Competition in Pricing Algorithms*

Zach Y. Brown
University of Michigan†

Alexander MacKay
Harvard University‡

November 1, 2019

Abstract

Increasingly, retailers have access to better pricing technology, especially in online markets. Through pricing algorithms, firms can automate their response to rivals' prices. What are the implications for price competition? We develop a model in which firms choose algorithms, rather than prices. Even with simple (i.e., linear) algorithms, competitive equilibria can have higher prices than in the standard simultaneous Bertrand pricing game. Using hourly prices of over-the-counter drugs from five major online retailers, we document evidence that these retailers possess different pricing technologies. In addition, we find pricing patterns consistent with competition in pricing algorithms. A simple calibration of the model suggests that pricing algorithms lead to meaningful increases in markups, especially for firms with superior pricing technology.

*We thank Matt Grennan and Scott Kominers for helpful comments. We also thank seminar and conference participants at Harvard Business School and the IIOC. We are grateful for the research assistance of Alex Wu and Pratyush Tiwari.

†University of Michigan, Department of Economics. Email: zachb@umich.edu.

‡Harvard University, Harvard Business School. Email: amackay@hbs.edu.

1 Introduction

Increasingly, retailers have access to better pricing technology, especially in online markets. In particular, pricing algorithms allow retailers to commit to pricing strategies that depend on the prices of competitors. In this paper, we investigate how the presence of pricing algorithms affects equilibrium outcomes. We introduce a game in which firms choose algorithms, rather than prices, and we show that competitive equilibria of this game tend to have higher prices than the competitive price-setting equilibrium. Further, the collusive (profit-maximizing) prices can be achieved with competitive algorithms that ensure no profitable unilateral deviation. This equilibrium can be maintained with simple algorithms that that are linear in rivals' prices.

There is a growing concern that pricing algorithms, which are becoming more prevalent, will lead to higher prices for consumers. In contrast to our analysis, the previous literature has focused on whether or not pricing algorithms can better enable collusion,[1] as has the popular press.[2] The idea that *competition* in pricing algorithms can lead to higher prices may be somewhat surprising. Simple pricing algorithms, such as those we examine in this paper, do not "look collusive." Moreover, in practice, many algorithms have a simple linear form.[3]

We also address the question of whether algorithms may lead to competitive prices. We show that it is not an equilibrium for all firms to choose algorithms that equal their price-setting best-response (Bertrand reaction) functions. The Bertrand equilibrium is achieved only if algorithms do not depend on rivals' prices. Therefore, if we observe firms using algorithms that include rivals' prices, we should not expect Bertrand prices in equilibrium. Our results indicate that the potential of algorithms to raise prices goes beyond the possibility of facilitating collusion.

To support our theoretical analysis, we provide an empirical analysis of prices for five large online retailers. Our novel dataset—the first to capture high-frequency prices for a product category across multiple websites—allows us to document new facts about heterogeneity in pricing technology and price dispersion. First, we show that firms vary in the frequency in which they can update prices. Second, we find that firms that have higher-frequency pricing have lower prices than their competitors. The second result is consistent with our theoretical model. Thus, while the previous literature has focused on the role of search frictions as a explanation for price dispersion, our model provides an alternative, complementary explanation: differences in

[1] Recent papers on the subject include Miklós-Thal and Tucker (2019), Schwalbe (2018), and Salcedo (2015).

[2] 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.

[3] It is typical for algorithms have a linear adjustment based on the average price of a set of competitors. In one interesting example, a retailer on Amazon.com set its price for a book to be 0.9983 times its rival's price, and the rival set its price to be 1.270589 times the retailers' price. The price of the book rose to nearly $24 million. This, we note, was not an equilibrium. See "How A Book About Flies Came To Be Priced $24 Million On Amazon," Wired, April 27, 2011. https://www.wired.com/2011/04/amazon-flies-24-million/

prices for the same product across websites can be driven by differences in pricing technology alone. Using a simple calibration exercise, we find that asymmetric pricing technology leads to higher prices for all retailers and exacerbates price differences among similar retailers.

To explain these results, it is helpful to first understand the different aspects of pricing algorithms. 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 a 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-run) commitment device.*
   When an algorithm calculates price based on the prices of other firms, it can react quickly to price changes in the market. This becomes a short-run 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 the adoption of pricing algorithms, provides a firm with an advantage relative to its peers, as well as the implications for competitive outcomes. Though economists have, thus far, focused on whether increased sophistication can lead to higher prices, we focus instead on the features of frequency and short-run commitment. We believe these latter two features are most affected by the use of algorithms and have the greatest potential to impact competition.

We find that algorithms' impact on frequency and short-run commitment 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 pricing frequency, e.g., changing prices once each week versus once each day. When firms have asymmetries in pricing frequency, the game closely resembles a Stackelberg leader-follower pricing game, leading to higher prices than the Bertrand equilibrium with differentiated products. Moreover, when firms can choose their pricing frequency, each firm has a unilateral profit incentive to choose either more frequent or less frequent pricing than their rivals. Therefore, the simultaneous Bertrand model is not an equilibrium outcome when firms can choose pricing frequency.

Second, we show that symmetric short-run commitments can also generate higher prices. To do so, we introduce a one-shot competitive game in which firms submit pricing algorithms, rather than prices (Section 3). To provide a conservative analysis on the impact of algorithms, we tie our hands. We provide two restrictions on the set of equilibrium strategies that rule out cooperate-or-punish equilibria and ensure that the algorithms, in some sense, look "competitive." 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 short-run commitment devices that enable firms to internalize the effect of a change in its strategy on the behavior of its rivals.[4] Pricing algorithms provide an economic mechanism to assure the consistency between beliefs and behavior for a multitude of equilibria.

Indeed, one may interpret algorithms as providing 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 arguments that are used to support various equilibria in repeated games. Our analysis provides subgame perfect results for the one-shot game. Our findings complement the work of Maskin and Tirole (1988), who show that when short-run commitments allow for alternating price-setting behavior by rival firms, prices will be higher in Markov perfect equilibria than in the Bertrand game. For a discussion of similar outcomes that arise in other game theory models and additional related literature, see the related literature below.

In our empirical analysis, we study prices for over-the-counter allergy medications for the five largest online retailers for the category.[5] Our novel dataset is described in Section 4. By studying prices at the hourly level, we are able to document heterogeneity in pricing technology. We find that two firms have within-the-hour ("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 associated 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 approximately 30 percent higher. We document these results, as well as evidence that price changes at high-frequency retailers are

[4]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 for competition in conjectural variations.

[5]Based on Google search data. All five retailers have similar search shares for "allergy" + retailer name.

more likely after a price change by a low-frequency retailer, in Section 5.

In Section 6, we perform a simple calibration exercise. We introduce an empirical model of spatial differentiation that is a generalization of the models of Hotelling (1929) and Salop (1979). We fit the model to average prices and market shares in our data, and we simulate the counterfactual equilibrium for simultaneous Bertrand price competition. Relative to the Bertrand equilibrium, the calibrated model predicts that algorithmic competition increases average prices by 5.2 percent across the five firms.

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 discount 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 these algorithms often include rivals' prices.[6] Further, our empirical analysis shows that the price patterns observed in the data are at least consistent with the model we analyze. Thus, if policymakers are concerned that algorithms will raise prices, then the concern is much more broad than that of collusion.

Of course, algorithms may have several benefits, such as the ability to more efficiently respond to time-varying demand. 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. We briefly discuss other policies in Section 7.

Related Literature

Our analysis complements the economics literature that has focused on the impact of algorithms on prices through 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). Calvano et al. (2019) show that this concern applies to standard reinforcement-learning algorithms by demonstrating that they converge to supracompetive prices. Two recent papers show that better demand forecasting can have ambiguous effects on consumer welfare. Miklós-Thal and Tucker (2019) and O'Connor and Wilson (2019) show that better demand prediction can increase the incentive for

[6] 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."

firms to deviate from the collusive price, resulting in lower prices to consumers.[7]

Another strand of literature deals with one-shot games where players choose contracts (or commitment devices) that condition their actions on the strategies of the other players (Tennenholtz, 2004; Kalai et al., 2010; Peters and Szentes, 2012). In this literature, (equilibrium) contracts are functions of the other players' contracts. Tennenholtz (2004) gives the example of submitting a computer program that reads the rivals' program and chooses an action accordingly. This approach can support any individually rational payoffs, as in repeated games.

By contrast, we analyze a game where firms' strategies (algorithms) may only be functions of rivals' actions (prices). This modeling choice reflects pricing algorithms actually used by firms, and it requires additional analysis, as the equilibrium strategies of the above games, which (in general) condition on strategies, are no longer permissible. Further, we concern ourselves only with strategies that look "competitive," in that the algorithms must be continuous. Thus, we restrict attention to strategies that might not raise antitrust concerns about collusion *prima facie*. Finally, we highlight that there is an economic mechanism that supports the nature of commitment found in our model, whereas the credibility of commitment devices more generally may be questioned.

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 generate a unique equilibrium (Bresnahan, 1981). 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 is an equilibrium in consistent conjectures with the price-setting game.

There are several strands of theoretical literature that describe games where the equilibria result in prices higher than in the (static) Bertrand equilibrium. For general repeated games, variations of the folk theorem have shown that any individually rational set of payoffs can be obtained in equilibrium (e.g., Fudenberg and Maskin, 1986; Benoit and Krishna, 1985). This could, for example, include the collusive outcome in model of oligopoly pricing. A critique of using this general argument to understand oligopoly behavior has been that many of the obvious strategies to enforce higher prices are not robust. In particular, threatening to drastically lower price in response to a rivals' deviation from the collusive equilibrium tends to not be credible. In response to this, Maskin and Tirole (1988) analyze a dynamic oligopoly pricing game and restrict attention to Markov perfect equilibria that are "renegotiation proof" for

[7]These two papers arrive at similar conclusions through different mechanisms. Miklós-Thal and Tucker (2019) is based on a Rotemberg and Saloner (1986) model and O'Connor and Wilson (2019) is based on an extension of Green and Porter (1984).

any possible actions by the competing firms. With these restrictions, Maskin and Tirole (1988) are able to provide conditions under which the monopoly price—the collusive outcome—is the unique equilibrium.

In Maskin and Tirole (1988), firms take turn setting prices, relying on the ability of firms to make short-run commitments. Likewise, a Stackelberg leader-follower model requires the same short-run commitment to generate higher prices. We argue that such commitments are credible, made possible by investments in differential pricing technology.[8] We find that, empirically, the most important variation is in the frequency of pricing technology. Moreover, the two firms with the same pricing frequency tend to make price changes at roughly the same time, rather than in alternating fashion as in Maskin and Tirole (1988). In equilibrium, we find that firms prefer an outcome with asymmetric pricing frequencies, where one firm can update its prices more often than its rival.

The theoretical analysis of oligopolistic pricing behavior has been complemented with empirical studies that document the coordination of firms on higher prices (e.g., Miller and Weinberg, 2017; Byrne and de Roos, 2019). A recent paper by Igami and Sugaya (2018) explicitly estimates the dynamic incentives to collude in a vitamin cartel. We are not aware of any empirical studies that attempt to address the critiques that are tackled by Maskin and Tirole (1988).

2 Demand and Pricing Frequency

2.1 Illustrative Model: Price Competition with Differentiated Products

To motivate our analysis, consider a simple spatial differentiation model.[9] Firm 1 and firm 2 are located 1 unit (e.g., one mile) 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. This is a variant of the Hotelling (1929) model, with fixed locations and an outside option.

Each consumer i receives utility α from the good and has disutility of τd_{ij} for the distance d_{ij} they travel to purchase from firm j. 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_{ij} for consumer i of choosing $j \in \{1, 2, 0\}$—

[8] Another related concept is the cartel punishment device of Osborne (1976). Here again we consider algorithms as an economic mechanism to make such commitments credible.

[9] We provide the generalization of this model to one more suitable for empirical applications in Section 6.

purchasing from firm 1, firm 2, or not purchasing (0)—is given by:

$$u_{i1} = \alpha - \tau\theta_i - p_1 \tag{1}$$
$$u_{i2} = \alpha - \tau(1-\theta_i) - p_2 \tag{2}$$
$$u_{i0} = 0 \tag{3}$$

Consumer i will prefer firm 1 over firm 2 if and only if:

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

We consider the following utility parameters: $\alpha = 1$ and $\tau = 1$. 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 a (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*. Thus, firm 1 can obtain 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.

This result has an intuitive logic: Firm 2 commits to "undercut" the price of firm 1, maximizing its own profits conditional on its rival's price. This softens firm 1's incentive to compete on price. The Bertrand-Nash logic uses a dynamic metaphor to rule out the above outcome: if firm 2's price is fixed at $\frac{5}{4}$, firm 1 has a unilateral incentive to reduce 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 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.[10] When firm 2 alone has higher-frequency pricing technology, then it has the capability to update prices without a response by firm 1. As we show in the following section, when firms can choose their pricing frequency, asymmetric

[10]This is true when firms set prices simultaneously. Maskin and Tirole (1988) provide conditions under which firms would choose to alternate in equilibrium and prices would be higher.

frequencies are the equilibrium outcome.

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 ($\frac{\partial p_2}{\partial p_1} = 0$), 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 Adoption and the Choice of Pricing Frequency

In this section, we consider the choice of firms of whether or not to adopt superior (or inferior) pricing technology, in terms of frequency. Both firms are better off if only one firm adopts the technology, though they would prefer themselves to be the one to adopt.[11] The resulting technology adoption problem is a classic coordination game. Therefore, it is not surprising that

[11] 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.

Figure 1: Adoption Game

		Firm 2: Low	Firm 2: Moderate	Firm 2: High
Firm 1	Low	$(1,1)$	$\left(\frac{9}{8}, \frac{25}{16}\right)$	$\left(\frac{9}{8}, \frac{25}{16}\right)$
	Moderate	$\left(\frac{25}{16}, \frac{9}{8}\right)$	$(1,1)$	$\left(\frac{9}{8}, \frac{25}{16}\right)$
	High	$\left(\frac{25}{16}, \frac{9}{8}\right)$	$\left(\frac{25}{16}, \frac{9}{8}\right)$	$(1,1)$

asymmetric pricing technologies might arise endogenously in equilibrium, as we find in our data (Section 5).

To wit: Consider a world where demand is static and suppose that firms can costlessly change their pricing frequency. Each firm has a unilateral incentive to choose either a higher frequency or a lower frequency than their rivals, as profits increase for both the faster firm and the slower firm compared to the symmetric frequency case. Therefore, symmetric pricing frequency—the Bertrand model of simultaneous price-setting—is not a (subgame perfect) equilibrium outcome. The result does not depend how frequently prices are set on average; in settings with stable demand, the model is identical whether rivals set prices each month or each second.

To illustrate this point, consider the three-by-three first-stage game where firms can choose pricing frequency (Figure 1). Firms know the profits for each subgame when they choose a high frequency, a moderate frequency, or a low frequency. The scenario where both firms choose a moderate frequency is not an equilibrium, because each firm has an incentive to deviate by choosing either a faster or a slower pricing frequency. The only equilibria of the game are asymmetric choices in which only one player chooses (High).[12]

While this model provides clear *unilateral* incentives to adopt asymmetric technologies; symmetric pricing frequencies do arise in the real world. We believe there are other factors that help to maintain symmetric pricing frequency in equilibrium. First, a potential benefit of frequent price changes is the ability to adapt to time-varying demand conditions (so-called "dynamic pricing"). Second, changing one's pricing frequency is not costless; technological or operational costs may maintain symmetric frequencies in equilibrium. Regardless, these features do not eliminate the profit incentive we identify here.

[12]This simple game also illustrates an "arms race" effect: (Moderate, Low) is not an equilibrium because the second player would like to leapfrog the first by choosing High instead of Low. Thus, one player always chooses High in equilibrium.

2.4 The Three-Firm Case

To extend the intuition of asymmetry in pricing technology beyond duopolistic competition, we consider the case of three firms. We simulate equilibrium prices in the model with the aim of comparing model predictions to our empirical results in Section 5.

The setup of the model remains similar to that of the model in section 2.1, but the three firms are now located at equidistant 1-unit intervals along a circle with circumference of 3. Thus, we use the Salop (1979) model to characterize demand. Each unit of the circle's circumference contains a mass of 1 consumers. Consumers maintain travel costs as before.

Where the utility from both goods is positive, the (local) demand for each good is:

$$q_1 = 1 - p_1 + \frac{1}{2}p_2 + \frac{1}{2}p_3 \tag{10}$$

$$q_2 = 1 - p_2 + \frac{1}{2}p_1 + \frac{1}{2}p_3 \tag{11}$$

$$q_3 = 1 - p_3 + \frac{1}{2}p_1 + \frac{1}{2}p_2. \tag{12}$$

As before, the Bertrand-Nash equilibrium is $p_1 = p_2 = p_3 = 1$, and the collusive price is $p_1 = p_2 = p_3 = \frac{3}{2}$.

Now assume that there are three levels of pricing technology. Firm 1 has inferior pricing technology and can set prices only at the beginning of the period. Firm 2 has more frequent pricing, allowing it to update twice: at the beginning of the period and halfway through the period. Firm 3 has superior technology and can update three times during the period, with the last opportunity occurring after firm 2's second update. As demand is realized only at the end of the period, only the final price-setting order matters, rather than the exact frequency. As in Section 2.2, this situation mirrors multi-step Stackelberg competition in which firm 1 is the leader.[13]

Figure 2 demonstrates the conclusions of the model. Firm 1, which has the slowest pricing technology, has the highest price. Firm 3, which has the fastest pricing technology, has the lowest price. The model implies that prices are monotonically decreasing in pricing algorithm frequency. Furthermore, all prices in the pricing algorithm equilibrium are higher than those from the Bertrand-Nash equilibrium. Firms with inferior technology choose to compete less aggressively, as firms with superior technology can credibly commit to offering lower prices. Within a pricing algorithm equilibrium, more frequent pricing is correlated with lower prices, but all prices are elevated relative to the case where all firms have the slowest technology.

[13]Alternatively, the model can be recast by assuming that in each period, firms are only able to change price a limited number of times. In particular, assume firm 1 can set only one price in each period, firm 2 can set a price and change it at most once, and firm 3 can set a price and change it at most twice. Furthermore, demand is realized throughout the period but firms receive zero profit until they post a price. In equilibrium, firm 1 will set a price at the very beginning of the period. Firm 2 and then firm 3 will follow instantaneously.

Figure 2: Simulated Pricing Algorithm Equilibrium

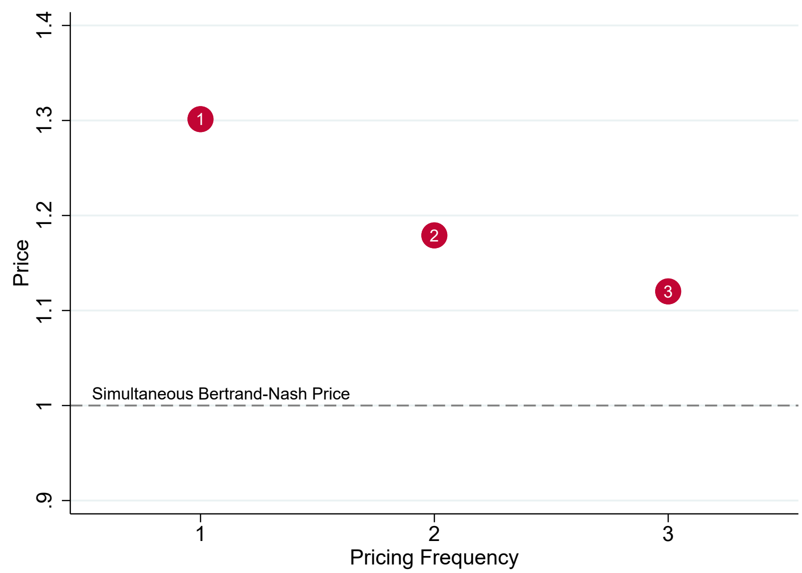

Notes: Figure displays the prices for three firms from a simulation of competition in pricing algorithms with heterogenous pricing technology. Pricing frequency of 3 is the superior (i.e. fastest) pricing technology. Marker labels indicate the firm.

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. This game maps more closely to the pricing technology used by firms in practice. The asymmetric pricing algorithm 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 ($\varepsilon \to 0$).

The implications of pricing algorithm model are relevant when prices update at a greater frequency than the algorithms. The pricing algorithm game allows us to represent firm behavior when two firms converge to high-frequency algorithms, yet simultaneous price-setting does not accurately capture behavior. If all firms could adjust their algorithms at the same frequency that their prices update, then the pricing frequency logic would prevail, and symmetric frequencies would result in the Bertrand equilibrium. Our discussion is in terms of a two-player game, but the concepts readily generalize. Throughout, we focus our attention on pure strategies.[14]

3.1 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 (Miklós-Thal and Tucker, 2019) 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 3.4.

After receiving the pricing algorithms, the central coordinator solves the system of equations set by the algorithms to determine prices. 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 joint 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.[15] In the real world, this selection procedure reflects pro-consumer market mechanisms to discipline firms.

Of course, even subject to the profit-minimizing coordinator, many 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).*

[14]We have yet to encounter mixed-strategy algorithms that are employed in practice.

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

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 collusion 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: in 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 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 in the frequency game. Neither firm can do better with a unilateral deviation. Thus, this asymmetric 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. This example reflects a general (negative) result of the model:

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:

$$\sigma_1(p_2) = 1 + \frac{1}{3}p_2 \tag{13}$$

$$\sigma_2(p_1) = 1 + \frac{1}{3}p_1. \tag{14}$$

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

Thus, reasonable-looking algorithms are capable of supporting many different price outcomes as equilibria, including the collusive outcome. These strategies, which are linear functions of rivals' prices, may not raise competitive concerns *prima facie*. 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 3.4.

3.2 A Multitude of Equilibria

It is possible to show that a multitude of equilibria can exists when firms compete in algorithms. To demonstrate this, we further restrict the class of algorithms to a special case: algorithms that are linear in other firms' prices. Even with these straightforward algorithms, we can show that many equilibria exist:

Proposition 3 *When firms compete in a one-shot game by submitting pricing algorithms, any price vector can be supported by algorithms that are linear functions of rivals' prices, provided the derivatives of profits with respect to prices exist at those prices.*

> **Proof:** For the two-firm case, consider the price vector $\hat{p} = (\hat{p}_1, \hat{p}_2)$. Recall that, in equilibrium, it must be the case that a firm cannot do better by reverting to price-setting behavior. Firm 1's equilibrium price-setting first-order condition can

be rewritten as:

$$\left.\frac{d\pi_1}{dp_1}\right|_{\hat{p}} = \left.\frac{\partial\pi_1}{\partial p_1} + \frac{\partial\pi_1}{\partial p_2}\frac{\partial\sigma_2}{\partial p_1}\right|_{\hat{p}} = 0 \tag{15}$$

$$\implies \left.\frac{\partial\sigma_2}{\partial p_1}\right|_{\hat{p}} = -\left.\frac{\partial\pi_1/\partial p_1}{\partial\pi_1/\partial p_2}\right|_{\hat{p}} \tag{16}$$

Likewise, $\frac{\partial\sigma_1}{\partial p_2} = -\frac{\partial\pi_2/\partial p_2}{\partial\pi_2/\partial p_1}$ when evaluated at $\hat{p}$. To support the prices $(\hat{p}_1, \hat{p}_2)$ with algorithms that are linear in rivals' prices, one can solve the system of equations so that beliefs and strategies are consistent:

$$\hat{p}_1 = \sigma_1(\hat{p}_2) = a_1 + b_{12}\hat{p}_2 \tag{17}$$

$$\hat{p}_2 = \sigma_2(\hat{p}_2) = a_2 + b_{21}\hat{p}_1 \tag{18}$$

It is apparent that the solution has $b_{12} = -\left.\frac{\partial\pi_2/\partial p_2}{\partial\pi_2/\partial p_1}\right|_{\hat{p}}$ and $b_{21} = -\left.\frac{\partial\pi_1/\partial p_1}{\partial\pi_1/\partial p_2}\right|_{\hat{p}}$. Thus, each equation has one unknown, and the system has a unique solution for the parameters a_1 and a_2. It is straightforward to extend the argument to many firms.[16]

Despite this multiplicity result, we expect algorithms to result in higher prices than the Bertrand-Nash equilibrium for two reasons. First, when algorithms have positive slope coefficients on rivals' prices, higher prices result. Imposing this restriction on firms' choices seems reasonable a priori because prices are strategic complements. In other words, prices that are lower than Bertrand-Nash are supported only when an algorithm treats the rival prices a strategic substitutes, despite the complementarity.

Second, many of these equilibria are "knife-edge" cases. To examine which equilibria are, in some sense, more robust, we simulate a simple learning process in Section A of the Appendix. Firms experiment with linear algorithms, updating the parameters if profits increase. From a starting point of randomly-chosen algorithms, firms disproportionately arrive at equilibria that are bounded from below by their best-response functions and bounded from above by the profit Pareto frontier. Our simulation shows that higher prices result.

Thus, we have demonstrated two properties of an equilibrium in pricing algorithms. First, it is not (in general) an equilibrium for firms to submit their price-setting best-response functions. Second, any price vector can be supported by strategies that are linear functions of rivals' prices. We now turn our attention to the asymmetric case.

[16] For example, one solution to the J-firm problem would be to allow each firm's algorithm to depend only on one other firm's price: $R_j(p) = a_j + b_{jk}p_k$, where $k = j+1 \forall j < J$ and $k = 1$ if $j = J$. The solution is $b_{jk} = -\left.\frac{\partial\pi_j/\partial p_j}{\partial\pi_j/\partial p_k}\right|_{\hat{p}}$ and $a_j = \hat{p}_j - b_{jk}p_k$.

3.3 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 that responds to demand shocks and cost shocks, or other observables. In the absence of such features, i.e., when demand is stable, its algorithm reduces to standard price-setting behavior.

The asymmetric game is of particular interest because the real world features asymmetry in the ability of firms to adjust prices. Thus, characterizing the equilibrium may help us understand real-world phenomena. The asymmetric game extends the pricing frequency game in Section 2.2, allowing the firm with superior technology to commit to a pricing function.

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 4 *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, an 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 3.

Proposition 5 *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:

Corollary 6 *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.*

3.4 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

Figure 3: Equilibrium in the Bertrand Game and
the Asymmetric Pricing Algorithm Game

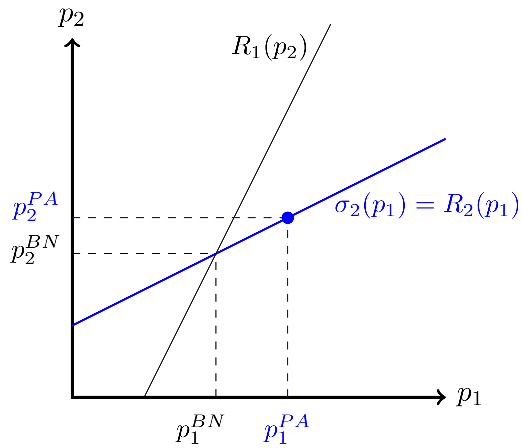

Notes: Figure plots the best-response functions $R_1(\cdot)$ and $R_2(\cdot)$ for price competition with differentiated products. The intersection of these functions produces the Bertrand-Nash (BN) equilibrium. PA indicates an equilibrium of the asymmetric pricing algorithm game where firm 2 submits $\sigma_2(\cdot) = R_2(\cdot)$ as its algorithm.

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.

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). 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

Table 1: Price Observations by Website and Brand

Retailer	Allegra	Benadryl	Claritin	Flonase	Nasacort	Xyzal	Zyrtec	Total
A	309,554	208,422	509,404	104,634	68,858	108,854	234,903	1,544,629
B	125,095	58,270	144,098	46,584	12,517	34,177	75,096	495,837
C	89,477	99,608	171,782	80,772	34,633	32,508	90,858	599,638
D	112,281	68,459	128,394	50,130	2,411	47,321	128,123	537,119
E	71,061	47,799	125,171	51,732	38,051	23,185	62,600	419,599
Total	707,468	482,558	1,078,849	333,852	156,470	246,045	591,580	3,596,822

needed. The supported equilibria are subgame perfect and outcomes of a one-shot game.

4 Data

For our novel dataset, we collected hourly prices of over-the-counter allergy drugs from five online retailers. These retailers are the five largest in the allergy category based on Google search data and are among the largest retailers overall by e-commerce revenues.[17] 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 seven brands of allergy drugs that are sold by all five retailers: Allegra, Benadryl, Claritin, Flonase, Nasacort, Xyzal, and Zyrtec. We define a product to be a drug-brand-form-(variant-)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. In addition, most brands offer variants with different amounts of the active drug, targeted for children, 12-hour or 24-hour use. There are also versions of the drug that are combined a decongestant. These varieties are captured by the *variant* variable. Finally, we distinguish products that are sold in a twinpack, so that twinpack of 12 tablets is a different product than a single pack of 24 tablets.[18] 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 10, 2018 through October 1, 2019, resulting in 3,596,822 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

[17]According to ecommerceDB (https://ecommercedb.com/), these five retailers combined for $6 billion in e-commerce revenues for personal care, which includes medicine, cosmetics, and personal care products.

[18]We drop multipacks that are of greater size than a twinpack, as they are not common across retailers.

Figure 4: Observed Products Over Time

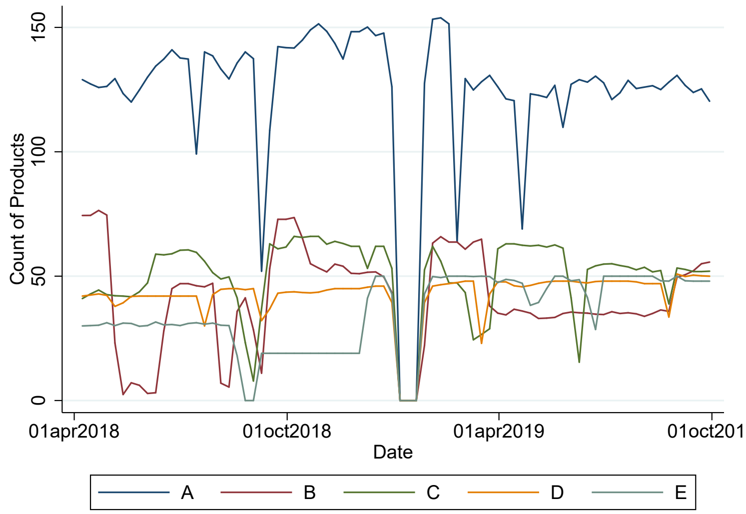

Notes: Figure displays the average daily count of observed products in our sample by week and by retailer.

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 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 if the gap of missing prices is fewer than six hours. 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.[19]

Figure 4 displays the count of products in our sample over time. The figure illustrates the challenge of capturing high-frequency price data over an extended period. Dips in the data correspond to changes to the retailer website and issues with the researchers' servers. We note that we have several periods of many thousands of observations for which we have a consistent sample, and the periods of missing data do not meaningfully affect our results once we account for period fixed effects. We also include specifications using only data from July 1, 2019 through October 1, 2019, which are the most recent three months and for which we have a fairly consistent panel. Retailers *A* and *B* offer significantly more product varieties than the other retailers. 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 124 products each day for retailer *A*, compared to 41 products for retailer *B*. Across all retailers,

[19] Overall, 7.8 percent of the prices are imputed in our analysis sample.

Table 2: Summary Statistics

	A	B	C	D	E	All
Daily Mean						
Count of Products	124.2	41.1	49.9	42.5	35.1	58.6
Daily Mean per Product						
Observations	20.9	20.4	19.0	21.1	19.1	20.1
Count of Reviews	101.1	231.9	258.5	241.3	302.1	219.4
Price Statistics						
Mean Price	27.354	16.185	17.628	20.925	21.742	20.750
Mean Abs. Price Change	1.358	2.236	1.124	3.281	3.063	1.891
Count of Price Changes	1.858	0.285	0.008	0.021	0.025	0.441
Any Price Change	0.372	0.088	0.008	0.020	0.024	0.103

Notes: Statistics are calculated by website by day.

we observe the price for each product in 20 out of 24 hours on average. The mean price for these products is $20.75, with a mean (absolute) price change of $1.89. The table indicates stark differences in the frequency of price changes. Retailer *A* changes the prices of 37 percent of its products in a given day, with an average count of 1.9 price changes per product. At the other extreme, retailer *C* only changes the price of 0.8 percent of its products each day, making a single change when it does so.

Figure 5 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.

5 Pricing Technology and Prices

5.1 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. Whether or not a price changes may not reflect the underlying capability to change price; even retailer *A* has periods of stable prices in the data.

However, examining the data further reveals that the five retailers possess very different

Figure 5: Example Time Series of Prices

(a) Xyzal, Tablets, 80 Count (b) Zyrtec, Liquid Gels, 40 Count

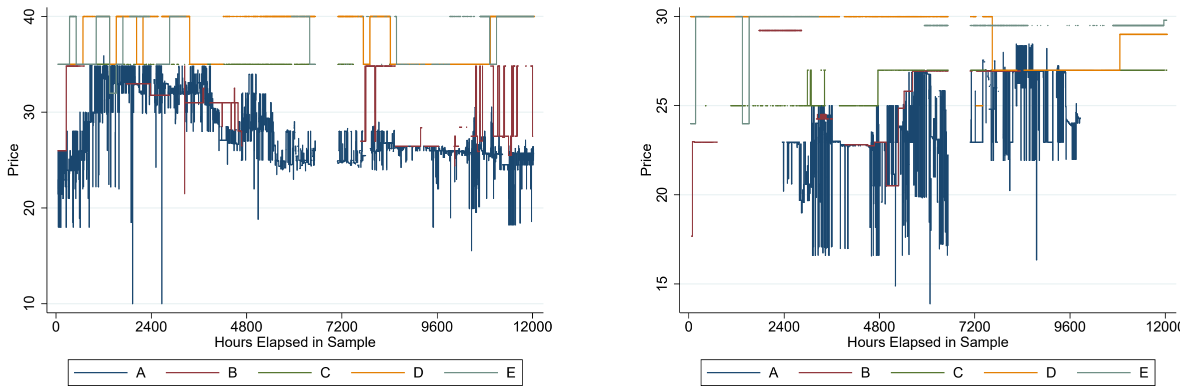

Notes: Figure displays the time series of hourly prices in our dataset for two example products across five retailers. Panel (a) displays the prices for an 80-count package of Xyzal tablets. Panel (b) displays the prices for a 40-count package of Zyrtec liquid gels.

pricing technologies. Figure 6 displays the heterogeneity in price changes by day of the week and hour of the day. First, panel (a) of Figure 6 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, panel (b) presents the distribution of price changes across hours of the day. Retailers *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.[20] 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* changes the prices of less than one 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 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 differently,

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

Figure 6: Heterogeneity in Pricing Technology

(a) Daily Price Changes, by Retailer and Day of Week

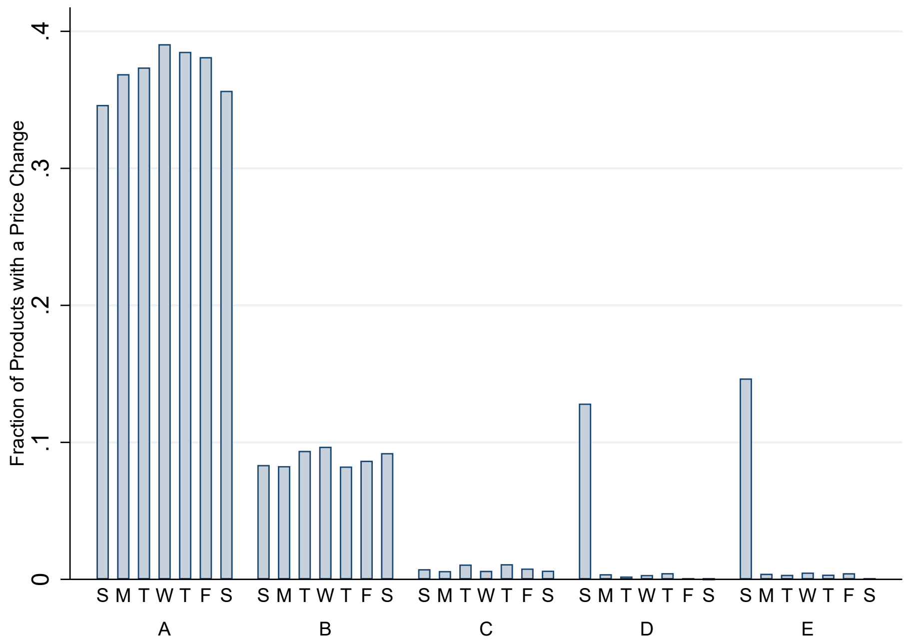

(b) Hourly Price Changes, by Retailer

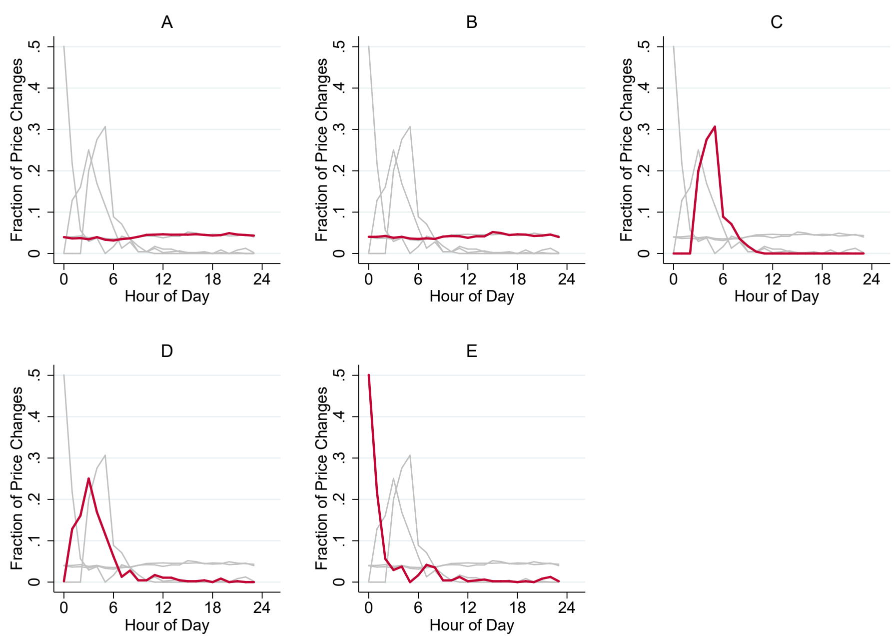

Notes: Panel (a) displays the fraction of products with a price change in each day of the week, by day of week and retailer. Panel (b) displays the fraction of all price changes that occur at a given hour of the day, by retailer. Hours are reported in Eastern Time.

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	1:00 AM to 6:00 AM EDT
E	Weekly on Sunday	12:00 AM to 2:00 AM EDT

Notes: Table summarizes the pricing technology of the five retailers in our data.

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.

5.2 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 to 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-variant-size) sold across different retailers.

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 6.8 percent (0.066 log point) premium at *B* and a 9.5 percent (0.091 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 28 percent and 33 percent, respectively. We observe the same qualitative patterns if we vary our estimation sample. Models (2) and (4) use observations from the most recent three months of the data, 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 we restrict the sample.

We plot the (scaled) coefficients from specification (3) against a measure of pricing technology in 7. The *x*-axis captures the pricing frequency, which increases along the *x*-axis. We report the frequency as the median number of hours between a pricing update on each website, so the axis values are reversed so that superior (more frequent) technology is to the right. Firm *E* has

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

	(1)	(2)	(3)	(4)
B	0.066***	0.047***	0.146***	0.117***
	(0.000)	(0.001)	(0.000)	(0.001)
C	0.091***	0.107***	0.171***	0.187***
	(0.000)	(0.001)	(0.000)	(0.001)
D	0.249***	0.289***	0.307***	0.337***
	(0.000)	(0.001)	(0.000)	(0.001)
E	0.284***	0.366***	0.340***	0.419***
	(0.000)	(0.001)	(0.000)	(0.001)
Product FEs	X	X	X	X
Period FEs	X	X	X	X
Sold at All Retailers			X	X
On or After Jul 1 2019		X		X
Observations	3,596,822	673,771	1,186,534	234,696

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

a median approximately equal to the number of hours in a week (168), whereas firm *A* has a median of 1. 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*). These firms 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.

Consistent with the model, we find that higher-frequency technology is correlated with lower prices. However, our model also predicts that all of these prices are elevated relative to the simultaneous price-setting game, and they may be elevated by a substantial amount. In the example plotted in Figure 2, the firm with the superior technology has prices 12 percent greater than the Bertrand equilibrium, and *all* prices would fall to this Bertrand level with no asymmetries in technology.

Further, 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 firms' algorithms depend on rivals' prices, then we should expect a price change by a low-frequency firm to increase the probability of a price change by the high-frequency firms. If, instead, firms were responding to commonly-observed

Figure 7: Price Differences Across Retailers

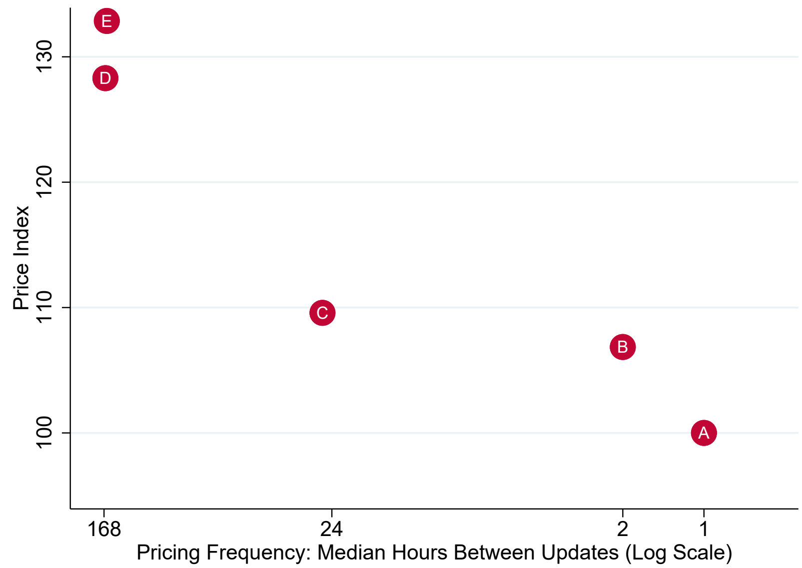

Notes: Figure displays the relative prices (Firm $A = 100$) plotted against the pricing frequency of each retailer. We report the frequency as the median number of hours between pricing updates. 168 hours corresponds to one week. The relative prices are obtained from the estimated coefficients in specification (3) of Table 4.

demand shocks and not rivals' prices, then price changes at higher-frequency pricing technology firms would anticipate those of a slower rival.[21]

To measure the reaction of prices to other firms, we take price changes occurring at retailer *E*, one of the two firms with weekly pricing technology, as the impulse function. We observe 374 price changes in our data occurring between midnight and 5 AM on Sunday. We partition the weeks into Friday through Thursday blocks, giving us a two-day pre period and a five-day post period around each price change. We then measure cumulative price changes of the same product occurring at rival retailers during each week. We capture "treated" product-weeks in which the product changed its price at retailer *E* and "control" weeks in which the product did not change its price.

Figure 8 plots the cumulative price changes before and after midnight on Sunday across each product-week. The blue line corresponds to treated product-weeks, i.e., weeks in which the price of a particular product changed at retailer *E*. The dashed line corresponds to control product-weeks that had no price change. The gap between the blue line and the dashed line is the marginal increase in price changes when a price change occurs at retailer *E*, and is analogous to a difference-in-differences estimate. We adjust for pre-period differences so the gap is zero in period 0.

Retailers *A* and *B* have an increased probability of a price change within 48 hours after

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

Figure 8: Price Changes in Response to a Price Change by Retailer *E*

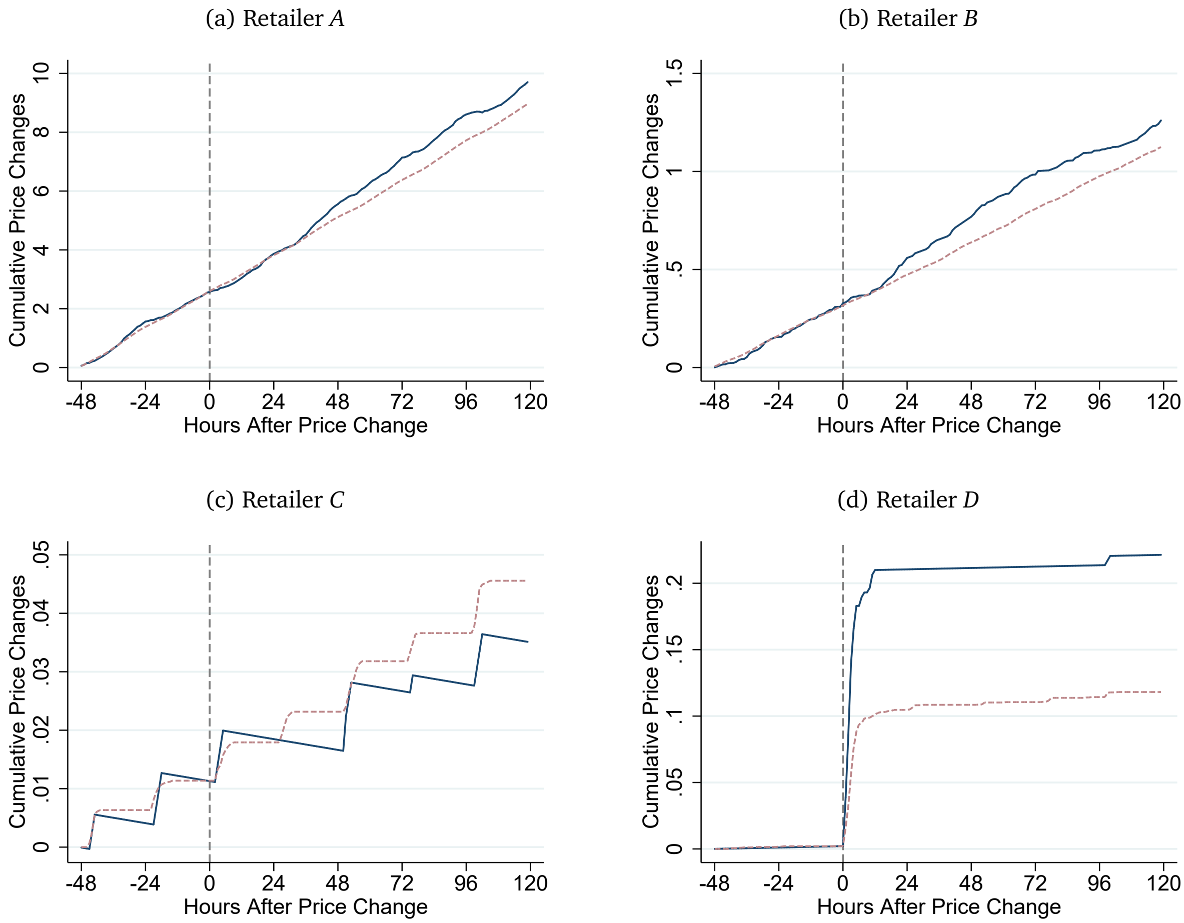

Notes: Figure displays the cumulative price changes of four retailers in response to a price change occurring at retailer *E*. The blue line displays the cumulative price change when retailer *E* changes a price of the same product in that week. The dashed line plots the cumulative price changes when the product at retailer *E* does not have a price change. The pre-period differences are netted out so that the difference is zero at period 0.

a price change at retailer *E*. Retailer *A* realizes approximately 1 additional price change on average, whereas *B* realizes roughly an additional 0.1 price change. The baseline rate of price changes at *A* is approximately 10 times that of *B*, so the proportional increase is roughly the same across the two retailers.

Retailer *D* is twice as likely to change the price of its product when the price changes at *E*. Since they update their prices only a few hours after *E*, it is likely that these changes are determined by a common unobserved factor, though it is plausible that *D* has technology in place that allows it to response to a change at *E*. The point estimates for retailer *C* indicate that it is slightly less likely to change its price after a change at *E*. However, the estimate is not precise because price changes at *C* occur at a much lower rate. We only observe 16 price changes for *C* during a product-week where *E* realizes a price change.

The evidence suggests that the two retailers with the most frequent pricing technology, *A* and *B*, are responding to other firms' prices. There is not enough data to conclude about retailer *C*. Retailer *D*, the second slowest firm, may be responding to price changes that occur a few hours earlier at *E*, or they may be determined simultaneously by an unobserved factor, such as a wholesale cost shock or a demand shock.

For the high-frequency firms, it is quite possible that this analysis underreports the degree to which they respond to *E*'s prices. The above figure captures an increased rate of price changes. The high-frequency retailers may plausibly react through the magnitudes of the price changes, while maintaining the same rate of price changes.

6 Calibration Exercise: The Impact of Algorithmic Competition

To examine the impact of algorithms on prices in our data, we perform a simple calibration exercise. We generalize the spatial demand model used in Sections 2 and 3 to allow for an arbitrary number of firms and flexible substitution patterns among firms. We then apply the model to empirically examine the five firms in our sample, taking into account the pricing technology of each firm. We find the model demand parameters that best match the moments of the data, namely prices and shares. We use these parameters to study how equilibrium prices may change if firms competed via simultaneous Bertrand competition.

One potential challenge for the empirical analysis of algorithmic competition is that the game can become computationally intractable, as the solution for one firm is an input into another firm's problem. A feature of our spatial demand model is that it generates analytical solutions for both the algorithm game and the simultaneous Bertrand game. This allows us to feasibly match the model predictions to the data and simulate alternative forms of competition.

6.1 Demand with Spatial Differentiation

We introduce a model of demand for products that are spatially differentiated. Consumers vary in their proximity to each firm, therefore the "travel" costs associated with each firm varies across consumers. The model is a generalization of the Hotelling (1929) line with two firms. Importantly, the model can capture flexible substitution patterns even with pure horizontal differentiation (i.e., homogeneous products). This is an advantage over models of vertical differentiation, such as the logit model, which restrict the horizontal substitution patterns to be symmetric across firms. A feature of the model is that it allows for positive markups when products are homogeneous.

Each firm j lies in a $(J-1)$-dimensional space. A mass of consumers μ_{jk} lie along the line segment connecting j to k.[22] The distance between each firm is 1 unit. Each firm sells a single

[22]Demand can be represented by a graph. The graph is complete if $\mu_{ij} > 0$ for all $\{i, j\}$.

product, which consumers value at $\alpha_j > 0$, and each firm chooses a price p_j. Each firm also has a mass of consumers on a line segment of distance D_0 connecting to an outside option $(j = 0)$, with $p_0 = 0$ and $\alpha_0 = 0$. Consumers lie on these segment with density μ_{j0} and mass $\mu_{j0}D_0$. D_0 may be arbitrarily large, so that the firm never captures the full segment.

Each consumer i is indexed by its location and bears a travel cost τd_{ij} for traveling a distance d_{ij} to firm j to purchase its product. A consumer along segment jk will choose j if $u_{ij} \geq 0$ and if

$$u_{ij} > u_{ik} \tag{19}$$

$$\implies (\alpha_j - p_j) - (\alpha_k - p_k) > \tau(d_{ij} - d_{ik}) \tag{20}$$

That is, the consumer will prefer j to k if the added value of product j is greater than the additional travel cost of visiting firm j. The consumer also has the option to stay home and get $u_{i0} = 0$, which he will do if $u_{ij} < 0$ and $u_{ik} < 0$.

Thus, the model flexibly captures horizontal differentiation: for all consumers that could choose product j, there are a fraction of consumers $\frac{\mu_{jk}}{\sum_{k'} \mu_{jk'}}$ that have product k as the next-best option. A mass of consumers μ_{j0} will substitute only between j and the outside option, though all consumers would choose not to buy if prices were high enough $(p_j > \alpha_j)$.

Consumers are distributed along each line segment connecting j to k according to a distribution F_{jk} with support $[0, 1]$. We assume that the distribution is symmetric about the midpoint of the segment. Symmetry implies $F_{jk} = F_{kj}$, so the direction of the connection is arbitrary. We also assume that the same distribution is applied to all segments: $F_{jk} = F$, though this could easily be relaxed. Demand along each segment can then be characterized by the distribution function F.

Noting that $d_{ik} = 1 - d_{ij}$ for a consumer on segment jk, we have that a consumer on this segment will choose j if

$$u_{ij} > u_{ik} \tag{21}$$

$$\implies \frac{1}{2} + \frac{1}{2\tau}\left((\alpha_j - p_j) - (\alpha_k - p_k)\right) > d_{ij} \tag{22}$$

and if

$$u_{ij} \geq 0 \tag{23}$$

$$\implies \frac{1}{\tau}\left(\alpha_j - p_j\right) \geq d_{ij} \tag{24}$$

Firm j receives customers for which d_{ij} satisfies both conditions above. Therefore, firm j receives a quantity of $\mu_{jk}F(y_{jk})$ from line segment jk, where

$$y_{jk} = \min\left\{\frac{1}{2} + \frac{1}{2\tau}\left((\alpha_j - p_j) - (\alpha_k - p_k)\right), \frac{1}{\tau}\left(\alpha_j - p_j\right)\right\}.$$

For the outside segments, $y_{j0} = \frac{1}{D_0}\frac{1}{\tau}(\alpha_j - p_j)$, as these segments have length D_0 instead of 1. The parameter D_0 can also be interpreted as the relative travel cost of choosing the outside option relative to an inside good, as the model has an isomorphic parameterization with outside travel costs $\tilde{\tau}_0 = D_0\tau$.

Overall, quantities are given by

$$q_j = \sum_{k \neq j} \mu_{jk} F(y_{jk}). \tag{25}$$

The flexibility in substitution patterns from this relatively parsimonious model comes primarily through the mass of consumers on each segment $\{\mu_{jk}\}$ and the choice of distribution F. In equilibrium, the consumers $\{\mu_{j0}\}$ that have no next-best alternative other than the outside option are also important in determining substitution patterns.

Contested and Uncontested Segments

We introduce some terminology to facility discussion of the model. When $\max(u_{ij}, u_{ik}) \geq 0$ for all i on segment jk and $y_{jk} < 1$, the segment is *contested*.[23] When some consumers prefer to stay home, rather than purchase, the segment is *uncontested*. If segment jk is uncontested, there is no consumer indifferent between j and k, so those firms have local monopoly power over a portion of consumers on that segment. That is, a change in the price of firm k does not affect demand for firm j at the margin. When all segments between firms (the "inside" segments) are contested, we say the market is *covered*. For a covered market, all consumers on inside segments purchase.

6.2 Calibration

For our calibration exercise, we parameterize F as a uniform distribution. When the market is covered, $F(y_{jk}) = y_{jk}$.[24] We also assume that the products are homogeneous (but for the travel costs), so that $\alpha_j = \alpha$ for all j except for the outside option, for which $\alpha_0 = 0$. Demand for retailer j is equal to

$$q_j = \sum_{k \neq j, 0} \mu_{jk}\left(\frac{1}{2} - \frac{1}{2\tau}(p_j - p_k)\right) + \mu_{j0}\frac{1}{\tau}(\alpha - p_j). \tag{26}$$

[23] When $y_{jk} \geq 1$, the segment is *dominated* by j.

[24] For the inside segments, y_{jk} must be between zero and one if the segments are contested. For the outside segment, y_{j0} is scaled by the distance D_0 so that it falls between zero and one.

Table 5: Calibrated Segment Weights

	A	B	C	D	E	Outside
A	0.00	11.22	2.05	0.52	0.52	0.00
B	11.22	0.00	2.05	0.52	0.52	1.79
C	2.05	2.05	0.00	0.52	0.52	1.42
D	0.52	0.52	0.52	0.00	0.52	3.02
E	0.52	0.52	0.52	0.52	0.00	3.93

Notes: Row j column k shows the mass of customers on the segment between firm j and k (μ_{jk}). The weights are symmetric; for convenience, they are displayed twice ($\mu_{jk} = \mu_{kj}$), representing the perspective of each firm. The outside segment weights represent the share of customers captured from the outside segments at the equilibrium prices.

Since we calibrate our model to an index of average prices and aggregate shares, p_j corresponds to j's average price. We also assume all firms have the same marginal cost, which we normalize to 1. Price-cost margins are determined by the calibrated prices in the model.

We allow firms to be asymmetric in their pricing technology, corresponding to our findings in Section 4. Retailers *D* and *E* set prices simultaneously, followed by retailer *C*, then *B*, and, finally, *A*. The sequence can be interpreted as arising from asymmetries in frequency (as in Section 2) or from asymmetric algorithms (as in Section 3). The model is analogous to a sequential price-setting game.

The unknown parameters to be recovered are the value of the product α, the travel cost parameter τ, and the relative weights on the segments $\{\mu_{jk}\}$. We parameterize the J by $(J+1)$ μ matrix with six parameters: $\{m_1, m_2, m_3, m_4, m_5, m_6\}$. While the fact that prices are negatively correlated with higher-pricing frequency is consistent with the model, this may also be due in part due to the fact that demand is not symmetric. In other words, consumers may have a preference for firms with lower pricing frequency. In the calibration, we allow for substitution patterns that could explain differential pricing across firms.

For the slower firms, *D* and *E*, we assume symmetric weights m_1 to all other retailers. The firm with daily pricing, *C*, has symmetric weights with the faster firms (m_2). The two fastest firms have a unique weight m_3. Finally, we give each firm a unique mass for the outside option, normalizing the mass for *E* to 1. We also set the mass along the outside option for *A* to zero.[25] This assumption is made because this retailer does not have any in-store sales for this market; we are imposing that the all of *A*'s marginal customers would substitute to one of the other four online retailers at the equilibrium prices.

We use the method of moments to choose the parameters $(\alpha, \tau, \{\mu_{jk}\})$ that best fit the relative prices and shares we observe in the data. We minimize the sum of squared deviations from relative average prices, taken from specification (1) of Table 4, and relative average shares. We measure shares using the average Google search frequency shares for the retailer and for

[25] Thus, $(\mu_{A0}, \mu_{B0}, \mu_{C0}, \mu_{D0}, \mu_{E0}) = (0, m_6, m_5, m_4, 1)$.

Figure 9: Calibration Exercise

(a) Markup (b) Shares

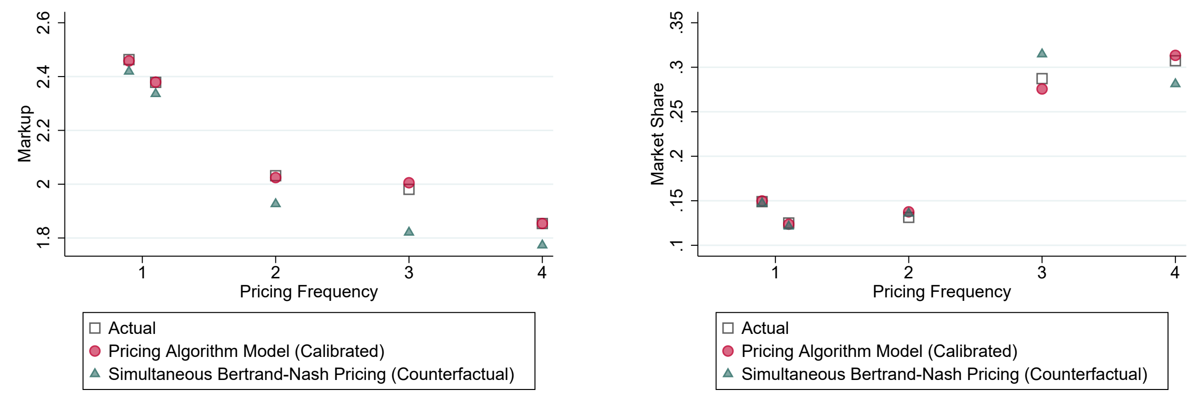

Notes: Figure displays the markups (panel (a)) and the relative shares (panel (b)) plotted against the pricing frequency of each retailer. Frequency is normalized to the relative sequence. The black squares indicate the data, and the red dots are the fitted prices from a calibration exercise. The relative prices are obtained from the estimated coefficients in specification (1) of Table 4. The markup level is pinned down by the calibrated model. The green triangles display the counterfactual simultaneous Bertrand markups at the calibrated parameters and the corresponding shares.

allergy products at that retailer.[26]

The calibrated parameters for the value of the product and travel costs are $\alpha = 5.09$ and $\tau = 0.67$. The calibrated segment weights are displayed in Table 5. These parameters generate an equilibrium mean price of 2.07. As marginal costs are normalized to 1, prices may be interpreted as markups (price over cost). Mean realized travel costs are 0.61. Thus, we estimate that, net of travel costs, willingness to pay is roughly twice the equilibrium price.

The fit of the calibration exercise is displayed in Figure 9. In panel (a), squares indicate the relative prices in the data; these prices are translated to markups based on the calibrated model. The *x*-axis displays the pricing frequency in terms of the relative sequence. The red dots indicate the markups from the calibrated model. Likewise, the black squares in panel (b) represent observed shares, and the red dots indicated the predicted shares from the model. Our eight-parameter model fits prices and shares quite well. Though we fit relative prices among the firms, underlying marginal costs play an important role in determining equilibrium in the model. Marginal costs are pinned downed by the first-order conditions, allowing us to recover an estimate of markups. The calibrated parameters imply reasonable price-cost margins between 0.461 (retailer *A*) and 0.593 (retailer *E*).

[26]In calibration, we impose a penalty if the parameters result in a firm capturing more than 95 percent of the consumers on a given segment. This ensures that the counterfactual simultaneous Bertrand prices have an interior solution. The resulting penalty is small and the constraint does not meaningfully affect our estimates. Our counterfactual effects are robust to alternative share definitions that are based on category revenues or a combination of revenues and search data.

Table 6: Own-Price and Cross-Price Demand Elasticities

	Price				
Share	A	B	C	D	E
A	-2.17	1.84	0.34	0.10	0.10
B	1.93	-2.81	0.39	0.12	0.12
C	0.71	0.77	-2.18	0.23	0.24
D	0.20	0.22	0.22	-1.76	0.27
E	0.17	0.18	0.18	0.21	-1.72

Notes: Row j column k shows $(\partial q_j/\partial p_k)(p_k/q_j)$.

Table 6 shows a matrix of elasticity of demand estimates from the model. Own-price elasticities ranges from -1.7 to -2.8, roughly consistent with other estimates from online goods.[27] Our estimated cross-price elasticities indicate that, when the price of a product at a firm increases, consumers are more likely to substitute towards to similar firms, e.g., consumers from retailer *A* are more likely to substitute to *B* and consumers from retailer *E* are more likely to substitute to *D*.

6.3 Counterfactual

To illustrate the potential impact of pricing algorithms on prices, we use our calibrated model to predict the equilibrium price if all firms were instead setting prices simultaneously. The simultaneous Bertrand equilibrium prices and shares are displayed with green triangles in Figure 9. Our model indicates that algorithmic competition increases the average price by 5.2 percent above the counterfactual Bertrand equilibrium. These price changes differ across firms. Firms *D* and *E* realize more modest price changes of 1.9 and 1.6 percent. Based on our calibration, these firms receive a greater relative share of consumers from outside segments, rendering their behavior closer to that of a (local) monopolist. Competition for customers is more intense between the other three firms, who realize price increases between 4.5 and 10.1 percent as a result of algorithmic competition.

The results from the counterfactual exercise are presented in Table 7. Algorithmic competition has the biggest impact on shares for firm *B*, which sees a 3.9 percentage point (12 percent) decline in market share relative to the counterfactual Bertrand environment. The majority of this shift in share accrues to Firm *A*, which increases market share by 3.2 percentage points. The remaining 0.7 percent lost by Firm *B* result in modest increases for the other three firms. The differential effects on prices and quantities generate heterogeneous effects on firm profits. Because retailer *A* realizes meaningful increases in both price and quantity as a result of algorithmic competition, it sees the largest gain in profits (22 percent). Despite lower quantities, retailer *B*'s price increase is great enough to generate a 6 percent increase in profits. By con-

[27] See, for instance, De los Santos et al. (2012).

Table 7: Counterfactual Effects on Markups and Profits

	Simultaneous Bertrand			Algorithmic Competition			Percent Change		
Firm	Markup	Share	Profit	Markup	Share	Profit	Markup	Share	Profit
A	1.77	0.281	6.4	1.85	0.313	7.8	4.6	11.5	22.0
B	1.82	0.315	7.6	2.01	0.276	8.1	10.1	-12.4	6.3
C	1.93	0.136	3.7	2.02	0.138	4.1	5.1	1.3	11.1
D	2.34	0.121	4.8	2.38	0.124	5.0	1.9	2.0	4.4
E	2.42	0.147	6.1	2.46	0.150	6.4	1.6	1.8	3.7
Aggregate	1.97	1	28.6	2.07	1	31.3	5.2	0	9.6

Notes: Table displays the implied markups, shares, and profits from the calibrated model. The first three columns report the counterfactual estimates with simultaneous Bertrand price-setting behavior. The middle three columns report the predicted values from the model of algorithmic competition that is fitted to the data. The final three columns report the percent changes of moving from simultaneous Bertrand to algorithmic competition. Profits are arbitrarily scaled so that 1 unit corresponds to $100 million of e-commerce in the Personal Care category.

trast, retailers *D* and *E* realize profit gains of about 4 percent from more modest increases in both price and quantity. Consistent with the stylized results in Section 2, all firms profit as a result of algorithmic competition.

Our model predicts that algorithmic competition results in a modest decline in quantity purchased of 0.9 percent. As all segments are contested in the algorithm equilibrium, marginal consumers are obtained via substitution away from the outside option into the market. This limited substitution outside of the market means that effects on total welfare are small (a decline of 0.3 percent). Algorithmic competition in our calibrated model serves primarily as a transfer between firms and consumers: consumer surplus falls by 4.1 percent, and firm profits increase by 9.6 percent. To assign a dollar value to these effects, we can do a rough back-of-the-envelope calculation. These five firms have annual e-commerce revenues of approximately $6 billion in the category of Personal Care. If we assume that our estimated price effects apply to the entire category, then consumer surplus for the category would improve by $300 million annually by moving from algorithmic competition to simultaneous Bertrand price-setting.

7 Discussion

In this paper, we provide 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. Our calibrated model suggests that these asymmetric pricing algorithms raise all prices relative to the simultaneous Bertrand-Nash equilibrium.

While these results are subject to a number of caveats, we view the results as initial evidence that pricing algorithms can raise prices. When it comes to pricing algorithms, policy makers should not just be concerned with collusion, rather 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. Enforcement of this may be challenging, but the burden could be substantially less than trying to divine whether algorithms have reached an "agreement" to collude.

What is special about prices? Because retail prices are public and immediately available to rival firms, they allow for short-run commitment that shapes the nature of competition. If firms were prohibited from using rivals' prices, one could imagine firms using algorithms based on rivals' quantities, inventories, or other factors. However, these data are rarely made public at the frequency necessary to support a short-run commitment.

Other possible policies include limiting the scraping of rival firms' prices, prohibiting the storage of recent prices by other firms, or regulating that price updates occur with an industry-standard (symmetric) frequency, such as once per day.[28] A hybrid solution that combines two or more approaches may be a more feasible long-run solution. For example, firms could participate in industry-standard pricing, or they could opt out and subject themselves to random reviews of their algorithms, to demonstrate that they do not depend on rivals' prices. Though the regulatory environment to enable these policies does not yet exist, one can envision a future agency that is empowered to enforce the above policies.

Of course, algorithms may have several benefits, such as the ability to more efficiently respond to time-varying demand. These beneficial considerations should be taken into account when considering any policy prescription.

[28]The standard could be jointly determined by market participants and a regulatory authority, per a suggestion from Fiona Scott Morton.

References

Benoit, J.-P. and V. Krishna (1985). Finitely Repeated Games. *Econometrica 53*(4).

Bresnahan, T. F. (1981). Duopoly Models with Consistent Conjectures. *The American Economic Review 71*(5).

Byrne, D. P. and N. de Roos (2019). Learning to Coordinate: A Study in Retail Gasoline. *The American Economic Review 109*(2), 591–619.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2019). Artificial Intelligence, Algorithmic Pricing and Collusion. *SSRN Electronic Journal*.

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

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review 102*(6), 2955–80.

Fudenberg, D. and E. Maskin (1986). The Folk Theorem in Repeated Games with Discounting or with Incomplete Information. *Econometrica 54*(3), 533–554.

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

Green, E. J. and R. H. Porter (1984). Noncooperative Collusion under Imperfect Price Information. *Econometrica 52*(1), 87–100.

Holt, C. A. (1985). An experimental test of the consistent-conjectures hypothesis. *The American Economic Review 75*(3), 314–325.

Hotelling, H. (1929, March). Stability in Competition. *The Economic Journal 39*(153), 41.

Igami, M. and T. Sugaya (2018). Measuring the Incentive to Collude: The Vitamin Cartels, 1990-1999. *SSRN Electronic Journal*.

Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet (2010, May). A commitment folk theorem. *Games and Economic Behavior 69*(1), 127–137.

Kamien, M. and N. L. Schwartz (1983). Conjectural Variations. *The Canadian Journal of Economics 16*(2), 191–211.

Maskin, E. and J. Tirole (1988, May). A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles. *Econometrica 56*(3), 571–599.

Miklós-Thal, J. and C. Tucker (2019). Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers? *Management Science 65*(4), 1552–1561.

Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica 85*(6), 1763–1791.

O'Connor, J. and N. E. Wilson (2019). Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition.

Osborne, D. K. (1976). Cartel Problems. *The American Economic Review 66*(5), 835–844.

Peters, M. and B. Szentes (2012). Definable and Contractible Contracts. *Econometrica 80*(1), 363–411.

Rotemberg, J. J. and G. Saloner (1986). A Supergame-Theoretic Model of Price Wars during Booms. *The American Economic Review, 76*(3).

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

Salop, S. C. (1979). Monopolistic Competition with Outside Goods. *The Bell Journal of Economics 10*(1), 141.

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

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

A Equilibrium Selection

Figure 10: Equilibrium Selection with Pricing Algorithms

(a) Firm 2 Only

(b) Firm 1 and Firm 2

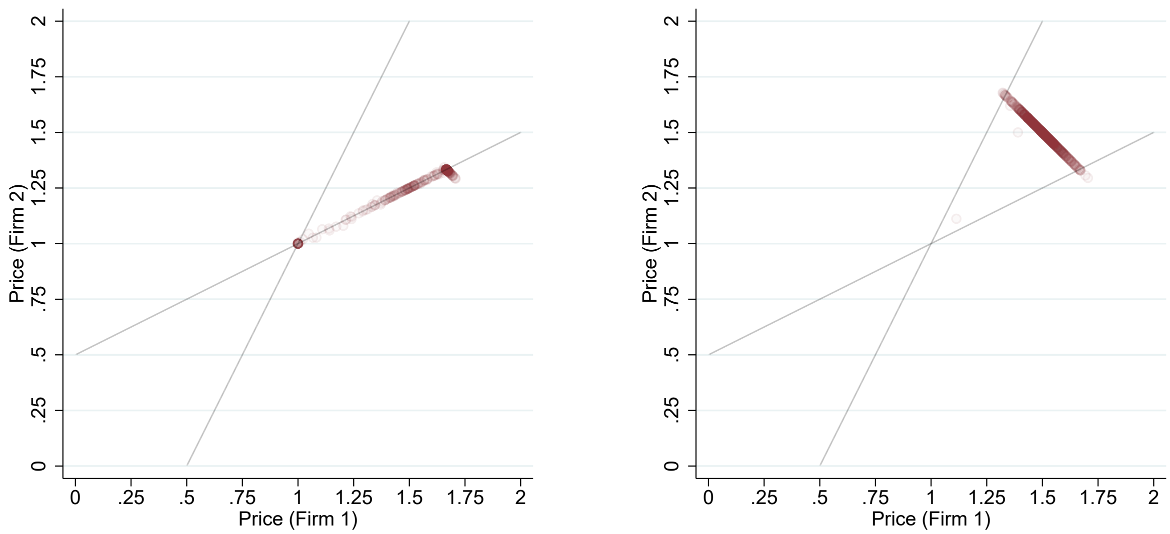

Notes: Figure displays the resulting prices from 500 simulated duopoly markets when firms use a simple learning rule to update their prices or pricing algorithms. Each firm will update its algorithm if a random deviation in the algorithm parameters improve profits. Any stable point in simulation is an equilibrium (no profitable deviation exists). Each point displays the prices after 10,000 experiments. Panel (a) displays the results from the asymmetric algorithm game (firm 1 chooses price). Panel (b) displays the results from the game where both have algorithms. The plotted lines indicate the two price-setting best-response functions; their intersection is the unique Bertrand-Nash equilibrium.

Our results show that *any* price vector can be supported by algorithms that are linear functions of other firms' prices (Proposition 3). However, we think it likely that algorithmic competition would result in higher prices. This is supported by the simple intuition that firms only have the incentive to adopt these algorithms if it would improve profits above the price-setting equilibrium.

To test this intuition, we simulate a simple learning process to select equilibria. We follow the duopoly setup of Section 2.1 and allow firms to choose linear algorithms: $p_{jt} = a_{jt} + b_{jt}p_{kt}$. We initialize each firm with random parameters a_{j0} and b_{j0}. Each period, one (randomly-chosen) firm runs an experiment, modifying their parameters: $\tilde{a}_{jt+1} = a_{jt} + \varepsilon_t^1$ and $\tilde{b}_{jt+1} = b_{jt} + \varepsilon_t^2$. If this experiment improves profits, the firm updates their benchmark to the new parameters ($(a_{jt+1}, b_{jt+1}) = (\tilde{a}_{jt+1}, \tilde{b}_{jt+1})$), otherwise, they revert to the previous parameters ($(a_{jt+1}, b_{jt+1}) = (a_{jt}, b_{jt})$).

A "rest point" of this game is an equilibrium, i.e., where no unilateral deviation exists. To find the rest points, we simulate 10,000 experiments in each of 500 duopoly markets. The resulting prices are displayed in Figure 10. Panel (a) displays the results from the asymmetric game in which firm 1 is a price-setter and firm 2 chooses an algorithm. The resulting prices, as would be expected, lie along firm 2's best-response function and are (weakly) higher than the simultaneous Bertrand-Nash equilibrium, $(1, 1)$. There is a mass at the Bertrand-Nash equilibrium, at firm 1's optimal choice conditional on the best-response of firm 2, and at the joint profit-maximizing point along firm 2's best-response function.

Panel (b) shows the resulting prices from the game in which both firms have pricing algorithms. The prices are centered around the collusive equilibrium ($(1.5, 1.5)$) and lie along the profit Pareto frontier. The equilibria are bounded by the two firms' best-response functions.

Our simulation of a simple learning process selects equilibria with higher prices. The resulting prices are bounded from below by each firm's best-response function and bounded from above by the profit Pareto frontier.