UNIVERSITY OF SUSSEX

ECONOMICS

# Working Paper Series

## No. 03-2017

# Eliciting strategies in indefinitely repeated games of strategic substitutes and complements*

**Matthew Embrey[a], Friederike Mengel[b] and Ronald Peeters[c]**

a. Department of Economics, University of Sussex

m.embrey@sussex.ac.uk

b. Department of Economics, University of Essex

fr.mengel@gmail.com

c. Department of Economics, Maastricht University

r.peeters@maastrichtuniversity.nl

**Abstract:** We introduce a novel method to elicit strategies in indefinitely repeated games and apply it to games of strategic substitutes and complements. We find that out of 256 possible unit recall machines (and 1024 full strategies) participants could use, only five machines are used more than 5 percent of the time. Those are "static Nash", "myopic best response", "Tit-for-Tat" and two "Nash reversion" strategies. We compare outcome data with "hot" treatments and find that the fact that we elicit strategies did not affect the path of play. We also discuss applications to IO literature and compare insights to previous literature on strategy elicitation mostly focused on the prisoner's dilemma.

**JEL classification:** C7, C9

**Key words:** Indefinitely Repeated Games, Strategy Elicitation, Experiments.

* 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, GAMES 2016), New York (ESA 2012), Vienna (Exp. Econ. workshop) and Boston (IIOC 2013) for valuable comments and the Netherlands Organisation for Scientific Research (NWO) and the European Union (FP7-ERG-277026) for financial support.

# 1 Introduction

Indefinitely repeated games are pervasive in our everyday life and of central relevance in many economic applications. They give rise to incentives that are fundamentally different from those of one-shot interactions and we experience these differences throughout our daily lives.[1]

Despite their importance and a great volume of (especially theoretical) research in the area, there are still surprisingly big gaps in our knowledge about these games. In particular, little is known about which *strategies* people use in these games, with the exception maybe of the prisoner's dilemma, research on which we discuss below. Understanding strategies can be important for a number of reasons. It can inform (applied) theorists trying to use game theory more predictively. Folk theorems, which provide characterizations of payoffs but not strategies, are often of limited use for applications. Researchers in evolutionary biology and game theory who try to understand which strategies survive evolutionary selection face similar problems. Given the infinite amount of possible strategies, these researchers are forced to constrain the set of possible strategies ex ante. Understanding which strategies people actually use in such games can hence be useful to inform the modeling choices of these researchers and can help theorists to identify conditions under which certain strategies are played. Strategies are also needed to understand off-equilibrium path behaviour, to get a grip on the extent of path dependency and a sense of the potential for miscoordination. The elicitation of strategies can help us gain a better understanding of the motives underlying certain behaviours (for example, whether people cooperate for altruistic or strategic reasons, Breitmoser, 2015).

In this paper we contribute to a recent emerging literature trying to elicit strategies in indefinitely repeated games using lab experiments (Selten et al., 1997; Dal Bó and Fréchette, 2011b; Mengel and Peeters, 2011; Romero and Rosokha, 2016). There are two important differences between this literature and our paper. First, we propose a new elicitation method (differences to each of these are explained below). Second, we apply this method to study $4 \times 4$ games mimicking linear Bertrand and Cournot settings, while most of the literature has focused on $2 \times 2$ prisoner's dilemma games.[2]

[1] The latter point is nicely illustrated by the following example from Mailath and Samuelson (2005): Suppose, on taking your car for an oil change, you are told that an engine problem has been discovered and requires an immediate and costly repair. The confidence you will have on whether or not this diagnosis is accurate will likely depend on whether you regularly do business with the service provider or whether you are passing through on vacation.

[2] The exceptions are Selten et al. (1997) and Mengel and Peeters (2011) discussed below.

Studying $4 \times 4$ games allows us to distinguish between strategic complements and substitutes with its important applications in Industrial Organization. It also allows us to uncover new motives, since strategies can be distinguished using the $4 \times 4$ setting that cannot be distinguished in the prisoner's dilemma game. For example, consider the "myopic best response strategy" and the strategy "always Nash". In the prisoner's dilemma these are the same and involve choosing "defection" at all histories. The prisoner's dilemma also does not allow for partial cooperation and punishing strategies cannot be distinguished from strategies such as Nash reversion. In the $4 \times 4$ game all of these can be distinguished. Using a slightly larger game implies an exponentially larger space of possible strategies. This creates design challenges, but also allows us to see how participants reduce this complexity, not present in simpler $2 \times 2$ games.

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. There are 256 of such unit-recall machines in the $4 \times 4$ game and 1024 combinations of initial action and dynamic response machine. In addition, participants can revise their machines at any point during the experiment by paying a small cost. Participants can also use one-shot deviations, but have to pay for them each time. A machine revision allows participants to economize on such costs in future choices. Machine revisions and one-shot deviations are included so that participants have the full (infinite) strategy space of the repeated game available. Small costs are imposed in order to encourage careful programming of the initial machines. A hot variation where these costs were set to zero is used to see whether behavior (not strategies) is affected by the cost imposed on these deviations and machine revisions.

In terms of the strategies participants use, we find that most use machines that are familiar from the literature. Out of the 256 possible machines participants could program, the most prominent machines mimic the static Nash, myopic-best-response, tit-for-tat and Nash-reversion strategies. There are only five strategies in total that are used more than 5 percent of the time in any treatment. Given the total number of possible strategies this is remarkable. Participants pick strategies we are familiar with from theory even when there are many different options to pick, as is the case in the $4 \times 4$ game.

"Hot" treatments, where participants in each period could choose any action they like at no additional cost, reveal no behavioural differences in terms of the distribution of action choices over time nor the efficiency of outcomes and only very few differences in terms of the extent of collusion in later matches. This leads us to conclude that our elicitation method did not substantially distort participants' behaviour. Further evidence in this regard comes from our analysis of one-shot deviations and machine revisions. One-shot deviations are used about 5% of the time and machine revisions roughly 3% of the time, i.e. machines are revised approximately once every five matches. Participants are largely relying on unit-recall strategies.

## 1.1 Contribution to the Literature

Our design and results contribute to the growing literature on strategy elicitation in indefinitely repeated games. Most of the existing literature has focused on the prisoner's dilemma (Dal Bó and Fréchette, 2011b,a; Camera et al., 2012; Fudenberg et al., 2012; Breitmoser, 2015; Romero and Rosokha, 2016)

Dal Bó and Fréchette (2011b) and Romero and Rosokha (2016) elicit strategies directly. Both in Dal Bó and Fréchette (2011b) and in our study participants tend to use simple strategies that only condition on the previous period's outcomes. They also find, just as we do, that the distribution of strategies changes with the game parameters. The most common strategy in Dal Bó and Fréchette (2011b) is "Always Defect", and in Romero and Rosokha (2016) they are "Tit-for-Tat" and "Always Defect". Around 67% of our strategies would correspond to either of these if projected on the prisoner's dilemma.

Rather than eliciting strategies, other authors have focused on inferring strategies from choices (Dal Bó and Fréchette, 2011a; Camera et al., 2012; Fudenberg et al., 2012; Breitmoser, 2015) often via maximum likelihood methods. Most recently Breitmoser (2015) conducted a meta-study of four experiments and finds that a strategy called "Semi-Grim" (cooperate after mutual cooperation, defect after mutual defection, randomize otherwise) summarizes behaviour well across these studies. This is in line with our results, but does not do justice to the considerable heterogeneity in terms of strategies that we find. In line with our findings, Camera et al. (2012) find that Grim-Trigger does not describe individual play well and that there is considerable heterogeneity in which strategy fits best. Fudenberg et al. (2012) have found that in games with imperfect monitoring "forgiving" strategies, such as Nash Reversion or Forgiving Grim Trigger, are used more often. To sum up, most of the literature with perfect monitoring finds evidence for memory one or unit recall strategies, suggesting that by focusing on those the researcher may not miss too much (see, also, the meta-study of Dal Bó and Fréchette, 2016). One exception may be lengthy games (with a discount factor close to one) for which Romero and Rosokha (2016) have found that strategies with longer memory can be important to sustain cooperation.

By virtue of using the $4 \times 4$ game, we also uncovered results that cannot be obtained in the prisoner's dilemma. Across the prisoner's dilemma studies discussed above many players use "Always Defect". What these studies cannot say, however, is whether this is a result of playing "Always Nash" or, for example, "myopic best response" or "punishing strategies", all of which coincide with "Always Defect" in the prisoner's dilemma. Our results show that about equally many people choose one of the former two strategies, while very few participants choose "punishing strategies". Which of these strategies is played can have fundamentally different implications in social dilemma situations that have more shades of grey than the heavily studied prisoner's dilemma, and each of them are open to differing degrees to interventions aimed at improving efficiency.

Not all work on strategy elicitation has focused on the prisoner's dilemma. Mengel and Peeters (2011) elicited strategies in the public good game. The motivation behind their work was, however, not strategy elicitation per se, but to understand the differential dynamics in partner (repeated) versus stranger (random rematching) settings in the public good game. Selten et al. (1997) asked experienced participants to program strategies in PASCAL to compete in a finitely repeated Cournot duopoly. They found evidence for strategies reminiscent of "Tit-for-Tat". Participants in Selten et al. (1997) try to achieve cooperation by a "measure-for-measure policy" which reciprocates movements towards and away from the ideal point by similar movements.

Our paper also contributes methodologically to the literature by proposing a novel elicitation method. Just as in Dal Bó and Fréchette (2011b) we ask participants to specify unit-recall strategies, i.e. strategies with memory one or less. There are two differences to our approach and Dal Bó and Fréchette (2011b)'s method. First, while in Dal Bó and Fréchette (2011b) participants are truly restricted to memory one, in our paper one-shot deviations and machine revisions make the full strategy space to participants, albeit at a small cost. Second, while Dal Bó and Fréchette (2011b)'s unit recall strategies condition on (partial) histories, our unit recall strategies condition only on the opponent's past choices.[3] Our elicitation method coincides with that in Mengel and Peeters (2011) for a public good game with ten actions available, but they did not allow one-shot deviations or machine revisions.

In another strand of literature participants are asked to program strategies. The famous tournaments conducted by Axelrod (1984), where scholars were asked to submit computer programs to play the prisoner's dilemma are an early example, and also Selten et al. (1997) fits in this category. The downside of this straightforward elicitation of strategies is that it is likely too complex for non-experienced participants or for participants with little formal education. Romero and Rosokha (2016) have recently proposed a method to elicit strategies in the prisoner's dilemma that is better suited for use in experiments with non-experienced participants. Our paper complements theirs in the sense that it involves only minimal restrictions on the set of strategies that can be implemented, while being feasible to use with non-experienced subjects.

Our study also contributes to the experimental literature, often motivated by IO applications, trying to understand when and why people cooperate in games of strategic substitutes and complements. Among these one study particularly close to ours is Potters and Suetens (2009), who study collusion in a laboratory experiment using finitely repeated games of strategic substitutes and complements without eliciting strategies. They find more cooperation when actions exhibit strategic complementarities. Our results in terms of behaviour are in line with theirs. The fact that there is more cooperation in the complements game seems indeed prevalent. An analysis of Dal Bó and Fréchette (2016)'s meta-data on the indefinitely repeated prisoner's dilemma game reveals that, controlling for discount factor, supergame and stage, there is more cooperation in complements compared to substitutes.[4]

[3] Another difference is that Dal Bó and Fréchette (2011b) first let participants play a standard prisoner's dilemma game. Second, they have a phase which is akin to our "hot" treatment, where participants programme a strategy but then are not committed to it. Third, there is a phase where the strategy plays for them.

[4] Following the normalisation of Dal Bó and Fréchette (2016) (where the Nash outcome is 0, the cooperative outcome is 1, the temptation payoff is $1 + g$ and the sucker payoff is $-l$) strategic complementarity is given by the condition $g > l$, i.e. increasing differences. We regress a dummy indicating whether a participant cooperated in a given round on a dummy indicating whether the game

Having elicited strategies, we can shed light on one of Potters and Suetens (2009) conjectures where they argue that the differences between complements and substitutes is likely due to the differential dynamics when a cooperative player is matched with someone choosing a best response strategy. While both of these machine categories are prevalent in our experiment, a counterfactual simulation exercise where we take the elicited strategy pairs from one stage game and predict outcomes in the other, reveal that this, more mechanical, effect cannot explain the differences between substitutes and complements in our data. These results illustrate how knowledge of participants' strategies can distinguish between different interpretations of findings that are based on outcome data only.

This paper is organized as follows. Section 2 describes the experimental design and the strategy elicitation method. Section 3 shows the main results and Section 4 concludes.

# 2 The experiment

In this section, we describe the main details of the experimental design and procedures. We use data from two treatments of the experiment conducted for our companion paper, Embrey et al. (2016).[5] In Section 2.1 we describe the design of these two treatments and why it is well suited for the questions asked in this paper. In Section 2.2 we describe the experimental procedures.

## 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.*

is of strategic complements and the discount factor $\delta$ using their data. Controlling for round and supergame, both linearly and via fixed effects, we find about 14% higher cooperation rates in the complements games. The effect is statistically significant at the 1% level, using standard errors clustered at the session level. Note that this *strategic complementarity* effect does not always work in the same direction as the *size-of-the-basin-of-attraction* effect identified in Dal Bó and Fréchette (2016). In particular, fixing the loss parameter ($l$) and changing the stage game parameters to move from strategic substitutes to strategic complements would require increasing the difference between the cooperation payoff and the temptation payoff (i.e. increasing $g$). This change would require a larger belief in the other player playing the grim trigger strategy, rather than the always defect strategy, to sustain cooperation and thus making it more difficult according to the size-of-the-basin-of-attraction effect.

[5]Embrey et al. (2016) focuses on strategy revision opportunities. The treatments *Elicit* in the present study are called *unilateral revision opportunities* there. The *Hot* treatments in the present paper are not part of that study.

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. The underlying duopoly games were calibrated so that the majority of payoffs for key action pairs are identical across games: (i) the Nash equilibrium payoffs that result from both players playing action C are identical; (ii) the joint payoff maximizing payoffs that result from both choosing action A are identical; (iii) 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 for the defector and the sucker across games; (iv) 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.[6]

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 $\delta = 7/8$ in our experiment is above this level.[7] These design choices ensure that any difference in the distribution of strategies we may detect should be uniquely due to the complements/substitutes character of the game.

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, this optimal deviation action is located between the collusive action (A)

[6] 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.

[7] 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.

and the Nash action (C) while it is located beyond the Nash action in the substitutes game.[8] This difference in the location of these 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 refer to the actions A, B, C and D as respectively *Collusion*, *Dev.SC*, *Nash* and *Dev.SS*.

**Strategy elicitation** 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. In total there are 256 possible machines, which combined with 4 possible initial actions allows for 1024 possible intended strategies.

To allow participants access to the full infinite strategy space, in the treatment termed 'Elicit', they are allowed to take an action that differs from the one recommended by their machine. Such changes are referred to as one-shot deviations. Strategies, such as grim-trigger, become feasible to implement via one-shot deviations.[9] In

[8] 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.

[9] Indeed, such a trigger strategy can be used to sustain collusion for all discount rates above the $\delta_{min}$ calculated earlier for the standard repeated game (i.e. just action choices in each period). See

addition, we allow for strategy revisions by allowing participants to modify their machines after any stage of the repeated game. To provide participants with an incentive to program their (initial) machines (strategies) carefully, one-shot deviations and machine modifications have a small cost associated with them. Each one-shot deviation costs 3 ECU and each machine modification costs 1 ECU irrespective of the number of elements in the machine that are changed. In our second treatment termed 'Hot' these costs were set to zero. This treatment is used to show that behavior was not unduly affected by the small cost imposed on these deviations.

## 2.2 Procedures

The experiment was conducted in the BEElab at Maastricht University. 160 students were recruited using ORSEE (Greiner, Greiner) and participated in one of the two treatments.[10] Decisions were made in isolated cubicles and interaction was computerized using z-Tree (Fischbacher, 2007). Sessions lasted an hour and a half on average, including a twenty minute instruction period.[11] On average participants earned between 12 and 16 Euro.

For each of the Elicit treatments six matching groups were run and for each of the Hot treatments we had four matching groups. 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.

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.

Section B of the Supplementary Materials for details.

[10]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. In addition, we conducted some treatments where we studied variations on renegotiation and strategic commitment. Details on those are reported in Embrey et al. (2016). Other than the treatments mentioned we did not conduct any additional treatments.

[11]Due to the randomness in the length of the matches, subjects were recruited for up to 2 hours. No session lasted longer than 1 hour and 45 minutes.

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 gives the number of observations for each treatment. Table D.1 of the supplementary materials provides further details on the sequence of match-lengths for the different matching groups.

Table 1: Summary of treatments.

| | Match. Groups | Num. Subj. | Efficiency (%) | Incurred Costs Avg. (ECU) | Incurred Costs 1-Shot (%) | Incurred Costs Machine (%) |
|---|---|---|---|---|---|---|
| *Substitutes* | | | | | | |
| Elicit | 1-6 | 48 | 16.5 | 0.2 | 6.5 | 3.7 |
| Hot | 1,3,4,5 | 32 | 16.9 | | | |
| *Complements* | | | | | | |
| Elicit | 1-6 | 48 | 23.2 | 0.2 | 5.1 | 3.0 |
| Hot | 1,3,4,5 | 32 | 26.2 | | | |

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.

# 3 Results

This section contains our main results. In Section 3.1 we discuss the distribution of strategies. In Section 3.2 we compare behaviour in the Elicit treatments with behaviour in the Hot treatments and in Section 3.3 we show results on strategy revisions.

## 3.1 Strategies

Table 2 shows the distribution over machines in the Elicit treatments. The table lists all machines that were used at least 5% of the time in either the Substitutes or the Complements game as well as two machines that were not used as frequently but we feel deserves some attention.

Across the two treatments there are only five strategies (out of a possible 1024) that were used more than 5% of the time. Those strategies are (i) A-AAAA: a strategy of unconditional cooperation in each round; (ii) A-ABCD (or sometimes A-ABCC): a strategy of conditional cooperation or tit-for-tat, where players start out cooperatively and then mimic the action of their opponent in the previous period; (iii) A-ACCC: a strategy known as Nash reversion, where players start out cooperatively and revert to Nash as soon as the opponent chooses something else; (iv) C-CCCC: the strategy that always plays the stage-game Nash action and (v): D-DCCC (or sometimes C-DCCC) or B-BCCC: a strategy that always myopically best responds to the choice made by the opponent in the previous period starting out either with the myopic best response to collusion or in some cases with the Nash choice.

Table 2: Distribution of initial choice and machine categories (in percent).

| Prominent Machine Category | Strategy Initial-Machine | Substitutes | Complements |
|---|---|---|---|
| Unconditional cooperation | **AAAA** | **5** | **8** |
| | A- | *5* | *8* |
| Conditional cooperation | **ABC(C/D)** | **11** | **23** |
| | A- | *6* | *19* |
| Nash reversion | **ACCC** | **7** | **12** |
| | A- | *7* | *12* |
| Partial collusion + Nash rev. | **BBCC** | **2** | **1** |
| Nash | **CCCC** | **22** | **14** |
| | C- | *14* | *13* |
| Myopic best reponse | **DCCC** | **19** | |
| | C- | *7* | |
| | D- | *11* | |
| | **BCCC** | | **15** |
| | B- | | *10* |
| Punishing | **DDDD** | **3** | **0** |
| **Other** | ——— | **31** | **27** |
| All machines with a ... | | | |
| cooperative response to action A | | 35 | 52 |
| cooperative response to action B | | 29 | 39 |
| deviation response (D/B) to action A | | 38 | 32 |
| punishing response to deviation (D/B) | | 23 | 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 to 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.

Hence, despite the large number of strategies available, there are five intuitive strategies that participants seem to focus on. If researchers were able to focus on only these five strategies, there would be some benefits. It would allow for a full characterization of equilibria in the indefinitely repeated game and hence possibly allow for sharper predictions. On the other hand, the results also emphasize the importance of heterogeneity with different strategies being used in different matches.

The category "Other", which comprises all machines used less than 5% of the time in both treatments, accounts for about 30% of the machines. In Table D.2 we decompose this category, by assigning – if possible – the machines in this category to a machine explicitly listed in Table 2 as long as they respond to collusion with the same action (e.g. ADDD does not count towards DDDD) and they have Hamming distance of at most one.[12]

After this reassignment procedure, the category "Other" is reduced to 9% (13%) of the cases in the Substitutes (Complements) treatment and more than 75% of all strategies belong to one of the following four categories: (i) conditional cooperation; (ii) Nash reversion; (iii) Nash and (iv) myopic best responses. The remaining (less than 25%) strategies are split among (i) unconditional cooperation, (ii) the category "Other" and the two machines that we had highlighted in Table 2, as they appeared intuitive to us, but are used more than 5% of the time only after the decomposition of the category "Other". Those are (i) the partial collusion machine BBCC that partially colludes but reverts to Nash if the opponent chooses C or D and (ii) the punishing machine DDDD. No other machine reaches the 5% threshold.

The five most prominent machines identified in Table 2 are the same across the complements and substitutes conditions. However, there are some differences when it comes to the distribution over these five machines. Participants respond more often with a cooperative choice to actions A/B and they use punishing responses less often under complements (see Table D.3 in Appendix D.2). This is in line with behavioural evidence obtained in a hot treatment by Potters and Suetens (2009) who find more cooperation in the complements compared to the substitutes game (see also Embrey et al., 2016). We summarize our results obtained in this section as follows.

**Result 1** *In both treatments more than* 75% *of strategies belong to one of these four categories: conditional cooperation (tit-for-tat), Nash reversion, always Nash or myopic best responses.*

[12]The Hamming distance (Hamming, 1950) between two strings of equal length is the number of positions at which the corresponding symbols are different. It hence measures the minimum number of substitutions required to change one string into the other, or the minimum number of errors that could have transformed one string into the other.

## 3.2 The Elicitation Method and Outcomes: Elicit versus Hot

In this section we compare our Elicit treatments with the Hot treatments in order to understand whether our elicitation procedure might have had an effect on participants' choices per se. In other words, we investigate whether the elicitation mechanism has a significant effect on outcomes in the repeated game.

Figure 2 gives an overview of the distribution of action choices, in the later matches of a session, for all stages in the left-hand panel and for the first stage only in the right-hand panel. This figure provides at this aggregate level no evidence for subjects in the Elicit treatments being inclined to make different action choices than those in the Hot treatments. None of the sixteen comparisons shown in Figure 2 are statistically significant at the 5% level using two-sided ranksum tests on matching group averages.[13]

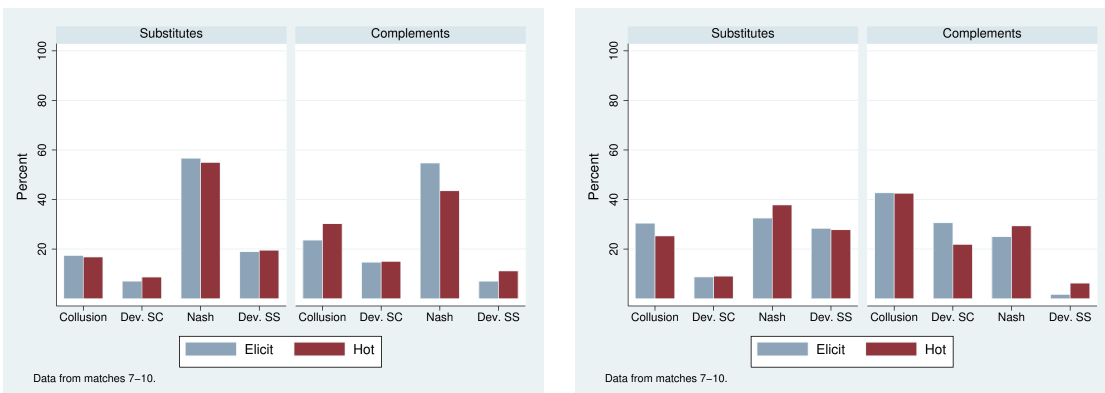


(a) In all stages.

(b) Initial choices only.

Figure 2: Average stage-game choices.

Two measures of outcomes are used to investigate further whether behaviour in the two treatments is empirically comparable: the first concerns the rate of collusion choice by subjects, while the second concerns the implied efficiency of the action choices made by a matched pair. Figure 3 illustrates the percent of collusive choices across matches (panel (a)) and stages (panel (b)). Figure D.3 in Appendix D.2 presents this information for the efficiency measure. As can be most clearly seen in the graphs, behaviour in the Hot treatment is remarkably similar to that in the Elicit treatment in all aspects. The only difference seems to be a 2 to 9 percentage point higher probability

[13]One difference is statistically significant at the 10% level. This is the difference in the proportion of initial choices of the myopic best-response action (Dev.SC) under strategic complements. See Table D.4 for further details.

to choose the collusive action in the Hot treatment, which also translates into somewhat higher efficiency.

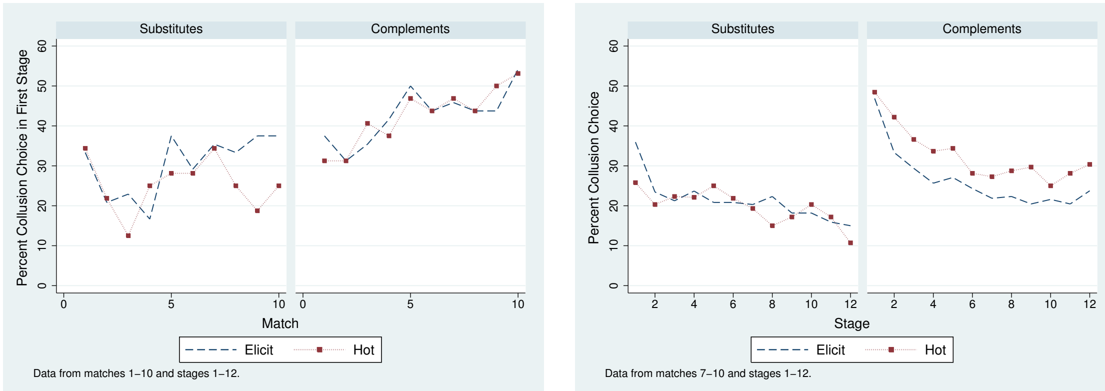


(a) Collusion choice across matches.

(b) Collusion choice within matches.

Figure 3: Collusion choice in the Elicit and Hot treatments.

Table 3 reports the results of a regression analysis on the probability of a collusive action choice, using the Elicit treatments as the baseline. Table D.5 in Appendix D.2 repeats the analysis for the efficiency measure. In the substitutes game (columns (1) and (2)) there is essentially no difference between the two treatments: The coefficient on Hot and the interactions Hot × Stage/Match are near zero and statistically insignificant. In the complements game there seems to be somewhat more collusion in the Hot treatment across matches 7-10, but no difference across all matches and no difference in the Hot × Stage/Match interactions. There is also no significant difference according to the efficiency measure.

Table 3: Random-effects logit regression of the probability of choosing the collusive action.

| | Substitutes | | | | Complements | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | | (2) | | (3) | | (4) | |
| Hot | −0.04 | (0.516) | −0.00 | (0.958) | 0.16** | (0.012) | 0.11 | (0.218) |
| Stage | −0.01*** | (0.002) | −0.01*** | (0.004) | −0.02*** | (0.000) | −0.02*** | (0.000) |
| Hot x Stage | −0.00 | (0.848) | 0.00 | (0.703) | −0.00 | (0.691) | 0.01 | (0.210) |
| Match | | | 0.02*** | (0.001) | | | 0.02** | (0.023) |
| Hot x Match | | | −0.00 | (0.771) | | | −0.01 | (0.642) |
| Matches | 7–10 | | 1–10 | | 7–10 | | 1–10 | |

Notes: Coefficients report the marginal effect on the probability of choosing the collusion action. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%. All regressions include only data for stages 1-12. All regressions include a set of controls for the match-stage composition.

**Result 2** *There are no statistically significant differences in the distribution of action choices nor in the efficiency of outcomes (mean and across stage/match evolution) and only few differences using a binary indicator of collusion between the Elicit and Hot treatments.*

This result makes us confident that our elicitation procedure works sufficiently well. In particular it shows that asking participants to specify strategies (and imposing small costs for one-shot deviations and revisions) does not lead to substantially different outcomes compared to the case where there are no costs and participants can choose actions freely in each period. Note, however, that this does not necessarily mean that the strategies described in Table 2 are very representative of the strategies participants really use. If participants revised strategies a great deal or relied extensively on one-shot deviations, then the true strategies could look quite different from those elicited. In the next subsection we will hence study how often participants revise machines and how often they engage in one-shot deviations.

## 3.3 Strategy Revisions and One-shot Deviations

Table 4 shows the average propensity of our participants to make one-shot deviations. The overall propensity to do so is below 10% and decreases to 5% across the last four matches. Given our average match length of about 8 stages and the fact that a one-shot deviation is not possible in the first stage of each repeated game, these 5% mean that participants made a one-shot deviation about once every three matches. Even focussing on just the subjects that made at least one one-shot deviation (over 45% did not make any in the last four matches), this rate rises only to on average just over three one-shot deviations over the last four matches. That is, less than one one-shot deviation per match by the end of the session.

The bottom part of Table 4 shows that all machines trigger one-shot deviations sometimes. Comparisons across machines are difficult because of the low number of instances of one-shot deviations. Table D.6 in Appendix D.2 shows what action participants choose after a one-shot deviation. The table shows that one-shot deviations are made both to be more and less cooperative; though, the numbers are somewhat hard to interpret because of the low number of instances of one-shot deviations.

In addition to one-shot deviations, participants in our experiment could also revise their machines. Table 5 shows how often machine changes occurred and what machines are changed most often. The table shows that machine changes are even rarer than

Table 4: Propensity to make one-shot deviations in the Elicit treatments.

| | | Substitutes Elicit | | Complements Elicit | |
|---|---|---|---|---|---|
| *By data subsample:* | | | | | |
| Matches 1–3 | | 10 | (72, 23) | 9 | (68, 21) |
| Matches 4–6 | | 9 | (91, 22) | 6 | (64, 17) |
| Matches 7–10 | | 6 | (99, 27) | 4 | (71, 25) |
| *By machine category (matches 7–10):* | | | | | |
| Unconditional cooperation | AAAA | 7 | (6, 3) | 2 | (2, 2) |
| Conditional cooperation | ABC(C/D) | 3 | (6, 3) | 8 | (29, 12) |
| Nash reversion | ACCC | 5 | (8, 5) | 3 | (8, 5) |
| Partial collusion + Nash rev. | BBCC | 0 | (0, 0) | 15 | (3, 2) |
| Nash | CCCC | 7 | (22, 8) | 1 | (2, 2) |
| Myopic best reponse | DCCC | 2 | (7, 5) | | |
| | BCCC | | | 5 | (12, 6) |
| Other | ——— | 9 | (50, 11) | 3 | (15, 9) |

Notes: Average propensity to make a one-shot deviation, as a percentage of decision periods in which such a change was possible. In parentheses, the total number of deviations and the number of distinct individuals deviating.

Table 5: Propensity to make machine changes in the elicit treatments.

| | | Substitutes Elicit | | Complements Elicit | |
|---|---|---|---|---|---|
| *By data subsample:* | | | | | |
| Matches 1–3 | | 5 | (37, 18) | 6 | (41, 15) |
| Matches 4–6 | | 5 | (52, 20) | 3 | (34, 16) |
| Matches 7–10 | | 3 | (61, 27) | 3 | (47, 22) |
| *By machine category (matches 7 –10):* | | | | | |
| Unconditional cooperation | AAAA | 6 | (5, 5) | 5 | (5, 3) |
| Conditional cooperation | ABC(C/D) | 1 | (2, 2) | 1 | (5, 5) |
| Nash reversion | ACCC | 3 | (4, 4) | 1 | (2, 2) |
| Partial collusion + Nash rev. | BBCC | 2 | (1, 1) | 10 | (2, 1) |
| Nash | CCCC | 4 | (13, 9) | 0 | (1, 1) |
| Myopic best reponse | DCCC | 3 | (8, 5) | | |
| | BCCC | | | 2 | (5, 4) |
| Punishing | DDDD | 7 | (3, 3) | 0 | (0, 0) |
| Other | ——— | 4 | (25, 13) | 6 | (27, 14) |

Notes: Average propensity to make a machine change, as a percentage of decision periods in which such a change was possible. In parentheses, the total number of deviations and the number of distinct individuals deviating.

one-shot deviations. Across the last four matches they occur in less than 3% of all possible cases. Given our average match length of about 8 stages and the fact that a machine revision is not possible in the first stage of each repeated game, these 3% mean that participants made a machine revision about once every five matches. Again, focussing just on subjects who made at least one machine revision (nearly 50% do not make any in the last four matches), this rate only rises to on average just over two machine revisions over the last four matches. The three machines changed most often in absolute numbers (across substitutes and complements) are the always Nash (CCCC), myopic best response (D/BCCC) and unconditional cooperation (AAAA) machines.

Table D.7 in Appendix D.2 shows what machines participants undertake. Because the low number of instances of machine revisions makes this table hard to interpret, Figures D.4 and D.5 in Appendix D.1 shows how these revisions impact the "average" machine. Interestingly, the "average machine" seems to become slightly more cooperative in the case of Substitutes and slightly less cooperative in the case of Complements after a revision. While the former seems in line with standard renegotiation arguments (Farrell and Maskin, 1989; McCutcheon, 1997; Aramendia et al., 2005), the latter is not. This effect can also be seen in Figure 4. The solid blue and dashed red lines compare observed behaviour (collusion choice and efficiency) in the last four matches to that predicted using the initial choice and dynamic response programmed in the first stage.[14] Overall, the two lines track each other closely, especially in the complements setting. In substitutes, the observed behaviour is slightly more cooperative than predicted behaviour, where as it is slightly less cooperative in complements. We summarize the results obtained in this subsection as follows.

**Result 3** *Participants rely on unit-recall machines most of the time. One-shot deviations are used in only* 5% *of the cases and machine revisions in only* 3% *of the cases across the last four matches.*

This result makes us confident that we can interpret the strategies elicited, listed in Table 2, as the core of the strategies actually used by participants.

[14]The third, dashed green line conducts a counter-factual exercise where, in the substitutes panels, we take the initially programmed behaviour of subjects from the complements treatment, adjust myopic best-response behaviour to be suitable for a strategic substitutes stage-game, and plot predicted behaviour of these subjects under substitutes. An analogous exercise is made for the complements panels. Figure D.6 re-runs the analysis, but using only machines that correspond to one of the prominent machines; this analysis is more conservative in identifying myopic best response behaviour compared to Figure 4. As can be seen, this *mechanical* adjustment (of subjects that programme initial choices and dynamic response machines consistent with myopic best response behaviour) is not able to capture the difference between behaviour under complements and substitutes. That is, initial choices and dynamic responses programmed under the complements setting still give rise to more cooperative behaviour than those under substitutes. The reason for this exercise is discussed in Section 4.

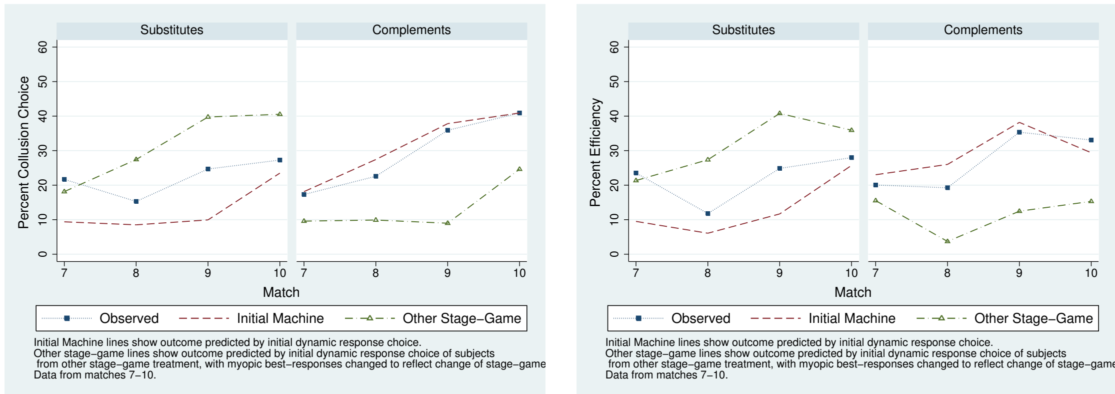


(a) Collusion Choice

(b) Efficiency

Figure 4: Observed cooperation measures versus dynamic-response prediction versus counterfactual.

# 4 Discussion and Conclusions

We elicited strategies in indefinitely repeated games of strategic substitutes and complements. To do so, we implemented a novel mechanism to elicit dynamic responses to others' choices in the repeated game. The mechanism circumvents a number of the difficulties that arise when trying to elicit strategies in complex environment such as repeated games, while providing the core of the strategies that subjects use. Our results find that participants largely rely on unit-recall strategies using one-shot deviations or machine revisions less than 5% of the time. Out of the 256 possible machines participants could program, the most prominent machines mimic the static Nash, myopic-best-response, tit-for-tat and Nash-reversion strategies. There are only five strategies in total that are used more than 5 percent of the time in any treatment.

An important feature of moving to stage games with a larger number of actions is that it allows dis-aggregation of a variety of strategies that would typically be always defect in the repeated prisoner's dilemma. In our games of strategic substitutes and complements with four actions in the stage games, we are able to distinguish between repeated play of the one-shot Nash equilibrium, myopic best response and punishing strategies. Among these strategies, we find consistent evidence for the former two in both games of strategic substitutes and complements, and little evidence for punishing strategies. However, it seems it is the always Nash strategy that responds more to the change in the strategic nature of the stage game, with a larger reduction in the proportion of this strategy in the complements stage game than is the case with the myopic best response strategy.

One of the reasons why eliciting strategies can be considered important is that it helps to understand off-path behaviour and the influence of path dependency. We illustrate this point using a conjecture by Potters and Suetens (2009) from their seminal paper showing that there is more collusion in games of strategic complements compared to strategic substitutes, a result that we replicate. They conjecture that these differences could be due to the differential dynamics when a cooperative player is matched with someone choosing a best response strategy (page 1139). Without having elicited strategies it is not possible to test this conjecture directly.[15] Since we have elicited information about subjects' strategies, we are able to consider this mechanism more directly. To do that, we undertake the following counter-factual exercise: Take the matched pairs and elicited strategies of subjects from the *strategic complements* sessions, adjust any myopic best-response behaviour to be suitable for a *strategic substitutes* stage-game, and plot the predicted behaviour of these pairs if they were to have played under strategic substitutes. An analogous exercise is conducted for matched pairs and elicited strategies from the strategic substitutes sessions.

This counter-factual, predicted behaviour is plotted in Figure 4 as the "other stage-game" line.[16] As can be seen, this *mechanical* adjustment (of subjects that programme initial choices and dynamic response machines consistent with myopic best response behaviour) is not able to capture the difference between behaviour under complements and substitutes in our data. That is, initial choices and dynamic responses programmed under the complements setting still give rise to more cooperative behaviour than those under substitutes, even after adjusting for the differential dynamics that can result when a conditional cooperator meets a myopic best-responder. This is one example of how knowledge about participants' strategies can help answer questions that would remain unresolved otherwise.

# References

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

[15] Potters and Suetens (2009) do find that the aggregate effect in their paper is driven by pairs who do not succeed in reaching full cooperation. This could, but need not be, concern matches between cooperative and best response players.

[16] Figure D.6 re-runs the analysis, but using only machines that correspond to one the prominent machines; this analysis is more conservative in identifying myopic best response behaviour compared to Figure 4. The conclusions are unaffected.

Axelrod, D. (1984). *The Evolution of Cooperation.* Basic Books.

Breitmoser, Y. (2015). Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma. *American Economic Review 105(9).*

Camera, G., M. Casari, and M. Bigoni (2012). Cooperative strategies in anonymous economies: An experiment. *Games and Economic Behavior 75*, 570–586.

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

Dal Bó, P. and G. Fréchette (2011b). Strategy choice in the infinitely repeated prisoners dilemma. Mimeo.

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

Embrey, M., F. Mengel, and R. Peeters (2016). Strategy revision opportunities and collusion. mimeo.

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.

Fudenberg, D., D. Rand, and A. Dreber (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review 102*, 720–749.

Greiner, B. Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association, year = 2015, volume = 1(1), pages = 114-125, owner = fuhrheiner, timestamp = 2016.10.11*.

Hamming, R. (1950). Error detecting and error correcting codes. *Bell System Technical Journal 29(2)*, 147–160.

Mailath, G. and L. Samuelson (2005). *Repeated Games and Reputation.* Oxford 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.

Romero, J. and Y. Rosokha (2016). Constructing strategies in the indefinitely repeated prisoner's dilemma game. mimeo.

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:[17] 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

[17] 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.[18] This difference in the location of these

[18]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.[19] 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.

[19] 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 primary treatments of the experiment. The purpose is to illustrate that eliciting dynamic responses, and including small costs for one-shot deviations and machine revisions, 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.

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

In the elicit treatments of the 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. Subject can also change their dynamic response for the following rounds at a cost of 1 ECU. Consequently, their action set in a round is

$$A^t = \left\{a^t, machine^t\right\} = \{A, B, C, D\} \times \{A, B, C, D\}^4$$

for all $t \geq 1$.

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, a strategy in the repeated game must specify an initial choice and dynamic response vector, plus an action choice and dynamic response 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:

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

where $\mathbb{I}\left(.\right)$ is an indicator function that takes value one if the statement in the argument is true and zero otherwise. In the experiment $c_{one} = 3$ and $c_{mach} = 1$. 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

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

Note that 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. In addition, this strategy implies that the player will never implement a machine change.

## B.2 Incentive Compatibility

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. To begin with, we will ignore the possibility of making machine changes after the first round, then show why the same logic follows through even with this possibility.

Given the nature of the elicitation setting, incentive compatibility in round one needs to be checked separately from incentive compatibility in later rounds, primarily because there are no costs for deviating from the action choice or dynamic response suggested by $S^{GT}$ in the first round. For later rounds, the only difference is that the specified choice, or the optimal deviation, might involve paying the small one-shot deviation cost of going against the recommended strategy.

Step 1: 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 $\delta$ that just makes the above inequality hold with equality is the same calculation for the minimum discount rate – denoted $\delta_{min}$ – as in the usual repeated game environment.

Step 2: 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

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

  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 $\delta$.

  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 $\delta$ since $\pi_i(A, C) - \pi^{Nash} > c_{one}$.

The following two observations imply that the above logic carries over when we add the possibility of making machine changes in later rounds:

- 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 $c_{one} \leq \pi^{Nash} - \pi_i\left(r_i^t, C\right)$ when $r_i^t \neq C$, it is preferable to pay the one-shot cost today to avoid the sucker payment.

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.

# C Example instructions: Complements Elicit

## 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 |
|---|---|---|---|---|---|
| | A | 43, 43 | 23, 54 | 14, 52 | 7, 47 |
| Your action | 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.

*Modifying the plan of intended actions*

16. One way to avoid incurring costs of choosing actions different from those prescribed by your plan in the future is to modify it. Every round, when you are asked to choose an action, you will be asked whether you would like to modify your current plan.

17. Whenever you declared that you would like to modify your plan, you will have the possibility to modify your plan. The same is true for the other participant.

18. Modifying the plan of intended actions costs 1 ECU. You have to modify your plan within the limit indicated on the screen

*At the end of each round*

19. 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**

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

21. 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 Additional material

## D.1 Figures

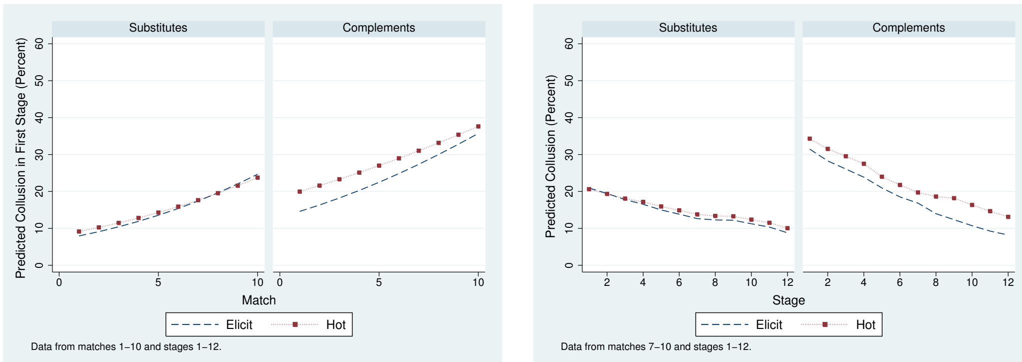


(a) Collusion choice across matches. (b) Collusion choice within matches.

Figure D.1: Collusion choice in the elicit and hot treatments: predicted values from a random-effects logit regression.

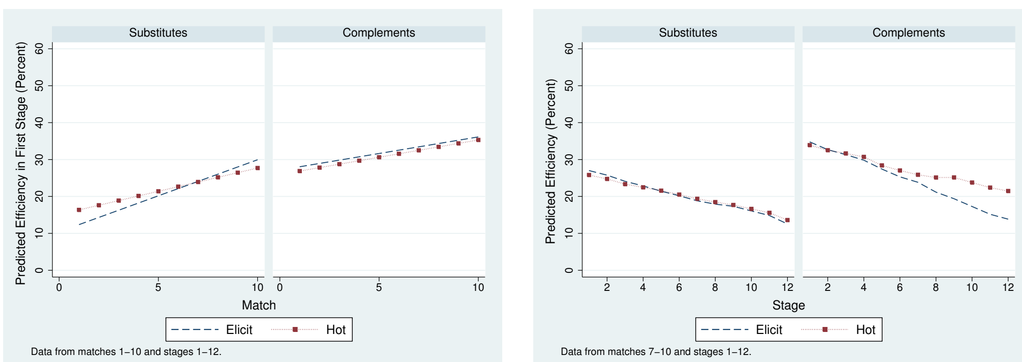


(a) Efficiency across matches. (b) Efficiency within matches.

Figure D.2: Efficiency in the unilateral and hot treatments: predicted values from a linear random-effects regression.

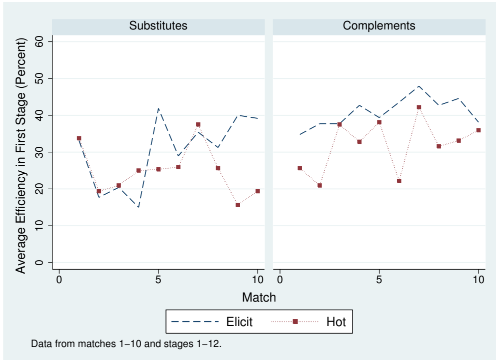


(a) Efficiency across matches.

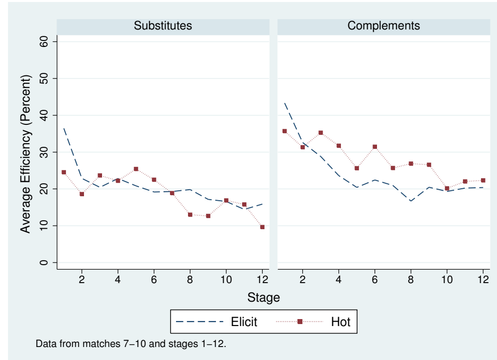


(b) Efficiency within matches.

Figure D.3: Efficiency in the unilateral and hot treatments.

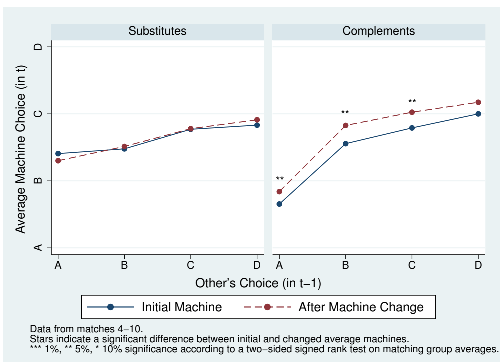


(a) Matches 4–10.

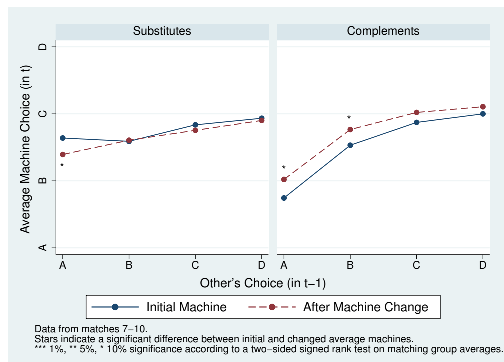


(b) Matches 7–10.

Figure D.4: The effect of machine changes: average changed machine before and after.

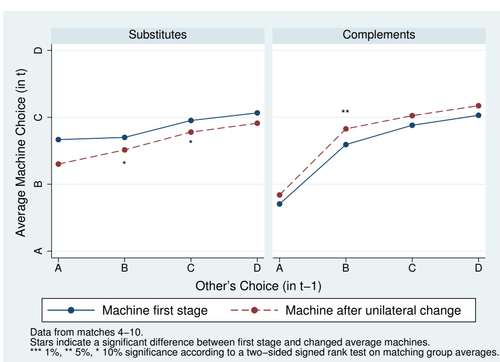


(a) Matches 4–10.

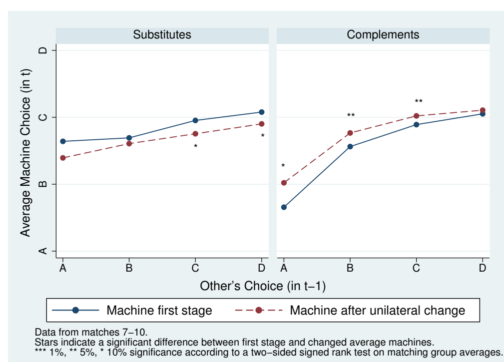


(b) Matches 7–10.

Figure D.5: The effect of machine changes: average stage-one machines versus average changed machines.

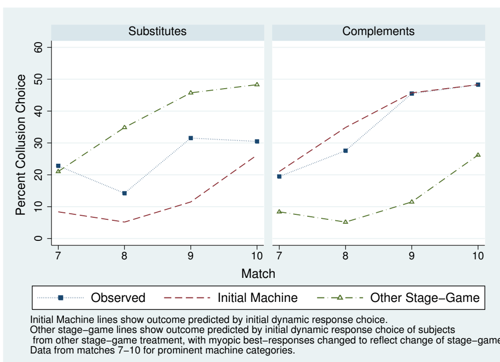


(a) Collusion Choice

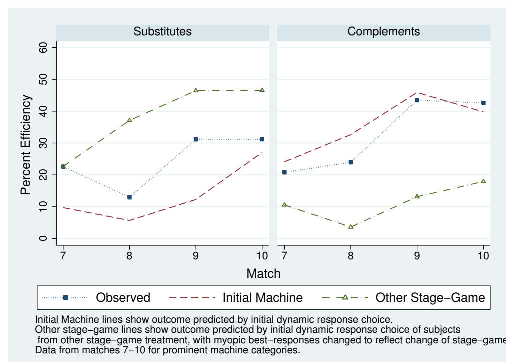


(b) Efficiency

Figure D.6: Observed cooperation measures versus dynamic-response prediction versus simulation – only prominent machine categorisations.

## D.2 Tables

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

| Matching Group | Match 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 |

Table D.2: Machine categorization (in percent): decomposing the "Other" category.

| | | Substitutes | | Complements | |
|---|---|---|---|---|---|
| | Machine | Elicit | Hot | Elicit | Hot |
| Unconditional cooperation | AAAA | 5 | 7 | 8 | 11 |
| Conditional cooperation | ABC(C/D) | 17 | 17 | 23 | 17 |
| Nash reversion | ACCC | 11 | 8 | 16 | 22 |
| Partial collusion + Nash rev. | BBCC | 3 | 0 | 4 | 3 |
| Nash | CCCC | 22 | 15 | 14 | 18 |
| Myopic best reponse | DCCC | 28 | 38 | | |
| | BCCC | | | 21 | 10 |
| Punishing | DDDD | 5 | 2 | 0 | 1 |
| Other | ——— | 9 | 13 | 13 | 18 |

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 D.3: 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 |
| Elicit | 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 |
| Elicit | 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.

Table D.4: Aggregate differences in action choice proportions between hot and elicit.

| | Action Choice | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Collusion (A) | | Dev. SC (B) | | Nash (C) | | Dev. SS (D) | |
| *In all stages:* | | | | | | | | |
| Substitutes | −0.01 | (0.911) | 0.01 | (0.644) | −0.01 | (0.877) | 0.01 | (0.777) |
| Complements | 0.06 | (0.491) | 0.00 | (0.988) | −0.10 | (0.131) | 0.04 | (0.185) |
| *Initial choices only:* | | | | | | | | |
| Substitutes | −0.05 | (0.509) | 0.00 | (0.932) | 0.05 | (0.312) | −0.01 | (0.936) |
| Complements | −0.00 | (0.977) | −0.09* | (0.057) | 0.04 | (0.536) | 0.05 | (0.173) |

Notes: The table uses data for matches 7-10. Significance based on a linear random-effects regression on a treatment indicator with standard errors clustered at the matching group level. *** 1%, ** 5%, *10% significance level.

Table D.5: Linear random-effects regression of payoff efficiency in the stage game.

| | Substitutes | | | | Complements | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | | (2) | | (1) | | (2) | |
| Hot | −0.09 | (0.337) | 0.01 | (0.926) | 0.11 | (0.196) | 0.04 | (0.728) |
| Stage | −0.01*** | (0.000) | −0.01*** | (0.009) | −0.02*** | (0.001) | −0.02*** | (0.000) |
| Hot x Stage | 0.00 | (0.954) | 0.00 | (0.786) | −0.00 | (0.546) | 0.01* | (0.089) |
| Match | | | 0.02** | (0.017) | | | 0.01 | (0.340) |
| Hot x Match | | | −0.01 | (0.589) | | | 0.00 | (0.979) |
| Matches | 7–10 | | 1–10 | | 7–10 | | 1–10 | |

Notes: Notes: Coefficients report the marginal effect on the efficiency of stage-game choices. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%. All regressions include only data for stages 1-12. All regressions include a set of controls for the match-stage composition.

Table D.6: Choice after a one-shot deviation at different prominent machines (matches 7–10).

| | Percent of Action Choices | | | | Number of | |
|---|---|---|---|---|---|---|
| Machine | Collusion | Dev.SC | Nash | Dev.SS | Obs. | Subj. |
| *Substitutes Elicit* | | | | | | |
| AAAA | 0 | 17 | 50 | 33 | 6 | 3 |
| ABC(C/D) | 50 | 50 | 0 | 0 | 6 | 3 |
| ACCC | 50 | 12 | 0 | 38 | 8 | 5 |
| CCCC | 45 | 23 | 0 | 32 | 22 | 8 |
| DCCC | 71 | 0 | 0 | 29 | 7 | 5 |
| Other | 24 | 24 | 28 | 24 | 50 | 11 |
| *Complements Elicit* | | | | | | |
| AAAA | 0 | 0 | 100 | 0 | 2 | 2 |
| ABC(C/D) | 28 | 41 | 24 | 7 | 29 | 12 |
| ACCC | 38 | 12 | 0 | 50 | 8 | 5 |
| BBCC | 67 | 0 | 0 | 33 | 3 | 2 |
| CCCC | 50 | 0 | 0 | 50 | 2 | 2 |
| BCCC | 33 | 50 | 0 | 17 | 12 | 6 |
| Other | 13 | 20 | 47 | 20 | 15 | 9 |

Notes: The percent-of-action-choices columns show the percent of choices that required paying the one-shot cost that correspond to each of the four possible actions. The last two columns give the total number of deviations and the number of distinct individuals deviating.

Table D.7: Machine choice after a machine change by prominent initial machines (matches 7–10).

| Initial Machine | AAAA | ABC (C/D) | ACCC | BBCC | CCCC | DCCC | DDDD | Other | Number of Obs. | Number of Subj. |
|---|---|---|---|---|---|---|---|---|---|---|
| *Substitutes Elicit* | | | | | | | | | | |
| AAAA | 0 | 20 | 0 | 0 | 20 | 40 | 0 | 20 | 5 | 5 |
| ABC(C/D) | 50 | 0 | 0 | 0 | 0 | 0 | 0 | 50 | 2 | 2 |
| ACCC | 0 | 0 | 0 | 0 | 0 | 25 | 25 | 50 | 4 | 4 |
| BBCC | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 100 | 1 | 1 |
| CCCC | 38 | 0 | 8 | 0 | 0 | 0 | 15 | 38 | 13 | 9 |
| DCCC | 0 | 0 | 50 | 12 | 0 | 0 | 0 | 38 | 8 | 5 |
| DDDD | 0 | 0 | 0 | 0 | 67 | 0 | 0 | 33 | 3 | 3 |
| Other | 4 | 4 | 4 | 4 | 4 | 20 | 0 | 60 | 25 | 13 |
| Initial Machine | AAAA | ABC (C/D) | ACCC | BBCC | CCCC | BCCC | DDDD | Other | Number of Obs. | Number of Subj. |
| *Complements Elicit* | | | | | | | | | | |
| AAAA | 0 | 0 | 0 | 0 | 80 | 0 | 0 | 20 | 5 | 3 |
| ABC(C/D) | 0 | 0 | 0 | 40 | 20 | 0 | 0 | 40 | 5 | 5 |
| ACCC | 0 | 0 | 0 | 0 | 0 | 50 | 0 | 50 | 2 | 2 |
| BBCC | 0 | 100 | 0 | 0 | 0 | 0 | 0 | 0 | 2 | 1 |
| CCCC | 0 | 100 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
| BCCC | 0 | 0 | 60 | 0 | 0 | 0 | 0 | 40 | 5 | 4 |
| Other | 0 | 7 | 0 | 0 | 4 | 15 | 4 | 70 | 27 | 14 |

Notes: The percent-of-changed-machines-changed-to columns show the percent of changed machines that were changed to each of the shown machine categories. The last two columns give the total number of machine changes and the number of distinct individuals changing their machine.