WORKING PAPERS

Cooperation in Indefinitely Repeated Games of Strategic Complements and Substitutes

Ayse Gül Mermer
Wieland Müller
Sigrid Suetens

May 2016

Working Paper No: 1603

DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

Cooperation in Indefinitely Repeated Games of Strategic Complements and Substitutes*

Ayşe Gül Mermer[†] Wieland Müller[‡] Sigrid Suetens[§]

May 30, 2016

Abstract

We report on experiments conducted to study the effect of strategic substitutability and strategic complementarity on the extent of cooperative behavior in indefinitely repeated two-player games. On average, choices in our experiment do not differ between the strategic complements and substitutes treatments. However, the aggregate data mask two countervailing effects. First, the percentage of joint-payoff maximizing choices is significantly higher under strategic substitutes than under strategic complements. We argue that this difference is driven by the fact that it is less risky to cooperate under substitutes than under complements. Second, choices of subjects in pairs that do not succeed in cooperating at the joint-payoff maximum tend to be lower (i.e. are less cooperative) under strategic substitutes than under strategic complements. We relate the latter result to non-equilibrium forces stemming from a combination of heterogeneity of subjects and differences in the slope of the response function between substitutes and complements.

JEL Classification numbers: L13, C72, C92.
Keywords: strategic substitutes, strategic complements, collusion, repeated games, experimental economics.

*We are grateful to Ralf Becker, Jan Boone, Dirk Engelmann, Charles Noussair, Jan Potters, Christoph Schottmüller, Randolph Sloof, and participants of seminars at Tilburg University, University of Manchester, of the European Association for Research in Industrial Organization Meetings 2015 in Munich, Royal Economic Society Meetings 2015 in Manchester, XXIX Jornadas de Economia Industrial 2014 in Barcelona, and ESA World Meetings 2013 in Zurich for their valuable comments and suggestions. We acknowledge the financial support from CentER, Tilburg University.

†Department of Economics, University of Manchester, School of Social Sciences, Arthur Lewis Building 3.005, Oxford Road, Manchester M13 9PL, United Kingdom. E-mail: *ayse.mermer@macnhester.ac.uk.*

‡Department of Economics & VCEE, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, and Department of Economics, Tilburg University, CentER & TILEC, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands. E-mail: *wieland.mueller@univie.ac.at.*

§Department of Economics, Tilburg University, CentER & TILEC, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands. E-mail: *s.suetens@uvt.nl.*

1. Introduction

The study of cooperation and its determinants has attracted a great deal of attention in the literature. It is well-known, for instance, that in indefinitely repeated games, cooperation can be supported in equilibrium if the discount factor is sufficiently high (Friedman, 1971). Not much is known, however, about how, empirically, the strategic environment—whether actions are strategic complements or substitutes—influences cooperative behavior in indefinitely repeated games. In this paper we report on experiments conducted to study the effect of strategic substitutability and strategic complementarity on the extent of cooperative behavior in indefinitely repeated games.

Strategic complementarity refers to the property that best-response functions slope upward, whereas under strategic substitutability best-response functions slope downward.[1] The complements/substitutes distinction is relevant in several important applications. For example, depending on whether firms in oligopolistic markets with homogeneous goods are engaged in price or quantity competition, actions are strategic complements or substitutes, and *vice versa* in markets with complementary goods. Also, depending on whether skills of members in teams are complementary or substitutable, efforts of team members are strategic complements or substitutes. Moreover, depending on whether the production of a public good is characterized by increasing or decreasing returns, contributions are strategic complements or substitutes. Finally, when spillovers are high or low, R&D competition is characterized by strategic complementarity or substitutability, respectively.

While real-world interactions in some of these applications might be best approximated by games with a finite and definite ending, in others decision makers might be uncertain about the number and the time horizon of interactions, so that these interactions might best be approximated and modeled by indefinitely repeated games. For the case of a finite and known number of repetitions, Potters and Suetens (2009) show that there is significantly more cooperation when actions are strategic complements rather than strategic substitutes. For the case of indefinitely repeated games such evidence is missing.[2] Moreover, as we

[1] A game is characterized by strategic complements (substitutes) if $\forall i, j$ and $i \neq j$: $\partial^2 \pi / \partial x_i \partial x_j > 0$ ($<$ 0), implying that the best-response functions are upward- (downward-) sloping (see Topkis, 1978; Bulow, Genakoplos and Klemperer, 1985; Fudenberg and Tirole, 1984).

[2] An important exception is Embrey, Mengel and Peeters (2014). This paper focuses on the effect of commitment on cooperation games of strategic substitutes and complements. We discuss how our paper and results relates to this paper in the discussion section.

explain further below, theory and earlier experimental results do not lead to an unambiguous prediction regarding the effect of the strategic environment (substitutes versus complements) on the extent of cooperative behavior in indefinitely repeated games. This is, hence, an additional motivation for our study, as experiments seem to be particularly well-suited to help understand which effects prevail if predictions are unclear.

In our experiment, pairs of subjects play games with an indeterminate final period that feature either strategic complementarity or strategic substitutability. The games are borrowed from Potters and Suetens (2009), henceforth referred to as PS. Across the two treatments, several variables are kept constant, namely, the actions and payoffs in the Nash equilibrium of the stage game and in the symmetric joint-payoff maximum, the optimal defection payoff and the absolute value of the slope of the stage-game best-response function. Subjects know that after each period the game proceeds to a next period with a fixed continuation probability. In order to allow for learning across games, subjects play at least 20 repeated games. After a repeated game ends, players are randomly re-matched to play another repeated game with the same continuation probability. The treatments are designed so that cooperation at the joint-payoff maximum can be sustained as a subgame-perfect Nash equilibrium. In particular, the treatments have the same critical discount factor above which such "full" cooperation is supported by, for example, a grim trigger strategy.

On average, choices in our experiment do not differ significantly between the strategic complements and substitutes treatments. This is in clear contrast to PS, who find in a finitely repeated game that an environment with strategic complements is more conducive to cooperation than one with strategic substitutes. However, our aggregate result masks two countervailing results that are in line with two distinct literatures. The first of these results is that the percentage of choices at the joint-payoff maximum is significantly higher under strategic substitutes than under strategic complements. This result fits well with the notion that strategic risk related to cooperation at the joint-payoff maximum is lower under substitutes than under complements. Recent theoretical and experimental studies on indefinitely repeated prisoner's dilemma games show that strategic risk is an important determinant of behavior. In particular, Blonski, Ockenfels and Spagnolo (2011) formalize the intuition that cooperation gets riskier, and thus less likely, the more it hurts to cooperate if the partner defects (that is, the lower the "sucker" payoff). In particular, they propose a threshold for the discount factor in an indefinitely repeated game above which cooperation at the joint-payoff maximum is supported in equilibrium, which is higher than the standard threshold based on e.g. grim-trigger strategies (see Blonski and Spagnolo, 2015). Blonski, Ockenfels and Spagnolo (2011) and Dal Bó and Fréchette (2011) provide experimental evidence showing that this threshold is necessary for cooperation in a prisoner's dilemma to increase to very high levels. This adjusted threshold is lower in games of strategic substitutes than in games of strategic complements, thus making it easier to cooperate in the former than in the latter case.

The second result in our experiment hidden at the aggregate level is that choices of subjects in pairs that do not succeed in cooperating at the joint-payoff maximum tend to be lower, i.e. are less cooperative, under strategic substitutes than under strategic complements. This finding squares well with theoretical and experimental findings on the differential effects of strategic substitutes and complements on cooperation in the presence of heterogeneous player types. To illustrate, if a cooperator is matched with a best-responder, the aggregate outcome in a pair will be less cooperative under strategic substitutes than under complements (Haltiwanger and Waldman, 1991, 1993; Camerer and Fehr, 2006). The reason is that a best-response to a cooperative choice is less cooperative under strategic substitutes than under complements in the sense that it deviates less from the static Nash equilibrium in the former than in the latter case.[3] Experimental evidence for this intuition in the context of a "long" finitely repeated dilemma game is provided by PS.[4]

The remainder of this paper is organized as follows. In Section 2 we introduce the experimental design and procedures. In Section 3 we develop the conjectures concerning predicted behavior in our experiment, focusing on the comparative static predictions between the treatments with complements and substitutes. In Section 4 we present the experimental results. In Section 5 we summarize and discuss our findings in the light of the existing literature.

[3] In contrast to theory, the estimated response functions in our two treatments have the same positive sign, but it still holds that the slope in the case of complements is larger than the one in the case of substitutes.

[4] See Haltiwanger and Waldman (1985) and Fehr and Tyran (2008) for applications where aggregate outcomes depend on the strategic environment if individuals are heterogeneous in the rationality of their expectations.

2. Experimental Design and Procedures

2.1. Experimental Design

Our experiment has two treatments: one where choices are strategic complements (Comp) and another where choices are strategic substitutes (Subs). In each treatment, subjects play an indefinite repetition of the same stage game. The stage game has a unique and Pareto dominated Nash equilibrium and a symmetric socially efficient (joint-payoff maximizing) outcome (JPM). The payoffs in each treatment are determined according to the following payoff functions (borrowed from PS):

$$\pi_i^{\mathsf{Comp}}(x_i, x_j) = -28 + 5.474x_i + 0.01x_j - 0.278x_i^2 + 0.0055x_j^2 + 0.165x_ix_j, \tag{1}$$

$$\pi_i^{\mathsf{Subs}}(x_i, x_j) = -28 + 2.969x_i + 2.515x_j - 0.082x_i^2 + 0.023x_j^2 - 0.0485x_ix_j. \tag{2}$$

The coefficients in the payoff functions are chosen in order to ensure a fair comparison between the two treatments. First, in both treatments the Nash-equilibrium choices are the same and the JPM-choices are the same. Second, the payoffs corresponding to the Nash equilibrium and the JPM are the same across the two treatments. Third, the payoff achieved by best responding to JPM play of the matched player, referred to as the defection payoff, is the same in the two treatments.[5] Lastly, the absolute value of the slopes of the best-response curves are the same in the two treatments to guarantee that the same speed of convergence is generated by best-response dynamics. Table 1 summarizes the main theoretical benchmarks of our design.

Table 1: Theoretical Benchmarks

	Comp	Subs
$Choice_{Nash}$	14.0	14.0
$Choice_{JPM}$	25.5	25.5
Π_{Nash}	27.71	27.71
Π_{JPM}	41.97	41.97
Π_{Defect}	60.14	60.14
$Slope\ of\ reaction\ function$	0.30	-0.30

Notes: This table shows the theoretical benchmarks regarding choices and payoffs in the experiment.

[5]The combination of the second and third condition mentioned above has as the consequence that payoffs on the best-response function are the same in the two treatments.

In order to allow for learning, in our experiment subjects played a series of one of the two games described above. We refer to each repeated game, that is, each sequence of periods determined by the continuation probability of 0.9, as a match. Once a match ended and depending on the time left, another one started. In each session, subjects participated in as many matches as possible such that at least 20 matches were played. If at least 20 matches had already been played, a session ended after one and a half hours of play. Subjects played with the same partner throughout a match. Once a match ended, subjects were randomly re-matched with another subject.

By using the payoff functions given in (1) and (2), we keep several actions and payoffs constant across treatments. We felt the same should be done with respect to the sequence of matches and their respective lengths. At the same time, because of possible order effects, we did not only want to have one sequence of matches to be played in each of the two treatments. We therefore decided to have five different draws of the lengths of matches prior to the start of the experiments, each of which was administered in one session for each of the two treatments Comp and Subs.[6] The length of each match in a draw was determined randomly with the continuation probability of 0.9. Figure VIII in the Web Appendix F shows the distribution of realized match lengths across all five draws.[7]

Since there is always the possibility of continuing to a next round, the randomization generates a game that is strategically equivalent to an indefinitely repeated game. In particular, the continuation probability δ is equivalent to the discount factor in an indefinitely repeated game assuming that within the time slot of an experiment, there is no discounting (Roth and Murnighan, 1978).

2.2. Experimental Procedures

The experiment consists of 10 sessions (five for each of the two treatments Comp and Subs) that were conducted at CentERlab at Tilburg University during September-October 2011.[8] A total number of 160 students participated in the experiment. Participants were recruited through an email list of students who are interested in participating in the experiments. In

[6] For instance, under draw number 1, the randomly determined lengths of the matches played was: 11, 5, 9, 5, 18, 33, 7, 7, 5, 12, 4, 16, 11, 1, 5, 4, 23, 9, 14, 6, 6, 10, 2, 7, 1.

[7] In an indefinitely repeated game with continuation probability $\delta = 0.9$, the expected number of periods in each match is 10.

[8] We used the experimental software toolkit *Z-Tree* to program and conduct the experiment (see Fischbacher, 2007).

each session, 16 subjects interacted anonymously in a sequence of matches, that is, indefinite repetitions of the same stage game. In each session subjects participated in between 20 and 25 matches. Each session lasted not more than two hours (including the time to read the instructions and payment of the subjects).

All participants were given the same instructions (see Web Appendix A). At the beginning of each match, subjects were randomly paired with each other. During a match, subjects played with the same partner. The matching rule was explained clearly before the experiment started. The identity of the partners was not revealed to subjects. It was explained to the subjects that their final earnings depended on their own choices and the choices of the matched participants. The subjects were asked to choose a number between 0.0 and 28.0 (up to one digit after the comma) in each round of a match. Subjects were provided an earnings calculator on the computer screen enabling them to calculate their earnings in points for any combination of hypothetical choices, and a payoff table for combinations of hypothetical choices that are multiples of two (see Figure I and Figure II in the Web Appendix A).

After choices were submitted in each round, subjects were informed about whether or not the match would continue to a next round. In the case the game continued to a next round, subjects received the message "*The match continues to the next round.*" on the computer screen. In the case the match ended, subjects received the message "*The match is over.*" on the computer screen. Once a match ended, another match would begin, depending on the time available. Moreover, after each round of a match subjects were provided with information of the previous round on the screen, namely their own choice and earnings and the matched partner's choice and earnings.

After subjects finished reading the instructions, we explained to them that the experiment itself would proceed for about 1.5 hours.

The payoffs in the experiment were expressed in points. At the end of the experiment, the sum of a subject's earnings in points in all rounds of all matches were converted into Euro at the exchange rate of 480 points = 1 Euro, and privately paid to subjects. The average earnings in the experiment was 16.45 Euro.

3. Predictions

A first prediction builds on the standard theory of infinitely repeated games. Based on simple grim-trigger strategies, this theory predicts that cooperation can be supported as a subgame perfect Nash equilibrium (SPNE) if the following condition holds:

$$\frac{\Pi_{JPM}}{1-\delta} \geq \Pi_{Defect} + \frac{\delta \Pi_{Nash}}{1-\delta}. \tag{3}$$

The left-hand side of (3) is the discounted sum of payoffs from cooperation, while the right-hand side is the discounted sum of payoffs from a one-time deviation followed by Nash equilibrium play forever after. By design, the JPM payoff, the defection payoff, and the static Nash equilibrium payoff are the same in both treatments. Rearranging condition (3) and using the numbers given in Table 1, we get

$$\delta \geq \underline{\delta} := \frac{\Pi_{Defect} - \Pi_{JPM}}{\Pi_{Defect} - \Pi_{Nash}} = \frac{60.14 - 41.94}{60.14 - 27.71} = 0.56 \tag{4}$$

for both treatments. We thus conclude that the critical discount factor above which cooperation at the joint-payoff maximum (full cooperation) is supported by a grim-trigger strategy is the same in both treatments.[9, 10]

A second prediction takes into account differences in the relative riskiness of cooperation in the two treatments. Inspecting the payoffs in Subs and Comp, one notices that if one player plays fully cooperatively, while the other player in the market defects optimally, the cooperating player's ("sucker") payoff is lower with complements than with substitutes. In addition, the payoff players get if they both optimally defect, is lower in Subs than in Comp. Intuitively, these two forces make it less attractive to choose actions that maximize joint payoffs in Comp than in Subs, because doing so is relatively more risky in the former than in the latter treatment.

[9] The range of actions that Pareto-dominate the static Nash equilibrium, and thus also the range of actions that can be sustained in equilibrium in an indefinitely repeated game, is larger under substitutes than under complements. This can be seen in Figure VII in the Web Appendix F that shows the iso-payoff contours in both cases. Given the findings of Gazzale (2009), we did not expect that this difference would lead to differences in the extent to which subjects succeed in fully cooperating. It may lead to larger variability in actions under substitutes than under complements, though.

[10] However, note the following. Any feasible and admissible average payoff vector "above" the NE of the stage game can be supported as a SPNE provided that δ is sufficiently high. The area of these payoff vectors for Comp is 386.648, while for Subs it is 403.246.

Table 2: A general and reduced PD games for the two treatments

(a) A general PD game

	C	D
C	c, c	a, b
D	b, a	d, d

(b) The reduced PD game for Comp

	C	D
C	41.94, 41.94	5.89, 60.14
D	60.14, 5.89	34.90, 34.90

(c) The reduced PD game for Subs

	C	D
C	41.94, 41.94	10.71, 60.14
D	60.14, 10.71	18.17, 18.17

Notes: This table illustrates the payoff matrices for a general PD game and the reduced PD games for Comp and Subs treatments.

Recently, this intuitive idea received formal support in Blonski, Ockenfels and Spagnolo (2011). These authors suggest an axiomatic approach to equilibrium selection in indefinitely repeated prisoner's dilemma (PD) games. They show that a set of five axioms leads to a discount factor δ^* that is strictly larger than the standard discount factor $\underline{\delta}$ derived above and that, more importantly for our purposes, also reflects the influence of the sucker payoff on the incidence of fully cooperative play.[11] In particular, given a PD stage game of the form shown in Panel (a) in Table 2 with $b > c > d > a$ and $2c > b + a$, Blonski, Ockenfels and Spagnolo (2011, Proposition 2) show that their five axioms imply the threshold $\delta^* = (b-c+d-a)/(b-a)$ above which a cooperation equilibrium is predicted to be played in the indefinitely repeated PD. Note that this threshold features the sucker payoff a, while the threshold $\underline{\delta}$ derived above does not (there $\underline{\delta} = (b-c)/(b-d)$). Note also that $\partial\delta^*/\partial a = -\left(c-d\right)/\left(a-b\right)^2 < 0$, so that a *lower* sucker payoff *increases* the threshold above which cooperation should be observed. Put differently, the lower the sucker payoff, the smaller the range of discount factors for which cooperation can be supported in equilibrium.

Blonski, Ockenfels and Spagnolo (2011) develop their approach in the context of a standard 2×2 PD game. Our stage game, however, has many more than just two actions. Still, we believe that the intuitive idea that a lower "sucker" payoff and higher "mutual optimal defection" payoff should *ceteris paribus* lead to less full cooperation is also relevant in the context of our stage games. A prediction that translates Blonski et al.'s approach to our games can be generated if one is willing to make the simplifying assumption that the action space of our stage games consists of just two strategies, say $C = Choice_{JPM}$ and $D = Choice_{Defect}$. Using

[11] The five axioms in Blonski, Ockenfels and Spagnolo (2011) are called (1) positive linear payoff transformation invariance; (2) δ-monotonicity, (3) boundary conditions (which is the crucial axiom that highlights the influence of the sucker payoff on the incidence of cooperation); (4) incentive independence; and (5) equal weight.

the payoff functions given in (1) and (2), these two choices lead to the two games shown in Panels (b) and (c) in Table 2.[12] It follows that $\delta^*_{\mathsf{Comp}} = 0.870$ and $\delta^*_{\mathsf{Subs}} = 0.518$, so that full cooperation can be sustained for a larger range of discount factors in treatment Subs than in treatment Comp.[13]

An alternative concept leading to the same comparative static prediction as the approach suggested by Blonski, Ockenfels and Spagnolo (2011) is based on the idea of the *basin of attraction* of a cooperative strategy in comparison to a defecting strategy (see Dal Bó and Fréchette (2011)). We provide details of this idea in Web Appendix B.

A third prediction is based on the literature that studies the interaction between the strategic environment (complements versus substitutes) and heterogeneity of players (Haltiwanger and Waldman, 1991; Camerer and Fehr, 2006), as well as its application to repeated-game experiments (see PS). The intuition goes as follows. In games of strategic complements a change in the matched player's choice gives a payoff-maximizing player an incentive to move in the *same* direction, while in games of strategic substitutes the incentive is to move in the *opposite* direction. Given that several experiments have shown that some individuals are (conditionally) cooperative in the sense that they try to induce cooperation and follow it when established by others, even when there is no future interaction, (see Fehr and Fischbacher, 2002; Clark and Sefton, 2001; Reuben and Suetens, 2009), it is plausible to assume that players are heterogeneous in their cooperativeness and defection strategies. Consider, for example, a cooperative player who is matched with a defector in the above-described games of complements and substitutes. If the cooperative player makes a cooperative choice (higher than the static Nash equilibrium), and the matched defector is an optimal defector in the sense that he best-responds to this move, then, in sum, choices will be higher (more cooperative) in Comp than in Subs. This is because in Comp, the best-response to a cooperative move is to (partly) follow the move and make a higher choice as well, whereas in Subs the best-response is to make a less cooperative choice. This mechanism may facilitate cooperation in Comp and may hamper it in Subs. In addition, a similar mechanism occurs when a cooperative player is matched with a spiteful defector who aims at maximizing the payoff difference between himself and the cooperator. In order to employ the same level of

[12]The choice $C = Choice_{JPM}$ is equal to 25.5 in both treatments, while $D = Choice_{Defect} = 17.42$ in Comp and $D = Choice_{Defect} = 10.64$ in Subs.

[13]In the case the two choices are $C = Choice_{JPM} = 25.5$ and $D = Choice_{Nash} = 14$ in both treatments, we get $\delta^*_{\mathsf{Comp}} = 0.7834$ and $\delta^*_{\mathsf{Subs}} = 0.664$, and so, again, $\delta^*_{\mathsf{Comp}} > \delta^*_{\mathsf{Subs}}$.

punishment (in payoff terms), a spiteful defector must choose much lower choices in the Subs treatment than in the Comp treatment. So here as well, choices will, on average, be higher, i.e. more cooperative in Comp than in Subs. PS provide evidence for this intuition in the context of a finitely repeated game.

Summarizing, based on theory and earlier experimental results no unambiguous prediction can be made regarding the higher prevalence of cooperation in our two treatments. Hence, we formulate the following research question:

Research Question In the context of an indefinitely repeated game, which environment is more conducive to cooperation: strategic substitutes or strategic complements?

4. Experimental Results

In this section we describe our main results. We analyze data from matches 1-20 for which we have observations in all sessions.

Averaged over all subjects, rounds and matches, the mean choice is 19.09 in treatment Subs and 18.70 in treatment Comp. In the last 10 matches the mean choice in the Subs treatment is 20.12 and that in the Comp treatment is 19.87.[14] The average choice is thus roughly the same in the two treatments.

Figure 1: Evolution of Average Individual Choices

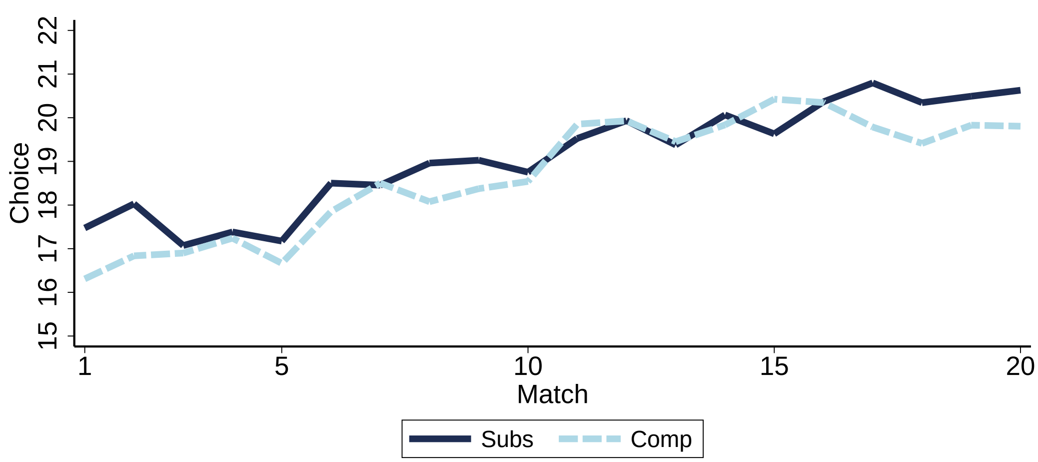

Notes: This figure shows the evolution of average individual choices across matches.

[14]The summary statistics for average choice is presented in Table V in the Web Appendix G.

Figure 1 illustrates the evolution of average choices over time under strategic complements and strategic substitutes. In both treatments, the average choice is increasing over the matches. However, there is no clear difference between the two treatments.[15] To formally quantify the difference between the two treatments, and to test whether it is statistically significant, we estimate the effect of strategic complementarity on the individual choice. We do so by regressing the choice of an individual on a treatment dummy, and clustering standard errors at the session level. Results are reported in column (1) of Table 3. The regression results confirm that the difference between the two treatments is small in size, and not statistically significant (the treatment dummy coefficient is -0.365 and statistically insignificant at $p = 0.679$).[16,17]

However, some properties of the data might be hidden when looking at aggregates. To analyze the data in more detail, in a next step we present the distribution of choices for strategic substitutes and complements. Figure 2 shows that choices in the Subs treatment are spread over the whole interval, while choices in the Comp treatment are somewhat more concentrated. Moreover, the modal choice in both treatments is a choice at or very close to the JPM level of 25.5. This is particularly accentuated in treatment Subs. To illustrate, in Subs almost 30% of the choices are at or very close to the JPM level of 25.5, whereas in Comp we only observe about 15% of such choices.

To further explore potential differences between Subs and Comp, we distinguish "fully-cooperative" and "non-fully cooperative" choices. We define a choice to be fully-cooperative if it lies within the interval $[25, 26]$, where 25.5 is the JPM choice in both treatments. We refer to a choice as non-fully cooperative if it lies outside the range $[25, 26]$.[18]

The left-hand panel of Figure 3 illustrates for both treatments the share of fully cooperative choices across matches. From this graph it becomes clear that the share of fully

[15]There is also no significant difference in payoffs between the two treatments as reported in Table IV in the Web Appendix G.

[16]The estimated treatment effect of strategic complementarity on individual choice becomes -0.386 at $p = 0.505$ when we control for the match and the interaction between treatment and match. No significant differences are obtained in payoffs either. This can be seen in Table IV in the Web Appendix G.

[17]Mann-Whitney-U tests based on independent observations yield similar results, both when the average choice is based on all matches or the last 10 matches ($p = 0.750$ in both cases).

[18]The choice of such a range is to some extent arbitrary, and one may argue that choices above 26 are also fully cooperative. For example, 28, which is the maximum choice possible, can serve as a focal point for subjects to coordinate on (almost) full cooperation. Enlarging the fully-cooperative interval to $[25, 28]$, does not affect any of our qualitative results in what follows. Choices above 26 correspond to 0.68 % of all choices in the experiment.

Table 3: Regression results on choice

VARIABLES	(1) Choice$_{it}$	(2) Choice$_{it}$	(3) Choice$_{it}$	(4) Choice$_{it}$
Comp	−0.365	−0.386	−2.421***	−2.190***
	(0.853)	(0.557)	(0.179)	(0.200)
Choice$_{jt-1}$			0.743***	0.734***
			(0.012)	(0.010)
Comp*Choice$_{jt-1}$			0.129***	0.126***
			(0.013)	(0.013)
Match		0.208***		0.0602***
		(0.041)		(0.015)
Comp*Match		0.001		−0.015
		(0.060)		(0.018)
Constant	19.210***	16.951***	5.005***	4.530***
	(0.703)	(0.487)	(0.118)	(0.096)
Observations	33,024	33,024	29,824	29,824
R-squared	0.001	0.043	0.604	0.607

Notes: This table reports results from linear regressions with standard errors (in parentheses) clustered at the session level. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level. The dependent variable is a subject's choice in all specifications.

Figure 2: Distribution of Choices

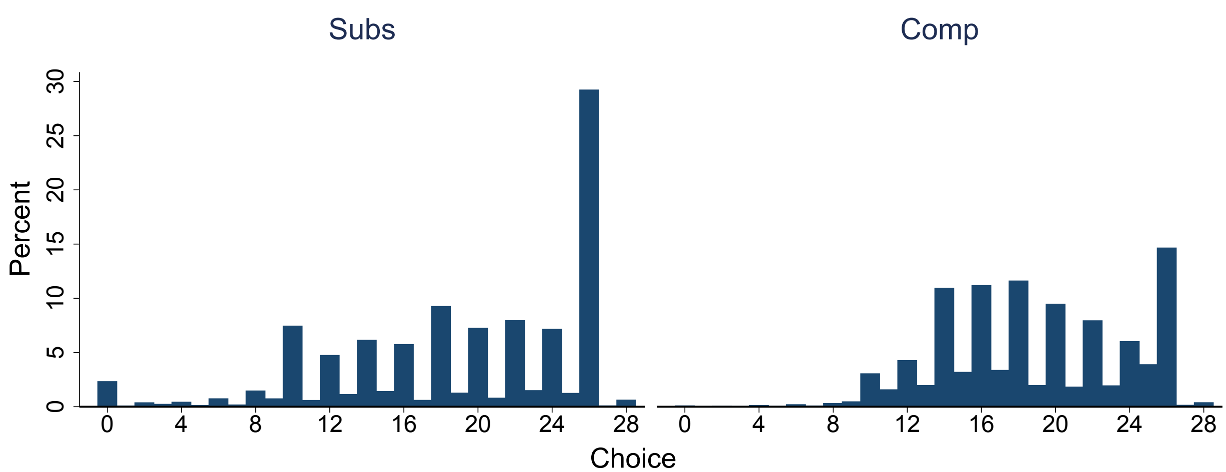

Notes: This figure shows the distribution of individual choices in the experiment.

cooperative choices is higher in Subs than in Comp. In addition, the share of fully cooperative choices increases in both treatments, but more so in Subs than in Comp. In the last 10 matches, the percentage of fully cooperative choices is around 40% in Subs, while it is around

Figure 3: Cooperative vs Non-Cooperative Behavior

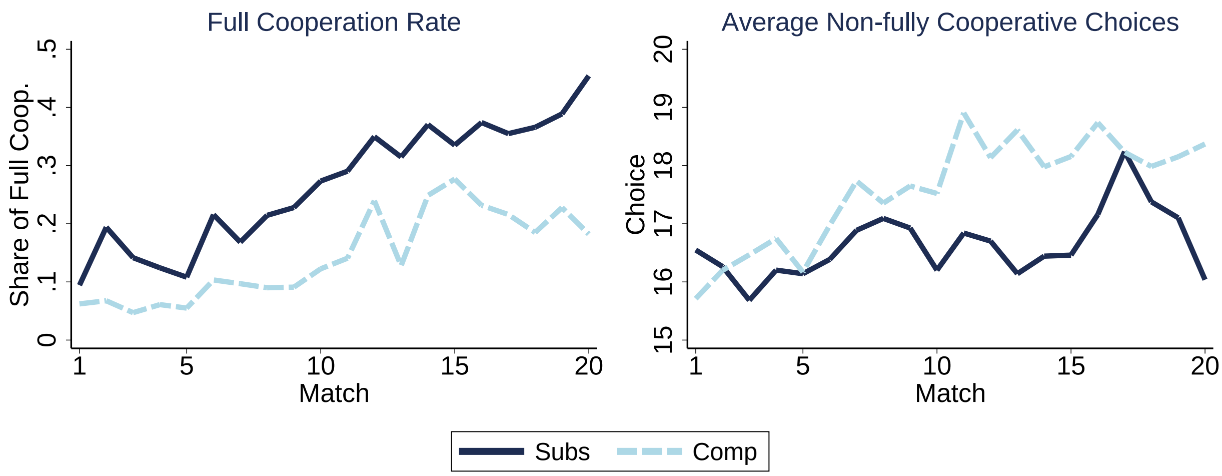

Notes: This figure shows the evolution of cooperative and non-cooperative behavior. The left-hand panel depicts the evolution of full cooperation rate across matches and the right-hand panel depicts the evolution of average non-fully cooperative choices across matches.

25% in Comp.[19]

The right-hand panel of Figure 3 depicts the evolution of averages of non-fully cooperative choices (those that fall outside the interval $[25, 26]$) across matches. Here we observe that the average choice of subjects is, overall, higher in Comp than in Subs. So it seems the effect of strategic complementarity on behavior switches—behavior is more cooperative because choices are higher—when we focus on non-fully-cooperative choices. To illustrate, averaged over subjects, rounds and matches, the mean non-fully cooperative choice is 16.65 in the Subs treatment and it is 17.59 in the Comp treatment. In the second half of the experiment, the average non-fully cooperative choice is 16.85 in Subs and 18.33 in Comp.[20]

In sum, although we do not observe a difference between the two treatments at the aggregate level, analyzing fully cooperative and non-fully cooperative behavior separately suggests that, overall, behavior is driven by two countervailing forces. On the one hand, subjects make choices at the fully cooperative level more frequently under Subs than under Comp. On the other hand, the average choice of subjects who do not make fully-cooperative choices is higher under Comp than under Subs. To understand which forces drive these two results, we analyze fully cooperative behavior in section 4.1 and non-fully cooperative

[19]For an in-depth analysis of the statistical significance of these observations see Section 4.1.

[20]For a more detailed analysis of the statistical significance of these results see Section 4.2.

behavior in 4.2 in more detail. Web Appendix C focuses on results at the pair level. Moreover, in view of the time trends visible in Figures 1 and 3, we explore learning across matches in Web Appendix D.

4.1. Full Cooperation Rates

In this section we take a closer look at full cooperation rates, that is, choices in the interval $[25, 26]$ at the level of subjects. In doing so, we examine the first and all rounds of a match separately since the cooperation rate might evolve within a match, depending on the number of rounds in that match (see Dal Bó and Fréchette, 2011). In addition, in the first rounds of each match subjects are playing with a new partner so that they do not have experience with their partners' behavior (or cannot recall it due to random matching). In this respect, subjects' behavior in the first round of each match is mainly driven by the fundamentals of the game they are playing (and possibly their experiences in the previous matches) and not by the current partners' behavior.

Figure 4 illustrates the evolution of the full cooperation rate across matches, in the left-hand panel for the first rounds and in the right-hand panel for all rounds of a match. The left-hand panel shows that in the first rounds of a match subjects make fully cooperative choices more frequently under Subs than under Comp. In addition, the first-round full cooperation rate follows an increasing trend in Subs, while in Comp it is more steady across matches. The full cooperation rate in the first match is almost the same for the two treatments, while towards the end of the experiment there is a considerable difference in full cooperation rates between the two treatments. Moreover, the first-round full cooperation rate reaches the level of about 25% in the Subs treatment by the end of the experiment, while it remains at around 5% in the Comp treatment.

In order to test whether these differences are statistically significant, we ran two specifications of a probit regression in which the dependent variable is a dummy referring to a subject making a fully cooperative choice or not. In the first specification shown in Table 4 we include as an independent variable a treatment dummy. In the second specification, next to the treatment dummy, we control for the match, and the interaction between treatment and match. As shown in Table 4, in both specifications the treatment dummy has a negative sign—the full cooperation rate in Subs is thus lower than the one in Comp—and is statistically

Figure 4: Full Cooperation Rate

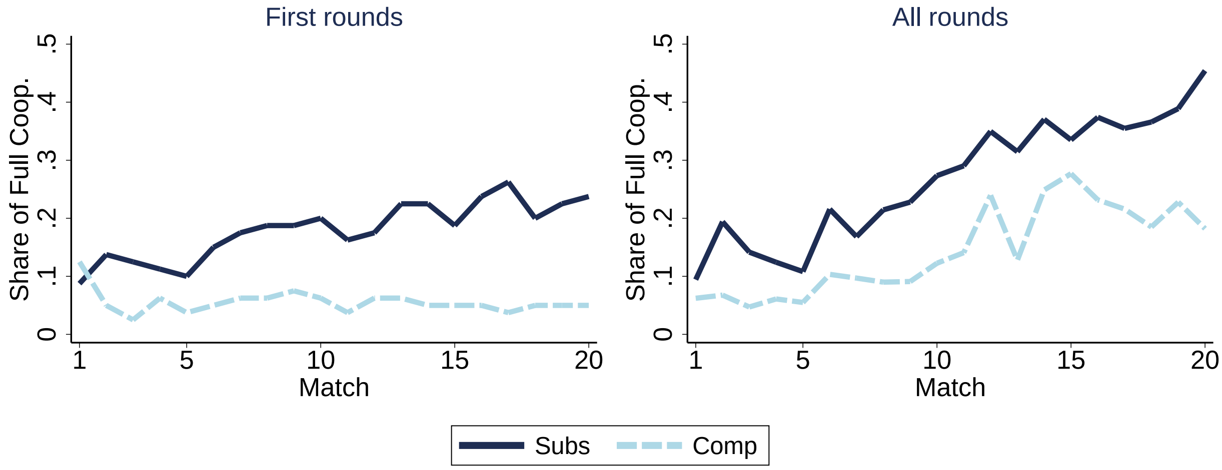

Notes: This figure shows the evolution of full cooperation rate across matches, on the left-hand panel for the first rounds only and on the right-hand panel for all rounds.

significant. The estimated marginal effect is -0.127 and -0.50, respectively.[21] In addition, column (2) shows that the first-round full cooperation rate significantly increases over the matches in Subs (marginal effect is 0.005, $p \leq 0.001$), but not so in Comp (marginal effect is -0.002, $p \leq 0.001$).

Next, we focus on the right-hand panel of Figure 4 and the remainder of Table 4. As illustrated in the figure, there is again a clear difference between the two treatments in the full cooperation rate. In contrast to the first rounds, the full cooperation rate now increases over matches in Comp as well. The full cooperation rate raises up to about 25% in Comp and up to about 45% in Subs.

The results of probit regressions, which we report in columns (3) and (4) of Table 4, indicate that the treatment effects are again statistically significant. Moreover, as shown in column (4), the full cooperation rate increases significantly over the matches in both treatments (marginal effect is 0.016, $p = 0.001$).

Summarizing, we find significantly more initiation of full cooperation at the beginning of a new match as well as more fully cooperative choices in general in Subs than in Comp. This result is in line with the discussion of the differences in the "riskiness of cooperation" in our two treatments in Section 3.

[21]The p-values in Mann-Whitney-U tests based on sessions averages are 0.016 if all matches are taken into account and 0.075 if only matches 11-20 are taken into account.

Table 4: Regression results on full cooperation

	First rounds		All rounds	
	(1)	(2)	(3)	(4)
VARIABLES	FullCoop_{it}	FullCoop_{it}	FullCoop_{it}	FullCoop_{it}
Comp	−0.127***	−0.050***	−0.115***	−0.178***
	(0.031)	(0.022)	(0.042)	(0.043)
Round				0.004***
				(0.002)
Comp*Round				0.004**
				(0.002)
Match		0.005***		0.016***
		(0.002)		(0.002)
Comp*Match		−0.007***		0.001
		(0.003)		(0.004)
Observations	3,200	3,200	33,024	33,024

Notes: This table reports marginal effects from probit regressions with delta-method standard errors (in parentheses) clustered at the session level. The dependent variable is a dummy which is equal to 1 if the choice is fully cooperative and 0 otherwise. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level. Specifications (1) and (2) are based on observations from the first rounds of matches only and specifications (3) and (4) are based on all observations.

4.2. Non-Fully Cooperative Behavior

In this section we analyze the effect of strategic complementarity on non-fully-cooperative behavior. In doing so, we focus on those data points that are not in the fully cooperative range of $[25, 26]$. Figure 5 depicts the evolution of the average non-fully-cooperative choice over matches, in the left-hand panel for the first rounds and in the right-hand panel for all rounds of a match.

The figure in the left-hand panel shows that in the first rounds of the matches there is no clear difference in non-fully-cooperative behavior between the two treatments. The figure also shows that in both treatments the average non-fully cooperative choice in the first rounds is initially above the static Nash equilibrium choice of 14 and increases over the matches. As shown in Table 5, presenting results from linear regressions where the average non-fully-cooperative choice is regressed on a treatment dummy, the treatment effect is small and not significant. In addition, as shown in column (2) of this table, the average choice significantly

Figure 5: Average Non-Fully Cooperative Choices

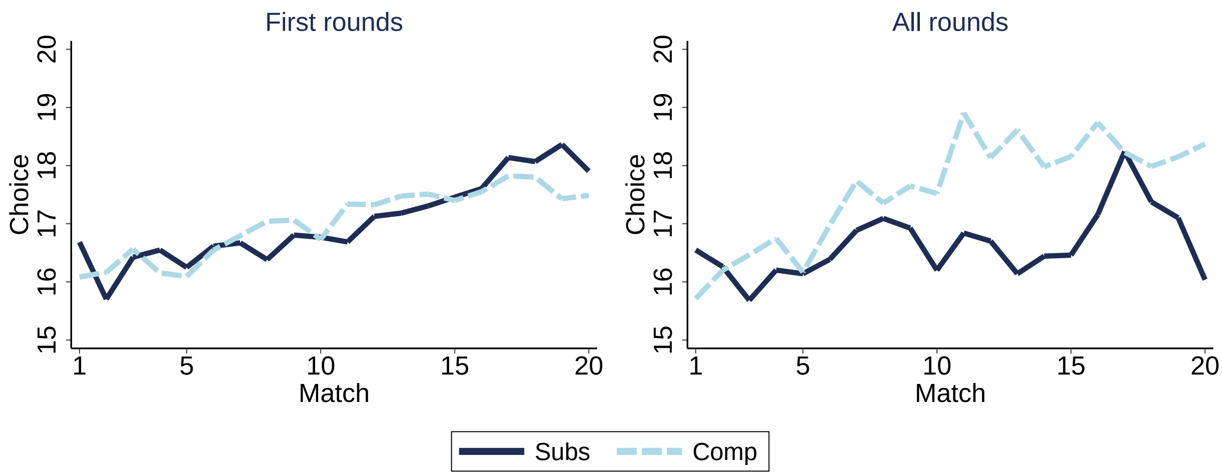

Notes: This figure shows the evolution of non-fully cooperative choices (i.e. choices outside the range $[25, 26]$) across matches, on the left-hand panel for the first rounds only and on the right-hand panel for all rounds.

increases over time.[22]

Next, we consider average non-fully-cooperative choices across all rounds. The evoloution of these choices across matches is shown in the right-hand panel of Figure 5.[23] Here, a different behavior emerges. When averages are taken across all rounds instead of just the first rounds of a match, the average non-fully-cooperative choice is higher in Comp than in Subs, although this difference is not significant ($p = 0.301$, see column (3) in Table 5).

Next we analyze the adjustments across rounds. During a match, subjects observe the past choice(s) of the matched subject and are likely to adjust their own behavior. If at least some of the subjects (noisily) best-respond it should be the case that in Comp the estimated response function has a higher slope than in Subs (see Table 1). Columns (4) and (5) of Table 5 report estimates of the observed response functions. The reported results come from linear regressions where the choice of a player is regressed on the choice of the matched player in the previous round (in the same match) as well as the interaction of the other subject's past choice and a treatment dummy. In column (5) additional controls are

[22] We also tested whether average choices of subjects who do not play fully cooperatively is the same in the two treatments by using a two-sided Mann-Whitney-U Test. The p-value of the null hypothesis that the average non-fully cooperative choice is the same in the two treatments is 0.25, for both the entire experiment and the second half of the experiment. So we fail to reject the null hypothesis.

[23] The right-hand panel of Figure 5 is the same as the right-hand panel of Figure 3.

Table 5: Regression results on non-fully cooperative choices

	First rounds		All rounds		
	(1)	(2)	(3)	(4)	(5)
VARIABLES	$Choice_{it}$	$Choice_{it}$	$Choice_{it}$	$Choice_{it}$	$Choice_{it}$
Comp	0.007	0.175	0.999	−2.830***	−2.968***
	(0.546)	(0.417)	(0.911)	(0.269)	(0.462)
$Choice_{jt-1}$				0.583***	0.582***
				(0.031)	(0.028)
Comp*$Choice_{jt-1}$				0.202***	0.192***
				(0.033)	(0.031)
Match		0.109***			0.017
		(0.018)			(0.033)
Comp*Match		−0.019			0.029
		(0.032)			(0.038)
Constant	17.016***	15.909***	16.334***	6.607***	6.453***
	(0.404)	(0.361)	(0.773)	(0.171)	(0.397)
Observations	2,823	2,823	25,061	22,238	22,238
R-squared	0.001	0.021	0.010	0.444	0.446

Notes: This table reports results from linear regressions with standard errors (in parentheses) clustered at the session level. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level. Specifications (1) and (2) are based on observations from the first rounds of matches only and specifications (3), (4) and (5) are based on all observations.

included for the match and the interaction between match and treatment. Both columns show that in both treatments subjects (partially) follow each other (i.e., an increase in the choice of the rival is followed by an increase in one's own choice in the current period), and the effect is statistically significant.[24] Importantly, the extent to which subjects follow each other is significantly greater in Comp than in Subs. To illustrate, an increase in the choice by a subject, increases the choice of the matched subject in the next round by 0.58 in Subs and by 0.78 in Comp. The effects are very similar when we control for the match and the interaction between match and treatment.

The positive effect of Comp shown in column (3) of Table 5 in combination with the result that the extent to which subjects follow each other is greater in Comp than in Subs (cf. columns (4) and (5) in Table 5), suggest that at least some subjects try to induce cooperation,

[24]Reaction functions being positively sloped in both treatments can be explained by endogenous complementarity that arises when subjects use reciprocal strategies (see also PS).

to which others (noisily) best-respond. For example, if a subject who increases its choice above the static Nash equilibrium, with the intention to move towards full cooperation, is matched with a (noisily) best-responding subject or a spiteful subject, choices in this pair will on average end up to be higher (more cooperative) in Comp than in Subs, which is exactly what we observe. This is the mechanism behind our prediction based on heterogeneity of subjects' types in Section 3.

Summarizing, when we focus on non-fully-cooperative choices, we find that behavior is in agreement with the mechanism based on heterogeneity of subjects, so that the average (non-fully-cooperative) choice tends to be higher in Comp than in Subs.

We also present the results of regressions of treatment effects and responses of subjects using all choices, so including those in the fully-cooperative range. Table 3 summarizes the results. Recall that the specification in column (1) shows the aggregate (non-significant) treatment effect on choices. The specifications in columns (3) and (4) show the estimated response of subjects to the matched subject's choice, as well as the treatment effect on this response (with and without controlling for the match). As can be seen, the estimated responses are qualitatively similar to those shown in Table 5. The size of the estimated response is larger now, because fully cooperative choices as well as subjects responding to full cooperation by fully cooperating themselves are included as well.

The heterogeneity explanation elaborated on above is supported by results of simulations assuming two different types of players: cooperative and non-cooperative players. A cooperative player is assumed to reciprocate full mutual cooperation and to induce full cooperation with a certain probability by playing fully cooperatively. This probability is positively related to the cooperativeness of the matched partner's cooperative response to his own cooperativeness in the previous period. Otherwise, this type of player either punishes by playing a best-reply or playing spitefully with a certain probability (where spiteful behavior consists of maximizing the difference in payoffs). A non-cooperative player plays either a best-reply or spitefully, each with a certain probability. These two types of players imply three types of possible matchings: those of two cooperative players, those of a cooperative and a non-cooperative player, and those of two non-cooperative players. To this basic setup (that draws on simulations presented in PS), we add the assumption of a differential propensity of initiating cooperation depending on the treatment, just as we found in our experiment. Finally, we simulate indefinitely repeated games in the same way we did in our experiment. The two versions of simulations ran replicate the key findings of our experiment (see Web Appendix E for details): On the one hand, we obtain similar average choices within a match in treatments Comp and Subs. On the other hand, we obtain full cooperation rates within a match, which are higher under Subs than under Comp, and average non-fully cooperative choices within a match, which are higher under Comp than under Subs.

5. Discussion

In our experiment subjects play indefinitely repeated dilemma games of strategic substitutes or complements. Our first result is that, on average, we find no significant difference in choices between the two environments. Thus we do not recover the result obtained in PS who find that average choices in a finitely-repeated game are higher (more cooperative) under strategic complementarity. However, an analysis based on averages masks two opposing forces that cancel out each other in the aggregate, which we refer to as our second and third result, respectively.

Our second result is is different from what PS find. Our data indicate that this is because under substitutes subjects more often take the risk to initiate full cooperation at the beginning of each repeated game. They do so more frequently the more repeated games they play. To illustrate, in the second half of the substitutes treatment the percentage of full cooperation in the first periods has increased to a level above 20%. In contrast, under complements, subjects rarely take this risk, and the percentage remains at about 5% in the second half.

Our third result is that if we focus on choices of subjects who do not succeed in fully cooperating, that is, who do not make joint-payoff maximizing choices, we find that, on average, choices tend to be *less* cooperative (lower) under strategic substitutes than under complements (although not statistically significantly so). Relatedly, we find that under complements, the slope of the estimated response function is (significantly) higher than under substitutes.

Our second result goes against PS, who find that, if anything, full cooperation is lower under substitutes. However, the result is in line with the idea that strategic risk has an effect on behavior in games. Loosely speaking, how much a player loses by cooperating in the case the other player defects has an impact on whether this player will choose to cooperate or not. In our games, it is less risky to fully cooperate or initiate full cooperation with strategic substitutes than with strategic complements. In this sense, the result is in line with theory and experiments on indefinitely repeated prisoner's dilemma games taking into account strategic risk (Blonski, Ockenfels and Spagnolo, 2011; Dal Bó and Fréchette, 2011; Blonski and Spagnolo, 2015) and on coordination games that have shown that payoff-dominant actions are chosen less frequently if they involve more strategic risk (Van Huyck, Battalio and Beil, 1990; Schmidt et al., 2003).

Our third result is in line with theory and experiments in the literature that studies the interaction between the strategic environment and heterogeneity of players (see also Haltiwanger and Waldman, 1991, 1993; Camerer and Fehr, 2006). This literature finds that if players are heterogeneous, aggregate outcomes tend to be different depending on the strategic environment. In particular, they tend to be more cooperative under strategic complements than under substitutes. This is what PS observe in finitely repeated games of strategic complements and substitutes.

Are the two opposing forces summarized above (Results 2 and 3) two "isolated" effects or are they in some way related to each other? Although there is some evidence that suggests the latter,[25] we argue that it is the former. It just seems to be the case that starting (and maintaining) collusion is easier under substitutes than under complements and that simple response dynamics (and the slopes of predicted and estimated response functions) are such that non-fully cooperative choices are on average lower under substitutes than under complements.

To the extent that our experimental results have implications outside the lab, e.g. for competition policies, the main message would be as follows. Our result regarding significant differences in fully-cooperative choices suggests that markets characterized by strategic substitutability are more prone to collusion than markets characterized by strategic complemen-

[25]The two forces could be related in the sense that episodes of collusion (which happen more often in treatment SUBS) are also followed by harsher punishments in the form of lower choices once collusion breaks down. To see whether this is the case, we analysed punishment choices of subjects when their partner deviates from a collusive episode that was sustained for at least three consecutive rounds. That is, we looked for example at the choices in period t of non-cheating subjects in pairs in which both players made JPM choices for at least three consecutive rounds (in rounds $t-4$, $t-3$ and $t-2$), and one player started to cheat in period $t-1$. We observe that subjects who make punishment choices typically choose actions around the static Nash equilibrium in treatment Comp, while they choose actions much lower than the static Nash equilibrium in treatment Subs. While this appears to be as conjectured at the beginning of this footnote, note that one needs to choose a lower action in Subs than in Comp in order to induce, *c.p.*, the same level of punishment for a cheating partner. Moreover, the number of incidences in which collusion breaks down and is followed by punishments is very low in both treatments, such that this mechanism cannot account for the two opposing forces observed in our data.

Figure 6: Collusion Index

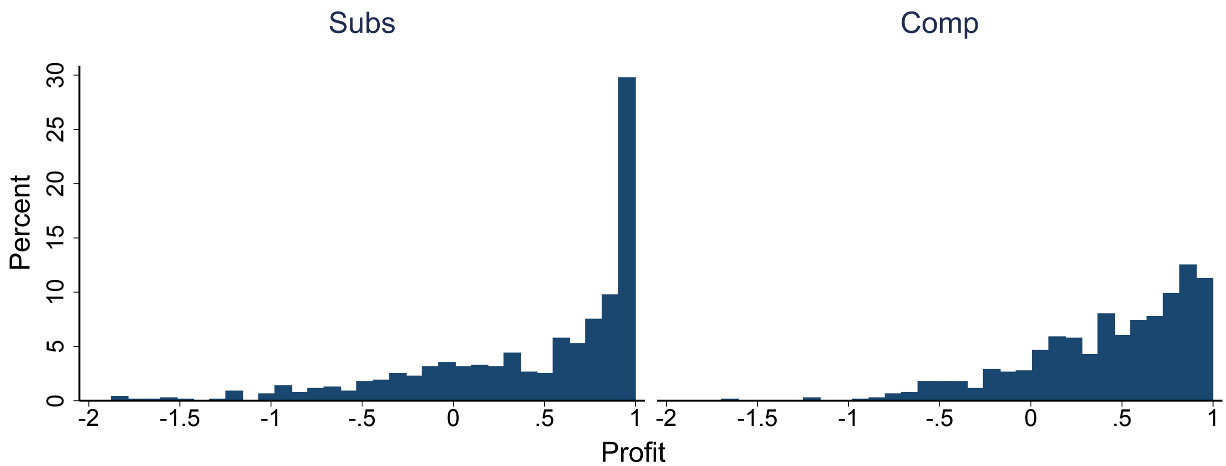

Notes: This figure shows the distributions of the average collusion index per individual match.

tarity. To illustrate this, define the collusion index $CI := (\pi_{Observed} - \pi_{Nash})/(\pi_{JPM} - \pi_{Nash})$, where $\pi_{Observed}$ is the average observed payoff of the two players in a match. Clearly, the collusion index equals 1 in the joint-profit maximum, and it is 0 if both players make Nash-equilibrium choices. Figure 6 shows histograms of the average collusion index per individual match in our two treatments. The most salient feature of this Figure is the large share of $CI = 1$ outcomes in treatment Subs. This would, if anything, justify increased monitoring of markets that are characterized by strategic substitutes.[26] Our result regarding differences in non-fully cooperative choices in pairs that do not succeed in cooperating at the joint-payoff maximum suggests that episodes resembling some kind of "price wars" could also be more likely under substitutes than under complements. These would be markets in which the collusion index is negative. However, Figure 6 shows that this is much less of a concern as the share of negative collusion indices is not too different across the two treatments (22.88 percent in treatment Subs versus 16.25 percent in treatment Comp, difference not significant).

Why do our results partly differ from those in PS? We speculate that the difference between our results and PS is driven by differences in the nature of the game. The repeated game in PS is a long finitely repeated game. It is played with the same partner for 30 rounds, and subjects know this. In contrast, in our experiment, subjects repeatedly play the repeated game with different partners and subjects do not know when each repeated game ends. The

[26] Ivaldi et al. (2003) discuss factors that are conducive to tacit collusion. The analysis above suggests that "strategic substitutes" should be added to this list of factors.

fundamentals of the interactions are thus very different. In the repeated game of PS, full cooperation, if it occurs, is typically built up gradually: subjects gradually increase their choice towards the level that maximizes joint payoffs. To illustrate, it often takes around 10 rounds to get to this level. In addition, subjects only participate in one first round, that is, at the very start of the repeated game, and they do not initiate full cooperation more frequently in the games with strategic substitutes than in those with complements in this first round. In our indefinitely repeated games, gradual build-up is difficult to obtain: subjects do not know how long the repeated game will last, and the expected length is much smaller (10 rounds versus 30 rounds). As compared to PS, full cooperation (if it occurs) hinges more on subjects taking the risk to fully cooperate in the first round of each repeated game. Therefore, we suspect that the higher strategic risk inherent in the games of strategic complements as compared to substitutes has played a fundamental role in our experiment, and not so in PS.

Next, consider our findings in relation to Embrey, Mengel and Peeters (2014). This paper studies in an experiment the effect of strategic commitment on cooperation in indefinitely repeated games of strategic complements and substitutes. Subjects choose an initial action and a strategy (a "machine") at the beginning of each repeated game. Treatments vary with respect to the level of commitment, that is, the costs at which strategies can be adjusted in each round of the repeated game. The treatments vary as well with respect to the strategic environment, with joint-payoff maximization being relatively more risky under strategic complements than under strategic substitutes.[27] Interestingly, subjects choose more often joint-payoff maximizing actions under strategic substitutes than under complements when the level of commitment is high, whereas the opposite holds when the level of commitment is low. Strategic risk thus seems to have a substantial impact on behavior when the level of commitment is high, but not so when it is low.

If we combine our findings with those of PS and Embrey, Mengel and Peeters (2014), then it seems that the effect of the strategic environment on cooperation in repeated games depends on the extent to which strategic risk has an effect on behavior of players. A testable hypothesis could be that in environments where strategic risk is an important factor (for example, relatively short games with an unknown end, or games with high levels of commitment), an environment of strategic substitutes is relatively more conducive of cooperation than an

[27]The difference in minimum thresholds above which full cooperation can be sustained between the two treatments is not as large as in our experiment, though. To illustrate, the minimum thresholds are $\delta^*_{\mathsf{Comp}} = 0.77$ and $\delta^*_{\mathsf{Subs}} = 0.58$ (compared to $\delta^*_{\mathsf{Comp}} = 0.870$ and $\delta^*_{\mathsf{Subs}} = 0.518$, in our experiment).

environment of strategic complements. Moreover, in games where strategic risk tends to be less important (for example, long repeated games with a known end, or games with low commitment), more cooperation can be expected under strategic complements than under strategic substitutes. A preliminary meta-analysis performed by one of the authors of this paper that uses data from Cournot and Bertrand experiments and other experiments with strategic complements or substitutes suggests that there indeed seems to be a positive and significant interaction between strategic complementarity and whether or not the game has a known end.

References

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

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.

Bulow, J. I., J. Genakoplos and D. Klemperer. 1985. "Multimarket Oligopoly: Strategic Substitutes and Complements." *Journal of Political Economy* 93:488–511.

Camerer, C. and E. Fehr. 2006. "When Does "Economic Man" Dominate Social Behavior?" *Science* 711:47–52.

Clark, K. and M. Sefton. 2001. "The sequential Prisoner's Dilemma: Evidence on Reciprocation." *Economic Journal* 111:51–68.

Dal Bó, P. and G. R. Fréchette. 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101:411–429.

Embrey, M., F. Mengel and R. Peeters. 2014. "Strategic Commitment and Cooperation in Experimental Games of Strategic Complements and Substitutes." Working Paper.

Fehr, E. and J. Tyran. 2008. "Limited Rationality and Strategic Interaction. The Impact of the Strategic Environment on Nominal Inertia." *Econometrica* 76:353–394.

Fehr, E. and U. Fischbacher. 2002. "Why Social Preferences Matter–The Impact of Non-Selfish Motives on Competition, Cooperation and Incentives." *Economic Journal* 112:C1–C33.

Fischbacher, U. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics* 10:171–178.

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

Fudenberg, D. and J. Tirole. 1984. "The Fat-cat Effect, the Puppy-dog Ploy, and the Lean and Hungry Look." *American Economic Review* 74:361–366.

Gazzale, R. S. 2009. Learning to Play Nash from the Best. Technical report University of Toronto. Working Paper No 2009-03.

Haltiwanger, J. and M. Waldman. 1985. "Rational Expectations and The Limits of Rationality: An Analysis of Heterogeneity." *American Economic Review* 75:326–340.

Haltiwanger, J. and M. Waldman. 1991. "Responders Versus Non-responders: A New Perspective on Heterogeneity." *Economic Journal* 101:1085–1102.

Haltiwanger, J. and M. Waldman. 1993. "The Role of Altruism in Economic Interaction." *Journal of Economic Behavior & Organization* 21:1–15.

Ivaldi, M., B. Jullien, P. Rey, P. Seabright and J. Tirole. 2003. "The Economics of Tacit Collusion." *Final Report for DG Competition, European Commission* .

Potters, J. and S. Suetens. 2009. "Cooperation in Experimental Games of Strategic Complements and Substitutes." *Review of Economic Studies* 76:1125–1147.

Reuben, E. and S. Suetens. 2009. "On The Cognitive and Motivational Components of Reciprocity." Working Paper.

Roth, A. E. and J. K. Murnighan. 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology* 17:189–198.

Schmidt, D., R. Shupp, J. Walker and E. Ostrom. 2003. "Playing Safe in Coordination Games: The Roles of Risk Dominance, Payoff Dominance, and History of Play." *Games and Economic Behavior* 42:281–299.

Topkis, Donald M. 1978. "Minimizing a Submodular Function on a Lattice." *Operations Research* 26:305–321.

Van Huyck, J. B., R. C. Battalio and R. O. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80:234–248.

Web Appendix

A. Instructions

You are participating in an experiment on decision making. You are not allowed to talk or try to communicate with other participants during the experiment. If you have a question, please raise your hand.

Description of the Experiment

In this experiment you will be asked to make a decision in several periods. You will be randomly paired with another participant for a sequence of periods. Each sequence of periods is referred to as a match.

The length of a match is randomly determined. After each round, there is a 90% probability that the match will continue for at least another round. So, for instance, if you are in round 2 of a match, the probability there will be a third round is 90 % and if you are in round 9 of a match, the probability there will be another round is also 90%.

Once a match ends, you will be randomly paired with another participant for a new match.

In each round you and the other participant you will be matched with (referred to as the "other") will be asked to choose a number between 0.0 and 28.0 (in 0.1 steps). The following table gives information about your earnings for some combinations of your and the other's choice. Every participant is given the same table.

You can calculate your and the other's earnings in more detail (for choices that are not multiples of 2 for instance) by using the EARNINGS CALCULATOR on your screen. By filling in a hypothetical value for your own choice and a hypothetical value for the other's choice you can calculate your and the other's earnings for this combination of choices.

Once you have made up your mind, you will enter your decision under DECISION ENTRY and then clicking the button ENTER. In each round you have about 1 minute to enter your decision.

Starting with round 2 of a match, you will be given information about the previous round on your screen. That is, you will be informed about your own and the other participant's choice and your own earnings in points in the previous round.

The identity of the other participants you will be matched with will be unknown to you.

At the end of the experiment you will be paid your earnings in cash and in private. Your total earnings in points are the sum of your earnings in points over all periods of all matches of the experiment. Your earnings in points will be converted into EUR according to the following rate: 300 points = 1 EUR.

Summary

The experiment will consist of a sequence of matches. Each match will consist of a sequence of periods. The number of periods of each match is determined randomly by the computer. After each round, with probability 90% the match continues to another round. You will interact with the same participant for an entire match. After a match is finished, you will be randomly matched with another participant. In each round of a match, you and the other participant you are matched with will choose a number between 0.0 and 28.0 simultaneously.

Payoff tables

Figure I: Payoff table handed out to subjects in the Comp treatment.

							The Other's Choice →									
		0.0	2.0	4.0	6.0	8.0	10.0	12.0	14.0	16.0	18.0	20.0	22.0	24.0	26.0	28.0
	0.0	-28.00	-27.96	-27.87	-27.74	27.57	-27.35	-27.09	-26.78	-26.43	-26.04	-25.60	-25.12	-24.59	-24.02	-23.41
	2.0	-18.16	-17.46	-16.72	-15.93	-15.09	-14.21	-13.29	-12.33	-11.32	-10.26	-9.16	-8.02	-6.84	-5.61	-4.33
	4.0	-10.55	-9.19	-7.78	-6.33	-4.84	-3.30	-1.72	-0.09	1.58	3.29	5.05	6.85	8.70	10.59	12.52
	6.0	-5.16	-3.14	-1.08	1.03	3.19	5.39	7.63	9.91	12.24	14.62	17.04	19.50	22.00	24.55	27.15
	8,0	-2.00	0.68	3.41	6.18	8.99	11.85	14.75	17.70	20.69	23.72	26.80	29.92	33.09	36.30	39.55
	10.0	-1.06	2.28	5.67	9.10	12.57	16.09	19.65	23.26	26.91	30.60	34.34	38.12	41.95	45.82	49.73
Your	12.0	-2.34	1.66	5.70	9.79	13.93	18.11	22.33	26.59	30.90	35.26	39.66	44.10	48.58	53.11	57.69
Choice	14.0	-5.85	-1.19	3.52	8.27	13.06	17.90	22.78	27.71	32.68	37.69	42.75	47.85	53.00	58.19	63.42
↓	16.0	-11.58	-6.26	-0.90	4.51	9.97	15.47	21.01	26.59	32.22	37.90	43.62	49.38	55.18	61.03	66.93
	18.0	-19.54	-13.56	-7.53	-1.46	4.65	10.81	17.01	23.26	29.55	35.88	42.26	48.68	55.15	61.66	68.21
	20.0	-29.72	-23.08	-16.39	-9.66	-2.89	3.93	10.79	17.70	24.65	31.64	38.68	45.76	52.89	60.06	67.27
	22.0	-42.12	-34.82	-27.48	-20.09	-12.65	-5.17	2.35	9.91	17.52	25.18	32.88	40.62	48.40	56.23	64.11
	24.0	-56.75	-48.79	-40.78	-32.73	-24.64	-16.50	-8.32	-0.09	8.18	16.49	24.85	33.25	41.70	50.19	58.72
	26.0	-73.60	-64.98	-56.32	-47.61	-38.85	-30.05	-21.21	-12.33	-3.40	5.58	14.60	23.66	32.76	41.91	51.11
	28.0	-92.68	-83.40	-74.07	-64.70	-55.29	-45.83	-36.33	-26.78	-17.90	-7.56	2.12	11.84	21.61	31.42	41.27

Figure II: Payoff table handed out to subjects in the Subs treatment.

		The Other's Choice →														
		0.0	2.0	4.0	6.0	8.0	10.0	12.0	14.0	16.0	18.0	20.0	22.0	24.0	26.0	28.0
	0.0	-28.00	-22.88	-17.57	-12.09	-6.42	-0.57	5.47	11.68	18.08	24.66	31.42	38.37	45.49	52.80	60.29
	2.0	-22.39	-17.46	-12.35	-7.06	-1.58	4.07	9.91	15.93	22.14	28.52	35.09	41.84	48.77	55.89	63.19
	4.0	-17.43	-12.70	-7.78	-2.69	2.59	8.06	13.70	19.53	25.54	31.73	38.11	44.66	51.40	58.32	65.43
	6.0	-13.13	-8.59	-3.87	1.03	6.12	11.39	16.84	22.47	28.29	34.29	40.47	46.83	53.37	60.10	67.01
	8,0	-9.48	-5.14	-0.61	4.10	8.99	14.07	19.32	24.76	30.38	36.19	42.17	48.34	54.69	61.23	67.94
	10.0	-6.49	-2.34	2.00	6.51	11.21	16.09	21.15	26.40	31.83	37.43	43.23	49.20	55.36	61.70	68.22
Your	12.0	-4.15	-0.19	3.95	8.27	12.77	17.46	22.33	27.38	32.61	38.03	43.63	49.41	55.37	61.51	67.84
Choice	14.0	-2.46	1.30	5.24	9.37	13.68	18.17	22.85	27.71	32.75	37.97	43.37	48.96	54.72	60.67	66.81
↓	16.0	-1.43	2.14	5.89	9.82	13.94	18.24	22.72	27.38	32.22	37.25	42.46	47.85	53.43	59.18	65.12
	18.0	-1.06	2.32	5.88	9.62	13.54	17.64	21.93	26.40	31.05	35.88	40.90	46.10	51.48	57.04	62.78
	20.0	-1.33	1.85	5.21	8.76	12.49	16.40	20.49	24.76	29.22	33.86	38.68	43.68	48.57	54.24	59.79
	22.0	-2.26	0.72	3.89	7.25	10.78	14.49	18.39	22.47	26.74	31.18	35.81	40.62	45.61	50.78	56.14
	24.0	-3.85	-1.05	1.92	5.08	8.42	11.94	15.64	19.53	23.60	27.85	32.28	36.90	41.70	46.68	51.84
	26.0	-6.09	-3.49	-0.71	2.26	5.40	8.73	12.24	15.93	19.81	23.86	28.10	32.52	37.13	41.91	46.88
	28.0	-8.98	-6.57	-3.99	-1.22	1.73	4.87	8.18	11.68	15.36	19.22	23.27	27.50	31.91	36.50	41.27

B. Prediction based on the basin of attraction

Based on Blonski, Ockenfels and Spagnolo (2011), in Section 3 we derived the prediction that cooperation can be sustained as an equilibrium for a larger set of discount factors in treatment Subs than in treatment Comp. Here we use the notion of a *basin of attraction* to derive the same comparative statics prediction. To understand the idea of the basin of attraction, assume (again, a strong assumption) that players either play "tit for tat" (a cooperative strategy) or "always defect" (a defective strategy) and nothing else in the repeated PD game and that this is common knowledge (see Dal Bó and Fréchette, 2011). Then a player needs to determine which of these two strategies generates the higher expected payoff given the belief that with probability p the other player plays "tit for tat" and with probability $1 - p$ plays "always defect". The basin of attraction of the cooperative strategy is the set of beliefs p for which playing this strategy gives a higher expected payoff than the defective strategy. In the context of the general game shown in Panel (a) in Table 2 in the main text, the expected payoff for the cooperative strategy is equal to

$$\begin{aligned} C(a,c,d,\delta) &= p(c+\delta c+\delta^2 c+...)+(1-p)(a+\delta d+\delta^2 d+...) \\ &= 1/(1-\delta)\,(a-a\delta+d\delta-ap+cp+ap\delta-dp\delta)\,, \end{aligned} \tag{5}$$

while the expected payoff for the defecting strategy is equal to

$$\begin{aligned} D(a,b,d,\delta) &= p(b+\delta d+\delta^2 d+...)+(1-p)(d+\delta d+\delta^2 d+...) \\ &= 1/(1-\delta)\,(d+bp-dp-bp\delta+dp\delta)\,. \end{aligned} \tag{6}$$

Equating the two expressions in (5) and (6) gives the threshold p^* above which playing the cooperating strategy is the payoff maximizing choice. That is, the lower p^* the larger the basin of attraction of the cooperative strategy and the more likely it is that subjects will choose to fully cooperate. For the games shown in Panel (b) and (c) in Table 2, we find $p^*_{\mathsf{Comp}} = 0.391$ and $p^*_{\mathsf{Subs}} = 0.038$, so that, again, full cooperation is predicted to emerge for a larger range of beliefs in Subs than in Comp.[28]

C. Cooperative versus Non-Cooperative Pairs

In this section we look at the experimental data from a different angle by focusing on the evolution of choices and cooperation at the pair level *within* matches. To do so, we divide the pairs into those in which the two players succeed in maximizing joint payoff and those in which the two players do not succeed in doing so (along the lines of PS). We classify a pair to be collusive if both subjects choose a number in the interval $[25, 26]$ in at least 60% of the rounds in their individual match. This threshold may look rather low, but if we do not choose the threshold sufficiently low, given that many pairs only play few rounds, they would easily be classified as non-JPM pairs.[29] For example, in order to classify pairs that only play 3 rounds in total in the indefinitely repeated game as JPM pairs if they maximize joint payoff in 2 out of 3 rounds, we need to put the threshold below 66.66%. In any case,

[28]If the 2×2 PD games are generated using the actions mentioned in footnote 13, we find $p^*_{\mathsf{Comp}} = 0.784$ and $p^*_{\mathsf{Subs}} = 0.664$, and so again $p^*_{\mathsf{Comp}} > p^*_{\mathsf{Subs}}$.

[29]Figure VIII in the Web Appendix F illustrates the distribution of realized match lengths in the experiment.

Figure III: JPM vs Non-JPM Pairs

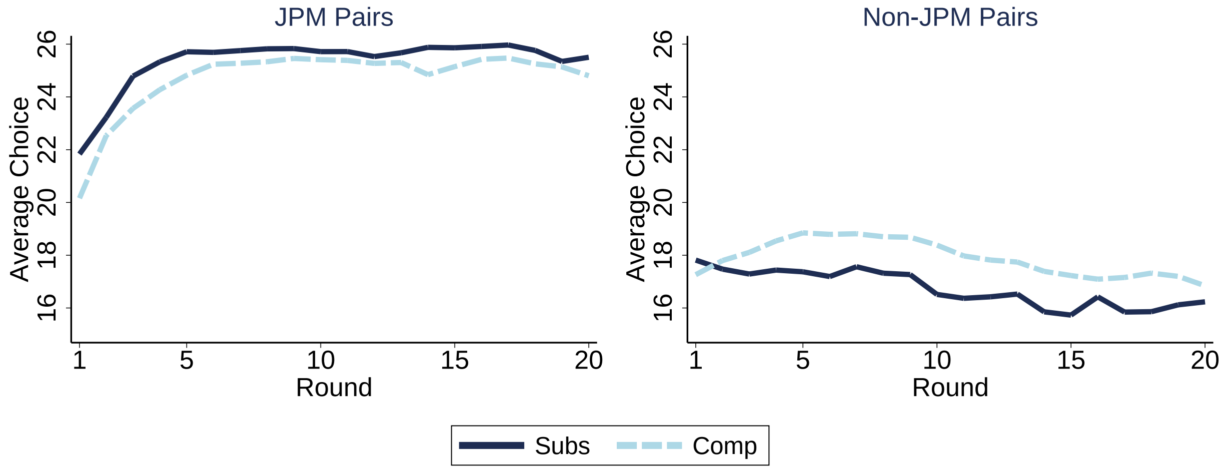

Notes: This figure shows the evolution of choices across matches for JPM and non-JPM pairs respectively on the right- and left-hand sides. A pair is referred to as JPM if both subjects makes a choice in the interval $[25, 26]$ in at least 60% of the rounds in their individual match.

any of the qualitative conclusions that are made in this section, are robust to changes in this threshold.

Figure III illustrates the evolution of average choices over time under strategic complements and strategic substitutes for JPM and non-JPM pairs respectively on the left- and right-hand panels. This graph suggests that different choice patterns emerge between JPM and non-JPM pairs. The left-hand panel of Figure III shows that in Subs pairs who succeed in full cooperation in at least 60% of a match, play higher choices than those in Comp. In the first rounds, the average choice of JPM pairs in Comp is 20, while it is 22 in Subs. As subjects gain experience over time the difference between the treatments disappears. That is, in both treatments once subjects reach the fully cooperative level they remain there. After round 5, the average choice in both treatments is around 25. (Table II summarizes the average choice for JPM and non-JPM pairs in the first and all rounds of the first match, all matches and the last 10 matches.)

The right-hand panel of Figure III illustrates the evolution of average choices of non-JPM pairs over time. Here we observe that the average choice is higher in Comp than in Subs, $p = 0.303$, see column (1) in Table I). In both treatments the average choice follows a decreasing trend over time. The estimated effect of round on the average choice is -0.061 with $p = 0.004$, see column (2) in Table I). We argue that the difference in the average choice

Table I: Regression results on average choice of non-JPM pairs

VARIABLES	(1) Choice$_{ij}$	(2) Choice$_{ij}$	(3) Choice$_{ij}$
Comp	0.972	0.797	0.183
	(0.890)	(0.853)	(0.665)
Round		−0.061***	−0.051**
		(0.015)	(0.018)
Comp*Round		0.020	0.027
		(0.028)	(0.027)
Match			0.098
			(0.071)
Comp*Match			0.051
			(0.095)
Constant	17.04***	17.58***	16.50***
	(0.754)	(0639)	(0.398)
Observations	26,500	26,500	26,500
R-squared	0.008	0.013	0.031

Notes: This table reports results from linear regression with standard errors (in parentheses) clustered at the session level. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level. A pair is referred to as JPM if both subjects makes a choice in the interval $[25, 26]$ in at least 60% of the rounds in their individual match.

Table II: Summary statistics at the pair level for JPM and non-JPM pairs

	Average choices in JPM pairs					
	First match		All matches		Last 10 matches	
	Comp	Subs	Comp	Subs	Comp	Subs
First round	-	16	19.63	21.96	19.33	22.30
All rounds	-	23.90	24.64	25.10	24.60	25.18

	Average choices in non-JPM pairs					
	First match		All matches		Last 10 matches	
	Comp	Subs	Comp	Subs	Comp	Subs
First round	17.55	18.26	17.26	17.83	17.55	18.26
All rounds	16.31	16.67	18.16	17.53	18.99	17.92

Notes: This table summarizes average choices for JPM and non-JPM pairs. A pair is referred to as a JPM pair if subjects in this pair sustain full cooperation in at least 60 % of the rounds in that match. We present the averages for the first round and all rounds of the first match, all matches and the last 10 matches separately.

of non-JPM pairs between the two treatments is due to the interaction between pairs. Since these subjects do not succeed in reaching full cooperation in 40% of the time in a match, their behavior depends on the choice of the matched subject, captured by the estimated coefficient of round. Again, choosing different thresholds of mutual cooperation (say, from from 65% to 80%), we observe almost the same patterns as the one discussed in this section.

D. Learning across matches

In this section we explore learning across matches. To do so, we study how the behavior of a subject in the first round of a match is affected by (a) the behavior of the partner in the previous match, (b) the length of the previous match, and (c) a subject's own behavior in the previous match (in the spirit of Dal Bó and Fréchette, 2011). We do this in two ways. We check how the variables just mentioned affect the average subject's probability to start a match fully cooperatively (i.e., by making a choice in the interval [25,26]), and how these variables affect a subject's level of choice. We present the results in Table III.

In column (1) and (2) of Table III we report, for each treatment separately, the results from probit regressions where the dependent variable is a dummy which equals 1 if the choice in the first round of a match is fully cooperative and 0 otherwise. In column (3) and (4) regression results are reported where the dependent variable is a subject's choice in the first round of a match. In all specifications we use the same independent variables: a dummy indicating whether or not the partner in the previous match made a fully cooperative choice in the first round of the previous match, the length of the previous match, and a dummy indicating whether a subject himself made a fully cooperative choice in the very first round of the experiment.

Consider first columns (1) and (2) of Table III. A subject who was matched with someone who played fully cooperatively in the first round of the previous match is more likely to start the current match fully cooperatively in both treatments. However, this effect is statistically significant only in treatment Comp in which cooperation is more risky. Furthermore, in both treatments there is a positive and significant relationship between the length of the previous match and subjects' likelihood of starting a match fully cooperatively. This suggests that after a longer match, during which mutual cooperation is more likely to develop, subjects more often take the risk to start the new match fully cooperatively than after a shorter

Table III: Learning across matches

	FullCoop_{it}		Choice_{it}	
	(1)	(2)	(3)	(4)
	Comp	Subs	Comp	Subs
Partner cooperated in round 1 of previous match	0.065***	0.038	1.609*	0.633
	(0.016)	(0.036)	(0.662)	(0.328)
Previous match length	0.001***	0.001**	0.013	0.044**
	(0.001)	(0.001)	(0.008)	(0.014)
Subject cooperated in round 1 of match 1	0.053	0.271***	−0.293	4.972***
	(0.059)	(0.098)	(0.465)	(0.911)
Constant			17.319***	17.673***
			(0.542)	(0.441)
Observations	1,520	1,520	1,520	1,520

Notes: Column (1) and (2) report marginal effects from probit regressions with delta-method standard errors (in parentheses) clustered at the session level. The dependent variable is a dummy which is equal to 1 if the choice is fully cooperatively and 0 otherwise. Column (3) and (4) report results from linear regression with standard errors (in parentheses) clustered at the session level. The dependent variable is a subject's choice. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level.

match. Lastly, a subject who fully cooperated in the first round of the first match in the experiment is more likely to start a new match by full cooperation than someone who did not start the experiment by full cooperation. However, this result is much more pronounced and statistically significant in treatment Subs.

Consider next columns (3) and (4) in Table III where we report the estimates from linear regressions in which the dependent variable is a subject's choice in the first round of a match. We find that subjects start a match with a higher choice if the partner in the previous match fully cooperated in the first round of the previous match, with the effect being more pronounced in treatment Comp. Also, subjects make higher or more cooperative choices after a longer previous match, where this effect is significant only in treatment Subs. Finally, a subject in treatment Subs makes significantly higher choices if he had chosen a fully cooperative choice himself in the first round of the first match. In treatment Comp, however, the effect is negative and insignificant.

To sum up, a partner's full cooperation in the previous match and a longer previous match increase the choice and the likelihood of full cooperation in the next match. Also, subjects who fully cooperate in the very first round of the experiment are significantly more likely to cooperate later in the experiment in Subs but not so in treatment Comp. Taken together, these results suggest that subjects' behavior is influenced by learning across matches as well as by the nature of the game being being played (complements or substitutes).

E. Simulation Results

In this section, we first present our experimental data in an alternative way by focusing on the evolution of choices and cooperation within matches and not across matches as we did in the main body of the paper.[30] This better enables us to compare our experimental results with those obtained by simulations. Second, we present results from two sets of simulations to show that the key features of our experimental data can be replicated.

The upper panel of Figure IV depicts the evolution of average choices within matches. We observe that average choices are roughly the same in the two treatments, similar to the observed behavior illustrated in Figure 1. The lower left-hand panel in Figure IV illustrates the evolution of full cooperation rates within matches. We observe that subjects make a fully cooperative choice more often under Subs than under Comp, similar to the left-hand panel of Figure 3. The lower right-hand panel of Figure IV depicts the evolution of average non-fully cooperative choices within matches. We see that subjects on average play higher choices under Comp than under Subs, similar to the right-hand panel of Figure 3.

Next, we report the results from two sets of simulations that build on those reported in PS, featuring cooperative and non-cooperative types of players. Basically, the first version of our simulation model is the one of PS to which we add a differential propensity of initiating cooperation depending on the treatment, and the adjustment to an indefinitely repeated game. The second version is a model that is based on all features of the first version, except that we use a redefinition of the non-cooperative type.

In both simulations models we assume that there are two types of players: a cooperative and a non-cooperative player. In the first simulation model, non-cooperative players are defined as follows. In the first period of a match, non-cooperative types always randomize over the entire action set, while in each other period of a match they condition their choice on the previous choice of the partner. They play a myopic best-reply with probability β and play spitefully with probability $1 - \beta$. Cooperative players are defined as follows. In

[30]In doing so, we only present the averages for the first 10 rounds, which is the expected length of a match in our experiment.

Figure IV: Observed Within-Match Behavior

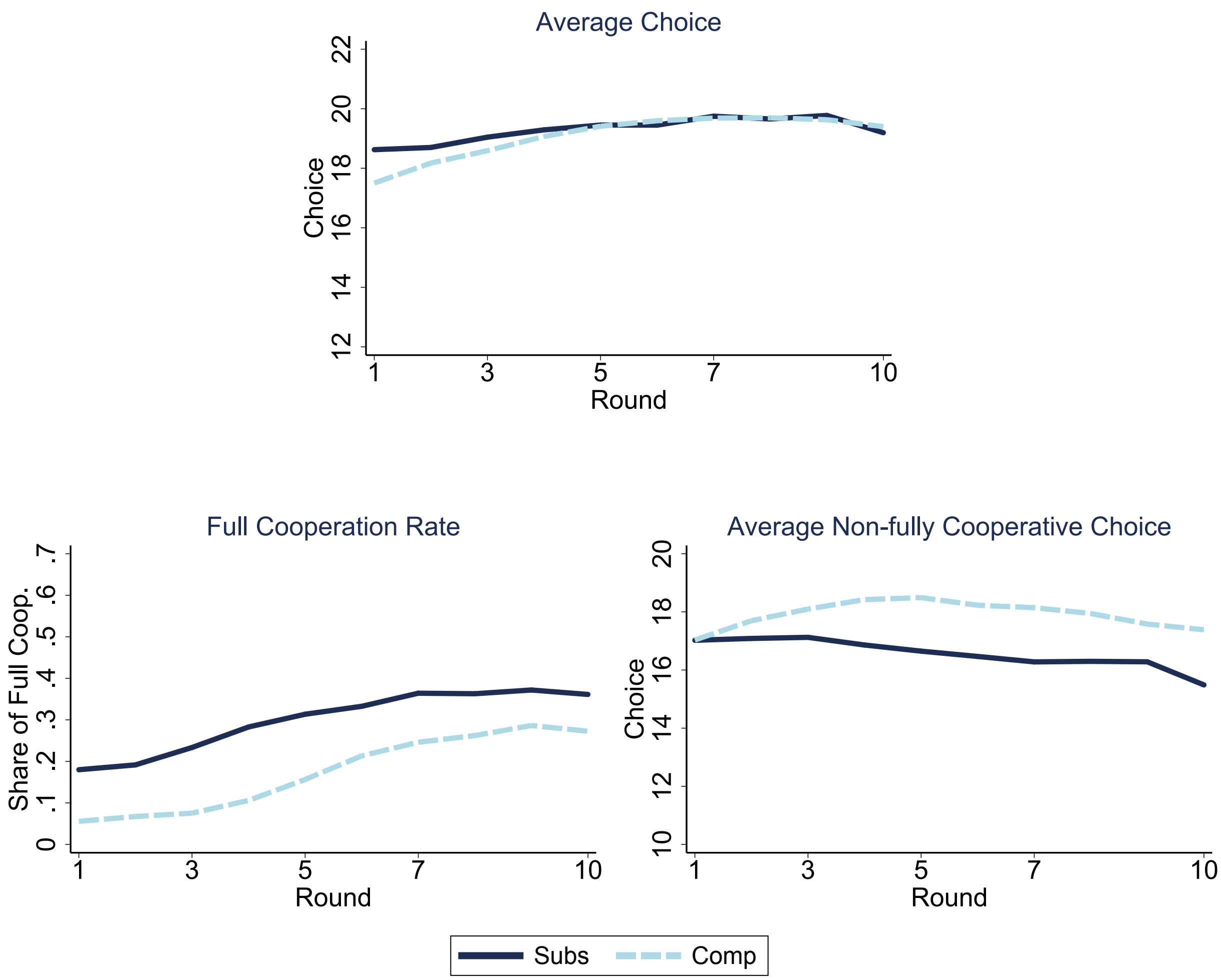

Notes: This figure shows the observed evolution of choices and cooperation within matches. In the upper panel for all choices, in the lower left-hand panel for full cooperation rates, and on the lower right-hand panel for non-fully cooperative choices (i.e. choices outside the range $[25, 26]$).

the first round, cooperative types play fully cooperatively with probability α and otherwise randomize over the entire action set. In any later period cooperative types always reciprocate full mutual cooperation and also try to induce it with a certain probability: if the pair reached full cooperation in round $t-1$, they play fully cooperatively in round t, otherwise the play fully cooperatively with probability α_t and non-cooperatively with probability $1-\alpha_t$. Here we take α_t to be history dependent to add the intuition that the probability of fully cooperative play depends on the response of the partner to one's own cooperation. Given one's own full cooperation in round $t-2$, we assume that α_t is the following positive function of the degree

Figure V: Simulation Results for Within-Match Behavior (Model 1)

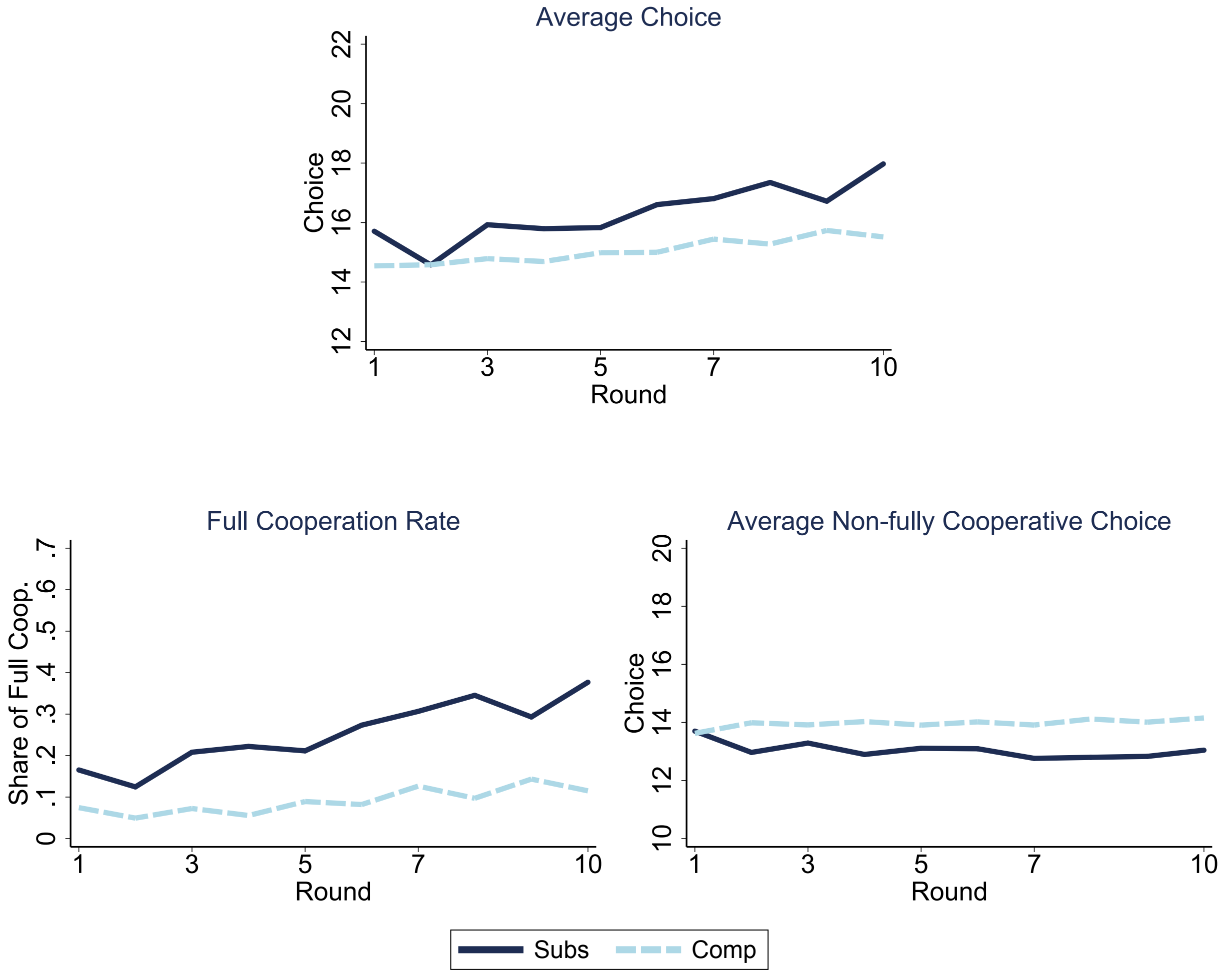

Notes: This figure shows the simulated evolution of choices and cooperation within matches for model 1. In the upper panel for all choices, in the lower left-hand panel for full cooperation rates, and on the lower right-hand panel for non-fully cooperative choices (i.e. choices outside the range $[25, 26]$).

of cooperativeness of the partner in round $t-1$: $\alpha_t = 1/(1+e^{-k(y_{t-1}-14)})$ (as in PS). In the first two rounds of play α_t is constant, say, equal to α. In our experiment we observe different levels of full cooperation in the first rounds. For this reason, in our simulations we employ different values of α for the two treatments: α^{Comp} and α^{Subs}, which are based on the full cooperation rates observed in the first rounds. The non-cooperative play of a cooperative type consists of myopic best-reply with probability $\beta(1-\alpha_t)$ and spiteful play with probability $(1-\beta)(1-\alpha)$.

We assume that the share of cooperative players is p and the share of non-cooperative players is $1 - p$. This implies that in our simulations a share of p^2 matches consists of pairs where both players are cooperative types, a share of $(1 - p)^2$ of matches consists of pairs where both players are non-cooperative types, and a share of $2(1 - p)p$ of matches consists of pairs with a cooperative player and a non-cooperative player. In our simulations, we set $p^2 = 0.4$ and note that the key features obtained in the simulations we present are robust to changes in p. Finally, we determine the length of a match on our simulations by using the same continuation probability as in our experiment ($\delta = 0.9$).

Figure V shows the evolution of average choices, full cooperation rates, and average non-fully cooperative choices across rounds for model 1, each based on 1000 simulations with $k = 0.5$, $\alpha^{\mathsf{Comp}} = 0.06$, $\alpha^{\mathsf{Subs}} = 0.18$, and $\beta = 0.95$. We see that the simulations replicate the key features of the experimental data, except for the fact that the averages for non-fully cooperative choices are higher in our experiment in comparison to those obtained from the simulation. This might be due to the experience human subjects gain throughout our experiment. More specifically, towards the end of the experiment, when subjects gain more experience, it might be that non-cooperative players are also playing cooperatively to some extend. Indeed we observe in our experimental data that the average level of choices increases across matches. To add this intuition to our simulations, we modify the definition of non-cooperative players and run a second simulation model.[31].

In this second model, we define non-cooperative players to not only play non-cooperatively, but to also try to induce cooperation with a certain probability. We do so by assuming that non-cooperative players randomly choose an action above the Nash equilibrium of the stage game with a certain probability. More precisely, non-cooperative players play spitefully with probability $(1-\beta)$, play myopic best-reply with probability $\beta(1-\gamma)$, and play the maximum of a myopic best-reply and random choice above the static Nash equilibrium with probability $\beta\gamma$.[32]

Figure VI shows the evolution of average choices, full cooperation rates and average non-fully cooperative choices across rounds for model 2, each based on 1000 simulations with $k = 0.5$, $\alpha^{\mathsf{Comp}} = 0.06$, $\alpha^{\mathsf{Subs}} = 0.18$, $\beta = 0.95$, and $\gamma = 0.4$. One can see that with the new definition of non-cooperative players, we obtain higher levels for non-fully cooperative

[31]The only difference between the two simulations we present is the definition of a non-cooperative type

[32]We take the maximum of the myopic best-reply and the random choice above Nash to avoid the case where the best-reply is a higher choice than the random choice above static Nash, as such a case would be in conflict with the intuition that the non-cooperative player is trying to induce cooperation.

Figure VI: Simulation Results for Within-Match Behavior (Model 2)

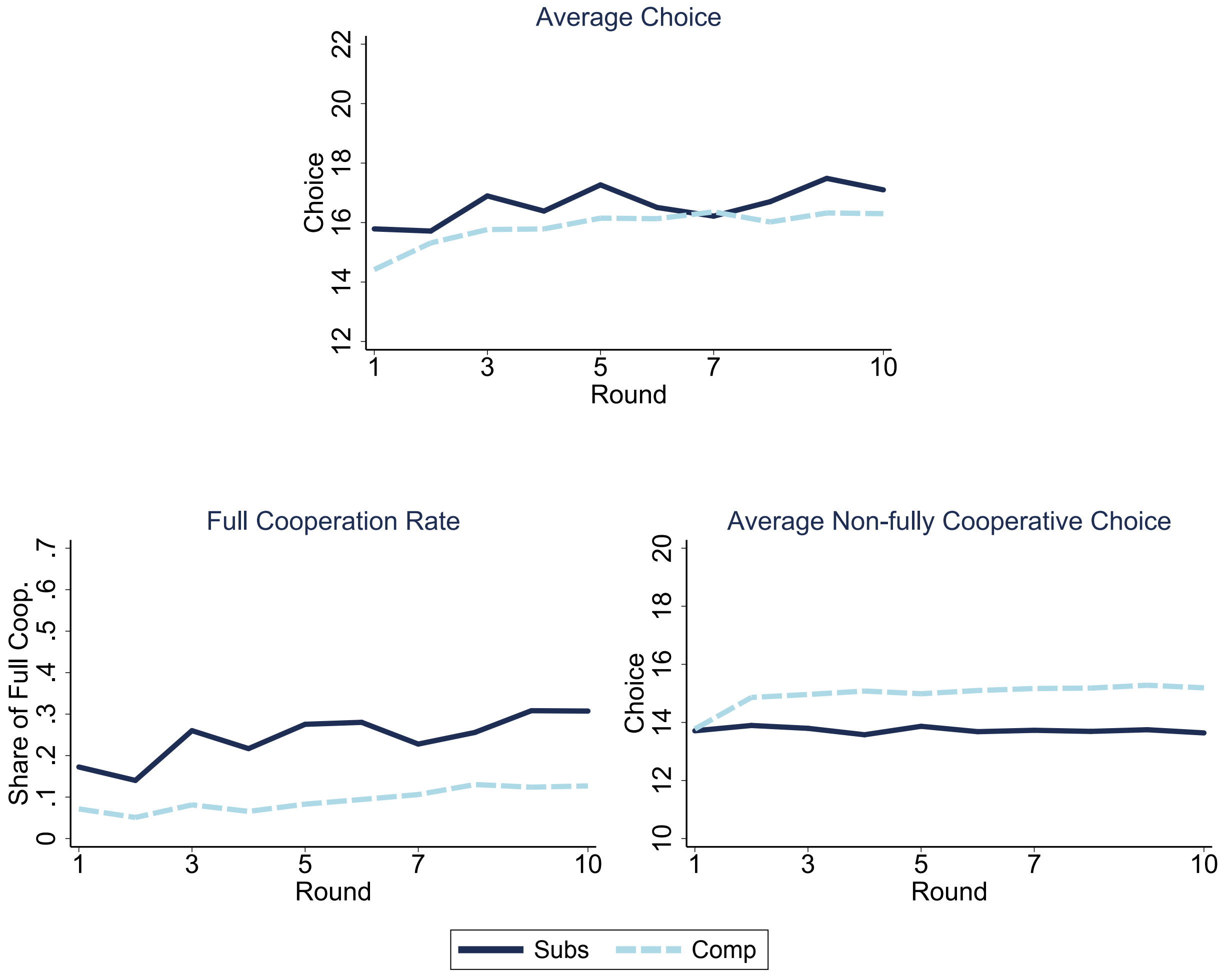

Notes: This figure shows the simulated evolution of choices and cooperation within matches for model 2. In the upper panel for all choices, in the lower left-hand panel for full cooperation rates, and on the lower right-hand panel for non-fully cooperative choices (i.e. choices outside the range $[25, 26]$).

choices, bringing them closer to the levels we observe in our experiment.

F. Additional Graphs

Figure VII: Equilibrium Range

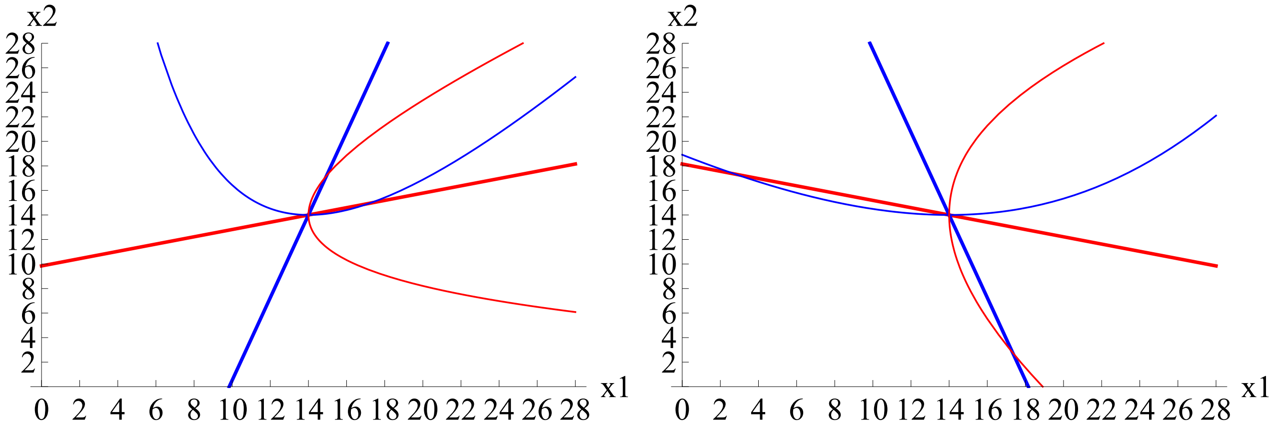

(a) Complements (b) Substitutes

Notes: This figure shows the best-reply functions and iso-payoff contours for our experimental games.

Figure VIII: Distribution of Match Lengths in the Experiment

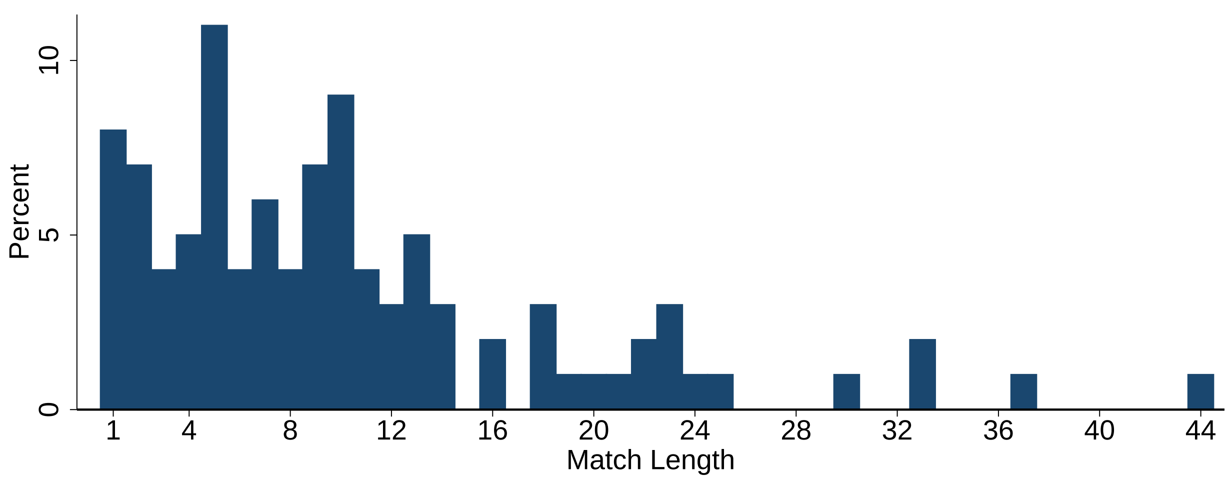

Notes: This figure shows the distribution of the randomly determined match lengths.

G. Additional Tables

Table IV: Regression results on payoffs

VARIABLES	(1) Payoff_i	(2) Payoff_i	(3) Payoff_i
Comp	0.208	−0.246	−1.306
	(1.640)	(1.424)	(1.256)
Round		−0.085	−0.059
		(0.051)	(0.049)
Comp*Round		0.050	0.059
		(0.062)	(0.060)
Match			0.255**
			(0.086)
Comp*Match			0.091
			(0.116)
Constant	33.68***	34.45***	31.45***
	(1.370)	(1.034)	(1.041)
Observations	33,024	33,024	33,024
R-squared	0.001	0.002	0.022

Notes: This table reports results from linear regression with standard errors (in parentheses) clustered at the session level. *** (**) [*] indicate that the estimated coefficient is significant at the 1% (5%) [10%] level.

Table V: Summary statistics at the individual level

	Average choices					
	First match		All matches		Last 10 matches	
	Comp	Subs	Comp	Subs	Comp	Subs
First round	17.33	17.55	17.50	18.63	17.90	19.37
All rounds	16.31	17.47	18.70	19.09	19.87	20.12

	Full Cooperation Rate					
	First match		All matches		Last 10 matches	
	Comp	Subs	Comp	Subs	Comp	Subs
First round	0.05	0.18	0.05	0.18	0.13	0.25
All rounds	0.06	0.09	0.14	0.27	0.20	0.36

	Average non-fully cooperative choices					
	First match		All matches		Last 10 matches	
	Comp	Subs	Comp	Subs	Comp	Subs
First round	16.08	16.69	17.02	17.03	17.51	17.58
All rounds	16.71	16.55	17.59	16.64	18.33	16.85

Notes: This table summarizes average choices (top panel), full cooperation rates (middle panel) and average non-fully cooperative choices (bottom panel). The results are reported for the first rounds and all rounds of the first match, all and the last 10 matches.