Risk of AI-enabled collusion with algorithmic trading in OTC markets

Foreign Exchange Contact Group

European Central Bank

Frankfurt, March 2026

Roel Oomen

Head of the Quantitative R&D Lab for Sales and Trading

Artificial Intelligence (AI)

In an increasingly digital world, AI and related techniques offer many benefits:

- speed of information processing and decision making (e.g. self-driving cars);
- scale of deployment (e.g. personalised recommendations, adaptive educational training);
- complex pattern detection (e.g. disease detection, natural language processing);
- strategic interaction (e.g. chess / go);
- adaptability & efficiency.

But several concerns have been raised:

- unpredictable or unintended behaviours arising from interactions amongst AIs (e.g. Hendrycks, Mazeika, and Woodside, 2023);
- tacit collusion amongst independently operated AIs (e.g. OECD, 2017; European Commission, 2017; Competition & Markets Authority, 2018, 2021, 2026; Löfström, Ralsmark, and Johansson, 2021);
- discriminatory biases, predatory price differentiation (e.g. The White House, 2015).

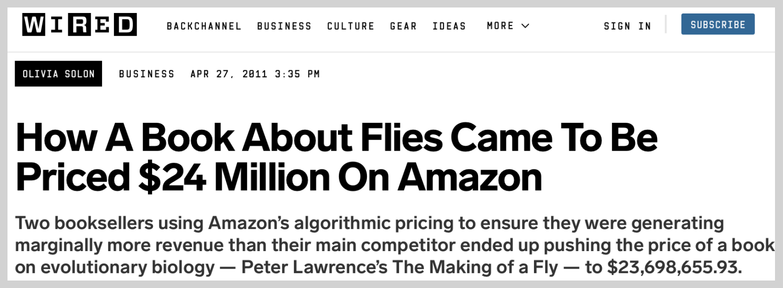

WIRED BACKCHANNEL BUSINESS CULTURE GEAR IDEAS MORE SIGN IN SUBSCRIBE

OLIVIA SOLON BUSINESS APR 27, 2011 3:35 PM

How A Book About Flies Came To Be Priced $24 Million On Amazon

Two booksellers using Amazon’s algorithmic pricing to ensure they were generating marginally more revenue than their main competitor ended up pushing the price of a book on evolutionary biology — Peter Lawrence’s The Making of a Fly — to $23,698,655.93.

Source: https://www.wired.com/2011/04/amazon-flies-24-million (see also https://www.michaeleisen.org/blog/?p=358)

So what happened? Only two sellers of the book, who update prices every day as follows:

- seller A, sets price at 0.99830 times best available price (undercut competitor on price);
- seller B, sets price at 1.270589 times best available price (charge premium based on better ratings).

➜ With the book price starting at $100, only 50 updates were needed for it to reach into the millions.

Illustration of risks | supra-competitive pricing & tacit collusion

Assad, Clark, Ershov, and Xu (2024) study the adoption of AI-powered algorithmic pricing software by German petrol stations and find:

- when all stations within an area adopt the software, average prices increase by 6 cents/litre, and margins increase by 3.2 cents/litre, or roughly 38%;
- "algorithms learn that undercutting will not be profitable, since lower prices will be followed".

Musolff (2022, 2026) study algorithmic pricing on Amazon Marketplace and find:

- use of "resetting" strategies that regularly raise prices in the hope that competitors will follow;
- price resets often done at night when sales volumes are low;
- adoption of such strategies led to an average 11% price increase.

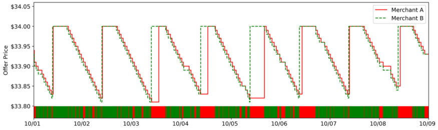

Source : Musolff (2026), Figure 2.

Competition vs cooperation | and the Prisoner's dilemma

Central question : *if you place multiple self-interested and self-learning agents in competition, will they learn to aggressively compete, or will they find a way to cooperate for mutual gain?*

Cooperation is often desirable, e.g.

- self-driving cars;
- nuclear disarmament;
- environmental pollution;
- international trade / tariffs.

But sometimes it is not, e.g.

- business;
- competitive sports.

(see, e.g., Axelrod, 1984)

LP1 payoff

		LP2 action	
		Quote tight (defect)	Quote wide (cooperate)
LP1 action	Quote tight (defect)	2 equal share at tight spread	4 undercut LP2 and win all flow
	Quote wide (cooperate)	0 undercut by LP2 and lose all flow	3 equal share at wide spread

"Prisoner's dilemma" pay-off profile leads to a Nash equilibrium where LPs act competitively quoting tight (not cooperatively quoting wide).

AI & Algorithmic Pricing | beating the Prisoner's dilemma?

- Some competitive scenarios may be "one-shot games" (e.g. mobile network spectrum auction) but in e-commerce it is typically a repeated game, many times over (e.g. airline tickets, hotel bookings, FX trading).

- In a repeated game, where competitors set prices using AI algorithms – potentially utilising a rich set of historical competitor pricing data and deal flow – can the Prisoner's dilemma be beaten and supra-competitive prices be charged?

- Algorithmic collusion requires
  - ✔ independent algorithms converge to set supra-competitive prices (higher than Nash);
  - ✔ the algorithms are not programmed to collude explicitly, i.e. it arises tacitly;
  - ✔ a reward-punishment retaliation mechanism[1] to sustain the equilibrium.

[1] This requires the algorithm to have memory. Without it, tacit collusion cannot arise but supra-competitive pricing may.

AI & Algorithmic Pricing | selected literature

There is an extensive recent literature showing that independent competitors can end up charging supra-competitive prices enabled by their use of AI-driven pricing algorithms.

- Tacit collusion – sustained by a learned price-trigger punishment – arising with reinforcement Q-learning algorithms (e.g. Waltman and Kaymak, 2008; Klein, 2021; Calvano, Calzolari, Denicolò, and Pastorello, 2020, 2021, 2023).
- Supra-competitive pricing – without reward-punishment mechanism – arising with reinforcement Q-learning algorithms due to
  - simultaneous learning and/or imperfect exploration (Abada and Lambin, 2023; Lambin, 2024);
  - asynchronous learning (Asker, Fershtman, and Pakes, 2024);
  - correlated experiments leading to misspecification of price sensitivity (Hansen, Misra, and Pai, 2021);
  - differences in price update frequency (Brown and MacKay, 2023).
- Tacit collusion arising with LLM-based pricing agents (Fish, Gonczarowski, and Shorrer, 2026).

AI & Algorithmic Pricing | selected literature (cont'd)

- Discussion from a competition law perspective
  (e.g. Mehra, 2016; Ezrachi and Stucke, 2017; Harrington, 2018; Schwalbe, 2019; Nazzini and Henderson, 2024; Doerr, 2025).

- Literature reviews and policy discussion
  (e.g. Abada, Harrington, Lambin, and Meylahn, 2025; Assad, Calvano, Calzolari, Clark, Denicolò, Ershov, Johnson, Pastorello, Rhodes, Xu, and Wildenbeest, 2021; Dorner, 2021; Deng, 2024; Bichler, Durmann, and Oberlechner, 2025).

- In mitigation of algorithmic collusion, market design can be adjusted, e.g. via dynamic price-directed prominence where high priced sellers are shown to fewer customers
  (as studied in Johnson, Rhodes, and Wildenbeest, 2023).

- Folk theorems can also be proven
  (e.g. Álvaro Cartea, Chang, Penalva, and Waldon, 2026; Askenazi-Golan, Mergoni Cecchelli, Plumb, and Possnig, 2026).

AI & Algorithmic Trading | in OTC financial markets

Much of the above literature	vs	OTC financial markets
Similar but differentiated products (e.g. a bike).	✘	Homogenous products (e.g. EURUSD).
Competitor prices are observed by all.	✘	Competitor prices are not observed.
Best price may not win (substitution & quality).	✘	Best price (typically) wins.
Production costs (and profit margin) are known.	✘	Fair price unknown (asym info & adverse selection).

- Cartea, Chang, Mroczka, and Oomen (2022) is – to the best of our knowledge – the first to study the risk of algorithmic collusion in a realistic "request for stream" OTC market model. The key findings of this paper are summarised in the remainder of this presentation.
- Related work includes:
  - Xiong and Cont (2021); Cont and Xiong (2024) study competing dealers in an RFQ setting;
  - Cartea, Chang, and Penalva (2025) study competing dealers in a central limit order book;
  - Dou, Goldstein, and Ji (2025) study competing informed traders in a Kyle-type model;
  - Colliard, Foucault, and Lovo (2026) study competing dealers in a Glosten-Milgrom model.

AI & Algorithmic Trading | the market model

We adopt the RFS "request for stream" OTC market model of Oomen (2017).

- ✔ N liquidity providers (LPs) compete for a trader's (random) order-flow.
- ✔ LPs stream liquidity at a spread around their estimate of the unobserved true mid-price.
- ✔ Trader executes with dealer showing the best price subject to a reservation price.

Within this model, we can derive the *competitive* equilibrium spread s^* and the *monopolistic* equilibrium spread $s^\dagger$.

This lets us evaluate the "efficiency" or competitiveness of any equilibrium s reached by the AI-algorithm as

$$\mathcal{E}_i(s) = \frac{\mathbb{V}_t(s^\dagger) - \mathbb{V}_t(s)}{\mathbb{V}_t(s^\dagger) - \mathbb{V}_t(s^*)},$$

where $\mathcal{E} = 100\%$ (0%) corresponds to competitive (monopolistic) equilibrium.

We consider class of "multi-armed bandit" (MAB) reinforcement learning algorithms.[2]

(i) discretise state space of candidate spreads, e.g. $s \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$.

(ii) pick a spread s (draw an "arm")
- with probability ϵ_t select a candidate spread at random to "explore"
- with probability $1 - \epsilon_t$ select spread with highest estimated award to "exploit"

(iii) receive the realised reward π and update estimate of $E(\pi|s)$

(iv) loop back to step (ii) for next price, potentially decaying exploration rate ϵ_t gradually to exploit more as estimates become more accurate.

Algorithm 1: ϵ-greedy

Parameters: $\epsilon \in [0,1]$ and a set of candidate spreads $\{s(1), \ldots, s(K)\}$.
Initialise: the running average reward $\overline{\pi}_1(s(k)) = 0$ and total number of draws $n_1(k) = 0$ for each arm $k = 1, \ldots, K$.
for $t = 1, 2, \ldots$ **do**
  Sample $u \sim \mathcal{U}(0,1)$.
  if $u > \epsilon$ **then**
    Pick the arm with the highest average reward, breaking ties randomly, i.e. $k_t^* = \arg\max_k \overline{\pi}_t(s(k))$.
  else
    Randomly pick an arm k_t^* with equal probability $1/K$.
  end
  Update: propagate previous states for all arms, and increment those for the selected arm as $\overline{\pi}_{t+1} = \frac{\overline{\pi}_t n_t + \pi_t}{n_t + 1}$, $n_{t+1} \leftarrow +1$ where π_t is the realised reward.
end

Source : Cartea, Chang, Mroczka, and Oomen (2022)

[2] Q-learning is unnecessary because in an RFS market there are no state variables (like past competitor prices) to condition on.

AI & Algorithmic Trading | model-free but extremely slow/inefficient

First, consider only one LP is optimising (with other LPs quoting the optimal spread).

dealer objective function

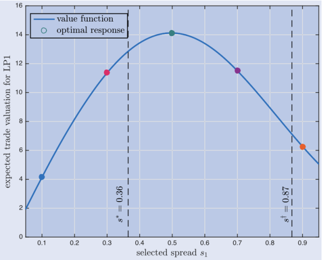

MAB convergence to optimal spread

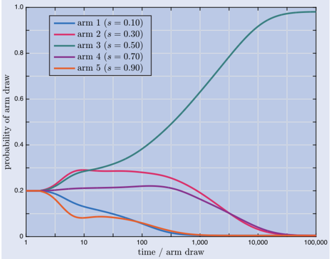

convergence of MLE

	Panel A : frequency of arm draws				
Sample size	$s = 0.1$	$s = 0.3$	$s = 0.5$	$s = 0.7$	$s = 0.9$
Scenario 1 : correct model specification					
$T = 10$	0.2%	14.6%	63.0%	16.9%	5.3%
$T = 100$	0.0%	0.0%	91.3%	8.7%	0.0%
$T = 1000$	0.0%	0.0%	99.5%	0.5%	0.0%
$T = 10,000$	0.0%	0.0%	100.0%	0.0%	0.0%
$T = 100,000$	0.0%	0.0%	100.0%	0.0%	0.0%

Source : Figure 4B, 5A, Table 1A Cartea, Chang, Mroczka, and Oomen (2022)

✔ AI correctly converges to the optimal spread of $s = 0.5$

✘ Model-free AI converges order(s) of magnitude slower than model-based estimator.

AI & Algorithmic Trading | model-free but extremely slow/inefficient

- In Cartea, Chang, Mroczka, and Oomen (2022), simulations indicate convergence required 50,000 trades or more. This implies an average client requires 20 years of optimisation by the dealer, assuming the competitive environment remains unchanged!

- Dou, Goldstein, and Ji (2025) find for more complex Q-learning optimisation "*convergence occurs within a range of approximately 20 million to 50 billion periods*"

- den Boer, Meylahn, and Schinkel (2024) find "*if the Q-learning algorithms converge to collusive equilibria, they do so intrinsically slowly, without hope of possibilities to speed them up by hyperparameter tuning*".[3]

[3] Bhole and Surana (2025) study alternative RL methods and find faster convergence compared to Calvano, Calzolari, Denicolò, and Pastorello (2020).

pricing efficiency by price granularity & number of competitors N

		Number of competing LPs					
		$N=2$	$N=3$	$N=4$	$N=5$	$N=6$	$N\geq 7$
Number of arms	$K=2$	0%	0%	0%	0%	0%	100%
	$K=3$	48%	48%	48%	48%	48%	100%
	$K=5$	74%	74%	74%	74%	74%	100%
	$K=10$	88%	88%	88%	100%	100%	100%
	$K=15$	93%	93%	100%	100%	100%	100%
	$K=20$	94%	100%	100%	100%	100%	100%
	$K=30$	96%	100%	100%	100%	100%	100%
	$K\geq 40$	100%	100%	100%	100%	100%	100%

Source : Table 4 Cartea, Chang, Mroczka, and Oomen (2022)

- Supra-competitive spreads can merely be an artefact of restrictive AI hyper-parameters.
  - As price/spread granularity is increased, pricing inefficiency vanishes.
  - Increasing the number of competitors is also effective at mitigating any such effects.
  - Also consistent with findings of Cartea, Chang, and Penalva (2025).

AI & Algorithmic Trading | pseudo collusion *can* arise but fragile

Example of UCB convergence to pseudo collusive equilibrium ($N = 2$)

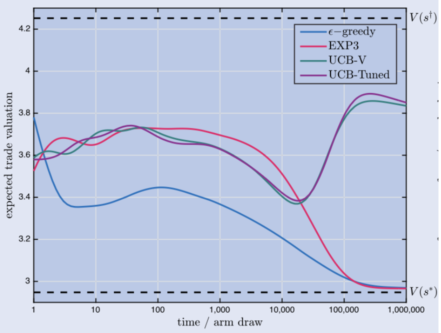

Source : Figure 11A Cartea, Chang, Mroczka, and Oomen (2022)

UCB pricing efficiency by number of competitors N

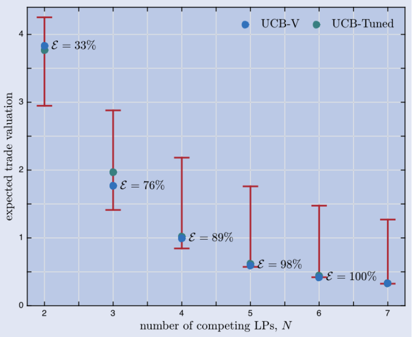

Source : Figure 14A Cartea, Chang, Mroczka, and Oomen (2022)

- Pseudo collusion *can* arise for certain (contrived) configurations and algorithms.
- Vanishes with increase in N and/or competitor using a different AI algorithm.

AI & Algorithmic Trading | asynchronous RL disrupts everything

LP1 spread selection

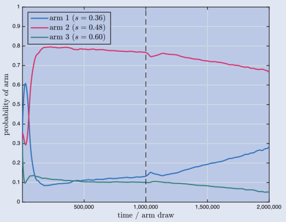

LP2 spread selection

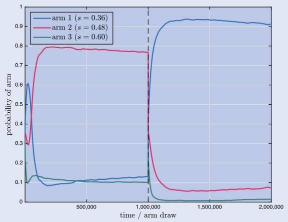

LP1 profitability

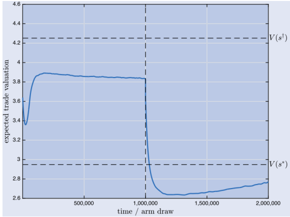

LP2 profitability

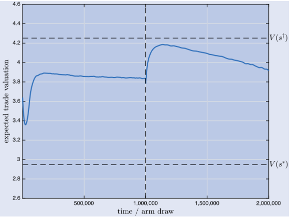

Source : Figure 13 Cartea, Chang, Mroczka, and Oomen (2022)

Scenario : LP 1 & 2 optimise spreads for 1m steps. Then LP2 resets their AI algorithm!

- ✔ With new cycle of exploration, LP2 quickly learns to undercut LP1, almost reaching monopolistic profit levels.
- ✘ LP1 is mostly exploiting (not exploring) in steady state and is slow to recognise – to its detriment – that the environment has changed.
- 💥 Any equilibrium (incl collusive ones) is disrupted with merely one LP reconfiguring their RL algorithm.

AI & Algorithmic Trading | conclusion

In Cartea, Chang, Mroczka, and Oomen (2022) we show that collusive effects can arise in theory.

In practice, however, it may require an "alignment of the stars" scenario for it to materialise.[4]

1. Unobservability of competitor prices in OTC markets is primary mitigant against collusion.
2. Even a modest number of competing LPs effectively mitigates any potential collusive effects.
3. Heterogeneity in approach (e.g. algorithm, training cycle, hyper-parameters) quickly eliminates any potential collusive effects. The common assumption of a homogeneous & synchronised training approach is exceedingly unlikely to occur with independent competitors.
4. Very slow convergence to any potential collusive state, i.e. with average trader executing 5-10 trades/day in the currency markets, it would take the MAB algorithm 20+ years to converge.
5. The trader has a information advantage over the LPs: any collusive effects are far easier for the trader to monitor and disrupt (e.g. rotate LPs) than it is for the LPs to effectuate.

[4] A number of our findings have recently been supported by subsequent studied, e.g den Boer, Meylahn, and Schinkel (2024); Douglas, Provost, and Sundararajan (2026); Keppo, Li, Tsoukalas, and Yuan (2026).

Suggested discussion points

1. Transparency. Price transparency is seen by many as an important and desirable feature of market design, but here it a the key enabler for tacit collusion.[5]

2. Measurability. Prior to AI, was the market pricing at the competitive Nash equilibrium? How to even establish any net-benefit/harm of AI, weighing any risk of collusion against granular and high frequency price competition?

3. Controls. Any potential algorithmic collusion appears easy to break by disrupting the fragile environment that sustains it, e.g. periodically rotate in/out LPs, vary execution style/channels, monitoring LP pricing patterns.

4. Counter-arguments. Competitors unlikely to be homogeneous & synchronised. However, they may do pre-training on common / correlated datasets? Use alternative faster algorithms, or access to more data?

[5] "*Courts and the enforcement agencies may be reluctant to restrict this free flow of information in the marketplace. Its dissemination, observed the Supreme Court, "is normally an aid to commerce" and "can in certain circumstances increase economic efficiency and render markets more, rather than less, competitive." Indeed, concerted action to reduce price transparency may itself be an antitrust violation*." Ezrachi and Stucke (2017, p 1797).

References

Abada, I., J. E. Harrington, X. Lambin, and J. M. Meylahn, 2025, "Algorithmic Collusion: Where Are We and Where Should We Be Going?," available at https://ssrn.com/abstract=4891033.

Abada, I., and X. Lambin, 2023, "Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided?," *Management Science*, 69(9), 5042–5065.

Álvaro Cartea, P. Chang, J. Penalva, and H. Waldon, 2026, "Algorithmic collusion and a folk theorem from learning with bounded rationality," *Games and Economic Behavior*, 157, 1–21.

Askenazi-Golan, G., D. Mergoni Cecchelli, E. Plumb, and C. Possnig, 2026, "The Bounds of Algorithmic Collusion: Q-learning, Gradient Learning, and the Folk Theorem," available at https://arxiv.org/abs/2411.12725v2.

Asker, J., C. Fershtman, and A. Pakes, 2024, "The impact of artificial intelligence design on pricing," *Journal of Economics & Management Strategy*, 33(2), 276–304.

Assad, S., E. Calvano, G. Calzolari, R. Clark, V. Denicolò, D. Ershov, J. Johnson, S. Pastorello, A. Rhodes, L. Xu, and M. Wildenbeest, 2021, "Autonomous algorithmic collusion: economic research and policy implications," *Oxford Review of Economic Policy*, 37(3), 459–478.

Assad, S., R. Clark, D. Ershov, and L. Xu, 2024, "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy*, 132(3), 723 – 771.

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

Bhole, B., and S. Surana, 2025, "Tacit Collusion by Pricing Algorithms," *Economic Inquiry*, 63(4), 1036 – 1065.

Bichler, M., J. Durmann, and M. Oberlechner, 2025, "Algorithmic Pricing and Algorithmic Collusion," *Business & Information Systems Engineering*, 67(6), 971–979.

Brown, Z. Y., and A. MacKay, 2023, "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics*, 15(2), 109–156.

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

——— , 2021, "Algorithmic Collusion with Imperfect Monitoring," *International Journal of Industrial Organization*, 79, 102712.

——— , 2023, "Algorithmic Collusion: Genuine or Spurious?," *International Journal of Industrial Organization*, 90, 102973.

Cartea, A., P. Chang, M. Mroczka, and R. Oomen, 2022, "AI driven liquidity provision in OTC financial markets," *Quantitative Finance*, 22(12), 2171–2204.

Cartea, A., P. Chang, and J. Penalva, 2025, "Algorithms and Supracompetitive Prices in Electronic Markets: The Impact of Tick Size," available at https://ssrn.com/abstract=4105954.

Colliard, J.-E., T. Foucault, and S. Lovo, 2026, "Algorithmic Pricing and Liquidity in Securities Markets," *The Review of Financial Studies*, forthcoming.

References

Competition & Markets Authority, 2018, “Pricing Algorithms,” available at https://www.gov.uk/government/publications/pricing-algorithms-research-collusion-and-personalised-pricing.

——— , 2021, “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.

——— , 2026, “AI and collusion: frontiers, opportunities and challenges,” available at https://competitionandmarkets.blog.gov.uk/2026/03/04/ai-and-collusion-frontiers-opportunities-and-challenges/.

Cont, R., and W. Xiong, 2024, “Dynamics of market making algorithms in dealer markets: Learning and tacit collusion,” *Mathematical Finance*, 34(2), 467–521.

den Boer, A. V., J. M. Meylahn, and M. P. Schinkel, 2024, “Artificial Collusion: Examining Supracompetitive Pricing by Q-Learning Algorithms,” Amsterdam Center for Law & Economics Working Paper No. 2022-06, available at https://ssrn.com/abstract=4213600.

Deng, A., 2024, “What Do We Know About Algorithmic Collusion Now? New Insights from the Latest Academic Research,” available at https://ssrn.com/abstract=4521959.

Doerr, A., 2025, *Algorithmic Tacit Collusion*. Mohr Siebeck, Tübingen, 1 edn.

Dorner, F. E., 2021, “Algorithmic collusion: A critical review,” available at https://arxiv.org/abs/2110.04740.

Dou, W. W., I. Goldstein, and Y. Ji, 2025, “AI-Powered Trading, Algorithmic Collusion, and Price Efficiency,” NBER Working Paper, available at https://www.nber.org/papers/w34054.

Douglas, C., F. Provost, and A. Sundararajan, 2026, “The Illusion of Collusion,” available at https://arxiv.org/html/2411.16574v2.

European Commission, 2017, “Algorithms and Competition,” Speech by Commissioner Margrethe Vestager at Bundeskartellamt 18th Conference on Competition, Berlin.

Ezrachi, A., and M. E. Stucke, 2017, “Artificial Intelligence and Collusion: When Computers Inhibit Competition,” *University of Illinois Law Review*, 5, 1775–1810.

Fish, S., Y. A. Gonczarowski, and R. I. Shorrer, 2026, “Algorithmic Collusion by Large Language Models,” available at https://arxiv.org/abs/2404.00806.

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

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

Hendrycks, D., M. Mazeika, and T. Woodside, 2023, “An Overview of Catastrophic AI Risks,” Center for AI Safety, available at https://arxiv.org/pdf/2306.12001.

Johnson, J. P., A. Rhodes, and M. R. Wildenbeest, 2023, “Platform Design When Sellers Use Pricing Algorithms,” *Econometrica*, 91(5), 1841–1879.

References

Keppo, J., Y. Li, G. Tsoukalas, and N. Yuan, 2026, "On the Fragility of AI Agent Collusion," available at https://ssrn.com/abstract=5386338.

Klein, T., 2021, "Autonomous Algorithmic Collusion: Q-Learning under Sequential Pricing," *RAND Journal of Economics*, 52(3), 538–558.

Lambin, X., 2024, "Less Than Meets the Eye: Simultaneous Experiments as a Source of Algorithmic Seeming Collusion," available at https://ssrn.com/abstract=4498926.

Löfström, T., H. Ralsmark, and U. Johansson, 2021, "Collusion in Algorithmic Pricing," Swedish Competition Authority, Uppdragsforskningsrapport 2021:3, available at https://www.konkurrensverket.se/globalassets/dokument/informationsmaterial/rapporter-och-broschyrer/uppdragsforskning/forsk-rapport_2021-3.pdf.

Mehra, S. K., 2016, "Antitrust and the Robo-Seller: Competition in the Time of Algorithms," *Minnesota Law Review*, 100, 1323–1375.

Musolff, L., 2022, "Algorithmic Pricing Facilitates Tacit Collusion: Evidence from E-Commerce," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, p. 32–33, New York, NY, USA. Association for Computing Machinery.

——— , 2026, "Algorithmic Pricing, Price Wars And Tacit Collusion: Evidence from E-Commerce," *Management Science*, forthcoming.

Nazzini, R., and J. Henderson, 2024, "Overcoming the Current Knowledge Gap of Algorithmic "Collusion" and the Role of Computational Antitrust," *Stanford Computational Antitrust*, 4, 1–32.

OECD, 2017, "Algorithms and Collusion: Competition Policy in the Digital Age," available at https://www.oecd.org/content/dam/oecd/en/publications/reports/2017/05/algorithms-and-collusion-competition-policy-in-the-digital-age_02371a73/258dcb14-en.pdf.

Oomen, R., 2017, "Execution in an aggregator," *Quantitative Finance*, 17(3), 383–404.

Schwalbe, U., 2019, "Algorithms, Machine Learning, and Collusion," *Journal of Competition Law & Economics*, 14(4), 568–607.

The White House, 2015, "Big data and differential pricing," available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf.

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

Xiong, W., and R. Cont, 2021, "Interactions of Market Making Algorithms: a Study on Perceived Collusion," conference proceedings ICAIF21, USA, available at https://dl.acm.org/doi/pdf/10.1145/3490354.3494397.

Disclaimer

This material was prepared in part by an employee in the Group Strategic Analytics function within Deutsche Bank AG and/or its affiliates ("DB"), and was not produced, reviewed or edited by the Research Department of DB (which is independent from the Sales or Trading function). Any opinions expressed herein may differ from the opinions expressed by other DB departments including the Research Department. Different functions within DB may be subject to additional potential conflicts of interest which the Research Department does not face.

This document is intended for discussion purposes only and does not create any legally binding obligations on the part of DB. The information contained in this document is based on material the authors believe to be reliable; however, neither the authors nor DB represent that it is accurate, current, complete, or error free. Assumptions, estimates and opinions contained in this document constitute a judgment as of the date of the document and are subject to change without notice.

DB may engage in transactions in a manner inconsistent with the views discussed herein. Employees of DB may be compensated in part based on the volume of transactions effected by them, the performance of DB in general or their applicable function.

DB SPECIFICALLY DISCLAIMS ALL LIABILITY FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL OR OTHER LOSSES OR DAMAGES INCLUDING, WITHOUT LIMITATION, LOSS OF PROFITS INCURRED BY YOU OR ANY THIRD PARTY THAT MAY ARISE FROM, OR IN CONNECTION WITH, ANY RELIANCE ON THIS DOCUMENT OR FOR THE RELIABILITY, ACCURACY, COMPLETENESS OR TIMELINESS THEREOF.