda \psi_{GT}^{i,m}}{\left((1+\lambda)\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)\right)^2} < 0.$$

Thus, the larger propensities are, the smaller the impact of realized payoffs on choice probabilities will be.

We next turn to reinforcement learning with updating of counterfactuals where i also updates the propensity for AD. The change in choice probability from match m to match $m+1$ is now given by:

$$\Delta p_{GT}^{i,m} = p_{GT}^{i,m+1} - p_{GT}^{i,m} = \frac{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m)}{\psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \psi_{AD}^{i,m} + \pi(AD, s^{-i,m})} - \frac{\psi_{GT}^{i,m}}{\psi_{GT}^{i,m} + \psi_{AD}^{i,m}}$$

Again we set $\psi_{AD}^{i,m} = \lambda \psi_{GT}^{i,m}$. Thus,

$$\Delta p_{GT}^{i,m} = \frac{\lambda \pi(GT, s^{-i,m}, T^m) - \pi(AD, s^{-i,m}, T^m)}{(1+\lambda)\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \pi(AD, s^{-i,m}, T^m)\right)}$$

Note that this change is positive provided that $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) > 1/\lambda = \psi_{GT}^{i,m}/\psi_{AD}^{i,m}$ and is non-positive otherwise. Thus, for an agent to increase her probability of choosing GT in round $m+1$, GT had to earn proportionally more than AD than it previously did (as expressed by the ratio of propensities). We again consider the change in $\Delta p_{GT}^{i,m}$ as the level of propensities $\psi_{GT}^{i,m}$ changes:

$$\frac{\partial \Delta p_{GT}^{i,m}}{\partial \psi_{GT}^{i,m}} = -\frac{\lambda \pi(GT, s^{-i,m}, T^m) - \pi(AD, s^{-i,m}, T^m)}{\left(\psi_{GT}^{i,m} + \lambda \psi_{GT}^{i,m} + \pi(GT, s^{-i,m}, T^m) + \pi(AD, s^{-i,m}, T^m)\right)^2}.$$

Provided that $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) > 1/\lambda$, so that the probability of choosing GT increases in match m, this expression is negative. Thus, later experience will lead to smaller changes in choice probabilities. In the case where $\pi(GT, s^{-i,m}, T^m)/\pi(AD, s^{-i,m}, T^m) \leq 1/\lambda$ we have $\Delta p^{i,m} \leq 0$ and $\partial \Delta p^{i,m}_{GT}/\partial \psi^{i,m}_{GT} \geq 0$. Thus, again later experience leads to less pronounced changes in choice probability.

Finally, consider the case of fictitious play. Unlike the reinforcement learning models propensities are not necessarily increasing over time but are given by the expected payoffs which are determined by beliefs on the play of others, $\sigma^{-i,m}$, and beliefs on match length realization, $\bar{T}^m$). Choice probability is then given by

$$p^{i,m+1}_{GT} = \frac{\hat{\pi}(GT, \sigma^{-i,m}, \bar{T}^m)}{\hat{\pi}(GT, \sigma^{-i,m}, \bar{T}^m) + \hat{\pi}(AD, \sigma{-i,m}, \bar{T}^m)}$$

One can verify that $\partial p^{i,m+1}_{GT}/\partial \bar{T} > 0$, implying that agents are more likely to choose GT for longer expected match lengths (keeping beliefs on behavior of others fixed). As early matches enter beliefs both directly and indirectly (since they affect what happens in later matches) this model also displays the power law of practice, though, since this is a second order effect, to a lesser extent than the other models.

A.2 Further details on simulations

The payoff parameters of the underlying PD games were chosen to correspond to the 25th, 50th and 75th percentile of the distribution of SizeBAD. We follow Erev and Roth (1998) and have normalized payoffs of the underlying PD games to ensure payoffs and propensities are positive. Table A.1 reports the corresponding parameter values of Dal Bó and Fréchette (2018) alongside the values for R and P when S and T were normalized to 0 and 1, respectively.

In each simulation run, 16 agents were matched against each other in a round robin tournament and match length realizations were drawn at the session level according to the discount factor $\delta = 0.75$, thus targeting the median group size and the most frequent discount factor in Dal Bó and Fréchette (2018). Our exercise contains 4000 simulated experiments played over 15 matches for each of the three payoff configuration leading to 960 000 indefinitely repeated PD games for each of the three learning models.

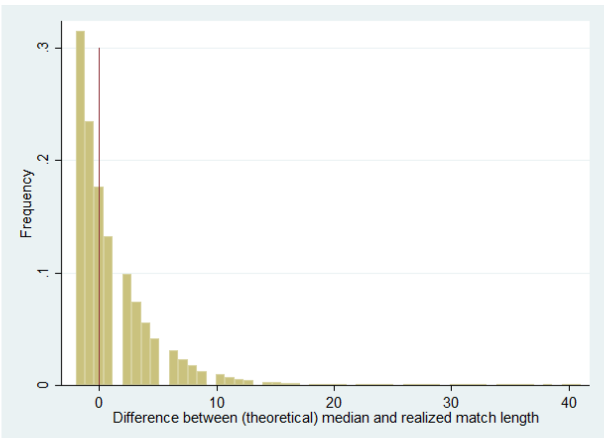

Figure A.1: Match Length Realization in our simulations.

The median match length in the simulated data is 3 and the mean 3.99 (as expected). The smallest realization is 1 and the largest 44. Appendix Figure A.1 shows the distribution of match lengths in our simulated data.

percentile	SizeBAD	ℓ	g	R	P
25	0.1625	0.5	0.4	0.789	0.263
50	0.2	0.56	1	0.609	0.219
75	0.667	1.85	2	0.588	0.381

Table A.1: Payoff parameters used for simulations.

In the reinforcement learning models, initial propensities for both strategies were set to the expected payoff when the match length corresponds to its mean and the opponent randomizes with equal probability among the two strategies, averaged across the two strategies, i.e. $\hat{\psi}_s^{i,0} = (\hat{\pi}(GT, \frac{1}{2}, T) + \hat{\pi}(AD, \frac{1}{2}, T))/2$. Note that for the reinforcement models this implies initial choice probabilities of $p_{GT}^{i,m} = p_{AD}^{i,m} = \frac{1}{2}$. Correspondingly, in the fictitious play model we started with initial belief on the strategy of the opponent of $\sigma^{-i,0} = 1/2$ and with initial belief on the match length of $\bar{T}^0 = \mathbb{E}[T]$.

B Additional Information Experiments

We conducted our own experiment to complete Case Study I (Section 3.1). In this Appendix we provide additional information on these experiments. Our experiments were conducted at Essex Lab at the University of Essex. The deterministic, replication and inverse treatments were conducted in February/March 2017 and the MatchStoch treatment in May 2019. Participants were students (and some non-students) who signed up for lab experiment at EssexLab at the University of Essex. They were recruited using recruitment software hroot. Table B.1 shows some demographics of our participants as well their answers to a post experimental questionnaire. We used exactly the same questionnaire as Bigoni et al. (2015), but replaced the question about whether people were born in Italy with whether they were born in the UK.

	Sample Characteristics			
	Rep	Inverse	MatchStoch	Det
Age	26.54	24.25	22.76	25.27
From UK (1=yes)	0.38	0.23	0.37	0.44
Gender (1=female)	0.47	0.44	0.42	0.50
Risk Attitude (0-10)	6.54	5.89	6.16	6.23
Trust (0-1)	0.38	0.39	0.33	0.42
Logic1 (0-1)	0.64	0.58	0.70	0.67
Logic2 (0-1)	0.33	0.31	0.52	0.40

Table B.1: Basic Sample Characteristics of Lab Experimental Sample used in Section 3.1. Mean Age, fraction of participants born in UK, fraction female, mean risk attitude (0, most risk averse, 10 least risk averse), fraction displaying high trust, fraction answering Logic 1 question correctly and fraction answering Logic 2 question correctly.

In all the experiments we followed exactly the same procedures used by Bigoni et al. (2015) including using the exact same Instructions and software (translated from Italian to English). The only change made to the software was to change the match length realizations (i) to be drawn at the match level in session MatchStoch and (ii) inversed in treatment "Inverse" (see Figure D.3). While in MatchStoch we draw match lengths within each session as the experiment progresses, we use the draws generated by Bigoni et al. (2015) in the "Replication" and "Inverse" treatments. As the interest of the study is match length realization, there was no other way to conduct an exact replication. Both procedures are in line with what was communicated to participants which is the following: "*How is a period duration established? The period may stop at every tick of 0.16 seconds. This event depends on the results of a random draw...*". No further information was given about when this draw took/will take place.

C Additional Tables

This Appendix contains additional tables.

C.1 Additional Tables for Section 2.2

In Table C.1 we consider separately the effect of match length realization of the 1st tenth, ninth, eight,..., half of matches on first-stage cooperation rates in the remaining matches. The table shows that there is a positive effect of match length realization in all these cases. However, the effect is smaller for the first tenth, ninth,...,sixth of matches compared to our baseline specification using the first third of matches. This shows that, while the very first matches are very important there is still learning and match length realizations become more important as more early matches are aggregated. Using the 1st half of matches, however, does not lead to a larger effect than using the 1st third.

1st...	(1) tenth	(2) ninth	(3) eight	(4) seventh	(5) sixth	(6) fifth	(7) fourth	(8) third	(9) half
Δ^{1st}_{above}	0.125**	0.081	0.020	0.055	0.023	0.078	0.181***	0.226***	0.165***
	(0.057)	(0.059)	(0.058)	(0.061)	(0.060)	(0.063)	(0.061)	(0.056)	(0.073)
SizeBAD	-0.592***	-0.692***	-0.781***	-0.815***	-0.829***	-0.744***	-0.628***	-0.539***	-0.688***
	(0.0726)	(0.0769)	(0.0967)	(0.0964)	(0.0867)	(0.116)	(0.120)	(0.0945)	(0.152)
SizeBAD $\times$ Δ^{1st}_{above}	-0.189**	-0.056	0.053	-0.066	-0.105	-0.028	-0.212***	-0.296***	-0.103
	(0.081)	(0.082)	(0.093)	(0.077)	(0.084)	(0.116)	(0.073)	(0.094)	(0.163)
Constant	0.852***	0.916***	1.003***	1.038***	1.051***	0.965***	0.843***	0.747***	0.886***
	(0.0574)	(0.0750)	(0.0977)	(0.0992)	(0.0901)	(0.118)	(0.123)	(0.0987)	(0.155)
Observations	45,869	45,469	44,773	43,955	43,481	41,873	38,778	34,319	25,795
R-squared	0.188	0.188	0.189	0.195	0.196	0.199	0.212	0.223	0.237

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.1: Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th of matches where above theoretical median length.

Table C.2 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. The coefficient Δ^{1st}_{above} is statistically significant in seven out of these nine different splits while all other dummies are statistically significant in fewer than 50 percent of the cases where they are included. Δ^{1st}_{above} also has the biggest coefficient size in all cases $X = 2, ..., 6$ and the second-biggest coefficient size when $X = 7, .., 10$, though it should be noted that these differences are not always statistically significant. Furthermore for all $X = 2, ..., 9$ the coefficient corresponding to the first X-th of matches is bigger than the one corresponding to the most recent (the $X - 1$th) X-th of matches. As the finer splits often involve only very few matches in each group we consider the cases $X = 2, ..., 6$ to be more meaningful. The table hence emphasizes the point that early match length realization is at least as important or more important than match length realization in more recent matches.

Table C.3 shows the results of a placebo test, where we regress cooperation in the 1st third of matches on Δ^{3rd}_{above} (as well as `SizeBAD` and an interaction). 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

1st...	(1) tenth	(2) ninth	(3) eight	(4) seventh	(5) sixth	(6) fifth	(7) fourth	(8) third	(9) half
Δ^{1st}_{above}	0.010	0.057*	0.048	0.082**	0.065*	0.077**	0.105***	0.101**	0.123***
	(0.037)	(0.031)	(0.033)	(0.035)	(0.034)	(0.039)	(0.036)	(0.041)	(0.044)
Δ^{2nd}_{above}	0.051	0.026	0.022	0.037	0.017	0.075**	0.097**	0.069*	
	(0.03)	(0.041)	(0.037)	(0.034)	(0.045)	(0.036)	(0.038)	(0.039)	
Δ^{3rd}_{above}	-0.014	-0.015	0.019	0.087**	0.048	0.076*	0.074**		
	(0.039)	(0.039)	(0.038)	(0.037)	(0.036)	(0.040)	(0.035)		
Δ^{4th}_{above}	0.070*	0.103***	0.017	0.003	0.050	0.075**			
	(0.038)	(0.035)	(0.037)	(0.034)	(0.037)	(0.036)			
Δ^{5th}_{above}	-0.000	0.001	0.091**	0.006	0.014				
	(0.035)	(0.032)	(0.040)	(0.045)	(0.044)				
Δ^{6th}_{above}	0.033	0.044	0.047	0.001					
	(0.033)	(0.043)	(0.039)	(0.039)					
Δ^{7th}_{above}	-0.004	0.012	0.022						
	(0.041)	(0.040)	(0.034)						
Δ^{8th}_{above}	-0.025	0.030							
	(0.033)	(0.037)							
Δ^{9th}_{above}	0.050								
	(0.037)								
SizeBAD	-0.824***	-0.814***	-0.806***	-0.806***	-0.809***	-0.805***	-0.805***	-0.809***	-0.769***
	(0.0808)	(0.0729)	(0.0772)	(0.0875)	(0.0796)	(0.0777)	(0.0734)	(0.0681)	(0.0670)
Constant	0.989***	0.966***	0.965***	0.969***	0.977***	0.973***	0.983***	0.994***	0.967***
	(0.0940)	(0.0789)	(0.0836)	(0.0926)	(0.0852)	(0.0823)	(0.0798)	(0.0747)	(0.0736)
Observations	5,905	7,338	7,506	8,296	9,148	11,721	13,573	18,536	25,795
R-squared	0.257	0.273	0.266	0.269	0.262	0.268	0.265	0.251	0.236

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.2: Initial (first stage) cooperation rate in remaining matches depending on whether at least $\frac{2}{3}$ of first X-th, second X-th, third X-th,... of matches where above theoretical median length.

not significant. This 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.

	Placebo test		
	(1)	(2)	(3)
Δ^{3rd}_{above}	0.038	-0.054	-0.053
	(0.055)	(0.061)	(0.060)
SizeBAD		-0.643***	-0.650***
		(0.061)	(0.078)
SizeBAD × Δ^{3rd}_{above}			0.009
			(0.055)
Constant	0.365***	0.908***	0.915***
	(0.045)	(0.070)	(0.085)
δ f.e.	NO	YES	YES
Observations	29,467	29,467	29,467
R-squared	0.002	0.141	0.141

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.3: Placebo test. Initial (first stage) cooperation rate in the *1st* third of matches explained by dummy Δ^{3rd}_{above} indicating whether more than $\frac{2}{3}$ of matches in the 3rd third of the experiment were longer than the theoretical median match length.

Appendix Table C.4 reproduces our main results reported in Table 1 using paper fixed effects.

In Appendix Table C.5 we use the share of matches above theoretical median length instead of a dummy variable.

	Main Result with paper fixed effects					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.115**	0.120***	0.267***	0.120*	0.119***	0.247***
	(0.055)	(0.040)	(0.053)	(0.060)	(0.040)	(0.058)
Δ^{2nd}_{above}				0.025	0.068	0.058
				(0.067)	(0.042)	(0.061)
`SizeBAD`		-0.746***	-0.497***		-0.808***	-0.573***
		(0.075)	(0.084)		(0.075)	(0.148)
`SizeBAD` $\times$ Δ^{1st}_{above}			-0.328***			-0.291***
			(0.084)			(0.094)
`SizeBAD` $\times$ Δ^{2nd}_{above}						-0.015
						(0.132)
Constant	0.826***	1.006***	0.757***	0.227***	1.101***	0.946***
	(0.0147)	(0.0782)	(0.0869)	(0.0215)	(0.120)	(0.161)
δ f.e.	NO	YES	YES	NO	YES	YES
paper f.e.	YES	YES	YES	YES	YES	YES
Test Δ^{1st}_{above}=Δ^{2nd}_{above}	-	-	-	0.3477	0.4220	0.0508
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.104	0.239	0.245	0.122	0.274	0.278

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.4: 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 141 sessions spread across 15 papers. With paper fixed effects.

	Share of matches above median					
	(1)	(2)	(3)	(4)	(5)	(6)
$Share^{1st}_{above}$	0.218***	0.206**	0.309***	0.072	0.210**	0.310**
	(0.068)	(0.079)	(0.113)	(0.112)	(0.080)	(0.120)
$Share^{2nd}_{above}$				0.235**	0.151	0.111
				(0.112)	(0.094)	(0.130)
`SizeBAD`		-0.764***	-0.446		-0.813***	-0.573
		(0.0759)	(0.383)		(0.0754)	(0.424)
`SizeBAD` $\times$ $Share^{1st}_{above}$			-0.441			-0.434
			(0.490)			(0.498)
`SizeBAD` $\times$ $Share^{2nd}_{above}$						0.104
						(0.281)
Constant	0.290***	1.023***	0.706*	0.250***	1.041***	0.801*
	(0.031)	(0.079)	(0.384)	(0.030)	(0.078)	(0.425)
δ f.e.	NO	YES	YES	NO	YES	YES
Test $Share^{1st}_{above}$=$Share^{2nd}_{above}$	-	-	-	0.4465	0.6338	0.2631
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.023	0.236	0.237	0.044	0.271	0.271

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.5: Columns (1)-(3): Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by share 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 share of matches longer than theoretical median length in 1st and 2nd third of matches. Standard errors clustered at session level. Observations stem from 141 sessions spread across 15 papers.

In Appendix Table C.6 we conduct the same analysis as Table 1 with the only difference that Δ^{1st}_{above} takes the value 1 if more than half (rather than $\frac{2}{3}$) of early matches are above median. It is hence a much weaker test as match length realizations with more or less than half of matches above median need not differ by much. In the most extreme case they might differ by one round only, which is not much if matches are long, i.e. δ is high. On the other hand it should allow us to better identify the effect of shorter match length when $\delta = 0.5$, as there will be more variation in the dummy in this case. Indeed we do find that the dummy defined in this way is more effective in capturing differences when $\delta = 0.5$. Overall the effect of shorter matches reported in column (1) is now a 38% increase in cooperation rates compared to the 43% increase identified in Table 1. Table C.7 goes the opposite direction and shows results when a dummy is used indicating whether more than $\frac{3}{4}$ of early matches were above median. Again results are similar, even though this dummy takes the value 1 much less often.

	Reduced Threshold					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.120**	0.142***	-0.095	0.051	0.143**	-0.102
	(0.050)	(0.051)	(0.085)	(0.058)	(0.057)	(0.091)
Δ^{2nd}_{above}				0.157**	0.064	0.076
				(0.062)	(0.054)	(0.094)
`SizeBAD`		-0.792***	-1.967***		-0.826***	-2.039***
		(0.066)	(0.286)		(0.067)	(0.322)
`SizeBAD` $\times \Delta^{1st}_{above}$			1.184***			1.239***
			(0.279)			(0.321)
`SizeBAD` $\times \Delta^{2nd}_{above}$						-0.018
						(0.163)
Constant	0.320***	1.000***	2.175***	0.277***	1.011***	2.225***
	(0.032)	(0.072)	(0.287)	(0.032)	(0.073)	(0.324)
δ f.e.	NO	YES	YES	NO	YES	YES
Test Δ^{1st}_{above}=Δ^{2nd}_{above}	-	-	-	0.3178	0.2876	0.1933
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.014	0.219	0.223	0.036	0.249	0.254

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.6: Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy Δ^{1st}_{above} indicating whether more than $\frac{1}{2}$ of matches in the 1st third of the experiment were longer than the theoretical median match length.

In Appendix Table C.8 we ask whether unusually short early matches or unusually long early matches have a larger effect. To this end we reproduce Table 1 using now a dummy Δ^{1st}_{below} which takes the value 1 if more $\frac{2}{3}$ of early matches are below the theoretical median length. For the sessions identified by this dummy cooperation rates are 32% lower in specification (1), indicating that the overall effect of match length realizations is broadly symmetric.

	Increased Threshold					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.192***	0.054	0.147**	0.149**	0.044	0.139**
	(0.067)	(0.038)	(0.059)	(0.070)	(0.038)	(0.059)
Δ^{2nd}_{above}				0.181***	0.079**	0.112*
				(0.062)	(0.038)	(0.063)
`SizeBAD`		-0.743***	-0.671***		-0.766***	-0.661***
		(0.067)	(0.077)		(0.068)	(0.090)
`SizeBAD` $\times \Delta^{1st}_{above}$			-0.215***			-0.221**
			(0.081)			(0.085)
`SizeBAD` $\times \Delta^{2nd}_{above}$						-0.060
						(0.092)
Constant	0.346***	0.951***	0.879***	0.297***	0.951***	0.846***
	(0.030)	(0.073)	(0.082)	(0.029)	(0.074)	(0.095)
δ f.e.	NO	YES	YES	NO	YES	YES
Test $\Delta^{1st}_{above}=\Delta^{2nd}_{above}$	-	-	-	0.3477	0.4220	0.050
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.031	0.215	0.218	0.064	0.249	0.252

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.7: Reduced Threshold. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy Δ^{1st}_{above} indicating whether more than $\frac{3}{4}$ of matches in the 1st third of the experiment were longer than the theoretical median match length.

	Short matches					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{below}	-0.144***	-0.042	-0.145**	-0.063	-0.045	0.092
	(0.049)	(0.051)	(0.061)	(0.080)	(0.057)	(0.074)
Δ^{2nd}_{below}				-0.138	-0.041	0.040
				(0.083)	(0.060)	(0.102)
`SizeBAD`		-0.759***	-0.750***		-0.799***	-0.775***
		(0.066)	(0.067)		(0.068)	(0.071)
`SizeBAD` $\times \Delta^{1st}_{below}$			-1.093***			-0.768***
			(0.188)			(0.239)
`SizeBAD` $\times \Delta^{2nd}_{below}$						-0.228
						(0.152)
Constant	0.442***	1.010***	1.907***	0.468***	1.071***	1.824***
	(0.036)	(0.097)	(0.150)	(0.039)	(0.127)	(0.185)
δ f.e.	NO	NO	YES	NO	NO	YES
Test $\Delta^{1st}_{below}=\Delta^{2nd}_{below}$	-	-	-	0.6325	0.9650	0.6977
Observations	34,319	34,319	34,319	18,536	18,536	18,536
R-squared	0.018	0.213	0.215	0.033	0.242	0.244

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.8: Focus on Short Matches. Initial (first stage) cooperation rate in the 2nd and 3rd third of matches explained by dummy Δ^{1st}_{below} indicating whether more than $\frac{2}{3}$ of matches in the 1st third of the experiment were shorter than the theoretical median match length.

C.2 Additional Tables for Section 2.3

	Simulated Reinforcement Learning					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.040***	0.040***	0.042***	0.039***	0.039***	0.039***
	(0.002)	(0.003)	(0.004)	(0.002)	(0.002)	(0.005)
Δ^{2nd}_{above}				0.027***	0.026***	0.035***
				(0.003)	(0.002)	(0.004)
`SizeBAD`		-0.139***	-0.138***		-0.161***	-0.159***
		(0.003)	(0.003)		(0.003)	(0.004)
$\Delta^{1st}_{above}\times$ `SizeBAD`			-0.004			-0.000
			(0.011)			(0.012)
$\Delta^{2nd}_{above}\times$ `SizeBAD`						-0.025**
						(0.010)
Constant	0.415***	0.463***	0.463***	0.419***	0.475***	0.474***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$	-	-	-	0.0063	0.0003	0.0989
Observations	1,920,000	1,920,000	1,920,000	960,000	960,000	960,000
R-squared	0.002	0.013	0.013	0.002	0.017	0.017

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.9: Simulated Reinforcement Learning Model. 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.

	Simulated Reinforcement Learning with Counterfactuals					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.030***	0.029***	0.037***	0.025***	0.024***	0.029***
	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)
Δ^{2nd}_{above}				0.016***	0.018***	0.024***
				(0.001)	(0.001)	(0.002)
`SizeBAD`		-0.181***	-0.178***		-0.184***	-0.180***
		(0.001)	(0.001)		(0.002)	(0.001)
$\Delta^{1st}_{above}\times$ `SizeBAD`			-0.023***			-0.017***
			(0.003)			(0.003)
$\Delta^{2nd}_{above}\times$ `SizeBAD`						-0.017***
						(0.005)
Constant	0.399***	0.461***	0.460***	0.400***	0.463***	0.462***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$	-	-	-	0.0001	0.0000	0.0000
Observations	1,920,000	1,920,000	1,920,000	960,000	960,000	960,000
R-squared	0.010	0.200	0.201	0.010	0.207	0.207

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.10: Simulated Reinforcement Learning Model with Counterfactuals. 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.

	Simulated Fictitious Play Learning					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^{1st}_{above}	0.027***	0.029***	0.038***	0.021***	0.023***	0.030***
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.002)
Δ^{2nd}_{above}				0.023***	0.022***	0.029***
				(0.001)	(0.001)	(0.002)
`SizeBAD`		-0.178***	-0.175***		-0.189***	-0.185***
		(0.001)	(0.001)		(0.001)	(0.002)
$\Delta^{1st}_{above} \times$ `SizeBAD`			-0.025***			-0.019***
			(0.003)			(0.005)
$\Delta^{2nd}_{above} \times$ `SizeBAD`						-0.022***
						(0.005)
Constant	0.392***	0.449***	0.448***	0.395***	0.456***	0.454***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Test $\Delta^{1st}_{above} = \Delta^{2nd}_{above}$	-	-	-	0.2481	0.4042	0.9592
Observations	1,920,000	1,920,000	1,920,000	960,000	960,000	960,000
R-squared	0.007	0.128	0.128	0.009	0.149	0.149

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.11: Simulated Fictitious Play Learning Model. 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.

C.3 Additional Tables for Section 3.1

	Average Cooperation Rates			
	(1)	(2)	(3)	(4)
Replication	-14.46***	-14.72***		
	(4.722)	(4.955)		
Inverse	-3.142	-3.582		
	(3.555)	(3.548)		
Match Stochastic	-6.641	-6.956	-6.923	-7.210
	(5.154)	(5.059)	(5.093)	(4.993)
L.Duration		-0.022		0.013
		(0.045)		(0.054)
Constant	54.04***	55.31***	42.03***	42.43***
	(3.454)	(3.497)	(8.895)	(8.445)
Demographics	NO	NO	YES	YES
Observations	6,624	6,336	4,416	4,224
Number of id	288	288	192	192

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.12: Random Effects OLS regression of average cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions.

	Initial Cooperation Rates			
	(1)	(2)	(3)	(4)
Replication	-22.06***	-22.44***		
	(3.921)	(3.756)		
Inverse	-0.589	-0.538		
	(3.367)	(3.451)		
Match Stochastic	-12.77**	-13.11**	-12.94**	-13.23**
	(6.306)	(6.391)	(6.429)	(6.533)
L.Duration		0.061		0.119*
		(0.041)		(0.065)
Constant	73.96***	73.68***	57.61***	56.58***
	(3.144)	(3.369)	(12.20)	(11.97)
Demographics	NO	NO	YES	YES
Observations	6,624	6,336	4,416	4,224
Number of id	288	288	192	192

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.13: Random Effects OLS regression of initial cooperation rate on treatment dummies and covariates (age, gender, birthplace, nationality). Standard errors clustered at session level. Columns (3) and (4) only use data from match stochastic and deterministic sessions.

C.4 Additional Tables for Section 3.2

Decision Setting	First Round Cooperation Rate Finite	Prob All	Prob S1	Prob S23	S1 vs S23
$N = 4$, MPCR= 0.3	24.9	28.8	16.56	34.95	
		$<$	$>^{*}$	$<^{**}$	$p = 0.002$
$N = 4$, MPCR= 0.6	44.0	47.7	38.88	51.00	
		$<$	$>$	$<$	$p = 0.224$
$N = 2$, MPCR= 0.6	52.4	44.3	42.56	45.46	
		$>$	$>$	$>$	$p = 0.678$
$N = 2$, MPCR= 0.6, Binary	76.1	57.8	54.28	59.23	
		$>^{***}$	$>^{*}$	$>^{***}$	$p = 0.613$

Table C.14: Initial 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.

C.5 Additional Tables for Section 3.3

with RO	All groups	Short first matches	Other Groups	Short vs Other
Substitutes	16.5	9.0	20.1	
Complements	23.2	34.9	17.3	
Treatment Effect	- 6.7	-25.2**	2.8	$p = 0.1280$
without RO	All groups	Short first matches	Other Groups	Short vs Other
Substitutes	17.6	20.9	15.9	
Complements	10.9	16.2	8.1	
Treatment Effect	6.7*	4.7	7.8	$p = 0.8949$

Table C.15: Efficiency measure from Table 2 (all matches) in 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.

C.6 Additional Tables for Section 4

	$\delta = 0.5$				$\delta = 0.75$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
mean duration	0.021		-0.001	0.022	0.015*		-0.004	-0.002
	(0.070)		(0.077)	(0.075)	(0.008)		(0.016)	(0.012)
median duration		0.104*	0.105	0.087		0.020**	0.024*	0.035**
		(0.063)	(0.082)	(0.105)		(0.008)	(0.016)	(0.016)
SizeBAD	-0.662***	-0.704***	-0.705***	-0.642***	-0.825***	-0.847***	-0.848***	-0.965***
	(0.113)	(0.113)	(0.103)	(0.089)	(0.071)	(0.065)	(0.064)	(0.077)
Constant	0.665***	0.519***	0.521***	0.527**	0.776***	0.753***	0.759***	0.575***
	(0.150)	(0.174)	(0.151)	(0.201)	(0.061)	(0.062)	(0.061)	(0.098)
paper fixed effects	NO	NO	NO	YES	NO	NO	NO	YES
Observations	10,786	10,786	10,786	10,786	12,380	12,380	12,380	12,380
R-squared	0.151	0.160	0.160	0.173	0.185	0.189	0.189	0.225

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.16: First round cooperation in the 2nd and 3rd third of matches regressed on mean and median match length in the 1st third. We selected the two discounted factors with most papers in the meta-study (8 papers for $\delta = 0.75$ and 4 papers for $\delta = 0.5$).

D Additional Figures

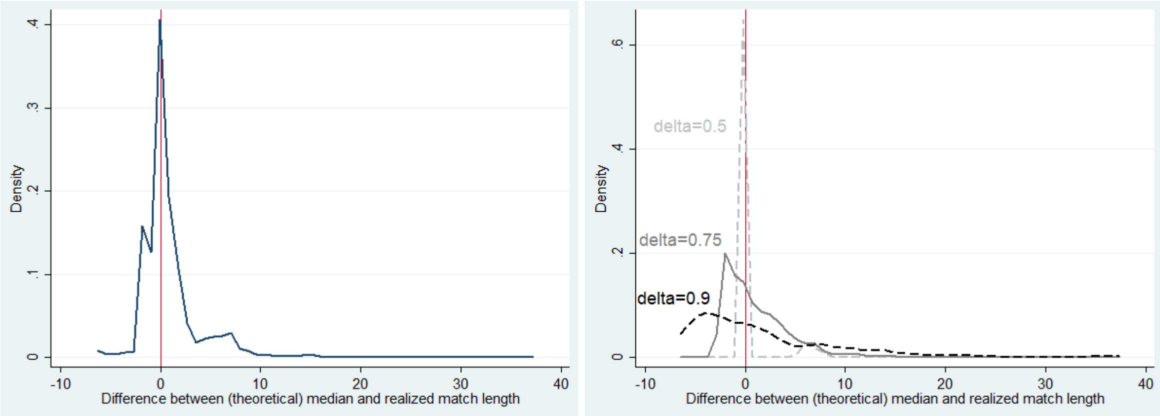

Figure D.1: Meta Data used in Section 2.2: Kernel density estimates of the difference between theoretical median match length and realized match length overall (left panel) and separately for $\delta = 0.5, 0.75$ and $\delta = 0.9$ (right panel).

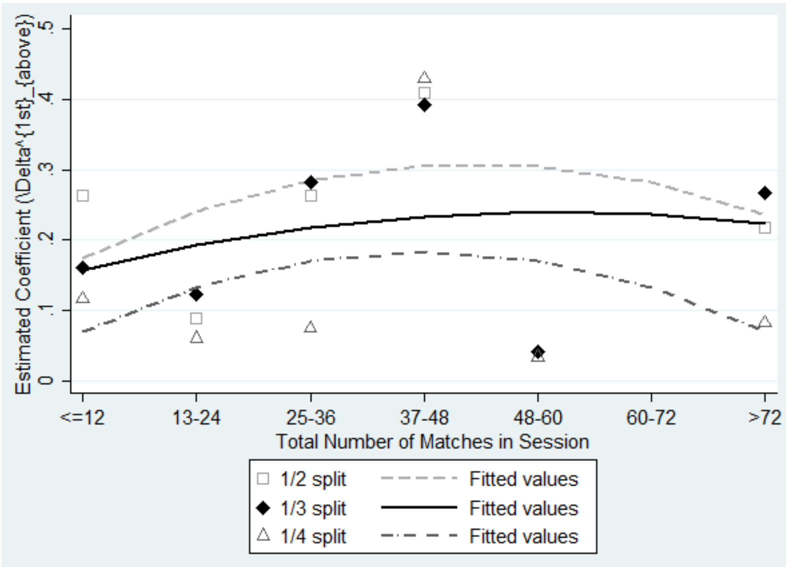

Figure D.2: Section 2.2 The Effect of Experience (Number of Matches Played in a Session) on Estimated Coefficient Δ^{1st}_{above}. Black diamonds show the estimated coefficient Δ^{1st}_{above} on cooperation in the last third of matches depending on the total number of matches played in the session. The black line shows fitted values using a square polynomial. Squares and light gray dashed line show the effect of the first half of matches on the second half and triangles and dark gray dashed line the effect of the first quarter on the last quarter of matches.

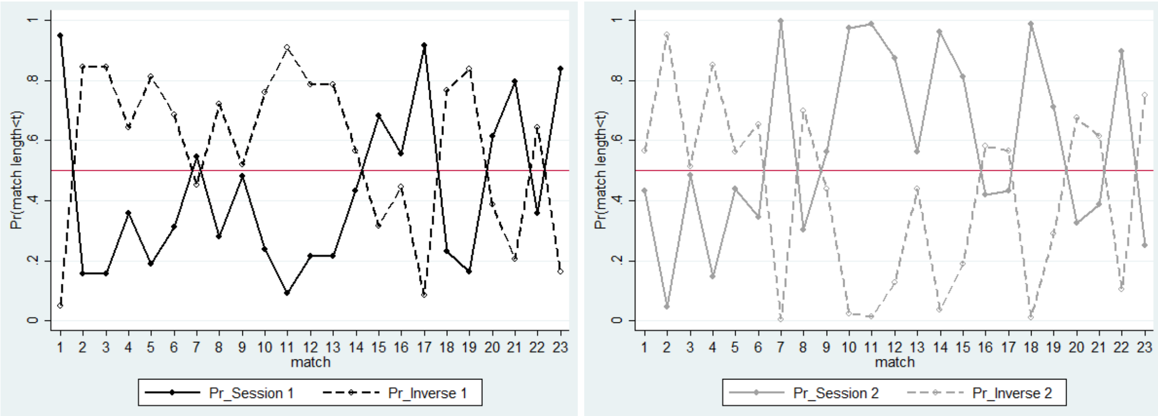

Figure D.3: Section 3.1: Illustration of how sequence of "inverse" match lengths is generated for the two sessions.

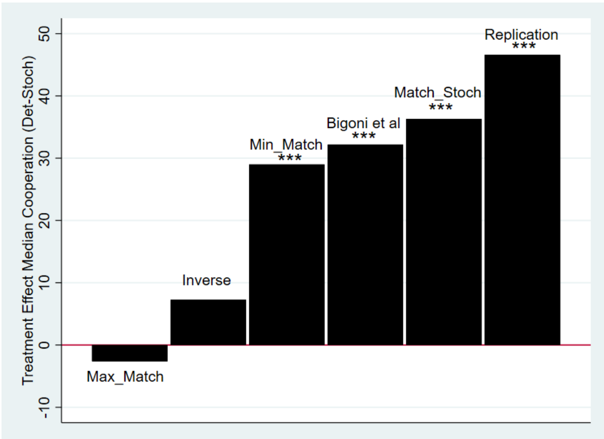

Figure D.4: Different Effect sizes Median Cooperation Frequency Det-Stoch.

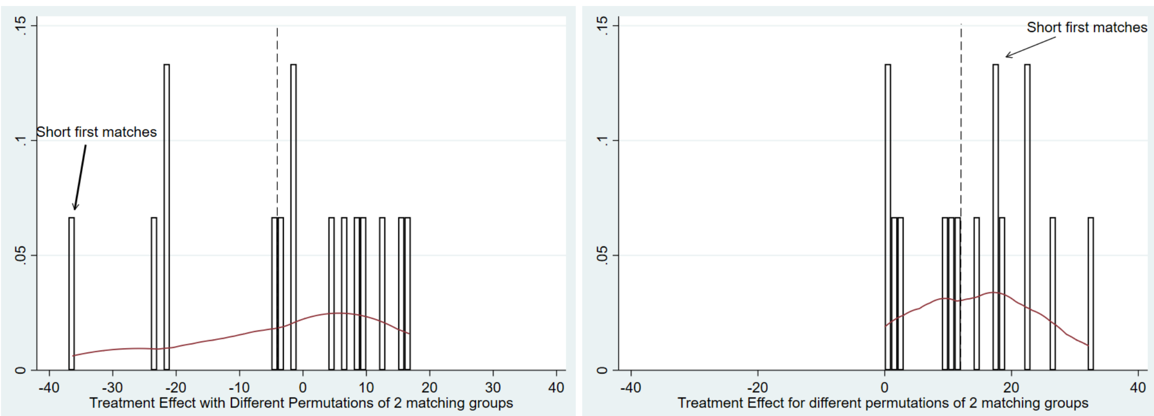

Figure D.5: Section 3.3: Illustration of treatment effect for different permutations of two matching groups. Histogram and estimated kernel density. Left Panel: treatments with revision opportunities. Right panel: treatments without revision opportunities.