

Dec 12th, 12:00 AM

Competition Between Human and Artificial Intelligence in Digital Markets: An Experimental Analysis

Andreas Schauer

Research Group Data Policies, University of Passau, andreas.schauer@uni-passau.de

Daniel Schnurr

University of Passau, daniel.schnurr@ur.de

Follow this and additional works at: <https://aisel.aisnet.org/icis2022>

Recommended Citation

Schauer, Andreas and Schnurr, Daniel, "Competition Between Human and Artificial Intelligence in Digital Markets: An Experimental Analysis" (2022). *ICIS 2022 Proceedings*. 6.
https://aisel.aisnet.org/icis2022/it_policy/it_policy/6

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ICIS 2022 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Competition Between Human and Artificial Intelligence in Digital Markets: An Experimental Analysis

Short Paper

Andreas Schauer

University of Regensburg
93040 Regensburg, Germany
andreas.schauer@ur.de

Daniel Schnurr

University of Regensburg
93040 Regensburg, Germany
daniel.schnurr@ur.de

Abstract

In digital markets business decisions are increasingly taken by artificial intelligence (AI). Especially in e-commerce, a growing share of retailers uses AI-driven algorithmic pricing, whereas remaining vendors rely on manual price setting. However, policymakers have raised concerns about anti-competitive tacit collusion between humans and AI that could allow firms to soften competition. Therefore, we empirically investigate outcomes that arise when humans and AI repeatedly interact in digital market environments. Based on an economic laboratory experiment in near real-time, we compare the degree of tacit collusion among humans and reinforcement learning algorithms to market settings where only humans or only algorithms compete. Preliminary findings demonstrate that tacit collusion emerges between humans and AI, although at lower levels than in settings with only humans or only algorithms. Altogether, our study sheds light on competition in digital markets where AI plays an increasingly important role and thus bears timely policy and managerial implications.

Keywords: Artificial Intelligence, Algorithmic Pricing, Human-Computer Interaction, Algorithmic Collusion, Reinforcement Learning, IT Policy, AI Regulation, Economics of IS

Introduction

Technological advances in computing power and data storage present growing opportunities for the use of artificial intelligence (AI) in organizational contexts. In consequence, business tasks and decisions are increasingly delegated to AI learning algorithms based on the premise that they solve more complex problems, derive more accurate predictions, and make more efficient decisions than their human counterparts (OECD 2017). A particular popular application area is algorithmic pricing, where algorithms determine and execute pricing strategies for goods and services in an automated manner (Calvano et al. 2020a; Chen et al. 2016).

Although algorithmic pricing has been used for air travel, hotel bookings, and retail since the 1980s (Calvano et al. 2020b; Ezrachi and Stucke 2017), the scope of application has recently expanded drastically, especially due to the growth in e-commerce (Chen et al. 2016). Thereby, prices are increasingly set by AI-driven algorithms that can process vast amounts of data and information (e.g., real-time data of customers and consumers), and learn how to set prices given pre-specified objectives (Calvano et al. 2020a). At the same time, widespread emergence of AI-driven algorithmic pricing has also raised significant concerns about its impact on competition in digital markets. Policymakers and scholars have questioned whether the use of pricing algorithms could lead to anti-competitive behavior through coordinated pricing that allows firms to soften competition, resulting in excessive consumer prices (Ezrachi and Stucke 2017; Mehra 2016; Schwalbe 2018). In particular, it is feared that AI-driven algorithmic pricing could facilitate collusion among firms without explicit communication, i.e., allow competing firms to implicitly coordinate their prices to achieve higher profits than in the competitive equilibrium (Harrington 2018; OECD 2017).

Thus, the questions of whether and under which conditions AI-based pricing algorithms indeed facilitate collusive outcomes have sparked a rapidly growing strand of academic literature. Calvano et al. (2020b), Klein (2021), and Waltman and Kaymak (2008) show that reinforcement learning algorithms are indeed able to learn collusive behavior by self-play in simulated market environments. Moreover, empirical findings by Assad et al. (2020) suggest that the adoption of pricing algorithms in local German gasoline markets has led to higher retail prices and softened competition. Despite this growing evidence that AI may be able to tacitly collude with other computer agents, far less is known about the outcomes and competitive processes in market settings, where algorithms compete with human decision makers, which we will refer to as *hybrid market settings* in this study. These hybrid market settings are very widespread in current practice. For example, about two-thirds of e-commerce vendors use computer programs to monitor and adjust prices, while the remaining vendors rely on manual price setting (European Commission 2017). Therefore, the high relevance of these market settings and the scant scientific knowledge call for more IS research on the interaction between humans and computers in the contexts of algorithmic pricing (Caro et al. 2022; Jain et al. 2018) and the regulation of AI technology (Gozman et al. 2020). To advance research on this topic, our study empirically investigates the outcomes that arise when humans and AI compete and repeatedly interact in dynamic market environments, focusing explicitly on e-commerce markets as an area of application. In an economic laboratory experiment, we compare competition in hybrid markets that include both humans and reinforcement learning algorithms to market settings where only human decision makers or only learning algorithms compete with each other, respectively. Based on a between-subjects experimental design, our study isolates the causal effects of AI-driven pricing algorithms on market outcomes and competition intensity in hybrid markets.

In particular, our findings aim to shed light on whether concerns about anti-competitive effects facilitated by AI arise in market settings where algorithms interact with human decision makers. Beyond existing studies on human-computer interaction (HCI) in AI contexts (e.g., Fügener et al. 2021), digital markets offer a novel and peculiar competitive setting where humans and algorithms compete against each other, but can also gain from mutual cooperation. Hence, our study explores a novel domain of HCI, where the decision to compete or to cooperate itself presents an endogenous learning challenge to AI and humans. Moreover, our study addresses highly relevant policy questions with direct implications for IT regulation (De Vaujany et al. 2018). If humans and AI were able to effectively coordinate on high price levels, this could significantly harm consumers and welfare in digital markets. Thus, our analysis has important ramifications on whether policy intervention and stricter regulation of AI-driven algorithmic pricing are warranted. Finally, our study also informs managers about the potential dangers of employing AI-driven pricing algorithms with respect to compliance and liability in the context of autonomous AI learning agents (Harrington 2018; Mehra 2016).

Background and Hypotheses

This study relates to research on (i) the effects of algorithmic pricing on competition and collusive behavior, (ii) empirical evidence of algorithmic collusion as identified by field studies and computer simulations and (iii) human-algorithm interaction in hybrid market settings investigated by experimental studies.

The Impact of Algorithmic Pricing on Competition

Algorithmic pricing has been found to have diverse effects on the competitive process and ensuing market outcomes. In particular, three pro-competitive effects can be identified based on the extant literature. First, as algorithms are able to react instantly to changes in the market, price adjustments can be made faster and more frequently (Calvano et al. 2020b; Chen et al. 2016). Second, the use of pricing algorithms can reduce firms' costs, as fewer (human) resources must be allocated to monitor the market as well as to manually determine and set prices (OECD 2017). Finally, smaller firms, especially online retailers, can benefit from lower barriers to market entry when algorithmic pricing is available as a service (see, e.g., Amazon Automate Pricing or SAP Dynamic Pricing) as they can cope with unknown demand structures and competition dynamics at lower costs (Schwalbe 2018). On the contrary, legal studies and policy reports have raised concerns about the anti-competitive effects of algorithmic pricing. In particular, these studies scrutinize the risk that algorithmic decision-making may facilitate (tacit) collusion and could thus lead to higher market prices. In this context, the taxonomy on algorithmic collusion, developed by Ezrachi and Stucke (2017), has

become popular among competition authorities (Competition and Markets Authority 2021; Monopolkommission 2018; OECD 2017). According to Ezrachi and Stucke, algorithms may facilitate explicit collusion through enhanced monitoring, control or prediction of explicit agreements among firms. However, they also highlight the novel challenges raised by autonomous learning algorithms that may be able to tacitly collude without neither an explicit agreement among competitors nor a clear intent to collude. Such implicit coordination, which is akin to human tacit collusion, is difficult to detect and prosecute as illegal collusive behavior. It is therefore also viewed by many competition authorities (e.g., the Competition and Markets Authority (2021) in the UK, the Monopolkommission (2018) in Germany, and the OECD (2017)) as the main current threat posed by AI-driven pricing algorithms. Thus, in this study, we focus on algorithmic collusion as tacit collusion and will use both terms interchangeably. Thereby, we follow the definition by Ivaldi et al. (2003), who established that "tacit collusion need not involve any "collusion" in the legal sense, and in particular need involve no communication between the parties" (Fn. 2), but rather indicates that market prices exceed the competitive outcome.

Evidence of Tacit Collusion among AI-driven Pricing Algorithms

Despite the rich literature on algorithmic pricing in general, empirical evidence on the impact of algorithmic pricing on competition, in particular on tacit collusion, remains scarce. Among the first studies, Assad et al. (2020) finds that the adoption of pricing algorithm in local German gasoline markets has facilitated tacit collusion, especially if algorithms were implemented by all market participants in duopolies. However, the exact pricing algorithms and their implementation within firms cannot be directly observed by the authors (Assad et al. 2020). As detection of tacit collusion in field data is notoriously hard, the effects of algorithmic pricing on competition have so far been mainly studied in computer simulations. Among the first studies, Waltman and Kaymak (2008) analyze AI-driven pricing algorithms in simulated oligopoly markets. Although, Waltman and Kaymak find that algorithms are able to earn profits above the theoretical competitive prediction, it has been questioned whether the observed actions are indeed indicative of collusive behavior or rather the result of algorithms' failure to learn an optimal strategy (Calvano et al. 2020b). Therefore, more recent simulation studies (Calvano et al. 2020b; Klein 2021) have focused on the analysis of strategies learned by AI algorithms in order to verify that outcomes are indeed driven by the algorithms' ability to tacitly coordinate on high prices. Most notably, Calvano et al. (2020b) show that algorithms can sustain price levels above the competitive market outcome by learning punishment strategies, which make it unprofitable for competitors to deviate from a collusive market outcome. These findings indicate that AI can achieve collusive price levels simply by being instructed to maximize their respective individual profit and without explicit communication (Calvano et al. 2020a; Calvano et al. 2020b; Klein 2021).

Human-Computer Interaction and Collusion in Hybrid Market Settings

Interaction between humans and algorithms is becoming prevalent in digital markets. For example, according to the European Commission (2017), two-thirds of e-commerce vendors now use computer programs to monitor and adjust prices whereas other vendors rely on manual price setting (see also Assad et al. 2020; Chen et al. 2016). Although there exists a rich literature on competition among human decision makers based on economic laboratory experiments (Potters and Suetens 2013), it is an open question whether findings from studies on human-human interaction in oligopoly markets carry over to settings where humans interact with AI-driven algorithms. First experimental studies to explicitly consider competition in hybrid market settings include Zhou et al. (2018) as well as Normann and Sternberg (2022), who employ simple rule-based algorithms that follow a cooperative strategy. Collusive outcomes are found in both studies when humans compete with algorithms. According to Zhou et al. (2018) collusive outcomes are achieved noticeable faster compared to human-human experiments. Normann and Sternberg (2022) find that the presence of a pricing algorithm significantly increases tacit collusion in markets with three competitors, but not in markets with four competitors. The most closely related study to our analysis is Werner (2022), who investigates competition between human decision makers and self-learning pricing algorithms in an economic laboratory experiment. The experimental findings indicate that for duopolies, market prices increase as more firms delegate pricing decisions to algorithms. Thus, in markets with two competitors, tacit collusion increases in the number of firms that employ a pricing algorithm.

Hypotheses

Based on the findings of collusion among AI-driven pricing algorithms in recent simulation studies (Calvano et al. 2020b; Klein 2021) and the established evidence on tacit collusion among human decision makers (Potters and Suetens 2013), we also expect collusion to tacitly arise in hybrid market settings.

Hypothesis 1. *Competition between humans and AI leads to collusive market outcomes, i.e., firms can tacitly coordinate on prices above the competitive outcome.*

Moreover, our study compares outcomes in hybrid markets to markets where only humans or only AI-driven algorithms compete, respectively. With regard to the latter market settings, simulation analyses of competition among AI-driven pricing algorithms have generally found high levels of tacit collusion (Calvano et al. 2020b; Klein 2021). Experimental analyses of competition among humans have also found significant levels of tacit collusion, however, with considerably more variance across individual market cohorts and with high dependence on the underlying market structure and market conditions (Horstmann et al. 2018; Potters and Suetens 2013). As algorithms can be trained on specific market conditions, we therefore conjecture that, in general, AI-driven algorithms will achieve a higher level of tacit collusion than humans for static market environments. Given the mix of human and algorithmic decision makers in hybrid market settings, we hypothesize that tacit collusion in these settings will then be on a medium level below collusive outcomes for algorithmic competition, but above collusive outcomes for human competition. However, there are additional behavioral effects and factors that could influence market outcomes for hybrid settings, such as algorithmic aversion or (missing) human trust in AI (see, e.g., Dietvorst et al. 2015). In consequence, tacit collusion between humans and AI could also fall below levels of human competition. Ultimately this is an empirical question, which we explicitly address by testing the following two hypotheses.

Hypothesis 2. *Competition between humans and AI yields a level of tacit collusion that is higher than for competition among humans.*

Hypothesis 3. *Competition between humans and AI yields a level of tacit collusion that is lower than for competition among AI-driven algorithms.*

Methodology

Economic laboratory experiments are particularly suited to investigate tacit collusion and competitive interactions between human and AI-based decision makers as they allow for maximum experimental control, thus guaranteeing high internal validity and also allowing the experimenter to exogenously match human and algorithmic decision makers. In the following, we describe the design, the computerized market environment, the reinforcement learning agents, and the procedures of our experimental study.

Experimental Design

To test our hypotheses, we vary the market setting by considering different types of decision makers in a market (humans and algorithms). To corroborate the robustness of our analyses, we also vary the market size (two and three firms). This yields a full-factorial 3x2 design with six treatment combinations, as illustrated by Table 1. The experiment employs a between-subjects design, i.e., each subject participates in exactly one treatment. This prevents possible carry-over effects and confounding effects of within-subjects designs. Subjects are aware of the treatment conditions that they are receiving, but they do not know that it is a treatment nor do they know about the other treatments. Treatments are randomized at the session level.

	Human vs. Human	Human vs. Algorithm	Algorithm vs. Algorithm
Duopoly	HH-2	HA-2	AA-2
Triopoly	HH-3	HA-3	AA-3

Table 1. Overview of the experimental treatments.

Market Environment

In our experiment, we consider the competition model by Singh and Vives (1984) and its generalization to more than two firms (Häckner 2000). This allows us to capture price competition with a parsimonious economic model that generalizes the original model of horizontally differentiated goods by Hotelling (1929). The competition model by Singh and Vives (1984) has proven to represent a suitable market environment for several other experimental oligopoly studies that have investigated tacit collusion (Horstmann et al. 2018). The model considers a market with $n \in \mathbb{N}$ symmetric firms. Each firm $i \in \{1, \dots, n\}$ produces and sells a single good. Marginal costs are assumed to be zero for all goods. Each firm sets the price p_i for its good. The prices of all firms in the market determine the respective quantity q_i sold per firm. The profit of each firm is given by $\pi_i = p_i q_i$. For a more detailed description of the competition model see Horstmann et al. (2018).

We consider two equilibrium predictions for the market outcome: (i) the competitive Nash equilibrium of the one-shot game, which is also the unique subgame perfect Nash equilibrium (SPE) of the finitely repeated game; In the Nash equilibrium, firm i maximizes its profit π_i with respect to its own price p_i taking into account competitors' optimal price setting p_{-i} . (ii) the collusive equilibrium, where firms maximize their joint profit like a hypothetical monopolist. In the collusive equilibrium, all n firms in the market employ joint profit maximization (JPM), i.e., they maximize $\sum_{i=1}^n \pi_i$ when choosing their price p_i . Joint profit maximization can be sustained as a SPE if the stage game is assumed to be infinitely repeated.

Based on these two equilibrium concepts, the *degree of tacit collusion* ϕ , can be measured as the relative deviation of the average price of all firms in the market $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$ from the Nash equilibrium p^{Nash} towards the collusive equilibrium p^{JPM} (Engel 2007; Horstmann et al. 2018). Formally this is $\phi = \frac{\bar{p} - p^{Nash}}{p^{JPM} - p^{Nash}}$. For $0 < \phi < 1$, the market outcome is considered to be collusive. For $\phi = 0$, the average market price corresponds to the Nash equilibrium p^{Nash} . For $\phi = 1$, firms behave like a hypothetical monopolist and the market is fully collusive. If the degree of tacit collusion $\phi < 0$, firms could increase prices and earn a higher profit in the competitive outcome. For $\phi > 1$, firms could increase the joint profit by lowering their prices.

Reinforcement Learning Agents

Depending on the treatment, the role of one or more firms in the experiment may be assumed by an autonomous software agent that implements a pricing algorithm. As we are interested in the idiosyncratic properties of AI when competing with human decision makers, we follow Calvano et al. (2020b), Klein (2021), and Werner (2022) in investigating reinforcement learning algorithms instead of rule-based pricing algorithms. Reinforcement learning is particularly suited for our experimental analysis, because the algorithm independently learns a strategy to maximize its received rewards, thus avoiding the need for training data (Sutton and Barto 2018). Reinforcement learning consists of an agent that continually interacts with its environment, which encompasses everything outside the agent. The agent learns through experience and chooses its action a_t based on the perceived current state of the environment s_t . After the agent has selected an action at step t , the environment proceeds to the new state s_{t+1} and returns a reward π_{t+1} (Sutton and Barto 2018). In our experiment, the action space is defined by integer prices to be chosen by firm i in the market, i.e., the price interval $[0, 100]$. The state is determined by the action(s) of firm i 's competitor(s) in the previous period. The reward obtained in period t is given by firm i 's profit π_i in the previous period.

Following Calvano et al. (2020b), Klein (2021), and Werner (2022), we employ Q-learning as a specific reinforcement learning algorithm. Q-learning has emerged as a popular algorithm for simulations and experiments in the literature, as it is widely used in practice and represents a particularly simple machine learning algorithm that requires to determine only few parameters ex ante (Calvano et al. 2020b). As with any reinforcement learning algorithm, Q-learning consists of two distinct interacting modules: a learning module and an action selection module (Klein 2021). The learning module is based on the agent's objective to maximize the sum of discounted future expected rewards which can be formally expressed by the Q-function

$$Q(s_t, a_t) = E \left[\pi_t + \delta \max_{a \in A} Q(s_{t+1}, a) \mid s_t, a_t \right]. \quad (1)$$

As both the action space A and the state space S are finite, the Q-function can be denoted as a $A \times S$ matrix.

In order to determine an optimal strategy, the Q-learning agent has to iteratively update the cells of this Q-matrix for each state-action combination (Werner 2022). The values of the Q-matrix are updated as follows:

$$Q_{t+1}(s, a) = \begin{cases} (1 - \alpha)Q_t(s, a) + \alpha(\pi_t + \delta \max_{a \in A} Q_t(s_{t+1}, a)) & \text{if } s = s_t \wedge a = a_t, \\ Q_t(s, a) & \text{otherwise.} \end{cases} \quad (2)$$

The discount factor $\delta \in [0, 1)$ denotes the time preference of the agent, while the learning rate $\alpha \in (0, 1]$ determines how newly experienced rewards and states are weighted against past experiences, i.e., past Q-values (Waltman and Kaymak 2008). For the action selection, the algorithm faces a trade-off between exploration and exploitation. In our experiment, we use the widely used *ϵ -greedy method* to deal with this trade-off in a simple manner (Sutton and Barto 2018). Formally, the action selection method is given by:

$$a_t = \begin{cases} a \sim U\{A\} & \text{with probability } \epsilon_t, \\ \arg \max_{a \in A} Q_t(s_t, a) & \text{with probability } 1 - \epsilon_t. \end{cases} \quad (3)$$

For exploration, an action from the action space A is chosen randomly with probability $\epsilon_t \in [0, 1]$. The random action is thereby chosen from a discrete uniform distribution over the action space A , i.e., $a \sim U\{A\}$. For exploitation, the greedy action with the highest expected reward is selected with probability $1 - \epsilon_t$ (cf. Calvano et al. 2020b; Klein 2021). This action selection method ensures that, in expectation, the agent incrementally updates the value of every state-action combination (Klein 2021).

For the experimental sessions, Q-learning agents were trained via self-play in ex-ante simulation runs in the same market environment as in the experiment. To select the most profitable agent for the experiment, the learning rate α and the discount factor δ were systematically varied with ten runs per parameter constellation and each run consisting of 50 million periods. Trained Q-learning agents exhibit reward-punishment strategies indicative of collusive behavior after converging to a steady-state (cf. Calvano et al. 2020b; Klein 2021). During the experiment, a Q-learning agent autonomously sets prices p_i in each period t and receives a reward, i.e., its profit π_i . Agents continue to learn during the experiment according to the updating rule given by (2), thereby adjusting its pricing strategy to the behavior of its competitor(s). However, the exploration parameter in (3) is set to zero, i.e., $\epsilon_t = 0 \forall t$, such that agents abstain from choosing random actions, which could irritate human competitors and introduce noisy behavior.

Experimental Procedures

The experiment was programmed in Java using the experimental software framework *Brownie* (Hariharan et al. 2017). Moreover, the implementation of the market environment by Horstmann et al. (2016), which supports immediate market interactions, is used. The study was approved by the German Association for Experimental Research. All subjects are fully informed about the procedure of a session before the start of the experiment.

Upon arriving at the laboratory, subjects are randomly assigned to a seat where they can neither see nor communicate with any other participant in the experiment. After all participants have been seated, the instructions for the experiment are handed out and then read aloud. Before the experiment, each participant must complete a computerized comprehension test that verifies whether subjects have fully understood the instructions. After all subjects have successfully completed the comprehension test, the experiment starts automatically. Before the payoff-relevant competition phase starts, subjects can make themselves familiar with the user interface and set prices in a test environment. The competition phase, which lasts exactly 30 minutes, starts when all participants have set an initial action in the test environment. In the competition phase, subjects set their actions in real-time and without interruptions. To this end, input and output variables are updated every 500ms, which is perceived as real-time by human participants. The user interface displays current prices, quantities, and profits numerically as well as historic prices, quantities, and profits of all firms graphically. After the end of the competition phase, subjects answer a short questionnaire on their pricing strategy and how they perceived the behavior of their competitor(s). At the end of the experiment, subjects receive a monetary payoff that is given by their firm's profit accumulated over the entire competition phase. Subjects are paid in private such that other participants cannot learn about their performance.

Preliminary Results

In the following, we present preliminary results for the first experimental sessions that have been conducted between December 2021 and March 2022 with students from the University of Passau. In total, 23 participants took part in three sessions that all employed price competition in a duopoly market. The sessions lasted about one hour. On average, subjects earned 15.69 EUR. In the analysis, we focus on treatment differences with respect to the degree of tacit collusion. In all treatments, the independent observation is at the market level, as competitors' actions depend on each other and are correlated over periods. Table 2 reports the average degree of tacit collusion $\bar{\phi}$ per treatment together with its standard deviation σ .

Treatment	Subjects	Observations	$\bar{\phi}$ (σ)
HH-2	14	7	0.62 (0.18)
HA-2	9	9	0.34 (0.17)
AA-2	-	1	1.00 (0.00)
AA-3	-	1	1.12 (0.00)

Table 2. Average degrees of tacit collusion across treatments.

Figure 1 visualizes average degrees of tacit collusion over the duration of the experiment across treatments with two firms. Across all treatments, prices exceed the competitive price level, as indicated by positive average degrees of tacit collusion in Table 2, which are significantly different from zero based on one sample t-tests ($p < 0.05$). In Figure 1, this is most clearly seen for later periods after about one third of the entire time horizon (i.e., after ten minutes). For treatments with human market participants, tacit collusion is significantly higher when only humans compete ($\bar{\phi}_{HH} = 0.62$) than when humans compete with algorithms ($\bar{\phi}_{HA} = 0.34$) based on a two-sample t-test ($p < 0.1$). Thus, replacing one of the human competitors with an AI-driven pricing algorithm reduces the degree of tacit collusion. However, in both treatments with human participation, the degree of tacit collusion is significantly below the fully collusive outcome that is achieved when AI-driven algorithms compete with themselves (two-sample t-test, $p < 0.1$).

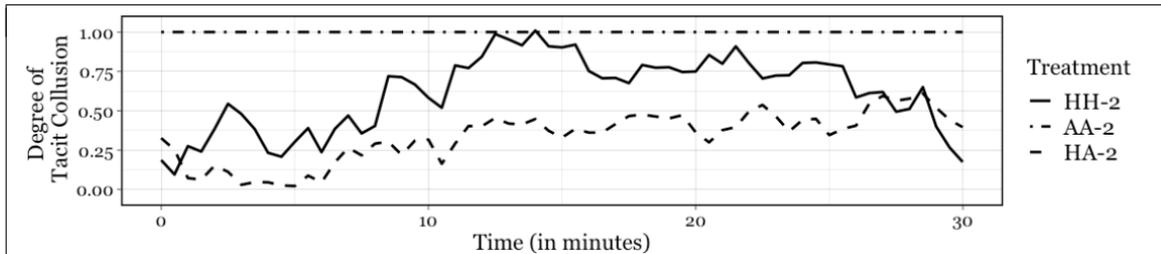


Figure 1. Average degrees of tacit collusion over time across treatments.

Based on these preliminary findings, one can conjecture that the use of AI-driven pricing algorithms facilitates tacit collusion compared to human decision makers, but only if algorithms solely compete with each other. When human decision makers are involved, tacit collusion decreases. Yet, the lowest degree of tacit collusion arises for hybrid market settings where humans and AI compete.

Discussion

The preliminary results would confirm Hypothesis 1 and Hypothesis 3, but reject Hypothesis 2. The latter finding would run against the hypothesized notion of a linear relationship between the degree of tacit collusion and the number of pricing algorithms in a market. The finding of a non-monotonic relationship would also be in contrast to results of previous experimental studies (Normann and Sternberg 2022; Werner 2022). In general, a lower degree of tacit collusion in hybrid market settings could be explained by (i) a lower intention of humans to collude with AI, (ii) a lower ability of humans to collude with AI, or (iii) both. With

respect to (i), algorithm aversion (Dietvorst et al. 2015) may lead to a lower willingness of humans to cooperate with AI and could thus explain a lower intention to collude in hybrid market settings. Responses to the questionnaire provide some anecdotal evidence for this explanatory approach. For example, some subjects stated that their ultimate goal was to outperform their algorithmic counterpart, which indicates a different goal function than in markets with only human competitors. With respect to (ii), humans and AI may follow fundamentally different strategies and behavioral rationales to achieve cooperation in competitive settings, which could possibly explain a lower ability to collude with AI. For example, some subjects stated that they perceived the behavior of the algorithm as rather noisy and non-transparent. This points to a failure of coordination, despite the intention to collude. In any case, further data collection is required before more conclusive statements can be drawn. This holds for the quantitative analysis of the degree of tacit collusion as well as the qualitative assessment of potential explanations based on subjects' intention and ability to collude.

Conclusion

This study is among the first to investigate digital market settings where artificial and human intelligence compete, which have so far received only little attention by IS and HCI research. To investigate market outcomes and tacit collusion among humans and AI, we design and conduct an oligopoly experiment in near real-time. We compare market outcomes in these hybrid markets to market settings where only humans or only algorithms compete. The preliminary findings demonstrate that our computerized market environment is well-suited to investigate the presented research questions. The findings also provide first indications that tacit coordination can emerge among humans and AI. However, more data collection (that is currently ongoing) is required before more definite conclusions can be drawn. To this end, our study also establishes a building block for future IS research on algorithmic collusion in hybrid markets that may evaluate the impact of additional factors such as market size, firm size, or product type. Such investigations will have important implications for policymakers with respect to IT and AI regulation as well as managers with respect to potential compliance and liability issues in the context of AI-driven algorithmic pricing. Besides general limitations of laboratory experiments with regard to external validity, the choice and training of the algorithm used in our experiment imposes potential limitations on the generalizability of our findings. Therefore, we plan to corroborate the robustness of our findings by also considering alternative training methods for AI-based pricing algorithms, such as online training or training on human data. A further limitation is the applicability to other market types (e.g., financial markets), which may incorporate different market mechanisms. Yet, the investigation of other market types with their corresponding dynamics is still open for future research.

Acknowledgments

The authors acknowledge generous funding by the Dr. Theo and Friedl Schöller Research Center for Business and Society as well as the Bavarian State Ministry of Science and the Arts (coordinated by bidt).

References

- Assad, S., Clark, R., Ershov, D., and Xu, L. 2020. *Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market*. Working Paper. <https://ssrn.com/abstract=3682021>.
- Calvano, E., Calzolari, G., Denicolò, V., Harrington, J. E., and Pastorello, S. 2020a. "Protecting consumers from collusive prices due to AI," *Science* (370:6520), pp. 1040–1042.
- Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. 2020b. "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review* (110:10), pp. 3267–3297.
- Caro, F., Colliard, J.-E., Katok, E., Ockenfels, A., Stier-Moses, N., Tucker, C., and Wu, D. 2022. "Management Science Special Issue on the Human-Algorithm Connection," *Management Science* (68:1), pp. 7–8.
- Chen, L., Mislove, A., and Wilson, C. 2016. "An empirical analysis of algorithmic pricing on Amazon Marketplace," in *Proceedings of the 25th International WWW Conference*, Montréal, CA, pp. 1339–1349.
- Competition and Markets Authority 2021. *Algorithms: How they can reduce competition and harm consumers*. <https://bit.ly/3OugrCd>. Accessed 28 April 2022.

- De Vaujany, F.-X., Fomin, V. V., Haefliger, S., and Lyytinen, K. 2018. “Rules, practices, and information technology: A triffecta of organizational regulation,” *Information Systems Research* (29:3), pp. 755–773.
- Dietvorst, B. J., Simmons, J. P., and Massey, C. 2015. “Algorithm aversion: People erroneously avoid algorithms after seeing them err.” *Journal of Experimental Psychology: General* (144:1), pp. 114–126.
- Engel, C. 2007. “How much collusion? A meta-analysis of oligopoly experiments,” *Journal of Competition Law & Economics* (3:4), pp. 491–549.
- European Commission 2017. *Final Report on the E-commerce Sector Inquiry*. https://ec.europa.eu/competition/antitrust/sector_inquiry_final_report_en.pdf. Accessed 28 April 2022.
- Ezrachi, A. and Stucke, M. E. 2017. “Artificial intelligence & collusion: When computers inhibit competition,” *University of Illinois Law Review* (2017:5), pp. 1775–1810.
- Fügener, A., Grahl, J., Gupta, A., and Ketter, W. 2021. “Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI,” *Management Information Systems Quarterly* (45:3), pp. 1527–1556.
- Gozman, D., Butler, T., and Lyytinen, K. 2020. “Regulation in the Age of Digitalization,” *Call for Papers on Special Issue of the Journal of Information Technology*, pp. 1–5.
- Häckner, J. 2000. “A note on price and quantity competition in differentiated oligopolies,” *Journal of Economic Theory* (93:2), pp. 233–239.
- Hariharan, A., Adam, M. T., Dorner, V., Lux, E., Mueller, M. B., Pfeiffer, J., and Weinhardt, C. 2017. “Brownie: A platform for conducting NeuroIS experiments,” *Journal of the Association for Information Systems* (18:4), pp. 264–296.
- Harrington, J. E. 2018. “Developing competition law for collusion by autonomous artificial agents,” *Journal of Competition Law & Economics* (14:3), pp. 331–363.
- Horstmann, N., Krämer, J., and Schnurr, D. 2016. *Oligopoly competition in continuous time*. Working Paper. <https://papers.ssrn.com/abstract=2630664>.
- Horstmann, N., Krämer, J., and Schnurr, D. 2018. “Number effects and tacit collusion in experimental oligopolies,” *The Journal of Industrial Economics* (66:3), pp. 650–700.
- Hotelling, H. 1929. “Stability in Competition,” *The Economic Journal* (39:153), pp. 41–57.
- Ivaldi, M., Jullien, B., Rey, P., Seabright, P., and Tirole, J. 2003. *The economics of tacit collusion*. Final Report for DG Competition, European Commission.
- Jain, H., Padmanabhan, B., Pavlou, P. A., and Santanam, R. T. 2018. “Call for papers – special issue of Information Systems Research – humans, algorithms, and augmented intelligence: The future of work, organizations, and society,” *Information Systems Research* (29:1), pp. 250–251.
- Klein, T. 2021. “Autonomous algorithmic collusion: Q-learning under sequential pricing,” *The RAND Journal of Economics* (52:3), pp. 538–558.
- Mehra, S. K. 2016. “Antitrust and the Robo-Seller: Competition in the Time of Algorithms,” *Minnesota Law Review* (100:4), pp. 1323–1375.
- Monopolkommission 2018. *Wettbewerb 2018: XXII. Hauptgutachten der Monopolkommission gemäß § 44 Abs. 1 Satz 1 GWB*. <https://bit.ly/3vIsbIS>. Accessed 29 April 2022.
- Normann, H.-T. and Sternberg, M. 2022. *Human-Algorithm Interaction: Algorithmic Pricing in Hybrid Laboratory Markets*. MPI Collective Goods Discussion Paper. <https://ssrn.com/abstract=3840789>.
- OECD 2017. *Algorithms and Collusion: Competition Policy in the Digital Age*. <https://bit.ly/3OrRzLp>. Accessed 23 April 2022.
- Potters, J. and Suetens, S. 2013. “Oligopoly experiments in the current millennium,” *Journal of Economic Surveys* (27:3), pp. 439–460.
- Schwalbe, U. 2018. “Algorithms, machine learning, and collusion,” *Journal of Competition Law & Economics* (14:4), pp. 568–607.
- Singh, N. and Vives, X. 1984. “Price and Quantity Competition in a Differentiated Duopoly,” *The RAND Journal of Economics* (15:4), pp. 546–554.
- Sutton, R. S. and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: The MIT Press.
- Waltman, L. and Kaymak, U. 2008. “Q-learning agents in a Cournot oligopoly model,” *Journal of Economic Dynamics and Control* (32:10), pp. 3275–3293.
- Werner, T. 2022. *Algorithmic and Human Collusion*. Working Paper. <https://ssrn.com/abstract=3960738>.
- Zhou, N., Zhang, L., Li, S., and Wang, Z. 2018. *Algorithmic Collusion in Cournot Duopoly Market: Evidence from Experimental Economics*. Working Paper. <https://arxiv.org/abs/1802.08061>.