Third Party Pricing Algorithms and the Intensity of Competition*

Joseph E. Harrington, Jr.
Department of Business Economics & Public Policy
The Wharton School
University of Pennsylvania
harrij@wharton.upenn.edu

3 November 2020

Abstract

This paper explores the properties of a pricing algorithm when it is designed by a third party. The setting is one where the pricing algorithm allows a firm to condition price on high-frequency demand information. I find that third party development of a pricing algorithm has an anticompetitive effect even when only one firm in a market adopts it, and the anticompetitive effect is greater when more firms adopt it.

*I appreciate the comments of Ai Deng, Timo Klein, Jeanine Miklós-Thal, and participants in the Notre Dame Theory Seminar.

1 Introduction

As a result of Big Data and AI, firms can condition prices on high frequency data, tailor prices to narrow submarkets or even individual consumers, and engage in more effective learning to discover the most profitable pricing rules. While there are potential efficiency benefits from these advances, concerns have been raised about possible consumer harm. Enhanced price discrimination fueled by customer-specific data may increase total welfare, but could result in a transfer of surplus from consumers to firms. Automated pricing with high frequency data could make markets more efficient by increasing the speed of response to demand changes but it is unclear how it will affect price competition. Learning algorithms could deliver more profitable pricing rules but that could be because they facilitate collusion. An active competition policy debate has arisen regarding algorithmic pricing and whether legal and enforcement regimes are equipped to deal with the associated challenges.[1] Helping to inform this debate, there is a growing theoretical literature which identifies and investigates possible anticompetitive effects. As later reviewed, this body of work has established reason for concern in that supracompetitive prices can indeed arise from pricing algorithms augmented by Big Data and AI.

One of the critical implications of Big Data and AI is that it is more attractive for a firm to outsource pricing. By having prices driven more by data and less by the judgment of those employees in the firm with the best soft information, pricing can be delegated to a third party or to a third party's pricing algorithm because the requisite information is portable. A third party is likely to have better pricing algorithms than would be developed internally by a firm because it has more expertise and experience, access to more data, and stronger incentives to invest in their development (as the pricing algorithm can be licensed to many firms). While there are then efficiencies from using a third party, it has also been noted that third party delegation could facilitate coordinated pricing between competitors. The UK's Competition & Markets Authority has expressed concern about the anticompetitive risk when "competitors decide ... that it is more effective to delegate their pricing decisions to a common intermediary which provides algorithmic pricing services" and notes that "[i]f a sufficiently large proportion of an industry uses a single algorithm to set prices, this could result in a ... structure that may have the ability and incentive to increase prices."[2] The German Monopolies Commission has described the concern of using a third party because it "sells an algorithm that it knows or accepts could contribute to a collusive market outcome [and] it is even conceivable that [they] see such a contribution as an advantage, as it makes the algorithm more attractive for users interested in profit maximization.[3] Finally, the OECD has warned: "concerns of coordination would arise if firms outsourced the creation of algorithms to the same IT companies and programmers. This might create a sort of 'hub and spoke' scenario where co-ordination is, willingly or not, caused by competitors using the same 'hub' for developing their pricing algorithms and end up relying on the same algorithms."[4]

The objective of this paper is to investigate the possible anticompetitive effects of firms

[1] Some of that debate can be found in Mehra (2016), Ezrachi and Stucke (2017), OECD (2017a), Deng (2018), Harrington (2018), Gal (2019), Schwalbe (2019), and Calvano et al (2020a).

[2] Competition & Markets Authority (2018), pp. 26-27.

[3] German Monopolies Commission (2018), p. 23.

[4] OCED (2017), p. 27.

using a pricing algorithm developed by a third party. The setting is one where Big Data and AI allow price to condition on high-frequency demand information. As described above, the existing view is that the potential for anticompetitive harm arises when two or more competitors adopt a third party's pricing algorithm.[5] The central finding of this paper is that supracompetitive prices occur even when only one firm adopts the third party's pricing algorithm. In addition to showing that average prices are higher, I also explore how the sensitivity of price to demand compares between when a pricing algorithm is developed externally and internally. To my knowledge, this is the first paper to examine the competitive implications of outsourcing the design of a pricing algorithm.[6]

Section 2 organizes and reviews the theoretical research which explores the effect of Big Data and AI on the intensity of competition; it should be of independent interest. The model is described in Section 3. Section 4 delivers the paper's main finding, and additional results are offered in Section 5 for the case of linear demand. Section 6 puts forth a policy proposal for consideration, and Section 7 concludes.

2 Literature Review

The literature review begins with a classification of the broad body of research and is followed with a more detailed discussion of the most relevant papers.

The theoretical literature examining the implications of Big Data and AI for pricing can be categorized along two dimensions: 1) the space of pricing algorithms; and 2) the criterion for selecting a pricing algorithm. The first dimension pertains to the modelling of how Big Data and AI enrich the feasible set of pricing algorithms. One branch of this literature focuses on personalized pricing. Adding to the voluminous literature on price discrimination, it allows price to condition on a customer's history of purchases or some other customer-specific data. A second branch, often referred to as "dynamic pricing," focuses on how Big Data and AI allow a firm to be more informed of demand when it is setting price. This can mean using past data to have a more accurate demand forecast or using high-frequency data to better tailor price to current market conditions. A third branch examines how pricing algorithms affect the way in which a firm's price responds to competitors' prices in terms of either the speed of response or committing to a particular response. As the first branch is not relevant to the current paper, we'll not discuss it further, while papers in the last two branches are reviewed below.[7]

The second dimension is the modelling of how a firm selects a pricing algorithm. The conventional approach characterizes equilibrium pricing algorithms for a well-defined game. An alternative approach specifies a learning algorithm; that is, how past data (prices, sales, profits) is used to identify a better performing pricing algorithm. Two classes of learning

[5] For example, Ezrachi and Stucke (2017) provide four general ways in which algorithms can be anticompetitive, and all require the pricing algorithm to be used by at least two firms in a market.

[6] The common agency literature covers a related distinction for it deals with the use of a third party's services (Bernheim and Whinston, 1985). In particular, Decarolis, Goldmanis, and Penta (2020) consider when multiple advertisers at a sponsored search auction delegate their bid decisions to a third party.

[7] With regards to the first branch, a sampling of relevant papers include Acquisiti and Varian (2005), Choudhary et al (2005), Chena and Zhang (2009), Zhang (2011), and Chen, Choe, and Matsushima (2020).

algorithms have been considered: estimation-optimization learning and reinforcement learning. The former embodies two distinct modules. The estimation module estimates the firm's environment and delivers predictions as to how the firm's price or quantity determines its profit or revenue. More specifically for our setting, past prices and sales are used to estimate a firm's demand function (where various papers have used OLS, Maximum Likelihood, and an artificial neural network), and thereby have an estimate of how price affects a firm's profit (or revenue). With that estimated environment, the optimization module selects price to maximize profit (or revenue) using the estimated demand function, while adding some randomness to generate experimentation. An example discussed below is Cooper, Homen-de-Mello, and Kleywegt (2015), while den Boer (2015) provides an overview of this work.

An estimation-optimization learning algorithm separately estimates the environment and then optimizes in the selection of an action for the estimated environment. In comparison, reinforcement learning fuses estimation and optimization by learning directly over actions; it seeks to identify the best action for a particular state based on how various actions have performed in the past for that state. Its approach is model-free in that it operates without any prior knowledge of the environment. One common method of reinforcement learning is Q-learning. With this approach, there is a value assigned to each action-state pair (e.g., an action is a price and a state is a history of prices) and these values are updated based on realized profit. Given the current collection of values and the current state, the action is chosen which yields the highest value. Recent papers using Q-learning are Calvano et al (2019) and Klein (2019).[8] Hansen, Misra, and Pai (2020) use the Upper Confidence Bound algorithm which, for each price, keeps track of the empirical average of the profit for that price and the number of times it was chosen. There is an index which is increasing in the empirical average profit and decreasing in the number of times a price was chosen. In any period, the price with the highest index is chosen, so a price is more likely to be selected when it has performed better and has been chosen less frequently.

Let me now turn to reviewing those papers that most directly examine how AI and Big Data affects market competition. The first four papers consider the impact of AI and Big Data on the propensity or extent of collusion. Firms interact in an infinitely repeated price game where pricing algorithms can arbitrarily condition on the history of past prices. Salcedo (2015) modifies the canonical perfect monitoring setting to allow for commitment to and observability of pricing algorithms. A pricing algorithm is a finite automaton which maps price histories into the set of feasible prices. A firm's pricing algorithm is a state variable in that it can be changed only during stochastic revision opportunities. At such an opportunity, a firm is assumed to know its rival's pricing algorithm. Thus, in selecting its pricing algorithm at a revision opportunity, a firm recognizes it will be committed to it until the next revision opportunity and, should its rival have a revision opportunity in the meantime, that rival will observe the firm's pricing algorithm and know it is committed to it. A striking result is derived: under certain conditions, all subgame perfect equilibria result in prices close to monopoly prices. However, a word of caution, for this result is erected on the untenable assumption that a firm observes a rival's pricing algorithm. The presumption

[8] Earlier papers using Q-learning in an environment where multiple firms choose prices or quantities include Tesauro and Kephart (2002), Xie and Chen (2004), Waltman and Kaymak (2008), Dogan and Güner (2015), and Hilsen (2016).

is that past price data would allow a firm to "decode" its rival's pricing algorithm, though that cannot generally be possible (e.g., when the number of observations are fewer than the number of states in the finite automaton).

Miklós-Thal and Tucker (2019) considers a duopoly with homogeneous goods where there is one consumer type with fixed demand. A consumer's maximum willingness-to-pay (WTP) can take two possible values and is *iid* over time. In each period, firms receive a common signal of the WTP prior to choosing price. There are two possible signals and $\rho \geq 1/2$ is the probability that the signal is accurate. The influence of Big Data and AI are captured by a higher value of ρ; hence, a firm has better demand information when it chooses price. The analysis focuses on grim trigger strategy equilibria under perfect monitoring. A higher value of ρ has two counteracting effects on the maximal collusive equilibrium price. More accurate demand information allows the cartel to better predict the joint profit-maximizing price which increases the collusive value and thus makes collusion less difficult. However, more accurate demand information also increases the maximal deviation profit by better informing a prospective deviator when deviation profit is high, which makes collusion more difficult. When the discount factor is sufficiently high, more accurate demand forecasting harms consumers. When the discount factor is sufficiently low, it is possible for consumers to benefit from firms being better informed of demand.

Closely related in motivation is O'Connor and Wilson (2019) which also considers the implications of enhanced demand forecasting though under imperfect monitoring. Without Big Data and AI, demand is affected by two unobservable demand shocks. With Big Data and AI, one of those demand shocks is observed so price can condition on that shock. As with Miklós-Thal and Tucker (2019), the deviation payoff is higher because of the improved demand information which makes collusion harder, but monitoring is more effective which makes collusion easier. The net effect on prices is ambiguous.

The final paper that explore the implications of Big Data and AI for collusion is Calvano et al (2020b). This paper assumes each firm uses Q-learning to discover its pricing algorithm. The central question is whether collusive pricing rules can emerge under Q-learning and, if so, how robust a phenomenon it is. For the infinitely repeated price game with differentiated products, they find it is quite common for prices to converge to levels well above static Nash equilibrium levels. Furthermore, pricing algorithms evolve to having properties of collusive pricing rules.[9] For example, one pricing algorithm that emerged has firms settle on a supracompetitive price and, in response to a rival undercutting it, firms' prices significantly drop and then gradually climb back up to supracompetitive levels. The paper consider many variants of the basic model in concluding that collusion is a robust outcome of Q-learning. Firms whose pricing algorithms are determined by a general form of reinforcement learning can learn to collude.

The remaining papers show how Big Data and AI can result in supracompetitive prices under static optimization. In Brown and MacKay (2020), the profit function is fixed and known, and they focus on the implications of firms being able to respond more rapidly to rivals' prices. In the context of a duopoly game with differentiated products, firms can be

[9] "Collusion is when firms use history-dependent strategies to sustain supracompetitive outcomes through a reward-punishment scheme that rewards a firm for abiding by the supracompetitive outcome and punishes it for departing from it." Harrington (2017), p. 1.

heterogeneous in the frequency with which they can change price. For example, one firm may be able to change price every hour, while the other firm can only change price once a day. This heterogeneity introduces commitment in that the firm which is locked into its price over a longer period is effectively a price leader with respect to its rival. Allowing firms to choose their pricing technologies, firms are shown to select different frequencies because creating a leader-follower relationship yields higher prices and profits for both firms compared to when they simultaneously choose prices (which, by the model's timing structure, occurs when they choose the same frequency). So as to ensure itself of being the follower (which is more profitable than being a leader), one of the firms chooses the most rapid pricing technology. By allowing firms to commit to a pricing frequency, Big Data and AI produce higher prices.

Cooper, Homen-de-Mello, and Kleywegt (2015) and Hansen, Misra, and Pai (2020) consider a duopoly setting with differentiated products, where firms do not know their demand or profit functions and are endowed with a learning algorithm. The only available data to a firm are its own past prices and profits which means, in estimating the relationship between its price and profit, the firm has a misspecified model that does not take account of the other firm's price. With an omitted variable that is endogenous to what the pricing algorithm does, estimates will be biased. For example, if, when a firm raises its price, the other firm also happens to raise its price then the firm's demand will be estimated to be less price-elastic than it actually is. Underestimating the price elasticity of demand would cause firms to set higher prices than would be achieved for a full-information equilibrium. Both papers find that this misspecification results in supracompetitive prices. Cooper, Homen-de-Mello, and Kleywegt (2015) assumes prices are set optimally given an OLS-estimated demand curve. Hansen, Misra, and Pai (2020) view it as a multi-armed bandit problem where a pricing algorithm is chosen to minimize statistical regret (i.e., the difference between average profit achieved with the algorithm and ex-post optimal profit). They find that when the signal-to-noise ratio for sales is high (i.e., sales are relatively more responsive to price changes than to demand shocks), firms' prices are supracompetitive and positively correlated. It is when learning results in a high positive correlation that a firm finds a high price relatively profitable because the rival also tends to set a high price.

The current paper focuses on how Big Data and AI allow price to condition on high-frequency demand information and, therefore, the market setting is along the lines of Miklós-Thal and Tucker (2019) and O'Connor and Wilson (2019). Like Brown and MacKay (2020), the current paper has the implication that Big Data and AI result in higher frequency price changes. However, price changes are more frequent in Brown and MacKay (2020) because Big Data and AI make it feasible, while here it is because information is arriving at a higher frequency and thus makes it optimal. The singular feature of the current paper is the source of the pricing algorithm. All preceding research assumes the pricing algorithm is designed to benefit the firm and, therefore, it could have been developed internally. In this paper's model, the pricing algorithm is designed to benefit the third party which, as we'll see in the next section, implies it is designed with a different objective.

3 Model

3.1 Market

Consider a duopoly with differentiated products.[10] Firm i's demand function is $D_i(p_1, p_2, a)$: $\Re^2_+ \times A$ where p_i is the price of firm $i \in \{1,2\}$. The variable $a \in A \equiv [\underline{a}, \overline{a}]$ is a demand shifter where $\underline{a} < \overline{a}$. Assume symmetric demand functions:

$$D_1(p', p'', a) = D_2(p'', p', a) \ \forall (p', p'', a). \tag{1}$$

Given symmetry, assumptions on firms' demand functions can be stated with respect to firm 1. D_1 is assumed to be twice continuously differentiable in (p_1, p_2, a) and, when $D_1 > 0$, is increasing in p_1, decreasing in p_2, and increasing in a. Furthermore, $\partial D_1(p_1, p_2, a)/\partial p_1$ is non-decreasing in p_2 and a. These assumptions are summarized here:

$$\frac{\partial D_1(p_1, p_2, a)}{\partial p_2}, \frac{\partial D_1(p_1, p_2, a)}{\partial a} > 0 > \frac{\partial D_1(p_1, p_2, a)}{\partial p_1} \tag{2}$$

$$\frac{\partial^2 D_1(p_1, p_2, a)}{\partial p_1 \partial p_2}, \frac{\partial^2 D_1(p_1, p_2, a)}{\partial p_1 \partial a} \geq 0. \tag{3}$$

Firms have a common constant marginal cost $c \geq 0$ so a firm's profit function is:

$$\pi_1(p_1, p_2, a) \equiv (p_1 - c) D_1(p_1, p_2, a). \tag{4}$$

$\pi_1(p_1, p_2, a)$ is assumed to be strictly concave in p_1:

$$\frac{\partial^2 \pi_1(p_1, p_2, a)}{\partial p_1^2} = 2\frac{\partial D_1(p_1, p_2, a)}{\partial p_1} + (p_1 - c)\frac{\partial^2 D_1(p_1, p_2, a)}{\partial p_1^2} < 0. \tag{5}$$

By (2), $\pi_1(p_1, p_2, a)$ is decreasing in p_2 and increasing in a. From (3), it follows:

$$\frac{\partial^2 \pi_1(p_1, p_2, a)}{\partial p_1 \partial p_2} = \frac{\partial D_1(p_1, p_2, a)}{\partial p_2} + (p_1 - c)\frac{\partial^2 D_1(p_1, p_2, a)}{\partial p_1 \partial p_2} > 0 \tag{6}$$

$$\frac{\partial^2 \pi_1(p_1, p_2, a)}{\partial p_1 \partial a} = \frac{\partial D_1(p_1, p_2, a)}{\partial a} + (p_1 - c)\frac{\partial^2 D_1(p_1, p_2, a)}{\partial p_1 \partial a} > 0. \tag{7}$$

Assuming $(p-c)D_1(p, p, a)$ is strictly concave in p, the joint profit-maximizing price exists and is unique:

$$p^M(a) \equiv \arg\max_p (p - c) D_1(p, p, a).$$

Further assume $p^M(a)$ is non-decreasing in a. Firm demand is assumed to be positive for a wide range of price pairs:

$$D_1(p^M(a), c, a) > 0 \ \forall a \in A.$$

[10] While the paper's insight is robust to there being more than two firms, the analysis becomes more complex.

Thus, a firm charging the joint profit-maximizing price has positive demand when the rival firm prices at cost. However, demand is zero when price is high enough:

$$\exists p^{\max} \text{ such that } D_1(p,p,a) = 0 \ \forall a \in A, \forall p > p^{\max}. \tag{8}$$

Note that (8) implies: $D_1(p^{\max}, p_2, a) = 0 \ \forall p_2 \leq p^{\max}$. Without loss of generality, we can restrict price to $[c, p^{\max}]$.

In defining the advantage that comes from a third party's pricing algorithm, we first need to define the environment that a firm faces when it is pricing on its own. In that situation, the demand variable a is assumed to be unobservable to the firm at the time that price is chosen. One can imagine that a is some demand shock that occurs at a higher frequency than the pricing decision so that price cannot condition on it. When deciding on price, firms have common beliefs on a given by the continuously differentiable cdf $F : A \rightarrow [0,1]$ with mean μ and variance σ^2. Define ψ as the best response function under this informational assumption:

$$\psi(p_2) \equiv \arg \max_{p_1 \in [c, p^{\max}]} \int \pi_1(p_1, p_2, a) F'(a) da.$$

Given $\pi_1(p_1, p_2, a)$ is strictly concave in p_1, then so is $\int \pi_1(p_1, p_2, a) F'(a) da$ and thus $\psi(p_2)$ exists and is uniquely defined. The symmetric Nash equilibrium is: $p^N = \psi\left(p^N\right)$. In the absence of a third party providing a pricing algorithm, firms will then price at p^N.

3.2 Third Party Pricing Algorithm Design Problem

The third party provides a pricing algorithm which allows a firm to condition its price on the high-frequency demand state a and thus be able to engage in "dynamic pricing". Let $\phi(\cdot) : A \rightarrow [c, p^{\max}]$ denote a generic pricing algorithm.[11] In modelling the design of the pricing algorithm, several simplifying assumptions are made. First, the third party cannot design the algorithm so that it conditions on whether another firm in the market also uses it. Clearly, such a feature could promote collusive pricing and possibly run afoul of antitrust/competition laws.[12] Second, the third party's objective in designing the algorithm is to maximize the algorithm's performance. The motivation is that the future demand for the algorithm will be enhanced when it has been shown to perform better for those who adopted it. While this is a heuristic for the third party, it is a plausible one for taking account of how future demand is impacted by the pricing algorithm's design. Third, the focus is on the third party's design decision which means we do not consider the licensing fee it charges a firm for use of the pricing algorithm. Fourth, the adoption decision is exogenous and stochastic; conditional on one firm adopting the pricing algorithm, β denotes the probability that both firms adopt the algorithm. As this is an initial investigation into the implications of a third party supplying a pricing algorithm, the last three assumptions are meant to simplify the analysis so as to gain some initial insight. In Section 5, this insight is explained to be robust to endogenizing adoption and the licensing fee.

[11] Implicitly, the pricing algorithm also conditions on the firm's other parameters. I only make explicit its dependence on the high frequency demand state.

[12] Bernheim and Whinston (1985) show that a third party providing marketing services can generate the collusive solution though it requires a third party's contract to condition on how many firms in market use that third party's services.

Let us now characterize the objective function of the third party. In adopting the pricing algorithm $\phi(\cdot)$, suppose firm 1 expects firm 2 to price at: i) p_2 when firm 2 does not adopt the pricing algorithm; and ii) $\phi(\cdot)$ when firm 2 does adopt the algorithm. In that case, a firm's expected profit from adopting the pricing algorithm is

$$(1-\beta)\int \pi_1(\phi(a), p_2, a) F'(a)da + \beta \int \pi_1(\phi(a), \phi(a), a) F'(a)da. \tag{9}$$

(9) measures the performance of the pricing algorithm, and the assumption is that the third party designs $\phi(\cdot)$ to maximize its performance.

A strategy profile for this game is a pair $(\phi(\cdot), p) : \Phi \times [c, p^{\max}]$ where Φ is the space of functions from A to $[c, p^{\max}]$. $\phi(\cdot)$ is the pricing algorithm designed by the third party and p is the price a firm sets when it does not adopt the pricing algorithm and its rival does adopt it. The extensive form is:

- Stage 1: Third party designs the pricing algorithm in order to maximize the expected profit of a firm that adopts it.
- Stage 2: Nature determines how many firms adopt the pricing algorithm.
- Stage 3: Nature determines demand state variable a and reveals it to the firm(s) with the pricing algorithm.
  - If both firms adopted the pricing algorithm then they price according to the pricing algorithm.
  - If one firm adopted the pricing algorithm then it prices according to the pricing algorithm and the other firm chooses price to maximize its expected profit given the other uses the pricing algorithm.
  - If no firms adopted the pricing algorithm then firms simultaneously choose price to maximize expected profit given the other firm's price.

In the last scenario, equilibrium has both firms price at p^N. Our analysis will focus on the design of the pricing algorithm and its implications for prices when one or both firms adopt the pricing algorithm.

Before moving on, it is worth noting that this formulation does not require the third party to know a firm's cost and demand function when it is developing the pricing algorithm. The third party can design a pricing algorithm so that it is tunable with respect to cost and demand parameters. If a firm knew its cost and demand parameters then the third party would construct the software so the firm would input the appropriate values. If a firm did not know its demand parameters then, as reviewed in Section 2, some pricing algorithms augment an optimization module with an estimation module to learn demand. The pricing algorithm characterized in this paper would be the product of the optimization module after the estimation module derived demand estimates. This description does presume estimation is conducted prior to price being chosen, while actual estimation-optimization learning algorithms have estimation occurring in real time along with optimization. The pricing algorithm derived here can then be thought of as the long-run pricing algorithm after the demand parameters have been learned.

4 General Case

A Nash equilibrium pricing algorithm-price pair $\left(\widehat{\phi}\left(\cdot\right),\widehat{p}\right)$ is defined by:

$$\begin{aligned}\widehat{\phi}\left(\cdot\right) &= \arg\max_{\phi(\cdot)\in\Phi}(1-\beta)\int\pi_1\left(\phi\left(a\right),\widehat{p},a\right)F'(a)da \\ &\quad +\beta\int\pi_1\left(\phi\left(a\right),\phi\left(a\right),a\right)F'(a)da \\ \widehat{p} &= \arg\max_{p\in[c,p^{\max}]}\int\pi_2\left(\widehat{\phi}\left(a\right),p,a\right)F'(a)da.\end{aligned} \tag{10}$$

Equivalently, (10) can be represented as:

$$\begin{aligned}\widehat{\phi}\left(a\right) &= \arg\max_{p\in[c,p^{\max}]}(1-\beta)\pi_1\left(p,\widehat{p},a\right)+\beta\pi_1\left(p,p,a\right),\forall a\in A \\ \widehat{p} &= \arg\max_{p\in[c,p^{\max}]}\int\pi_2\left(\widehat{\phi}\left(a\right),p,a\right)F'(a)da.\end{aligned} \tag{11}$$

Implicit in this equilibrium approach are some informational assumptions warranting discussion. First, a firm knows whether its rival adopted the pricing algorithm. That seems reasonable as it could be inferred from a firm's high-frequency price changes. Second, a firm who did not adopt the pricing algorithm knows the properties of the pricing algorithm. Though a strong assumption, it is consistent with standard assumptions made in oligopoly models. As shown in Section 5 for when demand is linear, it is sufficient for a firm to know the expected price charged by a firm using the pricing algorithm rather know than the algorithm itself. Given the empirical distribution on a rival's price, a firm would have an estimate in the form of the average price.

Theorem 1 *An equilibrium $\left(\widehat{\phi}\left(\cdot\right),\widehat{p}\right)$ exists.*[13]

Proofs are in Appendix A.

In assessing the effect on price levels from the adoption of a third party pricing algorithm, there are two effects at work. First, a firm that uses the pricing algorithm conditions its price on the high-frequency demand state. Effectively, it is able to engage in third-degree price discrimination. Second, the pricing algorithm is designed to maximize the algorithm's performance while recognizing it may be competing against itself. When $\beta=0$, this second effect is neutralized as, when designing the algorithm, the third party does not consider that more than one firm in a market may adopt it. Theorem 2 shows that prices are increasing in β which means third party development leads to higher prices when the prospect of the pricing algorithm facing itself in the market is given more weight in the design process.[14]

Theorem 2 *If $|\overline{a}-\underline{a}|$ is sufficiently small then, for the minimal equilibrium, $\widehat{\phi}\left(\cdot\right)$ and $\widehat{p}$ are increasing in β.*

[13] As existence is established using the Tarski Fixed Point Theorem, it is also the case that there exists a minimal and maximal equilibrium.

[14] For technical reasons, the result is shown when the demand variation is not too great, though that is not believed to be a necessary condition.

In response to a higher probability that the pricing algorithm will compete against itself, the third party programs the pricing algorithm to price higher; that is, if β is higher than $\widehat{\phi}(a)$ is higher for any value of a. The third party designs the pricing algorithm to be less aggressive because that will enhance its performance in the event that both competitors adopt the pricing algorithm. In responding to a higher price by a rival with the pricing algorithm, a non-adopting firm also prices higher - $\widehat{p}$ is increasing in β - because it anticipates (or learns) that its rival prices higher due to having adopted the pricing algorithm. While these effects are straightforward in light of the structure of the model, the next section will show they deliver a finding which runs contrary to existing understanding about the potential harm of third party pricing algorithms.

5 Linear Demand Case

5.1 Model and Benchmark

In order to derive some further properties of equilibrium, let us assume linear demand:

$$D_1(p_1, p_2; a) = a - bp_1 + dp_2$$

where $b > d \geq 0$ and $\underline{a} - (b-d)c > 0$. As a benchmark, the Nash equilibrium price when neither firm conditions on the high-frequency demand state is $p^N = \frac{\mu+bc}{2b-d}$, where μ is the expected value of a.[15] As a second benchmark, suppose both firms condition price on the high-frequency demand state but without the assistance of the third party; thus, the pricing algorithm is internally developed. In that situation, the Nash equilibrium price is $p^N(a) = \frac{a+bc}{2b-d}$, and again the expected equilibrium price is $\frac{\mu+bc}{2b-d}$. In fact, the expected equilibrium price is $\frac{\mu+bc}{2b-d}$ for both firms regardless of how many firms use an internally-developed pricing algorithm.

5.2 Equilibrium Prices

It is derived in Appendix B that the unique solution to (11) is:

$$\begin{aligned}\widehat{\phi}(a) &= \widehat{\gamma} + \widehat{\theta}a \qquad (12)\\ &= \frac{\mu(1-\beta)d(2(b-\beta d)+d) + c(b-\beta d)(4b(b-\beta d)+2bd(1-\beta))}{(b-\beta d)(8b(b-d)+2d(1-\beta)(4b-d))} + \left(\frac{1}{2(b-\beta d)}\right)a\end{aligned}$$

$$\widehat{p} = \frac{\mu(2(b-\beta d)+d) + c(2b+d)(b-\beta d)}{4b\beta(b-d) + (4b^2-d^2)(1-\beta)}. \qquad (13)$$

Given that average price is $p^N = \frac{\mu+bc}{2b-d}$ in the absence of a third party developer, we can see that the firm without the pricing algorithm is pricing higher:

$$\widehat{p} - p^N = \beta\left(\frac{d^2}{2b-d}\right)\left(\frac{\mu-(b-d)c}{4b\beta(b-d)+(4b^2-d^2)(1-\beta)}\right) > 0. \qquad (14)$$

[15] That p^N is the equilibrium price does presume $|\overline{a} - \underline{a}|$ is not too large so that firms' demands are always positive at a price of p^N.

Next note that, on average, the firm with the pricing algorithm prices higher than the firm without the pricing algorithm:

$$\widehat{\gamma} + \widehat{\theta}\mu - \widehat{p} = \frac{\beta d\left(\mu - (b-d)c\right)}{4b\beta(b-d) + (4b^2 - d^2)(1-\beta)} > 0. \tag{15}$$

(14)-(15) imply $\widehat{\gamma} + \widehat{\theta}\mu > p^N$. In sum, when the third party takes into account that its pricing algorithm may compete against itself ($\beta > 0$), average price of both firms is higher compared to competitive prices in the absence of the third party provision of a pricing algorithm, whether firms do or do not internally develop pricing algorithms. Furthermore, (15) implies average market price is higher when both firms adopt the pricing algorithm compared to when one firm adopts it. This is summarized as Property 1.

Property 1: If $\beta > 0$ then: i) when at least one firm adopts the pricing algorithm, both firms' average prices are higher compared to the competitive solution without third party provision of the pricing algorithm; and ii) the average market price is higher when more firms adopt the third party's pricing algorithm.

Property 1 runs contrary to existing understanding in that it shows a third party's pricing algorithm results in supracompetitive prices even when only one firm in a market adopts it. Recall that the pricing algorithm was prohibited from conditioning on the *ex post* adoption decisions. This restriction reflects the expressed concern of competition authorities that, when competitors adopt a pricing algorithm from the same third party, they could be designed to communicate and coordinate on higher prices. Our analysis shows that, under such a prohibition, the third party will *ex ante* design the pricing algorithm to be less competitive in order to take into account the possibility of competitors adopting it. Outsourcing pricing to a third party raises average prices.

Outsourcing of a firm's pricing rule also affects price variability. First, let us consider the sensitivity of average price to the average demand state, which is measured by μ. There could be predictable changes in μ such as higher demand during the holiday season. Without third party provision, a firm's average price is p^N whether or not it internally develops the pricing algorithm. It is straightforward to show:

$$\frac{\partial\left(\widehat{\gamma} + \widehat{\theta}\mu\right)}{\partial\mu} = \frac{2b + d(1-\beta)}{4b(b-\beta d) - d^2(1-\beta)} \gtreqqless \frac{1}{2b-d} = \frac{\partial p^N}{\partial\mu} \text{ as } \beta \gtreqqless 0$$

$$\frac{\partial\widehat{p}}{\partial\mu} = \frac{2(b-\beta d) + d}{4b\beta(b-d) + (4b^2 - d^2)(1-\beta)} \gtreqqless \frac{1}{2b-d} = \frac{\partial p^N}{\partial\mu} \text{ as } \beta \gtreqqless 0,$$

which delivers Property 2.[16]

Property 2: If $\beta > 0$ then, when at least one firm adopts a third party's pricing algorithm, both firms' average prices are more sensitive to the average demand state compared to when firms do not use a third party's pricing algorithm.

[16] The first property follows from: $\frac{\partial\left(\widehat{\gamma}+\widehat{\theta}\mu\right)}{\partial\mu} = \frac{\partial p^N}{\partial\mu}$ when $\beta = 0$ and $\frac{\partial^2\left(\widehat{\gamma}+\widehat{\theta}\mu\right)}{\partial\mu\partial\beta} \frac{2bd(2b+d)}{(d^2\beta+4b^2-d^2-4bd\beta)^2} > 0$, and the second property from: $\frac{\partial\widehat{p}}{\partial\mu} = \frac{\partial p^N}{\partial\mu}$ when $\beta = 0$ and $\frac{\partial^2\widehat{p}}{\partial\mu\partial\beta} = \frac{(2b+d)d^2}{(4b(b-d\beta)-d^2(1-\beta))^2} > 0$.

This property is another implication of the softening of price competition with the third party's pricing algorithm. In response to stronger demand, a firm raises its price but is constrained by losing demand to the other firm. However, the third party's pricing algorithm is designed to implicitly value the other firm's profit (as captured by β) in which case the latter effect is less constraining; hence, price rises more in response to stronger demand.

Comparing external and internal development of pricing algorithms with regards to price sensitivity of price to the high-frequency demand state yields ambiguous results. The source of that ambiguity comes from two counteracting effects. We have already noted one effect, which is that the third party's pricing algorithm is less aggressive which tends to make price more responsive to the demand state. A second effect is that the third party's pricing algorithm is prevented from conditioning on whether the other firm also adopted the pricing algorithm, and thus whether the other firm is also conditioning its price on a. In contrast, a firm which internally develops its pricing algorithm can have it depend on whether or not its rival is also conditioning price on the high-frequency demand state. When its rival also has internally developed a pricing algorithm, this will make a firm's price more sensitive to the demand state. Given the rival's price is higher when the demand state is higher then this firm's price will be higher due to strategic complements. This second effect tends to make price more responsive to the demand state when the pricing algorithm is internally developed.

To see these effects at work, let us begin by considering when only one firm adopts in which case the second effect is absent. When only one firm adopts, the internally developed pricing algorithm is $\widehat{\gamma} + \widehat{\theta} a$ when β is set at zero.[17] Next note

$$\frac{\partial(\widehat{\gamma} + \widehat{\theta} a)/\partial a}{\partial \beta} = \frac{d}{2\,(b - d\beta)^2} > 0. \tag{16}$$

Given that external development allows for $\beta > 0$, (16) implies price sensitivity is greater under external development.

Next consider when both firms have an algorithm that conditions price on the high-frequency demand state. Price is still more sensitive with outsourcing but now only when $\beta > 1/2$.

$$\frac{\partial \widehat{\phi}(a)}{\partial a} = \frac{1}{2b - 2\beta d} \gtreqless \frac{1}{2b - d} = \frac{\partial p^N(a)}{\partial a} \text{ as } \beta \gtreqless \frac{1}{2}.$$

The two effects are at work here. First, the firm that internally develops it conditions on its rival having also done so. This leads to more price sensitivity compared to external development which only assigns probability β to the rival also conditioning its price on the high-frequency demand state. Second, the externally-sourced pricing algorithm is also taking into account how price affects the rival's profit because that will (probabilistically) affect the pricing algorithm's performance. The second effect dominates the first effect when the third party attaches sufficient weight to the prospect of adoption by both firms.

Property 3: Third party development results in a pricing algorithm that is more sensitive to the high-frequency demand state when: i) one firm adopts the pricing algorithm;

[17] The pricing algorithm is $\frac{(2b-d)(a+bc)+d(\mu+bc)}{2b(2b-d)}$, which is a best response to the non-adopting firm using $\frac{\mu+bc}{2b-d}$.

or ii) both firms adopt the pricing algorithm and $\beta > 1/2$. Otherwise, internal development results in a pricing algorithm that is more sensitive to the high-frequency demand state.

5.3 Robustness to Endogenizing Adoption and the Licensing Fee

The main insight of this paper is that the third party's pricing algorithm softens price competition because a third party designs the algorithm to maximize the algorithm's performance. This result relies on the third party assigning positive probability to more than one firm in the market adopting its pricing algorithm; that is, $\beta > 0$ so the third party believes its pricing algorithm might face itself in the market. The analysis has treated adoption as exogenous and ignored the third party's licensing fee. In this section, I show that if an equilibrium with endogenous adoption and licensing fee exists then it must have $\beta > 0$.[18]

For some model of adoption, let us suppose, to the contrary, that there is an equilibrium in which at most one firm adopts so the third party sets $\beta = 0$ when designing the pricing algorithm. In that case, the pricing algorithm is

$$\widehat{\phi}(a) = \frac{(2b-d)(a+bc) + d(\mu + bc)}{2b(2b-d)}$$

and, in the event one firm does adopt, the non-adopting firm's price is

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

I will first show that, in the absence of any licensing fee, adopting the pricing algorithm is always profitable for a firm.

If the rival firm does not adopt the pricing algorithm, the expected profit from not adopting is

$$\int \left(\frac{\mu + bc}{2b-d} - c\right)\left(a - (b-d)\left(\frac{\mu + bc}{2b-d}\right)\right) F'(a)da = \frac{b\left(\mu - (b-d)c\right)^2}{(2b-d)^2},$$

and from adopting is

$$\begin{aligned} &\int \left(\frac{(2b-d)(a+bc) + d(\mu + bc)}{2b(2b-d)} - c\right) \times \\ &\left(a - b\left(\frac{(2b-d)(a+bc) + d(\mu + bc)}{2b(2b-d)}\right) + d\left(\frac{\mu + bc}{2b-d}\right)\right) F'(a)da \\ &= \frac{b\left(\mu - (b-d)c\right)^2}{(2b-d)^2} + \frac{\sigma^2}{4b}. \end{aligned}$$

Hence, adoption is optimal and the incremental expected profit gain is $\frac{\sigma^2}{4b} > 0$, where recall σ^2 is the variance of the high-frequency demand state a.

[18] It is left to future research to prove existence and characterize such an equilibrium.

If the rival firm adopts the pricing algorithm, the expected profit from not adopting is

$$\int \left(\frac{\mu + bc}{2b - d} - c \right) \left(a - b \left(\frac{\mu + bc}{2b - d} \right) + d \left(\frac{(2b - d)(a + bc) + d\,(\mu + bc)}{2b(2b - d)} \right) \right) F'(a)da$$

$$= \frac{b\,(\mu - (b - d)c)^2}{(2b - d)^2}$$

and from adopting is

$$\int \left(\frac{(2b - d)(a + bc) + d\,(\mu + bc)}{2b(2b - d)} - c \right) \times$$
$$\left(a - (b - d) \left(\frac{(2b - d)(a + bc) + d\,(\mu + bc)}{2b(2b - d)} \right) \right) F'(a)da$$

$$= \frac{b\,(\mu - (b - d)c)^2}{(2b - d)^2} + \frac{(b + d)\sigma^2}{4b^2}.$$

Hence, adoption is optimal and the incremental expected profit gain is $\frac{(b+d)\sigma^2}{4b^2} > 0$. Adoption is then a dominant strategy. Also note that the incremental expected profit from adoption is higher when the other firm adopts the pricing algorithm:

$$\frac{(b + d)\sigma^2}{4b^2} > \frac{\sigma^2}{4b} \Leftrightarrow d > 0. \tag{17}$$

If the third party were to design the pricing algorithm based on $\beta = 0$ and firms had accurate expectations on the profits from adopting the pricing algorithm then, as long as the licensing fee charged by the third party is not too high, both firms would adopt the pricing algorithm because it yields higher expected profits irrespective of the rival firm's adoption decision. That would contradict $\beta = 0$. If the third party charges the same licensing fee to all firms then, given the incremental expected profit from adoption is higher when the other firm adopts (as shown in (17)), either the fee is set so that both firms adopt or neither firm adopts. Clearly, the former is preferable for the third party.

I conclude that an equilibrium for a model which endogenizes adoption and the licensing fee would result in $\beta > 0$ in which case the anticompetitive effect from third party development of the pricing algorithm is present.

6 Policy Proposal

It is clear that a third party's pricing algorithm should be prohibited from being able to recognize when another firm is using the same pricing algorithm. If that was allowed then pricing algorithms could be programmed to "communicate" with each other and coordinate on setting higher prices. The analysis of this paper shows that such a prohibition is not enough to avoid anticompetitive effects. For if *ex post* recognition is prohibited then the third party will *ex ante* design the pricing algorithm to take into account the possibility that it will face itself in the market. As a third party wants to enhance the performance of its pricing algorithm, it will design it to soften price competition and, consequently, prices are higher even when only one firm adopts the pricing algorithm.

One possible policy response is to prohibit a firm from adopting a third party's pricing algorithm. Such a policy is not only extreme but its welfare implications are unclear and difficult to assess. While the prohibition would intensify price competition - as we showed average prices are higher with the third party's pricing algorithm - it would also prevent a firm from engaging in third-degree price discrimination. Prohibiting price discrimination could either raise or lower welfare, depending on the particular market conditions.[19] If third-degree price discrimination reduced welfare then prohibiting the supply of pricing algorithms by a third party would increase welfare by avoiding the softening of price competition as well as preventing price discrimination. However, it is generally beyond an antitrust or regulatory agency to determine with confidence when price discrimination reduces welfare which is why they have generally not been given such authority. A prohibition on pricing algorithms developed by a third party is excessive and inappropriate.

A more viable policy is to prohibit more than one firm in a market from using a pricing algorithm developed by the same third party. This prohibition would cause the third party to set $\beta = 0$ when developing the pricing algorithm, in which case it would no longer be designed to soften price competition. Thus, the adoption of the pricing algorithm would not be anticompetitive. However, there are some drawbacks of this policy. If there is only one third party supplier then this policy would limit one firm in a market to engaging in third-degree price discrimination; again, we are back to regulating the extent of price discrimination with its ambiguous welfare effects. Furthermore, the policy's implementation would face the challenge of defining the market so that competitors could be identified. Of course, while market definition can be difficult, it is a common exercise performed in the context of merger evaluation and other antitrust issues. On the positive side, this policy would incentivize other third parties to develop pricing algorithms, in which case the benefits of dynamic pricing would be delivered without the anticompetitive effect coming from a single third party supplier.

While it has its advantages and disadvantages, a prohibition on competitors adopting a pricing algorithm from the same third party seems to be a plausible option worthy of further examination.

7 Concluding Remarks

This paper is the first to explore the competitive implications of a third party developing a firm's pricing algorithm. While a firm would design its pricing algorithm to maximize the firm's performance, a third party will design it to maximize the algorithm's performance. This consideration was shown to result in the third party's pricing algorithm softening competition with higher prices and, contrary to existing understanding, this anticompetitive effect does not require multiple competitors to adopt it. The pricing algorithm is designed under the restriction that it cannot condition on the *ex post* adoption decisions of firms because of the expressed concern of competition authorities that, when competitors adopt

[19] See, for instance, Varian (1989), Bergemann, Brooks, and Morris (2015), and Cowan (2016).

a pricing algorithm from the same third party, the algorithm could be designed to communicate and coordinate on higher prices. Consequently, the third party will *ex ante* design around that restriction by making the pricing algorithm less competitive in order to take into account the possibility of the algorithm competing against itself.

Even before the development of this theory of harm, commentators have noted the challenge faced by competition authorities in responding to anticompetitive effects coming from firms independently using a third party's pricing algorithm or services. For the European Union, it has been noted:

> If the software developer decides to use the algorithm in an anticompetitive way to manipulate the market of the software users without their knowledge, neither of them is likely to be found liable under Article 101(1) TFEU.[20]

The Antitrust Division of the U.S. Department of Justice sees a similar lack of liability in the United States:

> [I]ndependent adoption of the same or similar pricing algorithms is unlikely to lead to antitrust liability even if it makes interdependent pricing more likely. For example, if multiple competing firms unknowingly purchase the same software to set prices, and that software uses identical algorithms, this may effectively align the pricing strategies of all the market participants.[21]

The legal challenge is that competitors' prices are coordinated without them necessarily having an illegal agreement. Our analysis showed that it is even more challenging than imagined because supracompetitive prices can emerge when only one firm adopts a third party's pricing algorithm. There is then a disconcerting loophole in antitrust/competition law.

In conducting this initial investigation, some simplifying assumptions were made. In particular, firms' adoption decisions were assumed to be exogenous and the focus was on the design of the pricing algorithm and not how it was priced by the third party. While the paper's main finding was explained to be robust to those assumptions, other questions will require developing a model of adoption and allowing the third party to choose both design and price. For example, what determines how widespread is adoption of a third party's pricing algorithm? That can only be addressed by deriving the demand curve for pricing algorithms and the third party's licensing fee. What is the effect of competition among third party providers? In comparing different design-fee options offered by third parties, adoption again will need to be endogenized. Having introduced multiple third party providers, we can also investigate the proposal to prohibit a third party from supplying more than one firm. Does it encourage third parties to incur the fixed cost of development or does it reduce competition among third parties and result in higher licensing fees? Research could also consider the third party development of pricing algorithms for another Big Data dimension: customer-specific data. How do personalized pricing algorithms differ when developed externally by a third party rather than internally by the firm? These are some of the many open questions related to the effect on market competition from the outsourcing of pricing algorithms.

[20] Marx, Ritz, and Weller (2019), p. 7.

[21] OECD (2017b), p. 6.

8 Appendix A: Proofs

8.1 Proof of Theorem 1

(11) can be cast as the following fixed-point problem. Define $\widetilde{\phi}(a,\widetilde{p})$ to be the optimal price for the pricing algorithm given demand state a and the other firm prices at $\widetilde{p}$ when it does not adopt the pricing algorithm:

$$\widetilde{\phi}(a,\widetilde{p}) = \arg\max_{p\in[c,p^{\max}]}(1-\beta)\pi_1(p,\widetilde{p},a) + \beta\pi_1(p,p,a). \tag{18}$$

Next define $\varphi(\widetilde{p})$ as the optimal price for a firm that did not adopt the pricing algorithm given the other firm did adopt it and the pricing algorithm is $\widetilde{\phi}(\cdot,\widetilde{p})$:

$$\varphi(\widetilde{p}) = \arg\max_{p_2\in[c,p^{\max}]}\int \pi_2\left(\widetilde{\phi}(a,\widetilde{p}),p_2,a\right)dF(a). \tag{19}$$

An equilibrium price for the non-adopting firm is then a fixed point to φ, $\widehat{p}=\varphi(\widehat{p})$, and the equilibrium pricing algorithm is the best response to that price, $\widehat{\phi}(a)=\widetilde{\phi}(a,\widehat{p})$. The plan is to prove $\varphi:[c,p^{\max}]\to[c,p^{\max}]$ is increasing and, therefore, a fixed point exists (Tarski, 1955).

As an initial step, let us show $\widetilde{\phi}(a,\widetilde{p})$ is continuous and increasing in $\widetilde{p}$. Given $\pi_1(p,\widetilde{p},a)$ and $\pi_1(p,p,a)$ are both strictly concave in p then $(1-\beta)\pi_1(p,\widetilde{p},a)+\beta\pi_1(p,p,a)$ is strictly concave in p.[22] That property, along with continuity of $(1-\beta)\pi_1(p,\widetilde{p},a)+\beta\pi_1(p,p,a)$ in $\widetilde{p}$, imply its optimum $\widetilde{\phi}(a,\widetilde{p})$ is continuous in $\widetilde{p}$. Next note, by strict concavity,

$$\frac{\partial\left((1-\beta)\pi_1(p,\widetilde{p},a)+\beta\pi_1(p,p,a)\right)}{\partial p} \gtreqless 0 \text{ as } p \lesseqgtr \widetilde{\phi}(a,\widetilde{p}). \tag{20}$$

Given $\partial\left((1-\beta)\pi_1(p,\widetilde{p},a)+\beta\pi_1(p,p,a)\right)/\partial p$ is increasing in $\widetilde{p}$ by (6) then: if $p''>p'$ then

$$\frac{\partial\left((1-\beta)\pi_1(p,p'',a)+\beta\pi_1(p,p,a)\right)}{\partial p} > 0\ \forall p \le \widetilde{\phi}(a,p')$$

which implies $\widetilde{\phi}(a,p'') > \widetilde{\phi}(a,p')$. Hence, $\widetilde{\phi}(a,\widetilde{p})$ is increasing in $\widetilde{p}$.

With this property, the next step is show φ is increasing. Given $\partial\pi_2(p_1,p,a)/\partial p_2$ is continuous and increasing in p_1 and $\widetilde{\phi}(a,\widetilde{p})$ is continuous and increasing in $\widetilde{p}$, then $\partial\pi_2\left(\widetilde{\phi}(a,\widetilde{p}),p,a\right)/\partial p_2$ is continuous and increasing in $\widetilde{p}$. Hence,

$$\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}(a,\widetilde{p}),p,a\right)}{\partial p_2}\right)F'(a)da$$

is continuous and increasing in $\widetilde{p}$. Given $\int\pi_2(p_1,p_2,a)F'(a)da$ is strictly concave in p_2 then

$$\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}(a,\widetilde{p}),p_2,a\right)}{\partial p_2}\right)F'(a)da \gtreqless 0 \text{ as } p_2 \lesseqgtr \varphi(\widetilde{p}).$$

[22] This is the only place where we need strict concavity rather than strict quasi-concavity.

Evaluate it at $\widetilde{p} = p'$,

$$\int \left(\frac{\partial \pi_2 \left(\widetilde{\phi}(a, p'), p_2, a \right)}{\partial p_2} \right) F'(a) da \gtreqless 0 \text{ as } p_2 \lesseqgtr \varphi(p').$$

It follows from $\int \left(\partial \pi_2 \left(\widetilde{\phi}(a, \widetilde{p}), p, a \right) / \partial p_2 \right) F'(a) da$ being increasing in $\widetilde{p}$ that: if $p' > p''$ then

$$\int \left(\frac{\partial \pi_2 \left(\widetilde{\phi}(a, p''), p_2, a \right)}{\partial p_2} \right) F'(a) da > 0 \ \forall \ p_2 < \varphi(p')$$

which implies $\varphi(p'') > \varphi(p')$. Hence, $\varphi(\widetilde{p})$ is increasing in $\widetilde{p}$.

Given φ is increasing on a compact set, it has a fixed point by the Tarski Fixed Point Theorem. ■

8.2 Proof of Theorem 2

By Theorem 1 in Milgrom and Roberts (1994), if φ is increasing in $\widetilde{p}$ and β then the minimal and maximal fixed points of φ are increasing in β. Unfortunately, I have not been able to generally establish that φ is increasing in β which prevents use of that theorem. However, the same end can be achieved through a bit of work. The proof strategy has three steps, where recall $p^M(a)$ is the joint profit-maximizing price which is non-decreasing in a.

1. If $|\overline{a} - \underline{a}|$ is sufficiently small then the minimal equilibrium has $\widehat{p} < p^M(\underline{a})$. (Note: $\widehat{p} < p^M(\underline{a})$ implies $\widehat{p} < p^M(a) \forall a \in [\underline{a}, \overline{a}]$.)

2. If $\widetilde{\phi}(a, \widetilde{p}) < p^M(a)$ then $\widetilde{\phi}(a, \widetilde{p})$ is increasing in β. By step 1, $\widetilde{\phi}(a, \widehat{p}) < p^M(a) \ \forall a \in [\underline{a}, \overline{a}]$ when $|\overline{a} - \underline{a}|$ is sufficiently small which then implies $\widetilde{\phi}(a, \widetilde{p})$ is increasing in β around an equilibrium.

3. If $\widetilde{\phi}(a, \widetilde{p})$ is increasing in β then $\widehat{p}$ and $\widehat{\phi}(\cdot)$ are increasing in β.

Lemma 3 *If $|\overline{a} - \underline{a}|$ is sufficiently small then the minimal equilibrium has $\widehat{p} < p^M(\underline{a})$.*

The proof strategy is to show, when $|\overline{a} - \underline{a}|$ is sufficiently small, $\varphi\left(p^N(\underline{a})\right) > p^N(\underline{a})$ and $\varphi\left(p^M(\underline{a})\right) < p^M(\underline{a})$. Given φ is increasing, there exists a fixed point in $\left(p^N(\underline{a}), p^M(\underline{a})\right)$ and, therefore, $\widehat{p} < p^M(\underline{a})$.

Assume $|\overline{a} - \underline{a}|$ is sufficiently small so the Nash equilibrium price for the highest demand state is less than the joint profit-maximizing price for the lowest demand state: $p^N(\overline{a}) < p^M(\underline{a})$. Consider $\widetilde{p} = p^N(\underline{a})$ and let us show $\varphi\left(p^N(\underline{a})\right) > p^N(\underline{a})$ when $|\overline{a} - \underline{a}|$ is sufficiently small. Define the full-information Nash equilibrium:

$$\frac{\partial \pi_1 \left(p^N(a), p^N(a), a\right)}{\partial p_1} = 0.$$

Let us show that if $\widetilde{p} = p^N(\underline{a})$ then the marginal expected profit for the third party objective is increasing for prices not exceeding $p^N(\underline{a})$:

$$\begin{aligned} & \frac{\partial\left((1-\beta)\pi_1\left(p, p^N(\underline{a}), a\right) + \beta\pi_1\left(p, p, a\right)\right)}{\partial p} \\ = \ & (1-\beta)\frac{\partial\pi_1\left(p, p^N(\underline{a}), a\right)}{\partial p_1} + \beta\left(\frac{\partial\pi_1\left(p, p, a\right)}{\partial p_1} + \frac{\partial\pi_1\left(p, p, a\right)}{\partial p_2}\right) > 0 \\ \forall p \ \leq \ & p^N(\underline{a}), \forall a \in \left[\underline{a}, \overline{a}\right]. \end{aligned} \tag{21}$$

To prove (21), first note that

$$\frac{\partial\pi_1\left(p, p^N(\underline{a}), \underline{a}\right)}{\partial p_1} \gtreqqless 0 \text{ as } p \lesseqqgtr p^N(\underline{a}).$$

Given $\partial^2\pi_1/\partial p_1\partial a > 0$, it follows

$$\frac{\partial\pi_1\left(p, p^N(\underline{a}), a\right)}{\partial p_1} \geq 0 \ \forall p \leq p^N(\underline{a}), \forall a \geq \underline{a}. \tag{22}$$

Given $p^N(\overline{a}) < p^M(\underline{a}) \left(\leq p^M(a)\forall a\right)$ then

$$\frac{\partial\pi_1\left(p, p, a\right)}{\partial p_1} + \frac{\partial\pi_1\left(p, p, a\right)}{\partial p_2} > 0 \ \forall p \leq p^N(\underline{a}), \forall a. \tag{23}$$

(22) and (23) imply (21). By strict concavity, it follows from (21) that the optimum to the third party's objective exceeds $p^N(\underline{a})$: $\widetilde{\phi}\left(a, p^N(\underline{a})\right) > p^N(\underline{a})\forall a \in \left[\underline{a}, \overline{a}\right]$.

Given

$$\frac{\partial\pi_2\left(p^N(\underline{a}), p^N(\underline{a}), \underline{a}\right)}{\partial p_2} = 0,$$

$\partial^2\pi_2/\partial p_2\partial a > 0$ implies

$$\frac{\partial\pi_2\left(p^N(\underline{a}), p^N(\underline{a}), a\right)}{\partial p_2} > 0, \forall a > \underline{a}. \tag{24}$$

Given $\partial^2\pi_2/\partial p_2\partial p_1 > 0$, (24) implies

$$\text{if } p > p^N(\underline{a}) \text{ then } \frac{\partial\pi_2\left(p, p^N(\underline{a}), a\right)}{\partial p_2} > 0, \forall a > \underline{a}. \tag{25}$$

Given $\widetilde{\phi}\left(a, p^N(\underline{a})\right) > p^N(\underline{a})$, (25) implies:

$$\int_{\underline{a}}^{\overline{a}} \frac{\partial\pi_2\left(\widetilde{\phi}\left(a, p^N(\underline{a})\right), p^N(\underline{a}), a\right)}{\partial p_2} dF(a) > 0. \tag{26}$$

From (26) and strict concavity, we have: $\varphi\left(p^N(\underline{a})\right) > p^N(\underline{a})$.

Next consider $\widetilde{p} = p^M(\underline{a})$ and let us show $\varphi\left(p^M(\underline{a})\right) < p^M(\underline{a})$ when $|\overline{a} - \underline{a}|$ is sufficiently small. Consider the first derivative of the objective function for the firm with the pricing algorithm when $\widetilde{p} = p^M(\underline{a})$ and evaluate it at $p = p^M(a)$:

$$\begin{aligned}\Gamma(a) &\equiv (1-\beta)\frac{\partial\pi_1\left(p^M(a), p^M(\underline{a}), a\right)}{\partial p_1} + \beta\left(\frac{\partial\pi_1\left(p^M(a), p^M(a), a\right)}{\partial p_1} + \frac{\partial\pi_1\left(p^M(a), p^M(a), a\right)}{\partial p_2}\right) \\ &= (1-\beta)\frac{\partial\pi_1\left(p^M(a), p^M(\underline{a}), a\right)}{\partial p_1}.\end{aligned}$$

Since

$$\Gamma(\underline{a}) = (1-\beta)\frac{\partial\pi_1\left(p^M(\underline{a}), p^M(\underline{a}), \underline{a}\right)}{\partial p_1} < 0$$

then $\widetilde{\phi}\left(\underline{a}, p^M(\underline{a})\right) < p^M(\underline{a})$. By continuity, $\Gamma(a) < 0$ for $a \simeq \underline{a}$ which implies: $\widetilde{\phi}\left(a, p^M(\underline{a})\right) < p^M(\underline{a})\ \forall a \in [\underline{a}, \overline{a}]$ when $|\overline{a} - \underline{a}|$ is sufficiently low.

Given $p^M(a)$ is non-decreasing in a, it follows: $\widetilde{\phi}\left(a, p^M(\underline{a})\right) < p^M(a)\ \forall a \in [\underline{a}, \overline{a}]$. I want to show:

$$\int_{\underline{a}}^{\overline{a}}\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a, p^M(\underline{a})\right), p^M(\underline{a}), a\right)}{\partial p_2}\right)F'(a)da < 0 \tag{27}$$

which, by strict concavity, would imply $\varphi\left(p^M(\underline{a})\right) < p^M(\underline{a})$. $\widetilde{\phi}\left(a, p^M(\underline{a})\right) < p^M(a)$ and $\partial^2\pi_2/\partial p_2\partial p_1 > 0$ imply

$$\frac{\partial\pi_2\left(\widetilde{\phi}\left(a, p^M(\underline{a})\right), p^M(\underline{a}), a\right)}{\partial p_2} < \frac{\partial\pi_2\left(p^M(a), p^M(\underline{a}), a\right)}{\partial p_2}. \tag{28}$$

Thus, if

$$\frac{\partial\pi_2\left(p^M(a), p^M(\underline{a}), a\right)}{\partial p_2} < 0 \tag{29}$$

then

$$\frac{\partial\pi_2\left(\widetilde{\phi}\left(a, p^M(\underline{a})\right), p^M(\underline{a}), a\right)}{\partial p_2} < 0. \tag{30}$$

Given

$$\frac{\partial\pi_2\left(p^M(\underline{a}), p^M(\underline{a}), \underline{a}\right)}{\partial p_2} < 0 \tag{31}$$

then, by continuity of $\partial\pi_2/\partial p_2$ in p_1 and a, if $a \simeq \underline{a}$ then (31) implies (29) and, therefore, (30) holds. In sum, if $|\overline{a} - \underline{a}|$ is sufficiently small then

$$\frac{\partial\pi_2\left(\widetilde{\phi}\left(a, p^M(\underline{a})\right), p^M(\underline{a}), a\right)}{\partial p_2} < 0\ \forall a \in [\underline{a}, \overline{a}]. \tag{32}$$

It follows that (27) is true. Hence, if $\widetilde{p} = p^M(\underline{a})$ then, for $|\overline{a} - \underline{a}|$ is sufficiently small, $\widetilde{\phi}\left(a, p^M(\underline{a})\right) < p^M(a)$. By (27) and strict concavity, the optimal value of p_2 is less than $p^M(\underline{a})$; that is, $\varphi\left(p^M(\underline{a})\right) < p^M(\underline{a})$.

In sum, if $|\overline{a} - \underline{a}|$ is sufficiently small then $\varphi\left(p^N(\underline{a})\right) > p^N(\underline{a})$ and $\varphi\left(p^M(\underline{a})\right) < p^M(\underline{a})$ which implies there is a fixed point in $\left(p^N(\underline{a}), p^M(\underline{a})\right)$. ■

Lemma 4 *If $\widetilde{\phi}(a,\widetilde{p}) < p^M(a)$ then $\widetilde{\phi}(a,\widetilde{p})$ is increasing in β.*

The first-order condition for

$$\widetilde{\phi}(a,\widetilde{p}) = \arg\max_p (1-\beta)\pi_1(p,\widetilde{p},a) + \beta\pi_1(p,p,a), \forall a \in A$$

is:

$$(1-\beta)\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{p},a\right)}{\partial p_1} + \beta\left(\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_1} + \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_2}\right) = 0,$$

which is equivalent to

$$-\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{p},a\right)}{\partial p_1} = \left(\frac{\beta}{1-\beta}\right)\left(\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_1} + \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_2}\right). \tag{33}$$

Consider the cross-partial derivative:

$$\frac{\partial^2\left((1-\beta)\pi_1(p,\widetilde{p},a) + \beta\pi_1(p,p,a)\right)}{\partial p \partial \beta} = \frac{\partial \pi_1(p,p,a)}{\partial p_1} + \frac{\partial \pi_1(p,p,a)}{\partial p_2} - \frac{\partial \pi_1(p,\widetilde{p},a)}{\partial p_1},$$

evaluate it at $p = \widetilde{\phi}(a,\widetilde{p})$, and use (33):

$$\begin{aligned}
&\frac{\partial^2\left((1-\beta)\pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{p},a\right) + \beta\pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)\right)}{\partial p \partial \beta} \\
= \;& \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_1} + \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_2} - \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{p},a\right)}{\partial p_1}, \\
= \;& \left(\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_1} + \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_2}\right)\left(\frac{1}{1-\beta}\right).
\end{aligned} \tag{34}$$

If $\widetilde{\phi}(a,\widetilde{p}) < p^M(a)$ then, by strict concavity of $\pi_1(p,p,a) + \pi_2(p,p,a)$,

$$\frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_1} + \frac{\partial \pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)}{\partial p_2} > 0$$

which, along with (34), implies

$$\frac{\partial^2\left((1-\beta)\pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{p},a\right) + \beta\pi_1\left(\widetilde{\phi}(a,\widetilde{p}),\widetilde{\phi}(a,\widetilde{p}),a\right)\right)}{\partial p \partial \beta} > 0. \tag{35}$$

Suppose $\beta = \beta'$. By strict concavity of $(1-\beta')\pi_1(p,\widetilde{p},a) + \beta'\pi_1(p,p,a)$ in p and that $\widetilde{\phi}(a,\widetilde{p},\beta')$ is the optimum, it follows:

$$\frac{\partial\left((1-\beta')\pi_1(p,\widetilde{p},a) + \beta'\pi_1(p,p,a)\right)}{\partial p} \gtreqqless 0 \text{ as } p \lesseqqgtr \widetilde{\phi}(a,\widetilde{p},\beta'). \tag{36}$$

Using (35), we have

$$\frac{\partial^2 \left((1-\beta')\pi_1\left(\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),\widetilde{p},a\right)+\beta'\pi_1\left(\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),a\right)\right)}{\partial p\partial\beta}>0. \tag{37}$$

(36)-(37) imply $\exists\varepsilon>0$ such that

$$\frac{\partial\left((1-\beta'-\varepsilon)\pi_1\left(p,\widetilde{p},a\right)+\left(\beta'+\varepsilon\right)\pi_1\left(p,p,a\right)\right)}{\partial p}<0\ \forall\ p\leq\widetilde{\phi}\left(a,\widetilde{p},\beta'\right)$$

which implies $\widetilde{\phi}\left(a,\widetilde{p},\beta'+\varepsilon\right)>\widetilde{\phi}\left(a,\widetilde{p},\beta'\right)$. Hence, if $\widetilde{\phi}\left(a,\widetilde{p}\right)<p^M(a)$ then $\widetilde{\phi}\left(a,\widetilde{p}\right)$ is increasing in β. ■

From Lemmas 3 and 4, it follows: if $|\overline{a}-\underline{a}|$ is sufficiently small then there is a minimal equilibrium with $\widehat{p}<p^M(\underline{a})$ and $\exists\varepsilon>0$ such that $\widetilde{\phi}\left(a,\widetilde{p}\right)$ is increasing in β $\forall\widetilde{p}<\widehat{p}+\varepsilon$. Suppose $|\overline{a}-\underline{a}|$ is sufficiently small and let $\left(\widehat{\phi}\left(\cdot,\beta'\right),\widehat{p}(\beta')\right)$ be a minimal equilibrium when $\beta=\beta'$. As it is a minimal equilibrium then

$$\varphi(\widetilde{p},\beta')\gtreqqless\widetilde{p}\text{ as }\widetilde{p}\lesseqqgtr\widehat{p}(\beta'). \tag{38}$$

Recall that

$$\varphi(\widetilde{p},\beta)=\arg\max_{p_2}\int\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta\right),p_2,a\right)F'(a)da$$

which means, by strict concavity,

$$\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta\right),p_2,a\right)}{\partial p_2}\right)F'(a)da\gtreqqless 0\text{ as }p_2\lesseqqgtr\varphi(\widetilde{p},\beta). \tag{39}$$

Given $\widetilde{p}<\widehat{p}(\beta')$ implies $\varphi(\widetilde{p},\beta')>\widetilde{p}$ by (38), it follows from (39):

$$\text{if }\widetilde{p}<\widehat{p}(\beta')\text{ then }\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),p_2,a\right)}{\partial p_2}\right)F'(a)da>0\ \forall p_2\leq\widetilde{p}(<\varphi(\widetilde{p},\beta')).$$

Hence,

$$\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),\widetilde{p},a\right)}{\partial p_2}\right)F'(a)da\gtreqqless 0\text{ as }\widetilde{p}\lesseqqgtr\widehat{p}(\beta'). \tag{40}$$

As $\partial\pi_2/\partial p_2$ is increasing in p_1 and $\widetilde{\phi}\left(a,\widetilde{p},\beta\right)$ is increasing in β $\forall\widetilde{p}<\widehat{p}\left(\beta'\right)+\varepsilon$ (where $\varepsilon>0$ and small) then $\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta\right),p_2,a\right)/\partial p_2$ is increasing in β $\forall\widetilde{p}<\widehat{p}\left(\beta'\right)+\varepsilon$. Hence, if $\beta''>\beta'$ then

$$\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta''\right),\widetilde{p},a\right)}{\partial p_2}\right)>\int\left(\frac{\partial\pi_2\left(\widetilde{\phi}\left(a,\widetilde{p},\beta'\right),\widetilde{p},a\right)}{\partial p_2}\right)\ \forall\widetilde{p}<\widehat{p}\left(\beta'\right)+\varepsilon. \tag{41}$$

(40)-(41) imply

$$\int \left(\frac{\partial \pi_2 \left(\widetilde{\phi}\left(a, \widetilde{p}, \beta''\right), \widetilde{p}, a \right)}{\partial p_2} \right) F'(a)da > 0 \ \forall \ \widetilde{p} \leq \widehat{p}(\beta'), \tag{42}$$

and, therefore, $\widehat{p}\left(\beta''\right) > \widehat{p}\left(\beta'\right)$. The minimal equilibrium is then increasing in β. Given $\widehat{\phi}\left(a, \beta\right) = \widetilde{\phi}\left(a, \widehat{p}\left(\beta\right), \beta\right)$ and $\widetilde{\phi}\left(a, \widetilde{p}, \beta\right)$ is increasing in $\widetilde{p}$ and β then $\widehat{p}\left(\beta''\right) > \widehat{p}\left(\beta'\right)$ implies

$$\left(\widehat{\phi}\left(a, \beta''\right) = \right) \widetilde{\phi}\left(a, \widehat{p}\left(\beta''\right), \beta''\right) > \widetilde{\phi}\left(a, \widehat{p}\left(\beta'\right), \beta'\right) \left(= \widehat{\phi}\left(a, \beta'\right) \right).$$

Therefore, $\widehat{\phi}\left(a, \beta''\right) > \widehat{\phi}\left(a, \beta'\right)$. ■

9 Appendix B: Solving for the Equilibrium in the Linear Demand Case

$\left(\widehat{\phi}\left(\cdot\right), \widehat{p} \right)$ are defined by:

$$\widehat{\phi}\left(a\right) = \arg\max_{p_1} (1-\beta)(p_1 - c)\left(a - bp_1 + d\widehat{p}\right) + \beta(p_1 - c)(a - (b-d)p_1), \ \forall a \in A \tag{43}$$

$$\widehat{p} = \arg\max_{p_2} \int (p_2 - c) \left(a - bp_2 + d\widehat{\phi}\left(a\right) \right) F'(a)da. \tag{44}$$

Using the first-order condition to solve (43) yields:

$$p_1 = \frac{(b(1-\beta) + \beta(b-d))c + d(1-\beta)\widehat{p}}{2(b - d\beta)} + \left(\frac{1}{2(b-d\beta)} \right) a = \gamma\left(\widehat{p}\right) + \theta a \tag{45}$$

where

$$\begin{aligned} \gamma\left(\widehat{p}\right) &\equiv \frac{(b(1-\beta) + \beta(b-d))c + d(1-\beta)\widehat{p}}{2(b-d\beta)} \\ \theta &\equiv \frac{1}{2(b-d\beta)}. \end{aligned}$$

Turning to (44), I can use (45) and the expectation of a so as to represent the objective function as

$$(p_2 - c)(\mu - bp_2 + d(\gamma\left(\widehat{p}\right) + \theta\mu)).$$

Maximizing it yields the optimum:

$$p_2 = \frac{\mu + bc + d(\gamma\left(\widehat{p}\right) + \theta\mu)}{2b}. \tag{46}$$

Using (45)-(46), $\widehat{p}$ is defined as the fixed point:

$$\widehat{p} = \frac{\mu + bc + d(\gamma\left(\widehat{p}\right) + \theta\mu)}{2b}$$

which yields:

$$\widehat{p} = \frac{\mu(2b - d + 2(1-\beta)d) + c(b - \beta d)(2b + d)}{4b^2 - d^2 - \beta(4b - d)d}. \tag{47}$$

Finally, insert (47) into (45) and re-arrange:

$$\begin{aligned} \widehat{\phi}(a) &= \widehat{\gamma} + \widehat{\theta} a \\ &= \frac{\mu(1-\beta)d(2\,(b-\beta d)+d) + c(b-\beta d)\,(4b(b-\beta d) + 2bd(1-\beta))}{(b - d\beta)\,(8b(b-d) + 2d(1-\beta)(4b-d))} + \left(\frac{1}{2(b-\beta d)}\right) a \end{aligned}$$

Note that this solution is valid as long as $|\overline{a} - \underline{a}|$ is not too great, so firms' demands are always positive at the equilibrium prices.

References

[1] Acquisti, Alessandro and Hal R. Varian, "Conditioning Prices on Purchase History," *Marketing Science*, 24 (2005), 367-381.

[2] Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, "The Limits of Price Discrimination," *American Economic Review*, 105 (2015), 921-957.

[3] Bernheim, B. Douglas and Michael D. Whinston, "Common Marketing Agency as a Device for Facilitating Collusion," *RAND Journal of Economics*, 16 (1985), 269-281.

[4] Brown, Zach Y. and Alexander MacKay, "Competition in Pricing Algorithms," Harvard Business School, April 2020.

[5] Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, Joseph E. Harrington, Jr., and Sergio Pastorello, "Protecting Consumers from Collusive Prices Due to AI," University of Bologna, October 2020a (*Science*, forthcoming).

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

[7] Chen, Zhijun, Chongwoo Choe, and Noriaki Matsushima, "Competitive Personalized Pricing," *Management Science*, Articles in Advance, 2020, pp. 1-21.

[8] Chena, Yuxin and Z. John Zhang, "Dynamic Targeted Pricing with Strategic Consumers," *International Journal of Industrial Organization*, 27 (2009), 43-50.

[9] Choudhary, Vidyanand, Anindya Ghose, Tridas Mukhopadhyay, and Uday Rajan, "Personalized Pricing and Quality Differentiation," *Management Science*, 51 (2005), 1015-1164."

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

[11] 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.

[12] Cowan, Simon, "Welfare-increasing Third-degree Price Discrimination," *RAND Journal of Economics*, 47 (2016), 326-340.

[13] Decarolis, Francesco, Maria Goldmanis, and Antonio Penta, "Marketing Agencies and Collusive Bidding in Online Ad Auctions," *Management Science*, Articles in Advance, 2020, 1-22.

[14] den Boer, Arnoud V., "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions," *Surveys in Operations Research and Management Science*, 20 (2015), 1-18.

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

[16] Dogan, Ibrahim and Ali R. Güner, "A Reinforcement Learning Approach to Competitive Ordering and Pricing Problem," *Expert Systems*, 32 (2015), 39-47.

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

[18] Gal, Michal S., "Algorithms as Facilitating Practices," *Berkeley Technology Law Journal*, 34 (2019), 67-118.

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

[20] Hansen, Karsten, Kanishka Misra, and Mallesh Pai, "Algorithmic Collusion: Supra-Competitive Prices via Independent Algorithms," Rice University, January 2020 (*Marketing Science*, forthcoming)

[21] Harrington, Joseph E., Jr., *The Theory of Collusion and Competition Policy*, Cambridge, Mass.: The MIT Press, 2017.

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

[23] Hilsen, Hans Olav Østbø, "Simulating Dynamic Pricing Algorithm Performance in Heterogeneous Markets," Masters Thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, June 2016.

[24] Klein, Timo, "Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing," Amsterdam Center for Law & Economics Working Paper No. 2018-05, April 2019.

[25] Lorenz Marx, Christian Ritz, and Jonas Weller, "Liability for Outsourced Algorithmic Collusion - A Practical Approximation," *Concurrences*, No. 2, 2019, 1-7.

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

[27] Miklós-Thal, Jeanine and Catherine Tucker, "Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?," *Management Science*, 4 (2019), 1552-1561.

[28] Milgrom, Paul and John Roberts, "Comparing Equilibria," *American Economic Review*, 84 (1994), 441-459.

[29] O'Connor, Jason and Nathan E. Wilson, "Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Collusion," FTC Bureau of Economics, Working Paper No. 341, June 2019.

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

[31] OECD, "Algorithms and Collusion - Note by the United States," DAF/COMP/WD(2017)41, 26 May 2017b.

[32] Salcedo, Bruno, "Pricing Algorithms and Tacit Collusion," University of Western Ontario, November 2015.

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

[34] Tarski, Alfred, "A Lattice-Theoretical Fixpoint Theorem and its Applications," *Pacific Journal of Mathematics*, 5 (1955), 285-309.

[35] Tesauro, Gerald and Jeffrey O. Kephart, "Pricing in Agent Economics Using Multi-Agent Q-Learning," *Autonomous Agents and Multi-Agent Systems*, 5 (2002), 289-304.

[36] Varian, Hal R., "Price Discrimination," in *Handbook of Industrial Organization*, Volume 1, Amsterdam: Elsevier, 1989.

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

[38] Xie, Ming and Jian Chen, "Studies on Horizontal Competition Among Homogeneous Retailers through Agent-based Simulations," *Journal of Systems Science and Systems Engineering*, 13 (2004), 490-505.

[39] Zhang, Juanjuan, "The Perils of Behavior-Based Personalization," *Marketing Science*, 30 (2011), 170-186.