Strategy Revision Opportunities and Collusion*

Matthew Embrey†
University of Sussex

Friederike Mengel‡
University of Essex and Lund University

Ronald Peeters§
University of Otago

February, 2018

Abstract

This paper studies whether and how strategy revision opportunities affect levels of collusion in indefinitely repeated two-player games. Consistent with standard theory, we find that such opportunities do not affect strategy choices, or collusion levels, if the game is of strategic substitutes. In games of strategic complements, by contrast, revision opportunities lead to more collusion. We discuss alternative explanations for this result.

JEL classification codes: C73, C92, D43.

Keywords: strategy revision opportunities, cooperation, repeated games, complements vs. substitutes.

*We thank Jose Apesteguia, Antonio Cabrales, Pedro Dal Bó, Guillaume Fréchette, Konrad Mierendorff, Hans Theo Normann, Carlos Oyarzun as well as seminar participants in Brisbane, Düsseldorf, Maastricht (BEElab meeting), New York (ESA 2012), Vienna (Exp. Econ. workshop), Boston (IIOC 2013) and Lisbon (EARIE 2016) for valuable comments and the Netherlands Organisation for Scientific Research (NWO) and the European Union (FP7-ERG-277026) for financial support.

†Department of Economics, University of Sussex, Jubilee Building, Falmer, BN1 9SL, United Kingdom (E-mail: m.embrey@sussex.ac.uk).

‡Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom *and* Department of Economics, Lund University, SE-220 07 Lund, Sweden (E-mail: fr.mengel@gmail.com).

§Department of Economics, University of Otago, PO Box 56, Dunedin 9054, New Zealand (E-mail: ronald.peeters@otago.ac.nz).

1 Introduction

Strategy revision opportunities describe possibilities for players to change a strategy during the play of a repeated game. The role of strategy revision opportunities is somewhat of a conundrum for economists. From a theory perspective, strategy revision opportunities should not affect behaviour, since behaviour strategies, which allow full flexibility during the course of play, and mixtures over strategies chosen at the start of the game are seen as equivalent in games of perfect recall (Kuhn, 1953; Aumann, 1964). Any revision a decision-maker may want to make to a strategy during the course of play can be encoded in a suitably specified mixed strategy as long as arbitrarily complex strategies are allowed. As soon as there is some limit to the complexity (number of states) of strategies, revision opportunities could become important.

While standard subgame perfect equilibria are unaffected by revision opportunities, there are two types of considerations that can give some role to revision opportunities. One such refinement is renegotiation proofness. While renegotiation proofness implies that revision opportunities cannot increase collusion – since with revision opportunities players may not be able to credibly commit to the required punishment paths – weak renegotiation proofness, for example, has no bite in oligopoly games (Farrell, 2000; Aramendia et al., 2005). Another consideration is miscoordination, which could play an important role given the large number of possible equilibria in indefinitely repeated oligopoly games.

In this paper we conduct a laboratory experiment to study the impact of strategy revision opportunities on strategy choices and levels of collusion. This experimental approach enables us to elicit information about the strategies participants use in the repeated game, allowing us to identify strategy revisions and observe intended behavior off the realized outcome path. Such data is crucial to understand the mechanisms through which revision opportunities affect collusive behavior, and typically would be impossible to obtain from field data. Our experiment systematically varies revision opportunities across treatments, while keeping other factors such as the timing of moves, the available strategy sets, the size of the stakes, the number of players and the incentives to deviate or cooperate constant.

In the experiment, participants play an indefinitely repeated game where the stage game is either a game of strategic substitutes or of strategic complements. The stage games are derived from linear duopoly games (Cournot and Bertrand, respectively) and reduced to symmetric, normal-form games in which both players have four actions to choose from. The demand systems and action sets are chosen so that the resulting payoff matrices are as close as possible: they have identical diagonal elements (including the collusion and Nash outcomes), as well as identical temptation and sucker payoffs. The games primarily differ in the location of the (myopic) best response to collusion. In the substitutes game, the best response to collusion is less cooperative than the Nash action, while in the complements game it is more cooperative than the Nash action.

At the beginning of a supergame, participants program a strategy by choosing an initial action choice and a dynamic response machine, which specifies a recommended action in response to their rival's previous action choice. Three treatment variations change the degree to which strategy revisions are possible. These variations are labeled the baseline, unilateral and bilateral variations. In the baseline treatment participants cannot change their dynamic response machine; that is, they lack revision possibilities. Under the unilateral variation, unilateral changes are possible, while under the bilateral variation mutual consent is required to change one's dynamic response. The unilateral variation allows for full revision opportunities, while the bilateral variation is designed to mimic orchestrated revisions, such as those covered by renegotiation theory. In all variations, participants can deviate from their recommendation using one-shot deviations – these deviations come at a small cost.[1] A machine revision allows participants to economize on such costs in future choices.

We find that the existence of revision opportunities has no effect on cooperation under strategic substitutes, yet has a significant, positive effect under strategic complements. The second effect is large enough to reverse the ranking of collusion rates between interaction types: there is more cooperation under strategic substitutes if revision possibilities are absent, while the opposite is true if unilateral revisions are possible. Neither standard risk dominance nor considerations of renegotiation can explain these results. Given the large multiplicity of equilibria in these games, it is intuitive that fear of miscoordination might play an important role. We define a notion of "fear of miscoordination", based on minmax regret, and show that it yields predictions consistent with our main results on the effect of revision opportunities.

Our paper contributes mainly to two literatures: (i) literature on strategy revision opportunities and (ii) the literature on cooperation in games of strategic substitutes and complements.

There is not much research on strategy revision opportunities per se, but there is some experimental literature on cooperation/collusion that explicitly investigates the role of communication and renegotiation. Fonseca and Normann (2012), Andersson and

[1]This possibility is included so that participants have the full strategy space of the repeated game available.

Wengström (2012) or Cooper and Kühn (2014), for example, study renegotiation with communication and find mixed results as to whether communication, and the timing of communication, leads to more collusion or not. Our setting and results provide insight into renegotiation when explicit communication is not possible.[2]

A seminal study on cooperation in games of strategic substitutes and complements was conducted by Potters and Suetens (2009). They find more cooperation when actions exhibit strategic complementarities. As all their treatments are within a framework of behavior strategies, their results are best compared to our unilateral variations where strategy revision is possible. Our results with strategy revision opportunities confirm theirs. However, without revision opportunities, we find the opposite is true: there is more collusion with strategic substitutes. We discuss the intuition behind these differences in detail in Section 4.

Finally, our results also relate to the experimental literature on indefinitely repeated games. Usually this literature either elicits strategies without revision possibilities (Selten et al., 1997; Dal Bó and Fréchette, 2017) or lets subjects play the game in a "hot" setting without eliciting strategies (Dal Bó, 2005; Casari and Camera, 2009). Our results show that which setting is chosen can potentially affect behavior, at least when the game is one of strategic complements. One exception is Mengel and Peeters (2011), who have a "semi-hot" treatment (hot but with small costs) and a treatment without revision opportunities in a study comparing contributions by partners and strangers in a repeated public good game. Their setting is not suitable to study strategy revisions, however, since, although participants are allowed to deviate from pre-programmed strategies in some treatments, they are not allowed to revise strategies.

The paper is organized as follows: Section 2 outlines the experimental design and experimental procedures. Section 3 establishes our primary empirical results, in particular that revision opportunities have a positive effect under strategic complements and that this effect is large enough to reverse the ranking of collusion rates between the interaction types. Section 4 provides a discussion of the alternative explanations for our results, and formalises the concept of fear of miscoordination. A final section concludes.

[2]Note that the theory of renegotiation does not require *any* explicit communication between players. However, it does require both players willingness to change strategies. Our bilateral treatment provides a minimal signal (that is, allowing the other player to modify their machine) of willingness to "renegotiate" the current path of play.

2 The experiment

Designing experiments to understand strategic behaviour in indefinitely repeated games has two principal challenges. The first is the well-known theoretical problem of characterising the entire set of equilibrium strategies. The second concerns the complexity to directly elicit strategies due to the size and complexity of the strategy space, and constraints on participants' time, cognitive abilities and experience. Related to this second challenge, the existing experimental literature has usually limited the strategy space considerably in order to do so (Dal Bó and Fréchette, 2017). Such an approach has clear consequences for studying strategy revision opportunities since, under a restricted strategy space, the equivalence of behaviour and mixed strategies can break down. On the other hand, the impossibility to encode everything into a mixed strategy is one of the primary reasons why revision opportunities may matter in many real-life situations of interest.

Our design resolves these tensions by restricting participants to program a unit-recall dynamic response, but allowing them to deviate from the action proposed by the response to make available the full strategy space in all treatments. The experiment then studies revision opportunities in a 3×2 design with three levels of strategy revision opportunities and two types of strategic interaction.

2.1 Design

The two stage games. Participants play one of two possible games in Figure 1 that differ in the type of strategic interaction: *strategic substitutes* or *strategic complements*. Payoffs are in experimental currency units (ECU), which are converted to Euros at the end of the experiment.

	A	B	C	D
A	43, 43	31, 51	25, 52	23, 54
B	51, 31	36, 36	32, 40	29, 38
C	52, 25	40, 32	33, 33	31, 32
D	54, 23	38, 29	32, 31	30, 30

Strategic substitutes.

	A	B	C	D
A	43, 43	23, 54	14, 52	7, 47
B	54, 23	36, 36	32, 40	28, 37
C	52, 14	40, 32	33, 33	31, 32
D	47, 7	37, 28	32, 31	30, 30

Strategic complements.

Figure 1: The two stage games in the experiment.

The structure and payoffs of the games are designed so that, while each game has a natural duopoly analogue, the two are as identical as possible. To provide this analogue, the substitutes game is a discretized version of a differentiated-goods linear Cournot duopoly and the complements game is a discretized version of a differentiated-goods linear Bertrand duopoly. In both cases, the duopolists produce differentiated-goods that are product substitutes. To ensure a fair comparison across games, the underlying duopoly games were calibrated so that the majority of payoffs for key action pairs are identical across games:

1. the Nash equilibrium payoffs (π^{Nash}) that result from both players playing action C are identical.

2. the joint payoff maximizing payoffs ($\pi^{Collusion}$) that result from both choosing action A are identical.

3. the optimal deviation against the co-player playing action A, which requires playing action B in the complements game and action D in the substitutes game, yields the same payoff (π^{Dev}) for the defector and the sucker across games.

4. the remaining actions in the games, action D for the complements game and action B for the substitutes game, are such that all diagonal elements are identical across games.[3]

Sustaining cooperation via trigger strategies requires

$$\tfrac{1}{1-\delta}\pi^{Collusion} \geq \pi^{Dev} + \tfrac{\delta}{1-\delta}\pi^{Nash}$$

to be satisfied. All the payoff parameters involved in this inequality are the same for both our complements and substitutes variations.[4] As a consequence of these choices, the minimal discount factor needed to sustain collusion via trigger strategies is the same in both games ($\delta_{min} = 0.8077$) and the chosen continuation probability of $\delta = 7/8$ is above this level.

The crucial difference between the two games is the location of the optimal deviation against the co-player playing the joint payoff maximizing action, which is action B with strategic complements and action D with strategic substitutes. For convenience, we will refer to the actions A, B, C and D as respectively *Collusion*, *Dev.SC*, *Nash* and *Dev.SS*.

[3]After rounding the payoffs to numbers, one unit of payoff was changed to some payoffs in order to avoid degeneracies that are caused by rounding. This is done in such a way that games become even more similar: for instance, this led to the box formed by actions B and C and that formed by actions C and D being identical across games. See Section A of the supplementary materials for the underlying demand systems of the two games, as well as a description of the process that generated the discretized versions.

[4]The same is also true for other collusive strategies, such as tit-for-tat. While such strategies are not subgame perfect, they can be implemented without one-shot deviations or machine changes.

Repeated game strategies. At the beginning of a repeated game, participants are asked to specify an *intended strategy*. This strategy consists of an *initial action*, to be played in the first stage, and a programmed *machine*, which recommends at each later stage an action conditional on their co-player's action in the previous stage. The machine is denoted by a quadruple $z^A z^B z^C z^D$ specifying which action $z^k \in \{A, B, C, D\}$ the machine is programmed to play if the opponent has chosen action $k \in \{A, B, C, D\}$ in the previous stage. An intended strategy is denoted by $z^{\emptyset} - z^A z^B z^C z^D$, where the first element refers to the initial action choice.

The most general strategy one can formulate in a repeated game maps any possible history of observed action profiles into actions. In this design, however, participants' intended strategies are restricted so that actions can only be conditioned on their co-player's action in the previous stage. Some examples of familiar strategies that can be programmed are: unconditional cooperation (A–AAAA), tit-for-tat (A–ABCD), (forgiving) Nash reversion (A–ACCC), and always Nash (C–CCCC). Also strategies such as myopic best responses can be programmed. A well-known machine that cannot be programmed is grim-trigger. The ACCC-machine that comes closest implements a forgiving grim-trigger; that is, it reverts to cooperation if the opponent chooses to cooperate in some stage following a deviation.

While the focus is on simple strategies with few states and hence lower complexity, we allow participants to play other strategies as well. In particular, in all treatments participants are allowed to take an action that differs from the one recommended by their machine. Such changes are referred to as one-shot deviations. Consequently, more general strategies, such as grim-trigger, become feasible to implement via one-shot deviations.[5] These trigger strategy can be used to sustain collusion for all discount rates above the δ_{min} calculated earlier for the standard repeated game (i.e. just action choices in each period).[6]

To provide participants with an incentive to program their machines (strategies) carefully, one-shot deviations are costly. Each one-shot deviation costs 3 ECU. Hence, we expect participants to rely mostly on unit recall strategies. However, if participants

[5]Even though one-shot deviations have not been used a lot during the experiments (see Table 2), the machine AAAA (unconditional cooperation) accounts for a large majority of the instances in which they were used. Conditional on using AAAA the frequency of one-shot deviations is 58% or 77% in the substitutes and complements variations of the baseline treatment, respectively. For all other machines the frequency is below 10%. While these one-shot deviations came from very few participants and hence should be interpreted with care, they were used exclusively to play Nash or to punish, where the punishment action coincides with the myopic best response in the case of strategic substitutes.

[6]See Section B of the supplementary materials for details.

have strong enough preferences to choose another strategy, the full strategy space is available for participants in all treatments.

Revision opportunities. There are no revision opportunities in the treatments labeled *baseline*. Here, participants keep their machines for the entire duration of the repeated game, and can only deviate from the recommendations of their machines via one-shot deviations. In the *unilateral* treatments strategy revisions are possible. Participants can modify their machines after any stage of the repeated game.

To provide participants with an incentive to program their machines (strategies) carefully, machine modifications also have a small cost associated with them. In particular each machine modification costs 1 ECU, irrespective of the number of elements of the machine that are changed. Choosing the costs in such a manner we hoped to ensure that playing with a poorly programmed strategy is more costly in the baseline (where one needs to rely on one-shot deviations) than with unilateral revision opportunities (where machine changes are possible).

Again, collusive equilibria can be supported for all discount rates above δ_{min} using the same trigger strategy (A–ACCC) as in the baseline, which only relies on one-shot deviations to ensure permanent Nash reversion – see Section B.3 of the supplementary materials for further details.[7] Under the third variation – labeled *bilateral* – participants can modify their machines if and only if consent to a modification has been given by their opponent.

2.2 Procedures

The experiment was conducted in the BEElab at Maastricht University during October–December 2011. 288 students were recruited using ORSEE (Greiner, 2015) and participated in one of the six treatments.[8] For each of our treatments we have 6 independent observations. During each session, three matching groups were run in parallel on separate z-Tree servers (Fischbacher, 2007). Sessions lasted an hour and a half on average,

[7]That behavior is not unduly affected by the small cost imposed on these deviations is also shown empirically using a hot variation, where both the costs of machine changes and one-shot deviations were set to zero. See Embrey et al. (2016) for details of this treatment, where the use of such dynamic response machines to elicit strategies in repeated games is described, along with a detailed investigation of the elicited strategies for the more commonly implemented environment with revision opportunities.

[8]Other than the treatments mentioned we did not conduct any additional treatments. We conducted one pilot session with a 6×6 game (and some other differences in design), after which we decided to switch to a 4×4 game to reduce complexity for participants and hence the duration of the experiment.

including a twenty minute instruction period. On average participants earned between 12.60 and 15.30 Euro for their participation.

For each treatment six matching groups were run. Each matching group comprised eight participants that all played the repeated game (of the same treatment) ten times. At the beginning of a *match*, as a single repeated game is referred to, participants within a matching group were randomly paired. At the end of a session, participants were paid in cash according to the amount of ECUs they earned in one randomly drawn match. Table 2 gives the number of observations for each treatment.

Participants were fully informed about all details of the decision task, the environment and procedures in the experimental instructions (see Section C of the supplementary materials for an example of the instructions). Participants were never informed of the machine employed by other participants, but instead observed the history of play. That is, after every stage they were informed of their own action and the action of the person they were matched with, as well as the resulting payoffs.

For all members in a matching group, any given match consisted of the same number of stages, but this number changed across matches. Across matching groups this sequence of match-lengths differed. However, to facilitate comparison between treatments, the sequences were generated at random upfront and the same sequences were used for the different matching groups of each treatment. Table 1 provides details on the sequence of match lengths for the different matching groups.

Table 1: Number of stages played in the ten matches for the six different matching groups.

Matching	Match										
Group	1	2	3	4	5	6	7	8	9	10	Total
1	13	8	1	4	1	5	20	7	8	2	69
2	1	4	3	4	15	18	15	6	2	2	70
3	10	10	10	8	3	2	2	13	11	12	81
4	9	5	8	10	9	4	12	12	18	4	91
5	2	1	9	14	15	14	3	8	20	6	92
6	6	4	6	8	3	11	8	26	19	7	98

3 Results

Table 2 provides a summary of the six treatments. In general participants had difficulty establishing more cooperative behavior, capturing on average less than 25% of the potential gains from cooperating in all treatments. The treatments with strategic complements provided both the least and the most cooperative behavior, with low levels of cooperation in the baseline and bilateral treatments and high levels in the unilateral treatment. In all treatments, participants incurred very low costs for deviating from or modifying their machines. One-shot deviations are observed in less than 11% of stage games. In the unilateral treatments, machine modifications were implemented after less than 4% of stage games, while in the bilateral treatment (in which mutual agreement was required), machine modifications were made after less than 1% of stage games.

Table 2: Summary of treatments.

	Match. Groups	Num. Subj.	Num. Matches	Num. Stages Per Match (Avg.)	Efficiency (%)	Deviations 1-Shot (%)	Deviations Machine (%)
Substitutes							
Baseline	1–6	48	10	8.35	17.6	9.6	
Bilateral	1–6	48	10	8.35	16.9	10.5	0.2
Unilateral	1–6	48	10	8.35	16.5	6.5	3.7
Complements							
Baseline	1–6	48	10	8.35	10.9	5.1	
Bilateral	1–6	48	10	8.35	10.1	6.8	0.4
Unilateral	1–6	48	10	8.35	23.2	5.1	3.0

Note: Efficiency = $100 \times \left(\frac{\text{average stage game payoff} - \pi^{Nash}}{\pi^{JPM} - \pi^{Nash}}\right)$, where the stage game payoff is averaged over all matches and all stages. The column 1-shot deviations shows the percentage of stage games in which one-shot deviations were observed and the column "Machine" shows the percentage of stage games in which machine changes are observed.

The majority of the subsequent analysis uses data from the last third of a session (matches 7–10). This sub-sample provides a reasonable trade-off between using the final matches, where subject behavior is most likely to have converged, and ensuring enough observations. Only when analysing the evolution of behavior across matches, or behavior in particular histories for which there is the need to expand the sample size, we increase the sub-sample to data from the last two-thirds of a session (matches 4–10). In terms of stages we use only data from stages twelve or earlier. The reason is that later stages did not exist in each match for each matching group (see Table 1). All reported regressions and statistical tests use cluster-robust standard errors, corrected for arbitrary correlation at the matching-group level (see Section E.3 of the supplementary materials for robustness checks using matching-group averages and non-parametric statistics).

Table 3: Linear regression of payoff efficiency in the stage game.

	Substitutes				Complements			
	(1)		(2)		(3)		(4)	
Bilateral	−0.06	(0.477)	0.08	(0.381)	−0.05	(0.392)	−0.18*	(0.063)
Unilateral	−0.03	(0.745)	0.02	(0.837)	0.19***	(0.003)	0.06	(0.294)
Stage	−0.02***	(0.000)	−0.01***	(0.000)	−0.01**	(0.040)	−0.01**	(0.044)
Bilateral x Stage	−0.00	(0.765)	−0.00	(0.551)	0.01	(0.148)	0.01	(0.105)
Unilateral x Stage	0.00	(0.730)	0.00	(0.804)	−0.01	(0.389)	−0.00	(0.517)
Match			0.06***	(0.001)			−0.00	(0.997)
Bilateral x Match			−0.05*	(0.061)			0.05**	(0.023)
Unilateral x Match			−0.02	(0.528)			0.05**	(0.010)
Constant	0.36***	(0.000)	0.21***	(0.000)	0.11**	(0.012)	0.12***	(0.009)
R2	0.03		0.04		0.10		0.12	

Notes: The baseline case is the baseline treatment. All regressions use data from matches 7–10 and stages 1–12 and include match-stage composition dummies. All regressions have 2352 observations across 18 matching groups (clusters). VCE clustered at the matching-group level. p-values are reported in parentheses. ***1%, **5%, *10% significance.

3.1 Impact of revision opportunities on cooperation

We consider two measures of cooperation. The primary measure is the actual surplus generated over and above the one-shot Nash equilibrium as a percentage of the maximum available surplus (efficiency). This measure aggregates the impact of all choices, including partial collusion and deviation choices. The secondary measure focusses on just the percent of (full) collusion choices by subjects.

Figure 2 shows the evolution of efficiency both across matches and within matches. As can be seen, across matches (Panel (a)) revision opportunities do not seem to affect first-stage efficiency with strategic substitutes while there is a clear separation of the unilateral variation from the other two conditions with strategic complements. All treatments show the same pattern of decline of efficiency across stages within matches (Panel (b)).

A linear regression on efficiency is used to formally quantify the effect of strategy revision opportunities and to understand statistical significance. Table 3 reports the results of this analysis separately for strategic substitutes and complements. In columns (1) and (3), a complete set of dummies for the various levels of revision opportunities is interacted with the stage number of the match.[9] This regression confirms the impres-

[9]The regression includes matching-group dummies that control for possible fixed effects resulting from the particular draw of the number of stages in each match. See Tables E.1, E.2 and E.3 of the supplementary materials for a series of robustness checks for these regressions.

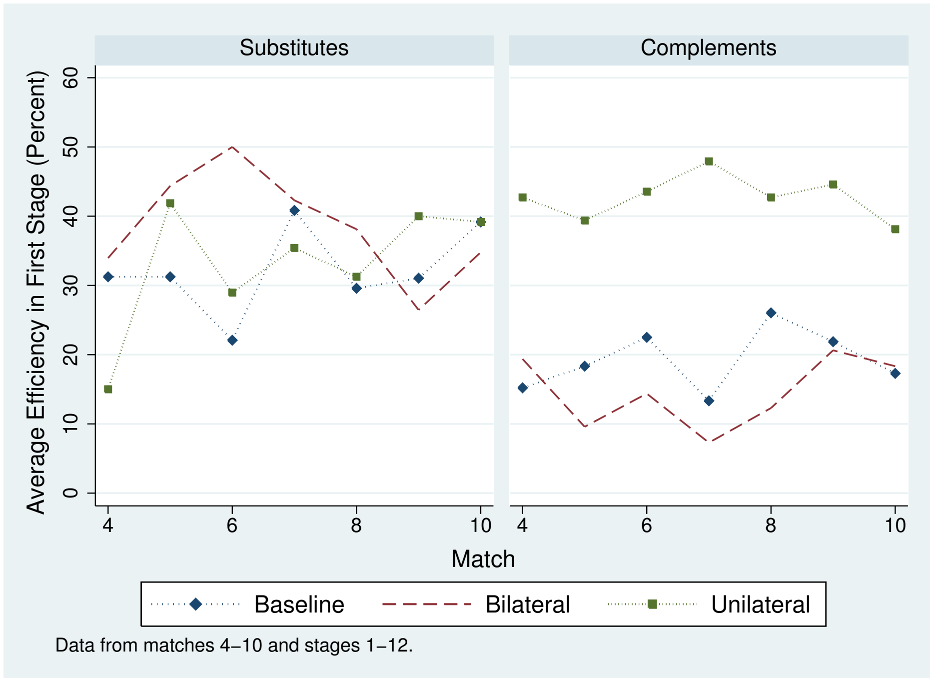

(a) Efficiency across matches.

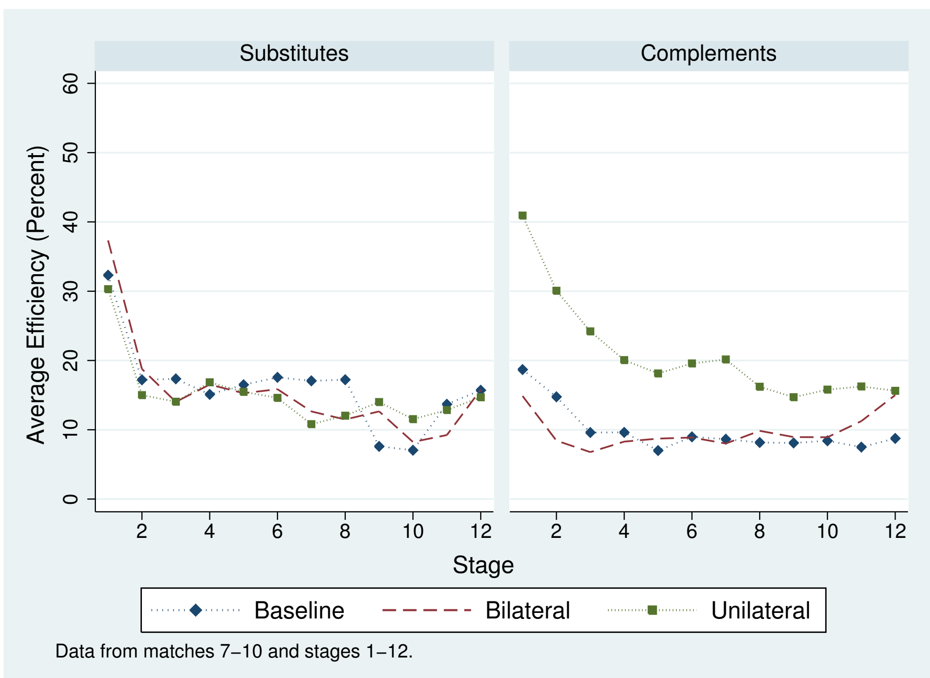

(b) Efficiency within matches.

Figure 2: The effect of strategy revision opportunities on efficiency with strategic complements and substitutes.

sion that strategy revision opportunities have no impact on rates of collusion under strategic substitutes, but have a significant impact under strategic complements. With strategic complements, the unilateral variation is associated with significantly higher rates of collusion (column (3)), while the baseline and bilateral variation have statistically similar rates. Note also that the R2 is substantially higher for the complements variation, arguably because revision opportunities can explain more of the observed variation in this setting.

Columns (2) and (4), expand the specification to include the match number as a dependent variable. Doing so reveals that strategy revision opportunities also work through the dynamics across matches. These regressions also show significant effects for the bilateral variation in the treatments with strategic complements. Here the mean is about the same as in the baseline but the dynamics are quite different, as can be seen in column (4). These patterns are most easily seen in Figure E.1 of the supplementary materials, which graphs the predicted efficiency from the linear regression.

Figure 3 shows the evolution of choices for the collusive action over time. Panel (a) shows this time trend across matches for choices in the first stage. With strategic substitutes (left panel), initial stage collusive choices are increasing over matches for the baseline and unilateral variations, while no clear trend is seen when bilateral consent is needed to modify machines. An obvious ranking of levels of revision opportunities, on the extent to which it generates collusion, cannot be made on the basis of this graph. With strategic complements (right panel), the graph illustrates a clear separation of treatments. Collusion rates are highest under unilateral. Under the baseline and bilateral variations, collusion rates are lower. There is a trend for collusion rates to increase over time in the unilateral treatment. However, no such trend is evident for the other treatments. Panel (b) of Figure 3 shows how the rates of collusion change within a match. It displays the typical pattern of collusion decreasing sharply after the first few stages, then remaining approximately constant at a lower level. A logit regression using choice of action A as the dependent variable can be found in Table E.4 of the supplementary materials. The results of the logit regression show the same qualitative patterns in terms of treatment comparisons as the OLS regressions on efficiency shown in Table 3.

To sum up, we find that with strategic substitutes, strategy revision opportunities do not affect average collusion, while with strategic complements, strategy revision opportunities have a positive effect on collusion. This conclusion is based on the following Result:

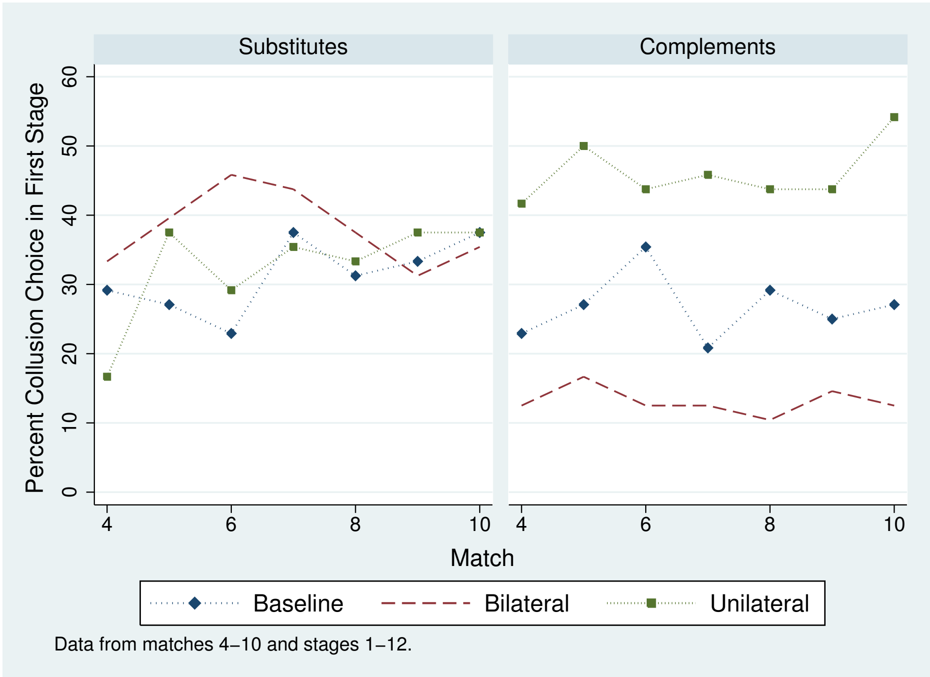

(a) Collusion choice across matches.

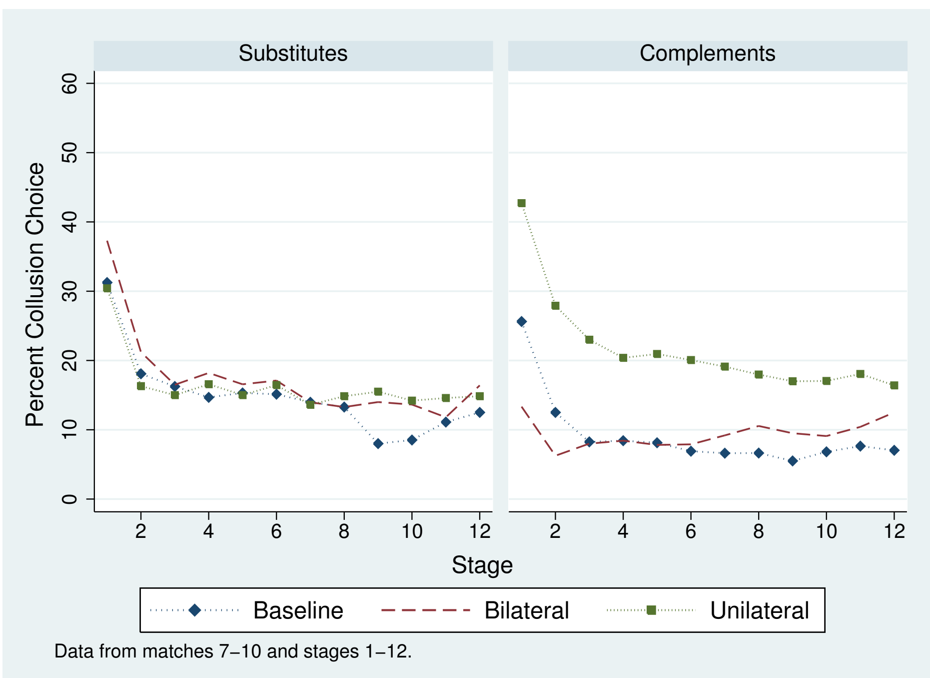

(b) Collusion choice within matches.

Figure 3: The effect of strategy revision opportunities on collusion with strategic substitutes and complements: percentage of collusion.

Result 1 *(i) Under strategic substitutes there is no treatment difference in average efficiency or collusion rates across the baseline, unilateral and bilateral variations. (ii) Under strategic complements average efficiency as well as average collusion rates are higher under the unilateral compared to the other variations.*

3.2 Difference between complements and substitutes

Strategy revision opportunities also affect the comparison between complements and substitutes. With fewer or no revision opportunities there is more collusion with strategic substitutes than with strategic complements: 35.2% versus 19.6% for first stage choice under the baseline variation (p-value = 0.08), and 35.4% versus 14.6% under the bilateral variation (p-value < 0.01).[10] However, with revision opportunities there is more collusion with strategic complements: 36.5% versus 43.5% (p-value = 0.07) under the unilateral variation.

Table 4: Linear regression of payoff efficiency in the stage game – strategic substitutes versus complements.

	Combined regression			
	(1)		(2)	
Complements	-0.15^{**}	(0.044)	0.02	(0.791)
Complements x Bilateral	0.01	(0.927)	-0.27^{**}	(0.036)
Complements x Unilateral	0.22^{**}	(0.040)	0.04	(0.770)
Complements x Stage	0.00	(0.668)	0.00	(0.885)
Complements x Bilateral x Stage	0.01	(0.188)	0.02^{*}	(0.092)
Complements x Unilateral x Stage	-0.01	(0.360)	-0.01	(0.503)
Complements x Match			-0.07^{***}	(0.005)
Complements x Bilateral x Match			0.11^{***}	(0.002)
Complements x Unilateral x Match			0.07^{*}	(0.078)
Constant	0.31^{***}	(0.000)	0.16^{***}	(0.007)
R2	0.05		0.07	

Notes: The baseline case is the baseline treatment with strategic substitutes. All regressions use data from matches 7–10 and stages 1–12 and include match-stage composition dummies. Both regressions have 4704 observations across 18 matching groups (clusters). VCE clustered at the matching-group level. *p*-values are reported in parentheses. ***1%, **5%, *10% significance. The table reports just the results for regressors involving the complements dummy variable plus the constant term; see Table E.5 of the supplementary materials for the complete output.

This strategic interaction effect is further quantified using analogous regression specifications to those reported in Table 3, except the data from both game types are

[10]See Table E.10 of the supplementary materials for a complete set of comparisons across game types.

pooled and additional dependent variables are added – a strategic complements indicator interacted with the levels of revisions opportunities and stage and with the level of revision opportunities and match. Table 4 reports the results of this exercise.[11] The results show that in the baseline treatment, average payoff efficiency is higher under substitutes as illustrated by the negative coefficient on the complements dummy in column (1) (p-value 0.044). In the unilateral treatment, by contrast, payoff efficiency is higher under complements, as shown by the sum of coefficients on "complements" and "complements $\times$ unilateral" (p-value 0.073). These results also confirm the overall message, with respect to the comparison across game types that is visible in Figure 2. Namely, there is a significant effect of the type of strategic interaction on the development of collusion across matches. The development of collusion within a match is comparable across game types.

Result 2 *With revision opportunities (unilateral) there is (weakly) more collusion in first-stage choices and (weakly) higher average payoff efficiency with strategic complements. Without revision opportunities (baseline) there is (weakly) more collusion in first-stage choices and (weakly) higher average payoff efficiency with strategic substitutes.*

3.3 Individual behavior

The previous subsections dealt with the impact of revision opportunities and the type of strategic interaction by assessing outcomes along the realized path of play. To understand further what drives these realized paths, this subsection analyses individuals' intended strategies. To this end, Table 5 gives the distribution of the machines programmed at the beginning of matches 7–10, along with a breakdown of the initial choices associated with each machine.

In principle there are 256 different machines – that is, the dynamic response part of the intended strategy – that participants could use. However, only seven types were used with a frequency of at least 5 percent in at least one of the treatments.[12] These prominent types of machine seem reasonable, with the majority corresponding to strategies that are commonly seen either in prior experimental studies or from the

[11] See Table E.6 of the supplementary materials for the analogous analysis using the probability of choosing the collusive action.

[12] 20–36 percent of programmed machines do not fall into one of the seven prominent categories. Table E.7 in the supplementary materials decomposes this "Other" category to show that the seven prominent machines give a fair characterization: many machines categorized as "Other" in Table 5 are minor deviations from these. Using this decomposition, the "Other" category drops to 6–19 percent.

theory of repeated games. The first four – AAAA, ABC(C/D), ACCC and BBCC – attempt to establish some collusion, either unconditionally or conditionally.[13] Of the non-cooperative machines, there is the static Nash machine (CCCC), the myopic best response machine (DCCC under substitutes and BCCC under complements) and the punishing machine (DDDD).

Table 5: Distribution of initial choice and machine categories (in percent).

Prominent Machine Category	Strategy Initial-Machine	Substitutes Base	Substitutes Bil	Substitutes Uni	Complements Base	Complements Bil	Complements Uni
Unconditional cooperation	AAAA	**1**	**1**	**5**	**1**	**1**	**8**
	A-	*1*	*1*	*5*	*0*	*1*	*8*
Conditional cooperation	ABC(C/D)	**18**	**16**	**11**	**10**	**18**	**23**
	A-	*12*	*8*	*6*	*6*	*6*	*19*
	B-	*1*	*2*	*2*	*2*	*6*	*4*
Nash reversion	ACCC	**12**	**16**	**7**	**16**	**5**	**12**
	A-	*10*	*13*	*7*	*12*	*2*	*12*
Partial collusion + Nash rev.	BBCC	**1**	**2**	**2**	**5**	**7**	**1**
	B-	*0*	*0*	*1*	*4*	*6*	*1*
Nash	CCCC	**17**	**10**	**22**	**22**	**26**	**14**
	C-	*11*	*9*	*14*	*21*	*19*	*13*
Myopic best reponse	DCCC	**16**	**12**	**19**			
	C-	*7*	*4*	*7*			
	D-	*7*	*7*	*11*			
	BCCC				**25**	**21**	**15**
	B-				*5*	*10*	*10*
	C-				*17*	*10*	*4*
Punishing	DDDD	**0**	**7**	**3**	**0**	**0**	**0**
	D-	*0*	*7*	*3*	*0*	*0*	*0*
Other	———	**35**	**36**	**31**	**20**	**22**	**27**
All machines with a ...							
cooperative response to action A		40	48	35	33	30	52
cooperative response to action B		28	28	29	22	33	39
deviation response (D/B) to action A		31	33	38	35	35	32
punishing response to deviation (D/B)		22	37	23	3	2	4

Notes: Distribution of prominent machines programmed at the beginning of matches 7–10 (in bold), along with a breakdown of the prominent initial choices associated with each machine (in italics). Machine combinations that were used with a frequency below 5 percent in every treatment are categorized as "Other". A cooperative response to A is any machine that chooses A in response A; a cooperative response to B is any that chooses A or B in response to B; a punishing response to the deviation action is any that chooses D in response to D under substitutes and D in response to B under complements.

The treatment comparisons in Table 5 mirror those seen in the data from the real-

[13] AAAA is unconditional cooperation, ABCD is tit-for-tat or also imitation (Apesteguia et al., 2007), ACCC is Nash reversion and BBCC could be interpreted as cautious or partial cooperation.

ized path of play. Under strategic substitutes there is no consistent effect of revision opportunities on intended strategies. In particular, there is no significant difference between the baseline treatment and the unilateral treatment in the likelihood of programming a collusive dynamic response – that is, either responding cooperatively to the other player having played A (40% versus 35%; p-value = 0.520) or to the other player having played B (28% versus 29%; p-value = 0.785).[14] Indeed the only significant difference under substitutes is that subjects are more likely to respond cooperatively to action A with bilateral revision opportunities than unilateral ones. This more collusive response under bilateral is counter-balanced by a greater use of punishing responses, especially after the other player has played the deviation action (although in isolation this difference between bilateral and unilateral, 37% versus 23%, is not significantly different, p-value = 0.125).

By contrast, revision opportunities have a significant and positive effect on the cooperativeness of intended strategies under strategic complements. In the unilateral treatment, subjects are significantly more likely to respond in a collusive manner after the other player chose A than in both the baseline (52% versus 33%; p-value = 0.019) and the bilateral (52% versus 30%; p-value < 0.001) treatments. In addition, subjects are also more likely to respond with a collusive action after action B in the unilateral than in the baseline treatment (39% versus 22%; p-value = 0.012). As seen in Table 5, intended strategies are generally composed of the "intuitive" initial choice to go with the chosen machine. Furthermore, as seen in Figures 2 and 3, stage-one choices and efficiency respond strongly to the change in revision opportunities under strategic complements.[15]

Result 3 *(i) With strategic substitutes, strategy revision opportunities have no clear effect on intended strategies. (ii) With strategic complements, strategy revision opportunities lead to an increase in the use of more cooperative dynamic responses.*

In summary, strategic complementarity appears to induce more collusive outcomes

[14] These p-values are based on a linear random-effects regression on revision-opportunities treatment dummies. The p-values quoted in this and the subsequent paragraphs are based on analogous tests. See Table E.11 of the supplementary materials for a complete set of pairwise tests for the four summary statistics reported at the bottom of Table 5, as well as the p-values obtained via a non-parametric test on matching-group averages that were run as a robustness check.

[15] These observations lead to the question of whether this effect of revision opportunities is driven entirely by initial choice. Since our experimental design collects data on both the realized path of play and intended strategies, it is possible to analyse the role of such path dependency. This analysis is reported in Section D of the supplementary materials (see Mengel and Peeters, 2011, for an earlier example of such analysis). In summary, both initial choices and dynamic responses appear to be important.

when players have the opportunity to revise their initial intended strategies for two reasons. The first, and most direct effect, is that intended strategies are more collusive, both in terms of more efficient initial choices and more cooperative dynamic responses. However, there is a second reinforcing effect in how participants respond to myopic best response actions (action B under strategic complements and action D under strategic substitutes), which is revealed by the numbers in the bottom part of Table 5. The myopic best response action triggers a punishing response in 22-37% of the cases under strategic substitutes and in only 2-4% of the cases under strategic complements. Under complements, a participant playing the myopic best response is still quite likely to iterate to at least a partially cooperative outcome; under substitutes, such a strategy most likely instigates a spiral towards a Nash or punishment outcome.

4 Discussion

Our primary result is that, while strategy revision opportunities have no effect on collusion under strategic substitutes, they have a significant positive effect under strategic complements. What could explain the role of revision opportunities, given that the standard prediction from game theory suggests it should play no role? In what follows, we discuss two popular concepts from the theoretical and experimental literature on repeated games (see, for example, Farrell and Maskin, 1989; Fonseca and Normann, 2012; Blonski et al., 2011; Dal Bó and Fréchette, 2011). Neither will be able to provide a satisfactory explanation for our results. We define a notion of fear of miscoordination that yields predictions in line with the observed effect of changing the availability of revision opportunities.

4.1 Renegotiation

The observed ranking of collusion rates across treatments goes against the intuition delivered by the renegotiation literature. In particular, with fewer revision opportunities, and hence reduced concerns for renegotiation, we observe less collusion under strategic complements.

Although renegotiation should never happen in equilibrium – whether collusion is weak renegotiation proof or not – it is reasonable to expect the strategic forces that drive the concept would need to be learnt by experience. Consequently, there is still the possibility that subjects engaged in something like renegotiation, but that such efforts did not feed back into reduced collusion at the beginning of a match. We can look at our data from the bilateral treatment to see whether many attempts to "renegotiate" were made. Bilateral modifications take place very rarely. In particular, during matches 7–10 there were only 6 bilateral deviations (out of over 1200 interactions) with strategic substitutes and 9 with strategic complements from respectively 5 and 7 machines. Consequently, there are few instances where participants succeeded in coordinating on a mutual modification of their machines.

Nonetheless, the data collected on strategic decisions throughout the experiment allows some analysis regarding which paths are "renegotiated", and if so, how. To do so, we study when and how machines are (attempted to be) modified conditional on the last outcome of the realized path of play. We classify paths into three categories: (i) "failed collusion" (outcomes (B,A) and (A,B)), (ii) "miscoordination" (from the perspective of a cooperative agent; outcomes (A,C), (A,D), (B,C) and (B,D)) and (iii) "punishment paths" (outcomes (C,D), (D,C) and (D,D)). After a "miscoordination" stage, participants mostly try to modify cooperative machines into more punishing machines. Along "punishment paths", participants mostly (want to) modify non-cooperative machines, but rarely modify them into more collusive machines. Renegotiation theory, though, would say that participants modify punishing machines into more cooperative machines that allow them to leave a punishment stage.[16] Hence, in addition to the treatment comparisons, there is also no direct evidence that participants engaged in something like renegotiation.

4.2 Risk-dominance

Given the large number of possible equilibria in these indefinitely repeated games, it seems intuitive that considerations of renegotiation might be overshadowed by concerns of coordination on one of the different equilibria. Hence, it seems reasonable to look at risk-dominance, since it gives a role for a fear of equilibrium miscoordination. As discussed in Dal Bó and Fréchette (2011), extending the idea of risk-dominance to infinitely repeated games poses a number of difficulties, even with only two actions for each player. These difficulties include extending the definition to repeated-game strategies and the issue that two repeated-game strategies can generate equivalent outcome paths. To these difficulties, our environment also adds the issue of extending the definition of risk dominance to more action choices in the stage game.

Blonski et al. (2011) consider an extension of the concept to the repeated prisoners' dilemma that involves only the strategies permanent Nash reversion and static Nash.

[16]Table E.8 of the supplementary materials provides more details.

Translating this approach to our environment results in the prediction that Nash reversion, with an initial choice of cooperation, is the risk-dominant strategy for both types of strategic interaction and all levels of strategy revision opportunities.[17] In what follows, we discuss how a definition of fear of miscoordination, which does not restrict itself to equilibrium miscoordination, can accommodate the observed behavior.

4.3 Fear of miscoordination

Since neither renegotiation nor risk dominance considerations are in line with our results, we consider a refinement based on a notion of "fear of miscoordination". Given the large number of possible equilibria in indefinitely repeated games, fear of miscoordination seems a particularly relevant concern and its effect may overshadow any potential effect of renegotiation. Intuitively, players will be less concerned about miscoordinating if they have the possibility to revise their strategy during the course of play. Hence fear of miscoordination delivers exactly the opposite intuition compared to renegotiation in terms of how revision opportunities should affect collusion.

We formalize this idea as follows. For a player i using repeated game strategy s_i^*, the maximal regret possible for this strategy is given by

$$F_{NR} = \Pi_i(s_i^*, s_i^*) - \min_{s_{-i}} \Pi_i(s_i^*, s_{-i}). \tag{1}$$

F_{NR} is constructed as the difference between the payoff i expects when choosing s_i^*, while expecting her opponent to do the same, and the worst possible payoff that she could obtain by choosing this strategy. Notice that this definition is formulated from the perspective of a symmetric equilibrium in the context of two-player games, which is sufficient for the purpose of this study, but not crucial for the result we will state below.[18,19] Equation (1) describes the fear of miscoordination in cases where strategy revision is not possible.

[17] When focussing only on the Nash reversion and static Nash strategies, there is only one difference between the strategic complement and strategic substitute games: the (A,C) payoff in the complements game is lower than that in the substitutes game (14 rather than 25). With a discount rate of 7/8 and the uniform prior as the belief of the opponent's strategy, this difference is too small to result in different predictions for the risk-dominance concept. One could consider allowing for a non-uniform prior. However, only a relatively small range of beliefs would result in the static Nash machine being selected in the complements game, whereas the Nash reversion machine is selected in the substitutes game. The weight on the opponent choosing the static Nash machine would need to be at least 77% and no more than 88%. There is no support in the data for such a range of beliefs.

[18] Defining this notion for all equilibria would require knowledge about which equilibria can be supported by any given strategy.

[19] Chassang (2010) has used fear of miscoordination in a weaker sense. He studies dynamic global games and refers to fear of miscoordination as the possibility of miscoordination arising from a lack

To define fear of miscoordination if strategy revision is possible, define a sequence of pure strategies $(\mathbf{s}_\mathbf{i}^\tau)_{\tau=0,\dots,\infty}$ with the interpretation that at t the action prescribed by s_i^t given the history induced by $(\mathbf{s}_\mathbf{i}^\tau)_{\tau=0,\dots,t-1}$ is chosen. Note that any fixed sequence of strategies can simply be expressed as a (potentially very complex) strategy itself and that constant sequences are possible as well. Denote by $\mathbf{s}$ a pair of such sequences $\mathbf{s} = ((s_1^\tau)_{\tau=0,\dots,\infty}, (s_2^\tau)_{\tau=0,\dots,\infty})$ and by $\Pi_i(\mathbf{s})$ player i's discounted average payoff under $\mathbf{s}$. Fear of miscoordination F can then be defined as,

$$F((\mathbf{s}^*)) = \Pi_i((\mathbf{s}_\mathbf{i}^*), (\mathbf{s}_\mathbf{i}^*)) - \min_{(\mathbf{s}_{-\mathbf{i}})} \Pi_i((\mathbf{s}_\mathbf{i}^*), (\mathbf{s}_{-\mathbf{i}})). \tag{2}$$

Of course the question arises, why would a player ever want to revise their strategy? After all, any situation under which a player would want to make a revision can be encoded in initial strategies as we have discussed in the introduction (Kuhn, 1953). However, such strategies may be relatively complex, for example being automata involving many states. If there is any arbitrarily small but fixed cost of using automata with more states, agents would not want to use these additional states to encoding revisions for zero probability events (that is, histories that they do not expect to be reached).

Table 6: Fear of miscoordination F for the most used machines.

		Baseline		Unilateral	
Strategy	Target	Strat. subst.	Strat. compl.	Strat. subst.	Strat. compl.
A–AAAA	(A,A)	$20\frac{1}{1-\delta}$	$36\frac{1}{1-\delta}$	$20+1+12\frac{\delta}{1-\delta}$	$36+1+12\frac{\delta}{1-\delta}$
A–ABCC	(A,A)	$20+12\frac{\delta}{1-\delta}$	$36+12\frac{\delta}{1-\delta}$	$20+12\frac{\delta}{1-\delta}$	$36+12\frac{\delta}{1-\delta}$
A–ABCD	(A,A)	$20+13\frac{\delta}{1-\delta}$	$36+13\frac{\delta}{1-\delta}$	$20+1+12\frac{\delta}{1-\delta}$	$36+1+12\frac{\delta}{1-\delta}$
A–ACCC	(A,A)	$20+12\frac{\delta}{1-\delta}$	$36+12\frac{\delta}{1-\delta}$	$20+12\frac{\delta}{1-\delta}$	$36+12\frac{\delta}{1-\delta}$
B–BBCC	(B,B)	$7+5\frac{\delta}{1-\delta}$	$8+5\frac{\delta}{1-\delta}$	$7+5\frac{\delta}{1-\delta}$	$8+5\frac{\delta}{1-\delta}$
C–CCCC	(C,C)	$2\frac{1}{1-\delta}$	$2\frac{1}{1-\delta}$	$2\frac{1}{1-\delta}$	$2\frac{1}{1-\delta}$
D–DCCC	(C,C)	$2\frac{\delta}{1-\delta}$		$2\frac{\delta}{1-\delta}$	
C–DCCC	(C,C)	$2\frac{1}{1-\delta}$		$2\frac{1}{1-\delta}$	
B–BCCC	(C,C)		$5+2\frac{\delta}{1-\delta}$		$5+2\frac{\delta}{1-\delta}$
C–BCCC	(C,C)		$2\frac{1}{1-\delta}$		$2\frac{1}{1-\delta}$
D–DDDD	(D,D)	0	0	$1+(-1)\frac{\delta}{1-\delta}$	$1+(-1)\frac{\delta}{1-\delta}$

Notes: Maximal regret is obtained if the opponent plays D–DDDD. The second column labeled "target" is the outcome on which the strategy aims to coordinate on. The maximal regret is computed relative to this target.

of common knowledge and in particular arbitrarily small amounts of private knowledge. His characterization of sequentially rationalizable strategies is related to risk dominance in the one-shot game and thus quite different from ours.

Table 6 shows the level of fear of miscoordination for prominent strategies in our main treatments, where "prominent strategies" are those that are used in at least 5% of the cases in at least one treatment.[20] As can be seen, fear of miscoordination is higher if there are no revision opportunities. The biggest difference in fear of miscoordination is identified when we compare the baseline and the unilateral variation for strategies with (A,A) target ($\frac{1-25\delta}{1-\delta}$, see Table 6). Here we also see in our results that participants do use more cooperative machines in the unilateral variation compared to the baseline variation (52% versus 33%; p-value = 0.019, see Table E.11) at least for strategic complements.

Fear of miscoordination is also higher for collusive strategies with complements than with substitutes. In particular, for strategies that target (A,A), the difference in fear of miscoordination between strategic complements and substitutes is $16\frac{1}{1-\delta}$ in the baseline condition and 16 under the unilateral variation. Thus, fear of miscoordination can also explain why revision opportunities (that decrease fear of miscoordination) have a bigger effect on the incidence of cooperative strategies under strategic complements compared to strategic substitutes.

4.4 Comparison to Potters and Suetens (2009)

Despite some design differences our result of greater collusion under strategic complements than strategic substitutes *in the presence* of strategy revision opportunities replicates Potters and Suetens (2009)'s seminal finding.[21] Potters and Suetens (2009)

[20]The numbers for the strategy A–AAAA under strategic substitutes are derived as follows. In the baseline treatment, where revision is not possible, a player expects to receive a payoff 43 every stage when expecting his opponent to apply the same strategy, while in the worst case he only receives a payoff of 23 every stage. The fear of miscoordination is here the difference between the present values of these two constant streams of payoffs: $20\frac{1}{1-\delta}$. In the unilateral treatment, the player again expects to receive a payoff of 43 every stage, but now is in the worst case able to adapt his strategy after realizing that the opponent did not use the same cooperative strategy. Then, the player receives a payoff of 23 in the first stage and continues to receive a payoff of 31 in all later stages. Taking into account the one unit cost for a modification of the machine, this situation renders the player a present value of $23-1+31\frac{\delta}{1-\delta}$. Hence, the fear of miscoordination is given by $43\frac{1}{1-\delta}-(23-1+31\frac{\delta}{1-\delta}) = 20+1+12\frac{\delta}{1-\delta}$. The other numbers in the table are derived in a similar manner.

[21]The first main design difference is that we repeat the stage game indefinitely. Interestingly, Gül Mermer et al. (2014) contest the robustness of the Potters and Suetens (2009) result in an indefinitely repeated games framework. They find that while on average there is not much difference between complements and substitutes, the percentage of fully cooperative choices is significantly higher under substitutes. In contrast, analysis of Dal Bó and Fréchette (2016)'s meta-data on the indefinitely repeated prisoner's dilemma game finds significantly more cooperation in strategic complements compared to strategic substitutes (controlling for discount factor, supergame and stage; see Embrey et al., 2016, for details). The second main design difference is that we used a smaller choice set in the stage-games. Along with changing the salience of important stage-game actions, and potentially

explain their finding by suggesting that in response to a collusive move "best responding players will move in the same direction if choices are complements and in the opposite direction if choices are substitutes." In other words, conditional on using for example myopic best responses, we should see more collusion under complements.

Our design allows us to also identify effects on intended strategies which tend to be more collusive under complements with revision opportunities (Section 3.3). The effect of revision opportunities on intended strategies can also explain why we find the opposite ranking than Potters and Suetens (2009) *in the absence* of strategy revision opportunities. Fewer cooperative strategies are used under strategic complements compared to strategic substitutes if revision opportunities are missing. This can, for example, be explained with fear of miscoordination as described in the previous subsection.

Finally, one can ask which environment (presence or absence of revision opportunities) seems more relevant in applications. The answer to this question can depend, among others, on the attention the decision-maker gives to the interaction, whether there are organisational constraints that allow or impede quick revisions or whether human actors or computers (for example, price bots) execute a strategy. The importance of these factors in applications will determine whether they are best thought of as settings with or without strategy revision opportunities.

5 Conclusion

We have studied the effect of strategy revision opportunities on collusion in infinitely repeated games and found that, while revision opportunities do not affect collusion with strategic substitutes, they have a positive effect with strategic complements. The latter effect is strong enough that, although there is more collusion with substitutes if there are few or no revision opportunities, there is more collusion with complements if there are revision opportunities.

Our findings challenge received wisdom in the repeated game literature. This literature has largely ignored strategy revision opportunities due to the theoretical equivalence between behaviour strategies, which allow full flexibility during the course of play, and mixtures over strategies chosen at the start of the game are seen as equivalent in games of perfect recall. Our results show that revision opportunities matter, despite

learning dynamics, the reduced action space allows us to hold some key marginal incentives constant: $\pi^{Collusion}$, π^{Nash}, but also π^{Dev} and the corresponding sucker payoff (see Section 2.1). By contrast, the latter payoff differs across game types in Potters and Suetens (2009).

the fact that participants have a "large" strategy space available and could encode most standard behaviour strategies at the beginning of the game.

Our results could be of importance in a number of possible applications. Revision opportunities are often set by higher-level management and play an important role in the design of organizations. Managers at higher levels of an organization have to make day-to-day decisions on strategic oversight – that is, how much flexibility to give to lower-level management to make choices, or revise initially set strategies, as market conditions unfold (Daily et al., 2003). The extent to which corporations regulate franchises "top down" or allow revisions to, for example, pricing strategies is only one example of such decisions. Policy makers should also be concerned about revision opportunities. Legal restrictions on meetings and agreements between market participants do affect the possibility of (orchestrated) strategy revisions and potentially affect prices and competition in local markets (see, for example, McCutcheon, 1997). In some markets, including those for gasoline and ready-mixed concrete, competitors are sometimes forced by authorities to announce their prices upfront, leaving little room for spontaneous price revisions. In some other markets computers (for example, price bots) are typically executing strategies and to which extent they are regulated differs across settings. Our findings suggest that in markets where competition exhibits strategic complementarity, any policy which complicates strategy revisions might have an additional preventative effect on collusion. By contrast, such policies seem less likely to have an effect in a market of strategic substitutes. Future research could explore strategy revision in more applied designs to allow us to gain more insight into the importance of these questions in these applications.

References

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

Apesteguia, J., S. Huck, and J. Oechssler (2007). Imitation: Theory and empirical evidence. *Journal of Economic Theory 136(1)*, 217–235.

Aramendia, M., C. Larrea, and L. Ruiz (2005). Renegotiation in the repeated cournot model. *Games and Economic Behavior 52(1)*, 1–19.

Aumann, R. (1964). Mixed and behaviour strategies in infinite extensive games. In

M. Dresher, L. Shapley, and A. Tucker (Eds.), *Advances in Game Theory*, Volume 52, pp. 627–650. Princeton University Press.

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(3)*, 164–192.

Casari, M. and G. Camera (2009). Cooperation among strangers under the shadow of the future. *American Economic Review 99(3)*, 979–1005.

Chassang, S. (2010). Fear of miscoordination and the robustness of cooperation in dynamic global games with exit. *Econometrica 78(3)*, 973–1006.

Cooper, D. and K. Kühn (2014). Communication, renegotiation and the scope for collusion. *American Economic Journal: Microeconomics 6(2)*, 247–278.

Daily, C., D. Dalton, and A. Cannella (2003). Corporate governance: Decades of dialogue and data. *Academy of Management Review 28(3)*, 371–382.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review 95(5)*, 1591–1602.

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

Dal Bó, P. and G. Fréchette (2016). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature in press*.

Dal Bó, P. and G. Fréchette (2017). Strategy choice in the infinitely repeated prisoners dilemma. Mimeo.

Embrey, M., F. Mengel, and R. Peeters (2016). Eliciting strategies in repeated games using dynamic response machines. Available at SSRN: https://ssrn.com/abstract=2896321.

Farrell, J. (2000). Renegotiation in repeated oligopoly interaction. In G. Myles and P. Hammond (Eds.), *Incentives, Organization, and Public Economics: Papers in Honour of Sir James Mirrlees*, pp. 303–322. Oxford University Press.

Farrell, J. and E. Maskin (1989). Renegotiation in repeated games. *Games and Economic Behavior 1(4)*, 327–360.

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

Fonseca, M. and H. Normann (2012). Explicit vs. tacit collusion – the impact of communication in oligopoly experiments. *European Economic Review 56(8)*, 1759–1772.

Fréchette, G. (2011). Session effects in the laboratory. *Experimental Economics 15(3)*, 485–498.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association 1(1)*, 114–125.

Gül Mermer, A., W. Müller, and S. Suetens (2014). Cooperation in indefinitely repeated games of strategic complements and substitutes. Mimeo.

Huck, S., H. Normann, and J. Oechssler (1999). Learning in cournot oligopoly – an experiment. *Economic Journal 109(454)*, 80–95.

Huck, S., H. Normann, and J. Oechssler (2000). Does information about competitors' actions increase or decrease competition in experimental oligopoly markets. *International Journal of Industrial Organization 18(1)*, 39–57.

Kaynak, E. (1991). *International Marketing: Sociopolitical and Behavioral Aspects*. Routledge.

Kuhn, H. (1953). Extensive form games and the problem of information. In H. Kuhn and A. Tucker (Eds.), *Contributions to the Theory of Games II*. Princeton: Princeton University Press.

McCutcheon, B. (1997). Do meetings in smoke filled rooms facilitate collusion? *Journal of Political Economy 105(2)*, 330–350.

Mengel, F. and R. Peeters (2011). Strategic behavior in repeated voluntary contribution experiments. *Journal of Public Economics 95(1)*, 143–148.

Potters, J. and S. Suetens (2009). Cooperation in experimental games of complements and substitutes. *Review of Economic Studies 76(3)*, 1125–1147.

Selten, R., M. Mitzkewitz, and G. Uhlich (1997). Duopoly strategies programmed by experienced traders. *Econometrica 65(3)*, 517–555.

SUPPLEMENTARY MATERIALS: FOR ONLINE PUBLICATION ONLY

Contents

List of Figures

List of Tables

A Stage game details

This section provides further details of the two differentiated-goods linear duopoly markets that underlie the two stage games implemented in the experiment. To provide a natural duopoly analogue for the strategic substitutes matrix, a discretized version of a discretized version of a differentiated-goods linear Cournot duopoly was used; for the strategic complements matrix, a discretized version of a differentiated-goods linear Bertrand duopoly. Note that, in order to ensure the incentives to cooperate are balanced across the games it was necessary to choose different demand systems under price competition and under quantity competition.

We started with the continuous strategy space version of the games and calibrated the parameters so that the payoffs from three key outcomes were constant across the two duopoly markets:[22] the Nash equilibrium payoffs, the joint payoff maximizing payoffs and the optimal deviation against the co-player playing the joint payoff maximizing action. In each matrix, action A corresponds to the joint profit maximizing quantity/price, and action C the Nash equilibrium quantity/price. In the substitutes game, action D corresponds to the optimal deviation to the other player choosing the joint payoff maximizing action, while in the complements game this is action B. To complete the action choices, action B in the substitutes game corresponds to the quantity in which, if both players chose this quantity, the payoff would be the same as the payoff in the Bertrand game when both players choose the optimal deviation price – i.e. the payoff when both players choose action B in the complements game. An analogous calculation is used to find the price that corresponds to action D in the complements game.

This calibration and selection procedure lead to the following duopoly games:

Strategic substitutes Demand given by $p_i = 53.1997 - q_i - 0.303175 \cdot q_{-i}$ and costs given by $c(q_i) = 500$. The four quantities used for actions A through D are $q^A = 20.41157$, $q^B = 22.35949$, $q^C = 23.09842$ and $q^D = 23.50571$. The left-hand matrix in Table A.1 is the resulting matrix game without rounded payoff numbers.

Strategic complements Demand given by $q_i = 11.4288 - p_i - 0.613878 \cdot p_{-i}$ and costs given by $c(q_i) = 8.6562 \cdot q_i$. The four quantities used for actions A through D are $p^A = 19.12757$, $p^B = 15.91350$, $p^C = 14.49007$ and $p^D = 13.47479$. The

[22]Since the payoffs were to be rounded for the implemented matrices, we only constrained the payoffs to be the same up to the nearest integer.

right-hand matrix in Figure A.1 is the resulting matrix game without rounded payoff numbers.

	A	B	C	D
A	42.94	30.89	26.32	23.80
B	51.20	38.00	32.99	30.23
C	52.35	38.71	33.54	30.68
D	52.52	38.64	33.37	30.47

Strategic substitutes.

	A	B	C	D
A	42.34	21.68	12.53	6.00
B	52.67	38.35	32.01	27.48
C	50.64	39.13	34.03	30.40
D	46.72	37.21	33.00	30.00

Strategic complements.

Figure A.1: The two stage games before rounding (payoffs shown are for the row player).

To implement the stage games in the laboratory all the payoffs were first rounded to the nearest integer. After rounding, some payoffs were increased or decreased by one unit in order to avoid degeneracies that are caused by rounding. This is done in such a way that games become even more similar: for instance, this led to the box formed by actions B and C and that formed by actions C and D being identical across games. The implemented stage games are repeated in Figure A.2 for convenience.

	A	B	C	D
A	43	31	25	23
B	51	36	32	29
C	52	40	33	31
D	54	38	32	30

Strategic substitutes.

	A	B	C	D
A	43	23	14	7
B	54	36	32	28
C	52	40	33	31
D	47	37	32	30

Strategic complements.

Figure A.2: The two stage games, after rounding, implemented in the experiment (payoffs shown are for the row player).

The crucial difference between the two games is the location of the optimal deviation against the co-player playing the joint payoff maximizing action, which is action B with strategic complements and action D with strategic substitutes. In games of strategic complements as my opponent "increases" her action, I would like to do the same. Consequently, the optimal deviation action is located between the collusive action (A) and the Nash action (C) in the complements game, whereas it is located beyond the Nash action in the substitutes game, where I would like to respond to an "increase" in my opponent's action by a "decrease" myself.[23] This difference in the location of these

[23]In continuous market games, the type of strategic interaction is determined by the second cross-derivative of player i's payoff function with respect to the actions of i and $-i$. This type is one of complements (substitutes) if this cross-derivative is positive (negative). In our discretized versions, the positive (negative) cross-derivative for complements (substitutes) is reflected in the (myopic) best response to the collusive action being "close to" ("far from") the collusive action itself.

actions is the primary difference between the games; a difference that will prove to have a significant interaction with the level of strategy revision opportunities. For convenience, we will refer to the actions A, B, C and D as respectively *Collusion*, *Dev.SC*, *Nash* and *Dev.SS*. Note that our games are designed such that, theoretically, collusion can be sustained as an equilibrium for both game types in all treatment variations. The necessary and sufficient conditions on discount factors (continuation probabilities) for trigger strategies to support collusion are identical.[24] This is a consequence of the restrictions imposed when designing the games, namely that the joint payoff maximizing payoffs, Nash payoffs and optimal deviation payoffs are the same in both games.

[24]The same is also true for other collusive strategies, such as tit-for-tat. While such strategies are not subgame pefect, they can be implemented without one-shot deviations or machine changes.

B Extending the Trigger Strategy to the Elicitation Setting

This section considers formally extending the grim trigger strategy (i.e. permanent reversion to Nash) to the initial-choice-plus-dynamic-response elicitation setting of the experiment. The purpose is to illustrate that eliciting dynamic responses, and including small costs for one-shot deviations from recommended action choices, does not fundamentally change this standard strategy – one that is commonly used in the theory of repeated games to show that cooperation can be sustained as part of a subgame-perfect equilibrium. Indeed, the minimum discount factor needed to sustain cooperation using this natural extension of the grim trigger does not change. Furthermore, allowing machine changes, as in the unilateral and bilateral variations, does not change this conclusion. The same extension of the grim trigger supports collusion with the same minimum discount factor as in both the baseline experiment and the standard repeated-game setting. Thus, machine changes in the unilateral and bilateral treatments are not necessary to implement collusive strategies.

B.1 Extending the Set-up to the Elicitation Setting

In the baseline experiment, subjects are asked in the first round to choose an initial action and a dynamic response vector, where the latter determines their recommended action in all future rounds as a function of their partner's choice in the previous round. In all rounds after the first one, subjects must choose an action for that round with all choices that do not correspond to their recommended action incurring a cost of 3 ECUs. Consequently, their action sets are

$$A^1 = \{A, B, C, D\}^5$$

in the first round and

$$A^t = \{A, B, C, D\}$$

for all $t > 1$.[25]

In the repeated game, the history at round $t \geq 1$ is a list of all the choice pairs, $\left(a_i^s, a_{-i}^s\right)$, that player i and their match, player $-i$, have made for all rounds $s < t$, with the understanding that all histories at round 1 are empty (null). Consequently,

[25] In an abuse of notation, $a^1 \in A^1$ will be used to denote the action choice alone, as well as the complete action choice plus dynamic response vector as defined above, as long as there is no ambiguity in doing so.

a strategy in the repeated game must specify an initial choice and dynamic response vector, plus an action choice for any possible history with $t > 1$.

Let r_i^t be the recommended strategy for player i in round t. Then the payoff from period t is given by the stage game outcome minus any one-shot deviation costs:

$$\pi_i^t = \pi_i\left(a_i^t, a_{-i}^t\right) - c_{one} \cdot \mathbb{I}\left(a_i^t \neq r_i^t\right)$$

where $\mathbb{I}\left(a_i^t \neq r_i^t\right)$ takes value one if the action choice is different to the recommended choice, and zero otherwise. In the experiment $c_{one} = 3$. Payoffs from the repeated game are evaluated at time $t \geq 1$ using the discounted sum

$$\Pi_i^t = \sum_{s=0}^{\infty} \delta^s \pi_i^{t+s}$$

for $\delta \in [0, 1)$.

The natural way to define the grim trigger strategy in this setting is as follows: S^{GT} is the strategy that specifies

$$a_i^t = \begin{cases} (A - ACCC) & \text{, if } t = 1 \\ A & \text{, if } t > 1 \text{ and } \left(a_i^s, a_{-i}^s\right) = (A, A)\, \forall s \leq t \\ C & \text{, otherwise} \end{cases}$$

Note that the last row implies that player i will pay c_{one} in the event that the recommended strategy is something other than C. This situation will only arise if either player chose something other than A in a round before the previous round, but the other player chose A in the previous round. The strategy says that the player i will pay c_{one} to play C instead of the recommended A in such a circumstance.

B.2 Incentive Compatibility in the Baseline Treatment

Next we check whether the pair $\left(S^{GT}, S^{GT}\right)$ can form a sub-game perfect Nash equilibrium of the repeated game. In doing so, we will demonstrate that the minimum discount rate for cooperation to be sustainable using the grim trigger strategy is the same for the initial choice plus dynamic response game as it is for the standard repeated game.

Suppose player $-i$ is playing according to S^{GT}. Then we will show that playing S^{GT} is a best response for player i using the standard procedure of checking single deviations. Given the nature of the elicitation setting, incentive compatibility in round one needs to be checked separately from incentive compatibility in later rounds. This is because, the optimal deviation in round one can involve not just changing the current choice, but also changing the dynamic response vector. For later rounds, the only difference to the standard setting is that the specified choice, or the optimal deviation, might involve paying the small one-shot deviation cost of going against the recommended strategy.

<u>Step 1</u>: Incentive compatibility in round one ($t = 1$). Given $-i$ plays according to S^{GT}, the optimal deviation is to choose $D - CCCC$ in the substitutes game, and $B - CCCC$ in the complements game. This gives a deviation payoff of

$$\begin{aligned}\Pi_i^1 (dev) &= \pi^{dev} + \delta \left(\pi^{Nash} + \delta\pi^{Nash} + ...\right) \\ &= \pi^{dev} + \left(\frac{\delta}{1-\delta}\right)\pi^{Nash}\end{aligned}$$

Continuing with the strategy S^{GT} gives

$$\begin{aligned}\Pi_i^1 \left(S^{GT}\right) &= \pi^{JPM} + \delta \left(\pi^{JPM} + \delta\pi^{JPM} + ...\right) \\ &= \pi^{JPM} + \left(\frac{\delta}{1-\delta}\right)\pi^{JPM}\end{aligned}$$

Consequently, for $\left(S^{GT}, S^{GT}\right)$ to be part of an sub-game perfect Nash equilibrium it must be that

$$\delta \left(\pi^{JPM} - \pi^{Nash}\right) \geq (1-\delta)\left(\pi^{dev} - \pi^{JPM}\right)$$

Solving for the smallest δ that just makes the above inequality hold with equality is the same calculation for the minimum discount rate – denoted δ_{min} – as in the usual repeated game environment.

<u>Step 2</u>: Incentive compatibility in later rounds ($t > 1$). This step is split into two cases that essentially correspond to a reward path and a punishment path.

- Reward path: Suppose the history is such that $\left(a_i^s, a_{-i}^s\right) = (A, A)$ for all $s < t$. Then the recommendation for both players will be $r_i^t = r_{-i}^t = A$. The optimal deviation for player i is to pay c_{one} and play D in the substitutes game and B in the complements game. Thus, for incentive compatibility it must be that

$$\begin{aligned}\pi^{dev} + \delta \left(\pi^{Nash} + \delta\pi^{Nash} + ...\right) - c_{one} \quad \leq \\ \pi^{JPM} + \delta \left(\pi^{JPM} + \delta\pi^{JPM} + ...\right)\end{aligned}$$

  That is, it must be that

$$\delta \left(\pi^{JPM} - \pi^{Nash}\right) \geq (1-\delta)\left(\pi^{dev} - c_{one} - \pi^{JPM}\right)$$

  Given the additional cost c_{one}, this inequality holds strictly for any $\delta \geq \delta_{min}$.

- Punishment path: Suppose the history is such that $\left(a_i^s, a_{-i}^s\right) \neq (A, A)$ for some $s < t$. Here we need to check that it is incentive compatible for player i to play C. The punishment path has two sub-cases to consider that depend on whether the recommended action is also to play C or not. In either case, player i knows that $a_{-i}^t = C$ and $\left(a_i^{t'}, a_{-i}^{t'}\right) = (C, C)$ for $t' > t$, since player $-i$ is following S^{GT} and player i is only considering a single deviation today from S^{GT}.

  1. Suppose $r_i^t = C$. Given that C is the best response to C in the one-shot game and there is no possibility to change the future outcomes to anything other than repeated play of the one-shot Nash equilibrium, $a_i^t = C$ is player i's best-response for any δ.

  2. Suppose $r_i^t \neq C$. Here we are essentially checking whether it is worth paying c_{one} to play C, as required by S^{GT}:

  $$\pi_i\left(r_i^t, C\right) + \delta\left(\pi^{Nash} + \delta\pi^{Nash} + ...\right) \leq \pi^{Nash} + \delta\left(\pi^{Nash} + \delta\pi^{Nash} + ...\right) - c_{one}$$

  That is, it must be that

  $$c_{one} \leq \pi^{Nash} - \pi_i\left(r_i^t, C\right)$$

  Given that the only recommendation under S^{GT} other than C is A, this holds for any δ since $\pi_i(A, C) - \pi^{Nash} > c_{one}$.

In summary, moving from the standard repeated-game set up to the initial-choice-plus-dynamic-response set up introduces two changes into the incentive compatibility check:

1. Incentive compatibility needs to be checked separately for the initial round and for later rounds.

2. On the punishment path, players must be willing to play the Nash action even if they need to pay the one-shot deviation cost to do so.

The former has no implications for the minimal discount factor needed to sustain cooperation using the grim trigger. The latter introduces an additional condition that requires the one-shot deviation cost be not too large; a condition that is unrelated to the discount factor and is met in the stage games we implement in the experiments.

B.3 Accommodating Dynamic Response Changes

In the unilateral and bilateral treatments, subjects (may) have the opportunity to change their dynamic response during a match. Such changes take effect from the next period onwards (via potentially giving a different recommended action) and come at a small cost, c_{mach}. Consequently, to extend the above analysis to include this case, it is necessary to include in the strategy what dynamic response the player will specify, if they have the opportunity to change their machine, and whether they would permit the other to change their machine in the bilateral treatment. The per period payoff function becomes:

$$\pi_i^t = \pi_i\left(a_i^t, a_{-i}^t\right) - c_{one} \cdot \mathbb{I}\left(a_i^t \neq r_i^t\right) - c_{mach} \cdot \mathbb{I}\left(machine_i^{t+1} \neq machine_i^t\right)$$

The grim trigger strategy is extended to the unilateral and bilateral settings as follows:

$$a_i^t = \begin{cases} A - ACCC & \text{if } t = 1 \\ A - ACCC & \text{if } t > 1 \text{ and } \left(a_i^s, a_{-i}^s\right) = (A, A)\, \forall s \leq t \\ C - ACCC & \text{otherwise} \end{cases}$$

Note that this strategy implies that the player will not implement a machine change even if given the opportunity to do so. Furthermore, the player is indifferent between allowing or not allowing the other player to make a strategy change in the bilateral treatment.

Checking incentive compatibility for the above versions of S^{GT} follows very similar lines as that shown for the baseline case. As before incentive compatibility should be checked separately for the initial choices, where there are no one-shot or machine-change costs, and for all later periods. Again, it will be the initial choices that provide the binding constraints that define δ_{min}.

The following observations ensure that the basic logic from the baseline case carries over to the unilateral and bilateral cases:

- On the initial or reward path, if $\delta \geq \delta_{min}$ then $\left(\pi^{JPM}, \pi^{JPM}, ...\right)$ is preferred $\left(\pi^{dev}, \pi^{Nash}, ...\right)$. This is the case both today, when considering a $t = 1$ deviation in action or paying c_{one} to deviate in action from a recommendation for $t > 1$, and tomorrow, when considering a $t = 1$ deviation in the dynamic response or paying c_{mach} to deviate in dynamic response at $t > 1$. Furthermore, this holds for any cost of one-shot or machine-change cost, as long as the costs are greater than or equal to zero.

- On the punishment path, if the other player is playing the S^{GT} strategy and we are only considering deviations today from the S^{GT} strategy, then there is no reason to pay $c_{mach} > 0$ to switch the dynamic response for tomorrow from $ACCC$ to $CCCC$ if the action C is being played today. This is because the S^{GT} strategies from tomorrow onwards will ensure that C is played in all future periods. Furthermore, as long as $c_{one} \leq \pi^{Nash} - \pi_i(r_i^t, C)$ when $r_i^t \neq C$, it is preferable to pay the one-shot cost today to avoid the sucker payment.

In summary, essentially the same grim trigger strategy can be used to support collusion in the unilateral and bilateral treatments, using the same minimum discount rate and without the need to change the dynamic response during a match.

C Example instructions: Baseline – strategic complements

Part 1

Welcome!

You are about to participate in a session on interactive decision-making. Thank you for agreeing to take part. The session should last 90 to 120 minutes.

You should have already turned off all mobile phones, smart phones, mp3 players and all such devices by now. If not, please do so immediately. These devices must remain switched off throughout the session. Place them in your bag or on the floor besides you. Do not have them in your pocket or on the table in front of you.

The entire session, including all interaction between you and other participants, will take place through the computer. You are not allowed to talk or to communicate with other participants in any other way during the session.

You are asked to abide by these rules throughout the session. Should you fail to do so, we will have to exclude you from this (and future) session(s) and you will not receive any compensation for this session.

We will start with a brief instruction period. Please read these instructions carefully. They are identical for all participants in this session with whom you will interact. If you have any questions about these instructions or at any other time during the experiment, then please raise your hand. One of the experimenters will come to answer your question.

Compensation for participation in this session

In addition to the 3 participation fee, what you will earn from this session will depend on your decisions, the decisions of others and chance. In the instructions and all decision tasks that follow, payoffs are reported in Experimental Currency Units (ECUs). At the end of the experiment, the total amount you have earned will be converted into Euros using the following conversion rate:

1 ECU = 4 Eurocents.

The payment takes place in cash at the end of the experiment. Your decisions in the experiment will remain anonymous.

General instructions

The session is structured as follows:

1. This session consists of 10 *matches*. At the beginning of each match, you will be randomly paired with another participant.

2. During the match, you will interact repeatedly with this same participant for a number of *rounds*.

3. The number of rounds is randomly determined. After each round, there is an 87.5% chance that the match will continue for at least another round. This is as if we were to roll an 8-sided die and end if the number 1 came up and continue if 2 through 8 came up. Notice that, if you are in round 2, the probability that there will be a third round is 87.5% and if you are in round 9, the probability that there will be a tenth round is also 87.5%. That is, at any point in the match, the probability that there will be at least one more round is 87.5%. This means that, in expectation, another 8 rounds will follow, irrespective of the number of rounds you have just completed.

4. Once a match ends, you will be matched with a randomly drawn participant for the next match.

Description of a match

5. During a match you will repeatedly interact with the same participant for a number of rounds. Each round consists of the same decision situation.

6. In this decision situation, you will be asked to choose an action. There are four possible actions: A, B, C or D. The participant you are matched with will also be asked to choose an action. The set of possible actions to choose from is the same for both of you.

7. Your payoff for the round depends on your action and the action of the participant you are matched with. For each possible combination of actions, the table below displays the payoffs for you and the other participant. The rows, which correspond to your action, are labeled in capital letters; the columns, which correspond to the other's action, are labeled in lower-case letters. In each cell your payoff is first (in the darker font) and the other participant's payoff is second (in

the lighter font). For example, if your action is B and the other participant's action is c, your payoff is 32 ECU and the other participant's payoff is 40 ECU.

		Other's action			
		a	b	c	d
Your action	A	43, 43	23, 54	14, 52	7, 47
	B	54, 23	36, 36	32, 40	28, 37
	C	52, 14	40, 32	33, 33	31, 32
	D	47, 7	37, 28	32, 31	30, 30

8. To summarize, in a match you interact **repeatedly** with the **same participant** for an unknown number of rounds in the decision situation described above. As described in point 3 above, after every round, there is a 87.5% chance of another round in this match.

Your decisions (How actions are chosen)

At the beginning of a match

9. At the very beginning of every match, you will be asked to specify your *initial action* and to provide a *plan of intended actions*. The initial action is the action you choose in the first round of this match. The plan of intended actions determines for each subsequent round which action you intend to choose in response to each possible action choice of the other participant in the previous round.

10. The table below presents an example of a plan of intended actions, as it will be visualized on your screen. In this example, the plan prescribes you to take action D in all rounds immediately following one in which the other participant has taken action a (action D is checked in column a). In periods immediately following one in which the other participant has chosen action b, the plan prescribes you to take action B (action B is checked in column b) and so forth.

Your plan (example)	a	b	c	d
	A ○	A ○	A ○	A ○
	B ○	B ●	B ○	B ○
	C ○	C ○	C ○	C ●
	D ●	D ○	D ●	D ○

Notice that the table above is just one example of a plan. In the experiment you will be asked to design your own plan.

11. Since it will be costly (see point 15 below) to choose an action different from the one prescribed by your intended plan of action, you are advised to think carefully about how to design your plan.

12. Once you and the participant you are matched with have made your choice of initial action and plan of intended actions, the first round of the sequence of decision situations described above will begin.

During round 1

13. In the first round, your action choice will be the initial action you just chose.

During later rounds

14. At the beginning of any subsequent round you will be told the prescribed action from your plan of intended actions.

15. You will then be asked to choose your action for the current round. It is possible to choose an action different from the one prescribed by your plan of intended actions. However, doing so will cost 3 ECU. Note also that you will need to select this action and click on the "OK" button within the time limit shown on your screen; otherwise your prescribed action will be chosen.

At the end of each round

16. At the end of each round you will receive feedback on your action chosen, the action chosen by the other participant and your payoffs as well as about any costs incurred for deviating from the plan of intended actions.

The end of the session

17. After a match is finished, you will be randomly paired for a new match. This session consists of 10 such matches.

18. In each of the 10 matches, your payoff starts at 0 and from there accumulates until the end of your match. At the end of the session – after the tenth match – one match will be selected at random. The payoff you gained during the selected match will be used to calculate your final payoff.

Control Questions

Please read through the following and answer the questions. When you have finished answering these questions, please raise your hand.

Assume you specified action A as initial action and the following plan of intended actions:

Your plan	a	b	c	d
	A ●	A ○	A ●	A ○
	B ○	B ●	B ○	B ●
	C ○	C ○	C ○	C ○
	D ○	D ○	D ○	D ○

Suppose that the other participant chooses action b in the first round.

1. What is your payoff in the first round?
2. What is the other participant's payoff in the first round?
3. Which action does your plan prescribe you to choose in the second round?

Assume that you choose the prescribed action in the second round. Suppose that the other participant chooses action d in the second round.

4. What is your payoff in the second round?
5. Which action will you be prescribed to choose in the third round?

True or False?

Please answer whether the following statements are true or false:

6. The longer a match has been going on the more likely it is to end.
7. Each round I can choose the action I want.
8. I can modify my plan of intended action after each round within a match.
9. I am matched with the same participant during the entire session.
10. I am matched with the same participant during each match.

Part 2

Control Questions – Answers

1. In the first round, if you choose A and the other participant chooses b, then your payoff is 23.

2. In the first round, if you choose A and the other participant chooses b, then the other participant's payoff is 54.

3. The other participant chose b in the first round. Reading column b of your plan gives you a prescribed action of B.

4. In the second round, if you choose B and the other participant chooses d, then your payoff is 28.

5. The other participant chose d in the second round. Reading column d of your plan gives you a prescribed action of B.

True or False? – Answers

6. False: at any point in the match, the probability that there will be at least one more round is 87.5%.

7. True: In any round, you can choose the action you would like. In particular, it is possible to choose an action different from the one prescribed by your plan of intended actions. However, doing so will cost 3 ECU.

8. False: you can only modify your plan of intended actions at the beginning of a match.

9. False: once a match ends, you will be matched with a randomly drawn participant for the next match.

10. True: in a match you interact repeatedly with the same participant

Summary

1. At the beginning of a match, you choose your initial action and your plan of intended actions.

2. Every round (except for the first round) your plan prescribes an action that depends on the action of the other participant in the previous round.

3. In any round (except for the first round), you can either choose the prescribed action or choose another action. Choosing an action which is different from your prescribed action has a cost of 3 ECU.

4. The length of a match is randomly determined. After each round, there is an 87.5 % chance that the match will continue for at least one more round. You will play with the same person for the entire match.

5. After a match is finished, you will be randomly paired for a new match. This session consists of 10 such matches.

D Controlling for path dependency

In order to isolate the impact of the responsive part of the strategies from initial choices, we compute the invariant distribution over actions of the Markov chain as specified by the dynamic responses. The invariant distribution tells us with which probabilities the actions will be chosen in the long run if play continues as it is observed in our experimental sessions, regardless of the initial choices. We compute these invariant distributions using the transition probabilities as defined by (1) the average initial machine, (2) the average over chosen actions in response to rival's actions, and (3) the average chosen actions in response to outcomes. Note that by taking averages over the average behavior of an individual in a match (of the last four matches only), we disregard the heterogeneity that is clearly present in our data (which is more pronounced between individuals than between matches). Although, in principle it is better to properly control for heterogeneity in behavior, it is certainly computationally more involved and, we believe, not needed for our illustrative purposes. The invariant distributions obtained for the different treatments using the three different specifications of the Markov chain are presented in Table D.1.

Table D.1: Invariant distribution.

		Substitutes			Complements		
Transitions	Action	Base	Uni	Bil	Base	Uni	Bil
Initial Machines	A	0.0372	0.0916	0.0630	0.0086	0.1416	0.0153
	B	0.0760	0.0523	0.0709	0.0633	0.1146	0.0701
	C	0.7834	0.6889	0.6414	0.8769	0.6531	0.8410
	D	0.1034	0.1672	0.2247	0.0511	0.0907	0.0736
Realized Actions (actions)	A	0.1252	0.1430	0.1210	0.0561	0.0989	0.0528
	B	0.0997	0.0543	0.0560	0.1105	0.1062	0.1021
	C	0.6361	0.6235	0.5879	0.7735	0.6691	0.7703
	D	0.1390	0.1792	0.2351	0.0599	0.1258	0.0748
Realized Actions (outcomes)	A	0.1258	0.1095	0.1076	0.0473	0.0552	0.0471
	B	0.0968	0.0585	0.0630	0.0887	0.0948	0.0646
	C	0.6580	0.6826	0.6130	0.7881	0.7006	0.7669
	D	0.1194	0.1494	0.2163	0.0759	0.1494	0.1215

The regressions presented in Tables 3 and E.4 (see also Figures 2(a) and 3(a)) show that, under strategic complements, a higher level of collusion is obtained when unilateral modifications of the dynamic response are allowed. The invariant distributions show a consistent picture and that this effect is not only caused by initial choices (see also Figures 2(b) and 3(b)).

Apart from Complements–Unilateral, the weight on the top two actions (A and B) is larger and that on the lower two actions (C and D) is smaller using the realized actions than using the initial machines. This observation indicates that one-shot deviations and machine modifications are mainly used to escape profit eroding states in favor for collusive ones.

Table D.2 presents the persistence of the collusive outcome, the Nash outcome, the top two actions and the lower two actions. For the single state outcomes the persistence is the probability that the state is not left in the next stage given that it is reached. For the sets of states these probabilities are weighted by the invariant distribution of the Markov process over outcomes (to properly account for all states within a set not being more likely being reached). Again, we derive the probabilities in the three alternative ways as before.

Table D.2: Direct recurring transitions.

		Substitutes			Complements		
Transitions	Outcome	Base	Uni	Bil	Base	Uni	Bil
Initial Machines	(A,A)	0.1567	0.1218	0.2296	0.1111	0.2659	0.0913
	(C,C)	0.6944	0.5703	0.5393	0.7932	0.6104	0.7565
	(AB,AB)	0.1106	0.1258	0.1690	0.0782	0.4036	0.1529
	(CD,CD)	0.8369	0.7948	0.8262	0.8914	0.7645	0.8895
Realized Actions (actions)	(A,A)	0.2004	0.2440	0.1757	0.1494	0.3285	0.1715
	(C,C)	0.5339	0.5405	0.5133	0.6934	0.5989	0.6980
	(AB,AB)	0.1619	0.2217	0.1854	0.2014	0.3008	0.2891
	(CD,CD)	0.6834	0.7568	0.7701	0.7918	0.7806	0.8377
Realized Actions (outcomes)	(A,A)	0.8422	0.7204	0.5814	0.7008	0.8580	0.8789
	(C,C)	0.6809	0.7514	0.5993	0.7987	0.7022	0.8010
	(AB,AB)	0.7071	0.5950	0.5530	0.4774	0.6647	0.7321
	(CD,CD)	0.7924	0.8670	0.8062	0.8763	0.8427	0.9207

It is apparent that the persistence probabilities of the collusive outcome are substantially larger when applying the data of how outcomes are translated into actions. This shows us that one-shot deviations and dynamic response modifications are used to sustain cooperation. Application of outcome to action data also increases the persistence probabilities of the static Nash outcome when dynamic responses can be revised, but not when only one-shot deviations are possible. This indicates that, in a setting with more revision opportunities, and hence better possibilities to renegotiate, punishments are sustained by means of one-shot deviations and dynamic response modifications. If anything, this provides evidence against the idea of renegotiation playing an important role.

E Additional material

E.1 Figures

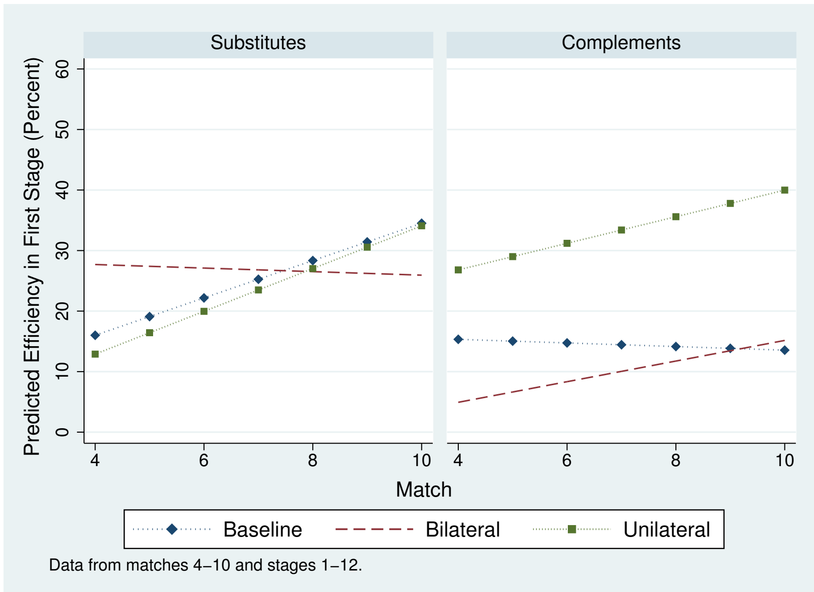

(a) Efficiency across matches.

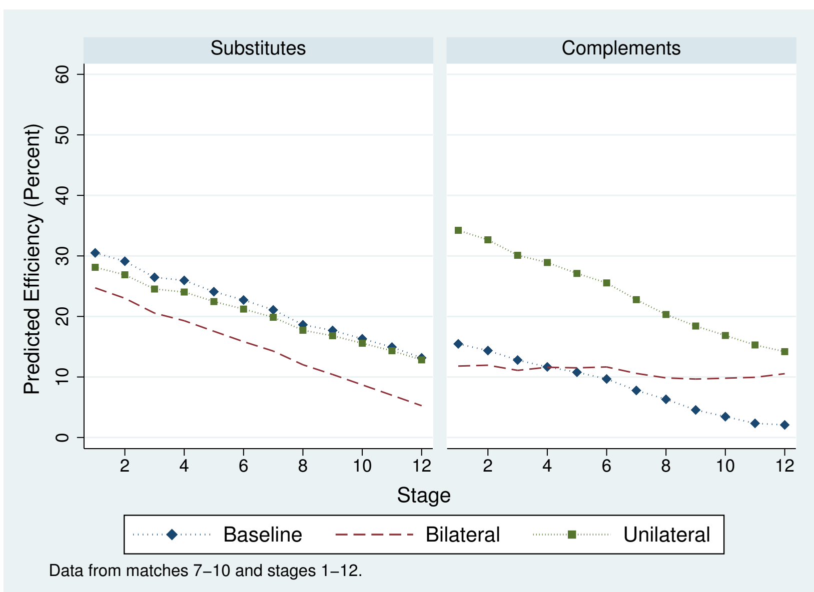

(b) Efficiency within matches.

Figure E.1: The effect of strategy revision opportunities on efficiency with strategic complements and substitutes: predicted values from a linear regression.

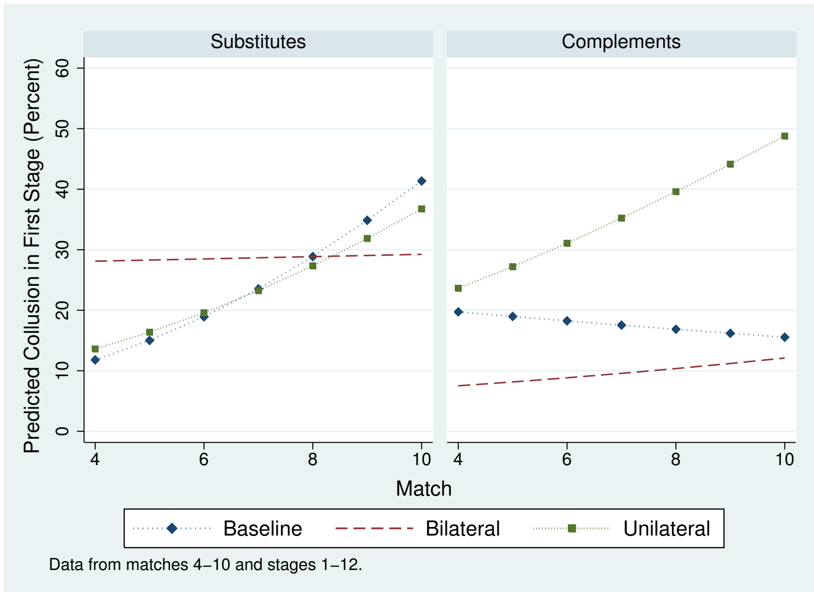

(a) Collusion choice across matches.

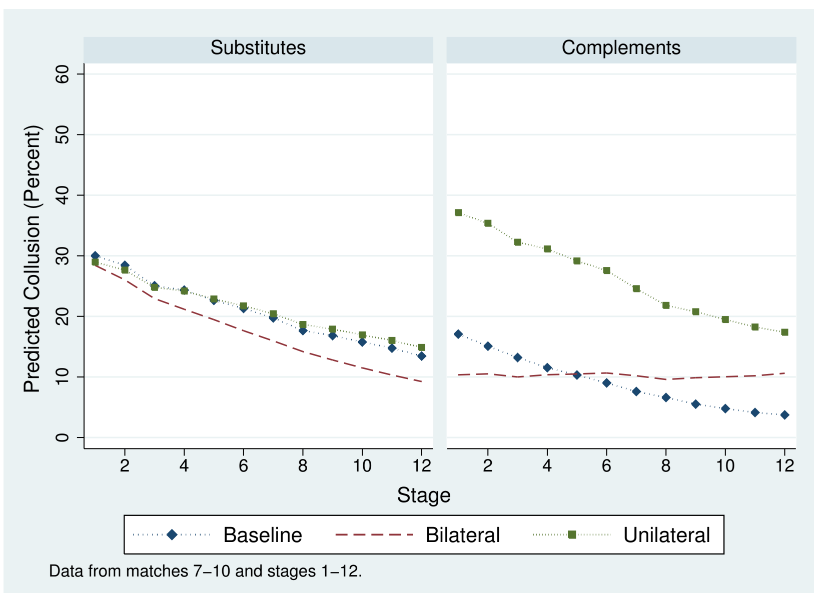

(b) Collusion choice within matches.

Figure E.2: The effect of strategy revision opportunities on collusion choice with strategic complements and substitutes: predicted values from a logit regression.

E.2 Tables

Table E.1: Robustness check on the inclusion of subject random-effects: Linear regression of payoff efficiency in the stage game.

	Substitutes				Complements			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Bilateral	−0.06	−0.02	0.08	0.02	−0.05	−0.06	−0.18*	−0.13
	(0.477)	(0.791)	(0.381)	(0.728)	(0.392)	(0.377)	(0.063)	(0.104)
Unilateral	−0.03	−0.02	0.02	−0.03	0.19***	0.25***	0.06	0.08
	(0.745)	(0.861)	(0.837)	(0.749)	(0.003)	(0.000)	(0.294)	(0.260)
Stage	−0.02***	−0.01***	−0.01***	−0.01***	−0.01**	−0.01**	−0.01**	−0.01**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.040)	(0.040)	(0.044)	(0.049)
Bilateral x Stage	−0.00	−0.00	−0.00	−0.00	0.01	0.01	0.01	0.01
	(0.765)	(0.896)	(0.551)	(0.778)	(0.148)	(0.228)	(0.105)	(0.194)
Unilateral x Stage	0.00	0.00	0.00	−0.00	−0.01	−0.01	−0.00	−0.01
	(0.730)	(0.980)	(0.804)	(0.954)	(0.389)	(0.103)	(0.517)	(0.239)
Match			0.06***	0.05*			−0.00	0.01
			(0.001)	(0.097)			(0.997)	(0.398)
Bilateral x Match			−0.05*	−0.02			0.05**	0.03
			(0.061)	(0.622)			(0.023)	(0.162)
Unilateral x Match			−0.02	0.01			0.05**	0.07**
			(0.528)	(0.711)			(0.010)	(0.014)
Constant	0.36***	0.35***	0.21***	0.23***	0.11**	0.12**	0.12***	0.09
	(0.000)	(0.000)	(0.000)	(0.000)	(0.012)	(0.011)	(0.009)	(0.134)

Notes: The baseline case is the baseline treatment. Specifications 2 and 4 are the same as the reported specifications, 1 and 3, but include subject random effects (linear two-way crossed effects model). *p*-values are reported in parentheses. ***1%, **5%, *10% significance.

Table E.2: Robustness check on the inclusion of match-stage composition dummies: Linear regression of payoff efficiency in the stage game.

	Substitutes				Complements			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Bilateral	−0.06	−0.06	0.08	0.08	−0.05	−0.05	−0.18*	−0.18*
	(0.477)	(0.538)	(0.381)	(0.410)	(0.392)	(0.447)	(0.063)	(0.088)
Unilateral	−0.03	−0.03	0.02	0.02	0.19***	0.19**	0.06	0.06
	(0.745)	(0.750)	(0.837)	(0.838)	(0.003)	(0.020)	(0.294)	(0.593)
Stage	−0.02***	−0.02***	−0.01***	−0.01***	−0.01**	−0.01**	−0.01**	−0.01**
	(0.000)	(0.002)	(0.000)	(0.001)	(0.040)	(0.027)	(0.044)	(0.034)
Bilateral x Stage	−0.00	−0.00	−0.00	−0.00	0.01	0.01	0.01	0.01
	(0.765)	(0.786)	(0.551)	(0.599)	(0.148)	(0.154)	(0.105)	(0.126)
Unilateral x Stage	0.00	0.00	0.00	0.00	−0.01	−0.01	−0.00	−0.00
	(0.730)	(0.740)	(0.804)	(0.791)	(0.389)	(0.400)	(0.517)	(0.559)
Match			0.06***	0.05**			−0.00	−0.01
			(0.001)	(0.013)			(0.997)	(0.345)
Bilateral x Match			−0.05*	−0.05*			0.05**	0.05**
			(0.061)	(0.084)			(0.023)	(0.011)
Unilateral x Match			−0.02	−0.02			0.05**	0.05*
			(0.528)	(0.573)			(0.010)	(0.082)
Constant	0.36***	0.32***	0.21***	0.18***	0.11**	0.17***	0.12***	0.20**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.012)	(0.006)	(0.009)	(0.029)
Matches	7–10	7–10	7–10	7–10	7–10	7–10	7–10	7–10
Stages	1–12	1–12	1–12	1–12	1–12	1–12	1–12	1–12
M-S Dummy	Yes	No	Yes	No	Yes	No	Yes	No
Cluster VCE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The baseline case is the baseline treatment. Specifications 2 and 4 are the same as the reported specifications, 1 and 3, but without the match-stage composition dummies. *p*-values are reported in parentheses. ***1%, **5%, *10% significance.

Table E.3: Robustness check on the data sub-sample: Linear regression of payoff efficiency in the stage game.

	Substitutes			Complements		
	(1)	(2)	(3)	(1)	(2)	(3)
Bilateral	0.08	0.16*	0.02	−0.18*	−0.13*	−0.06
	(0.381)	(0.076)	(0.788)	(0.063)	(0.067)	(0.541)
Unilateral	0.02	−0.03	−0.12	0.06	0.10	0.28**
	(0.837)	(0.605)	(0.214)	(0.294)	(0.289)	(0.014)
Stage	−0.01***	−0.01***	−0.01***	−0.01**	−0.01*	−0.00
	(0.000)	(0.002)	(0.003)	(0.044)	(0.057)	(0.978)
Bilateral x Stage	−0.00	−0.01	−0.01**	0.01	0.01	0.00
	(0.551)	(0.102)	(0.022)	(0.105)	(0.156)	(0.672)
Unilateral x Stage	0.00	−0.00	−0.01	−0.00	−0.01**	−0.03***
	(0.804)	(0.889)	(0.474)	(0.517)	(0.021)	(0.000)
Match	0.06***	0.03	−0.02	−0.00	−0.00	0.02
	(0.001)	(0.102)	(0.268)	(0.997)	(0.727)	(0.611)
Bilateral x Match	−0.05*	−0.03	0.06	0.05**	0.02*	0.00
	(0.061)	(0.121)	(0.120)	(0.023)	(0.056)	(0.972)
Unilateral x Match	−0.02	0.00	0.06	0.05**	0.02	−0.03
	(0.528)	(0.827)	(0.145)	(0.010)	(0.173)	(0.582)
Constant	0.21***	0.18**	0.26***	0.12***	0.11**	0.04
	(0.000)	(0.013)	(0.002)	(0.009)	(0.042)	(0.610)
Matches	7–10	4–10	4–6	7–10	4–10	4–6
Stages	1–12	1–12	1–12	1–12	1–12	1–12
M-S Dummy	Yes	Yes	Yes	Yes	Yes	Yes
Cluster VCE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The baseline case is the baseline treatment. Specifications 2 and 3 are the same as the reported specification, 1, except use the final two-thirds and the middle third sub-samples. VCE clustered at the matching-group level. p-values are reported in parentheses. ***1%, **5% ,*10% significance.

Table E.4: Logit regression of the probability of choosing the collusive action.

	Substitutes				Complements			
	(1)		(2)		(1)		(2)	
Bilateral	−0.05	(0.879)	0.99**	(0.038)	−0.76	(0.241)	−1.91	(0.109)
Unilateral	−0.07	(0.873)	0.24	(0.670)	1.05**	(0.037)	−0.28	(0.670)
Stage	−0.09***	(0.000)	−0.08***	(0.000)	−0.15*	(0.087)	−0.15*	(0.098)
Bilateral x Stage	−0.03	(0.233)	−0.04*	(0.070)	0.16	(0.115)	0.17	(0.109)
Unilateral x Stage	0.02	(0.648)	0.01	(0.718)	0.05	(0.557)	0.07	(0.470)
Match			0.35***	(0.000)			−0.10	(0.361)
Bilateral x Match			−0.40**	(0.021)			0.44*	(0.093)
Unilateral x Match			−0.12	(0.435)			0.52***	(0.000)
Constant	−0.60***	(0.002)	−1.45***	(0.000)	−2.25***	(0.000)	−1.92***	(0.003)

Notes: The baseline case is the baseline treatment. All regressions use data from matches 7–10 and stages 1–12 and include match-stage composition dummies. VCE clustered at the matching-group level. *p*-values are reported in parentheses. ***1%, **5%, *10% significance.

Table E.5: Linear regression of payoff efficiency in the stage game – strategic substitutes versus complements.

	Combined regression			
	(1)		(2)	
Complements	−0.15**	(0.044)	0.02	(0.791)
Bilateral	−0.06	(0.475)	0.08	(0.326)
Unilateral	−0.03	(0.765)	0.02	(0.858)
Complements x Bilateral	0.01	(0.927)	−0.27**	(0.036)
Complements x Unilateral	0.22**	(0.040)	0.04	(0.770)
Stage	−0.02***	(0.000)	−0.01***	(0.000)
Complements x Stage	0.00	(0.668)	0.00	(0.885)
Bilateral x Stage	−0.00	(0.773)	−0.00	(0.555)
Unilateral x Stage	0.00	(0.738)	0.00	(0.809)
Complements x Bilateral x Stage	0.01	(0.188)	0.02*	(0.092)
Complements x Unilateral x Stage	−0.01	(0.360)	−0.01	(0.503)
Match			0.07***	(0.001)
Complements x Match			−0.07***	(0.005)
Bilateral x Match			−0.05**	(0.046)
Unilateral x Match			−0.02	(0.593)
Complements x Bilateral x Match			0.11***	(0.002)
Complements x Unilateral x Match			0.07*	(0.078)
Constant	0.31***	(0.000)	0.16***	(0.007)

Notes: The baseline case is the baseline treatment with strategic substitutes. All regressions use data from matches 7–10 and stages 1–12 and include match-stage composition dummies. VCE clustered at the matching-group level. *p*-values are reported in parentheses. ***1%, **5%, *10% significance. The table reports just the results for regressors involving the complements dummy variable.

Table E.6: Logit regression of the probability of choosing the collusive action – strategic complements versus substitutes.

	Combined regression			
	(1)		(2)	
Complements	−0.66	(0.186)	0.49	(0.508)
Complements x Bilateral	−0.70	(0.320)	−2.86**	(0.024)
Complements x Unilateral	1.08*	(0.087)	−0.52	(0.568)
Complements x Stage	−0.06	(0.484)	−0.08	(0.405)
Complements x Bilateral x Stage	0.18*	(0.064)	0.21**	(0.047)
Complements x Unilateral x Stage	0.04	(0.681)	0.06	(0.563)
Complements x Match			−0.45***	(0.008)
Complements x Bilateral x Match			0.83***	(0.004)
Complements x Unilateral x Match			0.62***	(0.007)
Constant	−0.96***	(0.001)	−1.80***	(0.000)

Notes: The baseline case is the baseline treatment with strategic substitutes. All regressions use data from matches 7–10 and stages 1–12 and include match-stage composition dummies. VCE clustered at the matching-group level. *p*-values are reported in parentheses. ***1%, **5%, *10% significance. The table reports just the results for regressors involving the complements dummy variable plus the constant term.

Table E.7: Machine categorization (in percent): decomposing the "Other" category.

		Substitutes			Complements		
	Machine	Base	Uni	Bil	Base	Uni	Bil
Uncond. coop.	AAAA	1	5	1	1	8	1
Cond. coop.	ABC(C/D)	19	13	16	11	23	18
Nash reversion	ACCC	15	13	22	20	14	8
Part. coll. + Nash rev.	BBCC	1	2	3	8	4	8
Static Nash	CCCC	20	22	10	25	14	31
Myopic best repsonse	DCCC	28	28	19			
	BCCC				29	21	25
Punishing	DDDD	0	3	8	0	0	0
Other	———	14	13	19	6	13	9

Notes: Distribution of machines in initial stages of matches 7–10 across treatments. For each category mentioned, machines (i) with Hamming distance at most 1 from the machine mentioned, (ii) that are not named explicitly in the second column (e.g. CCCC does *not* count towards DCCC) and (iii) do respond to collusion with the same action as the machine mentioned (e.g. ADDD does *not* count towards DDDD) are counted. Each machine is counted once and in case multiple categories apply, they are counted with equal weight in these categories. The category "Other" includes all machines satisfying none of the properties above.

Table E.8: When and how machines are modified.

	Previous stage outcome: Label	Previous stage outcome: Outcomes	Machine modification: From	Machine modification: To	Total
Substitutes					
Uni	Failed collusion	(A,B), (B,A)	—	—	—
	Miscoordinate	(A,C), (A,D), (B,C), (B,D)	AAAA	CCCC	(3,2)
		(A,C), (A,D), (B,C), (B,D)	AAAA	DCCC	(2,2)
	Punish path	(C,D), (D,C), (D,D)	CCCC	DDDD	(2,2)
			DDDD	CCCC	(3,3)
Bil	Failed collusion	(A,B), (B,A)	—	—	—
	Miscoordinate	(A,C), (A,D), (B,C), (B,D)	AAAA	—	(11,3)
			ABCD	—	(3,2)
			ACCC	—	(3,3)
			BCCC	—	(3,2)
	Punish path	(C,D), (D,C), (D,D)	AAAA	—	(11,3)
			ABCD	—	(5,4)
			ACCD	—	(6,3)
			DDDD	—	(5,4)
Complements					
Uni	Failed collusion	(A,B), (B,A)	—	—	—
	Miscoordinate	(A,C), (A,D), (B,C), (B,D)	AAAA	CCCC	(3,2)
	Punish path	(C,D), (D,C), (D,D)	—	—	—
Bil	Failed collusion	(A,B), (B,A)	—	—	—
	Miscoordinate	(A,C), (A,D), (B,C), (B,D)	ABCD	—	(2,2)
			BBBB	—	(6,2)
	Punish path	(C,D), (D,C), (D,D)	—	—	—

Notes: Only prominent machines that were (attempted to be) modified by at least two different individuals for a given type of history during matches 4–10 are included. In the final column (total), the first number is the number of times the particular machine was (attempted to be) modified; the second number is the number of different individuals that contributed to that number.

E.3 Further Statistical Robustness Checks

Table E.9: Pairwise treatment comparisons of the effect of revision opportunities on efficiency – further robustness checks for Result 1.

	Substitutes			Complements		
	Baseline	Bilateral	Unilateral	Baseline	Bilateral	Unilateral
First Stage						
Baseline	35.2	(0.982, 0.936)	(0.910, 0.749)	19.6	(0.489, 0.873)	(0.012, 0.025)
Bilateral		35.4	(0.924, 0.873)		14.6	(0.000, 0.004)
Unilateral			36.5			43.3
All Stages						
Baseline	23.3	(0.328, 0.522)	(0.821, 0.873)	9.7	(0.762, 0.337)	(0.044, 0.037)
Bilateral		16.6	(0.361, 0.423)		11.0	(0.038, 0.016)
Unilateral			21.7			25.8

Notes: The table uses data for matches 7-10. The statistics on the diagonal report the average efficiency for the respective treatment. The statistics in the upper diagonal give p-values for tests of the difference between the two treatments. The first p-value is based on a linear regression. The second p-value is the result of a ranksum non-parametric test on matching group averages. The first three columns replicate the result that there is no effect of revision opportunities under substitutes. The last three columns replicate the result that collusion is significantly higher with unilateral revision opportunities under complements.

Table E.10: Pairwise treatment comparisons of the effect of interaction type on efficiency – further robustness checks for Result 2.

	First Stage			All Stages		
	Substitutes	Complements	p-value	Substitutes	Complements	p-value
Baseline	35.2	19.6	(0.076, 0.150)	23.3	9.7	(0.148, 0.109)
Bilateral	35.4	14.6	(0.001, 0.010)	16.6	11.0	(0.046, 0.037)
Unilateral	36.5	43.3	(0.065, 0.055)	21.7	25.8	(0.212, 0.262)

Notes: The table uses data for matches 7-10. The statistics in the "p-value" columns give p-values for tests of the difference between the two treatments. The first p-value is based on a linear regression on interaction-type treatment dummies. The second p-value is the result of a ranksum non-parametric test on matching group averages. The first three columns show that, with no revision opportunities, there is significantly more collusion, at least in first-stage choices, under substitutes than under complements.

Table E.11: Pairwise treatment comparisons of the effect of revision opportunities on intended strategies – further robustness checks for Result 3.

	Substitutes			Complements		
	Baseline	Bilateral	Unilateral	Baseline	Bilateral	Unilateral
All machines with a cooperative response to action A						
Baseline	39.6	(0.203, 0.076)	(0.520, 0.684)	33.3	(0.672, 0.746)	(0.019, 0.064)
Bilateral		47.9	(0.003, 0.016)		30.2	(0.000, 0.008)
Unilateral			34.9			51.6
All machines with a cooperative response to action B						
Baseline	27.6	(0.933, 0.872)	(0.785, 0.870)	22.4	(0.079, 0.148)	(0.012, 0.054)
Bilateral		28.1	(0.841, 0.936)		33.3	(0.336, 0.370)
Unilateral			29.2			39.1
All machines with a punishment response to the deviation action (D/B)						
Baseline	22.4	(0.125, 0.150)	(0.861, 0.686)	3.1	(0.288, 0.434)	(0.823, 1.000)
Bilateral		37.0	(0.125, 0.226)		1.6	(0.305, 0.601)
Unilateral			23.4			3.6

Notes: The table uses data for matches 7-10. The statistics on the diagonal report the average percentage of intended strategies for the respective treatment. The statistics in the upper diagonal give p-values for tests of the difference between the two treatments. The first p-value is based on a linear random-effects regression on revision-opportunities treatment dummies. The second p-value is the result of a ranksum non-parametric test on matching group averages.

Table E.12: Pairwise treatment comparisons of the effect of interaction type on intended strategies.

	Collusive response to action A			Collusive response to action B		
	Substitutes	Complements	p-value	Substitutes	Complements	p-value
Baseline	39.6	33.3	(0.489, 0.466)	27.6	22.4	(0.449, 0.421)
Bilateral	47.9	30.2	(0.000, 0.005)	28.1	33.3	(0.356, 0.520)
Unilateral	34.9	51.6	(0.005, 0.023)	29.2	39.1	(0.082, 0.076)
	Deviation response (D/B) to action A			Punishment response to deviation (D/B)		
	Substitutes	Complements	p-value	Substitutes	Complements	p-value
Baseline	30.7	35.4	(0.562, 0.569)	22.4	3.1	(0.000, 0.027)
Bilateral	32.8	35.4	(0.791, 1.000)	37.0	1.6	(0.000, 0.007)
Unilateral	38.0	31.8	(0.428, 0.420)	23.4	3.6	(0.000, 0.003)

Notes: The table uses data for matches 7-10. The statistics in the "p-value" columns give p-values for tests of the difference between the two treatments. The first p-value is based on a linear random-effects regression on interaction-type treatment dummies. The second p-value is the result of a ranksum non-parametric test on matching group averages.