Peak-Load Pricing in Platform Markets

Takeaki Sunada, Diego Aparicio, and Yufeng Huang*

July 7, 2025

Abstract

We study welfare and profit impacts of peak-load pricing in the context of a dining reservation platform that allows restaurants to set variable prices. Using unique data on reservations and a measure of restaurant traffic, we estimate a model where equilibrium prices respond to both time-varying consumer price sensitivity and restaurants' capacity constraints. We find peak-load pricing, rather than intertemporal price discrimination, is the primary driver behind the observed price variation. We show that variable pricing increases social welfare by 8.6%, demonstrating the vital role of pricing in improving efficiency in the context of platform markets. We also find variable pricing can reduce the profit of the platform, which thus might lack the incentive to provide the variable-pricing technology despite its welfare benefits. Notably, this supply-side incentive misalignment is most salient when the platform and firms share their joint profit, thus offsetting the relative efficiency of profit-sharing contracts over per-unit-fee contracts.

Key words: Peak-Load Pricing, Price Discrimination, Platform Economy

JEL codes: D43, L11, L22

*Takeaki Sunada, University of Rochester (tsunada@simon.rochester.edu), Diego Aparicio, IESE Business School (daparicio@iese.edu), Yufeng Huang, University of Rochester (yufeng.huang@simon.rochester.edu).

1 Introduction

Peak-load pricing is a pricing strategy where firms charge higher prices during peak periods to redirect consumers to off-peak periods that are less capacity-constrained. This strategy is commonly studied in the context of public infrastructure services, such as electricity supply and highway toll pricing. However, many profit-maximizing firms also face time-varying demand and a limited capacity, and their use of peak-load pricing and its welfare consequences are not well understood. Airlines, hotel chains, and theme parks employ sophisticated algorithms to adjust prices across days and times. Uber's surge pricing, higher airfare during holiday periods, and Disney tickets whose price depends on the day of week, are typical examples. In contrast, many platforms maintain inflexible pricing despite their sellers facing significant demand fluctuations and capacity constraints—notable examples include dining platforms (OpenTable), salon booking apps (StyleSeat), and taxi-hailing services (Curb).

We study the welfare-improving role of peak-load pricing in a platform market with profit-maximizing firms. Peak-load pricing diverts demand from capacity-constrained peak periods to off-peak periods, thereby relaxing capacity constraints and improving welfare through increased overall supply. This pro-social implication contrasts with intertemporal price discrimination, in which the firm extracts surplus by inducing sorting among consumers with heterogeneous preferences across times, a form of second-degree price discrimination that often leads to a distorted pricing schedule. However, because the firm's peak-load pricing and intertemporal price discrimination incentives often coexist, their overall welfare impact remains theoretically ambiguous. We separate the two mechanisms by identifying whether firms' variable pricing is primarily driven by capacity constraints or by heterogeneous consumer price sensitivities. We show substantial welfare gains can result when the incentive for price variability primarily arises from capacity constraints.

We also show that peak-load pricing in a platform market can reduce the profit of the platform, which thus might lack the incentive to allow variable pricing even when it benefits participating firms and consumers. In many platform markets, participating firms serve customers both through the platform and through a direct sales channel: typical examples include restaurants serving Opentable reservations and walk-in diners, and taxi companies serving Curb reservations and street-hailed taxis. In such markets where limited firm capacity is shared between platform and off-platform transactions, peak-load pricing on the platform can shift the platform sales to relatively unprofitable off-peak times, and divert more profitable peak sales to off-platform markets. This sales diversion does not benefit the platform, creating a supply-side incentive misalignment that has welfare consequences. In particular, we show that this incentive misalignment is most pronounced under profit-sharing arrangements between the platform and firms, but it largely disappears when the platform instead charges fixed per-unit fees. Thus, this incentive misalignment under capacity constraints partly offsets the well-known efficiency gains from profit-sharing contracts in avoiding double-marginalization between firms and the platform.

Our empirical environment is a dining reservation platform. The platform is marketed to restaurants as a new technology to generate incremental revenue by filling otherwise underutilized capacity during off-peak hours. Specifically, the platform serves a small segment of price-sensitive consumers that is separate from the restaurant's existing customer base. Restaurants can offer discounts to platform users that vary by time of day and day of week. Platform users who make reservations receive the percentage discount applied to their entire bill. The platform earns a commission as a fixed fee per reservation. Because only platform users receive the time-varying discounts, peak-load pricing in this market corresponds to reallocating the platform demand across times for a given inflow of walk-in (off-platform) diners. Restaurants set the platform discount rates ahead of time, managing their time-varying capacity constraints that arise from variable walk-in traffic, while also attempting to price discriminate between platform customers with potentially time-varying price elasticities. We obtain internal data on consumer reservations and restaurants' discount rate schedules, and we augment these data with a measure of walk-in traffic for each restaurant and time slot.

This market environment provides two key advantages to the study. First, the observed volume of walk-in traffic allows us to infer how the restaurants' time-varying capacity constraints affect their platform pricing. Accounting for capacity constraints is crucial to separate firm incentives for peak-load pricing from price discrimination, because, absent data on capacity, both incentives create similar intertemporal price patterns (Hendel and Nevo, 2013). Second, the demand for dining reservation lacks dynamics. In other markets where peak-load pricing is employed, capacity often gradually fills over time (e.g., airlines, hotels), and firms employ dynamic inventory control, which significantly complicates the analysis.[1] In contrast, the vast majority of demand in our focal market consists of walk-

[1]Most existing papers on hotel and airline pricing assume fixed consumer travel time for model tractability (Cho et al., 2018, Williams, 2022), thereby assuming away the core mechanism of peak-load pricing — using intertemporal price variation to affect consumers' choices between peak and non-peak times.

in customers and last-minute reservations. The static nature of demand provides a unique laboratory-like setup in which the firm's capacity control problem collapses to a single-period model, where it optimizes platform reservations given the expected walk-in traffic.

We start by documenting substantial price variation on the platform across times of the day. For an average restaurant, the gap between the highest and lowest price on a typical day is about 20% of the menu price. We demonstrate that the intra-day price patterns are highly correlated with the walk-in traffic pattern of the restaurant, and that platform customers are willing to move their reservation time to a lower-priced, off-peak hour within the same day, suggesting that peak-load pricing is the primary driver of intra-day price variation. We also document modest heterogeneity in customer price sensitivity, suggesting that price discrimination between platform customers is limited.

We next construct and estimate a model of demand and supply to separate the firm motives for peak-load pricing and intertemporal price discrimination. In the demand model, platform customers arrive first and make a reservation decision. Each platform customer has an ideal dining time and can be motivated to switch to other dining times and/or restaurants by lower prices. We also allow consumers' price sensitivities to vary across time, thereby creating restaurants' intertemporal price discrimination incentives. Given platform customers' reservations, walk-in customers arrive exogenously and fill the available tables. Because restaurants must honor reservations, if the capacity constraint binds, platform reservations displace walk-in customers who would otherwise pay the regular price, creating an opportunity cost for the restaurant. This opportunity cost gives the restaurant an incentive to use prices to redirect platform customers to a less-capacity-constrained dining time.

We estimate the model using data on reservations, prices, and restaurant traffic and identify both consumers' time-varying price sensitivity and restaurants' time-varying opportunity cost due to sales displacement. The fact that restaurants' capacity constraints arise from variable walk-in traffic, which is separate from the platform market, allows us to utilize preference externality (Waldfogel, 2003), greatly simplifying identification. We find a price elasticity of around -10, with little intertemporal variability. Consumers are more willing to substitute within a restaurant across times, rather than substituting to other restaurants. On the supply side, we find restaurants face meaningful capacity constraints. Approximately 40% of the restaurants expect their peak capacity to bind with at least 50% probability, and their platform prices account for the possible revenue loss due to this binding capacity. These findings indicate that the platform is an ideal market for exercising peak-load pricing, matching capacity-constrained restaurants with users who are willing to shift dining times and not substitute to competitors. We find that counterfactually eliminating the capacity constraints decreases the intertemporal price variability of a restaurant by 88%, suggesting that peak-load pricing is the primary driver of the observed price patterns.

We use the model to evaluate the efficiency gains from peak-load pricing. Specifically, we compare the welfare the platform generates under variable pricing against that of a counterfactual scenario in which each restaurant is constrained to a uniform discount rate within a day. We find that with uniform pricing, reservations concentrate more on peak periods and displace a larger number of walk-in customers than variable pricing. Each reservation displaces 0.21 walk-in customers (or displace one 21% of the time) under the uniform-pricing scenario, whereas the displacement rate drops to 0.15 under the variable-pricing scheme. Peak-load pricing hence decreases the displacement of walk-in customers by 30%. Due to more efficient capacity allocation, variable pricing increases the total welfare generated by the platform, and the profit of the restaurant and the platform, by 8.6%, 8.0% and 9.6%, respectively. Peak-load pricing thus provides substantial efficiency gains and increases welfare, even when the firm's motivation is profit maximization. Notably, our findings indicate that the per-unit-fee contract between restaurants and the platform observed in the data produces a roughly even split of the gains among market participants.

We next show that alternatively considering profit-sharing contracts between restaurants and the platform can lead to a supply-side incentive misalignment problem. Specifically, we show that while a profit-sharing contract substantially enhances the efficacy of peak-load pricing, it incentivizes the platform less to allow such a pricing strategy. On the one hand, we find that, under a profit-sharing scheme, variable pricing increases the overall welfare and the restaurants' profit by 11.5% and 12.5%, respectively, which amounts to 130% - 150% of the gains under the observed per-unit-fee scheme. By lowering the restaurants' marginal costs (akin to eliminating double marginalization), a profit-sharing scheme allows restaurants to reduce off-peak prices further, creating a larger intertemporal substitution and improving capacity utilization. This finding thus adds to our conventional wisdom on the relative efficiency of profit-sharing contracts over per-unit-fee contracts. On the other hand, we also find that the same profit-sharing contract limits the platform's gains from variable pricing to only 3.1%. In particular, variable pricing on the platform *reduces* the platform's profit from 53% of the restaurants. These restaurants primarily use price variability to shift platform sales to off-peak times and increase off-platform sales during the peak time, reducing their platform profits. Our results suggest that the platform may not allow these restaurants to use variable pricing despite the welfare benefits: excluding these restaurants, we find the welfare gain from variable pricing is limited at 7.7%. Thus, the incentive misalignment can offset a substantial portion of the relative efficiency of profit-sharing contracts over per-unit-fee contracts.

Our study is among the first to present empirical evidence that the ability to flexibly adjust prices, even if it is employed by profit-maximizing firms, can result in substantial welfare improvements when firms face capacity constraints. Our unique environment allows for clear separation of peak-load pricing incentive from intertemporal price discrimination, making this welfare evaluation feasible. Because costly capacity investments often serve as an entry barrier, markets in which pricing power is a welfare concern often exhibit firm capacity constraints (e.g., airlines and hotels). Our result suggests the same capacity constraint can encourage firms to implement pro-social pricing, thus contributing to policy debates on firm pricing power in such markets.

We also present evidence of previously unexplored supply-side incentive misalignment in platform markets. Although profit-sharing contracts avoid double marginalization from per-unit-fee contracts, in markets with external capacity constraints they also create incentive misalignment problems in which the platform may not allow the full welfare benefits to materialize. These findings highlight the importance of accounting for the *platform's* incentive in implementing a socially desirable pricing strategy: in many platform markets, sellers only serve a small market niche and hence lack scale to justify the adoption of flexible pricing technology outside of the platform (e.g., short-term rentals, restaurants). Thus, the welfare consequence can hinge exclusively on whether the platform is willing to offer the pricing flexibility (e.g., Airbnb) or not (e.g., OpenTable). Our finding that the profit-sharing contract, a contract that we tend to consider as most efficient, creates misaligned incentives with majority restaurants, indicates that a substantial welfare loss can arise from improperly designed contracts in platform markets.

Related literature. Our primary contribution is to the literature on peak-load pricing, particularly in the context of profit-maximizing firms operating in platform markets. The closest literature studies firms' dynamic pricing strategies under capacity constraints, including Williams (2022) on airlines, Cho et al. (2018) on hotels, Sweeting (2012) on concert tickets, and Sanders (2024) on fresh food. Whereas these markets implicitly have peak-load pricing considerations—the opportunity costs of selling now varies with how likely a unit is to sell in the future—the existing models rarely capture such an incentive.

A key limitation is that models of dynamic inventory control typically assume that consumers do not time their purchases strategically, thereby removing the channel through which prices may influence purchase timing and help balance capacity utilization. In contrast, our paper studies a market environment where dynamic pricing is absent, permitting a static model that explicitly characterizes how prices affect consumers' choices between peak and off-peak periods, an essential channel in firms' peak-load pricing considerations.

Our paper also builds on the extensive literature on peak-load pricing in public infrastructure contexts. The long-standing theory literature has outlined peak-load pricing as a way to balance demand and supply for electricity, water, and other publicly-supplied goods facing a capacity constraint (Williamson, 1966, Turvey, 1968). More recent works document consumer responses and social outcomes when peak-load pricing regimes are implemented in regulated markets of electricity and natural resources (Joskow, 2012, Borenstein, 2012, Brecko and Hartmann, 2023). We extend this literature's insights to private markets, where prices are set by profit-maximizing firms and the pricing technology is provided by a platform.

Finally, we highlight a novel mechanism in platform markets where the platform may lack incentives to provide flexible pricing technology despite its welfare benefits. The previous works on pricing technology primarily focuses on pricing algorithms as a way to automate price-setting decisions (Aparicio and Misra, 2023, MacKay et al., 2023, Castillo and Mathur, 2023). The specific incentive misalignment we identify regarding the provision of variable pricing technology highlights market inefficiencies in platform-mediated transactions.

2 Empirical Setting

2.1 A Dining Reservation Platform

Our research setting is a dining reservation platform in Germany. The platform is marketed to restaurants as a new technology to generate incremental revenue by utilizing otherwise unused capacity during off-peak hours. Specifically, the platform attracts a small segment of price-sensitive consumers that is separate from the restaurants' existing customer base. Participating restaurants can offer discounts that vary by both day of week and time of day, in 30-minute intervals. For each day and time slot (e.g., 8:30pm on Saturdays), restaurants choose a discount rate ranging from 0% to 50%, in 10 percent increments. When a customer books a table through the platform, the selected discount is applied to the total bill with no upper limit. For example, if a customer reserves a table with a 10% discount, they may order freely from the regular menu, and the final bill — covering all food and most drinks — is reduced by 10%. In this sense, the discount is equivalent to a uniform price discount across items on the regular menu.[2] The platform earns a commission as a fixed fee per reservation. Although restaurants' existing customers can also substitute to booking through the platform, we show below that most diners do not seem aware of the platform, leading to limited cross-market substitution.

Restaurants' primary pricing objective on the platform is to accommodate the price-sensitive platform users during off-peak times with a lower price, while limiting their access during the peak times by raising the price. Because platform users are more price sensitive than regular diners, accepting their reservations during the peak time, potentially crowding out regular dining demand, is socially wasteful. Thus, motivating the price-sensitive segment to less capacity-constrained times through time-varying prices — "peak-load pricing" in our focal market — not only increases profit, but can also enhance welfare.

Restaurants can also use variable pricing for intertemporal price discrimination among platform users. If platform users themselves exhibit heterogeneous price sensitivities and different preference across times, variable pricing can induce consumer sorting, in the form of second-degree price discrimination. Notably, unlike the peak-load pricing incentive, this price-discrimination incentive does not require restaurant capacity constraints. Separating out these two pricing incentives — by measuring how prices reflect the intertemporal variation in off-platform traffic and in the price sensitivity of platform users — is the main identification task of this paper. We exploit the unique nature of the market that these two incentives arise from different markets.

The platform operates in multiple cities in Germany and the U.K., with the largest presence in Berlin. By March of 2020 (pre-pandemic period), 231 restaurants in Berlin operated with the platform. The vast majority of participants are lower to mid-range, small-scale restaurants not affiliated with a chain. The platform only recruits restaurants above a certain Google Review score of approximately 4.1. Because restaurants themselves often possess limited ability to vary prices due to their high menu costs, the platform essentially serves as the sole provider of flexible pricing technology in the market.[3] By nature of the

[2]Each restaurant may impose exclusion (e.g., high-end liquor items). For the vast majority of cases, such exclusions only apply to a small fraction of items in the menu.

[3]To the best of our knowledge, no other platforms offer the pricing ability at this level of flexibility. Although many restaurants also offer coarse discounts in the form of happy hours, the discounts on the

platform, participating restaurants are not a random subset of the restaurants in the city, but are ones that benefit the most from variable pricing. The set of platform users are also more price sensitive than average diners. We discuss implications of the non-representative market in the result section.

2.2 Data

The original data covers the universe of reservations made through the platform for restaurants in Berlin between October 2018, which is the inception of the platform, and November 2020. For each reservation, we observe consumer and restaurant identifiers, along with the date and time of the reservation. Additionally, we observe the full set of time slots and discount rates offered by all participating restaurants on each day—regardless of whether a transaction occurs. We also have access to several restaurant characteristics, including name, location, cuisine type (e.g., American, Asian), regular price per person, and average Google rating. The regular price reflects the average amount a typical diner would pay without a discount. This metric is collected by the platform and displayed to customers when browsing the platform.

The prices are not determined at the time of reservation, because they depend on the actual items ordered. We define the price consumers expect to pay for each time slot as $P = \text{Regular Price} \times (1 - \text{Discount Rate})$. We assume consumers make their reservation decisions based on this expected price.

To complement the platform data, we construct a measure of restaurant traffic using data from Google Maps. Each restaurant's Google Maps page displays a plot indicating the volume of traffic by day of the week and hour of the day. This metric is created from aggregated Google device location data and is scaled between 0 and 100, where 0 represents either the restaurant is closed or no guests are dining, and 100 represents the busiest time during the week. Our data contains a snapshot of the traffic measure for each restaurant $\times$ hour of day $\times$ day of week tuple. We use this traffic measure as the measure of the restaurant's "walk-in" traffic, that is, the set of diners either without a reservation or from other reservation channels, excluding platform users. Although the measure also includes customers from the platform, we show below that the volume of platform transactions is quite small compared to a typical restaurant's size. It is thus safe to assume the platform customers' traffic has a negligible impact on this measure.[4] Notably, walk-in customers

platform are often deeper in magnitude and carry far less exclusions.

[4]This traffic measure is not a direct measure of capacity constraints, as it does not account for the restau-

do not face the platform discounts (because most are not aware of the platform), explicitly ruling out a reverse-causality from the discount to the traffic measure.

To consider periods with stable demand and supply conditions, we remove from the sample the first 11 months of observations (the platform's launch periods), as well as all observations after March 8th, 2020 that exhibit a substantial impact of COVID-19. We also restrict our focus to the set of restaurants that received at least 10 reservations during the 6-month sample period, as well as the set of customers who made at least three reservations.[5] Finally, we consider time slots between 3pm and 9:30pm (i.e., the maximum of 14 time slots per day with a 30-minute increment), which include the peak dinner times and off-peak times around them and include approximately 75% of the reservations. Details of our sample selection process are available in Appendix B.

2.3 Summary Statistics

The data cover 6,035 reservations by 1,326 customers across 80 restaurants during 190 days between September 1st, 2019 and March 8th, 2020. Panel (A) of Table 1 shows summary statistics at the restaurant level. On average, restaurants offer 12.5 half-hour time slots per day among the 14 slots we consider in the analysis, indicating that most restaurants are available throughout the day. 73% of restaurant-date pairs show all 14 slots listed as available on the platform.

Restaurants set discount rates ranging from 0% to 50%, with an average rate of 26%, indicating that the discounts offered are economically meaningful. The resulting average discounted price per person is €13.2, with prices ranging from €4.30 to €54.00. Approximately 80% of available options are priced below €15 per person, suggesting that participating restaurants are concentrated at the lower end of the price spectrum. On average, a restaurant receives one reservation every two days, although there is considerable heterogeneity. The most popular restaurant receives an average of 3.28 reservations per day, while the least popular receives just 0.06. The relatively low volume of platform-based reservations suggests that the platform accounts for a small share of the restaurants'

rant's seating capacity. We use our supply-side model to uncover the time-varying occupancy rates that best rationalize the observed prices. The usefulness of this traffic measure is that it provides the crucial information on how capacity constraints *shift* over time, even though we need to estimate the (time-invariant) seating capacity. We discuss further details of the measure in Appendix A.

[5]We focus on repeat users, because we exclude each user's first reservation from the likelihood. Including the first reservation causes a potentially misspecified likelihood, because we do not observe whether or not the user was active on the platform prior to her first reservation. Consumers with three or more reservations account for approximately 22% of consumers and 54% of reservations in the data.

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
(A) Restaurant-level:				
Number of Available Slots (per Day)	12.638	2.811	2	14
Discount Rate	0.256	0.102	0	0.5
Price (per Person, euro)	13.208	7.598	4.3	54
Number of Reservations (per Day)	0.509	0.940	0	9
Walk-in Traffic	48.927	24.190	0	100
(B) Customer-level:				
Number of Reservations	4.551	3.318	1	41
Average Discount Rate Booked	0.317	0.052	0.1	0.5
Average Price Booked (per Person)	10.681	3.169	5.7	37.6
Number of Restaurants	80			
Number of Customers	1,326			
Number of Reservations	6,035			

Note: In panels (A) and (B), the unit of observation is a distinct restaurant - date - time combination and a distinct consumer, respectively, unless otherwise noted.

overall traffic.[6] Finally, our walk-in traffic data show that restaurants operate at nearly half (48.9%) the volume relative to their highest level throughout a typical week.

Panel (B) of Table 1 shows summary statistics at the consumer level. On average, consumers make 4.5 reservations during the 6-month period. The average discount obtained in their bookings is 32%. Additional details on other variables used for demand estimation and histograms of prices and bookings are available in Appendix C.

3 Descriptive Evidence

In this section, we document empirical patterns in pricing and reservation behavior observed in the data. We first show that platform discounts are closely aligned with fluctua-

[6]The scale of the platform does not undermine the validity of our analysis. We evaluate how walk-in traffic influences firms' behavior on the platform, rather than how the platform affects total traffic. Restaurants flexibly adjusting discounts across times in the data also indicates that they indeed make rational decisions on the platform. The fact that they raise prices during the peak time indicates that even a small number of platform reservations may crowd out walk-in customers at the margin, allowing us to infer that many restaurants operate near full capacity.

tions in walk-in traffic — i.e., variable pricing seems to be responsive to expected capacity utilization. Second, we document that platform consumers primarily react to price variability by shifting their dining times within the same day, rather than across days. Lastly, we provide evidence indicating a limited scope for intertemporal price discrimination among platform users, as their responses to discounts appear relatively homogeneous.

3.1 Restaurants' Pricing Behavior

Panel (a) in Figure 1 shows the intra-day variation in average platform discount rates (solid line) and average walk-in traffic volume (dashed line). We observe a substantial variation in the discount rates within the day, from 15% at 7pm to 34% at 4pm. The lowest discount rates are observed between 7pm and 8pm, henceforth referred to as "peak time". The discount rate also reflects the volume of expected foot traffic: the negative correlation implies restaurants lower the discount rates when they expect higher volume of walk-in traffic. A simple regression of the discount rates (in %) on the traffic measure and restaurant fixed effects yields a coefficient of -0.12 that is statistically significant: a one-unit increase in the measure of walk-in traffic is associated with a 0.12 percentage point (pp) lower discount rate for a given restaurant. This negative correlation is consistent with the hypothesis that pricing reflects capacity constraints.[7]

Panel (b) depicts the average number of reservations per restaurant by time of day, plotted against the average discount rates. Contrary to the walk-in traffic, platform sales peaks at 6:30pm, just before 7pm when the discount rate drops sharply. There are two other local peaks of reservations: at 4pm when the discount rate is highest, and at 9pm just after the sharp increase in the discount rates. These initial observations are consistent with the practice of peak-load pricing. That is, restaurants use within-day price variation to drive platform reservations to less occupied time slots and allocate their peak time slots to walk-in customers, who pay the regular price.

These figures also present suggestive evidence that the platform discounts do not draw meaningful substitution from the market of walk-in diners. Panel (b) shows that despite an average discount of 30% during off-peak times, restaurants only receive about 0.4 off-peak reservations per day, indicating limited awareness of the platform among general diners. All but one restaurants even offer discounts during the peak times — hence walk-in diners

[7]Because walk-in customers do not face online discounts, this correlation explicitly rules out reverse causality from the price to the walk-in traffic. However, the relationship may still not be purely causal, because the volume of walk-in traffic is likely correlated with the demand on the platform.

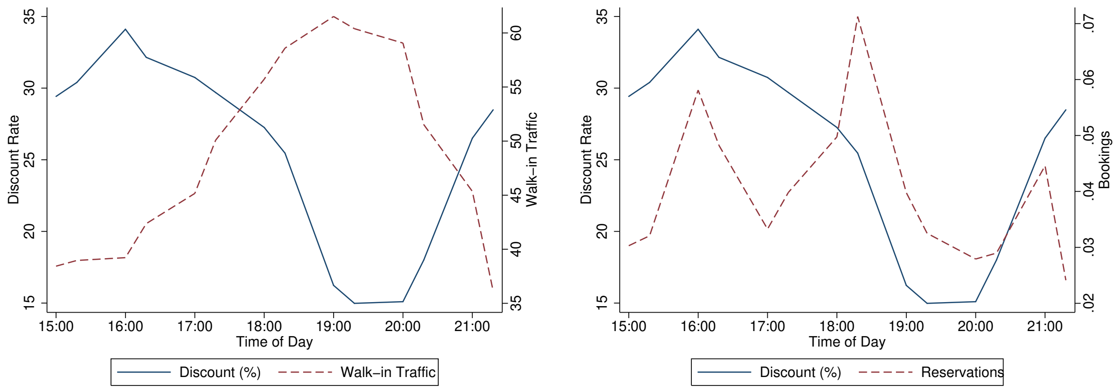

(a) Traffic (b) Reservations

Figure 1: Discounts, Traffic and Reservations by Hour of Day

Note: The lines correspond to per restaurant and time, averaged across restaurants and dates.

should be strictly better off by switching to platform reservation while still dining at their preferred time — and yet receive an average of only 0.1 peak-time reservations. Thus, we view the peak-time discount as restaurants exercising price discrimination between walk-in diners and the (small) price-sensitive segment of platform users, in the absence of substitution from walk-in to platform due to the platform's lack of public awareness.[8]

Restaurants also vary discount rates across days of week, but the magnitude of such variation is much less pronounced than that of intra-day variation. The largest gap in discount rates across dates, observed at 4pm between Sundays and Mondays, is approximately 7.5 percentage points, which is about half of the intra-day price gaps. The smaller price variability across days seems to indicate that the restaurants' primary motivation is to move customers within a day to other slots, rather than moving them to other dates. In Appendix D, we present the price variability across days.

Restaurants rarely adjust their price schedule. Among 80 restaurants in the sample, 20 never adjusted their price schedule, and among the remaining 60 restaurants, the average frequency of any price adjustment is once in 73 days. In contrast, the vast majority of reservations are made last minute. In particular, 88.5% of the reservations are made on the day of dining, and less than 3% of the reservations are made more than three days ahead. This comparison suggests that, unlike hotels and airlines, restaurants' discount rates are set

[8]We account for this third-degree price discrimination between markets in the model, but we keep this incentive fixed throughout the paper as we study the implication of time-varying prices, which is the focus of this paper.

far in advance of reservations and are not adjusted dynamically as capacity fills. Moreover, nearly all time slots remain available on the platform until the last minute (e.g., 30 minutes before the dining time), suggesting that reservations alone — whether through the focal platform or other (unobserved) reservation channels — saturate the full restaurant capacity. Thus, a reservation can crowd out a walk-in customer, but not another customer with a reservation.

3.2 Consumer Responses to Price Variation

We next examine how platform customers respond to price variation by shifting their reservations toward time slots offering higher discounts. This form of intertemporal substitution is central to the effectiveness of peak-load pricing, which aims to redistribute demand away from congested peak hours toward less busy, lower-priced periods.

To this end, we estimate a linear regression of restaurant-level reservations on its effective prices across time slots:

$$\log(R_{j,t,\tau}) = \alpha_0 + \alpha_1 \log(P_{j,t,\tau}) + \alpha_2 \log(\bar{P}_{j,t,\tau\pm1}) + \alpha_3 \log(\bar{P}_{j,t\pm1,\tau}) + \lambda_{j,dow(t),\tau} + \epsilon_{j,t,\tau}. \quad (1)$$

where $R_{j,t,\tau}$ represents the number of reservations for restaurant j on day t at time slot τ, $P_{j,t,\tau}$ denotes the price, $\bar{P}_{j,t,\tau\pm1}$ is the average of the prices at two adjacent time slots, and $\bar{P}_{j,t\pm1,\tau}$ is the average of the prices at the same time on two adjacent days. α_1, α_2 and α_3 represent own elasticity, intra-day cross elasticity to nearby slots, and inter-day cross elasticity, respectively.

Because the platform is marketed as a tool for restaurants to adjust prices in response to variable demand, observed prices may be endogenous to unobserved demand shifts. Identifying instruments for prices is also challenging, as most price variation appears to be sorting consumers across time slots.[9] To address this endogeneity concern, we exploit the observed stickiness in restaurants' variable pricing schedules. In practice, restaurants rarely fine-tune their discount rates on a day-to-day basis. Instead, they typically assign a fixed pricing schedule to each day of the week and maintain it over extended periods, often several months. We argue that such infrequent adjustments are likely responses to gradual, long-run shifts in demand, and the precise timing of any schedule update is plausibly uncorrelated with daily demand shocks. Following this identification strategy, we

[9] Our walk-in traffic measure is not a valid instrument due to its correlation with the ideal dining time among platform customers.

include restaurant $\times$ day-of-week $\times$ time-of-day fixed effects, denoted by $\lambda_{j,\text{dow}(t),\tau}$. These fixed effects absorb persistent heterogeneity in pricing and demand at the restaurant, day-of-week and time-of-day level, allowing us to exploit the occasional changes in effective prices across calendar days within each day-of-week and time-of-day combination.

Table 2 reports the regression results. Column (1) presents the benchmark specification without fixed effects, which may reflect potential confounding due to endogenous pricing decisions. Column (2) includes fixed effects and shows an own-elasticity of -9.27. Column (3) accounts for the substitution effect within a day and across dates. The results reveal that consumers exhibit a meaningful willingness to shift their dining times within a given day, but show limited responsiveness across days. The cross elasticity to the adjacent slot of the same day is 2.09, whereas that to the same time slot to the adjacent day is 0.02. Combined with the relatively limited variation in prices across dates, these results suggest that the primary channel through which peak-load pricing operates is intra-day substitution, rather than inter-day reallocation of demand.

Table 2: The Effect of Prices on Reservations

	(1)	(2)	(3)	(4)
$P_{j,t,\tau}$	-0.840	-9.272	-10.453	-0.039
	(0.027)	(0.421)	(1.039)	(0.006)
$\bar{P}_{j,t,\tau\pm1}$			2.092	0.006
			(0.869)	(0.005)
$\bar{P}_{j,t\pm1,\tau}$			0.021	0.0003
			(0.998)	(0.005)
j-dow(t)-τ FE	No	Yes	Yes	Yes
Observations	149,899	149,878	111,977	111,977

Note: The dependent variable is $R_{j,t,\tau}$. j, t and τ correspond to restaurant, date and time, respectively. In columns 1-3, both the dependent and the independent variables are in logarithms. We add 10^{-10} to $R_{j,t,\tau}$ to keep $\log(R_{j,t,\tau})$ well-defined in case of zero reservations. Standard errors are reported in parentheses.

Because a large share of observations report zero reservations, the log specification of elasticity may be noisy. As a robustness check, we replicate the regression using levels rather than logarithms for both the dependent and independent variables. The results, shown in Column (4), yield quantitatively similar patterns. The implied own-price elasticity, evaluated at the sample mean, is -12.85, and the intra-day cross-price elasticity for adjacent time slots is 1.99.

These results collectively demonstrate that the observed price variation effectively prompts platform customers to shift their reservations from peak to off-peak hours. This behavior aligns with the patterns illustrated in panel (b) of Figure 1. For restaurants facing peak-time capacity constraints, such intra-day time shifting helps increase the overall capacity utilization. Appendix F reports similar elasticity estimates utilizing other sources of identification, including a specification accounting for substitution beyond adjacent slots.

3.3 Possible Role of Price Discrimination

In addition to peak-load pricing, restaurants can also use the variable-pricing ability for intertemporal price discrimination. Higher peak-time prices allow restaurants to extract surplus when platform customers exhibit heterogeneous price sensitivities — for instance, less price-sensitive platform customers may have stronger preference for peak hours than more price-sensitive ones. In this environment, price differences reflect heterogeneity in consumer willingness to pay across time slots, in addition to restaurants' incentives for capacity management.

A direct implication of the price discrimination hypothesis is that consumers who reserve during peak versus off-peak times should differ in their sensitivity to price. To test this, we compare the price sensitivity of peak-time and off-peak-time customers, using the (leave-one-out) average discount rate chosen by the customer. This measure captures the customer's revealed preference for discounts, excluding the focal reservation to avoid mechanical correlation. Panel (A) of Figure 2 displays the distributions. Uncolored bars represent the distribution of average discount rates (across other reservations) for customers who booked a peak-time slot (defined as 7–8pm), while colored bars show the corresponding distribution for customers who booked off-peak times. We observe suggestive evidence that off-peak diners are more price-sensitive: on average, they select higher-discount options (32.4%) in their other bookings, compared to peak-time diners (30.6%).[10]

The observed difference in price sensitivity arises from consumer sorting between peak and off-peak time slots. We find that, conditional on making an off-peak reservation, a customer books another off-peak time slot in 84% of their other reservations, and a peak time slot in only 16%. In contrast, among those who make a peak-time reservation, the corresponding probabilities are 75% for another peak-time booking and 25% for an off-

[10] Kolmogorov–Smirnov tests confirm statistically significant differences between the two distributions. Results are robust to alternative specifications that include the focal reservation in the customer's average discount rate. See Appendix Figure 15.

peak one.

These results serve as some evidence that variable pricing induces consumer sorting across time slots, indicating a possible role of intertemporal price discrimination. However, we also note the economic magnitude is relatively modest: the difference in average discount rates chosen by off-peak versus peak-time diners is only 1.8 percentage points.

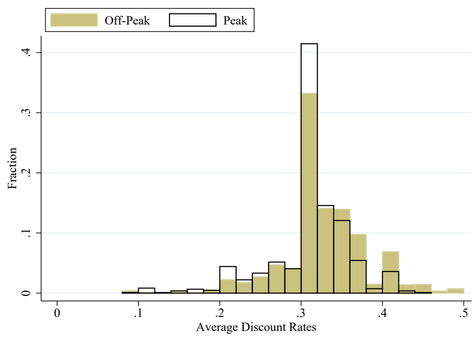

(a) Average Chosen Discount Rates

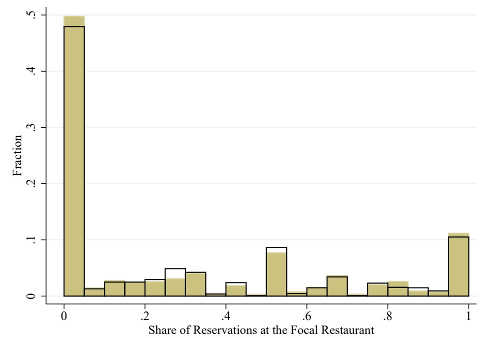

(b) Share of Reservations at the Focal Restaurant

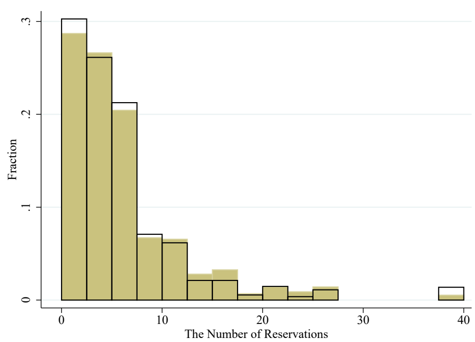

(c) The Number of Reservations

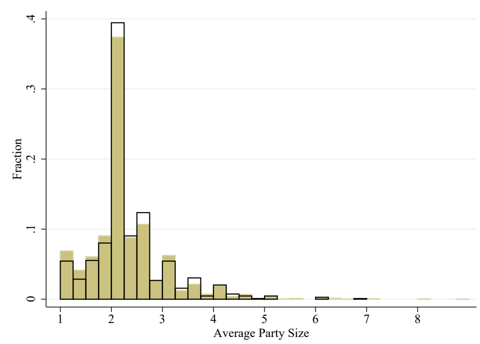

(d) Average Party Size

Figure 2: Distribution of Average Characteristics Between Consumers Choosing Peak and Off-peak Times

Note: Colored and uncolored bars represent the distribution of consumer characteristics by reservations at off-peak and peak slots, respectively. Kolmogorov-Smirnov test statistics are 0.15 (average chosen discount), 0.02 (concentration of reservations), 0.03 (total reservations), and 0.05 (party size).

In panels (B) through (D) of Figure 2, we further examine additional customer characteristics to assess potential differences between peak-time and off-peak-time diners. Specifically, we analyze: (i) the concentration of a customer's reservations at the focal restaurant, shown in Panel (B), which serves as a proxy for the breadth of their choice set; (ii) the total number of reservations made by the customer (Panel C); and (iii) their typical party size (Panel D). In general, there are minimal differences in observable characteristics across customer groups. Kolmogorov-Smirnov tests find statistically significant differences only in party size.

Overall, these descriptive statistics suggest possible, but limited scope for intertemporal price discrimination on the platform. Several factors unique to this platform may explain the relatively homogeneous nature of the platform's user base. Customers who join the platform are likely those actively seeking deals and, on average, more price-sensitive than typical walk-in diners. Additionally, it is possible that the restaurants participating on the platform tend to attract a similar clientele—further reinforcing the homogeneity in consumer responsiveness to price. In our structural model, we allow for heterogeneity in both price sensitivities and time preferences to explicitly account for price discrimination as a potential alternative mechanism. We discuss the theoretical framework next.

4 Model

We develop a model of customer dining reservations and restaurant pricing behavior on the platform. On each day, platform users arrive first and make a reservation decision. Because the data suggests lack of cross-market substitution, we treat the platform as a separate market from that of walk-in customers and assume no substitution between the two. Each restaurant also receives an exogenous and stochastic flow of walk-in customers in each time slot. Restaurants set platform prices while accounting for capacity constraints—specifically, a platform reservation may displace a walk-in customer who would otherwise pay the regular (full) price, creating an opportunity cost for the restaurants.

4.1 Platform Demand

On each date t, consumer i arrives at the platform and decides whether to make a reservation by choosing among all available restaurant-time slot combinations (which we refer to as "products"). The utility of choosing product j is given as follows:

$$u_{ijt} = X_{jt}\beta_i^X - p_{jt}\beta_{i,\tau_{it}}^p - (\tau_j - \tau_{it})^2\beta_{i,\tau_{it}}^\tau + \zeta_{ir_jt} + (1-\sigma_i)\varepsilon_{ijt},$$

where X_{jt} are product characteristics, p_{jt} is the price, τ_j is the time slot of product j (e.g., 4pm), and τ_{it} is the most preferred dining time of consumer i on that day. The cost of switching dining time is quadratic in the distance between the chosen and the ideal time slots, with coefficient $\beta^{\tau}_{i,\tau_{it}}$. We normalize the utility of the outside option to zero, which represents not making a reservation on the platform for that day.

We assume that consumer preferences over restaurant characteristics (e.g., cuisine type) are fixed across days and drawn from a finite set of discrete types, with the probability mass function $G(\beta_i)$. In contrast, consumers' ideal dining times are allowed to vary across days and are orthogonal to their restaurant preferences. These time preferences follow the distribution $F_t(\tau_{it})$, varying by day of the week, to capture systematic temporal patterns in consumer habits during the week. We let consumer's price sensitivity ($\beta^{p}_{i,\tau_{it}}$) and cost of deviating from their ideal dining time ($\beta^{\tau}_{i,\tau_{it}}$) be unique to each combination of the discrete preference type and the ideal dining time, allowing flexible correlations between a consumer's price sensitivity and her preferred dining time. This specification thus allows time-varying price elasticity — for instance, a consumer can be more price sensitive and is more willing to switch dining time when she prefers to dine early, because an early dinner is typically associated with a casual dining occasion with some scheduling flexibility.[11]

We assume that ε_{ijt} follows a Type 1 Extreme Value distribution, and that ζ_{ir_jt} is i.i.d. at the restaurant level (which we denote by r_j) and follows a distribution such that $\zeta_{ir_jt} + (1-\sigma_i)\varepsilon_{ijt}$ also follows a Type 1 Extreme Value distribution (Cardell, 1997). Thus, the choice probability of a customer is given by the following nested logit choice probability, where each nest corresponds to a restaurant:

$$Pr_{it}(j|X_t, p_t, \tau_t) = \frac{\exp(\delta_{ijt}/(1-\sigma_i))\left(\sum_{j'\in r}\exp(\delta_{ij't}/(1-\sigma_i))\right)^{-\sigma_i}}{1+\sum_r(\sum_{j'\in r}\exp(\delta_{ij't}/(1-\sigma_i)))^{(1-\sigma_i)}},.$$

where $\delta_{ijt} = X_{jt}\beta^X_i - p_{jt}\beta^p_{i,\tau_{it}} - (\tau_j - \tau_{it})^2\beta^{\tau}_{i,\tau_{it}}$, and X_t, p_t and τ_t are the collections of observed characteristics, prices, and time-of-day's for all products available on day t.

Among M platform customers, the total demand for product j is $q_{\tau_jt} = M \times Pr_t(j|X_t, p_t, \tau_t)$, where $Pr_t(j|X_t, p_t, \tau_t)$ is the per-customer choice probability, integrated over consumer preferences and the ideal dining time, as follows:

$$Pr_t(j|X_t, p_t, \tau_t) = \int_{\beta_i}\int_{\tau_{it}} Pr_{it}(j|X_t, p_t, \tau_t)dF_t(\tau_{it})dG(\beta_i).$$

[11]Note that even without time-varying $\beta^{\tau}_{i,\tau_{it}}$, consumer preference for each time slot is heterogeneous across segments. $\beta^{\tau}_{i,\tau_{it}}$ thus serves as an extra layer of intertemporal variation to further increase the model flexibility to capture the restaurant incentive for intertemporal price discrimination.

4.2 Restaurant Pricing Problem

We model the restaurants' pricing problem on the platform as a static Bertrand competition among multi-product firms. On each calendar day t, each restaurant r offers a set of time slots denoted by J_{rt}. We treat this set J_{rt} as given to the restaurant, because the presence of the platform does not affect restaurant operating hours. We define a restaurant's total profit as the sum of its profit from platform reservations and from walk-in customers, as follows:

$$\pi_{rt} = \sum_{j \in J_{rt}} \left(q_{\tau_j t}(p_{jt} - c_r - c_p) + \mathbb{E}[\min\{n_{\tau_j t}, L_r - q_{\tau_j t}\}](p_r^{regular} - c_r) \right), \tag{2}$$

where c_r is the marginal cost of serving a customer, c_p is a per-reservation fee the restaurant pays to the platform, $n_{\tau_j t}$ is the size of walk-in demand for time slot τ_j on day t, L_r is the total capacity of the restaurant, and $p_r^{regular}$ is the regular price charged to walk-in customers, assumed as given and constant across time slots.[12]

The first term of expression (2) is the restaurant's profit from platform customers, net of the fees it pays to the platform. The second term corresponds to the profit from walk-in customers. Because customers with a reservation are prioritized, the number of walk-in guests the restaurant can serve at each τ_j is the minimum of the walk-in demand $(n_{\tau_j t})$ and the available capacity ($L_r - q_{\tau_j t}$). This sales crowd-out creates the restaurant's incentive to raise prices (lower discounts) on the platform in time slots where capacity constraints are more likely to bind.

Because $n_{\tau_j t}$ is unknown at the time of pricing, the restaurant takes expectation over its realizations.[13] Let $x_{\tau_j t} = L_r - n_{\tau_j t}$ be the difference between the restaurant capacity and the number of walk-in guests at time τ_j, which represents the number of tables the platform can sell on the platform without crowding out walk-in customers. We assume that $x_{\tau_j t}$ is independent across τ_j and t, and follows a Poisson distribution with rate $\lambda_{\tau_j t}$. In our

[12]Prices are measured per person, and hence a unit sales in Expression (2) corresponds to one guest. This formulation is without loss of generality: In Appendix G, we show that even when the average party size exceeds one and potentially differs between the platform and walk-ins, it only affects the interpretation of L_r and $n_{\tau_j t}$.

[13]In principle, the platform reservation is also a random variable whose realization is unknown at the time of pricing. We include its expected value as part of the capacity constraint term ($L_r - q_{\tau_j t}$) and abstract from calculating its full expectation jointly with the expectation on $n_{\tau_j t}$. The latter approach creates a computationally intractable first-order condition with respect to p_{jt}.

empirical application, we parameterize $\lambda_{\tau_j t}$ as follows:

$$\lambda_{\tau_j t} = L_r - \underbrace{(\gamma_{1r} g_{\tau_j t} + \gamma_{2r} g^2_{\tau_j t})}_{\bar{n}_{\tau_j t}} - \sum_{\tau'=1}^{2} \underbrace{(\gamma_{3r\tau'} g_{\tau_j+\tau',t} + \gamma_{4r\tau'} g^2_{\tau_j+\tau',t})}_{\bar{n}_{\tau_j+\tau',t}}, \tag{3}$$

where $g_{\tau_j t}$ is our measure of walk-in traffic at period τ_j on date t, as constructed using Google Maps data. We assume that the expected walk-in demand (denoted by $\bar{n}_{\tau_j t}$) is quadratic in the concurrent value of the traffic measure. In addition, if guests occupy the table for more than 30 minutes, a reservation at time τ_j may also displace walk-in customers in later periods. Although fully accounting for this intertemporal dependence in capacity constraints is computationally costly, we allow for this possibility in a parsimonious manner, by allowing $\lambda_{\tau_j t}$ to also depend on expected walk-in demand in subsequent periods (denoted by $\bar{n}_{\tau_j+\tau',t}$). Accounting for subsequent traffic volume allows the model to capture the restaurant incentive to proactively adjust the price in anticipation of traffic at later times. We assume that restaurants consider up to two subsequent 30-minute periods and each $\bar{n}_{\tau_j+\tau',t}$ is a quadratic function of its concurrent value of $g_{\tau_j+\tau',t}$.

Each restaurant chooses prices (denoted by $\{p_{jt}\}_{j \in J_{rt}}$), subject to the constraint that platform prices do not exceed the regular price. By structure, the per-unit margin from platform customers is weakly lower than that from walk-in customers. Solving the maximization problem yields the following first-order conditions:

$$q_{\tau_j t} + \sum_{j' \in J_{rt}} \frac{\partial q_{\tau_{j'} t}}{\partial p_{jt}} (p_{j't} - c_r - c_p) - \sum_{j' \in J_{rt}} Pr(q_{\tau_{j'} t} > x_{\tau_{j'} t}) \frac{\partial q_{\tau_{j'} t}}{\partial p_t} (p_r^{regular} - c_r) = 0, \tag{4}$$
$$\forall j \in J_{rt} \text{ s.t. } p_{jt} \leq p_r^{regular}.$$

The first two terms are equivalent to the standard condition that equates the marginal revenue to the marginal cost. The third term represents the opportunity cost of accepting one additional reservation. With probability $Pr(q_{\tau_{j'} t}(p_t) > x_{\tau_{j'} t})$, the capacity binds at time $\tau_{j'}$ and the restaurant needs to turn away one walk-in customer, leading to a profit loss equal to $p_r^{regular} - c_r$. The Bertrand equilibrium is established when all restaurants' prices for all time slots satisfy Expression (4) given the prices of other restaurants.

This pricing model captures the two key incentives underlying variable pricing. First, consumers whose most preferred dining time is τ_j may exhibit different price sensitivities than those who prefer another time $\tau_{j'}$, creating an incentive for intertemporal price discrimination. Second, variation in the expected residual capacity $x_{\tau_j t}$ across time slots creates an incentive for peak-load pricing. Specifically, restaurants have an incentive to raise prices during peak hours, while lowering them at nearby, less-congested times, to induce consumer substitution away from peak times.

5 Parametrization, Identification, and Estimation

5.1 Parametrization and Identification

On the demand side, we assume that the preference distribution $G(\beta_i)$ is bimodal, representing two discrete consumer types for preferences over restaurant characteristics. We identify these types and their shares by leveraging the systematic difference in reservation patterns across users.

We also assume that consumers' ideal dining time distribution (denoted by $F_t(\tau_{it})$) mirrors that of walk-in consumers. Specifically, we proxy $F_t(\tau_{it})$ using the empirical distribution of our walk-in traffic measure across time for each day of week, aggregated across restaurants. We argue the realized walk-in dining times match with the ideal dining times, because walk-in customers likely choose their most preferred time in the absence of platform discounts.[14] We discuss more details on the construction of $F_t(\tau_{it})$ in Appendix A.

To identify the price coefficients, we exploit the stickiness in restaurant pricing, the strategy we used in Section 3.2. We assume that the control variables (X_{jt}) accounts for all predictable demand shifts across restaurants and times, and that the residual price variation is orthogonal to contemporaneous demand shocks: the changes in pricing schedules are infrequent and likely are a response to a gradual, long-term shift in demand rather than its daily fluctuations. Although the identification was established with multiple interactive fixed effects in Section 3.2, we aim to mimic the identification strategy with a smaller set of variables in X_{jt} for computational performance.

We include in X_{jt} the following variables: restaurant Google review rating, regular price, the restaurant's relative popularity at each day of week $\times$ time of day, local restaurant density in the neighborhood, cuisine type fixed effects, and district fixed effects. In addition, we include fixed effects for year-month pairs and for weekend days. To measure

[14] The realized dining time may underestimate the true dining demand at peak times due to restaurant capacity constraints. To address this concern, we reestimate the demand exclusively using the traffic distribution among restaurants that we identify as not facing capacity constraints. Our results are robust to this modification.

a restaurant's relative popularity at a given time slot, we use the ratio of its specific traffic value ($g_{\tau_j t}$) to the aggregated market-level traffic at that time. Local restaurant density is measured by the number of restaurants within a 1 km radius, using auxiliary data from Google Maps that includes restaurants both on and off the platform.[15] In Section 6.1, we show that including these control variables produces structural elasticity estimates similar to the reduced-form results in Section 3.2, providing empirical support for utilizing these control variables.

On the supply side, the parameters of interest are the marginal cost c_r and the expected available capacity at each time slot, $\lambda_{\tau_j t}$, which is parameterized with the intercept, L_r, and the coefficients of the traffic measure, γ_r. We allow these parameters to be restaurant-specific. We assume c_p, the fee to the platform, is known and fix it at the (restaurant-specific) value we learned through the conversation with the platform. We identify c_r using periods in which the capacity constraint does not bind: the third term of Expression (4) is near zero, and the expression collapses to the standard pricing first-order condition in which c_r is the only unknown term. The terms L_r and γ_r are identified by how prices vary with walk-in traffic, $g_{\tau_j t}$ and $g_{\tau_j+\tau',t}$.[16] Intuitively, when a restaurant sets a price that is different from what would be justified solely by demand elasticity and the residual prices are correlated with walk-in traffic, we interpret this pricing behavior as indicative of peak-load pricing incentive.

This identification strategy exploits the uniqueness of our market structure, in which the two pricing incentive arises from separate channels. On the one hand, variable walk-in demand affects peak-load pricing incentives, but does not create intertemporal price discrimination, because walk-in diners themselves do not face platform discounts. On the other hand, heterogeneous platform consumers can create an intertemporal price discrimination incentive, but the small platform demand by itself does not create peak-load incentives. Our identification is hence akin to preference externality (Waldfogel, 2003), in which restaurant pricing in a "small" (platform) market is influenced by demand shocks in "large" (walk-in) markets, but prices do not affect the outcomes of large markets. This structure substantially simplifies an otherwise nontrivial separation task: if such an external shift in capacity does not exist (e.g., suppose instead no external market exists and peak-load

[15]This term is included to control for aggregate demand shifters in the area and we do not intend to structurally interpret its coefficient. The coefficient is negative if higher density results in a higher competitive pressure, whereas it could be positive if higher restaurant density suggests a more active business landscape.

[16]The restaurant-specific coefficients on these terms also account for the fact that $g_{\tau_j t}$ is scaled at the restaurant level. Although $g_{\tau_j t}$ is not directly comparable across restaurants, our estimates of $\lambda_{\tau_j t}$ are.

pricing is also triggered by the volume of platform users), separating these two incentives typically requires separation between the intertemporal shift in demand levels (affecting capacity constraints), and that in demand slopes (affecting price discrimination incentives), rendering the separation less transparent and more data-intensive.

5.2 Estimation

We estimate the model in two steps. First, we estimate the demand with maximum likelihood. Formally, the likelihood function is given as follows:

$$L(\theta) = \prod_i \left(\int_{\beta_i} \prod_{t=\underline{t}_i}^{189} \left(\int_{\tau_{it}} Pr_{it}(j_{it}|X_t, p_t, \tau_t) dF_t(\tau_{it}) \right) dG(\beta_i) \right).$$

For each consumer i, the likelihood corresponds to the joint probability that the consumer takes the observed action (j_{it}) on each day during the sample period. We denote by $\underline{t}_i$ the first day that consumer i is observed in the data, which we define as the day following their first reservation.[17]

We estimate the supply-side parameters using generalized method of moments (GMM). For each restaurant, we stack its first-order conditions across days and time slots and minimize the squared sum. Estimation is separable across restaurants because both the set of supply-side parameters and the set of first-order conditions are restaurant-specific.[18]

6 Estimation Results

6.1 Demand

Table 3 reports the parameter estimates of the demand model. We estimate two consumer segments. Segment 1 represents the majority—accounting for 75.6% of the customer base. Both segments show low intercepts, driven by the fact that an average consumer makes only four reservations over the six-month period. For each segment, we estimate four price

[17]We exclude periods prior to the first reservation because we cannot observe whether the consumer was active on the platform during that time. Additionally, we drop the first reservation itself because, by construction, the first observed choice always corresponds to an "inside good."

[18]We choose the number of γ_r terms unique to each restaurant, to balance model fit and statistical power at the restaurant level. We discuss further details of the estimation in Appendix I.

coefficients and the cost of switching time slots, corresponding to the consumer's most preferred dining time being each 60-minute interval, except for the late night in which we consolidate coefficients into a single 120-minute interval.

	Segment 1		Segment 2	
Intercept	-6.779	(0.472)	-12.275	(1.435)
Price (τ_{it} ∈[3pm, 4pm])	-0.395	(0.007)	-0.712	(0.017)
Price (τ_{it} ∈[4:30pm, 5:30pm])	-0.420	(0.009)	-0.763	(0.042)
Price (τ_{it} ∈[6pm, 7pm])	-0.348	(0.005)	-0.625	(0.015)
Price (τ_{it} ∈[7:30pm, 9:30pm])	-0.373	(0.006)	-0.612	(0.017)
Switching cost (τ_{it} ∈[3pm, 4pm])	-0.409	(0.069)	-0.036	(0.026)
Switching cost (τ_{it} ∈[4:30pm, 5:30pm])	-0.392	(0.100)	-10.92	(228.6)
Switching cost (τ_{it} ∈[6pm, 7pm])	-0.273	(0.038)	-0.961	(0.265)
Switching cost (τ_{it} ∈[7:30pm, 9:30pm])	-0.445	(0.056)	-1.132	(0.397)
Review Ratings	-0.301	(0.088)	0.240	(0.294)
Popularity	0.087	(0.025)	0.240	(0.065)
Regular Price	0.227	(0.001)	0.522	(0.002)
Nest Parameter	0.696	(0.019)	0.061	(0.076)
Fraction Segment 1	0.756	(0.014)		

Note: Standard errors are reported in parentheses. The specification also includes fixed effects for year-month pairs, for weekend days, for restaurant locations (at district level), cuisines and for quartile bins of restaurant density in the neighborhood. The coefficients of these fixed effects are segment-specific. In addition, we include interactions between restaurant locations and cuisines whose coefficients are common across segments.

Table 3: Demand Parameter Estimates

Among segment 1 consumers, we find that price sensitivity is mostly flat during the day. The differences across time periods are economically small though statistically significant. The estimated cost of switching between time slots is lowest when τ_{it} is between 6 and 7 pm. The willingness to accept to switch to the adjacent time slot—given by the ratio between the switching cost parameter and the price coefficient— have a relatively narrow range within a day, between €0.78 (for τ_{it} between 6 and 7 pm) and €1.19 (for τ_{it} after 7:30 pm). Segment 2 has a large negative intercept, indicating this segment rarely dines out. This segment exhibits a slightly higher price sensitivity during early times, with larger intertemporal variation in switching costs. Their lowest switching cost is observed when their preferred time is early, representing a near-zero cost to switch to an adjacent slot.

These consumers also do not substitute to other slots when their preferred time is between 4:30 and 5:30 pm.[19] Overall, both segments show willingness to shift their dining time in response to discounts (with the exception of segment 2 around 5 pm), with approximately €1 in savings (or about 22% of the average discount offered in the data) being sufficient to induce a 30-minute shift in their dining time.

The coefficients on control variables largely align with expectations. The estimated coefficients for restaurant popularity and regular (pre-discount) price are both positive and statistically significant. Conditional on the actual discounted price, the regular price likely proxies for perceived restaurant quality, thereby increasing consumer utility. The coefficient on review ratings for segment 1 shows an unexpected sign. This may be due to the limited variation in the data: the platform only recruits restaurants with high Google ratings, leading to the rating's interquartile range of just 0.2 (from 4.2 to 4.4). As a result, review scores may carry limited informational value for consumers. Finally, the nest parameters indicate that segment 1 consumers mostly switch to other slots of the same restaurant, whereas segment 2 are more willing to substitute across restaurants.

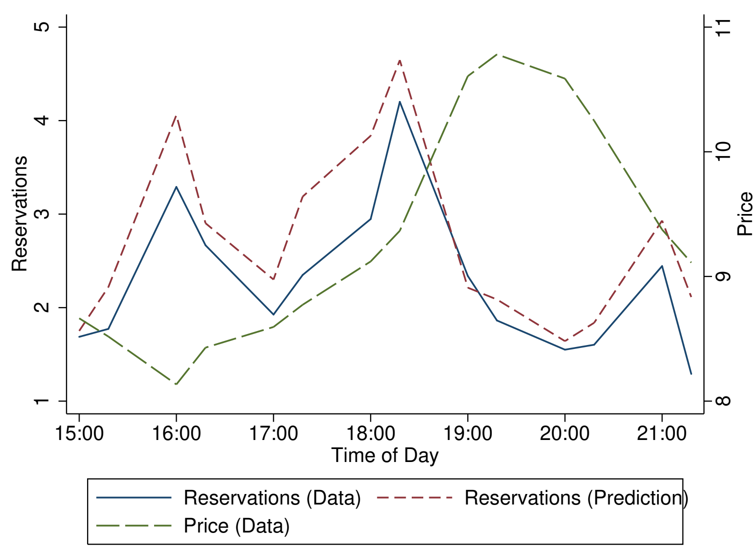

Figure 3: Demand Model Fit

Note: The lines correspond to the average across restaurants and dates.

Figure 3 shows the model fit for the within-day reservations. Overall, the model tracks the key patterns of the data well, including the timing and magnitude of peak and off-peak

[19]The large standard error is because any switching cost that is consistent with lack of substitution is rationalized.

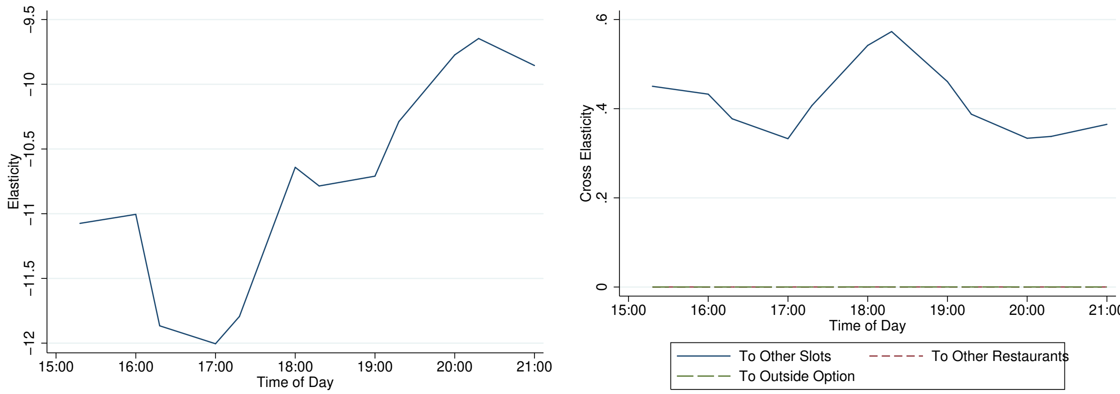

(a) Own Elasticity

(b) Cross Elasticity

Figure 4: Across-time Elasticity Variations within a Day

Note: In Panel B, "To Other Slots" corresponds to the average cross elasticity to other time slots within the restaurant at each date, averaged across dates, "To Other Restaurants" is the average cross elasticity to other restaurant - time slot pairs, averaged across restaurants and dates, and "To Outside Option" is the cross-elasticity to the outside option, averaged across dates.

periods.

In Panel A of Figure 4, we show the average own-price elasticity within the day and across time slots. We estimate an own-price elasticity between -10 and -12, indicating that consumers on the platform are quite price sensitive, likely more so than average diners. The model-implied price elasticities are close to the reduced-form estimates presented in Table 2, suggesting that the model's control variables can capture unobserved demand relevant to pricing decisions, to a similar degree as the granular interacted fixed effects in the reduced-form specification. We also conduct further robustness checks of the selection of control variables in Appendix H. We also find declining elasticity over time: Diners during the early hours are more price sensitive than ones in later hours, a variation that explains the drop in reservations around 5 pm.[20] As we show in Section 7.1, this shift in elasticity is not large enough to enable meaningful intertemporal price discrimination.

Panel B of Figure 4 displays the average cross-price elasticities with respect to (i) other time slots at the same restaurant, (ii) other restaurants, and (iii) the outside option. The results show that nearly all substitution occurs toward other time slots within the same

[20]The model mechanically generates lower elasticity at the start and the end of the day, because these slots lack one adjacent slot. Figure 4 excludes these corner periods. The lower elasticity at the corner does not affect our other results, because these are the periods with small demand and sales.

restaurant (as per the solid line), while substitution to other restaurants is negligible. Although the estimated elasticity with respect to the outside option is small, substitution to the outside good is large in absolute magnitude because consumers choose the outside option approximately 98% of the time.

6.2 Supply

In Figure 5, we show the distribution of marginal cost estimates. The average marginal cost is €7.55, and the standard deviation is 6.63. The average profit margin is 59.8% of the regular price or 45.1% of the discounted price. Our marginal cost estimates are statistically significant for 77 out of the 80 restaurants.

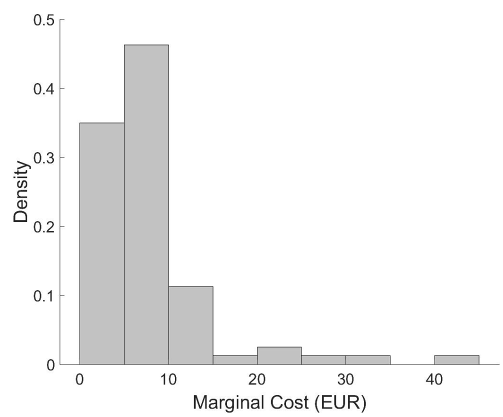

Figure 5: Marginal Cost

Note:

Figure 6 shows the predicted capacity constraint at each time slot, $\lambda_{\tau_j t}$. The value of $\lambda_{\tau_j t}$ represents the expected number of tables the restaurant can sell on the platform without crowding out the walk-in demand. Panel A displays the histogram across the whole sample, excluding restaurant-time pairs that indicate less than one percent chance of binding capacity.[21] In our data, 79% of observations exhibit more than one percent of binding capacity, i.e., an informative $\lambda_{\tau_j t}$ value. The average and the median value of $\lambda_{\tau_j t}$ are 1.79 and 1.57, respectively, with the standard deviation of 0.96.

[21]Because each restaurants receive an average of only 0.5 reservations per day, most values of $\lambda_{\tau_j t} > 4$ are as good as completely unbounded capacity and do not affect the platform prices. The estimates of $\lambda_{\tau_j t}$ in this range is obtained as a pure parametric extrapolation and is completely uninformative and irrelevant for our study.

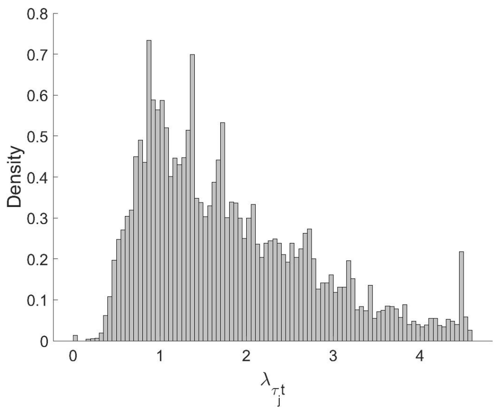

(a) Overall distribution

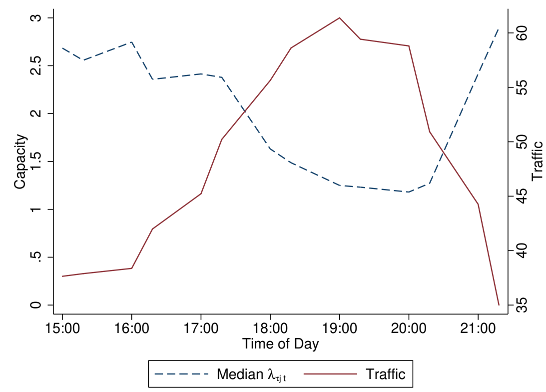

(b) Across-time Variations

Figure 6: Capacity Constraints ($\lambda_{\tau_j t}$)

Note: In Panel A, $\lambda_{\tau_j t}$ estimates that imply less than one percent chance of binding capacity are not informative and hence are dropped. In Panel B, the dashed line corresponds to the median value (across restaurants and dates) in each time slot, and the solid line corresponds to the average traffic volume across restaurants and dates.

Panel B displays the intra-day variations in $\lambda_{\tau_j t}$ evaluated at the sample median. $\lambda_{\tau_j t}$ is negatively correlated with the concurrent value of the traffic measure—the model correctly captures restaurants' peak-load pricing incentives by linking higher observed prices to periods of higher traffic volume.

In Figure 7, we present the sellout probability for each time slot—defined as the probability that the demand exceeds available capacity, namely $Pr(q_{\tau_j t} > x_{\tau_j t})$. Panel A displays the overall distribution of sellout probabilities (colored bars), overlaid with the distribution of each restaurant's maximum sellout probability at the busiest time (uncolored bars). On average, the sellout probability across all observations is 0.188, with a standard deviation of 0.177. Focusing on the busiest time for each restaurant, we find that 32 restaurants (40% of the sample) experience a sellout probability greater than 50%, and 2 restaurants exceed 80%. Panel B displays the within-day variation in sellout probability, evaluated at the sample median. During peak hours between 7pm and 8pm, restaurants face an average 30% chance of hitting their capacity limit. These findings underscore the importance of peak-load pricing: during high-demand periods, offering platform discounts incurs a substantial opportunity cost for restaurants, reinforcing the incentive to raise prices when capacity is tight.

Figure 8 shows the model fit for the intra-day pricing patterns. The solid line repre-

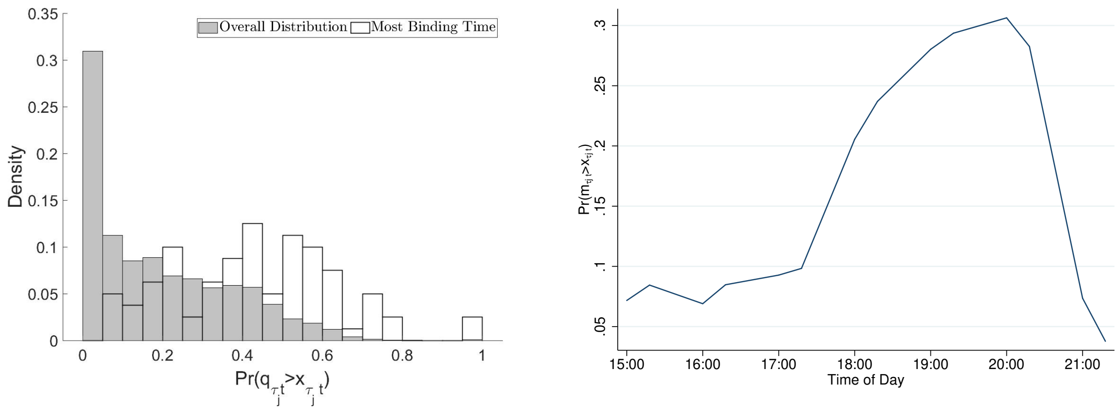

(a) Distribution

(b) Across-time Variations

Figure 7: Probability of Binding Capacity ($Pr(q_{\tau_j t} > x_{\tau_j t})$)

Note: In Panel A, "Most Binding Time" corresponds to the maximum value of $Pr(q_{\tau_j t} > x_{\tau_j t})$ for each restaurant during the sample period. Panel B corresponds to the median (across restaurants) of the average probability of binding capacity of each restaurant in each time slot.

sents the observed average prices across time slots, and the dashed line shows the model predicted prices, computed by solving for the optimal prices for each restaurant at the estimated parameter values. Overall, the predicted prices closely track the observed variation throughout the day, although the model predicts the restaurants should raise the prices earlier than observed in the data - as seen in panel (a) of Figure 1, the observed change in discount rates lags the surge of the traffic. The average predicted price is €13.72, compared to an average observed price of €13.21.

Overall, we find that this platform provides an ideal environment for restaurants to exercise peak-load pricing. The platform matches restaurants facing substantial capacity constraints with price-sensitive consumers who are willing to change their dining times with relatively small compensation. Moreover, small cross-restaurant substitution reduces the restaurant's risk of losing customers to competitors due to variable pricing. Although we model the supply side as a Bertrand equilibrium, the optimal prices resemble that of a single-agent decision, and strategic interactions between restaurants are not relevant when evaluating the welfare implications of variable pricing. Thus, the platform serves as an ideal ground to evaluate the welfare-improving role of peak-load pricing in markets with profit-maximizing firms, although we acknowledge that the our substantive findings should be viewed as a best-case scenario rather than the representative of the overall dining markets,

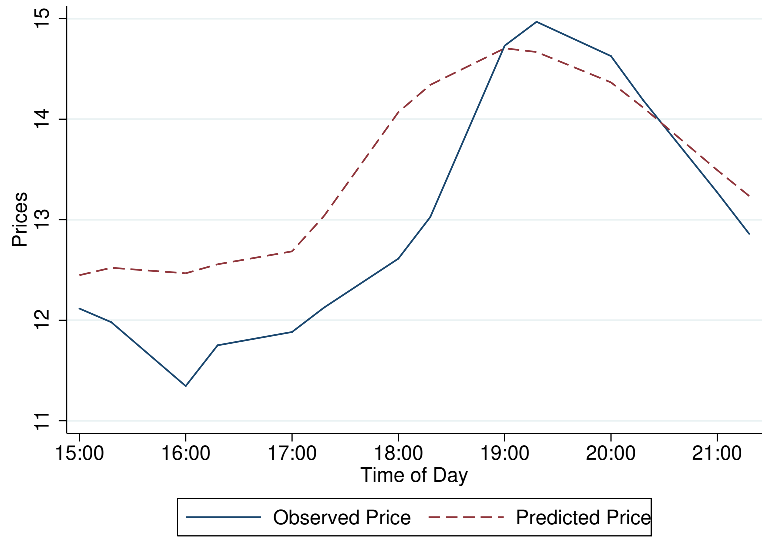

Figure 8: Model Fit for Prices

Note: The lines correspond to the average across restaurants and dates.

due to the selection of users and restaurants.

7 Welfare and Profit Implications of Peak-load Pricing

We now use the model to analyze the mechanisms and welfare implications of variable pricing in our focal market. In Section 7.1, we show that the primary driver of observed price variation is peak-load pricing rather than intertemporal price discrimination. In Section 7.2, we show that variable pricing improves overall welfare and the profit of restaurants and the platform, thus providing evidence that peak-load pricing improves welfare even when the firm's objective is profit maximization. In Section 7.3, we consider a profit-sharing scheme between restaurants and the platform and show that it can create misaligned incentives between them. We show that, although peak-load pricing under profit-sharing scheme can produce a substantially larger welfare gain than the observed fixed-fee contract, the platform has less incentive to allow the use of such a pricing strategy. Thus, the incentive misalignment partly offsets the relative efficiency of profit-sharing contracts over per-unit-fee contracts.[22]

[22]In all counterfactual results, we recompute the equilibrium prices under the assumed regime.

7.1 The Role of Capacity Constraints in Determining Prices

We begin by conducting two sets of counterfactual analyses to identify the relative magnitude of peak-load pricing and intertemporal price discrimination as drivers of the restaurants' pricing behavior. In the first scenario, we eliminate intertemporal price discrimination by assuming time-invariant consumer price sensitivity. Specifically, for each consumer segment, we fix the price coefficient across all time slots at the average of the four estimated coefficients for that segment. In the second scenario, we remove the incentive for peak-load pricing by assuming that restaurants have unlimited capacity. These two counterfactuals isolate the distinct contributions of price discrimination and capacity constraints.

Figure 9 displays the intra-day price variability under each scenario. The solid line represents the average prices under the baseline (factual) model, the short-dashed line shows prices under time-invariant price sensitivity, and the long-dashed line depicts prices under the assumption of unlimited capacity.

We find that capacity constraints have a substantial influence on the pricing schedule. When we remove the capacity constraint, the average price drops by 15%, from €13.72 to €11.62. This decline is more pronounced during peak periods than off-peak hours, leading to a significantly flatter intra-day pricing profile. At the restaurant level, the standard deviation of prices across time slots falls by 88% on average, representing a marked reduction in price dispersion. In contrast, eliminating the intertemporal variation in consumer price sensitivity has a more limited impact. Figure 9 shows the price patterns from this counterfactual are virtually identical to the factual case. This result is consistent with our demand estimates, which reveal limited variation in price sensitivity across times of day. Taken together, these findings indicate that the observed intra-day price variability arises almost entirely from restaurants' peak-load pricing incentives, rather than from price discrimination between platform users.

7.2 Welfare Effects of Price Variability

We now formally quantify the welfare impact of peak-load pricing. Specifically, we compare the welfare between the observed pricing regime and the counterfactual scenario in which restaurants are constrained to a uniform discount rate across all time slots of the day (allowing prices to vary across days).[23] We hold other market conditions fixed, such

[23] Uniform discount scheme exclusively eliminates time-varying prices, while holding fixed the restaurants' ability for third-degree price discrimination between the platform and walk-ins.

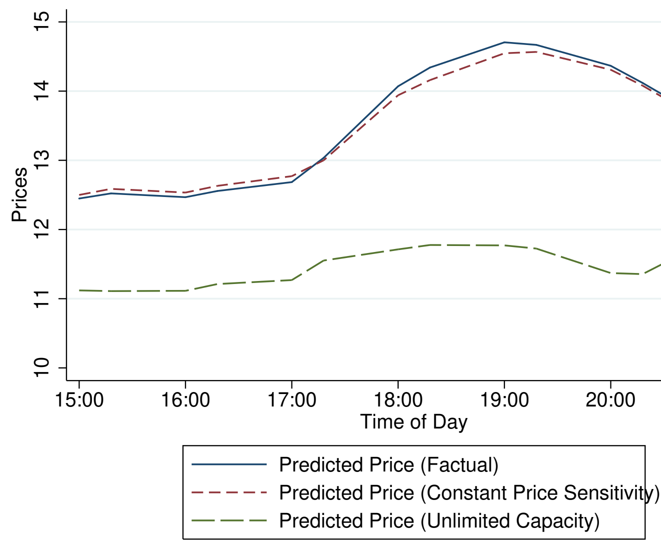

Figure 9: Predicted Prices

Note: The lines correspond to the average across restaurants and dates.

as platform participation of users and restaurants, and restaurants' opening hours.[24] Our welfare measure is defined as the sum of platform customers' surplus, restaurant profits and the platform profit. The restaurant profit is given in Expression (2) and includes profits from both platform reservations and walk-in customers. The platform profit consists of a fixed per-unit fee it collects from the restaurants.[25] We assume the per-unit fee, c_p, remains unchanged in the counterfactual. Note that the surplus of walk-in customers is excluded from our welfare calculation because it is unobservable. Because peak-load pricing always (weakly) improves the surplus of walk-in customers through better table availability, our analyses present a lower bound for the actual welfare improvement.[26]

The results are reported in Table 4. Columns (1) and (2) correspond to the variable- and uniform-pricing scenario, respectively. Panel A reports the overall welfare outcomes, documenting the welfare-improving role of variable pricing. Under variable pricing, the

[24]The former can be restrictive if the platform differentiates itself by this variable-pricing ability. The data suggests that the platform offers restaurants access to new consumer segments, and hence restaurants are likely better off staying even in the absence of variable pricing. Consumers are also strictly better off by making a reservation with discounts than showing up as a walk-in, thus unlikely to substitute away merely due to the lack of time variability in discounts.

[25]Formally, the platform profit from restaurant r on day t is given by $\pi_{platform,rt} = c_p \sum_{j \in J_{rt}} q_{\tau_j t}$.

[26]The level of the profits and welfare this platform generates is relatively small, due to the small scale of the platform. Because our goal is to document the welfare benefits of peak-load pricing, we focus on the percentage change in profits and welfare between the variable- and uniform-pricing regimes.

	(1) Variable	(2) Uniform
Panel (A): Gains Generated by the Platform		
Total Welfare Generated by the Platform (per Day, €)	182.5	168.1
Restaurant Profits from Participation	69.85	64.66
Platform Profit	37.60	34.28
Platform Users' Surplus	75.02	69.20
Panel (B): Effect on The Platform		
Daily Reservations	30.57	28.23
Average Price Booked (€)	13.12	13.85
Fraction of Peak-Time Reservations (7-8 pm)	0.217	0.295
Restaurant Profit per Reservation (€)	4.060	4.789
Surplus per Reservation (€)	6.970	7.695
Panel (C): Effect on Walk-in Customers		
Displaced Walk-in Customers per Reservation	0.151	0.209
Lost Profit due to Displacement per Reservation (€)	1.633	2.345

Note: The second through fourth rows of panel (A) sum to the first row. "Displaced walk-in customers per reservation" is given by the difference in the expected number of walk-in customers with and without the platform reservations, divided by the number of platform reservations. "Lost profit due to displacement" is calculated analogously, but using the difference in the expected profit from walk-in customers as the variable of interest. Formal definition of all the metrics is available in Appendix K.

Table 4: Impact of Variable Pricing

platform generates the welfare of approximately €183 per day, compared to €168 under uniform pricing. This represents a lift of roughly 8.6%.

Almost all the restaurants benefit from adopting variable pricing. The total profit of restaurants from platform participation — that is, the profit from platform reservations net of displaced walk-in profit, summed across restaurants — is €69.9 per day under variable pricing, compared to €64.7 under uniform pricing, representing a 8% increase. Figure 10 displays the distribution of the percentage profit change across restaurants. Approximately 60% of the restaurants see a modest profit increase of up to 10% relative to their uniform-pricing baseline, while two restaurants experience profit gains of over 100%. One restaurant exhibits a slight profit decline, likely due to intensified price competition, but the magnitude of the loss is negligible (0.13%). These minor losses suggest that any competitive disadvantages from adopting variable pricing are rare and economically small.

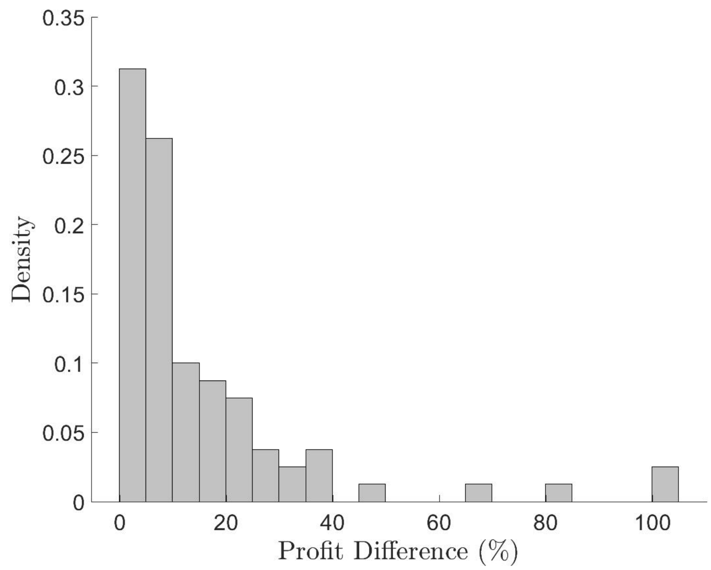

Figure 10: Profit Increase due to Variable Pricing (% of Uniform Pricing Profit)
Note: The unit of analysis is a restaurant - date pair. The rightmost bar denotes an increase above 100%.

The platform also benefits from variable pricing. The total fee the platform collects increases from €34.3 to €37.6 per day, a relative 9.6% increase. As we show in Panel (B), variable pricing substantially increases the number of reservations on the platform by lowering average prices, thereby increasing the fees the platform collects. The users' surplus also increases by 8.4% from €69.2 to €75.0. Overall, these findings indicate that variable pricing benefits both restaurants and the platform in a roughly equal manner under the per-unit-fee contract in the data.

Panels B and C of Table 4 decompose the total welfare gain into on- and off-platform transactions, thereby illustrating how peak-load pricing improves welfare through capacity reallocation. Panel B focuses on outcomes within the platform. We find that, under uniform pricing, 28.2 reservations are made daily across all the restaurants, at an average price of €13.85. In contrast, under variable pricing, 30.6 daily reservations are made at a lower average price of €13.12. Under uniform pricing, approximately 29.5% of all reservations occur during peak hours (7 to 8pm), whereas this share declines to 21.7% under variable pricing. Thus, the observed variable pricing redistributes the demand away from congested time slots through lower off-peak prices, accompanied by a 8.3% net increase in reservation volume.

We find this shift to off-peak times reduces average restaurant profit per reservation from €4.79 to €4.06, reflecting the lower price point booked. This decline is consistent with the restaurant's incentives to preserve capacity for walk-in customers, albeit at the cost of sales substitution to lower-priced times on the platform. In particular, we find that for the majority restaurants, their *total* profit from platform customers decreases even after accounting for the increased reservation volume. This substitution of peak-period sales to off-platform markets can affect the platform's incentive to allow variable pricing under alternative platform compensation schemes, a question we turn to in the next section. Overall, the surplus a reservation generates on the platform (not accounting for displaced walk-ins) decreases slightly: from €7.70 under uniform pricing to €6.97 under variable pricing.

Panel C reports the welfare impact of platform reservations on walk-in customers. Under uniform pricing, each reservation displaces, on average, 0.21 walk-in customers, resulting in a lost profit of €2.35 per reservation. This displacement accounts for approximately 30.5% of the surplus a reservation generates on the platform, indicating meaningful inefficiency due to crowding out of walk-in demand. In contrast, under variable pricing, the average number of displaced walk-in customers falls to 0.15 per reservation, and the corresponding profit loss drops to €1.63, or 23.4% of the welfare generated on the platform. As a result, time-varying pricing reduces the inefficiency from demand displacement by approximately 28% in terms of lost walk-in volume and by 30% in terms of displaced profit, per reservation. Because these efficiency gains per reservation are quite large, we find variable pricing reduces the total volume of displacement at the day level by 10.0%, even after accounting for the larger reservation volume. Once again, because our measure of welfare loss does not include the forgone surplus of walk-in customers, the welfare gains we report from variable pricing represent a conservative estimate of the actual magnitude in this market.

Our counterfactual analyses highlight the crucial role of peak-load pricing in improving efficiency by reducing the extent to which walk-in customers are crowded out by platform bookings, while simultaneously increasing the platform users' surplus through lower prices. This reallocation of demand also increases profit of restaurants and the platform. Collectively, the results offer new empirical evidence that profit-maximizing behavior by firms can contribute to welfare improvements in markets characterized by capacity constraints.

7.3 The Role of Platform Compensation and Supply-Side Incentive Misalignment

We next examine the role of platform compensation scheme in determining the gains from variable pricing. In practice, our focal platform charges restaurants a flat fee per reservation, which corresponds to an increased marginal cost for the restaurants to serve platform customers. The conventional wisdom posits that this form of compensation structure typically distorts the restaurant prices upward — because restaurants are "taxed" on the quantity sold but not on the price, a mechanism akin to the double-marginalization problem — and thus produces an inefficient outcome. Forming a profit-sharing scheme between the platform and each restaurant typically eliminates this inefficiency, because fees proportional to the restaurant profit do not distort the restaurant's optimal prices.[27] In this section, we show that a profit-sharing contract also substantially enhances the efficiency gains from adopting peak-load pricing, thus adding to the conventional wisdom on the relative efficiency of profit-sharing contracts over per-unit-fee contracts. However, we also document that under a profit-sharing contract, adoption of variable pricing can *reduce* the platform's profit, creating an incentive misalignment problem in which the platform may not allow the full welfare benefits to materialize.

To formally characterize the restaurant pricing problem under this (counterfactual) profit-sharing scheme, we define the restaurant's profit as follows:

$$\pi_{rt} = \sum_{j \in J_{rt}} \left((1-\rho) \times q_{\tau_j t}(p_{jt} - c_r) + \mathbb{E}[\min\{n_{\tau_j t}, L_r - q_{\tau_j t}\}](p_r^{regular} - c_r) \right),$$

where ρ corresponds to the fraction of profits from platform customers that is shared to the platform.[28] Correspondingly, the platform's profit from the restaurant is given as follows:

$$\pi_{platform,rt} = \rho \sum_{j \in J_{rt}} q_{\tau_j t}(p_{jt} - c_r).$$

[27] In the context of vertical contracts such a profit-sharing contract corresponds to a two-part tariff, in which the linear price is zero (which is the platform's marginal cost) and the fixed portion corresponds to a fraction of the restaurant profit. Because the marginal price is zero, this compensation scheme does not affect the restaurant's pricing at the margin.

[28] Because only the profit on the platform is shared and not the walk-in profit, this profit sharing scheme is not fully efficient, though as we show below it is still substantially more efficient than the observed per-unit-fee contract.

In what follows, we fix ρ at the value such that the platform's profit under the baseline case of uniform pricing is identical to the factual linear contract. In Table 5, we present the market outcomes under this profit-sharing contract. Column 1 and 2 correspond to the outcomes from variable- and uniform-pricing regimes under the profit-sharing contract. Column 3 and 4 correspond to the factual per-unit fee structure and are identical to the two columns of Table 4, presented for comparison. The metrics shown in the three rows are identical to those shown in panel (A) of Table 4. Comparing columns 2 and 4, we first confirm the conventional wisdom that, in the absence of variable pricing, a profit-sharing contract outperforms a per-unit-fee contract. Holding the platform profit constant, the profit-sharing contract generates welfare of €209.2 per day, as compared to €168.1 under the per-unit fee, or an increase of 24%. The restaurants' collective profits also increase by approximately 13.6% from €64.7 to €73.5 per day. Thus, we find that a profit-sharing contract can lead to a substantial efficiency gain by eliminating the distortion in the restaurants' pricing incentives.

	Profit Share		Per-Unit Fee (Factual)	
	Variable	Uniform	Variable	Uniform
	(1)	(2)	(3)	(4)
Welfare Generated by the Platform	233.2	209.2	182.5	168.1
Restaurants' Profit	82.73	73.50	69.85	64.66
Platform Profit	35.35	34.28	37.60	34.28

Note: The metrics shown in the three rows are identical to the first three rows of Table 4. All the metrics are per day, aggregated across restaurants.

Table 5: The Role of Platform Compensation

We next document the role of profit-sharing contract in improving the efficacy of peak-load pricing. Comparing columns 1 and 2, we find the overall welfare increases by 11.5% from €209.2 to €233.2, accompanied with a 12.5% increase in restaurant profits from €73.5 to €82.7. These increases amount to 134% and 156% of the corresponding increases under the per-unit-fee structure. Thus, we find a profit-sharing scheme not only improves the baseline efficiency, but also complements the use of peak-load pricing to substantially enhance its efficiency gains. The reason behind this benefit is again the elimination of double-marginalization: For a peak-load pricing to be most effective, the off-peak prices need to be sufficiently low to attract intertemporal substitution. A profit-sharing contract, by lowering the restaurants' effective marginal costs, allows restaurants to offer a deeper discount than a per-unit-fee contract.

We next show, however, that the platform's gain from variable pricing *decreases* under the profit-sharing regime, as compared to the case of per-unit fee. The platform's gain from variable pricing is limited to 3.1% (from €34.3 to €35.4) under profit-sharing scheme, which is only 32% of the gains under the per-unit-fee scheme (9.6%). The smaller gains arises because, unlike the case of per-unit fees, the platform's profit reflects the prices that users book, which declines substantially due to peak-load pricing. Figure 11 displays the distribution of the changes in platform profit from each restaurant due to variable pricing under the profit-sharing scheme. We find that the platform's profit decreases among 42 restaurants, or 53% of the sample. This result indicates that the platform potentially lacks incentive to provide variable pricing to a large fraction of restaurants despite the welfare benefits.[29] The welfare impact can be large: if the platform does not allow these 42 restaurants to exercise variable pricing, the total welfare gain from variable pricing is limited at 7.7%, which is comparable to the increase under the per-unit-fee contract. These results thus indicate the potential loss due to incentive misalignment can largely offset the incremental gains from eliminating double-marginalization in the context of peak-load pricing.

The reason behind the substantial heterogeneity in the platforms' gains across restaurants (ranging from -10% to 300+%) is because restaurants face heterogeneous capacity constraints. Consider a restaurant that faces a strict peak-time capacity constraint and the platform enforces a uniform-discount policy. In this case, the restaurant is forced to accommodate the price-sensitive platform users during the constrained peak time. In this case, the restaurant sets a uniform discount close to zero, effectively shutting down the platform sales altogether. In this case, the platform is better off, often substantially in percentage terms (because the denominator is close to zero), by offering variable pricing, thereby creating incremental sales during the off-peak periods without cannibalizing the restaurants' peak-time sales to regular diners. On the other hand, the platform loses money from offering variable pricing to restaurants with modest capacity constraints: because these restaurants can serve platform users even under the uniform-pricing regime, the aforementioned sales diversion becomes the first order effect, reducing the platform profit. In the data, restaurants from which the platform gains have an average peak-time sellout probability of 39.5%, whereas the corresponding number for those from which the platform loses is

[29] For comparison, this incentive misalignment is not a major issue under the observed per-unit-fee structure, because the platform's profit only depends on the quantity sold, which often increases with peak-load pricing. In the factual environment, the platform's profit decreases from only one restaurant.

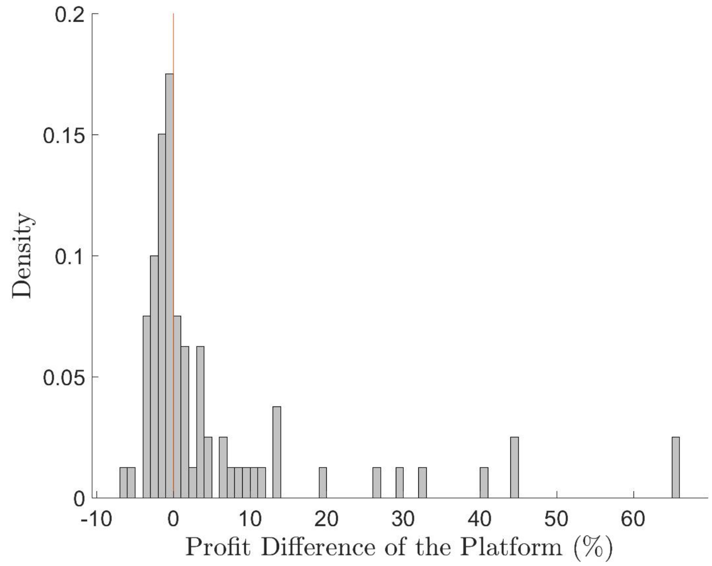

Figure 11: Changes in Platform's Profit

Note: Each observation corresponds to a restaurant, averaged across dates. The rightmost bar corresponds to the increase in platform profit of 65% and above. The maximum value is 360%.

15.5%.

Overall, these results uncover the crucial role of platform compensation in materializing the welfare benefits of peak-load pricing. On the one hand, a profit-sharing contract eliminates the upward pricing distortion and thus allows restaurants to offer deep off-peak discounts, substantially improving the efficacy of peak-load pricing. On the other hand, peak-load pricing under profit-sharing contracts often hurts the platform through lower average transaction prices, reducing the platform's incentive to offer such a pricing strategy despite the welfare benefits. Our findings that the profit-sharing contract, a contract that would create the largest efficiency gains, does not incentivize the platform to allow variable pricing among 53% of the restaurants, indicate that a substantial portion of the relative efficiency of profit-sharing contracts over per-unit-fee contracts may be lost in markets with capacity constraints.

8 Conclusion

We study the role of peak-load pricing in improving the welfare in a platform market with profit-maximizing firms. We find that price variability improves the welfare by 8.6%, a substantial welfare gain over the case in which firms must set uniform prices, even when the firm's motivation is profit maximization. The results contribute to policy debates on regulating firm pricing power, because markets in which the firm exercises price discrimination are often ones with firm capacity constraints.

We also present a novel piece of evidence of incentive misalignment in platform markets that may prevent the implementation of the welfare-improving pricing policy as a market outcome. The relevance of our findings extends well beyond restaurants. Many modern services—such as ridesharing platforms, fitness memberships, and coworking spaces—operate under similar hybrid models with fixed short-run supply, fluctuating demand, and a mix of walk-in and reserved access. In these settings, peak-load pricing offers clear benefits: it can reduce congestion during busy periods, smooth demand across time, and improve access for consumers who are more flexible—without requiring firms to expand costly capacity. However, realizing these benefits often depends on coordinated price-setting, which no single actor is incentivized to implement when the gains are distributed unevenly. Our results suggest that even modest frictions in how prices are set—especially when individual firms do not capture the full value of market-level improvements—can lead to significant inefficiencies and leave large welfare gains unrealized. To fully harness the benefits of dynamic pricing, platforms must be designed not only to optimize market-level outcomes but also to align incentives across participating firms. Without mechanisms that reward firms for contributing to efficiency, platforms risk falling short of their potential to improve how modern markets operate.

References

APARICIO, D. AND K. MISRA (2023): "Artificial intelligence and pricing," Artificial intelligence in marketing, 103–124.

BORENSTEIN, S. (2012): "The redistributional impact of nonlinear electricity pricing," American Economic Journal: Economic Policy, 4, 56–90.

BRECKO, K. AND W. R. HARTMANN (2023): "Marketing and Economics of Harm Reduction: Evidence from Conservation Field Experiments," .

CARDELL, S. N. (1997): "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," Econometric Theory, 13, 185–213.

CASTILLO, J. C. AND S. MATHUR (2023): "Matching and Network Effects in Ride-Hailing," AEA Papers and Proceedings, 113, 244–247.

CHO, S., G. LEE, J. RUST, AND M. YU (2018): "Optimal dynamic hotel pricing," .

HENDEL, I. AND A. NEVO (2013): "Intertemporal price discrimination in storable goods markets," American Economic Review, 103, 2722–2751.

JOSKOW, P. L. (2012): "Creating a smarter US electricity grid," Journal of Economic Perspectives, 26, 29–48.

MACKAY, A., D. SVARTBÄCK, AND A. G. EKHOLM (2023): "Dynamic Pricing, Intertemporal Spillovers, and Efficiency," Tech. rep., Harvard Business School Strategy Unit Working Paper.

SANDERS, R. E. (2024): "Dynamic pricing and organic waste bans: A study of grocery retailers' incentives to reduce food waste," Marketing Science, 43, 289–316.

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

TURVEY, R. (1968): "Peak-load pricing," Journal of Political Economy, 76, 101–113.

WALDFOGEL, J. (2003): "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," The RAND Journal of Economics, 34, 557–568.

WILLIAMS, K. R. (2022): "The welfare effects of dynamic pricing: Evidence from airline markets," Econometrica, 90, 831–858.

WILLIAMSON, O. E. (1966): "Peak-load pricing and optimal capacity under indivisibility constraints," The American Economic Review, 56, 810–827.

NOTE: This appendix contains older figures and estimates and may not be 100% consistent with the current main text. The edits are in progress.

A The Traffic Measure ($g_{\tau_j t}$)

In this section, we discuss details of the measure of restaurant traffic, $g_{\tau_j t}$, collected from Google Map. Google map shows a histogram of traffic volume for each restaurant at each day of week. An example screenshot can be seen in Figure 12. The traffic measure represents the total volume of foot traffic at each hour of day, constructed from Google devices' location data.[30] We scraped the numbers underlying the histogram in late 2022. Our data hence include a snapshot of the traffic measure at each restaurant at each day of week - time of day pair, scaled between 0 and 100. 0 means the restaurant is either closed or no guest is dining, and 100 means it is the busiest time during the week. Because our data of reservation is available for each 30-minute increment, we predict the traffic at each 30-minute increment with a piecewise interpolation. We use the imputed measure as our measure of traffic, $g_{\tau_j t}$. Note the traffic data is not concurrent with the reservation data. We assume that the traffic pattern remains the same between the two data periods.

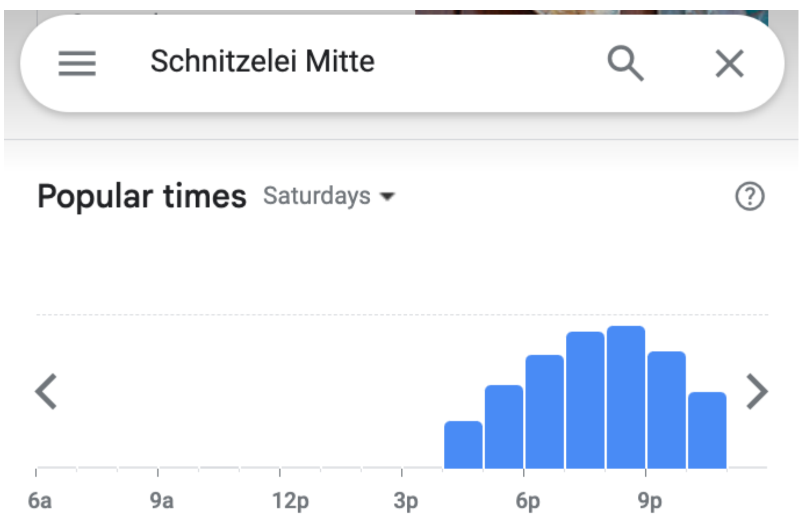

Figure 12: Histogram of Traffic on Google Maps

In various parts of the estimation procedure, we also use the market-level aggregate

[30]The exact formula Google uses to aggregate the geolocation data is unknown. It is reasonable to assume that it considers the weighted average of past several days of realizations.

measure of traffic volume, which we denote by $g_{\tau t}$. For instance, we use the empirical distribution of $g_{\tau t}$ as that of each consumer's most preferred dining time, $F_t(\tau_{it})$.[31] We also use the ratio $g_{\tau_j t}/g_{\tau t}$ (i.e., the traffic volume of a given restaurant relative to the aggregate traffic volume) as the measure of restaurant popularity in the vector of restaurant characteristics.

To aggregate $g_{\tau_j t}$ up to the market level (and recall that $g_{\tau_j t}$ is scaled restaurant by restaurant, so that we cannot simply take its average across restaurants). We take its weighted average, in which the weights are given by the maximum number of tables each restaurant ever made available on the platform during the sample period, which corresponds to a rough measure of restaurant capacity. By weighting $g_{\tau_j t}$ by a measure of restaurant capacity, we can account for the fact that the variation in $g_{\tau_j t}$ of a large restaurant is more informative for the aggregate distribution of dining time than that of a small restaurant.

B Sample Selection Criteria

The original data consists of the universe of 76,486 reservations on the platform between October 2018, the inception of the platform, and November 2020. We drop the following observations:

- Reservations between Oct 2018 and Aug 2019 (23,912 obs.) because the platform was growing rapidly with unstable demand and price patterns
- Reservations after Mar 2020 (26,420 obs.) because of Covid shocks
- Missing variables, data error (760 obs.)
- Canceled reservations (4,940 obs.)
- Reservations outside of the 3 pm-9:30 pm window (5,410 obs.)
- Restaurants with no walk-in traffic measure available (3,035 obs.)

These criteria leave approximately 12,760 reservations by 6,140 consumers across 101 restaurants. However, majority of these consumers (60%) only make one reservation throughout the sample period, and some restaurants receive very few reservations. Because we

[31] In practice, the realized dining time likely underestimates the consumer preference for peak times due to firm capacity constraints. We argue that the measure is nevertheless the most useful proxy for the ideal dining time due to its wide availability across all restaurants, and its granularity at hour - day level variation.

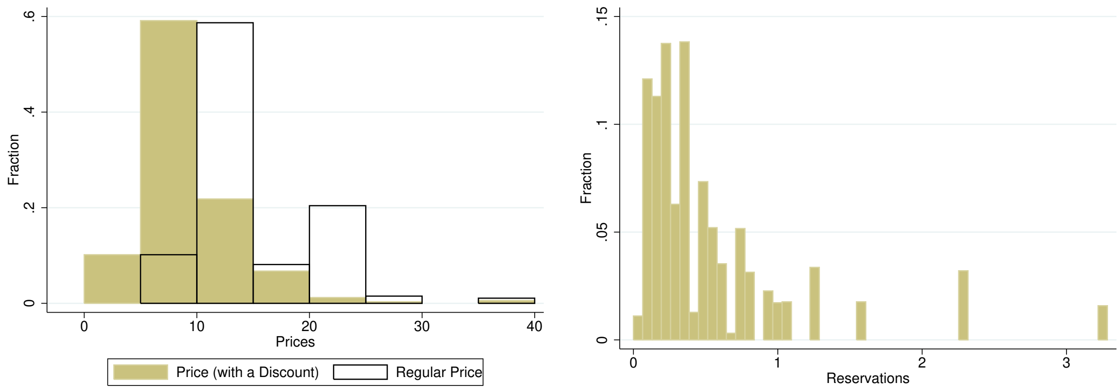

(a) Prices

(b) Reservations

Figure 13: Histograms of Prices and Reservations

Note: In panel (a), the distribution of discounted prices is over date - time - restaurant tuple, whereas that of regular price is over the restaurants. Panel (b) is the histogram of average daily reservations of restaurants.

need to exclude each consumer's first reservation from the likelihood (because their platform participation is unobservable prior to it), we only keep consumers with at least three reservations. We also keep restaurants that receive at least 10 reservations. At the end of these steps, our final dataset is comprised of 6,035 reservations made by 1,326 customers across 80 restaurants during 189 days between September 1st, 2019 and March 8th, 2020.

C Further Details on Data

C.1 Histograms of Prices and Quantities

In panel (a) of Figure 13, we present the histogram of discounted prices (colored bars) and the regular price of participating restaurants (uncolored bars). The regular price corresponds to the prices that an walk-in guest will pay without the online discount. In panel (b) of Figure 1, we present the histogram of average daily reservations across restaurants.

C.2 Variables Used as Controls for Demand Shifters

Besides the main variables we discussed in Section 2, the data include a set of restaurant characteristics we use as controls for demand shifters. In Table 6, we present the summary statistics for non-categorical variables. Regular price corresponds to the average price a walk-in customer pays, a measure collected by the platform. The average price is 12.4 Euro per person, with the standard deviation of 5.7. Review ratings correspond to ratings on Google. Most restaurants are rated around 4.3 with relatively small variation around it, partly because the platform exclusively recruits restaurants with high ratings. Maximum Table Offered corresponds to the maximum number of tables the restaurant ever offered on the platform during the sample period, a measure of restaurant size we use for demand estimation. On average, restaurants offer 18 tables with the standard deviation of 8.6.

Categorical variables we use for estimation are the restaurant location at the district level and restaurant cuisine type. We observe 7 districts in Berlin (Charlottenburg, Friedrichshain, Kreuzberg, Mitte, Neukölln, Prenzlauer Berg and Schöneberg) and 5 cuisine types (American, Asian, European, Middle Eastern and others).

Table 6: Summary statistics

	Mean	Std. Dev.	Min	Max
Regular Price	12.438	5.681	5	40
Review Rating	4.343	0.243	3.6	4.8
Maximum tables offered	17.975	8.620	4	50

Note: Each unit of observation is a restaurant.

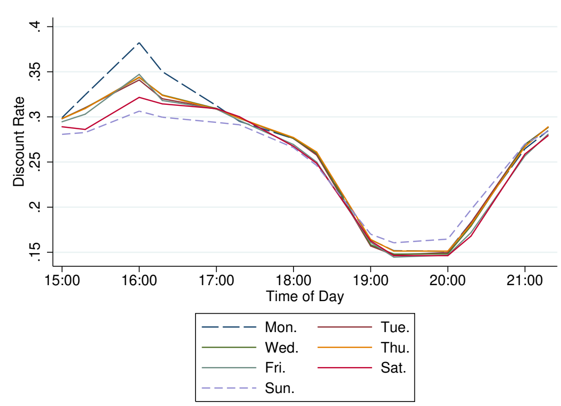

Figure 14: Price Schedule across Days

Note: The lines correspond to the average across restaurants and dates (within day of week).

D Price Schedule across Days

In Figure 14, we present the discount schedule at each day of week. The long-dashed line corresponds to Monday with the largest intra-day variability, and the short-dashed line corresponds to Sunday with the smallest variability. We find virtually identical price schedules across days. The largest price gap (and hence the ones that most likely trigger consumers substitution across days) is the one between Monday and Sunday at 4pm, with the average gap of 7.5 percentage points in discount rate. Because the gap is less than half of the average intra-day shift in discount rates, we argue that the primary objective of variable pricing is to move consumers within a day, rather than across dates.

E Average Observed Consumer Characteristics by Peak and Off-peak Reservations

Figure 15 replicates Figure 2 without excluding the focal reservation in computing the customer-level average. The distribution look largely the same and our findings in Figure 2 are robust to this modification.

F Robustness of the Identification Strategy of the Price Coefficient

Our main identification strategy for the price coefficient is that within a given day of week - time of day pair, price variations are independent of the demand shock of the particular day on which the price adjustment took place, because the price adjustment is often made in response to a more gradual, longer-term shift in dining patterns, providing a fuzzy regression-discontinuity design. We check robustness of this strategy regarding two possible threats to identification, using our reduced-form regressions (Expression (1)) in Section 3.2.

We first focus more exclusively on local variations around the date of price change. Although Expression (1) partly exploits the aforementioned variations by exclusively focusing on variations across calendar days within restaurant - day of week - time of day tuple, it is not equivalent to the ideal regression-discontinuity design: for instance, it does not exclude the variation in prices between two calendar days that are far apart, as long as

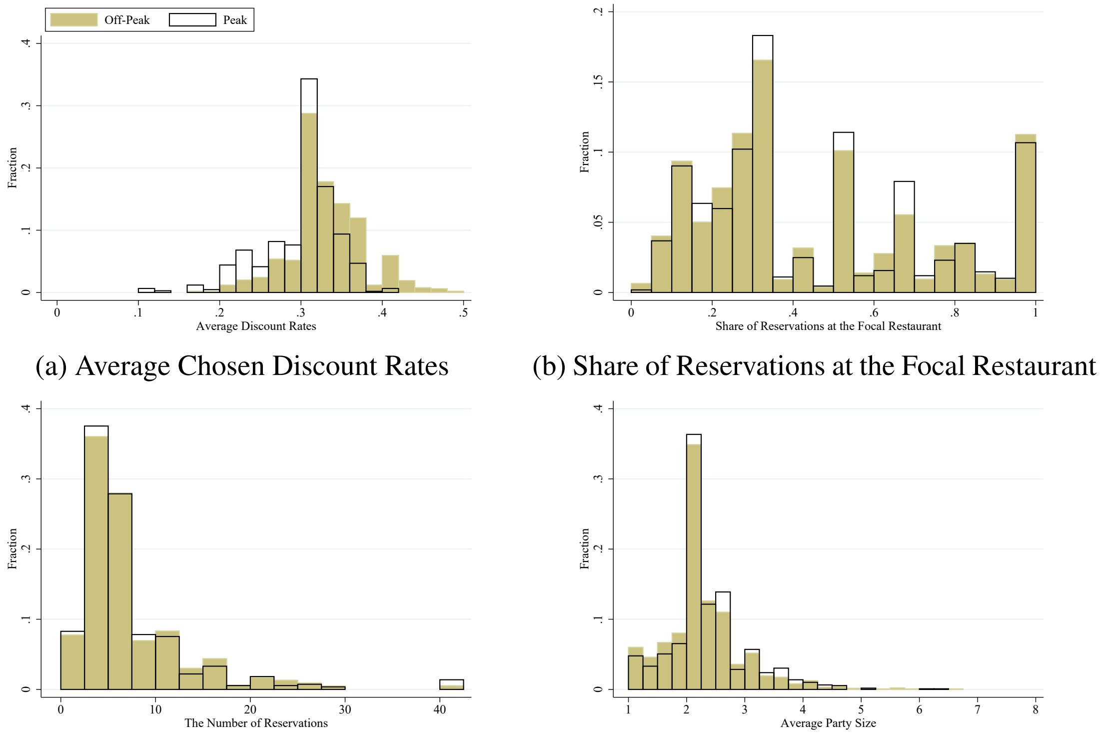

(a) Average Chosen Discount Rates (b) Share of Reservations at the Focal Restaurant

(c) The Number of Reservations (d) Average Party Size

Figure 15: Distribution of Average Characteristics Between Consumers Choosing Peak and Off-peak Times

Note: The construction of these figures are identical to Figure 2, except that we do not exclude the focal reservation.

they belong to the same tuple. Hence, any ”gradual, longer-term shift in demand”, if it is indeed the source of the price change, still confounds the estimate of price coefficient. To address this concern, we now exclusively use observations within two weeks of the date of price change for estimation. For instance, if restaurant r implemented a price adjustment on Wed, Nov 13, 2019, for future spots on Saturday at 4pm, then we keep observations from Saturdays, Nov 2, 9, 16 and 23 at 4pm in the sample. Although these sample selection criteria are extremely limiting and do not allow us to estimate cross-elasticity terms due to lack of sample size, they most closely follow our ideal variation for identifying the own elasticity. We present the results in Table 7. The first column is our main result and is identical to column 3 of Table 2. The second column presents the estimate from the more restrictive specification. We find the own elasticity of -11.7, which is approximately 12% larger than that of the main specification. The difference is likely due to the sample selection that the focal data include a larger fraction of observations around 6:30 pm and 7 pm (9.51% and 10.09% of the sample, respectively) than the data used in column 1 (7.43% and 7.41%, respectively), which are the time slot in which we predict a higher price elasticity (see Figure **??** in Section 3.3). Besides the sample selection, we argue that these two specifications produce similar elasticity estimates.

The other threat to identification is the possibility that restaurants do set prices at calendar-day level, so that price difference between two calendar days (within a restaurant - day of week - time of day tuple) still reflects calendar-day specific demand shifts. To address this issue, we consider a specification in which we add two sets of fixed effects to Expression 1. The first set of fixed effects is unique to each date - time of day pair. This fixed effect accounts for any date - time of day specific shift in demand that affects prices of all restaurants (e.g., national holidays, city-wide events). The second set of fixed effects is at restaurant - calendar day - hour of day level. The observations are restaurant - calendar day - 30 minute level, so this extra layer of fixed effects allows us to exclusively focus on comparing between observations from the top of the hour against ones from the bottom of the hour within the same restaurant - date - hour tuple. The intuition behind exclusively focusing on price variations within each 30-minute interval is that, any demand shift that is large enough to create price variability at the calendar-day level should likely persist across multiple time slots (e.g., local events likely create larger foot traffic throughout the entire evening), and the timing of price change within a given day (e.g., when the restaurant wants to increase its discount rate after the peak time, whether it does so at 8:00pm or 8:30pm) should be mostly exogenous to the day-to-day fluctuation in demand. This motivation is akin to another dimension of fuzzy regression discontinuity, in which we utilize discontinuous jump in discount rate within a given day, whereas the intra-day demand variability does not necessarily spike.

We present the results in column 3 of Table 2. We find the own elasticity of 11.56, or approximately 11% higher than our main results. The difference is not statistically significant, partly due to the spike in the standard error that arises from the substantial drop in degree of freedom. The intra-day cross-elasticity term also remains within the same ballpark of the main specification. Given that all three specifications produce similar price coefficients despite the substantial differences in the source of identification and in their respective degree of freedom, we argue that the true price coefficient in this market is also located within the same ballpark of these estimates, which we also recover in the structural analysis (between 10 to 10.7: see panel (a) of Figure 4).

Table 7: Various Sources of Identification

	$\log(R_{j,t,\tau})$		
$\log(P_{j,t,\tau})$	-10.453	-11.694	-11.559
	(1.039)	(0.949)	(2.292)
$\log(\bar{P}_{j,t\pm1,\tau})$	2.092	-	2.756
	(0.869)		(2.035)
$\log(\bar{P}_{j,t,\tau\pm1})$	0.021	-	0.570
	(0.998)		(2.508)
j-dow(t)-τ FE	Yes	Yes	Yes
t-τ FE	No	No	Yes
j-t-hour(τ) FE	No	No	Yes
$\pm$ 2 wks of price change	No	Yes	No
Observations	111,977	2,755	92,480
Degree of Freedom	106,977	2,088	43,241

Note: Standard errors are reported in parentheses. j, t and τ correspond to restaurant, date and time, respectively.

We next account for substitution to non-adjacent time periods. Specifically, we consider a specification in which we include the average price of time slots that are one hour (two time slots) away from the focal slot within the same day, as well as average price of two adjacent time slots of the two adjacent days.[32] The result is presented in Table 8. Note that the price terms from nearby slots are extremely highly correlated, causing substantial noise in the estimates. Nevertheless, we find the own elasticity virtually remains unchanged from the main specification, and ”within the same day, adjacent time slot” cross elasticity is still positive and statistically significant. The other substitution terms, namely, within-day substitution beyond 30 minutes and across-date substitution, do not seem to play a meaningful role in the focal market.

[32] For instance, if the focal time slot is Saturday at 6pm, the former corresponds to the average price of Saturday at 5pm and 7pm, and latter corresponds to the average of the prices on Friday and Sunday at 5:30pm and 6:30pm.

Table 8: Substitution to Non-adjacent Slots

	$\log(R_{j,t,\tau})$
$\log(P_{j,t,\tau})$	-10.209
	(1.507)
$\log(\bar{P}_{j,t\pm1,\tau})$	3.857
	(1.919)
$\log(\bar{P}_{j,t\pm2,\tau})$	0.096
	(1.036)
$\log(\bar{P}_{j,t,\tau\pm1})$	-1.230
	(1.674)
$\log(\bar{P}_{j,t\pm1,\tau\pm1})$	-1.045
	(2.005)
j-dow(t)-τ FE	Yes
Observations	91,385

Note: Standard errors are reported in parentheses. j, t and τ correspond to restaurant, date and time, respectively.

G Accounting for Multiple Guests per Reservation

For estimation, we use per-guest price as our measure of p_{jt} and $p_r^{regular}$, essentially assuming each reservation and walk-in demand consists of one guest. In this section, we show this assumption is without loss of generality. Suppose first N guests arrive per party (common across platform and walk-in guests), then the objective function is given as follows:

$$\begin{aligned}\pi_{rt} &= \sum_{j\in J_{rt}} \big(Nm_{\tau_j t}(X_t,p_t,\tau_t)(p_{jt}-c) + \mathbb{E}(\min\{Nn_{\tau_j t}, L - Nm_{\tau_j t}(X_t,p_t,\tau_t)\})(p_{full}-c)\big) \\ &= N\sum_{j\in J_{rt}} \left(m_{\tau_j t}(X_t,p_t,\tau_t)(p_{jt}-c) + \mathbb{E}\left(\min\left\{n_{\tau_j t}, \frac{L}{N} - m_{\tau_j t}(X_t,p_t,\tau_t)\right\}\right)(p_{full}-c)\right).\end{aligned}$$

Hence, allowing multiple guests per reservation simply changes the scale of L, leaving everything else unchanged. Because we estimate L based off of the first-order conditions, the value of the estimated L should be interpreted as the restaurant capacity in terms of the number of parties it can accommodate (i.e., if the restaurant has 50 seats with the average party size of 2, then our estimate of L equals 25).

We next consider the possibility that the average number of guests per party differs between platform customers and walk-in guests. Denoting the average number of guests per party for platform and walk-in by N_1 and N_2, respectively, the objective function in this case is given as follows.

$$\begin{aligned}\pi_{rt} &= \sum_{j\in J_{rt}} \left(N_1 m_{\tau_j t}(X_t, p_t, \tau_t)(p_{jt} - c) + \mathbb{E}(\min\{N_2 n_{\tau_j t}, L - N_1 m_{\tau_j t}(X_t, p_t, \tau_t)\})(p_{full} - c)\right) \\ &= N_1 \sum_{j\in J_{rt}} \left(m_{\tau_j t}(X_t, p_t, \tau_t)(p_{jt} - c) + \mathbb{E}\left(\min\left\{\frac{N_2}{N_1} n_{\tau_j t}, \frac{L}{N_2} - m_{\tau_j t}(X_t, p_t, \tau_t)\right\}\right)(p_{full} - c)\right).\end{aligned}$$

This time, the $n_{\tau_j t}$ we estimate represents $\frac{N_2}{N_1} n_{\tau_j t}$ - it is the volume of walk-in demand *measured in terms of the demand from the platform*. For example, if the walk-in demand on average consists of 4 people whereas that from the platform consists of 2 people, $n_{\tau_j t}$ is estimated doubled compared to the case of same N, in order to reflect the fact that two reservations from the platform will displace one walk-in party. Again, besides the difference in the interpretation of the estimated L and $n_{\tau_j t}$, everything else in the model remains unchanged.

H The Use of Popularity Measure to Account for Price Endogeneity

We show in Appendix F that our identification strategy does produce a robust estimate for the reduced-form price coefficient. However, for estimation of the structural model, the same set of fixed effects is not usable due to model complexity and we instead opt to use a variety of continuous variables to mimic the same identification strategy. In particular, to account for the restaurant - time specific demand shifters, we use the restaurant popularity measure, which is defined as the ratio between the restaurant-specific traffic volume and the aggregate traffic volume, in place of restaurant - date - time fixed effect. Whether or not a single-dimensional continuous variable can adequately replace the slew of fixed effects is an obvious concern. In this section, we show robustness in our structural estimates with respect to alternative specifications in the restaurant popularity measure.

Table 9 shows the estimates of price coefficients and costs for switching time slots under the specification in which we include the popularity term as a polynomial in the utility, rather than as a linear function in our main specification. We find stable parameter estimates across specifications. The vast majority of the estimates are within $\pm$ 10% range of the ones from the linear specification presented in Table 3, and the difference is not statistically significant. Although several coefficients, such as the price coefficients for those that prefer early time slots, shift more than 10% from our main results, we find that these shifts cause little to no impact on our estimate of price elasticity. In Figure 16, we show the within-day evolution in own- (left figure) and cross-elasticity (right figure) when we include the third-degree polynomial of the popularity term. Compared to the results in Figure 4, we find that the own elasticity is within 5% of our main specification, and the cross elasticity remains virtually identical. These results show that the linear specification we employ in the main text is sufficient to capture the within-restaurant, across-time variation in demand.

	Popularity in Quadratic				Popularity in Cube			
	Segment 1		Segment 2		Segment 1		Segment 2	
Price (Early)	-0.573	(0.025)	-0.582	(0.074)	-0.613	(0.026)	-0.513	(0.051)
Price (Peak)	-0.539	(0.023)	-0.557	(0.076)	-0.563	(0.024)	-0.546	(0.055)
Price (Late)	-0.522	(0.021)	-0.538	(0.076)	-0.545	(0.022)	-0.519	(0.049)
Switching cost (Early)	-0.502	(0.063)	-0.538	(0.165)	-0.492	(0.069)	-0.743	(0.201)
Switching cost (Peak)	-0.535	(0.073)	-0.175	(0.070)	-0.577	(0.074)	-0.173	(0.060)
Switching cost (Late)	-0.393	(0.057)	-0.379	(0.099)	-0.420	(0.060)	-0.359	(0.102)

Note: Standard errors are reported in parentheses. ”Early”, ”Peak” and ”Late” stand for the most preferred dining time between 3pm and 4:30pm, 5pm and 6:30pm and 7pm and 9:30pm, respectively. The specification also includes fixed effects for year-month pairs, for weekend days, for restaurant locations (at district level), cuisines and for quartile bins of restaurant density in the neighborhood. The coefficients of these fixed effects are segment-specific. In addition, we include interactions between restaurant locations and cuisines whose coefficients are common across segments.

Table 9: Demand Parameter Estimates

I Parameterization of $\lambda_{\tau_j t}$

Recall that in our supply-side model, restaurants face a random arrival of available capacity at each time slot, $x_{\tau_j t}$, which follows a Poisson distribution with parameter $\lambda_{\tau_j t}$ which is

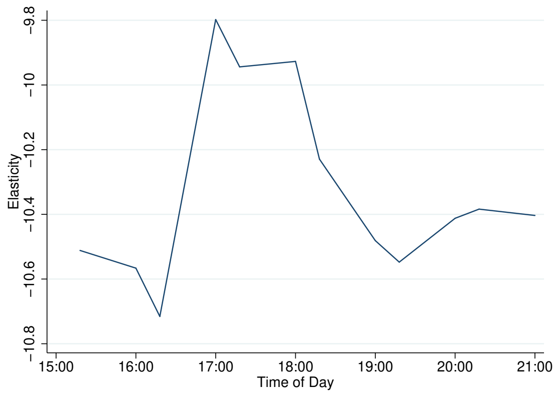

(a) Own Elasticity

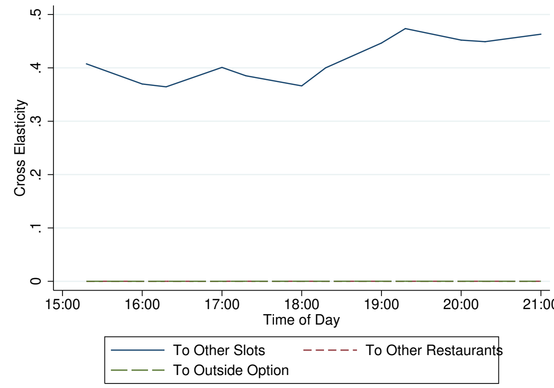

(b) Cross Elasticity

Figure 16: Across-time Elasticity Variations within a Day

Note: The number of reservations corresponds to the average number of reservations for each time slot, aggregated across all the restaurants.

parameterized as follows.

$$\lambda_{\tau_j t} = L_r - \underbrace{(\gamma_{1r} g_{\tau_j t} + \gamma_{2r} g_{\tau_j t}^2)}_{\bar{n}_{\tau_j t}} - \underbrace{\sum_{\tau'=1}^{2} (\gamma_{3r\tau'} g_{\tau_j+\tau',t} + \gamma_{4r\tau'} g_{\tau_j+\tau',t}^2)}_{\bar{n}_{\tau_j+\tau',t}},$$

where L_r, $\bar{n}_{\tau_j t}$ and $\bar{n}_{\tau_j+\tau',t}$ correspond to total available capacity, the expected walk-in demand for the current and future periods, respectively. Because we estimate these parameters restaurant by restaurant through the inversion of the respective first-order conditions, we can allow different parameterizations across restaurants to balance the model fit and statistical power at the restaurant level. As a baseline, we use the following specification:

$$\lambda_{\tau_j t} = L_r - \max\{0, (\gamma_{1r} g_{\tau_j t} + \gamma_{2r} g_{\tau_j t}^2)\} - \rho_{r\tau} \sum_{\tau'=1}^{2} \max\{0, (\gamma_{3r\tau'} g_{\tau_j+\tau',t} + \gamma_{4r\tau'} g_{\tau_j+\tau',t}^2)\}. \tag{5}$$

Besides the total capacity term L_r, we allow for the linear and quadratic term for the current traffic value that are constant across time (γ_{1r} and γ_{2r}), the linear and quadratic terms for the traffic one-period and two-period ahead (γ_{3r1}, γ_{3r2}, γ_{4r1}, γ_{4r2}), and a "discount factor", $\rho_{r\tau} \in [0, 1]$ that can vary across each 90-minute interval (i.e., one $\rho_{r\tau}$ term for 3pm-4pm, 4:30pm-5:30pm, etc. with 6:00pm-7:00 pm term fixed at 1 for normalization, totaling 4 parameters to estimate per restaurant). Including L_r and c_r, we have 12 parameters for each restaurant, which is estimated by 98 price variations (14 time slots x 7 days of week).[33] The reason we include the flexible discount factor terms is because restaurant responses against future (expected) traffic are often asymmetric (e.g., restaurants increase pre-peak prices substantially when they expect future peak traffic, but do not decrease prices even when the post-peak drop in traffic is imminent), indicating the way they account for future traffic in pricing is different across time of day. We also assume that $\bar{n}_{\tau_j t}$ and $\bar{n}_{\tau_j+\tau',t}$ are both nonnegative, based on our microfoundation that these numbers correspond to current and future consumer arrivals.[34]

J Estimate of L_r

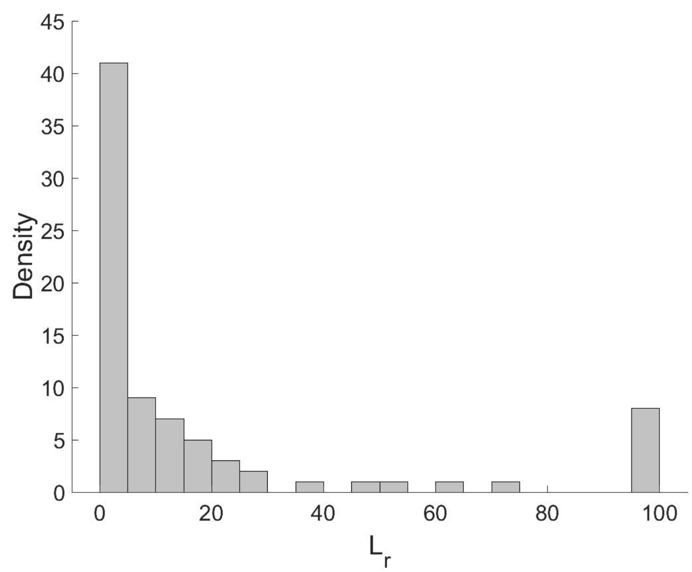

Figure 17: Restaurant Total Capacity

Note: The rightmost bin corresponds to restaurants with 100 tables and above.

Figure **??** shows the histogram of L_r, the intercept of restaurant capacity constraints.

[33]Recall that our traffic measure $g_{\tau_j t}$ only varies across days of week and time of day, and not across calendar day within each pair.

[34]For one restaurant in the data, the specification above does not produce sufficient statistical power because the restaurant does not offer all 98 price variations. We assume $\rho_{r\tau} = 1$ across all times for this restaurant.

K Definition of Welfare Metrics

In this section, we formally define the welfare metrics we used in Table 4. All the metrics are defined at date (t) level, and Table 4 shows their average across dates.

- Total Welfare Generated by the Platform:

  Total welfare generated by the platform is given as follows:

$$W = \underbrace{M * CS}_{\text{Users' Surplus}} + \underbrace{\rho \sum_{r,j} q_{\tau_j t} p_{jt}}_{\text{Platform Profit}} + \underbrace{\sum_{r} \pi_{rt} - \sum_{r} \sum_{j \in J_{rt}} \mathbb{E}[\min\{n_{\tau_j t}, L_r\}](p_r^{regular} - c_r)}_{\text{Restaurants' Profit from Participation}},$$

  where M is the number of platform users, CS is per-user (ex-ante) consumer surplus derived from the discrete choice framework, and π_{rt} is the restaurant profit (Expression (2) in Section 4). The first three terms correspond to the social welfare when the platform exists. The last term corresponds to the firm profit without the platform, which equals the total walk-in profit in the absence of customers with a reservation. The difference between the two equals the welfare generated by the platform. Each element with an underbrace corresponds to the metrics presented in the second through fourth rows of Table 4.

- The Number of Reservations:

  The number of reservations for each restaurant in each day is $\sum_{j \in J_{rt}} q_{\tau_j t}$. We sum across restaurants to calculate the daily total reservation on the platform.

- The Average Price Booked:

  The average price per reservation is $\sum_{j \in J_{rt}} q_{\tau_j t} p_{jt} / \sum_{j \in J_{rt}} q_{\tau_j t}$. We take its average to compute the average price on the platform.

- Restaurant Profit per Reservation (on the platform):

  Profit per reservation on the platform for each restaurant is $\sum_{j \in J_{rt}} (q_{\tau_j t}((1-\rho)p_{jt} - c_r) / \sum_{j \in J_{rt}} q_{\tau_j t}$. The numerator corresponds to the restaurant profit from platform users, excluding walk-ins. We present its average across restaurants in the table.

- Platform Surplus per Reservation:

  Platform surplus per reservation is given by the sum of consumer surplus per reservation, restaurant profits from platform users per reservation and the platform profit per reservation. The formula for the restaurant profits and the platform profit is given above.

Consumer surplus per reservation corresponds to the expected utility generated from the reservation, conditional on the reservation made (the utility from reservation exceeds all the other alternatives). More formally, for each customer with realized type β_i and τ_{it} it is given by $\mathbb{E}[u_{ijt}|u_{ijt} > u_{ij't}]$. The exact formula is available in Arcidiacono and Miller (2011), Lemma 3. To compute its market average, we take its expectation over all customer types to compute the average consumer surplus from each option j, and take their weighted average (across j) in which the weights are given by the frequency of reservation, $q_{\tau_j t}$.

- Displaced Walk-in Customers per Reservation:

  The number of displaced walk-in customers for each restaurant in each day is given by $\sum_{j \in J_{rt}} (\mathbb{E}[\min\{n_{\tau_j t}, L_r\}] - \mathbb{E}[\min\{n_{\tau_j t}, L_r - q_{\tau_j, t}\}])$. We divide it by the number of reservations to make it per reservation.

- Lost Profit due to Displacement per Reservation:

  Lost profit due to demand displacement for each restaurant in each day is given by $\sum_{j \in J_{rt}} (\mathbb{E}[\min\{n_{\tau_j t}, L_r\}] - \mathbb{E}[\min\{n_{\tau_j t}, L_r - q_{\tau_j, t}\})])(p_r^{regular} - c_r)$. We divide it by the number of reservations to make it per reservation.