

ORIGINAL PAPER

Algorithmic price recommendations and collusion: Experimental evidence

Matthias Hunold¹  and Tobias Werner² 

¹Faculty III, Universität Siegen, Unteres Schloß 3, Siegen, NW, Germany

²Center for Humans and Machines, Max Planck Institute for Human Development, Lentzeallee 94, Berlin, BE, Germany

Corresponding author: Tobias Werner; Email: tobias.felix.werner@gmail.com

(Received 31 August 2023; Revised 23 July 2024; Accepted 4 April 2025)

Abstract

This paper investigates the collusive and competitive effects of algorithmic price recommendations on market outcomes. These recommendations are often non-binding and common in many markets. We develop a theoretical framework and derive two algorithms that recommend collusive pricing strategies. Utilizing a laboratory experiment, we find that sellers condition their prices on the recommendation of the algorithms. The algorithm with a soft punishment strategy lowers market prices and has a pro-competitive effect. The algorithm that recommends a subgame perfect equilibrium strategy increases the range of market outcomes, including more collusive ones.

Keywords: Bertrand Oligopoly; collusion; experiment; human–machine interaction

JEL Codes: C92; D43; L13; L41

1. Introduction

We explore the potential of non-binding algorithmic price recommendations to facilitate tacit collusion among three sellers in a Bertrand game. The algorithms recommend collusive pricing strategies and vary in the severity of punishment phases after deviations from the recommended price. We focus on whether these recommendations can coordinate pricing behavior despite being non-binding and allowing for deviations.

There are various examples of non-binding price recommendations by a common intermediary in practice. For instance, various platforms provide sellers on their marketplaces with an algorithmic recommendation on how to set the price for the products they sell based on historical market data¹. While it can increase market efficiency, platforms can have incentives to foster tacit collusion in the downstream market, depending on their monetization scheme, as they usually receive a share of the revenues from the seller market². Besides sales platforms, there exists a market for third-party

¹For instance, Airbnb uses price recommendation algorithms that utilize historical and geographical data and combine machine learning methods with human intuition. Furthermore, the algorithmic price recommendation changes daily for the upcoming dates. See Hill (2015) for details.

²For instance, price recommendations might also reduce information asymmetries between the platform and sellers, which makes them potentially attractive from a business perspective beyond collusion. For instance, platforms may have better demand information than individual sellers, and they can use price recommendations to transmit this information

providers of price recommendations, where the same developer provides the same recommendation algorithm to multiple competing firms. Those developers can have an incentive to make the recommendations from the algorithm less competitive as it increases the profits of the supplied firms and thereby increases the demand for the algorithm itself (Harrington, 2022).

Competition authorities are concerned that price recommendation algorithms by a common intermediary can dampen competition as it might make coordination between sellers easier (Competition and Markets Authority, 2021, Bundeskartellamt & Autorité de la concurrence, 2019). For example, according to reporters of ProPublica, price recommendation software allegedly led to coordination effects in the U.S. rental market, especially in regions where a few property managers control a large share of the apartments³. Moreover, recent investigations by the Federal Trade Commission and the U.S. Department of Justice suggest that hotels in New Jersey were using the same third-party price recommendation tools to coordinate their behavior (Federal Trade Commission and U.S. Department of Justice, 2024). The hotels argued that the recommendations were non-binding, but enforcers noted that even non-binding recommendations can facilitate collusion.

To investigate the effects of price recommendations on the prices of competing sellers, we derive two rule-based algorithms from economic theory and behavioral insights⁴. These algorithms aim at increasing sellers' profits and recommend collusive strategies with different punishment mechanics. They provide the same price recommendation to all sellers in a market during a given round and adapt their recommendations across rounds based on the past pricing decisions of the sellers. The recommendations are non-binding and do not change the game's action space and payoff functions, as they do not provide fundamentally new information. In practice, however, collusive outcomes can be challenging to achieve without communication among competitors or another form of coordination (see, for instance, Fonseca & Normann, 2012).

The theory-based algorithm recommends actions consistent with a collusive trigger strategy and Nash reversion. If a seller undercuts the collusive price level, it recommends competitive prices for several periods to all sellers until it returns to recommending collusive prices. All firms following the recommendations constitutes a subgame perfect Nash equilibrium. Additionally, we consider an algorithm that is motivated by behavioral findings whereby firms often do not use harsh punishment strategies (see, for instance, Wright, 2013). This algorithm recommends brief punishment phases with prices at the level of the deviating price and returns to recommending high prices when all sellers comply.

The subjects in our laboratory experiments resemble competing sellers who repeatedly set prices and receive price recommendations from an algorithm in each round. Across treatments, we vary whether participants receive a recommendation or not, as well as the type of recommendation algorithm. Subjects in the experiment are informed that the recommendation algorithm aims to symmetrically maximize sellers' profits without favoring any particular seller.

The algorithmic price recommendations positively influence individual pricing decisions in the sense that higher recommended sales prices induce sellers to set higher individual prices. The estimated "pass-on rate" from recommended prices to sales prices is between 0.22 and 0.57, depending on the recommendation algorithm. The pass-on rate is higher for the theory-based algorithm recommending a collusive trigger strategy with temporary Nash reversal.

The effects on the realized market prices and profits differ sharply between the distinct recommendation algorithms. We find insightful price patterns for the theory-based algorithm even though

(see Pavlov & Berman, 2019, Lefez, 2021). For a discussion on the overall incentives of platforms to foster collusion in the seller market, see, for instance, Teh (2022), Schlütter (2022), and Online Appendix D.

³See Vogell, Coryne & Little, "Technology Rent Going Up? One Company's Algorithm Could Be Why," <https://www.propublica.org/article/yieldstar-rent-increase-realtor-rent>, last accessed March 17, 2025.

⁴Recent studies by Musolff (2022) and Hanspach et al. (2024) highlight that price algorithms on sales platforms often follow simple rule-based logic. Similarly, even complex reinforcement learning algorithms often converge to strategies that simple rules can describe (see, for instance, Werner, 2022).

the average market prices do not differ from the control treatment without any price recommendation. The substantial heterogeneities can explain the absence of an average treatment effect in market outcomes. The collusive effect of the algorithm depends on the seller's characteristics. In markets where sellers have low levels of negative reciprocity, the recommendation algorithm decreases market prices. Thus, if participants are usually unwilling to punish unfair behavior, the recommendation lowers market prices.

The behaviourally motivated algorithm also recommends the monopoly price but differs in the reaction to the deviation of a seller. For this algorithm, we find lower market prices than without any recommendation. Participants repeatedly deviate downwards from the recommendation, which triggers a downward spiral that leads to lower prices. We find no evidence that this algorithm fosters collusion for any subgroup. This is particularly interesting against the backdrop of observations where humans prefer soft punishments for deviations from collusion in experiments (see, for instance, Wright, 2013).

Related literature.

Our article relates to the literature on the collusive effects of algorithmic pricing. There exists evidence that algorithms can foster collusion and lead to anti-competitive prices (Klein, 2021, Calvano et al., 2020, Hansen et al., 2021, Brown & MacKay, 2023). Johnson et al. (2023) focus on tacit collusion among self-learning algorithms on sales platforms and discuss how the platform's design choices influence it. Ezrachi and Stucke 2017 and Harrington (2022) discuss the possibility that multiple competing firms use the same pricing algorithm and how it can lead to seller coordination. Normann and Sternberg (2023) and Werner (2022) show experimentally that algorithms may raise market prices even above the price level usually observed in human markets. We differ from this approach as we consider algorithms that only give recommendations but do not compete with the sellers.

Our study also relates to the literature on recommended retail prices. These are pricing recommendations that a manufacturer provides to its retailers. In theory, they can act as a coordination device (Faber & Janssen, 2019, Buehler & Gärtner, 2013) and can make markets more collusive (Foros & Steen, 2013, Gill & Thanassoulis, 2016). Furthermore, they can also influence demand by setting a reference point for the consumers (Bruttel, 2018), and manufacturers may use them to influence and guide the search process of consumers (see, for instance, Lubensky (2017) and Janssen and Reshidi (2022)). Algorithmic price recommendations, while similar in providing common recommendations to competing firms, differ significantly. Unlike the relatively static recommended retail prices, which are often distributed in print format, algorithmic recommendations are highly dynamic and digitally distributed. Furthermore, they are not visible to consumers, influencing demand solely through the pricing decisions of sellers.

Recommendations and requests influence the decisions of participants in various experimental games. They can increase contributions to public goods, reduce tax evasion, and facilitate coordination in games with correlated equilibria (Silverman et al., 2014, Croson & Marks, 2001, Cadsby, Maynes and Trivedi, 2006, Duffy & Feltovich, 2010). Various papers also study the effect of price announcements on collusion. Those announcements can be understood as a recommendation of one participant to other market participants. While they can temporarily foster collusion, the effect usually fades as the game progresses (e.g., Holt & Davis, 1990, Harstad et al., 1998, Harrington et al., 2016). Sonntag and Zizzo (2015) show that authoritarian quantity recommendations can lower quantities in a Cournot market game⁵. Schotter and Sopher (2003) shows that intergenerational advice passed from one participant to the next in a sequence can help coordinate actions and improve outcomes

⁵While we consider neutrally framed dynamic recommendation algorithms that depend on the history of the game, we also conducted a static recommendation treatment, similar to this aspect of the study by Sonntag and Zizzo (2015). We document it in Appendix C.4.

across various games and experimental setups. Our approach is distinct as an algorithm provides dynamic recommendations instead of previous subjects.

The remainder of the article is structured as follows. Section 2 introduces the experimental design and discusses the rule-based algorithms we consider. Furthermore, we derive our hypotheses. In Section 3, we present the results. We discuss the implications of our results in Section 4 and provide avenues for future research. Online Appendix A contains the theoretical derivations underlying our algorithms. Online Appendix B contains the instructions and post-experimental questionnaire. We document various robustness checks and further algorithm variations in Online Appendix C. In Online Appendix D, we demonstrate that a monopoly platform can benefit from collusive price recommendations, highlighting that it can be a suitable example for collusive price recommendations in practice.

2. Experimental design

2.1. Market environment

To experimentally investigate the collusive effect of price recommendations, we consider a market setup similar to Dufwenberg and Gneezy (2000) and Fonseca and Normann (2012). There are $n = 3$ sellers, each represented by a participant. The market size is chosen such that tacit collusion is unlikely without any recommendation (Huck et al., 2004)⁶. All firms sell the same homogeneous good and have no capacity constraints⁷. The demand side consists of $k = 30$ computerized consumers with perfectly inelastic unit demand. Firms compete in an indefinitely repeated game with a discount rate of 95%. In each round of the game, all participants choose their prices from the set of integers in the range between the competitive Nash price of $p^N = 1$ and the monopoly price of $p^M = 10$ independently. The seller with the lowest price in a given period supplies the entire market. If multiple sellers have the lowest price, they share the market equally. There is no direct communication between the participants.

2.2. Price recommendations and treatments

We consider two additional treatments next to a BASELINE treatment in which we do not provide any price recommendation to the participants. Across those treatments, we vary which type of algorithm provides the price recommendations. The algorithms are designed to foster tacit collusion. While they do not provide fundamentally new information to the seller, they could help coordinate market behavior. Each participant in a market receives a price recommendation at the beginning of each round. Crucially, this price recommendation is identical for all participants within the same market in a given round. After each participant selects a price, the participants receive information about the pricing decision of the other participants and their payoff in the given round. Furthermore, the recommended price is again shown to the respective treatment participants.

It is commonly understood that for collusion to be successful, firms have to anticipate that collusion takes place. The fact that an algorithm provides collusive price recommendations can support this anticipation. Moreover, the firms must obtain a common understanding of the collusive strategy. The collusive price is not necessarily the monopoly price of p^M but could be any price above the competitive price p^N . Additionally, collusion also depends on a shared understanding of how to punish

⁶Horstmann et al. (2018) provide a meta-study of the relationship between the number of firms in the market and tacit collusion. They find that collusion is more challenging to sustain when moving from two to three-firm markets. Also, the meta-study by Engel (2015) highlights that market prices are decreasing in the number of firms. Freitag et al. (2021) report no evidence of collusive outcomes within a multi-market environment with three firms.

⁷Differentiation between products or sellers is a typical feature in many markets. As is common in many models and experiments, however, we abstract from this for the sake of simplicity. Qualitatively analogous results for the recommendation algorithms can be obtained, at least theoretically, when introducing differentiation. The detailed derivations are in Online Appendix A.

deviations from the collusive price. It includes a punishment price but also an understanding of how many periods this price is set before, possibly, the firms return to a collusive price. Recommendations can act as a coordination device that addresses all these issues. The idea behind a recommendation algorithm is that sellers may expect other sellers to behave according to the recommendation. It makes it incentive-compatible to do the same⁸.

We focus on rule-based algorithms as they are highly tractable and allow us to derive clear, theoretically guided hypotheses. Furthermore, in digital markets, many algorithms are simple as well. Hanspach et al. (2024) and Musolff (2022) show that real-world pricing algorithms are often rule-based and follow straightforward conditional processes. Moreover, although alternative methods like reinforcement learning algorithms have more complex routines to learn a pricing strategy, they eventually often converge to strategies that simple rules can describe (Werner, 2022, Klein, 2021, Kasberger et al., 2023)⁹. Hence, our focus on those algorithms is attractive from a methodological perspective and realistic regarding the tools used in actual markets.

In the following, we outline the recommendation algorithms we consider. We provide further details on the theoretical foundations of the algorithms and the different equilibria they induce in Online Appendix A¹⁰.

2.2.1. Algorithm that recommends Nash equilibrium actions

The following algorithm, labeled RECTHEORY, recommends prices according to a trigger strategy. It relies on two state variables, “collusive” and “punitive”, which are determined by the sellers’ past pricing behavior. Based on these states, the algorithm recommends either the monopoly price (p^M) or the stage game Nash equilibrium price (p^N). It operates as follows:

Algorithm 1. (RECTHEORY)

- At the beginning of period $t = 1$, initialize:
 - Set the state variable S_t to “collusive”.
 - Set the punishment counter τ_t to 0.
- At the beginning of any period $t > 1$, check for a state transition:
 - If $S_{t-1} = \text{“collusive”}$:
 - * If all firms set a price of p^M in period $t - 1$, set S_t to “collusive”.
 - * Otherwise:
 - Set S_t to “punitive”.
 - Set the punishment counter to be equal to the punishment duration T ($\tau_t = T$).
 - If $S_{t-1} = \text{“punitive”}$:
 - * Decrease the punishment counter by 1 ($\tau_t = \tau_{t-1} - 1$).
 - * If the punishment counter equals 0 ($\tau_t = 0$), set S_t to “collusive”.
- In each period t , provide a recommendation based on S_t :
 - If $S_t = \text{“collusive”}$, recommend a price of p^M .
 - If $S_t = \text{“punitive”}$, recommend a price of p^N .

⁸We show that it is incentive-compatible to follow the recommendations in our setup in Online Appendix A.

⁹The Q-learning algorithms in Klein (2021) punish for a certain number of periods before reverting to the monopoly price. In Werner (2022), they learn one-period punishment strategies similar to a win-stay lose-shift strategy where algorithms revert to the stage game Nash equilibrium for punishment.

¹⁰We consider two additional mechanisms as a robustness check that we report in Online Appendix C.4. In RECSATIC, the algorithm provides a static price recommendation at the monopoly price similar to Sonntag and Zizzo (2015). Additionally, we analyze an algorithm similar to RECTHEORY but with a shorter punishment phase. Both additional algorithms do not foster collusion compared to BASELINE.

Following our parameterization of the environment with $p^M = 10$ and $p^N = 1$, RECTheory recommends actions that constitute a subgame perfect Nash equilibrium for a punishment phase of $T \geq 3$ (see Online Appendix A). We choose the minimal feasible punishment length of $T = 3$ for the experiments. In summary, RECTheory recommends the monopoly price to all firms. If all firms continue to play the monopoly price, the recommendation remains unchanged in the next period. If one or more firms deviate from this recommendation, the algorithm recommends punishing this deviation jointly for three periods and reverting to the monopoly price afterward.

2.2.2. Behaviourally motivated soft punishment algorithms

Empirical and experimental evidence indicates that punishment is often less harsh than in theory models with trigger or even grim-trigger strategies. For instance, Wright (2013) finds that only a small fraction of subjects in market experiments use optimal or grim punishment strategies. Most punishment strategies are softer and more gradual. It concerns both the punishment length, as well as by how much prices are reduced in a punishment phase. Similarly, Dal Bó and Fréchette (2019) show that humans often use tit-for-tat strategies in the iterated prisoners' dilemma, which is strategically similar to our stylized market environment.

To reflect these practices, we set up a behaviourally motivated recommendation algorithm. It works as follows:

Algorithm 2. (RECSoft)

- At the beginning of period one, recommend a price of p^M .
- At the beginning of any period after period one:
 - If all sellers set the same price in the previous period, recommend the monopoly price of p^M in the current period.
 - If not all sellers set the same price in the previous period, recommend a punitive price equal to the lowest price from the previous period (e.g., $\min(10, 10, 9) = 9$).

In view of the behavioral insights cited above, such a recommendation mechanism may be superior to the algorithm implementing a subgame perfect Nash equilibrium with trigger strategies. The recommendation is similar to a tit-for-tat algorithm as it mimics the firms' decisions in the previous period. However, it also proactively tries to increase prices after all firms arrive at the same price level by recommending the monopoly price again to everyone.

2.3. Hypothesis

We test the following hypotheses in the experiment.

Hypothesis 1. Recommendations positively influence individual prices. A higher recommended price leads to higher individual prices.

As the recommendation may act as a coordination device, we expect that firms factor it into their pricing decision, and we hypothesize that higher recommendations lead to higher individual prices. This is necessary for any sensible algorithm to have a collusive effect.

Hypothesis 2. The RECTheory recommendation algorithm leads to higher market prices than the absence of a recommendation algorithm.

Hypothesis 2 expresses the rationale that RECTheory acts as a coordination device among the firms and thereby indeed facilitates collusion.

Similarly, but for different reasons, RECSoft can have a pro-collusive effect. Following the recommendation might be behaviourally attractive as no harsh punishment needs to be implemented compared to RECTheory. With k -level reasoning, for instance, a seller might rationalize that other sellers prefer to punish if it bears little cost and it yields an expected price soon after. Suppose sellers anticipate punishment under the current soft punishment algorithm. In that case, it may deter them from departing from the collusive price¹¹. Furthermore, if sellers deviated in the past, the algorithm promotes cooperation as it again recommends the monopoly price once sellers agree on a joint price level. These arguments lead to

Hypothesis 3. The RECSoft recommendation algorithm leads to higher market prices than the absence of a recommendation algorithm.

It is noteworthy that, in contrast to RECTheory, following the recommendations from RECSoft does not constitute a subgame perfect Nash equilibrium at the parameters used in the experiment. We show this formally in Online Appendix A. Following the soft recommendation algorithm may nevertheless be more attractive than the recommendation involving Nash reversion in punishment phases. It depends on the willingness of the sellers to implement drastic and longer-lasting punishments and their beliefs about the behavior of the other market participants. Nevertheless, sellers might find the soft punishment not harsh enough. Which recommendation algorithm performs better thus remains an ex-ante open question.

2.4. Procedure

The experiments were conducted between February 2020 and August 2021 in the University of Duesseldorf DICE Lab. We used ORSEE (Greiner, 2015) to recruit the subject for the experiments. The experiment was programmed in oTree (Chen et al., 2016). We utilized a between-subject design, and each subject participated only once.

At the beginning of each experiment session, participants were randomly assigned to a computer in the lab and could read the instructions on the computer screen. Additionally, the participants received a printed version of the instructions. The instructions were the same for each subject. The original and the translated version are in Online Appendix B. After the subjects read the instructions, they answered several control questions to ensure they understood the setup¹². In case a participant failed to answer all control questions correctly, the software asked the participant to reread the instructions and allowed the participant to ask the experimenter clarifying questions in private.

In RECTheory and RECSoft, the instructions describe the objective of the algorithms to the participants. Specifically, we explain that the recommendation algorithm aims at increasing the sellers' joint long-term profits. We also explain that the recommendations reflect this long-run objective rather than achieving the highest possible profit in any single round. One control question specifically assesses whether participants comprehend the design purpose of the algorithm. The answers are affirmative and confirm that the participants have the same understanding of the algorithm's objective of maximizing the sellers' joint profits.

Following the design of the algorithm described in Section 2.2, all firms in a market receive the same recommendation in a given round. We emphasize this clearly in the instructions to ensure participants understand this central aspect of the experimental setup. The recommended price is identical for all firms, both when the algorithm suggests collusion and when it recommends punishment. The instructions also point out that the price recommendation is non-binding, so each subject can choose a different price. This approach is motivated by the price suggestions that sellers receive in online marketplaces.

¹¹We also consider a recommendation algorithm without any punishment in the Online Appendix C.4.

¹²All control questions are in Online Appendix B.3.

Table 1. Number of observations by treatment

Treatment	Number of participants	Number of independent observations		
		Supergame 1	Supergame 2	Supergame 3
BASELINE	54	18	6	6
RECSOFT	54	18	6	6
RECTHEORY	54	18	6	6

Note: The number of independent observations in later supergames is determined by the matching group size which always consists of nine participants.

In real-world markets, sellers usually do not know how the algorithms work exactly and how they are programmed. We mimic this information structure in our experiment and do not describe the exact strategy of the algorithm besides the explanation that it aims at maximizing the joint profits of all sellers. Crucially, participants can learn the algorithm's behavior over time. Hence, participants can learn that following the recommendation may benefit them in the long run. This again imitates the setup in real markets.

To mimic an infinitely repeated game, each round of the game has a continuation probability of 95%. Thus, with a probability of 5% each game terminates after a given round. Within this setup, the continuation probability is equivalent to the discount rate of $\delta = 0.95$ (Roth & Murnighan, 1978). The game is repeated for three supergames to observe possible learning effects¹³. Within each supergame, the group composition is fixed. We use a perfect stranger matching scheme across supergames. Hence, the participants know they will meet each participant only once during the entire experiment. It rules out possible reputation effects that might arise otherwise. At the end of the experiment, the participants answered a post-experimental questionnaire (see Online Appendix B.3).

In total, we allocated 162 participants evenly across the three main treatments¹⁴. The market and matching group sizes determine the number of independent observations. In the first supergame, there are no spillovers from one market to another. Hence, each of the 18 markets per treatment is an independent observation. In later supergames, participants are rematched with other participants from the same matching group. Each perfect stranger matching group consists of nine subjects. The markets are not independent anymore due to possible spillovers created by previous supergames. Therefore, the number of independent observations in later supergames is lower. To account for this dependency, we either cluster the standard errors at the matching group level or aggregate the respective outcome variable at the matching group level if we use nonparametric tests. Table 1 contains an overview of the number of independent observations.

We used an experimental currency unit (ECU) with an exchange rate of 100 ECU = EUR 1. On average, the participants received a payoff of EUR 10.73 plus a show-up fee of EUR 4¹⁵. The average session length was 45 minutes.

3. Results

3.1. The influence of price recommendations on individual prices

Hypothesis 1 states that price recommendations influence individual prices as participants base their pricing decisions on the recommendations. To test this hypothesis, we regress the individual prices (p_t^i) on the recommended prices (p_t^R). The results of the linear regressions are shown in Table 2.

¹³The exact number of rounds was pre-drawn with a random number generator to allow for the same supergame length across different experimental sessions. The round numbers are 27 (Supergame 1), 8 (Supergame 2), and 18 (Supergame 3).

¹⁴For details on the additional control treatments see Online Appendix C.4. The control treatments highlight the effects of algorithms without any punishment or fewer punishment periods.

¹⁵During the COVID-19 pandemic, we paid each participant an additional EUR 4. This bonus was announced after the end of the session. Thus, it does not influence the behavior in the experiment itself.

Table 2. Individual prices explained by the recommendation in a linear regression

Dependent Variable:		Individual price (p_t^i)			
Model:	(1)	(2)	(3)	(4)	(5)
<i>Variables</i>					
p_t^R	0.376*** (0.075)	0.203*** (0.026)	0.385*** (0.083)	0.224*** (0.038)	0.307*** (0.079)
p_{t-1}^i		0.554*** (0.029)			
p_{t-2}^i		0.223*** (0.013)			
RECTHEORY				0.348 (0.713)	1.48 (1.09)
$p_t^R \times \text{RECTHEORY}$				0.346*** (0.078)	0.152 (0.117)
Further controls:			Yes	Yes	Yes
Sub-sample:	All	All	All	All	Supergame 3
<i>Fixed-effects</i>					
Round			Yes	Yes	Yes
Supergame			Yes	Yes	
<i>Fit statistics</i>					
Observations	5,724	5,076	5,724	5,724	1,944
R ²	0.160	0.655	0.268	0.365	0.324
Within R ²			0.199	0.305	0.293

Note: The coefficients are from a linear regression. Model specifications without fixed effects are estimated with a constant. For the fixed effects regression, we include dummies for each of the rounds and supergames. Signif. codes: *** < 0.01, ** < 0.05, * < 0.1; Clustered (matching group) standard errors in parentheses.

In all five columns, price recommendations positively and significantly affect individual prices. The effect remains significant when we control for lagged prices (column 2) and time-fixed effects (column 3). Furthermore, the effect size is more extensive for RECTHEORY than for RECSTOFT (column 4)¹⁶. In specifications (3) and (4), we furthermore control for a set of individual-specific control variables¹⁷.

Notably, the additional effect of RECTHEORY diminishes in the final supergame (column 5). After the first two supergames, 97% of participants observed punishment by the algorithms at the Nash equilibrium stage game at least once and learned their mechanisms. Therefore, once participants understand how the recommendations work, both algorithms exhibit a similarly positive impact on prices. This supports Hypothesis 1, indicating that recommendations positively influence prices.

Result 1. Sellers condition their prices on the recommendation of the algorithms. Price recommendations positively influence the individual sales prices of the participants.

In all regression specifications, the coefficient of the price recommendation is below one. It indicates that the price recommendation only translates partially into the individual price. Increasing the price recommendation by one only increases the individual price by 0.20 to 0.57, depending on

¹⁶The effect is not driven by a difference in the average price recommendation across both treatments. We discuss this in the following section.

¹⁷These include economic preferences and measures of socioeconomic status. We provide a list in the Online Appendix B.3.

the model specification and treatment. Thus, although prices change with the recommendations, it appears that, on average, participants do not fully follow them.

In both treatments, the modal action is to pick the recommended price. However, the average differences between the recommended and individual prices are 2.06 in REC*SOFT* and -0.53 in REC*THEORY*. Thus, in REC*SOFT*, participants tend to undercut the recommendation, and in REC*THEORY*, they do not always adhere to the punishment recommendation of the algorithm.

Furthermore, adherence to the algorithm differs depending on whether the recommendation is punitive or collusive. To illustrate the magnitude of these deviations for each state of the algorithms, we define an indicator variable that equals one if the recommended price is 10 (collusive state) and zero otherwise (punitive state). We regress the difference between the recommended and individual prices on this indicator variable. The results of this additional analysis, outlined in detail in Appendix C.1, show that participants choose a price 3.55 below the recommendation in the collusive state and exceed the recommendation by 1.77 in the punitive state. As such, in absolute terms, participants deviate from the recommendation more when the algorithms recommend colluding than when they recommend punishing. This difference across states is highly statistically significant in the first supergame ($p < 0.01$) and statistically significant at the 10% level across all supergames. However, the effect fades and is not statistically significant at any conventional level in the last supergame (separate Wald tests)¹⁸. Overall, this additional analysis indicates that participants tend to undercut recommendations when collusion is recommended but exceed them during punishment. It highlights that although participants follow the recommendations, they do so imperfectly.

3.2. Collusive effects of price recommendations

Building on the finding that subjects use the algorithms' recommendations for their pricing decisions, we now investigate whether the recommendations effectively foster collusion. Therefore, we compare the mean market prices in the treatments featuring recommendations with outcomes in the baseline treatment of no price recommendations. Note that the market price has a 1:1 relation with the sum of the sellers' profits, so an analysis of the market price is equivalent to an analysis of the aggregate profits.

According to Hypotheses 2 and 3, price recommendations foster collusion as they provide a common reference point and simplify coordination on common punishment strategies after the deviation of a firm.

Figure 1 shows the mean market prices by treatment pooled across supergames¹⁹. The average market prices in BASELINE and REC*THEORY* are similar. There are no statistically significant differences (p -value = 0.818, two-sided Mann–Whitney U test). Thus, we find no evidence that, on average, the REC*THEORY* recommendation algorithm fosters collusion.

One of our initial conjectures was that the REC*THEORY* recommendation algorithm, while constituting a subgame perfect Nash equilibrium, might feature too harsh punishments from the perspective of human players. We, therefore, designed the softer recommendation algorithm REC*SOFT*. On balance, this algorithm, however, does not foster collusion either. In fact, the market prices are on average *lower* than in BASELINE (p -value < 0.05, two-sided Mann–Whitney U test). In other words, the algorithm makes competitive market outcomes more likely even though the initial design objective was to make markets more collusive. It is in contrast to the consideration provided in Section 2.2 and to Hypotheses 2 and 3. Furthermore, average market prices in REC*SOFT* are also lower than in REC*THEORY* (p -value < 0.1, two-sided Mann–Whitney U test).

Result 2. REC*SOFT* leads to statistically significantly lower prices than BASELINE and REC*THEORY*. We do not find a statistically different mean price of REC*THEORY* relative to BASELINE.

¹⁸Further details on the statistical tests are provided in Appendix C.1.

¹⁹Furthermore, we provide an overview of the development of market prices across time in Figure 2.




Figure 1. Market price for the main treatments. The error bars represent 95% confidence intervals.

Result 2 summarizes the findings about the average treatment effects. While we find support for the hypothesis that recommendations influence individual prices (Result 1), we find no evidence that, on average, price recommendations foster collusion²⁰. This effect also remains robust if we use linear regression of the market price on the different treatment indicators, including time-fixed effects or further control variables²¹. Furthermore, the average treatment effects are similar across supergames and remain present after participants have had the chance to learn the game and how the algorithms work. We report those additional results in Table C.10 and C.11 in the online appendix.

We do find evidence that price recommendations can lead to lower prices. Notably, this result is not driven by the fact that RECSOFT leads to lower price recommendations than RECTHEORY. The average recommendation is similar in both treatments (5.40 vs. 5.94, p -value = 0.31, two-sided Mann-Whitney U test)²². Thus, there appears to be a difference in participants' responses to the different recommendation mechanisms. We explore the dynamics that drive behavior in the following section.

Our results indicate that fostering collusion purely through price recommendations may be difficult. On the contrary, price recommendations can make markets more competitive and lower market prices relative to a baseline without any recommendations. In the following section, we explore the mechanism for the price-decreasing effects of the soft recommendation algorithm. Furthermore, we discuss heterogeneous treatment effects for the RECTHEORY treatment. While the difference between BASELINE and RECTHEORY could potentially become statistically significant with greater statistical power, the heterogeneity analysis highlights that the lack of significance mainly arises from opposing responses to the recommendations across markets.

3.3. Exploratory analysis: Heterogeneity and mechanisms

In this section, we explore heterogeneity among participants to uncover some of the mechanisms driving the results that we report in the previous section. In other words, we move beyond comparing treatment differences and conduct exploratory analyses within specific population subgroups

²⁰Moreover, algorithms without any punishment mechanism or with fewer punishment periods do not increase market prices. See Online Appendix C.4 for details.

²¹The control variables were elicited on an individual level using a post-experimental questionnaire (see Online Appendix B.3). We aggregate them on the group level by calculating the mean across all group members.

²²We also provide a plot of the distribution of recommendations in Figure C.2 in the online appendix.

Table 3. Market price statistics by treatment

	BASILINE	RECTHEORY	RECSOFT
$\bar{p}_{max}$	4.94	7.96	3.5
$\bar{p}_{median}$	4.42	4.71	2.77
$\bar{p}_{min}$	2.62	1.45	1.71

based on the questionnaire items elicited at the end of the experiment. First, we document the heterogeneity in market outcomes across treatments and highlight specific stylized facts contributing to this heterogeneity. Then, we examine the price patterns that emerge in the different treatments.

3.3.1. Heterogeneous response to recommendations

There are substantial differences in market outcomes in RECTHEORY across matching groups. While the variance in average market prices in BASELINE ($\sigma^2 = 0.65$) and RECSOFT ($\sigma^2 = 0.34$) is small, there exists a large variation in RECTHEORY ($\sigma^2 = 4.85$). Those differences in variances are statistically significant ($p < 0.05$, two separate Bartlett-tests)²³. This indicates that the recommendation algorithm RECTHEORY, which recommends strategies that constitute a subgame perfect Nash equilibrium, fosters more heterogeneous market outcomes.

To study the origin of the differences in variances, we show the maximal, minimal, and median average market price across matching groups for each treatment in Table 3. In line with the previous analysis, the median market price in RECSOFT is small, and the maximal price is even below the median of the other treatments. Interestingly, although the median prices in BASELINE and RECTHEORY are similar, market prices are more spread out in RECTHEORY than in BASELINE. The recommendations in RECTHEORY make specific markets more collusive, whereas they make others more competitive.

We confirm this by dividing the observations for each treatment into subgroups that are above (HIGH) and below (LOW) the treatment-specific median market price at the matching group level across all rounds. We observe that the average market prices for the RECTHEORY-HIGH subgroup are statistically significantly higher than in BASELINE-HIGH, although only at the 10% level (two-sided MWU test). Also, the market prices in BASELINE-LOW are higher than in RECTHEORY-LOW. Nevertheless, those differences are not statistically significant, likely due to a lack of statistical power because of the sample split ($p=0.4$, two-sided MWU test).

Result 3. The variance in market outcomes is larger in RECTHEORY compared to RECSOFT and BASELINE.

Result 3 summarizes these findings. It provides some context for the lack of statistically significant differences between RECTHEORY and BASELINE. While, on average, there is no significant difference in market prices between the two treatments, this does not imply that the recommendation itself has no effect. Instead, the average treatment effect masks an increase in heterogeneity, where some markets become more collusive while others become more competitive compared to the baseline treatment.

3.3.2. Relationship between negative reciprocity and the effect of recommendations

To understand the heterogeneity in market outcomes, we regress market prices on negative reciprocity. In market games, negative reciprocity measures the willingness to punish deviations from a set price level, which is critical for sustaining collusion. Especially with recommendations that actively

²³ As in the previous analysis, we aggregate the market prices at the matching group level. Thereby, we account for dependencies that arise by rematching participants at the end of each supergame. It allows for correct statistical inference. We provide an overview of the number of independent observations in Table 1.

Table 4. Market price explained by negative reciprocity and treatments

Dependent Variable:		Market price		
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
NEG. REC.	2.03 (1.40)	−3.40** (1.54)	−3.40** (1.55)	−7.11** (3.18)
RECTHEORY		−2.44** (0.936)	−2.44** (0.941)	−5.16* (2.53)
RECSOFT		−5.51*** (1.22)	−5.51*** (1.22)	−7.38** (2.94)
NEG. REC. × RECTHEORY		4.54** (2.12)	4.54** (2.13)	8.89* (4.53)
NEG. REC. × RECSOFT		6.85*** (2.04)	6.85*** (2.05)	9.91* (5.19)
Sub-sample:	All	All	All	Supergame 3
<i>Fixed-effects</i>				
Round			Yes	Yes
Supergame			Yes	Yes
<i>Fit statistics</i>				
Observations	2,862	2,862	2,862	972
R ²	0.009	0.068	0.105	0.102
Within R ²			0.071	0.086

Note: The coefficients are from a linear regression. Model specifications without fixed effects are estimated with a constant. For the fixed effects regression, we include dummies for each of the rounds and supergames. Signif. Codes: *** < 0.01, ** < 0.05, * < 0.1; Clustered (Matching group) standard errors in parentheses.

promote punishment, participants may show varying sensitivity to the recommendation based on different levels of negative reciprocity. The observed heterogeneity appears to stem from differences in participants’ responses to the treatment, conditional on distinct levels of negative reciprocity, rather than from overall issues with randomization. In the online appendix, we provide randomization checks in Table C.2.

We elicited the economic preferences on the subject level at the end of the experiment using the validated post-experimental questionnaire by Falk et al. (2023)²⁴. We apply a min-max normalization to the variables on the individual level. Thus, all measures are between zero and one. Furthermore, we average them on the market level for the subsequent analysis.

Differences in negative reciprocity lead to vastly different market outcomes across treatments (see Table 4). In the BASELINE treatment without any price recommendations, higher degrees of negative reciprocity lead to lower market prices, as indicated by the negative coefficient of NEG. REC. in model specification 2. In other words, markets tend to exhibit lower prices if the participants are more inclined to punish each other when they feel maltreated. The ability to coordinate on collusive prices after a punishment phase may be missing for successful collusion.

For the treatments with price recommendations, this pattern is indeed different. The coefficients of the interaction terms with negative reciprocity are positive and statistically significant. Thus, as the

²⁴Next to negative reciprocity, the questionnaire also includes positive reciprocity, time preferences, risk aversion, and measures of altruism and trust. We report the results regarding those variables in Online Appendix C.3.

degree of negative reciprocity increases, market prices increase in REC*THEORY* and REC*SOFT*²⁵. This suggests that the collusive recommendations and the willingness to punish deviations are complementary. Interestingly, the price level is lower for lower levels of negative reciprocity in REC*THEORY* and REC*SOFT* compared to BASELINE, as indicated by the negative coefficients of REC*THEORY* and REC*SOFT*.

We interpret negative reciprocity as a willingness to punish deviations in this context. The recommendations harm collusion in markets, with sellers usually unwilling to punish. Possibly, the recommendations lead to harsh punishments that would not have happened without them. If participants are unable to recover from the punishment, the recommendations reduce the market prices below the level that is observed in markets without recommendations but with similarly low levels of negative reciprocity. The postulated mechanisms can explain around 10.5 % of the variation in the data, as indicated by the R-squared in column 3. We conclude that those heterogeneous treatment effects can explain lower prices than in BASELINE for the treatments with a price recommendation. Result 4 summarizes our findings regarding negative reciprocity and the effectiveness of the recommendations.

Result 4. Variations in negative reciprocity can explain the heterogeneous market outcomes. Recommendations reduce prices in markets where sellers have low negative reciprocity.

Those differences make intuitive sense and explain the considerable heterogeneity in market outcomes discussed in Result 3. The result emphasizes that algorithms can be pro-collusive for particular subgroups, even though we do not find statistically significant pro-collusive effects on average. Hence, if the algorithm's provider understands their users and targets the recommendation to the specific population of sellers, the algorithm could increase the price level.

3.3.3. Price patterns across treatments

In Figure 2, we plot the market prices for each treatment by supgame and round. In the initial round, market prices in REC*SOFT* and REC*THEORY* are higher than in BASELINE following the recommendation of $p_{t=1}^R = 10$ (p-value=0.052 & $p < 0.05$, two-sided Mann–Whitney U tests)²⁶. Yet, there are deviations from the recommendation in 86.1% of all markets in the first round. As a result, the treatment-specific punishment recommendations are triggered in the subsequent round.

Let us focus first on the pattern of REC*THEORY* in Figure 2. In response to deviations in the first round, the market prices drop for the following three rounds. At the end of this punishment phase, the prices increase sharply as the algorithm reverts to recommending the monopoly price. However, the prices do not stabilize entirely at this level. In the following rounds, there are reoccurring deviations after a recommendation of the monopoly price. This results in clearly visible spikes in the price pattern. In the second and third supgames, the spikes become less frequent, and the price patterns are more similar to BASELINE.

The recurring deviations in REC*THEORY* are almost entirely driven by matching groups with below median market prices (REC*THEORY*-LOW) as discussed in Section 3.3.1. This becomes clear when assessing the market price patterns for REC*THEORY* for both subgroups separately. Whereas there are deviations in both subgroups in the first round, the market prices stabilize in REC*THEORY*-HIGH after the initial punishment phase. In REC*THEORY*-LOW, the share of markets with deviations from the collusive recommendations is significantly higher after the first round, which results in price spikes

²⁵The average marginal effect of the treatment dummies is not statistically significant at the 10%-level if NEG. REC. is equal to one (see Table C.8 in the online appendix). Note that one is the maximal value that NEG. REC. can take due to the normalization we apply.

²⁶Market prices in REC*THEORY* and REC*SOFT* are similar in the first round following the same initial recommendation ($p=0.25$, two-sided Mann–Whitney U test), which confirms that randomization into treatments worked.




Figure 2. Market price for each treatment by supergame and round.

(p -value < 0.05, two-sided Mann–Whitney U test)²⁷. We visualize those price patterns in the online appendix in Figure C.3²⁸.

Matching groups with below-median market prices show repeated deviation patterns in the first supergame. They do not recover from this experience as average market prices remain lower throughout the rest of the experiment²⁹. Hence, we find suggestive evidence that the recommendation in REC THEORY works as expected for specific subgroups. However, other participants repeatedly deviate from the recommendation, which leads to lower market prices than in BASELINE for this subgroup.

For REC SOFT, the price patterns in Figure 2 are also interesting. We designed this recommendation algorithm to be forgiving to slight deviations as it does not immediately punish at the stage game Nash equilibrium price of $p^N = 1$ (see Section 2.2). We expected the punishment to be softer and, possibly, short compared to REC THEORY. However, the data does not support this claim. After an initial deviation from the recommended price, the following recommendation is usually above the stage game Nash equilibrium ($\bar{p}_{t=2}^R = 5.33$). Hence, in contrast to REC THEORY, there are again profitable deviations from the recommendation in the following round. Participants repeatedly deviate from the recommendation. This triggers a downward spiral as the recommendation for the next

²⁷We test this by restricting the data to the first supergame and to rounds in which the monopoly price was recommended. Then, we calculate for each market in REC THEORY-HIGH and REC THEORY-LOW the share of rounds in which at least one participant deviated from the recommendation. We test for differences in this variable across the two subgroups. Rematching only occurs after the first supergame, so each market constitutes an independent observation, allowing correct inference.

²⁸For the respective analysis for REC SOFT see Figure C.4.

²⁹Similarly, Dal Bó and Fréchette (2018) show that participants' initial experience in the infinitely repeated Prisoners Dilemma is essential for their cooperation behavior in subsequent supergames.

period is again the deviation price. There are, on average, 5.44 rounds with a recommendation below the monopoly price after the first deviation in the first supergame. This initial punishment period is significantly longer than in RECTHEORY, which always punishes for three periods ($p\text{-value} < 0.05$, one-sided one-sample t -test). Due to those frequent deviations from the recommendation, market prices deteriorate in the first rounds and only recover insufficiently in the subsequent rounds. As a result, the average prices in the treatment RECSOFT are low, and markets are even more competitive than in BASELINE. It appears that the algorithms lead the participants to learn pro-competitive behavior as the effect remains present even in the last supergame as visualized in the most right panel of Figure 2 and respective regression in Table C.11.

In RECSOFT, for $p_t^R > 5$, the mode of the difference between the recommended and individual prices is 9, while it is 0 for $p_t^R \leq 5$. This makes intuitive sense, given the discussed price patterns. Although the recommendation algorithm attempts to increase the price level by suggesting the monopoly price after all firms choose the same price, participants eventually disregard these upward recommendations. Many markets sustain a joint price of 1 after the initial deviations and punishments. As all prices are the same, the algorithm continues to recommend $p_t^R = 10$ despite the low price level. This leads to a significant discrepancy between the recommendations and prices for these high price recommendations³⁰.

Result 5. Repeated deviations from the recommendation in RECTHEORY lead to lower market prices for specific markets. The recommendation in RECSOFT offers repeated deviation opportunities that decrease market prices.

Result 5 again emphasizes the adverse effects of recommendations for a third party that likes to foster collusion if they are not designed appropriately. Furthermore, it suggests that recommendations can be used to decrease sellers' prices. It can be attractive in specific scenarios, for instance, to avoid excessive double marginalization. Sales intermediaries like platforms could specifically design a recommendation algorithm to foster competition among the sellers. Our results suggest those price recommendations are feasible using recommendation patterns as in RECSOFT.

4. Concluding remarks

We derive two rule-based recommendation algorithms and study their effects on seller collusion in a stylized Bertrand market environment. Both algorithms aim to foster collusion compared to a baseline without any recommendation. The recommendation of the RECTHEORY algorithm uses harsh punishment phases after deviations from the recommended price and aims at implementing a subgame perfect Nash equilibrium. Motivated by experimental evidence, we also design a recommendation algorithm (RECSOFT) that recommends softer punishments after a seller deviates from the collusive price. We test both algorithms in a laboratory experiment where each participant represents a seller.

We find evidence that the recommendations influence the sales prices in the sense that higher recommended sales prices induce sellers to set higher individual prices. The algorithm RECTHEORY, which recommends collusive trigger strategies with temporary Nash reversal, does not lead to higher prices on average. However, we find extensive and interesting heterogeneity in market outcomes. The variance in market prices increases in RECTHEORY compared to the baseline treatment, suggesting that while some markets are more collusive, others become more competitive. This indicates that the null effect is not primarily due to a lack of statistical power but instead reflects a divergence into opposing market outcomes, which masks the impact of the recommendations in the

³⁰Note that in RECTHEORY, the modal price is always the recommendation, independent of whether the recommendation is to pick a high or low price. Further analyses on the deviation from the recommendation are provided in Appendix C.1.

average. Furthermore, REC*THEORY* lowers market prices in markets with low levels of negative reciprocity among sellers. For the behaviourally motivated algorithm REC*SOFT*, which can recommend brief punishment phases with moderate price levels, we find lower market prices compared to the case without any recommendation. Participants frequently deviate from the recommendation, which starts a downward spiral that lowers prices. Similarly to REC*THEORY*, market prices are lower than in BASELINE for markets with sellers that have low negative reciprocity. There is no evidence that REC*SOFT* facilitates collusion for any subgroup.

Our findings contrast with previous findings from targeted punishment in Cournot (Roux & Thöni, 2015) or public goods games (see, for instance, Fehr & Gächter, 2000). In those experiments, participants can punish specific participants by explicitly reducing their payoff. This is highly effective at fostering cooperation. In our setup, the possibility of punishment works through coordination on a punishment strategy instead, an inherently less targeted approach. At the same time, the recommended trigger strategies theoretically have a considerable potential to foster collusion and are a standard modeling assumption when studying collusion. Our findings are surprising in this sense and suggest that coordination on non-targeted punishment strategies is insufficient to foster collusion.

While the results are, on balance, not alarming regarding the collusive risks of recommendation algorithms, we do provide reasons for potential concern. We find that recommendation algorithms can facilitate seller collusion in certain circumstances. Recommendations may foster collusion and harm consumers if the algorithm provider understands the sellers' characteristics and targets the recommendation based on these characteristics³¹. Furthermore, in practice, recommendations can provide additional information on demand or help with pricing more generally, which could make sellers more likely to follow them. We chose to abstract from these factors in order to isolate the pure coordination effect. However, we suspect that the collusive potential of recommendations may be higher when there are other reasons for sellers to follow them. Therefore, we believe our experiment is relatively conservative in terms of demonstrating collusive effects. We consider it fruitful for future research to study the collusive effects of recommendations that incorporate these additional factors.

In other cases, we find that recommendation algorithms may even decrease prices despite being designed and intended to facilitate collusion. The finding is consistent with the theoretical insight that all players following this behaviorally motivated algorithm does not constitute a subgame perfect Nash equilibrium. One interpretation of the price-reducing effects is that platforms may be able to use recommendation algorithms to make the sellers' offers more competitive. Under certain circumstances, such as excessive double marginalization or a dynamic pricing strategy, this could be in the interest of a sales platform. A caveat applies, as we told our experiment participants that the algorithm would aim to increase prices and profits, which aligned with our expectations. On average, the opposite, however, turned out to be the case for this algorithm. Over time, sellers may thus lose trust in following the algorithm's recommendations. More research in this regard would be desirable.

In the experiment, we abstract from the contracts between the recommendation provider, the sellers, and the developers, thus excluding their incentives. From the perspective of vertical relations theory, it raises the question of why a provider or developer would need to facilitate seller collusion when other mechanisms could influence sales prices. For instance, from an ex-ante perspective, it is unclear why a sales platform would provide collusive price recommendations instead of charging higher commission rates. In Online Appendix D, we study contracting between such an intermediary and the sellers theoretically, showing that an intermediary can benefit from collusive recommendations. Future research could explore the role of the recommendation provider and developer as strategic players more explicitly in experiments.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/eec.2025.9>.

³¹For example, accommodation platforms may gather more and more data about their hosts and guests over time and thus could condition their recommendations on seller characteristics in specific local markets to make them more effective.

Acknowledgements. We thank participants of the ESA World 2022, DICE Brown Bag, EARIE 2023 and, in particular, Lisa Bruttel, Johannes Muthers, and Hans-Theo Normann for their helpful comments and suggestions. We thank Robin Bitter, Leon Heidelberg, Vera Konrad, and, in particular, Mara Siegesmund for excellent research assistance. The replication material for the study is available at <https://doi.org/10.17605/OSF.IO/V2TMG>.

Statements and Declarations. The authors declare that they have no competing financial or non-financial interests.

References

- Bó, D., Pedro, & Fréchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1), 60–114.
- Bó, D., Pedro, & Fréchette, G. R. (2019). Strategy choice in the infinitely repeated Prisoner's Dilemma. *American Economic Review*, 109(11), 3929–52.
- Brown, Z. Y., & MacKay, A. (2023). Competition in pricing algorithms. *American Economic Journal: Microeconomics*, 15(2), 109–156.
- Bruttel, L. (2018). The effects of recommended retail prices on consumer and retailer behaviour. *Economica*, 85(339), 649–668.
- Buehler, S., & Gärtner, D. L. (2013). Making sense of nonbinding retail-price recommendations. *American Economic Review*, 103(1), 335–59.
- Bundeskartellamt and Autorité de la concurrence. (2019) Algorithms and competition. Available at, https://www.bundeskartellamt.de/SharedDocs/Publikation/EN/Berichte/Algorithms_and_Competition_Working-Paper.pdf.
- Cadsby, C. B., Maynes, E., & Umashanker Trivedi, V. (2006). Tax compliance and obedience to authority at home and in the lab: A new experimental approach. *Experimental Economics*, 9(4), 343–359.
- Calvano, E., Calzolari, G., Denicolo, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–97.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree-An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Competition and Markets Authority, Algorithms: How they can reduce competition and harm consumers, Available at: <https://www.gov.uk/government/publications/algorithms-how-they-can-reduce-competition-and-harm-consumers> 2021.
- Croson, R., & Marks, M. (2001). The effect of recommended contributions in the voluntary provision of public goods. *Economic Inquiry*, 39(2), 238–249.
- Duffy, J., & Feltovich, N. (2010). Correlated equilibria, good and bad: An experimental study. *International Economic Review*, 51(3), 701–721.
- Dufwenberg, M., & Gneezy, U. (2000). Price competition and market concentration: An experimental study. *International Journal of Industrial Organization*, 18(1), 7–22.
- Engel, C. (2015). Tacit collusion: The neglected experimental evidence. *Journal of Empirical Legal Studies*, 12(3), 537–577.
- Ezrachi, A., & Stucke, M. E. (2017). Artificial intelligence & collusion: When computers inhibit competition. *U. Ill. L. Rev.*, 2017, 1775.
- Faber, R. P., & Janssen, M. C. W. (2019). On the effects of suggested prices in gasoline markets. *The Scandinavian Journal of Economics*, 121(2), 676–705.
- Falk, A., Becker, A., Dohmen, T., Huffman, D., & Sunde, U. (2023). The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science*, 69(4), 1935–1950.
- Federal Trade Commission and U.S. Department of Justice (2024). Statement of Interest of the United States. *Case No.*, 1(23), March :-cv-02536-KMW-EAP.
- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4), 980–994.
- Fonseca, M. A., & Normann, H. -T. (2012). Explicit vs. tacit collusion-The impact of communication in oligopoly experiments. *European Economic Review*, 56(8), 1759–1772.
- Foros, O., & Steen, F. (2013). Vertical control and price cycles in gasoline retailing. *The Scandinavian Journal of Economics*, 115(3), 640–661.
- Freitag, A., Roux, C., & Thöni, C. (2021). Communication and market sharing: an experiment on the exchange of soft and hard information. *International Economic Review*, 62(1), 175–198.
- Gill, D., & Thanassoulis, J. (2016). Competition in posted prices with stochastic discounts. *The Economic Journal*, 126(594), 1528–1570.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125.
- Hansen, K. T., Misra, K., & Pai, M. M. (2021). Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1), 1–12.
- Hanspach, P., Sapi, G., & Wieting, M. (2024). Algorithms in the marketplace: An empirical analysis of automated pricing in e-commerce. *Information Economics and Policy*, 69, 101111.

- Harrington, J. E. (2022). The effect of outsourcing pricing algorithms on market competition. *Management Science*, 68(9), 6889–6906.
- Harrington, R. H. G., & Kujal, P. (2016). The relative efficacy of price announcements and express communication for collusion: Experimental findings. *Journal of Economic Behavior & Organization*, 128, 251–264.
- Harstad, R., Martin, S. and Normann, H. -T. (1998). *Intertemporal Pricing schemes: Experimental Tests of Consciously Parallel Behavior in oligopoly* Applied Industrial Organization. Cambridge University Press.
- Hill, D. (2015). How much is your spare room worth?. *IEEE Spectrum*, 52(9), 32–58.
- Holt, C. A., & Davis, D. (1990). The effects of non-binding price announcements on posted-offer markets. *Economics letters*, 34(4), 307–310.
- Horstmann, N., Krämer, J., & Schnurr, D. (2018). Number effects and tacit collusion in experimental oligopolies. *The Journal of Industrial Economics*, 66(3), 650–700.
- Huck, S., Normann, H. -T., & Oechssler, J. (2004). Two are few and four are many: number effects in experimental oligopolies. *Journal of Economic Behavior & Organization*, 53(4), 435–446.
- Janssen, M., & Reshidi, E. (2022). Regulating recommended retail prices. *International Journal of Industrial Organization*, 85, 102872.
- Johnson, J. P., Rhodes, A., & Wildenbeest, M. (2023). Platform Design When Sellers Use Pricing Algorithms. *Econometrica*, 91(5), 1841–1879.
- Kasberger, B., Martin, S., Normann, H. -T., & Werner, T. (2023). Algorithmic Cooperation. Available at SSRN 4389647.
- Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3), 538–558.
- Lefez, W. (2021) Price recommendations and the value of data: A mechanism design approach. Working Paper.
- Lubensky, D. (2017). A model of recommended retail prices. *The RAND Journal of Economics*, 48(2), 358–386.
- Musolf, L. (2022). Algorithmic pricing facilitates tacit collusion: Evidence from e-commerce. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, July (pp. 32–33). <https://dl.acm.org/doi/abs/10.1145/3490486.3538239>
- Normann, H. -T., & Sternberg, M. (2023). Human-algorithm interaction: Algorithmic pricing in hybrid laboratory markets. *European Economic Review*, 152, 104347.
- Pavlov, V. and Berman, R. (2019) Price manipulation in Peer-to-Peer markets and the sharing economy. Working Papers 19-10, NET Institute September.
- Roth, A. E., & Keith Murnighan, J. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17(2), 189–198.
- Roux, C., & Thöni, C. (2015). Collusion among many firms: The disciplinary power of targeted punishment. *Journal of Economic Behavior & Organization*, 116, 83–93.
- Schlütter, F. (2022) Managing seller conduct in online marketplaces and platform most-favored nation clauses. Working Paper.
- Schotter, A., & Sopher, B. (2003). Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of Political Economy*, 111(3), 498–529.
- Silverman, D., Slemrod, J., & Uler, N. (2014). Distinguishing the role of authority “in” and authority “to”. *Journal of Public Economics*, 113, 32–42.
- Sonntag, A., & John Zizzo, D. (2015). Institutional authority and collusion. *Southern Economic Journal*, 82(1), 13–37.
- Teh, T. -H. (2022). Platform governance. *American Economic Journal: Microeconomics*, 14(3), 213–254.
- Werner, T. (2022). Algorithmic and human collusion. Available at SSRN, 3960738.
- Wright, J. (2013). Punishment strategies in repeated games: Evidence from experimental markets. *Games and Economic Behavior*, 82, 91–102.