Ghidoni, R., & Suetens, S. (2022). The effect of sequentiality on cooperation in repeated games. *American Economic Journal: Microeconomics*, *14*(4), 58-77.

https://doi.org/10.1257/mic.20200268

American Economic Journal: Microeconomics 2022, 14(4): 58–77
https://doi.org/10.1257/mic.20200268

The Effect of Sequentiality on Cooperation in Repeated Games[†]

By Riccardo Ghidoni and Sigrid Suetens*

Sequentiality of moves in an infinitely repeated prisoner's dilemma does not change the conditions under which mutual cooperation can be supported in equilibrium relative to simultaneous decision-making. The nature of the interaction is different, however, given that sequential play reduces strategic uncertainty. We show in an experiment that this has large consequences for behavior. We find that with intermediate incentives to cooperate, sequentiality increases the cooperation rate by around 40 percentage points, whereas with very low or very high incentives to cooperate, cooperation rates are respectively very low or very high in both settings. (*JEL* C72, C73)

Folk theorems show that both opportunism and cooperation can be sustained in a prisoner's dilemma game when the interaction is repeated and players are sufficiently patient (Fudenberg and Maskin 1986). A remarkable property of this setup is that whether players move simultaneously or sequentially in the stage game does not affect the conditions that support mutual cooperation in equilibrium.[1] In both cases, mutual cooperation can be sustained if the discount factor is above a

* Ghidoni: Department of Economics, University of Bologna, and Department of Economics, CentER, Tilburg University (email: riccardo.ghidoni@unibo.it); Suetens: Department of Economics, CentER, Tilburg University (email: s.suetens@tilburguniversity.edu). Leeat Yariv was coeditor for this article. We acknowledge financial support from the Netherlands Organization for Scientific Research (NWO) through the VIDI program (project number 452-11-012). Data and codes are provided in Ghidoni and Suetens (2022). An earlier version of this paper was circulated as "Empirical Evidence on Repeated Sequential Games." We thank the three anonymous reviewers for their excellent comments and encouragement. We also thank Maria Bigoni, Matthew Embrey, Dan Friedman, Eline van der Heijden, Jan Potters, and Boris van Leeuwen, seminar audiences at the University of Bologna, the University of the Düsseldorf, the University of Heidelberg, the University of Maastricht and Tilburg University, and participants of iSee 2018 at NYU Abu Dhabi, SEET 2018 at the University of Salento, LEG 2018 at Lund University, the 2018 ESA world meeting in Berlin, the 2018 RECent conference at the University of Modena, IMEBESS 2019 at Utrecht University, M-BEES 2019 at Maastricht University, EEA-ESEM 2019 at the University of Manchester, and the GRASS 2019 Workshop at the University of Milano-Bicocca, for insightful discussions and comments.

[†] Go to https://doi.org/10.1257/mic.20200268 to visit the article page for additional materials and author disclosure statements or to comment in the online discussion forum.

[1] Sequential moves, whereby the first mover's choice is revealed to the second mover before the latter makes a choice, are common in the context of trust (Kreps 1990), borrower-lender relations (Thomas and Worrall 1990; Kehoe and Levine 1993), employer-employee relations (Akerlof 1982; Fehr, Kirchsteiger, and Riedl 1993), and trade (Greif 1993; Brown, Falk, and Fehr 2004).

threshold that depends on the parameters of the game (Wen 2002).[2] Yet, given that sequentiality reduces strategic risk for the player who moves second, it creates a very different strategic environment. Specifically, by cooperating if and only if the first mover cooperates, the second mover can reap the benefits of cooperation and at the same time avoid being betrayed. If the first mover understands this, then the strategic risk he faces is also lower than that of a player in a simultaneous game. Consequently, one might plausibly expect that sequentiality is a key determinant of cooperation. This paper reports on a controlled experiment that studies whether and under what conditions sequentiality leads to more cooperation. The paper is relevant for understanding cooperation across a wide range of applications (e.g., trade, employer-employee relations, borrower-lender relations) and contributes to the literature that investigates the determinants of cooperation.

Strategic uncertainty has been highlighted as a crucial determinant of behavior within the class of repeated simultaneous prisoner's dilemmas (PDs). As summarized by Dal Bó and Fréchette (2018), the more money a player might lose by cooperating, the less she is willing to cooperate.[3] Two distinct but related approaches formalize the role of strategic uncertainty: Blonski, Ockenfels, and Spagnolo (2011) and Blonski and Spagnolo (2015), who apply the concept of risk dominance to the repeated PD, and Dal Bó and Fréchette (2011), who appeal to the basin of attraction of repeated-game strategies. These approaches help to formalize the intuition that sequentiality of moves may facilitate cooperation. A key element is that the second mover in a repeated sequential PD can, unlike a player in a repeated simultaneous PD, avoid ending up with the *sucker* payoff, by conditionally cooperating. This leads to the prediction that second movers conditionally cooperate and first movers cooperate whenever mutual cooperation is supported in equilibrium, and otherwise, they defect.[4] In contrast, with simultaneous decision-making, the approaches predict a smooth relation between payoffs and the likelihood of cooperation, conditional on mutual cooperation being supported in equilibrium. In summary, if strategic uncertainty is taken into account, the cooperation rate in sequential PDs is predicted to be (weakly) higher than that in simultaneous PDs in games in which mutual cooperation is supported in equilibrium.

In our experiment, participants play a series of indefinitely repeated sequential or simultaneous PDs. In each round, players proceed to the next round with a constant and known continuation probability δ.[5] The experiment covers six parameter configurations that vary between subjects, as in Dal Bó and Fréchette (2011). In one configuration, cooperation cannot be sustained in equilibrium because δ is below

[2] This builds on the use of the grim trigger strategy as a cooperative strategy (Friedman 1971). Since that strategy leads to minimax payoffs (equal to static Nash payoffs) independently of sequentiality, it is the harshest punishment strategy in both settings (Fudenberg and Maskin 1986).

[3] Strategic uncertainty is also an important factor in finitely repeated PDs (Embrey, Fréchette, and Yuksel 2018), repeated entry games (Calford and Oprea 2017), and dynamic dilemma games (Vespa and Wilson 2019).

[4] The prediction is reminiscent of a case discussed by Camera, Casari, and Bigoni (2013) in relation to a game where strangers decide whether to help one another in exchange for fiat money. In this case, the only two stable population configurations are a population of defectors and a population of conditional cooperators (traders), with basins of attraction depending on the parameters of the game.

[5] Building upon the assumption that participants do not discount the future in the short period of time they are in the lab, δ has the same role as that of the rate at which risk-neutral players discount the future in an infinitely repeated game (Roth and Murnighan 1978).

the threshold of the standard theory of infinitely repeated games, while in the other configurations, δ is above the theoretical threshold. We formulate predictions taking into account strategic uncertainty. In the treatment in which mutual cooperation cannot be sustained in equilibrium, no difference is predicted between the sequential and simultaneous versions. In the other treatments, sequentiality is predicted to (weakly) increase the cooperation rate to above that in the simultaneous equivalent, with the largest effect predicted for the games where δ is closest to the theoretical threshold. The reason for this is that strategic risk is largest in the simultaneous version of the latter games.

The experimental results show strong support for these predictions. In the treatments that are characterized by relatively high strategic risk in the simultaneous version, sequentiality increases the cooperation rate by 40 percentage points. In the treatments with relatively little strategic risk, sequentiality has no significant effect on the cooperation rate; the cooperation rate is close to 1 then when mutual cooperation is sustainable, and close to 0 otherwise.

Other experimental studies have compared sequential and simultaneous social dilemma games. Evidence from one-shot experiments, in which repeated-game incentives are absent, indicates that the effect of sequentiality on cooperation appears to depend on the game's parameters and the subject pool (Ahn, Ostrom, and Walker 2003; Ahn et al. 2007; Khadjavi and Lange 2013). Oskamp (1974), who compares repeated sequential- and simultaneous-move PDs with different payoff *levels* but otherwise the same repeated-game incentives, finds evidence for an interaction between sequentiality of moves and payoff levels. In sequential-move games, cooperation rates tend to be less responsive to a change in the payoff level than in simultaneous-move games.[6] Furthermore, there is a literature on leading-by-example where a leader is modeled as the first mover in a voluntary-contributions setting. For example, Potters, Sefton, and Vesterlund (2005) find that exogenously imposed sequentiality of moves increases contributions relative to a simultaneous-move setting if the leader has private information about the game's parameters. Yet, results are mixed in full information settings (for example, Andreoni, Brown, and Vesterlund 2002; Güth et al. 2007).[7] Finally, Kartal and Müller (2021) compare simultaneous and sequential infinitely repeated PDs in an experiment inspired by a model with heterogeneity in cooperation preferences and private information. They focus on a case in which cooperation cannot be sustained in equilibrium and find that sequentiality increases the cooperation rate by about 20 percentage points.

The remainder of the paper is organized as follows. Section I provides the theoretical background. Section II includes the experimental design, procedures, and hypotheses. Section III presents the main results, and Section IV concludes.

[6] In these experiments it was announced that the repeated game would last for 60 rounds but was actually ended after 50 to avoid end-game effects.

[7] See also Clark and Sefton (2001), who study the effect of stakes and subject pool on the cooperation rate in one-shot sequential PDs; Engle-Warnick and Slonim (2006), who study behavior in infinitely repeated trust games; and Reuben and Suetens (2012), who elicit stage-game strategies of players in infinitely repeated sequential PDs in which players can condition their strategy on whether they are playing the last round.

Table 1—Stage Game of a Simultaneous PD

	Cooperate	Defect
Cooperate	c, c	s, t
Defect	t, s	d, d

Note: $t > c > d > s$ and $2c > t + s$.

I. Theoretical Background

In a repeated simultaneous PD with a stage game as shown in Table 1, the standard theory of infinitely repeated games prescribes that mutual cooperation can be supported as an equilibrium outcome if $\delta \geq \delta^* \equiv (t - c)/(t - d)$ (see Proposition 4 in Friedman 1971). Both players playing grim trigger (GT) strategies constitutes an equilibrium then.[8] If the PD is played sequentially, then the theory predicts that mutual cooperation can be supported in equilibrium under the same condition as in the simultaneous PD, that is, if $\delta \geq \delta^*$. Likewise, GT leads to the harshest possible punishment and both players using a GT strategy constitutes an equilibrium (see online Appendix Section C.1 for calculations).[9] In summary, according to standard game theory, cooperation rates should not be different in sequential PDs than in simultaneous ones: if $\delta < \delta^*$, the only equilibrium is one in which both players defect, and if $\delta \geq \delta^*$, cooperative and noncooperative equilibria exist in both settings.

More precise predictions can be obtained by appealing to risk dominance (Blonski, Ockenfels, and Spagnolo 2011) or to the basin of attraction of repeated-game strategies (Dal Bó and Fréchette 2011). These approaches help to identify under which conditions players are more likely to coordinate on a mutually cooperative equilibrium in games with $\delta \geq \delta^*$. A key element is that the relative cost of cooperating with a partner who defects becomes an important determinant of behavior for players who do not know with certainty whether their partner will defect. In particular, consider a simplification of the repeated game to a game in which players choose at the beginning of the repeated game between the always defect strategy (AD) and a conditionally cooperative strategy (CC) à la GT.[10] We assume that the payoffs in the reduced game represent utilities and that they are common knowledge. The size of the basin of attraction of AD versus CC (referred to as SizeBAD) is defined as the maximum probability of the partner choosing CC such that playing AD is still a best response. SizeBAD turns out to be highly useful in understanding how behavior in sequential PDs might differ from that in simultaneous PDs. In what follows, we

[8] GT is defined as follows: start by cooperating and continue to do so if both players cooperate, and if one of the players defects, switch to defection forever.

[9] For a second mover, GT is implemented as follows: cooperate if the first mover cooperates, and if one of the players defects, switch to defection forever.

[10] Since players are assumed to choose their strategy at the beginning of the repeated game, tit-for-tat (TFT) or another conditionally cooperative strategy with limited punishment would also qualify as CC.

explain the intuition. The detailed calculations are presented in online Appendix Section C.2.

Consider first a repeated simultaneous PD. If $\delta \geq \delta^*$, the reduced game in which players choose between AD and CC is a game with two pure-strategy equilibria: (AD, AD) and (CC, CC). Players are more likely to choose CC and thus to end up in equilibrium (CC, CC) if the expected payoff of CC exceeds that of AD. This holds true if they believe that their partner will choose CC with a sufficiently high probability, namely with a probability that exceeds $\frac{d-s}{c+d-t-s+\frac{\delta(c-d)}{1-\delta}} \equiv \bar{p}$. The threshold belief $\bar{p}$, which we refer to as SizeBAD, depends on the game's parameters and decreases *ceteris paribus* as c or δ increases. Thus, it is predicted that for $\delta \geq \delta^*$, the likelihood that participants cooperate depends on the game's parameters. It is predicted to be higher, the higher is c or δ. For $\delta < \delta^*$, the cooperation rate is predicted to be zero.

Consider next a repeated sequential PD. If $\delta \geq \delta^*$, the expected payoff for the second mover of choosing CC is (weakly) larger than that of choosing AD for *all* possible beliefs about the strategy of the first mover. This is because, in contrast to a player in a simultaneous PD, a second mover who uses CC avoids the sucker payoff. She prefers CC if the discounted payoff of CC is higher than that of AD, namely if $\delta \geq \delta^*$, and plays AD if $\delta < \delta^*$.[11] The first mover is not confronted with strategic uncertainty either, because he anticipates that the second mover will conditionally cooperate (due to the assumption that the PD's payoffs represent the utilities of the players and that is common knowledge). Therefore, the first mover imitates the strategy of the second mover and also plays CC if $\delta \geq \delta^*$ and AD if $\delta < \delta^*$. Therefore, it is predicted that the cooperation rate will be equal to 100 percent if $\delta \geq \delta^*$ and 0 otherwise.[12] In summary, the cooperation rate in a repeated sequential PD is predicted to be (weakly) higher than in the repeated simultaneous PD with corresponding parameters. In Section II, we formulate more precise comparative-static predictions for the parameters used in the experiment.

Finally, allowing for heterogeneity of players, for example in terms of other-regarding preferences, does not change the core prediction that the cooperation rate in a sequential PD is (weakly) higher than in the simultaneous version.[13] However, if players have heterogeneous preferences, then the threshold above which CC is preferred over AD is player-specific. For example, sufficiently prosocial players would prefer CC over AD in the role of second mover in a sequential PD even if $\delta < \delta^*$, whereas relatively spiteful players would need a larger δ than δ^* to prefer CC over AD. Thus, for a given distribution of selfish, prosocial, and spiteful players in the population, the cooperation rate depends on the parameters of the game, even

[11] She is indifferent if $\delta = \delta^*$.

[12] Notice that the same predictions hold in the limit of a quantal response equilibrium, as noise completely vanishes (Turocy 1995). If noise has not vanished, then a smooth relation is predicted between the parameters of the game and the cooperation rate, even in Seq if $\delta > \delta^*$ (see online Appendix Section C.3 for predictions based on quantal responses).

[13] A large literature shows that players are heterogeneous in that at least some of them hold pro- or antisocial preferences (e.g., Levine 1998; Fehr and Schmidt 1999; Charness and Rabin 2002). For them, payoffs in PDs do not represent utilities. Ahn, Ostrom, and Walker (2003) and Ahn et al. (2007) illustrate how heterogeneity models help to understand cooperation in one-shot simultaneous and sequential PDs.

in sequential PDs with $\delta > \delta^*$. In online Appendix Section C.4, we illustrate the effect of heterogeneity using a Charness and Rabin (2002) utility function without a reciprocity component. A heterogeneity model with privately informed players can be found in Kartal and Müller (2021).

II. The Experiment

A. *Design and Procedures*

Participants in the experiment play 50 repeated games. The number of periods in a repeated game (referred to as rounds) is stochastic and ex ante unknown to both the participants and the experimenter. In each round, the (known) probability that the game proceeds to the next round is δ. At the beginning of each repeated game, participants are randomly divided into pairs within matching groups of ten. They remain matched with the same counterpart for all rounds of a repeated game. In the sequential PDs, participants are randomly allocated the role of first or second mover at the beginning of each repeated game. We expected that letting participants play in both roles helps them understand the strategic nature of the game.[14] The software had a built-in history box that participants could use to review all previous actions in the current repeated game.

We use the same parameters and between-subjects treatment variations as in the simultaneous PD experiment conducted by Dal Bó and Fréchette (2011)—henceforth, DBF: $d = 25$, $t = 50$, $s = 12$; $c = 32$, $c = 40$ *or* $c = 48$; and $\delta = 0.5$ *or* $\delta = 0.75$. These parameters cover a variety of settings ranging (in expectation) from short games with low cooperation gains to longer games with high cooperation gains. Table 2 presents an overview of the treatments, where Sim and Seq refer to the treatments with simultaneous and sequential moves, respectively. As can be seen from the table, both the average lengths of the repeated games and the share of repeated games that last just one round are in line with expectations.

The experiment was programmed with zTree (Fischbacher 2007) and conducted at the LINEEX lab in Valencia between July 2017 and April 2018. Sessions lasted 106 minutes on average and participants earned on average €22.7. The procedures are described in more detail in online Appendix Section A, and an English translation of the instructions can be found in online Appendix Section B.[15]

B. *Predictions*

Table 3 provides an overview of the values of SizeBAD for all treatments based on the assumption that PD payoffs represent utilities. The larger the difference in SizeBAD between two particular treatments, the larger is the expected difference in cooperation between them. Taking into account that DBF have already shown that

[14] Reassigning roles at the beginning of each repeated game also ensures that contagion effects à la Kandori (1992) are constant across simultaneous and sequential treatments.

[15] We also ran treatments in which the strategy method was used to elicit choices of second movers, and we plan to use these data in a future paper that compares hot and cold decision-making in sequential PDs.

TABLE 2—THE TREATMENTS

	Sim						Seq						Total
	$\delta = 0.5$			$\delta = 0.75$			$\delta = 0.5$			$\delta = 0.75$			
$c =$	32	40	48	32	40	48	32	40	48	32	40	48	
Number of participants	30	30	30	30	30	30	60	60	60	60	60	60	540
Number of matching groups	3	3	3	3	3	3	6	6	6	6	6	6	54
Number of repeated games	50	50	50	50	50	50	50	50	50	50	50	50	600
Number of rounds (mean)	1.8	1.9	1.9	4.1	4.1	4.1	1.8	1.8	1.8	4.3	4.3	3.3	–
One-round games (share)	0.60	0.54	0.54	0.24	0.26	0.26	0.54	0.54	0.54	0.25	0.25	0.26	–

Notes: Sessions were conducted with 40, 50, or 60 participants and treatments were distributed across several sessions. Apart from one exception, matching groups in a session faced the same δ and the same style of decision-making but a different c.

TABLE 3—SizeBAD BY TREATMENT

	Panel A. Sim					*Panel B. Seq*			
		c					c		
		32	40	48			32	40	48
δ	0.5	1	0.72	0.38	δ	0.5	1	0	0
	0.75	0.81	0.27	0.16		0.75	0	0	0

Notes: The table indicates the basin of attraction of AD (SizeBAD) in the different treatments. SizeBAD is defined as the maximal probability of the partner following a CC strategy that makes AD optimal.

the cooperation rate is close to 1 in Sim with $c = 48, \delta = 0.75$, and close to 0 with $c = 32, \delta = 0.5$, we can summarize the predictions as follows:

(i) The cooperation rate is expected to be close to 0 in Sim and Seq in treatment $c = 32, \delta = 0.5$.

(ii) The cooperation rate is expected to be close to 1 in Sim and Seq in treatment $c = 48, \delta = 0.75$.

(iii) The cooperation rate is expected to be (weakly) higher in Seq than in Sim in the other treatments, and the difference in cooperation rate is expected to (weakly) increase with the difference in SizeBAD between Seq and Sim: $c = 32, \delta = 0.75 \leq c = 40, \delta = 0.5 \leq c = 48, \delta = 0.5 \leq c = 40, \delta = 0.75$.

III. Results

A. *Effect of Sequentiality on Cooperation Rates*

This section reports the treatment effects of sequentiality on cooperation rates. We focus on cooperation rates across first rounds because (i) repeated games may have different lengths and (ii) the adopted theoretical framework involves the choice of whether to use a cooperative or noncooperative strategy at the beginning of the repeated game.[16] We first focus on comparative-static results after learning has taken place and then discuss dynamic patterns.

Figure 1 shows first-round cooperation rates across the last 20 repeated games. We find that the difference between Sim and Seq is small in the treatments with the lowest or highest incentive to cooperate ($p = 0.200$ and $p = 0.635$, respectively).[17] The cooperation rate is respectively close to 0 and close to 1 in these two treatments. In the Seq treatments with intermediate incentives to cooperate, the cooperation rate is substantially higher than in the corresponding Sim treatments. In particular, in treatments $\delta = 0.5, c = 40$; $\delta = 0.5, c = 48$; and $\delta = 0.75, c = 32$, sequentiality increases the post-learning cooperation rate by 38 to 41 percentage points ($p < 0.001$, $p < 0.001$, and $p = 0.015$, respectively). In the Seq treatment with $\delta = 0.75, c = 40$, the cooperation rate is somewhat higher than in the corresponding Sim treatment but the difference is not statistically significant ($p = 0.639$). Therefore, patterns of cooperation are overall closely in line with the SizeBAD predictions.[18]

The results are robust to controlling for individual-level variables, such as proxies for other-regarding preferences, risk preferences and proneness to mistakes, and experienced length of the first ten repeated games (see online Appendix Table F.2).[19] The results are also robust to a re-estimation of treatment effects on the basis of a dataset in which our data are merged with the data of DBF (see online Appendix Table F.3).

With respect to the effects of c and δ on the cooperation rate, Figure 1 shows that we have replicated the result of DBF that an increase in c or δ generally leads to an increase in the cooperation rate in Sim after learning. A similar effect is also

[16] Statistics and graphs based on all rounds are included in online Appendix Sections F and G, respectively. Patterns are generally very similar to those reported in the main text.

[17] Unless otherwise mentioned, the statistics reported in the results section are based on pairwise treatment comparisons of behavior in the last 20 repeated games using probit regressions. The regressions take the choice to cooperate in the first round of a repeated game as the dependent variable and include a treatment dummy as an independent variable. Standard errors are clustered at the matching-group level. Estimated treatment effects on the cooperation rate are presented in detail in online Appendix Tables F.1 and F.4.

[18] See online Appendix Figure G.1 for graphs that include predicted cooperation rates in Seq and cooperation rates observed in DBF's simultaneous PDs (data provided in Dal Bó and Fréchette 2019). As can be seen, DBF cooperation rates generally fall within 95 percent confidence intervals of the cooperation rates in Sim in our experiment, suggesting that the patterns are robust to changes in language, subject pool, and small differences in the procedures.

[19] Overall, we find a positive relation in the first rounds between prosociality and risk-loving on the one hand and cooperation on the other whereas our proxy for proneness to mistakes is less related to cooperation. We also find that, in line with, for example, Engle-Warnick and Slonim (2006) and Dal Bó and Fréchette (2018), the difference between expected and median realized length of the first ten repeated games has a positive effect on cooperation after learning.

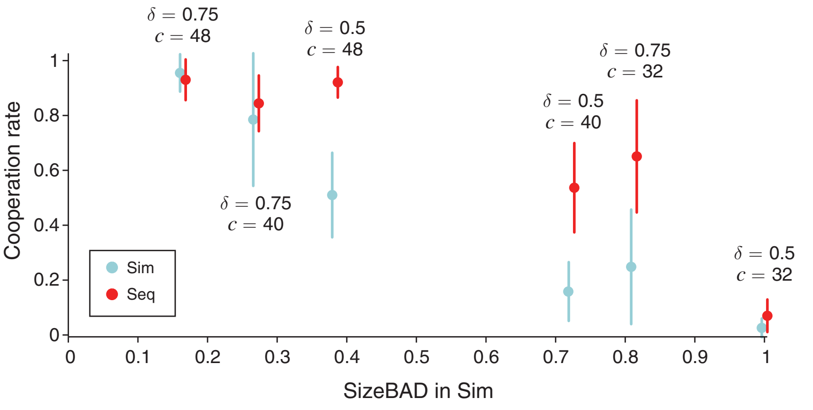

Figure 1. Cooperation Rates

Notes: The graph shows first-round cooperation rates and 95 percent confidence intervals based on predictions from probit regressions ran on treatment indicators with clustered standard errors at the matching-group level. Based on the last 20 repeated games.

observed in Seq, even when focusing solely on the treatments with $\delta > \delta^*$. In both Sim and Seq with $\delta > \delta^*$, the effect of c and δ on cooperation is statistically significant ($p \leq 0.01$ in probit regressions; see online Appendix Table F.5). Although such an effect is not predicted in Seq among rational payoff-maximizing players, it is consistent with the notion that players make mistakes, as in a quantal response equilibrium. It is also consistent with players being heterogeneous, for example in terms of social preferences, as outlined at the end of Section I.

We now turn to the learning dynamics. Figure 2 illustrates how first-round cooperation rates evolve across the fifty repeated games for each treatment. The graphs show that some learning is necessary before the treatment effects reported above set in. In the treatment with $\delta = 0.5, c = 32$, in which cooperation cannot be sustained in equilibrium, the cooperation rate is first well above 0 and then sharply declines to a rate close to 0, whereas in the treatments in which SizeBAD predicts a cooperation rate of 1, the cooperation rate increases across games. In Sim, the cooperation rate increases substantially only in treatments $\delta = 0.75, c = 40$ and $\delta = 0.75, c = 48$, which are both characterized by a low SizeBAD, and shows a decaying trend in the treatments with a higher SizeBAD.[20]

[20] Probit regressions with standard errors clustered at the matching-group level corroborate the result. For each treatment, we regress the first-round cooperation choice on a time trend. In Seq, the average marginal effect is positive and statistically significant for $\delta > \delta^*$ ($p \leq 0.021$) and negative and significant for $\delta < \delta^*$ ($p < 0.001$). In Sim, a positive and significant effect is obtained for $\delta = 0.75, c = 40$ and $\delta = 0.75, c = 48$ ($p = 0.021$ and $p < 0.001$, respectively), while the effect is negative and significant for $\delta = 0.5, c = 32$ and $\delta = 0.5, c = 40$ ($p < 0.001$ and $p < 0.001$, respectively). The effect is not statistically significant for $\delta = 0.5, c = 48$ and $\delta = 0.75, c = 32$ ($p = 0.182$ and $p = 0.748$, respectively). Patterns by matching group are shown in online Appendix Figure G.3.

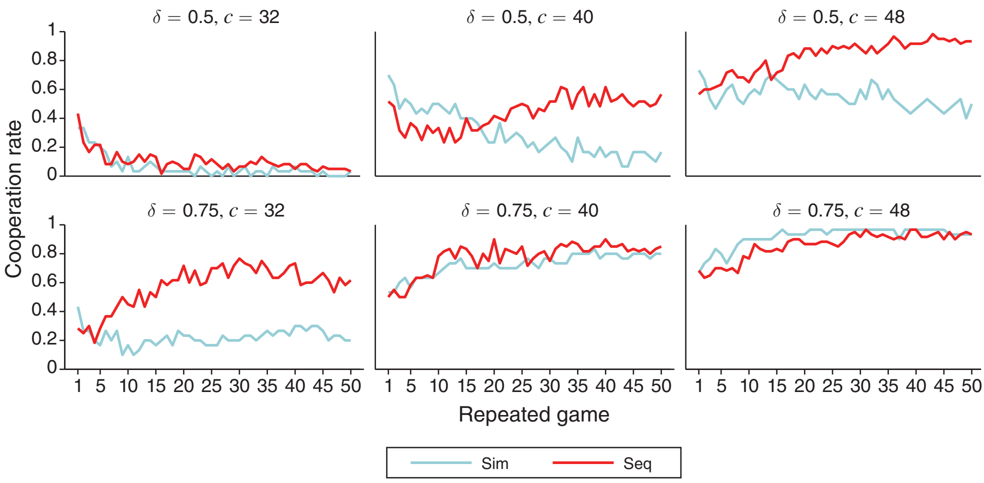

Figure 2. Evolution of Cooperation Rates

Note: The graphs show cooperation rates across first rounds of repeated games by treatment.

B. *Cooperation Rates by Role in the Sequential PDs*

In this section, we further study what drives cooperation in the sequential PDs after learning. Figure 3 splits up the cooperation rate in Seq by role according to the first-mover cooperation rate, the second-mover cooperation rate conditional on cooperation by the first mover (which we shall refer to as the conditional cooperation rate), and the second-mover cooperation rate conditional on defection by the first mover. The first observation is that the conditional cooperation rate among second movers ranges from 43.9 to 95.4 percent depending on the treatment, and it is in all treatments significantly higher than the cooperation rate conditional on the first mover defecting ($p < 0.001$). Overall, second movers rarely cooperate if the matched first mover defects. This provides support for our focus on conditional cooperation as the most important cooperative strategy for second movers.

The second observation is that in the three treatments in which the difference in SizeBAD between Seq and Sim is highest, the first-mover cooperation rate and the second-mover conditional cooperation rate in Seq are both higher than the cooperation rate in Sim ($p \leq 0.007$ and $p < 0.001$, respectively). This supports a key feature of the SizeBAD predictions, namely that sequentiality does not just reduce strategic uncertainty for second movers relative to players who move simultaneously, but also for first movers. Such an effect is not observed in treatments $\delta = 0.75, c = 48$ and $\delta = 0.75, c = 40$, in which differences in SizeBAD between Sim and Seq are low ($p \geq 0.357$ for first movers and $p \geq 0.320$ for second movers). In treatment $\delta = 0.5, c = 32$, in which cooperation cannot be supported in equilibrium, the second-mover conditional cooperation rate is substantially higher than the cooperation rate in Sim ($p < 0.001$), while the first-mover cooperation rate is only weakly higher ($p = 0.079$).

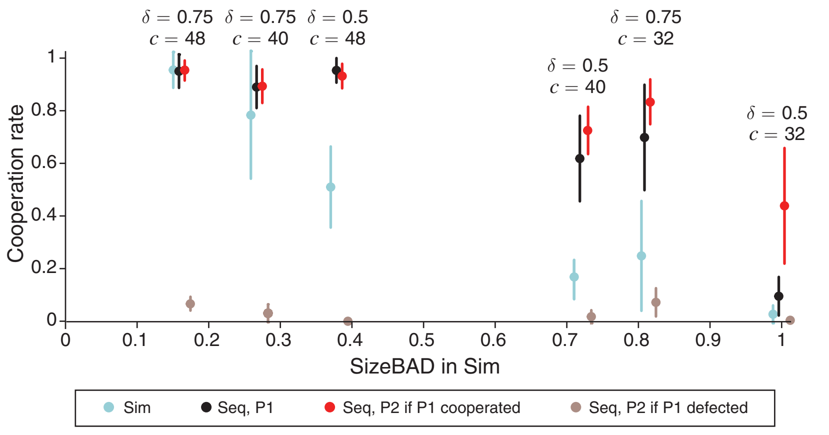

FIGURE 3. COOPERATION RATES BY ROLE

Notes: The graph shows first-round cooperation rates of P1, cooperation rates of P2 conditional on P1 defecting or cooperating, and cooperation rates in Sim, and 95 percent confidence intervals based on predictions from probit regressions ran on treatment and role indicators with clustered standard errors at the matching group level. Based on the last 20 repeated games.

The third observation, also in line with the SizeBAD predictions, is that both the first-mover cooperation rate and the second-mover conditional cooperation rate are higher in the treatments with $\delta > \delta^*$ than in the treatment with $\delta < \delta^*$ ($p < 0.001$ and $p < 0.001$, respectively).[21]

If we focus on whether first- and second-mover choices are aligned, then three other noteworthy patterns emerge from Figure 3. First, in treatment $\delta = 0.5$, $c = 32$, the conditional cooperation rate of second movers is well above the cooperation rate of first movers ($p < 0.001$). Second, in the treatments with $\delta > \delta^*$, the first-mover cooperation rate and the second-mover conditional cooperation rate are relatively well aligned.[22] Third, both cooperation rates are positively related to c and δ, even for $\delta > \delta^*$ ($p \leq 0.012$ for both c and δ in probit regressions excluding the treatment $\delta = 0.5, c = 32$). These patterns cannot be explained based on a strict interpretation of the SizeBAD predictions, but are consistent with a quantal response explanation or with the notion that players are heterogeneous. In the next section, we examine the results more closely at the individual and matching-group level and provide evidence that supports a heterogeneity interpretation.

[21] If we compare $\delta = 0.5, c = 32$ to $\delta = 0.5, c = 40$, then we get respectively $p < 0.001$ and $p = 0.014$, while if we compare $\delta = 0.5, c = 32$ to $\delta = 0.75, c = 32$, we get $p < 0.001$ and $p = 0.001$. For an overview of the statistical test results of treatment comparisons, see online Appendix Table F.4. Moreover, as shown in online Appendix Figure G.5, with $\delta < \delta^*$ the first-mover cooperation rate tends to decrease over time (negative linear trend with $p = 0.004$) while the second-mover conditional cooperation rate shows no trend ($p = 0.980$), whereas with $\delta > \delta^*$, both cooperation rates increase over time (positive linear trend with $p \leq 0.054$ and $p \leq 0.029$, respectively).

[22] Specifically, $p = 0.013$ in $\delta = 0.5, c = 40$, $p = 0.516$ in $\delta = 0.5, c = 48$, $p = 0.017$ in $\delta = 0.75, c = 32$, $p = 0.850$ in $\delta = 0.75, c = 40$, and $p = 0.816$ in $\delta = 0.75, c = 48$.

C. *Disaggregated Analysis*

Second Movers.—We have shown that the conditional cooperation rate of second movers is well above 0 in treatment $\delta = 0.5, c = 32$ (with $\delta < \delta^*$) and well below 1 in treatments $\delta = 0.5, c = 40$ and $\delta = 0.75, c = 32$ (with $\delta > \delta^*$). This implies either that *some* second movers *often* behave differently than a rational payoff-maximizer (consistent with a heterogeneity interpretation) or that *most* second movers *sometimes* behave differently than a rational payoff-maximizer (consistent with quantal response behavior). In order to differentiate between these two explanations, we examine the frequency with which each subject cooperates in the role of second mover, conditional on the first mover cooperating. If second movers are homogeneous in the extent to which they deviate from the predicted choice, as is the case in representative-player models like the quantal response model, then the share of conditionally cooperative choices should be similar across subjects in a given treatment. Alternatively, if second movers are heterogeneous in the sense that some of them systematically deviate from the rational payoff-maximizing benchmark, then the share of conditionally cooperative choices should differ across subjects in a given treatment.

As can be seen in Figure 4, most of the conditional cooperation choices in treatment $\delta = 0.5, c = 32$ can be attributed to just a few subjects.[23] These subjects can be viewed as conditional cooperation types; subjects who conditionally cooperate because they have a preference to do so. In the treatments with $\delta > \delta^*$, where conditional cooperation types cannot be identified because they pool with payoff maximizers, many more subjects always or almost always conditionally cooperate.

Moreover, Figure 4 shows that the opposite pattern emerges in treatments $\delta = 0.75, c = 48$; $\delta = 0.5, c = 48$; and $\delta = 0.75, c = 40$. Here, very few subjects are responsible for the majority of defection choices. Given that in these treatments, the decision to defect is more costly for second movers than in the other treatments, these subjects seem to have a strong taste for defection. We conclude therefore that a representative-player model does not suffice to explain disaggregated patterns of behavior of second movers. Instead, it appears to be necessary to allow for heterogeneity. This is further backed up by an analysis which statistically compares distributions of observed choices shown in Figure 4 to i.i.d. choices (see online Appendix Section D for details). Overall, the findings closely align with the notion that second movers are heterogeneous with respect to their cooperation preference. This is illustrated in online Appendix Section C.4, in which we show that the data are well represented by a heterogeneity model with payoff-maximizing, prosocial, and spiteful types.

First Movers.—Building on the insight that second movers come in types, we now focus more closely on behavior of first movers. Although the theoretical framework we use to formulate hypotheses builds on common knowledge of utilities, this assumption seems unrealistic if players are heterogeneous, especially in the

[23] For identification purposes, all analyses reported in this section include data from the first rounds of *all* the repeated games. Focusing on the last 20 repeated games would leave little power to perform disaggregated analyses.

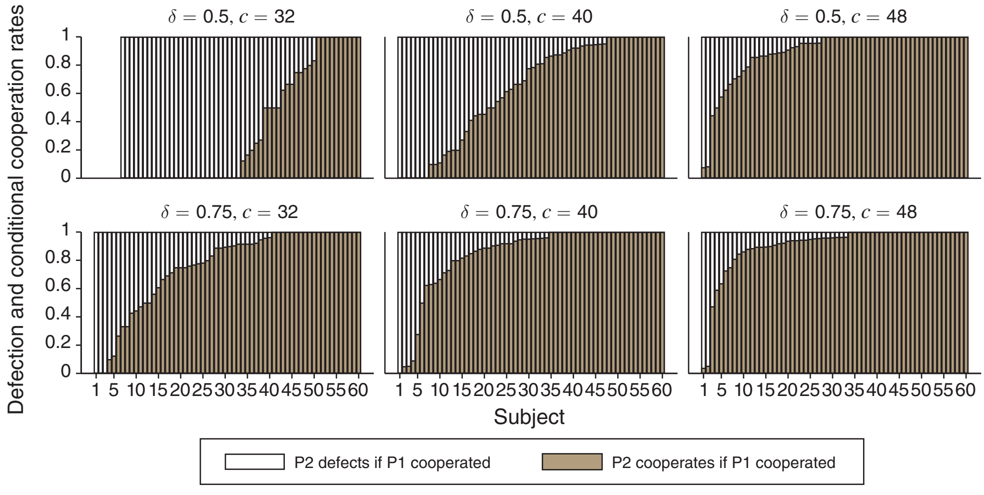

Figure 4. Conditional Cooperation Rates by Subject

Notes: The graphs show first-round defection and cooperation rates in the role of second mover by subjects in Seq, conditional on the matched first mover cooperating. In $\delta = 0.5, c = 32$, 6 second movers never encountered cooperation by the first mover, while the remaining 54 second movers encountered cooperation by the first mover between 1 and 12 times with a median of 3. In the other treatments, all second movers encountered cooperation by the first mover at least 4 times with the median ranging from 12.5 to 22 across the 5 treatments.

anonymous context of a lab experiment. We therefore assume that participants learn the distribution of second-mover types in their matching group during the course of the experiment, but do not know the specific type of their game partner (as in Kartal and Müller 2021).[24] With this in mind, we can compare observed choices of first movers to choices that expected-payoff maximizers would make if they were faced with the same second-mover choices.

For each first mover, we first compute the conditional cooperation rate she encountered in her matching group across first rounds of all repeated games. Panel A of Figure 5 shows these encountered conditional cooperation rates by treatment and matching group. The dashed horizontal lines refer to the conditional cooperation rate that leaves an expected-payoff-maximizing first mover indifferent between a conditionally cooperative strategy and always defect. As can be seen, there is substantial variation across matching groups and treatments in the extent to which the conditional cooperation rate encountered by first movers deviates from the indifference threshold. Taking the encountered conditional cooperation rate as given, we calculate for each first mover the (normalized) difference between the expected payoff of the cooperative strategy and that of the defection strategy. A risk-neutral first mover is better off cooperating (defecting) when the difference is positive (negative) and is indifferent when the difference is zero. We then plot the first-mover cooperation rates aggregated by matching group as a function of the (normalized)

[24] Recall that at the start of each repeated game, participants are randomly allocated partners within matching groups and randomly assigned roles. Thus, in a sense, each matching group constitutes a different "population" of players.

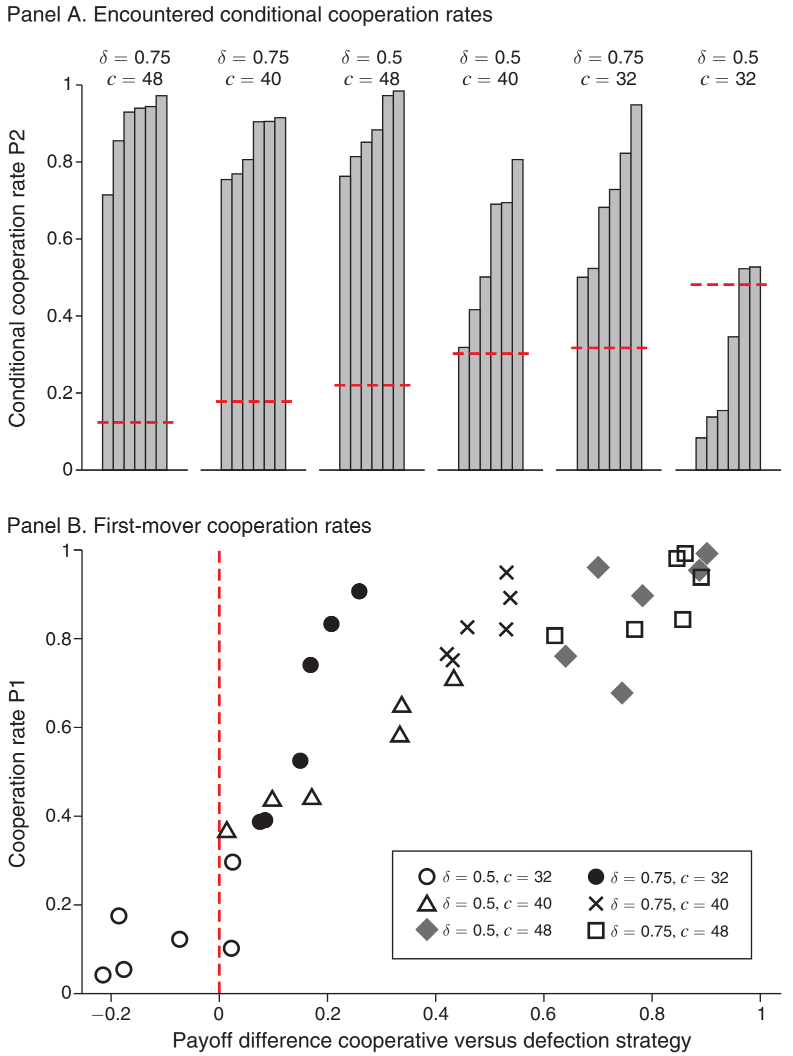

Figure 5. Cooperation Rates by Matching Groups

Notes: Panel A shows the conditional cooperation rates encountered by first movers across first rounds of all repeated games by treatment and matching group. The horizontal lines represent the conditional cooperation rate that leaves a payoff-maximizing first mover indifferent between defection and cooperation. Treatments are ordered by the SizeBAD in Sim. Panel B shows first-mover cooperation rates across first rounds of all repeated games as a function of the normalized difference between the expected payoffs from cooperation and defection, given the encountered conditional cooperation rate. Each dot in the graph corresponds to a matching group, and the 6 different shapes correspond to the 6 parametrizations in the experiment.

payoff difference. If all first movers would be expected-payoff maximizers, then their cooperation rates would jump straight to one when the indifference threshold is crossed. Panel B of Figure 5 shows that the cooperation rate of first movers is close to 0 in matching groups where the payoff difference is negative (in four of the six matching groups in treatment $\delta = 0.5, c = 32$) and that it increases as the payoff difference increases. Once cooperation is much more profitable than defection, then the cooperation rate stays close to 1. We conjecture that the lack of a sudden jump at the threshold in Panel B is due to heterogeneity of first movers. For example, the pattern is consistent with a substantial fraction of first movers being averse to disadvantageous inequality (see online Appendix Section C.4).

Within-Subject Analysis.—Given that subjects make choices in both roles, additional insights related to heterogeneity can be obtained by investigating choice patterns of subjects across the two roles. We focus on the correlation between their conditional cooperation rate as a second mover and the extent to which their cooperation rate as a first mover differs from their personal optimal choice. We calculate this optimal choice as the choice that maximizes the expected payoff taking into account the conditional cooperation rate encountered in first rounds in one's matching group, as introduced in the previous paragraph. In most cases, the optimal choice is to defect in $\delta = 0.5, c = 32$, and to cooperate in the other treatments. Scatter plots by treatment are shown in online Appendix Figure G.6.

The first finding is that in the treatments with $\delta > \delta^*$, the correlation is overall positive and strong ($p \leq 0.018$). Players thus tend to cooperate as a first mover to almost the same extent as they conditionally cooperate as a second mover. We conjecture that this result is largely due to payoff maximizers having an incentive to cooperate in both roles, which makes them behave similarly to conditional cooperation types. The second finding is that no positive correlation is detected in treatment $\delta = 0.5, c = 32$, in which $\delta < \delta^*$. This result is consistent with the fact that payoff maximizers now have no incentive to conditionally cooperate as a second mover, nor to cooperate as a first mover. Any choice other than defection in $\delta = 0.5, c = 32$ can thus be attributed to behavior that differs from rational payoff maximization (such as, for example, other-regarding behavior or quantal responses). Given that as a first mover one is faced with higher strategic risk than as a second mover, there is no reason to expect that players who prefer to conditionally cooperate in $\delta = 0.5, c = 32$ as a second mover also prefer to cooperate as a first mover.

To further illustrate how players in $\delta = 0.5, c = 32$ make choices in different roles, we split up conditional cooperation types according to their behavior as a first mover. For simplicity, players are defined as conditional cooperation types if they conditionally cooperate more than half of the time when encountering cooperation from the matched first mover. We find that 78 percent of them (14 out of 18) cooperate less frequently as a first mover than what is optimal and 17 percent (3 out of 18) cooperate more frequently than what is optimal. Among the other players, the percentages are 53 percent (19 out of 36) and 42 percent (15 out of 36), respectively, indicating a more balanced distribution. Although power is too low to provide conclusive statistical support, these findings suggest that conditional cooperation types tend to be more averse to disadvantageous inequality (or more risk-averse) than other players.

IV. Conclusion

Failure to coordinate on efficient outcomes is largely due to individuals avoiding strategic risk (Van Huyck, Battalio, and Beil 1990, 1991). A similar logic applies with respect to cooperation in repeated games. Cooperation rates are highest in games where conditionally cooperative strategies involve little risk (Blonski, Ockenfels, and Spagnolo 2011; Dal Bó and Fréchette 2011). We use this insight to predict that introducing sequentiality in games that are characterized by substantial strategic risk may facilitate cooperation by reducing that risk. The experiment we carry out shows that the prediction is borne out by the data. In games where it is difficult for players to achieve mutual cooperation—even though it can be supported in equilibrium—introducing sequentiality increases the cooperation rate by around 40 percentage points. In games where cooperation is not supported in equilibrium or where it is supported but strategic risk is particularly low, cooperation rates are close to 0 or 100 percent respectively, independent of sequentiality. We thus conclude that individuals strongly react to sequentiality in environments with coordination problems that are the result of substantial strategic risk.

When modeling decision-making it is not always clear whether a simultaneous-move setting or a sequential-move setting is most appropriate. We show that behavior strongly depends on the setting, implying that possible policy implications may strongly depend on whether a simultaneous-move or sequential-move setting is ultimately chosen. The results also have implications for behavioral mechanism design. If a designer's goal is to achieve and sustain high efficiency levels, it is optimal that players decide sequentially and that second movers have information about the decision of the first mover. Consider, for instance, the issue of climate change, in which long-run incentives are arguably large enough for it to be optimal that countries engage in a cooperative mitigation of greenhouse gas emissions (Dutta and Radner 2004; Calzolari, Casari, and Ghidoni 2018). If a country commits to a policy of reducing emissions in anticipation that other countries will follow suit, then those other countries will indeed have an increased incentive to do so because the risk of free-riding has been reduced. This may be good news for environmental policymakers because convincing one country or even a small group of countries to commit to environmentally-friendly actions is arguably easier to achieve than convincing all countries. Sequentiality might therefore help countries coordinate to achieve socially optimal outcomes. The same is true for other contexts, such as trade and employer-employee relations. Nevertheless, it is an open question as to whether the strong efficiency-enhancing effect of sequentiality is also achieved if the game's parameters are uncertain, which is a more realistic assumption in most applications. The result of Wilson and Vespa (2020), that cooperation does not predominate in a sequential-move setting with asymmetric information about payoffs, suggests that this is not necessarily the case.

An alternative instrument that can in principle reduce strategic uncertainty is pre-play communication (see, for example, Arechar et al. 2017) and it appears that sequentiality can overcome some of the disadvantages associated with communication. First, given that communication is not consequential for monetary payoffs, it has no effect on predictions based on equilibrium refinements or on concepts such as the basin of attraction of a particular strategy (Crawford 1998). In contrast, sequentiality does affect monetary payoffs because it allows the second mover to avoid the sucker payoff. Second, the efficacy of communication in increasing coordination appears to be quite sensitive to the communication protocol, which makes its implementation less straightforward than introducing sequentiality (see, for example, Cooper et al. 1992; Andersson and Wengström 2012, for evidence from simple coordination games).[25]

Our results have implications for the interpretation of behavior in PD games played in (quasi-)continuous time (see, for example, Friedman and Oprea 2012; Bigoni et al. 2015). Cooperation rates in (quasi-)continuous time are typically very high but the reasons are not entirely understood. These games differ in at least three respects from discrete-time simultaneous PDs: (i) the frequency of the (albeit shorter) interactions is higher in each repeated game; (ii) players move de facto sequentially, i.e., they observe the partner's choice before making a choice; and (iii) players choose the timing of their moves. Friedman and Oprea (2012) show that frequency of interaction increases the cooperation rate in discrete-time PDs. However, our experiment shows that sequentiality may on its own lead to a substantial increase in cooperation, provided that cooperation is sustainable in equilibrium. The sequential-move nature of games played in (quasi-)continuous time may thus be one of the structural characteristics that leads to the higher cooperation rate. This is consistent with the results of an experiment in which strategic uncertainty is removed by freezing choices for a few seconds, which is shown to increase cooperation (Calford and Oprea 2017). Strategically, a sequential PD is similar to a simultaneous PD in which the choice of one of the players is frozen for one period.

Our analysis builds on a framework in which it is assumed that payoff-maximizing players choose between always defecting and conditional cooperation under common knowledge. This makes it possible to construct a simple measure of the degree of strategic uncertainty and helps to formalize the difference between sequential-move and simultaneous-move PDs. Thus, the approach is not meant to provide an accurate description of how individuals play. There are at least two ways in which behavior can be plausibly expected to deviate from the assumptions. First, players may follow strategies other than always defect or conditional cooperation. Results of strategy estimations show that by far the majority of the cooperative strategies involve conditional cooperation à la grim trigger or tit-for-tat (see online Appendix Section E).[26] This, and the fact that we are dealing with relatively short games, gives us confidence that the simplification of the repeated games to binary-choice games is not overly simplistic.[27]

[25] That said, it also holds that pre-play communication can trigger behavioral responses that go beyond removing strategic uncertainty and can foster cooperation even if this is not an equilibrium outcome, for example by appealing to honesty (Gneezy 2005) or inducing guilt aversion (Charness and Dufwenberg 2006). To illustrate, pre-play chat has been shown to increase cooperation in one-shot interactions (see Balliet 2010, for a meta-analysis) or in repeated simultaneous games in which cooperation cannot be sustained in equilibrium (Kartal and Müller 2021).

[26] An exception is the strategy to first defect and then switch to tit-for-tat (D-TFT), which is particularly popular among first movers and to some extent in the case of simultaneous moves, in the game in which cooperation cannot be sustained in equilibrium. We speculate that this may have to do with the fact that D-TFT protects a player from the sucker payoff if matched with a defecting partner and at the same time achieves mutual cooperation if the partner is lenient.

[27] Notice that if the second mover believes that the first mover either always defects or always cooperates, then it would be optimal for her to always defect, even if $\delta > \delta^*$, and this may explain why the conditional cooperation rate is well below 1 in the treatments with intermediate gains from cooperation. However, given that in these treatments less than 2 percent of the first movers are estimated to always cooperate, holding such a belief would be largely irrational. We therefore feel that this is not a sufficient explanation.

Second, players may not all be perfect payoff maximizers with common knowledge. We have shown that some form of heterogeneity is needed in order to explain all aspects of the data. To do so, we have used an example on other-regarding preferences but a similar intuition holds if there is heterogeneity in risk preferences or in the strength of quantal responses.[28] A key element is that the heterogeneity introduces individual-specific trade-offs between a conditional cooperation strategy and an always-defect strategy, which leads to smoothness into the aggregate effect of a game's parameters on the cooperation rate, even in sequential-move games with $\delta > \delta^*$. A promising model that incorporates strategic risk and at the same time predicts smoothness is that of Kartal and Müller (2021). The model provides a microeconomic foundation for strategic uncertainty by assuming that players have heterogeneous and unobservable tastes.

REFERENCES

Ahn, Toh-Kyeong, Myungsuk Lee, Lore Ruttan, and James Walker. 2007. "Asymmetric Payoffs in Simultaneous and Sequential Prisoner's Dilemma Games." *Public Choice* 132 (3–4): 353–66.

Ahn, Toh-Kyeong, Elinor Ostrom, and James Walker. 2003. "Heterogeneous Preferences and Collective Action." *Public Choice* 117 (3): 295–314.

Akerlof, George A. 1982. "Labor Contracts as Partial Gift Exchange." *Quarterly Journal of Economics* 97 (4): 543–69.

Andersson, Ola, and Erik Wengström. 2012. "Credible Communication and Cooperation: Experimental Evidence from Multi-stage Games." *Journal of Economic Behavior and Organization* 81 (1): 207–19.

Andreoni, James, Paul M. Brown, and Lise Vesterlund. 2002. "What Makes an Allocation Fair? Some Experimental Evidence." *Games and Economic Behavior* 40 (1): 1–24.

Arechar, Antonio A., Anna Dreber, Drew Fudenberg, and David G. Rand. 2017. "'I'm Just a Soul Whose Intentions Are Good': The Role of Communication in Noisy Repeated Games." *Games and Economic Behavior* 104: 726–43.

Balliet, Daniel. 2010. "Communication and Cooperation in Social Dilemmas: A Meta-Analytic Review." *Journal of Conflict Resolution* 54 (1): 39–57.

Bigoni, Maria, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo. 2015. "Time Horizon and Cooperation in Continuous Time." *Econometrica* 83 (2): 587–616.

Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo. 2011. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence." *American Economic Journal: Microeconomics* 3: 164–92.

Blonski, M., and G. Spagnolo. 2015. "Prisoners' Other Dilemma." *International Journal of Game Theory* 44 (1): 61–81.

Brown, Martin, Armin Falk, and Ernst Fehr. 2004. "Relational Contracts and the Nature of Market Interactions." *Econometrica* 72 (3): 747–80.

Calford, Evan, and Ryan Oprea. 2017. "Continuity, Inertia, and Strategic Uncertainty: A Test of the Theory of Continuous Time Games." *Econometrica* 85 (3): 915–35.

Calzolari, Giacomo, Marco Casari, and Riccardo Ghidoni. 2018. "Carbon Is Forever: A Climate Change Experiment on Cooperation." *Journal of Environmental Economics and Management* 92: 169–84.

Camera, Gabriele, Marco Casari, and Maria Bigoni. 2013. "Money and Trust among Strangers." *Proceedings of the National Academy of Sciences* 110 (37): 14889–93.

Charness, Gary, and Martin Dufwenberg. 2006. "Promises and Partnership." *Econometrica* 74 (6): 1579–1601.

Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117 (3): 817–69.

[28] For example, risk-averse (risk-seeking) players will prefer conditional cooperation less (more) than always defect.

Clark, Kenneth, and Martin Sefton. 2001. "The Sequential Prisoner's Dilemma: Evidence on Reciprocation." *Economic Journal* 111 (468): 51–68.

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

Crawford, Vincent. 1998. "A Survey of Experiments on Communication via Cheap Talk." *Journal of Economic Theory* 78 (2): 286–98.

Dal Bó, Pedro, and Guillaume R. Fréchette. 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101 (1): 411–29.

Dal Bó, Pedro, and Guillaume R. Fréchette. 2018. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature* 56 (1): 60–114.

Dal Bó, Pedro, and Guillaume R. Fréchette. 2019. "Replication data for: The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review*. http://doi.org/10.3886/E112401V1.

Dutta, Prajit K., and Roy Radner. 2004. "Self-Enforcing Climate-Change Treaties." *Proceedings of the National Academy of Sciences* 101 (14): 5174–79.

Embrey, Matthew, Guillaume R. Fréchette, and Sevgi Yuksel. 2018. "Cooperation in the Finitely Repeated Prisoner's Dilemma." *Quarterly Journal of Economics* 133 (1): 509–51.

Engle-Warnick, Jim, and Robert L. Slonim. 2006. "Learning to Trust in Indefinitely Repeated Games." *Games and Economic Behavior* 54 (1): 95–114.

Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics* 108: 437–59.

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.

Friedman, Daniel, and Ryan Oprea. 2012. "A Continuous Dilemma." *American Economic Review* 102 (1): 337–63.

Friedman, James W. 1971. "A Non-cooperative Equilibrium for Supergames." *Review of Economic Studies* 38 (1): 1–12.

Fudenberg, Drew, and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (3): 533–54.

Ghidoni, Riccardo, and Sigrid Suetens. 2022. "Replication data for: The Effect of Sequentiality on Cooperation in Repeated Games." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E127781V1.

Gneezy, Uri. 2005. "Deception: The Role of Consequences." *American Economic Review* 95 (1): 384–94.

Greif, Avner. 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition." *American Economic Review* 83 (3): 525–48.

Güth, W., M. V. Levati, M. Sutter, and E. van der Heijden. 2007. "Leading by Example with and without Exclusion Power in Voluntary Contribution Experiments." *Journal of Public Economics* 91 (5–6): 1023–42.

Kandori, Michihiro. 1992. "Social Norms and Community Enforcement." *Review of Economic Studies* 59 (1): 63–80.

Kartal, Melis, and Wieland Müller. 2021. "A New Approach to the Analysis of Cooperation under the Shadow of the Future: Theory and Experimental Evidence." https://ssrn.com/abstract=3222964.

Kehoe, Timothy J., and David K. Levine. 1993. "Debt-Constrained Asset Markets." *Review of Economic Studies* 60 (4): 865–88.

Khadjavi, Menusch, and Andreas Lange. 2013. "Prisoners and Their Dilemma." *Journal of Economic Behavior and Organization* 92: 163–75.

Kreps, David M. 1990. "Corporate Culture and Economic Theory." In *Perspectives on Positive Political Economy*, edited by J. Alt and K. Shepsle, 90–142. Cambridge, UK: Cambridge University Press.

Levine, David. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics* 1 (3): 593–622.

Oskamp, Stuart. 1974. "Comparison of Sequential and Simultaneous Responding, Matrix, and Strategy Variables in a Prisoner's Dilemma Game." *Journal of Conflict Resolution* 18 (1): 107–16.

Potters, Jan, Martin Sefton, and Lise Vesterlund. 2005. "After You—Endogenous Sequencing in Voluntary Contribution Games." *Journal of Public Economics* 89 (8): 1399–1419.

Reuben, Ernesto, and Sigrid Suetens. 2012. "Revisiting Strategic versus Non-strategic Cooperation." *Experimental Economics* 15 (1): 24–43.

Roth, Alvin E., and J. Keith Murnighan. 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology* 17 (2): 189–98.

Thomas, Jonathan, and Tim Worrall. 1990. "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem." *Journal of Economic Theory* 51 (2): 367–90.

Turocy, Theodore L. 1995. "A Dynamic Homotopy Interpretation of the Logistic Quantal Response Equilibrium Correspondence." *Games and Economic Behavior* 51 (2): 243–63.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80 (1): 234–48.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1991. "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games." *Quarterly Journal of Economics* 106 (3): 885–910.

Vespa, Emanuel, and Alistair J. Wilson. 2019. "Experimenting with the Transition Rule in Dynamic Games." *Quantitative Economics* 10 (4): 1825–49.

Wen, Quan. 2002. "A Folk Theorem for Repeated Sequential Games." *Review of Economic Studies* 69 (2): 493–512.

Wilson, Alistair J., and Emanuel Vespa. 2020. "Information Transmission under the Shadow of the Future: An Experiment." *American Economic Journal: Microeconomics* 12 (4): 75–98.