

# ALGORITHMIC COLLUSION: SUPRA-COMPETITIVE PRICES VIA INDEPENDENT ALGORITHMS ONLINE APPENDIX

KARSTEN T. HANSEN<sup>A</sup>, KANISHKA MISRA<sup>B</sup>, AND MALLESH M PAI<sup>C</sup>

## APPENDIX A. ONLINE APPENDIX

### A.1. *Some almost correct intuition*

For a first intuition about why this phenomenon obtains recall that in our simulations the true data generation process is that demand as given by a linear demand model:

$$d_{j,t} = \alpha - \beta p_{j,t} + \gamma p_{-j,t} + \varepsilon_{j,t}.$$

Demand is observed with an additional small sample econometric error ( $\varepsilon_{j,t}$ ). The key complication in our setup is that the firm does not observe the competitive firm's price ( $p_{-j,t}$ ). Therefore the firm cannot distinguish the econometric error ( $\varepsilon_{j,t}$ ) from the competing firm's price ( $p_{-j,t}$ ).

In a departure from the bandit procedure we considered before, suppose the firm estimates a linear demand model

$$d_{j,t} = \hat{\alpha} - \hat{\beta} p_{j,t} + \varepsilon_{j,t},$$

and then prices optimally given the estimated model. Notice that the firm's statistical model is different from the true data generating process. This can result in a bias in the estimate of the main parameter of interest, i.e.  $\beta$ . The magnitude of this bias depends on the historical correlation between the prices of the two firms ( $p_{j,t}$  and  $p_{-j,t}$ ). If these are highly positively correlated then bias will be an upward bias, leading to higher (lower in absolute value) estimated price sensitivity and supra-competitive prices.

To see this consider two thought experiments:

- (1) Past prices are perfectly correlated: As  $t$  grows large, since past prices are perfectly correlated, by a law of large numbers, we have that  $\hat{\alpha} = \alpha$  and  $\hat{\beta} = \beta - \gamma$ . The

---

<sup>A</sup> RADY SCHOOL OF MANAGEMENT, UNIVERSITY OF CALIFORNIA, SAN DIEGO. [K4HANSEN@UCSD.EDU](mailto:k4HANSEN@UCSD.EDU)

<sup>B</sup> RADY SCHOOL OF MANAGEMENT, UNIVERSITY OF CALIFORNIA, SAN DIEGO. [KAMISRA@UCSD.EDU](mailto:KAMISRA@UCSD.EDU)

<sup>C</sup> DEPARTMENT OF ECONOMICS, RICE UNIVERSITY. [MALLESH.PAI@RICE.EDU](mailto:MALLESH.PAI@RICE.EDU)

Date: August 17, 2020.

optimal price of firm  $j$  given these estimates is therefore

$$p_j^* = \frac{\hat{\alpha}_j}{2\hat{\beta}_j} = \frac{\alpha}{2(\beta - \gamma)}.$$

- (2) Past prices are uncorrelated: As  $t$  grows large, by a law of large numbers, we have  $\hat{\alpha} = \alpha + \gamma\bar{p}_{-j,t}$  and  $\hat{\beta} = \beta$ . The optimal price of firm  $j$  given these estimates is therefore

$$p_j^* = \frac{\hat{\alpha}_j}{2\hat{\beta}_j} = \frac{\alpha + \gamma\bar{p}_{-j,t}}{2\beta}.$$

In steady-state, therefore,  $\bar{p}_{-j,t} = p_{-j}^* = \frac{\alpha + \gamma p_j^*}{2\beta}$ . Substituting, we have

$$p_j^* = \frac{\alpha + \gamma \frac{\alpha + \gamma p_j^*}{2\beta}}{2\beta},$$

$$\implies p_j^* = \frac{\alpha}{2\beta - \gamma}.$$

Note that in the second case, despite each firm's model being misspecified, the resulting long-run steady prices are exactly the full-information Nash equilibrium prices. In the first case, both firms' prices are exactly the full-information collusive/ joint-monopoly prices.

While this intuition is suggestive of the mechanism, it is not fully consistent: the "long run" outcome here depends on the correlation of prices in the past. However, these correlations are endogenously determined by the prices charged in the past by both firms' algorithms. So which of these two cases is more representative of the outcome given the exogenous environment and the algorithms used by the firms?

## APPENDIX B. CONNECTIONS TO THE LITERATURE

Beyond the related literature mentioned in the body, our paper touches on several strands of literature. Given the space limitations in the main text, we discuss these connections here. We also explain the relationship to some more closely connected papers in further detail in this section.

### B.1. Existing Literature on Algorithmic Collusion

As we pointed out in the introduction, there is a small but growing literature pointing to the possibility of collusion in settings where firms set prices by algorithm. Here we list the work that we are aware of, and highlight key differences. As we mentioned earlier, a fuller review of existing work and connections to the legal literature on antitrust can be found in [Harrington \(2018\)](#).

A seminal recent paper by [Calvano et al. \(2020\)](#) studies a setting where competing sellers set prices by algorithm based on a history of past prices. All sellers use a machine learning algorithm known as “Q-learning”. They show that the competing sellers’ algorithms all “learn” to collude, i.e. the Q-learning algorithms learn to play collusive repeated-game strategies. Similar to our paper, therefore, even neutrally programmed algorithms may nevertheless result in supra-competitive prices. A key point of difference is that firms in their paper observe competitors’ prices. Our work shows that such an effect may carry over even if the firms do not observe this. That is to say that while the conclusions are similar, the channels through which collusive prices are established are completely different.

Another recent paper is that of [Miklós-Thal and Tucker \(2019\)](#). They study a setting where competitors have incomplete information about the market fundamentals. They look to understand how information about the market fundamentals affects the possibility of collusion. Their main tension is whether this assists collusion (competitors know better the collusive prices to aim for), or hinders it (superior information means firms can better evaluate deviations from collusion, and hence may be more tempted to do so). They show conditions under which the latter effect dominates. Subsequent work by [O’Connor and Wilson \(2019\)](#) studies a similar tension and points out that superior information also increases the value of collusion, so the net effects may be ambiguous. Our model is quite different from these papers as we consider explicit process of experimentation to learn information.

[Brown and MacKay \(2020\)](#) consider the implications of timing of pricing updating on equilibrium prices. In their model, firms know the true demand model, but unlike the canonical pricing model, they allow firms to update prices asynchronously. They find that these asynchronously responses can lead to non-Nash prices in equilibrium. A related work is the paper of [Salcedo \(2015\)](#). He considers a model where firms “commit” to an algorithm in the short-run—this partial commitment can again allow for supra-competitive prices to emerge. These models capture a different mechanism to our paper as we focus on a setting where the firms do not know the underlying demand curve (and do not observe competitive prices). Studying demand learning with asynchronously pricing algorithms in competitive settings is an interesting question for future research.

## B.2. *Competition with miss-specification*

There is a small literature that like us, considers the case of competing sellers who each erroneously behave as if they are monopolists. The earliest examples of this are from the theoretical industrial organization/ economics literature, starting with [Kirman \(1975\)](#).

As we mentioned previously, the paper closest to us that we are aware of is the paper of [Cooper et al. \(2015\)](#). To understand the difference between the two papers more closely,

here is a brief summary of the paper. In their model, two firms compete by setting prices. They assume that the true demand realized by each firm is a linear function of both prices (intuitively, increasing in competitor’s price and decreasing in own price). However, each firm erroneously thinks they are a monopoly, in every period runs a linear regression of past observed demand on past own prices. It then sets optimal monopoly prices based on the regression parameters. They prove three main sets of results. First, if each firm correctly knows the true slope of own price on own demand, and is only regressing to learn the intercept, then prices converge to the static Nash Equilibrium (Proposition 1). Conversely, if firms know the intercept but do not know the slope, then instead they converge on the joint monopoly prices (Theorem 1). Finally, if both are slope and intercept are unknown, then the result depends on the starting prices and they show an “anything goes” theorem (Proposition 3), i.e. any limit point (both above and below the competitive levels) is achievable from some well chosen starting point.

Proposition 3 and Theorem 1 is similar to our main takeaway (prices can be supra-competitive). It is important then to highlight the differences. First, observe that our paper and theirs lie in different frameworks. Theirs is in what may be referred to as a “estimation-optimization” framework, where there are two modules that operate independently, an estimation module that estimates the underlying environment, and an optimization module that optimally sets a price given the estimates. By contrast, like e.g. [Calvano et al. \(2020\)](#), and modern machine learning methods more broadly, we consider a reinforcement learning setting where both are done simultaneously. In doing so we model the price experimentation process explicitly. In terms of takeaways, we highlight that the contrast between Proposition 1, Theorem 1 and Proposition 3 suggests that the possibility of collusive seeming prices depends on whether the intercept and slope are commonly known (this cannot be empirically testable). By contrast, our results suggest that it hinges on the underlying signal-to-noise ratio in the demand, i.e. the stochasticity of underlying demand (this is empirically testable).

### B.3. *Dynamic Pricing*

There is a large and influential literature investigating dynamic pricing, and relatedly, personalized pricing by firms. As we pointed out previously, a defining characteristic of this literature is that it imagines the firm is a monopoly, or equivalently, ignores competitive response by competitors. Within economics the focus has been on characterizing and understanding the comparative statics of fully optimal policies in stylized models—see e.g. [Keller and Rady \(2003\)](#) for an example. The literature has considered both the “estimation-optimization” paradigm (e.g. described above) and reinforcement learning approaches (see e.g. [Calvano et al. \(2020\)](#) or [Misra et al. \(2019\)](#) for examples of this).

Keskin and Zeevi (2018) show that the estimation optimization paradigm can lead to “incomplete learning” – a certainty equivalent control may get stuck at a fixed point without having learned the true underlying parameter. Forced experimentation (including reinforcement learning) can help resolve this. Interestingly, while UCB does not suffer from the problem of incomplete learning if the firm was truly a monopoly, our results may be interpreted as saying that incomplete learning is a possibility when such algorithms are used in competition.

A paper at the intersection of these two approaches is Keskin and Zeevi (2014), who show that well designed small variations on the “estimation-optimization” paradigm can result in asymptotically policies. In particular, instead of the optimization module doing what is greedily/ myopically optimal (i.e. optimal viewing the estimated parameters as the true parameters, also known as certainty equivalent control), they design simple policies they term as “semi-myopic”. They show that such policies can be close to asymptotically optimal.

At a higher level, a reason to consider dynamic pricing is that the underlying demand parameters are both unstable/ time varying, and unknown (e.g., demand side dynamics Bergemann and Valimaki (1996); Nair (2007)). Of course in our paper and a majority of the works we list above, only the latter is modeled (indeed, we explicitly suppose the demand is stable and time invariant). There is a literature which considers how a seller can price in this setting:— see e.g. Rustichini and Wolinsky (1995) or Keller and Rady (1999) in the economics literature or Keskin and Zeevi (2017) in operations.

#### B.4. *Learning in Games*

There is a large literature on learning in games, spanning the literature in economic theory, computer science and statistics. We refer the reader to textbook expositions in Cesa-Bianchi and Lugosi (2006) (computer science/ statistics) and Fudenberg and Levine (1998) (economics). The underlying literature begins with the celebrated results on regret minimizing heuristics— seminal papers Foster and Vohra (1997) and Hart and Mas-Colell (2000)). A key point of difference is that this is very much a normative literature. That is to say, the results are of the form “*if  $x$  solution concept is of interest in  $y$  class of games, then both players using  $z$  heuristic is sufficient*”—in particular, the literature was motivated by providing learning foundations to game theoretic solution concepts. Ours is, by contrast, a positive question motivated by current industry practice: “*if both players use a popular machine learning algorithm used in practice, what can we expect as outcomes?*”. None of the existing results apply, because index algorithms considered in our paper are not guaranteed to be a regret-minimizing procedure in a competitive setting.

Regarding collusion, there is a large theoretical literature on repeated games argues that collusive pricing can result in equilibrium (see e.g. Mailath and Samuelson (2006)

for a textbook exposition), and a long line of empirical literature starting with [Green and Porter \(1984\)](#) suggests that such strategies are indeed observed in practice.

### B.5. Behavioral Game Theory

Finally, there has been an increased interest in understanding outcomes in settings of interest where agents are somehow misspecified. For instance in [Spiegler \(2006\)](#) or [Spiegler \(2013\)](#) society misunderstands the relationship between outcomes and the actions of strategic agents, which affects the actions the latter take in equilibrium and resulting outcomes (in the former, in the context of a market for quacks, in the latter with implications to the reforms taken by a politician); see [Spiegler \(2011\)](#) for a textbook overview in the context of industrial organization. [Liang \(2018\)](#) studies outcomes in games of incomplete information where agents behave like statisticians and have limited information. There is a larger literature which studies the outcomes when agents are modeled as statisticians or machine learners, e.g., [Al-Najjar \(2009\)](#), [Al-Najjar and Pai \(2014\)](#), [Acemoglu et al. \(2016\)](#) and [Cherry and Salant \(2018\)](#). Finally, the paper of [Olea et al. \(2019\)](#) studies a setting where agents with different misspecified models compete in an auction after attempting to learn from a common dataset, and tries to understand how the misspecification helps/ handicaps them in terms of auction outcomes. We contribute to this literature both by considering a new form of misspecification motivated by complexity/ computational rather than behavioral concerns, and studying outcomes in this context.

### B.6. Algorithmic Bias

Related to our high-level question on possible downsides of “management by algorithm” is an emerging literature seeks to document so-called “algorithmic bias” (e.g., [Lambrrecht and Tucker \(2019\)](#), [Hajian et al. \(2016\)](#))—the notion that algorithms may discriminate among users based on race, age, sexual orientation etc. If this is true, regulators may want to force companies to “de-bias” their algorithms. We refer interested readers to [Kearns and Roth \(2019\)](#) for a textbook overview.

## APPENDIX C. ROBUSTNESS TO TIME VARYING SHOCKS

Figure 1 shows the robustness of our results to non-stationary error. In our main model we consider a DGP for profit as  $\pi_{j,t} = \pi^*(p_{j,t}, p_{-j,t}) + \varepsilon_{j,t}$  with  $\varepsilon_{j,t} \sim U[-\frac{1}{\delta}, \frac{1}{\delta}]$ . While in our model  $\delta$  is fixed over rounds in our simulation. In this section, we consider a model where  $\delta$  varies across rounds. Here we consider two variations, first with decreasing noise (increasing SNR)  $\frac{1}{\delta_r} = \frac{1}{\delta} r^{-\frac{1}{2}}$  (as in section 6.3 of [Keskin and Zeevi \(2017\)](#)) and second with increasing noise (decreasing SNR),  $\frac{1}{\delta_r} = \frac{1}{\delta} \log(1 + r)$ .

Importantly we note these simulation have the caveat that the UCB algorithm (assumptions are inconsistent with i.i.d. error assumption in [Auer \(2002\)](#)) is no longer guaranteed to result in ex-post optimal monopoly prices. Broadly, these simulations we do replicate the main results in our paper – with high SNR prices are supra-competitive. In addition, the simulations suggest that with decreasing (increasing) errors the results are further amplified (muted). Overall, we believe these simulations provide additional support for our theory model – the critical variation in our results is driven by initial experimentation when the firms have limited data. Importantly we note that in these simulation the UCB does not result in ex-post optimal monopoly prices with increasing noise (bottom chart) and small SNR (left light grey bars). This is consistent with caveat that the UCB is not guaranteed under the data generating process in this simulation.

[Figure 1 about here.]

#### REFERENCES

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): “Fragility of asymptotic agreement under Bayesian learning,” *Theoretical Economics*, 11, 187–225.
- AL-NAJJAR, N. I. (2009): “Decision makers as statisticians: Diversity, ambiguity, and learning,” *Econometrica*, 77, 1371–1401.
- AL-NAJJAR, N. I. AND M. M. PAI (2014): “Coarse decision making and overfitting,” *Journal of Economic Theory*, 150, 467–486.
- AUER, P. (2002): “Using Confidence Bounds for Exploitation-Exploration Trade-offs,” *Journal of Machine Learning Research*, 397–422.
- BERGEMANN, D. AND J. VALIMAKI (1996): “Market Experimentation and Pricing,” Cowles Foundation Discussion Paper 1122.
- BROWN, Z. AND A. MACKAY (2020): “Competition in Pricing Algorithms,” *Available at SSRN 3485024*.
- CALVANO, E., G. CALZOLARI, V. DENICOLÒ, AND S. PASTORELLO (2020): “Artificial intelligence, algorithmic pricing and collusion,” *American Economic Review*, Forthcoming.
- CESA-BIANCHI, N. AND G. LUGOSI (2006): *Prediction, learning, and games*, Cambridge university press.
- CHERRY, J. AND Y. SALANT (2018): “Statistical Inference in Games,” Tech. rep., mimeo.
- COOPER, W. L., T. HOMEM-DE MELLO, AND A. J. KLEYWEGT (2015): “Learning and pricing with models that do not explicitly incorporate competition,” *Operations research*, 63, 86–103.
- FOSTER, D. P. AND R. V. VOHRA (1997): “Calibrated learning and correlated equilibrium,” *Games and Economic Behavior*, 21, 40.

- FUDENBERG, D. AND D. K. LEVINE (1998): *The theory of learning in games*, vol. 2, MIT press.
- GREEN, E. J. AND R. H. PORTER (1984): "Noncooperative collusion under imperfect price information," *Econometrica: Journal of the Econometric Society*, 87–100.
- HAJIAN, S., F. BONCHI, AND C. CASTILLO (2016): "Algorithmic bias: From discrimination discovery to fairness-aware data mining," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, 2125–2126.
- HARRINGTON, J. E. (2018): "Developing Competition Law for Collusion by Autonomous Artificial Agents," *Journal of Competition Law & Economics*, 14, 331–363.
- HART, S. AND A. MAS-COLELL (2000): "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, 68, 1127–1150.
- KEARNS, M. AND A. ROTH (2019): *The ethical algorithm: The science of socially aware algorithm design*, Oxford University Press.
- KELLER, G. AND S. RADY (1999): "Optimal experimentation in a changing environment," *The review of economic studies*, 66, 475–507.
- (2003): "Price dispersion and learning in a dynamic differentiated-goods duopoly," *RAND Journal of Economics*, 138–165.
- KESKIN, N. B. AND A. ZEEVI (2014): "Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies," *Operations Research*, 62, 1142–1167.
- (2017): "Chasing demand: Learning and earning in a changing environment," *Mathematics of Operations Research*, 42, 277–307.
- (2018): "On incomplete learning and certainty-equivalence control," *Operations Research*, 66, 1136–1167.
- KIRMAN, A. P. (1975): "Learning by firms about demand conditions," in *Adaptive economic models*, Elsevier, 137–156.
- LAMBRECHT, A. AND C. TUCKER (2019): "Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads," *Management Science*.
- LIANG, A. (2018): "Games of Incomplete Information Played by Statisticians," *Working paper, University of Pennsylvania*.
- MAILATH, G. J. AND L. SAMUELSON (2006): *Repeated games and reputations: long-run relationships*, Oxford university press.
- MIKLÓS-THAL, J. AND C. TUCKER (2019): "Collusion by algorithm: Does better demand prediction facilitate coordination between sellers?" *Management Science*, 65, 1552–1561.
- MISRA, K., E. M. SCHWARTZ, AND J. ABERNETHY (2019): "Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments," *Marketing Science*.
- NAIR, H. (2007): "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US Market for Console Video-Games," *Quantitative Marketing and*

*Economics*, 5, pp. 239–292.

O’CONNOR, J. AND N. WILSON (2019): “Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition,” *FTC Bureau of Economics, Working Paper*.

OLEA, J. L. M., P. ORTOLEVA, M. M. PAI, AND A. PRAT (2019): “Competing Models,” *arXiv preprint arXiv:1907.03809*.

RUSTICHINI, A. AND A. WOLINSKY (1995): “Learning about variable demand in the long run,” *Journal of economic Dynamics and Control*, 19, 1283–1292.

SALCEDO, B. (2015): “Pricing algorithms and tacit collusion,” *Manuscript, Pennsylvania State University*.

SPIEGLER, R. (2006): “The market for quacks,” *The Review of Economic Studies*, 73, 1113–1131.

——— (2011): *Bounded rationality and industrial organization*, Oxford University Press.

——— (2013): “Placebo reforms,” *American Economic Review*, 103, 1490–1506.

## Figures

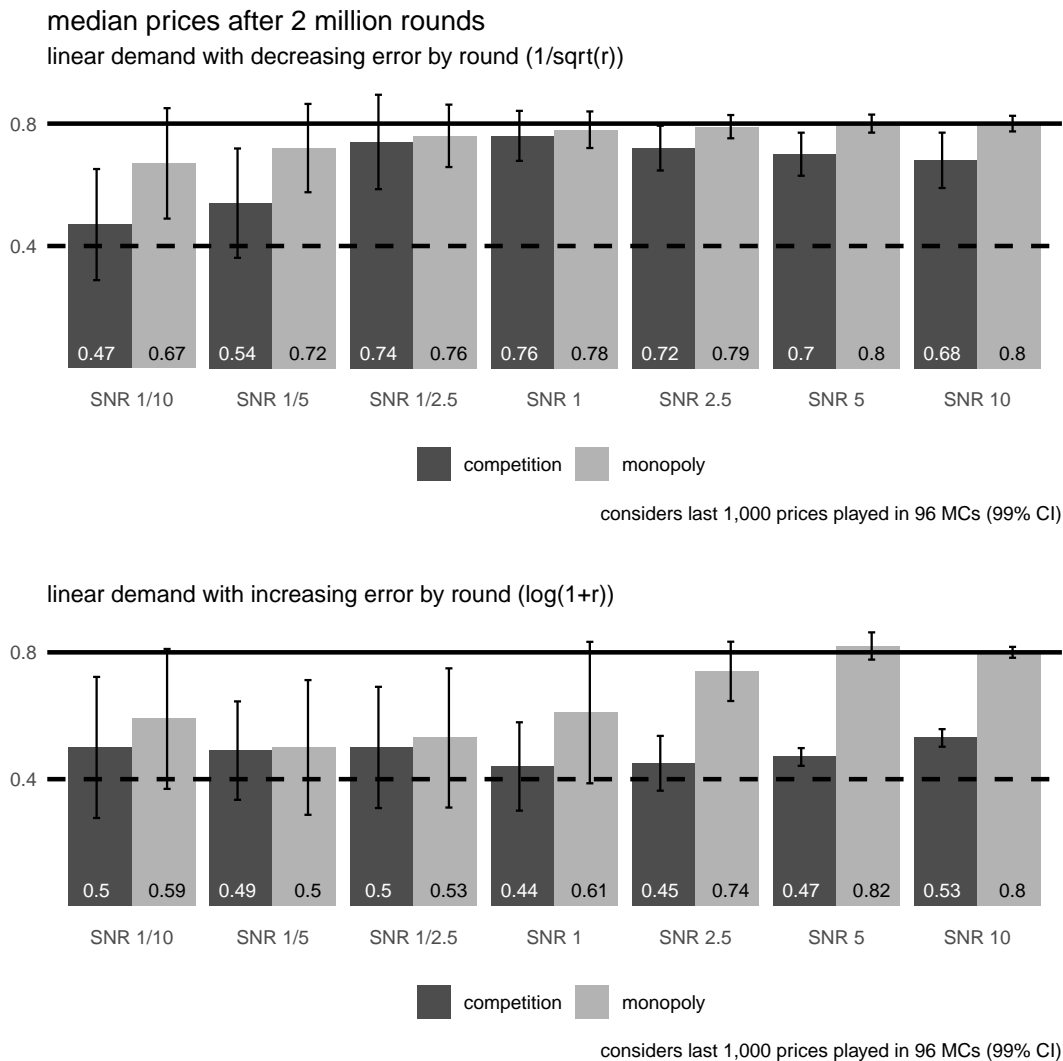


FIGURE 1. Estimated median prices by algorithm and level of competition. In this simulation we consider a model where experimental noise decreases (or increases by round). In the top chart we consider decreasing noise (increasing SNR) [ $\frac{1}{\delta_r} = \frac{1}{\delta} r^{-\frac{1}{2}}$ ] and in the bottom chart we consider increasing noise (decreasing SNR) [ $\frac{1}{\delta_r} = \frac{1}{\delta} \log(1+r)$ ]. The error bars represent 99% confidence interval for the estimated median across MC simulations. The dark gray bars represent a setting where two firms are running simultaneous algorithms, while the light gray bars represent a setting when a monopolist is jointly pricing the products. For each simulation, we consider the median price charged in the last 1,000 rounds out of 10 million rounds. The dashed lines reflect the competitive equilibrium prices; the solid lines reflect monopoly prices. [Algorithm used: UCB tuned]