Competition in Pricing Algorithms*

Zach Y. Brown
University of Michigan†

Alexander MacKay
Harvard University‡

January 15, 2019

PRELIMINARY AND INCOMPLETE:
NOT FOR GENERAL DISTRIBUTION

Abstract

Increasingly, retailers have access to better pricing technology, especially in online markets. In particular, retailers may determine prices with a pricing algorithm. When firms choose algorithms, competitive equilibria can have higher prices than in the Bertrand game, in which firms choose prices. Using hourly prices of over-the-counter allergy medications from five major online retailers, we document evidence that these retailers possess different pricing technologies. In addition, we observe pricing patterns that are consistent with the model. Overall, our analysis suggests that pricing algorithms may lead to higher prices, even in the absence of collusion.

*We thank seminar participants at Harvard Business School for constructive comments.
†University of Michigan, Department of Economics. Email: zachb@umich.edu.
‡Harvard University, Harvard Business School. Email: amackay@hbs.edu.

1 Introduction

There is a growing concern that pricing algorithms, which are becoming more prevalent, will lead to higher prices to consumers. This concern has been expressed by the popular press in terms of enabling firms to collude on prices and form cartels.[1] Indeed, whether or not pricing algorithms can better enable collusion has been the focus of the economics literature as well.[2]

In this paper, we take a different approach to the topic of pricing algorithms. Instead of examining the effect of pricing algorithms on collusion, we investigate how the presence of pricing algorithms affects *competitive* outcomes. We define a competitive game where firms choose algorithms, rather than prices, and we show that this game tends to result in higher prices in equilibrium. As a particular result, we show that it is not an equilibrium for all firms to choose algorithms that reflect their price-setting best-response (Bertrand reaction) functions.

To support our theoretical analysis, we provide an empirical analysis of prices for five large online retailers. We document heterogeneity in pricing technology—in particular, firms vary in their capabilities to update prices—and we show that the price differences we observe in the data are consistent with the theory. We find that firms that have higher-frequency pricing have lower prices than their competitors. We believe this is the first paper to document these empirical findings.

To explain these results, it is helpful to first define what exactly an algorithm is. An algorithm is *a set of instructions to perform a calculation given an initial state.* A pricing algorithm, therefore, is a formula to determine prices, as a function of other (observable) variables that define the state.

A pricing algorithm, compared to (human) agent, has three significant features:

1. *An algorithm lowers the cost of sophistication in pricing behavior.*
   In principle, a pricing algorithm can incorporate all of the knowledge that would be available to the most sophisticated price-setting agent. Algorithms may enable more sophisticated price-setting behavior than a single human agent through the ability to combine information more efficiently and from multiple sources.

2. *An algorithm lowers the cost of updating prices more frequently.*
   By writing the set of instructions to software, the solution to a difficult pricing problem can be found in less time and with less error than if done by a human agent.

3. *An algorithm provides a (short-term) commitment device.*
   When an algorithm calculates price based on the prices of other firms, it can react quickly

[1] See, for example, "Price-Bots Can Collude Against Consumers," The Economist, May 6, 2017. "When Bots Collude," The New Yorker, April 25, 2015. "How Pricing Bots Could Form Cartels and Make Things More Expensive," Harvard Business Review, October 27, 2016. "Policing the Digital Cartels," Financial Times, January 8, 2017. We thank Schwalbe (2018) for this list of articles.

[2] Recent papers on the subject include Miklos-Thal and Tucker (2018), Schwalbe (2018), and Salcedo (2015).

to price changes in the market. This becomes a short-term commitment device, as the algorithm is typically updated at a lower frequency than it is used to set prices.

We seek to understand whether pricing technology, in terms of more sophisticated and higher-frequency pricing algorithms, provides a firm with an advantage relative to its peers, and also what are the implications for competitive outcomes. The analysis of algorithms undertaken in the economics literature have primarily focused on the feature of increased sophistication. Economists have argued that increasing sophistication can be used to support collusion by encoding cooperation into the algorithm, either implicitly (Salcedo, 2015) or explicitly (Tennenholtz, 2004). Miklos-Thal and Tucker (2018) show that increased sophistication, in terms of better demand forecasting, can lead to lower prices.

Our focus is, instead, on the latter two features. One reason to abstract away from increased sophistication is because the economics literature often assumes that agents are maximally sophisticated, especially when it comes to the analysis of game-theoretic equilibrium. Here, we hope to highlight the equilibrium outcomes when no profitable deviations exist.

Moreover, we find that the latter two features alone can lead to increased prices in equilibrium. We develop these results using a simple spatial differentiation model, which we introduce in Section 2. First, variation in pricing technology can lead to asymmetries in terms of pricing frequency. Asymmetric pricing frequency, when products are substitutes and strategic complements (in prices), leads to higher prices. In effect, when firms have asymmetries in pricing frequency, the game closely resembles a Stackelberg leader-follower pricing game, which has higher prices than the Bertrand equilibrium with differentiated products.

Second, we show that the short-term commitment feature can also generate higher prices. To do so, we introduce a one-shot competitive game in which firms submit pricing algorithms, rather than prices. To provide a conservative analysis of the effect of algorithms on prices, we introduce two restrictions on the set of equilibrium strategies. First, the authority selects the profit-minimizing solution to the system of equations implied by the algorithms, if multiple solutions exist. Thus, the authority acts in favor of consumers to discipline prices. Second, the authority insists that the algorithms are continuous, therefore ruling out obvious punishment strategies.

Even with these restrictions, which result in reasonable-looking pricing strategies in equilibrium, competition in pricing algorithms can support higher prices than in the Bertrand price-setting game. In our spatial differentiation example, the collusive equilibrium is supported by firms submitting symmetric (continuous) algorithms. In rough terms, these algorithms serve as as short-term commitment device that enable firms to internalize the effect of a change in its strategy on the behavior of its rivals. The resulting game resembles competition with conjectural variations. Competition with pricing algorithms can support any conjectural variation equilibrium, even when the equilibrium would not satisfy the consistency condition of conjectural variations (were firms instead setting prices with consistent conjectures). Pricing algorithms provide an economic mechanism to assure the consistency between beliefs and behavior for a multitude of equilibria.

Indeed, one interpretation for this paper is to provide an economic mechanism to generate outcomes (or equilibria) from models that were previously difficult to reconcile with real-world pricing behavior. Perhaps surprisingly, these results do not depend on folk theorem style arguments that are used to support various equilibria in repeated games. Our analysis provides subgame perfect results for the one-shot game.

In our empirical analysis, we study a large dataset of prices for over-the-counter allergy medications for five large online retailers. By studying prices at the hourly level, we are able to document heterogeneity in pricing technology. We find that two firms have hourly pricing technology, one firm has daily pricing technology, and the remaining two have weekly pricing technology, updating their prices early every Sunday morning. This high-degree of asymmetry is correlated with asymmetric prices. Relative to the firm with the most flexible pricing technology, the firm with daily pricing technology sells the same products at prices that are 10 percent higher, whereas the firms with weekly technology sell those products at prices that are over 30 percent higher.

Intuitively, these outcomes are supported by the following logic: The high-frequency firm commits to "beat" (best respond) to whatever price is offered by its low-frequency rivals, and its technology enables it to commit to being a follower. Since this commitment is credible, the low-frequency rivals price less aggressively.

Overall, we find that pricing algorithms facilitate higher-price equilibria, even when firms act competitively. The empirical evidence suggests that these algorithms are becoming more widespread (Cavallo, 2018), and at least some firms include rivals' prices in their algorithms.[3] Further, our empirical analysis shows that the price patterns were observe in the data are at least consistent with model we analyze. Thus, if policymakers are concerned that algorithms will raise prices, then the concern is much more broad than that of collusion.

There exists a simple policy prescription if policymakers are interested in keeping prices closer to the Bertrand equilibria: policymakers could insist that firms cannot make their pricing algorithms functions of rivals' prices. Firms would remain free to have frequent price updates as a function of other factors, such as demand shocks, but explicitly accounting for rivals prices would be forbidden. As we show, algorithms that incorporate rivals' prices can support the collusive outcome in equilibrium.

[3] See, for example, Amazon's "Match Low Price" feature it offers to third-party sellers, documented in the Competition & Markets Authority's 2018 report, "Pricing Algorithms."

2 Pricing Frequency, Pricing Algorithms, and Equilibrium

2.1 Model: Price Competition with Differentiated Products

Consider a simple spatial differentiation model. Firm 1 and firm 2 are located 2 units (e.g., miles) apart. A mass of 2 customers are distributed uniformly on a line segment connecting the two firms. Firms are symmetric and sell a single good at zero cost. They set prices at the beginning of the period. At the end of the period, consumers decide to travel to one of the two firms to buy the good, or to stay home. The game ends after one period (e.g., day).

Consumers receive utility α from the good and have disutility of τD for the distance D they travel to purchase. Utility is linear in income and is normalized so that the marginal utility of income is 1. Consumers have type θ, which indexes their location at the beginning of the period. Thus, the utility u_{ia} for consumer i of choosing action $a \in \{1, 2, 0\}$—purchasing from firm 1, firm 2, or not purchasing (0)—is given by:

$$u_{i1} = \alpha - \tau\theta_i - p_1 \tag{1}$$

$$u_{i2} = \alpha - \tau(2 - \theta_i) - p_2 \tag{2}$$

$$u_{i0} = 0 \tag{3}$$

Consumer i will prefer firm 1 over firm 2 iff:

$$1 - \frac{1}{2\tau}(p_1 - p_2) > \theta_i. \tag{4}$$

We consider the following utility parameters: $\alpha = 2$ and $\tau = \frac{1}{2}$. Where the utility from both goods is positive, the (local) demand for each good is:

$$q_1 = 1 - p_1 + p_2 \tag{5}$$

$$q_2 = 1 - p_2 + p_1. \tag{6}$$

Price-Setting Equilibrium

Firms set prices to maximize profits. Conditional on the price of the other firm, the first-order conditions for optimality yield the following best-response (or "reaction") functions:

$$R_1(p_2) = \frac{1}{2}(1 + p_2) \tag{7}$$

$$R_2(p_1) = \frac{1}{2}(1 + p_1). \tag{8}$$

Thus, the Bertrand-Nash equilibrium of the game is $p_1 = p_2 = 1$, yielding $q_1 = q_2 = 1$ and profits $\pi_1 = \pi_2 = 1$. Firms split the consumers equally, and a consumer in the middle ($\theta_i = 1$)

receives positive utility $u_{11} = u_{21} = \frac{1}{2}$. That is, all consumers purchase a good.

In this game, the collusive outcome is one in which prices are $p_1 = p_2 = \frac{3}{2}$. At these prices, industry profits are maximized at $\pi_1 + \pi_2 = 3$. At prices above this level, some consumers opt to stay home, and the marginal loss on these consumers is higher than the inframarginal gain of higher prices. As we focus on a one-shot game, the collusive outcome is not an (subgame perfect) equilibrium.

2.2 Price Competition and Pricing Frequency

Now suppose that firm 2 adopts a new pricing technology. This technology allows the firm to update its price after some interval, ε, but before the end of the period. Firm 1, who does not use the new pricing technology, continues to set price only at the beginning of the period.

What happens to equilibrium prices?

First, suppose that firm 2 adopts and announces the technology after the "price-setting" phase at the beginning of the period. In this case, there is no effect on prices. Firm 1 and firm 2 initially play the Bertrand-Nash equilibrium, $p_1 = p_2 = 1$. After firm 2 adopts the technology, it remains firm 2's best response to keep the price at $p_2 = 1$. The outcome remains unchanged.

Suppose instead that firm 2 adopts and announces the technology before the price-setting phase. Firm 1 now considers the impact of the technology on the pricing behavior of its competitor. Firm 1 knows that firm 2 will have the ability to update its price before demand is realized (at the end of the period).

In this case, firm 1 recognizes that, whatever price it chooses, the outcome will lie along firm 2's best-response function (the optimal "reaction" function). Firm 1 can now choose the price that maximizes its own profits *conditional on firm 2's best-response function*. Firm 1 can make higher profits by choosing a higher price. In this example, firm 1's optimum strategy is to price at $\frac{3}{2}$. This leads Firm 2 to price at $\frac{5}{4}$. Quantities are $(\frac{3}{4}, \frac{5}{4})$, and profits are $(\frac{9}{8}, \frac{25}{16})$, which give both firms higher profits than the Bertrand-Nash solution.

What's the intuition here? Firm 2 will "undercut" the price of firm 1, maximizing its own profits conditional on its rival's price. The Bertrand-Nash logic says that firm 1 would to react by lowering prices, which would then induce a reaction by firm 2, and so on until the Bertrand-Nash equilibrium is obtained. Though both firms may recognize that they would be better off by not undercutting the competitor, they cannot credibly commit not to (especially in a static game). However, since firm 2 is able to undercut firm 1's price through more frequent pricing, firm 1 is able to internalize firm 2's reaction and maintain prices that are above the Bertrand competitive equilibrium.

This outcome mirrors results for Stackelberg price competition with differentiated products. A typical critique of the Stackelberg pricing model is that it is difficult to reconcile with real-world behavior. Real firms are able to change their price more than once, making it difficult for the Stackelberg leader to commit to not changing its price in response to the follower. For example, if firms take turns setting prices for a finite number of periods, then the Stackelberg equilibrium unravels and Bertrand-Nash prices result.

We argue that asymmetries in pricing frequency create a natural Stackelberg game. Firms are able to commit to a leader-follower order via asymmetries in technology that determine pricing frequency. Here, asymmetry is essential to generating higher prices. If firm 1 adopts technology that enables it to update prices at the same frequency as firm 2, then the equilibrium prices return to the Bertrand-Nash equilibrium.[4] When firm 2 alone has higher-frequency pricing technology, then it has the capability to update prices without a response by firm 1.

Thus, when firms can choose whether to adopt higher-frequency pricing, an asymmetric equilibrium results. Both firms are better off if only one firm adopts the technology, but they would prefer themselves to be the one to adopt.[5] The resulting technology adoption problem is the classic coordination game. Therefore, it is not surprising that asymmetric pricing technologies might arise endogenously in equilibrium, as we find in our data (Section 4).

One can consider the impact of pricing frequency itself, rather than the asymmetry alone, by modifying the model so that consumers decide to purchase throughout the period, rather than at the end. In this case, firm 1 will internalize the realized profit before firm 2 has a chance to update its prices. For a small ϵ, the initial phase provides an infinitesimal share of profits, and the outcome resembles the game described above, with $(p_1, p_2) = (\frac{3}{2}, \frac{5}{4})$. The longer the interval ϵ, the more weight firm 1 will put on this initial phase, and the lower the equilibrium prices, approaching the Bertrand equilibrium. Thus, equilibrium prices lie somewhere between $(p_1, p_2) = (1, 1)$ and $(p_1, p_2) = (\frac{3}{2}, \frac{5}{4})$, as a function of how frequently firm 2 can update its price.

We conclude this section by showing that higher prices resulting from asymmetric pricing frequency are a general result for a large class of problems. Consider a typical case where the products are substitutes (i.e., $\frac{\partial q_1}{\partial p_2} > 0$) and prices are strategic complements (with upward-sloping best-response functions in the price-setting game, $\frac{\partial R_2}{\partial p_1} > 0$).

Proposition 1 *When firms produce substitute goods and prices are strategic complements, then, when firms have asymmetries in pricing frequency, both firms realize higher prices compared to the price-setting (Bertrand-Nash) equilibrium.*

Proof: Denote the firm with less-frequent pricing firm 1. Consider its first-order condition to maximize profits (π):

$$\frac{d\pi_1}{dp_1} = \frac{\partial \pi_1}{\partial p_1} + \frac{\partial \pi_1}{\partial p_2}\frac{\partial p_2}{\partial p_1} = 0 \tag{9}$$

In the one-shot price-setting equilibrium, firm 1 takes firm 2's price as given ($\frac{\partial p_2}{\partial p_1} =$

[4] This is true regardless of whether they set prices simultaneously or in alternating fashion.

[5] This is generalized by the results of Gal-Or (1985), who shows that a follower is better off when the best-response functions are upward-sloping.

0), and $\frac{\partial \pi_1}{\partial p_1} = 0$. When firm 1 accounts for firm 2's technology, firm 1 recognizes that $\frac{\partial p_2}{\partial p_1} = \frac{\partial R_2}{\partial p_1} > 0$ (by strategic complementarity) and $\frac{\partial \pi_1}{\partial p_2} > 0$ (because the products are substitutes). Therefore, relative to the Bertrand-Nash prices, firm 1 has an incentive to raise its price: $\frac{d\pi_1}{dp_1} > 0$. Higher prices for both firms result from strategic complementarity.

2.3 Competition in Pricing Algorithms

The previous section discussed outcomes in which firms have asymmetries in pricing frequency. We believe the above model captures one of the essential features of pricing algorithms in the real world: namely, the ability to update prices on a more frequent basis.

Nevertheless, we also think it is valuable to characterize a one-shot game of a different nature. Instead of choosing prices, suppose that each firm chose a formula for prices—an "algorithm"—that calculates its price as a function of other variables, including (potentially) the prices of rival firms. The asymmetric game, in which one firm submits an algorithm and the other firm submits a price, is a natural extension of the pricing frequency game described above. In many respects, it resembles that game when the interval between pricing updates goes to zero ($\varepsilon \to 0$). Our discussion is in terms of a two-player game, but the concepts readily generalize. Throughout, we focus our attention on pure strategies.

Definition and Equilibrium Concept

We now define a competitive game—*competition in pricing algorithms*—and its equilibrium concept. Firms compete in pricing algorithms by submitting a pricing strategy $\sigma(\cdot)$, or "algorithm", to a central coordinator. The algorithm may be a function of variables that are observable to the firm, including rival prices, but are not functions of other player's algorithms.

We focus on a one-shot game to illustrate that the equilibrium outcomes do not depend on collusive behavior (Miklos-Thal and Tucker, 2018) or folk-theorem style results from dynamic games. Higher prices are a competitive outcome. Indeed, our results have a flavor of dynamic games and conjectural variations, as the algorithms map closely to dynamic considerations. We discuss these comparisons in more detail in Section 2.5.

After receiving the pricing algorithms, the central coordinator solves the system of equations set by the algorithms to determine prices. Of course, without further restrictions, the game thus far described may suffer from an *indeterminacy problem*: there may be multiple solutions to the system of equations set by the algorithms. To determine the equilibrium, we provide the coordinator with a simple selection rule. When multiple solutions are possible, the coordinator picks the solution that minimizes the profits of the firms. If multiple such solutions exist, the coordinator randomizes among them.

Restriction 1 (Profit-Minimizing Coordinator) *In the pricing algorithm game, the coordinator selects the solution to the system of equations set by the algorithms that minimizes the join profit of the firms.*

This selection rule is a natural choice for us because it provides "conservative" results regarding prices. We wish to show that it is possible for such a solution to generate higher prices than Bertrand-Nash price setting, so we focus on the least-profitable case. If firms submit algorithms that both have punishment features, the punishment outcome will be chosen by the coordinator.[6] In the real world, this selection procedure reflects pro-consumer market mechanisms to discipline firms.

Of course, even subject to the profit-minimizing coordinator, a multitude of equilibrium strategies are profitable. We add a second restriction to focus attention on a more narrow class of equilibrium strategies.

Restriction 2 (Continuity) *In the pricing algorithm game, the algorithms must be continuous. Otherwise, the firms receives zero profits (the coordinator shuts down the market).*

We motivate this restriction for the same reason as the profit-minimizing coordinator: any observed algorithm that has a discontinuity is likely to raise suspicions about market power and could potentially trigger antitrust scrutiny. This restriction rules out discontinuous punishment features from algorithms, including "take-it-or-leave-it" equilibria where one firm threatens to set a punishment price unless the other firm cooperates. Thus, we further restrict our attention to algorithms that may look reasonable (or "competitive").

We now define the equilibrium concept for the algorithm-setting game:

> **Equilibrium definition:** *When firms compete in pricing algorithms, an equilibrium of the game is one in which each firm's algorithm maximizes its own profit, conditional on the algorithms submitted by the other firms and subject to a coordinator that minimize the joint profits when multiple solutions to the algorithms exist.*

Note that any equilibrium of the pricing algorithm game has the following property: at the equilibrium, no firm can do better by submitting a single price, conditional on the algorithms of its rivals. Therefore, any equilibrium lies at the intersection of modified best-response functions for price, where the best-response functions take into a account the algorithms of the rivals.

Given the equilibrium concept, we now illustrate some of the similarities and differences to the pricing frequency game described in the previous section. Consider the model of demand from Section 2.1, and suppose that both firms have the capability to submit pricing algorithms. If firm 1 submits $\sigma_1(\cdot) = \frac{3}{2}$, and firm 2 submits $\sigma_2(p_1) = R_2(p_1)$, the solution is $(p_1, p_2) = (\frac{3}{2}, \frac{5}{4})$, as before. Neither firm can do better with a unilateral deviation. Thus, this asymmetric

[6]Further, this has a nice parallel to the minimax rule.

equilibrium—where one firm submits the price, and the other a function of that price—is an equilibrium of the game.

If firm 1 instead were to submit its best-response function from the price-setting game, $\sigma_1 = R_1(p_2)$, the unique solution is $(p_1, p_2) = (1, 1)$. Thus, as in Section 2.2, firm 1 can do strictly better by submitting $\sigma_1(\cdot) = \frac{3}{2}$ instead of $\sigma_1 = R_1(p_2)$. Therefore, $(\sigma_1, \sigma_2) = (R_1, R_2)$ is not an equilibrium of the algorithm-setting game. Even when firms are competing in algorithms, the algorithms will not reflect the price-setting best-response functions in equilibrium. Thus we obtain a negative result regarding the equilibria of the game:

Proposition 2 *When firms compete in a one-shot game by submitting pricing algorithms, it is (in general) not an equilibrium for each firm to submit their price-setting best-response function.*

> **Proof:** By the above reasoning, individual firms can realize a profitable deviation by submitting a price that lies along their rival's best-response function and results in greater profits to the firm. QED.

Interestingly, the Bertrand-Nash solution is still an equilibrium of the game. For example, $(p_1, p_2) = (1, 1)$ is an equilibrium when each firm decides to submit those prices and not have the algorithm be a function of the other firm's price (e.g., simple price-setting).

Thus, the Bertrand-Nash and Stackelberg price-setting solutions can be supported when firms submit pricing algorithms. What about the collusive outcome? Our restrictions rule out the typical strategies to sustain collusive behavior. Consider a case in which both σ_1 and σ_2 depend on the other firm's price. The usual cooperate or punish strategy is ruled out by the fact that the algorithms must be continuous. And, even if a continuous strategy could generate cooperate or punish as the two solutions, the profit-minimizing coordinator would select punishment.

However, a collusive equilibrium can be supported by algorithms that satisfy the two restrictions. For the model of demand in Section 2.1, the collusive equilibrium is $(p_1, p_2) = (\frac{3}{2}, \frac{3}{2})$. This outcome is an equilibrium with the following strategies:

$$\sigma_1(p_2) = 1 + \frac{1}{3}p_2 \tag{10}$$

$$\sigma_2(p_1) = 1 + \frac{1}{3}p_1. \tag{11}$$

It is straightforward to verify that, conditional on these algorithms, the collusive price maximizes profits for each firm.

Indeed, reasonable-looking algorithms are capable of supporting many different price outcomes as equilibria, including the collusive outcome. This result is reminiscent of competition in conjectural variations. Indeed, one interpretation for this paper is to provide a game-theoretic foundation for such analysis. When considering the technology available to each firm, algorithms act as a commitment device that can generate prices higher than the Bertrand-Nash prices, even in a one-shot game. We discuss this in greater detail in Section 2.5.

2.4 Asymmetric Competition in Pricing Algorithms

We now focus on a special subset of equilibria for the two-firm game in which firm 2 submits an algorithm that is a function of the other firm's price, and firm 1 does not. We call this game the "asymmetric game" to refer to the asymmetry in the nature of the algorithms. In general, firm 1 may have an algorithm of demand shocks and cost shocks, or other observable. In the absence of such features, its algorithm reduces to standard price-setting behavior. The firms may be asymmetric or symmetric, yet we still can obtain asymmetric equilibria (see the example above).

The asymmetric game is of particular interest because it reflect the pricing frequency game in Section 2.2. In the real world, there is asymmetry in the ability of firms to adjust prices, so a one-shot game in which both firms submit algorithms at the same time is not necessarily a realistic model. (Though this may accurately capture dynamic behavior in the real world.) For one-shot games, the asymmetric game reflects that one firm has better pricing technology vis-a-vis its rival. This more closely resembles firm behavior, and characterizing the equilibria in the game may be helpful to understand real-world phenomena.

In a two-firm asymmetric equilibrium, σ_2 depends on p_1, but σ_1 does not depend on p_2. In such an equilibrium, it is optimal for σ_2 to mirror firm 2's best-response function.

Lemma 3 *In an asymmetric algorithm-setting equilibrium, the firm that submits a price-dependent algorithm cannot do better than submitting its Bertrand best-response function as its algorithm. Therefore, one equilibrium of the game has that firm submitting its best-response function.*

The proof follows immediately. Of course, there are many possible equilibria where firm 2 has an algorithm that, local to the equilibrium, the algorithm maps to the best-response function. Since off-equilibrium play is not restricted, there are few limitations on how the algorithm looks away from the equilibrium.

Our second proposition is a result for the equilibrium where firm 2 submits its best-response function. This equilibrium and the Bertrand best-response functions are illustrated in Figure 1.

Proposition 4 *There exists an equilibrium to the asymmetric game in which one firm submits its best-response function as its algorithm. The other firm submits a price that maximizes its own profit along the best-response function of the other firm.*

It is apparent that no profitable deviation exists. This outcome mirrors that of Section 2.2. Indeed, we present our second result for this section as a corollary to Proposition 1:

Figure 1: Equilibrium in the Bertrand Game and the Asymmetric Pricing Algorithm Game

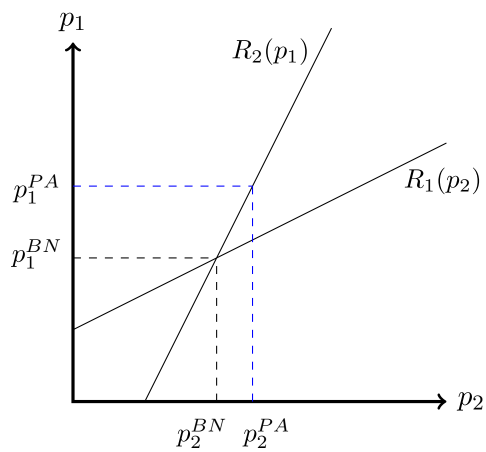

Notes: Figure plots the best-response functions $R_1(\cdot)$ and $R_2(\cdot)$ for price competition with differentiated products. The intersection of these functions produces the Bertrand-Nash (BN) equilibrium. PA indicates an equilibrium of the asymmetric pricing algorithm game where firm 2 submits $R_2(\cdot)$ as its algorithm.

Corollary 5 *When firms produce substitute goods and prices are strategic complements, then, in the asymmetric equilibrium where one firm submits its best-response function as its algorithm, both firms realize higher prices compared to the price-setting (Bertrand-Nash) equilibrium.*

2.5 Discussion

We have shown that asymmetries in pricing technologies—through pricing frequency or pricing algorithms—are sufficient to generate higher prices than the in the simultaneous price-setting equilibrium. Many of the equilibria generated from these supply-side assumptions mirror equilibria that arise in other models (i.e., Stackelberg price competition and competition with conjectural variations). Indeed, one interpretation of our paper is to highlight economic mechanisms that arise in the real world that have similar properties to those models. With asymmetry in pricing frequency, firms are committed to a pricing order, and the Stackelberg results are obtainable. Thus, technologies that result in different frequencies serve as a commitment device to attain the more profitable equilibrium. Moreover, firms prefer the asymmetry; if both firms realize the same pricing frequency, the Stackelberg result is lost, and profits fall. Thus, we should not be surprised to see asymmetries arise in the real world.

Our model of competition in pricing algorithms generates results that closely parallel the analysis of conjectural variations (Kamien and Schwartz, 1983). Competition with conjectural variations can support many different equilibria, including the collusive outcome. Often, the set of equilibria are restricted by imposing that the conjectural variations are consistent with the beliefs and actions of the other players. Under certain assumptions, consistent conjectures can

Table 1: Price Observations by Website and Brand

Retailer	Allegra	Benadryl	Claritin	Flonase	Xyzal	Zyrtec	Total
A	104,534	128,549	153,307	66,916	62,749	117,297	633,352
B	147,093	31,999	69,683	41,598	22,808	41,812	354,993
C	48,080	41,937	62,778	43,719	24,092	42,343	262,949
D	29,461	19,397	47,785	26,948	9,770	24,955	158,316
E	57,815	31,536	45,357	23,295	25,248	62,616	245,867
Total	386,983	253,418	378,910	202,476	144,667	289,023	1,655,477

generate a unique equilibrium (Bresnahan, 1981), as it would in the model of Section 2.1. This provides a notable divergence for our notion of competition in pricing algorithms. In the equilibria of the game in pricing algorithms, firm's beliefs are consistent with the pricing strategies (algorithms) played by other firms, yet *any* conjectural variation equilibrium may be supported, regardless of whether it would be consistent in the price-setting game. Thus, competition in algorithms expands the possibility of multiple equilibria in a way that is consistent with rational beliefs of other firms' behavior.

One can further understand our results by considering two different interpretations of the conjectural variations model. The first interpretation is to take seriously the one-shot nature of the model with conjectural variations. A criticism of this approach is that the resulting equilibria are not subgame perfect. That is, in a price-setting game, a firm would prefer to choose a different price, conditional on the strategy (price) of the other firm. By contrast, the strategy in the algorithm game is a function. Thus, firms are committed to responding to a deviation from their rivals, and the resulting equilibria are subgame perfect.

The second interpretation is to consider the conjectural variations model as an approximation to a repeated dynamic game (Holt, 1985). In repeated dynamic games, many possible equilibria may be supported (e.g. Fudenberg and Maskin, 1986; Benoit and Krishna, 1985). These folk theorem arguments may be translated to conjectural variations models and used to justify the stability of the equilibria. An important distinction arising in our model is that no notion of dynamics are needed. The supported equilibria are subgame perfect and outcomes of a one-shot game.

3 Data

We collected hourly price data for over-the-counter allergy drugs from five large online retailers. In this working paper, we have kept the identities of the retailers anonymous, calling them *A*, *B*, *C*, *D*, and *E*. We hope to be able to release the identities in a future version. For each of these retailers, allergy drugs represent an important product category. All five retailers sell products in many other categories, and four of the five have a large in-store presence in addition to their online channel.

For this analysis, we focus on the six brands of allergy drugs that are sold by all five retailers: Allegra, Benadryl, Claritin, Flonase, Xyzal, and Zyrtec. We define a product to be a drug-brand-form-size combination, e.g. Loratadine-Claritin-Tablet-20. Each of the retained brands specializes in one drug, but they often offer the products in multiple forms (e.g., Liquid Gels, Liquid, or Tablets). Each brand offers many different size options, so there are several products per brands. When a retailer sells multiple versions of the same product, we select the most popular version by retaining the version that has the greatest number of reviews, on average, in our sample. Our dataset spans April 22, 2018 through January 4, 2018, resulting in 1,655,477 price observations across the five websites. See Table 1 for a tabulation.

Obtaining online prices can be challenging, as updates to price information may take a while to propagate through the network, retailers can have complicated websites that take time to load, and the websites tend to change over time. These features are reflected in our raw data, and we have taken steps to eliminate measurement error. First, we have focused on a subset of brands that are high-volume brands, helping to ensure the availability of price information. Second, we use supplemental information obtained at the time of our price sample to rule out price changes brought about by a lag in the website. For example, we can see if the description of the product is consistent over time. Third, we impute missing prices by filling in missing prices with the most recently observed price, but only if we have another price observation on the same date. That is, if we do not observe a price on date for a particular product on a particular website, it is dropped from our sample. Finally, for the three retailers that do not change prices hourly, we smooth over single-period blips in price that revert back to the earlier price.[7]

Figure 2 displays the count of products in our sample over time. With the exception of retailer *B* and a few short-lived blackout periods, the count of products we observe each day remains roughly constant over time. These blackout periods do not meaningfully affect our results. We condition on the number of observed products for most of our baseline statistics, and they are not sensitive to the product set. We also perform robustness checks using only data on or after October 1, 2018, when we have a more consistent panel. Retailers *A* and *B* offer significantly more product varieties than the other retailers. Again, this is primarily due to the number of size options offered for each brand.

Summary statistics for our data are presented in Table 2. On average, we observe 98 products each day for retailer *A*, compared to 55 products for retailer *B*. Across all retailers, we observe the price for each product in 19.6 out of 24 hours on average. The mean price for these products is \$24.09, with a mean (absolute) price change of \$2.20. The table indicates stark differences in the frequency of price changes. Retailer *A* changes the prices of approximately half of the products in a given day, with an average count of 2.62 price changes per product.

[7] In our analysis sample, 282,332 observations out of 1,655,477 are imputed.

Figure 2: Observed Products Over Time

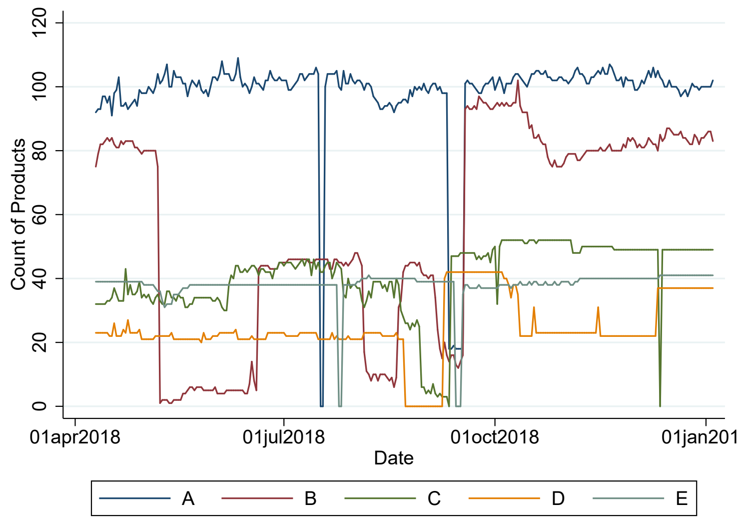

Figure 3: Example Time Series of Prices

(a) Xyzal, Tablets, 80 Count

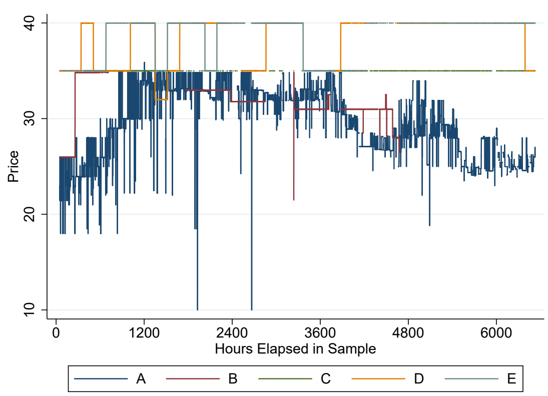

(b) Zyrtec, Liquid Gels, 40 Count

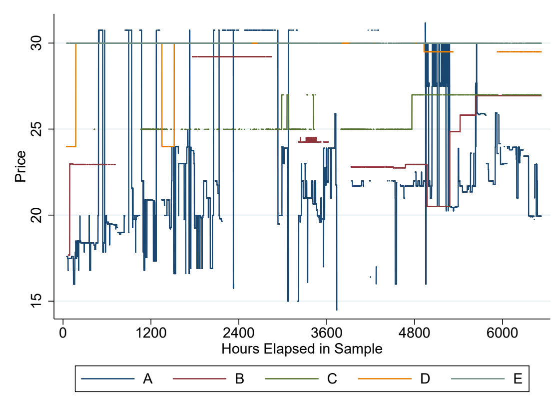

Table 2: Summary Statistics

	A	B	C	D	E	All
Daily Mean						
Count of Products	97.8	54.9	40.6	24.5	38.0	51.2
Daily Mean per Product						
Observations	21.1	19.9	16.2	21.2	19.7	19.6
Count of Reviews	83.5	165.6	257.1	230.7	183.7	167.8
Price Statistics						
Mean Price	25.793	33.739	18.250	21.600	20.842	24.094
Mean Abs. Price Change	1.368	3.243	0.997	3.248	3.052	2.199
Count of Price Changes	2.623	0.191	0.015	0.035	0.026	0.585
Any Price Change	0.495	0.085	0.015	0.031	0.026	0.132

Note: Statistics are calculated by website by day.

At the other extreme, retailer *C* only changes the price of 1.5 percent of its products each day, making a single change when it does so.

Figure 3 displays example time series for two products in our sample: Xyzal-Tablet-80 and Zyrtec-Liquid Gel-40. These two examples illustrate fundamentally different pricing patterns across the five retailers. Retailer *A* has frequent price changes of a large magnitude, but prices that are on average lower than its competitors. Retailer *B* has price movements that are closer to *A*, though less frequent, whereas *C*, *D*, and *E* tend to have more similar prices.

4 Pricing Technology and Prices

4.1 Heterogeneity in Pricing Technology

The previous section showed that there is variation across the five retailers in terms of how frequently they change prices. This fact alone is not sufficient evidence to demonstrate that the retailers possess different technologies, in terms of the capability to change prices quickly. Indeed, a price change is an outcome, and even retailer *A* has periods of stable prices in the data.

However, examining the data further reveals that the five retailers possess very different pricing technologies. First, Figure 4 presents the occurrence of price changes by day of the week. Though *A*, B, and C have roughly equal amounts of price changes throughout the week, retailers *D* and *E* realize nearly all of their price changes on Sunday.

Second, Figure 5 presents the distribution of price changes across hours of the day. Retailers

Figure 4: Daily Price Changes, by Retailer and Day of Week

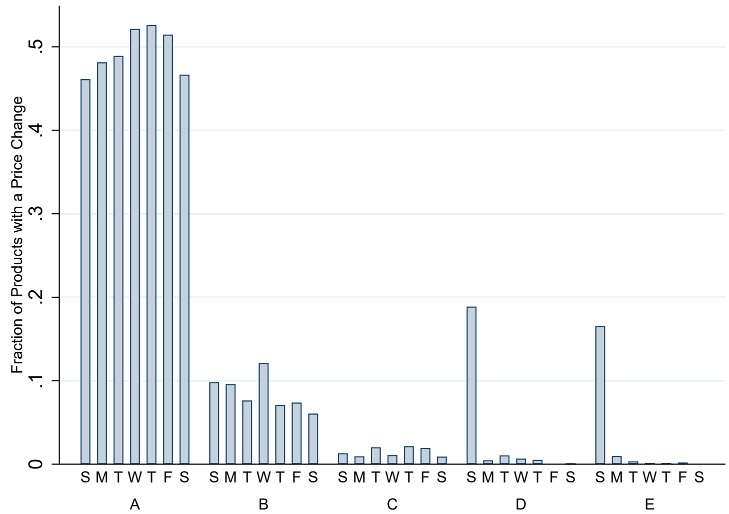

Table 3: Pricing Frequency by Online Retailers

Retailer	Frequency	Period
A	Hourly	Any time
B	Hourly	Any time
C	Daily	3:00 AM to 6:00 AM EDT
D	Weekly on Sunday	12:00 AM to 2:00 AM EDT
E	Weekly on Sunday	1:00 AM to 6:00 AM EDT

A and *B* have price changes well-dispersed across the 24 hours of the day. In contrast, *C*, *D*, and *E* have nearly all of the observed price changes occurring in a period of a few hours in the morning.[8] Firms *D* and *E* begin their price update script around midnight EDT. Thus, we observe that *A* and *B* have pricing technology that allows for updates at any hour of the day, *C* has technology that allows for a daily update each morning, and *D* and *E* have technology that allow them to update their prices once per week (on Sundays). Table 3 summarizes these findings.

Technology, in the sense of this paper, is the capability to change prices on the online store. We highlight the stark differences in the distribution of observed price changes as pointing to heterogeneity in these capabilities. Though firms do not use every opportunity to change prices (recall that firm *C* only changes the price of 1.5 percent of its products each day), we find the consistency in the times that they do change prices as compelling evidence of technological constraints.

Of course, our definition of technology is not merely the set of hardware and software that

[8]Several of the changes that occur away from these peaks are likely due to measurement error.

Figure 5: Hourly Price Changes, by Retailer

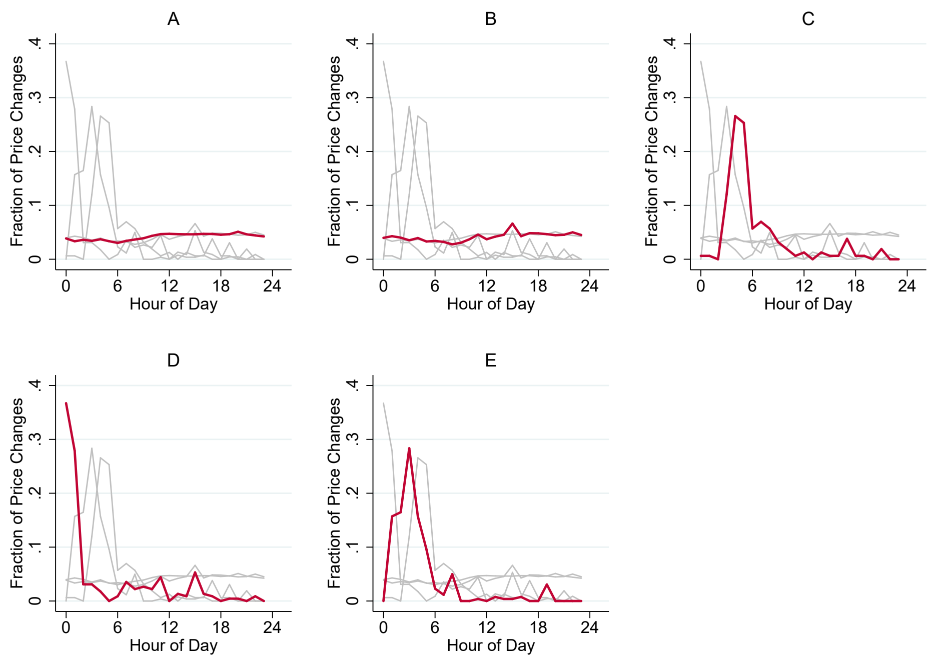

Notes: Hours are reported in Eastern Daylight Time, which is offset from Coordinated Universal Time by -4 hours. Times are not adjusted for daylight saving time.

functionally updates a price on website. Technology also includes managerial and operational constraints that restrict a firm from updating prices on a more frequent basis. Put different, even if firm C had access to the same hardware and software as A, it would take significant operational changes to enable the firm to update the prices as frequently.

4.2 Evidence of Competitive Effects

Having established that the five retailers in our data have different technologies affecting the frequency at which they update prices, we now examine the pricing patterns in more detail to determine whether the data are consistent with the model of Section 2. The theory generates a stark prediction: firms that have higher-frequency pricing technology will have lower prices. Again, the intuition is higher-frequency pricing allows a firm to commit to meet its rivals' best price; as a consequence, the rival prices less aggressively.

We examine this prediction. To compare prices, we use a regression in order account for differences in product assortment in the cross-section and over time. More specifically, we regress log prices on indicators for each retailer, while including product and period (hourly) fixed effects. The resulting coefficients reflect the average difference in (log) price for identical products (brand-drug-form-size) sold across different retailers.

Table 4: Log Price Differences Across Retailers, Relative to *A*

	(1)	(2)	(3)	(4)
B	0.054***	0.093***	0.149***	0.134***
	(0.001)	(0.000)	(0.001)	(0.001)
C	0.093***	0.079***	0.158***	0.120***
	(0.001)	(0.000)	(0.001)	(0.000)
D	0.293***	0.269***	0.391***	0.322***
	(0.001)	(0.000)	(0.001)	(0.001)
E	0.271***	0.253***	0.339***	0.298***
	(0.001)	(0.000)	(0.001)	(0.000)
Product FEs	X	X	X	X
Period FEs	X	X	X	X
Sold at All Retailers			X	X
On or After Oct 1 2018	X		X	
Observations	697,690	1,655,477	213,160	527,324

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4 presents the results. Retailer *A* serves as a baseline, so the coefficients reflect the average difference in log price relative to *A*. Relative to retailer *A*, products are typically sold at a 5.5 percent (0.054 log point) premium at *B* and a 9.7 percent (0.093 log point) premium at *C*. These same products are sold at a substantial premium at retailers *D* and *E*, who have average price differences of 34 percent and 31 percent, respectively! Our focal period is for observations on or after October 1, 2018, for which we have a more consistent panel of retailer *B*. We observe the same qualitative patterns if vary our estimation sample. Models (2) and (4) use observations from our entire sample, and models (3) and (4) includes only products sold by all five retailers. The results remain qualitatively similar, though the price differences between *A* and the rest increase when only using products sold by all five retailers.

The resulting price patterns are consistent with the model described in Section 2. Firm *A* has implemented a pricing technology that enables them to perform frequent updates, and *A* has the lowest prices. This is in line with the prediction that a faster pricing algorithm enables a firm to best respond to its competitors, resulting in a lower equilibrium price. The pattern also holds up if we look at firms with weekly pricing technology (*D* and *E*), who sell at a price substantially higher than (*B* and *C*), who have a more frequent pricing technology. Lastly, we note that *E* sells at a slightly lower price than *D*, and it updates its prices a few hours later.

In addition to describe these differential prices on average, we look at the time series of pricing responses to see if it is plausible that firms are indeed basing their prices in response to other firm's prices. The nature of an algorithm generates a second set of predictions: If

Figure 6: Price Response to a Price Change by Retailer *D*

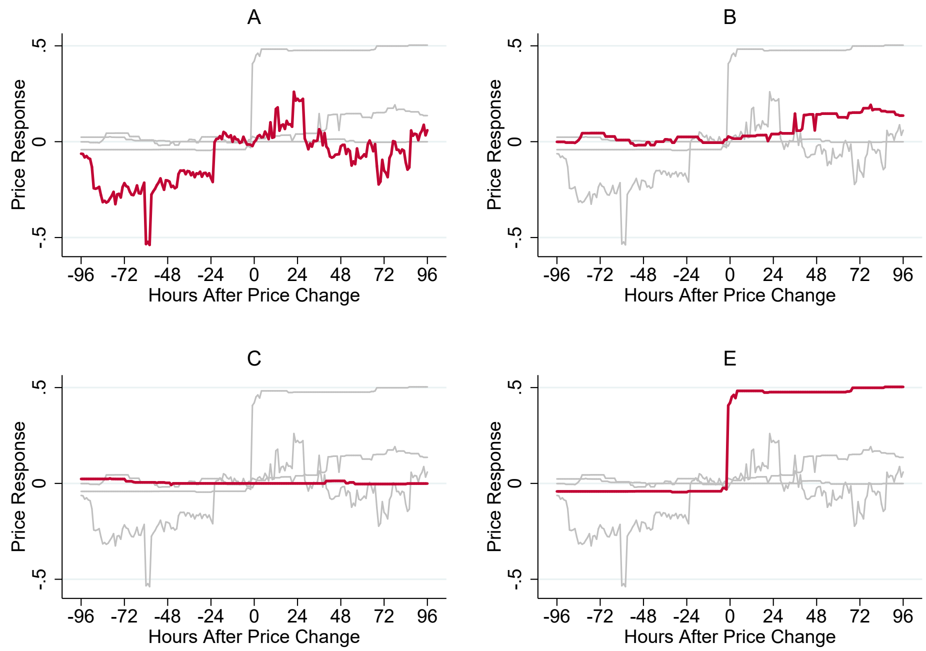

firms were responding to commonly-observed demand shocks, then we should not expect firms with higher-frequency technology to react after a price change by a lower-frequency rival.[9] We should expect, instead, that firms with higher-frequency pricing technology reflect these demand shocks earlier.

Figure 6 plots the average price response by retailers to price change by retailer *D*, where the price change is normalized to 1. We selected *D* as the baseline because it has the slowest frequency (weekly), and *E* updates its prices a few hours later. We plot the average prices in the 96 hours before and after a price change by *D*, and we normalize the price in the preceding 12 hours to 0.

The evidence is somewhat mixed, but it suggests that retailers may be responding to other firms' prices. In particular, *A* appears to raise its price within 24 hours of a price increase by *D*.[10] In the 24 hours before the price change, prices for *A* are stable, though they had increased over the previous 72 hours, indicating that both *A* and *D* could be responding to common shocks occurring earlier that week. Strikingly, *E* raises its price within a few hours of a price change by *D*, and it reflects half of *D*'s change. *B* also changes its price after a price change by *D*, but the response is more muted and occurs 24–28 hours later.

[9] The predictions get murkier if firms can anticipate demand shocks. We ignore this feature for now.

[10] Price changes are normalized to $|1|$, so these also capture price decreases.

5 Implications

In this paper, we have provided theoretical results for how changes to pricing technology—through asymmetric pricing frequency and commitment through algorithms—can lead to higher equilibrium prices relative to the Bertrand equilibrium. We then document how five large online retailers appear to have asymmetric pricing technology, and we show that the pricing patterns in the data are consistent with the predictions of our theoretical results. In particular, we find that firms with higher-frequency pricing technology have, on average, substantially lower prices.

There are, of course, several possible reasons for the correlation between pricing frequency and price levels we observe in our data. By no means are the firms in our data symmetric in terms of other factors. However, firms are increasingly adopting more sophisticated pricing technologies. If policymakers are concerned that algorithms will raise prices, then the concern is much more broad than that of collusion. Indeed, higher prices from algorithms can arise as competitive outcomes. Moreover, our negative result in Proposition 2—that the Bertrand best-response functions are *not* equilibrium strategies when firms compete in algorithms—should raise concerns that the Bertrand equilibrium will be the exception, rather than the rule.

There exists a simple policy prescription if policymakers are interested in keeping prices closer to the Bertrand equilibria: policymakers could insist that firms cannot make their pricing algorithms functions of rivals' prices. Firms would remain free to have frequent price updates as a function of other factors, such as demand shocks, but explicitly accounting for rivals prices would be forbidden. As we have shown, algorithms that incorporate rivals' prices can support the collusive outcome in equilibrium.

References

Benoit, J.-P. and V. Krishna (1985). Finitely Repeated Games. *Econometrica*.

Bresnahan, T. F. (1981). Duopoly Models with Consistent Conjectures. *The American Economic Review*.

Cavallo, A. (2018). More Amazon Effects: Online Competition and Pricing Behaviors. Working paper.

Fudenberg, D. and E. Maskin (1986). The Folk Theorem in Repeated Games with Discounting or with Incomplete Information. *Econometrica*.

Gal-Or, E. (1985, October). First Mover and Second Mover Advantages. *International Economic Review 26*(3), 649.

Holt, C. A. (1985). An experimental test of the consistent-conjectures hypothesis. *The American Economic Review*, 179–198.

Kamien, M. and N. L. Schwartz (1983). Conjectural Variations. *The Canadian Journal of Economics*.

Miklos-Thal, J. and C. Tucker (2018). Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers? Working paper.

Salcedo, B. (2015). Pricing Algorithms and Tacit Collusion. Working paper.

Schwalbe, U. (2018). Algorithms, Machine Learning, and Collusion. Working paper.

Tennenholtz, M. (2004, November). Program equilibrium. *Games and Economic Behavior 49*(2), 363–373.