The effect of sequentiality on cooperation in repeated games

Riccardo Ghidoni
University of Milano-Bicocca, CefES, Department of Economics, Management, and Statistics
Tilburg University, CentER, Department of Economics
e-mail: riccardo.ghidoni@unimib.it

Sigrid Suetens
Tilburg University, CentER, Department of Economics
e-mail: s.suetens@tilburguniversity.edu

November 2, 2020

Abstract

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. In particular, 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 Codes: C70, C90, D70

Keywords: cooperation, infinitely repeated game, prisoner's dilemma, sequential moves, strategic uncertainty, experiment, conditional cooperation

1 Introduction

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 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, the second mover can reap the benefits of cooperation and at the same time avoid being betrayed by cooperating if and only if the first mover cooperates. 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 determinants of cooperation.

The role of 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 et al. (2011) and Blonski and Spagnolo (2015) who apply the concept of risk dominance to the repeated PD and Dal Bó

[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 et al., 1993), and trade (Greif, 1993; Brown et al., 2004).

[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 worst possible punishment strategy in both settings (Fudenberg and Maskin, 1986).

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

and Fréchette (2011) who appeal to the basin of attraction of repeated-game strategies. These approaches also help to formalize the intuition that the 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. Otherwise, they defect.[4] With simultaneous decision-making, in contrast, 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 the threshold of the standard theory of infinitely repeated games, while in the others, δ is above the theoretical threshold. We formulate predictions while 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 in the simultaneous version of the latter games, strategic risk is largest.

The experimental results show strong support for the predictions that take into account strategic risk. In the treatments that are characterized by relatively high strategic

[4]The prediction is reminiscent of a case discussed by Camera et al. (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).

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 one then when mutual cooperation is sustainable, and close to zero when it is not.

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 et al., 2003, 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. In this literature, (exogenously imposed) sequentiality of moves increases contributions relative to a simultaneous-move setting if the leader has private information about the game's parameters (Potters et al., 2005) but leads to mixed results in full information settings (for example Andreoni et al., 2002; Güth et al., 2007).[7] Finally, Kartal and Müller (2018) compare simultaneous and sequential infinitely repeated PDs in an experiment inspired by their 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 2 provides the theoretical background; Section 3 describes the experimental design and procedures and presents predicted effects of sequentiality in our experiment; Section 4 presents the main results, with focus on the treatment effect of sequentiality and on the behavior of

[6] In these experiments, it was announced that the repeated game would last for 60 rounds but was 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 or not 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

11*Notes:* $t > c > d > s$ and $2c > t + s$.

first and second movers in the sequential games; Section 5 concludes.

2 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 Section C.1 of the Appendix for calculations).[9] In summary, standard game theory does not provide a specific reason why cooperation rates should 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 non-cooperative equilibria exist in both cases.

More precise predictions can be obtained by appealing to risk dominance (Blonski et al., 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 or not their partner will defect. In particular, consider a simpli-

[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] Since a second mover never moves first, the implementation of GT is as follows: the second mover should cooperate if the first mover cooperates and if she himself cooperated in the previous move, and switch to defection forever after a defection of one of the two players.

fication 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 basin of attraction of AD versus CC (referred to as SizeBAD) is defined as the maximum probability that the partner will follow the CC strategy that makes AD optimal, which is based on the set of beliefs about the partner's strategy that makes AD optimal. 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 explain the intuition. The detailed calculations are presented in Section C.2 of the Appendix.

First, consider 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 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.

Next, consider 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

[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 punishment would also qualify as CC.

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

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% if $\delta \geq \delta^*$ and zero 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 3, 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 pro-social 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, pro-social, and spiteful players in the population, the cooperation rate depends on the parameters of the game, even in sequential PDs with $\delta > \delta^*$. In Section C.4 of the Appendix 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 (2018).

3 The experiment

3.1 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

[12]Notice that the same predictions hold in the limit of a quantal response equilibrium, since 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 Section C.3 of the Appendix 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 anti-social preferences (e.g. Levine, 1998; Fehr and Schmidt, 1999; Charness and Rabin, 2002). For them, payoffs in PDs do not represent utilities. Ahn et al. (2003) and Ahn et al. (2007) illustrate how heterogeneity models help to understand cooperation in one-shot simultaneous and 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	
# Participants	30	30	30	30	30	30	60	60	60	60	60	60	540
# Matching groups	3	3	3	3	3	3	6	6	6	6	6	6	54
# Repeated games	50	50	50	50	50	50	50	50	50	50	50	50	600
# 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.

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 also randomly allocated the role of first or second mover at the beginning of each repeated game. We expect 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 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 wide range of settings ranging (in expectation) from short games with low gains to mutual cooperation to longer games with high gains to mutual cooperation. 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

[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.

minutes on average and participants earned on average of €22.7. The procedures are described more in detail in Section A of the Appendix, and an English translation of the instructions can be found in Section B of the Appendix.[15]

3.2 Predictions

The predictions are based on the basin-of-attraction approach discussed in Section 2. 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 the cooperation rate is close to one in Sim with $c = 48, \delta = 0.75$, we can summarize the predictions as follows:

1. The cooperation rate is expected to be close to zero in Sim and Seq in treatment $c = 32, \delta = 0.5$.

2. The cooperation rate is expected to be close to one in Sim and Seq in treatment $c = 48, \delta = 0.75$.

3. 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:
   $c = 32, \delta = 0.75 \leq c = 40, \delta = 0.5 \leq c = 48, \delta = 0.5 \leq c = 40, \delta = 0.75$.

4 Results

4.1 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 (a) repeated games may have different lengths and (b) the adopted theoretical framework involves the choice of

[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 3: SizeBAD by treatment

(a) Sim

		c		
		32	40	48
δ	0.5	1	0.72	0.38
	0.75	0.81	0.27	0.16

(b) Seq

		c		
		32	40	48
δ	0.5	1	0	0
	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.

Figure 1: Cooperation rates

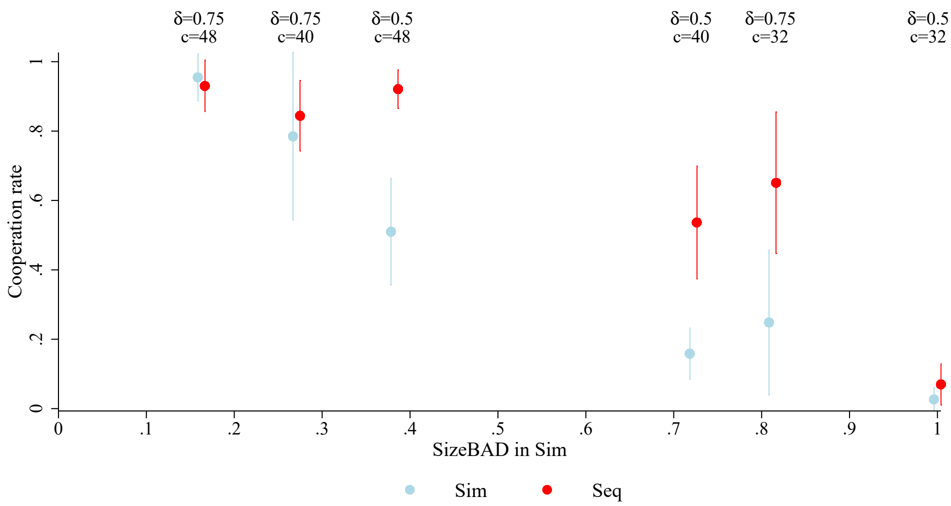

Notes: The graph shows first-round cooperation rates and 95% confidence intervals across the last 20 repeated games depending on the SizeBAD (including treatment labels). Estimates and confidence intervals are based on predictions from probit regressions run on a treatment dummy with clustered standard errors at the matching-group level.

whether to use a cooperative or non-cooperative 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 twenty repeated games.[17] We find that the difference between Sim and Seq is small in the treatments with the

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

[17]Figure G.1 in the Appendix includes predicted cooperation rates in Seq and cooperation rates observed in DBF's simultaneous PDs. As can be seen, DBF cooperation rates generally fall within 95% 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 changes in procedure.

lowest or highest incentive to cooperate ($p = 0.200$ and $p = 0.635$, respectively).[18] The cooperation rate is respectively close to zero and close to one 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.

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 Table F.2 in the Appendix).[19] The results are also robust to a re-estimation of treatment effects on the basis of a dataset in which our data is merged with that of DBF (see Table F.3 in the Appendix).

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 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 Table F.5 in the Appendix). 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

[18] 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 Tables F.1 and F.4 in the Appendix.

[19] Overall, we find a positive relation in the first rounds between pro-sociality 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 2: Evolution of cooperation rates

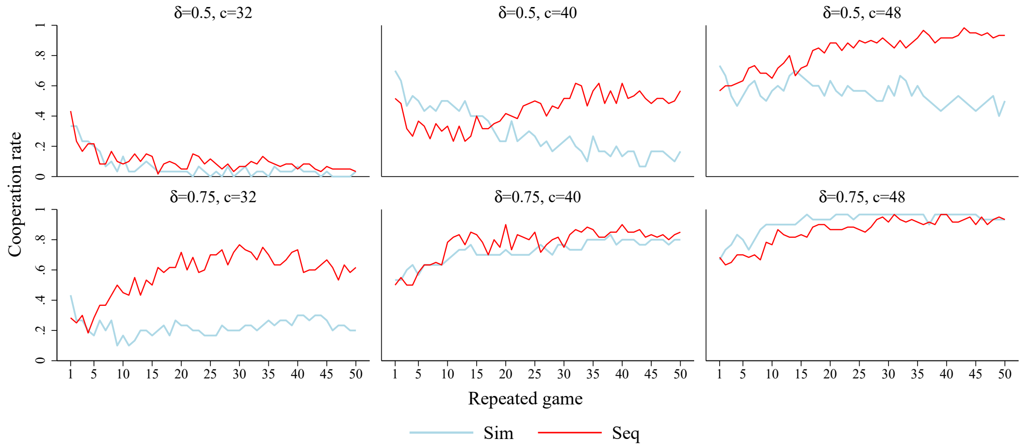

Notes: The graphs show cooperation rates across first rounds of repeated games by treatment.

the end of Section 2.

We now turn to the learning dynamics. Figure 2 illustrates how first-round cooperation rates evolve across the fifty repeated games for each treatment.[20] The graphs show that some learning is necessary before the above-reported treatment effects 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 zero and then sharply declines to a rate close to zero, whereas in the treatments in which SizeBAD predicts a cooperation rate of one, 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.[21]

[20] Patterns by matching group are shown in Figure G.3 in the Appendix.

[21] 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).

4.2 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: first-mover cooperation rate; second-mover cooperation rate conditional on cooperation by the first mover (which we shall refer to as the conditional cooperation rate); and 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 not ($p = 0.079$).

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).[22]

[22]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 Table F.4 in the Appendix. Moreover, as shown in Figure G.5 in the Appendix, with $\delta < \delta^*$ the first-mover cooperation rate tends to decrease over time

Figure 3: Cooperation rates by role

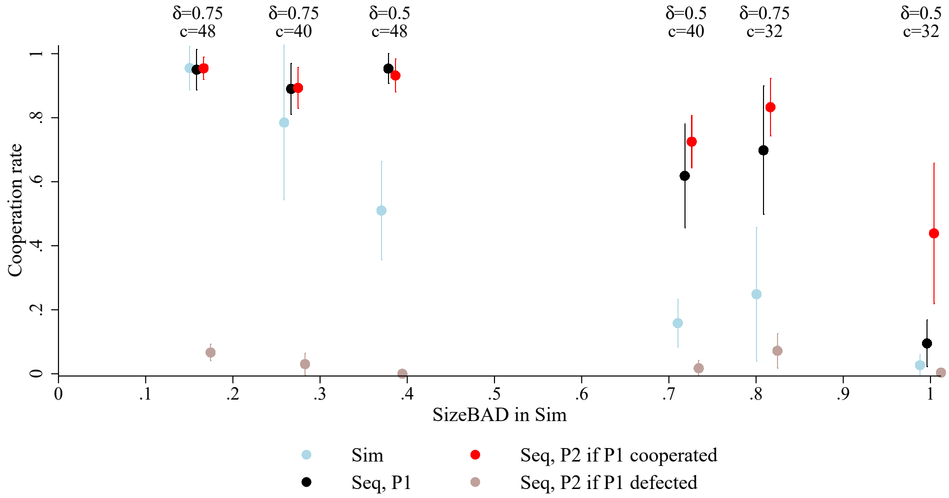

Notes: The graph shows first-round cooperation rates of P1, cooperation rates of P2 conditional on P1 defecting or cooperating, and cooperation rates in Sim, and 95% confidence intervals across the last 20 repeated games depending on the SizeBAD (including treatment labels). Estimates and confidence intervals are based on predictions from probit regressions run on treatment-role dummies with clustered standard errors at the matching group level.

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.[23] 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 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 levels and provide evidence that supports a heterogeneity interpretation.

(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).

[23]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 $\delta = 0.75, c = 48$.

Figure 4: Conditional cooperation rates by subject

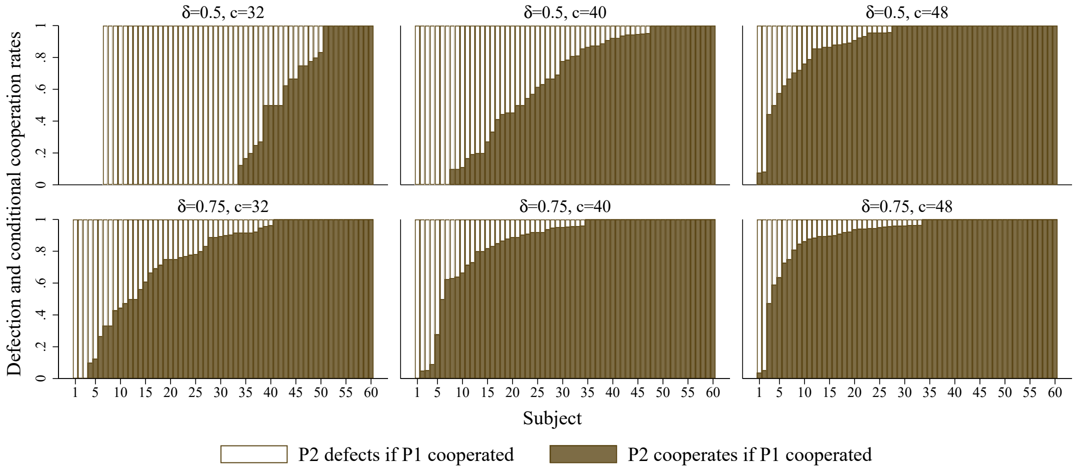

Notes: The graphs show first-round defection and cooperation rates in the role of second mover by subjects in Seq, conditional on the matched first mover cooperating. In $\delta = 0.5, c = 32$, 6 second movers never encountered cooperation by the first mover, while the remaining 54 second movers encountered 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.

4.3 Disaggregated analysis

4.3.1 Second movers

We have shown that the conditional cooperation rate of second movers is well above zero 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.[24] These subjects can be viewed as conditional cooperation types; types 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 backed up by an analysis which statistically compares distributions of observed choices shown Figure 4 to *iid* choices (see Section D in the Appendix for a detailed description). Overall, the findings closely align with the notion that second movers are heterogeneous with respect to their cooperation preference. This is illustrated in Section C.4 in the Appendix, in which we show that the data are well-represented by a heterogeneity model with payoff-maximizing, pro-social, and spiteful types.

4.3.2 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 anonymous context of our lab experiment. Instead, we 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, 2018).[25] With this in

[24]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.

[25]Recall that at the start of each repeated game participants are randomly allocated partners within matching groups and they randomly switch roles. Thus, in a sense each matching group constitutes a

Figure 5: Cooperation rates by matching groups

(a) Encountered conditional cooperation rates **(b) First-mover cooperation rates**

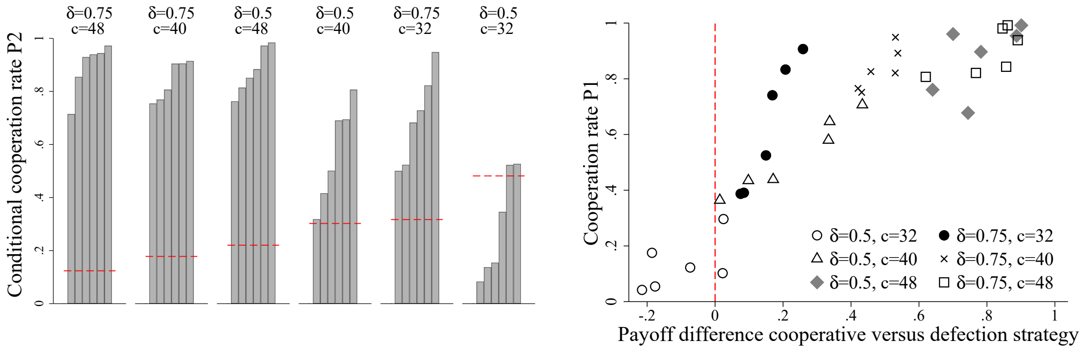

Notes: Panel (a) shows the conditional cooperation rates encountered by first movers across first rounds of all repeated games by treatment and matching group. The horizontal lines represent the conditional cooperation rate that leaves a payoff-maximizing first mover indifferent between defection and 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.

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. Figure 5a 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 the repeated-game strategies of defection and cooperation. 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

different 'population' of players.

group as a function of the (normalized) 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. Figure 5b shows that their cooperation rates are close to zero 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 one. We conjecture that the lack of a sudden jump at the threshold 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 Section C.4 in the Appendix).

4.3.3 Within-subject analysis

Given that subjects make choices in both roles, additional insights related to heterogeneity can be obtained by investigating the choice patterns within subjects. We focus on the correlation between the conditional cooperation rate as a second mover, on the one hand, and the extent to which the cooperation rate as a first mover differs from the optimal first-mover cooperation rate, on the other hand. We define this optimal rate as the cooperation rate of a first mover who maximizes expected payoff while taking into account the conditional cooperation rate she encountered, as introduced in 4.3.2. In most cases, it is equal to zero for $\delta = 0.5, c = 32$ and to one for the other treatments. Scatter plots by treatment are shown in Figure G.6 in the Appendix.

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 that 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 the time when encountering cooperation from the matched first mover. We find that 78% of them (14 out of 18) cooperate less frequently as a first mover than what is optimal and 17% (3 out of 18) cooperate more frequently than what is optimal. Among the other players, the percentages are 53% (19 out of 36) and 42% (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 than other players.

5 Conclusion

Failure to coordinate on efficient outcomes is largely due to individuals avoiding strategic risk (Van Huyck et al., 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 et al., 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 zero 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.

In 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 et al., 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 policy makers 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 of 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 preplay communication (see, for example, Arechar et al., 2017) and it appears that sequentiality can overcome some of its disadvantages. First, given that communication is not consequential on 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 implementation less straightforward than introducing sequentiality (see, for example, Cooper et al., 1992; Andersson and Wengström, 2012, for evidence from simple coordination games).[26]

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: (a) the frequency of the (albeit shorter) interactions is higher in each repeated game; (b) players move *de facto* sequentially, i.e. they observe the partner's choice before making a choice; and (c) 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 on its own may also 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

[26] 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, 2018).

cooperation à la grim trigger or tit-for-tat (see Section E in the Appendix).[27] 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.[28]

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.[29] A key element is that the heterogeneity introduces individual-specific trade-offs between a conditional cooperation strategy and an always-defect strategy, which introduce smoothness into the aggregate effect of the 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 (2018). The model provides a microeconomic foundation for strategic uncertainty by assuming that players have heterogeneous and unobservable tastes.

References

Ahn, T.-K., M. Lee, L. Ruttan, and J. Walker (2007). Asymmetric payoffs in simultaneous and sequential prisoner's dilemma games. *Public Choice 132*(3-4), 353–366.

Ahn, T.-K., E. Ostrom, and J. Walker (2003). Heterogeneous preferences and collective action. *Public Choice 117*(3), 295–314.

[27] 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.

[28] One might argue that the finding that the conditional cooperation rate of second movers is well below 100% even if $\delta > \delta^*$ is related to the beliefs of second movers. In particular, 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. However, given that in the treatments with $\delta > \delta^*$ and intermediate gains from cooperation less than 2% of the first movers is estimated to always cooperate, holding such a belief would be largely irrational. We therefore feel that this is not a sufficient explanation.

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

Akerlof, G. (1982). Labor contracts as partial gift exchange. *Quarterly Journal of Economics 97*(4), 543–569.

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

Andreoni, J., P. Brown, and L. Vesterlund (2002). What makes an allocation fair? Some experimental evidence. *Games and Economic Behavior 40*, 1–24.

Arechar, A. A., A. Dreber, D. Fudenberg, and D. 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*, 726–743.

Balliet, D. (2010). Communication and cooperation in social dilemmas: A meta-analytic review. *Journal of Conflict Resolution 54*(1), 39–57.

Bigoni, M., M. Casari, A. Skrzypacz, and G. Spagnolo (2015). Time horizon and cooperation in continuous time. *Econometrica 83*(2), 587–616.

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

Blonski, M. and G. Spagnolo (2015). Prisoners' other dilemma. *International Journal of Game Theory 44*, 61–81.

Brown, M., A. Falk, and E. Fehr (2004). Relational contracts and the nature of market interactions. *Econometrica 72*, 747–780.

Calford, E. and R. Oprea (2017). Continuity, inertia, and strategic uncertainty: A test of the theory of continuous time games. *Econometrica 85*(3), 915–935.

Calzolari, G., M. Casari, and R. Ghidoni (2018). Carbon is forever: A climate change experiment on cooperation. *Journal of Environmental Economics and Management 92*, 169–184.

Camera, G., M. Casari, and M. Bigoni (2013). Money and trust among strangers. *Proceedings of the National Academy of Sciences 110*(37), 14889–14893.

Charness, G. and M. Dufwenberg (2006). Promises and partnership. *Econometrica 74*(6), 1579–1601.

Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics 117*(3), 817–869.

Clark, K. and M. Sefton (2001). The sequential prisoner's dilemma: Evidence on reciprocation. *Economic Journal 111*, 51–68.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1992). Communication in coordination games. *Quarterly Journal of Economics 107*(2), 739–771.

Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory 78*(2), 286–298.

Dal Bó, P. and G. R. Fréchette (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review 101*(1), 411–429.

Dal Bó, P. and G. R. Fréchette (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature 56*(1), 60–114.

Dal Bó, P. and G. R. Fréchette (2019). Strategy choice in the infinitely repeated prisoner's dilemma. *American Economic Review 109*(11), 3929–52.

Dutta, P. K. and R. Radner (2004). Self-enforcing climate-change treaties. *Proceedings of the National Academy of Sciences 101*(14), 5174–5179.

Embrey, M., G. R. Fréchette, and S. Yuksel (2018). Cooperation in the finitely repeated prisoner's dilemma. *Quarterly Journal of Economics 133*(1), 211–224.

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

Fehr, E., G. Kirchsteiger, and A. Riedl (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics 108*, 1912–1917.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics 114*, 817–868.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Friedman, D. and R. Oprea (2012). A continuous dilemma. *American Economic Review 102*(1), 337–63.

Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies 38*(1), 1–12.

Fudenberg, D. and E. Maskin (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica 54*(3), 533–554.

Fudenberg, D., D. G. Rand, and A. Dreber (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review 102*(2), 720–49.

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review 95*(1), 384–394.

Greif, A. (1993). Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition. *American Economic Review 83*(3), 525–548.

Güth, W., M. 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*, 1023–1042.

Kandori, M. (1992). Social norms and community enforcement. *Review of Economic Studies 59*(1), 63–80.

Kartal, M. and W. Müller (2018). A new approach to the analysis of cooperation under the shadow of the future: Theory and experimental evidence. Available at SSRN: `https://ssrn.com/abstract=3222964`.

Kehoe, T. J. and D. K. Levine (1993). Debt-constrained asset markets. *Review of Economic Studies 60*(4), 865–888.

Khadjavi, M. and A. Lange (2013). Prisoners and their dilemma. *Journal of Economic Behavior & Organization 92*(2013), 163–175.

Kreps, D. M. (1990). Corporate culture and economic theory. In J. Alt and K. Shepsle (Eds.), *Perspectives on positive political economy*. Cambridge University Press.

Levine, D. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics 112*, 593–622.

McKelvey, R. D., A. M. McLennan, and T. L. Turocy (2016). Gambit: Software tools for game theory, version 16.0.1. Available online at `www.gambit-project.org`.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior 10*, 6–38.

Oskamp, S. (1974). Comparison of sequential and simultaneous responding, matrix, and strategy variables in a prisoner's dilemma game. *Journal of Conflict Resolution 18*(1), 107–116.

Potters, J., M. Sefton, and L. Vesterlund (2005). After you—Endogenous sequencing in voluntary contribution games. *Journal of Public Economics 89*(8), 1399–1419.

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

Romero, J. and Y. Rosokha (2018). Constructing strategies in the indefinitely repeated prisoner's dilemma game. *European Economic Review 104*, 185–219.

Roth, A. E. and J. K. Murnighan (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology 17*(2), 189–198.

Thomas, J. and T. Worrall (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory 51*(2), 367–390.

Turocy, T. L. (1995). A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior 51*(2), 243–263.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review 80*(1), 234–248.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *Quarterly Journal of Economics 106*(3), 885–910.

Vespa, E. and A. J. Wilson (2019). Experimenting with the transition rule in dynamic games. *Quantitative Economics 10*(4), 1825–1849.

Wen, Q. (2002). A folk theorem for repeated sequential games. *Review of Economic Studies 69*(2), 493–512.

Wilson, A. J. and E. Vespa (2020). Information transmission under the shadow of the future: An experiment. *American Economic Journal: Microeconomics 12*(4), 75–98.

Online appendix

A Experimental procedures

In all sessions, participants first took part in ten binary dictator games (DG). The dictator could either choose to keep 50 points for oneself and give 12 points to the recipient, or to keep c points and give c points to the recipient. Across the ten DGs c ranged from 30 to 48 points with increments of two points. All participants made choices in the role of dictator. They were explained that at the end of the session pairs would be randomly formed and roles would randomly assigned to determine the payment.

The experiment then proceeded to the main part. The Seq treatments covered seven sessions of 60 participants and the Sim treatments covered two sessions with 40 participants and two with 50 participants. All participants played 50 repeated games.[30] We chose a large number of repeated games per session because previous studies have shown the importance of learning (Dal Bó and Fréchette, 2011). At the beginning of the session participants were randomly assigned to matching groups of ten. Participants were not aware of the matching groups. At the beginning of each repeated game pairs were randomly formed within the matching groups. Within the same session, participants in different matching groups faced different mutual cooperation payoffs. This minimizes possible session effects that may for example stem from a correlation between tendency to cooperate and preferences for a particular time slot.

The level of understanding of the instructions was tested through a set of non-incentivized control questions about the PD's parameters (payoffs and continuation probability), and about the matching protocols within and across repeated games. Participants were not time constrained and were allowed to proceed with the experiment only after they correctly answered all questions. We kept record of the number of times that a participant submitted the answers to the control questions with at least one mistake. Below we report an English translation of the control questions:

1. How many points do you earn in a round if both you and the other participant choose A?

2. How many points do you earn in a round if both you and the other participant choose B?

3. How many points do you earn in a round if you choose A and the other participant chooses B?

[30] In one of the sessions the laboratory assistant accidentally implemented 51 repeated games.

4. How many points do you earn in a round if you choose B and the other participant chooses A?

5. You are in round two of a match. What is the probability that the match ends after this round?

6. True or false? I will be paired with the same participant in all rounds of a match.

7. True or false? I will be paired with the same participant across all matches.

At the end of the session we conducted a short survey, collecting information on gender, origin, age, and educational background. We also asked participants to self-report how risk averse they are. In particular, we asked :"How much of a risk taker would you evaluate yourself on a scale from 1 to 6?"

The final payment of a participant was determined by her earnings in one of the ten randomly drawn DGs plus the total earnings over all rounds of a randomly drawn repeated game. Each point earned during the experiment was worth 0.1€. Participants received a show-up fee of 6€ in the treatments with $\delta = 0.5$ and of 7€ in the treatments with $\delta = 0.75$. The show-up fee was larger in the latter treatments because these sessions took longer.

B Translated instructions

B.1 Instructions simultaneous games, $\delta = 0.5$

General instructions

Welcome to the experiment.

All participants receive the same instructions. Please read them carefully.

Do not communicate with any of the other participants during the entire experiment and turn off your cell phone. If you have questions, raise your hand, and wait until the experimenter comes to you to answer your question in private.

You receive a show-up fee of €6. The amount of money you earn on top of this depends on decisions made by you and other participants. Earnings are expressed in points during the experiment. Points convert to Euros in the following way: 10 points = €1. You will be paid your earnings in cash at the end of the experiment. The experiment is anonymous. Your identity will not be revealed to other participants and the identity of others will not be revealed to you.

The experiment consists of two parts. For both parts, you will get a separate set of instructions. The instructions for Part 1 are on the next page. The instructions for Part 2 will be distributed after Part 1 has finished.

Instructions Part 1

There are two players: player 1 and player 2. The task of player 1 is to choose between UP or DOWN. Player 2 is a passive player.

- If player 1 chooses UP, both players earn the same amount, which will be between 30 and 48 points depending on the scenario.
- If player 1 chooses DOWN, player 1 earns 50 points and player 2 earns 12 points.

In the experiment, there will be 10 scenarios. The figures below show the payoffs in points for both players in each of these 10 scenarios. P1 refers to player 2 and P2 to player 2.

For each scenario we ask you the following: if you have the role of player 1, what do you choose, UP or DOWN?

In order to calculate your earnings, the computer program randomly divides all participants into pairs of a player 1 and a player 2. You will be paid for the role of player 1 or player 2. At the point of decision-making, you don't know which role you have. Also, the computer program randomly selects 1 out of the 10 scenarios that will be used for payment. At the point of decision-making, you don't know which scenario is selected.

At the end of the experiment, you will be informed about the role for which you are paid, the randomly selected scenario and your earnings in Part 1.

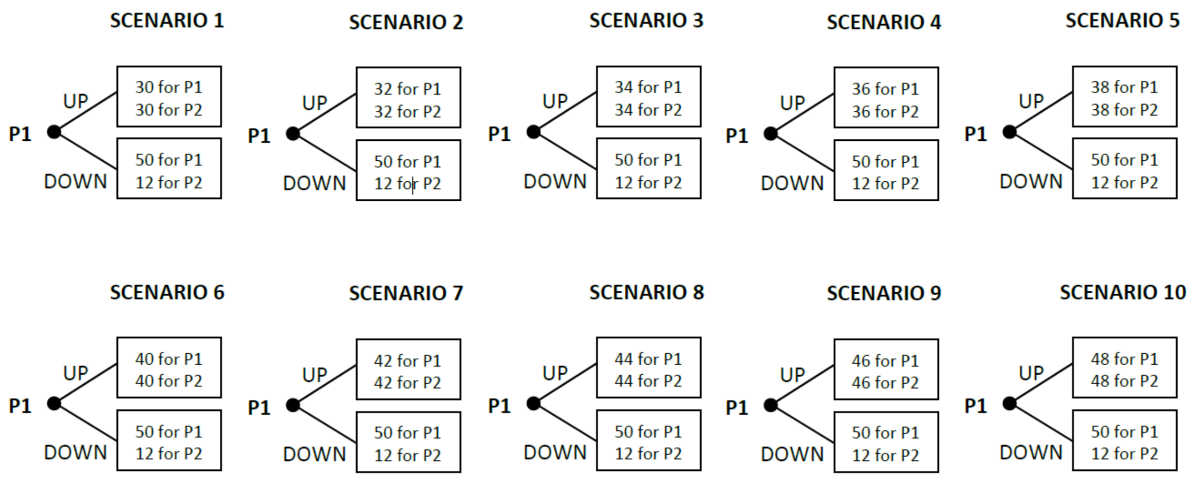

Instructions Part 2

In Part 2 you will make decisions in several sequences of rounds. Each sequence of rounds is referred to as a **match**.

In each match, you will be paired with another participant for one or more rounds. Within a match, pairs remain the same. There will be 50 matches in total, and after each match you are randomly paired with another participant for a new match.

At the end of the experiment, one of the 50 matches will be randomly selected for payment by the computer program. Your payment depends on the total points you have earned in that match.

Choices and earnings

In each round of a match, you and the paired participant make a choice between option A and option B.

Earnings in a round will be indicated on your computer screen in a table like the one below with Z>Y>X>W:

- If both of you choose A, you both earn Y points.
- If you choose A and the other chooses B, you earn W points and the other earns Z points.
- If you choose B and the other chooses A, you earn Z points and the other earns W points.
- If both of you choose B, you both earn X points.

The table is the same for all participants you will be paired with, and remains the same throughout Part 2.

Table: Earnings in points with Z>Y>X>W

Both choose A	Y	Y
You choose A and other chooses B	W	Z
You choose B and other chooses A	Z	W
Both choose B	X	X

At the end of each round, you will get to see the choice of the paired participant and your earnings in points in that round. You will also get to see the history of choices within the current match.

Number of rounds in a match

The number of rounds in a match is determined **randomly**. At the end of each round, there is a 50% probability that the match continues for at least another round. The computer virtually tosses a fair coin (50% probability of landing on heads and 50% probability of landing on tails) and the outcome of the coin toss will appear on your screen at the end of each round. If the outcome of the coin toss is heads, the match continues to a next round. If the outcome of the coin toss is tails, the match ends.

Control questions

Before decision-making in Part 2 starts, you will be asked to answer a number of control questions on the computer screen. Once everyone has answered all questions correctly, Part 2 starts.

B.2 Instructions sequential games, $\delta = 0.5$

General instruction and instructions for Part 1 are identical to those in Section B.1.

Instructions Part 2

In Part 2 you will make decisions in several sequences of rounds. Each sequence of rounds is referred to as a **match**.

In each match, you will be paired with another participant for one or more rounds. Within a match, pairs remain the same. There will be 50 matches in total, and after each match you are randomly paired with another participant for a new match.

At the end of the experiment, one of the 50 matches will be randomly selected for payment by the computer program. Your payment depends on the total points you have earned in that match.

Choices and earnings

In each match you will make decisions in the role of player 1 or player 2. Before each match starts your role will be randomly selected by the computer program and indicated on the screen. It will remain the same throughout that match.

In each round, player 1 and player 2 make a choice between option A and option B.

If you are player 1, you make this choice unconditionally, so you simply choose between A and B.

If you are player 2, you can condition your choice on the choice of player 1. This means you will observe the choice of the other before making your choice between A and B.

Earnings in a round will be indicated on your computer screen in a table like the one below with Z>Y>X>W:

- If both of you choose A, you both earn Y points.
- If you choose A and the other chooses B, you earn W points and the other earns Z points.
- If you choose B and the other chooses A, you earn Z points and the other earns W points.
- If both of you choose B, you both earn X points.

The table is the same for all participants you will be paired with, and remains the same throughout Part 2.

Table: Earnings in points with Z>Y>X>W

Both choose A	Y	Y
You choose A and other chooses B	W	Z
You choose B and other chooses A	Z	W
Both choose B	X	X

At the end of each round, you will get to see your earnings in points in that round. Participants with the role of player 2 will get to see the choice of the paired player 1 in that round, and participants with the role of player 1 will get to see the choice of the paired player 2 in that round. You will also get to see the history of choices within the current match.

Number of rounds in a match

The number of rounds in a match is determined **randomly**. At the end of each round, there is a 50% probability that the match continues for at least another round. The computer virtually tosses a fair coin (50% probability of landing on heads and 50% probability of landing on tails) and the outcome of the coin toss will appear on your screen at the end of each round. If the outcome of the coin toss is heads, the match continues to a next round. If the outcome of the coin toss is tails, the match ends.

Control questions

Before decision-making in Part 2 starts, you will be asked to answer a number of control questions on the computer screen. Once everyone has answered all questions correctly, Part 2 starts.

C Supplement on theory

C.1 Standard theory of a repeated sequential PD

We illustrate that the threshold for mutual cooperation to be an equilibrium outcome is the same under sequential decision-making than under simultaneous decision-making by comparing the expected payoff of a grim trigger strategy (GT) to that of *always defect* (AD). GT is generally defined as follows: "choose C on the first move and continue to do so on future moves as long as both players choose C; if one of the players chooses D, then switch to D forever after" (see for example Dal Bó and Fréchette, 2011). This strategy can be implemented as follows for the first mover in a sequential PD: "choose C in round 1 and continue to do so in round $t > 1$ as long as both players chose C in round $t-1$; if one of the players chose D in round $t-1$, choose D in t and forever after." For the second mover, a GT strategy is implemented as follows: "choose C (D) in round t if the first mover chooses C (D) in round t; choose D unconditionally in round t and forever after if one of the players chose D in round $t-1$".

Both players playing GT constitutes a subgame perfect equilibrium (SPE) if the rate at which players discount the future is sufficiently low, that is, if discount factor δ is sufficiently high (see Propositions 4 and 5 in Friedman, 1971). Given that the first mover plays GT, the second-mover expected payoff of GT is higher than that of AD if:

$$\begin{aligned} c + \delta c + \delta^2 c + \ldots &\geq t + \delta d + \delta^2 d + \ldots \\ c + \frac{\delta}{1-\delta} c &\geq t + \frac{\delta}{1-\delta} d \\ \delta &\geq \frac{t-c}{t-d} \equiv \delta^*. \end{aligned}$$

For the first mover the expected payoff of GT is higher than AD, given that the second mover plays GT, if:

$$\begin{aligned} c + \delta c + \delta^2 c + \ldots &\geq d + \delta d + \delta^2 d + \ldots \\ \frac{c}{1-\delta} &\geq \frac{d}{1-\delta}, \end{aligned}$$

which holds by definition. The condition thus reduces to $\delta \geq (t-c)/(t-d) \equiv \delta^*$ (see also Wen, 2002, who proves a folk theorem for repeated sequential games in general).

C.2 Basin of attraction

We follow Dal Bó and Fréchette (2011) and simplify the repeated simultaneous PD to a game with two strategies, namely always defect (AD) and a conditionally cooperative strategy (CC) à la GT. The basin of attraction of AD is calculated as the maximum probability of the partner using a CC strategy that makes it optimal for a player to always defect. If we assume that p is the probability that the partner uses CC, then the expected payoff of CC is larger than that of using the AD strategy if:

$$p(c+\delta c+\dots)+(1-p)(s+\delta d+\dots) > p(t+\delta d+\dots)+(1-p)(d+\delta d+\dots)$$
$$p\left(c+\frac{\delta c}{1-\delta}\right)+(1-p)\left(s+\frac{\delta d}{1-\delta}\right) > p\left(t+\frac{\delta d}{1-\delta}\right)+(1-p)\left(d+\frac{\delta d}{1-\delta}\right)$$
$$p > \frac{d-s}{c+d-t-s+\frac{\delta(c-d)}{1-\delta}} \equiv \bar{p}. \quad (1)$$

It can easily be seen that if $\delta < (t-c)/(t-d) \equiv \delta^*$, $\bar{p} > 1$, which implies that AD is the optimal strategy then. If $\delta > \delta^*$, there exists a $0 < \bar{p} < 1$ so that CC is optimal for $p > \bar{p}$.

To use the concept of the basin of attraction in the sequential PD, we also simplify the repeated sequential PD to a game with two strategies, namely AD and CC. We start with calculating the basin of attraction of AD for the second mover by calculating the maximum probability of the first mover using a CC strategy that makes it optimal for the second mover to always defect. If we assume that p_1 is the probability that the first mover uses CC, then the expected payoff of CC for the second mover is larger than that of using the AD strategy if:

$$p_1(c+\delta c+\dots)+(1-p_1)(d+\delta d+\dots) > p_1(t+\delta d+\dots)+(1-p_1)(d+\delta d+\dots)$$
$$c+\frac{\delta c}{1-\delta} > t+\frac{\delta d}{1-\delta}$$
$$\delta > \frac{t-c}{t-d} \equiv \delta^*. \quad (2)$$

The second mover will thus be "fully attracted" to AD if $\delta < \delta^*$ and to the CC strategy if $\delta > \delta^*$. The implication for the first mover (in a complete information environment) is that he will also be "fully attracted" to AD if $\delta < \delta^*$ and to the CC strategy if $\delta > \delta^*$. The same calculations hold if instead of using a CC strategy, the second mover would use a TFT strategy or another strategy with limited punishment.

C.3 Quantal-response predictions

Assume that the probability that a player chooses l is equal to $P_l = \frac{e^{\lambda} E\pi_l}{\sum_l e^{\lambda} E\pi_l}$ with expected payoff $E\pi_l$ calculated on the basis of choice probability P_l. Parameter $\lambda \in (0, \infty)$ stands for the degree of precision of decision-making, so is inversely related to the degree of noise (McKelvey and Palfrey, 1995). Quantal-response equilibrium predictions for the reduced repeated games, with k either equal to AD or to CC, are shown in Figure C.1. In the figure, the likelihood of AD is shown as a function of λ. On the one hand, the figure shows that the predicted effect of sequentiality on the likelihood of AD, thus on the cooperation rate, is in line with the prediction based on SizeBAD. On the other hand, the figure shows that the likelihood that first and second movers in Seq choose AD is predicted to decrease as c or δ increases even if $\delta > \delta^*$.

Figure C.1: Quantal-response predictions.

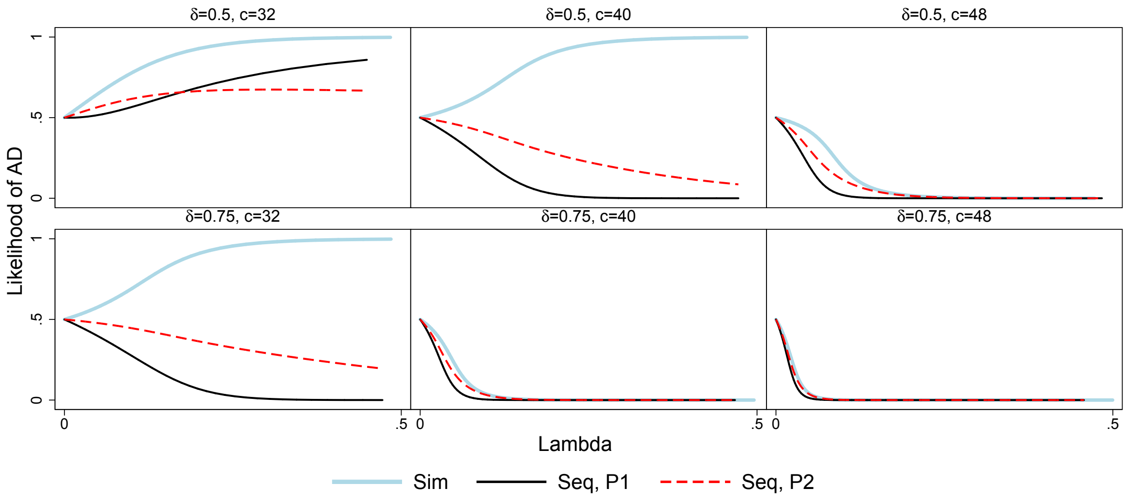

Notes: Calculated using Gambit version 15 (McKelvey et al., 2016).

C.4 Heterogeneity in other-regarding preferences

Assume that player i has a commonly known utility of the form

$$U_i = \begin{cases} \pi_i - \rho_i(\pi_i - \pi_j) & \text{if } \pi_i > \pi_j, \\ \pi_i + \sigma_i(\pi_j - \pi_i) & \text{if } \pi_i < \pi_j, \\ \pi_i & \text{otherwise,} \end{cases} \tag{3}$$

with $\rho_i = \rho_j$ and $\sigma_i = \sigma_j$, and $\sigma_i \leq \rho_i \leq 0$, $\sigma_i \leq 0 \leq \rho_i \leq 1$ or $0 \leq \sigma_i \leq \rho_i \leq 1$. Parameter ρ_i indicates player i's preference in cases in which she earns more than her partner, and σ_i indicates her preference in cases in which she earns less than her partner. The utility function corresponds to that of Charness and Rabin (2002), but without the reciprocity component. Restricting σ_i and ρ_i as mentioned above allows for three types of players, namely competitive types, difference averse types, and types concerned about efficiency. To calculate SizeBAD, we assume again that a player i is faced with a choice between strategy AD and strategy CC at the start of a repeated PD.

For the simultaneous-move PD, the condition under which the expected payoff of CC is larger than that of AD now becomes:

$$p\left(c + \frac{\delta c}{1-\delta}\right) + (1-p)\left(s + \frac{\delta d}{1-\delta} + \sigma_i(t-s)\right) > p\left(t + \frac{\delta d}{1-\delta} - \rho_i(t-s)\right) + (1-p)\left(d + \frac{\delta d}{1-\delta}\right)$$

$$p > \frac{d - s - \sigma_i(t-s)}{c + d - t - s + \frac{\delta(c-d)}{1-\delta} + (\rho_i - \sigma_i)(t-s)} \equiv \tilde{p}(\rho_i, \sigma_i). \tag{4}$$

As can be seen from condition 4, the threshold above which CC is preferred over AD is a negative function of ρ_i for a given σ_i. If $\rho_i > 0$, which implies that player i dislikes having more money than player j, player i prefers CC more easily than if $\rho = 0$. If, in addition $\sigma_i = 0$, then it holds that $\tilde{p} < \bar{p}$. Instead, if $\rho_i < 0$, which would imply that player i prefers to have a higher payoff than the partner, player i prefers CC less easily than if $\rho_i = 0$. For example, if $\sigma_i = 0$, it holds that $\tilde{p} > \bar{p}$. If we focus on the effect of σ_i on $\tilde{p}$, it can be shown that for a given ρ_i, $\tilde{p}$ decreases as σ_i increases. As σ_i increases, player i is thus relatively more inclined to choose CC than AD.

For second movers in the sequential-move PD, the condition under which the expected payoff of CC is larger than that of AD only depends on ρ_i because π_{P2} is never lower than π_{P1}. Player i chooses CC in the role of second mover if:

$$c + \frac{\delta c}{1-\delta} > t + \frac{\delta d}{1-\delta} - \rho_i(t-s)$$
$$\delta > \frac{t - c - \rho_i(t-s)}{t - d - \rho_i(t-s)} \equiv \tilde{\delta}^*(\rho_i). \tag{5}$$

If $\rho_i > 0$, then $\delta^*(\rho_i) < \delta^*$, implying that the second mover prefers CC over AD more easily than if $\rho_i = 0$. If $\rho_i < 0$, then $\delta^*(\rho_i) > \delta^*$, implying that CC is now less easily preferred. If we assume that ρ_i is distributed in interval $[-1, 1]$, the implication is that the conditional cooperation rate in a population of players can now be in between 0 and 100 percent. Moreover, for a given distribution of ρ_i, condition 5 is more easily satisfied, the higher c or the higher δ.

Finally, we focus on the first mover. Under complete information, which refers to players being informed about each other's type, the cooperation rate of first movers and thus the overall cooperation rate corresponds one-to-one with the second-mover conditional cooperation rate. Under the more realistic assumption that players are uncertain about each other's type, the cooperation rate of first movers does not necessarily correspond to the conditional cooperation rate (e.g. Kartal and Müller, 2018). Assume, for example, that first movers are uncertain about the type of second movers but know the distribution of types in the population. Intuitively, removing the information about the second-mover type exposes the first mover to strategic risk. This has two implications. First, the first mover's choice to use CC becomes a matter of comparing the expected payoff with that of AD rather than merely copying the second-mover strategy. Second, σ_i enters the trade-off. Specifically, if p_2 represents the probability that the second mover uses CC, then player i prefers CC over AD in the role of first mover if:

$$p_2\left(c+\frac{\delta c}{1-\delta}\right)+(1-p_2)\left(s+\frac{\delta d}{1-\delta}+\sigma_i(t-s)\right)>d+\frac{\delta d}{1-\delta}$$

$$p_2>\frac{d-s-\sigma_i(t-s)}{c-s+\frac{\delta(c-d)}{1-\delta}-\sigma_i(t-s)}\equiv\tilde{p_2}(\sigma_i). \quad (6)$$

Condition 6 is more easily satisfied, implying that player i is more likely to choose CC in the role of first mover, the higher σ_i. The implication for the aggregate cooperation rate of first movers is that it will be between 0 and 1 depending on the distribution of σ_i and the game parameters.

Example

Take as an example the six parameterizations in the PD games in our experiment and consider the following distribution of 50% (near-)payoff-maximizing, 20% spiteful and 30% pro-social players: $\rho_i \in (-0.13, 0.29)$ for 50% of the players, $\rho_i \in (-0.55, -0.13)$ for 20% of the players, and $\rho_i \in (0.29, 0.55)$ for 30% of the players. With this distribution, the conditional cooperation rate of second movers would be equal to 0.30 in treatment $\delta = 0.5, c = 32$, to 0.80 in treatments $\delta = 0.75, c = 32$ and $\delta = 0.5, c = 40$, and to 1 in the three other treatments. Moreover, with this distribution an expected-payoff-maximizing first mover (i.e. with $\sigma_i = 0$) finds it optimal to choose AD in treatment $\delta = 0.5, c = 32$ and CC in all other treatments. To illustrate, Table C.1 gives an overview of predicted cooperation rates. As can be seen, comparative statics are much in line with the cooperation rates observed in the experiment. In particular, c and δ have a positive effect on the cooperation rate in Seq, and the conditional cooperation rate can be below 100% even if $\delta > \delta^*$. A major discrepancy left between predicted and observed cooperation rates is that the first mover cooperation rate is well below the second-mover conditional cooperation rate in treatment $\delta = 0.5, c = 32$. This is because the predictions do not take into account that, in the experiment, there is strategic uncertainty left for the first mover in Seq, related to not knowing the type of the second mover. If this is taken into account along the lines of solving (6), then the predicted first-mover cooperation rate would be 0.

Table C.1: Predicted cooperation rates with social preferences.

	$\delta = 0.5$			$\delta = 0.75$		
	$c = 32$	$c = 40$	$c = 48$	$c = 32$	$c = 40$	$c = 48$
Predicted						
Sim	[0-30%]	[0-80%]	[0-100%]	[0-80%]	[0-100%]	[0-100%]
P2 Seq	30%	80%	100%	80%	100%	100%
P1 Seq	30%	80%	100%	80%	100%	100%
Seq	30%	80%	100%	80%	100%	100%
Observed						
Sim	2.7%	15.8%	51%	24.8%	78.5%	95.5%
P2 Seq	43.9%	72.5%	93.2%	83.3%	89.3%	95.4%
P1 Seq	9.5%	61.8%	95.3%	69.8%	89%	95%
Seq	7%	53.7%	92.1%	65.1%	84.4%	93%

Notes: The table shows predicted cooperation rates and conditional cooperation rates (P2 Seq) under common knowledge if it is assumed that $\sigma_i = 0$ and $\rho_i \in (-0.13, 0.29)$ for 50% of the players, $\rho_i \in (-0.55, -0.13)$ for 20% of the players, and $\rho_i \in (0.29, 0.55)$ for 30% of the players. Cooperation rates observed in the first rounds of the last twenty repeated games of the experiment are included as well.

D Representative players versus heterogeneity

We evaluate how likely it is that observed distributions of individual-level conditional cooperation and defection rates, as shown in Figure 4 of the main text, can stem from individuals making random choices with different probabilities in different treatments. To do so, we first compare observed distributions of first-round conditional cooperation rates with simulated *iid* distributions using the following procedure:

1. For each treatment, consider first the N subjects who encountered cooperation by the matched first mover in the role of second mover across $m_i > 0$ first rounds, where i is a subject-specific identifier. Let $M = \sum_{i=1}^{N} m_i$ be the total number of choices made by the N subjects in the treatment.

2. For each treatment, simulate $N \times M$ conditional cooperation choices by drawing from a binomial distribution characterized by a success probability (i.e. simulated conditional cooperation equal to 1) corresponding to the overall conditional cooperation rate observed in that treatment.

3. Calculate for subject 1 to N in each treatment the simulated conditional cooperation rate based on the simulated decisions obtained in step 2.

4. Run for each treatment an OLS regression without a constant term and with standard errors clustered at the matching group level of the observed conditional cooperation rate on the simulated conditional cooperation rate.

5. Test for each treatment using a Wald test whether the coefficient estimated in step 4 is statistically significantly different from 1 and store the test's p-value.

6. Repeat steps 2 to 5 200 times for each treatment.

Figure D.1 shows an example of individual-level simulated conditional cooperation rates obtained from the above-described procedure. The simulated rates in Figure D.1 can be easily compared to the observed rates in Figure 4. It can be seen that the distributions of the simulated rates are generally smoother than the respective distributions of the observed rates, with smaller differences in conditional cooperation rates across subjects under both $\delta < \delta^*$ and $\delta > \delta^*$.

If the coefficient estimated in step 5 is not statistically significantly different from 1, we can conclude that the observed conditional cooperation rates' distribution is not significantly different from the simulated one. Instead, if the estimated coefficient is statistically different

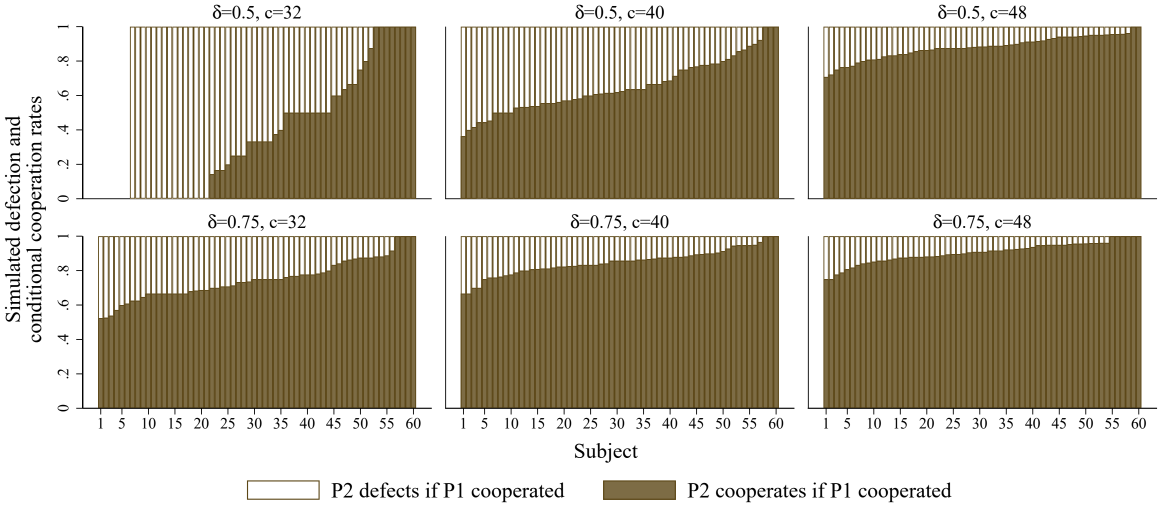

Figure D.1: Simulated conditional cooperation rates by subject.

Notes: The figure illustrates distributions of simulated conditional cooperation rates obtained with the above-described procedure. Patterns tend to be stable across different simulations.

from 1, then it can be concluded that the simulated distribution does not well approximate the observed distribution. Concerning defection rates observed in the first rounds, we follow the same procedure but replace conditional cooperation with defection. Figure D.2 reports for each treatment the distribution of the 200 p-values obtained from step 6 of the procedure.

We first focus on conditional cooperation rates. Figure D.2a shows that the distribution of p-values in $\delta = 0.5, c = 32$ is skewed towards values which are well below the 10% significance level. In particular, 97% of the p-values in our simulation is lower than 10%. This result implies that the observed distribution of conditional cooperation choices in $\delta = 0.5, c = 32$ is substantially different from the simulated distributions. We hereby confirm that consistent with a heterogeneity interpretation, the observed distribution of conditional cooperation choices in this treatment follows from a few subjects being highly motivated to conditionally cooperate rather than from many subjects randomizing. Instead, in the treatments with $\delta > \delta^*$, none of the p-values is below 10%. This result is consistent with a heterogeneity explanation because next to pro-social types also rational payoff maximizers have an incentive to conditionally cooperate if $\delta > \delta^*$.

Next, we consider defection rates. Figure D.2b shows that not just in treatment $\delta = 0.5, c = 32$ but also in treatments $\delta = 0.5, c = 48$, $\delta = 0.75, c = 40$ and $\delta = 0.75, c = 48$, a substantial proportion of the p-values in our simulation is now below 10%. The percentages are 18.5%, 16%, and 30.5%, respectively. We take this as additional evidence in favor of heterogeneity. The reason is that in these last three treatments, defection types can quite easily be separated from

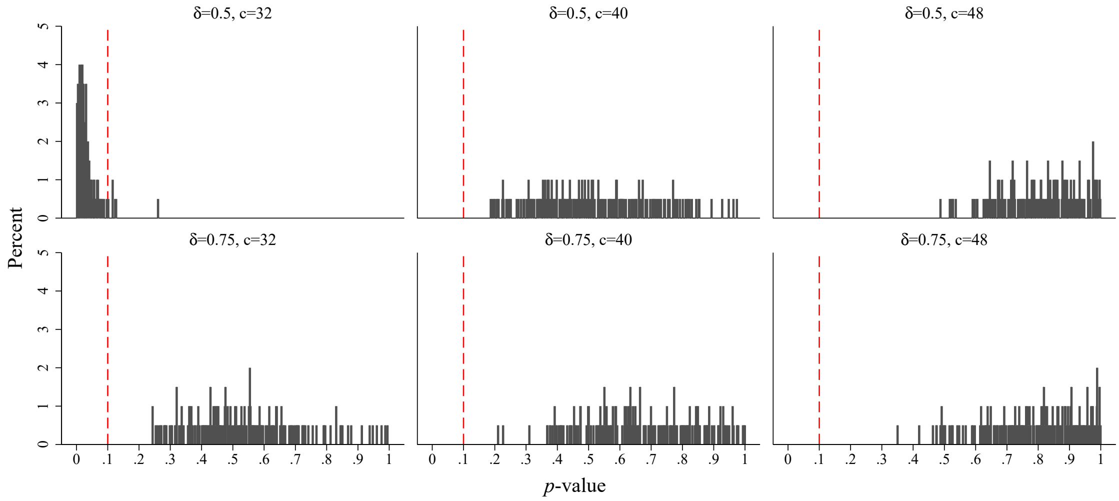

(a) Conditional cooperation

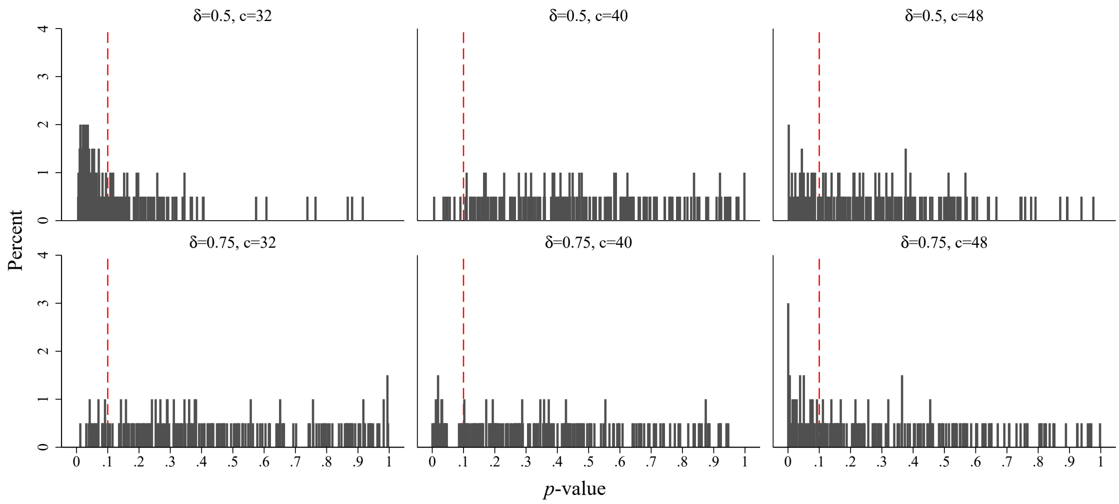

(b) Defection

Figure D.2: Observed versus simulated *iid* choices of second movers.

Notes: The figure reports distributions of p-values obtained from step 5 of the above-described procedure. Notice that the conditional cooperation rates used to simulate *iid* distributions in the six treatments are 0.40 in $\delta = 0.5, c = 32$, 0.64 in $\delta = 0.5, c = 40$, 0.89 in $\delta = 0.5, c = 40$, 0.78 in $\delta = 0.75, c = 32$, 0.85 in $\delta = 0.75, c = 40$, and 0.91 in $\delta = 0.75, c = 48$. The vertical dash lines highlight the 10% significance level.

rational payoff maximizers and pro-social types. The result that much more than 10% of the p-values in these treatments is lower than 10% shows that defection choices in these treatments are not the outcome of randomization but are well-motivated decisions. In $\delta = 0.5, c = 40$, $\delta = 0.75, c = 32$, in which behavior does not translate easily into types, the distribution of p-values is closer to uniform and defection types cannot be easily separated from other types.

E Estimation of Repeated-Game Strategies

E.1 General Description

We have used a framework that builds upon the assumption that players choose AD or CC at the start of the repeated game. By estimating the frequencies of repeated-game strategies that participants have used in our experiment, we provide evidence that the simplification is justified.[31] Theoretically, the threshold δ^* above which mutual cooperation is supported in a Nash equilibrium is different if other cooperative strategies than CC are used. For example, it can be shown that the first mover in Seq using the strategy "Defect in round 1 and then tit-for-tat" (D-TFT) in combination with the second mover applying "Cooperate in round 1 and then TFT" (C-TFT) or "Cooperate in round 1 and then GT" (C-GT) constitutes an equilibrium leading to mutual cooperation for the parameters in our experiment. In contrast, equivalent combinations of strategies in Sim can at best constitute an equilibrium with partial cooperation.[32]

The empirical identification of repeated-game strategies is notoriously challenging because the experimenter only observes choices. By using maximum likelihood, we estimate the relevance of a set of predetermined strategies within each treatment using an approach common in the literature (e.g. Dal Bó and Fréchette, 2011; Fudenberg et al., 2012; Bigoni et al., 2015). See Section E.2 for a detailed description. In an initial step, we estimated strategies AD, GT, TFT, and *always cooperate* (AC) for all players, strategy D-TFT for players in Sim and first movers in Seq, and strategies C-TFT and C-GT for second movers in Seq. Given that the latter two strategies were estimated to have a frequency close to zero, we decided to leave them out and focus on the first five strategies.[33] Table E.1 reports the results for the treatment with $\delta < \delta^*$ and jointly for all treatments in which $\delta > \delta^*$. It is important to mention that heterogeneity in strategies across treatments with $\delta > \delta^*$ is quite substantial. Nevertheless, we decided to pool data from these treatments because this simplification is not crucial for our discussion.[34] Estimations are based on choices from all rounds of the last 20 repeated games. Results are qualitatively similar if based on all repeated games (see Table E.2).

The first observation that can be made in the table is that the estimated share of cooperative

[31] Romero and Rosokha (2018) and Dal Bó and Fréchette (2019) directly elicit strategies and show that these correspond to a large extent to the estimated strategies.

[32] The reason is that in Seq, there is no punishment of the first mover's initial defection so that from round 2 onwards both players cooperate. In Sim, the player using TFT punishes in round 2 the first-round defection by the partner by defecting oneself in round 2, which sets in a series of switching back and forth between cooperation and defection. To reach full cooperation against a player who uses D-TFT in Sim, one more round of patience is needed.

[33] We also estimated more complex strategies for players in Sim and first movers in Seq, including several memory-two strategies, but these exercises do not give us much additional insight.

[34] Estimates by treatment are available upon request.

Table E.1: Estimated repeated-game strategies.

	$\delta < \delta^*$			$\delta > \delta^*$		
	Sim	Seq, P1	Seq, P2	Sim	Seq, P1	Seq, P2
AD	0.727	0.404	0.588	0.336	0.073	0.111
	(0.000)	(0.046)	(0.000)	(0.033)	(0.202)	(0.088)
AC	0.000	0.015	0.000	0.028	0.059	0.000
	(0.356)	(0.332)	(0.481)	(0.294)	(0.310)	(0.500)
GT	0.000	0.001	0.247	0.296	0.222	0.328
	(0.356)	(0.471)	(0.057)	(0.036)	(0.136)	(0.063)
TFT	0.000	0.034	0.165	0.201	0.566	0.562
	(0.356)	(0.204)		(0.085)	(0.012)	
D-TFT	0.273	0.546	–	0.139	0.080	–
γ	0.262	0.350	0.234	0.351	0.312	0.356
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Coop. strat.	0.273	0.596	0.412	0.664	0.927	0.889

Notes: The table shows estimates from maximum likelihood based on data from all rounds of the last 20 repeated games (with p-values in parentheses). Parameter $\gamma \in [0, \infty)$ captures the quality of fit between observed and prescribed behavior; the higher γ, the worse the fit. The share of cooperative strategies (coop. strat.) is equal to 1 minus the share of AD.

strategies of second movers out of all strategies in the game where $\delta < \delta^*$ is very close to the conditional cooperation rate based on all rounds of the last 20 repeated games (see Table F.4 in the Appendix). Likewise, for $\delta > \delta^*$ the estimated share is in the range of conditional cooperation rates reported for the five treatments in which $\delta > \delta^*$. These findings show that most of the cooperative strategies of second movers fall within the category of CC. A second observation is that in both Sim and Seq AD is overall used less frequently and (conditionally) cooperative strategies are used more frequently in the games with $\delta > \delta^*$ than in the game with $\delta < \delta^*$, which is consistent with the evidence on cooperation rates reported in the previous section.

If we focus on the effect of sequentiality on the type of strategies adopted, we see that cooperative strategies are generally used more frequently by first and second movers in Seq than by players in Sim. This finding does not come as a surprise for $\delta > \delta^*$ and is consistent with results reported on cooperation rates. However, for $\delta < \delta^*$ the finding is not trivial because the observed cooperation rate does not differ between Sim and Seq. Remarkably, the strategy most frequently used by first movers in the game where $\delta < \delta^*$ is D-TFT (54.6% of the time) instead of AD (40.4% of the time). This finding contrasts to Sim, where AD is more common than D-TFT (72.7% versus 27.3%).

Other insights from the estimations are related to strategies used within the set of conditionally cooperative strategies. Comparing the set of conditionally cooperative strategies between Sim and Seq reveals that different types of such strategies are used. Overall, TFT is more prominent in Seq than in Sim, among both first and second movers. In Sim, GT and TFT are roughly equally popular for $\delta > \delta^*$ and D-TFT is most common for $\delta < \delta^*$. In Seq, first movers tend to prefer D-TFT over TFT if $\delta < \delta^*$ and TFT over D-TFT if $\delta > \delta^*$. Second movers tend to use GT more frequently than TFT if $\delta < \delta^*$ (24.7% versus 16.5%) and swap if $\delta > \delta^*$ (32.8% versus 56.2%).

Table E.2: Estimated repeated-game strategies over all repeated games.

	$\delta < \delta^*$			$\delta > \delta^*$		
	Sim	Seq, P1	Seq, P2	Sim	Seq, P1	Seq, P2
AD	0.695	0.285	0.640	0.302	0.075	0.123
	(0.000)	(0.050)	(0.000)	(0.034)	(0.153)	(0.123)
AC	0.000	0.017	0.000	0.012	0.025	0.000
	(0.438)	(0.320)	(0.212)	(0.306)	(0.313)	(0.493)
GT	0.000	0.017	0.283	0.288	0.165	0.220
	(0.438)	(0.193)	(0.002)	(0.019)	(0.205)	(0.041)
TFT	0.000	0.017	0.077	0.251	0.615	0.657
	(0.438)	(0.192)		(0.033)	(0.002)	
D-TFT	0.305	0.665	–	0.146	0.121	–
γ	0.356	0.421	0.289	0.417	0.400	0.391
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Coop. strat.	0.305	0.715	0.360	0.698	0.925	0.877

Notes: Estimates from maximum likelihood based on all rounds of all repeated games. p-values are in parentheses. The share of cooperative strategies (coop. strat.) is equal to 1 minus the share of AD.

E.2 Details about the Methodology

We provide a description of the Strategy Frequency Estimation Method (SFEM) proposed by Dal Bó and Fréchette (2011), which we use to estimate shares of repeated-game strategies. We postulate a strategy set $\mathcal{S}$ and assume that participants in the experiment could only choose their actions according to a strategy $s \in \mathcal{S}$. For each participant, the sequence of actions prescribed by all postulated strategies is contrasted to the sequence of actions observed in the experiment. To define a prescribed sequence of actions, the behavior of each participant's partner in the previous round in a repeated game (or the previous stage for second movers in Seq) is taken as given. The maximum likelihood estimation takes into account that, in each round, participants might make a mental error in the implementation of their strategy, thereby deviating from the prescribed action.

More precisely, let $d^i_{gt}(\mathbf{h})$ and $s^i_{gt}(\mathbf{h})$ be respectively participant i's observed action and participant i's action as prescribed by strategy s in round t of repeated game g for a given history $\mathbf{h}$. In any round, the probability that the observed action is equal to the prescribed one is modeled as follows:

$$\Pr\left(d^i_{gt}(\mathbf{h}) = s^i_{gt}(\mathbf{h})\right) = \frac{1}{1+\exp\left(\frac{-1}{\gamma}\right)} \equiv \beta. \tag{7}$$

Thus, $1-\beta$ can be interpreted as the probability of making a mental error. The parameter $\gamma > 0$, which is to be estimated, captures the quality of fit between observed and prescribed behavior. As $\gamma \to 0$, $\beta \to 1$, implying that the action prescribed by strategy s fits the experimental observation perfectly. Conversely, as $\gamma \to \infty$, $\beta \to 0.5$, i.e., a random draw fits perfectly the experimental observation. Starting from the comparison between observed and prescribed actions in each round (Equation 7), we can extend the comparison to all rounds of interest. Let y^i_{gt} be an indicator equal to 1 if prescribed and observed actions coincide, and 0 otherwise. Given Equation 7, the likelihood of observing strategy s for participant i is

$$p_i(s) = \prod_g \prod_t \left(\frac{1}{1+\exp\left(\frac{-1}{\gamma}\right)}\right)^{y^i_{gt}} \left(\frac{1}{1+\exp\left(\frac{1}{\gamma}\right)}\right)^{1-y^i_{gt}}. \tag{8}$$

Finally, we aggregate at the population level and obtain the log-likelihood $\sum_i \ln(\sum_s \phi_s p_i(s))$, with ϕ_s representing the frequency of strategy s in the experimental data. Maximum likelihood allows to estimate both parameters γ and ϕ_s.

F Supplementary tables

Table F.1: Cooperation rate by treatment

		Round 1					All rounds				
		$c = 32$		$c = 40$		$c = 48$	$c = 32$		$c = 40$		$c = 48$
Repeated games 1 to 50											
	Sim	6.2	$\lll$	29.1	$\ll$	55.4	6.1	$\lll$	23.9	$\lll$	46.5
		(0.7)		(7.6)		(8.7)	(0.3)		(4.6)		(7.5)
$\delta = 0.5$		$\wr$		$\wr$		⤊	$\wr$		⤊		⤊
	Seq	10.0	$\lll$	44.1	$\lll$	83.0	8.9	$\lll$	41.9	$\lll$	81.6
		(3.0)		(5.6)		(5.4)	(2.6)		(4.9)		(5.2)
	Sim	22.3	$\lll$	72.4	$\ll$	91.9	14.7	$\lll$	59.0	$\ll$	81.3
		(10.2)		(10.7)		(3.3)	(6.4)		(7.6)		(4.7)
$\delta = 0.75$		$\Uparrow$		$\wr$		$\wr$	⤊		$\wr$		$\wr$
	Seq	57.5	$\ll$	77.5	$\approx$	86.0	47.0	$\lll$	69.4	$<$	81.5
		(9.0)		(3.8)		(4.0)	(6.8)		(4.8)		(4.1)
Repeated games 31 to 50											
	Sim	2.7	$\lll$	15.8	$\lll$	51.0	2.6	$\lll$	12.8	$\lll$	43.5
		(1.8)		(3.9)		(8.1)	(1.4)		(1.1)		(5.0)
$\delta = 0.5$		$\wr$		⤊		⤊	$\wr$		⤊		⤊
	Seq	7.0	$\lll$	53.7	$\lll$	92.1	6.3	$\lll$	48.5	$\lll$	88.2
		(3.0)		(8.2)		(2.8)	(2.3)		(6.5)		(3.7)
	Sim	24.8	$\lll$	78.5	$\approx$	95.5	16.8	$\lll$	66.4	$<$	86.5
		(11.0)		(12.7)		(3.6)	(8.0)		(11.8)		(4.9)
$\delta = 0.75$		$\Uparrow$		$\wr$		$\wr$	$\wedge$		$\wr$		$\wr$
	Seq	65.1	$<$	84.4	$\approx$	93.0	46.9	$\ll$	74.1	$\ll$	89.1
		(10.3)		(5.1)		(3.7)	(7.8)		(6.5)		(4.0)

Notes: The unit of observation is a participant in a round. Differences between treatments are tested using probit regressions with standard errors clustered at the matching group level (in parentheses). $<$, $\ll$, and $\lll$ refer to $p < 0.1$, $p < 0.05$. and $p < 0.01$, respectively.

Table F.2: Cooperation rate by treatment including individual-level controls

			Round 1					All rounds				
			$c = 32$		$c = 40$		$c = 48$	$c = 32$		$c = 40$		$c = 48$
Repeated games 1 to 50												
	$\delta = 0.5$	Sim	6.5	$\ll$	30.8	$\lll$	55.7	6.3	$\ll$	25.6	$\lll$	46.6
			(2.7)		(6.8)		(10.5)	(2.2)		(5.3)		(10.0)
			≀		≀		≀	≀		∧∧		∧∧
		Seq	10.1	$\lll$	46.4	$\lll$	82.9	8.9	$\lll$	44.1	$\lll$	81.5
			(3.8)		(7.9)		(6.6)	(3.2)		(7.4)		(6.5)
	$\delta = 0.75$	Sim	22.4	$\lll$	75.6	$\lll$	91.6	14.8	$\lll$	61.7	$\lll$	81.1
			(8.5)		(10.3)		(6.6)	(5.7)		(7.8)		(7.3)
			∧∧∧		≀		∨∨∨	∧∧∧		≀		≀
		Seq	57.6	$\ll$	79.4	$\ll$	86.0	47.0	$\lll$	70.8	$\ll$	81.5
			(8.5)		(5.9)		(4.5)	(7.5)		(7.0)		(4.4)
Repeated games 31 to 50												
	$\delta = 0.5$	Sim	2.6	$<$	17.7	$\lll$	51.2	2.6	$<$	14.8	$\lll$	43.7
			(1.2)		(6.6)		(10.6)	(1.2)		(4.0)		(9.2)
			∨∨		∧∧∧		∧∧∧	∨		∧∧∧		∧∧∧
		Seq	7.1	$\lll$	56.3	$\lll$	91.8	6.3	$\lll$	51.4	$\lll$	88.1
			(3.8)		(10.4)		(4.2)	(3.0)		(8.9)		(5.0)
	$\delta = 0.75$	Sim	25.2	$\lll$	80.9	$<$	94.9	17.1	$\lll$	68.3	$\lll$	86.1
			(11.3)		(9.8)		(6.2)	(7.8)		(7.7)		(6.7)
			∧∧∧		≀		∨∨∨	∧∧∧		≀		≀
		Seq	65.2	$\approx$	86.5	$\approx$	93.0	46.9	$\ll$	75.7	$<$	89.1
			(11.3)		(8.0)		(4.3)	(9.2)		(9.4)		(4.4)

Notes: The unit of observation is a participant in a round. Differences between treatments are tested using probit regressions with standard errors clustered at the matching group level (in parentheses) and including individual-level controls for other-regarding preferences, risk preferences, proneness to mistakes and experienced length of repeated games. Other-regarding preferences are proxied by a pro-sociality indicator taking value 1 if the participant chooses an equal distribution in a dictator game with the same parameters of the stage-game PD, and value 0 if (s)he chooses the selfish option (see also Section A). Risk preferences are proxied by a continuous variable ranging from 1=very risk averse to 6=very risk seeking elicited through self-reports. Proneness to mistakes is proxied by a continuous variable counting the number of times that the participant submitted answers that had at least one mistake in the quiz with control questions ran before the experiment started. Experienced length of repeated games is measured by taking the difference between expected and median realized length of the first ten repeated games. $<$, $\ll$, and $\lll$ refer to $p < 0.1$, $p < 0.05$. and $p < 0.01$, respectively.

Table F.3: Cooperation rate by treatment including data from Dal Bó and Fréchette (2011)

		Round 1					All rounds				
		$c = 32$		$c = 40$		$c = 48$	$c = 32$		$c = 40$		$c = 48$
Repeated games 1 to 50											
	Sim	8.8	$\lll$	22.3	$\lll$	43.5	8.1	$\lll$	19.6	$\lll$	38.5
		(1.6)		(4.5)		(5.8)	(1.8)		(2.7)		(4.4)
$\delta = 0.5$		≀		⤊		⤊	≀		⤊		⤊
	Seq	10.0	$\lll$	44.1	$\lll$	83.0	8.9	$\lll$	41.9	$\lll$	81.6
		(3.0)		(5.6)		(5.4)	(2.6)		(4.9)		(5.2)
	Sim	23.9	$\lll$	67.1	$\lll$	88.6	17.6	$\lll$	58.9	$\lll$	78.8
		(5.3)		(9.2)		(2.5)	(3.4)		(6.1)		(3.3)
$\delta = 0.75$		⤊		≀		≀	⤊		≀		≀
	Seq	57.5	$\ll$	77.5	$\approx$	86.0	47.0	$\lll$	69.4	$<$	81.5
		(9.0)		(3.8)		(4.0)	(6.8)		(4.8)		(4.1)
Repeated games 31 to 50											
	Sim	4.4	$\lll$	17.1	$\lll$	41.4	3.9	$\lll$	16.3	$\lll$	36.9
		(1.5)		(3.2)		(6.6)	(1.5)		(2.5)		(4.6)
$\delta = 0.5$		≀		⤊		⤊	≀		⤊		⤊
	Seq	7.0	$\lll$	53.7	$\lll$	92.1	6.3	$\lll$	48.5	$\lll$	88.2
		(3.0)		(8.2)		(2.8)	(2.3)		(6.5)		(3.7)
	Sim	24.5	$\lll$	71.0	$\ll$	96.1	16.2	$\lll$	65.3	$\ll$	88.3
		(9.9)		(11.5)		(3.0)	(7.0)		(8.9)		(4.2)
$\delta = 0.75$		⇑		≀		≀	⤊		≀		≀
	Seq	65.1	$<$	84.4	$\approx$	93.0	46.9	$\ll$	74.1	$\ll$	89.1
		(10.3)		(5.1)		(3.7)	(7.8)		(6.5)		(4.0)

Notes: Data on Sim include data from the first up to 50 repeated games played in the experiment of Dal Bó and Fréchette (2011). Differences between treatments are tested using probit regressions with standard errors clustered at the matching group level (in parentheses). $<$, $\ll$, and $\lll$ refer to $p < 0.1$, $p < 0.05$. and $p < 0.01$, respectively.

Table F.4: Cooperation rates in Seq by role and treatment

		Round 1					All rounds				
		$c = 32$		$c = 40$		$c = 48$	$c = 32$		$c = 40$		$c = 48$
Repeated games 1 to 50											
	$\delta = 0.5$	13.2	$\lll$	52.4	$\lll$	87.2	12.2	$\lll$	47.7	$\lll$	84.8
		(3.6)		(5.2)		(5.0)	(3.2)		(4.9)		(4.6)
P1		∧∧∧		∧∧∧		≀	∧∧∧		∧∧∧		≀
	$\delta = 0.75$	63.1	$\ll$	83.1	$\approx$	89.7	50.0	$\lll$	72.3	$<$	83.5
		(8.6)		(3.2)		(3.3)	(7.1)		(4.6)		(3.9)
	$\delta = 0.5$	39.9	$\ll$	64.1	$\lll$	88.9	37.2	$\lll$	71.6	$\lll$	90.9
		(9.0)		(5.2)		(3.2)	(6.9)		(3.7)		(2.8)
P2		∧∧∧		∧∧∧		≀	∧∧∧		∧∧∧		≀
	$\delta = 0.75$	78.0	$\approx$	85.2	$\approx$	90.6	81.2	$\lll$	90.9	$\approx$	93.9
		(5.2)		(2.6)		(2.6)	(1.8)		(1.6)		(1.4)
Repeated games 31 to 50											
	$\delta = 0.5$	9.5	$\lll$	61.8	$\lll$	95.3	9.1	$\lll$	53.9	$\lll$	90.6
		(3.7)		(8.2)		(2.3)	(3.0)		(6.7)		(3.0)
P1		∧∧∧		∧∧∧		≀	∧∧∧		∧∧		≀
	$\delta = 0.75$	69.8	$\ll$	89.0	$\approx$	95.0	49.7	$\ll$	76.3	$\ll$	90.8
		(10.1)		(4.0)		(3.2)	(8.2)		(6.3)		(3.7)
	$\delta = 0.5$	43.9	$\ll$	72.5	$\lll$	93.2	35.8	$\lll$	78.4	$\lll$	94.0
		(11.1)		(4.1)		(2.6)	(7.6)		(2.5)		(2.3)
P2		∧∧∧		∧∧∧		≀	∧∧∧		∧∧∧		≀
	$\delta = 0.75$	83.3	$\approx$	89.3	$<$	95.4	82.4	$\lll$	93.3	$\approx$	95.6
		(4.5)		(3.2)		(1.7)	(1.4)		(1.2)		(1.4)

Notes: For first movers (P1) cooperation rates are reported and for second movers (P2) conditional cooperation rates are reported. Differences between treatments are tested using probit regressions with standard errors clustered at the matching group level (in parentheses). In round 1 of repeated game 50 standard errors could not be computed for P1 in $\delta = 0.5, c = 48$ nor for P2 in $\delta = 0.75, c = 32$ because of perfect fit. $<$, $\ll$, and $\lll$ refer to $p < 0.1$, $p < 0.05$. and $p < 0.01$, respectively.

Table F.5: Effect of c and δ on cooperation under $\delta > \delta^*$

	Round 1		All rounds	
	(1) Sim	(2) Seq	(3) Sim	(4) Seq
c	0.041***	0.022***	0.038***	0.031***
	(0.000)	(0.000)	(0.000)	(0.000)
δ	0.438***	0.168**	0.347***	0.131*
	(0.000)	(0.010)	(0.000)	(0.067)
Observations	3000	6000	9480	21120

Notes: The table reports marginal effects from probit regressions with standard errors clustered at the matching group level. The variable δ is a dummy taking value 1 when $\delta = 0.75$, and 0 otherwise. The variable c is a continuous variable ranging from 32 to 48. Data are based on the last 20 repeated games of the treatments with $\delta > \delta^*$. p-values are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

G Supplementary figures

Figure G.1: Effect of sequentiality on cooperation rate with expected outcomes

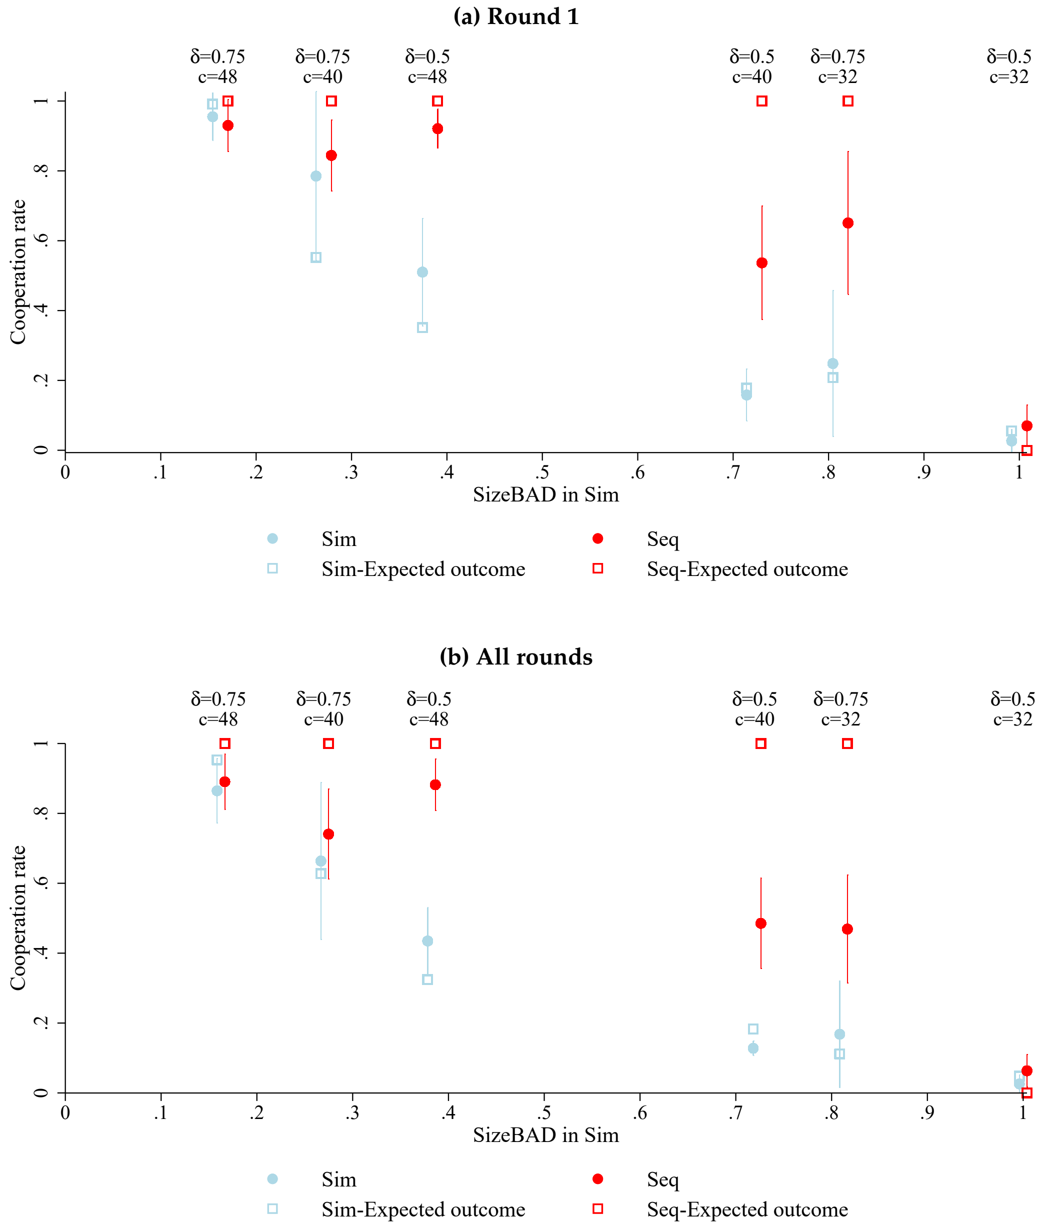

Notes: The graphs show cooperation rates and 95% confidence intervals across the last 20 repeated games depending on the SizeBAD (including treatment labels and expected outcomes). Estimates and confidence intervals are based on predictions from probit regressions ran on treatment dummy with clustered standard errors at the matching group level. In Sim expected outcomes are calculated as the cooperation rates in Dal Bó and Fréchette (2011) using data from repeated game 31 to the highest available repeated game smaller or equal to 50.

Figure G.2: Evolution of cooperation rate by treatment–All rounds

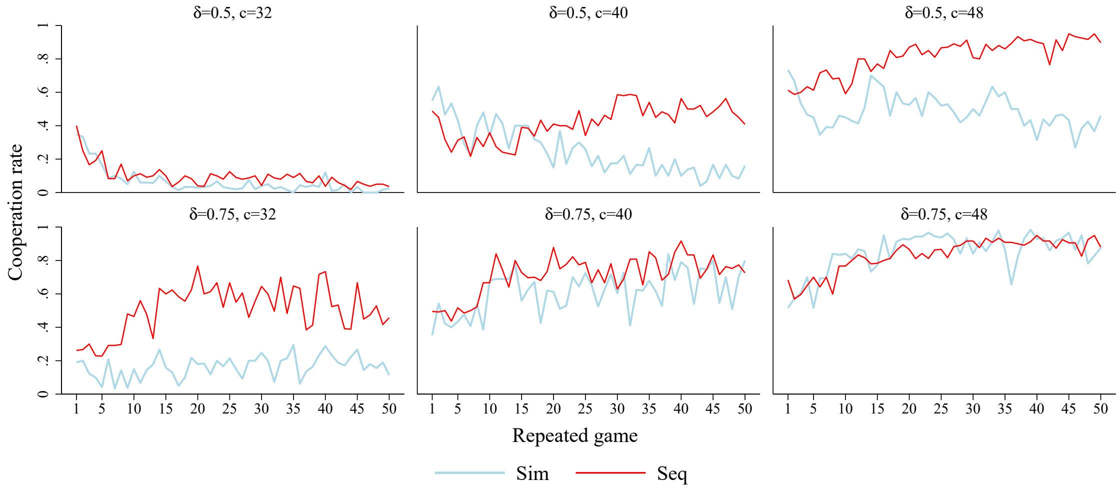

Notes: The graphs show cooperation rates across repeated games by treatment. The unit of observation is a participant's decision in a round.

Figure G.3: Evolution of cooperation rate by treatment by matching group

(a) Round 1

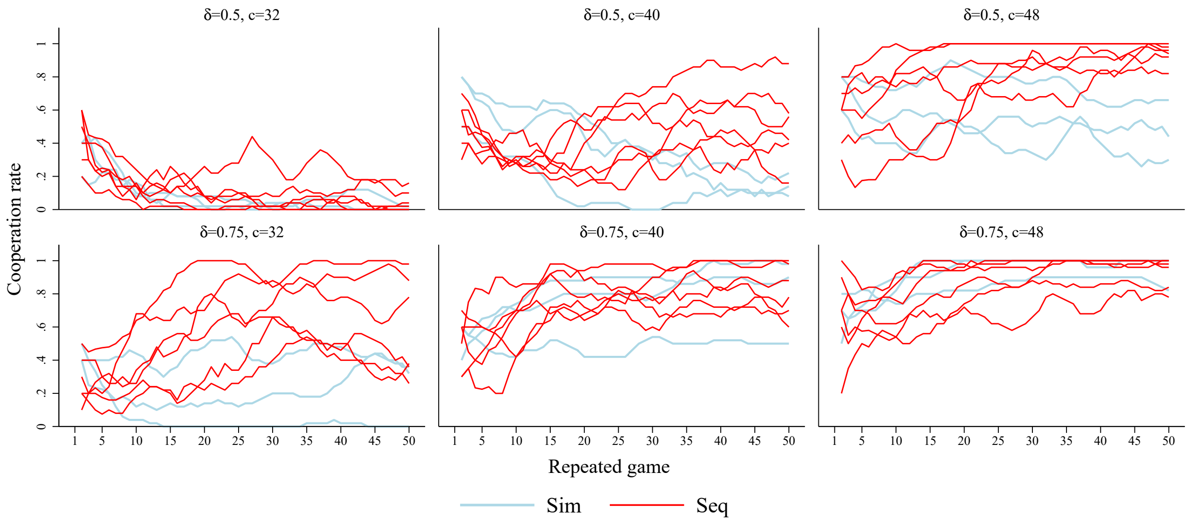

(b) All rounds

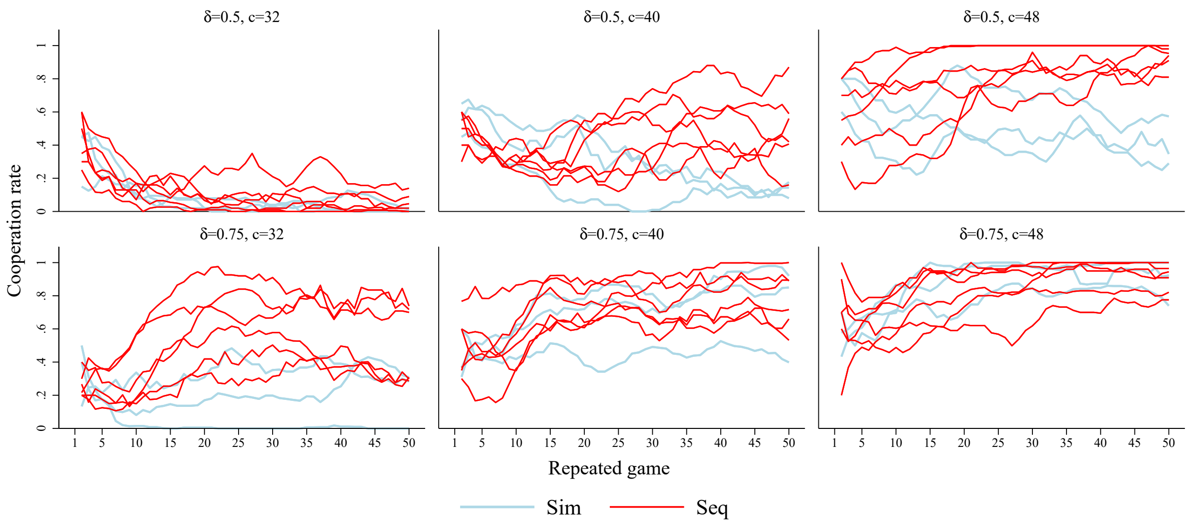

Notes: The graphs show five-repeated game moving averages of cooperation rate by repeated game and by treatment. Each line depicts a matching group. The unit of observation is a participant's decision in a round.

Figure G.4: Cooperation rate by role–All rounds

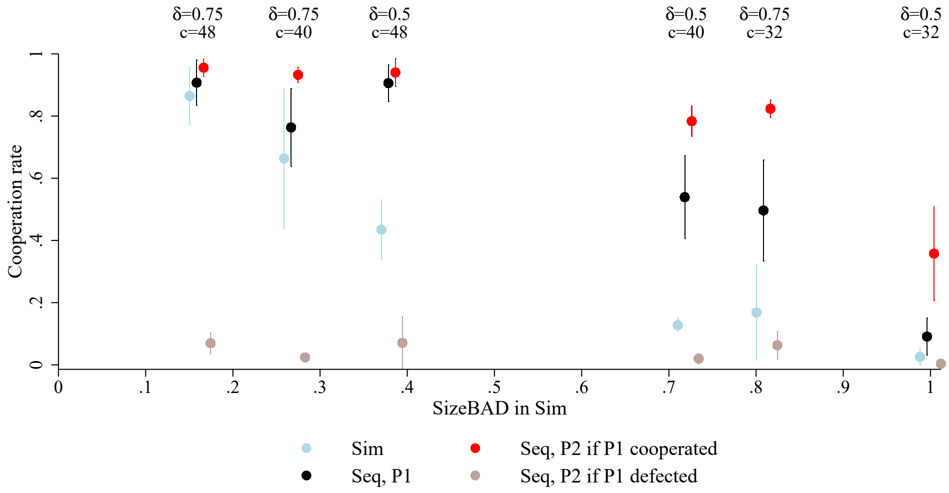

Notes: The graph shows cooperation rates of P1, cooperation rates of P2 conditional on P1 defecting or cooperating, and cooperation rates in Sim, and 95% confidence intervals, across the last 20 repeated games depending on the SizeBAD (including treatment labels). Estimates and confidence intervals are based on predictions from probit regressions ran on treatment-role dummies with clustered standard errors at the matching group level.

Figure G.5: Evolution of cooperation rate in Seq by treatment and role

(a) Round 1

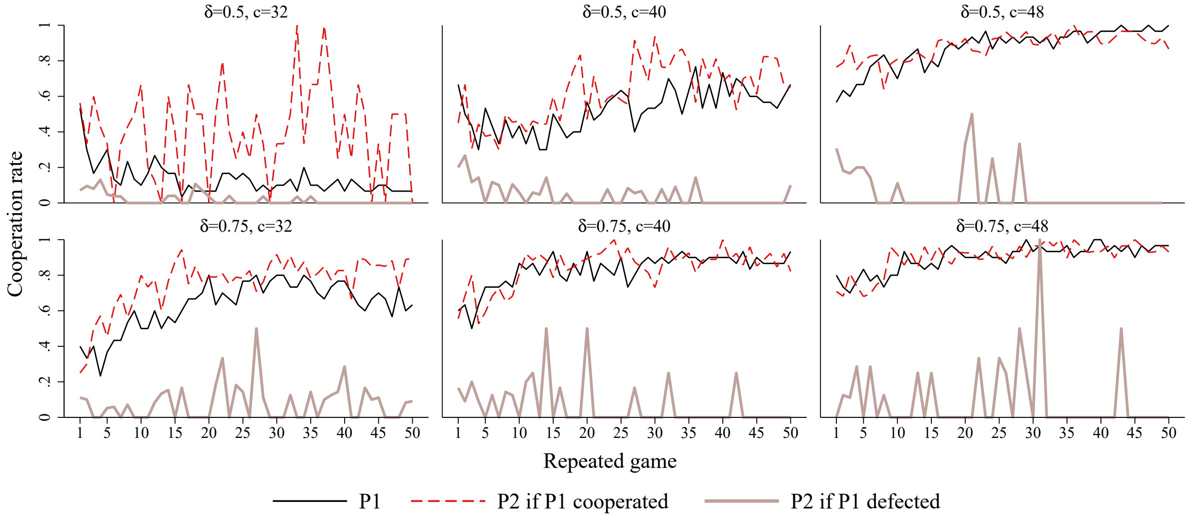

(b) All rounds

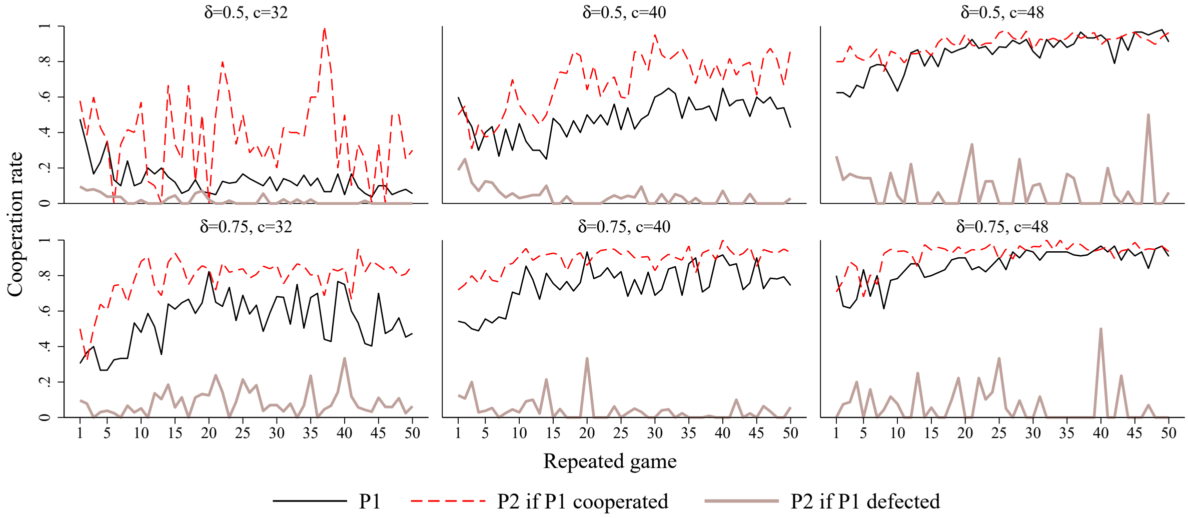

Notes: The graphs show cooperation rates across repeated games by treatment.

Figure G.6: Cooperation as first and second mover by subject

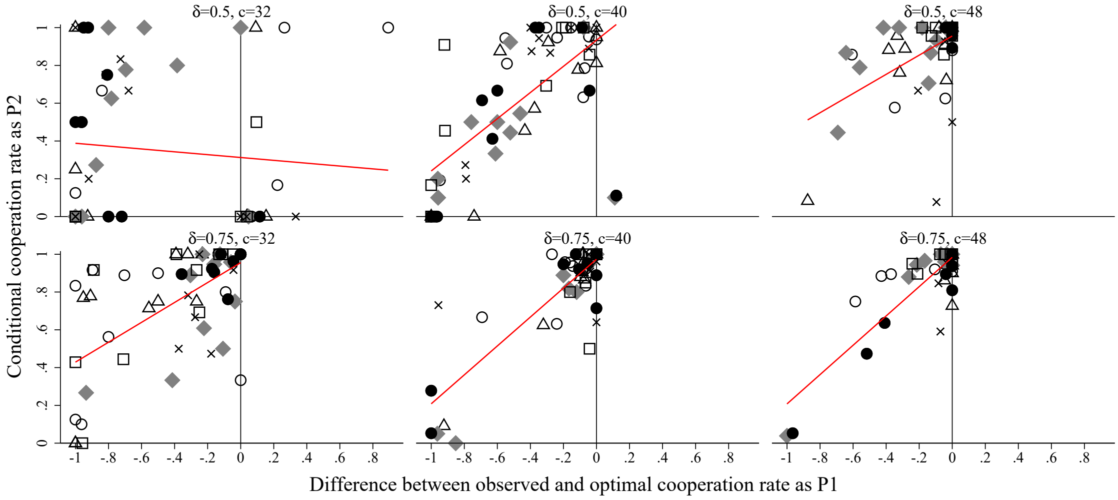

Notes: The graphs show the conditional cooperation rate in the role of second mover across first rounds as a function of the difference between the first-round cooperation rate in the role of first mover and the first-mover optimal cooperation rate. The first-mover optimal cooperation is equal to 1 if the expected payoff from the cooperative strategy is greater or equal to the expected payoff from the defection strategy given the encountered conditional cooperation rate. In $\delta = 0.5, c = 32$ 6 second movers never encountered cooperation by the first mover, and the remaining 54 second movers encountered cooperation by the first mover 1 to 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 between 12.5 and 22 across the 5 treatments.