Working Paper 23-007

Dynamic Pricing and Demand Volatility: Evidence from Restaurant Food Delivery

Alexander MacKay
Dennis Svartbäck
Anders G. Ekholm

Dynamic Pricing and Demand Volatility: Evidence from Restaurant Food Delivery

Alexander MacKay
Harvard Business School

Dennis Svartbäck
Priceff Ltd

Anders G. Ekholm
Priceff Ltd; LUT University

Working Paper 23-007

The data for this project were provided by Priceff Ltd. through an agreement with one of its client companies. Ekholm is a founding partner of Priceff and Svartbäck is an employee of Priceff. No party received compensation from another party for this research. Funding for this research was provided in part by Harvard Business School.

Dynamic Pricing and Demand Volatility: Evidence from Restaurant Food Delivery*

Alexander MacKay
Harvard University†

Dennis Svartbäck
Priceff Ltd

Anders G. Ekholm
Priceff Ltd and
LUT University

March 7, 2023

Abstract

Pricing technology that allows firms to rapidly adjust prices has two potential benefits. Prices can respond more rapidly to demand shocks, leading to higher revenues. On the other hand, time-varying prices can be used to smooth out demand across periods, reducing costs in markets with capacity constraints. Using data from the staggered adoption of a pricing algorithm, we measure the impacts of time-varying pricing in the context of restaurant food delivery. On average, the pricing algorithm reduced prices, though it led to substantial variation in prices within and across days. We find that the adoption of time-varying pricing reduced demand volatility, resulting in a relative increase in the share of transactions occurring during low-demand periods. Consumers appear to strategically time purchases at high frequencies and also across days of the week. We estimate that the volatility semi-elasticity, which we define to reflect the relationship between time-series variation in quantities and prices, is -1.96. Our results point to the potential efficiency gains of time-varying pricing when firms face capacity constraints.

Keywords: Pricing Algorithms, Dynamic Pricing, Demand Volatility, Delivery Services
JEL Classification: D4, L1, L81, L86

---

*We thank Zach Brown, Juan Camilo Castillo, Rob Clark, Jim Dana, Kwabena Donkor, Dan Ershov, Charlie Murry, Nate Miller, and Jesse Shapiro for helpful comments, as well as seminar and conference participants at Harvard Business School and the NBER Workshop of Digital Economics. The authors wish to express their gratitude to Johanna Kuosmanen and Jens Liman for their comments and support. The data for this project were provided by Priceff Ltd. through an agreement with one of its client companies. Ekholm is a founding partner of Priceff and Svartbäck is an employee of Priceff. No party received compensation from another party for this research.

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

1 Introduction

Pricing algorithms enable high-frequency automated price changes in response to time-varying market conditions. This feature—often called dynamic pricing—is generally understood to increase efficiency. By raising prices in periods of high demand (or low supply) and lowering prices in periods of low demand (or high supply), algorithms can increase output and overall welfare relative to prices that are held fixed. In this way, dynamic pricing shares parallels with other forms of flexible pricing, such as prices that vary geographically or personalized pricing.

What is less commonly understood are the implications of intertemporal spillovers resulting from dynamic pricing. Dynamic pricing differs from other forms of price discrimination is that individual consumers are often able to substitute across periods with little cost. Consumers may act strategically, timing their purchases to avoid periods with higher prices. Thus, with dynamic pricing algorithms, demand will endogenously respond to the time-varying nature of prices. Further, by shifting demand away from periods with high demand, firms that face capacity constraints or other discrete operational decisions can also increase welfare by reducing costs.

In this paper, we provide an empirical study of the adoption of a dynamic pricing algorithm in a setting with time-varying demand and capacity constraints. We exploit the staggered rollout of the algorithm to evaluate the impact of the algorithm on demand and welfare. We find that dynamic pricing significantly reduced demand volatility and increased the share of transactions during low-demand periods. Consumers appear to have responded to both across-day and within-day price variation, even at very high frequencies. On average, prices fell and output increased. Our results suggest that the adopting firm had more efficient utilization as a result of dynamic pricing, lowering costs and increasing overall welfare.

Our paper makes three distinct contributions. First, we document the adoption of a dynamic pricing technology and its impacts in practice. Though there is a growing literature on the theoretical effects of pricing algorithms, less is known about the design and impacts of algorithms that firms use in various contexts, such as retail markets.[1] Second, we present evidence that consumers strategically time their purchases, generating intertemporal demand spillovers at high frequencies and across days of the week. Third, we highlight how these spillovers may influence dynamic pricing decisions. We define and estimate a parameter, the volatility semi-elasticity (VSE), that reflects the response of time-series variation in quantities to time-series variation in prices. The parameter provides guidance on how much time-series variation in prices is needed to obtain a target reduction in demand volatility.

To motivate the analysis, we provide a stylized model in Section 2. The model introduces the interaction of time-varying prices and operational costs that arise from (e.g.,) capacity constraints. Through the model, we show how intertemporal substitution affects the incentives of a dynamic pricing mechanism. By reducing price in the low-demand period, consumer de-

[1] Important exceptions include studies of advanced-purchase markets, like airlines and hotels, and recent studies on rail-hailing platforms, as discussed with the related literature below.

mand is pulled away from the high-demand period. This is costly in that it reduces sales to high-margin consumers, but it can be beneficial if it reduces operational costs. Because of these two mechanisms, intertemporal spillovers have ambiguous effects on the desired amount of time-series price variation. Further, with intertemporal spillovers, it is not necessary to raise prices in periods of high demand to avoid capacity constraints. This is motivated by a feature of our setting: in practice, firms may place limits on the maximum price that can be charged in response to increased demand.

Our empirical setting is the staggered adoption of a dynamic pricing algorithm in a large restaurant chain in the Nordic countries. We observe high-frequency transaction data for each restaurant, which we describe in Section 3. With the adoption of dynamic pricing, the restaurant chain switched from a flat (or uniform) delivery fee to a time-varying delivery fee. Delivery represents a substantial portion of the restaurant chain's overall business and delivery fees represent approximately 15 percent of the average transaction price.[2] The algorithm that sets the delivery fee updates prices every ten minutes based on residual demand, historical orders, and target quantities.[3] Our high-frequency dataset covers two and a half years, from January 2020 through June 2022.

In Section 4, we examine how consumers respond to dynamic pricing. First, we use a pre-post analysis to look at the demand response across higher-demand and lower-demand periods. We show that, in every hour of the week, within-hour demand became less volatile after the implementation of dynamic pricing. Second, we show that consumers appear to strategically time their purchases at very high frequencies. At identical prices, quantities are higher than expected when the price 5 minutes prior was higher, and quantities are lower than expected if the previous price was lower. In other words, consumers seem to be following price changes at a high frequency and responding accordingly. Consistent with this, we document that website traffic increased substantially—more than the increase in order volume—after the introduction of dynamic pricing, and it increased by more during lower-demand days of the week. This suggests that consumers engaged in active price shopping after the adoption of the algorithm.

We evaluate the impact of the dynamic pricing algorithm on prices, quantities, and demand volatility by utilizing the staggered timing of adoption. The algorithm was rolled out in three distinct phases: to an initial group of restaurants in September 2020, to a second group in June 2021, and to the remaining restaurants in November 2021. We exploit this variation in timing by matching early adopters to late adopters based on empirical sales patterns. Due to the long panel of our data, we are able to match restaurants based on periods in which all restaurants employ the same technology: before any restaurant used dynamic pricing (July-August 2020) and also after all restaurants employed dynamic pricing (January-June 2022). We detail our empirical strategy in Section 5.

[2]Delivery is done in-house by employees and is not outsourced to a third party.
[3]The algorithm does not directly incorporate the prices of other restaurants.

Section 6 presents results. In our setting, the adoption of dynamic pricing lowered average transaction prices and created a great deal of variation within and across days. Average prices were higher during certain periods than others, but even within an hour prices could fluctuate up and down multiple times. We find that dynamic pricing reduced demand volatility across hourly periods and days of the week. The within-week coefficient of variation for hourly transactions fell by 10 percent, and the share of transactions occurring during the weekend (the peak demand days) fell by more than 3 percentage points. We estimate that the volatility semi-elasticity—defined as the percentage change in the volatility of quantities in terms of a change in the volatility of prices—is -1.96. Thus, a relatively modest increase in the time-series variation in prices led to a meaningful decline in demand volatility. We find that the total number of transactions increased, as well as revenues, though the revenue results are weaker. Overall, our findings of lower average prices, higher output, and reduced demand volatility suggest that dynamic pricing improved utilization for the adopting firm and increased consumer welfare.

Section 7 introduces an empirical model of consumer demand that builds on the conceptual framework from Section 2. Our estimates provide additional support for high-frequency consumer substitution. Further, the demand model allows us to measure consumer welfare, which we estimate to have increased in the dynamic pricing regime.

Section 8 concludes. Overall, our analysis highlights the implications of strategic consumers for dynamic pricing, including how frequently prices should change and how much volatility in prices is desirable.

Related Literature Our paper complements several empirical papers that study the effects of time-varying pricing and algorithms in various contexts. Our paper is unique in that we observe the switch from uniform pricing to a time-varying pricing algorithm in our data, and we use this data to document the empirical effects of a dynamic pricing algorithm. A related idea is studied by Assad et al. (2022), who explore the adoption of pricing algorithms in the context of retail gasoline. A key difference is that, in their context, retail gasoline stations engage in time-varying pricing whether or not they adopt an algorithm, though the algorithm may increase the frequency of price changes.

Empirical work on high-frequency dynamic pricing includes studies of advance-purchase markets, such as airlines (Lazarev, 2013; Williams, 2022; Aryal, Murry and Williams, 2022; Hortaçsu et al., 2021), hotels (Cho et al., 2018), and event tickets (Sweeting, 2012). In the context of ride hailing, Castillo (2020) studies the welfare impacts of real-time pricing that shifts both supply and demand. These papers typically model the arrival of consumers at a particular point in time as an exogenous process. A key contribution of our work, relative to these papers, is to explicitly consider intertemporal substitution and the endogenous impact of prices on demand volatility. The results in our paper indicate that consumers strategically time their purchases, and we show how firms can benefit from this behavior through dynamic pricing.[4]

Intertemporal spillovers in demand have been considered in settings without high-frequency dynamic pricing, such as video games (Nair, 2007; Lee, 2013), printers (Melnikov, 2013), and camcorders (Gowrisankaran and Rysman, 2012). These studies analyze contexts with durable goods where consumers can choose to delay purchases and prices are infrequently updated. Hendel and Nevo (2013) study weekly price changes (sales) in the context of storable goods. Though the goods in our setting are not storable, consumers who shift demand to other periods function in a similar way to storage. We complement this work by focusing on how intertemporal spillovers may affect costs, rather than screening among customer types.[5]

Finally, there is a growing literature studying the potential (anti)competitive effects of pricing algorithms (Calvano et al., 2020; Brown and MacKay, 2021). The algorithm we study does not explicitly incorporate the price of rivals, which is a key feature in the literature of algorithmic competition. In our analysis, we do not consider oligopoly interactions. However, we believe our study highlights additional considerations for future research on dynamic pricing in a competitive context.

2 Conceptual Framework

To provide a conceptual framework for our analysis, we introduce a stylized two-period model where a monopolist faces time-varying demand and capacity constraints. First, we examine the baseline scenario with no intertemporal spillovers in demand. We then show the implications of intertemporal spillovers for optimal pricing, even when the firm faces constraints on its ability to raise prices.

2.1 Setup

Consider a monopolist that faces demand across two periods $t \in \{1, 2\}$. These periods can be conceptualized as capturing shorter (e.g., 10 minute) or longer (e.g., daily) intervals. Let demand be linear and take the following form:

$$q_1 = \alpha_1 - \beta p_1 + \gamma p_2 \tag{1}$$

$$q_2 = \alpha_2 - \beta p_2 + \gamma p_1 \tag{2}$$

[4]A comparison can also be made to the analysis of real-time pricing in electricity markets. Jessoe and Rapson (2014) provide evidence that consumers exhibit intertemporal complementarities in this context; i.e., they *reduce* consumption across unaffected periods in response to higher prices during peak times.

[5]The theoretical literature considering time-varying pricing is more extensive, going back to at least Coase (1972) and Stokey (1979) in the context of durable goods. Dana (1999) studies intertemporal spillovers when demand peaks are unknown, which shares parallels with our empirical context. Dana (1999) considers a pricing schedule based on cumulative purchases to account for uncertain demand realizations, which is typical in advance-purchase markets and rare in retail settings.

Quantity demanded in each period is given by q_t. The parameter α_t incorporates demand shocks, which may vary across periods. Without loss of generality, let $\alpha_2 > \alpha_1$ so that period 2 is the high-demand period. Intertemporal spillovers are captured by the coefficient $\gamma \in [0, \beta)$. Thus, an increase in p_2 will reduce demand in period 2 and increase demand in period 1, and γ/β gives the share of marginal consumers from period 2 that are diverted to period 1. This ratio can be interpreted as the fraction of consumers that strategically time their purchases.

The monopolist faces constant marginal costs c in each market, as well as operational costs that are a function of quantity demanded, $\psi(q_1, q_2)$. We assume that these costs are increasing with quantity in each period, $\psi_1(q_1, q_2) \geq 0$, and $\psi_2(q_1, q_2) \geq 0$. We will impose a specific functional form below.

The monopolist maximizes profits by choosing a single price across both periods or by choosing different prices in each period (using a dynamic pricing algorithm). For clarity, we assume complete information, though the key tradeoffs persist in settings with incomplete information. Thus, the firm's objective is

$$\max_{p_1, p_2} q_1(p_1 - c) + q_2(p_2 - c) - \psi(q_1, q_2). \tag{3}$$

In the absence of dynamic pricing, the firm is subject to the constraint $p_1 = p_2 = \bar{p}$. We assume that the firm has perfect information over the realization of demand shocks α_t, though this is not necessary for the results.[6]

We assume that operational costs are discrete capacity constraints. For any quantity above a threshold $\bar{q}$, the firm bears an additional per-period operational cost ϕ for increased capacity. After paying for these costs, the marginal cost of production remains the same. These costs can be though of as the cost of hiring an additional employee to provide service in periods of high demand.[7] We will consider two possibilities:

- **Flexible operational costs**: The firm makes the decision to use higher capacity (at cost ϕ) separately in each period.

- **Inflexible operational costs**: The firm must make a single decision about capacity across both periods (paying ϕ in each period).

Thus, the operational cost function takes two possible forms:

$$\psi(q_1, q_2) = \begin{cases} \phi \mathbb{1}[q_1 > \bar{q}] + \phi \mathbb{1}[q_2 > \bar{q}] & \text{Flexible operational costs} \\ \\ 2\phi \mathbb{1}[\max(q_1, q_2) > \bar{q}] & \text{Inflexible operational costs} \end{cases}$$

[6] Due to the linearity of the demand system, maximizing the profits for the expected (mean) demand shock is equivalent to maximizing across the realizations.

[7] The costs can be generalized to incur an additional cost of ϕ for each multiple of $\bar{q}$ exceeded.

Figure 1: Dynamic Pricing Across Independent Periods

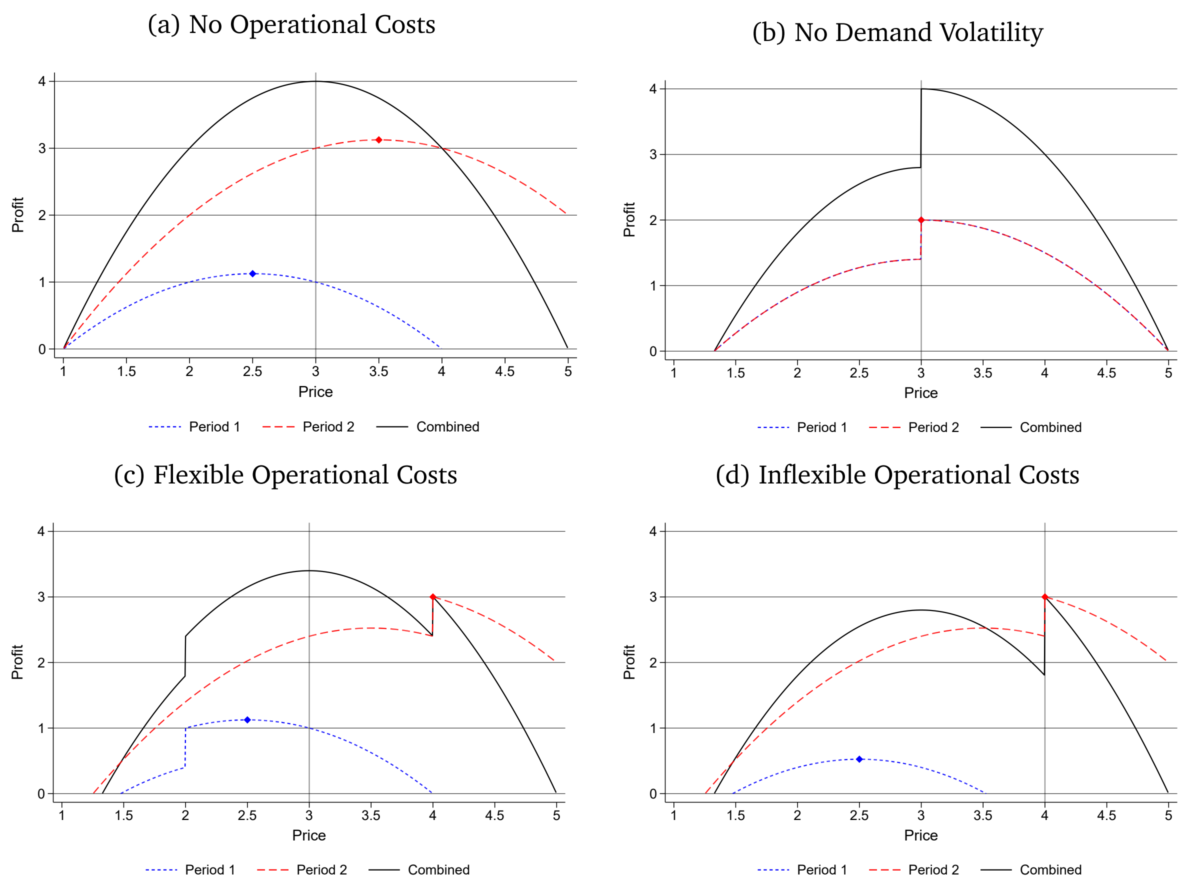

Notes: Period 1 profits are plotted with a blue dotted line, period 2 profits are plotted with a red dashed line, and combined profits are plotted with a solid black line. The vertical line indicates the profit-maximizing price when the firm must choose the same price in each period, and the diamond markers indicate the profit-maximizing values when different prices can be set in each period.

where $\mathbb{1}[\cdot]$ is the indicator function. In the second case, the firm must pay for higher capacity if it uses it in any period.

2.2 Dynamic Pricing with No Spillovers

In the case of no intertemporal spillovers ($\gamma = 0$), the two periods function as independent markets. We will evaluate this case as a baseline for price discrimination with and without operational costs.

Figure 1 plots the profits as a function of prices for the parameter values $(\alpha_1, \alpha_2, \beta, c) = (2, 3, 0.5, 1)$. Period 1 profits are plotted with a blue dotted line, period 2 profits are plotted with a red dashed line, and combined profits are plotted with a solid black line. The vertical line indicates the profit-maximizing price when the firm must choose the same price in each period, and the diamond markers indicate the profit-maximizing values when different prices can be set in each period. Four different scenarios are considered.

Panel (a) presents the case with no capacity constraints ($\phi = 0$). If the firm sets a single price across both markets, it is as if it faces aggregate demand $q = (\alpha_1 + \alpha_2) - 2\beta\bar{p}$. Given these parameters, the profit-maximizing price is 3. If the firm can use a dynamic pricing algorithm to respond to time-varying demand shocks, the firm would optimally set a lower price in the low-demand period and a higher price in the high-demand period, $(p_1, p_2) = (2.5, 3.5)$. With no operational costs, this simply reflects the willingness to pay of consumers in different periods and is a form of third-degree price discrimination.

Now we consider scenarios in which the firm faces operational costs ($\phi = 0.6$) for providing a quantity greater than $\bar{q} = 1$. As a baseline, panel (b) presents the case in which operational costs are present but there is no demand volatility. At the profit-maximizing price, the firm produces just at the capacity constraint.

Panels (c) and (d) indicate how capacity constraints interact with demand volatility and influence optimal pricing. In panel (c), the firm has flexible operational costs that can be turned on or off in each period. In this example, the presence of these costs does not influence the optimal uniform price ($\bar{p} = 3$), but it increases the optimal variation in prices when using dynamic pricing. Relative to panel (a), which shows no operational costs, the optimal price in the high demand period increases (from 3.5 to 4). The higher price enables the firm to stay within the capacity constraint and avoid additional operational costs. The optimal price in the low-demand period is unchanged.

Panel (d) provide the case in which operational costs are inflexible, in that they must be borne for both periods, if at all. In this example, the optimal prices for dynamic pricing are the same as in panel (c), at 2.5 and 4. However, this scenario indicates another potential effect of operational costs. If the firm does not have dynamic pricing technology, the firm would prefer to set a single price of 4 and only operate in the high-demand periods. At a price of 3, the firm must pay the operational costs in both periods, even though the extra capacity is only used in high-demand periods.

These scenarios indicate the potential benefits of dynamic pricing in the absence of intertemporal spillovers. In the presence of fixed operational costs, firms might use dynamic pricing to reduce the incidence of these costs (panel (a)) or operate in periods/markets that would not otherwise be profitable (panel (b)).

Note that dynamic pricing can lead to higher or lower average prices. In this example, with no operational costs or flexible operational costs (panels (a) and (c)), dynamic pricing leads to higher average prices because higher prices are charged during peak demand periods. The average price is 3.13 in panel (a) and 3.37 in panel (c). In panel (d), the average price is also 3.37, but this is lower than the uniform price with inflexible operational costs (4.00).

2.3 Dynamic Pricing with Intertemporal Spillovers

We now consider the case in which consumers may substitute across periods ($\gamma > 0$). With these spillovers, the firm now faces two additional considerations that have opposite implications for the optimal levels of price volatility. The first is the direct effect of the demand spillovers. If the firm charges a lower price in the low-demand period, it pulls demand away from the high-demand (high-margin) period. This effect would tend to reduce the optimal variation in prices across periods, relative to an environment with no spillovers in demand. In the absence of operational costs, the optimal pricing strategy would balance the marginal benefit of lower prices in the low-demand period with the cost of fewer consumers in the high-margin period.

The second consideration is the indirect effect resulting from capacity constraints. If the firm can pull enough demand away from the high-demand period, it can reduce its operational costs. This incentive may be strong enough to have the firm lower its price by more than would otherwise be optimal. This effect can lead to greater variation in prices across periods.

Figure 2 illustrates these incentives. Similar to Figure 1, period 1 profits are plotted with a blue dotted line, period 2 profits are plotted with a red dashed line, and combined profits are plotted with a solid black line. However, the figure has two key differences. First, period 2 prices are held fixed at the optimal uniform price that yields profitability in both periods ($p_2 = 3$), which is indicated by the vertical line. The x-axis thus only represents the change in the period 1 price, holding fixed the price in period 2.

Panel (a) indicates the incentives with no operational costs. The price the maximizes period 1 profits, conditional on $p_2 = 3$, is $p_1 = 2.2$. However, taking into account the intertemporal spillovers, the firm instead would price at $p_1 = 2.5$, which maximizes the combined profits (black line).

Panel (b) introduces flexible operational costs. This provides an additional incentive for the firm to lower price in period 1, as it could pull enough demand from period 2 in order to avoid paying for extra capacity in that period. Thus, the optimal price in period 1 is even lower, approximately 1.67. This yields a discrete jump in combined profits (black line) while period 1 profits change continuously. Lowering the price further than this does not benefit the firm because of the direct effect on period 2 profits and the potential indirect effect of having to pay higher operational costs in period 1.

Panel (c) presents the same scenario with inflexible operational costs. The optimal pricing decision is the same as in panel (b). The only difference to note is that the period 1 and overall profit incentives are aligned, because the firm also sees a discrete jump in profits in period 1 at $p_1 = 1.67$.

Figure 2: Dynamic Pricing with Intertemporal Spillovers

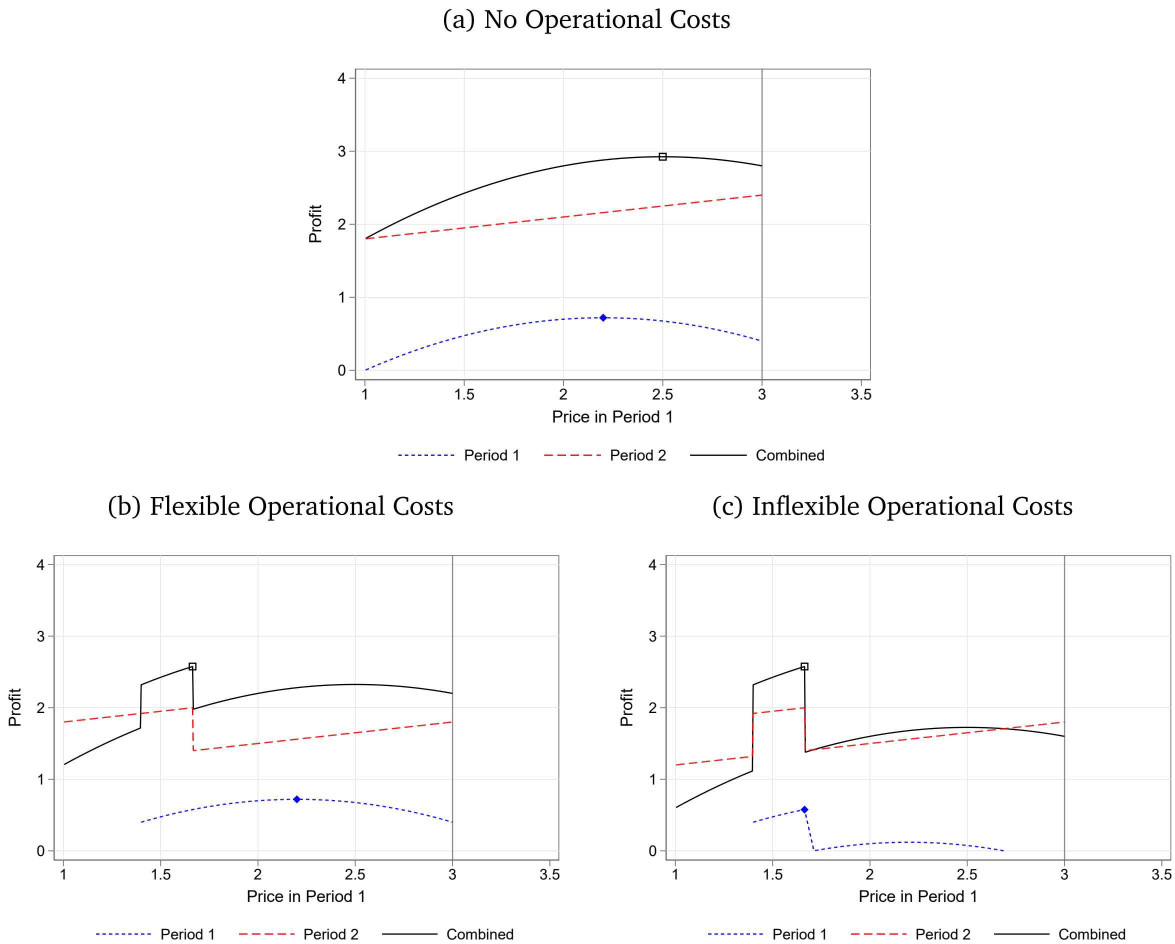

(a) No Operational Costs
(b) Flexible Operational Costs
(c) Inflexible Operational Costs

Notes: Period 1 profits are plotted with a blue dotted line, period 2 profits are plotted with a red dashed line, and combined profits are plotted with a solid black line. The vertical line indicates the profit-maximizing price when the firm must choose the same price in each period, and the diamond markers indicate the profit-maximizing values when different prices can be set in each period.

2.4 Discussion

The stylized model captures several of the key considerations in dynamic pricing. First, firm may use time-varying pricing as a form of third-degree price discrimination to increase revenues and profits. Second, time-varying pricing may also be used to avoid operational costs that would otherwise have to be borne with invariant/uniform pricing. When these costs are inflexible, dynamic pricing may even enable the firm to operate at times that it would not otherwise choose to do (i.e., stay open longer hours).

Finally, the presence of intertemporal spillovers has additional implications for dynamic pricing. The direct effect of these spillovers is to mitigate the difference in prices between low-demand and high-demand periods. A firm that prices too low in the low-demand period gives up too much profit in the high-demand period. However, the presence of operational costs may provide an incentive to increase the time-series variation in prices. Firms may want to price even lower in the low-demand period in order to pull demand away from the high-demand period and stay within capacity constraints. Overall, these two forces show that the presence of intertemporal spillovers has ambiguous effects on the optimal levels of price volatility.

Our analysis in Section 2.3 focused on a scenario in which the high-demand period price was held fixed. Though this simplification is useful to illustrate the incentives and tradeoffs, it also was motivated by real-world considerations and our empirical context.

In some settings, the adoption of dynamic pricing may come with an added constraint of a price cap. Indeed, in our empirical setting, the maximum price charged was constrained by the adopting firm to a level that was similar to that prior to the adoption of dynamic pricing. In our setting, we understand that the firm did not want to risk consumer backlash from dynamic pricing, so placing limits on the maximum price was a key consideration. In addition to avoiding consumer backlash, firms in other context may wish to avoid the scrutiny of policymakers. For example, surge pricing for Uber has received a lot of negative attention.[8]

We highlight that our analysis shows that even without the ability to raise prices above the optimal uniform price, dynamic pricing can pull demand away from high-demand periods and reduce operational costs through the presence of intertemporal spillovers.

3 Empirical Setting

3.1 Data

Our primary data comes from a single restaurant chain located in the Nordic countries. The chain consists of approximately 300 restaurants. It is the largest chain in its category and one of the most popular restaurant brands. The chain offers dine in, take away, and home delivery. The vast majority of orders are for take away and home delivery. Home delivery is produced in-house with a proprietary online ordering system, and it is serviced with branded vehicles by restaurant employees. During our sample period, 228 of the restaurants offered the home delivery option. For these restaurants, take away and home delivery have roughly equal revenue shares.

We obtained high-frequency transaction-level sales data for each restaurant. The data come from the restaurant chain's online ordering system, and they include revenue, type of order (delivery, take away, and eat in), and a timestamp. The online ordering system is the primary means of ordering delivery and take away, and these two order types constitute the vast majority of revenues in our sample. We combine this data with the full panel of delivery fees, allowing us

[8]See, e.g., "Uber forced to suspend surge pricing in Delhi," https://www.theverge.com/2016/4/21/11477038/uber-surge-pricing-delhi-india. Of course, a change in the reference price can shift the classification of high-price periods to the baseline price and low-price periods to discounts.

to observe the prices even when no purchases were made. For much of our analysis, we collapse the data to the restaurant-week level, which allows us to construct within-week measures of volatility. Our dataset consists of 287 restaurants and spans two and a half years, from January 1, 2020 through June 28, 2022.

We supplement this data with data on website traffic obtained from SimilarWeb. The data contains daily page views on the restaurant chain's website, which includes the online ordering system for delivery.

3.2 Adoption of a Dynamic Pricing Algorithm

Before implementing dynamic pricing, the price of a home delivery order was fixed at EUR 5.90. Demand for home delivery varied over time. Demand was typically higher from Friday to Sunday than from Monday to Thursday. During weekdays, demand typically peaked around noon ("lunch") and late afternoon ("after work"). Despite these general patterns, demand at any particular point of the week could vary substantially week-to-week.

As discussed in Section 2, the uniform pricing rule potentially led to substantial operational costs. The restaurant chain maintained a capacity that could meet higher levels of demand, yet this capacity was idle during periods of lower demand. For example, one key operational decision was how many employees to have at each location. Employees could be an inflexible input cost when hired on a full-time basis.

In order to address the costs associated with demand volatility, the restaurant chain adopted a dynamic pricing algorithm. The chain used a staggered rollout in order to pilot the technology and observe the impact of the algorithm on sales and profits. The algorithm was turned on for an initial group of 49 restaurants on September 7, 2020. A second group of 16 was activated on June 14, 2021, and the technology was rolled out to the remaining 58 qualifying restaurants on November 19, 2021.[9] Each group covers a distinct geographic region.

Restaurants that adopted the algorithm switched from a fixed delivery fee to one that was updated every ten minutes. Initially, prices were allowed to vary from EUR 2.90 to 5.90 Monday-Thursday and from 2.90 to 7.90 for weekends (Friday-Sunday). In October 2021, some restaurants increased the upper price limit to EUR 8.90 for weekends.[10]

3.3 Description of the Algorithm

The dynamic pricing algorithm adopted by the restaurant chain was procured from Priceff, Ltd., a third-party provider. Though we cannot disclose the exact calculations of the algorithm, we

[9] A handful restaurants did not activate the technology on these exact dates; we ignore them in our analysis.

[10] In April 2022, three restaurants located at skiing resorts further increased the upper limit to EUR 11.90. These restaurants are excluded from our analysis sample.

Figure 3: Example of Algorithm Prices

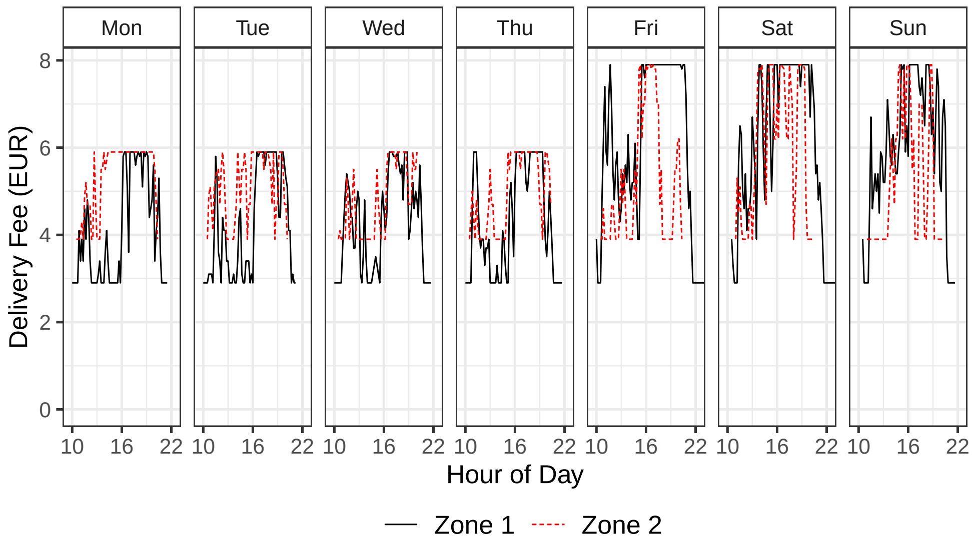

Notes: Figure displays the time series of delivery fees for two different pricing zones over a single week in January 2022. Delivery fees are updated every 10 minutes by the pricing algorithm and represent roughly 15 percent of the total transaction price.

can describe the features of the algorithm qualitatively.[11]

The pricing algorithm updates prices in response to residual demand. A seller (in this case, an individual restaurant) sets a target quantity to sell within a set timeframe. If quantity demanded is too low, the algorithm will decrease the price, and vice versa if quantity demanded is too high. The algorithm makes these judgments based on recent purchases (including those within the previous ten minutes), along with historical purchases, previous price changes, characteristics of the data sample, and other factors. The algorithm forms a prediction of how demand will evolve in order to meet the target, and it adjust the price accordingly. The pricing algorithm has the capability to evolve over time. Notably, the algorithm does not directly take into account the prices of competing restaurants, which is a key consideration in studies of algorithmic competition (Calvano et al., 2020; Brown and MacKay, 2021).

In principle, the algorithm allows for restaurant-specific dynamic pricing. However, restaurants in the same city were put into the same pricing zones, as the restaurant chain did not want customers in the same city to see different prices. The algorithm pooled data across restaurants within the same zone.

Figure 3 provides examples of the time series delivery fees as set by the algorithm for one week of January 2022 and two different pricing zones. The figure illustrates that there is substantial within-day variation in the delivery fee. On average, higher delivery fees are realized

[11]The algorithm builds on Ekholm (2019), which presents a "Method, system, and computer program product for dynamically pricing perishable goods."

on the weekend. Within a day, there are typical times with higher fees, though there is variation in fees during these periods across weeks and restaurants.

The rollout of the algorithm was transparent to affected consumers that visited the restaurant's website. Before placing an order, a consumer had to enter a delivery address. If the consumer was in a zone with dynamic pricing, the consumer saw a statement that delivery fees could change every 10 minutes, the range of possible fees, and the current fees in their location.[12]

4 How Do Consumers Respond to Dynamic Pricing?

In this section, we explore how consumers respond to dynamic pricing using high-frequency transaction data. First, we examine how the algorithm affected quantities and volatility across periods of higher and lower demand. Second, we look at changes in order volume in short intervals around the timing of price changes. Third, we look at how website traffic volume changed after the implementation of the algorithm. The evidence suggests the consumers strategically time their purchases at a very high frequency (within ten minutes) and across days of the week.

4.1 Hour of the Week Analysis

For this analysis, we conduct a pre-post analysis across the 65 typical hours of business for the restaurants in our sample: 11 AM to 8 PM from Sunday to Thursday, and 11 AM to 9 PM on Friday and Saturday. For each of these hours, we calculate statistics across restaurants and weeks for two time periods. The "pre" period represents January through August 2020, before any restaurant had dynamic pricing. The "post" period corresponds to January through June 2022, when all of the restaurants in our sample had dynamic pricing. We aggregate across restaurants in all three groups in our sample.

To illustrate how consumer behavior shifted within the week and within hours in the week, we examine quantities, prices, and within-hour volatility. The within-hour volatility is constructed as follows: we partition each hour into six 10-minute blocks, and calculate the number of transactions in each block. We then calculate the within-hour coefficient of variation (standard deviation divided by the mean) for each of the 65 hours by restaurant-week.

Figure 4 presents the results. On the x-axis, the 65 hours of business are sorted by the mean transaction volume in the pre period (2020). The pre-period values serve as a proxy for demand; given no per-period variation in prices, moving toward the right indicates that demand was higher during those hours of the week (e.g., Friday and Saturday nights). Post-period values are plotted with a solid black line, and pre-period values are plotted with a gray

[12] During the sample, the restaurant chain engaged in occasional promotional periods with free or steeply discounted delivery. These promotions affected all restaurants, and we exclude the affected weeks from our analysis. Promotional activity did not vary materially across restaurants.

Figure 4: Pre-Post Impact by Hour of Week

(a) Transactions

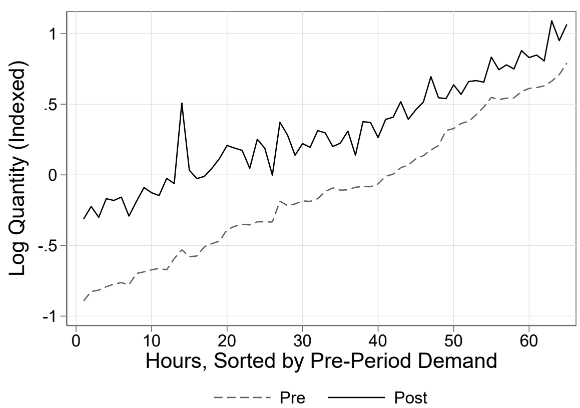

(b) Delivery Fees

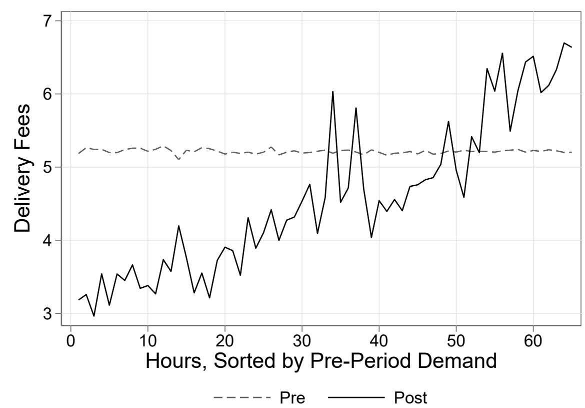

(c) Within-Hour Volatility

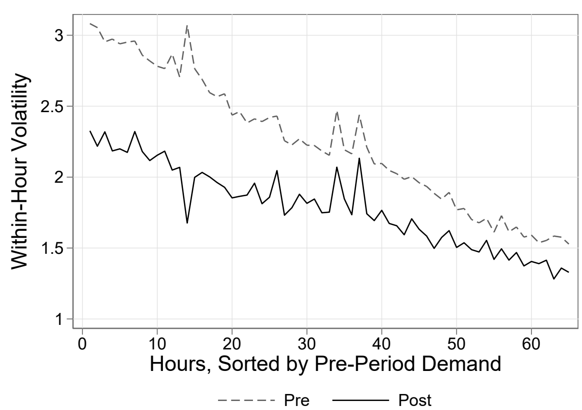

Notes: Figure displays the pre-period (dashed line) and post-period values (solid line) for outcomes across all treated restaurants in our sample. The *x*-axis corresponds to the 65 typical business hours of the week. These hours are sorted by mean quantities from low-to-high. Panel (a) plots the transactions, panel (b) plots the delivery fees, and panel (c) plots the within-hour volatility, which is defined as the coefficient of variation across the six 10-minute blocks within each restaurant-week-hour.

dashed line. Panel (a) shows that, by 2022, quantities increased for every hour of the week. Some of the growth in transactions was likely due to overall restaurant growth and recovery from the initial measures of the COVID-19 pandemic. However, the growth was not uniform across hours of the week. The increase in quantities was proportionally higher for hours that initially had lower volume.

Panel (b) shows that the mean delivery fees were flat across hours of the week in the pre period.[13] Dynamic pricing, as shown by the post-period values (solid line), decreased prices during lower-demand periods and increased prices on average during higher-demand periods.

[13]The mean value does not equal 5.90 exactly because of promotional periods.

On average, prices fell. Even during high-demand periods, where prices increased on average, the algorithm was able to generate lower prices during high-frequency windows when demand was lower than average. As indicated by Figure 3, the price often fell below 4.00, even during peak times. This high-frequency response of prices allowed for additional transactions to occur during high demand periods, and may explain in part why the periods where prices increased on average also saw an increase in quantity.

Panel (c) provides more direct evidence of a demand-smoothing effect. In the post period, within-hour volatility fell for every hour of the week. The reduction in volatility was greater for low-demand periods, consistent with the notion that the algorithm was able to shift demand to times when the restaurants were less capacity constrained. The reduction in volatility during high-demand periods suggests that the growth in transactions during those periods was facilitated by high-frequency variation in prices that smoothed out demand.

4.2 High-Frequency Strategic Timing

Above, we show that within-hour volatility falls for every hour in the week, suggesting that the algorithm can smooth out demand within the hour. Here, we present additional evidence that consumers strategically time their purchases at a high frequency.

The dynamic pricing algorithm we study updates prices every 10 minutes. On average, the algorithm increases prices during periods of higher demand. Thus, as can be seen in panels (a) and (b) of Figure 4, prices and quantities are positively correlated in the data. However, if some consumers are strategically timing their purchases, we could observe a discrete change in behavior around a change in price. All else equal, we might expect that consumers that want to pay lower prices are more likely to purchase immediately after a price decrease than immediately after a price increase.

To explore this possibility, we divide our data into 10-minute periods around every potential price change (e.g., 10:55 to 11:05). We then measure the number of transactions occurring in the 5 minutes prior to an opportunity to change prices (10:55 to 11:00) and the 5 minutes afterward (11:00 to 11:05), as well as the delivery fees. We normalize the mean number of transactions to 100.

We then calculate the mean number of transactions occurring for a given delivery fee (e.g., EUR 4.90) based on whether the fee in the preceding 5 minutes (10:55 to 11:00) was higher, lower, or the same value. If demand were completely static, the fee in the previous time period should have no effect on the quantity demanded in the current period. If consumers were not strategic, then, based on the objective of the pricing algorithm, we would instead expect that price increases would correspond to higher quantity demanded. If consumers are strategic, this could mitigate the overall price-quantity correlation generated by the algorithm, or even reverse the sign when looking at a small enough window.

For each delivery fee from 2.90 to 7.90, we calculate the change in transactions before and

Figure 5: High-Frequency Intertemporal Substitution

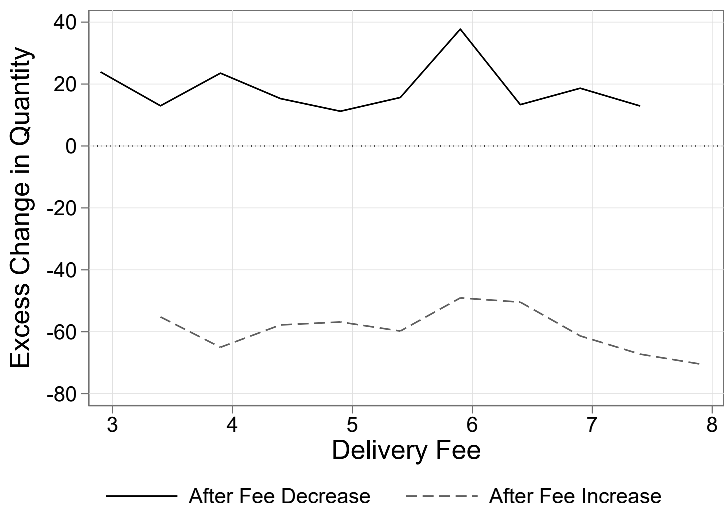

Notes: Figure displays the excess change in quantity for the 5-minute period after each price change opportunity relative to the quantity in the 5-minute period prior to each price opportunity. The data are presented by the current delivery fee (*x*-axis) and are relative to the (fee-specific) change in quantity after no price change. The solid line plots the excess change in quantity following price decreases, and the dashed line plots the excess change in quantity following a price increase.

after a price change opportunity when no price change actually occurs. We then calculate the excess change in quantity as the change in quantities relative to these fee-specific values. Figure 5 plots the results. The solid line shows that, for any current price level, the excess change in quantity is positive after a delivery fee decrease. In other words, more consumers make purchases at identical prices when the price 5 minutes prior was higher. Likewise, the dashed line shows that consumers make fewer purchases at a given price when the price immediately prior was lower. The magnitudes are economically significant: quantities increase by roughly 20 percent of the mean rate of transactions after a price decrease and fall by roughly 60 percent of the mean rate following a price increase.[14]

We use these data to calculate the high-frequency price elasticity. We regress the log change in quantity on the log change in prices for each 10-minute block in the data. We obtain an elasticity of -13.0 (std. err. 0.791), indicating that consumers are very sensitive to high frequency changes.[15]

This evidence is consistent with strategic timing by consumers. Within 10-minute intervals, consumers appear to be less likely to purchase when prices have increased and more likely to purchase when prices have decrease, even when these consumers face the same prices. We note

[14]The differential responses to increases and decreases should not be interpreted as a purely asymmetric behavioral response. They reflect, in part, that the average fee decrease may be smaller than the average fee increase.

[15]We obtain a delivery fee elasticity of -1.65 (std. err. 0.097). We find some differences in elasticities after positive (-11.3) and negative (-14.8) price changes.

Figure 6: Website Activity Increases with Algorithm Adoption

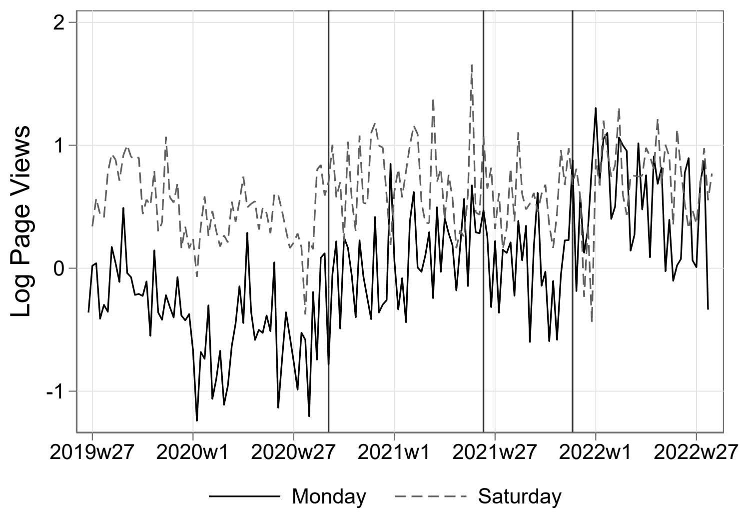

Notes: Figure displays the daily website traffic for Mondays (solid line) and Saturdays (dashed line), which are the days with the lowest and the highest demand, respectively. The y-axis captures log page views. The vertical lines indicate the staggered adoption dates for the pricing algorithm. By the end of the sample, all eligible restaurants had dynamic pricing.

that the high-frequency elasticity may be quite different from a demand elasticity measured at a lower frequency. For example, consumers may be more willing to "shop" for a 30-minute period than they are to substitute away from the restaurant entirely.

4.3 Website Traffic

If consumers are strategically timing their purchases, we should expect their behavior to change after the implementation of dynamic pricing. Before, prices were fixed over the course of the week, and there was no opportunity to take advantage of variation in prices.

We use supplementary data on website traffic to look at changes in consumer behavior after the adoption of the algorithm. Specifically, we obtain daily page views on the website (which includes the online ordering system). We then split the data by the day of the week. Here, we focus on Mondays and Saturdays, which are the days with the lowest and highest demand, respectively.

Figure 6 plots the log daily page views for each week on Mondays (solid line) and Saturdays (dashed line) from July 2019 through July 2022. The logged values are indexed to the mean log page views across all days of the week from July 2019 through August 2020, before dynamic pricing was adopted by any restaurant. Prior to dynamic pricing, website traffic on Mondays was, on average, 0.40 log points below the mean, while traffic on Saturdays was 0.48 log points above the mean. Across days of the week, website traffic is positively correlated with the number of delivery orders.

The vertical black lines in the figure indicate the staggered rollout dates for dynamic pricing. As more restaurants gained access to dynamic pricing, website traffic overall increased, and the gap between Monday and Saturday traffic decreased. By the end of the sample, Monday and Saturday traffic had nearly converged. In 2022, Monday traffic averaged 0.55 log points above the per-period mean (an increase of 0.95 log points), and Saturday traffic had an average log value of 0.78 (an increase of 0.30).

Overall, website traffic increased substantially after the adoption of dynamic pricing. The observed patterns indicate that consumers do engage in strategic timing. The increase in website traffic was greater than the increase in transactions, i.e., consumers spend more time shopping on the website without purchasing. Given the previous findings, this suggests that the website activity is due to consumers looking for lower prices. Moreover, the substantial increase in website traffic during lower-demand days indicates that consumers are timing their purchases across days. Overall, this analysis provides additional evidence for intertemporal substitution by consumers.

5 Impacts of Dynamic Pricing: Empirical Strategy

5.1 Overview

To evaluate the impact of the adoption of a dynamic pricing algorithm on sales and demand, we exploit the fact that the technology was adopted by different restaurants at different times. The staggered rollout allows us to compare restaurants with dynamic pricing to those without it at the same point in time.

To account for heterogeneity across groups based on adoption timing, we use a matching estimator. We match each "treated" restaurant to similar restaurants that adopted dynamic pricing at a different period. Importantly, we determine similarity between restaurants based on sales patterns during periods when all restaurants use the same technology. Specifically, we make use of both data from 2020, prior to the adoption of dynamic pricing by any restaurant, and data from 2022, after all restaurants use dynamic pricing. We use the staggered rollout over the two-year window between these periods to evaluate the impact of dynamic pricing.

5.2 Matching Procedure

To measure changes in outcomes due to dynamic pricing, we exploit the staggered rollout of the algorithm and match restaurants to other restaurants that adopt the algorithm at a different period but have similar revenue patterns during the matching windows. We calculate the impact of dynamic pricing by examining the window where one group ("treated") has dynamic pricing but the matched group ("control") does not.

For each treated restaurant, we use weekly sales patterns to identify the most comparable control restaurants. Specifically, we use sales from periods July 6, 2020 through September 6, 2020 (nine weeks) and January 10, 2022 through June 27, 2022 (23 weeks). The first window ends immediately before the adoption of dynamic pricing by the first group of restaurants. We use a shorter window in 2020 due to interruptions to business by the COVID-19 pandemic, though, as seen in later figures, our matched controls track the treated groups fairly well throughout the first half of 2020. Our results are not sensitive to the particular matching window we employ.

Within each window, we calculate the mean log weekly revenues and the deviation from each year-specific mean. This yields two means and 32 weekly "seasonality" matching variables. For each treated restaurant and potential control, we calculate the difference between each of these variables. To account for idiosyncratic closures, we drop two seasonality variables with the largest difference between each potential pair. We then calculate distance as the sum of squared differences across the remaining 32 matching variables, where we place equal weights on seasonality and the means (i.e., 0.25 on the 2020 mean, 0.25 on the 2022 mean, and 0.50 on week-to-week variation from the means). We use this distance to select the three nearest neighbors for each treated restaurant, allowing control restaurants to be matched to multiple treated restaurants.

5.3 Measuring Effects

We use mached controls from the nearest-neighbor procedure to construct counterfactual outcomes. Let Y_{it} denote an outcome of interest (e.g., revenue) for restaurant i in period t, where $t = 0$ corresponds to the week that dynamic pricing was implemented. Let $Y_{it}(1)$ indicate the outcome with dynamic pricing and $\widehat{Y}_{it}(0)$ indicate estimated counterfactual without dynamic pricing. Given $Y_{it}(1)$ and $\widehat{Y}_{it}(0)$, we can obtain a restaurant-specific estimate of the effect of dynamic pricing on the outcome, $\widehat{\Delta Y}_{it}$:

$$\widehat{\Delta Y}_{it} = Y_{it}(1) - \widehat{Y}_{it}(0). \tag{4}$$

We observe the outcome $Y_{it}(1)$ for the restaurants in our data. The counterfactual outcome, $\widehat{Y}_{it}(0)$, is unobserved and is calculated as follows. For each treated restaurant i, we select three nearest neighbors using the above procedure. We calculate the counterfactual outcome, $\widehat{Y}_{it}(0)$, as the average value of $Y_{it}(0)$ across the three matched control restaurants. This provides a restaurant-specific estimated effect for each period. In contrast to typical "difference-in-differences" specifications, we do not normalize $\widehat{Y}_{it}(0)$ to zero out pre-period differences. Instead, we report the unadjusted means, which are fairly close in most cases.

To quantify the average impact of dynamic pricing across restaurants, we take the average of the restaurant-specific treatment effects:

$$\widehat{\overline{\tau}}_t = \frac{1}{N_g} \sum_i \widehat{\Delta Y}_{it}, \tag{5}$$

where N_g is the number of restaurants in the treatment group. We plot the mean values of $Y_{it}(1)$, $\widehat{Y}_{it}(0)$, and $\widehat{\overline{\tau}}_t$ to diplay the trends in the data before and after the adoption of dynamic pricing.

To report overall effects, we use the regression specification

$$Y_{it} = \beta a_{it} + \delta_i + \gamma_t + \varepsilon_{it} \tag{6}$$

where a_{it} is an indicator of whether or not restaurant i has dynamic pricing in period t, and δ_i and γ_t provide restaurant and period fixed effects, respectively. The parameter β provides an estimate for the effect of treatment on the treated for a particular outcome of interest. We use a weighted regression where the weights are equal to 1 for each treated restaurant. For control restaurants, weights are proportional to the number of times it is matched to a treated restaurant and are scaled so that treated and control groups receive equal total weight. Thus, we use nearest-neighbor matching to assign weights for the regression specification. Restaurants that are not similar to the treated restaurants are assigned a weight of zero.

We note that this "two-way fixed effects" design that we employ has come under scrutiny when applied to settings with variation in treatment timing (see, e.g., Goodman-Bacon, 2021; Baker, Larcker and Wang, 2022). We report the matching event-study estimates based on equation (5), and these estimates suggest that the effects of algorithms are fairly constant over time, mitigating concerns of potential bias as described by Baker, Larcker and Wang (2022). Second, we exploit the discrete timing of the staggered rollout to generate distinct subsamples for analysis, as we describe below.

To summarize the response of demand volatility to time-varying prices, we define and estimate a volatility semi-elasticity (VSE). Let V_Q denote the within-week coefficient of variation for quantities, which we calculate separately for each restaurant-week as the standard deviation of the number of transactions in each hour divided by the mean. Let V_P denote the within-week coefficient of variation in prices faced by consumers. The VSE is given by:

$$VSE = \frac{dlnV_Q}{dV_P}. \tag{7}$$

This measure reflects the causal response of demand volatility to an increase in the time-series variation of prices.[16] As with analogous measures like demand elasticities, the parameter reflects the specific scope of the measurement. For example, the weekly VSE likely has a different

[16] We use a semi-elasticity rather than an elasticity because, in many cases, the initial value of V_P is 0 and any percent change from 0 would be undefined.

Table 1: Summary Statistics

Variable	Group A	Group B	Group C	Matched Controls: A	Matched Controls: B
Indexed Revenue	100.0	99.4	86.0	103.7	98.8
Indexed Revenue, Jul–Sep 2020	72.5	69.1	50.4	68.2	67.1
Revenue per Transaction	33.35	32.21	32.07	31.41	31.58
Delivery Fees per Transaction	5.06	5.32	4.73	4.67	4.67
Share of Revenue on Fri–Sun	0.62	0.61	0.61	0.61	0.60
Number of Restaurants	49	16	57	44	35

Notes: Table displays the summary statistics for the three groups in our analysis sample. Each group is determined by the date they implemented the dynamic pricing algorithm. Reported statistics are the mean across restaurants and weeks. The last two columns provide statistics for the matched control groups, which are sampled with replacement from group C.

value than the daily VSE. The parameter also reflects the specific nature of the variation in prices. In our context, the estimate reflects how prices are updated by the dynamic pricing algorithm, which likely yields a different demand response than would, say, the introduction of random price variation of an equivalent magnitude. The VSE can be interpreted as how much a marginal increase in coefficient of variation for prices affects the coefficient of variation for transactions, while holding fixed the rules for determining price variation.

5.4 Exploiting the Staggered Rollout

The staggered rollout consisted of three distinct phases: an initial group of restaurants (group A) that adopted dynamic pricing in September 2020, a second group of restaurants (group B) that adopted in June 2021, and the remaining restaurants (group C) that adopted in November 2021. We define restaurants that belong to in groups A and B as belonging to the treated group of restaurants, and restaurants that belong to group C as controls.

We exclude restaurants that (a) did not offer delivery during our sample period, (b) started the pricing algorithm outside of the three key dates identified above, and (c) had 4 or more weeks with delivery revenues less than EUR 100 in our sample period. After these filters, we retain 122 restaurants: 65 in the treated group (49 in A and 16 in B), and 57 in the control group.

Table 1 provides summary statistics for the three groups. All statistics reflect the mean weekly values for 2022 (January through June), except for the one row which reports mean weekly revenues from July through September 2020. To mask exact revenue figures, we have indexed the values so that the group A mean in 2022 is equal to 100.

Restaurants in group A have, on average, similar 2022 revenues as those in group B. The average restaurant in group C is fourteen percent smaller, in terms of revenues, than the other two groups. All groups have seen revenue growth since 2020. In percentage terms, revenue growth has been faster for group C than groups A and B. We attempt to mitigate these differences by using both 2020 and 2022 sales data in our matching procedure. In terms of other statistics, all groups look similar in 2022. The average transaction is roughly EUR 32 in revenue, with 15 percent of revenues coming from delivery fees (EUR 5 per transaction). Three-fifths of the revenue comes from the weekend period of Friday through Sunday.

The last two columns report summary statistics for the matched control restaurants, which are selected from group C. The matched controls are much more similar to the treated groups. The fifth column shows that the matched control revenues for group B are very similar to the mean group B (indexed) values of 99 in 2022 and 69 in 2000. Likewise, the mean revenues in the matched controls for group A are 68 in 2020 and 104 in 2022, which are much closer to the group A mean than the group C overall average. Despite improving the match, the matched controls for A still realize faster revenue growth than group A. This shows up in some of our results but does not seem to impact our estimated coefficients much, perhaps because the key identifying variation is in the middle of the sample. Moreover, we obtain similar results when using group B and its matched controls only.

For our regression analysis, we restrict the sample to range from September 7, 2020 through June 26, 2022. We do this for two reasons: first, after this restriction, group A has dynamic pricing for the entire sample, and the only identifying variation within the matched cohort occurs when C adopts dynamic pricing in November 2021. Thus, this restriction creates a subsample with a single treatment timing, eliminating concerns from staggered timing in two-way fixed effects models. For robustness, we report the group A and group B effects separately in addition to reporting the pooled estimates. The second reason is that the onset of the COVID-19 pandemic created some disruptions in the business, which we wish to avoid when reporting overall effects.

6 Impacts of Dynamic Pricing: Results

6.1 Price Effects

First, we evaluate the effects on delivery fees and overall prices. Panel (a) of Figure 7 plots the mean delivery fee by group in our data. After the adoption of dynamic pricing, average delivery fees in most weeks were lower than the uniform fee charged by restaurants without dynamic pricing. This was in part by design, as the Monday-Thursday fee could only go lower than the uniform fee, and the Friday-Sunday fee was capped at a maximum value of EUR 7.90 to 8.90. The two large downward spikes in 2022 reflect promotional activity by the restaurant chain; during these periods, the average delivery fee chosen by the algorithm was temporarily set to be much lower.[17]

[17] In addition, gaps in the time series before 2022 reflect promotional periods of zero delivery fees across all restaurants.

Figure 7: Delivery Fees by Treatment Group

(a) Mean

(b) Standard Deviation

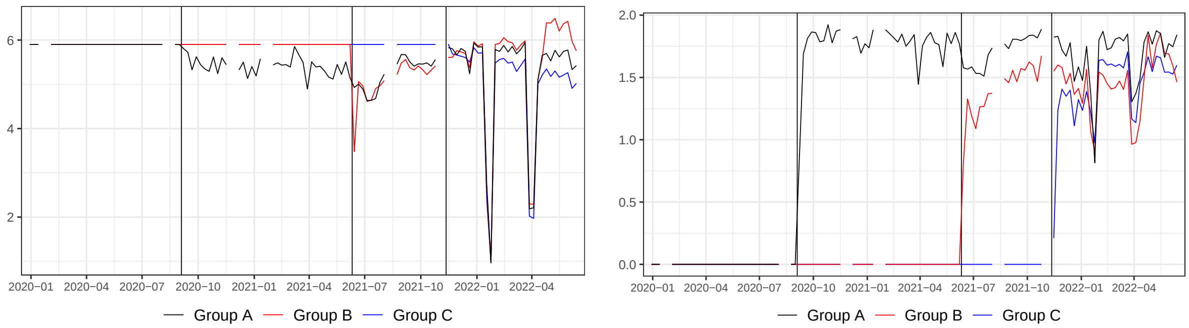

Notes: Figure displays statistics of delivery fees by treatment group over the sample period. Panel (a) displays the weekly mean, and panel (b) displays the average within-week standard deviation across restaurants in each group. Gaps in the time series correspond to promotional periods with zero delivery fees across all restaurants.

Panel (b) plots the average within-week standard deviation of fees across restaurants within each treatment group. Prior to the adoption of dynamic pricing, this is exactly zero. Immediately after the adoption of dynamic pricing, the standard deviation of fees within a week jumps to nearly EUR 2.00, and typically fluctuates between EUR 1.00 and 2.00 over the rest of the sample. The mean weekly standard deviation of fees for restaurants using dynamic pricing is 1.60.

Since there are independent pricing zones within each group, the above figure masks a good deal of within-group heterogeneity. In Appendix Figure A.1, we provide a version of the figure at the restaurant level, instead of the group level. The figure illustrates that the within-group heterogeneity is often larger than the across-group heterogeneity.

Table 2 summarizes the overall effect on prices following regression equation (6). Observations are at the restaurant-week level. Columns (1)-(3) display the results when the dependent variable is the mean log transaction price, including delivery fees. Columns (4)-(6) show the within-week standard deviation of log transaction price. As described in the previous section, we show results separately for the pooled treatment and for groups A and B separately. Columns (1) and (4) correspond to the pooled treatment across groups A and B. Columns (2) and (5) correspond to group A and columns (3) and (6) to group B.

Overall, we estimate that the algorithm led to a 5 percent decrease in average transaction prices, while the standard deviation of log prices increased by 0.03. We find that our estimates are similar across treatment groups and highly significant. As describe above, a reduction in prices from Monday through Thursday was anticipated, because the delivery fee could only fall relative to the uniform fee. Average transaction prices fell by 8.3 percent during these days.[18] Over the weekend, fees could go higher than the uniform fee, but we find that average delivery

[18]Regression results for weekdays versus weekends separately are reported in Appendix Table A.1.

Table 2: Price Effects of the Dynamic Pricing Algorithm

	(1) Mean	(2) Mean	(3) Mean	(4) Std. Dev.	(5) Std. Dev.	(6) Std. Dev.
Post × Dynamic Pricing	-0.049***	-0.051***	-0.060***	0.032***	0.032***	0.044***
	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(0.004)
Treatment Group	Pooled	A	B	Pooled	A	B
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Outcome Mean	3.393	3.395	3.384	0.484	0.485	0.480
Observations	10,404	8,640	4,735	10,404	8,640	4,735

Notes: Table displays the coefficient estimates for regressions of price outcomes on an indicator for dynamic pricing, restaurant fixed effects, and time period fixed effects. The dependent variable in the first three columns is the (mean) log transaction price. The dependent variable in the last three columns is the standard deviation of the log transaction price. Observations are at the restaurant-week level. Standard errors are in parentheses and are clustered at the restaurant level. Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

fees also fall for this group, ranging from 2.1 to 2.7 percent, depending on the subsample. The fact that we find this change when focusing on only on weekend days points to the fact that within-day demand variation is an important consideration for the algorithm.

6.2 Effects on Demand Volatility

Next, we examine the impacts of the pricing technology on demand volatility. For this analysis, we employ two measures of variation in quantity demanded. First, we construct a within-week volatility measure as the standard deviation of the number of transactions per hour divided by the mean number of transactions within each week.[19] Second, we use the share of transactions within each week that occur over the weekend (Friday-Sunday), which is the period of high demand.

Figure 8 presents the time series of the means for each outcome for group B restaurants and matched controls. Panel (a) plots the within-week coefficient of variation in number of transactions, which has declined since the introduction of dynamic pricing. Prior the adoption of dynamic pricing, and also after all restaurants employed dynamic pricing, volatility for group B (black line) is similar to the volatility for its matched controls (red dashed line). The period where group B employs dynamic pricing but its controls do not is bracketed by vertical lines. During this period, volatility was on average lower than its matched controls. After the adoption of dynamic pricing by the control restaurants, volatility in the matched controls fell to the level of the group B restaurants. The difference between the two lines is plotted in panel (b). These results indicate that demand volatility fell after the introduction of dynamic pricing, and the reduction was similar for group B and its controls.

[19] For this measure, we use only the hours between 12 PM and 8 PM across restaurants. Some restaurants are open outside of these windows, but this constitutes a low portion of transactions and revenues.

Figure 8: Dynamic Pricing and Demand Volatility

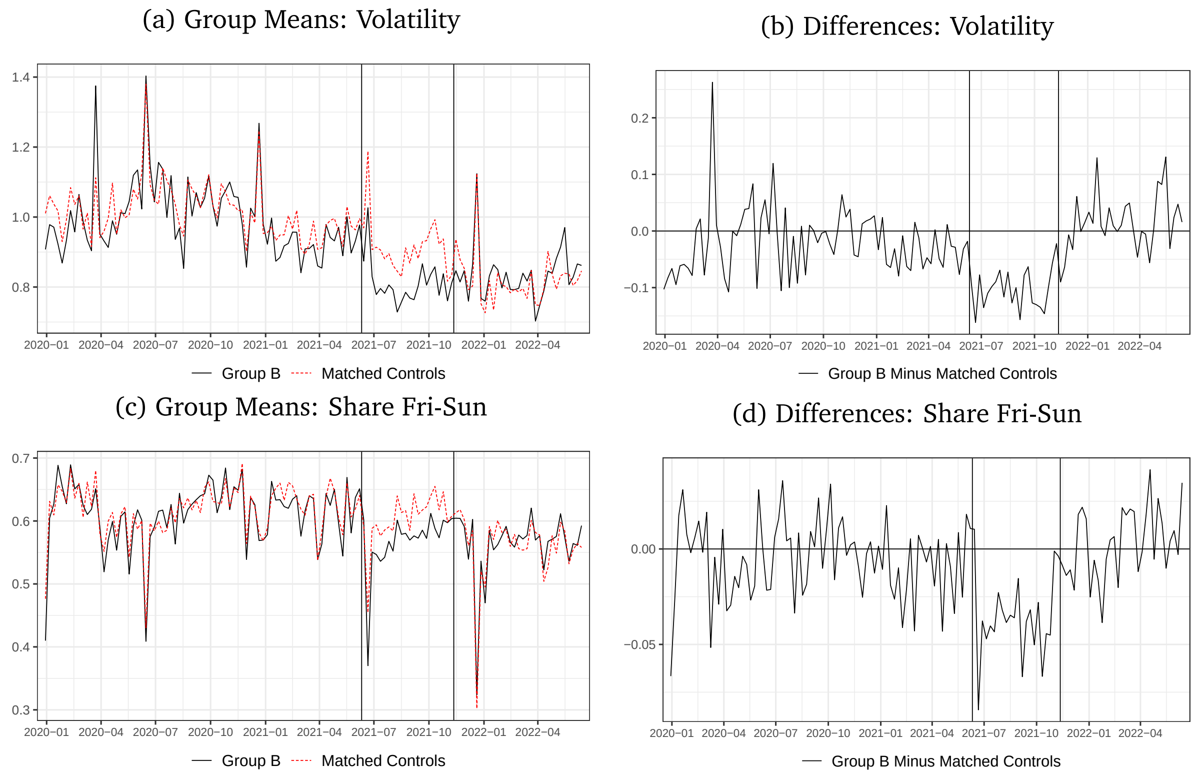

Notes: Figure displays the within-group means for our measures of volatility. Panel (a) plots the within-week coefficient of variation in number of transactions. Panel (c) plots the share of transactions that occur between Friday and Sunday. Each panel plots group B (black line) and its matched controls (dashed red line). Panels (b) and (d) present the differences between the two, providing a week-by-week event study estimate. The period where group B had dynamic pricing and its controls did not is bracketed by the two vertical lines.

Panels (c) and (d) show a similar pattern. The results indicate that the relative share during the weekend declined after the adoption of pricing algorithms, and that the decline coincided with the particular timing of adoption for each group. The corresponding plots for group A are presented in Appendix Figure A.2. The plots show similar decreases in volatility during the treatment windows, though group A has higher levels of hourly volatility in the pre-adoption and 2022 periods.[20]

Table 3 summarizes the overall effect on volatility following regression equation (6). As before, observations are at the restaurant-week level, and results are reported for the pooled treatment and for groups A and B separately. Columns (1)-(3) display the results when the dependent variable is transaction volatility. Columns (4)-(6) report the results when the dependent variable is the share of transactions on the weekend.

[20] In these figures, the means are not normalized: to construct a "difference-in-differences" estimate one would have to net out the difference between the vertical lines from a baseline period (e.g., January 2022 through June 2022). The means happen to be very similar for group B and its matched controls when they have the same pricing technology.

Table 3: Demand Volatility Effects of the Dynamic Pricing Algorithm

	(1) Volatility	(2) Volatility	(3) Volatility	(4) Share Fri–Sun	(5) Share Fri–Sun	(6) Share Fri–Sun
Post × Dynamic Pricing	-0.086*** (0.016)	-0.080*** (0.017)	-0.099*** (0.021)	-0.034*** (0.006)	-0.033*** (0.008)	-0.034*** (0.008)
Treatment Group	Pooled	A	B	Pooled	A	B
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Outcome Mean	0.890	0.885	0.904	0.594	0.593	0.595
Observations	10,404	8,640	4,735	10,404	8,640	4,735

Notes: Table displays the coefficient estimates for regressions of volatility outcomes on an indicator for dynamic pricing, restaurant fixed effects, and time period fixed effects. The dependent variable in the first three columns is a measure of within-week volatility, calculated as the standard deviation of the number of transactions per hour divided by the mean transactions per hour. The dependent variable in the last three columns is the share of transactions occurring during the weekend (the peak demand period). Observations are at the restaurant-week level. Standard errors are in parentheses and are clustered at the restaurant level. Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

Overall, we estimate that the algorithm led to decrease in the within-week transaction volatility (coefficient of variation) of -0.086. Given the outcome mean of 0.890, this represents a decline in weekly volatility of approximately ten percent. We find that our estimates are similar across treatment groups, obtaining a coefficient of -0.080 for group A only and -0.099 for group B only. Likewise, we find a substantial decrease in the relative share of transactions occurring on the weekend. The share declines by 3.4 percentage points. All of the estimates are statistically significant at the 1 percent level.

Figure 9 examines the source of the relative share decline for weekend orders. Panels (a) and (b) reflect weekday (log) transactions, and panels (c) and (d) reflect weekend (log) transactions. The figure indicates that transactions are growing over the sample period, and that the trend is similar for the matched controls. However, the pattern of growth appears to be affected by the presence of the pricing algorithm. The differences between group B and its matched controls are similar in panels (b) and (d) prior to the adoption of dynamic pricing by group B. After adoption, group B saw a relative increase in the number of weekday transactions and approximately no change in the number of weekend transactions. As weekend prices fell by 2.7 percent, the fact that transactions did not increase during the weekend period could reflect that some demand was shifted to other days of the week, which had even lower prices. The corresponding plots for group A are presented in Appendix Figure A.3. These plots show that the matched controls are growing faster than those in group A, consistent with the summary stats reported in Table 1. Adjusting for the differential trends, Figure A.3 indicates an increase in weekday transactions, and, unlike the above figures, a reduction in weekend transactions.

We summarize these effects with an estimate of the volatility semi-elasticity (VSE). We use the values from column (1) of Table 3 to estimate the percent change in demand volatility as

Figure 9: Dynamic Pricing and Quantity Demanded

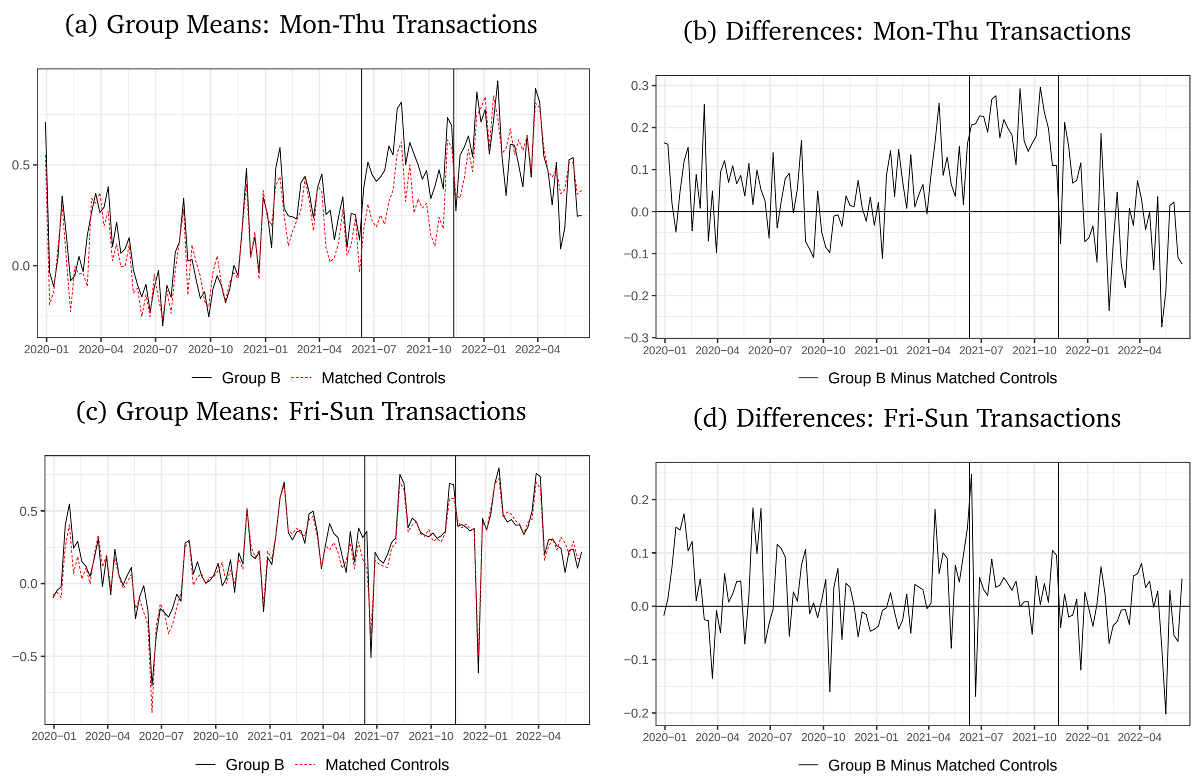

Notes: Figure displays the within-group means for the log weekly number of transactions. Panel (a) plots the transactions occurring from Monday through Thursday, and Panel (c) plots transactions that occur between Friday and Sunday. Each panel plots group B (black line) and its matched controls (dashed red line). In these panels, values are relative to the Jul-Sep 2020 mean value for group B. Panels (b) and (d) present the differences between the two, providing a week-by-week event study estimate. The period where group B had dynamic pricing and its controls did not is bracketed by the two vertical lines.

$dlnV_Q = dV_Q/V_Q \approx -0.086/0.890 = -0.097$. To calculate the change in the coefficient of variation in prices, we use the fact that the algorithm generated a mean change in the weekly standard deviation of delivery fees of 1.60. Because this change is the direct effect of adopting the algorithm,[21] we construct the change in the coefficient of variation as this value divided by the mean transaction price of 32.60 across the three groups of restaurants. This yields a change in the coefficient of variation of prices of $dV_P = 0.049$. We obtain an estimate of the VSE of -1.96, with a standard error of 0.365. This value indicates that demand is fairly responsive to the time-varying prices introduced by the algorithm, and it suggests that if the firm were to allow greater variation in list prices (e.g., time-varying prices for menu items, in addition to delivery fees), demand volatility could be reduced further.

[21] The other components of the transaction price might also endogenously change in response to changing delivery fees. On average, the transaction price net of fees was slightly smaller in 2022 than in 2020. This could indicate that consumers are more willing to place smaller orders when the delivery fee is lower. However, other factors could also be driving this change.

Figure 10: Dynamic Pricing and Revenues

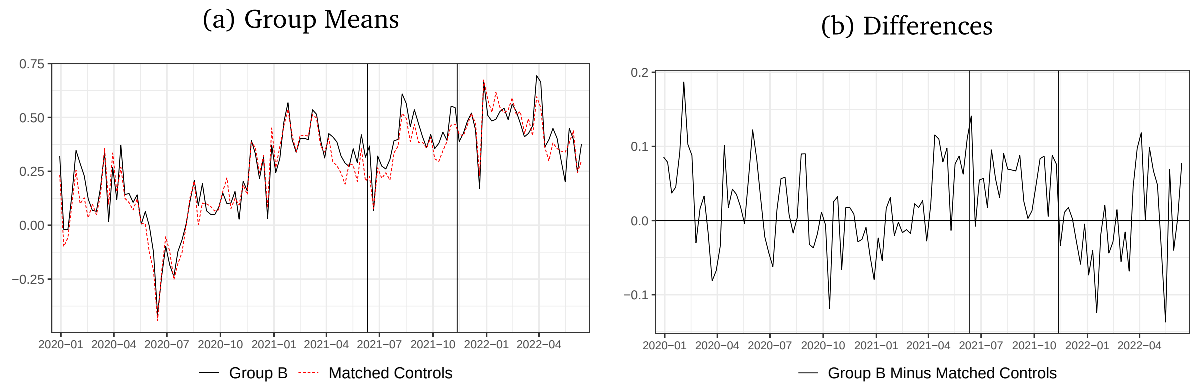

Notes: Figure displays the within-group means for the log weekly revenues. Panel (a) plots group B (black line) and its matched controls (dashed red line). In this panel, values are relative to the Jul-Sep 2020 mean value for group B. Panels (b) presents the differences between the two, providing a week-by-week event study estimate. The period where group B had dynamic pricing and its controls did not is bracketed by the two vertical lines.

Overall, these results suggest that the dynamic pricing algorithm led to reduced demand volatility. These results point to the potential mechanism of shifting demand from high-demand periods over the weekend to low-demand periods during the week, especially for restaurants that may have been closer to their capacity constraints.

6.3 Effects on Revenues

Finally, we examine the impact of dynamic pricing on overall transactions and revenues. Figure 10 presents the corresponding group mean plots and event study differences for revenues. In panel (a), group B is plotted with the solid black line, and its matched controls are plotted with a red dashed line. There appears to be similar revenues before adoption and a relative increase in revenues during the algorithm window, though there is a decent amount of week-to-week variation. Appendix Figure A.4 provides the same plot for group A restaurants. Consistent with our earlier findings, the revenues suggest that matched controls are growing faster than the group A restaurants over this period. In this figure, it is hard to tease out visually whether there is an impact on revenues due to the pricing algorithm.

Table 4 summarizes the overall effect on transactions and revenues following regression equation 6. As before, observations are at the restaurant-week level, and results are reported for the pooled treatment and for groups A and B separately. Columns (1)-(3) display the results when the dependent variable is the log number of weekly transactions. The dependent variable in columns (4)-(6) is the log weekly revenues.

Overall, we estimate that the algorithm led to a 10 percent increase in transactions. As discussed in the previous section, this increase was primarily driven by increases during the weekday period. We find roughly similar coefficients for each subsample separately, and all are

Table 4: Effects of the Dynamic Pricing Algorithm on Transactions and Revenues

	(1) Transactions	(2) Transactions	(3) Transactions	(4) Revenue	(5) Revenue	(6) Revenue
Post × Dynamic Pricing	0.097***	0.084**	0.092***	0.057**	0.040	0.051*
	(0.028)	(0.035)	(0.028)	(0.027)	(0.034)	(0.029)
Treatment Group	Pooled	A	B	Pooled	A	B
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,404	8,640	4,735	10,404	8,640	4,735

Notes: Table displays the coefficient estimates for regressions of price outcomes on an indicator for dynamic pricing, restaurant fixed effects, and time period fixed effects. The dependent variable in the first three columns is the log number of transactions. The dependent variable in the last three columns is the log revenue. Observations are at the restaurant-week level. Standard errors are in parentheses and are clustered at the restaurant level. Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

significant at the 5 percent level.

The impacts on revenues are weaker. Given that transaction prices have declined, we should expect smaller coefficients relative the impact on transactions. However, we still find positive point estimates on the impact of dynamic pricing on revenues. We estimate a statistically significant pooled increase in revenues of 6 percent. This estimate is less robust to the subsample analysis than our other findings. The subsample estimates are smaller and only the group B estimate is marginally significant.

6.4 Discussion of Intertemporal Spillovers, Operational Costs, and Welfare

Our setting provides a valuable case study to understand how dynamic pricing works in practice and the incentives for firms to adopt such technology. The adoption of this dynamic pricing algorithm reduced average prices, led to high-frequency price changes with large swings within an hour, and reduced the volatility of demand.

Our findings are consistent with the presence of intertemporal spillovers in demand. In our context, potential consumers can generate intertemporal spillovers in two ways. Consumers can observe the current delivery fee and simply wait 10 minutes (or more) to see if the delivery fee falls. Such price changes happen with a non-zero probability, even during peak demand hours. On the other hand, consumers that look ahead can shift their consumption habits so that they plan to order from this restaurant chain during periods when fees are usually lower. By reacting to the current or expected price, both types of consumers can reduce demand volatility. This can reduce costs as discussed in Section 2.3.

Our evidence also indicates that total transactions and revenues increased as a result of dynamic pricing. Though we do not have measures of costs, a reduction in demand volatility can increase profits when a restaurant faces meaningful operational costs (even with lower average prices and lower variable profits). Using a rough back-of-the envelope calculation, we estimate that profits increased if dynamic pricing enabled a reduction in operational costs equivalent to 0.5 percent of revenues.[22] Operational costs are a key consideration for the partner firm in our study and a motivation for the adoption of dynamic pricing.

Finally, the fact that prices are, on average, lower, suggests that consumers might have gained from the adoption of dynamic pricing. Whether or not consumers gain can depend on the distribution of price changes across periods. Further, the benefits of lower prices may be offset by the costs of ordering in another period and/or finding a substitute meal option, which we do not attempt to measure in this paper. However, lower prices do have a first-order benefit for consumers.

The combination of these findings—lower prices, higher output, and reduced demand volatility—are all consistent with the possibility that the adoption of dynamic pricing increased overall welfare. In the following section, we explore an empirical demand model that allow us to estimate consumer welfare directly.

7 Empirical Model of Consumer Demand

In this section, we construct an empirical model of consumer demand that builds on the conceptual framework from Section 2 and allows for intertemporal substitution across periods. We use the data obtained from the staggered rollout of the dynamic pricing algorithm to estimate the parameters of the model. The model allows us to quantify the impact to consumer surplus while accounting for the fact that prices increased in some periods and decreased in others. Using the model, we calculate that consumer surplus increased under dynamic pricing.

We model an individual consumer's decision as a discrete choice problem. A consumer i decides to purchase from the restaurant in one of the 65 typical business hours (h) of the week. Let each day be divided into lunch and dinner periods, and hours be grouped into each of these 14 weekly blocks (b).[23]

Let the utility of consumer i for a purchase in hour h be

$$u_{ih} = \alpha p_h + \xi_h + \zeta_{ib} + (1-\sigma)\varepsilon_{ih} \tag{8}$$

where p_h is the average price, ξ_h provides the mean utility across consumers for that hour of the week, ζ_{ib} is the idiosyncratic utility the consumer has for eating in the particular block, and ε_{ih} is the idiosyncratic utility for eating at that particular hour. Let ε_{ih} be independent identically distributed type 1 extreme value (T1EV) demand shocks, and ζ_{ib} follow the conjugate

[22] For this calculation, we assume that the price levels were optimal before dynamic pricing, and we calculate the reduction in variable profits under the assumptions of linear demand and constant marginal costs, using the observed impacts to prices and transactions.

[23] Lunch is defined as 11 AM to 4 PM. Dinner is 4 PM to 8 PM Sun-Thu and 4 PM to 9 PM Fri-Sat.

distribution such that $\zeta_{ib}+(1-\sigma)\varepsilon_{ih}$ is also T1EV. The consumer purchases during the hour that maximizes her utility, based on her preferences and the expected price charged in each hour of the week. The consumer can also choose not to buy in a week ($h=0$).

These assumptions yield the standard nested logit formulation as in Berry (1994). Rather than choosing across products, the consumer choice here is across periods. This model has the advantage of allowing the consumer's decision to depend on the prices in all other periods in a parsimonious way. Further, the nested logit model allows consumer preferences to have a greater correlation within nests. That is, consumers may have a stronger preference to order from a restaurant during a particular block (e.g., dinner on Friday) and thus display greater substitutability across hours within the block. This decision is a simplification of the overall problem facing the consumer—for example, consumer taste shocks are likely to be more closely correlated for adjacent periods. Nevertheless, we feel that it provides a useful way to capture the potential for intertemporal substitution at different frequencies.

Aggregating across consumers, the nested logit model yields the linear relationship

$$\ln s_h - \ln s_0 = \alpha p_h + \sigma \ln s_{h|b} + \xi_h, \tag{9}$$

where $\sigma \in [0,1)$ indicates the within-block correlation of idiosyncratic preference shocks and $s_{h|b}$ is the conditional share for hour h within block b. When $\sigma=0$, the standard logit model is obtained. As $\sigma \to 1$, hours within a block approach perfect substitutes.

To identify the parameters of this model, we use the exogenous variation induced by the staggered rollout of dynamic pricing. We obtain data on the average transactions (quantities), delivery fees, and prices for each hour of the week by restaurant group (A, B, and C) and for time periods bracketed by each rollout event. Specifically, we use differences between group A and group B over the period September 2020 through June 2021, when group A restaurants had dynamic pricing and group B restaurants did not. We use this comparison because the restaurants in these two groups were quite similar in pre period (January through August 2020), as shown in Table 1.[24]

For each group, we construct (average) within week shares for each hour of the week over the post period, and we construct the within-block conditional shares. To calculate shares, we assume that the market size is four times the average total weekly quantities during the pre period. The dataset consists of 2 (aggregated) observations for every hour of the week, one from group A with dynamic pricing and one from group B with static pricing. We then run a regression for equation (9) to recover the parameters of interest, while putting in hour-of-the-week fixed effects to account for the time-varying demand, ξ_h.

Because we only include the post-period treatment window and use hour-of-week fixed effects, the parameters are identified based on the quasi-experimental variation introduced by

[24]Employing nearest-neighbor matching would, in principle, allow us to make group C a better comparison as well. We currently do not have the use of the matched data at the hourly level.

Table 5: Consumer Demand Model Estimates

Coefficient	(1) Logit	(2) Nested Logit
Price (α)	-0.061*** (0.008)	-0.059*** (0.007)
Nest (σ)		0.898*** (0.148)
Group A Indicator	0.053*** (0.012)	0.054*** (0.010)
Own-Price Elasticity	-2.02	-15.38
Block Elasticity		-1.92
Hour-of-Week FEs	Yes	Yes
R^2 (Within)	0.727	0.828
R^2	0.993	0.996
Observations	130	130

Notes: Table displays the coefficient estimates for the discrete choice demand specification. Column (1) corresponds to a logit model and column (2) corresponds to a nested logit model. Standard errors are in parentheses. Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

the adoption of dynamic pricing. This shows up directly with a change in p_h for group A and indirectly through a shift in the conditional shares, $s_{h|b}$. We also include an indicator for the treatment group (group A), to account for any aggregate demand effects from dynamic pricing.

One potential threat to identification is that there may be correlation between the magnitude of price changes (or conditional shares) and the demand shocks. This is substantially mitigated by the use of fixed effects for ξ_h, which flexibly account for time-varying demand. The inclusion of the control set of restaurants, group B, allows us to account for the mean demand shock in each period, as the control restaurants realized variation in demand that was not due to varying prices.

Table 5 displays the results. Column (1) reports a standard logit regression (without including $\ln s_{h|g}$, and column (2) reports the nested logit estimates. The estimated parameter values are reasonable, and the model has a very good fit. The within R^2 is 0.828 for the nested logit model, compared to 0.727 for the standard logit. Overall, the parsimonious model can predict nearly all of the variation in the data.

The price coefficient does not change much from the logit to the nested logit, but the nesting parameter is economically large (0.9). This indicates that there is a good deal of correlation in idiosyncratic preferences across hours within a block. This translates to very elastic demand: for the nested logit model, we obtain an own-price elasticity of -15.4. This can be reconciled to the (smaller) logit elasticity with an appropriate level of aggregation. When aggregating to the block level, we obtain an elasticity of -1.92, similar to the logit elasticity of -2.02. In other words, the model indicates that consumers are very elastic to relative price changes within a block, but they much less likely to substitute out of a block if the prices for every hour in the block increase.

Thus, consistent with our earlier reduced-form findings, the estimates from the model indicate a high degree of consumer substitution across relatively fine time intervals. In this case, our data permit us to evaluate the substitution across hours within the same meal block. Our estimated elasticity of -15 is similar to the elasticity of -13 we found in Section 4. Though these elasticities reflect different margins and different prices (high-frequency realized changes in Section 4, expected prices in this section), both findings reflect the fact that consumers may strategically time purchases at high frequencies. Demand models that do not account for this (e.g., the standard logit above) may fail to capture the degree of consumer response to a given price change.

Finally, our estimated demand model also enables us to calculate consumer surplus. Though average prices fell, prices increased during high demand periods, where consumer surplus had been highest. The nested logit model allows us to aggregate welfare effects across periods while accounting for differences in demand and prices. Using a modified version of the Small and Rosen (1981) formula to account for the nesting structure, our estimates indicate that consumer welfare is higher under the dynamic pricing regime. The increase in consumer surplus is equal to 1.5 percent of baseline revenue.

8 Conclusion

This paper analyzes the impacts of adopting a dynamic pricing algorithm in the context of restaurant food delivery. Using high-frequency transaction data, we present evidence that consumers strategically time their purchases in response to time-varying pricing. Intertemporal substitution appears to occur across days of the week and also at high frequency (within an hour). We then measure the impact of dynamic pricing on prices, quantities, and demand volatility. Our empirical strategy utilizes the staggered adoption of the technology by different groups of restaurants within the same chain.

Our analysis provides a case study of how time-varying prices are implemented in practice and the potential welfare gains. In our context, we find that the adoption of time-varying prices led to lower average prices, greater transactions, and a reduction in within-week demand volatility. These results suggest that the firm was more efficiently able to use capacity, as time-varying prices smoothed out demand within and across days. Our findings indicate that consumers strategically time their purchases, and we highlight how firms can benefit from this strategic behavior through high-frequency dynamic pricing.

References

Aryal, Gaurab, Charles Murry, and Jonathan W. Williams. 2022. "Price Discrimination in International Airline Markets." Working paper.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2022. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." Working paper.

Baker, Andrew C, David F Larcker, and Charles CY Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics*, 144(2): 370–395.

Berry, Steven T. 1994. "Estimating discrete-choice models of product differentiation." *The RAND Journal of Economics*, 242–262.

Brown, Zach Y, and Alexander MacKay. 2021. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics*. Forthcoming.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. 2020. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review*, 110(10): 3267–97.

Castillo, Juan Camilo. 2020. "Who benefits from surge pricing?" Working paper.

Cho, Sungjin, Gong Lee, John Rust, and Mengkai Yu. 2018. "Optimal dynamic hotel pricing." Working paper.

Coase, Ronald H. 1972. "Durability and monopoly." *The Journal of Law and Economics*, 15(1): 143–149.

Dana, James D. 1999. "Using Yield Management to Shift Demand When the Peak Time Is Unknown." *The RAND Journal of Economics*, 30(3): 456.

Ekholm, Anders Goran. 2019. "Method, system and computer program product for dynamically pricing perishable goods." US Patent 10,430,842.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Gowrisankaran, Gautam, and Marc Rysman. 2012. "Dynamics of Consumer Demand for New Durable Goods." *Journal of Political Economy*, 120(6): 1173–1219.

Hendel, Igal, and Aviv Nevo. 2013. "Intertemporal price discrimination in storable goods markets." *American Economic Review*, 103(7): 2722–51.

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams. 2021. "Organizational Structure and Pricing: Evidence from a Large US Airline." Working Paper.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417–1438.

Lazarev, John. 2013. "The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets."

Lee, Robin S. 2013. "Vertical Integration and Exclusivity in Platform and Two-Sided Markets." *American Economic Review*, 103(7): 2960–3000.

Melnikov, Oleg. 2013. "Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market." *Economic Inquiry*, 51(2): 1277–1298.

Nair, Harikesh. 2007. "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games." *Quantitative Marketing and Economics*, 5(3): 239–292.

Small, Kenneth A, and Harvey S Rosen. 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica*, 105–130.

Stokey, Nancy L. 1979. "Intertemporal price discrimination." *The Quarterly Journal of Economics*, 355–371.

Sweeting, Andrew. 2012. "Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets." *Journal of Political Economy*, 120(6): 1133–1172.

Williams, Kevin R. 2022. "The Welfare Effects of Dynamic Pricing: Evidence From Airline Markets." *Econometrica*, 90(2): 831–858.

Appendix

A Additional Tables and Figures

Figure A.1: Delivery Fees by Restaurant

(a) Mean

(b) Standard Deviation

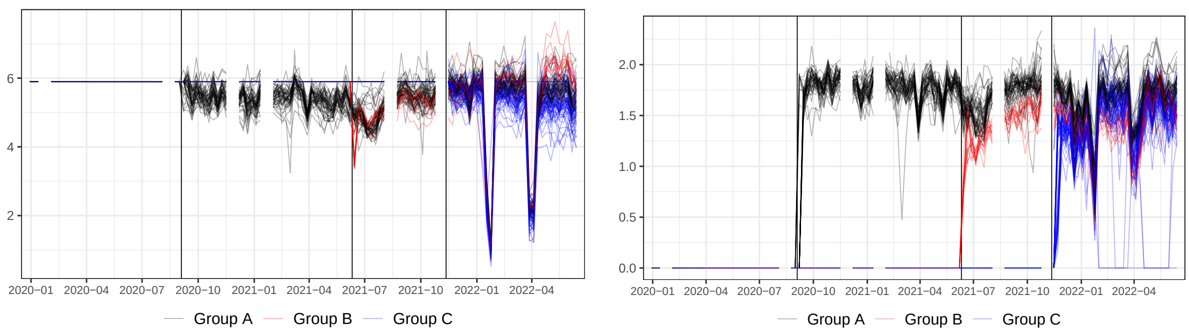

Notes: Figure displays statistics of delivery fees by restaurant over the sample period. Panel (a) displays the weekly mean, and panel (b) displays the within-week standard deviation for each restaurant. Restaurants in group A are in gray, group B are in blue, and group C are in red. Gaps in the time series correspond to promotional periods with zero delivery fees across all restaurants.

Table A.1: Price Effects: Weekdays versus Weekends

(a) Monday through Thursday

	(1) Mean	(2) Mean	(3) Mean	(4) Std. Dev.	(5) Std. Dev.	(6) Std. Dev.
Post $\times$ Dynamic Pricing	-0.083***	-0.083***	-0.097***	0.034***	0.032***	0.061***
	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)
Treatment Group	Pooled	A	B	Pooled	A	B
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,273	8,528	4,715	10,273	8,528	4,715

(b) Friday through Sunday

	(1) Mean	(2) Mean	(3) Mean	(4) Std. Dev.	(5) Std. Dev.	(6) Std. Dev.
Post $\times$ Dynamic Pricing	-0.021***	-0.025***	-0.027***	0.022***	0.025***	0.022***
	(0.003)	(0.004)	(0.005)	(0.004)	(0.006)	(0.005)
Treatment Group	Pooled	A	B	Pooled	A	B
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,390	8,629	4,728	10,390	8,629	4,728

Notes: Table displays the coefficient estimates for regressions of price outcomes on an indicator for dynamic pricing, restaurant fixed effects, and time period fixed effects. Panel (a) uses only data from Monday through Thursday, and panel (b) uses only data from Friday through Sunday. The dependent variable in the first three columns in each panel is the (mean) log transaction price. The dependent variable in the last three columns is the standard deviation of the log transaction price. Observations are at the restaurant-week level. Standard errors are in parentheses and are clustered at the restaurant level. Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

Figure A.2: Dynamic Pricing and Demand Volatility: Group A

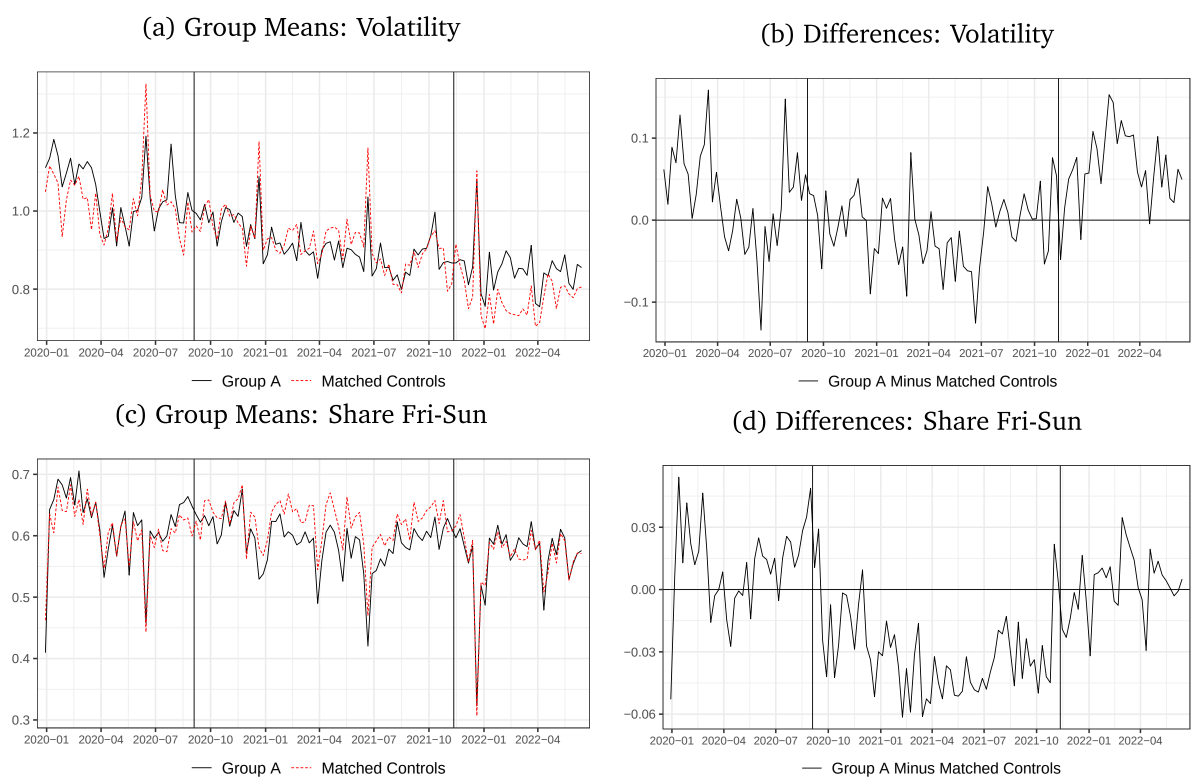

(a) Group Means: Volatility

(b) Differences: Volatility

(c) Group Means: Share Fri-Sun

(d) Differences: Share Fri-Sun

Notes: Figure displays the within-group means for our measures of volatility. Panels (a) plots the within-week coefficient of variation in number of transactions. Panel (c) plots the share of transactions that occur between Friday and Sunday. Each panel plots group A (black line) and its matched controls (dashed red line). Panels (b) and (d) present the differences between the two, providing a week-by-week event study estimate. These differences do not also remove the pre-period differences; as can be seen in the figure, group A had higher mean volatility before the adoption of the algorithm. The period where group A had dynamic pricing and its controls did not is bracketed by the two vertical lines.

Figure A.3: Dynamic Pricing and Quantity Demanded: Group A

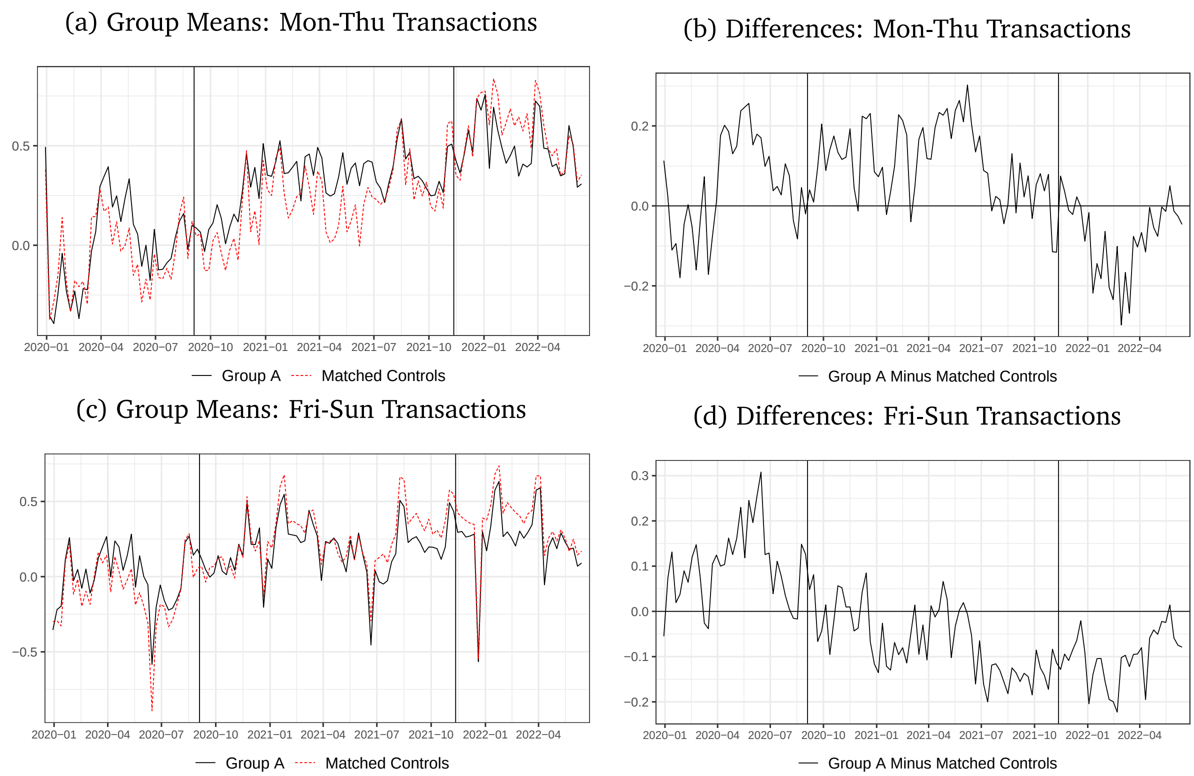

Notes: Figure displays the within-group means for the log weekly number of transactions. Panel (a) plots the transactions occurring from Monday through Thursday, and Panel (c) plots transactions that occur between Friday and Sunday. Each panel plots group A (black line) and its matched controls (dashed red line). In these panels, values are relative to the Jul-Sep 2020 mean value for group A. Panels (b) and (d) present the differences between the two, providing a week-by-week event study estimate. These differences do not also remove the pre-period differences; as can be seen in the figure, group A had greater transactions before the adoption of the algorithm. The figure also indicates the slightly higher growth rates by the matched controls for group A. The period where group A had dynamic pricing and its controls did not is bracketed by the two vertical lines.

Figure A.4: The Impacts of Dynamic Pricing on Revenues: Group A

(a) Group Means: Revenues

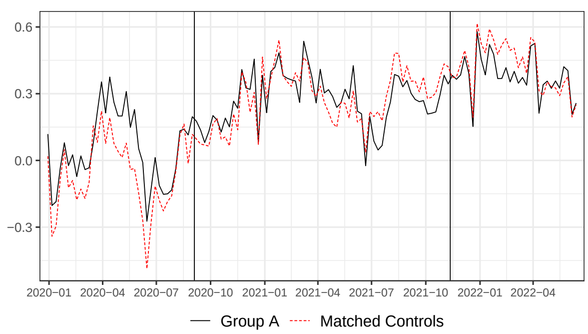

(b) Differences: Revenues

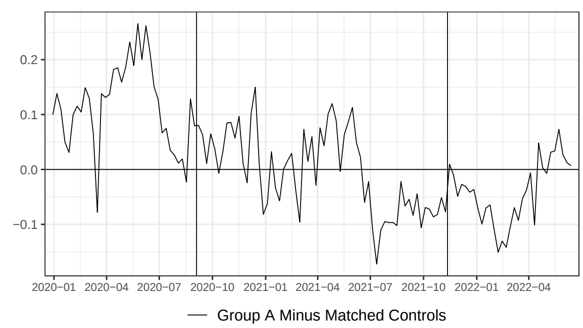

Notes: Figure displays the within-group means for the log weekly revenues. Panel (a) plots group A (black line) and its matched controls (dashed red line). In this panel, values are relative to the Jul-Sep 2020 mean value for group A. Panels (b) presents the differences between the two, providing a week-by-week event study estimate. These differences do not also remove the pre-period differences; as can be seen in the figure, group A had greater transactions before the adoption of the algorithm. The figure also indicates the slightly higher growth rates by the matched controls for group A. The period where group A had dynamic pricing and its controls did not is bracketed by the two vertical lines.