

Competition in Pricing Algorithms

Zach Y. Brown
Alexander MacKay

Working Paper 20-067



Competition in Pricing Algorithms

Zach Y. Brown
University of Michigan

Alexander MacKay
Harvard Business School

Working Paper 20-067

Copyright © 2019, 2020, 2021 by Zach Y. Brown and Alexander MacKay.

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Funding for this research was provided in part by Harvard Business School.

Competition in Pricing Algorithms*

Zach Y. Brown **Alexander MacKay**
University of Michigan[†] Harvard University[‡]

April 29, 2021

Abstract

Increasingly, retailers have access to better pricing technology, especially in online markets. Using hourly data from five major online retailers, we show that retailers set prices at regular intervals that differ across firms. In addition, faster firms appear to use automated pricing rules that are functions of rivals' prices. These features are inconsistent with the standard assumptions about pricing technology used in the empirical literature. Motivated by these facts, we consider a model of competition in which firms can differ in pricing frequency and choose pricing algorithms rather than prices. We demonstrate that, relative to the standard simultaneous price-setting model, pricing technology with these features can increase prices in Markov perfect equilibrium. A simple counterfactual simulation implies that pricing algorithms lead to meaningful increases in markups in our empirical setting, especially for firms with the fastest pricing technology.

Keywords: Pricing Algorithms, Pricing Frequency, Online Competition
JEL Classification: L40, D43, L81, L13, L86

*We thank John Asker, Emilio Calvano, Giacomo Calzolari, Matt Grennan, George Hay, Scott Kominers, Fernando Luco, Nate Miller, Marc Rysman, Mike Sinkinson, Konrad Stahl, and Ralph Winter for helpful comments. We also thank seminar and conference participants at Harvard Business School, the IIOC, the ASSA Meeting (Econometric Society), the Toulouse Digital Economics Conference, the NYU Law/ABA Antitrust Scholars Conference, the Winter Business Economics Conference, the NBER Economics of Digitization meeting, Brown University, the FTC Microeconomics Conference, Monash University, the MaCCI Annual Conference, and Stanford University. We are grateful for the research assistance of Pratyush Tiwari and Alex Wu.

[†]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 are becoming more prevalent. Algorithms can change pricing behavior by enabling firms to update prices more frequently and automate pricing decisions. Thus, firms can commit to pricing strategies that react to price changes by competitors. This may have important implications for price competition relative to standard oligopoly models in which firms set prices simultaneously. Do pricing algorithms lead to higher prices?

In this paper, we present new facts about pricing behavior that highlight the above features of pricing algorithms. Using a novel dataset of high-frequency prices from large online retailers, we document pricing patterns that are (i) consistent with the use of automated software and (ii) inconsistent with the standard empirical model of simultaneous price-setting behavior. Retailers update prices at regular intervals, but these intervals differ across firms, allowing some retailers to adjust prices at higher frequencies than their rivals. Firms with faster pricing technology quickly respond to price changes by slower rivals, indicating commitment to automated strategies that depend on rivals' prices. Finally, we examine price dispersion, and we show that price differences across retailers are related to asymmetries in pricing technology.

Motivated by these facts, we introduce a new model of price competition that incorporates increased pricing frequency and short-run commitment through the use of algorithms. Our model also allows for asymmetric technology among firms. We show that asymmetry in pricing technology can fundamentally shift equilibrium behavior: if one firm adopts superior technology, both firms can obtain higher prices. If both firms adopt high-frequency algorithms, collusive prices can be supported without the use of traditional collusive strategies. Further, we demonstrate that algorithms that depend on rivals' prices will not deliver competitive Bertrand prices in equilibrium. Thus, we illustrate how pricing algorithms can generate supracompetitive prices through novel, non-collusive mechanisms.¹ Our results show that the competitive impacts of algorithms can be quite broad. Frequency, commitment, and asymmetry in pricing technology allow firms to support higher prices in competitive (Markov perfect) equilibrium.

We use our model to analyze the potential empirical implications of differences in pricing technology. The model can rationalize why firms that have higher-frequency pricing have lower prices than their competitors for identical products, even when the firms are otherwise symmetric. Thus, our model provides a supply-side explanation for price dispersion, presenting a complementary alternative to the demand-side explanations that are emphasized in the literature, such as the presence of search frictions. We use a counterfactual simulation to quantify these impacts, finding that asymmetric pricing technology leads to higher prices for all retailers and exacerbates price differences among similar retailers.

¹The existing literature has focused on whether algorithms can facilitate collusion, almost exclusively assuming that firms have symmetric, price-setting technology (e.g., Calvano et al., 2020; Miklós-Thal and Tucker, 2019; Salcedo, 2015).

We begin by highlighting the key features of pricing algorithms used by online retailers (Section 2). We present three stylized facts using high-frequency price data for over-the-counter allergy medications for the five largest online retailers in the category. First, we document heterogeneity in pricing technology. Two firms have high-frequency algorithms that change prices within an hour, one firm updates prices once per day, and the remaining two have weekly pricing technology, updating their prices early every Sunday morning. Second, we show that the fastest firms quickly react to price changes by slower rivals, consistent with the use of automated pricing algorithms that monitor rivals' prices and follow a pre-specified strategy. Third, we show that asymmetric pricing technology is associated with asymmetric prices. Relative to the firm with the fastest pricing technology, the firm with daily pricing technology sells the same products at prices that are 10 percent higher, whereas the firms with weekly pricing technology sell those products at prices that are approximately 30 percent higher. These facts are inconsistent with the widespread assumption that firms have essentially symmetric price-setting technology in online markets.

We introduce an economic framework to capture these features of online price competition. We study competitive equilibria when firms may have high-frequency algorithms that condition on rivals' prices. To illustrate key mechanisms and build intuition, we introduce features of the model sequentially. We first consider asymmetry across firms through differences in pricing frequency alone. We then allow for asymmetry in the ability of firms to autonomously react to rivals' prices. Finally, we consider the case where all firms have high-frequency algorithms that condition on rivals' prices.

We examine differences in pricing frequency in Section 3. We show that the model generates prices that lie between the simultaneous (Bertrand) and sequential (Stackelberg) equilibria and nests both as special cases. When prices are strategic complements, as is typical in empirical models of demand, the faster firm has lower prices and higher profits than the slower firm. Thus, our model provides a supply-side explanation for the price dispersion observed in the data. Moreover, when firms can choose their pricing frequency, asymmetric frequencies are the equilibrium outcome. Each firm has a unilateral profit incentive to choose either more frequent or less frequent pricing than their rivals. Therefore, the simultaneous price-setting model is not an equilibrium outcome when pricing frequency is endogenous.

We develop a more general model in Section 4. In the model, algorithms enable firms to differ in their pricing frequency and also to commit to a pricing strategy for future price updates. This model nests the pricing frequency game developed earlier. Further, we show that a model with asymmetric commitment—i.e., when only one firm can condition its algorithm on its rival's price—closely parallels the model of asymmetric frequency. Equilibrium prices lie between the simultaneous and sequential equilibria, depending on how quickly the algorithm can react. When pricing algorithms react instantaneously, the model generates the sequential (Stackelberg) equilibria where the algorithmic competitor is the follower.

We then analyze the case where all firms can condition on rivals' prices, deriving a one-shot competitive game in which firms submit pricing algorithms, rather than prices. We use the one-shot game to show that symmetric short-run commitments, in the form of automated pricing, can also generate higher prices. To demonstrate the significant implications of this dimension of algorithmic competition, we focus on equilibrium pricing strategies that, in some sense, "look competitive." That is, we eliminate collusive strategies that rely on cooperate-or-punish schemes. Even with these restrictions, pricing algorithms can increase prices relative to the Bertrand-Nash equilibrium. Supracompetitive prices, including the fully collusive prices, can be supported with algorithms that are simple linear functions of rivals' prices.² In this way, algorithms fundamentally change the pricing game, providing a means to increase prices without resorting to collusive behavior.

We also address the question of whether pricing algorithms can arrive at competitive (Bertrand) prices. Our model provides a stark negative result: all firms will not choose price-setting best-response (Bertrand reaction) functions in equilibrium. Further, if any firm uses an algorithm that depends on a rival's price, Bertrand prices do not arise in equilibrium. Intuitively, our results are supported by the following logic: A superior-technology firm commits to best respond to whatever price is offered by its rivals, and its investments in frequency or automation makes this commitment credible. The rivals take this into account, softening price competition. Our model nests several different theoretical approaches that were developed prior to the advent of pricing algorithms and have largely been dismissed in the modern literature, including conjectural variations. We highlight these connections below.

The empirical literature on price competition and firm markups has almost exclusively assumed that firms play a simultaneous pricing game. As a first step toward quantifying the role of heterogeneous pricing technology, we compare observed prices to a counterfactual equilibrium in which firms have simultaneous price-setting technology (Section 5). We introduce a model of demand that allows for flexible substitution patterns among retailers and provides a tractable empirical approach to modeling supply-side competition with algorithms. Using the observed pricing technology of the retailers as an input, we fit the model to average prices and market shares in our data. We then use the estimated demand parameters to 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. This corresponds to a 9.6 percent increase in profits and a 4.1 percent decrease in consumer surplus. The effect on markups and profits is especially large for firms with superior pricing technology, i.e., those with the ability to quickly

²In practice, it is typical for algorithms to 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/>

adjust prices.

Online markets have allowed retailers to gather high-frequency data on rivals' prices and react quickly through the use of automated software. Indeed, these are key features advertised by third-party providers of pricing algorithms.³ Evidence suggests that algorithms are becoming more widespread as online retailing continues to grow (Cavallo, 2018). This growing prevalence of pricing algorithms has drawn significant attention from competition authorities.⁴

Overall, our results imply that pricing algorithms can support higher-price equilibria, even when firms act competitively. Our empirical analysis shows price patterns consistent with the model and suggests that pricing algorithms can have an economically meaningful effect on markups. Thus, if policymakers are concerned that algorithms will raise prices, then the concern is more broad than that of collusion. Of course, algorithms may also have several benefits, such as the ability to more efficiently respond to time-varying demand. In light of these issues, we briefly discuss implications for policymakers in Section 6. Though we focus on competitive equilibria, our study also has implications for collusion. By increasing competitive prices and profits, algorithms may make punishment less severe in a collusive scheme, reducing the likelihood of collusion. Additionally, our model explicitly features a new dimension in the strategy space, allowing firms to change pricing technology as an either a substitute or a complement to the pursuit of collusion.

Related Literature

We contribute to the nascent literature studying the impacts of algorithms on prices. We present a new model of price competition to capture features of algorithms—frequency and commitment—that have not been studied previously. The existing literature has focused on the price effects of *learning* algorithms (Salcedo, 2015; Calvano et al., 2020; Asker et al., 2021) or *prediction* algorithms (Miklós-Thal and Tucker, 2019; O'Connor and Wilson, 2019) in the context of a standard simultaneous price (or quantity) game. This literature focuses on how learning or prediction algorithms affect the sophistication of players and their ability to collude.⁵ The equilibria of the environments studied by these papers have been extensively studied. By contrast, we examine how pricing algorithms change the nature of pricing game, focusing on Markov perfect equilibria as in Maskin and Tirole (1988b).⁶ Our model generates a new set of

³For instance, ChannelAdvisor advertises its automated pricing product as “constantly monitoring top competitors on the market.” Repricer.com “reacts to changes your competitors make in 90 seconds.” Intelligence Node allows retailers to “have eyes on competitor movements at all times and...automatically update their prices.”

⁴See, for instance, the U.K. Competition and Markets Authority's 2018 report, “Pricing Algorithms” and Germany's “Twenty-second Biennial Report by the Monopolies Commission.” Thus far, government authorities have focused on the potential for algorithms to facilitate collusion.

⁵Klein (2019) considers the same question but in the alternating-move setting of Maskin and Tirole (1988b).

⁶Maskin and Tirole (1988b) show that higher prices can result in a duopoly game where firms set prices in alternate periods using strategies that rely exclusively on payoff-relevant variables. Our analysis complements their work by showing how higher prices may be obtained in Markov perfect equilibrium in a different economic environment—one in which algorithms provide variation in pricing frequency and enable short-run commitment.

equilibrium strategies and outcomes that can be supported by algorithms.

There has been little empirical evidence on the pricing strategies used by major online retailers. Using surveys and case studies, competition authorities have noted that online firms may collect information on prices of competitors and use the information to adjust their own prices.⁷ Some studies in the computer science literature have examined pricing rules employed by third-party sellers that use rivals' prices as an input (Chen et al., 2016). Our novel high-frequency dataset allows us to document, systematically, new empirical facts about the pricing behavior of online competitors. In an offline context, a recent paper by Assad et al. (2020) examines whether algorithms change pricing strategies and increase prices in retail gasoline markets.⁸

We provide a new framework for quantifying the effects of pricing technology on prices, contributing to the empirical literature that studies supracompetitive prices (e.g., Porter, 1983; Nevo, 2001; Miller and Weinberg, 2017; Byrne and de Roos, 2019). Our model and empirical results suggest that the mode of competition can lead to meaningful price increases without the need for collusion. Previous empirical studies of supracompetitive prices have exclusively considered stage games with symmetric technology where firms choose actions (price or quantity) simultaneously; this framework has been the basis for antitrust analysis as well.⁹ Our analysis takes a first step toward incorporating heterogeneous pricing technology and quantifying its implications.

Our findings also contribute to the broader literature on price dispersion in online markets by providing an explanation for differences in prices for identical products across firms. Despite the fact that online competition is thought to reduce search costs and expand geographic markets, substantial price dispersion has been documented (e.g., Baye et al., 2004; Ellison and Ellison, 2005). An empirical literature has focused on demand-side features such as search frictions, but little attention has been paid to firm conduct.¹⁰ One exception is Ellison et al. (2018), who examine managerial inattention and price dispersion in an online marketplace in 2000 and 2001, prior to the widespread use of pricing algorithms. Our results suggest that differences in pricing technology across firms leads to persistent differences in prices for identical products.

We argue that a key feature of pricing algorithms is the ability to condition on the prices of rivals. This mechanism relates to a large class of models where firms internalize the reactions of their rivals, including conjectural variations (Bowley, 1924) and the classic Stackelberg model. The real-world applicability of these models has been subject to a long debate (e.g., Fellner, 1949). The conjectural variations model has fallen out of favor, likely because consistent conjectures other than Cournot are difficult to rationalize (Daughety, 1985; Lindh, 1992). Models

⁷See, for instance, the European Commission's 2017 report, "E-commerce Sector Inquiry."

⁸Assad et al. (2020) find evidence for price effects only when both firms in duopoly markets adopt superior pricing technology, which suggests that the mechanism in their setting may be collusion or symmetric commitment.

⁹See, for instance, "Commentary On The Horizontal Merger Guidelines" by the U.S. Department of Justice.

¹⁰Work examining online search frictions includes Hong and Shum (2006), Brynjolfsson et al. (2010), and De los Santos et al. (2012).

with sequential behavior have been dismissed as unrealistic for empirical settings because it requires the assumption that one firm can honor a (sub-optimal) commitment to an action or strategy while the other reacts. For this reason, applied researchers and antitrust authorities have almost universally assumed that firms play a simultaneous Bertrand or Cournot game. We argue that such commitments are credible, made possible by investments in differential pricing technology. Algorithms provide a natural mechanism for the type of technological commitment discussed in Maskin and Tirole (1988a). Thus, one interpretation of our model is that it provides a new foundation for theoretical results arising in this older literature. By nesting these models under a common structure, we also provide a framework for firms to choose among different models of competition by changing their pricing technology.

The logic of how pricing algorithms leads to higher prices is similar to that of price-matching guarantees, which some have argued can be anticompetitive (Salop, 1986; Hay, 1981; Moororthy and Winter, 2006). Both are predicated on commitment, which software makes possible in online markets.¹¹ We show that price-matching guarantees are not chosen in equilibrium in our model. There are also parallels between our model and previous literature focused on commitment in other settings. Grossman (1981) and Klemperer and Meyer (1989) study supply function equilibrium in which firms simultaneously decide on quantities in response to a (endogenously-determined) market price in a setting with homogeneous products. Lazarev (2019) shows that higher prices can result when firms first commit to a restricted set of prices, then choose from among those prices in a second stage. Conlon and Rao (2019) find that wholesalers can set the collusive price when they can commit to a price schedule. The game-theoretic notion of commitment ties into a broader literature on strategic delegation that has been applied in diverse settings.¹² We consider algorithms to be an economic mechanism to make such commitments credible.¹³ Moreover, we are the first to link pricing algorithms to models with these features.

¹¹Hal Varian discussed the appeal of price matching in online markets in the August 24, 2000 *New York Times* article “When commerce moves online, competition can work in strange ways.” In a set of lab experiments, Deck and Wilson (2000, 2003) find that subjects that use automated price-matching strategies obtain higher profits than those that manually set prices.

¹²Fershtman and Judd (1987) and Sklivas (1987) show that, by giving managers a mixture of revenue-based and profit-based incentives, owners can commit to behavior that is not profit maximizing, leading to higher prices. Bannano and Vickers (1988) show that manufacturers can soften price competition by selling through an independent retailer, rather than one that is vertically integrated.

¹³A related 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’ computer program and chooses an action accordingly. Another related concept is the cartel punishment device of Osborne (1976).

2 Algorithms and Pricing Behavior: Evidence

2.1 What is an Algorithm?

Broadly speaking, an algorithm is a set of instructions that maps inputs to a desired set of output. Pricing algorithms used by online retailers can be characterized as a formula to determine prices that is pre-specified by a computer program. Many online retailers consider rivals prices' to be a key input in these calculations. In general, an algorithm may depend on variables related to past, present, and future supply and demand conditions, including the past play of rivals or the outcomes of experiments. By using automated programs to collect this information and compute prices, firms can update prices at a higher frequency and place rules on pricing behavior. We wish to investigate two key features of pricing algorithms that may change the nature of the pricing game relative to a human agent.

First, an algorithm lowers the cost of updating prices and facilitates a regular pricing frequency. Typically, firms use software to schedule pricing updates at regular intervals, e.g., once per day or every 15 minutes. The frequency with which a firm can update prices depends on investments in pricing technology, which may differ across firms. Algorithms facilitate both regular pricing updates and more frequent updates, as software can better monitor rivals' prices and can find the solution to a difficult pricing problem more efficiently than a human agent. For numerical calculations, human agents can be slow and error-prone, and they cannot be expected to maintain a regular pricing frequency.¹⁴ Large online retailers sell several thousand products; relying on humans to update all prices at regular intervals would be extremely costly.

Second, an algorithm provides a short-run commitment device to a pricing strategy. When an algorithm depends on rivals' prices, it can autonomously react to price changes by rivals according to the formula encoded by the computer program. The program itself is typically updated at a lower frequency than it is used to set prices. Thus, in between updates to its algorithm, the firm changes prices based on a fixed set of rules. It is widely thought that humans lack this sort of commitment power (e.g., Maskin and Tirole, 1988a). In other words, we typically expect human agents to be bound by an incentive compatibility constraint at every opportunity to set prices.

Below, we present new empirical facts about pricing technology that demonstrate the importance of these two features of algorithms. We show that firms differ in the frequency with which they change prices and that faster firms react to rivals' price changes. We also find that faster firms have lower prices than slower firms. In the following sections, we provide an economic framework to capture these features and examine the effects on equilibrium prices. We introduce the features of frequency and commitment sequentially to build intuition. In particular, we introduce a model of (asymmetric) pricing frequency in Section 3 and a more general model that allows for short-run commitment in Section 4.

¹⁴The study by Ellison et al. (2018) provides empirical evidence of human inefficiency along these dimensions.

Table 1: Daily Statistics for Hourly Price Data

Statistic	Retailer A	Retailer B	Retailer C	Retailer D	Retailer E	All Retailers
Count of Products	124.9	41.3	49.9	42.5	35.1	58.7
Observations per Product	20.9	20.4	19.0	21.1	19.1	20.1
Price: Mean	27.18	16.88	17.63	20.93	21.74	20.86
Price: 10th Percentile of Products	9.75	6.93	5.53	6.88	7.50	7.32
Price: 90th Percentile of Products	51.11	28.95	33.30	38.21	39.65	38.21
Mean Absolute Price Change	1.35	2.31	1.12	3.28	3.06	1.91
Price Changes per Product	1.89	0.28	0.01	0.02	0.03	0.45
Share of Products with a Price Change	0.373	0.089	0.008	0.020	0.024	0.103

Notes: Table displays the daily averages (means) for each statistic for each website.

2.2 Data

For our empirical analysis, we collect a dataset of hourly prices for over-the-counter allergy drugs from five online retailers in the United States. The 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.¹⁵ We have kept the identities of the retailers anonymous, calling them *A*, *B*, *C*, *D*, and *E*. 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.

It is important to note that the retailers do not simply set uniform prices across both online and brick-and-mortar channels. For example, Cavallo (2017) finds that online prices at drug-stores differ 62 percent of the time from observed offline prices, and they are on average 1 percent lower. While prices may differ across a retailer’s brick-and-mortar stores, prices on the websites are set uniformly for online shoppers across the country.¹⁶

In order to collect high-frequency price data, we focus on the seven brands of allergy drugs that are sold by all five retailers: Allegra, Benadryl, Claritin, Flonase, Nasacort, Xyzal, and Zyrtec.¹⁷ We collect price information for all versions of the allergy drugs and define a product to be a drug-brand-form-(variant-)size combination, e.g., Loratadine-Claritin-Tablet-20. Using this definition, the average retailer sells 59 distinct allergy products on average. This set of products provides a relatively straightforward set of competing products in which to examining pricing technology in detail, however we believe our analysis of firms’ pricing technology applies more broadly to other products sold by the retailers.

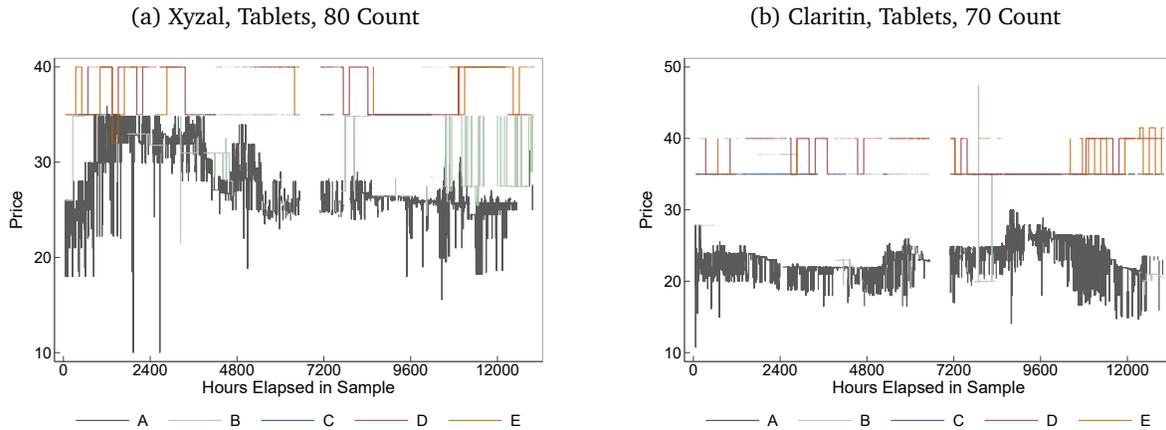
Our sample spans approximately one and a half years, from April 10, 2018 through October

¹⁵E-commerce revenue is obtained from eCommerceDB. Overall, these five retailers accounted for \$6 billion in e-commerce revenues for personal care, which includes medicine, cosmetics, and personal care products.

¹⁶Cavallo (2017) presents evidence of geographic variation across brick-and-mortar stores within the same retail chain. Any geographic variation in prices across stores would guarantee that some customers will face different prices online and offline.

¹⁷Our sample consists of products sold directly by retailers and not products in which a third-party seller sets the price. Third-party sellers are less popular for allergy products.

Figure 1: Example Time Series of Prices for Identical Products Across Retailers



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 70-count package of Claritin tablets.

1, 2019. Collecting high-frequency price data can be challenging. Websites change over time, there can be errors loading pages, and there are often other technical issues. During our sample period, we have relatively good coverage and observe the price for each product in 20 out of 24 hours on average. We take some steps to impute missing prices and identify outliers, which we describe in Appendix A. Our final dataset has 3,606,956 price observations across the five websites. Appendix Table 7 provides a tabulation of price observations for each retailer and brand.

Daily summary statistics of our data are presented in Table 1. On average, we observe 59 products each day on each website, though retailer A carries more products than the other four retailers. While retailer E only sells 35 products in the category on average, retailer A sells 125. Prices vary across retailers, though it is important to note that the raw average in the table reflects differences in available products. All of the retailers make large price adjustments over the sample period, with an average absolute price change of \$1.91. However, some retailers change prices more often than others. On an average day, retailer A changes the prices of 37 percent of its products while retailer C only changes the price of 0.8 percent of its products. Retailers D and E change the price of 2 percent of products each day. These stark differences indicate that there may be differences in the underlying pricing technology.

2.3 Three Facts About Online Prices

We now use a descriptive analysis of our dataset to document three stylized facts about pricing behavior in online markets.

Stylized Fact 1: Online retailers update prices at regular intervals. These intervals differ widely across firms.

To understand the pricing technology used by online retailers, we start by examining the time series for individual products. Figure 1 shows prices for Xyzal-Tablet-80 and Claritin-Tablet-70. These two examples illustrate fundamentally different pricing patterns across the five retailers. Retailer *A* often has high frequency price changes of a large magnitude. Retailer *B* also has high-frequency price changes, although less often. Retailer *C* appears to adjust prices at lower frequency while *D* and *E* tend to have prices that remain constant for long periods.

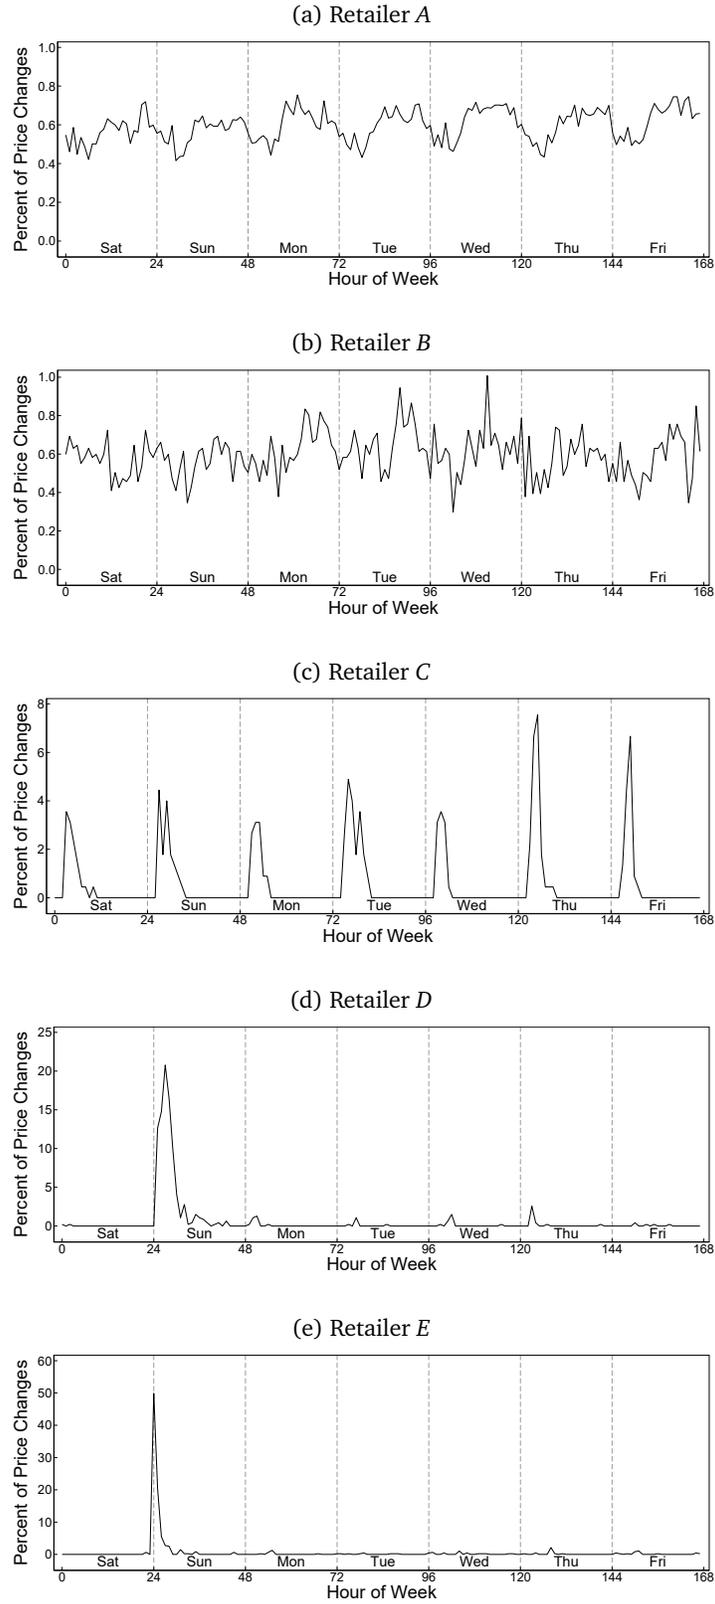
The differences in frequency are systematic across all products offered by the retailers. To capture variation in each firm’s underlying pricing technology, we plot the density of price changes across all products by hour of the week in Figure 2. The results show important differences in when firms are able to update prices. Retailers *A* and *B* have price changes that are relatively uniformly distributed across all hours of the week. In fact, anecdotal evidence suggests that these retailers are able to adjust prices multiple times within an hour, with Retailer *A* able to adjust prices at the highest frequency. The other retailers show regular patterns of price changes that are consistent with each firm running a pricing update script at pre-specified intervals. Retailer *C* adjusts prices daily between 3:00 AM to 6:00 AM EDT, whereas retailers *D* and *E* adjust prices weekly just after midnight EDT on Sunday.¹⁸ Thus, the figure documents stark differences in pricing frequencies among competing retailers, including weekly, daily, and near “real-time” pricing technology.

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 price changes occur as compelling evidence of technological constraints. Firms face several costs to upgrade their pricing technology, including new systems to gather and process higher-frequency input data, software to solve for the optimal higher-frequency prices, and new hardware that enables the algorithms to run at a higher frequency. It is important to note that pricing technology is not exclusively defined by software and hardware. Technology may also include managerial or operational constraints that prevent a firm from updating a price on a more frequent basis. For example, higher-frequency prices changes may be inconsistent with a retailer’s marketing strategy or make inventory management more challenging. Even if slower firms had access to the same hardware and software as retailers *A* or *B*, it would likely take significant organizational changes to enable the firms to update their prices as frequently.

The pricing patterns imply that, for the majority of hours in the week, only a subset of firms have pricing technology that allows for a price change. Only for a brief period once a week, on Sundays, do all firms simultaneously set prices. In other words, heterogeneous pricing technology is inconsistent with the simultaneous move assumption in standard models of competition.

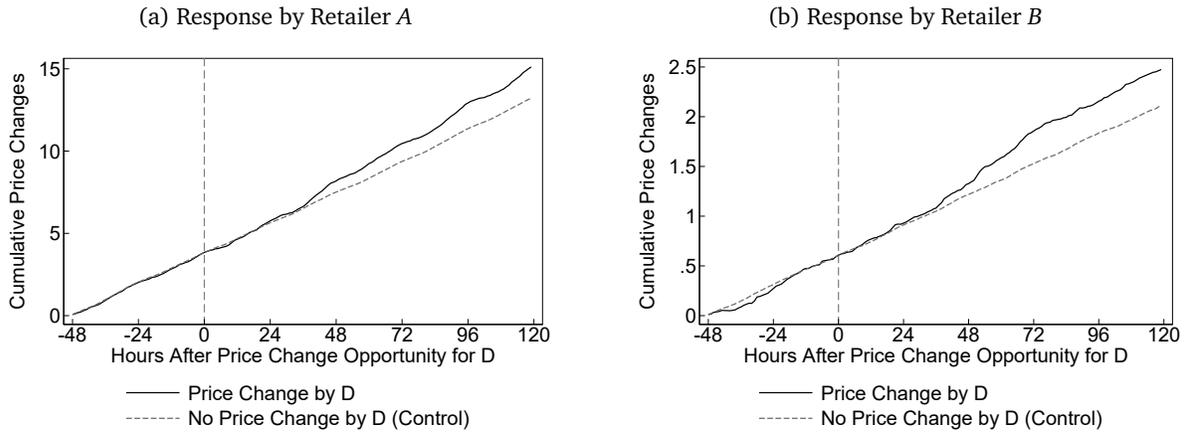
¹⁸Many of the price changes that occur outside of these times are likely due to measurement error.

Figure 2: Heterogeneity in Pricing Technology by Hour of the Week



Notes: Displays percent of each retailer's price changes in each hour of the week. Hours are reported in Eastern Time (UTC-5).

Figure 3: Price Changes by Fastest Retailers in Response to Price Change by Retailer *D*



Notes: Figure displays the cumulative price changes for high-frequency retailers *A* and *B* in response to a price change occurring at retailer *D*, which adjusts prices only once per week. The solid line displays the cumulative price change when retailer *D* changes a price of the same product in that week. The dashed line plots the cumulative price changes when the product at retailer *D* does not have a price change. The solid line is adjusted by the pre-period difference in rates so that the lines coincide at period -1.

Stylized Fact 2: Retailers with the fastest pricing technology quickly react to price changes of slower rivals, consistent with the use of automated pricing algorithms.

If algorithms depend on rivals' prices, then we should expect high-frequency firms to quickly react to price changes by low-frequency firms. High-frequency firms may change prices for many reasons, including cost shocks, demand shocks, and experimentation. In order to isolate the response to rivals' prices, we analyze the timing of price changes by high-frequency firms in weeks with and without a price change by a slower rival. A slow firm may be spurred to change prices due to an idiosyncratic cost shock arising from, e.g., shipping delays or low inventory. If the faster firm's algorithm is a function of the slower firm's prices, we may observe additional price changes by the faster firm after the slower firm changes its price.¹⁹

To examine the reaction of prices to other firms, we start by taking price changes occurring at retailer *D*, one of the firms with weekly pricing technology, as the impulse. We observe 348 price changes in our data occurring between midnight and 6 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. While retailer *D* runs their price update script once per week, not all prices are updated each week. We capture "treated" product-weeks in which the product changed its price at retailer *D* and "control" weeks in which the product did not change its price, despite the fact that retailer *D* had the opportunity to adjust prices.

¹⁹If both firms are responding to common shocks (to demand or supply), we would typically expect the price changes at the faster firm to happen *before* those of a slower rival.

Table 2: Effect of Price Change by Slower Retailers on Price Changes by Faster Rivals

	Price Change by D			Price Change by E		
	(1) Retailer A	(2) Retailer B	(3) Retailer C	(4) Retailer A	(5) Retailer B	(6) Retailer C
$Post_{h(t)} \times PriceChange_{w(t)}$	0.770*** (0.207)	0.319*** (0.109)	-0.005 (0.017)	0.667*** (0.189)	0.291** (0.127)	-0.013 (0.017)
Product \times Week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Outcome Mean	5.709	0.927	0.027	5.709	0.927	0.027
Observations	1,115,035	353,873	426,905	1,115,035	353,873	426,905

Notes: Results from OLS regressions in which the outcome is an indicator for whether the faster retailer changed its price. We include 48 hours before and 72 hours after each opportunity for a price change by the slow retailers, which occur Sunday at midnight. Therefore, the sample includes Friday through Wednesday of each week. The outcome is scaled by 72 so the rate change can be interpreted as cumulative changes over the three-day post period. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3 plots the cumulative price changes before and after midnight on Sunday across each product-week. The solid line corresponds to treated product-weeks, i.e., weeks in which the price of a particular product changed at retailer *D*. The dashed line corresponds to control product-weeks that had no price change. The solid line is adjusted by the pre-period difference in rates so that the lines coincide at period -1. The gap between the solid line and the dashed line is the marginal increase in price changes when a price change occurs at retailer *D*.

Based on Figure 3, it is clear retailers *A* and *B* have an increased probability of a price change after a price change at retailer *D*. The fast retailers respond to a price change by retailer *D* within about 48 to 72 hours. The fact that we do not observe a differential increase before the price changes by the slow retailer is evidence that the faster firms are responding to retailer *D* rather than to a common shock. By the end of the week, the fast retailers realize roughly 20 percent more price changes over the baseline. In Appendix Figure 13 we examine the results for retailer *E*, the other retailer with weekly pricing technology, and find very similar results.

In order to further quantify the effect on faster retailers of a price change by the two slowest retailers, we use a difference-in-difference specification given by

$$y_{it} = \beta(Post_{h(t)} \times PriceChange_{w(t)}) + \gamma_{i,w(t)} + \gamma_{h(t)} + \varepsilon_{it} \quad (1)$$

where y_{it} is an indicator for whether the faster retailer changed its price for product i in hour t . We use a 48-hour period before and a 72-hour period after the slow firm adjusts prices, and we scale the dependent variable by 72 so that the rate change can be interpreted as cumulative changes over the three-day post period. $Post_{h(t)}$ is an indicator for whether the hour of the week, $h(t)$, is after an opportunity for the slow firm to adjust price. $PriceChange_{w(t)}$ is an

indicator for whether the slow firm adjusted prices in week $w(t)$.²⁰ We include product-week fixed effects, $\gamma_{i,w(t)}$, to control for product-specific time-varying factors that are common across retailers, such as a demand shock that causes both retailers to adjust prices, with the faster firm able to respond first. Finally, we include hour-of-week fixed effects, $\gamma_{h(t)}$, to account for time-varying factors within the week. In this way, β can be interpreted as the effect of the slow retailer’s price change on cumulative price changes by the faster retailer. Identification exploits two sources of variation: variation across weeks in which the slow firm does or does not adjust the prices for a product and variation within each week before and after the opportunity for the slow firm to adjust prices.

Table 2 reports regression results analyzing the response of the faster retailers, A , B , and C , to the slower retailers, D and E . Results indicate that when retailer D changes the price of a product, retailer A has 0.8 additional price changes for the same product within 72 hours. Retailer B has 0.3 additional price changes. Relative to the average number of price changes over the same period—5.7 for retailer A and 0.9 for retailer B —the coefficients correspond to a 13 percent and a 34 percent increase in the rate of price changes, respectively. Results estimating the effect of a price change by retailer E are similar, and all the estimated responses by A and B are statistically significant. We observe many fewer price changes by retailer C , and therefore the estimated effects for retailer C are imprecise.

These results imply that the two retailers with the most frequent pricing technology, A and B , are responding to price changes of lower frequency rivals within a relatively short period. Given the large number of prices that these firms update and the speed at which prices are updated, the results are consistent with firms using automated pricing algorithms that are a function of rivals’ prices. As we discuss in the introduction, anecdotal and survey evidence indicates that this is a common practice for online retailers. To the extent that these algorithms are updated at lower frequency than prices are adjusted, this implies short-run commitment to a pricing strategy. Slower firms may anticipate the pricing algorithm used by faster firms, which is inconsistent with the standard Bertrand-Nash assumption.

Stylized Fact 3: Firms with faster pricing technology have persistently lower prices for identical products.

We now examine the relationship between pricing frequency and prices for identical products across different retailers. By using a high-frequency pricing algorithm, firms may commit to best-respond to their rivals. As we formalize later, this best response is often to undercut rivals’ prices, implying that high-frequency firms set lower prices than slower rivals.

In order to account for differences in product assortment across retailers and over time, we regress log prices on indicators for each retailer while controlling for product and hour-day

²⁰Note that $w(t)$ and $h(t)$ map the hour t to week and hour of the week respectively.

Table 3: Price Differences for Identical Products Relative to Retailer A

	(1)	(2)	(3)	(4)
Retailer B	0.064*** (0.000)	0.047*** (0.001)	0.146*** (0.000)	0.117*** (0.001)
Retailer C	0.092*** (0.000)	0.107*** (0.001)	0.171*** (0.000)	0.187*** (0.001)
Retailer D	0.249*** (0.000)	0.289*** (0.001)	0.307*** (0.000)	0.337*** (0.001)
Retailer E	0.284*** (0.000)	0.366*** (0.001)	0.340*** (0.000)	0.419*** (0.001)
Product FEs	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes
Sold at All Retailers			Yes	Yes
On or After Jul 1 2019		Yes		Yes
Observations	3,606,956	677,650	1,186,571	234,696

Notes: Results from OLS regressions in which outcome is log price. Coefficients show price difference relative to retailer A. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

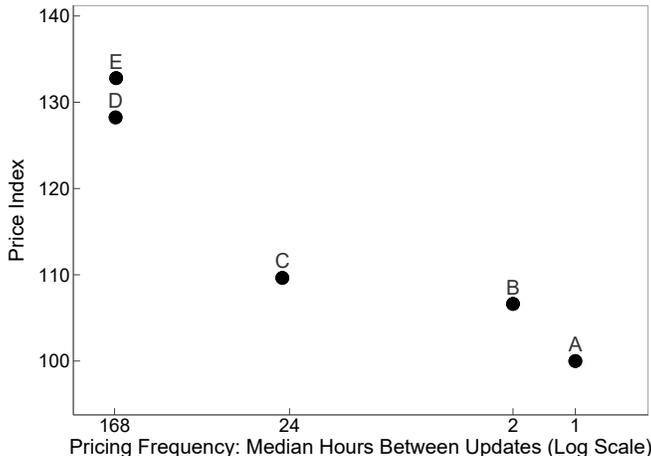
fixed effects. The resulting coefficients reflect the average difference in (log) price for identical products (brand-drug-form-variant-size) sold across different retailers at the same point in time.

Table 3 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 6.6 percent (0.064 log point) premium at B and a 9.6 percent (0.092 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. Specifications (2) and (4) use observations from the most recent three months of the data (July 1, 2019 through October 1, 2019), the period with the most stable panel. Specifications (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 (1) against a measure of pricing technology in Figure 4. The x-axis captures the pricing frequency, which increases along the x-axis. We report the frequency as the median number of hours between any pricing update on each website; the axis values are reversed so that superior (more frequent) technology is to the right. Firm E has a median approximately equal to the number of hours in a week (168), whereas firm A has a median of 1.

The large degree of price dispersion in online markets has largely been attributed to search frictions. Yet, the robust correlation between pricing technology and average prices suggests that pricing technology may play a role. High-frequency pricing algorithms may allow firms to commit to undercutting slower rivals, softening competition and implying retailers with high-frequency pricing have lower prices in equilibrium. It is important to note that there are other

Figure 4: Price Index for Identical Products by Retailer Pricing Frequency



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 (1) of Table 3.

reasons why prices could be higher for firms with low-frequency pricing, such as differences in marginal costs or asymmetric demand. We discuss these issues in our empirical exercise in Section 5.

3 Competition with Pricing Frequency

Motivated by the fact that retailers update prices at different intervals, we begin by modeling pricing frequency. We show that enabling firms to choose different pricing frequencies has important implications, and it provides some intuition for a richer model where firms can also commit to a pricing strategy in the short run. We present this more general model in Section 4.

3.1 Infinite Horizon Model

Consider two firms with the ability to change prices at different frequencies. Both firms initially set prices at $t = 0$. Firm 1 can update its price at discrete points after each interval of time T_1 , and firm 2 can likewise update its price after intervals of length T_2 . We assume that $T_1 = \theta T_2$, where $\theta \in \mathbb{N}$. This implies that firm 2 has (weakly) superior technology, allowing it to change its price at least as frequently as firm 1. For example, T_1 may equal one week, while T_2 equals one day ($\theta = 7$). Without loss of generality, we normalize $T_1 = 1$, i.e., we define units of time in terms of the period between firm 1’s potential price changes.

In the next section, we formalize the link of this model to a more general model of competition in algorithms. The implicit assumption we make in this section is that firms can revise their

algorithms whenever they have the ability to update prices, i.e., they completely re-solve for the optimal price. In other words, firms cannot commit to a fixed pricing rule in intermediate periods. Pricing frequency therefore corresponds to the frequency that firms can update their algorithms. We assume that firms are fully sophisticated when it comes to monitoring current prices and understanding rivals' algorithms.²¹ Under the assumption of no commitment, it suffices to analyze the pricing game. We focus on the two-firm case, but our results readily extend to multiple firms.

Demand arrives in continuous time, with a measure $m(t) \geq 0$ of consumers arriving at t . The distribution of consumers is stable over time, so that demand looks identical at any instant t except for the size of the market. Given demand and prices (p_1, p_2) , firm j realizes instantaneous profit flow $\pi_j(p_1, p_2)$. We assume the profit functions are quasiconcave and have a unique maximum with respect to a firm's own price. Firms discount the future exponentially at rate ρ and have an infinite horizon.

Firms choose a sequence of prices to maximize profits, conditional on the flow of consumers $m(t)$, the profit flows π_j , and the behavior of the rival firms. Let $p_1(t)$ and $p_2(t)$ denote the prices of each firm over time, and let P_1 be the discrete sequence of prices chosen by firm 1 at $t = \{0, 1, 2, \dots\}$. For timing purposes, we assume that P_{1s} is relevant for demand over the period $(s, s + 1]$. Firm 1's problem can be written as:

$$\max_{P_1} \sum_{s=0}^{\infty} \int_s^{s+1} e^{-\rho t} \pi_1(P_{1s}, p_2(t)) m(t) dt. \quad (2)$$

Because firm 2 can change its price at every point $s \in \{0, 1, \dots, \infty\}$ in addition to intermediate times, the problem can be expressed as a sequence of single-period stage games. We restrict our attention to subgame perfect equilibrium in each stage game. The resulting equilibrium is the unique (pure-strategy) Markov perfect equilibrium of the infinite horizon problem.

3.2 Stage Game Analysis

As we have shown, the repeated game can be expressed as a sequence of single-period stage games. Firm 1's problem in stage game s is

$$\max_{p_1} \int_s^{s+1} e^{-\rho t} \pi_1(p_1, p_2(t)) m(t) dt. \quad (3)$$

We now analyze the behavior of firm 2 in each period. Firm 2's pricing behavior will satisfy

²¹In practice, firms may use machine learning and experimentation to learn about the pricing algorithms of their rivals. Our environment can be considered the limiting case of an arbitrary (but consistent) learning process. We do limit sophistication in strategies by focusing on Markov perfect equilibria where firms cannot condition on past prices. Our analysis can be contrasted with the literature on algorithmic collusion in which firms employ history-dependent strategies, allowing them to sustain collusion.

the following two properties in equilibrium: (1) firm 2's price will be constant over the period (despite its ability to update prices), and (2) firm 2's price will lie along its Bertrand best-response function. The first property is a result of $\pi_2(\cdot)$ being time-invariant and p_1 being fixed in the period. The second property arises from the fact that it is optimal for firm 2 to price along the Bertrand best-response function when it is pricing simultaneously with its rival ($t = s$) and also in any later pricing update (e.g., $t = s + 1/\theta$). The Bertrand best-response function for firm 2 treats p_1 as fixed, which is a Nash equilibrium condition at $t = s$ and is literally true at any other point when firm 2 can update its price. Let $R_2(p_1, s)$ denote firm 2's reaction function in period s .

We return to firm 1's problem. Without loss of generality, we focus on the first period ($s = 0$). Let p_2 now denote the price of firm 2, which is time-invariant (in the stage game) in equilibrium, and let $R_2(p_1) = R_2(p_1, 0)$. Firm 1 chooses p_1 recognizing that p_2 can react to its price after a period of $1/\theta$. Firm 1's problem can be expressed as:

$$\max_{p_1} \int_0^{\frac{1}{\theta}} e^{-\rho t} \pi_1(p_1, p_2) m(t) dt + \int_{\frac{1}{\theta}}^1 e^{-\rho t} \pi_1(p_1, R_2(p_1)) m(t) dt. \quad (4)$$

Because the profit flow function is time-invariant, we can write firm 1's stage game problem as:

$$\max_{p_1} \underbrace{(1 - \alpha) \pi_1(p_1, p_2)}_{\text{Simultaneous Pricing Incentive}} + \underbrace{\alpha \pi_1(p_1, R_2(p_1))}_{\text{Sequential Pricing Incentive}} \quad (5)$$

where $\alpha = \left(\int_0^1 e^{-\rho t} m(t) dt \right)^{-1} \int_{\frac{1}{\theta}}^1 e^{-\rho t} m(t) dt$. The value $1 - \alpha$ describes the relative weight that firm 1 places on the initial period $(0, 1/\theta]$, which is a function of ρ , $m(t)$, and θ .²² In the initial price-setting phase, the usual Nash-in-price logic holds: firm 1 treats firm 2's price as given over the period $(0, 1/\theta]$. After $t = 1/\theta$, firm 1 recognizes that firm 2 will price optimally against its chosen price when it has the opportunity to update. Therefore, the sequential pricing logic holds in this second phase.

There are two special cases of this pricing model that we now highlight. When $\alpha = 0$, firm 1 considers only the current price of firm 2. Roughly speaking, firm 1 places zero weight on the ability of firm 2 to react to a price change by firm 1. This can arise when $\theta = 1$, i.e., when firms have symmetric technology and set prices simultaneously. Thus, our model nests the usual Bertrand-Nash equilibrium assumption that firm set prices while holding fixed the prices of rivals.

The second special case is when $\alpha = 1$. In this case, firm 1 only considers its profits after firm 2 has a chance to update its price. Roughly speaking, firm 1 fully internalizes the reaction of its rival. This can arise when $\theta \rightarrow \infty$, i.e., when firm 2 has much faster pricing technology

²²When the stage game interval is small, it is reasonable to assume that demand arrives uniformly and that $\rho = 0$, in which case we have the simple expression $\alpha = \frac{\theta-1}{\theta}$.

than firm 1. The result is equivalent to a sequential pricing model, where first firm 1 chooses a price and then is followed by firm 2. In this way, our model provides a foundation for the sequential pricing game—i.e., the Stackelberg pricing model—analyzed in the theory literature but rarely in applied work.

Depending on the underlying parameters, the model can capture both simultaneous and sequential price-setting behavior. More generally, the asymmetric technology allowed for in our model provides a foundation for a rich set of equilibrium outcomes that capture a mix of the incentives in these games. We now provide our first proposition, which describes the set of equilibrium outcomes for any value of α :

Proposition 1. *In the pricing frequency game, the equilibrium prices will lie on the faster firm’s Bertrand best-response function between the Bertrand equilibrium and the sequential pricing equilibrium.*

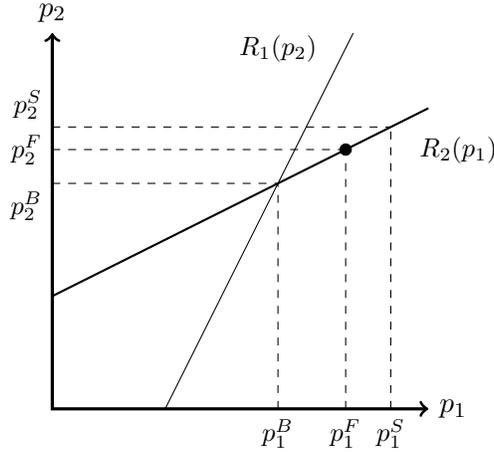
Proof: We have established that firm 2’s price will lie along its Bertrand best-response function, as it always treats firm 1’s price as given. When $\alpha = 0$, the problem is equivalent to a simultaneous Bertrand pricing game. Note that this is obtained when $\theta = 1$, in which case the game corresponds exactly to simultaneous price setting. Denote the optimal price in this game p_1^B . When $\alpha = 1$, the game is equivalent to a sequential price-setting game, where firm 1 is the leader and firm 2 is the follower, with optimal price p_1^S . Because the profit function is quasiconcave, the price that maximizes the weighted sum of $\pi_1(p_1, p_2)$ and $\pi_1(p_1, R_2(p_1))$ lies in between p_1^B and p_1^S . QED.

Figure 5 illustrates the equilibrium of the game. When firms are very impatient or most consumers arrive before firm 2 can update its price, the equilibrium will resemble Bertrand (p^B). When firms are patient and all consumers arrive after firm 2 can update its price, the equilibrium resembles sequential price setting (p^S). The equilibrium prices p^F can fall anywhere between these points, depending $m(t)$, θ , ρ , and the profit functions. Note that p^F is not necessarily a linear combination of p^B and p^S ; it is in the figure because the best-response function is linear.

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, $\frac{\partial R_2}{\partial p_1} > 0$). Under these conditions, the sequential price-setting equilibrium will have higher prices than the Bertrand equilibrium. Thus, we obtain our second proposition:

Proposition 2. *Suppose firms produce substitute goods and prices are strategic complements. In the pricing frequency game, both firms realize higher prices compared to the simultaneous price-setting (Bertrand-Nash) equilibrium.*

Figure 5: Equilibrium in the Pricing Frequency Game



Notes: Figure plots the best-response functions $R_1(\cdot)$ and $R_2(\cdot)$ for simultaneous price competition with differentiated products. The intersection of these functions produces the Bertrand-Nash equilibrium (p_1^B, p_2^B) . The point (p_1^S, p_2^S) indicates the equilibrium of the sequential pricing game. The point (p_1^F, p_2^F) is the equilibrium of a pricing frequency game, which lies between (p_1^B, p_2^B) and (p_1^S, p_2^S) .

Proof: Above, we have demonstrated that firm 1's price lies between the Bertrand price p_1^B and the sequential equilibrium price p_1^S . It suffices to show that $p_1^B < p_1^S$, in which case the optimal price lies on $[p_1^B, p_1^S]$.

Consider firm 1's first-order condition to maximize profits (π):

$$\frac{d\pi_1}{dp_1} = \frac{\partial\pi_1}{\partial p_1} + \frac{\partial\pi_1}{\partial p_2} \frac{\partial p_2}{\partial p_1} = 0 \quad (6)$$

In the simultaneous price-setting equilibrium, firm 1 takes firm 2's price as given ($\frac{\partial p_2}{\partial p_1} = 0$), and $\frac{\partial\pi_1}{\partial p_1} = 0$. In the sequential game, 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 in the sequential game: $\frac{d\pi_1}{dp_1} > 0$. Firm 1's optimal price will be strictly greater than p_1^B when $\alpha > 0$ and the profit function is well-behaved. Higher prices for both firms result from strategic complementarity. QED.

In typical models of differentiated products, prices are strategic complements (Tirole, 1988). If prices are instead strategic substitutes, then the equilibrium will have one firm with higher prices and one firm with lower prices, and the net effect on prices may be ambiguous.

3.3 Pricing Frequency Game: Example

We have described above conditions under which a dynamic game of price competition with asymmetric pricing frequency can be broken down into single-period stage games. We now

provide an example to help fix ideas. In this game, firms compete for demand over a single period. Each firm produces a single product and set prices to maximize profits. Firms initially set prices at the beginning of the period, and, depending on the technology, can update prices throughout the period.

We assume that demand is such that products are (imperfect) substitutes and prices are strategic complements. In particular, we use a variant of the Hotelling (1929) model, with fixed locations and an outside option.²³ Where the utility from both goods is positive, the (local) demand for each good has the convenient linear form:

$$q_j(t) = \frac{1}{2}m(t)(1 - p_j + p_{-j}).$$

We assume $\int_0^1 m(t)dt = 2$. Because equilibrium prices are invariant throughout the period, we can integrate over t to obtain $q_j = 1 - p_j + p_{-j}$ for each firm.

As above, firm 1 sets its price at the beginning of each period, whereas firm 2 can update its at a frequency of $\theta \in \mathbb{N}$, corresponding to elapsed intervals of $T_2 = 1/\theta$. Firm 2's price will lie along its best-response function. Firm 1 will internalize the reaction by firm 2, choosing its price to maximize the profit function given by equation (5). In this example, the equilibrium prices are given by

$$\begin{aligned} p_1 &= \frac{3}{3 - \alpha} \\ p_2 &= \frac{6 - \alpha}{6 - 2\alpha}, \end{aligned} \tag{7}$$

where $\alpha = \left(\int_0^1 e^{-\rho t} m(t) dt\right)^{-1} \int_{\frac{1}{\theta}}^1 e^{-\rho t} m(t) dt$. in general, prices depend on the relative level of technology of firm 2 (θ), as well as the discount rate ρ and the arrival rate of consumers $m(t)$.²⁴ Note that, even with linear demand, equilibrium prices may have a nonlinear relationship with α or θ .

To illustrate the impact of pricing technology in this example, we consider three cases. First, consider the standard case where firms have symmetric technology, i.e., $\theta = 1$. This corresponds conceptually to a game in which firms use human agents to set prices. In this case, $\alpha = 0$, and thus equilibrium prices, $p_1 = p_2 = 1$, and profits, $\pi_1 = \pi_2 = 1$, are equivalent to the simultaneous Bertrand-Nash equilibrium.

Now consider the case in which firm 2 adopts new pricing technology and is able to adjust prices at a higher frequency than firm 1. This implies that $\theta > 1$ and $\alpha > 0$. From equation

²³Each consumer i receives utility v from consuming the good and has disutility of τd_{ij} for the distance d_{ij} they travel to purchase from firm j . We set $v = 2$ and $\tau = 1$. Utility is linear in income and is normalized so that the marginal utility of income is 1. Consumer locations are uniformly distributed and the value of not purchasing is normalized to have zero utility.

²⁴When demand arrives uniformly throughout the period and $\rho = 0$, we can represent equilibrium prices as function of the faster firms technology, θ : $p_1 = \frac{3\theta}{1+2\theta}$ and $p_2 = \frac{1+5\theta}{2+4\theta}$.

(7), we can see that firm 1 and firm 2 increase their prices, but firm 2 chooses a lower price than firm 1. 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. For example, when $\alpha = \frac{1}{2}$ (which may correspond to $\theta = 2$), firm 1 chooses a price of 1.2 and firm 2 chooses a price of 1.1. Firm 1 loses market share to firm 2, as equilibrium quantities are (0.9, 1.1), but profits are (1.08, 1.21), which are higher for both firms than in the Bertrand equilibrium.

Finally, consider the case in which firm 2’s technology is much more advanced, allowing them to update prices “in real time.” In our model, this corresponds to $\theta \rightarrow \infty$ and $\alpha = 1$. Firm 1 now fully internalizes the reaction of firm 2 and chooses a price of 1.5. This leads firm 2 to price at 1.25. Quantities are (0.75, 1.25), and profits are (1.125, 1.5625), resulting in an equivalent outcome to the sequential (Stackelberg) pricing game.

The Bertrand-Nash logic uses a dynamic metaphor to rule out the above outcomes: if firm 2’s price is fixed at either 1.1 or 1.25, 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 one-shot 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. In this way, the model provides a foundation for commitment; such commitment is necessary to generate higher prices than the Bertrand game.

3.4 Endogenous Pricing Technology

We have characterized a pricing game in which firms may differ in their pricing technologies. 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. For this reason, firm 1 has a disincentive to upgrade its technology to match that of firm 2.

Thus, when firms can choose the pricing frequency in this model, asymmetric frequencies are the equilibrium outcome. We formalize this result by modeling a first-stage adoption decision in Appendix B, but the result is quite intuitive. Whenever firms choose the same technology, Bertrand prices result. Each firm has a unilateral incentive to move away from symmetric technology, and they would do so if the cost to change technology were not prohibitively high. A firm may adopt costly technology even if its rival gains more from the outcome, as the firm prefers this outcome to the world in which neither firm adopts. Conversely, a firm may even pay to *downgrade* its technology to avoid the Bertrand outcome. In other words, firms may be willing to disadvantage themselves relative to their rivals to gain the benefits of softened price competition. For these reasons, we might not expect simultaneous price-setting behavior to

hold in equilibrium.²⁵

We have shown in Section 2 that, consistent with the incentives described above, asymmetric pricing technology is a key feature of major online retailers. In other settings, factors outside of the model may allow firms to maintain symmetric frequencies in equilibrium, such as the benefits of adapting to time-varying demand conditions (so-called “dynamic pricing”) or severe technological constraints.

4 Algorithms with Commitment

The previous section discussed outcomes in which firms have asymmetries in pricing frequency. The frequency model corresponds to a game where the algorithms employed by firms are continually revised so that they are optimal at every moment. In practice, this maps to an environment where firms are able update their algorithms whenever there is an opportunity to update prices, so that the encoded algorithm is not fixed and does not provide commitment.

Here, we provide a generalization of this game where we allow firms to choose algorithms that determine future prices. Firms may update these algorithms at different frequencies. When prices are updated at higher frequencies than algorithms, an algorithm serves as a short-run commitment device. Roughly speaking, the algorithm enables commitment to a pricing rule that may not be optimal in the short run. From the same general model, we derive a “one-shot” game of competition in algorithms.

4.1 Setup

Two rival firms choose pricing algorithms, which they may update at different frequencies. Both firms can update their algorithms at $t = 0$. Each firm j can update its algorithm at regular intervals given by $1/\theta_j$, where $\theta_j \in \mathbb{N}$. The minimum update frequency $\theta_j = 1$ corresponds to an update at the beginning of each period. At the time a firm updates its algorithm, it may also change its price. Algorithms enable additional pricing updates at higher frequencies (γ_1, γ_2) . We assume that $\gamma_j = a_j\theta_j$, where $a_j \in \mathbb{N}$.²⁶ For the additional price changes, the firm delegates the pricing decision to a rule determined by the previous update to the algorithm. An illustration of timing in this game can be seen in Appendix Figure 14.

We assume that firms use pricing algorithms that are a function of the current price of rivals, i.e., the “payoff-relevant” price, although firms may respond with a lag. We will show that firms need not condition on past prices to sustain supracompetitive prices in equilibrium. Formally, an algorithm is a function $p_j = \sigma_j(\hat{p}_{-jt}, x_t)$, where \hat{p}_{-jt} is the most recently observed price of the rival firm. Non-price observables, such as cost shocks or the entire history of play, may be

²⁵Hamilton and Slutsky (1990) show similar incentives in a two-stage game where firms first choose whether to move first or second. They do not address how a firm may commit to only moving once.

²⁶This assumption provides expositional clarity. For other values of a_j , similar qualitative results may be obtained.

captured by a state vector, x_t . One can interpret our equilibrium analysis as conditional on any realization of the state, therefore, we suppress x_t in our notation and simply write algorithms as $\sigma_j(\hat{p}_{-j})$.

At $t = 0$, both firms have the ability to flexibly change their algorithm, σ_j . Each firm's strategy at $t = 0$ consists of $(p_{j0}, \sigma_{j0}(\cdot))$, where p_{j0} is the price determined while updating the algorithm and $\sigma_{j0}(\cdot)$ is the pricing rule at future opportunities. Firm 2 submits a new strategy $(p_{j1}, \sigma_{j1}(\cdot))$ when $t \in \{0, 1/\theta_2, 2/\theta_2, \dots\}$. The strategy space captures the fact that whenever a firm can make a revision to its algorithm, its rival does not take the commitment to that algorithm to be credible in that instant.

Firms choose a sequence of prices and algorithms to maximize profits, conditional on the flow of consumers $m(t)$, the profit flows π_j , and the behavior of the rival firms. Let $p_1(t)$ and $p_2(t)$ denote the prices of each firm over time, and let $S_1 = \{(p_{1t}, \sigma_{1t})\}$ be the sequence of strategies chosen by firm 1 at $t = \{0, 1/\theta_1, 2/\theta_1, \dots\}$. Demand adheres to the same conditions as the previous section.

When pricing updates correspond to algorithm updates ($\gamma_1 = \theta_1$ and $\gamma_2 = \theta_2$), we obtain the pricing frequency game of Section 3. In this game, there is no opportunity to rely on the pricing rule $\sigma_j(\cdot)$ to set prices.

In this paper, we focus on two additional special cases of the model. These special cases capture the key features of pricing technology that we observe in real-world environments and highlight the role of short-run commitment.²⁷

- **Asymmetric Commitment:** We can consider a game with asymmetric commitment, where only one firm has an algorithm that commits to automatic updates as a function of its rival's price ($\gamma_1 = \theta_1 = 1$ and $\gamma_2 > \theta_2$). This game closely corresponds to the pricing frequency model. We discuss this game and the connections to the frequency game in Section 4.2.
- **Symmetric Commitment:** We consider a case with symmetric short-run commitment, which allows us to highlight the role of commitment in algorithmic pricing. We turn our attention to this case in Section 4.3.

In each case, we restrict attention to Markov perfect equilibria. Because of the synchronous nature of the updates, it suffices to analyze subgame perfect equilibrium of a single-period stage game. Using these cases, we illustrate how the changes to frequency and commitment brought about by algorithms can lead to higher prices in competitive equilibrium.

²⁷The general setup admits many cases that cannot be neatly summarized by a single representation.

4.2 Asymmetric Competition in Pricing Algorithms

We first focus on the case in which firm 2 can commit to an algorithm that conditions on the price of firm 1, but firm 1 does not have this capability. We call this game the *asymmetric commitment game* to refer to the asymmetry in the nature of the algorithms. Though firm 1 does not automate its response to firm 2's prices, it may, in general, 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 given the differences in the ability of firms to monitor rivals and adjust prices documented in Section 2. The model differs from the frequency game of Section 3 by allowing the firm with superior technology to commit to a pricing function. It is a case of the general model with $\theta_1 = 1$, $\gamma_1 = 1$, $\theta_2 = 1$, and $\gamma_2 > 1$.

Conditional on firm 2's strategy $S_2 = (p_2, \sigma_2)$, firm 1's problem in the first period can be expressed as:

$$\max_{p_1} \int_0^{\frac{1}{\gamma_2}} e^{-\rho t} \pi_1(p_1, p_2) m(t) dt + \int_{\frac{1}{\gamma_2}}^1 e^{-\rho t} \pi_1(p_1, \sigma_2(p_1)) m(t) dt. \quad (8)$$

As before, we can write firm 1's stage game problem as a weighted average of the pre-update period $(0, 1/\gamma_2]$ and the post-update period $(1/\gamma_2, 1]$:

$$\max_{p_1} (1 - \alpha) \pi_1(p_1, p_2) + \alpha \pi_1(p_1, \sigma_2(p_1)) \quad (9)$$

where $\alpha = \left(\int_0^1 e^{-\rho t} m(t) dt \right)^{-1} \int_{\frac{1}{\gamma_2}}^1 e^{-\rho t} m(t) dt$. In the asymmetric commitment game, σ_2 depends on p_1 . The duration $\frac{1}{\gamma_2}$ represents the time lag between firm 1's pricing decision and the response of the algorithm by firm 2.

As in the asymmetric frequency case, the model provides an incentive for firm 1 to deviate from the competitive price. As long as $\partial \sigma_2(p_1) / \partial p_1 \neq 0$, then firm 1 will not set a price consistent with its Bertrand best-response function.

In this game, it is a (weakly) dominant strategy for σ_2 to mirror firm 2's best-response function. We use this result to highlight a special equilibrium where firm 2 submits its best-response function.

Proposition 3. *There exists an equilibrium to the asymmetric commitment game in which the second firm submits its best-response function as its algorithm. This strategy is weakly dominant. The first firm submits a price that maximizes its own profit along the second firm's best-response function.*

It is readily apparent that no profitable deviation exists. The firm that submits a price-dependent algorithm cannot do better than submitting its Bertrand best-response function as

its algorithm, regardless of the price chosen by firm 1. Thus, this is the unique equilibrium after eliminating weakly dominated strategies.²⁸ At this equilibrium, equation (9) is equivalent to (5). Thus, the asymmetric commitment game mirrors the asymmetric pricing frequency game from Section 3. In particular, the asymmetric commitment game obtains an identical equilibrium to the asymmetric frequency game when firm 2 chooses this weakly dominant strategy and γ_2 is equal to θ . Indeed, we present our second result for this section as a corollary to Proposition 2:

Corollary. *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.*

We have shown that asymmetries in pricing technologies are sufficient to generate higher prices than the in the simultaneous price-setting equilibrium. The results from this section highlight a somewhat surprising result: asymmetries arising from either frequency or commitment generate the same outcomes in equilibrium. Thus, understanding the exact nature of the pricing strategies may matter less than accounting for asymmetries. One can model a firm with a superior algorithm that conditions its rival’s price as simply having the ability to update prices more frequently.

As we show next, these similarities end when considering symmetric technology. Symmetric pricing frequency leads uniquely to Bertrand prices. By contrast, when both firms have algorithms with short-run commitment, firms are able to realize higher prices and profits than the Bertrand equilibrium, despite possessing symmetric technology.

4.3 Symmetric Competition in Pricing Algorithms

We now consider the case in which both firms have algorithms that can depend on the prices of rivals. Further, these algorithms update prices at a higher frequency than the frequency which firms can update their strategies, generating short-run commitment to the strategies. Without loss of generality, we consider the first period, $t \in (0, 1]$. Our objective is to characterize the equilibrium strategies that would be chosen by both firms.

Suppose that firm 1 and firm 2 can both update their algorithms with equal frequency, which we normalize to one ($\theta_1 = \theta_2 = 1$). Firms are also able to commit to an algorithmic pricing rule for future price updates, which occur simultaneously, with $\gamma_1 = \gamma_2 = \gamma$. Thus, initial price-setting behavior determines prices until $t = 1/\gamma$, after which the algorithms determine prices. For expositional clarity, we assume that there is no mass point in $m(t)$ at $t = 1/\gamma$ and that algorithms instantaneously converge to the “steady-state” prices, so the transition has no

²⁸There are many Nash equilibria where firm 2 has an algorithm that, local to the equilibrium, the algorithm maps to the best-response function. There are fewer limitations on how the algorithm looks away from the equilibrium.

impact on profits. In other words, we allow the dynamic process of tâtonnement to play out in every instant.²⁹

As before, we can write firm 1's stage game problem as a weighted average of the pre-update period $(0, 1/\gamma]$ and the post-update period $(1/\gamma, 1]$:

$$\max_{p_1, \sigma_1} (1 - \alpha)\pi_1(p_1, p_2) + \alpha\pi_1(\sigma_1, \sigma_2) \quad (10)$$

where $\alpha = \left(\int_0^1 e^{-\rho t} m(t) dt\right)^{-1} \int_{\frac{1}{\gamma}}^1 e^{-\rho t} m(t) dt$.³⁰ The value $1 - \alpha$ describes the relative weight that firm 1 places on the initial period $(0, 1/\gamma]$, which is a function of ρ , $m(t)$, and θ_2 . In the initial price-setting phase, the usual Nash-in-price logic holds: firm 1 treats firm 2's price as given over the period $(0, 1/\gamma]$. After $t = 1/\gamma$, firm 1 recognizes that firm 2's algorithm will control the pricing updates, and it will choose σ_1 optimally with that in mind.

As in the asymmetric game, each firm chooses a strategy that maximizes a weighted average of two profit components. As before, the first component is equivalent to the profit function for the Bertrand model. The second component is different, as firm 1 chooses σ_1 while taking into account the choice of σ_2 . To make progress on understanding the equilibria of the general setup, we analyze the equilibria of the subgame in which firms choose algorithms (σ_1, σ_2) . We can treat this component as a subgame because our setup is equivalent to a model in which firms first choose prices at $t = 0$ and then choose (σ_1, σ_2) at $t = 1/\gamma$.

This subgame merits special attention because it captures the equilibrium of the full model when both firms have high-frequency algorithms (as $\gamma \rightarrow \infty$, $\alpha \rightarrow 1$). We consider the case of $\alpha = 1$ to be a fair approximation to price competition when both firms have very high-frequency algorithms. Below, we examine the equilibria of this subgame.

4.4 Stage Game Analysis

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 market coordinator. The algorithms may condition directly on the prices of rivals. The algorithm may also be a function of variables that are observable to the firm, but they cannot be functions of other player's algorithms. This game captures price competition when both firms have very high-frequency algorithms.

²⁹Alternatively, one could explicitly model this process over discrete pricing updates determined by γ . Our focus for the symmetric commitment model is when γ is large; for this case, the process has no impact on firm profits or strategies.

³⁰The simplification is possible because the profit flow function is time-invariant. The full problem is

$$\max_{p_1, \sigma_1} \int_0^{\frac{1}{\gamma}} e^{-\rho t} \pi_1(p_1, p_2) m(t) dt + \int_{\frac{1}{\gamma}}^1 e^{-\rho t} \pi_1(\sigma_1, \sigma_2) m(t) dt.$$

After receiving the pricing algorithms, the market coordinator solves the system of equations set by the algorithms to determine prices. Based on the general model developed above, the market coordinator may be thought of as the process of tâtonnement arising from an initial price vector. 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. For example, consider the case where both firms submit an algorithm of the form

$$\sigma(p_{-j}) = \begin{cases} p^C, & \text{for } p_{-j} = p^C \\ p^B, & \text{otherwise} \end{cases} \quad (11)$$

where p^C is the collusive price and p^B is the punishment (Bertrand) price. Both (p^B, p^B) and (p^C, p^C) are equilibria of the system, depending on the initial price vector.

To resolve the issue of multiple solutions, we provide a modification to the general game that results in a unique solution conditional on algorithms. When multiple solutions are possible, the market coordinator picks the solution that minimizes the profits of the firms. If multiple such solutions exist, the coordinator randomizes among them. Effectively, we allow an adversarial market coordinator to choose the initial price vector.

Restriction 1 (Profit-Minimizing Coordinator). *In the pricing algorithm game, the market coordinator selects the solution to the system of equations determined by the algorithms that minimizes joint profits. Formally, the market coordinator chooses $p = (p_1, p_2)$ to solve*

$$\begin{aligned} \min_p \quad & \sum_{j \in \{1,2\}} \pi_j(\sigma_j(p_{-j}), \sigma_{-j}(p_j)) \\ \text{s.t.} \quad & p_j = \sigma_j(p_{-j}) \quad \forall j. \end{aligned} \quad (12)$$

If no solution exists, all firms earn zero profits.

A second and related issue is that cooperate-or-punish strategies like the one above would raise immediate antitrust concerns if made public. We wish to analyze, fundamentally, the impact of algorithmic competition on prices. Do they lead to higher prices in the absence of behavior that looks collusive? It is possible for firms to employ strategies with discontinuous punishments at the collusive price but that generate a unique solution for the coordinator. To remove all “obviously collusive” strategies from consideration, we also require firms to submit strategies that are continuous.

Restriction 2 (Continuity). *Firms must submit algorithms that are continuous functions of rivals’ prices, otherwise all firms earn zero profits.*

These restrictions provide conservative results regarding prices. We tie our own hands, eliminating equilibria that mirror typical collusive strategies, in order to demonstrate the power of

commitment. In the real world, these restrictions reflect pro-consumer market mechanisms to discipline firms. These mechanisms may be employed by antitrust authorities, savvy consumers, or a platform seeking to maximize consumer welfare.

We now define the equilibrium concept for the algorithm-setting game. In equilibrium, each firm's algorithm maximizes its own profit, conditional on the algorithms submitted by the other firms and subject to a market coordinator that minimize the joint profits when multiple solutions to the algorithms exist. We formalize this below.

Equilibrium definition: *When firms compete in pricing algorithms, equilibrium algorithms $\{\sigma_j^*\}$ satisfy*

$$\begin{aligned} \sigma_j^* &= \arg \max_{\sigma_j | \sigma_{-j}^*} \pi_j(\sigma_j(p_{-j}^*), \sigma_{-j}^*(p_j^*)) \quad \forall j & (13) \\ \text{s.t. } p^* &= \arg \min_{p \in \tilde{P}} \sum_{j \in \{1,2\}} \pi_j(\sigma_j^*(p_{-j}), \sigma_{-j}^*(p_j)) \\ \tilde{P} &\equiv \{p : p_j = \sigma_j^*(p_{-j}) \quad \forall j\}, \end{aligned}$$

resulting in equilibrium prices $p^ = (p_1^*, p_2^*)$.*

Even subject to the profit-minimizing coordinator, many equilibrium strategies can be supported. 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. Formally,

$$\pi_j(\sigma_j^*(p_{-j}^*), \sigma_{-j}^*(p_j^*)) \geq \pi_j(p_j, \sigma_{-j}^*(p_j^*)) \quad \forall p_j, j. \quad (14)$$

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 asymmetric commitment game from Section 4.2. Consider a scenario in the pricing algorithm game in which firm 1 submits algorithm $\sigma_1(\cdot) = p_1^S$ and firm 2 submits algorithm $\sigma_2(p_1) = R_2(p_1)$, where $p_1^S = \arg \max_{p_1} \pi_1(p_1, R_2(p_1))$ and $R_2(\cdot)$ is firm 2's best-response function. Recall that p_1^S is equivalent to the equilibrium price of the first-mover in a sequential pricing game. As in Section 4.2, neither firm can do better with a unilateral deviation. Thus, this asymmetric case—where one firm submits the price, and the other a function of that price—is an equilibrium of a game even when both firms have the technology to condition on the prices of rivals.

If both firms were to instead submit their best-response functions from the price-setting game, $\sigma_j(p_{-j}) = R_j(p_{-j})$, the unique price vector that satisfies both algorithms is the Bertrand equilibrium. Thus, as in Section 4.2, firm 1 can do strictly better by submitting $\sigma_1(\cdot) = p_1^S$

instead of $\sigma_1(\cdot) = R_1(p_2)$. Therefore, $(\sigma_1, \sigma_2) = (R_1, R_2)$ is not an equilibrium of the algorithm-setting game. This is a central negative result of our model.

Proposition 4. *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.

When firms compete in algorithms, the algorithms will not reflect the price-setting best-response functions in equilibrium. That is, if any firm’s algorithm depends on its rival’s price, the algorithms cannot be “competitive” in equilibrium. Further, if any firm adopts an algorithm that depends on a rival’s price, competitive prices are not obtained in equilibrium. Bertrand-Nash prices are possible only when the algorithms *do not depend* on rivals’ prices.³¹ This is a straightforward implication of the incentives illustrated in the previous section.

Though we can show that all firms choosing Bertrand best-response functions is not an equilibrium, the symmetric commitment game still admits a multitude of possible equilibria. We analyze the full set of equilibria and provide a discussion of equilibrium selection in Appendix C. Despite this result, we expect algorithms to result in higher prices than the Bertrand-Nash equilibrium for three reasons. First, when algorithms have positive slope coefficients on rivals’ prices, higher prices result. Imposing this restriction on firms’ choices seems reasonable a priori when prices are strategic complements. In other words, prices that are lower than Bertrand-Nash are supported only when an algorithm treats the rival prices as 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 Appendix C. Firms experiment with algorithms that are linear functions of rivals’ prices, 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 simulations show that higher prices result.

4.5 Algorithms, Supracompetitive Prices, and Collusive Prices

We have, thus far, address two questions related to the use of algorithms and supracompetitive prices. First, we have demonstrated that asymmetries in frequency and commitment—key features of pricing algorithms—lead to higher prices than the Bertrand equilibrium. Thus, by

³¹For example, $p^B = (p_1^B, p_2^B)$ is obtained in equilibrium if both firms resort to simple price-setting technology, with algorithms $\sigma_j(p_{-j}) = p_j^B$. More generally, when $\sigma_j(\cdot)$ is differentiable at p_{-j}^B , a necessary condition to obtain p^B in equilibrium is that $\partial\sigma_j(p_{-j})/\partial p_{-j} = 0 \ \forall j$. Otherwise, the reaction by rivals creates an incentive to deviate from the Bertrand price.

unilaterally changing one’s pricing technology, a firm can increase its prices and profits above the usual competitive benchmark. In other words, technology provides firms with a means to increase profits without resorting to collusion.

Second, we have shown that algorithms that depend on rivals’ prices *cannot* be competitive in equilibrium. Thus, if all firms use algorithms that condition on rivals’ prices, we might expect supracompetitive prices to result. As discussed above, sensible restrictions on equilibrium strategies do result in higher prices. Simulations that provide firms with a simple reinforcement learning rule to select strategies provide additional support for this conclusion.

We now address a third question: Can algorithms be used to obtain collusive outcomes in competitive equilibrium? In other words, are collusive profits possible in Markov perfect equilibrium? We again focus on the symmetric commitment game when both firms have very high-frequency algorithms ($\alpha \rightarrow 1$). In Markov perfect equilibrium, one-shot mechanics prevail so that each firm commits to an algorithm that is optimal conditional on the algorithm of its rival.

As discussed above, our restrictions rule out the typical strategies to sustain collusive behavior. However, the collusive outcome can be supported by algorithms that satisfy the restrictions. For example, in the model of demand in Section 3.3, the collusive outcome is $(p_1, p_2) = (\frac{3}{2}, \frac{3}{2})$. This is an equilibrium with the following strategies:

$$\sigma_j(p_{-j}) = 1 + \frac{1}{3}p_{-j} \tag{15}$$

It is straightforward to verify that, conditional on these algorithms, no firm wishes to deviate in its algorithm and the collusive price results. In fact, the collusive outcome $p^C = (p_1^C, p_2^C)$ can be achieved in equilibrium in general with simple linear algorithms. These algorithms take the form

$$\sigma_j(p_{-j}) = p_j^C + b_j(p_{-j} - p_{-j}^C), \tag{16}$$

where b_j is chosen to eliminate any incentive for the rival firm ($-j$) to deviate in prices.³²

The previous literature has argued that sophisticated pricing strategies employing artificial intelligence can learn to collude. However, when firms simultaneously set pricing algorithms with short-run commitment, simple linear strategies can support fully collusive prices. Importantly, these strategies do not rely on the history of prices and do not feature “severe” punishments that characterize traditional models of collusion (Harrington, 2018). Rather, collusive outcomes can be supported by marginal changes that, without detailed knowledge of demand, are indistinguishable from competitive reaction functions.

The intuition behind higher prices in our model is related to the logic of how price-matching guarantees may lead to higher prices: if a firm (credibly) commits to match the price of its rival,

³²Specifically, $b_j = - \frac{\partial \pi_{-j} / \partial p_{-j}}{\partial \pi_{-j} / \partial p_j} \Big|_{p^C}$. See derivation in Appendix C.

then the rival has a reduced incentive to lower its price. Our model allows price matching as a possible strategy, and it is straightforward to show that pure price-matching algorithms do not arise in equilibrium. If one firm chooses the price-matching algorithm $\sigma(p_{-j}) = p_{-j}$, the other will pick the collusive price. But, conditional on the second firm's price, the first firm will want to deviate along its best-response function. If both firms choose price-matching algorithms, then the adversarial market coordinator is free to pick any price that delivers the lowest profits.

Our model of symmetric commitment is also related to the analysis of conjectural variations. One important distinction is that the conjectural variations literature has attempted to restrict the set of equilibria to those in which the conjectural variations are consistent with the beliefs and actions of the other players (e.g., Bresnahan, 1981; Kamien and Schwartz, 1983; Daughety, 1985; Lindh, 1992). In the equilibria of our model of 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.

Thus, our general model unifies several different pricing games (e.g., Bertrand, sequential pricing, conjectural variations) under the same set of primitives. We view algorithms as providing a real-world foundation for many classic models of price competition. By nesting these models under a common structure, we also provide a framework for firms to choose among different models of competition by changing their pricing technology. Our model also provides a basis for more flexible assumptions of price competition that can be adapted to empirical settings. We demonstrate the importance of accounting for pricing technology when examining competition empirically in Section 5.

4.6 Pricing Algorithm Game: Oligopoly Example

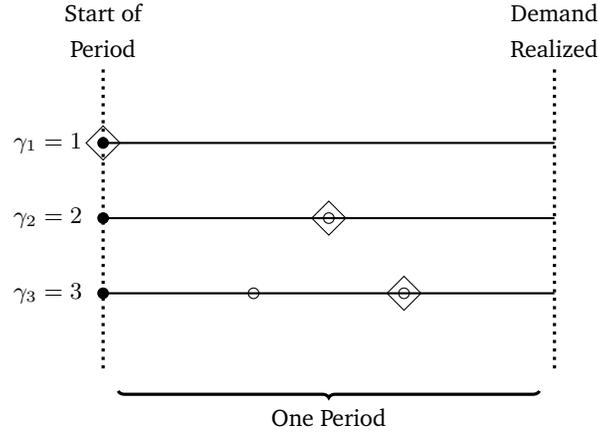
To extend the intuition of asymmetry in pricing algorithms beyond duopolistic competition, we consider an oligopoly setting with three firms. We simulate equilibrium prices in the model with the aim of comparing model predictions to our empirical findings in Section 2. Similar qualitative results can be obtained for any number of firms.

Demand remains similar to that of the model in section 3.3, 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:

$$q_j = 1 - p_j + \frac{1}{2} \sum_{k \neq j} p_k \quad (17)$$

As before, the Bertrand-Nash equilibrium is $p_j = 1$ and the collusive price is $p_j = \frac{3}{2}$ (for all j).

Figure 6: Timing for Oligopoly Example



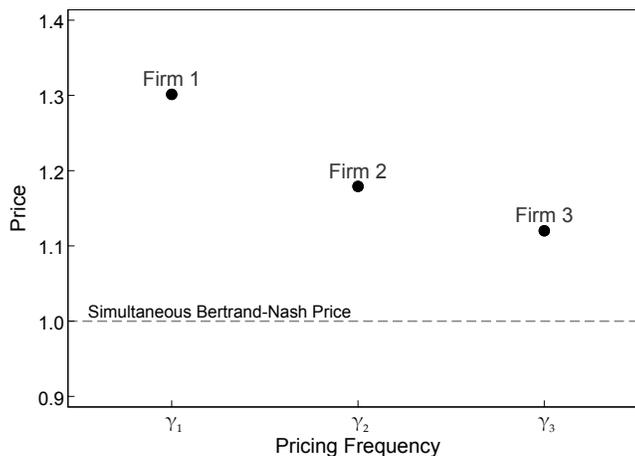
Notes: Solid black markers represent opportunities to adjust algorithms and update prices. Open circles indicate opportunities to update prices based on the previously-determined algorithm. Algorithm updates are governed by $\theta = 1$ and pricing updates are governed by γ . Diamonds indicate consequential opportunities to adjust prices when all pricing updates can occur before demand is realized.

Now assume that there are three levels of pricing technology. Firm 1 has inferior pricing technology and can update prices only at the beginning of the period. Firm 2 has more frequent pricing, allowing it to react to firm 1 before the end of the period. Firm 3 has superior technology and can update prices in response to both firm 1 and firm 2. In particular, assume $\gamma_3 > \gamma_2 > \gamma_1 = \theta$ and $\theta_j = \theta \forall j$. Furthermore, assume the differences in pricing frequency are large so that it is as if firms with faster algorithms react instantly to slower rivals. In other words, the faster algorithms can react before demand is realized.

Figure 6 illustrates how timing works in this oligopoly example. When the algorithms can react faster than demand is realized, any set of technology satisfying $\gamma_3 > \gamma_2 > \gamma_1$ will have equivalent strategic effects. In the figure, we show the edge case when $(\gamma_1, \gamma_2, \gamma_3) = (1, 2, 3)$ and all demand is realized at the end of the period. Open circles indicate pricing updates determined by the algorithms for firms 2 and 3. Diamonds indicate pricing decisions that are consequential for the realized demand. Effectively, firms with superior technology have a last-mover advantage for price. Variation in pricing technology can sort firms into a sequential pricing game, with the pricing order given by γ_j . Thus, pricing frequency provides a simple economic mechanism for firms to commit to a specific sequence, even in oligopoly settings.

Figure 7 demonstrates the equilibrium prices of the model compared to the simultaneous price-setting benchmark. 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 tech-

Figure 7: Simulated Pricing Algorithm Equilibrium



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

nology can credibly commit to offering lower prices. Within a pricing algorithm equilibrium, more frequent pricing is correlated with lower prices, consistent with the results from Section 2. In addition, the model implies all prices are elevated relative to the case where all firms have the slowest technology.

5 Quantifying the Impact of Algorithmic Competition

While previous empirical work has assumed that firms have symmetric price-setting technology, we find that differences in pricing technology is an important feature of the market we examine. As a first step towards quantifying the impact of algorithmic technology on prices, we perform a counterfactual exercise in which we study how equilibrium prices would change if firms competed via simultaneous Bertrand competition. The exercise also suggests that fitting a (misspecified) Bertrand model could generate biased estimates of markups in online markets.

To calculate counterfactual prices, we calibrate a demand system that allows for differentiation across firms and flexible substitution patterns. We apply the model to the five firms in our sample, taking into account the pricing technology of each firm. We then simulate the alternative of Bertrand competition using our calibrated model.

5.1 Demand

We introduce a tractable demand system that allows us to capture two relevant features of the market we study. First, we wish to allow for flexible substitution patterns that reflect heterogeneous demand conditions across retailers. Second, in algorithmic competition, the supply-side

optimization problem for one firm may be an input into another firm's problem. This can render estimation and simulation computationally intractable. Our demand model 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.

We model demand through the lens of spatial differentiation. Each consumer is located between two firms; these two firms represent each consumer's first and second choice at equilibrium prices. Consumers vary in their proximity to each firm, therefore the "travel" costs associated with each firm varies across consumers. In our setting, travel costs represent psychological costs and hassle costs of visiting each website. This may roughly be interpreted as search costs, though we provide no formal connection.

The model is a generalization of the Hotelling (1929) line. Unlike the circle model of Salop (1979), firms compete with all other firms, not just their closest neighbors. In this way, the model is related to the pyramid model of von Ungern-Sternberg (1991) and the spokes model of Chen and Riordan (2007). Unlike previous models, our approach allows for the mass of consumers on each segment to be different, including the mass of consumers on segments that link to an outside option. This feature is important since it allows for flexible substitution patterns that could explain differences in prices across retailers. This is also an advantage over models of vertical differentiation, such as the logit model, which restrict the horizontal substitution patterns to be symmetric across firms.

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 .³³ The distance between each firm is 1 unit. Each firm sells a single product, which consumers value at $v_j > 0$, 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 $v_0 = 0$. Consumers lie on these segments with 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 $u_{ij} > u_{ik}$, or

$$(v_j - p_j) - (v_k - p_k) > \tau(d_{ij} - d_{ik}). \quad (18)$$

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 $u_{ik} < 0$.

For our calibration exercise, we assume that consumer locations are distributed uniformly within each segment. We also assume that the products are homogeneous (but for the travel costs), so that $v_j = v$ for all j except for the outside option, for which $v_0 = 0$. Finally, we assume that consumer valuations are sufficiently high that all consumers on the inside segments

³³Demand can be represented by a graph. The graph is complete if $\mu_{jk} > 0$ for all $\{j, k\}$.

purchase a product.³⁴ Demand for retailer j is equal to

$$q_j = \sum_{k \neq j, 0} \mu_{jk} \left(\frac{1}{2} - \frac{1}{2\tau} (p_j - p_k) \right) + \mu_{j0} \frac{1}{\tau} (v - p_j). \quad (19)$$

The model flexibly captures horizontal differentiation through the distribution of consumers across segments: 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 ($p_j > v$). For additional details of this model, see Appendix D.

5.2 Supply

We consider supply-side assumptions that approximate observed pricing behavior for the five retailers examined in Section 2.2. Retailers D and E set prices simultaneously (once per week). Given the relative pricing frequency of the other firms and the fact that faster retailers respond quickly to slower retailers, we assume this is followed by retailer C , then B , and, finally, A . The sequence can be interpreted as arising from asymmetries in frequency (as in Section 3) or from asymmetric commitment (as in Section 4). The key assumption is that the faster firms can change their prices in response to slower rivals before rivals realize meaningful demand.

Retailers maximize profits given constant marginal costs, c . Under these assumptions, the firms' best-response functions are:

$$\begin{aligned} R_A(p_B, p_C, p_D, p_E) &= \arg \max_{p_A} (p_A - c) q_A(p_A, p_B, p_C, p_D, p_E) \\ R_B(p_C, p_D, p_E) &= \arg \max_{p_B} (p_B - c) q_B(R_A(\cdot), p_B, p_C, p_D, p_E) \\ R_C(p_D, p_E) &= \arg \max_{p_C} (p_C - c) q_C(R_A(\cdot), R_B(\cdot), p_C, p_D, p_E) \\ R_D(p_E) &= \arg \max_{p_D} (p_D - c) q_D(R_A(\cdot), R_B(\cdot), R_C(\cdot), p_D, p_E) \\ R_E(p_D) &= \arg \max_{p_E} (p_E - c) q_E(R_A(\cdot), R_B(\cdot), R_C(\cdot), p_D, p_E). \end{aligned}$$

Equilibrium prices are determined by the solution to this system of equations. A key advantage of the demand system in equation (19) is that it admits an analytical solution for prices.³⁵

5.3 Calibration

The goal of the calibration exercise is to find demand parameters in order to match each retailer's price index, p_j , and aggregate shares, q_j . Each firm's price index is calculated by av-

³⁴In a slight abuse of notation, we omit the arrival rate of consumers $m(t)$.

³⁵The expressions for prices are several pages long and are available upon request.

Table 4: Calibrated Segment Weights

		Retailer k					Outside
		A	B	C	D	E	
Retailer j	A	0.00	11.44	2.10	0.54	0.54	0.00
	B	11.44	0.00	2.10	0.54	0.54	1.74
	C	2.10	2.10	0.00	0.54	0.54	1.45
	D	0.54	0.54	0.54	0.00	0.54	3.07
	E	0.54	0.54	0.54	0.54	0.00	3.99

Notes: Row j column k shows the mass of customers on the segment between retailer 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.

eraging over the price of all products and then constructing an index relative to retailer A as in Figure 4. A key challenge in online markets is that market shares for individual products are rarely observed by researchers. We construct a proxy for aggregate market shares using the share of Google searches for the retailer name and the word “allergy.”³⁶ In order to help validate this measure of market share, we also obtain market shares of online personal care products for the retailers from ecommerceDB. Appendix Table 8 shows that the implied market shares are quite similar. We also assume firms have identical marginal cost, which we normalize to 1.³⁷ Price-cost margins are determined by the calibrated prices in the model.

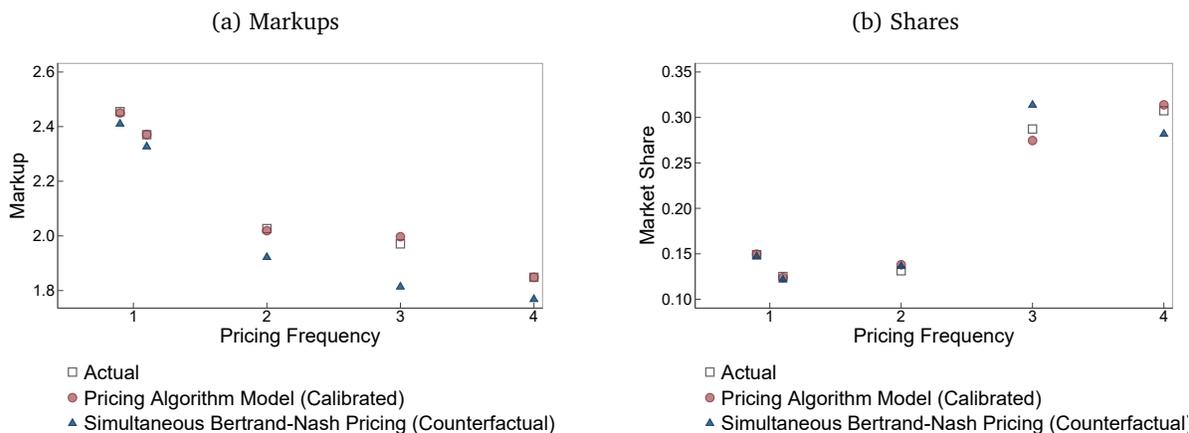
The unknown parameters to be recovered are the value of the product v , the travel cost parameter τ , and the relative weights on the segments $\{\mu_{jk}\}$. We parameterize the J by $(J+1)$ μ matrix with six parameters: $\{m_1, m_2, m_3, m_4, m_5, m_6\}$. While the fact that prices are negatively correlated 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 substitution patterns that could explain differential pricing across firms. Thus, we can use our model to capture the impacts of both preferences and pricing technology on price differences across firms.

Specifically, we choose a parameterization for the segment weights so that differences in

³⁶We use the average of Google searches for the retailer name alone as well as the retailer name in addition to “allergy.” See Appendix Table 8. The data were obtained from Google Trends (trends.google.com). Recent evidence suggests that a primary motivation for brand-specific searches is to navigate to a particular website in lieu of typing in a URL (Golden and Horton, 2021). The greater the extent that retailer-specific searches serve this navigational purpose and that conversion rates are similar across websites, the better our proxy captures aggregate shares.

³⁷In the context of allergy drugs, we argue that differences in marginal costs across retailers for identical products are relatively small. As in Ellison et al. (2018), we take wholesale costs to be common across retailers. All five retailers sell large quantities of these brands across online and brick-and-mortar channels. Shipping costs may differ among retailers, but shipping costs are a relatively small portion of the total price. The average price ranges from \$16 to \$27 across retailers, and the products are small and light. Overall, differences in marginal cost are unlikely to generate the price differences seen in Figure 4.

Figure 8: Calibration Fit for Markups and 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 3. 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.

preferences can account for differences in prices and quantities we observe in the data. For the slower firms, D and E , we constrain the segment weights so that substitution is symmetric to all other retailers: $m_1 = \{\mu_{AD}, \mu_{BD}, \mu_{CD}, \mu_{AE}, \mu_{BE}, \mu_{CE}\}$. The firm with daily pricing, C , has symmetric weights with the faster firms, $m_2 = \{\mu_{AC}, \mu_{BC}\}$. The two fastest firms have a unique weight $m_3 = \mu_{AB}$. Finally, we give each firm a unique mass for the outside option, though we set the mass for the outside option for A to zero.³⁸ 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.

The calibrated parameters for the value of the product and travel costs are $v = 5.09$ and $\tau = 0.67$. The calibrated segment weights are displayed in Table 4. 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.

We use the method of moments to choose the parameters $(v, \tau, \{\mu_{jk}\})$ 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 3, and relative average shares

³⁸We normalize the density along the outside option segment for E to equal 1, which pins down the value of the distance D_0 . Thus, $(\mu_{A0}, \mu_{B0}, \mu_{C0}, \mu_{D0}, \mu_{E0}) = (0, m_6, m_5, m_4, 1)$, generating the outside option consumer mass vector $(0, m_6 D_0, m_5 D_0, m_4 D_0, D_0)$.

Table 5: Own-Price and Cross-Price Demand Elasticities

		Retailer Price				
		A	B	C	D	E
Retailer Share	A	-2.18	1.84	0.34	0.10	0.11
	B	1.95	-2.83	0.39	0.12	0.12
	C	0.71	0.77	-2.18	0.23	0.24
	D	0.20	0.22	0.22	-1.77	0.27
	E	0.17	0.18	0.18	0.22	-1.72

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

using our proxy for quantities.³⁹

The fit of the calibration exercise is displayed in Figure 8. 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. Allowing for flexible substitution patterns is important; if we had instead assumed symmetric demand, we would not be able to rationalize the data. 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 down by the first-order conditions, allowing us to recover an estimate of markups. The calibrated parameters imply reasonable price-cost margins between 0.46 (retailer *A*) and 0.59 (retailer *E*).

Table 5 shows a matrix of elasticity of demand estimates from the model. Own-price elasticities range from -1.7 to -2.8 , consistent with other estimates from online goods.⁴⁰ Our estimated cross-price elasticities indicate that, when the price of a product increases, consumers are more likely to substitute towards 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*.

5.4 Counterfactual

To illustrate the potential impact of pricing algorithms on prices, we use our calibrated model to predict equilibrium prices if all firms instead had simultaneous price-setting technology. The Bertrand equilibrium prices and shares are displayed with green triangles in Figure 8. Our model indicates that algorithmic competition increases the average price by 5.2 percent above

³⁹In 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.

⁴⁰See, for instance, De los Santos et al. (2012).

Table 6: Counterfactual Effects on Markups and Profits

Firm	Simultaneous Bertrand			Algorithmic Competition			Percent Change		
	Markup	Share	Profit	Markup	Share	Profit	Markup	Share	Profit
A	1.77	0.282	6.5	1.85	0.314	7.9	4.5	11.4	22.0
B	1.81	0.314	7.6	2.00	0.275	8.1	10.1	-12.4	6.4
C	1.92	0.136	3.8	2.02	0.138	4.2	5.1	1.3	11.1
D	2.33	0.121	4.8	2.37	0.124	5.0	1.9	2.0	4.5
E	2.41	0.147	6.2	2.45	0.150	6.4	1.7	1.8	3.8
Aggregate	1.97	1	28.9	2.07	1	31.7	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.

the counterfactual Bertrand equilibrium. These price changes differ across firms. Firms *D* and *E* realize more modest price changes of 1.9 and 1.7 percent. Based on our calibrated demand parameters, 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 6. 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 from asymmetric technology. By contrast, 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 3, all firms profit as a result of algorithmic competition.

Our model predicts that algorithmic competition results in a modest decline in market-level quantities of 0.9 percent. This limited substitution to the outside option 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.

6 Conclusion

Online markets were initially expected to usher in “frictionless commerce” and intensify competition among firms (Ellison and Ellison, 2005). Our results demonstrate how advances in pricing technology can have the opposite effect, generating higher prices and exacerbating price dispersion. By employing high-frequency pricing algorithms, firms can soften competition and increase profits in equilibrium, even if the firms are otherwise identical. In our theoretical examples and our counterfactual simulation, the largest gains accrue to a dominant firm with the most advanced technology and the largest market share. While standard models in microeconomics and macroeconomics often assume symmetric pricing technology across firms, we show that accounting for this asymmetry can be quite important.

Our findings suggest that the Bertrand equilibrium may be the exception in online markets, rather than the rule. This raises new considerations for future policies about digital markets, as the potential role of algorithms is much more broad than facilitating collusion. As we show, simple pricing algorithms can increase prices in competitive equilibrium and may even obtain the fully collusive outcome. To prevent such price increases, policymakers would have to limit the ability of firms to react to rivals’ prices.⁴¹ One solution would be to prohibit algorithms from directly conditioning on rivals’ prices, while still allowing firms to have frequent price updates as a function of other factors, such as demand shocks. Besides prohibiting the behavior, policymakers could limit the scraping of rival firms’ prices or restrict the storage of recent prices by other firms; either of these policies may be more feasible to implement and yield similar results.⁴² However, enforcement measures along these lines do not fit neatly into existing regulatory and antitrust frameworks in most countries. Thus, the growing use of algorithms raise conceptual and legal challenges that merit further consideration.

Though we focus on competitive equilibria, our study also has important implications for collusion. First, the competitive equilibrium is typically used as “punishment” in a collusive equilibrium. In our model, pricing algorithms can support a competitive equilibrium with

⁴¹In our analysis, rivals’ prices play a special role. Retail prices are public and immediately available, allowing firms to respond to changes in real time. 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 a frequency that would be useful to the algorithm. Furthermore, the use of rival-specific measures (prices) provides firms with several instruments to discipline price competition.

⁴²Alternatively, policymakers could regulate the frequency with which firms update their algorithms and their prices. This could restore simultaneous pricing and limit the ability of rival firms to react. Pricing frequency regulation has been applied to retail gasoline markets in Austria and Australia.

higher profits than the Bertrand equilibrium. Thus, pricing algorithms can make punishment less severe, reducing the likelihood of collusion. On the other hand, our model explicitly considers the ability of firms to increase their pricing frequency. In addition to making collusive strategies more feasible, high-frequency pricing also gives firms the ability to obtain collusive profits with linear, non-collusive strategies.

Online sales represent an increasing share of many diverse markets, including insurance, accommodations, and automobiles, in addition to retail goods. In all of these sectors, the shift online coincides with an increased availability of publicly posted prices and pricing technology that uses these prices as inputs. Offline markets are increasingly adopting pricing algorithms as well, and similar issues arise if brick-and-mortar stores adopt these methods. Though we view the issues raised in this paper as quite general, there is a large scope for future research that incorporates other features of these markets and examines additional implications of competition in pricing algorithms.

References

- Asker, J., C. Fershtman, and A. Pakes (2021). Artificial Intelligence and Pricing: The Impact of Algorithm Design. Working paper, National Bureau of Economic Research.
- Assad, S., R. Clark, D. Ershov, and L. Xu (2020). Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. CESifo Working Paper 8521.
- Baye, M. R., J. Morgan, and P. Scholten (2004). Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site. *The Journal of Industrial Economics* 52(4), 463–496.
- Bonanno, G. and J. Vickers (1988). Vertical Separation. *The Journal of Industrial Economics* 36(3), 257.
- Bowley, A. L. (1924). *Mathematical Groundwork of Economics*. California Press, Oxford.
- Bresnahan, T. F. (1981). Duopoly Models with Consistent Conjectures. *The American Economic Review* 71(5).
- Brynjolfsson, E., A. A. Dick, and M. D. Smith (2010). A Nearly Perfect Market? *Quantitative Marketing and Economics* 8(1), 1–33.
- 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 (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110(10), 3267–97.
- Cavallo, A. (2017). Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers. *American Economic Review* 107(1), 283–303.
- Cavallo, A. (2018). More Amazon Effects: Online Competition and Pricing Behaviors. Working paper, National Bureau of Economic Research.
- Chen, L., A. Mislove, and C. Wilson (2016). An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. In *Proceedings of the 25th International Conference on World Wide Web*, pp. 1339–1349.
- Chen, Y. and M. H. Riordan (2007). Price and Variety in the Spokes Model. *The Economic Journal* 117(522), 897–921.
- Conlon, C. T. and N. Rao (2019). The Price of Liquor is Too Damn High: Alcohol Taxation and Market Structure. NYU Wagner Research Paper 2610118.
- Daughety, A. F. (1985). Reconsidering Cournot: The Cournot Equilibrium is Consistent. *The Rand Journal of Economics*, 368–379.
- 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.

- Deck, C. A. and B. J. Wilson (2000). Interactions of Automated Pricing Algorithms: An Experimental Investigation. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pp. 77–85.
- Deck, C. A. and B. J. Wilson (2003). Automated Pricing Rules in Electronic Posted Offer Markets. *Economic Inquiry* 41(2), 208–223.
- Ellison, G. and S. Ellison (2005). Lessons about Markets from the Internet. *Journal of Economic Perspectives* 19(2), 139–158.
- Ellison, S. F., C. Snyder, and H. Zhang (2018). Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market. Working paper, National Bureau of Economic Research.
- Fellner, W. (1949). *Competition Among the Few*. Kelley, New York, NY.
- Fershtman, C. and K. L. Judd (1987). Equilibrium Incentives in Oligopoly. *The American Economic Review*, 927–940.
- Golden, J. and J. J. Horton (2021). The Effects of Search Advertising on Competitors: An Experiment Before a Merger. *Management Science* 67(1), 342–362.
- Grossman, S. J. (1981). Nash Equilibrium and the Industrial Organization of Markets with Large Fixed Costs. *Econometrica* 49(5), 1149–1172.
- Hamilton, J. H. and S. M. Slutsky (1990). Endogenous Timing in Duopoly Games: Stackelberg or Cournot Equilibria. *Games and Economic Behavior* 2(1), 29–46.
- Harrington, J. E. (2018). Developing Competition Law for Collusion by Autonomous Artificial Agents. *Journal of Competition Law & Economics* 14(3), 331–363.
- Hay, G. A. (1981). Oligopoly, Shared Monopoly, and Antitrust Law. *Cornell L. Rev.* 67, 439.
- Hong, H. and M. Shum (2006). Using Price Distributions to Estimate Search Costs. *The RAND Journal of Economics* 37(2), 257–275.
- Hotelling, H. (1929). Stability in Competition. *The Economic Journal* 39(153), 41.
- Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet (2010). 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.
- Klein, T. (2019). Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing. Working paper, SSRN 3195812.
- Klemperer, P. D. and M. A. Meyer (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica* 57(6), 1243–1277.
- Lazarev, J. (2019). Getting More from Less: Understanding Airline Pricing. Working paper.

- Lindh, T. (1992). The Inconsistency of Consistent Conjectures: Coming back to Cournot. *Journal of Economic Behavior & Organization* 18(1), 69–90.
- Maskin, E. and J. Tirole (1988a). A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs. *Econometrica* 56(3), 549–569.
- Maskin, E. and J. Tirole (1988b). 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.
- Moorthy, S. and R. A. Winter (2006). Price-Matching Guarantees. *The RAND Journal of Economics* 37(2), 449–465.
- Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica* 69(2), 307–342.
- O'Connor, J. and N. E. Wilson (2019). Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition. Working paper, SSRN 3406834.
- 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.
- Porter, R. H. (1983). A Study of Cartel Stability: The Joint Executive Committee, 1880-1886. *The Bell Journal of Economics*, 301–314.
- Salcedo, B. (2015). Pricing Algorithms and Tacit Collusion. Working paper.
- Salop, S. (1986). Practices that (Credibly) Facilitate Oligopoly Coordination. In J. Stiglitz and F. Mathewson (Eds.), *New Developments in the Analysis of Market Structure*. MIT Press, Cambridge, MA.
- Salop, S. C. (1979). Monopolistic Competition with Outside Goods. *The Bell Journal of Economics* 10(1), 141–156.
- Sklivas, S. D. (1987). The Strategic Choice of Managerial Incentives. *The RAND Journal of Economics* 18(3), 452–458.
- Tennenholtz, M. (2004). Program Equilibrium. *Games and Economic Behavior* 49(2), 363–373.
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.
- von Ungern-Sternberg, T. (1991). Monopolistic Competition on the Pyramid. *The Journal of Industrial Economics*, 355–368.

Appendix

A Details on High-Frequency Price Data

In this section, we provide further details on the collection of high-frequency price data and product definitions.

We focus on the main seven brands of allergy drugs. 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 with a decongestant. These varieties are captured by the *variant* of the drug. 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.⁴³ 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. 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.

Due to the technological challenges involved in collecting high-frequency data, there is concern about measurement error. We address this in a few ways. 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.⁴⁴ Table 7 displays the count of observations by brand and retailer.

Figure 9 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

⁴³We drop multipacks that are of greater size than a twinpack, as they are not common across retailers.

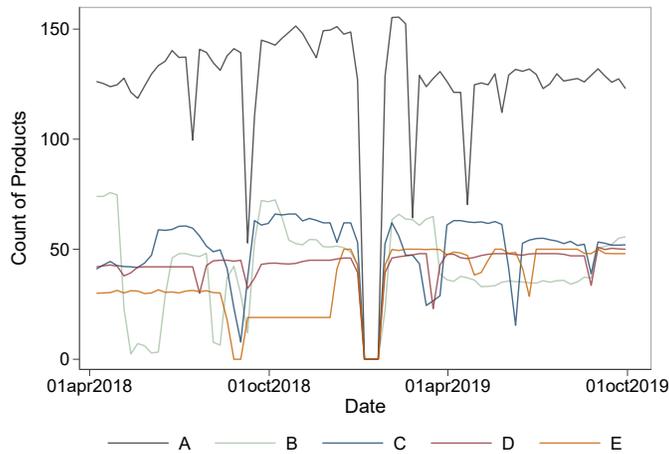
⁴⁴Overall, 7.8 percent of the prices are imputed in our analysis sample.

Table 7: Price Observations by Website and Brand

Retailer	Allegra	Benadryl	Claritin	Flonase	Nasacort	Xyzal	Zyrtec	Total
A	309,554	219,098	508,768	104,634	66,178	108,854	236,044	1,553,130
B	126,738	58,270	144,098	46,584	12,517	34,177	75,096	497,480
C	89,477	99,608	171,782	80,772	34,633	32,508	90,858	599,638
D	112,273	68,466	128,385	50,130	2,411	47,321	128,123	537,109
E	71,061	47,799	125,171	51,732	38,051	23,185	62,600	419,599
Total	709,103	493,241	1,078,204	333,852	153,790	246,045	592,721	3,606,956

Notes: Count of price observations for the sample period from April 10, 2018 through October 1, 2019.

Figure 9: Observed Products Over Time



Notes: Figure displays the average daily count of observed products in our sample by week and by retailer. Dips in the data correspond to changes to the retailer website and issues with the researchers' servers. 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.

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.

B Endogenous Pricing Frequency

B.1 Adoption Game

In this appendix, we provide a two-stage game in which firms can initially choose their pricing technology, before choosing prices. Firms are characterized by pricing technology $\theta_j \in \{1, 2, 3, \dots, \bar{\theta}\}$, where a higher value represents superior technology and $\bar{\theta}$ represents the best available technology. Firms can adopt $\theta_j = 1$ at zero cost or pay an adoption cost A to choose any other feasible technology. Firms compete in the pricing game after determining their technology.

In the model, the profits do not depend directly on the technology each firm has, but rather on their relative order. Denote the profits for the superior technology firm as π^H , the profits for the inferior technology firm as π^D , and the profits for when they have the same technology as π^S . Following the results from the main text, $\pi^H > \pi^D > \pi^S$. We assume that $\pi^H - \pi^S > A$, so that it can be profitable for one firm to adopt costly technology.

We now characterize equilibria of the game. Without loss of generality, let firm 2 represent the firm with (weakly) superior technology in equilibrium. To characterize the equilibria, there are two relevant cases to consider.

Case 1: $\pi^H - \pi^D \geq A$. Under these conditions, a pure-strategy equilibrium is for firm 2 to choose the best available technology ($\theta_2 = \bar{\theta}$) while firm 1 chooses $\theta_1 = 1$. It must be both profitable for firm 2 to adopt a superior technology, relative to symmetric technologies (this is true by assumption), and firm 2 must choose a technology so that firm 1 would not want to “leapfrog” firm 2’s choice. As the adoption cost is the same for any technological improvement, firm 2 must choose the best possible technology. The firm with superior technology has higher profits.

Case 2: $\pi^H - \pi^D < A$. The pure-strategy equilibria are characterized by firm 2 adopting any technology $\theta_2 > 1$ and by firm 1 choosing $\theta_1 = 1$. Firm 2 is indifferent to the exact level of technology because firm 1 has no incentive to invest in superior technology in equilibrium. In fact, the firm with inferior technology has higher profits (net of adoption costs) in this scenario. Thus, the firm that adopts superior technology is only motivated to do so to break the symmetric outcome, in which both realize lower profits. Though it competes more aggressively and realizes higher profits in the pricing game, it would prefer to be in firm 1’s position.

The pure strategy equilibria result in higher prices and higher profits for both firms, compared to the simultaneous price-setting equilibrium. As a corollary, any mixed strategy equilibrium also has higher expected prices and profits than the simultaneous price-setting equilibrium. Firms have a positive profit incentive to endogenously sort into asymmetric pricing technologies.

To illustrate this point, consider the three-by-three first-stage game where firms can choose pricing frequency and adoption is costless ($A = 0$). Firms know the profits for each subgame

Figure 10: Example Pricing Frequency Adoption Game

		Firm 2		
		Low	Moderate	High
Firm 1	Low	(1.00, 1.00)	(1.08, 1.21)	(1.08, 1.21)
	Moderate	(1.21, 1.08)	(1.00, 1.00)	(1.08, 1.21)
	High	(1.21, 1.08)	(1.21, 1.08)	(1.00, 1.00)

when they choose a low frequency, a moderate frequency, or a high frequency ($\theta \in \{1, 2, 3\}$). Figure 10 presents the payoffs based on the illustrative model in Section 3.3 when $\alpha = 0.5$. Any scenario where both firms choose the same frequency—low, moderate, or high—is not an equilibrium, because each firm has an incentive to deviate by choosing either a faster or a slower pricing technology. The only equilibria of the game are asymmetric where only one player chooses the highest frequency.

B.2 Adoption with an Initial Endowment of Technology

To further highlight the motivation for firms to make asymmetric choices in technology, we now consider a variant of the game above where both firms are initially endowed with technology $\theta^e > 1$. To change to a different technology, firms pay an adoption cost A as before, but they may costlessly retain their endowment or costlessly switch to $\theta = 1$. The costs for the initial endowment are sunk, so there is no salvage value for the endowed technology.

Without loss of generality, suppose that firms are initially endowed with $\theta^e = 2$. If $\pi^H - \pi^D \geq A$, then, similarly to case 1 above, the equilibrium has firm 2 choosing $\bar{\theta}$, while firm 1 keeps its initial endowment $\theta_1 = \theta^e$.⁴⁵

Now suppose that $\pi^H - \pi^D < A$, so that surpassing your rival with costly investments is not profitable. In this scenario, the unique pure-strategy equilibrium is for firm 1 to *downgrade* its technology to $\theta_1 = 1$ and for firm 2 to maintain its endowment. Here, firms willingly choose inferior technology to generate asymmetry. This is profitable for both firms, but it is less profitable for the firm that gives up its initial endowment. Perhaps surprisingly, this result holds even when there is some cost to downgrade (a), provided that the asymmetric outcome is still more profitable for firm 1 than the symmetric outcome ($\pi^D - a > \pi^S$, and also $\pi^D - a > \pi^H - A$).

B.3 Discussion

The simple adoption game highlights a few properties of the price competition when firms vary in pricing frequency. First, the incentive to have asymmetric technologies is quite robust. A firm may adopt costly technology even if its rival gains more from the outcome, as the firm

⁴⁵If firm 1 were to costlessly reduce its technology to $\theta_1 = 1$, firm 2 would prefer to keep its initial endowment. But this is not an equilibrium because firm 1 would then optimally leapfrog firm 2.

prefers this outcome to the world in which neither firm adopts. A firm may even pay a cost to downgrade its technology, if the firm and its rival are endowed with similar technology to begin with. Thus, though the most salient case for asymmetry is one in which the investing firm gains vis-a-vis its rivals, firms may even be willing to disadvantage themselves relative to their rivals to gain the benefits of softened price competition.

The above equilibrium results also apply if technology adoption is costless. Thus, if firms can choose their pricing technology at costs that are not prohibitively high, then we should not expect simultaneous price-setting behavior to hold in equilibrium. This raises some interesting considerations for empirical researchers, where a simultaneous price-setting behavior is the standard assumption.

When extending the analysis to dynamic settings, the model provides potentially interesting interpretations of observed phenomena. In the first case discussed above, we have one firm adopting the best available technology, and the other firm choosing to not invest at all in costly technology. Thus, this model has flavor of a one-sided “arms race,” where the superior technology firm over-invests in technology to prevent being bested by its rival. This over-investment can be quantified in a more general model where the cost of adoption depends on the technology level, i.e., as a (weakly increasing) function, $A(\theta)$. We omit an exposition of the model here, as it can complicate the analysis by eliminating all pure-strategy equilibria.

Over multiple periods, it would be possible to observe an arms race if the best-available technology were increasing over time, and firms maintained their technology from the previous period. With an increase in $\bar{\theta}$ from one period to the next, firm 1 would find it profitable to leapfrog firm 2, and, if the positions switch, a future increase in $\bar{\theta}$ would allow firm 2 to again overtake firm 1.

C Equilibrium Selection

C.1 A Multitude of Equilibria

It is possible to show that a multitude of equilibria can exist 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 5. *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:

$$\left. \frac{d\pi_1}{dp_1} \right|_{\hat{p}} = \frac{\partial \pi_1}{\partial p_1} + \frac{\partial \pi_1}{\partial p_2} \frac{\partial \sigma_2}{\partial p_1} \Big|_{\hat{p}} = 0 \quad (20)$$

$$\implies \left. \frac{\partial \sigma_2}{\partial p_1} \right|_{\hat{p}} = - \left. \frac{\partial \pi_1 / \partial p_1}{\partial \pi_1 / \partial p_2} \right|_{\hat{p}} \quad (21)$$

Likewise, $\left. \frac{\partial \sigma_1}{\partial p_2} \right|_{\hat{p}} = - \left. \frac{\partial \pi_2 / \partial p_2}{\partial \pi_2 / \partial p_1} \right|_{\hat{p}}$ 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:

$$\hat{p}_1 = \sigma_1(\hat{p}_2) = a_1 + b_1 \hat{p}_2 \quad (22)$$

$$\hat{p}_2 = \sigma_2(\hat{p}_1) = a_2 + b_2 \hat{p}_1 \quad (23)$$

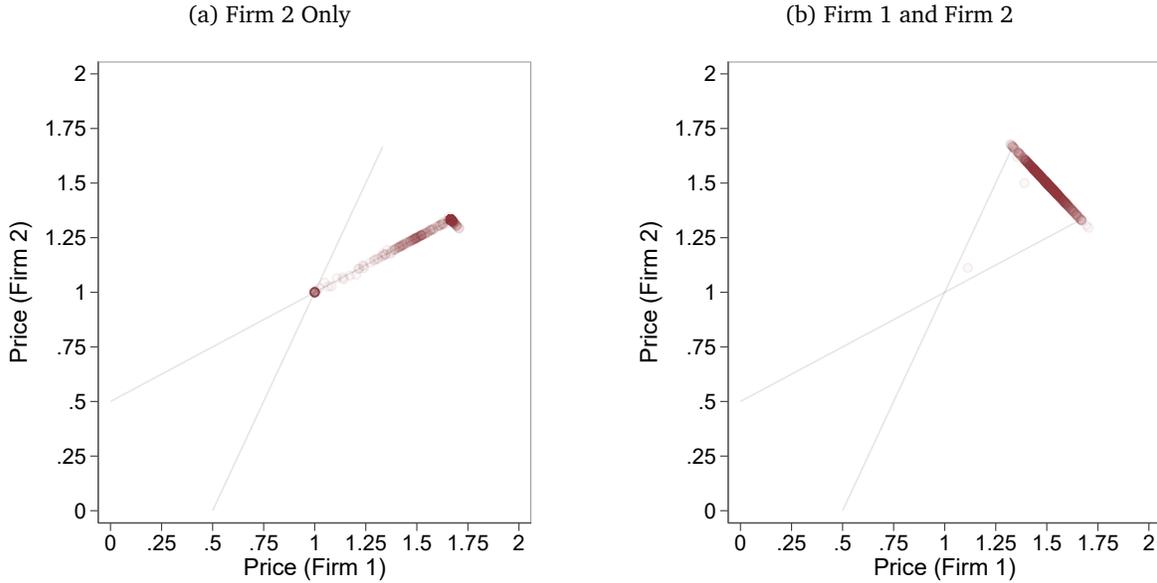
It is apparent that the solution has $b_1 = - \left. \frac{\partial \pi_2 / \partial p_2}{\partial \pi_2 / \partial p_1} \right|_{\hat{p}}$ and $b_2 = - \left. \frac{\partial \pi_1 / \partial p_1}{\partial \pi_1 / \partial p_2} \right|_{\hat{p}}$. Thus, each equation has one unknown, and the system has a unique solution for the parameters a_1 and a_2 . It is straightforward to extend the argument to many firms.⁴⁶

C.2 Simulations

Despite this multiplicity result, we expect algorithms to result in higher prices than the Bertrand-Nash equilibrium. We discuss these reasons in the main text. Here, we highlight one of the

⁴⁶For 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 $a_j = \hat{p}_j - b_{jk} \hat{p}_k$.

Figure 11: Equilibrium Selection with Pricing Algorithms



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.

reasons: many of these equilibria are “knife-edge” cases. To examine which equilibria are, in some sense, more robust, we simulate a simple learning process. We allow firms to 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 than those of the Bertrand equilibrium.

To test this intuition, we simulate a simple learning process to select equilibria. We follow the duopoly setup of Section 3.3 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 11. 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. Some simulations arrive at the Bertrand-Nash equilibrium because the firms never realize more profitable algorithms strategies. The second mass point corresponds to the equilibrium of the sequential pricing game.

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. 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.

D Details of Spatial Differentiation Model

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. In this section, we present additional formal details about the model. For further motivation, see Section 5.1.

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 .⁴⁷ The distance between each firm is 1 unit. Each firm sells a single product, which consumers value at $v_j > 0$, 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 $v_0 = 0$. 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. Figure 12 provides a visual representation of the demand system for the case of three firms.

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 $u_{ij} > u_{ik}$, or

$$(v_j - p_j) - (v_k - p_k) > \tau(d_{ij} - d_{ik}). \quad (24)$$

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 $u_{ik} < 0$.

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 .

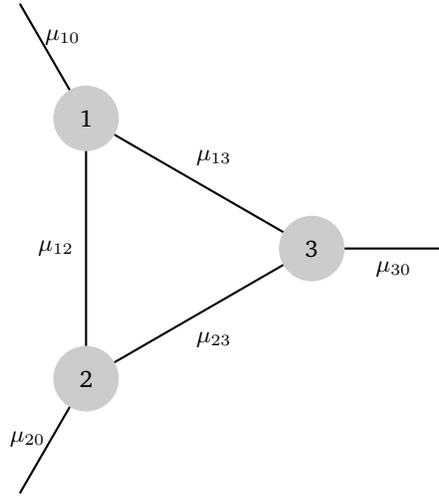
Noting that $d_{ik} = 1 - d_{ij}$ for a consumer on segment jk , a consumer on this segment will choose j if $u_{ij} > u_{ik}$ and if $u_{ij} \geq 0$, i.e., $\frac{1}{2} + \frac{1}{2\tau} ((v_j - p_j) - (v_k - p_k)) > d_{ij}$ and $\frac{1}{\tau} (v_j - p_j) \geq d_{ij}$. Firm j receives customers for which d_{ij} satisfies both conditions. Therefore, firm j receives a quantity of $\mu_{jk}F(y_{jk})$ from line segment jk , where

$$y_{jk} = \min \left\{ \frac{1}{2} + \frac{1}{2\tau} ((v_j - p_j) - (v_k - p_k)), \frac{1}{\tau} (v_j - p_j) \right\}. \quad (25)$$

For the outside segments, $y_{j0} = \frac{1}{D_0} \frac{1}{\tau} (v_j - p_j)$, 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 $\tilde{\tau}_0 = D_0\tau$.

⁴⁷Demand can be represented by a graph. The graph is complete if $\mu_{jk} > 0$ for all $\{j, k\}$.

Figure 12: Spatial Differentiation Model with Three Firms



Notes: Example of demand for three firms with an outside option. The mass of consumers along each segment is given by μ_{jk} . The segments with mass μ_{10} , μ_{20} , and μ_{30} represent consumers whose next-best alternative to the linked firm is the outside option.

Overall, quantities are given by

$$q_j = \sum_{k \neq j} \mu_{jk} F(y_{jk}). \quad (26)$$

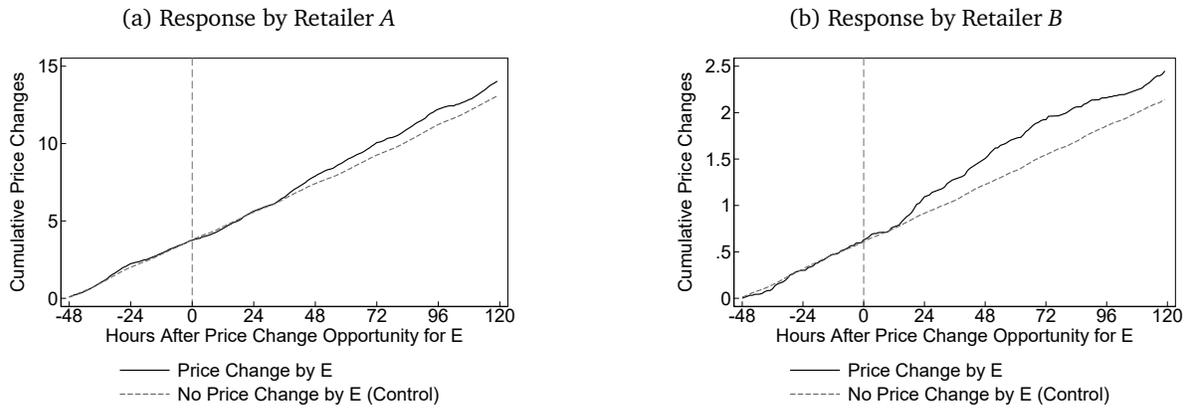
The flexibility in substitution patterns from this relatively parsimonious model comes primarily through the mass of consumers on each segment $\{\mu_{jk}\}$ and the choice of distribution F . In equilibrium, the consumers $\{\mu_{j0}\}$ that have no next-best alternative other than the outside option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

We introduce some terminology to facilitate 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*.⁴⁸ 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.

⁴⁸When $y_{jk} \geq 1$, the segment is *dominated* by j .

E Additional Tables and Figures

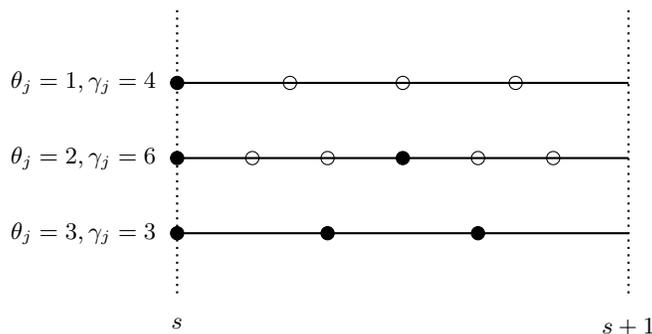
Figure 13: Price Changes by Fastest Retailers in Response to a Price Change by Retailer *E*



Notes: Figure displays the cumulative price changes of high-frequency retailers *A* and *B* in response to a price change occurring at retailer *E*. The solid 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.

Figure 13 shows the reaction of high-frequency firms (retailers *A* and *B*) to price changes by low-frequency retailer *E*. The charts imply that the high-frequency firms respond to a price change by Retailer *E* within about 48 hours.

Figure 14: Timing with Pricing Technology (θ, γ)



Notes: Solid black markers represent opportunities to adjust algorithms and update prices. Open circles indicate opportunities to update prices based on the previously-determined algorithm. Algorithm updates are governed by θ and pricing updates are governed by γ .

Figure 14 illustrates the timing of pricing decisions in period s of the repeated pricing algorithm game. Pricing technology for firm j is governed by the frequency with which the firm can update its algorithm (θ_j) and the frequency that it can update prices (γ_j). When $\gamma_j > \theta_j$, the firm has a short-run commitment to update prices according to the previously-determined algorithm, $\sigma_j(\cdot)$.

Table 8: Measures of Retailer Market Share

Retailer	Share of Online Personal Care	Google Search		
		“Retailer name”	“Retailer name” + Allergy	Mean
A	0.338	0.427	0.188	0.307
B	0.252	0.311	0.263	0.287
C	0.084	0.139	0.123	0.131
D	0.119	0.062	0.188	0.125
E	0.207	0.061	0.237	0.149

Notes: Share of personal care category reflect 2019 revenue figures from ecommerceDB.com. This includes online sales of medical, pharmaceutical, and cosmetic products for each of the retailers, including sales through mobile channels. Google search figures refer to the searches over the sample period as a share of total searches for all of the five retailers. Google search data are obtained from Google Trends (trends.google.com).

Table 8 provides measures of aggregate shares for the retailers in our data. We calibrate our model to Google search shares, using the mean of search shares for the retailer name and search shares for the retailer name along with the word “allergy.” We cross-check these shares against revenue shares provided by ecommerceDB.com. The measures of online revenue shares are obtained for the category of personal care, which includes all medical, pharmaceutical, and cosmetic products. Four of our retailers are in the top five for the personal care category by revenue, and all are in the top ten. The other retailers in the top ten have a focus on cosmetics.