Algorithmic Price Recommendations and Collusion: Experimental Evidence

Matthias Hunold* and Tobias Werner†

February 2023

Abstract

This paper investigates the collusive and competitive effects of algorithmic price recommendations on market outcomes. 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. Variations in economic preferences lead to heterogeneous treatment effects and explain the results.

JEL Classification: C92, D43, L13, L41

Keywords: Collusion, Experiment, Human–Machine Interaction, Bertrand Oligopoly

*Universität Siegen, Unteres Schloß 3, 57068 Siegen, Germany; E-mail: matthias.hunold@uni-siegen.de

†Düsseldorf Institute for Competition Economics (DICE), University of Duesseldorf, Universitätsstr. 1, 40225 Duesseldorf, Germany, werner@dice.hhu.de

We thank participants of the ESA World 2022, DICE Brown Bag, and, in particular, Lisa Bruttel, Johannes Muthers, and Hans-Theo Normann for their helpful comments and suggestions. We thank Robin Bitter, Leon Heidelbach, Vera Konrad, and, in particular, Mara Siegesmund for excellent research assistance.

1 Introduction

Price recommendations are prevalent on many digital marketplaces. Platforms like Airbnb, Expedia, and eBay provide the sellers on their marketplaces with a recommendation on how to set the price for the product they sell. Those recommendations are typically created by algorithms based on historical market data.[1] Importantly, recommendations are typically non-binding in the sense that sellers can nevertheless freely choose their prices.

Various explanations for the use of price recommendations exist in practice. For instance, price recommendations might reduce information asymmetries between the platform and sellers (see Pavlov and Berman, 2019). This makes them potentially attractive from a business perspective as, for instance, platforms may have better demand information than individual sellers. However, competition authorities are concerned that price recommendation algorithms by a common intermediate can also dampen competition as it might make coordination between sellers easier (Federal Trade Commission, 2021; Bundeskartellamt and Autorité de la concurrence, 2019). For example, according to reporters of ProRebublica, price recommendation software allegedly led to coordination effects in the U.S. rental market, especially in regions where few property managers control a large share of the apartments.[2]

This article examines whether platforms can use algorithms that provide sellers with non-binding price recommendations to make markets more collusive. As digital sales platforms often receive a share of the seller's revenues through commission rates, they can benefit from higher prices if there is seller competition.[3] Even if the platform's income

[1]For 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 for which the accommodation is available. See Hill (2015) for details.

[2]See Vogell, Coryne & Little, "Technology Rent Going Up? One Company's Algorithm Could Be Why.", https://www.propublica.org/article/yieldstar-rent-increase-realpage-rent, last accessed November 25, 2022

[3]A natural question is why a platform cannot simply increase the commission rate to induce higher sales prices. We theoretically demonstrate that steering the sales prices with commission rates can be impossible or insufficient for sales platforms for various reasons, so collusive price recommendations may be desirable. See Appendix A for details.

does not directly depend on the sellers' prices, using algorithmic pricing software may increase sellers' profits and thus their willingness to pay for joining the platform and using the software. This, in turn, can increase the platform's profits. These arguments motivate the question of whether recommendation algorithms can indeed raise prices and, if so, merit closer examination by competition authorities and regulators.

We derive two rule-based algorithms from economic theory and behavioral insights to investigate the effects of a platform's recommendation algorithm on the prices of competing sellers.[4] These algorithms recommend collusive strategies with different punishment mechanics. The recommendations are non-binding and do not change the game's strategy 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 and Normann, 2012). We argue that the recommendation algorithms have the potential to facilitate coordination and, thus, collusion.

The theory-based algorithm recommends actions that are consistent with a collusive trigger-strategy and Nash reversion. If a seller undercuts the collusive price level, it recommends competitive prices for several periods until it returns to recommending collusive prices. All players 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 the sellers comply. The recommendations in our experiment provide no fundamentally valuable information, such as the state of demand, to sellers. Instead, their only purpose is to coordinate sellers. We abstract from other factors that could make following the recommendations desirable in order to isolate the pure coordination effect of recommendations.

[4]Recent studies by Musolff (2022) and Wieting and Sapi (2021) 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).

We derive several testable hypotheses based on our theoretical model and test those in laboratory experiments. Subjects 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. We inform the subjects in the experiment that the recommendation algorithm's objective is to maximize industry profits symmetrically without favoring a particular seller. Recommending high sales prices is consistent with the incentives of a platform that receives part of the sellers' revenues through commission rates.[5]

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 that recommends collusive trigger strategies 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 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 the 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 not willing to punish unfair behavior, the recommendation leads to lower market prices. Furthermore, if sellers are relatively impatient, the recommendations make markets more collusive. In other words, the recommendations increase market prices in groups of sellers which are usually too impatient for collusive strategies to be sustainable.

[5]See Appendix A for details.

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 and profits 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. Hence, the algorithm makes markets more competitive. It 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 and MacKay, 2022). Johnson et al. (2020) focus on tacit collusion among self-learning algorithms on sales platforms and discuss how the platform's design choices influence it. 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 other firms in the market.

Our work relates to hub-and-spoke collusion with algorithms. Ezrachi and Stucke (2017) discuss the possibility that multiple competing firms use the same pricing algorithm. They argue verbally that this can lead to coordination among sellers either because the sellers ask a developer to provide them with a collusive algorithm or simply because a generic algorithm induces similar behavior among sellers. We offer evidence for how it could work in practice, even if sellers do not use a specific algorithm software themselves but only receive a recommendation. Similarly, Harrington (2022) provides a theoretical model of the design and the adoption decision of such third-party algorithms and shows that they can decrease consumer welfare. Our approach differs from this as we consider non-binding recommendations instead of outsourcing the pricing decision. Furthermore, we do not model the adoption decision of firms, but the recommendation is always provided, as is commonly the case on many online sales platforms.

We also relate 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 and Janssen, 2019; Buehler and Gärtner, 2013) and can make markets more collusive (Foros and Steen, 2013). 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)). Platform price recommendations share similarities with recommended retail prices as multiple competing firms receive a joint recommendation. There are, however, apparent differences. Manufacturer's recommended retail prices typically stay mostly the same and are traditionally distributed in a printed format, whereas the digital price recommendations in online platforms may change rapidly. Notably, the recommendations on platforms are unobservable to consumers. Hence, they cannot influence demand directly but only through the pricing decision of the sellers.

Various papers study experimentally the effect of price announcements on collusion (e.g., Holt and Davis, 1990; Harstad et al., 1998; Harrington et al., 2016). Here, participants can announce prices and observe the announcement of the competitors before making the actual pricing decision. While price announcements can temporarily foster collusion, the effect usually fades, and prices decline to the level without any announcements. This reduced form of communication can be considered a recommendation by a firm in the market to its direct competitors. Our approach is distinct, as recommendations come from an algorithm that does not compete with the firms in the market.

Furthermore, recommendations and requests influence the decision of participants in various experimental games. They can increase contributions to public goods (Silverman et al., 2014; Croson and Marks, 2001), reduce tax evasion (Cadsby et al., 2006), and facilitate coordination in games with correlated equilibria (Duffy and Feltovich, 2010). Schotter and Sopher (2003) show that intergenerational advice provided by previous populations of experimental subjects can help to coordinate behavior. The result is robust to different games and experimental setups (see Schotter, 2003, for a literature review). Our approach is different as an algorithm instead of previous subjects provides the recommendations.

Sonntag and Zizzo (2015) consider static quantity requests in a Cournot market game. They vary the degree of authority with which the requests are communicated to the participants across treatments. They find that this type of authoritarian recommendation can lower quantities and, thus, make markets more collusive. We consider a setup similar to theirs. However, we concentrate on neutral recommendations by an algorithm without explicitly framing the recommendation as a request. Furthermore, we go beyond static quantity recommendations and focus on dynamic recommendation algorithms that depend on the history of the game.[6]

There are few articles that directly consider price recommendations in platform markets. Pavlov and Berman (2019) consider their effects in a cheap-talk model where the platforms possess superior information about demand. They find that recommendations can be desirable compared to centralized pricing, especially if the variance of the aggregate demand is large. Lefez (2021) focuses on how platforms use price recommendations to disclose information to sellers. The potential collusive effect of price recommendations by offering a coordination device is not explicitly discussed in either paper.

The remainder of the article is structured as follows. Section 2 discusses the theoretical framework and the rule-based algorithms we consider. Furthermore, we derive our hypothesis. Then, we introduce the experimental design in Section 3 and present the results in Section 4. We discuss the implications of our results in Section 5. In Appendix A, we demonstrate why a monopoly platform can benefit more from making collusive price recommendations for sellers than from only adjusting its commission rates. Appendix B contains the instructions and survey questions. We document various robustness checks and further algorithm variations in Appendix C.

[6]We also conducted a static recommendation treatment which we document in Appendix C.3.

2 Theoretical framework and predictions

We first set up a stylized pricing game with n sellers and solve for equilibria of the one-shot game and the infinitely repeated game. The game is similar to the framework discussed by Dufwenberg and Gneezy (2000) and describes the experimental setup that we introduce in Section 3. We then argue that a recommendation algorithm can induce different Nash equilibria by acting as a coordination device. Finally, we motivate and explain an alternative algorithm that recommends a softer punishment scheme.

In this section, we treat the platform as a black box and abstract from the contracts between the platform and the sellers to focus on how sellers react to collusive algorithms. In Appendix A, we instead focus on contracting between a platform and various sellers to analyze how the optimal price level depends on the commission rates and costs. We show that a platform can benefit from collusive recommendations even if it has the bargaining power to choose the commission rate.

2.1 Setting and Nash equilibria

We consider an infinitely repeated Bertrand game with $n \geq 2$ symmetric sellers denoted by A, B, and so on. Each seller aims to maximize its profit and discounts future profit flows with a discount factor of δ.

In each period, each seller chooses its price from the integers in the set $P = \{p^N, p^N + 1, ..., p^M\} \subset \mathbb{Z}^+$. There are k consumers who are willing to buy one unit of the good each and are willing to pay p^M per unit. 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.

The sellers have no costs and no capacity constraints. Note that abstracting from costs does not change the insights from this analysis. We would get qualitatively the same results if we explicitly modeled costs which, in reality, may include commission payments.

What matters for the analysis is that there is a range of prices between the relatively low competitive price level and a collusive price at which all firms make strictly higher profits.[7]

Nash equilibrium of the stage game. Suppose that, except for seller A, all sellers set prices larger than p^N and at least at a level of p. Notice that seller A makes zero profits for any price higher than p, whereas setting a price of p yields a profit of $p \cdot k/n$. On the other hand, a deviation to $p-1$ yields a profit of $(p-1) \cdot k$. Undercutting the lowest price of a competitor, p, by one unit is the best response if

$$\begin{aligned} &(p-1) \cdot k > p \cdot k/n \\ \implies &(p-1) > p \cdot 1/n \\ \implies &p \cdot (1 - 1/n) > 1 \\ \implies &p > n/(n-1). \end{aligned}$$

We define p^N as the integer weakly below $n/(n-1)$. At this price, no firm has an incentive to undercut, such that each firm setting a price of p^N and making a profit of $p^N \cdot k/n$ is a Nash equilibrium. For $n = 3$, there is a strict incentive to undercut any price larger than 1.5, such that $p^N = 1$. As $n/(n-1)$ is decreasing in n, it follows that $p^N = 1$ for any market with $n > 3$. For $n = 2$, both a symmetric price of 1 and a symmetric price of 2 constitute a Nash quilibrium.[8]

Collusive equilibrium of the repeated game. We now construct a collusive subgame perfect Nash equilibrium of the infinitely repeated game with trigger strategies. In line with the Folk Theorem, multiple collusive equilibria potentially exist. Variations are possible in the collusive price level and the punishment scheme. For instance, any price above

[7] See also Appendix A for our analysis of commission payments and, in particular, equation (2) which shows how commission payments affect competitive prices.

[8] For $n = 2$, there is a strict incentive to undercut any (integer) price larger than 2, such that $p^N = 2$. A symmetric price of 1 is also a Nash equilibrium, but there is no strict incentive to undercut a symmetric price of 2 either as $1 \cdot k = 2 \cdot k/2$.

the competitive price can potentially be supported as a collusive outcome. We focus on the highest and most profitable collusive price of p^M. Among the equilibria with Nash-reversion, we focus on the equilibrium with the shortest possible punishment length. As we explain below, behavioral evidence indicates that punishments are often relatively soft.

Suppose the collusive strategy is as follows:

- If the regime is collusive in the current period, set a price of p^M.
- If the regime is punitive in the current period, set a price of p^N.
- In period one, start in the collusive regime.
- If the regime was collusive in the previous period and everyone set a price of p^M, continue with the collusive regime in the current period.
- If, in the previous period, the regime was collusive, but someone set a price below p^M, switch to the punishment regime for T periods and switch back to the collusive regime afterward.

This yields the stability condition

$$\pi^M \cdot (1 + \delta + \delta^2 + ...) \geq \pi^D + \sum_{t=1}^{T} \delta^t \pi^N + \pi^M \cdot (\delta^{T+1} + \delta^{T+2} ...),$$

where π^M is the collusive period profit, π^D the deviation profit and π^N the static Nash profit as the punishment profit.

Rearranging yields

$$\pi^M \cdot (1 + \delta + \delta^2 + ... + \delta^T) \geq \pi^D + \pi^N \cdot \sum_{t=1}^{T} \delta^t$$
$$\Leftrightarrow \sum_{t=1}^{T} \delta^t \geq \frac{\pi^D - \pi^M}{\pi^M - \pi^N}.$$

Parameter values in the experiment. In the experiment, we have $p^M = 10$, k=30, n=3, and consequently $p^N = 1$. We use a value of 0.95 for the discount factor δ as this equals the continuation probability in our experiment, which we introduce in Section 3. To determine the shortest punishment length T that makes collusion stable for these parameters, we plug in the values for the profits:

$$\pi^M = 10 \cdot 30/3 = 100;$$
$$\pi^D = 9 \cdot 30 = 270;$$
$$\pi^N = 1 \cdot 30/3 = 10.$$

This yields

$$\sum_{t=1}^{T} \delta^t \geq \frac{270 - 100}{100 - 10} \approx 1.89.$$

As $\delta = .95$; $\delta + \delta^2 = 1.85$; $\delta + \delta^2 + \delta^3 = 2.59$, three punishment periods are necessary and sufficient for the stability condition to hold, which constitutes a subgame perfect Nash equilibrium in trigger strategies of the infinitely repeated stage game.

2.2 Algorithm that recommends Nash equilibrium actions

An algorithm that gives non-binding recommendations to all market participants does not change the game's action space and payoff functions. The recommendations do not provide fundamentally new information either and are non-binding.

However, collusion can be challenging to attain without a coordination device (Engel, 2007; Fonseca and Normann, 2012). First of all, the players have to anticipate that collusion takes place. The fact that an algorithm provides collusive price recommendations can support this anticipation. Moreover, the players must obtain a common understanding of the collusive strategy. Within our setup, the collusive price is not necessarily a price of

p^M but could be any price above p^N. Collusion also depends on a shared understanding of how to punish 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 players 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 same.

The following algorithm, labeled RecTheory, recommends prices according to the trigger strategy derived above.

Algorithm 1 (RecTheory)**.**

- *If the regime is collusive in the current period, recommend a price of* p^M.
- *If the regime is punitive in the current period, recommend a price of* p^N.
- *In period one, start in the collusive regime.*
- *Afterwards, if the regime was collusive in the previous period and everyone set a price of* p^M*, continue with the collusive regime in the current period.*
- *If in the previous period the regime was collusive but someone set a lower price than* p^M*, switch to the punishment regime for* T *periods and switch back to the collusive regime afterward.*

It is sensible for the competing sellers to follow the recommendations, provided it is individually rational. We inform the subjects in the experiment that the recommendation algorithm's objective is to maximize industry profits symmetrically, that is, without favoring a particular seller. Inducing high sales prices is consistent with the incentives of a platform that receives part of the sellers' revenues through commission payments. If a seller expects the other two sellers to follow the algorithm's recommendations, then it is best off in doing the same, as this constitutes a subgame perfect Nash equilibrium. A deviation from RecTheory is not profitable provided that the other sellers follow the recommendation and play the static Nash price of p^N in the punishment periods, which is again a mutually best response.

We test the following hypotheses in the experiment based on those considerations.

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

Hypothesis 1 states that firms' prices are increasing in the recommendation. 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. It is a minimal requirement 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 builds on the rationale that RecTheory acts as a coordination device among the firms and thereby indeed facilitates collusion.

2.3 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)**.**

- *Start with a recommendation of the monopoly price of p^M and continue with this recommendation in future periods as long as all sellers adhere to the recommendation.*
- *In case of a deviation, recommend a punitive price equal to the lowest price from the previous period (e.g. min(10,10,9)=9).*
- *If all sellers play the same price in a given period, recommend the monopoly price of p^M in the next period.*

In line with 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 firms agree on a joint price level.

Following the recommendation might be behaviourally attractive as no harsh punishment needs to be implemented. With k-level reasoning, for instance, a seller might rationalize that other sellers prefer to punish if it bears little costs 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 price.[9] 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. Since collusion at the monopoly price is the long-run objective of the algorithm, we argue that

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

It is noteworthy that following these recommendations does not constitute a subgame perfect Nash equilibrium. To see this, suppose that all sellers follow the recommendations throughout the game. If seller A follows the recommendations, the per-period profit is

[9] We also consider a recommendation algorithm without any punishment in the Appendix C.3.

$p^M \cdot k/n$ in each period, yielding a profit stream of

$$p^M \cdot k/n \cdot (1 + \delta + \delta^2 + \delta^3 + ...).$$

Consider a one-shot deviation of setting a price of $p_A < p^M$ while the algorithm recommends a price of p^M. The profit in the deviation period equals $p_A \cdot k$. The algorithm recommends a price of p_A in the next period. All sellers that follow the recommendation receive a profit of $p_A \cdot k/n$. Afterward, the algorithm reverts to the monopoly price of p^M. Thus, the deviating seller obtains a deviation profit of $p_A \cdot k$ for one period and a punishment profit of $p_A \cdot k/n$ for another period. Hence, the profit stream is

$$p_A \cdot k + p_A \cdot k/n \cdot \delta + p^M \cdot k/n \cdot (\delta^2 + \delta^3 + ...),$$

which is highest for the highest feasible deviation price of $p_A = p^M - 1$. The difference between the deviation profit stream and the collusion profit stream is

$$k \cdot (p_M \cdot (1 - 1/n) - 1 - 1/n \cdot \delta).$$

Thus, deviating from the recommendation is profitable if

$$p^M > \frac{n + \delta}{n - 1}.$$

For $n \geq 2$ and $\delta < 1$ the condition holds for any $p^M > 2$. Thus, following the recommendation does not constitute a subgame perfect Nash equilibrium for the parameters used in the experiment as $p^M = 10$.

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 belief 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. We, therefore, do not postulate a hypothesis in this regard.

3 Experimental design

To experimentally investigate the collusive effect of price recommendations, we consider a market setup that mimics the theoretical framework in Section 2. Each of the $n = 3$ sellers in a market is represented by a participant. The market size is chosen such that tacit collusion is unlikely without any recommendation (Huck et al., 2004). The demand side consists of $k = 30$ computerized consumers. The participants play a repeated game. In each round of the game, all participants choose their prices independently. There is no direct communication between the participants. Across treatments, we vary whether participants receive a price recommendation and which type of algorithm provides this recommendation. Each participant in a market receives the same price recommendation. 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.

3.1 Treatments

There exists a Baseline treatment in which we do not provide any price recommendation to the participants. Furthermore, we consider two treatments with rule-based price recommendation algorithms that are motivated by our theoretical considerations (Section 2). In the RecTheory treatment, the initial price recommendation is the monopoly price of $p^M = 10$. Any deviation from the recommendation by any participants triggers a punishment phase in which the stage game Nash equilibrium p^N is recommended. The punishment phase lasts for three periods. Afterward, the algorithm recommends the monopoly price again. Following the analysis in Section 2, RecTheory recommends actions that constitute a subgame perfect Nash equilibrium. In the RecSoft treatment, the algorithm also recommends a price of $p^M = 10$ in the initial round. However, after a deviation, the recommended price is the lowest price from the previous period. If all participants choose the same price in a given round, the algorithm recommends the monopoly price again. In addition to the main recommendation algorithms (RecTheory and RecSoft), we consider two additional mechanisms as a robustness check. In RecStatic, 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, and we only discuss them in the Appendix.

We focus on rule-based algorithms as they are highly tractable and allow us to derive clear, theoretical guided hypotheses that we developed Section 2. Furthermore, in digital platform markets, many algorithms are simple as well. Wieting and Sapi (2021) 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).[10] Hence, our focus on those algorithms is attractive from a methodological perspective but also realistic regarding the tools used in actual markets.

3.2 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

[10] The 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.

experiments. The experiment was programmed in oTree (Chen et al., 2016a). We utilized a between-subject design, and each subject only participated 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. A translated version of it is in Appendix B.1. After the subjects read the instructions, they answered several control questions to ensure they understood the setup.[11] 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. To be precise, we explain that the recommendation algorithm aims to increase long-term industry profits. 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.

Furthermore, the instructions emphasize that the price recommendation is non-binding so that each subject can choose a different price. This approach is motivated by the price suggestions that sellers receive in popular online marketplaces.

To mimic an infinitely repeated game as outlined in the theory (Section 2), 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 and Murnighan, 1978). The game is repeated for three supergames to observe possible learning effects.[12] Within each supergame, the group composition is fixed. We use a perfect stranger matching scheme across supergames.

[11] All control questions are in Appendix B.2.

[12] The 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).

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 different survey questions that are listed in Appendix B.2.

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 consist of nine participants.

In total, we distributed 162 participants evenly across the three main treatments.[13] 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 4.[14] The average session length was 45 minutes.

[13]For details on the additional control treatments see Appendix C.3.

[14]During 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.

4 Results

In this section, we discuss the experiment's results and test the hypotheses we derived in Section 2.

4.1 The influence of price recommendations on individual prices

Hypothesis 1 states that price recommendations influence individual prices as participants base their pricing decision 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 in Table 2.

In all four columns, price recommendations positively and significantly affect individual prices. The effect is maintained 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 RecSoft (column 4). In specifications 3 and 4, we furthermore control for a set of individual-specific control variables.[15] We conclude that the recommendations positively influence the prices. This is in support of Hypothesis 1.

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

[15] These include economic preferences and measures the socioeconomic status. We provide a list in the Appendix B.2.

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

Dependent Variable:	Individual price (p_t^i)			
Model:	(1)	(2)	(3)	(4)
Variables				
(Intercept)	2.77***	-0.201		
	(0.406)	(0.153)		
p_t^R	0.376***	0.203***	0.385***	0.224***
	(0.075)	(0.026)	(0.083)	(0.038)
p_{t-1}^i		0.554***		
		(0.029)		
p_{t-2}^i		0.223***		
		(0.013)		
RecTheory				0.348
				(0.713)
$p_t^R \times$ RecTheory				0.346***
				(0.078)
Further controls:			Yes	Yes
Fixed-effects				
Round			Yes	Yes
Supergame			Yes	Yes
Observations	5,724	5,076	5,724	5,724

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

4.2 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 industry profits, so an analysis of the market price is equivalent to an analysis of the 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.

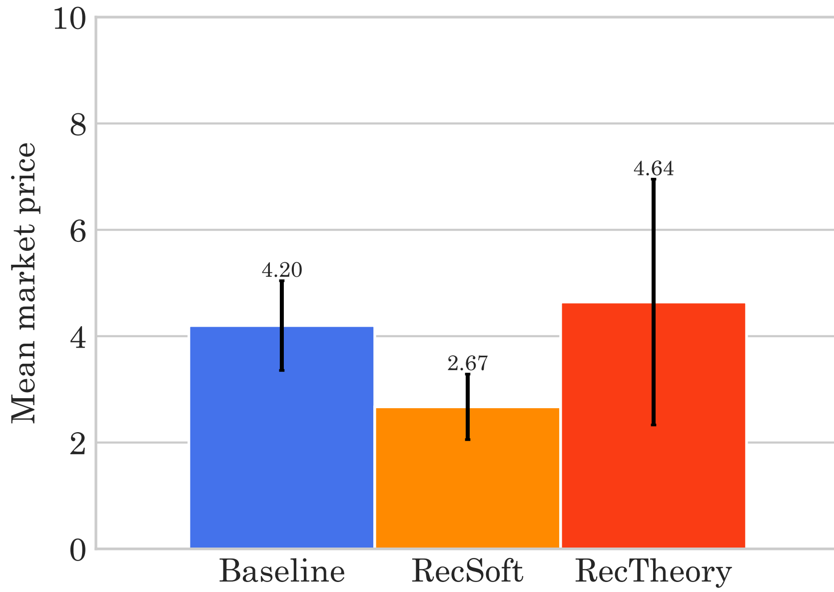

Figure 1: Market price for the main treatments. The error bars represent 95% confidence intervals.

Figure 1 shows the mean market prices by treatment pooled across supergames.[16] The average market prices in BASELINE and RECTHEORY 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 RECTHEORY recommendation algorithm fosters tacit collusion.

One of our initial conjectures was that the RECTHEORY 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 RECSOFT. On balance, this algorithm, however, does not foster tacit collusion either. In fact, the market prices are on average *smaller* than in BASELINE (p-value< 0.05, two-sided Mann–Whitney U test). In other words, the algorithm makes

[16]The results by supergame do not differ substantially and are provided in Table C.6 in the appendix. Furthermore, we provide an overview of the development of market prices across time in Figure 2.

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 and to Hypotheses 2 and 3. Furthermore, average market prices in RecSoft are also smaller than in RecTheory (p-value< 0.1, two-sided Mann–Whitney U test). Thus, the game theory based algorithm is preferred if an upstream firm wants to use a price recommendation algorithm to foster collusion in the downstream market.

Table 3: Linear regression for the treatment effects

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	4.20***		
	(0.308)		
RecSoft	-1.53***	-1.53***	-1.45**
	(0.382)	(0.384)	(0.530)
RecTheory	0.442	0.442	0.823
	(0.899)	(0.904)	(1.00)
Further controls			Yes
Fixed-effects			
Round		Yes	Yes
Supergame		Yes	Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 3 displays the results from a linear regression of the market price on the different treatment indicators. The average market prices in RecTheory are higher than in Baseline, but the standard errors are relatively large, so the differences are not statistically significant. In line with the results from the non-parametric test, the market prices are lower in RecSoft than in Baseline. The effect is robust to the inclusion of time-fixed effects and to the use of different aggregated survey measures.[17]

[17]The survey measures were elicited on an individual level and listed in Appendix B.2. We aggregate them on the group level by calculating the mean across all group members.

Result 2. RecSoft *leads to statistically significantly lower prices than* Baseline. *Price recommendations in* RecTheory *do not foster tacit collusion.*

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 tacit collusion. On the contrary, price recommendations can make markets more competitive and lower market prices relative to a baseline without any recommendations. Note that lowering the sales prices might be a desirable strategy for a sales platform if double marginalization is an issue (see Appendix A for a theoretical illustration). There is empirical evidence whereby the sales platform Amazon appears to make offers with low prices more prominent in certain circumstances (Chen et al., 2016b; Hunold et al., 2022). It could also be part of a dynamic business strategy that attempts to invest in a large consumer base first to possibly charge higher prices on many locked-in consumers later. 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.

4.3 Heterogeneity and mechanisms

On average, the treatment effects do not confirm that price recommendations raise market prices. Interestingly, we one of the algorithms even has pro-competitive effects on average. In the following, we discuss potential mechanisms that can explain those findings. We first consider heterogeneity in market outcomes across treatments and highlight specific stylized facts that drive the heterogeneity. Then, we examine price patterns that arise in the different treatments.

4.3.1 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).[18] 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 4. 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.

Table 4: Market price statistics by treatment

	Baseline	RecTheory	RecSoft
$\overline{p}_{max}$	4.94	7.96	3.5
$\overline{p}_{median}$	4.42	4.71	2.77
$\overline{p}_{min}$	2.62	1.45	1.71

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. 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-

[18] 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.

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.

4.3.2 Relationship between seller preferences and the effect of recommendations

To understand the origins of the heterogeneity in market outcomes, we regress the market prices on the different economic preference measures and interact them with the treatment variables. We focus on two variables that are critical for collusion to be sustainable from a theoretical perspective. First, we consider negative reciprocity. In the context of collusion in a market game, it is a natural measure to understand the willingness to punish deviations from a certain price level. Secondly, we analyze how time preferences interact with our treatments. Firms must be sufficiently patient for collusive strategies to be sustainable, as they have to value the long-run profits more than the short-term gain from deviating. Importantly, any heterogeneity is not driven by a lack of randomization but rather by differences in response to the treatment, conditional on distinct levels of those social preferences. We provide randomization checks in Table C.1 in the Appendix.

We elicited the economic preferences on an individual subject level at the end of the experiment using the validated survey questions by Falk et al. (2021).[19] We apply a min-max normalization to all economic preferences 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 5). In the Baseline treatment without any price recommendations,

[19]Next to negative reciprocity and time preferences, the survey also includes positive reciprocity, time preferences, risk aversion, and measures of altruism and trust. We report the results regarding those variables in Appendix C.2.

Table 5: Market price explained by negative reciprocity and treatments

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	2.63***	6.38***	
	(0.810)	(0.829)	
Neg. Rec.	2.03	-3.40**	-3.40**
	(1.40)	(1.54)	(1.55)
RecTheory		-2.44**	-2.44**
		(0.936)	(0.941)
RecSoft		-5.51***	-5.51***
		(1.22)	(1.22)
Neg. Rec. × RecTheory		4.54**	4.54**
		(2.12)	(2.13)
Neg. Rec. × RecSoft		6.85***	6.85***
		(2.04)	(2.05)
Fixed-effects			
Round			Yes
Supergame			Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

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 market prices if the participants are more inclined to punish each other when they feel maltreated. For the treatments with price recommendations, this pattern is different. While the price level is lower for small levels of negative reciprocity in RecTheory and RecSoft, as indicated by the negative coefficients of RecTheory and RecSoft, the coefficients of the interaction terms with negative reciprocity are positive and statistically significant. Thus, as the degree of negative reciprocity increases, market prices in Baseline become similar to the outcomes in RecTheory and RecSoft.[20] We interpret negative reciprocity as a willingness to punish deviations in this context. Thus, the recommendations harm collusion in markets with sellers that are 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. Those heterogeneous treatment effects can explain lower prices than in Baseline for the treatments with a price recommendation.

Furthermore, the differences in time preferences amongst sellers lead to distinct market outcomes (see Table 6 where a higher level of TIME corresponds to more patience). In Baseline and RecSoft, the prices are more collusive for markets with more patient participants. This makes intuitive sense. For collusion to be sustainable, participants must value the long-run profits more than any short-term gains from possible deviations. This is arguably the case for groups of sellers who are more patient. For RecTheory, on the other hand, market prices are higher than in Baseline if market participants are impatient. In other words, the recommendations foster collusion in situations where participants tend to deviate more due to their lack of patience. As the effect of Time is negative in this treatment, the effect wears off for more patient participants, and market prices become similar to Baseline for values of Time close to one. For large values of Time, the recommendation even has a negative effect on market prices compared to Baseline.[21] Evidently, the recommendations lead to lower prices if sellers are particularly patient. It is possible that participants who are particularly patient would not have punished in the first place without the recommendation. Small deviations may lead to harsher punishment

[20]The average marginal effect of the treatment dummies is not statistically significant at the 10%-level if Neg. Rec. is equal to one (Model specification 2 in Table 5). Note that one is the maximal value that Neg. Rec. can take due to the normalization we apply.

[21]The average marginal effect of RecTheory is negative and significant at the 5% level for Time being equal to one.

Table 6: Market price explained by time preferences and treatments

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	0.595	-1.99	
	(1.61)	(2.52)	
TIME	4.57*	9.06**	9.06**
	(2.22)	(3.56)	(3.57)
RECTHEORY		8.25*	8.25*
		(4.18)	(4.20)
RECSOFT		0.154	0.154
		(3.47)	(3.48)
TIME × RECTHEORY		-11.2**	-11.2**
		(5.21)	(5.24)
TIME × RECSOFT		-2.65	-2.65
		(5.01)	(5.03)
Fixed-effects			
Round			Yes
Supergame			Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

than usual. Result 4 summarizes our findings regarding negative reciprocity and time preferences.

Result 4. *Variations in economic preferences of negative reciprocity and patience can explain the heterogeneous market outcomes. Low negative reciprocity among market participants who receive price recommendations reduces collusion. The recommendation algorithm* RECTHEORY *makes markets more collusive if the sellers are impatient.*

Especially the markets with the RECTHEORY algorithm, which is motivated by our game theoretical considerations, outcomes depend on the sellers' degrees of negative reciprocity and patience. 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 platforms understand their users and target the recommendation to the specific population of sellers, the algorithm could increase the price level. As online sales platforms gather more and more data about their users, recommendations are more likely to be tailored to specific markets. Our results suggest that this could lead to an increased risk of collusion.

4.3.3 Price patterns across treatments

In Figure 2, we plot the market prices for each treatment by supergame and round. In the initial round, market prices in RecSoft and RecTheory are higher than in Baseline following the recommendation of $p^R_{t=1} = 10$ (p-value=0.052 & p<0.05, two-sided Mann–Whitney U tests).[22] Yet, there are deviations from the recommendation in 86.1% of all markets in the first round. As a result, the treatment-specific punishment mechanisms take effect in the subsequent round.

Let us focus first on the pattern of RecTheory 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 completely at this level. In the following rounds, there are reoccurring deviations after a recommendation at the monopoly price. This results in clearly visible spikes in the price pattern. In the second and third super games, the spikes become less frequent, and the price patterns are more similar to Baseline.

The recurring deviations in RecTheory are almost entirely driven by matching groups with below median market prices (RecTheory-Low) as discussed in Section 4.3.1. It becomes clear when assessing the market price patterns for RecTheory for both subgroups

[22]Market prices in RecTheory and RecSoft 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.

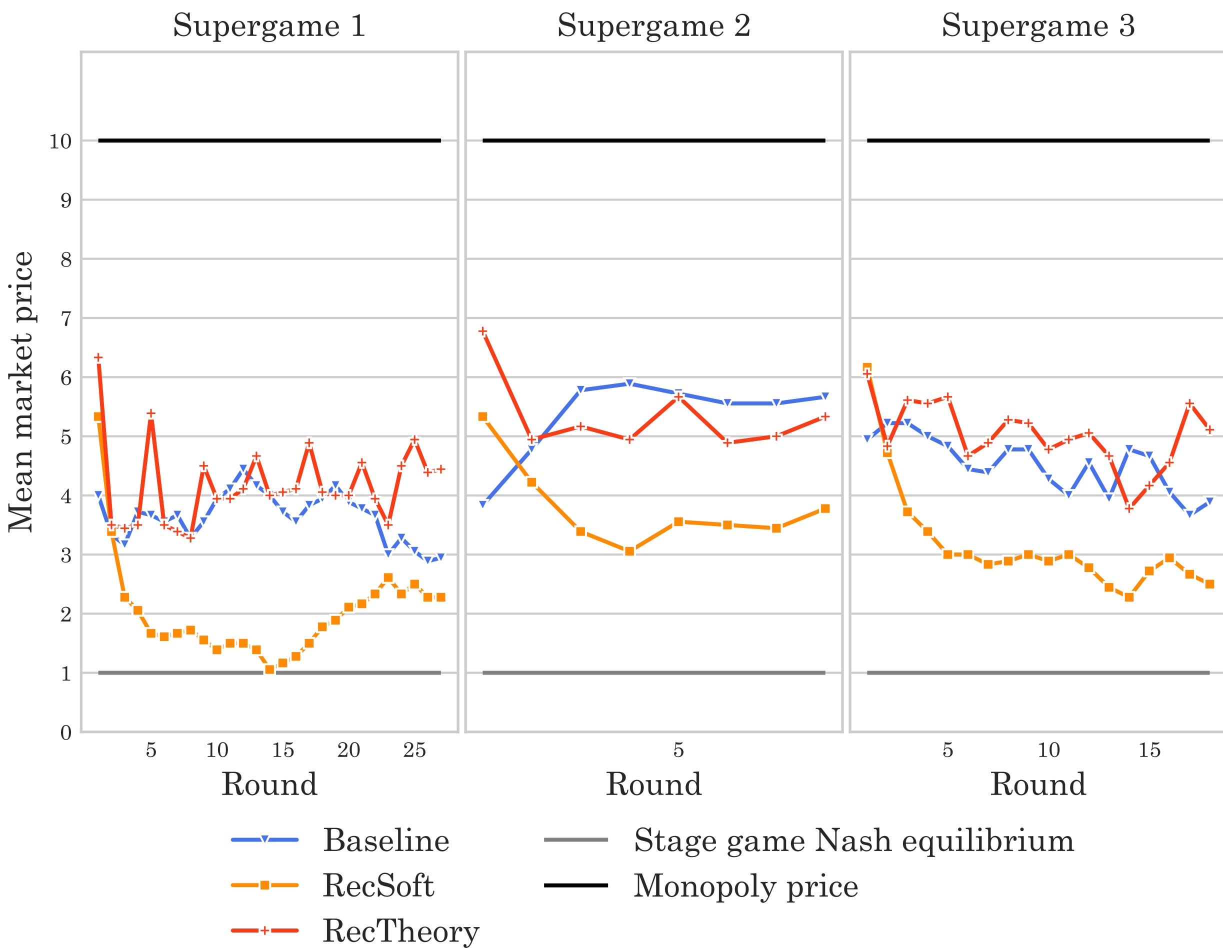

Figure 2: Market price for each treatment by supergame and round.

separately in Figure 3.[23] Whereas there are deviations in both subgroups in the first round, the market prices stabilize in RecTheory-High after the initial punishment phase. In RecTheory-Low, the share of markets with deviations from the collusive recommendations is significantly higher after the first round, which results in price spikes (p-value<0.05, two-sided Mann–Whitney U test).[24]

[23] For the respective analysis for RecSoft see Figure C.2.

[24] We 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 RecTheory-High and RecTheory-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.

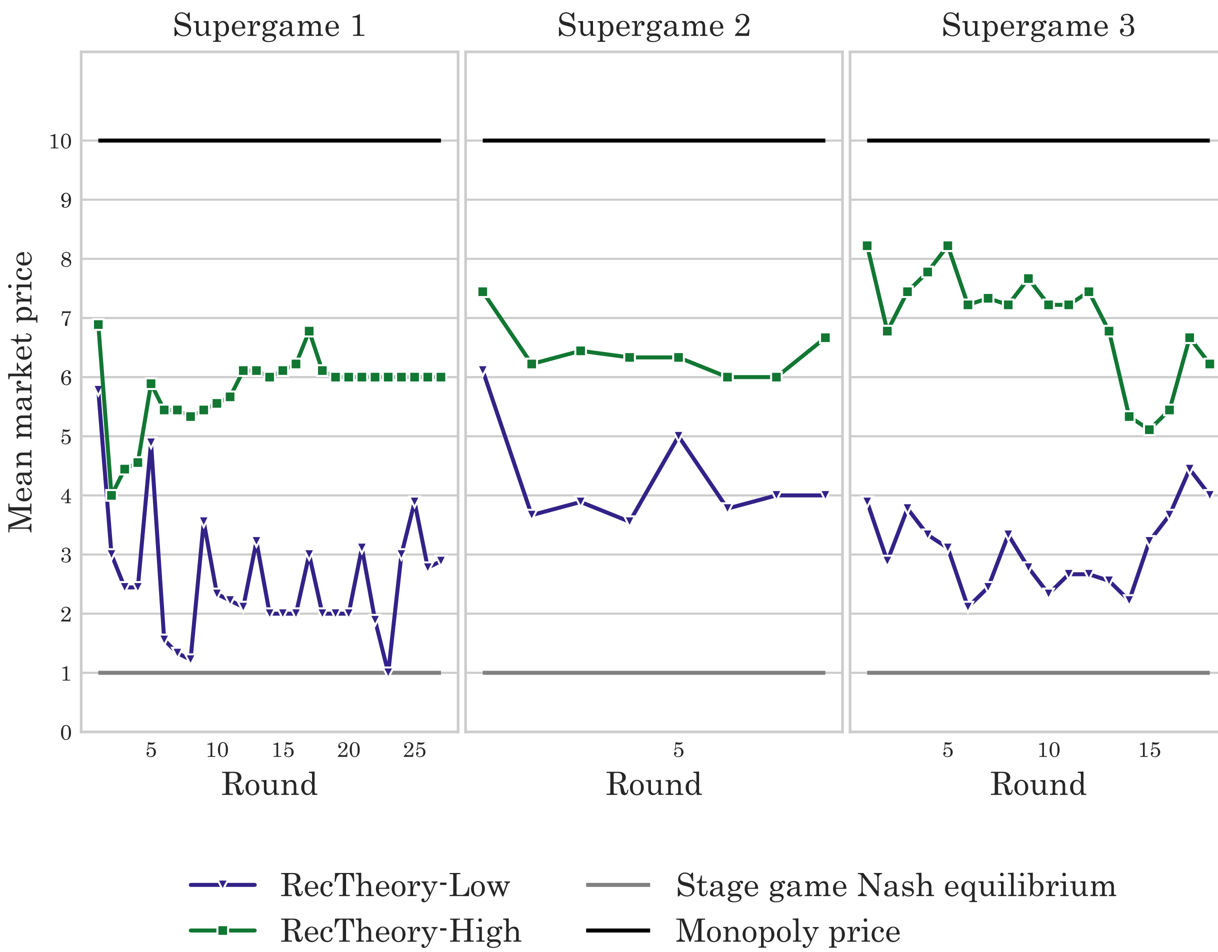

Figure 3: Market price in RecTheory for matching groups above (High) and below (Low) the median market price by supergame and round.

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 experiment.[25] Hence, we find suggestive evidence that the recommendation in RecTheory 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 RecSoft, 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

[25] Similarly, 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.

punish at the stage game Nash equilibrium price of $p^N = 1$ (see Section 2). We expected the punishment to be softer and, possibly, short compared to RecTheory. Yet, 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 ($\overline{p}^R_{t=2} = 5.33$). Hence, in contrast to RecTheory, 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 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-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.

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 drive the market prices down.*

Result 5 again emphasizes the adverse effects that recommendations can have for a platform if they are not designed appropriately. Furthermore, it suggests that platforms can use recommendations to decrease sellers' prices. In specific scenarios, this can be attractive, for instance, to avoid excessive double marginalization. Platforms could specifically design a recommendation algorithm to foster competition among the sellers. Our results suggest those price recommendations are feasible by using recommendation patterns as in RecSoft.

5 Concluding remarks

Price recommendations are a vital feature of many online markets. Companies like Airbnb, Expedia, and eBay give non-binding price recommendations to the sellers operating on their platforms. Furthermore, also in traditional markets, firms use recommendation algorithms to optimize their pricing decisions. While those pricing algorithms may be beneficial for the firms using them, competition authorities are concerned that those recommendations can dampen competition by helping firms to coordinate on non-competitive prices (Federal Trade Commission, 2021; Bundeskartellamt and Autorité de la concurrence, 2019). For instance, a recent report on the U.S. rental market suggests that such algorithms may lead to higher rental prices by enabling coordination between landlords.[26]

We derive two rule-based recommendation algorithms and study their effects on seller collusion in a stylized Bertrand market environment. Both algorithms have the objective 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 in which each participant represents a seller.

We find clear evidence that the recommendations influence the sales prices 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. This pass-on rate is higher for RecTheory. However, the effects on the realized market prices differ sharply between the different recommendation algorithms.

[26] See Vogell, Coryne & Little, "Technology Rent Going Up? One Company's Algorithm Could Be Why.", `https://www.propublica.org/article/yieldstar-rent-increase-realpage-rent`, last accessed November 25, 2022

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 heterogeneities in market outcomes. The collusive effects depend on the seller's characteristics. RecTheory lowers market prices in markets with low levels of negative reciprocity among sellers. Moreover, we find evidence that the recommendation can make markets more collusive if sellers are too impatient to commit to a collusive pricing strategy.

For the behaviourally motivated algorithm RecSoft, which can recommend brief punishment phases with moderate price levels, we find lower market prices and profits compared to the case without any recommendation. Participants frequently deviate from the recommendation, which starts a downward spiral that lowers prices. Similarly to RecTheory, market prices are lower than in Baseline for markets with sellers that have low negative reciprocity. There is no evidence that RecSoft facilitates collusion for any subgroup. Yet, it can be used to foster competitive outcomes and lower market prices for consumers.

We view our research as one of the first steps in understanding the effects of price recommendation algorithms on seller competition in digital sales platforms and other marketplaces. We demonstrate theoretically and experimentally that recommendations can benefit a platform by influencing sellers' pricing in the platform's favor. However, the experimental results indicate that it depends on the type of algorithm and sellers whether the effects are collusive. We find experimentally that recommendation algorithms can facilitate seller collusion if designed appropriately and if the sellers are rather impatient. Thus, recommendations may in particular foster collusion and harm consumers if sales platforms understand the sellers' characteristics and target the recommendation based on these characteristics.[27]

In other cases, we find that recommendation algorithms may have no price effects or even decrease prices, despite being designed and intended to facilitate collusion. The find-

[27] For 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.

ing 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 of the algorithm with soft punishments 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 at increasing prices and profits, in line 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.

While the results are, on balance, not alarming regarding the collusive risks of recommendation algorithms, we do provide reasons for potential concern. It is important to note that, in our experiment, the only purpose of the recommendations was to coordinate sellers. 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, but 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.

References

Bó, Pedro Dal and Guillaume R Fréchette, "On the determinants of cooperation in infinitely repeated games: A survey," *Journal of Economic Literature*, 2018, *56* (1), 60–114.

_ **and** _ , "Strategy choice in the infinitely repeated Prisoner's Dilemma," *American Economic Review*, 2019, *109* (11), 3929–52.

Brown, Zach and Alexander MacKay, "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics (forthcoming)*, 2022.

Bruttel, Lisa, "The Effects of Recommended Retail Prices on Consumer and Retailer Behaviour," *Economica*, 2018, *85* (339), 649–668.

Buehler, Stefan and Dennis L Gärtner, "Making sense of nonbinding retail-price recommendations," *American Economic Review*, 2013, *103* (1), 335–59.

Bundeskartellamt and Autorité de la concurrence, "Algorithms and Competition," Available at: `https://www.bundeskartellamt.de/SharedDocs/Publikation/EN/Berichte/Algorithms_and_Competition_Working-Paper.pdf` 2019.

Cadsby, C. Bram, Elizabeth Maynes, and Viswanath Umashanker Trivedi, "Tax compliance and obedience to authority at home and in the lab: A new experimental approach," *Experimental Economics*, 2006, *9* (4), 343–359.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review*, 2020, *110* (10), 3267–97.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Chen, Le, Alan Mislove, and Christo Wilson, "An empirical analysis of algorithmic pricing on amazon marketplace," in "Proceedings of the 25th International Conference on World Wide Web" International World Wide Web Conferences Steering Committee 2016, pp. 1339–1349.

Croson, Rachel and Melanie Marks, "The effect of recommended contributions in the voluntary provision of public goods," *Economic Inquiry*, 2001, *39* (2), 238–249.

Duffy, John and Nick Feltovich, "Correlated equilibria, good and bad: an experimental study," *International Economic Review*, 2010, *51* (3), 701–721.

Dufwenberg, Martin and Uri Gneezy, "Price competition and market concentration: an experimental study," *International Journal of Industrial Organization*, 2000, *18* (1), 7–22.

Engel, Christoph, "How much collusion? A meta-analysis of oligopoly experiments," *Journal of Competition Law & Economics*, 2007, *3* (4), 491–549.

Ezrachi, Ariel and Maurice E Stucke, "Artificial intelligence & collusion: When computers inhibit competition," *U. Ill. L. Rev.*, 2017, p. 1775.

Faber, Riemer P and Maarten CW Janssen, "On the effects of suggested prices in gasoline markets," *The Scandinavian Journal of Economics*, 2019, *121* (2), 676–705.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde, "The preference survey module: A validated instrument for measuring risk, time, and social preferences," *Management Science (forthcoming)*, 2021.

Federal Trade Commission, "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.

Fonseca, Miguel A and Hans-Theo Normann, "Explicit vs. tacit collusion—The impact of communication in oligopoly experiments," *European Economic Review*, 2012, *56* (8), 1759–1772.

Foros, Øystein and Frode Steen, "Vertical control and price cycles in gasoline retailing," *The Scandinavian Journal of Economics*, 2013, *115* (3), 640–661.

Greiner, Ben, "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

Hansen, Karsten T, Kanishka Misra, and Mallesh M Pai, "Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms," *Marketing Science*, 2021, *40* (1), 1–12.

Harrington, Joseph E. Jr., "The effect of outsourcing pricing algorithms on market competition," *Management Science*, 2022.

_ , Roberto Hernan Gonzalez, and Praveen Kujal, "The relative efficacy of price announcements and express communication for collusion: Experimental findings," *Journal of Economic Behavior & Organization*, 2016, *128*, 251–264.

Harstad, Ronald, Stephen Martin, and Hans-Theo Normann, *Intertemporal pricing schemes: experimental tests of consciously parallel behavior in oligopoly* Applied Industrial Organization, Cambridge University Press, 1998.

Hill, Dan, "How much is your spare room worth?," *IEEE Spectrum*, 2015, *52* (9), 32–58.

Holt, Charles A and Douglas Davis, "The effects of non-binding price announcements on posted-offer markets," *Economics letters*, 1990, *34* (4), 307–310.

Huck, Steffen, Hans-Theo Normann, and Jörg Oechssler, "Two are few and four are many: number effects in experimental oligopolies," *Journal of Economic Behavior & Organization*, 2004, *53* (4), 435–446.

Hunold, Matthias, Ulrich Laitenberger, and Guillaume Thébaudin, "Bye-box: An Analysis of Non-Promotion on the Amazon Marketplace," Working Paper 2022.

Janssen, Maarten and Edona Reshidi, "Regulating recommended retail prices," *International Journal of Industrial Organization*, 2022, *85*, 102872.

Johnson, Justin, Andrew Rhodes, and Matthijs R Wildenbeest, "Platform design when sellers use pricing algorithms," *Available at SSRN 3753903*, 2020.

Klein, Timo, "Autonomous algorithmic collusion: Q-learning under sequential pricing," *The RAND Journal of Economics*, 2021, *52* (3), 538–558.

Lefez, Willy, "Price recommendations and the value of data: A mechanism design approach," Working Paper 2021.

Lubensky, Dmitry, "A model of recommended retail prices," *The RAND Journal of Economics*, 2017, *48* (2), 358–386.

Musolff, Leon, "Algorithmic Pricing Facilitates Tacit Collusion: Evidence from E-Commerce," in "Proceedings of the 23rd ACM Conference on Economics and Computation" 2022, pp. 32–33.

Normann, Hans-Theo and Martin Sternberg, "Human-algorithm interaction: Algorithmic pricing in hybrid laboratory markets," *European Economic Review*, 2023, *152*, 104347.

Pavlov, Vladimir and Ron Berman, "Price Manipulation in Peer-to-Peer Markets and the Sharing Economy," Working Papers 19-10, NET Institute September 2019.

Roth, Alvin E and J Keith Murnighan, "Equilibrium behavior and repeated play of the prisoner's dilemma," *Journal of Mathematical Psychology*, 1978, *17* (2), 189–198.

Schlütter, Frank, "Managing Seller Conduct in Online Marketplaces and Platform Most-Favored Nation Clauses," Working Paper 2022.

Schotter, Andrew, "Decision making with naive advice," *American Economic Review*, 2003, *93* (2), 196–201.

_ **and Barry Sopher**, "Social learning and coordination conventions in intergenerational games: An experimental study," *Journal of Political Economy*, 2003, *111* (3), 498–529.

Silverman, Dan, Joel Slemrod, and Neslihan Uler, "Distinguishing the role of authority "in" and authority "to"," *Journal of Public Economics*, 2014, *113*, 32–42.

Sonntag, Axel and Daniel John Zizzo, "Institutional authority and collusion," *Southern Economic Journal*, 2015, *82* (1), 13–37.

Teh, Tat-How, "Platform Governance," *American Economic Journal: Microeconomics (forthcoming)*, 2022.

Werner, Tobias, "Algorithmic and human collusion," *Available at SSRN 3960738*, 2022.

Wieting, Marcel and Geza Sapi, "Algorithms in the marketplace: An empirical analysis of automated pricing in e-commerce," *Available at SSRN 3945137*, 2021.

Wright, Julian, "Punishment strategies in repeated games: Evidence from experimental markets," *Games and Economic Behavior*, 2013, *82*, 91–102.

Appendix A Additional theoretical results

In Section 2 we use a stylized model that treats the platform as a black box and abstracts from contracts between the platform and the sellers to focus on how sellers react to collusive algorithms. In this section, we instead focus on contracting between a platform and various sellers to analyze how the optimal price level depends on the commission rates and costs. We show that a platform can benefit from collusive recommendations even if it has the bargaining power to choose the commission rate.

High commission rates versus collusive recommendations. Let us demonstrate why a monopoly platform can achieve higher profits through collusive price recommendations for sellers than when just charging a high commission rate to the sellers. This comparison is relevant as many online sales platforms use commission rates and a natural question is whether a sales platform needs recommendations to achieve the desired seller price level. This analysis adds to the developing literature on seller collusion on online platforms, which focuses on settings where commission rates are sufficient for achieving high prices (Schlütter, 2022; Teh, 2022).[28] Teh (2022) has a related finding whereby it can be optimal for a platform to increase the seller margins through platform design, such as entry regulations if that is value-generating. He does, however, not consider seller collusion and uses a different modeling approach.

There are $n \geq 2$ symmetric sellers who sell differentiated products with marginal costs c. Each seller i makes a profit of

$$\pi_i = (p_i - r \cdot p_i - c)q_i(p_i, p_{-i})$$

when selling on the platform which charges a commission rate of r. Demand q_i has the usual properties and, in particular, decreases in the own price p_i and increases in the price(s)

[28] Schlütter (2022) primarily studies price parity clauses in a market where sellers can alternatively sell only via their direct sales channel.

p_{-i} of the competitors. A seller's (opportunity) costs of being active on the platform are $I \geq 0$, such that the seller participates if and only if

$$\pi_i \geq I. \tag{1}$$

We focus on situations where the platform wants to ensure that all sellers participate, so that condition (1) holds for all i.

A seller's first order condition with respect to its sales price is

$$(p_i - r \cdot p_i - c)\frac{\partial q_i(p_i, p_{-i})}{\partial p_i} + (1-r)q_i(p_i, p_{-i}) = 0$$

and can be written as

$$(p_i - \frac{c}{(1-r)})\frac{\partial q_i(p_i, p_{-i})}{\partial p_i} + q_i(p_i, p_{-i}) = 0. \tag{2}$$

Let $p^*(r,c)$ denote the symmetric Nash equilibrium price that solves the above equation when all n sellers compete.

At the competitive sales prices, the platform makes a profit of

$$\Pi(r) = r \cdot p^*(r,c) \cdot n \cdot q_i(p^*, p^*).$$

Platform profit maximization when sellers compete. For $c = 0$ the platform cannot influence the price level with r as it disappears in the first order condition (2).

For $c > 0$ the implicit function theorem on the first order condition (2) for the case of symmetric sales prices $p_i = p_{-i} = p^*$ yields $\partial p^*/\partial r > 0$ under the standard assumptions of a strictly concave seller profit π_i in p_i and decreasing demand ($\partial q_i/\partial p_i < 0$). The platform can thus raise the price level as long as selling remains profitable for the sellers.

Suppose that the sellers make lower profits if their common input costs increase. This is consistent with economic intuition and holds under standard demand assumptions. A sufficient condition for this is that the sellers' profit margin decreases as the costs increase: $\partial p^*/\partial k < 1$ with $k = c/(1-r)$. This is, for instance, the case with linear demand.

For illustration, suppose that $r = 0$ and that the resulting seller profits equal zero:

$$(p^* - c)q_i(p^*, p^*) - I = 0.$$

It is thus not feasible for the platform to charge a positive commission rate as the sellers would not break even. This argument generalizes to the case where break-even occurs at a positive commission rate that yields a price level $\hat{p}$ which is below the level which maximizes the industry profit. The platform is then restricted in the setting of the commission rate and thus cannot maximize the industry profit.[29]

Conversely, it might be that the platform achieves the industry profit maximizing price at a commission rate where the sellers make positive profits ($\pi_i > I$). This occurs if I is small enough. The platform then leaves more profits than necessary for participation to the sellers. It would thus be optimal for the platform to charge a higher commission rate while keeping the sales prices constant. This relates to the problem of double marginalization.

Platform profit maximization when sellers collude. For simplicity, assume that the platform can implement any price level p through recommendations. The platform can thus implement a price p and set r such that

$$(p \cdot (1-r) - c) \cdot q_i(p, p) = I.$$

The platform can thus implement the industry maximizing price

$$p^M = \arg\max_p \, (p - c) \cdot \sum_{i=1}^{n} q_i(p, p) - n \cdot I$$

[29]Fixed fees might solve the problem. However, in particular transfers to sellers might not work in practise. For instance, they might incentivize people to register as sellers just to obtain the transfers.

and extract through r all seller revenues up to $I + c \cdot q_i(p^M, q^M)$ per seller.

In summary, this analysis shows that a platform can benefit from prices recommendations even if it can change its commission rates for one of the following reasons:

- The sellers have (opportunity) costs of selling on the platform, such that a high commission rate is not acceptable but would be necessary for achieving high sales prices of competing sellers.
- The platform charges a commission rate and the sellers do not have marginal costs other than the commission payment, so that the commission rate does not affect the sellers' pricing.
- A high commission rate is optimal to extract the seller profits but yields too high sales prices (excessive double marginalization). In this case recommendations below the competitive level can be optimal.
- In addition to the above formal analysis, a platform might desire to charge the same commission rate across different markets to maintain a simple transparent policy albeit different seller price levels are optimal.

Appendix B Instructions and survey questions

B.1 Instructions

Hello and welcome to our experiment. In the next hour, you will make decisions on a computer. Please read the instructions carefully. All participants will receive the same instructions. You will also find a printed copy of these instructions at your seat. You will remain completely anonymous to us and to the other experiment participants. We will not save any data associated with your name.

Particularly important: Do not talk to your neighbors, do not use your cell phone, and keep quiet throughout the experiment. If you have any questions, please let us know. We will then come to your site and help.

In this experiment, you will repeatedly make pricing decisions. These allow you to earn real money. How much you earn depends on your decisions and on those of your fellow players. **Regardless, you will receive 4.00 euros for participating.**

In the experiment, we use a fictional monetary unit called ECU. After the experiment, the ECU will be converted to euros and paid to you. **Here, 100 ECU equal one euro.** The euro amounts are rounded to the first decimal place.

Example:
Participant A earned 465 ECU in the experiment. Converted, this is equal to 4.65 euros. Rounded to the first decimal place, Participant A is paid 4.70 euros. **Explanations**

In this game you represent a company in a virtual product market. In the market, two other companies sell the same product as you do. These companies are represented by two other experiment participants. The game has several rounds. You will meet the same companies (i.e. experiment participants) in each round of the game.

All companies decide again **independently and simultaneously** in each round,for how many ECU you want to sell your product. You can sell your product for a price of 1, 2, ... or 10 ECU to sell(whole units only). **There are no costs of production.** Your profit is the product of price and the number of units sold. In formal terms:

Profit = price x units sold.

The market has 30 identical customers. Each customer wants to buy **one unit** of the product as cheaply as possible in each round of a game. Each customer is willing to spend up to 10 ECU for that unit of the product.

The company with the **lowest price in the respective round** sells its products. **So the lowest price is the market price in that round**. Firms with a **price greater than the market price do not sell any products in that round and therefore receive a profit of zero.** If two or all three firms want to sell their product for the same market price, the demand is split evenly between the two or three firms.

Examples Exampe 1 Firm A sets a price of 4, Firm B sets a price of 4, Firm C sets a price of 6. Thus, Firms A and B together have set the lowest price. Firms A and B both sell the same amount of products, both firms have 15 customers each and thus get the same profit of 60 ECU. Firm C sells nothing and has a profit of 0.

	Firm A	Firm B	Firm C
Prices	4	4	6
Profits	60	60	0

Example 2: Firm A sets a price of 7, Firm B sets a price of 7, Firm C sets a price of 7. Thus, Firms A, B and C together have set the lowest price. They all sell the same amount of products (10 each) and thus get the same profit of 70 ECU.

	Firm A	Firm B	Firm C
Prices	7	7	7
Profits	70	70	70

Example 3: Firm A sets a price of 1, firm B sets a price of 4, firm C sets a price of 10. Thus, firm A has set the lowest price. Firm A is the only one that sells the product at a price of 1 to all 30 customers and thus gets a profit of 30 ECU. Firms B and C both sell nothing and have a profit of 0.

	Firm A	Firm B	Firm C
Prices	1	4	10
Profits	30	0	0

Price recommendations

Before you choose your price in each round, you receive a specific **price recommendation** from a computer algorithm. **All three firms in the market receive the same price recommendation.**

The algorithm aims to **maximize the total profits of all firms across all rounds**. Therefore, you will be given a recommendation that will allow all firms to make the highest possible profit in the long run. This means that the algorithm does **not** necessarily recommend a price that achieves the highest possible profit in a single round. **It recommends prices that achieve a high total profit over the entire game.**
The algorithm itself is not a market participant and cannot generate profits, **it only serves as information for all participants**.

Note: The recommended price is only a **proposal**. You are free to set any other price than the recommended one.

Duration of the experiment

After each round, all firms are informed about the chosen prices of all three firms and their own profits. In the next round, each firm has again the opportunity to choose their price. You interact with the same participants in each round within a game.

After each round, a random mechanism decides whether another round is played or the game ends. The probability that another round will be played is 95%. The game therefore ends after each round with a probability of 5%.
In other words, the computer throws a virtual dice with 20 sides before each possible further round. The result decides whether another round is played or not. With the number 20, the game is over, with all other numbers, another round is played.

Note:

You play the described game a total of three times. After each game, you will be put together with new participants to form a new market. This means that in each of the three games you interact with other participants.

After all games are finished, it will be randomly decided which of the three games will be paid out. You will receive your payoff after the experiment. You will also receive an additional 4.00 euros for participating in this experiment.

As a help we display a virtual calculator, with which you can calculate your profits in each round. **Comprehension Questions**

Question 1: How many consumers are in the market who want to buy the product?

- 25
- 35
- 30
- 40

Question 2: What is the probability of playing another round after completing one?

- 95%
- 5%
- 50%

Question 3: You are firm A and choose a price of 2, firm B chooses a price of 10, firm C chooses a price of 9. What is your profit in ECU in this round?

Question 4: You are firm A and choose a price of 8, firm B chooses a price of 8, firm C chooses a price of 8. What is your profit in ECU in this round?

Question 5: You have a profit of 650 ECU, what is your profit in euros?

Questioni 6: What is the objective of the algorithm?

- Maximizing profits for all firms in a single round
- Maximizing total profits for all firms across all rounds
- Maximizing total profits for a single firm across all rounds
- Maximizing profits of a single firm in a single round

B.2 Survey questions

Gender: What is your gender?

- Male
- Female
- Diverse
- No specification

Experiments: In how many economic experiments have you (approximately) already participated?

GPA (School): What was the final grade of your last school diploma (1.0 - 4.0)?

Math Grade: What was your last math grade (1.0 - 6.0)?

Budget: How much money do you have available each month (after deducting fixed costs such as rent, insurance, etc.)?

Spending: How much money do you spend each month (after deducting fixed costs such as rent, insurance, etc.)?

Risk: Are you generally a person who is willing to take risks or do you try to avoid risks? Please indicate your answer on a scale of 0 to 10, where 0 means not willing to take risks at all and 10 means very willing to take risks.

Time: Compared to others, are you generally willing to give up something today in order to benefit from it in the future, or are you unwilling to do so compared to others? Please indicate your answer on a scale of 0 to 10, where 0 means not willing to give up at all and 10 means very willing to give up something.

Trust: As long as I am not convinced of the opposite, I always assume that other people only have the best in mind. How strongly do you agree with this statement? Please indicate your answer on a scale of 0 to 10, where 0 means not true at all and 10 means very true.

Neg. Rec.: Are you someone who is generally willing to punish unfair behavior, even if it comes at a cost for you, or are you unwilling to do so? Please indicate your answer on a scale of 0 to 10, where 0 means not willing to punish at all and 10 means very willing to punish.

Pos. Rec.: If someone does me a favor, I'm willing to return it. How strongly do you agree with this statement? Please indicate your answer on a scale of 0 to 10, where 0 means not true at all and 10 means very true.

Altruism: Imagine the following situation: You won 1,000 € in a prize competition. How much would you donate to charity in your current situation?

Appendix C Further results

C.1 Randomization checks

Table C.1 provides the average outcome for different survey measures for the main treatments. Furthermore, we test for differences in those measures across treatments using Kruskal-Wallis tests. A complete list of the different survey questions we ask the participants is provided in Appendix B.2. There are few negligible differences in the control variables across treatments. Only the budget participants have each month differs between treatment at the 5%-level. Importantly, controlling for those survey measures does not influence the main outcomes (see Table 2 and 3). Thus, we conclude that randomization into treatments worked as expected.

Table C.1: Survey measures by treatment

	Risk	**Time**	**Trust**	**Neg. Rec.**	**Pos. Rec**	**Altruism**	**Woman**
BASELINE	0.53	0.68	0.39	0.64	0.92	0.12	0.48
RECSOFT	0.49	0.70	0.41	0.52	0.88	0.14	0.50
RECTHEORY	0.54	0.74	0.49	0.62	0.84	0.10	0.54
P-values	0.73	0.35	0.18	0.09	0.19	0.84	0.84

	Experiments	**Math Grade**	**GPA (School)**	**Budget**	**Spending**
BASELINE	7.24	1.99	1.99	414.98	300.15
RECSOFT	11.52	2.47	2.31	378.61	275.63
RECTHEORY	10.69	2.12	2.05	523.89	339.87
P-values	0.38	0.07	0.05	0.04	0.39

Note: The preferences measures are based on the survey questions by Falk et al. (2021) and scaled between zero and one. The p-values are based on Kruskal-Wallis-tests.

C.2 Economic preferences and recommendations

In Section 4, we discuss the influence of negative reciprocity on market prices when participants receive a price recommendation (see Table 5). Here, we provide the same analysis for the other economic preferences measures. All measures have been normalized to be between zero and one. Furthermore, as the measures have been elicited on the individual level, we aggregated them by calculating the group specific mean.

Altruism and trust do not influence market prices. Interestingly, there is also no significant effect of positive reciprocity on market outcomes. In other words, while differences in negative reciprocity lead to vastly different prices within and between treatment, it is not the case for positive reciprocity.

Table C.2: Market price explained by altruism and treatments

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	4.00***	4.18***	
	(0.575)	(1.11)	
Altruism	-1.34	0.191	0.191
	(3.34)	(8.56)	(8.61)
RecTheory		-0.209	-0.209
		(1.65)	(1.65)
RecSoft		-1.10	-1.10
		(1.14)	(1.14)
Altruism × RecTheory		6.43	6.43
		(10.9)	(10.9)
Altruism × RecSoft		-3.14	-3.14
		(8.65)	(8.70)
Fixed-effects			
Round			Yes
Supergame			Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table C.3: Market price explained by positive reciprocity and treatments

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	0.833	-1.57	
	(3.33)	(4.97)	
Pos. Rec.	3.41	6.26	6.26
	(3.62)	(5.53)	(5.56)
RecTheory		3.82	3.82
		(5.54)	(5.57)
RecSoft		-3.13	-3.13
		(7.41)	(7.44)
Pos. Rec. × RecTheory		-3.43	-3.43
		(6.12)	(6.15)
Pos. Rec. × RecSoft		2.07	2.07
		(8.31)	(8.36)
Fixed-effects			
Round			Yes
Supergame			Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table C.4: Market price explained by risk preferences and treatments

Dependent Variable:		Market price	
Model:	(1)	(2)	(3)
Variables			
(Intercept)	2.46**	1.84	
	(0.964)	(1.07)	
Risk	2.65	4.45*	4.45*
	(1.55)	(2.16)	(2.17)
RecTheory		4.21	4.21
		(3.03)	(3.05)
RecSoft		-0.654	-0.654
		(1.56)	(1.57)
Risk × RecTheory		-7.07	-7.07
		(4.81)	(4.83)
Risk × RecSoft		-1.45	-1.45
		(3.13)	(3.15)
Fixed-effects			
Round			Yes
Supergame			Yes
Observations	2,862	2,862	2,862

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table C.5: Market price explained by trust preferences and treatments

Dependent Variable:	Market price		
Model:	(1)	(2)	(3)
Variables			
(Intercept)	3.32***	4.60**	
	(0.863)	(1.71)	
TRUST	1.20	-1.01	-1.01
	(1.73)	(4.13)	(4.15)
RECTHEORY		-0.835	-0.835
		(3.06)	(3.07)
RECSOFT		-2.16	-2.16
		(1.77)	(1.78)
TRUST × RECTHEORY		2.79	2.79
		(5.58)	(5.61)
TRUST × RECSOFT		1.58	1.58
		(4.38)	(4.40)
Fixed-effects			
Round			Yes
Supergame			Yes
Fit statistics			
Observations	2,862	2,862	2,862
R^2	0.00319	0.05525	0.09202
Within R^2			0.05736

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

C.3 Additional control treatments

We consider two additional control treatments. In RECSTATIC, participants receive a static price recommendation at the monopoly in each period. Also, after deviations from the recommended price, the algorithm recommends the monopoly price, and there is no punishment mechanism. Furthermore, we consider the RECNASH algorithms. Similar to RECTHEORY, after any deviation from the monopoly price, the stage game Nash equilibrium is recommended in the subsequent period. Yet, the algorithm reverts back to the monopoly after one punishment round and is thus, in contrast to RecTheory, not incentive compatible. We consider for both algorithms only 36 subjects, which yields considerably less power than in the main treatments. The results are provided in Figure C.1. Both treatments yield similar market prices as in Baseline and RecTheory.

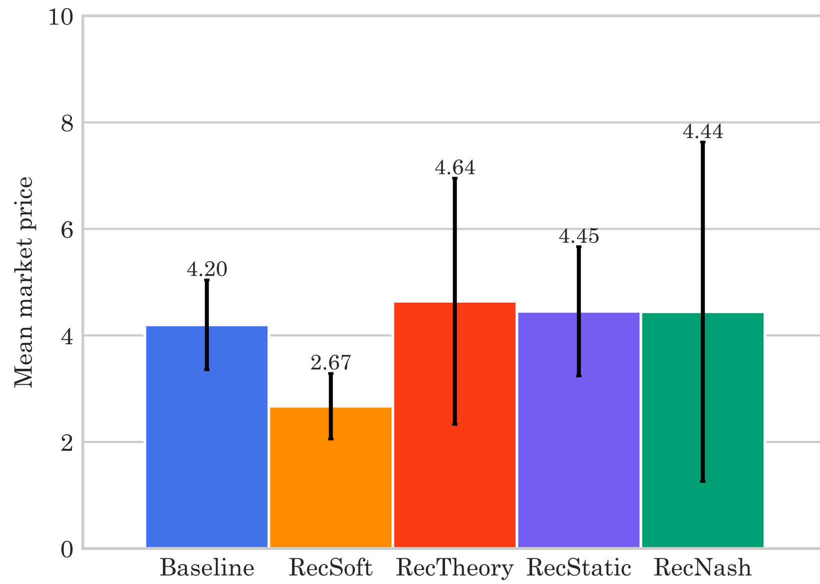

Figure C.1: Market price for all treatments. The error bars represent 95% confidence intervals.

C.4 Further figures and tables

Table C.6: Linear regression with the treatment effects by supergame

Dependent Variable:		Market price	
	(Supergame 1)	(Supergame 2)	(Supegame 3)
Variables			
RecSoft	-1.63**	-1.56*	-1.36*
	(0.572)	(0.881)	(0.740)
RecTheory	0.539	-0.007	0.497
	(0.990)	(1.27)	(1.26)
Fixed-effects			
Round	Yes	Yes	Yes
Fit statistics			
Observations	1,458	432	972
R^2	0.08748	0.04145	0.05861
Within R^2	0.07274	0.03805	0.04175

Clustered (Matching group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

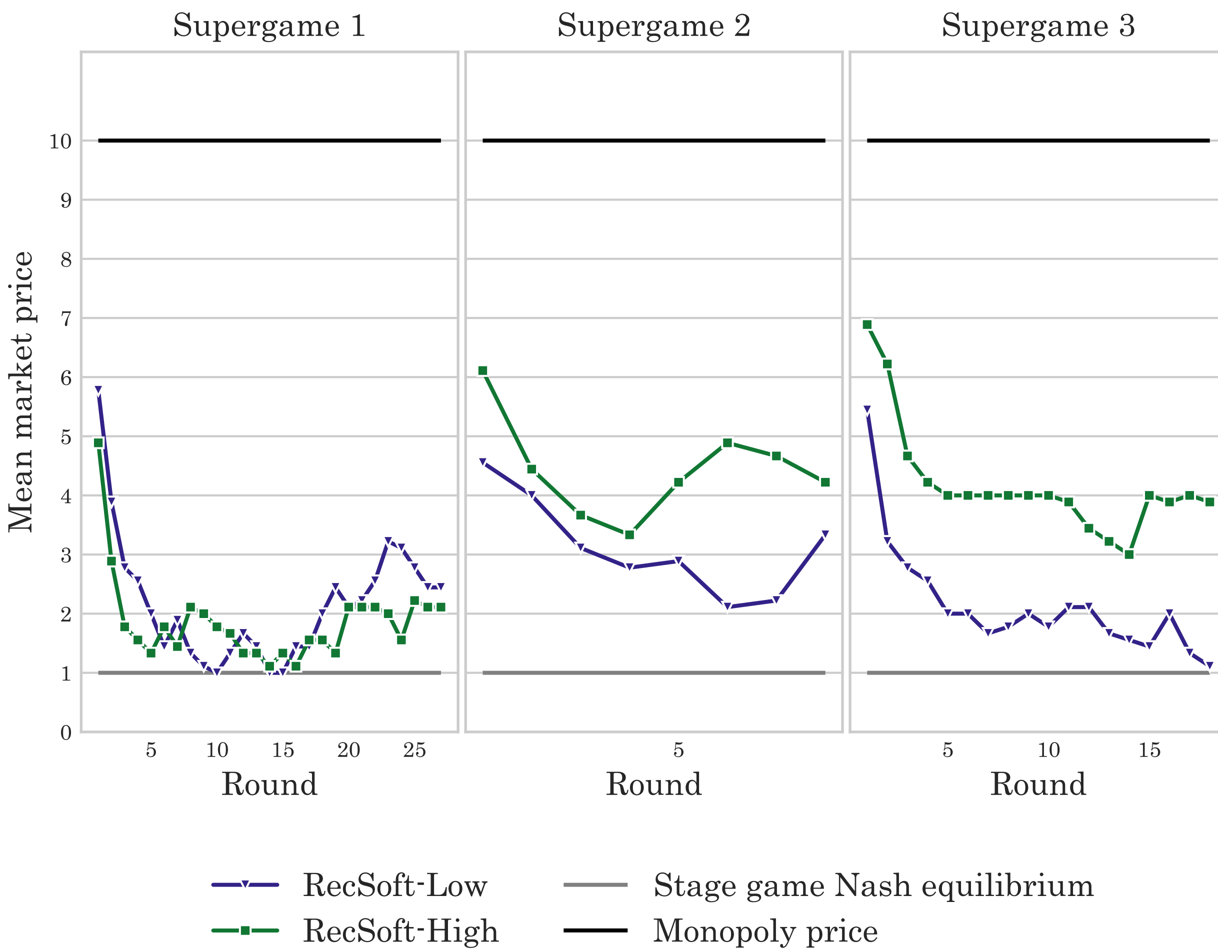

Figure C.2: Market price in RecSoft for matching groups above (High) and below (Low) the median market price by supergame and round.