Algorithmic Pricing and Consumer Sensitivity to Price Volatility

Diego Aparicio[1], Dean Eckles[2], and Madhav Kumar*[2]

[1]IESE Business School

[2]MIT

July 1, 2021

Abstract

Algorithmic pricing can be broadly defined as a formula to set prices by a computer. It is typically associated with a lower cost of changing prices and a greater frequency of price changes. While commonly observed in ride-sharing, lodging, and airline tickets, there has been recent evidence of its implementation in pharmaceutical drugs, gasoline, and online retail. However, little is known about how consumers respond to encountering frequently changing prices for goods. Here we use detailed clickstream data from an online retailer that varied pricing methods to examine how exposure to the frequently-changing prices feature of algorithmic pricing affects purchase behavior. Aggregate analysis at the product-week level, before-and-after event studies around adoption time, and granular user-level models, all show a consistent pattern — exposure to price volatility increases price sensitivity. This is economically consequential because, even if implementing algorithmic pricing is profitable, it triggers unintended side effects that modify consumer behavior in ways that may undermine those gains. We complement these empirical findings with laboratory experiments and provide evidence for a key underlying mechanism—price salience.

Keywords— algorithmic pricing, price volatility, price sensitivity, salience, experiments

---

*To whom correspondence may be addressed: madhavk@mit.edu

Author names are in alphabetical order.

The authors thank Sinan Aral, Rahul Bhui, Sebastien Brion, Bart Bronnenberg, Glenn Ellison, Isabelle Engeler, Pedro Gardete, Els Gijsbrechts, John Hauser, Pranav Jindal, Anja Lambrecht, Donald Lehmann, Greg Lewis, Alex MacKay, Carl Mela, Sridhar Moorthy, Johannes Muller-Trede, Thomas Otter, Drazen Prelec, David Rand, Elena Reutskaja, Sampsa Samila, Duncan Simester, Avner Strulov-Shlain, Catherine Tucker, Birger Wernerfelt, Jeremy Yang, Gokhan Yildirim, Juanjuan Zhang, and seminar participants at 2021 ZEW ICT Conference, 2021 Marketing Science Conference, 2021 Theory+Practice Marketing Conference, and 2021 EMAC Conference.

1 Introduction

$6.19 at 10:30 pm on Sunday, $6.39 at 3:28 am on Monday, $5.99 at 3:42 am, $2.99 at 4:28 am, $4.26 at 4:44 am, $3.99 at 8:40 am, and $4.47 at 12:21 pm. One may be forgiven for assuming these are prices for a stock listed on the stock exchange. These are, in fact, seven distinct prices of a single regular carbonated cola drink over the course of only two days in an online grocery retailer in the United States. How do consumers react when they see prices changing frequently?

Algorithmic pricing has been expanding across industries and channels. What perhaps used to be a specialized feature of airline tickets (McAfee and Te Velde, 2006) has now been documented in ride-sharing platforms (Chen, 2016; Cohen, Hahn, Hall, Levitt, and Metcalfe, 2016), gasoline markets (Assad, Clark, Ershov, and Xu, 2020), allergy drugs in online retailers (Brown and MacKay, 2019), and Amazon's durable goods marketplace (Chen, Mislove, and Wilson, 2016).

However, as with artificial intelligence or other forms of automation technologies (Brynjolfsson and McAfee, 2014; Ford, 2015; Agrawal, Gans, and Goldfarb, 2018), algorithmic pricing is an intangible concept that is not easy to dissect. The most salient feature identified in the literature is the striking price volatility, as measured by the number of price changes, over time (Chen, Mislove, and Wilson, 2016; Calvano, Calzolari, Denicolò, and Pastorello, 2019; Brown and MacKay, 2019; Assad, Clark, Ershov, and Xu, 2020). Research has shown that sellers that adopt algorithmic pricing are found to update prices several times per day. For example, Amazon is known to change product prices ~2.5 million times a day or, equivalently, the price for a product listed on Amazon changes every 10 minutes on average (Business Insider, 2018). Comparable examples from other industries include *Smart Pricing* by Airbnb (Airbnb, 2017) and *Surge Pricing* by Uber (Dholakia, 2015). In Uber's case, prices change as frequently as every three to five minutes (Washington Post, 2015).

We obtain clickstream data from an online retailer that contains abundant examples similar to the one in the introductory paragraph. Figure 1 shows visually compelling evidence. We can distinguish between periods of stable prices initially versus those of extremely volatile prices later on.

A primary focus of the literature has been on studying whether, and how, the adoption of algorithmic pricing can alter competition incentives across rival firms (Miklós-Thal and Tucker, 2019; Calvano, Calzolari, Denicolò, and Pastorello, 2019; Brown and MacKay, 2019; Hansen, Misra, and Pai, 2021; Asker, Fershtman, and Pakes, 2021). However, despite its increasing prevalence, we know little about how *consumers* react when they are exposed to the strikingly high price volatility of machine algorithms. We contribute to this discussion by studying whether and how consumers' price sensitivity reacts to heightened price volatility, as measured by frequency of price changes and exposure to multiple unique prices, caused by pricing algorithms. While we are mindful that price sensitivity covers just one dimension of the greater realm of consumer behavior, price sensitivity has been a fundamental question in the economics and marketing literature. To illustrate, early papers on advertising were in fact absorbed about the connection between advertising and price sensitivity (Dorfman and Steiner, 1954; Becker and Murphy, 1993). And to date, this question remains contested (Sethuraman, Tellis, and Briesch, 2011).

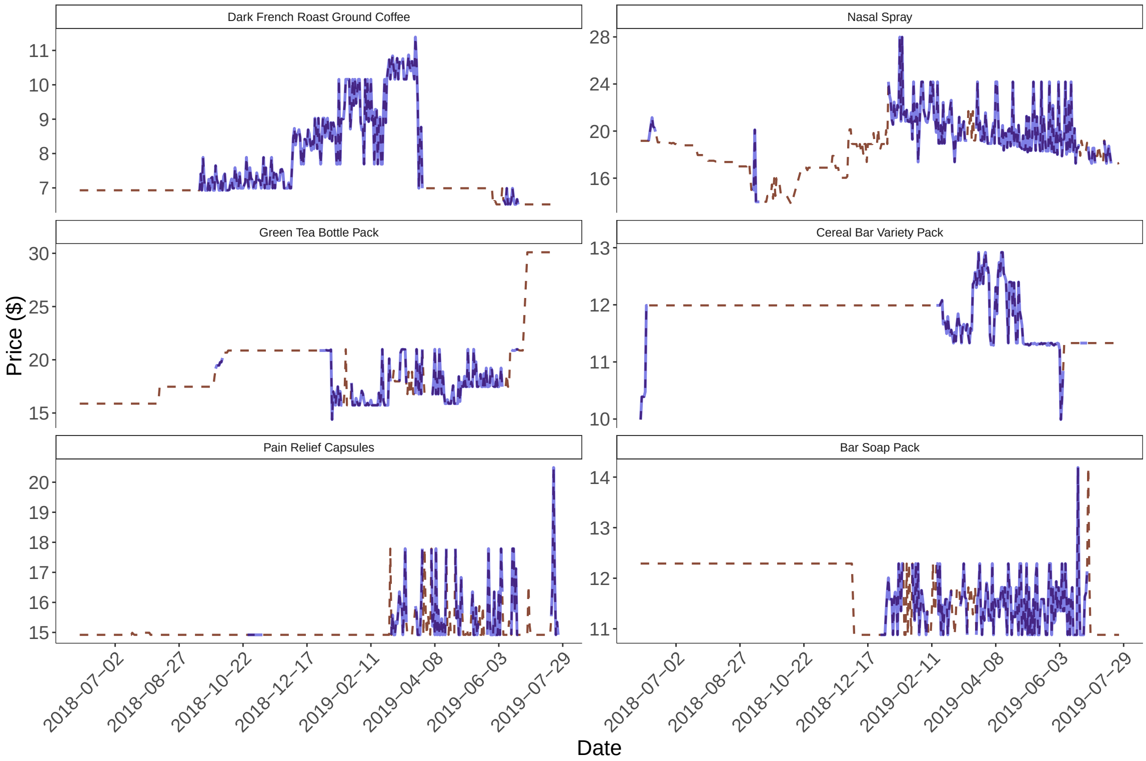

Figure 1: Examples of Price Variation: Daily Price with Periods of High Volatility Highlighted in Blue

Important exceptions in the area of consumer behavior are the studies of Haws and Bearden (2006) and Weisstein, Monroe, and Kukar-Kinney (2013), which show that unusual price differences may evoke feelings of unfairness and thereby reduce willingness-to-pay.[1] These findings are obtained in the context of laboratory experiments, and therefore highlight the need to understand more "realistic shopping environments and under conditions of higher involvement" (Haws and Bearden, 2006). Although not in the context of algorithmic pricing, prior studies have shown that deep price promotions can trigger customer antagonism or incentivize promotion-seeking behaviors (Mela, Gupta, and Lehmann, 1997; Anderson and Simester, 2004; Hendel and Nevo, 2004; Rotemberg, 2005; Anderson and Simester, 2010; Elberg, Gardete, Macera, and Noton, 2019). Reflecting upon this collection of evidence, it is reasonable to assume that consumer behavior will not be indifferent to algorithmic pricing; but it is not immediately clear in which direction.

We begin with a conceptual discussion of how algorithmic pricing affects consumers' price sensitivity in Section 2. The conceptual framework discusses two conflicting behavioral components. On the one hand, algorithmic pricing heightens price salience. We connect to Chetty, Looney, and Kroft (2009); Bordalo, Gennaioli, and Shleifer (2013, 2020), who study consumer choice in the context of boundedly rational consumers and salience effects. Relatedly, Finkelstein (2009); Busse, Lacetera, Pope, Silva-Risso, and Sydnor (2013); Hastings and Shapiro (2013); Busse, Pope, Pope, and Silva-Risso (2015); Aparicio and Rigobon (2020); Blake, Moshary,

[1]Some studies use the term "dynamic pricing". Throughout this work, we maintain the algorithmic pricing or machine pricing terminology.

Sweeney, and Tadelis (2021) find empirical support to the role of salience in offline and online markets. On the other hand, algorithmic pricing obfuscates the price anchor, namely "jams" the signal of good or bad deals. A number of conceptual and behavioral articles have examined limited price recall and constrained attention across attributes, such as Monroe (1973); Dickson and Sawyer (1990a); Lichtenstein, Ridgway, and Netemeyer (1993); Thomas, Simon, and Kadiyali (2010); Caplin and Dean (2015); Jung, Kim, Matejka, Sims, et al. (2019). The novelty of our work lies in conceptualizing the ambiguous effect of algorithmic pricing through the lens of consumer behavior; and more precisely, the connection of price sensitivity to the role of price salience and price anchors—two critical behavioral findings that are often treated separately.

With these ideas in mind, we proceed to study algorithmic pricing in a field setting. We collaborate with an online retailer in the United States that implemented algorithmic pricing. Overall, the data covers a subset of 2,044 products and more than 670,000 distinct consumers for 15 months. Critically, the clickstream dataset covers both search and purchases, allowing tracking patterns of visitation and purchases across and within users. This is important because, intuitively, we can exploit the fact that distinct consumers browsed and purchased the same product, but they had different exposure to prices (and price volatility).

Three core empirical strategies are used to estimate the effect of price volatility on price sensitivity. First, we build intuition by estimating aggregate models at the weekly-UPC level, allowing us to obtain comparable estimates to frequently used scanner data. We find own-price elasticities that are qualitatively similar to those in prior work (Anderson and Simester, 2008; Hitsch, Hortacsu, and Lin, 2019; DellaVigna and Gentzkow, 2019; Semenova, Goldman, Chernozhukov, and Taddy, 2017). Second, we consider a before-and-after event study around the time of adoption of algorithmic pricing in each product. Similar to Assad, Clark, Ershov, and Xu (2020), adoption dates can be recovered by observing unusual spikes in price volatility. Importantly, this experimentation varied across products and across categories over time, and is presumably exogenous to a customer's decision to visit the platform. Section 4 presents these results. Finally, in Section 5, we build upon these motivating signs to estimate a more stringent model of exposure to price volatility at the user level. We use two identification strategies - an instrumental variables approach and randomization inference to pin down the causal effect of frequently changing prices, and hence exposure to multiple unique prices *for the same product*, on purchase behavior. We find a consistent pattern throughout: price volatility makes demand more price sensitive.

We are mindful that, despite the granular clickstream dataset, it is not possible to exert complete control over a large-scale and long time-span field setting. Therefore, we conduct laboratory experiments to test the key effect of algorithmic pricing in a controlled environment. We implement a between-subject design in which participants are randomly assigned to two treatment conditions: stable pricing and algorithmic pricing. Participants in each cell are asked to simulate online purchases over a set of periods, and the price fluctuates from period to period. Importantly, the price series are calibrated with the real data, i.e. the stable prices and algorithmic prices mimic the prices of the online platform. Once again, and most importantly, participants are more price sensitive when exposed to higher price volatility. The experiment is implemented in two different subject pools—Amazon Mechanical Turk and MBA students. Section 6 presents the details.

Returning to the conceptual framework, we provide evidence that price *salience* is a key behavioral mechanism through which sensitivity to price volatility operates. We motivate a behavioral model with field data, obtained from a technology company, supporting the intuition that prices capture additional "bits" of attention (Jacoby (1984)). Using eye-tracking technology installed in digital screens (placed in physical stores), we find that showing prices captures more attention, compared to signage without prices. We interpret this evidence, albeit secondary, as a valuable step in the direction of price salience: if prices were to actually *change* on the screens, the attention (salience) would very likely be much greater. We then formally test for salience effects in the lab experiment. More precisely, following a vast tradition in the literature (Alba and Chattopadhyay (1986); Kissler, Herbert, Peyk, and Junghofer (2007); Finkelstein (2009); Kroft, Lange, and Notowidigdo (2013); Gaspelin, Leonard, and Luck (2015)), a series of recall questions in the lab indicate that price volatility exacerbates price salience. While we emphasize the role of salience, we make no claim that it precludes other processes to operate as well—an interesting question for future research.

Taken together, these findings shed light on the non-obvious side effects of algorithmic pricing. Consumers are not indifferent to price volatility: it modifies the shopping behavior and increases price sensitivity. Intuitively, a more price-sensitive demand "eats" some of the benefits that presumably could have been extracted had price elasticity remained unaffected by the extreme price fluctuation. While beyond the scope of our work, in theory, this side effect could be utilized as "input" to perfect the technology. Said differently, it speaks to the possibility of *personalizing* algorithmic pricing to mitigate behavioral reactions. For example, suppose that the machine determines that the optimal price is \$3.17. Moreover, suppose that a given consumer has recently visited that product twice, and on those occasions, the price was \$3.09 and \$3.99. Perhaps it is better, for this consumer, to coarsen the price to the already-seen \$3.09, rather than to show \$3.17, a new price for that consumer. In statistical parlance, one can think of this as adding a penalty or regularization term on the number of price changes the algorithm is allowed to make. The new price may be a "better" price but the trade-off needs to be judged after netting out the negative impact of increased price sensitivity.

2 Conceptual Framework

In a standard friction-less shopping process, a representative consumer decides whether to buy a single product based on two attributes, namely the quality and the price (Lilien, Kotler, and Moorthy (1995)). The decision is often summarized as $v = q - p$, and the outside option $v = 0$ implies no purchase. Algorithmic pricing alters the role of the price attribute, in ways that connect to various mechanisms studied in the literature.

Salience: Increasing the frequency of price changes makes the price attribute more salient, relative to the attributes that remain static (brand, package, features, etc.). The shift in relative salience can be thought of as changing the decision weights between a product's price and value (Bertini and Wathieu, 2008; Aparicio and Rigobon, 2020; Bordalo, Gennaioli, and Shleifer, 2020; Blake, Moshary, Sweeney, and Tadelis, 2021). Salience shifts irrespective of the sign of the price change, although the effect may be exacerbated when prices increase. In fact, Rotemberg (2005) and Anderson and Simester (2010) have shown that consumers develop antagonism when they realize that they have paid a higher price. That price variation might attract attention to prices can be indirectly related to evidence of how salient, visual attributes are over-weighted in the decision (Krider, Raghubir, and Krishna (2001); Folkes and Matta (2004)).

Signal to Noise: Consumers retrieve (or form) an anchor or reference price and compare it with the current price. Abundant research has explored how the price anchor is formed and the extent to which it can be manipulated by various forms of price presentation strategies (Kalyanaram and Winer (1995); Anderson and Simester (2003); Amaldoss and He (2018)). Typically, the price anchor is formed and updated upon exposure to past prices of the same product, advertised prices, or reference products in the category (Vanhuele and Drèze (2002); Jindal and Aribarg (2021); André, Reinholtz, and De Langhe (2021)). Echoing prior studies showing limited price recall (Monroe (1973); Dickson and Sawyer (1990a); Lichtenstein, Ridgway, and Netemeyer (1993)), algorithmic pricing exposes consumers to a complicated price path, often iterating between many distinct prices. This unstable path makes the price anchor noisier. Algorithmic pricing obfuscates the price anchor and, thereby, reduces sensitivity to notions of good or bad deals. Returning to the example in the Introduction: it is not obvious what the typical price for the carbonated cola should be.

This conceptual discussion captures two critical conflicting effects of algorithmic pricing. On the one hand, it increases price sensitivity by making the price a more salient element in the decision. The salience occurs as a result of shifting the relative variation between product attributes. Note that, interestingly, price sensitivity may be exacerbated even for consumers for whom the willingness-to-pay is greater than the actual price, i.e. an unnecessary side effect given that those consumers would have purchased this product regardless. On the other hand, it decreases price sensitivity by obfuscating the anchor price. Constant iterations between prices make the anchor price noisier (it "jams" the signal), thereby mitigating the reaction to a better or worse deal. We return to testable implications of this model in the context of lab experiments in Section 6.

To motivate the discussion that follows, we provide empirical evidence of price salience using novel experiment data from brick-and-mortar retailers. Digital screens are often placed in physical stores (e.g., supermarkets, gas stations, fashion stores) to advertise selected products of the assortment. We collaborate with a European marketing analytics company which manages the content of these campaigns with its partner retailers. Throughout a period of approximately two months, the company placed regular advertisements on those digital screens; in some cases with prices and in some other cases without prices. Importantly, the screen is equipped with an eye-tracking technology that records consumer-level eye views and time spent viewing.

While in practice it is infeasible to capture a price *change* (i.e., the price is constant in the digital screens), the eye-tracking sensor allows testing whether prices increase attention. Our empirical strategy resembles prior work in which salience of an attribute entails attention to that attribute (Duncan, 1984; Folkes and Matta, 2004), and time is used as a measure of attention (Townsend and Kahn, 2014; Cian, Krishna, and Elder, 2015).

We test whether showing prices in the screens captures additional signage attention (controlling for the number of eye-views). In total, the data includes 3,570,646 distinct eye-views and 42 digital campaigns throughout two months. The average per-person view time of a screen, conditional on viewing, is approximately 7 seconds. Let v_{it} be the total number of views to screen i on day t and let t_t be the total time spent viewing screen i on day t. Time is measured in milliseconds. The measure of interest is $\tau \equiv \frac{t_{it}}{v_{it}}$, i.e. the time spent per eye view. We then estimate the following model:

$$\tau_{it} = \beta_0 + \beta_1 PriceDisplayed_i + \delta_t + \gamma_s + \epsilon_{it} \tag{1}$$

where $PriceDisplayed_i$ is an indicator variable that takes value 1 when the screen i contains a price (and 0 otherwise); and δ_t and γ_s denote day- and store- fixed effects, respectively.

Table 1: Eye-Tracking and Price Salience

	Attention Time
Price Displayed	226.68***
	(33.03)
Fixed-effects	
Store	✓
Day	✓
Fit statistics	
Obs.	139,978
R^2	0.15

Table 1 shows the results. When the digital screens display prices, time spent viewing the screen increases by 227 milliseconds ($p < 0.01$). This evidence supports the idea that price is a product feature prone to be salient and thereby to capture additional cognitive attention. Furthermore, it complements the implications of algorithmic pricing. That is, because prices tend to capture attention, high-frequency price variation would presumably capture even more “bits” of attention and thereby heighten the role of price salience. Further research, perhaps in a laboratory setting with the availability of fMRI technology (like Karmarkar, Shiv, and Knutson (2015)), is needed to better examine the behavioral decision-making process. For our purpose, we find this evidence motivating to more keenly study how heightened price volatility, caused by algorithmic pricing, makes price more salient and hence may influence consumer price sensitivity.

3 Data and Empirical Setting

We use data from an online retailer in the United States to examine the scope and implications of algorithmic pricing. Throughout the relevant time period, the retailer tried out algorithmic pricing for thousands of products across a wide range of categories, departments, and price levels. This empirical setting is particularly well-suited to studying behavior in response to algorithmic pricing for two reasons. First, the data includes clickstream records at the user level, which allows us to observe the entire sequence of the click activity (e.g., image impressions, search queries, product views, add-to-carts, orders placed). Moreover, the online groceries context involves repeated purchases across users and products, as well as a relatively large assortment breadth.

Table 2 reports summary statistics on the data. We focus on a subset of products that experienced algorithmic pricing and, additionally, a minimum threshold of purchase records. Overall, the data covers 2,044 distinct products across groceries, household supplies, baby products, health and beauty, and pet supplies. The data covers 15 months, 673,677 distinct consumers, and over 2.6 million units sold.

Table 2: Data Description

		Summary Statistics
(1)	Distinct consumers	673,677
(2)	Categories	Household Supplies, Baby, Health & Beauty, Grocery, Pet Supplies
(3)	Distinct products	2,044
(4)	Time period	15 months
(5)	Units sold	$2,659,906$
(6)	Algorithmic periods	32%

There is no single definition of algorithmic pricing. In this work, we focus on one dimension of algorithmic pricing—a greater frequency of price changes. This echoes a growing literature that identifies price volatility as one of its most distinctive features (Chen, Mislove, and Wilson (2016); Brown and MacKay (2019); Assad, Clark, Ershov, and Xu (2020)). We operationalize this concept in the field data as follows. For each product and week pair, we compute the sum of absolute price changes and the number of unique prices; if the first measure in any given week is greater than its respective median values across all weeks, *and* if the second one is greater than three, then we classify that week as an algorithmic pricing period.[2] Return to Figure 1 for some visual examples using this definition. The pattern of results remains quantitatively similar under alternative thresholds and definitions, e.g. based on the standard deviation of prices, number of price changes, or absolute size of price changes. Robustness results are presented in Appendix G.

Summary statistics split by algorithmic pricing and stable pricing periods are shown in

[2]An important digression is helpful. Conceptually, it is possible that the output of the algorithm is to set a flat price, e.g., collusion between two rival firms (Miklós-Thal and Tucker (2019); Calvano, Calzolari, Denicolò, and Pastorello (2019)). This definition essentially implies that there the price changed frequently and substantially within a week. At least in our field setting, institutional knowledge strongly indicates that periods in which the price fluctuates intensively are driven by the implementation or experimentation of price algorithms (e.g., price matching, a grid of mark-up rules, inventory triggers).

Appendix A. Furthermore, a variance decomposition test, shown in Appendix B, indicates that the product-week indicator of algorithmic pricing significantly explains a large portion of the price variation.

4 Aggregate Purchase Behavior

We proceed with a series of models with data aggregated at the product-week level, which allows us to more directly contrast our demand estimates with those using scanner data (typically at the same aggregation level). Later in Section 5 we consider consumer-level exposure to price volatility. Before we proceed, it would help to fix notation. Users are indexed with $i = \{1, \ldots, I\}$, products with $j = \{1, \ldots, J\}$, product categories with $c = \{1, \ldots, C\}$, and time with $t = \{1, \ldots, T\}$. In much of our analysis cases, t is year-week unless specified otherwise. Y is the number of units purchased and P is price.

We initially consider a reduced form demand model similar to Hitsch, Hortacsu, and Lin (2019). We aggregate purchases at the weekly level, compute the unit-weighted price, and estimate the following baseline fixed-effects model:

$$\log(Y_{jt}) = \beta_0 + \beta_1 \log(P_{jt}) + \mu_j + \tau_t + \epsilon_{jt} \tag{2}$$

where Y_{jt} is the number of units sold for product j in week t and P_{jt} is the quantity-weighted price for product j at time t. μ & τ are product and time fixed effects respectively.

To understand the potential impact of algorithmic pricing on consumer behavior, we augment Equation 2 by including indicators for weeks during which the price for a product was highly volatile, as per the definition in the previous section. We simply call these weeks "algorithmic pricing weeks". Further, we interact these indicators with (log) price to test whether these high price-volatility periods influence consumer price sensitivity. The updated model is:

$$\log(Y_{jt}) = \beta_0 + \beta_1 \log(P_{jt}) + \beta_2 A_{jt} + \boldsymbol{\beta_3} \log(P_{jt}) \times A_{jt} + \mu_j + \tau_t + \epsilon_{jt} \tag{3}$$

where A_{jt} is a binary indicator that equals 1 if product j is under an algorithmic pricing week during year-week t.

The results are shown in Table 3. Our focus is on β_3; the interaction between price and the algorithmic pricing indicator. A negative β_3 indicates that demand is more price sensitive when exposed to high-frequency price variation. Consider the baseline own-price elasticity of -1.51 in column (2). In periods of algorithmic pricing, the price sensitivity increases 0.087, which represents a sizable 5.8% in relative terms.

If algorithmic pricing makes the demand more price-sensitive, one might imagine that a greater intensity of price volatility exacerbates price sensitivity. Indeed, we find that as we require a higher number of distinct prices in a given week to be classified as an algorithmic pricing week, the estimand of interest becomes more price-sensitive. For example, the effect increases from approximately -0.07 to -0.10 when the threshold of distinct prices increases from two to five, as shown in Appendix C.

To provide visual intuition for what these results imply, we draw the demand curves for three product categories – dog supplies, chips, and fresh produce. We estimate the demand separately during algorithmic and stable pricing weeks after residualizing the quantity and price. The results are shown in Figure 2. For all three categories, we find the demand curve becomes flatter, i.e., the demand becomes more price sensitive. Furthermore, the rotation in the demand curve varies across the categories indicating potential heterogeneity. For example, the demand for dog supplies and chips changes much more than the demand for fresh produce. We explore this heterogeneity later in the section.

Next, in column (3) of Table 3, we remove the holiday period (mid-November to mid-January), a time when retailers typically run multiple promotions, and re-estimate model 3. Here again, we find that even outside the holiday period, the effect is strong. We believe that this evidence is suggestive of important changes in consumer behavior as a result of the firm's pricing policy. We explore this hypothesis more in later sections and use consumer-level data to causally estimate the impact. In the remainder of this section, we provide more evidence for the aggregate result using multiple specifications.

First, as in Anderson and Simester (2008), we estimate a quasi-Poisson demand model as follows:

$$\mathbb{P}(Y_{jt} = y) = \frac{e^{-\lambda_{jt}}\lambda_{jt}^{q}}{q!}, \quad q = 0, 1, 2, \ldots \tag{4}$$

$$\log(\lambda_{jt}) = \beta_0 + \beta_1 \log(P_{jt}) + \beta_2 A_{jt} + \boldsymbol{\beta_3} \log(P_{jt}) \times A_{jt} + \mu_j + \tau_t + \epsilon_{jt} \tag{5}$$

The results are presented in column (4) in Table 3. Overall, we find very close estimates to those of the baseline model, reported in column (2).

Additionally, we consider recent methods in machine learning. We implement Semenova, Goldman, Chernozhukov, and Taddy (2017)'s approach of estimating own-price elasticities using orthogonal-machine learning. The methodology allows us to include a high-dimensional set of features as controls, including lagged values for price and purchases. The model is estimated in two-stages. In the first stage, we residualize the outcome (purchases) and the independent variable of interest (price) using lagged values of purchases, price, indicator for algorithmic pricing week, indicators for product, product category, and time. In the second stage, we regress the residualized outcome on the residualized price, the indicator for algorithmic pricing week, and the interaction of the two. To ensure unbiased estimates, regressions in the two stages are estimated on different sub-samples. More formally, the model is defined as follows:

Table 3: Aggregate Elasticity Estimates with Multiple Specifications

Dependent Variables:	Log units			Units	Log units	
	Gaussian	Gaussian		Poisson	Ortho ML	Mixed effects
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Elasticity	-1.525***	-1.508***	-1.545***	-1.558***	-1.052***	-1.424***
	(0.100)	(0.101)	(0.101)	(0.135)	(0.109)	(0.108)
Algo		0.260***	0.263***	0.232***	0.032***	0.285***
		(0.036)	(0.041)	(0.040)	(0.011)	(0.034)
Elasticity × Algo		-0.087***	-0.088***	-0.081***	-0.150**	-0.103***
		(0.015)	(0.017)	(0.018)	(0.062)	(0.017)
Fixed-effects						
Product	Yes	Yes	Yes	Yes	Yes	Yes
Year week	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	122,309	122,309	103,954	122,309	59,157	122,309
R^2	0.558	0.559	0.558	-	0.399	-
Log-Likelihood	-112,205	-111,987	-95,450	-591,338	-46,113	-116,096

Two-way (Product & Year week) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

First stage, estimated on sample S:

$$\log(Y_{jt}) = \beta_{0y} + g_y(\beta_y A_{jt},\ \delta_y \sum_{t-1}^{t-4} \log(Y_{jt}),\ \gamma_y \sum_{t-1}^{t-4} \log(P_{jt}),\ \mu_j,\ \tau_t) + \epsilon'_{jt} \tag{6}$$

$$\log(P_{jt}) = \beta_{0p} + g_p(\beta_p A_{jt},\ \delta_p \sum_{t-1}^{t-4} \log(Y_{jt}),\ \gamma_p \sum_{t-1}^{t-4} \log(P_{jt}),\ \mu_j,\ \tau_t) + \epsilon''_{jt} \tag{7}$$

Second stage, estimated on sample S' ($S \cap S' = \phi$):

$$\log(\tilde{Y}_{jt}) = \beta_0 + \beta_1 \log(\tilde{P}_{jt}) + \beta_2 A_{jt} + \beta_3 \log(\tilde{P}_{jt}) \times A_{jt} + \epsilon_{jt} \tag{8}$$

where $\tilde{Y}_{jt}$ are the residuals from Model 6 and $\tilde{P}_{jt}$ are the residuals from Model 7. $g_y()$ & $g_p()$ are functions that control for lagged features, product, category, and time effects. In our case, we use penalized l-1 regressions. Column (5) in Table 3 then shows the result from Equation 8. We again see that the results lead to the same conclusion as our baseline model, with the coefficient on the interaction term being negative.

4.1 Heterogeneity Across Product Categories

To unpack the heterogeneity across product categories, we estimate a hierarchical mixed-effects model (Gelman and Hill, 2006). Previous literature has used similar random effects models to study price elasticities and online consumer behavior (e.g. Hoch, Kim, Montgomery, and Rossi, 1995). Mixed-effects models allow for partial pooling of information across products and categories. In our case, we use them to efficiently estimate product-level elasticities and conduct sub-

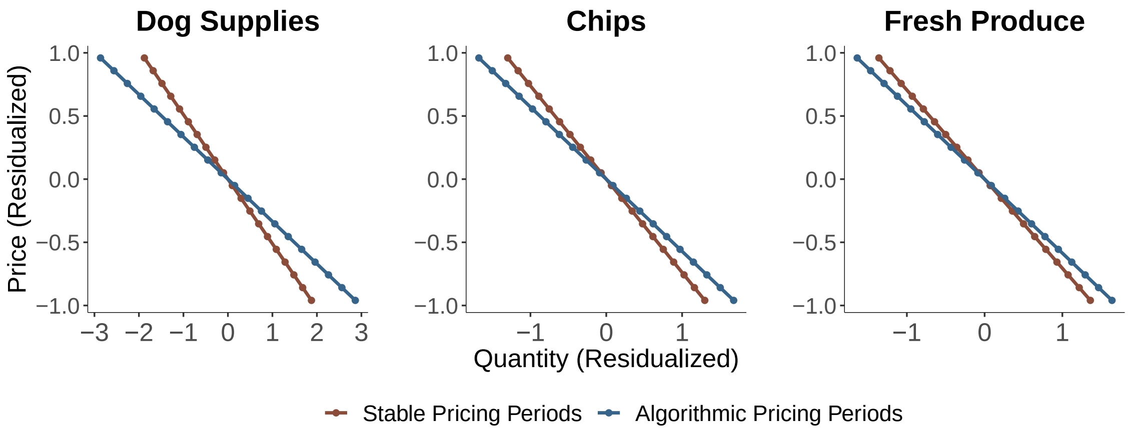

Figure 2: Price-Sensitive Demand Rotation

Notes: The graphs show demand curves across three categories after taking out the effects of product and time from log quantity and log price. The red demand curve is estimated separately by only considering periods of stable pricing. Analogously, the blue demand curve is estimated separately for periods of algorithmic pricing.

group analysis. We use a nested hierarchical approach where we allow the intercept and slopes to vary by category and by each product within that category; we also allow the intercept to vary over time. The model is estimated using Restricted Maximum Likelihood. Varying the price elasticities with categories and products allows pooling of information across levels and regularizes coefficients in a data-driven way. We estimate the following model:

$$\log(Y_{jt}) = \beta_{0cjt} + \beta_{1cj}\log(P_{jt}) + \beta_{2cj}A_{jt} + \beta_{3cj}\log(P_{jt}) \times A_{jt} + \epsilon_{jt} \quad (9)$$

where β_{0cjt} is the intercept that is allowed to vary by category, product, and year-week, and all three slope coefficients $\beta_{1cj}, \beta_{2cj}, \beta_{3cj}$ are allowed to flexibly vary by category and by product within a category.

Results from the mixed-effects model are shown in column (6) in Table 3. Once again, the results are qualitatively similar to the previous models. Reassuringly, the estimated own-price elasticities are qualitatively similar to recent studies using grocery data (Hitsch, Hortacsu, and Lin, 2019; Semenova, Goldman, Chernozhukov, and Taddy, 2017). See the product-level distribution of own-price elasticities in Appendix D. In Figure D.2, we show the distribution of elasticities during algorithmic pricing and stable pricing periods, as depicted. We see that during periods of algorithmic pricing, there is a significant shift in greater price sensitivity across most products.

As a motivation to understand the value of these flexible models, multilevel analysis of variance (ANOVA) supports the inclusion of varying intercept and slope parameters. The results in Appendix B show that varying slopes explain a significant portion of the purchase variation.

The mixed-effects models allow us to take a step further in decomposing the results across products categories. Figure 3 shows the percentage change in price elasticity during algorithmic pricing weeks split by product category. For simplicity, we visualize 20 categories, 10 with the smallest change in elasticity and 10 with the largest change in elasticity. The red dashed line is the global average across all categories. Overall, and interestingly, we observe some but not fundamental heterogeneity across products. The biggest change is seen in stockable snacks, cleaning products, and pet supplies. On the other hand, health and beauty products such as skin care, hair care, and digestion & nausea see the smallest changes. In Appendix E we test for heterogeneity using different sub-groups such as expensive and cheap products, popular and unpopular products, and perishable and non-perishable products.

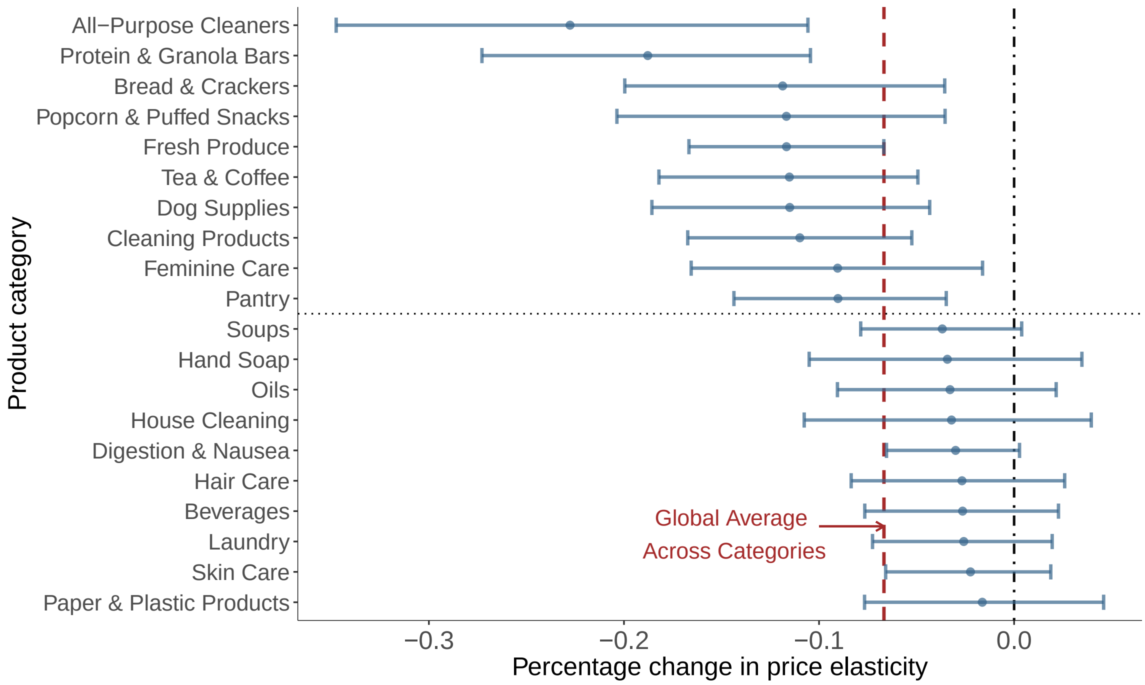

Figure 3: Top-10 and Bottom-10 Product Categories Based on Percentage Change in Price Elasticity During Algorithmic Pricing Weeks

Our set of analyses serves two purposes – 1) it facilitates comparison with previous research using scanner data and helps establish common ground with existing literature, 2) it provides motivation to explore implications of algorithmic pricing at a more granular level. The results from Table 3, while significant and robust to specifications, do not allow causal identification of the impact of algorithmic pricing on consumer behavior. Other than the potential impact of algorithmic pricing, current estimates could also be a reflection of either common demand shocks or compositional changes in demand, or a mix of the three. We conduct a more granular analysis with consumer visit-level data in the coming sections to identify this effect.

4.2 Before-and-After

We build upon the aggregate analysis with an event-study approach exploiting the first time algorithmic pricing was adopted in each product. We identify the date of the first algorithmic pricing week as per Section 3 and conduct an event study before-and-after the adoption. Importantly, we observe significant variation across products, i.e. different products had their first experimentation of algorithmic pricing in different dates. The same is true across product categories.

(Appendix F presents the distribution of the timestamps of the first events.)

We report two main specifications: one in which we pool observations at the product-week level, and a second in which we utilize observations at the user x product x week level. Appendix F reports robustness specifications. First, we estimate the following fixed-effects model:

$$\log(Y_{jt}) = \beta_0 + \beta_1 \log(P_{jt}) + \beta_2 Post_{jt} + \beta_3 \log(P_{jt}) \times Post_{jt} + \mu_j + \tau_t + \epsilon_{jt} \tag{10}$$

where $Post_{jt}$ is a binary indicator that takes the value 1 if $t' < t < t' + 8$ weeks for product j and t' is the first week when product j is under algorithmic pricing. To estimate the models, we use a window ± 8 weeks around the first week when a product adopted algorithmic pricing. However, to ensure robustness of our results, we consider three window cut-offs, i.e. adoption after 12 weeks of our sample start date, adoption after 20 weeks, and adoption after 28 weeks.

The results for products who adopted 20 weeks or after are shown in the first column of Table 4. Our main coefficient of interest is the one on the interaction between log price and the post-period indicator. We find that once the product switches to an algorithmic pricing regime, sensitivity to price changes increases.

The above result provides suggestive evidence towards increased price sensitivity. However, these results could partly be driven by compositional changes in the underlying user population. The granular scope of the data allows us to exclude this explanation by examining the same set of users who visit the product before and after the adoption of algorithmic pricing. We do this by using a different specification in which we estimate price sensitivity before-and-after the algorithmic pricing adoption at the user level. We use the same cut-offs as the aggregate model; however, we only consider users who browsed the product in both periods (i.e., before adoption and after adoption). More formally, we estimate the following fixed-effects model:

$$\log(Y_{ijt}) = \beta_0 + \beta_1 \log(P_{ijt}) + \beta_2 Post_{jt} + \beta_3 \log(P_{ijt}) \times Post_{jt} + \delta_{ij} + \epsilon_{ijt} \tag{11}$$

where Y_{ijt} is the number of units of product j purchased by user i at time t and P_{ijt} is the average price for product j seen by consumer i during time-period t. Note that here there are only two time-period observations per user-product pair, i.e. one before the product adopts algorithmic pricing and the second after the product adopts algorithmic pricing. Importantly, this allows to control for user-product fixed effects (δ_{ij}). We are interested in β_3, the coefficient on the interaction of average price and the indicator for the post-period. Similar to Equation 10, $Post_{jt}$ takes the value 1 if $t' < t < t' + 8$ weeks, where t' is the first week when product j is under algorithmic pricing. As in the aggregate case, to ensure the robustness of our results, we consider three window cut-offs for adoption – 12 weeks, 20 weeks, and 28 weeks.

The results are shown in the second column of Table 4. Since we control for user-product fixed effects, i.e., we estimate the coefficient using variation only within the same user-product pair. We see that after products adopt algorithmic pricing, consumers become more price sensitive. We posit that this effect is primarily driven by repeated exposure to different prices for the same product, which makes price more salient, making consumers put more weight on it during the purchase decisions. We explore this hypothesis in subsequent sections.

Table 4: Aggregate and User-Level Price Sensitivity Before and After Adoption of Algorithmic Pricing

Dependent Variables:	Log units	Units
Model:	(1)	(2)
Variables		
Log price	-1.03**	-1.12***
	(0.389)	(0.222)
Post period	0.238***	0.208***
	(0.084)	(0.058)
Log price × Post period	-0.101***	-0.072***
	(0.035)	(0.025)
Fixed-effects		
Product	Yes	
Year week	Yes	Yes
User-product		Yes
Fit statistics		
Observations	7,652	42,864
R^2	0.687	0.465
Log-Likelihood	-5,003	-57,195

Model (1): Two-way (Product & Year week) standard-errors
Model (2): Three-way (User, Product & Year week) standard-errors
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The first column is OLS regression of total weekly sales on average weekly price. Post period is a binary variable that takes the value 1 after a product adopts algorithmic pricing. Second column emulates the spirit of the first regression at the user-level. Each observation is at the user-product-period level, i.e., there is one observation for a user-product pair before algorithmic pricing adoption for that product and one observation after it. The dependent variable is the number of units purchased and price is average across all exposure for that user-product pair.

5 Consumer-Level Exposure to Price Volatility

The field data from the online retailer covers detailed browsing and shopping clicks at the consumer level. This granular data allows to estimate the effect of algorithmic pricing by exploiting consumer-level exposure to price variation while controlling for user-, product-, and time- fixed effects. Intuitively, consumers that had exposure to more distinct prices and a higher frequency of price changes, had a larger exposure to algorithmic pricing when shopping online groceries.

We introduce a model at the user level that bears resemblance to prior studies investigating the “stock” of advertising (Erdem, Keane, and Sun (2008); Shapiro, Hitsch, and Tuchman (2021)). In our case, we translate the model to one where, instead of advertising, we keep track of the stock of price volatility exposure. Said differently, the model accounts for spillovers of price volatility across products when accumulating the exposure over time. Conceptually thinking of algorithmic pricing through the lens of advertising and price sensitivity (Dorfman and Steiner (1954); Becker and Murphy (1993)) is helpful because it illustrates that algorithmic pricing cannot be reduced to a simple A/B test. Instead, its core effect must consider the accumulation of price volatility across products and over time.

With these ideas in mind, we define the following user-level model:

$$Y_{ijt} = f(P_{ijt}, A_{it}, X_{ijt}; \epsilon_{ijt}) \tag{12}$$

where Y_{ijt} is the number of units of product j purchased by user i at time t. A_{it} is the cumulative effect of algorithmic pricing that user i has accumulated till time t. It is the total number of unique prices that the consumer has seen over the past L days across all products that the consumer browsed. In our regressions, we use $L \in \{7, 15, 30, 60\}$ days to account for different intensities of exposure to algorithmic pricing. We refer to A_{it} as the algorithmic pricing stock or, succinctly, the algo-pricing stock. X_{ijt} are user history variables that account for the user's search and purchase intensity. We are interested in understanding how exposure to A_{ijt} modifies user behavior.

We use the following econometric specification for the model:

$$\begin{aligned}
\mathbb{P}(Y_{ijt} = y) &= \frac{e^{-\lambda_{ijt}} \lambda_{ijt}^{q}}{q!}, \quad q = 0, 1, 2, \ldots \\
\log(\lambda_{ijt}) &= \beta_0 + \beta_1 \log(P_{ijt}) + \beta_2 A_{it} + \beta_3 \log(P_{ijt}) \times A_{it} + \\
&\quad \delta X_{ijt} + \Gamma_i + \mu_j + \tau_t + \epsilon_{ijt}
\end{aligned} \tag{13}$$

Here again, we are interested in the coefficient β_3, which is the interaction between price and the algorithmic pricing stock. A negative value for β_3 indicates that consumers become more price sensitive after exposure to algorithmic pricing. In the regression, we measure the stock using the total number of unique prices that a consumer has been exposed to for the products that she visited *more than once*. As an example, say the consumer visited the product page for *Nutella* thrice in the past 15 days and saw two different prices. In addition, she visited the page for *Diet Coke* once and hence just saw one price for it. Her algorithmic pricing stock A_{it} is two. The price for *Diet Coke* is not counted since she only visited the product once. X_{ijt} are controls that account for the browsing intensity of the consumer. They include the total number of visits that the consumer made in the past L days, the total number of products browsed per visit, and the total number of purchases made. Γ_i are user-level fixed effects, μ_j are product fixed effects and τ_t are week fixed effects.

The motivation behind this model is to investigate whether consumers who are exposed to more unique prices for the same product, i.e. they have a larger A_{it}, tend to become more price sensitive. Arguably, one may worry that since A_{it} is not randomly assigned but rather dictated by the user's search process, the effect of exposure to algorithmic pricing on purchase behavior is not identified. For example, a given consumer who tends to search more may intrinsically be more price-sensitive and hence may repeatedly visit the retailer's website to fetch a good deal. As a consequence of their repeated visits, they naturally get exposed to different prices for the same product. Hence, the effect we estimate is just an artifact of browsing intensity and not necessarily a change in behavior.

Fortunately, the granular nature of clickstream data allows to finely control for time-varying browsing and purchase intensity of consumers. Estimation of the model then critically depends upon conditional variation in A_{it}. More specifically, we need variation in A_{it} conditional on the number of visits, i.e., we need users who visited the same product the same number of times but were exposed to a different number of prices. Figure 4 presents this variation. Panel (a) shows the marginal distribution of algorithmic pricing stock aggregated at the user-date level. Panel (b) shows the conditional distribution of algorithmic pricing stock aggregated at the user-product-date level. Each facet in Panel (b) conditions on the number of visits made by users for the same product. For example, the top-right facet of Panel (b) shows that users who visited a product five times in the past 30 days could have been exposed to anywhere between one and five unique prices. This variation allows us to control for the browsing intensity of users. To control for unobserved time-invariant user and product heterogeneity, we use user- and product- fixed effects.

Finally, to causally pin down the effect of algo-pricing stock A_{it}, we use two identification strategies – one based on instrumental variables and the second is based on randomization inference. Both strategies crucially depend upon the variation in the timing of user's visits to the website. Specifically, we assume that the exact time a user visits the retailer's website is as-good-as-random. Then, if she would have visited a few hours earlier or later, then she may have seen a different price for the products she visits. Consequently, this changes her exposure to algorithmic pricing, i.e., it changes the number of unique prices she ends up seeing. We explain both approaches in the sections below.

5.1 Instrumental Variables

To capture observed time-varying heterogeneity we include detailed user and user-product level controls such as the number of products searched, the number of total purchases made in the past, and the number of purchases for this particular product made in the past. Further, to causally pin down the effect of algo-pricing stock A_{it}, we exploit variation in the timing of user visits and calculate the number of unique prices the user *could have seen*, had she come at a different time, but *did not see*. This gives us an instrument for A_{it}. The intuition behind the instrument is that the purchase decision of a consumer naturally depends on the prices seen, but does not depend upon the prices not seen. However, prices for a particular product are correlated across time. Hence, the prices the consumer *did not see* cannot influence the outcome directly, except through their correlation with the prices she did see.

To make things concrete, consider two users A and B who both visited the retailer's website thrice in the past 15 days, albeit on different days or at different times on the same day. For simplicity, assume that both saw the same product three times. During the past 15 days, the price for this product was fluctuating independently of these two users' visits (because of competitor effects, inventory state, and/or aggregate demand). Because of the difference in timing of their visits, user A was exposed to only one unique price for the product whereas user B was exposed to three unique prices. The purchase decision that both the users are to make today depends upon the current price as well as the history of prices observed. However, it does not depend upon the

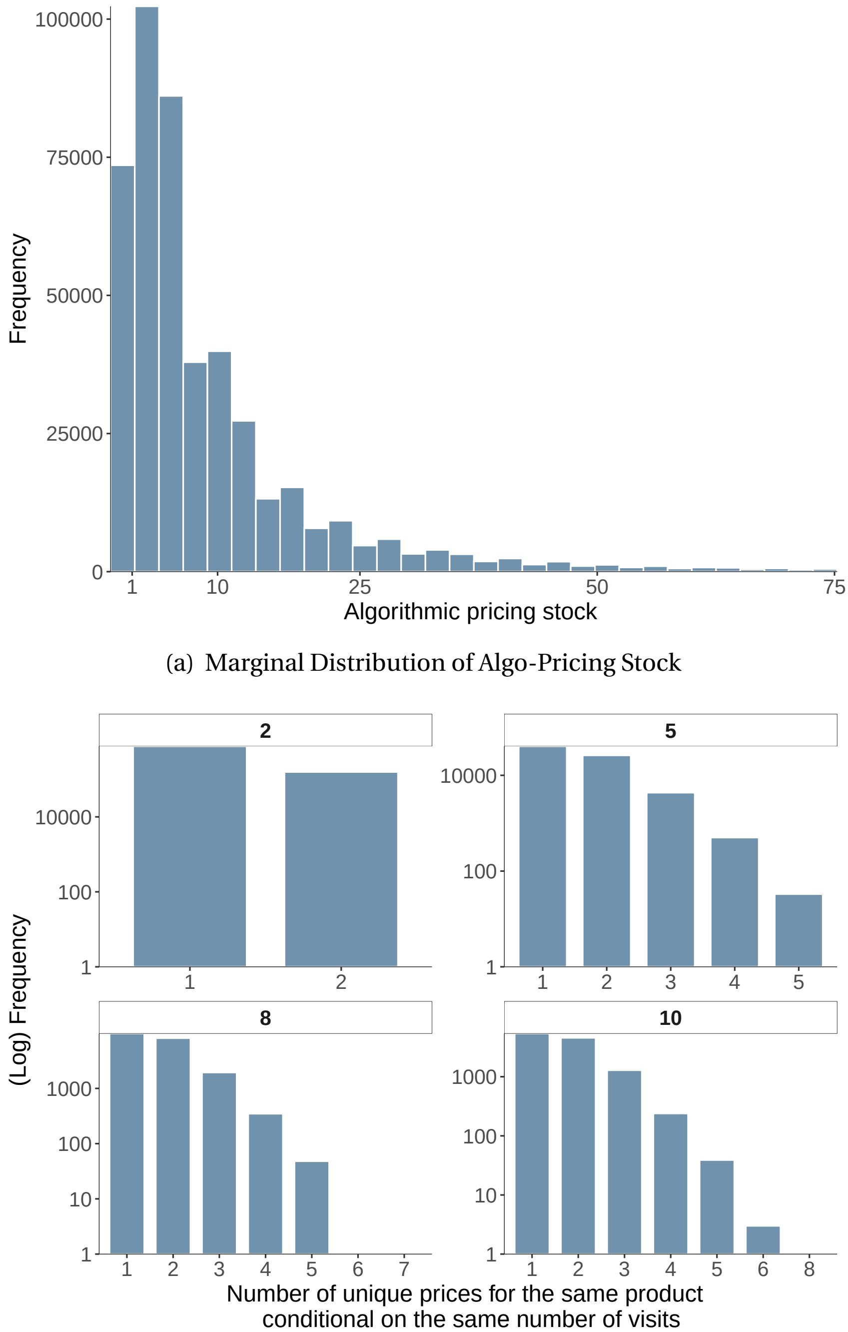

(a) Marginal Distribution of Algo-Pricing Stock

(b) Distribution of Algo-Pricing Stock Conditional on the Number of Visits. Each panel is a different number of visits for the same product.

Figure 4: Marginal and Conditional Distributions of Algorithmic Pricing Stock over a 30-day Period

prices not observed, except through their correlation with the observed prices. This makes the prices for the same product during the same time period that the consumer could have seen but did not see a valid instrument.

It is worth pointing out that this is a non-causal instrument. We don't observe exogenous shocks to prices at the system level. Rather we carefully isolate independent variation for each user based on their historical visit times. While not directly comparable, this instrument is similar in spirit to the one used by Assad, Clark, Ershov, and Xu (2020) and Ellison and Ellison (2009). Assad, Clark, Ershov, and Xu (2020) first identify station-level adoption of algorithmic pricing software in the gasoline markets using changes in high-frequency markers of prices[3] and then identify brand-level adoption using the proportion of the brand's stations who have adopted. Their instrument is non-causal as well and works on the assumption that brand-level adoption decisions are independent of local station-level shocks. Ellison and Ellison (2009) use prices of products from one category with prices from another category to estimate price elasticities for PC-RAM modules.

Table 5: User-Level Price Sensitivity Estimates using Two-Stage Control Functions

Dependent Variable:	Units			
	Baseline		Control Function	
Model:	(1)	(2)	(3)	(4)
Variables				
Log price	-0.98***	-0.96***	-0.92***	-0.84***
	(0.08)	(0.08)	(0.08)	(0.08)
Log price x Algo pricing stock	-0.06***	-0.07***	-0.09***	-0.14***
	(0.006)	(0.005)	(0.01)	(0.01)
Algo pricing stock	0.27***	0.22***	0.37***	0.77***
	(0.01)	(0.01)	(0.03)	(0.09)
# of sku per visit		-2.2***		-0.66**
		(0.06)		(0.33)
# of total visits		-0.11***		-0.43***
		(0.01)		(0.07)
# of prior purchases		-0.28***		-0.30***
		(0.008)		(0.009)
First stage residual - 1			-0.21***	-0.60***
			(0.03)	(0.09)
First stage residual - 2			0.04***	0.08***
			(0.01)	(0.01)
Fixed-effects				
User	Yes	Yes	Yes	Yes
Product	Yes	Yes	Yes	Yes
Year week	Yes	Yes	Yes	Yes
Fit statistics				
Observations	4,621,411	4,621,411	4,621,411	4,621,411
Pseudo R^2	0.220	0.229	0.220	0.229
Log-Likelihood	-1,864,529.7	-1,842,059.6	-1,864,059.6	-1,841,719.4

Two-way (User & Product) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table shows baseline estimates (columns 1 and 2) and two-stage control function estimates (columns 3 and 4) for the consumer-level model in Equation 13. Algorithmic pricing stock and time-varying history variables are calculated over a 30-day period. The dependent variable in all models is the number of units of a particular product, purchased by a user during a single visit. All models control for user, product, and year-week fixed effects. Standard errors for columns 3 & 4 are estimated using clustered bootstrap to account for first-stage estimation error.

[3] We follow a similar idea in our before-and-after analysis in Section 4.2.

For estimation, we use a two-stage control function approach as described in Petrin and Train (2010) where we instrument for the algo-pricing stock (A_{it}) and its interaction with price ($\log(P_{ijt}) \times A_{it}$). In the first stage, we run two regressions – 1) A_{ij} on the two instruments, exogenous controls (X_{ijt}) and the fixed effects from Equation 13, and 2) $\log(P_{ijt}) \times A_{it}$ on both instruments, exogenous controls, and fixed effects. In the second stage, we run the regression from Equation 13 by including the residuals from the two first-stage regressions. This control function is equivalent to running a TSLS procedure for instrumental variables when the outcome is linear. In our case, since our model for user purchases is a Poisson regression, we use the control function approach instead. Finally, to take into account the uncertainty from the first-stage estimation, we use clustered bootstrap to estimate standard errors.

The results are shown in Table 5. The first two columns show the baseline model where we don't use the control functions. After controlling for the user's browsing and purchase intensity, plus the user, product, and year-week fixed effects, we find that consumers who were exposed to more unique prices do become more price sensitive. In columns 3 and 4, we use a control function approach which accounts for potential endogeneity in algorithmic pricing stock. We find that, after correcting for endogeneity, the effect of algorithmic pricing stock almost doubles in absolute value. These results provide direct evidence of consumers becoming more sensitive due to heightened volatility in prices caused by algorithmic pricing. In all models, the algorithmic pricing stock and user history variables are calculated over a 30-day period. In the appendix, we provide robustness checks where user history is calculated over 7, 15, and 60-day periods. Furthermore, we also test for the robustness of functional form in Equation 13 by running vanilla OLS and correcting for endogeneity using two-stage least squares. Across all specifications, we unanimously find that exposure to more unique prices for the same product makes consumers more price sensitive.

5.1.1 Heterogeneity by Consumer Type

We investigate how the change in sensitivity varies by consumer type. We use two measures of consumer heterogeneity – historical purchases and tenure. For the first one, we calculate the number of total purchases made by a user in the past 60 days and split the consumer base at the median. Similarly, we calculate the tenure of the user on the retailer's platform from the date of the user's first visit and split at the median. We then re-estimate Model 13 separately for each sub-group using the two-stage control function approach described above.

The results are shown in Table 6. Columns (1) and (2) show the estimates for high-value and low-value customers, as defined by historical purchases. Overall, we find the effect of exposure to algorithmic pricing on price sensitivity to be strong and negative, i.e., both "high-value" and "low-value" consumers become more price sensitive. However, as a proportion of their baseline elasticity, high-value consumers, on average, experience twice as large of a change in price elasticity as compared to low-value consumers. This comparatively larger change also holds for consumers with a longer tenure with the retailer.

Table 6: User Level Sub-Group Analysis using Two-Stage Control Functions

Dependent Variable:	Units			
	Purchases >= Median	Purchases < Median	Tenure >= Median	Tenure < Median
Model:	(1)	(2)	(3)	(4)
Variables				
Log price	-0.657***	-1.02***	-0.686***	-0.984***
	(0.090)	(0.121)	(0.086)	(0.090)
Log price x Algo pricing stock	-0.150***	-0.168***	-0.146***	-0.137***
	(0.014)	(0.016)	(0.013)	(0.015)
Algo pricing stock	0.726***	1.54***	0.715***	1.27***
	(0.105)	(0.314)	(0.103)	(0.313)
# of sku per visit	-0.678*	-2.30**	-0.600	-3.34***
	(0.381)	(1.17)	(0.372)	(1.22)
Log # of total visits	-0.450***	0.603**	-0.427***	0.195
	(0.076)	(0.242)	(0.075)	(0.244)
Log # of prior purchases - total	-0.283***	-1.42***	-0.264***	-1.28***
	(0.010)	(0.048)	(0.010)	(0.043)
First stage residual - 1	-0.499***	-1.49***	-0.523***	-1.17***
	(0.106)	(0.315)	(0.104)	(0.314)
First stage residual - 2	0.079***	0.123***	0.081***	0.100***
	(0.017)	(0.018)	(0.016)	(0.016)
Fixed-effects				
User	Yes	Yes	Yes	Yes
Product	Yes	Yes	Yes	Yes
Year week	Yes	Yes	Yes	Yes
Fit statistics				
Observations	2,040,975	1,473,688	2,345,376	1,919,419
Log-Likelihood	-911,753.6	-506,671.9	-956,591.2	-791,023.1

Two-way (User & Product) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

5.2 Haphazard Visitation Timing

Our second identification strategy is based on the assumption that the exact time users visit is as-good-as random. Consequently, the actual price they end up seeing depends on the time they visit. Because prices change frequently, had they come a few hours earlier or later, they may have seen a different price. We build on this thought experiment and posit a randomization scheme that allows us identification and inference — using Fisherian randomization inference — of the effect of algorithmic pricing stock on consumer behavior.

5.2.1 An Ideal Experiment of Exposure to Algorithmic Pricing

It is instructive to ponder what an ideal experiment for exposure to algorithmic pricing, i.e. exposure to high volatility in prices would look like. One may presume that an A/B test at the user or product level could help us achieve a clear identification of the impact of price volatility on consumer purchase behavior. However, notice that even if such a test were possible, the real exposure cannot be captured in a single event, rather it accumulates over time and this stock would be heavily driven by the user's browsing intensity. That is, a simple A/B test can be understood as an encouragement design (Holland, 1988) that indirectly induces variation in exposure to varied prices.

We try to emulate an "ideal" experiment with observational data using randomization inference. The idea is similar to how Donnelly, Ruiz, Blei, and Athey (2019) use surrounding weekly in-store price changes to estimate price elasticities for groceries. However, in our case, there are complex dependencies in the data since consumers visit multiple times to purchase multiple products. Furthermore, price changes occur at many different times. Independently of any particular user's visit, prices for products are changing due to different factors such as inventory or competitive pressures. Hence, the actual price a consumer sees conditional on visit depends upon her time of visit. If she were to visit a few hours earlier or later, then she may end up seeing a different price for the same product. We use the idea that the actual visit time of a particular user is as-good-as-random to generate counterfactual distributions of price exposure and algorithmic pricing stock. Subsequently, we use these counterfactual exposures to causally pin down the effect of exposure to multiple prices on consumer behavior.

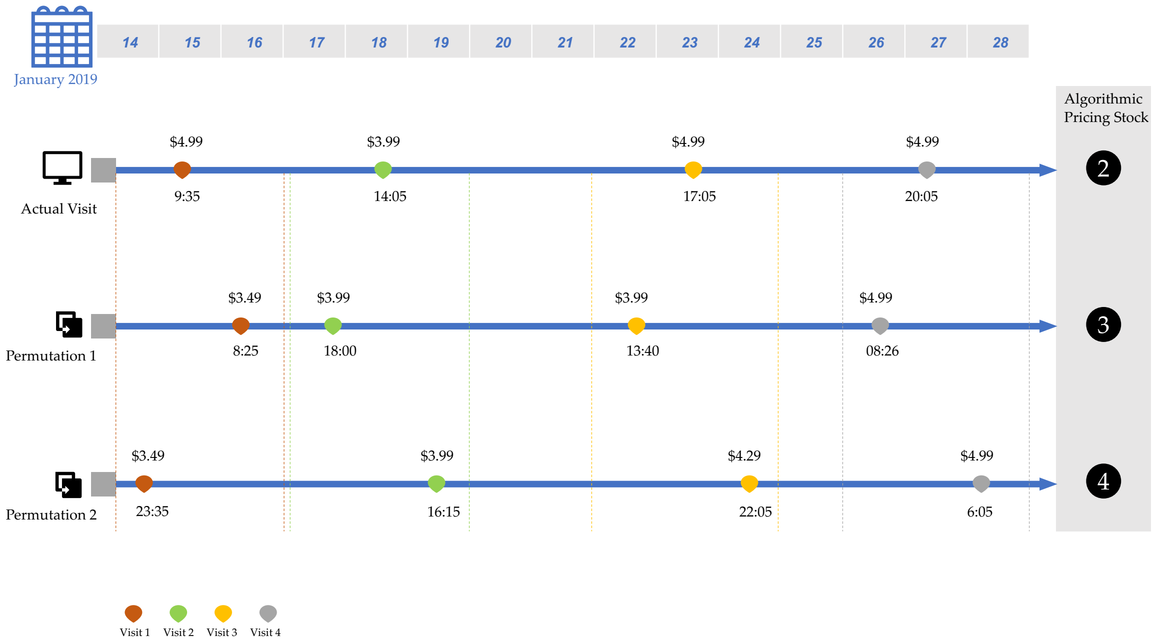

Figure 5: Permuted user visits by re-drawing user visit times within a ± 48 hour window of actual visit time.

Figure 5 helps build intuition behind this procedure. For simplicity, consider a single user who visits a single product four times during a two-week period. The dates and times the user visits are shown in the first row of the figure. Across these four visits, the user sees two distinct prices, and hence the total algorithmic pricing stock is 2. Consider the first visit of the user on Jan-15 at 9:15 AM. Suppose that instead of 9:15 AM on Jan-15, she visited the product at 8:25 AM on Jan-16. Then, all else equal, she would have been exposed to three prices, and her algorithmic pricing stock would be 3.

We generalize this idea and shuffle all visits for a user within ± 48 hours of the original visit, keeping the total number of visits and the products visited each time fixed. Since the prices are fluctuating independently of this user's visit and not personalized, she could get exposed to a different price in each permutation. Consequently, each permutation of a user's visit to the retailer's website creates a counterfactual price exposure and algorithmic pricing stock. The collection of all permutations for a user gives a vector of counterfactual price exposures and algo-pricing stocks. We use these counterfactual exposures and algo-pricing stocks for identification and inference. To give a sense of how the permuted assignments look like, Figure 6 shows the probability distribution of observed algo-pricing stock and permuted algo-pricing stock averaged across permutations. Overall we find that randomizing user's visit time does expose them to different prices for the same product. The distribution of observed algo-pricing stock exposure is shown in Figure I.1 in the Appendix.

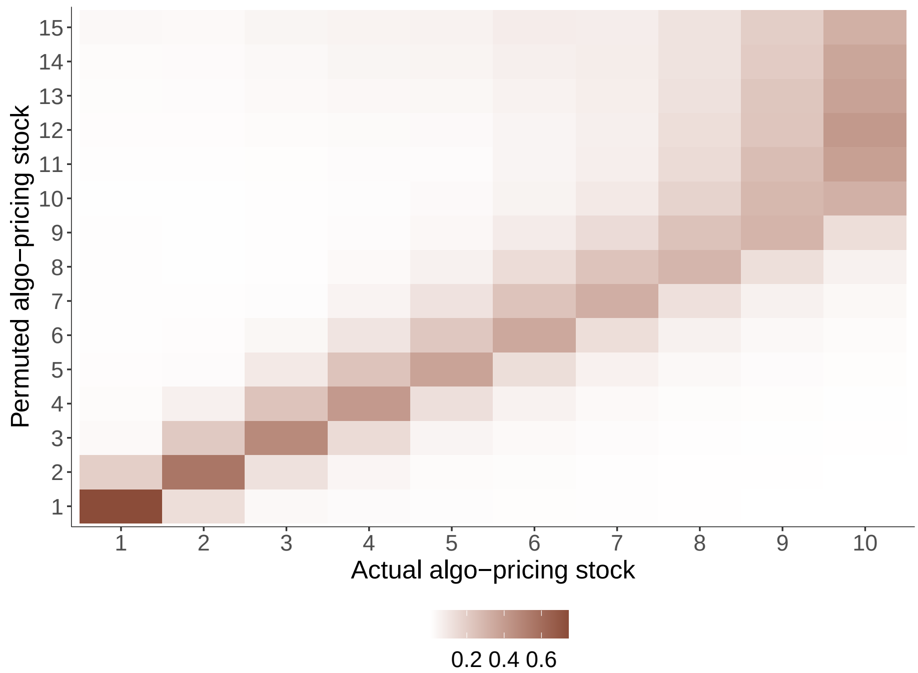

Figure 6: Conditional Distribution of Permuted Algo-Pricing Stock Given Observed Algo-Pricing Stock

It is important to test the validity of the assumption that consumer visit times are random. We present analysis similar to Donnelly, Ruiz, Blei, and Athey (2019) in which we compare the coefficients obtained by running the model on the actual data with the coefficients obtained by running the model on the permuted data. Since our data has more complex dependencies at the user and product level, we estimate the full fixed effects model specified in Equation 14 for each counterfactual exposure. For each permutation, we calculate the algo-pricing stock and user history variables over a 30-day period.

$$
\begin{aligned}
\mathbb{P}(Y_{ijt} = y) &= \frac{e^{-\lambda_{ijt}}\lambda_{ijt}^{q}}{q!}, \quad q = 0, 1, 2, \dots \\
\log(\lambda_{ijt}) &= \beta_0 + \beta_1 \log(P_{ijt}) + \beta_2 A_{it} + \beta_3 \log(P_{ijt}) \times A_{it} + \\
&\quad + \overline{\log(P_{ijt})} + \overline{A_{it}} + \overline{\log(P_{ijt}) \times A_{it}} \\
&\quad + \Gamma_i + \mu_j + \tau_t + \epsilon_{ijt}
\end{aligned} \tag{14}
$$

where $\overline{\log(P_{ijt})}$ is the average log price for a user, product, visit combination across all permutations, $\overline{A_{it}}$ is the average algo-pricing stock across all permutations, and $\overline{\log(P_{ijt}) \times A_{it}}$ is the average value of the interaction between log price and algo-pricing stock across all permutations. If the treatments have no effect, then on average across permutations, we would expect their effect to be centered at zero. The results from this test are shown in Figure 7 where we plot the distribution of z-stats from each permuted regression. We do indeed find that the treatments affect the outcome. The red line in each panel is the observed z-stat from the actual regression run using the observed exposures.

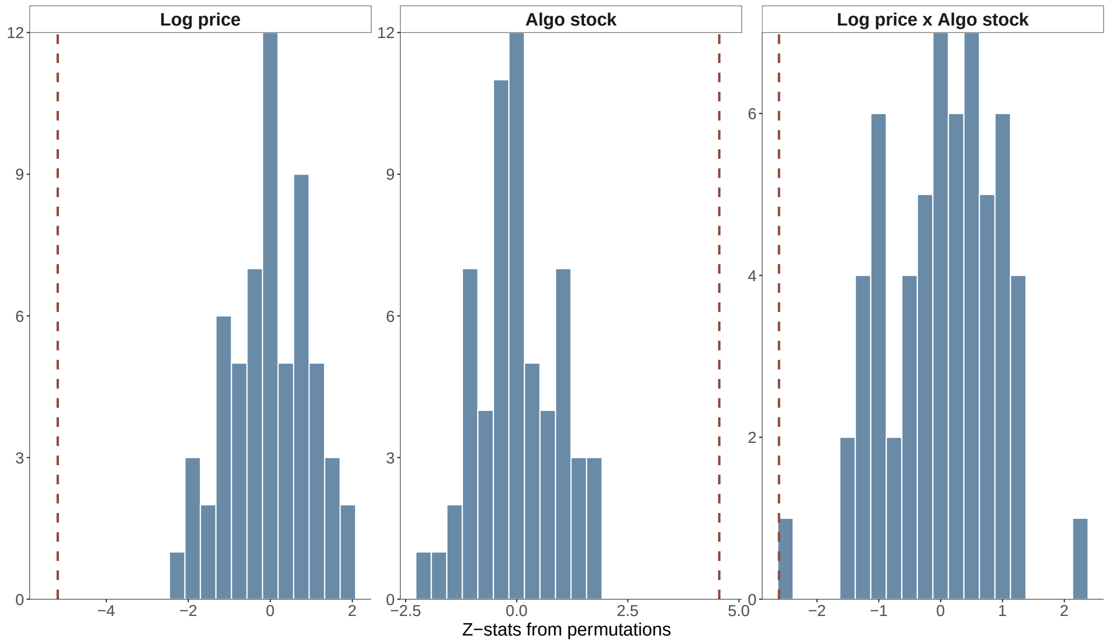

Figure 7: Placebo Tests using Counterfactual Price Exposure and Algo-Pricing Stock

Note: The blue histogram shows the distribution of z-scores obtained by running Model 14 on permuted data. The dashed red line is the z-score from the regression using observed data. All permutations and models are estimated on a random sample of 30,000 customers and the algo-pricing stock is calculated over a 30-day period.

5.2.2 Estimation Results

Finally, for inference, we estimate Equation 14 using the observed values of log price, algorithmic pricing stock, and their interaction while still controlling for the mean value of the independent variables across permutations. The results are shown in Table 7. As with the permutations, we use a 30-day window to calculate the observed algorithmic pricing stock. We estimate the model for a random sample of 30,000 consumers. As with the instrumental variables approach, we find that exposure to multiple prices for the same product does make the consumers substantially more price sensitive.

Table 7: User-level Price Sensitivity Estimates using Randomization Inference

Dependent Variable:	Units
Model:	(1)
Variables	
Log price	-1.06***
	(0.205)
Log price × Algo stock	-0.110***
	(0.042)
Algo stock	0.425***
	(0.093)
Mean price across perm.	0.034
	(0.203)
Mean algo stock across perm.	-0.126
	(0.081)
Mean price x Algo stock across perm.	-0.003
	(0.039)
Fixed-effects	
User	Yes
Product	Yes
Year week	Yes
Fit statistics	
Observations	505,425
Pseudo R^2	0.23673
Log-Likelihood	-192,254.8

Two-way (User & Product) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table shows the estimates from Model 14 estimated using the observed values of price, algo-pricing stock, and their interaction. The algo-pricing stock is computed over a 30-day period and the model is estimated for a random sample of 30,000 users.

6 Lab Experiments

The analysis in Section 5 shows that exposure to more number of unique prices increases price sensitivity. We are mindful that in a field setting it is not possible to exert full control over the unobserved reasons consumers decide to visit the online grocery platform. To address this limitation, we conduct a laboratory experiment to test the effect of price volatility in a controlled environment. The experimental design is simple: we ask participants to simulate purchase decisions, i.e. participants must click how many units they intend to buy each period. Participants are randomly assigned to two treatment conditions: stable prices and algorithmic prices. To the best of the authors' knowledge, there are no studies in the literature that have explored algorithmic pricing and price sensitivity in a laboratory experiment.

6.1 Experiment Design

The online shopping simulation involves a single product (e.g., a Nutella or Nestle's Cocoa), it lasts 12 periods, and the price might fluctuate from period to period. Participants decide how many units to buy (from 0 to 5) each period. They receive a budget at the beginning which is automatically adjusted based on the units that they have bought so far. Answers are sequential, i.e. participants answer period 1, then period 2, etc. Responses are incentivized by offering a bonus payout that depends on the total units bought and total savings, i.e. users that buy more (less) units when the price is low (high) receive a larger payout. Finally, many series of 12 prices are simulated based on the *real* data according to two pricing regimes: stable pricing and algorithmic pricing. In particular, the price sequences across conditions have a very similar average price (but different volatility and frequency of price changes).

For example, four periods under stable pricing might be ($5.98, $5.98, $5.76, $5.76); while the same periods under algorithmic pricing might be ($6.01, $5.88, $5.63, $5.91). Importantly, the price variation closely resembles the online grocer's strategy of stable pricing and algorithmic pricing, respectively. In periods of algorithmic pricing prices fluctuate frequently and often in tiny amounts, and in periods of stable pricing, prices are fairly stationary with small infrequent jumps.

We run two lab experiments – one on Amazon Mechanical Turk[4] and the other with MBA students from a large European university. Among the MBA students, the study covers 139 distinct users (self-reported average age 29.8 and 68% male), 52% randomly assigned to stable pricing, and 48% randomly assigned to algorithmic pricing. Additional methodological details and robustness specifications are reported in Appendix J. Reassuringly, we find similar results from the lab experiment on MTurk (Appendix J).

6.2 Lab Experiment Results

We estimate the following model:

$$\log(Y_{it}) = \beta_0 + \beta_1 \log(P_{it}) + \beta_2 Algo_i + \beta_3 \log(P_{it}) \times Algo_i + \epsilon_{it} \tag{15}$$

where Y_{it} and $P_{i,t}$ denote the quantity and price, respectively; $Algo_i$ is an indicator variable that takes value 1 if user i was assigned to the online shopping simulation with algorithmic prices (and 0 otherwise). As with the models before, we are interested in β_3, the coefficient on the interaction.

The results from the lab experiment are presented in Table 8. Consistent with the findings in Sections 4 and 5, participants exhibit a more price sensitive demand when exposed to prices with high volatility.

[4]MTurk has become a standard platform to conduct lab experiments in pricing (e.g., Wadhwa and Zhang (2015))

Table 8: Price Elasticity – Lab Experiment with MBA Students

Dependent Variables:	Log units		Units
	Gaussian	Gaussian	Poisson
Model:	(1)	(2)	(3)
Variables			
(Intercept)	1.70***	1.23***	1.38***
	(0.101)	(0.135)	(0.162)
Elasticity	-0.651***	-0.371***	-0.722***
	(0.059)	(0.076)	(0.091)
Algo		0.863***	1.61***
		(0.195)	(0.266)
Algo x Elasticity		-0.516***	-0.962***
		(0.114)	(0.160)
Fit statistics			
R^2	0.03901	0.04511	
Log-Likelihood	-1,479.5	-1,474.2	-2,632.7

One-way (User) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

6.3 Price Salience and Recall

We now revisit the conceptual framework discussed in Section 2 to conceptualize potential mechanisms underlying the effects. A critical role in that framework is price salience: variation in prices shifts attention to the price attribute thereby making consumers more price sensitive.

We make progress showing that salience indeed is a relevant mechanism. After the 12-period simulated shopping trip, we show participants in the lab experiment a different product (Oreo) for 5 seconds. Along with the standard product packaging, the image includes the product's price. We then ask participants in both groups to recall the price and size of the Oreo. We hypothesize that, if algorithmic pricing makes prices more salient, then a larger proportion of users in that treatment condition will be able to better recall the price of Oreo's. We operationalize salience as recall, consistent with the tradition in behavioral sciences, economics, and marketing. See, for example, Alba and Chattopadhyay (1986); Kissler, Herbert, Peyk, and Junghofer (2007); Finkelstein (2009); Kroft, Lange, and Notowidigdo (2013); Gaspelin, Leonard, and Luck (2015).

We test whether the proportion of correct responses is higher in the algorithmic pricing condition. Thus, we first classify a participant with a correct response if their answer was within \$0.05 of the correct price. More formally, consider a participant i who answers X_i; we define recall R_i as follows:

$$R_i = \begin{cases} 1, & \text{if} \quad |X_i - P^*| \leq 0.05 \\ 0, & \text{otherwise} \end{cases} \tag{16}$$

where P^* is the correct prices of Oreo's. Therefore, the proportion of correct respondents in each treatment condition is:

$$p_{algo} = \frac{1}{n_{algo}} \sum_i R_i$$
$$p_{stable} = \frac{1}{n_{stable}} \sum_i R_i$$

We then test the null whether $p_{algo} \leq p_{stable}$ using a two-sample proportions test. The results are shown in Figure 8. We find that participants in the algorithmic pricing condition are more likely to correctly recall the price of Oreo ($p-value < 0.03$). The proportions test is robust to using different cut-offs, e.g. $\{\$0.02, \$0.03, \$0.04\}$. Furthermore, we find no difference between the two conditions when asked to recall the size of the Oreo's package. Taken together, this evidence supports that a process through which high volatility increases price sensitivity is price salience. We emphasize that further research should examine the existence of additional mechanisms.

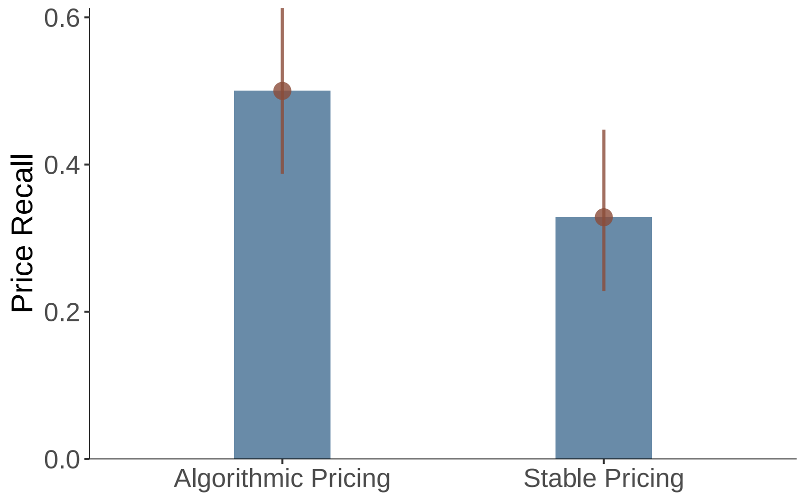

Figure 8: Testing Price Salience using Recall

Notes: Blue bars indicate the proportion of respondents who were able to recall the price of a second product (Oreo's) within $0.05 of the correct price. Red confidence bands are the 95% interval for the point estimate.

7 Implications and Discussion

Industry practitioners often express the concern of "lagging behind" in the race of adopting state-of-the-art pricing technology. Improvements in this technology typically include some form of machine-based tool to set or update prices. In this paper, we show that a stylized distinctive feature of algorithmic pricing, namely price volatility, modifies consumer behavior. We show, in the field and in the lab, that price volatility makes demand more price sensitive. Once again, the effect is identified within-consumers: greater exposure to algorithmic pricing makes a given consumer more price-sensitive. Our findings also indicate that a key mechanism through which this happens is salience in the price attribute.

This set of results encourages scholars to further connect technical innovations and consumer behavior. In light of the role of price salience, consumers are not indifferent to how online retailers change prices. Therefore, methodological improvements in the back-end (e.g., speed of optimization, high-dimensional inputs, price matching) are not sufficient in isolation; their connection to front-end user experiences is extremely relevant. Even if some form of machine-based pricing tool is profitable, it may trigger or shift salience to prices. And it is not obvious that all retailers benefit from price salience.

Thinking more broadly, which businesses want their customers to become more price-sensitive? The answer is probably very few. Perhaps price aggregator platforms or everyday low prices (EDLP) retailers might stand to benefit, but in general, businesses would like to avoid this side effect of algorithmic pricing. Said differently, while a retailer would not want to shut down algorithmic pricing, it would like to avoid the negative effect on price sensitivity. Our work suggests that price algorithms could be improved by accounting for consumer-level sensitivity to price volatility—a dimension often overlooked.

Finally, there are promising paths to further expand the analysis presented here. While we focus on understanding one process mechanism, namely salience, we are aware that it does not exclude other processes. Future work could explore the moderating role of price knowledge (Dickson and Sawyer, 1990b), price fairness (Xia, Monroe, and Cox, 2004; Anderson and Simester, 2008; Allender, Liaukonyte, Nasser, and Richards, 2021), fairness to machine algorithms (Lee, 2018), limited memory (Chen, Iyer, and Pazgal, 2010), or the formation of price cues.

It is also interesting to differentiate short-term reactions from long-term implications. Algorithmic pricing technology is fairly new and even specialized AI vendors are continually experimenting and updating their models. Studying the long-term impact of this new-age pricing technology on consumer behavior and market structure will help inform both business strategies and regulatory policies. Another important dimension, which is beyond the scope of the current paper, to consider is competition. Firms are not using pricing algorithms in isolation and a key input to these algorithms is competitor price. Characterizing the equilibrium effects between consumers and firms when multiple players in the market adopt algorithmic pricing is a promising, albeit challenging, avenue to pursue.

References

AGRAWAL, A., J. GANS, AND A. GOLDFARB (2018): *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

AIRBNB (2017): "What's smart about Smart Pricing?," https://blog.atairbnb.com/smart-pricing/.

ALBA, J. W., AND A. CHATTOPADHYAY (1986): "Salience effects in brand recall," *Journal of Marketing Research*, 23(4), 363–369.

ALLENDER, W. J., J. LIAUKONYTE, S. NASSER, AND T. J. RICHARDS (2021): "Price Fairness and Strategic Obfuscation," *Marketing Science*, 40(1), 122–146.

AMALDOSS, W., AND C. HE (2018): "Reference-dependent utility, product variety, and price competition," *Management Science*, 64(9), 4302–4316.

ANDERSON, E. T., AND D. I. SIMESTER (2003): "Effects of 9 price endings on retail sales: Evidence from field experiments," *Quantitative Marketing and Economics*, 1(1), 93–110.

——— (2004): "Long-run effects of promotion depth on new versus established customers: Three field studies," *Marketing Science*, 23(1), 4–20.

——— (2008): "Research note—does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes," *Marketing Science*, 27(3), 492–500.

——— (2010): "Price stickiness and customer antagonism," *The Quarterly Journal of Economics*, 125(2), 729–765.

ANDRÉ, Q., N. REINHOLTZ, AND B. DE LANGHE (2021): "Can Consumers Learn Price Dispersion? Evidence for Dispersion Spillover Across Categories," *Journal of Consumer Research*.

APARICIO, D., AND R. RIGOBON (2020): "Quantum Prices," *NBER Working Paper No. 26646*.

ASKER, J., C. FERSHTMAN, AND A. PAKES (2021): "Artificial Intelligence and Pricing: The Impact of Algorithm Design," *NBER Working Paper No. w28535*.

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*.

BECKER, G. S., AND K. M. MURPHY (1993): "A simple theory of advertising as a good or bad," *The Quarterly Journal of Economics*, 108(4), 941–964.

BERTINI, M., AND L. WATHIEU (2008): "Research note—Attention arousal through price partitioning," *Marketing Science*, 27(2), 236–246.

BLAKE, T., S. MOSHARY, K. SWEENEY, AND S. TADELIS (2021): "Price salience and product choice," *Marketing Science*, forthcoming.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2013): "Salience and consumer choice," *Journal of Political Economy*, 121(5), 803–843.

——— (2020): "Memory, attention, and choice," *The Quarterly Journal of Economics*, 135(3), 1399–1442.

BROWN, Z., AND A. MACKAY (2019): "Competition in Pricing Algorithms," *Available at SSRN 3485024*.

BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies.* WW Norton & Company.

BUSINESS INSIDER (2018): "Amazon changes prices on its products about every 10 minutes — here's how and why they do it," https://www.businessinsider.com/amazon-price-changes-2018-8.

BUSSE, M. R., N. LACETERA, D. G. POPE, J. SILVA-RISSO, AND J. R. SYDNOR (2013): "Estimating the effect of salience in wholesale and retail car markets," *American Economic Review*, 103(3), 575–79.

BUSSE, M. R., D. G. POPE, J. C. POPE, AND J. SILVA-RISSO (2015): "The psychological effect of weather on car purchases," *The Quarterly Journal of Economics*, 130(1), 371–414.

CALVANO, E., G. CALZOLARI, V. DENICOLÒ, AND S. PASTORELLO (2019): "Artificial intelligence, algorithmic pricing and collusion," *Available at SSRN 3304991.*

CAPLIN, A., AND M. DEAN (2015): "Revealed preference, rational inattention, and costly information acquisition," *American Economic Review*, 105(7), 2183–2203.

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, M. K. (2016): "Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform," in *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 455–455.

CHEN, Y., G. IYER, AND A. PAZGAL (2010): "Limited Memory, Categorization, and Competition," *Marketing Science*, 29(4), 650–670.

CHETTY, R., A. LOONEY, AND K. KROFT (2009): "Salience and taxation: Theory and evidence," *American Economic Review*, 99(4), 1145–77.

CIAN, L., A. KRISHNA, AND R. S. ELDER (2015): "A sign of things to come: behavioral change through dynamic iconography," *Journal of Consumer Research*, 41(6), 1426–1446.

COHEN, P., R. HAHN, J. HALL, S. LEVITT, AND R. METCALFE (2016): "Using big data to estimate consumer surplus: The case of uber," *NBER*.

DELLAVIGNA, S., AND M. GENTZKOW (2019): "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 134(4), 2011–2084.

DHOLAKIA, U. M. (2015): "Everyone Hates Uber's Surge Pricing – Here's How to Fix It," *Harvard Business Review*.

DICKSON, P. R., AND A. G. SAWYER (1990a): “The price knowledge and search of supermarket shoppers,” *The Journal of Marketing*, pp. 42–53.

DICKSON, P. R., AND A. G. SAWYER (1990b): “The Price Knowledge and Search of Supermarket Shoppers,” *Journal of Marketing*, 54(3), 42–53.

DONNELLY, R., F. R. RUIZ, D. BLEI, AND S. ATHEY (2019): “Counterfactual Inference for Consumer Choice Across Many Product Categories,” .

DORFMAN, R., AND P. O. STEINER (1954): “Optimal advertising and optimal quality,” *The American Economic Review*, 44(5), 826–836.

DUNCAN, J. (1984): “Selective attention and the organization of visual information.,” *Journal of Experimental Psychology: General*, 113(4), 501.

ELBERG, A., P. M. GARDETE, R. MACERA, AND C. NOTON (2019): “Dynamic effects of price promotions: Field evidence, consumer search, and supply-side implications,” *Quantitative Marketing and Economics*, 17(1), 1–58.

ELLISON, G., AND S. F. ELLISON (2009): “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, 77(2), 427–452.

ERDEM, T., M. P. KEANE, AND B. SUN (2008): “The impact of advertising on consumer price sensitivity in experience goods markets,” *Quantitative Marketing and Economics*, 6(2), 139–176.

FINKELSTEIN, A. (2009): “E-ztax: Tax salience and tax rates,” *The Quarterly Journal of Economics*, 124(3), 969–1010.

FISHER, M., S. GALLINO, AND J. LI (2018): “Competition-based dynamic pricing in online retailing: A methodology validated with field experiments,” *Management science*, 64(6), 2496–2514.

FOLKES, V., AND S. MATTA (2004): “The effect of package shape on consumers’ judgments of product volume: attention as a mental contaminant,” *Journal of Consumer Research*, 31(2), 390–401.

FORD, M. (2015): *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.

GASPELIN, N., C. J. LEONARD, AND S. J. LUCK (2015): “Direct evidence for active suppression of salient-but-irrelevant sensory inputs,” *Psychological science*, 26(11), 1740–1750.

GELMAN, A., AND J. HILL (2006): *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Analytical Methods for Social Research. Cambridge University Press.

HANSEN, K. T., K. MISRA, AND M. M. PAI (2021): “Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms,” *Marketing Science*.

HASTINGS, J. S., AND J. M. SHAPIRO (2013): “Fungibility and consumer choice: Evidence from commodity price shocks,” *The Quarterly Journal of Economics*, 128(4), 1449–1498.

HAWS, K. L., AND W. O. BEARDEN (2006): "Dynamic pricing and consumer fairness perceptions," *Journal of Consumer Research*, 33(3), 304–311.

HENDEL, I., AND A. NEVO (2004): "Intertemporal substitution and storable products," *Journal of the European Economic Association*, 2(2-3), 536–547.

HITSCH, G. J., A. HORTACSU, AND X. LIN (2019): "Prices and promotions in us retail markets: Evidence from big data," *NBER*.

HOCH, S. J., B.-D. KIM, A. L. MONTGOMERY, AND P. E. ROSSI (1995): "Determinants of store-level price elasticity," *Journal of marketing Research*, 32(1), 17–29.

HOLLAND, P. W. (1988): "Causal Inference, Path Analysis, and Recursive Structural Equations Models," *Sociological Methodology*, pp. 449–484.

JACOBY, J. (1984): "Perspectives on information overload," *Journal of Consumer Research*, 10(4), 432–435.

JINDAL, P., AND A. ARIBARG (2021): "The Importance of Price Beliefs in Consumer Search," *Journal of Marketing Research*, forthcoming.

JUNG, J., J.-H. KIM, F. MATEJKA, C. A. SIMS, ET AL. (2019): "Discrete actions in information-constrained decision problems," *The Review of Economic Studies*, 86(6), 2643–2667.

KALYANARAM, G., AND R. S. WINER (1995): "Empirical generalizations from reference price research," *Marketing Science*, 14(3_supplement), G161–G169.

KARMARKAR, U. R., B. SHIV, AND B. KNUTSON (2015): "Cost conscious? The neural and behavioral impact of price primacy on decision making," *Journal of Marketing Research*, 52(4), 467–481.

KISSLER, J., C. HERBERT, P. PEYK, AND M. JUNGHOFER (2007): "Buzzwords: early cortical responses to emotional words during reading," *Psychological Science*, 18(6), 475–480.

KRIDER, R. E., P. RAGHUBIR, AND A. KRISHNA (2001): "Pizzas: π or square? Psychophysical biases in area comparisons," *Marketing Science*, 20(4), 405–425.

KROFT, K., F. LANGE, AND M. J. NOTOWIDIGDO (2013): "Duration dependence and labor market conditions: Evidence from a field experiment," *The Quarterly Journal of Economics*, 128(3), 1123–1167.

LEE, M. K. (2018): "Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management," *Big Data & Society*, 5(1), 2053951718756684.

LICHTENSTEIN, D. R., N. M. RIDGWAY, AND R. G. NETEMEYER (1993): "Price perceptions and consumer shopping behavior: a field study," *Journal of Marketing Research*, pp. 234–245.

LILIEN, G. L., P. KOTLER, AND K. S. MOORTHY (1995): *Marketing models*. Prentice Hall.

McAfee, R. P., and V. Te Velde (2006): "Dynamic pricing in the airline industry," *Handbook on economics and information systems*, 1, 527–67.

Mela, C. F., S. Gupta, and D. R. Lehmann (1997): "The long-term impact of promotion and advertising on consumer brand choice," *Journal of Marketing research*, 34(2), 248–261.

Miklós-Thal, J., and C. Tucker (2019): "Collusion by algorithm: Does better demand prediction facilitate coordination between sellers?," *Management Science*, 65(4), 1552–1561.

Monroe, K. B. (1973): "Buyers' subjective perceptions of price," *Journal of Marketing Mesearch*, pp. 70–80.

Petrin, A., and K. Train (2010): "A Control Function Approach to Endogeneity in Consumer Choice Models," *Journal of Marketing Research*, 47(1), 3–13.

Rotemberg, J. J. (2005): "Customer anger at price increases, changes in the frequency of price adjustment and monetary policy," *Journal of Monetary Economics*, 52(4), 829–852.

Semenova, V., M. Goldman, V. Chernozhukov, and M. Taddy (2017): "Estimation and Inference about Heterogeneous Treatment Effects in High-Dimensional Dynamic Panels," *arXiv e-prints*, pp. arXiv–1712.

Sethuraman, R., G. J. Tellis, and R. A. Briesch (2011): "How well does advertising work? Generalizations from meta-analysis of brand advertising elasticities," *Journal of Marketing Research*, 48(3), 457–471.

Shapiro, B. T., G. J. Hitsch, and A. E. Tuchman (2021): "TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands," *Working Paper*.

Thomas, M., D. H. Simon, and V. Kadiyali (2010): "The price precision effect: Evidence from laboratory and market data," *Marketing Science*, 29(1), 175–190.

Townsend, C., and B. E. Kahn (2014): "The "visual preference heuristic": The influence of visual versus verbal depiction on assortment processing, perceived variety, and choice overload," *Journal of Consumer Research*, 40(5), 993–1015.

Vanhuele, M., and X. Drèze (2002): "Measuring the price knowledge shoppers bring to the store," *Journal of Marketing*, 66(4), 72–85.

Wadhwa, M., and K. Zhang (2015): "This number just feels right: The impact of roundedness of price numbers on product evaluations," *Journal of Consumer Research*, 41(5), 1172–1185.

Washington Post (2015): "How Uber surge pricing really works," https://www.washingtonpost.com/news/wonk/wp/2015/04/17/how-uber-surge-pricing-really-works/.

WEISSTEIN, F. L., K. B. MONROE, AND M. KUKAR-KINNEY (2013): "Effects of price framing on consumers' perceptions of online dynamic pricing practices," *Journal of the Academy of Marketing Science*, 41(5), 501–514.

XIA, L., K. B. MONROE, AND J. L. COX (2004): "The Price is Unfair! A Conceptual Framework of Price Fairness Perceptions," *Journal of Marketing*, 68(4), 1–15.

A Algorithmic Pricing Periods

Table A.1 shows additional descriptive statistics for periods (weeks) identified with and without algorithmic pricing. Overall, the statistics indicate that the measure of algorithmic pricing does capture periods in which the price of a product experienced intense variation.

Table A.1: Summary Statistics During Algorithmic and Non-Algorithmic Weeks

	Stable pricing	Algo. pricing
Observations	81,337	40,972
Std. price	0.08	0.33
Mean price	9.14	9.23
Min price	9.05	8.82
Max price	9.27	9.75
Weekly price changes	0.7717	3.536

Figure A.1 visualizes, in the form of a heat map, the percent of products across categories that experience algorithmic pricing over time. Reassuringly, there is variation in exposure to algorithmic pricing across products within and across categories, as well as throughout the time-series of the data.

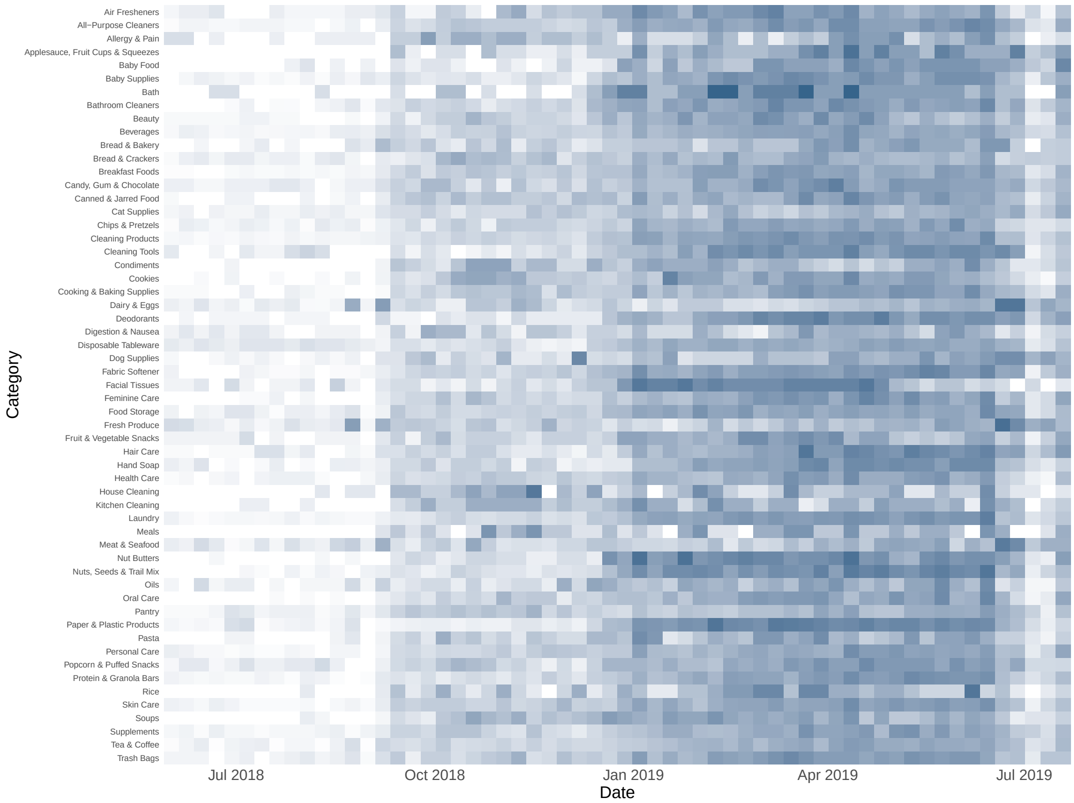

Figure A.1: Share of Products per Category under Algorithmic Pricing

Notes: Each cell is a category-week combination

If prices were dynamically updated according to a mark-up rule (e.g., $p_{it} = m * c_{it}$), changes in price (p_{it}) might be triggered by changes in cost (c_{it}). However, and interestingly, periods of algorithmic pricing do not appear to be driven by high-frequency changes in costs. Panels (a) and (b) of Figure A.2 show the distribution of the cost and price changes in algorithmic pricing weeks and in non algorithmic pricing weeks, respectively. Summary statistics are reported in Table A.2. Overall the evidence indicates, similar to Fisher, Gallino, and Li (2018), that algorithmic or dynamic pricing is not primarily driven by cost shifters.

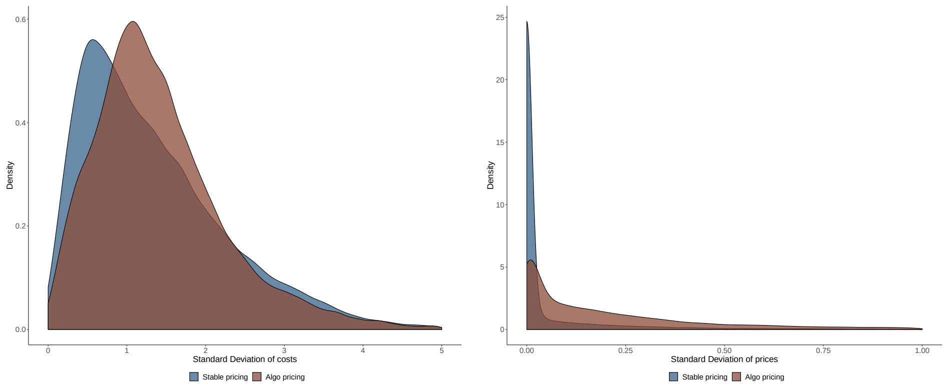

(a) Variation in cost (b) Variation in prices

Figure A.2: Costs and Price Changes during Algorithmic and Non-Algorithmic Pricing Periods

Table A.2: Variation in Costs and Prices during Algorithmic and Non-Algorithmic Pricing Periods

	Cost Std.	Price Std.
Stable pricing	1.419	0.078
Algorithmic pricing	1.476	0.379

B Statistical Tests of the Algorithmic Pricing Indicator

Table B.1 shows the results of a χ^2 test for the significance of the algorithmic pricing indicador, as defined in Section 3. We find that both the algorithmic pricing indicator and its interaction with price capture a statistically significant portion of the purchase variation.

Model	Log Lik.	χ^2	p.value
Only price	-112,205.29	-	-
Price + algorithmic pricing indicator	-112,074.71	256.95	< 0.001
Price + algorithmic pricing indicator + interaction	-112,004.51	138.06	< 0.001

Table B.1: χ^2 Test for Algorithmic Pricing

Additionally, Table B.2 shows the results of the ANOVA test for varying intercepts by product and by category. The results support the existence of individual differences in price elasticities across products and categories.

Model	Log Lik.	χ^2	p.value
Varying intercept by product	-118514.4	-	-
Varying intercept and slope by product	-114814.5	7399.65	< 0.001
Varying intercept and slope by category and product (nested)	-114565.3	498.44	< 0.001

Table B.2: ANOVA Test for Mixed Effects

C Definition of Algorithmic Pricing

Does more intensity in algorithmic pricing magnify the price sensitivity? Figure C.1 shows that that it does. More precisely, a more stringent definition of algorithmic pricing (i.e., a product-week pair is required to exhibit a more intense price variation to be classified as an algorithmic pricing week) increases the interaction with the price elasticity. The intensity is measured by the minimum number of unique prices in a given product-week.

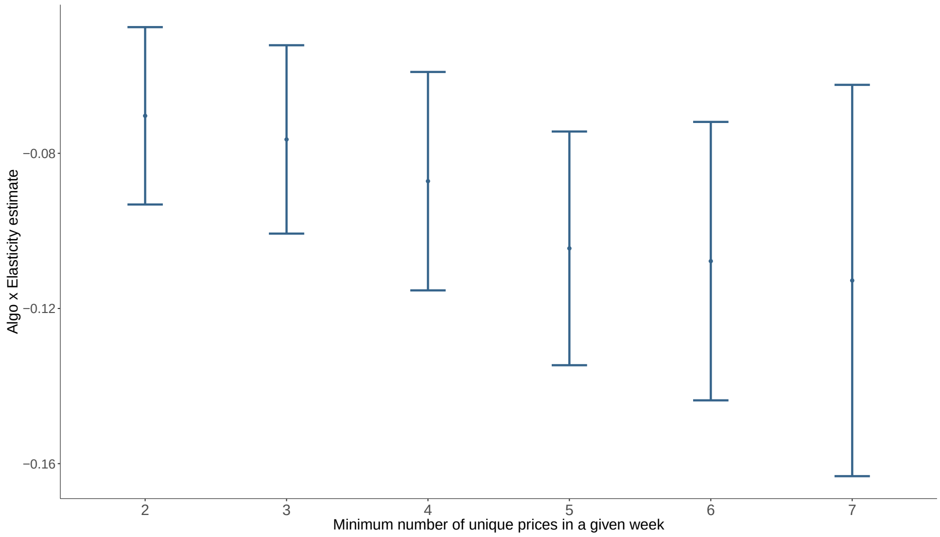

Figure C.1: Change in Price Elasticity with Increasing Variation in Prices

D Own-Price Elasticities

Figure D.1 shows the distribution of the own-price elasticities computed at the product-level. Price elasticities are restricted to those significant at the 10%. Panel (a) depicts the distribution using a separate linear regression for each product, and Panel (b) depicts the distribution using a multilevel model allowing the elasticity to vary by product.

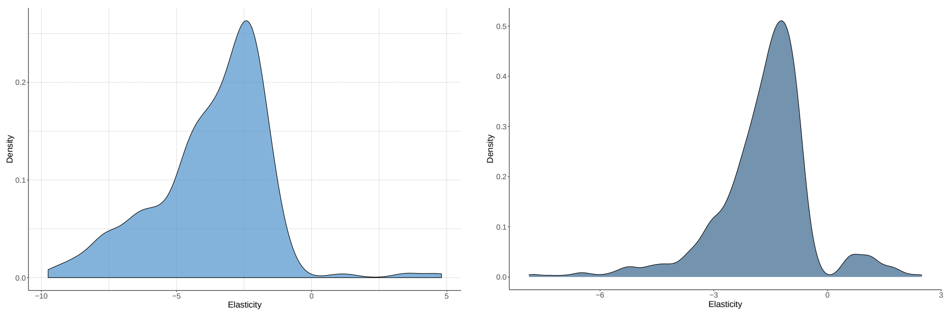

(a) OLS (b) Mixed Effects

Figure D.1: Distribution of Own-Price Elasticities

Figure D.2 shows the distribution of the product-level own-price elasticities during algorithmic and stable pricing weeks.

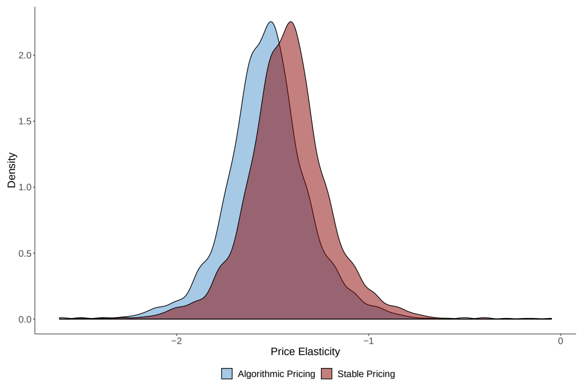

Figure D.2: Product-level distribution of price elasticities during stable and algorithmic pricing periods estimated using mixed-effects model

E Types of Products

We examine whether the effect of algorithmic pricing on price sensitivity varies across types of products. We consider three classifications of products: cheap and expensive, high-revenue and low-revenue, perishable or non-perishable. The results are shown in Figures E.1, E.2, and E.3, respectively. In all graphs, the dashed red line in the center is global average across categories.

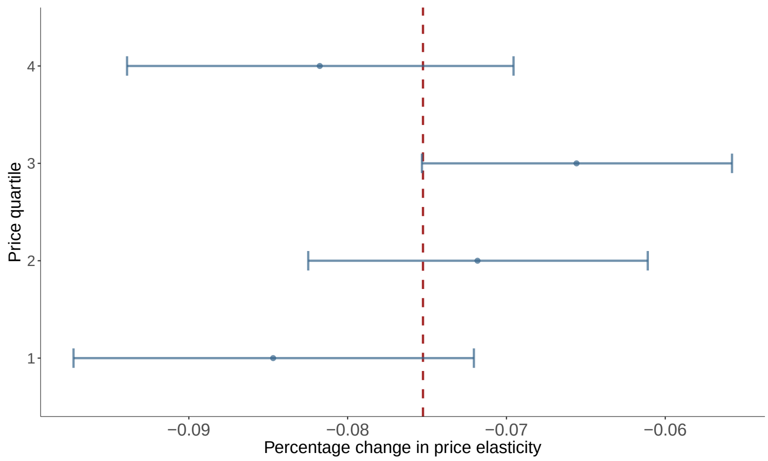

Figure E.1: Estimated price elasticity across products split by price quartile.

We use first 16 weeks of our sample to calculate the average price for each product and categorize them into quartiles based on average price. Elasticiies are estimated using mixed-effects model similar to Equation 9 on the remaining sample.

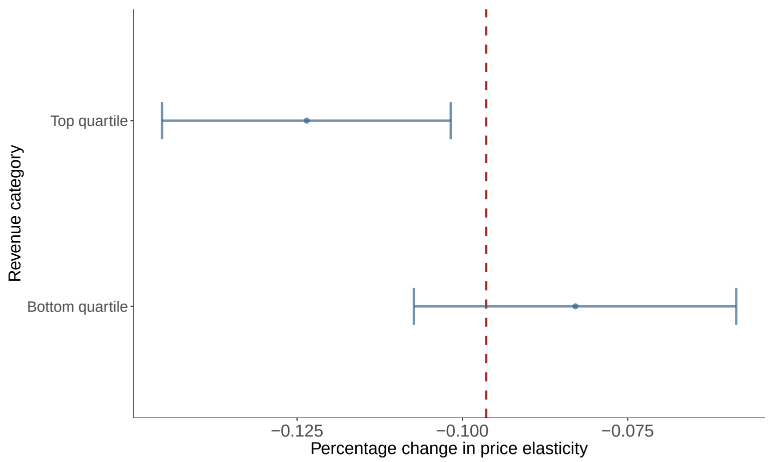

Figure E.2: Estimated price elasticity across less popular and more popular products.

We use first 16 weeks of our sample to calculate the total revenue for each product and categorize them into quartiles based on total revenue. Elasticiies are estimated using mixed-effects model similar to Equation 9 on the remaining sample. We shope the top and bottom quartile for clarity. The middle two quartiles were centered at the global average.

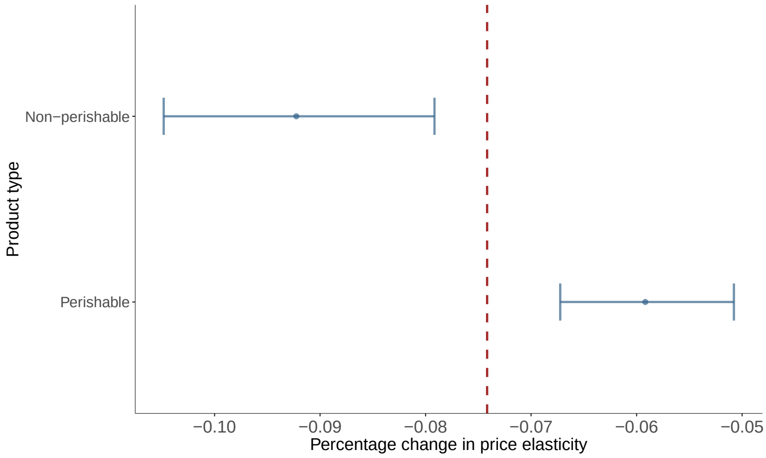

Figure E.3: Estimated price elasticity across perishable and non-perishable products.

Perishable products include categories such as dairy & eggs, meat & seafood, and fresh produce. Elasticiies are estimated using mixed-effects model similar to Equation 9 on the remaining sample.

F Before-and-After Events

We estimate the effect of algorithmic pricing through a before-and-after event study, exploiting variation in the timestamp different products had their first algorithmic pricing week. Figure F.1 shows that there is considerable variation in the timing of adopting algorithmic pricing across products. See also Figure A.1 for heatmap split by category.

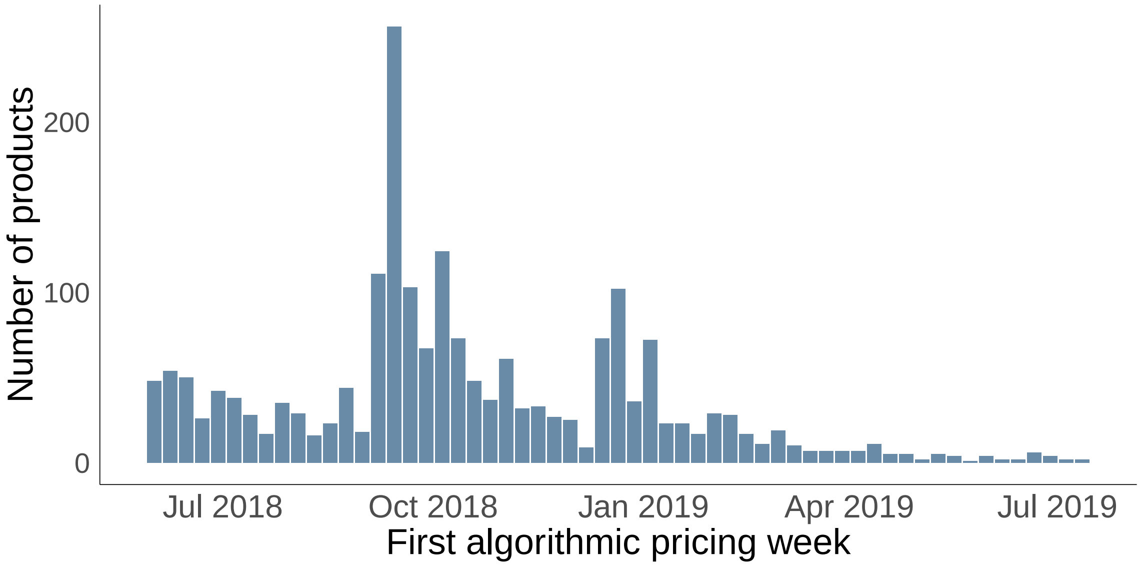

Figure F.1: Histogram of the First Event of Algorithmic Pricing for All 2,044 Products

Table F.1: Robustness checks for aggregate elasticity estimates before and after algorithmic pricing periods

Dependent Variable:	Log units	
	12 weeks	28 weeks
Model:	(1)	(2)
Variables		
Elasticity	-1.25***	-1.03**
	(0.258)	(0.389)
Post period	0.154***	0.238***
	(0.057)	(0.084)
Elasticity × Post period	-0.065***	-0.101***
	(0.024)	(0.035)
Fixed-effects		
Product	Yes	Yes
Year week	Yes	Yes
Fit statistics		
Observations	19,607	7,652
R^2	0.644	0.687
Log-Likelihood	-14,557	-5,003

Two-way (Product & Year week) standard-errors
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table F.2: Robustness tests for user-level price sensitivity before and after algorithmic pricing periods

Dependent Variable:	Units				
	12 weeks	28 weeks	12 weeks	20 weeks	28 weeks
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Log price	-1.13***	-0.953***	-0.950***	-0.964***	-0.877***
	(0.181)	(0.287)	(0.140)	(0.172)	(0.233)
Post period	0.187***	0.245***	0.184***	0.199***	0.240***
	(0.052)	(0.073)	(0.052)	(0.061)	(0.073)
Log price × Post period	-0.049**	-0.084***	-0.049**	-0.070**	-0.083**
	(0.024)	(0.032)	(0.024)	(0.027)	(0.032)
Fixed-effects					
User			Yes	Yes	Yes
Product			Yes	Yes	Yes
User-product	Yes	Yes			
Year week	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	60,404	27,264	99,008	65,894	38,784
Pseudo R^2	0.450	0.496	0.429	0.448	0.479
Log-Likelihood	-79,801	-36,816	-114,626	-78,589	-48,176

Three-way (User & Product & Year week) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

G Aggregate Robustness Checks

We test whether the effect of algorithmic pricing on price sensitivity is robust to the definition of price sensitivity. In Table G.1 we repeat the models of Table 3 using the following definition of algorithmic pricing week:

A product is under algorithmic pricing in a given week if:

1. The total absolute change in price in that week is greater than the median total absolute change acorss all weeks; **AND**,
2. The number of changes in price in that week is greater than the median number of changes in prices across all weeks

Note that in both the Tables, 3 and G.1, the idea is to capture high frequency price variation in prices; what differs is how we quantify those changes. Table G.1's results show that our results are robust to a different definition of algorithmic pricing.

Table G.1: Elasticity estimates with multiple specifications and different algo-week indicator

Dependent Variables:	Log units	Units	Log units	
	Gaussian	Poisson	Gaussian	Ortho ML
Model:	(1)	(2)	(3)	(4)
Variables				
Elasticity	-1.518***	-1.581***	-1.371***	-2.533***
	(0.100)	(0.135)	(0.128)	(0.044)
Algo	0.1051***	0.094***		0.029**
	(0.0328)	(0.034)		(0.014)
Elasticity × Algo	-0.031**	-0.032**		-0.273***
	(0.014)	(0.015)		(0.063)
Post period			0.202***	
			(0.056)	
Elasticity × Post period			-0.062**	
			(0.024)	
Fixed-effects				
Product	Yes	Yes	Yes	
Year week	Yes	Yes	Yes	
Fit statistics				
Observations	122,309	122,309	122,309	90,316
59,157				
R^2	0.558		0.567	0.326
Log-Likelihood	-112,172.42	-592,395.11	-82,303.68	-49,620.81

Product & Year week standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

H Consumer Level Robustness Checks

Table H.1: Robustness checks for user level elasticity estimates using OLS and TSLS

Dependent Variable:	Log units	
	(1)	(2)
Model:	OLS	TSLS
Variables		
Log price	-0.093***	-0.088***
	(0.007)	(0.008)
Log price x Algo pricing stock	-0.001**	-0.003***
	(0.0006)	(0.001)
Algo pricing stock	0.006***	0.036***
	(0.001)	(0.011)
# of sku per visit	-0.227***	-0.131***
	(0.006)	(0.042)
# of total visits	-0.018***	-0.037***
	(0.001)	(0.009)
# of prior purchases	-0.025***	-0.026***
	(0.0007)	(0.0008)
Fixed-effects		
User	Yes	Yes
Product	Yes	Yes
Year week	Yes	Yes
Fit statistics		
Observations	4,621,411	4,621,411
Log-Likelihood	-292,530.2	-295,076.2

Two-way (User & Product) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table shows OLS and two-stage least square estimates for the consumer-level model in Equation 13. Algorithmic pricing stock and time-varying history variables are calculated over a 30-day period. The dependent variable is the log number of units of a particular product, purchased by a user during a single visit.

Table H.2: Robustness checks for user level elasticity estimates using different number of days for historical data

Dependent Variable:	Units					
	7 days		15 days		60 days	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Log price	-1.05***	-0.837***	-1.01***	-0.840***	-0.948***	-0.854***
	(0.083)	(0.091)	(0.081)	(0.087)	(0.079)	(0.083)
Log price x Algo pricing stock	-0.074***	-0.203***	-0.079***	-0.170***	-0.063***	-0.104***
	(0.008)	(0.018)	(0.007)	(0.013)	(0.005)	(0.010)
# of sku per visit	-1.74***	-0.472	-1.92***	-0.422	-2.78***	-1.39***
	(0.071)	(0.315)	(0.064)	(0.319)	(0.074)	(0.363)
Log # of total visits	-0.150***	-0.434***	-0.142***	-0.455***	0.013	-0.285***
	(0.015)	(0.071)	(0.013)	(0.066)	(0.014)	(0.076)
Log # of prior purchases - total	-0.265***	-0.279***	-0.248***	-0.268***	-0.366***	-0.386***
	(0.010)	(0.010)	(0.008)	(0.009)	(0.010)	(0.011)
Algo pricing stock	0.269***	0.925***	0.241***	0.853***	0.182***	0.648***
	(0.022)	(0.098)	(0.017)	(0.088)	(0.014)	(0.095)
First stage residual - 1		-0.711***		-0.667***		-0.517***
		(0.097)		(0.087)		(0.094)
First stage residual - 2		0.153***		0.115***		0.063***
		(0.020)		(0.015)		(0.011)
Fixed-effects						
User	Yes	Yes	Yes	Yes	Yes	Yes
Product	Yes	Yes	Yes	Yes	Yes	Yes
Year week	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	2,672,388	2,672,388	3,680,937	3,680,937	5,206,155	5,206,155
Pseudo R^2	0.235	0.235	0.231	0.231	0.228	0.228
Log-Likelihood	-1,029,052.5	-1,028,759.8	-1,443,524.4	-1,443,161.3	-2,086,407.0	-2,086,154.9

Two-way (User & Product) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table shows two-stage control function estimates for the consumer-level model in Equation 13. Algorithmic pricing stock and time-varying history variables are calculated over 7-day, 15-day, and 60-day periods. The dependent variable in all models is the number of units of a particular product, purchased by a user during a single visit. All models control for user, product, and year-week fixed effects. Standard errors for columns 3 & 4 are estimated using clustered bootstrap to account for first-stage estimation error.

I Exposure to Algorithmic Pricing Stock

Figure I.1 shows the distribution of user level algorithmic pricing stock over a 30-day period.

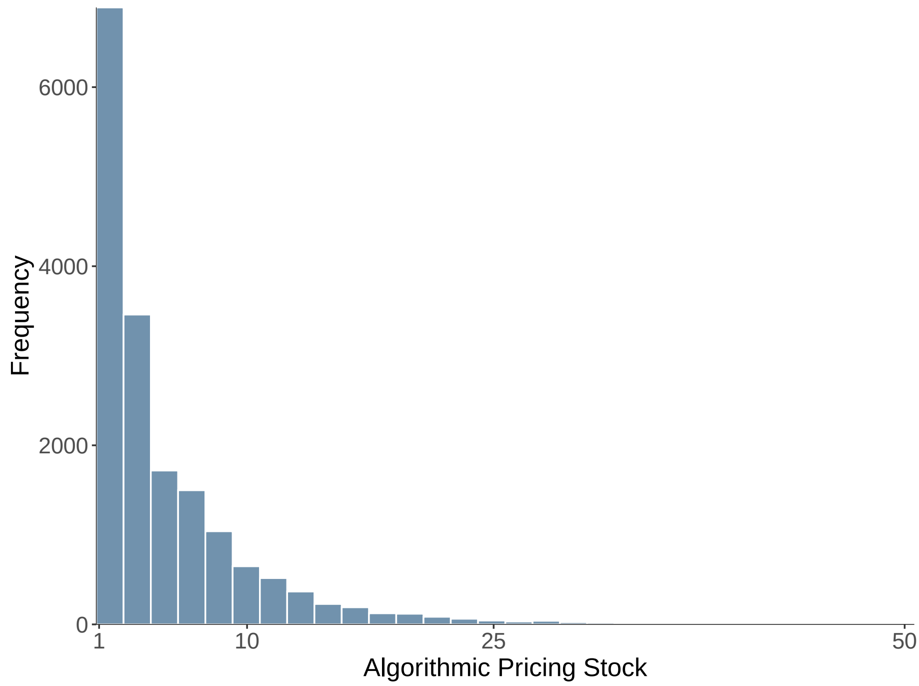

Figure I.1: Observed 30-day Algo-Pricing Stock Across Users

J Lab Experiment

We conduct two lab experiments – one on Amazon Mechanical Turk and the other with MBA students from a large European university. Participants are randomly assigned to two pricing regimes: stable pricing or algorithmic pricing. The set-up and task are similar in both studies, only the prices are re-drawn randomly. As described in Section 6, the algorithmic pricing condition is characterized by high frequency price changes (calibrated using real data from the online retailer). Figure J.1 shows one pair of price series used in the experiment (we used 4 pairs in total).

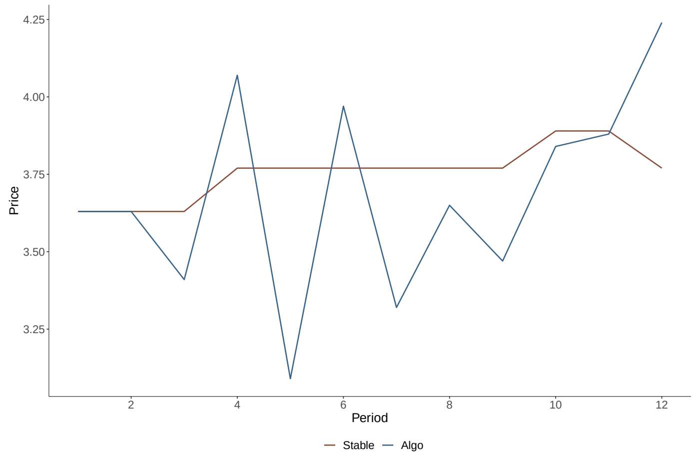

Figure J.1: Sampled price series for lab experiment

Note: A user was randomly assigned to either of these pricing conditions and then their purchase behavior was recorded.

J.1 MBA Students

Table J.1 shows the summary statistics across the two conditions. Important to note is that the average price is similar but the standard deviation of prices is much higher in the algorithmic pricing condition.

Table J.1: Summary Statistics from Lab Experiment with MBA Students

	Algo	Stable	*p*
Obs.	864	804	
Users	72	67	
Mean price	4.42	4.49	0.66
SD price	0.33	0.11	
Units purchased	14.9	14.2	0.08
	(3.66)	(4.04)	
Spend	64.7	62.6	0.20
	(9.52)	(9.54)	

The third columns shows the p-value from a t-test testing the difference in means across the two conditions.

Table J.4 in the main text shows the results from the main model. Here, we do a robustness check by including user fixed effects in the estimation. In the lab experiment, since a given participant is only allocated to one of the conditions and hence the main effect of the pricing regime is not identified. Hence we run separate regressions for the stable pricing and algorithmic pricing users, while accounting for user-level fixed effects. The results are presented in Table J.2 and, once again, show that price sensitivity is higher in the algorithmic pricing regime.

Table J.2: Robustness check for lab experiment with MBA students using fixed effects estimation

Dependent Variable:	Log units		
	Combined	Algo	Stable
Model:	(1)	(2)	(3)
Variables			
Elasticity	-3.23***	-3.68***	-1.25
	(0.259)	(0.269)	(0.759)
Fixed-effects			
User	Yes	Yes	Yes
Fit statistics			
R^2	0.16332	0.25041	0.07662
Log-Likelihood	-1,364.0	-681.39	-673.43

One-way (User) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

J.2 MTurk Experiment

Table J.3 shows the summary statistics across the two conditions from the MTurk experiment. Again the average prices are quite similar, only the standard deviation of price varies.

Table J.3: Summary Statistics from MTurk Lab Experiment

	Algo	Stable	p
Obs.	504	564	
Users	42	47	
Price	6.17	6.40	0.18
SD Price	2.83	1.21	
Units purchased	11.2	10.2	0.25
	(3.66)	(4.04)	
Spend	60.1	56.2	0.041
	(8.1)	(9.71)	

The third columns shows the p-value from a t-test testing the difference in means across the two conditions. Standard errors in parantheses.

We repeat the exercises of estimating elasticity using two specifications for the MTurk lab experiment as well. The results are in Tables J.4 and J.5. Again, in both specifications we find that consumers in the algorithmic pricing condition exhibit greater price sensitivity.

Table J.4: Price Elasticity – Lab Experiment

Dependent Variables:	Log units		Units
	Gaussian	Gaussian	Poisson
Model:	(1)	(2)	(3)
Variables			
(Intercept)	1.22***	1.12***	1.72***
	(0.053)	(0.085)	(0.186)
Log price	-0.385***	-0.340***	-1.00***
	(0.025)	(0.038)	(0.091)
Algo		0.207**	0.500**
		(0.102)	(0.228)
Log price x Algo		-0.099**	-0.244**
		(0.047)	(0.118)
Fit statistics			
R^2	0.064	0.066	-
Log-Likelihood	-815.50	-814.75	-1,384.3

One-way (User) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table J.5: Robustness check for lab experiment results using fixed effects estimation

Dependent Variable:	Log units		
	Combined	Algo	Stable
Model:	(1)	(2)	(3)
Variables			
Elasticity	-1.51***	-1.54***	-1.38**
	(0.292)	(0.338)	(0.680)
Fixed-effects			
User	Yes	Yes	Yes
Fit statistics			
R^2	0.35	0.37	0.33
Log-Likelihood	-620.59	-295.86	-324.62

One-way (User) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*