An Economic Test for an Unlawful Agreement to Adopt a Third-Party's Pricing Algorithm

Joseph E. Harrington, Jr.*

Department of Business Economics & Public Policy

The Wharton School

University of Pennsylvania

harrij@wharton.upenn.edu

11 July 2024

Abstract

AI has helped fuel a growing market in the supply of pricing algorithms by software developers. While there is an efficiency rationale for outsourcing pricing, anticompetitive concerns have been expressed when competitors in a market adopt the same pricing algorithm. These concerns have resulted in private litigation claiming a third-party company (who developed the pricing algorithm) and firms (who adopted it) had an unlawful agreement. This study develops an empirical test for determining whether firms' adoption decisions are coordinated. If adoption decisions are coordinated then adopters' average price is increasing in the number of adopting firms, while if adoption decisions are independent then adopters' average price does not depend on the number of adopting firms. This test could provide economic evidence to support a claim of an unlawful agreement between a third-party developer and adopting firms.

*I appreciate the comments of Espen Moen, Sander Onderstal, Geza Sapi, the anonymous referees, the editors, and participants at the Learning in Games and Algorithmic Collusion Workshop at the Oxford-Man Institute (October 2023), the Toulouse School of Economics (April 2024), the BECCLE Competition Policy Conference (May 2024), and the CRESSE Conference (July 2024). After this study was initiated, I was retained by a defendant in private litigation associated with the use of a data analytics company by hotels. This study is not funded in whole or in part, either directly or indirectly, by any person or entity. No client or other interested party has a right to review, or has reviewed, this paper.

1 Introduction

Awash in data, firms face the challenge of figuring out how to most effectively use it to augment decision-making and do so in real time. AI is playing an increasing role in addressing that challenge whether it involves managing inventories, the supply chain, or customer relations (Haan, 2023). Access to the benefits of AI is not limited to large companies for an advantage of data-driven decision-making is that it can be outsourced. A firm can share its data with a third party who, with superior programming expertise, computing power, and data, is better equipped to use it to assist decision making in a cost-effective manner. Just like any "make or buy" decision, a firm may find it more efficient to buy the input - here it is support for decision-making - rather than do it themselves. That creates an opportunity for IT service providers to invest and innovate to attract more business customers. Leaving these third-party vendors unfettered in how they design their services and to whom they sell would seem to be surest way to realize the greatest benefits.

The matter is more complicated, however, when the service that an IT company is supplying is guidance on how a firm should price. Through its design and implementation of a pricing algorithm, a third party can make it possible for a firm to adjust its price more rapidly to high-frequency demand and supply shocks, which will increase transaction volume and social welfare. By tailoring price to more narrowly-defined market segments, a third party's pricing algorithm allows for more precise price discrimination which increases profits and benefits those consumers (who are typically lower income) who pay a lower price. These efficiencies are attractive but are also the source of a conundrum. An IT company who successfully innovates is rewarded with widespread adoption of its service which could mean competitors in a market are using a pricing algorithm supplied by the same third party. Given an input supplier will seek to create value for its customers - as doing so will allow it to charge more and sell more - an IT service provider with a customer base comprising competitors may be inclined to design the pricing algorithm to reduce competition and thereby raise adopters' profits. Concerns about these possible anticompetitive effects have been widely and regularly expressed:

> OECD: "[C]oncerns of co-ordination could arise if firms outsourced the creation of algorithms to the same IT companies and programmers."[1]
>
> United Kingdom's Competition & Markets Authority: "If a sufficiently large proportion of an industry uses a single algorithm to set prices, this could result in ... the ability and incentive to increase prices."[2]
>
> Testimony before the U.S. Congress: "[Another] area of concern with the use of pricing algorithms seems more subtle and hard to detect: Companies avoiding price competition by using the same third-party vendor to collect data on supply and demand and 'recommend' pricing or output behaviors that facilitate price

[1] OECD, "Algorithms and Collusion - Background Note by the Secretariat," DAF/COMP(2017)4, 9 June 2017, para. 68.

[2] United Kingdom Competition & Markets Authority, "Pricing Algorithms: Economic Working Paper on the Use of Algorithms to Facilitate Collusion and Personalised Pricing," 8 October 2018, para. 5.21.

coordination."[3]

> German Monopolies Commission: "It may happen that such an IT service provider sells an algorithm that it knows or accepts could contribute to a collusive market outcome. It is even conceivable that individual IT service providers see such a contribution as an advantage, as it makes the algorithm more attractive for users interested in profit maximization."[4]

There is some evidence, and claims of evidence, that these anticompetitive effects have occurred. data analytics company a2i Systems created a pricing algorithm to assist retail gasoline companies in the pricing of their product. This is a natural market for a third-party developer to enter because the complexity of pricing in this market is well documented (Eckert, 2013). Gasoline pricing can involve high-frequency price changes - often several times over the course of a day - and price cycles where prices sharply rise then gradually fall with that pattern repeating itself. There is then likely to be room for a third party to develop a more effective pricing rule than is currently being used. a2i Systems's software was widely adopted in Germany and a recent study by Assad, Clark, Ershov, and Xu (2023) found evidence of anticompetitive effect: "Adoption increases margins, but only for non-monopoly stations. In duopoly and triopoly markets, margins increase only if all stations adopt, suggesting algorithmic pricing has a significant effect on competition."[5]

Though the claims are still to be evaluated by the judicial process, plaintiffs in several legal cases in the United States argue that an IT service provider and competitors in a market had an illegal agreement to restrain competition. In the market for apartment rentals, data analytics companies RealPage and Yardi designed and sold pricing algorithms to be used by apartment property owners, and are now being sued in separate litigation for violating Section 1 of the Sherman Act.[6] According to RealPage, many owners have historically used simple heuristics in setting rents. If that is the case, there is room for a third party to develop a more profitable pricing rule. But plaintiffs claim that RealPage went farther as it delivered value to adopters by reducing competition between them, and adopters were complicit.[7] A

[3] Written Testimony of Bill Baer, U.S. Senate Committee on the Judiciary, Subcommmittee on Competition Policy, Antitrust, and Consumer Rights, Hearings on "The New Invisible Hand? The Impact of Algorithms on Competition and Consumer Rights," December 13, 2023, p. 2. <https://www.judiciary.senate.gov/imo/media/doc/2023-12-13_pm_-_testimony_-_baer.pdf> (downloaded January 17, 2024) Bill Baer is a former Assistant Attorney General of the Antitrust Division of the U.S. Department of Justice.

[4] German Monopolies Commission, XXII. Biennial Report, Chapter on "Algorithms and Collusion," 2018, para. 263.

[5] Abstract of Assad, Clark, Ershov, and Xu (2023).

[6] The complaint against RealPage is Bason, et al. v. RealPage, Inc., et al., No. 3:22-cv-01611-WQH-MDD, U.S. District Court, Southern District of California, Oct. 18, 2022; and against Yardi is Duffy, et al. v. Yardi Systems, Inc., et al., No. 2:2023-cv-0139, U.S. District Court, Western District of Washington at Seattle, Sept. 8, 2023. It has also been reported that "[t]he Department of Justice's Antitrust Division has opened an investigation into whether rent-setting software made by a Texas-based real estate tech company is facilitating collusion among landlords, according to a source with knowledge of the matter." ("Department of Justice Opens Investigation Into Real Estate Tech Company Accused of Collusion with Landlords," by Heather Vogell, ProPublica, November 23, 2022 https://www.propublica.org/article/yieldstar-realpage-rent-doj-investigation-antitrust)

[7] "RealPage and participating Lessors have provided one another with such mutual assurances, agreeing

recent study by Calder-Wang and Kim (2023) offers evidence for both procompetitive and anticompetitive effects: "Our findings suggest that algorithm adoption helps building managers set more responsive prices: buildings with the software increase prices during booms but lower prices during busts, compared to non-adopters in the same market. However, we also find evidence that greater algorithm penetration can lead to higher prices, raising rents among both adopters and non-adopters in the same market."[8]

In the market for hotel rooms, third-party developer Rainmaker developed a pricing algorithm to assist hoteliers in pricing more profitably.[9] As with retail gasoline and apartment rentals, there is a plausible efficiency basis, for hotel demand is highly dynamic as it is subject to demand variation (e.g., seasonal fluctuations and demand-increasing events such as conventions) and hotel supply is a perishable good (e.g., a hotel room for August 1, 2023 is lost forever after that date). Consequently, profit-maximizing prices change as a date approaches and do so in a potentially complex way. There is then a legitimate basis for a data analytics company to invest in creating a superior pricing rule and for firms to adopt it. Nevertheless, there may be an incentive for it to design a collusive pricing algorithm and to facilitate an agreement among competitors to adopt it, which is what is being claimed by plaintiffs who bought a hotel room in Las Vegas or Atlantic City during the time that some hotels were using Rainmaker's pricing algorithm.[10]

In light of what is going on in the markets for apartments, gasoline, and hotels, the antitrust problem of competitors adopting a third-party developer's pricing algorithm is here and now. From a welfare perspective, how do we determine when competitors utilizing the same pricing algorithm provided by a third party is harmful to consumers? How do we permit adoption on efficiency grounds while preventing it when it is anticompetitive? Unless one is prepared to prohibit firms in the same market from adopting a pricing algorithm from the same third party - which is extreme and unjustified by the existing evidence - we face the challenge of distinguishing between adoptions that are anticompetitive and those that are not. From a legal perspective, how do we determine when competitors adopting a third-party developer's pricing algorithm is a violation of competition law? How do we know when competitors' adoption decisions were coordinated not independent, and the third-party developer facilitated those joint adoption decisions?

This study addresses the challenge of determining when a third party (who developed

among themselves not to compete on price for the sale of multifamily residential real estate leases. They have effectuated their agreement through two mutually reinforcing mechanisms. First, participating Lessors have agreed to set prices using RealPage's coordinated algorithmic pricing. Second, participating Lessors have agreed to stagger their lease renewal dates through RealPage, to avoid (otherwise natural) oversupplies in rental properties." Bason, et al. v. RealPage, Inc., et al., para. 46.

[8] Abstract of Calder-Wang and Kim (2023).

[9] For the Las Vegas market, the complaint is Richard Gibson, et al. v. MGM Resorts International, et al., No. 2:23-cv-00140, U.S. District Court, District of Nevada, January 25, 2023; and for the Atlantic City market, the complaint is Altman, et al. v. Caesars Entertainment, Inc., et al., No. 2 :23-cv-02536, U.S. District Court, District of New Jersey, May 9, 2023.

[10] "In furtherance of the contract, combination, or conspiracy, Defendants have committed one or more of the following acts: a) provided information to be used in the operation of the pricing algorithms; b) created and operated algorithms that provided pricing recommendations to Defendants; c) knowingly used algorithms that incorporated information from other Defendants in setting pricing recommendations; and/or d) set prices based in whole or in part on pricing recommendations provided by Rainmaker Group." Richard Gibson, et al. v. MGM Resorts International, et al., para. 89.

the pricing algorithm) and competitors in a market (who adopted it) have an agreement in violation of competition law. Obviously, evidence of documented communications to coordinate their conduct would be most convincing to the courts. This could involve direct and express communications between competitors and the third party or, as with a hub-and-spoke cartel, the third party using bilateral communications with competitors in order to establish a common understanding to adopt the pricing algorithm. However, let us suppose evidence of communications is insufficient, perhaps because it is not express or is not private (such as a third-party developer engaging in public announcements which could facilitate coordination though argued by them to be nothing more than marketing).[11] In that situation, circumstantial evidence is required.[12] One class of circumstantial evidence is referred to as "plus factors" in the United States and the concept is relevant to other jurisdictions.[13] One of the more compelling plus factors is conduct against a firm's independent self-interest. Examples are a firm adopting a higher price or choosing to no longer negotiate price with customers. It is not difficult to argue that raising prices or prohibiting discounts off of list prices would be unprofitable unless rival firms also chose to do so. However, the argument is more challenging with the adoption of a third-party developer's pricing algorithm because it can be profitable to adopt even if rival firms do not. If the pricing algorithm allows a firm to more effectively respond to demand shocks or engage in price discrimination, that can enhance profit while holding fixed rival firms' pricing rules. It is then more challenging to use this plus factor when the claimed agreement is adopting a third-party developer's pricing algorithm because the pricing algorithm may enhance efficiency as well as restrain competition.

The contribution of this study is to offer a plus factor for proving an unlawful agreement between the third party who developed the pricing algorithm and competitors in a market who adopted it.[14] If a third-party developer is part of an agreement with competitors, it will then design the pricing algorithm differently than if there is no agreement. More specifically, if adoptions are coordinated then adopters' prices will be increasing in the adoption rate (i.e., the fraction of firms who adopt) and, on average, adopters will price higher than non-adopters. In contrast, if firms' adoption decisions are independent then adopters' prices do not change with the adoption rate and, on average, adopters and non-adopters price the same. While the prediction under coordinated adoptions is expected - for the third party is effectively acting as a cartel manager - the prediction under independent adoptions is not. Even with independent adoptions, one might think that a third party would reduce competition in order to create more value from adoption; however, I show that is not the

[11] As reviewed in Harrington (2022a), U.S. courts have been reluctant to find an unlawful agreement based only on public announcements, even egregious ones.

[12] "Circumstantial evidence can come in several forms, including evidence of communications between rivals and economic evidence. Economic evidence consists of firm conduct, market structure, and evidence of facilitating practices. All types of evidence can be useful in a case and they should be employed together." OECD, "Prosecuting Cartels without Direct Evidence" DAF/COMP/G(2006)7, 11 September 2006, p. 18.

[13] "One formulation, developed in the United States in civil cases ..., requires that there exist certain 'plus factors,' which prove that agreement is more likely the cause of the parallel conduct than independent action. ... Other jurisdictions do not use the same terminology as US courts, but it seems that the analysis that they apply is similar." OECD, "Prosecuting Cartels without Direct Evidence" DAF/COMP/G(2006)7, 11 September 2006, p. 29.

[14] Identifying plus factors for collusion involving algorithms is discussed in Gal (2019).

case. Higher average prices result only when adoption decisions are coordinated. Therefore, the properties of the pricing algorithm can inform us when the pricing algorithm is designed for firms who have an agreement to adopt it.

In considering the body of research pertaining to algorithms and collusion, there is an extensive and growing theoretical literature on learning algorithms resulting in supracompetitive prices.[15] While this practice has not yet been documented in actual markets, it would create a serious antitrust challenge should it occur.[16] That situation is to be contrasted with when a third party develops the pricing algorithm. The fundamental economic difference is that a firm designing its own pricing algorithm does so in order to maximize its profit from *using* it, while a third-party developer designs a pricing algorithm in order to maximize its profit from *selling* it. A crucial implication is that a third-party developer will recognize that its pricing algorithm will "compete against itself" when adopted by competitors in a market. What this means for the design of the pricing algorithm and possible anticompetitive conduct is largely not understood, for the literature on the provision of pricing algorithms by a third party is scarce. As mentioned above, there are two empirical studies: Assad et al (2023) and Calder-Wang and Kim (2023). As regards theoretical analysis, there is only previous work by the author (Harrington, 2022b). The analysis of the current paper extends Harrington (2022b) from the case of two firms to when there are many firms so that predictions can be derived regarding how prices depend on how many firms adopt.[17]

Section 2 describes the model along with some preliminary analysis. In Section 3, a third-party developer's pricing algorithm is characterized when firms' adoptions decision are coordinated and when they are independent. Section 4 compares the pricing algorithm under coordinated and independent adoptions. A description of the plus factor for establishing an unlawful agreement is provided in Section 5 along with a discussion of its empirical implementation. Section 6 offers an auditing approach to develop economic evidence of an agreement. Section 7 concludes.

[15] Work on algorithmic collusion (either adopting supracompetitive prices or collusive pricing rules) includes Waltman and Kaymak (2008), Kimbrough and Murphy (2009), Cooper, Homen-de-Mello, and Kleywegt (2015), Calvano, Calzolari, Denicolò, and Pastorello (2020, 2021), Eschenbaum, Mellgren, and Zahn (2021), Hansen, Misra, and Pai (2021), Kastius and Schlosser (2021), Klein (2021), Musolff (2021), Abada, Lambin, and Tchakarov (2022), Banchio and Mantegazza (2022), van den Boer, Meylahn, and Schinkel (2022), Cartea, Chang, Penalva, and Waldon (2022), Leisten (2022), Meylahn and den Boer (2022), Sanchez-Cartas and Katsamakas (2022), Abada and Lambin (2023), Asker, Fershtman, and Pakes (2023), Brown and MacKay (2023), Possnig (2023), Abada, Harrington, Lambin, and Meylahn (2024), and Epivent and Lambin (2024).

[16] For a discussion of the legal challenges, see Mehra (2016), Ezrachi and Stucke (2017), Johnson (2017), Oxera (2017), Deng (2018), Harrington (2018), Gal (2019), Schwalbe (2019), Calvano, Calzolari, Denicolò, Harrington, and Pastorello (2020), Beneke and Mackenrodt (2021), Asil and Wollmann (2023), Ezrachi and Stucke (2023), (2023), and Inderst and Thomas (2024).

[17] A distinct but related set of issues arises with competitors who contract a third party to submit bids on their behalf at online advertising auctions; see Decarolis, Goldmanis, and Penta (2020).

2 Model

2.1 Model and Equilibrium without a Third-Party Developer

Consider a market where firms have symmetrically differentiated products and a common and constant marginal cost c. For tractability purposes, demand is assumed to be linear: $a - bp_i + dP_{-i}$ where p_i is a firm's own price, P_{-i} is the average price of rival firms, and $b > d > 0$. So as to allow for multiple firms in a tractable manner, there is assumed to be a continuum of firms.[18] A critical feature is demand variation with respect to a which is assumed to have a continuously differentiable cdf $G : [\underline{a}, \overline{a}] \rightarrow [0, 1]$ with mean μ and variance σ^2. The higher is a, the stronger is firm demand in the sense that, for any price vector, the amount of demand is larger and demand is more price-inelastic. Assume $\underline{a} - (b - d)c > 0$ so, at the static Nash equilibrium, demand is positive for all demand states. The demand variable a has two interpretations. A firm could be facing a single demand curve and a is a demand shock with distribution G. In that case, price may condition on the current demand shock. Alternatively, a firm faces a collection of market segments represented by G. In that case, price may condition on the market segment a.

In the absence of a third party, it is supposed that the demand shock a occurs at a higher frequency than a firm's pricing decisions or, when G represents a collection of market segments, the firm cannot distinguish among those segments when pricing. In that case, a firm is incapable of conditioning price on the demand state. A symmetric Nash equilibrium price p^N is defined by:

$$p^N \equiv \arg \max_{p_i \in \Re_+} \int (p_i - c) \left(a - bp_i + dp^N\right) G'(a)da \Leftrightarrow p^N = \frac{\mu + bc}{2b - d}.$$

The associated expected profit is

$$\pi^N \equiv \int \left(\frac{\mu + bc}{2b - d} - c\right) \left(a - (b - d) \left(\frac{\mu + bc}{2b - d}\right)\right) G'(a)da = \frac{b \left(\mu - (b - d)c\right)^2}{(2b - d)^2}.$$

A useful benchmark for subsequent analysis is when a fraction θ of firms can condition price on the demand state and the remaining firms cannot and thus set a uniform price. It is shown in the Appendix that the pricing rule for the firms who can condition price on the demand state is[19]

$$p^F(a) \equiv \frac{d(1 - \theta)\mu + (2b - d\theta)bc}{(2b - d\theta)(2b - d)} + \left(\frac{1}{2b - d\theta}\right) a, \tag{1}$$

and the firms who set a uniform price choose:

$$\frac{\mu + bc}{2b - d}.$$

[18] This demand specification is the extension of a common specification - see, for example, Vives (1999), p. 146 - to when there is a continuum of firms. The advantage of a continuum of firms is that it simplifies the analysis because, as we will see, an individual firm's adoption decision does not affect the third party's pricing algorithm. That property is a good approximation for when there are many firms in the market such as many small apartment property owners.

[19] The superscript F refers to the pricing algorithm being designed by the f(irm) rather than the third party.

Note that

$$E_a\left[p^F(a)\right] = \frac{\mu + bc}{2b - d}$$

so average price is the same for all firms.

2.2 Model with a Third-Party Developer

Let us now introduce a single third-party developer whose business model is to design and sell pricing algorithms. The efficiency delivered by the third-party developer is that its pricing algorithm can track the high-frequency demand shock (or market segment) a and condition price on it. An algorithm is denoted $\phi : [\underline{a}, \overline{a}] \rightarrow \Re_+$.

The setting is modelled as a simultaneous-move game involving the third-party developer and firms. The third-party developer chooses a design ϕ and a fee f (for a firm to be able to use ϕ) in order to maximize its profit given firms' expected adoption decisions. Firms make adoption decisions that are optimal given their expectations on (ϕ, f). After adoption decisions are made, the demand state is realized and firms choose prices. If a firm adopted ϕ then it prices at $\phi(a)$, and if a firm did not adopt ϕ then it chooses an optimal price that does not condition on a. The solution concept is Nash equilibrium.[20]

Given (ϕ, f), the third-party developer's profit (or, equivalently, revenue, as there is assumed to be no marginal cost to selling a pricing algorithm) is simply the number of firms that buy ϕ multiplied by f. As explicitly shown in Harrington (2022b) for the duopoly case, the third-party developer's optimal choice is characterized as follows. First, the fraction of firms that will adopt the pricing algorithm, which is denoted θ, is conjectured. I elaborate below how θ may be determined. Second, the third-party developer designs ϕ to maximize the willingness-to-pay (WTP) for an adopter, conditional on a fraction θ of firms adopting. Third, the third-party developer chooses f to equal the WTP (as that will maximize its profit) and results in a fraction θ of firms adopting.

Regarding the determination of θ, if adoption is only constrained by a firm finding it profitable to adopt then the third-party developer will choose θ to maximize its profit taking into account how it will affect the fee it can charge, which is the assumption made in Harrington (2022b). Alternatively, some firms may, for some unmodelled reason, not be willing to delegate their pricing to an algorithm, which then puts an upper bound on θ. Or firms could have some heterogenous firm-specific adoption cost - perhaps from integrating it into their system - which could result in some firms adopting and some not. Or the third-party developer may require marketing and training for adoptions to occur which could limit how fast it sells the pricing algorithm, which would also constrain θ. As the analysis will focus on how the design decision depends on θ, I can be agnostic about the determination of θ. What is critical is that the pricing algorithm is designed by the third-party developer for when a fraction θ of firms adopt it.

[20]In taking an equilibrium approach, it is implicitly assumed firms learn the profit from adopting and not adopting the pricing algorithm, and non-adopting firms and the third party learn the optimal price and pricing algorithm, respectively. These assumptions are admittedly strong but nevertheless useful for identifying relevant incentives when it comes to the design of the third party's pricing algorithm and its adoption by firms.

The model here differs from that in Harrington (2022b) in two ways. First, as has been mentioned, there is a continuum of firms rather than two firms. Second, the model in Harrington (2022b) allows an adopter to observably commit to its pricing rule prior to non-adopters choosing their prices which means it acts as a price leader. That commitment effect is not present in the current model. I do not believe commitment is essential to the third-party developer's design decision, nor is it clear which assumption is more appropriate. I chose not to allow commitment here because it simplifies the analysis. Furthermore, when market variation σ^2 is not too low, the key properties of the pricing algorithm in Harrington (2022b) and here (when adoption decisions are independent) are identical so the change in these two assumptions does not alter the operative forces in the model. In sum, no commitment is an assumption of convenience which is believed to be without consequence for our results, and the assumption of many firms allows us to explore comparative statics with respect to the adoption rate.

In preparation for the analysis in Section 3, I offer some preliminary results here. Given the linearity of demand, an optimal pricing algorithm will be an affine function of a: $\phi(a) = \alpha + \gamma a$ for some (α, γ).[21] Given a fraction θ of firms adopt $\alpha + \gamma a$, let us characterize the price of the remaining $1 - \theta$ firms who are non-adopters. Given the pricing algorithm of the adopters, the non-adopters are assumed to choose an equilibrium (uniform) price:

$$\begin{aligned} p^{NA} &\equiv \arg\max_{p} \int (p-c)\left(a - bp + d\left(\theta(\alpha+\gamma a) + (1-\theta)p^{NA}\right)\right) G'(a)da \\ &= \arg\max_{p} (p-c)\left(\mu - bp + d\left(\theta(\alpha+\gamma\mu) + (1-\theta)p^{NA}\right)\right). \end{aligned}$$

The first-order condition (FOC) is:

$$\mu - bp^{NA}(\theta) + d\left(\theta(\alpha+\gamma\mu) + (1-\theta)p^{NA}\right) - bp^{NA} + bc = 0,$$

which we can solve for the symmetric equilibrium non-adopters' price when a fraction θ of firms adopt the pricing algorithm $\alpha + \gamma a$:

$$p^{NA} = \frac{\mu + bc + d\theta(\alpha+\gamma\mu)}{2b - d(1-\theta)}.$$

3 Design of the Pricing Algorithm

3.1 Coordinated Adoption Decisions

We begin with the case of coordinated adoptions which is the situation of concern for competition authorities and is the claim of plaintiffs in recent private litigation. Suppose a fraction θ of firms have agreed to adopt the third party's pricing algorithm but only if all do so. The role of the third party can be described in either of two ways. First, it is approached by the colluding firms who instruct the third party to design the pricing algorithm to maximize their joint profit (and the third party will be properly compensated). Second, the third

[21] As shown in Section 3.1, the solution must be affine when firms make coordinated adoption decisions. For the case of independent adoption decisions in Section 3.2, I focus on the case of an affine solution.

party is knowledgeable of the agreement and designs the pricing algorithm to maximize the WTP of the adopting firms.[22] As their adoption decisions are joint - either all adopt or none adopt - their WTP for the pricing algorithm is an adopter's profit when a fraction θ of firms adopt minus a non-adopter's profit when no firms adopt. When adoption decisions are coordinated, the third-party developer's design problem is then:

$$\max_{(\alpha,\gamma)} \int (\alpha + \gamma a - c) \left(a - b(\alpha + \gamma a) + d\left(\theta(\alpha + \gamma a) + (1-\theta)p^{NA}\right)\right) G'(a)da - \pi^N. \quad (2)$$

π^N is the Nash equilibrium profit when no firm adopts and, as it does not depend on (α, γ), maximizing the WTP in (2) is the same as maximizing aggregate adopters' profit or, equivalently, an adopter's profit.[23] In other words, the third-party developer acts like a cartel manager where the cartel encompasses a fraction θ of firms.

Theorem 1 *Under coordinated adoptions, if a fraction θ of firms adopt then the third-party developer's pricing algorithm is*

$$\alpha = \frac{(b-d\theta)\,c\,(2b(2b-d(1+\theta)) + d^2\theta(1-\theta)) + d(1-\theta)\,(2b(\mu+bc) - d\theta(\mu+(b+d\theta)c))}{2(b-d\theta)\,(2b(2b-d(1+\theta)) + d^2\theta(1-\theta))}$$

$$\gamma = \frac{1}{2\,(b-d\theta)}.$$

Proof. The third-party developer's optimal pricing algorithm is one that maximizes an adopter's profit given the expected price of the non-adopters. If non-adopters are expected to choose the uniform price p^{NA} then, for any a, the pricing algorithm is designed to set price to maximize:

$$(p-c)\left(a - bp + d\left(\theta p + (1-\theta)\,p^{NA}\right)\right).$$

Note that p is chosen by all adopting firms, which reflects their joint adoption decisions. The FOC is:

$$0 = a - 2(b-d\theta)p + d(1-\theta)p^{NA} + (b-d\theta)\,c,$$

which can be solved for price to yield:

$$p^C = \frac{a + (b-d\theta)\,c + d(1-\theta)p^{NA}}{2(b-d\theta)}. \quad (3)$$

The superscript C refers to "coordinated" adoption decisions.

[22] As just noted, it is assumed the third party is aware of an agreement among the firms. In the context of the model, there would be no reason for firms to have an agreement to adopt and implement a third party's pricing algorithm unless it were to affect the design of the pricing algorithm but then that requires the third party to be informed of the agreement. Consistent with hub-and-spoke collusion, it is also possible the third party is not just aware of the agreement but is partly responsible for it as it facilitates it. Also relevant to this assumption is that the plaintiffs in private litigation in the markets for apartment rentals and hotels claim there is an agreement between the data analytics company and subscribing firms.

[23] Note that p^{NA} is taken as fixed. Thus, the third-party developer does not act as a leader by supposing its choice of (α, γ) will affect p^{NA}. This is the assumption of no commitment.

The non-adopters will settle on an equilibrium price given their expectation of the adopters' price $E_a[p^C]$. p^{NA} is defined by:

$$p^{NA} = \arg\max\,(p-c)\left(\mu - bp + d\left(\theta E_a[p^C] + (1-\theta)p^{NA}\right)\right)$$

with FOC:

$$0 = \mu - bp^{NA} + d\left(\theta E_a[p^C] + (1-\theta)p^{NA}\right) - bp^{NA} + bc \Leftrightarrow$$

$$p^{NA} = \frac{\mu + bc + d\theta E_a[p^C]}{2b - d(1-\theta)}. \tag{4}$$

From (3), we have:

$$E_a[p^C] = \frac{\mu + (b - d\theta)\,c + d(1-\theta)p^{NA}}{2(b-d\theta)}. \tag{5}$$

Using (4), substitute for p^{NA} in (5):

$$E_a[p^C] = \frac{\mu + (b-d\theta)\,c + d(1-\theta)\left(\frac{\mu+bc+d\theta E_a[p^C]}{2b-d(1-\theta)}\right)}{2(b-d\theta)},$$

and solve for $E_a[p^C]$:

$$E_a[p^C] = \frac{2b(\mu + (b-d\theta)c) + d^2\theta(1-\theta)c}{2b(2b - d(1+\theta)) + d^2\theta(1-\theta)}. \tag{6}$$

Inserting (6) into (4), a non-adopter's price is:

$$p^{NA} = \frac{\mu + bc + d\theta\left(\frac{2b(\mu+(b-d\theta)c)+d^2\theta(1-\theta)c}{2b(2b-d(1+\theta))+d^2\theta(1-\theta)}\right)}{2b-d(1-\theta)} = \frac{2b(\mu+bc) - d\theta(\mu + (b+d\theta)c)}{2b(2b-d(1+\theta)) + d^2\theta(1-\theta)}. \tag{7}$$

Using (7) in (3), an adopter's price is:

$$\begin{aligned} p^C &= \frac{a + (b-d\theta)\,c + d(1-\theta)\left(\frac{2b(\mu+bc)-d\theta(\mu+(b+d\theta)c)}{2b(2b-d(1+\theta))+d^2\theta(1-\theta)}\right)}{2(b-d\theta)} \\ &= \frac{(b-d\theta)\,c\,(2b(2b-d(1+\theta)) + d^2\theta(1-\theta)) + d(1-\theta)\,(2b(\mu+bc) - d\theta(\mu+(b+d\theta)c))}{2(b-d\theta)\,(2b(2b-d(1+\theta)) + d^2\theta(1-\theta))} \\ &\quad + \left(\frac{1}{2\,(b-d\theta)}\right)a, \end{aligned}$$

which proves the theorem. ∎

Let us derive some properties of the pricing algorithm in Theorem 1. From the proof of Theorem 1, we find that, on average, adopters price higher than non-adopters:

$$\begin{aligned} &E_a[p^C] - p^{NA} \\ =\; &\frac{2b(\mu + (b-d\theta)c) + d^2\theta(1-\theta)c}{2b(2b-d(1+\theta)) + d^2\theta(1-\theta)} - \frac{2b(\mu+bc) - d\theta(\mu+(b+d\theta)c)}{2b(2b-d(1+\theta)) + d^2\theta(1-\theta)} \\ =\; &\frac{d\theta(\mu - (b-d)c)}{2b(2b-(1+\theta)d) + d^2\theta(1-\theta)} > 0. \end{aligned}$$

Due to the recommendations of the third-party developer, adopters are pricing to maximize their joint profit, while each non-adopter is pricing to maximize its own profit. Given the adopters are setting supracompetitive prices, a non-adopter maximizes its profit by setting its uniform price to undercut adopters' average price.

Next we turn to assessing the effect of the fraction of firms that adopt the pricing algorithm. The average price for adopters and non-adopters are both increasing in the adoption rate θ :

$$\frac{\partial E_a[p^C]}{\partial \theta} = \frac{2bd(2b-d+2d\theta)\,(\mu-(b-d)c)}{\left(4b^2-2bd(1+\theta)+d^2\theta(1-\theta)\right)^2} > 0$$

$$\frac{\partial p^{NA}}{\partial \theta} = \frac{d^2\theta\,(4b-d\theta)\,(\mu-(b-d)c)}{\left(4b^2-2bd(1+\theta)+d^2\theta(1-\theta)\right)^2} > 0.$$

As more firms agree to adopt the pricing algorithm, the third-party developer designs the pricing algorithm to charge higher prices because the cartel of adopting firms is more inclusive; hence, there are fewer non-adopting firms acting as a competitive constraint. In response to higher prices for adopters, the non-adopters price higher, too.

The preceding properties are summarized in Corollary 2.

Corollary 2 *Under coordinated adoptions: i) the average price of adopters and non-adopters are increasing in the adoption rate; and ii) on average, the price of adopters exceeds the price of non-adopters.*

Having characterized the third party's pricing algorithm under coordinated adoptions, let us conclude with a brief discussion of how such coordination may occur. The canonical device is express communication, either between the firms and the third party or with the third party acting as a hub and engaging in bilateral communication with each firm to achieve a common understanding that all are to adopt the pricing algorithm. A method not involving express communication is for the third party to commit to a supply or fee schedule that necessarily implies joint adoption among firms. One scheme is to offer firms a contract that states the pricing algorithm will be supplied only when some minimum participation threshold is satisfied. Though there is no express invitation to collude, such a scheme would seem to meet the evidentiary standard for proving an antitrust violation as it is equivalent to conveying to each firm: "I will sell my pricing algorithm to you if and only if enough other firms also buy it." That is inviting an agreement and a firm's acceptance of such a contract is acceptance of that invitation. As such a contract would seem to be sure to draw antitrust scrutiny, a more subtle scheme is for the third party to have a fee schedule that makes the fee charged to a firm decreasing in how many firms adopt. While subsuming the first scheme (where there is a reasonable fee when some minimum fraction of firms adopt and an exorbitant fee otherwise), the fee schedule could be set so as not to be so blatantly anticompetitive.[24]

[24] While our analysis does not consider the option of firms colluding on their own, there are variety of advantages from engaging an IT service provider. First and foremost, collusion is more profitable because firms can achieve the first-best monopoly price discrimination solution. There are other advantages which are common to any hub-and-spoke cartel. Garrod, Harrington, and Olczak (2021) describe how the hub - such as an upstream supplier (e.g., an IT service provider) - can aid in coordinating on an outcome, monitoring

3.2 Independent Adoption Decisions

Now suppose firms make independent adoptions. This means that a firm anticipates its adoption decision will not affect the adoption decisions of other firms. Focusing on affine pricing algorithms, the third-party developer's design problem is to choose the design parameters (α, γ) to maximize a firm's WTP for the pricing algorithm:[25]

$$\max_{(\alpha,\gamma)} \quad \int (\alpha + \gamma a - c) \left(a - b(\alpha + \gamma a) + d \left(\theta(\alpha + \gamma a) + (1-\theta)p^{NA}\right)\right) G'(a)da \qquad (8)$$
$$- \int (p^{NA} - c) \left(a - bp^{NA} + d \left(\theta(\alpha + \gamma a) + (1-\theta)p^{NA}\right)\right) G'(a)da.$$

The first term is the expected profit of adopting the pricing algorithm $\alpha + \gamma a$ and the second term subtracted from it is the expected profit from not adopting and pricing at p^{NA} (and not conditioning price on the demand state). Note that the third party's pricing algorithm affects both terms. For example, if (α, γ) is increased so that prices are closer to the monopoly level, that will raise the profit from adopting (first term) but also the profit from not adopting (second term). As adoption decisions are independent (and there are many firms), the adoption rate for the industry is θ in both cases.[26] A property of particular importance is that maximizing the WTP is not equivalent to maximizing an adopter's profit. This means that, in maximizing its profit from selling the pricing algorithm, the third-party developer will not act as a cartel manager. What it implies for the design of the pricing algorithm is described in the next theorem.[27]

Theorem 3 *Under independent adoptions, if a fraction θ of firms adopt then the third-party developer's pricing algorithm is*

$$\alpha = \frac{2bc(b - d\theta) + d\mu(1 - 2\theta)}{2\,(b - d\theta)\,(2b - d)}$$
$$\gamma = \frac{1}{2(b - d\theta)}.$$

for compliance with that outcome, and punishing for non-compliance. When the third party is supplying a pricing algorithm, these factors could manifest as coordinating adoptions through their communications with firms, monitoring firms that they charge the pricing algorithm's recommended price, and punishing by cancelling a firm's subscription to the pricing algorithm (and thereby denying it the ability to price discriminate). One disadvantage from colluding firms involving an IT service provider relates to enforcement. Generally, cartel membership is a closely-guarded secret, but an IT service provider's customer list is discoverable if not public information. Thus, it could be easier for a competition authority to determine who is part of an agreement when that agreement involves an IT service provider than when it only involves the firms themselves.

[25] In (8), p^{NA} refers to the non-adopter's price under independent adoptions and will differ from the expression for p^{NA} in the case of coordinated adoptions.

[26] With literally an infinite number of firms, an individual firm's adoption decision does not affect the adoption rate. I consider that as an approximation for when there are many finite firms and one firm's decision has a small effect on the adoption rate. Assuming this effect is zero simplifies the analysis while maintaining the key driving forces of the model.

[27] In response to a conjecture of Andrew Rhodes, I explored whether there is a piece-wise affine function solution. A 4-parameter class of pricing functions $(\alpha', \alpha'', \eta, \gamma)$ was considered where α' is the intercept for $a < \eta$, α'' is the intercept for $a > \eta$, and γ is the slope parameter. It was proven the affine solution in Theorem 3 is also a solution to that problem with a larger choice set, and numerical analysis did not find any other solutions. Thus, I find support for the (unproven) claim that the affine solution is the unique solution.

Proof. For the objective in (8), consider the FOC with respect to α:[28]

$$\frac{\partial}{\partial\alpha} = \theta\left(\mu - (b-d\theta)(\alpha+\gamma\mu) + d(1-\theta)p^{NA}\right) - \theta\left(\alpha - c\right)(b-d\theta) - \theta\gamma\mu(b-d\theta) - \theta\left(p^{NA}-c\right)d\theta = 0,$$

and solve for α:

$$\alpha = \frac{\mu + bc + dp^{NA} - 2d\theta p^{NA} - 2b\gamma\mu + 2d\theta\gamma\mu}{2b - 2d\theta}. \tag{9}$$

Consider the FOC with respect to γ:

$$\frac{\partial}{\partial\gamma} = \theta\left(\sigma^2 + \mu^2 + bc\mu + dp\mu - 2b\alpha\mu - 2b\sigma^2\gamma - 2b\gamma\mu^2 + 2d\theta\sigma^2\gamma + 2d\theta\gamma\mu^2 - 2dp\theta\mu + 2d\theta\alpha\mu\right) = 0. \tag{10}$$

Use (9) to substitute for α in (10) and then solve for γ:

$$\gamma = \frac{1}{2(b-d\theta)}. \tag{11}$$

Substitute (11) for γ in (9) and solve for α:

$$\alpha = \frac{bc + dp^{NA} - 2d\theta p^{NA}}{2b - 2d\theta}.$$

We then have three unknowns $\left(\alpha, \gamma, p^{NA}\right)$ and three equations:

$$\begin{aligned} \alpha &= \frac{bc + dp^{NA} - 2d\theta p^{NA}}{2b - 2d\theta} \\ \gamma &= \frac{1}{2(b-d\theta)} \\ p^{NA} &= \frac{\mu + bc + d\theta(\alpha+\gamma\mu)}{2b - d(1-\theta)}. \end{aligned}$$

The solution to this system comprises the values for (α, γ) in the statement of the theorem and

$$p^{NA} = \frac{\mu + bc}{2b - d}.$$

The third-party developer's pricing algorithm is then

$$p^I(a) = \frac{2bc(b-d\theta) + d\mu(1-2\theta)}{2\left(b-d\theta\right)\left(2b-d\right)} + \left(\frac{1}{2(b-d\theta)}\right)a,$$

where the superscript I refers to "independent" adoption decisions. ■

A crucial property of the pricing algorithm is that the average price it produces is the same as the Nash equilibrium price without a third-party developer and, consequently, is independent of the adoption rate:

$$E_a\left[p^I\right] = \frac{2bc(b-d\theta) + d\mu(1-2\theta)}{2\left(b-d\theta\right)\left(2b-d\right)} + \left(\frac{1}{2(b-d\theta)}\right)\mu = \frac{\mu + bc}{2b - d}.$$

[28] It is straightforward to show that second-order conditions are satisfied.

Similarly, the non-adopters' price is also $\frac{\mu+bc}{2b-d}$. Hence, when firms make independent adoption decisions, there is no anticompetitive effect with regards to the average price: competitors using the third-party developer's pricing algorithm price the same, on average, as firms in the absence of a third-party developer.[29]

Corollary 4 *Under independent adoptions: i) the average price of adopters and non-adopters are independent of the adoption rate; and ii) on average, the price of adopters equals the price of non-adopters.*

3.3 Comments

Note that an adopter's average price is higher with coordinated adoptions than with independent adoptions:

$$E_a[p^C] = \frac{2b(\mu + (b - d\theta)c) + d^2\theta(1-\theta)c}{2b(2b - d(1+\theta)) + d^2\theta(1-\theta)} > \frac{\mu + bc}{2b - d} = E_a[p^I].$$

However, the sensitivity of price to the demand state is the same:

$$\frac{\partial p^I}{\partial a} = \frac{1}{2(b - d\theta)} = \frac{\partial p^C}{\partial a}.$$

Furthermore, this sensitivity is greater when the adoption rate is higher:

$$\frac{\partial\left(\frac{1}{2(b-d\theta)}\right)}{\partial\theta} = \frac{d}{2\left(b - d\theta\right)^2} > 0.$$

These properties are explained in the next section.

In concluding this section, let us address the matter of the optimal adoption rate. Substituting the pricing algorithm from Theorem 3 into the third-party developer's objective function in (8), one can simplify it to find the WTP of a firm for the pricing algorithm is:[30]

$$\frac{\sigma^2}{4(b - d\theta)}.$$

Note that it is increasing in the adoption rate:

$$\frac{\partial\left(\frac{\sigma^2}{4(b-d\theta)}\right)}{\partial\theta} = \frac{d\sigma^2}{4\left(b - d\theta\right)^2} > 0.$$

[29] Even if firms' beliefs on the adoption rate differ from the actual adoption rate, Theorem 3 implies the average price is independent of the actual adoption rate. Theorem 3 characterizes the pricing algorithm for the third party that maximizes firms' WTP which depends on what firms believe is the adoption rate. Thus, θ in Theorem 3 is the adoption rate believed by firms when considering whether to adopt. Since the associated average price for adopters is $\frac{\mu+bc}{2b-d}$ and the uniform price for non-adopters is also $\frac{\mu+bc}{2b-d}$ then, regardless of the actual adoption rate, the average price of all firms is $\frac{\mu+bc}{2b-d}$.

[30] I would like to thank Jacques Cremer for finding a substantive error in the WTP expression.

Thus, the third-party developer designs the pricing algorithm so that adoptions are strategic complements: the more firms who are expected to adopt the pricing algorithm, the more attractive it becomes to adopt it.

If an adopter incurs no other cost than the fee charged by the third-party developer, the third-party developer will optimally charge a fee equal to the WTP of $\frac{\sigma^2}{4(b-d\theta)}$. Thus, its profit from selling the pricing algorithm to a fraction θ of firms is

$$\frac{\theta\sigma^2}{4(b-d\theta)}$$

which is increasing in the adoption rate:

$$\frac{\partial\left(\frac{\theta\sigma^2}{4(b-d\theta)}\right)}{\partial\theta}=\frac{b\sigma^2}{4\left(b-d\theta\right)^2}>0.$$

Hence, the third-party developer optimally designs and prices the pricing algorithm so that all firms adopt it.

Complete adoption might not occur for the various reasons mentioned earlier: the presence of some firm-specific adoption cost which is heterogeneous across firms, limitations on how fast the third-party developer can install the pricing algorithm, and some firms' resistance to delegating pricing authority (perhaps out of fear of antitrust litigation, as is occurring). Even if not all firms adopt, it will always be optimal for the third-party developer to design the pricing algorithm to maximize the WTP as that will maximize the price it can charge to the firms who do adopt.

4 Comparison of Third-Party Pricing Algorithms under Coordinated and Independent Adoptions

As shown in Section 3, if firms' decisions to adopt a third-party developer's pricing algorithm are independent then adopters' and non-adopters' average prices will not depend on how many firms adopt, and the average price of adopters and non-adopters will be the same. If instead firms' adoption decisions are coordinated then both adopters' and non-adopters' average prices will be higher when more firms adopt, and the average price of adopters will exceed the average price of non-adopters. In this section, the source of these differences is explained and their generality is discussed.[31]

Let us consider three scenarios under which a fraction θ of firms condition price on the demand state. $p^C(a)$ is the third party's pricing algorithm when firms are making coordinated adoption decisions. It is designed to maximize the aggregate profit of adopters or, equivalently, an adopter's profit. $p^I(a)$ is the third party's pricing algorithm when firms are making independent adoption decisions. It is designed to maximize the willingness-to-pay of an adopter which, with independent adoption decisions, is the profit from adopting minus the profit from not adopting while holding fixed that a fraction θ of rival firms are

[31] The underlying forces are the same as in Harrington (2022b) for the duopoly case though a different approach is used to explain them.

adopting. $p^F(a)$ is the pricing algorithm when a firm designs the algorithm itself and firms do so independently. Its closed-form solution is in (1). It is designed to maximize the firm's profit, given a fraction θ of rival firms are using $p^F(a)$, and is the pricing algorithm that would occur under competition if the firms could design it on their own.

As summarized in Table 1, these pricing algorithms differ in terms of the identity of the designer and the relationship between firms' adoption decisions

Table 1

	Firms' adoption decisions	
Designer of pricing algorithm	Independent	Coordinated
Third party	p^I	p^C
Firm	p^F	-

The three pricing algorithms are depicted in Figure 1.

Figure 1

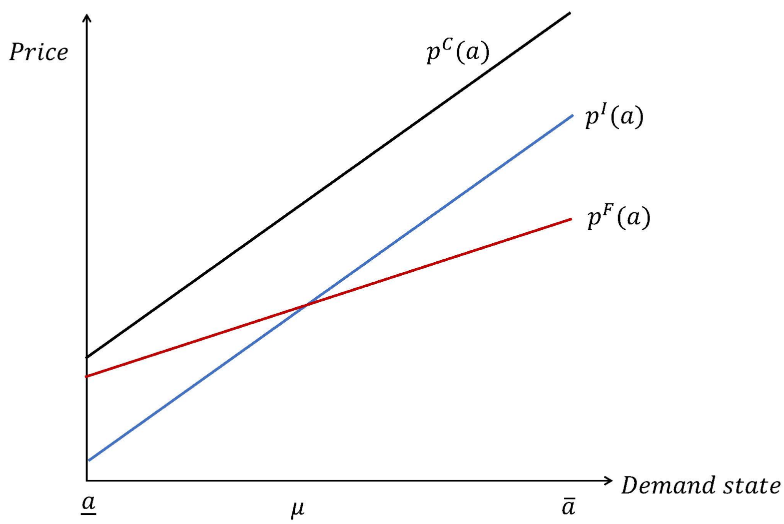

Let us compare these pricing algorithms. To begin, p^I has the same average price as p^F which is less than the average price of p^C,

$$E_a\left[p^C\right] > E_a\left[p^I\right] = \frac{\mu + bc}{2b - d} = E_a\left[p^F\right],$$

but p^I is more sensitive to the demand state than p^F,

$$\frac{\partial p^I}{\partial a} = \frac{1}{2(b - d\theta)} > \frac{1}{2b - d\theta} = \frac{\partial p^F}{\partial a}.$$

Next note that the third party's pricing algorithm has price respond to the demand state in the same way whether adoption decisions are independent or coordinated:

$$\frac{\partial p^I}{\partial a} = \frac{1}{2(b - d\theta)} = \frac{\partial p^C}{\partial a}.$$

As depicted in Figure 1, p^I is a parallel shift down of p^C.

Let me explain why the third party designs p^I to have these properties when firms' adoption decisions are independent. The third party chooses the design that maximizes firms' WTP which is the profit from adopting minus the profit from not adopting, given a fraction θ of firms are anticipated to adopt. Towards finding the best design, consider the thought experiment of creating a supracompetitive pricing algorithm by having price be higher than the competitive pricing algorithm p^F at each demand state, so it is $p^F(a)+\varepsilon(a)$ where $\varepsilon(a)$ is positive (and small enough that price does not exceed the monopoly price). Compared to $p^F(a)$, $p^F(a)+\varepsilon(a)$ raises profit both from adopting and not adopting because rival firms - specifically, those who adopt - will be setting higher prices. By adopting itself, the firm sets those same high prices but, and this is the key point, by not adopting the firm would set an average (more specifically, uniform) price that is the best response to adopters' high average price. That best response is below the average price of the pricing algorithm because the pricing algorithm is setting price above the competitive pricing algorithm p^F. Summing up, a pricing algorithm with a higher average price than p^F will raise both the profit from adopting and not adopting - because the adopting rivals of a firm now have a higher average price - but it raises the profit from not adopting *more* because a non-adopter sets the optimal average price in response to adopters' higher average price. Consequently, $p^F(a)+\varepsilon(a)$ lowers the WTP compared to $p^F(a)$. The general takeaway is that a third party will not want to raise the average price above the competitive pricing algorithm, which is why $E_a\left[p^I\right]=E_a\left[p^F\right]$.

Under the assumption of independent adoption decisions, a third party's optimal pricing algorithm will then set average price at the competitive level. In contrast, if adoption decisions are coordinated then the third party will maximize the profit from adopting - rather than the difference between adopting and not adopting - and that necessarily means a higher average price. With coordinated adoptions, the situation is equivalent to a partial cartel which encompasses a fraction θ of firms where the third party is acting like a cartel manager when it selects p^C. The higher is θ, the more inclusive is the cartel which implies p^C is higher and, therefore, average price is higher (as there are fewer non-cartel members to undercut the high average price of adopters). Thus, when adoptions are coordinated, p^C has average price increasing in the adoption rate θ, while with independent adoptions, the pricing algorithm's average price is at the competitive level and thus does not depend on θ.

For when firms' adoption decisions are independent, I have just explained why the third party will not have the pricing algorithm set a supracompetitive average price like a cartel manager would. However, it *will* have price respond to the demand state as a cartel manager would. p^I is designed to have price respond to the demand state in order to maximize an adopter's profit, which means making price more sensitive to the demand state than is the case under competition.[32] Increasing the sensitivity of price to the demand state raises the profit from adopting and, most importantly, does *not* affect the profit from not adopting (as long as the average price is unchanged). The reason is that a non-adopter's profit

[32] The reason that a cartel manager would make price more sensitive to the demand state is as follows. When the demand state is stronger, a firm's profit-maximizing price is higher. Given firms' prices are strategic complements, a firm optimally raises price more when other firms also raise their prices. A cartel manager will take that effect into account so firms' prices will rise more in response to stronger demand compared to when prices are set independently.

depends only on the average price of adopters, not how adopters respond to the demand state. Though in response to a high demand state, a non-adopter would like to set a high price that undercuts the high price of adopters, it is incapable of doing so because only by having the third party's pricing algorithm can it condition price on the demand state. In the case of independent adoptions, the secret to the third party maximizing the WTP (= profit from adopting minus the profit from not adopting) is to have price respond to the demand state in a way that maximizes the profit from adopting, while not raising the average price so the profit from not adopting does not rise.

Corollaries 2 and 4 follow from these properties of the pricing algorithms. Under independent adoption decisions, the third party maximizes the WTP by keeping the average price fixed compared to when there is no third party - in order to avoid raising the profit from not adopting - which implies an adopter's average price is independent of the adoption rate. Under coordinated adoption decisions, a third party maximizes adopters' profit and that implies a higher average price. The more firms that adopt, the stronger is that effect as the "adopters' cartel" is more inclusive so there are fewer non-adopting firms to undercut the adopters' supracompetitive average price. Hence, average price is increasing in the adoption rate when firms' adoption decisions are coordinated.

Though the model is highly structured with linear demand and symmetric firms, the forces determining the properties of the pricing algorithms seem quite general. This is clearly the case with coordinated adoptions as the third party is acting as a cartel manager - designing the pricing algorithm to maximize adopters' profits - and that will necessarily result in a higher average price when the cartel is more inclusive (i.e., the adoption rate is higher). Less immediate is the robustness of the third party's pricing algorithm under independent adoptions: average price is at the competitive level and is independent of the adoption rate. It is certainly possible that average price may not be exactly at the competitive level with non-linear demand functions and asymmetric firms. However, there is no reason to expect average price to systematically depend on the adoption rate. As explained above, the properties of the pricing algorithm under independent adoptions are based on the ability of a non-adopter to optimize with respect to the average price of adopters but not to how adopters' respond to the demand state. Those features are general and not specific to the particular assumptions of the model.

A second issue related to robustness is that we have only examined when either all adoption decisions are independent or all adoption decisions are coordinated. Suppose a fraction θ of firms adopt the pricing algorithm but only a fraction ω of them coordinate their adoption decisions. It seems reasonable to expect the third-party developer's optimal pricing algorithm to be some mixture of the optimal pricing algorithm when all firms make coordinated adoption decisions, p^C, and the optimal pricing algorithm when all firms make independent decisions, p^I. If that is true then the properties in Corollary 2 would still hold, it would just be the effect is not as strong. I would then expect the average price of adopters to be increasing in the adoption rate as long as at least some of the adoptions decisions are coordinated. While this evidence would tell us there is an agreement between the third-party developer and some adopters, it would not tell us whether the agreement encompasses all adopters or, when it does not not encompass all adopters, which adopters are part of the agreement. Nevertheless, with evidence of an agreement, other evidence could be used to determine which adopters are party to the agreement.

5 An Empirical Test for an Unlawful Agreement

The preceding analysis provides us with three candidate tests for providing evidence of an agreement between a third-party developer who designed a pricing algorithm and the firms who adopted it.

- Test #1: The average price of adopters is increasing in the adoption rate.
- Test #2: The average price of non-adopters is increasing in the adoption rate.
- Test #3: The average price of adopters exceeds the average price of non-adopters.

Test #1 is at the core of the incentives of a third-party developer when it comes to the design of the pricing algorithm. When adopters have an agreement to adopt the pricing algorithm, the third-party developer designs the pricing algorithm to maximize adopters' profit which implies supracompetitive prices. This supracompetitive markup is greater when more firms are part of the agreement which implies the average price of adopters is increasing in the adoption rate. This relationship is not predicted when firms make independent adoption decisions in which case adopters' average price is independent of the adoption rate.[33]

Test #2 is an implication of the conduct underlying Test #1. Given firms with an agreement to adopt the pricing algorithm set supracompetitive prices, non-adopters will optimally respond with higher prices (though, on average, undercutting adopters' prices). As the adopters' supracompetitive markup is higher with a higher adoption rate, non-adopters' supracompetitive markup is then also higher with a high adoption rate. A potential concern with this test is if non-adopters take time to learn the optimal response to adopters' prices, so adopters' average price is higher but non-adopters' average price has not yet responded to that higher price. Thus, it is possible there is an agreement which is picked up by Test #1 but not by Test #2. For this reason, Test #1 is superior to Test #2.

A possible concern with Test #3 is that it may not be robust to firm asymmetries. For example, suppose firms have different marginal costs and the pricing algorithm allows a firm to program in its marginal cost. While all adopters set supracompetitive prices, an adopter with a lower cost sets a lower price than an adopter with a higher cost. Further suppose, for whatever reason, a firm is more likely to adopt the pricing algorithm if its cost is lower. Even if firms have an agreement and the pricing algorithm is designed as the theory predicts, adopters' average price may not be higher than non-adopters' average price because

[33] It has been assumed that non-adopting firms are competing which seems reasonable since it isn't clear why a third party would be involved in an agreement that did not require participating firms to use its pricing algorithm. Regardless, consider some or all non-adopting firms being part of the agreement and let us distinguish between the adoption rate (i.e., the fraction of firms using the pricing algorithm) and the agreement rate (i.e., the fraction of firms who are colluding). The analysis of this paper assumes those two rates are the same but Test #1 should be effective as long as they move together. The only situation in which the test fails is if the agreement rate is independent of the adoption rate. For example, if all firms are part of the agreement then the third party will optimally design the pricing algorithm to be the monopoly pricing function and the non-adopting firms will set the monopoly uniform price. Consequently, the average price will equal the average monopoly price which does not depend on the adoption rate. That special case aside, the proposed test should work even when non-adopting firms are part of the agreement as long as the more firms that adopt the pricing algorithm, the more firms are part of the agreement.

the marginal cost of adopters is lower. In that case, Test #3 would not be satisfied even though there is an agreement. In conclusion, Test #1 is the most general and robust of the three tests and is the one recommended.[34]

When there is not direct evidence of communication supporting an agreement, evidence is circumstantial and obtaining a conviction often requires both economic and non-economic evidence. The test proposed here offers an approach to delivering economic evidence. Based on the European Court of Justice's recent *Eturas* decision, let me describe how this evidence along with certain non-economic evidence could be used to prove a violation of TFEU 101(1). In *Eturas*, travel agencies (firms) offered their services on the web site of Eturas (third party) who sent them all a message announcing it had imposed a cap on the discounts that they could offer (though it was possible for an agency to circumvent the cap). The ECJ concluded that "those economic operators may - if they were aware of that message - be presumed to have participated in a concerted practice ... unless they publicly distanced themselves from that practice, reported it to the administrative authorities or adduce other evidence to rebut that presumption, such as evidence of the systematic application of a discount exceeding the cap in question."[35] In that case, the third party imposed the practice and the ECJ viewed an informed firm as engaging in that practice unless they took clear action against it. In the case of a third party offering a supracompetitive pricing algorithm, an adopting firm is more culpable because the pricing algorithm is not imposed on them; a firm had to take explicit action to adopt it - entering into a contractual arrangement with the third party - and, in fact, paid for it. Thus, if firms knew the third party's pricing algorithm was made available to all firms and it was designed to reduce competition (as reflected in a higher average price when more firms adopt) then, with their adoption of the pricing algorithm, they can "be presumed to have participated in a concerted practice".

An empirical implementation of Test #1 would involve regressing adopters' prices on the adoption rate. The data required are: 1) whether (and possibly when) a firm adopted the third-party developer's pricing algorithm; 2) prices of the firms who adopted; and 3) other variables that could affect observed prices. At a minimum, a competition authority will need to know which firms adopted and their prices. If not publicly available, it is straightforward to obtain this information from the third party and adopting firms through the legal discovery process. The data set will need to have variation in the adoption rate, which could be across time or markets. Inter-market variation could come from a third-party developer offering

[34] While these tests are derived for when there is a continuum of firms, I expect them to hold with a finite number of firms. The test is based on two properties: 1) under coordinated adoptions, the average price is increasing in the fraction of adopting firms; and 2) under independent adoptions, the average price is constant with respect to the fraction of adopting firms. (2) is an implication of the property that the average price equals the static Nash equilibrium price without a price discriminating third party. That property is not only true for an infinite number of firms but also for the case of two firms, as shown in Harrington (2022b), and I conjecture it holds for any finite number of firms. Hence, it is generally expected that the average price is independent of the adoption rate under independent adoptions. (1) follows from the optimal cartel price being increasing in the fraction of firms that are cartel members which certainly holds when the number of firms is finite. Thought the tests are not applicable to the duopoly case, that is only because it is not possible to have variation in the adoption rate under coordinated adoptions; all firms must adopt. Contrasting that with when there is a triopoly, the adoption rate can be increased from 2/3 to 1 and it is predicted the average price rises (remains the same) when adoptions are coordinated (independent).

[35] Judgment of the Court (Fifth Chamber) of 21 January 2016 - "Eturas" UAB and Others v Lietuvos Respublikos konkurencijos taryba, para. 51.

its pricing algorithm in different geographic markets. One can use the variation in adoption rates across markets to estimate the relationship between adopters' average price and the adoption rate. Alternatively, there could be a single market with time-series variation in the adoption rate. While the theory is based on simultaneous adoption decisions, it applies as well to sequential adoption decisions as long as the third-party developer adjusts the pricing algorithm to changes in the adoption rate.

When applying this test, the data set should be a representative sample of the demand states used in designing the pricing algorithm. If, for example, all the demand states in the sample are above the mean demand state used when constructing the pricing algorithm then average price will be increasing in the adoption rate even with independent adoptions (because price is more sensitive to the demand state). However, as long as the mean of the demand states in the data set approximates the mean upon which the pricing algorithm is based then the average price will be increasing in the adoption rate only under coordinated adoptions. There are at least two scenarios where this condition is likely to hold. First, if the distribution of demand is relatively stable and the data set covers a sufficiently long period. For example, that condition is likely to apply to gasoline demand where demand variability typically occurs at the daily or weekly level so a data set comprising, say, a year would be a representative sample of the stationary distribution of demand states and thus of the distribution for which the pricing algorithm is likely to be based. Second, if the demand state varies across sub-markets and the distribution of those sub-markets is stable over time. For example, under normal circumstances, that condition is likely to hold with respect to geographic sub-markets for apartments within a metropolitan area. One way to test this assumption is to compare the distribution of demand states used to train the pricing algorithm with the distribution of demand states used to test the hypothesis that the average price is increasing in the adoption rate. Of course, the pricing algorithm is continuously being trained but one should be able to suitably partition the data set to conduct such a test.

A concern with this empirical approach is the possible endogeneity of the adoption rate. If firms are heterogeneous and that heterogeneity affects both their adoption decisions and the prices set by the pricing algorithm then the estimated coefficient on the adoption rate will be biased. For example, suppose firms differ in marginal cost and firms with higher marginal cost tend to adopt sooner. Even if there is an agreement, it is possible the average price of adopters is higher when the adoption rate is lower because it is the high marginal cost firms that adopt early which, by construction, is when the adoption rate is low. While such bias could occur, variation in firms' adoption decisions need not imply bias. For example, suppose variation in adoption decisions is due to some firm-specific cost to installing a pricing algorithm which varies across firms. Firms with a lower adoption cost are more likely to adopt (or adopt sooner) and that could be the source of variation in adoption rates. As the firm-specific adoption cost does not affect the prices produced by the pricing algorithm, estimates would not be biased. Or suppose the variation in adoption rates is due to the constrained capacity of the third-party developer to sell the pricing algorithm, perhaps due to limited personnel to market and install the program. That could introduce exogenous variation in adoption rates and again allow for unbiased estimates.

As proof of concept, the empirical test proposed here was conducted in Assad et al (2023) for retail gasoline markets in Germany. In doing so, they addressed the issue of the possible endogeneity of the adoption rate and provided a method for inferring when a firm adopted the pricing algorithm even when that information is confidential. The study found evidence supporting the hypothesis in Test #1: the average price of adopters is higher when more firms adopt. It is worth emphasizing that the empirical analysis was conducted using publicly-accessible data. Thus, this test and the empirical approach for implementing it could be used as an initial screen to identify a market of concern before there is any other evidence.

Addressing the Lucas critique, let me explain why the test is likely to be effective even it is known to be used as economic evidence of an agreement. To avoid a positive test result, a third party would have to design the pricing algorithm so the average price is not increasing in the adoption rate. The third party is then not designing the pricing algorithm to maximize firms' WTP which ultimately means having to charge a lower fee and that will lower the third party's profit. Hence, there is a cost to the third party reducing the test's efficacy. In optimally modifying their conduct in response to the test being used by competition authorities and courts, one would expect the third party to trade-off the lower chance of detection and conviction against the lower profit realized with a less optimal pricing algorithm. As it seems unlikely the optimal response is a corner solution which makes the test entirely ineffective, I would expect the test to still have power. The general principle at work is that, even under the Lucas critique, a test has power as long as it is costly for firms to reduce its effectiveness.[36]

In concluding, a caveat is offered. While the test for an unlawful agreement has been derived for a well-accepted model of imperfect competition and a plausible source of efficiency delivered by a third party, it is possible the test may not be supported by other specifications. Further research is needed to assess the universality of the test and to derive other methods for identifying when firms' adoptions decisions are coordinated rather than independent.

6 An Algorithm Audit Test for an Unlawful Agreement

The empirical approach described in the previous section is visually summarized in Figure 2. The distribution of demand states along with the adoption rate (and other variables such as cost) are used by the third party to design the pricing algorithm. Over time, the third party feeds demand states into the pricing algorithm to produce prices. We then have a sample distribution of demand states and an associated distribution of prices from which average price is calculated. With variation in the adoption rate, the hypothesis that the average

[36] This point is made as well in the context of cartel screening; see Section 6.3 of Harrington (2008) where it is also discussed how, in practice, many colluding firms fail to engage in conduct that would *costlessly* make detection less likely. Just because the Lucas critique could undermine a test, it is still an empirical question whether the parties are smart enough to do so.

price is increasing in the adoption rate can be tested.

Figure 2

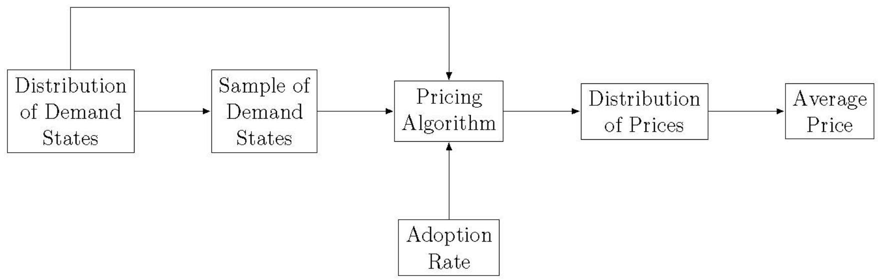

An alternative approach to assessing an algorithm's properties is to conduct an audit which involves "repeatedly and systematically querying an algorithm with inputs and observing the corresponding outputs in order to draw inferences about its opaque inner workings."[37] In the current situation, this means inputting demand states and then recording the resulting prices. Referring to Figure 2, an audit approach only requires the pricing algorithm and not data on demand states and prices.

For an audit approach to provide evidence of an unlawful agreement, it is necessary to identify some property of the output that distinguishes a pricing algorithm designed when adoptions are coordinated (so it maximizes adopting firms' profits) and when adoptions are independent (so it maximizes a firm's willingness-to-pay to adopt the pricing algorithm). Towards that end, note that, under the assumption of independent adoptions, price is increasing in the adoption rate if and only if the demand state is above the mean demand state:

$$\frac{\partial p^I(a,\theta)}{\partial \theta} = \frac{d(a-\mu)}{2\left(b-d\theta\right)^2} > 0 \text{ if and only if } a > \mu.$$

As shown in Figure 3,[38] a rise in the adoption rate causes the pricing algorithm to rotate around $a = \mu$; it becomes more responsive to the demand state but the average price is unchanged. Thus, price rises in response to a higher adoption rate only for a strict subset of demand states. For example, if the distribution on demand states is symmetric around the mean then the mean equals the median and price is increasing in the adoption rate for half

[37] Metaxa et al (2021), p. 288.

[38] The parameters used in Figure 3 are: $b = 1, d = 0.75, c = 10, \mu = 100$.

of the demand states.

Figure 3

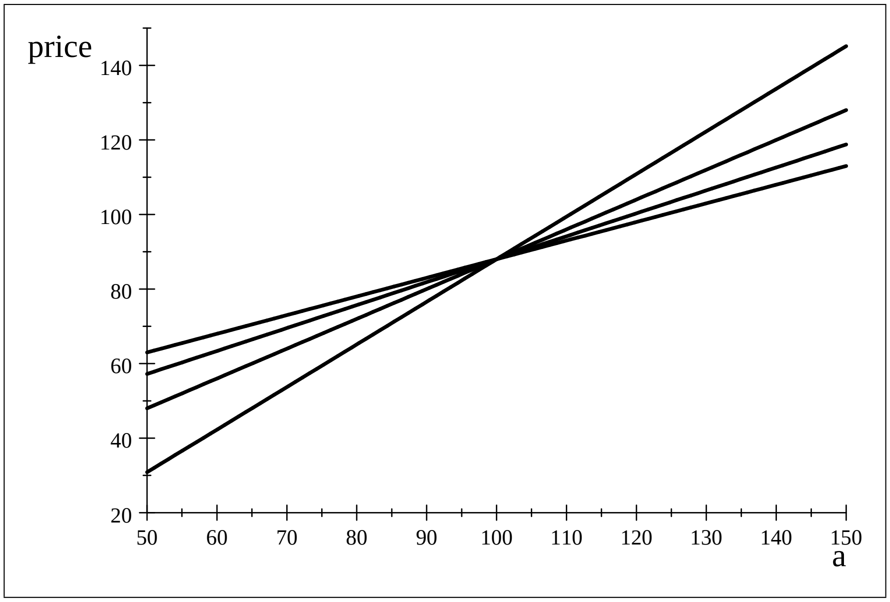

Pricing algorithm rotates with a steeper slope as θ is increased from 0 to 0.25 to 0.50 to 0.75

Next consider the pricing algorithm under coordinated adoptions, $p^C(a,\theta)$. It is straightforward to show:

$$p^C(a,\theta) = p^I(a,\theta) + \omega(\theta)$$

where

$$\omega(\theta) \equiv \frac{d(2b-d+d\theta)(\mu-bc+cd)\theta}{(2b-d)(2b(2b-d(1+\theta))+d^2\theta(1-\theta))}.$$

Given

$$\omega'(\theta) = \frac{2bd(\mu-(b-d)c)(2b-d+2d\theta)}{(4b^2-2bd\theta-2bd+d^2\theta(1-\theta))^2} > 0$$

then, for some $x > 0$,

$$\frac{\partial p^C(a,\theta)}{\partial\theta} = \frac{\partial p^I(a,\theta)}{\partial\theta} + \omega'(\theta) > 0 \ \forall a > \mu - x.$$

Thus, the set of demand states for which price is increasing in the adoption rate under coordinated adoptions is a superset of those demand states for which that property holds under independent adoptions. For the same parameter specification used for Figure 3, Figure 4 shows a higher adoption rate causes the pricing algorithm to shift up (with a steeper slope) and price is higher for all demand states.[39]

Figure 4

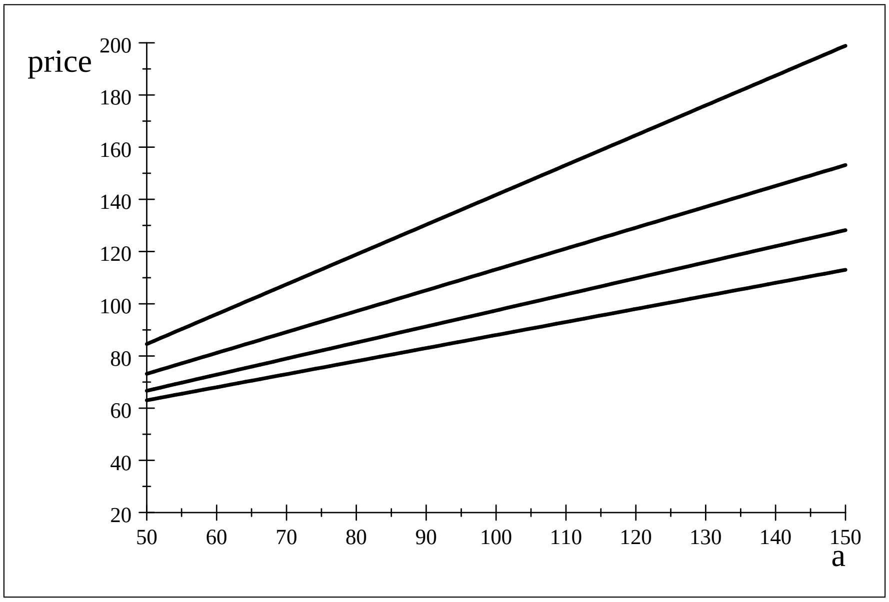

Pricing algorithm shifts up as θ is increased from 0 to 0.25 to 0.50 to 0.75

The preceding analysis is the basis for the following property.[40]

Property: If a higher adoption rate causes price to rise for a very large share of demand states then the pricing algorithm is likely to have been designed as part of an unlawful agreement.

An audit can test for this property by inputting a wide range of demand states, recording the outputted prices, and repeating this exercise for different adoption rates. For each pair of adoption rates, one identifies the set of demand states for which price is higher with the higher adoption rate. If that set comprises a very large share of the demand states used in the audit then it is evidence supportive of there being an unlawful agreement. There are two possible sources of variation in the adoption rate. When the adoption rate is a parameter in the pricing algorithm, it can be varied as part of the audit. But suppose the adoption rate is not an explicit parameter and instead the third party recalculates the pricing algorithm as the adoption rate changes. In that case, one would need pricing algorithms over time along with the adoption rate at the time the pricing algorithm was designed. One could then compare two pricing algorithms with different adoption rates and measure the set of demand states for which the pricing algorithm associated with a higher adoption rate produces a higher price.

[39] If we expand the set of demand states beyond $[50, 150]$, one finds price is decreasing in the adoption rate for sufficiently low demand states under coordinated adoptions. However, the general property is robust: price is increasing in the adoption rate for a very large set of demand states only when adoptions are coordinated.

[40] Other simulations provide similar results to those shown in Figure 4.

7 Concluding Remarks

This study has shown how the properties of a third-party developer's pricing algorithm depend on the manner in which adoption decisions are made by competitors in a market. If firms are independently deciding whether to adopt the pricing algorithm, the average price of adopting firms is predicted to be independent of the number of firms that adopt. If instead firms are coordinating their adoption decisions then the average price of adopting firms is predicted to be increasing in the number of firms that adopt. These findings provide a test for determining whether there is an unlawful agreement between a third-party developer and adopting firms.

The economic evidence provided by this test could be used in conjunction with other evidence to prove a violation of Article 101 of the Treaty on the Functioning of the European Union, Section 1 of the Sherman Act, or some other jurisdiction's competition law pertaining to collusion. In the United States, a case can be dismissed on the grounds that the evidence in a complaint does not make it sufficiently plausible that firms have an unlawful agreement.[41] Plaintiffs are then required to provide sufficient evidence prior to being permitted to go to discovery where they could obtain corporate documents and other critical information. The economic evidence delivered by the test developed here could help surmount that plausibility hurdle. Finally, this test could be used to screen possible cases for investigation and thereby avoid wrongly prosecuting data analytics companies who are not engaging in unlawful activity. Evidence that prices are higher when more firms adopt a third party's pricing algorithm is supportive of coordinated conduct and thus an investigation may be warranted.

[41] "[O]nly a complaint that states a plausible claim for relief survives a motion to dismiss." Ashcroft v. Iqbal, 556 U.S. 662, 679 (2009). That decision built on Bell Atlantic Corp. v. Twombly, 550 U.S. 544 (2007).

8 Appendix

Suppose a fraction $\theta \in [0,1]$ of firms can condition price on the demand state a and a fraction $1-\theta$ of firms cannot and thus set a uniform price which is denoted p'. $p^F(a)$ denotes the equilibrium pricing function for those firms who can condition price on the demand state and is defined by:

$$p^F(a) = \arg\max_p (p-c)(a - bp + d(\theta p^F(a) + (1-\theta)p')).$$

The first-order condition is

$$a - 2bp^F(a) + d(\theta p^F(a) + (1-\theta)p') + bc = 0,$$

which can be solved for:

$$p^F(a) = \frac{a + d(1-\theta)p' + bc}{2b - d\theta}. \tag{12}$$

The equilibrium price for a firm that sets a uniform price is defined by:

$$\begin{aligned} p' &= \arg\max_p \int (p-c)(a - bp + d(\theta p^F(a) + (1-\theta)p'))G'(a)da \\ &= \arg\max_p (p-c)(\mu - bp + d(\theta E_a[p^F(a)] + (1-\theta)p')). \end{aligned}$$

The first-order condition is

$$\mu - 2bp' + d(\theta E_a[p^F(a)] + (1-\theta)p')) + bc = 0,$$

which can be solved for:

$$p' = \frac{\mu + d\theta E_a[p^F(a)] + bc}{2b - d(1-\theta)}. \tag{13}$$

Substituting (12) into (13) gives us the expression for p' :

$$p' = \frac{\mu + d\theta \left(\frac{\mu + d(1-\theta)p' + bc}{2b - d\theta}\right) + bc}{2b - d(1-\theta)} \Leftrightarrow p' = \frac{\mu + bc}{2b - d}. \tag{14}$$

Using (14) in (12), we have the expression for $p^F(a)$:

$$p^F(a) = \frac{a + d(1-\theta)\left(\frac{\mu + bc}{2b-d}\right) + bc}{2b - d\theta} = \frac{d(1-\theta)\mu + (2b - d\theta)bc}{(2b - d\theta)(2b - d)} + \left(\frac{1}{2b - d\theta}\right) a.$$

References

[1] Abada, Ibrahim, Joseph E. Harrington, Jr., Xavier Lambin, and Janusz M. Meylahn, "Algorithmic Collusion: Where Are We and Where Should We Be Going?," working paper, July 2024.

[2] Abada, Ibrahim and Xavier Lambin, "Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided?," *Management Science*, 69 (2023), 5042-5065.

[3] Abada, Ibrahim, Xavier Lambin, and Nikolay Tchakarov, "Collusion by Mistake: Does Algorithmic Sophistication Drive Supra-competitive Profit?," working paper, December 2022.

[4] Asil, Aslihan and Thomas G. Wollmann, "Can Machines Commit Crimes under US Antitrust Laws?," *The University of Chicago Business Law Review*, 3 (2023), 1-38.

[5] Asker, John, Chaim Fershtman, and Ariel Pakes, "The Impact of Artificial Intelligence Design on Pricing," *Journal of Economics & Management Strategy*, 2023, 1-29.

[6] Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," working paper, February 2023 (*Journal of Political Economy*, forthcoming).

[7] Banchio, Martino and Giacomo Mantegazza, "Adaptive Algorithms and Collusion via Coupling," working paper, November 2022.

[8] Beneke, Francisco and Mark-Oliver Mackenrodt, "Remedies for Algorithmic Tacit Collusion," *Journal of Antitrust Enforcement*, 9 (2021), 152–176.

[9] van den Boer, Arnoud, Janusz M. Meylahn, and Maarten Pieter Schinkel, "Artificial Collusion: Examining Supracompetitive Pricing by Q-Learning Algorithms," working paper, December 2022.

[10] Brown, Zach Y. and Alexander MacKay, "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics*, 15 (2023), 109-156.

[11] Calder-Wang, Sophie and Gi Heung Kim, "Coordinated vs Efficient Prices: The Impact of Algorithmic Pricing on Multifamily Rental Markets," working paper, March 2023.

[12] Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, "Artificial Intelligence, Algorithmic Pricing and Collusion," *American Economic Review*, 110 (2020), 3267-3297.

[13] Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, "Algorithmic Collusion with Imperfect Monitoring," *International Journal of Industrial Organization*, 79 (2021), 102712.

[14] Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, Joseph E. Harrington, Jr., and Sergio Pastorello, "Protecting Consumers from Collusive Prices Due to AI," *Science*, 370 (27 November 2020), 1040-1042.

[15] Cartea, Álvaro, Patrick Chang, José Penalva, and Harrison Waldon, "The Algorithmic Learning Equations: Evolving Strategies in Dynamic Games," working paper, October 2022.

[16] Cooper, William L., Tito Homen-de-Mello, and Anton J. Kleywegt, "Learning and Pricing with Models That Do Not Explicitly Incorporate Competition," *Operations Research*, 63 (2015), 86-103.

[17] Francesco Decarolis, Maris Goldmanis, and Antonio Penta, "Marketing Agencies and Collusive Bidding in Online Ad Auctions," *Management Science*, 66 (2020), 4359-4619.

[18] Deng, Ai, "What Do We Know About Algorithmic Tacit Collusion?," *Antitrust*, 33 (2018), 88-95.

[19] Eckert, Andrew, "Empirical Studies of Gasoline Retailing: A Guide to the Literature," *Journal of Economic Surveys*, 27 (2013), 140-166.

[20] Epivent, Andréa and Xavier Lambin, "On Algorithmic Collusion and Reward- Punishment," *Economics Letters*, 237 (2024), 111661.

[21] Eschenbaum, Nicolas, Filip Mellgren, and Philipp Zahn, "Robust Algorithmic Collusion," working paper, December 2021.

[22] Ezrachi, Ariel and Maurice E. Stucke, "Artificial Intelligence and Collusion: When Computers Inhibit Competition," *University of Illinois Law Review*, 2017 (2017), 1775-1810.

[23] Ezrachi, Ariel and Maurice E. Stucke, "Algorithmic Tacit Collusion," in *Research Handbook on Cartels,* P. Whelan, ed., Edward Elgar, 187-212.

[24] Gal, Michal S., "Algorithms as Illegal Agreements," *Berkeley Technology Law Journal*, 34 (2019), 67-118.

[25] Gal, Michal S., "Limiting Algorithmic Collusion," *Berkeley Technology Law Journal*, 38 (2023), 173-229.

[26] Garrod, Luke, Joseph E. Harrington, Jr., and Matthew Olczak, *Hub-and-Spoke Cartels: Why They Form, How They Operate, and How to Prosecute Them*, The MIT Press, 2021.

[27] German Monopolies Commission, XXII. Biennial Report, Chapter on "Algorithms and Collusion," 2018.

[28] Haan, Katherine, "How Businesses are Using Artificial Intelligence in 2024," *Forbes*, April 24, 2023 <www.forbes.com/advisor/business/software/ai-in-business/> (downloaded January 17, 2024)

[29] Hansen, Karsten, Kanishka Misra, and Mallesh Pai, "Algorithmic Collusion: Supra-Competitive Prices via Independent Algorithms," *Marketing Science*, 40 (2021), 1-12.

[30] Harrington, Joseph E., Jr., "Detecting Cartels," in *Handbook of Antitrust Economics*, P. Buccirossi, ed., The MIT Press, 2008.

[31] Harrington, Joseph E., Jr., "Developing Competition Law for Collusion by Autonomous Artificial Agents," *Journal of Competition Law and Economics*, 14 (2018), 331-363.

[32] Harrington, Joseph E., Jr., "Collusion in Plain Sight: Firms' Use of Public Announcements to Restrain Competition," *Antitrust Law Journal*, 84 (2022a), 521-563.

[33] Harrington, Joseph E., Jr., "The Effect of Outsourcing Pricing Algorithms on Market Competition," *Management Science*, 68 (2022b), 6889-6906.

[34] Inderst, Roman and Stefan Thomas, "Algorithms and Antitrust: A Framework with Special Emphasis on Coordinated Pricing," working paper, May 2024.

[35] Johnson, Paul A., "Should We Be Concerned That Data and Algorithms Will Soften Competition?," *CPI Antitrust Chronicle*, 2 (2017), 10-15.

[36] Kastius, Alexander and Rainer Schlosser, "Dynamic Pricing under Competition using Reinforcement Learning," *Journal of Revenue and Pricing Management*, (2021), 1-14.

[37] Kimbrough, Steven O. and Frederic H. Murphy, "Learning to Collude: Tacitly on Production Levels by Oligopolistic Agents," *Computational Economics*, 33 (2009), 47-78.

[38] Klein, Timo, "Autonomous Algorithmic Collusion: Q-learning under Sequential Pricing," *RAND Journal of Economics*, 52 (2021), 538-558.

[39] Leisten, Matthew, "Algorithmic Collusion, with Humans," working paper, December 2022.

[40] Metaxa, Danaë, Joon Sung Park, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, and Christian Sandvig, "Auditing Algorithms," *Foundations and Trends in Human-Computer Interaction,* 14 (2021), 272-344.

[41] Meylahn, Janusz M. and Arnoud V. den Boer, "Learning to Collude in a Pricing Duopoly," *Manufacturing Service Operations Management*, 2 (2022), 1523-1614.

[42] Mehra, Salil K., "Antitrust and the Robo-Seller: Competition in the Time of Algorithms," *Minnesota Law Review*, 100 (2016), 1323-1375.

[43] Musolff, Leon, "Algorithmic Pricing Facilitates Tacit Collusion: Evidence from E-Commerce," working paper, September 2021.

[44] OECD, "Algorithms and Collusion - Background Note by the Secretariat," DAF/COMP(2017)4, 16 May 2017.

[45] Oxera, "When Algorithms Set Prices: Winners and Losers," Discussion Paper, June 2017.

[46] Possnig, Clemens, "Reinforcement Learning and Collusion," working paper, January 2023.

[47] Sanchez-Cartas, J. Manuel and Evangelos Katsamakas, "Artificial Intelligence, Algorithmic Competition and Market Structures," *IEEE Access*, 10 (2022), 10575-10584.

[48] Schwalbe, Ulrich, "Algorithms, Machine Learning, and Collusion," *Journal of Competition Law and Economics*, 14 (2019), 568-607.

[49] United Kingdom Competition & Markets Authority, "Pricing Algorithms: Economic Working Paper on the Use of Algorithms to Facilitate Collusion and Personalised Pricing," October 2018.

[50] Vives, Xavier, *Oligopoly Pricing: Old Ideas and New Tools*, The MIT Press, 1999.

[51] Waltman, Ludo and Uzay Kaymak, "Q-learning Agents in a Cournot Oligopoly Model," *Journal of Economic Dynamics and Control*, 32 (2008), 3275-3293.