Discussion Papers of the
Max Planck Institute for
Research on Collective Goods
2021/11

Hybrid Collusion: Algorithmic Pricing in Human-Computer Laboratory Markets

Hans-Theo Normann
Martin Sternberg

MAX PLANCK
SOCIETY

Discussion Papers of the
Max Planck Institute
for Research on Collective Goods 2021/11

Hybrid Collusion: Algorithmic Pricing in Human-Computer Laboratory Markets

Hans-Theo Normann / Martin Sternberg

May 2021

This version: June 2021

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, D-53113 Bonn
https://www.coll.mpg.de

Hybrid Collusion: Algorithmic Pricing in Human-Computer Laboratory Markets

Hans-Theo Normann
Düsseldorf Institute for Competition Economics (DICE) and Max Planck Institute for Research on Collective Goods, Bonn

Martin Sternberg
Max Planck Institute for Research on Collective Goods, Bonn

June 2021

Abstract: We investigate collusive pricing in laboratory markets when human players interact with an algorithm. We compare the degree of (tacit) collusion when exclusively humans interact to the case of one firm in the market delegating its decisions to an algorithm. We further vary whether participants know about the presence of the algorithm. We find that three-firm markets involving an algorithmic player are significantly more collusive than human-only markets. Firms employing an algorithm earn significantly less profit than their rivals. For four-firm markets, we find no significant differences. (Un)certainty about the actual presence of an algorithm does not significantly affect collusion.

JEL classification: C90, L41.

Keywords: algorithms, collusion, human-computer interaction, laboratory experiments

Contact information. Normann: Düsseldorf Institute for Competition Economics, Heinrich Heine University Düsseldorf, 40225 Düsseldorf, Germany, normann@dice.hhu.de. Sternberg: Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, 53113 Bonn, Germany, sternberg@coll.mpg.de.

Acknowledgements. We are grateful to Christoph Engel, Bernhard Ganglmair, Joe Harrington, Timo Klein, Marcel Schubert, and Frederike Zufall for helpful comments on our paper. We also received useful suggestions from seminar participants at Bayreuth University, the Competition and Regulation Day Mannheim 2020, DICE Winter School 2020, IMPRS Uncertainty Thesis Workshop 2019, 3rd Open Antitrust Doctoral Seminar Düsseldorf 2019, and the MPI Bonn Lab Meeting.

1 Introduction

Algorithms are increasingly taking over price decisions on behalf of the firms which employ them. Whereas pricing algorithms are also used in more traditional brick-and-mortar retailing, for example in supermarkets[1] or gasoline stations,[2] the strongly growing e-commerce[3] adds to their rapid dissemination. In its "E-commerce Sector Inquiry," the EU Commission (2017) reports that a majority of online firms track the prices of competitors and two thirds of them use algorithmic software. Hence, algorithms are on the rise.

A major problem with pricing algorithms is that they could facilitate tacit collusion. Algorithms are ideally suited to dealing with the wealth of data about competitors and customers and can immediately respond to any changes in their environment (see for example Brown and MacKay, 2019; Miklós-Thal and Tucker, 2019). In contrast to explicit agreements, they may leave no trace of collusive intent behind. Recently, it has been demonstrated that new self-learning algorithms manage to collude systematically and in a strikingly efficient manner (Abada and Lambin, 2020; Calvano et al., 2020b, 2021; Klein, 2020; Waltman and Kaymak, 2008). These new algorithms are not designed to collude. Instead, they learn to play sophisticated repeated-game strategies that maximize joint profits without explicitly being instructed to do so.[4] As Ezrachi and Stucke (2017, p. 1775) put it, "we are shifting from the world where executives expressly collude in smoke-filled hotel rooms to a world where pricing algorithms continually monitor and adjust to each other's prices and market data."

[1]See "Surge Pricing Comes To The Supermarket," The Guardian, June 4, 2017, available at: www.theguardian.com/technology/2017/jun/04/surge-pricing-comes-to-the-supermarket-dynamic-personal-data (last accessed on March 11, 2021).

[2]See Assad et al. (2020) and "Why Do Gas Station Prices Constantly Change? Blame the Algorithms," Wall Street Journal, May 8, 2017, available at: www.wsj.com/articles/why-do-gas-station-prices-constantly-change-blame-the-algorithm-1494262674? (last accessed on March 11, 2021).

[3]In 2020, 72% of internet users in the EU ordered goods or services online. See 2020 Eurostat Community Survey on ICT usage in households and by individuals, available at: www.ec.europa.eu/eurostat/statistics-explained/index.php?title=E-commerce_statistics_for_individuals (last accessed on March 11, 2021).

[4]See Klein (2020) for a comprehensive literature survey on self-learning algorithms, especially those that reply on Q-learning. See also Waltman and Kaymak (2008) for an early dynamic programming approach.

The problem of algorithmic pricing and collusion has been pointed out by many in the field and appears high on the agenda of academic scholars and competition authorities alike. Calvano et al. (2020a), Ezrachi and Stucke (2016, 2017), Harrington (2018, 2020), Mehra (2016), Oxera (2017), and others fear that algorithms are a threat to competition. The OECD Competition Committee's round table on "Algorithms and Collusion" concludes that "there is a clear risk that current changes in market conditions may facilitate anticompetitive strategies" (OECD, 2017, p. 24). Concerns about algorithmic collusion are being discussed by antitrust authorities around the world (British Competition and Markets Authority, 2018, 2021; Bundeskartellamt and Autorité de la Concurrence, 2019; Competition Bureau Canada, 2018). Likewise, the German Monopolies Commission (Monopolkommission, 2018) is calling for systematic investigations of markets with algorithmic-based pricing.

Other scholars are less concerned about the collusive potential of algorithms. Schwalbe (2018) argues that, currently, concerns about algorithmic collusion do not seem warranted. Aside from the lengthy time algorithms need to learn, he points out that the algorithms in Calvano et al. (2020b) do not achieve full maximum profits. Kühn and Tadelis (2018) claim that it is hard for algorithms to solve coordination problems without communication. Pointing to the experimental evidence (see below), both papers claim that algorithmic collusion should not raise concerns, as it would be unlikely to sustain collusion with a large number of firms (Kühn and Tadelis, 2018; Schwalbe, 2018).

This paper contributes to this debate by highlighting the importance of "hybrid" interaction between human players and algorithms. The levels of collusion self-learning algorithms can achieve in Abada and Lambin (2020), Calvano et al. (2020b, 2021), Klein (2020), and Waltman and Kaymak (2008) have been obtained when algorithms compete against algorithms. This raises the question of how such collusive algorithms fare against human players. Chen et al. (2016, p. 1348) emphasize that the effects of algorithmic pricing are not yet understood, "especially in heterogeneous markets that include competing algorithmic and non-algorithmic sellers," and indeed, Calvano et al. (2020b) note that useful extensions of their analysis would involve the case of player heterogeneity.

Our main contribution is an analysis of hybrid collusion in a laboratory experiment.[5] We investigate experimental oligopolies when human players interact with an algorithm. We compare the degree of (tacit) collusion when exclusively humans interact to the case when one firm in the market delegates its decisions to an algorithm. Our research question is whether the presence of an algorithm leads to an increase in prices.

To explore the role of human beliefs about algorithms, we further vary whether or not participants know about the presence of the algorithm. Do participants behave differently when they are aware they are facing an algorithm? This may indeed be the case: Farjam and Kirchkamp (2018) show in a laboratory asset market that humans trade differently if they expect algorithmic traders. As a possible explanation, they suggest that human traders perceive the algorithmic traders as behaving more rationally. De Melo et al. (2015) find that people tend to make different decisions depending on whether they are facing a human or a computer algorithm.[6] Thus, it seems warranted to test whether expectations about algorithms influence the behavior of participants.

In order to give collusion the best shot, we opted for a rather simple and transparent experimental design.[7] In three- and four-firm markets,[8]

[5]Collusion is notoriously difficult to detect and this is particularly true for tacit collusion. The impact algorithms have on the degree of collusion in the field is thus not easy to determine. Experimental economics can be useful. Laboratory experiments enable us to study the relevant factors in isolation. They allow for causal inference and thereby contribute to the understanding of economic phenomena such as collusion.

[6]For related findings, see Dijkstra et al. (1998), Weibel et al. (2008), Krach et al. (2008), Lee (2018) and Rilling et al. (2004).

[7]Conducting research on human-machine interaction, Crandall et al. (2018) provide helpful insights into what kind of algorithm might be suitable for experiments on hybrid markets. They show for stochastic two-player games that particularly complex algorithms do not establish cooperative relationships when it comes to human-machine interaction, whereas a non-trivial, but ultimately simple set of algorithmic mechanisms do achieve some level of cooperation.

[8]Experiments with three and four firms seem promising when it comes to identifying collusive effects in that duopolies can be collusive, whereas markets with four or more firms are usually not, see Engel (2015), Fonseca and Normann (2012), Huck et al. (2004), Potters and Suetens (2013). The evidence on cooperation in three-player groups is somewhat inconclusive (and hence a good starting point for us). While Horstmann et al. (2018) do find some collusion in $n = 3$ oligopolies with differentiated goods, Freitag et al. (2020) do not find any supracompetitive outcomes in a multimarket context benign to collusion. Already Marwell and Schmitt (1972) reported that three-person prisoner's dilemmas are substantially less cooperative than two-player experiments. Roux and Thöni (2015) demonstrate that larger oligopolies become collusive only when targeted

participants have two actions (high price, low price) available, so they play an n-player prisoner's dilemma. The algorithm we use is essentially static, a feature it shares with many algorithms in the field (EU Commission, 2017; Monopolkommission, 2018). It is a multiplayer generalization of tit-for-tat (Axelrod, 1984; Hilbe et al., 2015). This algorithm begins by cooperating, but subsequently adapts to the level of cooperation in the market. Using a relatively straightforward algorithm, we concur with Crandall et al. (2018) and the British Competition and Markets Authority (2018), who suspect that complex deep-learning algorithms may fail to collude when competing against humans.

We derive testable hypotheses for human-algorithm interaction with a simple theoretical model. To begin with, standard repeated-game analysis shows that algorithms can be an impediment to collusion. This may seem surprising, but the intuition is straightforward. The algorithm is more forgiving than, say, a grim-trigger player and thus increases the incentive to deviate from the collusive outcome. This raises the minimum discount factor required for collusion to be a subgame-perfect Nash equilibrium. However, if we take the strategic risk the players face in the presence of multiple equilibrium into account (Blonski et al., 2011; Blonski and Spagnolo, 2015; Dal Bó and Fréchette, 2011, 2018; Green et al., 2015), we predict that the algorithm indeed increases the scope for collusion. Our main hypothesis is that replacing a human player with an algorithm increases collusion. We further hypothesize that human players will collude less when they know an algorithm is present.

Our findings are as follows. The data suggest that three-firm markets involving an algorithmic player are significantly more collusive than human-only markets. While this raises profits for all firms in the industry, it turns out that those firms that employ the algorithm earn significantly less profit than their rivals. Thus, one firm must be willing to accept the set-up costs for the collusion. For four-firm markets, we find no significant differences. Beliefs about the actual presence of an algorithm appear not to affect collusion.

punishments are available.

		Other firms' prices		
		p_{high}, p_{high}	p_{high}, p_{low}	p_{low}, p_{low}
Own price	p_{high}	800	0	0
	p_{low}	1,440	720	480

Table 1: Payoff table ($n = 3$ treatments).

2 Experiment

2.1 Design

The stage game underlying the experiment is an n-player prisoner's dilemma, $n \in \{3, 4\}$, framed as a market interaction. Players choose a high price or a low price, so the action set for all players is $\{p_{high}, p_{low}\}$. Payoffs for $n = 3$ are as in Table 1 (which is similar to the one used in the experiment). These payoffs are derived from a Bertrand oligopoly model with inelastic demand and constant marginal costs of production.[9] For $n = 3$ players and actions $p_{high} = 100$ and $p_{low} = 60$, the payoffs in Table 1 result. The payoff table for the four-player treatments can be found below in Section 5.3.

We compare six different treatments. We run four treatments with $n = 3$ players, and we conduct two further treatments with $n = 4$ players. We vary treatments with and without algorithms and treatments with and without information on the presence of the algorithm. See Table 2.

In all experiments, groups of $n \in \{3, 4\}$ participants constitute one market. In the treatments labeled "Human_," there are n human players. In the treatments labeled "Algorithm_," there are $n - 1$ human players and one algorithm. In treatments involving an algorithm, the computer decides on behalf of one human; the n^{th} human is an experimental subject, but he or she is inactive and merely obtains the payoff earned by the algorithm.

[9]Suppose there are $m = 24$ consumers who demand one unit of the good up to a reservation price of 100. Each player can supply all consumers at production costs of zero. The player that charges the lowest price serves all consumers; if several players charge the lowest price, they split the profit equally.

n humans	$n-1$ humans 1 algorithm
Human_Uncertain_3	Algorithm_Uncertain_3
Human_Certain_3	Algorithm_Certain_3
Human_Certain_4	Algorithm_Certain_4

Table 2: Treatment design.

The second treatment dimension indicates whether the participants know the composition of the market.[10] These treatments were done with $n = 3$ players only. In the treatments labeled "_Certain," participants know from the instructions whether or not an algorithm is present. In the "_Uncertain" treatments, the participants do not know if they are part of the Human_Uncertain or the Algorithm_Uncertain treatment, so they do not know whether an algorithm is present. They are merely told that, with a probability of 50%, one of the three subjects' decisions is taken by an algorithm. We conducted the same number of sessions in both treatments. Thus, consistent with the instructions, there was a 50% chance that the participants were in the Algorithm_Uncertain_3 treatment.

The algorithm is programmed to play *proportional tit-for-tat*, or $pTFT$ (Hilbe et al., 2015). We wanted to use an algorithm that is cooperative, forgiving (to some extent), and that cannot be exploited. Such an algorithm is $pTFT$. It is an n-player generalization of tit-for-tat (Axelrod, 1984): Let t be the index for time. The algorithm begins by cooperating in the first period ($t = 0$) and later cooperates proportionally to the number of cooperators in the previous period. Accordingly, $pTFT$ chooses the high price with the following probabilities

$$prob.(p = p_{high}) = \begin{cases} 1 \text{ if } & t = 0 \\ \frac{j}{n-1} \text{ if } & t > 0 \end{cases}$$

where n is the number of players including the algorithm player and $j \in$

[10] Regarding this point, our design is similar to the one in Farjam and Kirchkamp (2018).

$\{0, 1, 2, ..., n-1\}$ is the number of rival players who chose p_{high} in the previous period. Subjects are not told how the algorithm is programmed. Nor are they told the algorithm's purpose.

The treatments are implemented as repeated games, and all treatments have three supergames. The subjects stay in the same market throughout the periods of the supergames. When a new supergame begins, subjects are randomly assigned to a new market. In other words, we have fixed matching within supergames and random matching across supergames. Each supergame lasts at least 20 periods. From the 20th period onward, a random rule with a continuation probability of 7/10 determines whether play continues. The number of periods in all three rounds was determined ex ante and is the same in all sessions (24, 20 and 21 periods for $n = 3$; and 22, 25, 21, for $n = 4$). Subjects knew they would play three supergames from the instructions and they also knew the termination probability.

3 Model and Hypotheses

3.1 Setup

Consider a three-player[11] game where players' action sets are the prices $\{p_{high}, p_{low}\}$. With p_i denoting player i's price, her payoff is generally denoted by $\pi_i(p_i, p_j, p_k)$, $\ i, j, k \in \{1, 2, 3\}$ where p_j and p_k are the prices of the rivals of player i, and $i \neq j$, $i \neq k$ and $j \neq k$ and where the identity of the rival players do not matter, that is, $\pi_i(p_i, p_j, p_k) = \pi_i(p_i, p_k, p_j)$. We further define the following shortcut notation:

$$\begin{aligned}
\pi^c &= \pi_i(p_{high}, p_{high}, p_{high}) = 800 \\
\pi^s &= \pi_i(p_{high}, p_{low}, p_{low}) = \pi_i(p_{high}, p_{low}, p_{high}) = 0 \\
\pi^d &= \pi_i(p_{low}, p_{high}, p_{high}) = 1440 \\
\pi^f &= \pi_i(p_{low}, p_{low}, p_{high}) = 720 \\
\pi^n &= \pi_i(p_{low}, p_{low}, p_{low}) = 480
\end{aligned}$$

[11]In Appendix A.5, we show that the results below also hold for the $n = 4$ case.

The numerical entries are those of the experiment and they are also reproduced in Table 1.

We now analyze an infinitely repeated version of this game. Let time be indexed by $t = 0, ..., \infty$. Future periods are discounted by the factor δ.

3.2 Repeated-game incentive constraint

Suppose the three players attempt to establish collusion on the high price, each following a 'grim-trigger' strategy (GT). When pursuing a GT strategy, a player chooses p_{high} in $t = 0$ and keeps charging p_{high} as long as no player has played p_{low} in any previous period. If any player deviates in t, a GT player charges p_{low}, the static Nash equilibrium price, from $t + 1, ..., \infty$. Expected payoffs are as follows. If player i chooses p_{high} in $t = 0$, she receives π^c from $t = 0, ..., \infty$. If she defects with p_{low}, she obtains π^d in $t = 0$ and, since she triggers the punishment path, π^n in $t = 1, ..., \infty$. Accordingly, playing GT is a subgame-perfect Nash equilibrium (SGPNE) if

$$\begin{aligned} \frac{\pi^c}{1-\delta} &\geq \pi^d + \frac{\delta \pi^n}{1-\delta} \\ \delta &\geq \frac{\pi^d - \pi^c}{\pi^d - \pi^n} = \frac{2}{3} \equiv \underline{\delta}_{GT} \end{aligned} \tag{1}$$

where the subscript GT indicates that all three participants are GT players, there is no algorithm.

Suppose now there are two players attempting to establish collusion via GT and the third player is an algorithm. The algorithm is committed to playing $pTFT$ (as defined above) and will thus not deviate from this strategy.[12] We analyze the incentives of a GT player to deviate. If a GT player chooses p_{high}, she receives π^c in $t = 0, ..., \infty$ in equilibrium. The profit from a one-off deviation is π^d, as before. The punishment payoff in $t = 1$ does change, however. If player i deviates in $t = 0$, the price vector reads $(p_{low}, p_{high}, p_{high})$. This prompts the $pTFT$ algorithm to cooperate with 50% in $t = 1$, so the price vectors $(p_{low}, p_{low}, p_{high})$ and $(p_{low}, p_{low}, p_{low})$

[12] Tit-for-tat strategies are often not subgame-perfect (for two-player cases, see Osborne (2006)). In our case, the algorithm itself will not deviate, as it is programmed to play $pTFT$, even if this is not a best response in some subgames.

and corresponding payoffs π^f and π^n, respectively, are equally likely. From $t = 2$ on, the two players and the algorithm choose p_{low} for the rest of the game. Thus, the incentive constraint becomes

$$\frac{\pi^c}{1-\delta} \quad \geq \quad \pi^d + \delta\left(\frac{\pi^f + \pi^n}{2}\right) + \frac{\delta^2 \pi^n}{1-\delta}$$

solving for δ for the values employed in the experiment[13]

$$\delta \quad \gtrsim \quad 0.69 \quad \equiv \quad \underline{\delta}_{pTFT} \tag{2}$$

where the subscript $pTFT$ indicates that one of the three players is the $pTFT$ algorithm.

To complete the proof for the subgame-perfectness of GT in the presence of the algorithm, consider additional out-of-equilibrium histories.[14] In histories ending in subgames where at least one player chooses p_{low} and at least one player selects the high price, the GT players will choose p_{low}, whereas the $pTFT$ algorithm will cooperate in $t+1$ with at least 50%.[15] A possible one-off deviation for a GT player would be to cooperate in the next period. But since the second GT player will defect in $t+1$, such a deviation would yield zero payoff , whereas sticking to GT (by defecting from period $t + 1$ on) would yield (at least) π^n. It follows that GT is subgame-perfect in the presence of the algorithm, provided (2) is met.

We summarize by comparing (1) and (2):

Proposition 1: The minimum discount factor required for collusion to be a SGPNE is lower for three GT players compared to two GT players and one $pTFT$ algorithm: $\underline{\delta}_{GT} < \underline{\delta}_{pTFT}$.

The intuition behind Proposition 1 is straightforward. GT and $pTFT$ are both cooperative strategies, but $pTFT$ is more forgiving and willing to

[13] A closed-form solution with general payoffs can be found in Appendix A.1.

[14] In the presence of the $pTFT$ algorithm, meeting the incentive constraint (2) is generally not sufficient for GT to be subgame-perfect.

[15] The set of subgames where at least one player deviates and at least one player cooperates includes $(p_{high}, p_{low}, \cdot)$, $(p_{low}, p_{high}, \cdot)$ and $(p_{high}, p_{high}, p_{low})$. In the latter case, the $pTFT$ algorithm cooperates 100% in $t + 1$. In two further possible subgames ($(p_{low}, p_{low}, p_{high})$ and $(p_{low}, p_{low}, p_{low})$), all players defect in $t + 1$, ensuring GT is subgame-perfect.

cooperate with a positive probability even when (exactly) one rival player defected in $t-1$. This raises the payoffs of a GT player after a defection and, accordingly, increases the minimum discount factor required for successful collusion.[16]

3.3 Strategic risk

The inequalities (1) and (2) are necessary conditions for collusion on p_{high} to be subgame-perfect. Other equilibria obviously exist as well. For example, all players always charging p_{low} is also a SGPNE of the repeated game, with and without the presence of the $pTFT$ algorithm. The inequalities (1) and (2) do not reflect the coordination problems players face in the presence of multiple equilibria.

Taking strategic risk into account is especially important when analyzing algorithms. The algorithm is committed to a strategy, whereas humans are not. That is, the algorithm reduces strategic uncertainty.[17] Merely to focus on incentives in a given collusion equilibrium and to ignore strategic risk would imply that we largely miss the collusive impact algorithms may have.

To deal with strategic uncertainty, a growing literature on repeated prisoner's dilemmas (Blonski et al., 2011; Blonski and Spagnolo, 2015; Dal Bó and Fréchette, 2011, 2018; Green et al., 2015) borrows from Harsanyi and Selten's (1988) concept of risk dominance which can easily be applied to symmetric coordination games with two strategies. A strategy is risk-dominant if it is a best response to the other players mixing with equal probability between the two strategies.

The approach can be adapted to repeated games. We follow Blonski et al. (2011), Blonski and Spagnolo (2015), Dal Bó and Fréchette (2011, 2018), and Green et al. (2015) in focusing on a simplified version of the game, the choice between two repeated-game strategies. We henceforth analyze the decision between the collusive GT and the non-cooperative

[16]Nevertheless, results from experiments with self-learning algorithms suggest that these algorithms learn to cooperate even after deviations and therefore pursue a more forgiving strategy than GT, see Calvano et al. (2020b, section IV. C.).

[17]The commitment effect of collusive algorithms is analyzed in the model of Brown and MacKay (2019).

'always defect' strategy (AD). That is, players' action sets are now the repeated-game strategies GT and AD.[18] Provided (1) and (2), respectively, hold, all players playing GT and all players playing AD are equilibria of this two-action game. Increasing δ reduces the riskiness of GT, and we solve for a new critical discount factor, δ^*, such that playing GT is the best response given the other players randomize with equal probability between the two strategies GT and AD. We then investigate how the presence of an algorithm affects δ^*.

Consider three players choosing between GT and AD and expecting their competitors to play GT or AD with equal probability. When playing GT, there are two contingencies for the profit of player i in period $t = 0$: Provided the other two players also play GT (which happens with a probability of $1/4$), i obtains π^c. If at least one other player defects (probability of $3/4$), i obtains $\pi^s = 0$ in period $t = 0$. If all players including i play GT in $t = 0$, i also obtains π^c in all future periods $t = 1, ..., \infty$. If at least one player defects in $t = 0$, i gets π^n in periods $t = 1, ..., \infty$. Thus, player i's expected payoff from playing GT is

$$\frac{1}{4}\left(\frac{\pi^c}{1-\delta}\right) + \frac{3}{4}\left(\pi^s + \frac{\delta\pi^n}{1-\delta}\right)$$

If player i instead plays AD, there are three possibilities. If both other players cooperate in $t = 0$ (which happens with a probability of $1/4$), i obtains π^d. If one rival player cooperates and the other defects (which happens with a probability of $1/2$), i obtains π^f. When both rival players defect (probability of $1/4$), i obtains π^n. In all three cases, i obtains π^n in $t = 1, ..., \infty$. Player i's expected payoff is

$$\frac{\pi^d}{4} + \frac{\pi^f}{2} + \frac{\pi^n}{4} + \frac{\delta\pi^n}{1-\delta}$$

Taking the difference in expected profits of GT and AD and solving for

[18] For the simplified version of the game with only two repeated-game strategies (GT and AD), Blonski and Spagnolo (2015) show that any collusive equilibrium is risk-dominant if GT is risk-dominant.

δ, we find that GT has a higher expected payoff than AD, if and only if

$$\delta \geq \frac{\pi^d + 2\pi^f - 3\pi^s + \pi^n - \pi^c}{\pi^d + 2\pi^f - 3\pi^s} = \frac{8}{9} \approx 0.89 \equiv \delta^*_{GT} \tag{3}$$

where $\delta^* \in (0,1)$ denotes the critical discount factor in the presence of strategic risk and δ^*. GT indicates that all three players are (potential) GT players. Note that $\delta^*_{GT} > \underline{\delta}_{GT}$ strictly. We further point out that the payoffs π^s and π^f play a role here – which is not the case for $\underline{\delta}$.

Now, one of the three market participants is an algorithm committed to playing $pTFT$, whereas the other two participants are rational players as before. We analyze the choice of these two players between GT and AD. The two players are expected to play GT and AD with equal probability. Suppose player i plays GT. Then there are only two contingencies: the other player plays either GT or she plays AD. Expected profits are accordingly

$$\frac{1}{2}\left(\frac{\pi^c}{1-\delta}\right) + \frac{1}{2}\left(\pi^s + \frac{\delta}{2}(\pi^f + \pi^n) + \delta^2 \frac{\pi^n}{1-\delta}\right)$$

If player i plays AD, she gets

$$\frac{1}{2}\left(\pi^d + \frac{\delta}{2}\left(\pi^f + \pi^n\right) + \frac{\delta^2 \pi^n}{1-\delta}\right) + \frac{1}{2}\left(\pi^f + \frac{\delta \pi^n}{1-\delta}\right)$$

We find that GT has a higher expected payoff if

$$\delta \geq \frac{\pi^d + \pi^f - \pi^c - \pi^s}{\pi^d + \pi^f - \pi^n - \pi^s} = \frac{17}{21} \approx 0.81 \equiv \delta^*_{pTFT} \tag{4}$$

Comparing (3) and (4), we obtain:

Proposition 2: The minimum discount factor required in the presence of strategic uncertainty is higher for three GT players compared to two GT players and one $pTFT$ algorithm: $\delta^*_{GT} > \delta^*_{pTFT}$.[19]

The above analysis is based on the presumption that players expect

[19]The reader can verify that $\delta^*_{GT} > \delta^*_{pTFT}$ not only for our experimental parameters, but in general: Note that both the numerator and the denominator of δ^*_{GT} exceed their δ^*_{pTFT} counterparts by $\pi^f + \pi^n - 2\pi^s > 0$, hence are increasing δ^*_{GT}.

the algorithm to play $pTFT$ with probability one. This is plausible when subjects learn during the repetitions of the stage game and the supergame that the algorithm is cooperative. Initially, however, subjects may hold different beliefs, and our experimental design is suitable to identify whether participants' beliefs about the algorithm matter.

To highlight the impact of beliefs, we introduce a general belief $\gamma \in (0, 1]$ as the probability with which players believe the algorithm is playing $pTFT$, such that $1 - \gamma$ is the belief that the algorithm will play AD. That is, players are uncertain of whether the algorithm is cooperative, but, if so, they correctly anticipate $pTFT$ play.[20] In Appendix A.1, we derive the expected payoffs of the players (as above) and solve for δ^*_{pTFT}, now as a function of γ. We show (perhaps unsurprisingly) that – if the algorithm is cooperative – playing GT becomes more attractive for players the higher the belief (γ) is:

$$\frac{\partial \delta^*_{pTFT}(\gamma)}{\partial \gamma} < 0 \tag{5}$$

Proposition 3: Suppose players believe with probability $\gamma \in (0, 1]$ that the algorithm will play $pTFT$, and AD otherwise. Then, in the game with two GT players and one $pTFT$ algorithm, δ^*_{pTFT} decreases in γ.

3.4 Hypotheses

The analysis of the repeated game suggests that an algorithm may affect participants' behavior via its actual behavior or *actions* (Proposition 1 and 2) or via participants' *beliefs* about the algorithm (Proposition 3). The two channels affect behavior differently and give rise to different hypotheses.[21]

[20] If players wrongly believe the algorithm is playing GT (with probability γ), this would not change the result below. However, in an experimental setting, players will quickly learn that the algorithm is more forgiving than GT.

[21] A third channel could be altered other-regarding preferences: Participants may feel inclined to defect when playing with an algorithm, but not with a human participant, especially if the money earned by the algorithm is kept by the experimenter. Our Algorithm$_{-}$ treatments, however, involved three human participants; the profit earned by the algorithm was paid out to a (passive) human participant. Therefore, altered other-regarding preferences should not play a role.

We begin with actions. At least in the long run, human subjects will probably be influenced by the algorithm's actual price-setting behavior and its responses, including the punishments it triggers, and so on. In other words, the algorithm's actions will matter.

Whereas propositions 1 and 2 imply contradicting effects, the existing experimental evidence overwhelmingly suggests that the minimum discount factor that takes strategic risk into account (δ^*) has more explanatory power than the standard minimum discount factor ($\underline{\delta}$). This is shown in the meta-study by Dal Bó and Fréchette (2018). Furthermore, Blonski et al. (2011) highlight the case where, for pairs of treatments, $\underline{\delta}$ and δ^* "change in opposite directions." This is the case in our experiment: the $pTFT$ algorithm increases $\underline{\delta}$, but reduces δ^*. The first experimental result in Blonski et al. (2011) is that, in this case, "the frequency of cooperation changes as predicted by changes in δ^*, contradicting predictions based on $\underline{\delta}$" (Blonski et al., 2011, p. 185).[22] Accordingly, based on Proposition 2, we expect that, ceteris paribus, the Algorithm_ treatments will be more collusive than their Human_ counterparts.

We turn to beliefs. Human subjects may expect the algorithm to play differently from other humans. Responding to this belief, humans adjust their behavior accordingly.[23] But in which direction will the belief be affected?

We hypothesize humans to be skeptical about the play of an algorithmic competitor, so they expect less cooperation when an algorithm is present. News about algorithms beating humans at Chess or Go demonstrate the power of machines in zero-sum games. This may suggest that humans also lose against the algorithm in the market domain – that is, firms run by humans earn less profit. Along these lines, Farjam and Kirchkamp (2018) find that algorithms are perceived as "more rational." This could correspond to skeptical expectations. From the subjects' perspective, "rationality" could imply that the algorithm will attempt to exploit human participants to gain

[22]Their analysis is often based on what they label as "class 2" data. In that class, the actual discount factor is above $\underline{\delta}$, but below δ^*, as is the case in our experiment.

[23]There is ample evidence that human subjects respond to beliefs about the action of others. In the prisoner's dilemma, there are two motives for defection (Ahn et al., 2001; Blanco et al., 2014; Charness et al., 2016). Subjects fear being exploited by others, but some may greedily also want to exploit others themselves.

higher profits through competitive behavior. Reports on competitive algorithmic price wars[24] and the fact that online shopping – often associated with algorithmic pricing – is considered a low-price alternative to brick-and-mortar purchases[25] could give rise to the notion that algorithms are particularly competitive. Furthermore, recall that participants are unaware of the algorithm's strategy, so there is little to suggest that subjects think it is pursuing long-run joint-profit maximization. We thus expect that subjects perceive algorithms as less cooperative than humans, and based on Proposition 3, this should, ceteris paribus, yield lower cooperation rates.

We now put our conjectures together and state our hypotheses about how (tacitly) collusive our treatments will be in comparison. We begin with the influence of the algorithm's action. The _Uncertain treatments are identical in terms of the instructions and the possibility of an algorithm being present, so the beliefs cannot matter. Here, however, the algorithm's actual play may have an impact. From Proposition 2, we hypothesize:

Hypothesis $\mathbf{H_1}$: Cooperation rates in Algorithm_Uncertain_3 are higher than those in Human_Uncertain_3.

For the _Certain treatments, Proposition 2, on the one hand, and skeptical beliefs together with Proposition 3, on the other, imply ambiguous effects of algorithms. We hypothesize that the use of algorithms will have a positive overall impact on collusion because we provide ample evidence of learning (three repeated games). Given these learning opportunities, subjects may update their beliefs and adjust them according to the more

[24]For example, CBNC reports on undercutting competition between Wal-Mart and Amazon through algorithmic pricing: Sarah Whitten,"Wal-Mart Scammed Into Selling PlayStation 4 for $90," CNBC November 18, 2014, available at: https://www.cnbc.com/2014/11/18/wal-mart-scammed-into-selling-playstation-4-for-90.html (last accessed on March 11, 2021); Also, the consultancy Simon-Kucher & Partners reports, in its 2019 Global Pricing Study, that 57% of the companies report they are currently involved in a price war. Available at: https://www.simon-kucher.com/en/about/media-center/global-pricing-study-2019 (last accessed on March 11, 2021).

[25]Prices play an important role in online shopping. Degeratu et al. (2000) find that online promotions are stronger signals for price discounts than offline promotions and the price sensitivity of consumers is higher online. A representative survey of German consumers has also shown that 52% of them are convinced that it is cheaper to buy products online, available at: https://www.bitkom.org/Presse/Presseinformation/Elf-Gruende-fuers-Online-Shopping (last accessed on March 11, 2021).

cooperative behavior of the algorithm. We hypothesize:

Hypothesis H_2: Cooperation rates in Algorithm_Certain_3 are higher than those in Human_Certain_3.

Next, we consider the two Algorithm_ treatments. The algorithm's actions are the same here, but in Algorithm_Certain, subjects know for sure they are facing an algorithm, whereas in Algorithm_Uncertain, they might still be competing with a human. Based on Proposition 3 and our presumption of skeptical beliefs, we hypothesize:

Hypothesis H_3: Cooperation rates in Algorithm_Uncertain_3 are higher than those in Algorithm_Certain_3.

For the Human_ treatments, it is the other way round. The third player is controlled by a human either way, but in Human_Uncertain, participants expect to meet an algorithm with 50% likelihood. So participants should be more optimistic in Human_Certain. We obtain:

Hypothesis H_4: Cooperation rates in Human_Certain_3 are higher than those in Human_Uncertain_3.

Taking hypotheses H_1 to H_4 together, we obtain an unambiguous and testable ranking for the cooperativeness of our treatments. We should observe that the levels of cooperation satisfy

$$Algorithm\_U\_3 > Algorithm\_C\_3 > Human\_C\_3 > Human\_U\_3 \qquad (6)$$

Finally, we look at markets with four competitors. Here, we do not maintain a directed hypothesis regarding Algorithm_Certain_4 and Human_Certain_4. On the one hand, Algorithm_Certain_4 should be more collusive than Human_Certain_4 due to the presence of the algorithm. In A.5 of the appendix, we prove that propositions 1 and 2 also hold for the $n = 4$ case, indicating the algorithm should facilitate collusion. On the other hand, markets with $n = 4$ firms may already be too competitive for significant collusion to occur at all. Collusive outcomes in market experiments are correlated with the number of firms, and prior experiments indicate that four competitors rarely reach collusive conduct (Engel, 2015; Fonseca and Normann, 2012; Horstmann et al., 2018; Huck et al., 2004; Potters and Suetens, 2013). For the sake of completeness, we state the Null here:

Hypothesis $\mathbf{H_5}$: (Null) Cooperation rates in Algorithm_Certain_4 do not differ from those in Human_Certain_4.

4 Procedures

Subjects were recruited from pools of subjects who had previously volunteered to participate in lab experiments. The experiments involved 429 participants in total. None of the subjects participated in more than one session. We had 24 sessions in total, four for each of the six treatments. Due to different market sizes (three and four players) and participants not showing up, session sizes varied between 12 and 30 participants. The experimental sessions were conducted at labs in Düsseldorf and MPI Bonn between August 2019 and October 2020. No sessions were conducted between early March and mid-July 2020, due to the pandemic. Sessions from mid-July 2020 on were conducted under common hygiene rules. See A.6 in the appendix for session details.

Upon arrival at the laboratory, subjects were randomly assigned to a cubicle, using tokens with the cubicle numbers. After a sufficient number of participants had arrived, the experiment started and participants received a hard copy of the instructions in German. While reading the instructions, subjects were allowed to ask questions privately in their cubicles. Afterwards, control questions made sure everyone had understood the task.

The decision-making parts were conducted as follows. We programmed the experiment in z-Tree (Fischbacher, 2007). In each period, the subjects had to decide by clicking a button whether they wanted to set p_{high} or p_{low}. After everyone had decided, an information screen displayed the choices of all three firms in the market and informed subjects about their payoff. At the end of a supergame, the individual overall payoff for that supergame was displayed and the subjects were informed that they would now be assigned to a new market, unless it was the last supergame.

We used an Experimental Currency Unit, where 1,000 ECU corresponded to 1 Euro. One of the three supergames was randomly chosen for payout. At the end of the third supergame, the subjects were informed about the supergame selected for payout and their total earnings.

In the _Uncertain treatments, we further asked participants whether they thought an algorithm was present in the experiment. This was done at the end after the last period of the last supergame. Subjects had to enter a number between zero and 100, expressing how confident they were that an algorithm was in the market. They were paid up to 2 euros for a correct guess: Given a guess $x \in \{0, 1, 2, ..., 100\}$ that an algorithm was present, the payoff was $2x/100$ if an algorithm was actually present and $2 - 2x/100$ if not. (Participants for whom the algorithm decided were paid 1 euro flat instead.)

The sessions lasted for about 60 minutes. The average payment was 16.8 euro, including a show-up fee.

5 Results

5.1 Overview

For the different treatments, Figure 1 shows how cooperation rates develop over time and supergames.[26] Generally, cooperation increases across supergames: In supergame 1, cooperation rates vary roughly between zero and less than 30%, whereas in supergame 3 they vary between 20 and more than 50%. It appears participants learn to collude tacitly with repetitions of the supergame, confirming the results of Bigoni et al. (2015), Dal Bó and Fréchette (2011, 2018), and Fudenberg et al. (2012).

A closer look reveals that cooperation rates improve for all treatments in

[26]To exclude the pronounced restart and endgame effects we observe, we focus on periods 6 to 19. The same graph including all periods can be found in A.7 of the appendix.

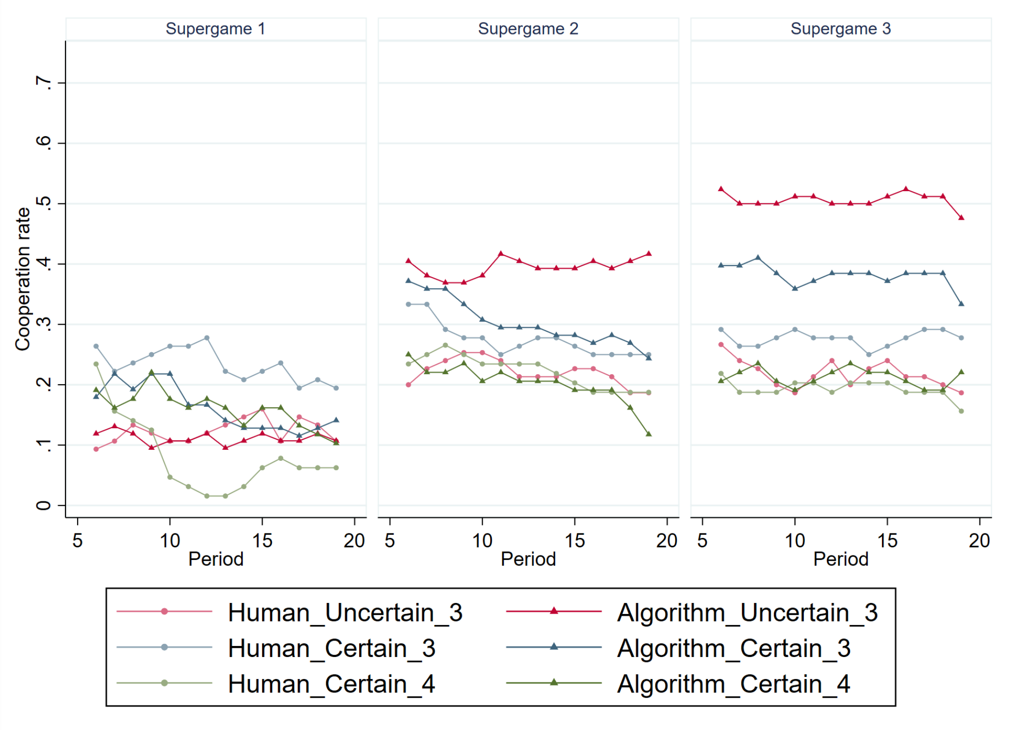

Figure 1: Cooperation rates over time (periods 6 to 19).

supergame 2, but when comparing supergames 2 and 3, only the treatments involving an algorithm increase substantially.[27] This is tentative evidence that the algorithm has a collusive impact.

Complementing Figure 1, Table 3 shows the cooperation rates averaged across periods 6 to 19. We note that the Algorithm_ treatments have higher averages than their Human_ counterparts for all treatments in (almost) all supergames.[28] Taking all supergames into account, the highest cooperation rate is observed in Algorithm_Uncertain_3 (0.337), followed by Algorithm_Certain_3 (0.282) which, in turn, exhibits more cooperation than Human_Certain_3 (0.262). We find higher cooperation in Human_Certain_3 than in Human_Uncertain_3 (0.187). This is exactly the ranking of treatments we hypothesize in (6). Also, in Algorithm_4 (0.191) and Human_4 (0.165), the cooperation rates are lower than in Algorithm_3 and Human_3. This order does not change if we include all periods or focus only on the de-

[27]There is a very minor increase of cooperation in Human_Certain_3 by 0.2 percentage points when comparing supergames 2 and 3. See Table 3.

[28]The exception is that Human_Certain_4 is slightly more cooperative than Algorithm_Certain_4 in supergame 2.

	Supergame 1	*Supergame 2*	*Supergame 3*	*All*
Human_U_3	0.123	0.221	0.218	0.187
	(0.328)	(0.415)	(0.413)	(0.390)
Algorithm_U_3	0.111	0.395	0.506	0.337
	(0.315)	(0.489)	(0.500)	(0.473)
Human_C_3	0.233	0.275	0.277	0.262
	(0.423)	(0.447)	(0.448)	(0.440)
Algorithm_C_3	0.162	0.303	0.381	0.282
	(0.369)	(0.460)	(0.486)	(0.450)
Human_C_4	0.0804	0.222	0.193	0.165
	(0.272)	(0.416)	(0.395)	(0.371)
Algorithm_C_4	0.160	0.202	0.212	0.191
	(0.366)	(0.401)	(0.409)	(0.393)

Standard deviations in parentheses.

Table 3: Average cooperation rates (periods 6 to 19).

cisions of human subjects (that is, if we exclude the algorithms' decisions). See A.7 and A.8 of the appendix for details.

How successful are the firms in actually establishing the collusive outcome? Figure 2 shows the percentages of three outcomes for the six treatments in the last supergame: "successful collusion" indicates (tacit) cooperation – all firms choose p_{high}; "competition" means that all firms charge p_{low}; and "failed collusion" occurs when at least one firm chooses p_{low} and at least one firm tried to collude – this is miscoordination. Again, it becomes clear that successful coordination on the high price occurs more often in Algorithm_Uncertain_3 and Algorithm_Certain_3. The two extremes are Algorithm_Uncertain_3 with a roughly 50% rate of successful collusion, whereas Human_Certain_4 involved almost 80% competition. The share of outcomes with miscoordination (failed collusion) is remarkably small in all treatments, meaning that subjects quickly coordinate on either the cooperative or the competitive outcome. This is also apparent from the quick drop in cooperation in the first five periods (see in A.7 of the appendix).

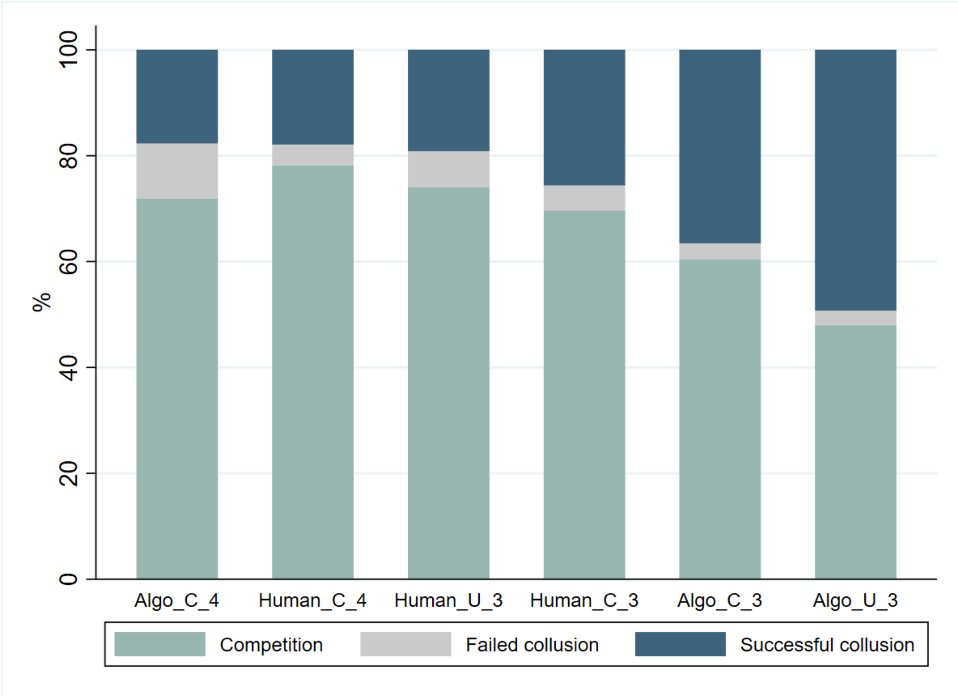

Figure 2: Collusive and competitive outcomes (supergame 3, periods 6 to 19).

We now systematically test our hypotheses and make statistically reliable statements about treatment effects. We begin with the three-firm markets (5.2). The results from markets with four firms follow in Section 5.3. Throughout, we take the possible dependence of observations into account using bootstrapping standard errors at the session level. See Cameron et al. (2008).

5.2 Three-firm treatments

5.2.1 Treatment differences

Table 4 shows the results of a linear probability model and highlights the main treatment effects we observe in the markets with three firms. Our dependent variable is whether or not a firm (participant or algorithm) cooperates in a given period. We include as explanatory variables dummies for the Algorithm_ treatments, the _Certain treatments, the interaction of the two, and for the initial and terminal periods of play. We report the results separately for the three supergames and jointly for all supergames where we add a cardinal variable for supergame. For the regression including all supergames, the constant reflects supergame 1.

The impact of the algorithm is positive and statistically significant from supergame 2 onward. When analyzing them jointly, the two Algorithm_ treatments cooperate better than the two Human_ treatments. The impact of _Certain is positive, small, and insignificant. When we add the interaction $algo \times certain$, the coefficient $algo$ becomes stronger and remains significant. This indicates that the cooperation rates in Algorithm_Uncertain_3 and Human_Uncertain_3 differ significantly (Hypothesis H_1). Comparing the two _Certain_3 treatments separately, we find no significant effect of the algorithm, so no support for Hypothesis H_2.[29]

Result 1. *In the $n = 3$ markets, the Algorithm_ treatments jointly exhibit significantly more collusion than the Human_ treatments. Cooperation rates are significantly higher in Algorithm_Uncertain_3 compared to Human_Uncertain_3. We find no statistically significant effects when comparing Algorithm_Certain_3 and Human_Certain_3.*

We hypothesize that human subjects play more competitively if they knowingly face or expect to face an algorithmic opponent. But neither the comparison of Algorithm_Uncertain_3 and Algorithm_Certain_3 (H_3), nor of Human_Uncertain_3 and Human_Certain_3 (H_4) shows significant effects. This suggests that expectations do not play a major role in these regressions.

Result 2. *We find no statistically significant effects between the _Uncertain_3 and the _Certain_3 treatments.*

One interpretation of Result 2 is that expectations do not matter much when subjects gain experience. Below, we report on period-one data, but even in the first period we cannot find much statistical support regarding differences between _Uncertain_3 and _Certain_3. It appears that even in the very first period of play (first period of the first supergame), when subjects are inexperienced, beliefs do not have a big impact.

[29] Across all supergames, the effect of the algorithm in the _Certain_3 treatments is statistically insignificant ($p > 0.1$).

	Supergame 1		*Supergame 2*		*Supergame 3*		*All*	
algorithm	-0.0225	0.00490	0.108**	0.172***	0.185**	0.255**	0.0847	0.137**
	(0.0528)	(0.0566)	(0.0526)	(0.0654)	(0.0813)	(0.110)	(0.0516)	(0.0556)
certain	0.0706	0.100	0.000116	0.0691	-0.0292	0.0466	0.0166	0.0733
	(0.0542)	(0.0966)	(0.0515)	(0.0839)	(0.0833)	(0.112)	(0.0523)	(0.0818)
algorithm × certain		-0.0564		-0.132		-0.145		-0.108
		(0.117)		(0.0991)		(0.170)		(0.105)
periods 1 to 5	0.0867***	0.0867***	0.129***	0.129***	0.116***	0.116***	0.110***	0.110***
	(0.0234)	(0.0235)	(0.0177)	(0.0173)	(0.0199)	(0.0198)	(0.0133)	(0.0132)
periods 20 to 25	-0.0524***	-0.0524***	-0.0522***	-0.0522***	-0.129***	-0.129***	-0.0821***	-0.0821***
	(0.0159)	(0.0159)	(0.0160)	(0.0159)	(0.0292)	(0.0289)	(0.0108)	(0.0107)
supergame							0.0968***	0.0968***
							(0.0216)	(0.0210)
Constant	0.133***	0.118***	0.245***	0.211***	0.268***	0.231***	0.120***	0.0923***
	(0.0405)	(0.0372)	(0.0518)	(0.0561)	(0.0625)	(0.0531)	(0.0401)	(0.0309)
Obs.	7,416	7,416	6,180	6,180	6,489	6,489	20,085	20,085
R^2	0.025	0.027	0.029	0.033	0.057	0.063	0.062	0.066

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Treatment effects, $n = 3$ variants, all periods, linear probability model.

Support for our beliefs hypothesis can nevertheless be detected. The mean cooperation rates correspond to our hypothesis. For the Algorithm_ treatments, we find (insignificantly) more collusion in _Uncertain than in _Certain in supergames two and three and all supergames. For the Human_ variants, subjects were more collusive in _Certain than in _Uncertain in all supergames, as predicted. Recall that we can rank our $n = 3$ treatments according to our hypotheses, see (6). Creating a cardinal rank variable (with 1 = Algorithm_Uncertain and 4 = Human_Uncertain) in our dataset allows for an ordered alternative hypothesis of the multiple independent samples jointly.[30] We see that this ranking variable has a significantly negative effect on the choice in the third supergame (linear probability model, $p < 0.05$, see Appendix A.2). The effect is also significant for the second supergame and over all supergames (both $p < 0.05$).

Result 3. *Consistent with our hypotheses, the variable* ranking *for the order of competitiveness of the treatments has a significant negative effect on the cooperation rate.*

Another piece of evidence in favor of our hypothesis on expectations comes from the incentivized guess in the _Uncertain treatments. This is what we analyze in detail next.

5.2.2 Beliefs about the presence of an algorithm

In the two _Uncertain treatments, we asked participants in an incentivized manner at the end of the experiment about their beliefs of whether one of the firms was equipped with an algorithm. Subjects had to state a probability (a number between zero and 100) that an "algorithm was present in the experiment."

It turns out that subjects in Human_Uncertain maintain an average belief of 58.05%, whereas those in Algorithm_Uncertain have a belief of 45.87%. Table 5 shows the results of a linear probability regression with data from Algorithm_Uncertain_3 and Human_Uncertain_3 and *algorithm* as an explanatory variable. The variable *algorithm* is significant at $p <$

[30] Similar to the non-parametric Jonckheere-Terpstra test, which is likewise highly significant.

	Guess	
algorithm	-12.19**	-8.663**
	(4.916)	(4.374)
sum miss-coordinated		1.303***
outcomes		(0.361)
Constant	58.06***	43.78***
	(3.920)	(5.216)
Obs.	131	131
R^2	0.027	0.072

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Incentivized guess about the presence of an algorithm in the _Uncertain treatments, linear probability model.

0.05. In other words, participants are significantly *more* inclined to believe an algorithm was present when this was *not* the case.

Result 4. *Guesses about an algorithm being present in the market are significantly lower in Algorithm_Uncertain_3 compared to Human_Uncertain_3.*

One possible explanation for this surprising finding is that participants associate cooperation with human behavior and not an algorithm.[31] Cooperation rates in Algorithm_Uncertain are significantly higher than in Human_Uncertain, and the lower performance of Human_Uncertain is clearly associated with a higher belief of an algorithm being present. When we add, as an explanatory variable to the regressions in Table 5, the number of miscoordinated outcomes (one or two firms chose p_{low}, whereas at least one firm chose p_{high}) which a participant experiences during the entire course of the experiment, this variable is highly significant ($p < 0.01$) and the magnitude of the *algorithm* coefficient decreases, but is still significant ($p < 0.05$). We take this as confirmation that the participants expect the algorithm to be more competitive than humans.

[31] According to Lee (2018), participants rate algorithmic decisions as less fair, trust algorithmic decisions less, and feel less positive about algorithmic decisions when it comes to tasks requiring human skills. With mechanical tasks, the fairness and trustworthiness of algorithms were attributed to their perceived efficiency and objectivity.

5.2.3 Differences between human and algorithmic play

One immediate effect of the $pTFT$ strategy is that it begins a supergame by choosing the high price with probability one, in contrast to the average human subject. Hence, a first attempt at finding differences between humans and the algorithm is to take a closer look at period-one decisions.

Table 6 shows details of regressions similar to those in Section 5.2.1 above, but truncating the data to the first period of each supergame. The algorithm has a substantial and significant effect on the cooperation rate in the first period throughout.[32] This is perhaps not surprising because of the way the algorithm is programmed, but it is important to state this effect formally because of the significance of period-one behavior for overall cooperation.

Result 5. *In the $n = 3$ markets, cooperation rates in the first period are significantly higher in the Algorithm_ treatments compared to the Human_ treatments. This also holds when comparing Algorithm_Uncertain_3 vs. Human_Uncertain_3, and Algorithm_Certain_3 vs. Human_Certain_3 separately.*

In the following periods, the $pTFT$ could also be a driver for collusion for two reasons. First, the algorithm faithfully adheres to collusion (provided both rival firms chose p_{high} previously), whereas some humans occasionally defect. Second, the algorithm is willing to tolerate episodes of defections by one rival by cooperating with 50%, whereas humans might not. The first argument could also turn out to be an impediment to collusion: Provided both rivals chose p_{low} previously, the algorithm will always charge p_{low}, whereas humans may experiment and switch to p_{high}.

[32] The effect is also significant comparing the _Uncertain_3 and _Certain_3 treatments separately. Across all supergames both p-values < 0.01.

	Supergame 1		*Supergame 2*		*Supergame 3*		*All*	
algorithm	0.166***	0.143**	0.160***	0.167**	0.202***	0.225***	0.176***	0.178***
	(0.0446)	(0.0668)	(0.0399)	(0.0724)	(0.0487)	(0.0676)	(0.0375)	(0.0600)
certain	0.0526	0.0278	0.0664*	0.0739	0.0360	0.0606	0.0517	0.0541
	(0.0436)	(0.0817)	(0.0399)	(0.0668)	(0.0485)	(0.0852)	(0.0369)	(0.0715)
algorithm × certain		0.0473		-0.0144		-0.0468		-0.00462
		(0.0889)		(0.0869)		(0.101)		(0.0789)
supergame							0.0777***	0.0777***
							(0.00911)	(0.00896)
Constant	0.321***	0.333***	0.444***	0.440***	0.465***	0.453***	0.332***	0.331***
	(0.0485)	(0.0639)	(0.0407)	(0.0540)	(0.0487)	(0.0634)	(0.0441)	(0.0601)
Obs.	309	309	309	309	309	309	927	927
R^2	0.031	0.031	0.030	0.030	0.043	0.044	0.050	0.050

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6: Cooperation in the first period, $n = 3$ treatments, linear probability model.

Figure 3 is an alluvial flow diagram that illustrates how humans compare to the algorithm with respect to such individual decisions. It is based on decisions by humans only, using data from all $n = 3$ treatments, periods 1 to 19 and all supergames.[33] The figure shows how participants' decisions in period $t - 1$ (left-hand side of the figure) map into market outcomes (middle), and how conditional on these outcomes decisions in period t emerge. The market outcome is defined as the number of p_{high} choices of all players in a market, including the subject herself and possibly the algorithm. Let us be more specific. Humans choose p_{high} at a rate of roughly 30% (light blue segment on the left) and, accordingly, p_{low} at 70% (dark blue segment). Due to the high degree of coordination in markets, outcomes labeled 0 ("all p_{high}") and 3 ("all p_{low}") result most frequently. If the coordination in markets fails, outcomes 1 ("one p_{high}, two p_{low}") and 2 ("two p_{high}, one p_{low}") result. The stream from the gray outcome boxes then indicates how humans decided conditional on outcome. Their own $t - 1$ decision can be identified by the color (light blue for p_{high} and dark blue for p_{low}).

The algorithm always chooses p_{high} if both competitors previously chose p_{high}—how do humans behave here? Overall, it turns out human participants are also highly likely to play p_{high} (92.7%). But there are substantial differences when the own prior choice is taken into account. Provided that they themselves previously played p_{high}, human subjects almost always play p_{high} again (99.1%).[34] When we look at the human subjects who played p_{low} while both their competitors chose p_{high} ("two p_{high}, one p_{low}"), we see that roughly 29.3% cooperate, whereas the algorithm would play 100% p_{high} here, too.

Differences between humans and the algorithm also become apparent in markets with mixed outcomes where one competitor chose p_{high} and the other one p_{low} in $t - 1$. The probability that the algorithm will play cooperatively is 50%, whereas that of the human subjects is only 26.2%.

[33] See in A.8 of the appendix, where we provide the same analysis for the individual treatments. Differences between treatments are minor and insignificant. We dropped the data from period 20 on because we are not specifically interested in the end-game behavior humans exhibit.

[34] In 26 out of 3,007 observations, these subjects chose p_{low}, which is too little to be visible in Figure 3.

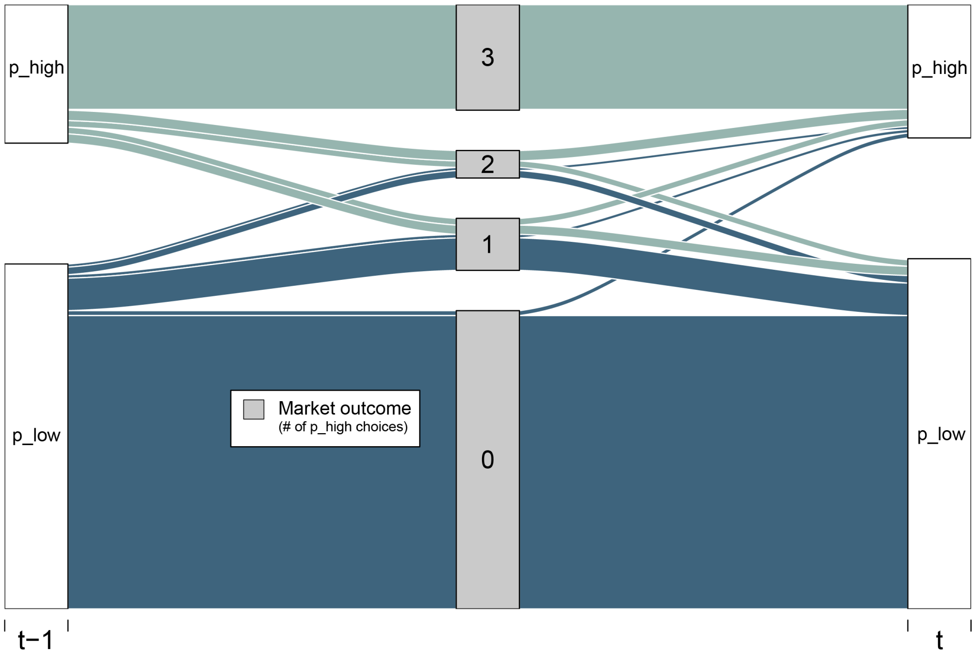

Figure 3: Alluvial flow diagram of choices by human subjects ($n = 3$ treatments, all supergames, periods 1 to 19).

Again, Figure 3 shows the differences between subjects who played p_{high} previously and those who chose p_{low}.[35] The cooperatively playing subjects stuck to their strategy with a probability of 61.2%. But such attempts to establish collusive conduct is hampered by the behavior of competitive rivals who rarely choose the high price (9.8%).

How about the potentially negative effect of the algorithm when both rival firms chose p_{low} previously? In this case, the algorithm would never choose p_{high}. But this does not differ much for human subjects who cooperate with 3.7 %. Conspicuously, the cooperative playing subjects continue their strategy with a relatively high probability (41.9%), while the competitive rivals play *p_high* only in very few cases (1.7%).

Overall, the probability of successful collusion, irrespective of the previous market outcome, is higher in Algorithm_ (27.4%) than in Human_ treatments (18.3%). The algorithm is less cooperative than the human subjects when it comes to attempts to establish a collusive outcome, but

[35] For two-player prisoner's dilemma experiments, Breitmoser (2015) suggests that subjects play a "semi-grim" strategy, such that subjects randomize across choices regardless of their own previous choice.

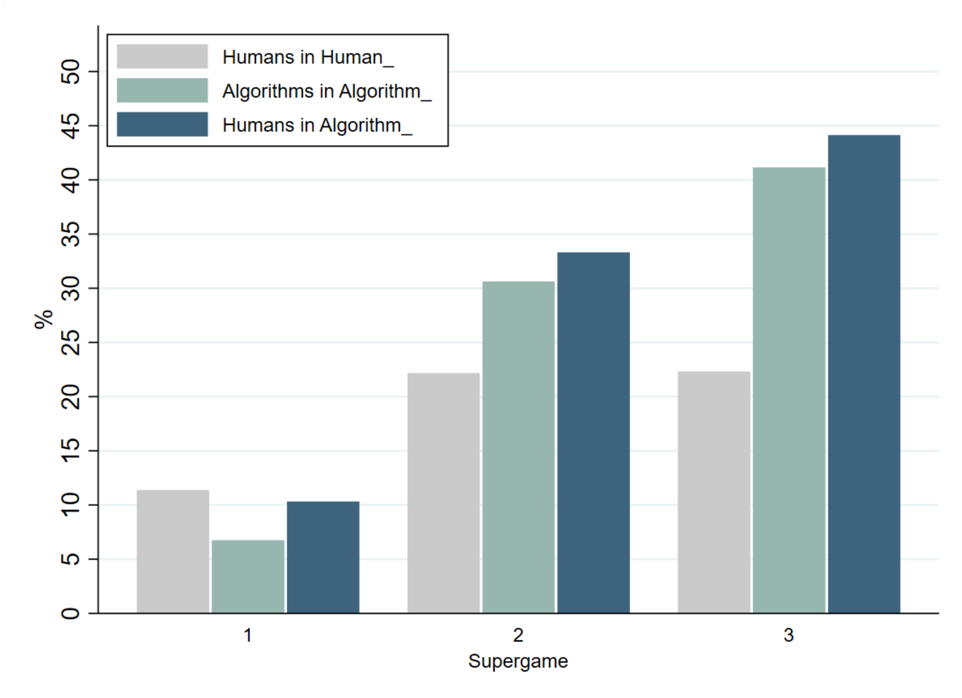

Figure 4: Profits in percent above Nash ($n = 3$ treatments, periods 6 to 19).

much more cooperative than subjects who chose p_{low} before. It seems that the human subjects rarely modify their strategy, trying instead to avoid a change in their price decision.

5.2.4 Profits

If the algorithm variants exhibit more cooperation, this suggests that all firms benefit in terms of higher profits. As we see more cooperation, the mean profits in Algorithm_ are actually higher than in Human_, so subjects earn more if an algorithm is present. In Appendix A.3, we analyze this systematically. The positive effect on profits is significant in the third supergame (linear probability model, $p < 0.05$).

By distinguishing between humans and algorithms, we can analyze who benefits most from the presence of the algorithm. Figure 4 measures profits relative to static Nash earnings (0%) and to perfect collusion (100%).[36] We see that subjects equipped with an algorithm earn substantially less than their competitors in every supergame. Taking all supergames into

[36] Formally, the index in Figure 4 is defined as $(\pi - \pi^n)/(\pi^c - \pi^n)$, where π is the observed profit.

account, the difference is statistically significant ($p < 0.01$).[37] Although the algorithm helps to increase the group's profit, it performs significantly worse than their competitors. This suggests a coordination problem in that no firm wants to adopt the algorithm first.

Result 6. *In the $n = 3$ markets, profits are significantly higher in the Algorithm_ treatments compared to the Human_ treatments. In the Algorithm_ treatments, participants represented by an algorithm earn significantly less than participants who decide themselves for their firm.*

5.3 Four-firm treatments

We now turn to the experiments with $n = 4$ firms. The payoffs underlying these sessions can be found in Table 7. Other than that, procedures are virtually identical to the three-firm treatments.

		Other firms' prices			
		All p_{high}	$p_{high}, p_{high}, p_{low}$	$p_{high}, p_{low}, p_{low}$	*All* p_{low}
Own price	p_{high}	600	0	0	0
	p_{low}	1,440	720	480	360

Table 7: Payoff table ($n = 4$ treatments)

Does the collusive effect of the algorithm extend to four-firm markets? On the one hand, the algorithm should promote collusion. On the other hand, we see that already three human subjects are finding it difficult to cooperate. Does it help when we add an algorithm as a fourth player?

Across all supergames, we indeed find more cooperation in Algorithm_Certain_4 (0.191) than in Human_Certain_4 (0.165). But the difference is small and the effect of the algorithm is not significant, see Appendix

[37] Appendix A.3 provides the results of a linear probability regression where the dependent variable is the profit subjects earn from period 6 to 19. Our explanatory variable for the type of player is role. The negative effect of role is also significant in the first ($p < 0.1$) and the third supergame ($p < 0.001$). The effect is not significant for the second supergame.

A.4. In Figure 2, we see that Algorithm_Certain_4 and Human_Certain_4 share similarly low weight on outcomes with successful collusion.[38] Consistent with Hypothesis H_5, we do not see much of a difference between the two treatments with $n = 4$ firms.

Result 7. *Cooperation rates in Algorithm_Certain_4 are not significantly higher than in Human_Certain_4.*

6 Conclusion

In this paper, we analyze the impact of algorithms on collusion in hybrid markets where humans interact with algorithms. The analysis of human-computer interaction is important because most markets in the field are heterogeneous and firms cannot be sure of whether their opponents are using algorithms for their pricing decision, nor do they know which type of algorithm competitors might use.[39] A recent and growing literature (Abada and Lambin, 2020; Calvano et al., 2021, 2020b; Klein, 2020; Waltman and Kaymak, 2008) shows that markets with exclusively firms using algorithmic pricing can become collusive. This raises the question of whether algorithms also have collusive impact when they interact with humans.

We study these issues in experimental markets with three or four firms where one firm is equipped with an algorithm. The algorithm, if present, plays proportional tit-for-tat (Axelrod, 1984; Hilbe et al., 2015) – a simple and cooperative strategy. We further vary whether the human participants know (in a non-deceptive way) about the presence of the algorithm. Participants of the experiments played three indefinitely repeated games.

We find that an algorithm significantly increases collusive play in the three-firm markets, but the effect disappears with four firms. Not knowing whether an algorithm is in play does not affect pricing much either way. The firms for which the algorithm decided earn significantly less profit.

[38]Subjects equipped with an algorithm still earn less than their competitors. The negative effect of role for the treatments with $n = 4$ firms is significant in the first ($p < 0.05$) and across all supergames ($p < 0.1$). The effect is not significant for the second and third supergame.

[39]Explicitly communicating and agreeing on the use of algorithms has been penalized as a violation of cartel law. See Poster Cartel case: US Department of Justice, Apr. 6, 2015, Press Release no. 15-421.

One possible rejoinder to our results is that replacing a human (who may or may not have collusive intentions) with a somewhat collusive algorithm must inevitably raise prices, at least weakly. Our rebuttal to this rejoinder follows four points. First, even if our results were bland in that sense, demonstrating this collusive effect empirically is still important because many people and competition authorities suspect that this is precisely what is happening in the field at the moment. One would hardly be concerned if human decision-makers were replaced with more competitive algorithms. Second, we reiterate that our algorithm is not ferociously committed to cooperating, but instead adapts to the level of collusion in the market. In fact, ignoring the initial periods of play, the cooperation rates of humans and algorithms in our data are virtually identical. Third, our data show that the collusive impact occurs only in three-firm markets and not when there are four players. Fourth, the fact that firms represented by an algorithm earn less profit suggests that firms want their rivals to adopt the algorithm first. In other words, firms face a severe coordination problem when it comes to delegating decisions to algorithms. Tacit collusion seems feasible if a firm is willing to accept the setup costs. All this suggests that the collusive impact of algorithms is not a foregone conclusion, but our data likewise indicate the anti-competitive potential algorithms have, even when interacting with humans.

Our results suggest promising topics for future research. Given that we do not observe higher levels of collusion with four firms, one avenue would be to further explore markets with more firms and a higher number of algorithms. Larger groups may also successfully collude if the share or the number of humans is not too high. Another extension would be to not impose the use of the algorithm exogenously (as we did), but to let subjects choose whether they want to employ algorithms. "Algorithm aversion" (Dietvorst et al., 2015) may kick in here. Demonstrating the force of algorithms may cure this reluctance; however, the aforementioned coordination problem might still be significant.

References

Abada, I. and Lambin, X. (2020). Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided? *Working Paper*.

Ahn, T. K., Ostrom, E., Schmidt, D., Shupp, R., and Walker, J. (2001). Cooperation in PD Games: Fear, Greed, and History of Play. *Public Choice*, 106:137–155.

Assad, S., Clark, R., Ershov, D., and Xu, L. (2020). Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. *CESifo Working Paper*, No. 8521.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Bigoni, M., Casari, M., Skrzypacz, A., and Spagnolo, G. (2015). Time Horizon and Cooperation in Continuous Time. *Econometrica*, 83(2):587–616.

Blanco, M., Engelmann, D., Koch, A. K., and Normann, H. T. (2014). Preferences and Beliefs in a Sequential Social Dilemma: A Within-Subjects Analysis. *Games and Economic Behavior*, 87:122–135.

Blonski, M., Ockenfels, P., and Spagnolo, G. (2011). Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence. *American Economic Journal: Microeconomics*, 3(3):164–192.

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

Breitmoser, Y. (2015). Cooperation, But No Reciprocity: Individual Strategies in the Repeated Prisoner's Dilemma. *American Economic Review*, 105(9):2882–2910.

British Competition and Markets Authority (2018). Pricing Algorithms: Economic Working Paper on the Use of Algorithms to Facilitate Collusion and Personalised Pricing.

British Competition and Markets Authority (2021). Algorithms: How They Can Reduce Competition and Harm Consumers.

Brown, Z. and MacKay, A. (2019). Competition in Pricing Algorithms. *Harvard Business School Working Paper*, no. 20-067.

Bundeskartellamt and Autorité de la Concurrence (2019). Algorithms and Competition. Discussion Paper.

Calvano, E., Calzolari, G., Denicolò, V., Harrington, J. E., and Pastorello, S. (2020a). Protecting Consumers from Collusive Prices due to AI: Price-Setting Algorithms can Lead to Noncompetitive Prices, but the Law is Ill Equipped to Stop It. *Science*, 370(6520):1040–1042.

Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2020b). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10):3267–3297.

Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2021). Algorithmic Collusion with Imperfect Monitoring. *International Journal of Industrial Organization (Pre-Proof)*.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90(3):414–427.

Charness, G., Rigotti, L., and Rustichini, A. (2016). Social Surplus Determines Cooperation Rates in the One-Shot Prisoner's Dilemma. *Games and Economic Behavior*, 100:113–124.

Chen, L., Mislove, A., and Wilson, C. (2016). An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. *WWW '16: Proceedings of the 25th International Conference on World Wide Web*, pages 1339–1349.

Competition Bureau Canada (2018). Big Data and Innovation: Key Themes for Competition Policy in Canada. Technical report, Competition Bureau Canada.

Crandall, J. W., Oudah, M., Tennom, Ishowo-Oloko, F., Abdallah, S., Bonnefon, J. F., Cebrian, M., Shariff, A., Goodrich, M. A., and Rahwan, I. (2018). Cooperating with Machines. *Nature Communications*, 9:233.

Dal Bó, P. and Fréchette, G. R. (2011). The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence. *American Economic Review*, 101(1):411–429.

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

De Melo, C. M., Gratch, J., and Carnevale, P. J. (2015). Humans versus Computers: Impact of Emotion Expressions on People's Decision Making. *IEEE Transactions on Affective Computing*, 6(2):127–136.

Degeratu, A. M., Rangaswamy, A., and Wu, J. (2000). Consumer Choice Behavior in Online and Traditional Supermarkets: The Effects of Brand Name, Price, and other Search Attributes. *International Journal of Research in Marketing*, 17(1):55–78.

Dietvorst, B. J., Simmons, J. P., and Massey, C. (2015). Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err. *Journal of Experimental Psychology: General*, 144(1):114–126.

Dijkstra, J. J., Liebrand, W. B., and Timminga, E. (1998). Persuasiveness of Expert Systems. *Behaviour and Information Technology*, 17:155–163.

Engel, C. (2015). Tacit Collusion: The Neglected Experimental Evidence. *Journal of Empirical Legal Studies*, 12(3):537–577.

EU Commission (2017). Final Report on the E-commerce Sector Inquiry. Technical report, European Commission, Brussels.

Ezrachi, A. and Stucke, M. E. (2016). Virtual Competition. *Journal of European Competition Law & Practice*, 7(9):585–586.

Ezrachi, A. and Stucke, M. E. (2017). Artificial Intelligence & Collusion: When Computers Inhibit Competition. *University of Illinois Law Review*, 2017(1):1775–1811.

Farjam, M. and Kirchkamp, O. (2018). Bubbles in Hybrid Markets: How Expectations About Algorithmic Trading Affect Human Trading. *Journal of Economic Behavior and Organization*, 146:248–269.

Fischbacher, U. (2007). Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10:171–178.

Fonseca, M. A. and Normann, H. T. (2012). Explicit vs. Tacit Collusion-The Impact of Communication in Oligopoly Experiments. *European Economic Review*, 56(8):1759–1772.

Freitag, A., Roux, C., and Thöni, C. (2020). Communication and Market Sharing: An Experiment on the Exchange of Soft and Hard Information. *International Economic Review*, 62(1):175–198.

Fudenberg, D., Rand, D. G., and Dreber, A. (2012). Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World. *American Economic Review*, 102(2):720–749.

Green, E. J., Marshall, R. C., and Marx, L. M. (2015). Tacit Collusion in Oligopoly. *The Oxford Handbook of International Antitrust Economics*, 2:464–522.

Harrington, J. E. (2018). Developing Competition Law for Collusion By Autonomous Artificial Agents. *Journal of Competition Law and Economics*, 14(3):331–363.

Harrington, J. E. (2020). Third Party Pricing Algorithms and the Intensity of Competition. *Working Paper*.

Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*.

Hilbe, C., Wu, B., Traulsen, A., and Nowak, M. A. (2015). Evolutionary Performance of Zero-Determinant Strategies in Multiplayer Games. *Journal of Theoretical Biology*, 374:115–124.

Horstmann, N., Krämer, J., and Schnurr, D. (2018). Number Effects and Tacit Collusion in Experimental Oligopolies. *Journal of Industrial Economics*, 66(3):650–700.

Huck, S., Normann, H. T., and Oechssler, J. (2004). Two are Few and Four Are Many: Number Effects in Experimental Oligopolies. *Journal of Economic Behavior and Organization*, 53(4):435–446.

Klein, T. (2020). Assessing Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing. *Tinbergen Institute Discussion Paper TI 2018-056/VII*.

Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., and Kircher, T. (2008). Can Machines Think? Interaction and Perspective Taking with Robots Investigated via fMRI. *PLoS ONE*, 3(7):e2597.

Kühn, K.-U. and Tadelis, S. (2018). The Economics of Algorithmic Pricing: Is Collusion Really Inevitable? *Unpublished*.

Lee, M. K. (2018). Understanding Perception of Algorithmic Decisions: Fairness, Trust, and Emotion in Response to Algorithmic Management. *Big Data and Society*, 5(1).

Marwell, G. and Schmitt, D. R. (1972). Cooperation in a Three-Person Prisoner's Dilemma. *Journal of Personality and Social Psychology*, 21(3):376–383.

Mehra, S. K. (2016). Antitrust and the Robo-Seller: Competition in the Time of Algorithms. *Minnesota Law Review, Paper No. 2015-15*, 100.

Miklós-Thal, J. and Tucker, C. (2019). Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers? *Management Science*, 65(4):1552–1561.

Monopolkommission (2018). Algorithmen und Kollusion. *Twenty-second Biennial Report by the German Monopolies Commission: Competition 2018*, pages 62–87.

OECD (2017). Algorithms and Collusion: Competition Policy in the Digital Age. Technical report, OECD.

Osborne, M. J. (2006). Strategic and Extensive Games. *University of Toronto, Department of Economics, Working Papers tecipa-231*.

Oxera (2017). When Algorithms Set Prices: Winners and Losers. *Oxera Consulting LLP. Discussion paper*, pages 1–37.

Potters, J. and Suetens, S. (2013). Oligopoly Experiments in the Current Millennium. *Journal of Economic Surveys*, 27(3):439–460.

Rilling, J. K., Sanfey, A. G., Aronson, J. A., Nystrom, L. E., and Cohen, J. D. (2004). The Neural Correlates of Theory of Mind Within Interpersonal Interactions. *NeuroImage*, 22(4):1694–1703.

Roux, C. and Thöni, C. (2015). Collusion Among Many Firms: The Disciplinary Power of Targeted Punishment. *Journal of Economic Behavior and Organization*, 116:83–93.

Schulz, J., Sunde, U., Thiemann, P., and Thöni, C. (2019). Selection Into Experiments: Evidence from a Population of Students. *Lund University, Department of Economics, Working Papers 2019:18.*

Schwalbe, U. (2018). Algorithms, Machine Learning, and Collusion. *Journal of Competition Law & Economics*, 14(4):568–607.

Waltman, L. and Kaymak, U. (2008). Q-Learning Agents in a Cournot Oligopoly Model. *Journal of Economic Dynamics and Control*, 32(10):3275–3293.

Weibel, D., Wissmath, B., Habegger, S., Steiner, Y., and Groner, R. (2008). Playing Online Games Against Computer- vs. Human-Controlled Opponents: Effects on Presence, Flow, and Enjoyment. *Computers in Human Behavior*, 24(5):2274–2291.

A Appendix

A.1 Proofs omitted in the main text

General closed-form solution for $\underline{\delta}_{pTFT}$

The minimum discount factor required for successful collusion in (2) has a general closed-form solution. The incentive constraints reads

$$\frac{\pi^c}{1-\delta} \quad \geq \quad \pi^d + \delta\left(\frac{\pi^f + \pi^n}{2}\right) + \frac{\delta^2 \pi^n}{1-\delta}$$

which is quadratic in δ. Solving for δ and taking the positive root, we obtain

$$\frac{\pi^f + \pi^n - 2\pi^d + \sqrt{z}}{2(\pi^f - \pi^n)}$$

where $z = 4\left(\pi^d\right)^2 + \left(\pi^f\right)^2 + (\pi^n)^2 - 8(\pi^c\pi^f) + 4(\pi^d\pi^f) + 8(\pi^c\pi^n) - 12(\pi^d\pi^n) + 2(\pi^f\pi^n)$. Plugging in the numerical values from the experiment, we obtain $\delta \gtrsim 0.69 \;\equiv\; \underline{\delta}_{pTFT}$ as in (2).

Derivation of equation (5)

Players choose between GT and AD and expect others to play GT and AD with equal probability. But here one of the three market participants is an algorithm committed to play either $pTFT$ or AD. Players believe with probability γ that the algorithm will play $pTFT$, so $1-\gamma_a$ is the belief that the algorithm is playing AD.

Suppose player i plays GT. Then, her expected profits are

$$\begin{aligned}
&\frac{\gamma}{2}\left(\frac{\pi^c}{1-\delta}\right) + \frac{\gamma}{2}\left(\pi^s + \frac{\delta}{2}\pi^f + \frac{\delta}{2}\pi^n + \delta^2\frac{\pi^n}{1-\delta}\right) \\
&+\frac{1-\gamma}{2}\left(\pi^s + \frac{\delta\pi^n}{1-\delta}\right) + \frac{1-\gamma}{2}\left(\pi^s + \frac{\delta\pi^n}{1-\delta}\right) \qquad (7)
\end{aligned}$$

If player i plays AD, she gets

$$\frac{\gamma}{2}\left(\pi^d + \frac{\delta}{2}\left(\pi^f + \pi^n\right) + \frac{\delta^2\pi^n}{1-\delta}\right) + \frac{\gamma}{2}\left(\pi^f + \frac{\delta\pi^n}{1-\delta}\right)$$
$$+\frac{1-\gamma}{2}\left(\pi^f + \frac{\delta\pi^n}{1-\delta}\right) + \frac{1-\gamma}{2}\left(\pi^n + \frac{\delta\pi^n}{1-\delta}\right) \quad (8)$$

The expected payoff from GT exceeds the one from AD if

$$\delta \geq \frac{\gamma(-\pi^c + \pi^d - \pi^n + \pi^s) + \pi^f + \pi^n - 2\pi^s}{\gamma(\pi^d - 2\pi^n + \pi^s) + \pi^f + \pi^n - 2\pi^s} \;\equiv\; \delta^*_{pTFT}(\gamma)$$

We take the derivative with respect to γ and obtain

$$\frac{\partial\delta^*_{pTFT}(\gamma)}{\partial\gamma} = \frac{-\left(\pi^c - \pi^n\right)\left(\pi^f + \pi^n - 2\pi^s\right)}{\left(\gamma(\pi^d - 2\pi^n + \pi^s) + \pi^f + \pi^n - 2\pi^s\right)^2} \;<\; 0$$

which is the statement in equation (5) as claimed.

A.2 Treatment ranking as explanatory variable

	Supergame 1	*Supergame 2*	*Supergame 3*	*All*
treatment ranking	0.00432	-0.0556***	-0.0879**	-0.0439**
	(0.0217)	(0.0203)	(0.0342)	(0.0191)
periods 1 to 5	0.0867***	0.129***	0.116***	0.110***
	(0.0234)	(0.0177)	(0.0199)	(0.0133)
periods 20 to 25	-0.0524***	-0.0522***	-0.129***	-0.0821***
	(0.0159)	(0.0160)	(0.0292)	(0.0108)
supergame				0.0968***
				(0.0216)
Constant	0.145**	0.438***	0.566***	0.280***
	(0.0568)	(0.0447)	(0.103)	(0.0521)
Obs.	7,416	6,180	6,489	20,085
R^2	0.016	0.033	0.062	0.065

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 8: Impact of treatment ranking, $n = 3$, periods 6 to 19, linear probability model.

A.3 Profits

	Supergame 1		*Supergame 2*		*Supergame 3*		*All*	
algorithm	62.63	116.0	832.5***	872.5***	1,348**	1,393**	714.8***	761.2***
	(274.8)	(274.6)	(300.2)	(303.6)	(600.2)	(599.1)	(265.9)	(266.5)
certain	491.2	491.2	249.6	249.6	289.1	289.1	351.6	351.6
	(539.1)	(539.1)	(385.1)	(385.1)	(591.3)	(591.3)	(402.5)	(402.5)
algorithm × certain	-330.3	-330.3	-770.9*	-770.9*	-857.9	-857.9	-636.3	-636.3
	(633.9)	(633.9)	(454.0)	(454.0)	(905.8)	(905.8)	(538.5)	(538.5)
role		-160*		-120		-133.3***		-139.1***
		(91.29)		(78.87)		(14.63)		(50.78)
supergame							521.1***	521.1***
							(106.7)	(106.7)
Constant	6,989***	6,989***	7,590***	7,590***	7,578***	7,578***	6,867***	6,867***
	(188.7)	(188.7)	(252.8)	(252.8)	(285.9)	(285.9)	(112.3)	(112.3)
Obs.	7,416	7,416	6,180	6,180	6,489	6,489	20,085	20,085
R^2	0.014	0.016	0.022	0.023	0.059	0.059	0.065	0.066

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 9: Total profits, $n = 3$ treatments, periods 6 to 19, linear probability model.

A.4 Treatment effects for n = 4 variants

	Supergame 1	*Supergame 2*	*Supergame 3*	*All*
algorithm	0.0725	0.0117	0.00604	0.0296
	(0.0873)	(0.118)	(0.0951)	(0.0910)
periods 1 to 5	0.136***	0.160***	0.164***	0.153***
	(0.0206)	(0.0368)	(0.0316)	(0.0170)
periods 20 to 25	-0.0455*	-0.153***	-0.0855***	-0.101***
	(0.0266)	(0.0526)	(0.0266)	(0.0247)
supergame				0.0413***
				(0.0158)
Constant	0.0839***	0.206**	0.200***	0.122***
	(0.0201)	(0.0891)	(0.0490)	(0.0418)
Obs.	2,904	3,300	2,772	8,976
R^2	0.042	0.066	0.034	0.051

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 10: Treatment effects, $n = 4$ variants, all periods, linear probability model.

Additional Material

A.5 Four-player model

Setup

Consider a four-player infinitely repeated game with $t = 0, ..., \infty$. Let the discount factor be δ. Players' action sets are prices $p_i \in \{p_{high}, p_{low}\}$ as before. Payoffs are generally denoted by

$$\pi_i(p_i, p_j, p_k, p_l)$$

where p_j, p_k and p_l are the prices of the rivals of player i, as in the three-firm case. We define

$$\begin{aligned}
\pi^c &= \pi_i(p_{high}, p_{high}, p_{high} \cdot p_{high}) = 600 \\
\pi^s &= \pi_i(p_{high}, p_{low}, \ \cdot \ ) = 0 \\
\pi^d &= \pi_i(p_{low}, p_{high}, p_{high}, p_{high}) = 1440 \\
\pi^f &= \pi_i(p_{low}, p_{low}, p_{high}, p_{high}) = 720 \\
\pi^g &= \pi_i(p_{low}, p_{low}, p_{low}, p_{high}) = 480 \\
\pi^n &= \pi_i(p_{low}, p_{low}, p_{low}, p_{low}) = 360 \qquad (9)
\end{aligned}$$

where $\cdot$ denotes any price, high or low.

Repeated-game incentive constraint n=4

Suppose the four players attempt to establish collusion on the high price by following a grim-trigger (GT) strategy. Playing GT is an SGPNE if

$$\begin{aligned}
\frac{\pi^c}{1-\delta} &\geq \pi^d + \frac{\delta \pi^n}{1-\delta} \\
\delta &\geq \frac{\pi^d - \pi^c}{\pi^d - \pi^n} = \frac{1440 - 600}{1440 - 360} = \frac{7}{9} = 0.778 \equiv \underline{\delta}_{GT} \qquad (10)
\end{aligned}$$

where the subscript GT indicates that all four players are GT players.

Suppose now there are three players attempting to establish collusion via GT and the fourth player is the $pTFT$ algorithm. We analyze the incentives of a GT player to deviate. If a GT player chooses p_{high}, she receives:

$$\begin{aligned}
\frac{\pi^c}{1-\delta} &\geq \pi^d + \delta\left(\frac{\pi^f}{3} + \frac{2\pi^n}{3}\right) + \frac{\delta^2\pi^n}{1-\delta} \\
\frac{600}{1-\delta} &\geq 1440 + \delta\left(\frac{480}{3} + \frac{2\cdot 360}{3}\right) + \frac{\delta^2 360}{1-\delta} \\
\delta &\gtrsim 0.784 \equiv \underline{\delta}_{pTFT}
\end{aligned}$$

We see that, as in the three-player case, the incentive constraint for cooperation to be SGP is more severe when one player is a $pTFT$ algorithm. We summarize

Proposition 4. The minimum discount factor required for collusion to be an SGPNE is lower for four *GT* players compared to three *GT* players and one $pTFT$ algorithm: $\underline{\delta}_{GT} < \underline{\delta}_{pTFT}$.

Strategic risk

Four *GT* players

Consider four players choosing between *GT* and *AD*. Player i's expected payoff from playing *GT* is

$$\begin{aligned}
& \left(\frac{1}{2}\right)^3\left(\frac{\pi^c}{1-\delta}\right) + \left(1-\left(\frac{1}{2}\right)^3\right)\left(\pi^s + \frac{\delta\pi^n}{1-\delta}\right) \\
= & \left(\frac{1}{2}\right)^3\left(\frac{600}{1-\delta}\right) + \left(1-\left(\frac{1}{2}\right)^3\right)\left(0 + \frac{\delta 360}{1-\delta}\right) \\
= & \frac{15}{1-\delta}(21\delta + 5)
\end{aligned}$$

If player i plays AD in period $t = 0$, player i's expected payoff is

$$\begin{aligned}
& \left(\frac{1}{2}\right)^3\pi^d + 3\left(\frac{1}{2}\right)^3\pi^f + 3\left(\frac{1}{2}\right)^3\pi^g + \left(\frac{1}{2}\right)^3\pi^n + \frac{\delta\pi^n}{1-\delta} \\
= & \left(\frac{1}{2}\right)^3 1440 + 3\left(\frac{1}{2}\right)^3 720 + 3\left(\frac{1}{2}\right)^3 480 + \left(\frac{1}{2}\right)^3 360 + \frac{\delta 360}{1-\delta} \\
= & \frac{1}{1-\delta}(675 - 315\delta)
\end{aligned}$$

Comparing the two

$$
\begin{aligned}
\frac{15}{1-\delta}(21\delta+5) &\geq \frac{1}{1-\delta}(675-315\delta) \\
\delta &\geq \frac{20}{21}=0.952
\end{aligned}
$$

Three GT players, one *pTFT* algorithm

We now compare the GT versus the AD strategy when one of the four players is an algorithm committed to playing $pTFT$ with probability one. Suppose player i plays GT. Then her expected profits are

$$
\begin{aligned}
&= \left(\frac{1}{2}\right)^2\left(\frac{\pi^c}{1-\delta}\right)+2\left(\frac{1}{2}\right)^2\left(\pi^s+\delta\left(\frac{2\pi^g}{3}+\frac{\pi^n}{3}\right)+\frac{\delta^2\pi^n}{1-\delta}\right) \\
&\quad +\left(\frac{1}{2}\right)^2\left(\pi^s+\delta\left(\frac{\pi^f}{3}+\frac{2\pi^n}{3}\right)+\frac{\delta^2\pi^n}{1-\delta}\right) \\
&= \left(\frac{1}{2}\right)^2\left(\frac{600}{1-\delta}\right)+2\left(\frac{1}{2}\right)^2\left(\delta\left(\frac{2\cdot 480}{3}+\frac{360}{3}\right)+\frac{\delta^2 360}{1-\delta}\right) \\
&\quad +\left(\frac{1}{2}\right)^2\left(\delta\left(\frac{480}{3}+\frac{2\cdot 360}{3}\right)+\frac{\delta^2 360}{1-\delta}\right) \\
&= \frac{10}{1-\delta}\left(-5\delta^2+32\delta+15\right)
\end{aligned}
$$

If player i plays AD in period $t=0$, her expected payoff is

$$
\begin{aligned}
&\left(\frac{1}{2}\right)^2\left(\pi^d+\delta\left(\frac{2\pi^g}{3}+\frac{\pi^n}{3}\right)+\frac{\delta^2\pi^n}{1-\delta}\right)+2\left(\frac{1}{2}\right)^2\left(\pi^f+\delta\left(\frac{\pi^g}{3}+\frac{2\pi^n}{3}\right)+\frac{\delta^2\pi^n}{1-\delta}\right) \\
&+\left(\frac{1}{2}\right)^2\left(\pi^g+\frac{\delta\pi^n}{1-\delta}\right) \\
= &\left(\frac{1}{2}\right)^2\left(1440+\delta\left(\frac{2\cdot 480}{3}+\frac{360}{3}\right)+\frac{\delta^2 360}{1-\delta}\right) \\
&+2\left(\frac{1}{2}\right)^2\left(720+\delta\left(\frac{480}{3}+\frac{2\cdot 360}{3}\right)+\frac{\delta^2 360}{1-\delta}\right) \\
&+\left(\frac{1}{2}\right)^2\left(480+\frac{\delta 360}{1-\delta}\right) \\
= &\frac{40}{1-\delta}\left(21-\delta^2-11\delta\right)
\end{aligned}
$$

Comparing the two, we obtain

$$\begin{aligned}\frac{10}{1-\delta}\left(-5\delta^2+32\delta+15\right) &\geq \frac{40}{1-\delta}\left(21-\delta^2-11\delta\right)\\ \delta &\geq 0.919\end{aligned}$$

which implies

Proposition 5. The minimum discount factor required in the presence of strategic uncertainty is higher for four GT players compared to three GT players and one $pTFT$ algorithm: $\delta^*_{GT} > \delta^*_{pTFT}$.

A.6 Session details

Date	**Session**	**# Part.**	**# Markets**	**Lab**[a]	**COVID-19**
August 28, 2019	31	21	7	Düsseldorf	0
September 11, 2019	41	24	8	Düsseldorf	0
September 11, 2019	32	21	7	Düsseldorf	0
September 12, 2019	42	18	6	Düsseldorf	0
February 19, 2020	51	16	4	Düsseldorf	0
February 19, 2020	52	20	5	Düsseldorf	0
March 04, 2020	33	24	8	Bonn	0
March 05, 2020	11	21	7	Bonn	0
March 05, 2020	12	21	7	Bonn	0
March 05, 2020	21	30	10	Bonn	0
July 06, 2020	22	18	6	Düsseldorf	1
July 15, 2020	61	20	5	Düsseldorf	1
August 05, 2020	13	18	6	Düsseldorf	1
August 17, 2020	62	20	5	Düsseldorf	1
September 02, 2020	23	15	5	Düsseldorf	1
September 22, 2020	43	18	6	Bonn	1
September 22, 2020	63	12	3	Bonn	1
September 22, 2020	53	16	4	Bonn	1
October 13, 2020	44	18	6	Bonn	1
October 13, 2020	64	12	3	Bonn	1
October 13, 2020	54	16	4	Bonn	1
October 14, 2020	24	12	4	Bonn	1
October 14, 2020	34	18	6	Bonn	1
October 16, 2020	14	12	4	Düsseldorf	1

Table 11: Session details

[a]"Corona" indicates that the session was conducted under the common hygiene rules of the pandemic. As a show-up fee, the participants received 5 euros in Bonn and 4 euros in Düsseldorf. During the COVID-19 pandemic, the fee was increased to 8 euros in Dsseldorf from mid-July 2020 on. This in line with Schulz et al. (2019), who find that moderately different show-up fees had no influence on the behavior of the participants.

A.7 Overview using data from all periods

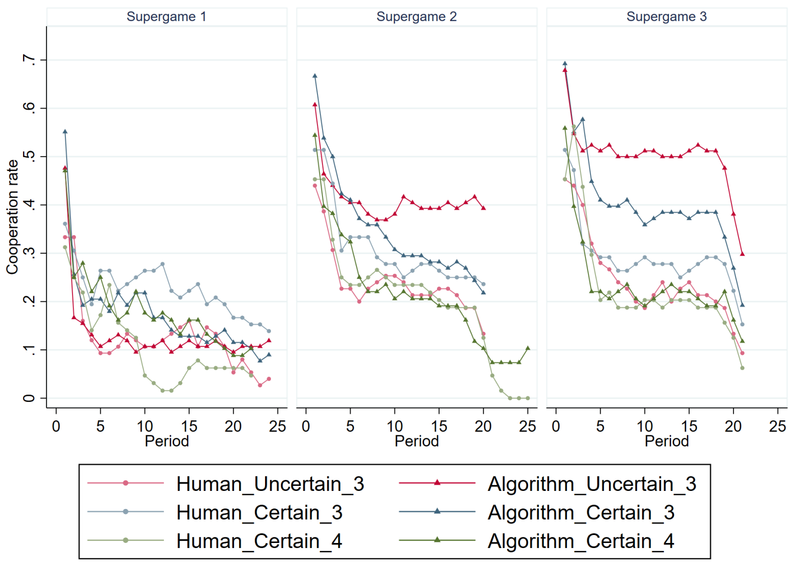

Figure 5: Cooperation rates over time (all periods).

	Supergame 1	*Supergame 2*	*Supergame 3*	*All*
Human_U_3	0.126	0.241	0.246	0.200
	(0.331)	(0.428)	(0.431)	(0.400)
Algo_U_3	0.130	0.412	0.502	0.337
	(0.337)	(0.492)	(0.500)	(0.473)
Human_C_3	0.226	0.310	0.293	0.273
	(0.418)	(0.463)	(0.455)	(0.446)
Algo_C_3	0.174	0.350	0.404	0.302
	(0.379)	(0.477)	(0.491)	(0.459)
Human_C_4	0.109	0.201	0.231	0.180
	(0.311)	(0.401)	(0.421)	(0.384)
Algo_C_4	0.181	0.212	0.237	0.210
	(0.385)	(0.409)	(0.425)	(0.407)

Standard deviations in parentheses.

Table 12: Average cooperation rates (all periods).

A.8 Data of human subjects only

	Supergame 1	*Supergame 2*	*Supergame 3*	*All*
Human_U_3	0.123	0.221	0.218	0.187
	(0.328)	(0.415)	(0.413)	(0.390)
Algorithm_U_3	0.112	0.395	0.504	0.337
	(0.316)	(0.489)	(0.500)	(0.473)
Human_C_3	0.233	0.275	0.277	0.262
	(0.423)	(0.447)	(0.448)	(0.440)
Algorithm_C_3	0.159	0.297	0.378	0.278
	(0.366)	(0.457)	(0.485)	(0.448)
Human_C_4	0.0804	0.222	0.193	0.165
	(0.272)	(0.416)	(0.395)	(0.371)
Algorithm_C_4	0.155	0.197	0.214	0.189
	(0.363)	(0.398)	(0.411)	(0.392)

Standard deviations in parentheses.

Table 13: Average cooperation rates (human subjects, periods 6 to 19).

	Rival behavior in $t-1$			
	Two Low	*High/Low*	*Two High*	*Total*
Human_U_3	0.0401	0.267	0.908	0.207
	(0.196)	(0.443)	(0.289)	(0.405)
Algorithm_U_3	0.0291	0.238	0.953	0.343
	(0.168)	(0.427)	(0.211)	(0.475)
Human_C_3	0.0405	0.258	0.948	0.278
	(0.197)	(0.438)	(0.221)	(0.448)
Algorithm_C_3	0.0347	0.281	0.886	0.301
	(0.183)	(0.450)	(0.318)	(0.459)
Total	0.0370	0.262	0.927	0.276
	(0.189)	(0.440)	(0.260)	(0.447)

Standard deviations in parentheses.

Table 14: Average cooperation rates with respect to the previous choices of rivals 1 and 2 (human subjects, $n = 3$ treatments, period 6-19).

A.9 Treatments effect with probit model

	Supergame 1		*Supergame 2*		*Supergame 3*		*All*	
algorithm	-0.0907	0.0284	0.304*	0.490**	0.509**	0.704**	0.258	0.436**
	(0.223)	(0.293)	(0.156)	(0.203)	(0.235)	(0.311)	(0.169)	(0.172)
certain	0.294	0.406	0.00233	0.210	-0.0760	0.147	0.0610	0.257
	(0.224)	(0.356)	(0.149)	(0.258)	(0.237)	(0.360)	(0.166)	(0.274)
algorithm × certain		-0.220		-0.378		-0.400		-0.357
		(0.491)		(0.296)		(0.501)		(0.337)
periods 1 to 5	0.320***	0.321***	0.349***	0.350***	0.309***	0.311***	0.320***	0.322***
	(0.0968)	(0.0973)	(0.0493)	(0.0492)	(0.0578)	(0.0572)	(0.0420)	(0.0410)
periods 20 to 25	-0.252***	-0.254***	-0.160***	-0.163***	-0.395***	-0.398***	-0.322***	-0.325***
	(0.0841)	(0.0841)	(0.0555)	(0.0562)	(0.0798)	(0.0806)	(0.0483)	(0.0476)
supergame							0.300***	0.301***
							(0.0647)	(0.0633)
Constant	-1.122***	-1.184***	-0.687***	-0.792***	-0.625***	-0.737***	-1.103***	-1.206***
	(0.175)	(0.169)	(0.159)	(0.185)	(0.186)	(0.183)	(0.147)	(0.119)
Obs.	7,416	7,416	6,180	6,180	6,489	6,489	20,085	20,085

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 15: Treatment effects, $n = 3$ variants, probit model.

A.10 Robustness check for impact of lab location and COVID-19 pandemic

	Supergame 1		*Supergame 2*		*Supergame 3*		*All*	
algorithm	-0.0225	-0.00359	0.108**	0.178*	0.185**	0.255	0.0847	0.136
	(0.0528)	(0.107)	(0.0526)	(0.106)	(0.0813)	(0.196)	(0.0516)	(0.109)
certain	0.0706	0.0346	0.000116	0.0712	-0.0292	0.0457	0.0166	0.0494
	(0.0542)	(0.114)	(0.0515)	(0.138)	(0.0833)	(0.219)	(0.0523)	(0.132)
algorithm × certain		0.0174		-0.138		-0.142		-0.0821
		(0.169)		(0.168)		(0.294)		(0.176)
corona[a]		-0.0277		0.0177		-0.00476		-0.00634
		(0.0928)		(0.0823)		(0.165)		(0.0881)
laboratory[b]		0.0368		-0.0110		0.0182		0.0161
		(0.0747)		(0.0733)		(0.131)		(0.0730)
periods 1 to 5	0.0867***	0.0896***	0.129***	0.124***	0.116***	0.116***	0.110***	0.110***
	(0.0234)	(0.0236)	(0.0177)	(0.0175)	(0.0199)	(0.0210)	(0.0133)	(0.0135)
periods 20 to 25	-0.0524***	-0.0512***	-0.0522***	-0.0534***	-0.129***	-0.134***	-0.0821***	-0.0842***
	(0.0159)	(0.0161)	(0.0160)	(0.0176)	(0.0292)	(0.0306)	(0.0108)	(0.0114)
supergame							0.0968***	0.105***
							(0.0216)	(0.0206)
Constant	0.133***	0.114	0.245***	0.208**	0.268***	0.224	0.120***	0.0795
	(0.0405)	(0.106)	(0.0518)	(0.102)	(0.0625)	(0.165)	(0.0401)	(0.101)
Obs.	7,416	7,128	6,180	5,940	6,489	6,237	20,085	19,305
R^2	0.025	0.025	0.029	0.034	0.057	0.064	0.062	0.074

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 16: Laboratory location and COVID-19 effects, n = 3 variants, linear probability model.

[a]corona = 1 if the session was conducted under hygiene rules of the pandemic.
[b]laboratory = 1 if the sessions was run in Bonn.

A.11 Instructions (_Uncertain treatments)

Welcome to the experiment

Thank you for your participation in this experiment. Please read the instructions carefully. For your participation in today's experiment, you will receive 5 euros. During the experiment, you will have the opportunity to earn an additional amount of money. The additional amount will depend on your decisions and the decisions of the other participants. A short questionnaire will follow the experiment. From now on, please stop any conversations with your neighbors. Turn off your cell phone and remove everything from your table that you do not need for the experiment. If you have any questions, please raise your hand and we will answer them one-on-one.

Instructions

In this experiment, you will take the role of a firm in a market. Each market consists of three firms. Each of the three firms is represented by a human participant. All firms offer 24 units of a comparable good with no cost of production, and with 24 consumers demanding one unit of the good. Consumers' willingness to pay for a good ranges from 1 to 100 ECU (Experimental Currency Units), where 1,000 ECU = 1 Euro. At the beginning of each period, all firms have the option to set a high price (100 ECU) or a low price (60 ECU) for their good. The company which alone has set the lowest price serves the entire demand. All other companies will not sell any of their units. If several companies have set the same lowest price, the demand is divided equally among them. The following three examples illustrate the mechanism of the market:

Example 1

You are firm A and you decide to charge a high price for the units of your good (100 ECU). Firm B makes the same decision, whereas C sets a low price (60 ECU). Firm C now has the cheapest sales offer and will serve the complete demand. Accordingly, firm C will earn (60 ECU $*24$ units sold =) $1,440$ ECU. Firms A and B will not sell any units and will therefore

	Both competitors choose the high price	One competitor chooses the high price, the other competitor chooses the low price	Both competitors choose the low price
You choose the high price (100 ECU)	800 ECU	0 ECU	0 ECU
You choose the low price (60 ECU)	1440 ECU	720 ECU	480 ECU

earn 0 ECU in this period.

Example 2

You are firm A and you decide to charge a low price for the units of your good (60 ECU). Firms B and C make the same decision. Firms A, B, and C have now all made the lowest sales offer and will each sell 1/3 of the demand, thus $24/3 = 8$ units of their goods. Accordingly, each firm will earn (60 ECU $*8$ units sold =) 480 ECU.

Example 3

You are firm A and you decide to charge a high price for the units of your good (100 ECU). Firms B and C make the same decision. Firms A, B, and C have now all made the most favorable sales offer and will each sell 1/3 of the demand, thus $24/3 = 8$ units of their goods. Accordingly, each firm will earn (100 ECU $*8$ units sold =) 800 ECU. Thus, your earnings depend on your own and the other firms' pricing decisions. This results in the following profit table for you:

After all the firms have made their choice, you will be informed about the chosen prices of the other two firms and about your profit.

Periods and rounds

In total, you will play at least 20 periods with the other two firms. Random chance will decide whether or not additional periods will be played in the sequel. With a probability of 70% the round will continue with another period; with a probability of 30% the round will end. The round continues until random chance determines the end. In each period of a round, you will be playing with the same participants in a market. At the end of these $20 + x$ periods, all participants are randomly assigned to new markets and a new round begins. The three participants in the new markets will then stay together again for $20 + x$ periods.

In total, you will play three rounds of $20+x$ periods. After three rounds, the experiment ends and a short questionnaire follows.

Market decisions by algorithms

In your markets, at least two participants decide for themselves the price for which they want to sell their goods for their firm and are paid the profit their firm makes in cash at the end of the experiment. **With 50% probability, the decisions for the third firm will also be made by one participant. Also with 50% probability, the third firm will be equipped with an algorithm in all rounds, which will make the necessary pricing decisions for the participant.** In this case, the participant does not make any decisions but still gets paid in cash the profit that her firm makes.

Payout

For your payout, one of the three rounds will be randomly selected. The ECU earned there will be paid to you additionally in euros.