Peak-Load Pricing in Platform Markets

Takeaki Sunada, Diego Aparicio, and Yufeng Huang*

Abstract

We study welfare and profit implications of variable pricing in a dining reservation platform where restaurants set time-varying discount rates (by day of the week and 30-minute time slots within the day). Using data on reservations and a measure of walk-in traffic, we estimate a model in which restaurants serve time-varying demand using fixed capacity. We find that peak-load pricing, rather than intertemporal price discrimination, is the primary driver of observed price variation. Variable pricing thus leads to substantial welfare gains through improving capacity utilization. Despite the benefits, variable pricing can reduce platform profits, weakening the platform's incentive to provide such technology. In particular, profit-sharing contracts between the platform and restaurants, although they eliminate double marginalization, can exacerbate this incentive misalignment, leading to inefficiently low provision of variable pricing.

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

*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). We are grateful for the insightful comments and suggestions by Juan Camilo Castillo, Sarah Moshary, and seminar and conference participants at Cornell-Rochester Marketing Workshop, IIOC 2025, SICS 2025, Marketing Science 2025, WEAI 2025. An alternative disclosure plan was submitted for this article. The raw data are not available; however, all the codes and disguised data are available in a repository to be determined.

1 Introduction

Prices that vary across time, location, or customer segments have become increasingly common in modern markets. Ride-hailing fares spike during busy hours, airline tickets cost more around holidays, and theme parks charge higher prices on peak days. These pricing practices have drawn growing regulatory scrutiny amid concerns that firms exploit individual data and market power to extract consumer surplus. In the United States, this scrutiny has taken multiple forms, including proposed legislation targeting "grossly excessive" price increases, investigations into so-called "surveillance pricing," and regulations aimed at discriminatory pricing practices.[1]

Despite this scrutiny, the welfare implications of time-varying prices remain theoretically ambiguous (Robinson, 1969, Varian, 1989). On the one hand, intertemporal price discrimination—a form of second-degree price discrimination in which firms distort the pricing schedule to induce consumer sorting across times—allows firms to extract additional surplus while potentially reducing welfare. On the other hand, when firms face capacity constraints, allowing prices to vary across time can improve allocative efficiency by shifting demand away from congested peak periods toward less crowded off-peak periods. This latter mechanism—referred to as peak-load pricing (Joskow, 2012)—can increase total surplus by relaxing binding capacity constraints. For example, Disney varies ticket prices by day of the week to steer price-sensitive visitors toward off-peak days, thereby easing congestion during peak periods.[2] Whether time-varying prices primarily serve as a tool for surplus extraction or instead improve efficiency through capacity reallocation is therefore an empirical question.

We study these welfare trade-offs and their unique implications to platform markets in a dining reservation platform where capacity-constrained restaurants offer discounts varying by time of day and day of week. Separating out firm incentives between price discrimination and peak-load pricing is often empirically challenging because most field data lacks an exogenous shifter of

[1] Examples include the Price Gouging Prevention Act of 2024; New York Bill A7062 and Washington SB 5600 restricting surge pricing; and the FTC's 6(b) study on surveillance pricing.

[2] As Disney executives explained, "we are adopting seasonal pricing on our one-day ticket to help better spread visitation throughout the year" (Time Magazine, February 27, 2016).

firm capacity constraints: given a time-invariant total capacity, price discrimination manifests as the firm response to intertemporal shifts in *consumer price sensitivity*, whereas peak-load pricing arises from those in demand *volume*. Teasing apart the volume and slope of the underlying demand from observed sales variations is known to be difficult.[3] We address this challenge by exploiting the unique environment in which restaurants use the shared capacity to serve both customers with a discounted reservation and walk-in diners who pay a regular price. Conditional on platform demand, variations in off-platform traffic serve as an excluded shifter of restaurant capacity constraints, thus allowing us to directly trace out how observed platform price variation reflects the restaurant's capacity considerations. We show that time-varying pricing, even when adopted by profit-maximizing firms, can generate substantial welfare gains when it primarily reflects capacity constraints rather than price discrimination.

We also show that peak-load pricing can reduce platform profits, weakening the platform's incentive to allow variable pricing even when such pricing benefits restaurants and consumers. When restaurants serve both platform and off-platform (walk-in) diners, platform discounts can shift reservations toward lower-priced off-peak periods while freeing peak-time capacity for off-platform demand. Since the platform is not compensated for off-platform sales, this diversion reduces its profits and generates incentive misalignment between the platform and participating restaurants. This misalignment is particularly pronounced under profit-sharing contracts between the platform and restaurants. Although such contracts eliminate double marginalization and are therefore conventionally viewed as efficient, they can further dampen the platform's incentive to offer variable pricing. Our results highlight a novel implication for platform contract design: when firms use shared capacity to serve both on- and off-platform markets, contracts that are efficient in isolation need not lead to efficient market outcomes.

Our empirical setting is a dining reservation platform operating in Germany and the U.K. on which restaurants offer advance-booking price discounts that vary by day of the week and by 30-

[3]The conventional discrete-choice framework requires the market size be given rather than estimated. An exception is the literature on dynamic pricing (Williams, 2022, Talluri and van Ryzin, 2004) in which the market size is estimated with auxiliary variations (e.g., firm optimality conditions).

minute time slots within a day. Discounts are highest during off-peak periods, when capacity is abundant, and lowest during peak periods, when tables are more likely to be fully occupied—e.g., 40% discount on Wednesday at 5:30 PM and 0% on Saturday at 7 PM.

The platform attracts demand primarily by offering time-varying discounts, which draws a more price-sensitive customer segment than restaurants' regular walk-in clientele. Customers who book through the platform receive the percentage discount applied to their entire bill, while the platform earns a fixed commission per reservation. Because restaurants' limited capacity is shared between platform and walk-in customers, restaurants have incentives to offer larger discounts during off-peak times and smaller discounts during peak times—a peak-load pricing incentive to shift platform users toward off-peak times and preserve peak-time capacity for full-price walk-in diners. Restaurants may also use variable pricing to price discriminate among platform customers with potentially time-varying price elasticities. We collaborate with the platform and observe detailed data on reservations and discount schedules, which we augment with measures of walk-in traffic throughout the day.

This setting provides a natural exclusion restriction that allows us to separate peak-load pricing incentives from intertemporal price discrimination. Restaurants' binding capacity constraints are driven primarily by peak-time traffic from off-platform diners, who do not face platform discounts. As a result, variations in off-platform traffic affect the platform prices, but not vise versa. We exploit this feature to identify peak-load pricing by examining how platform discount rates respond to changes in off-platform traffic. Moreover, an additional advantage is that pricing decisions are effectively static and therefore tractable. Roughly 89% of reservations are made on the day of dining, leaving little scope for dynamic capacity management, which is the focus among much of the existing literature (Williams, 2022, Cho et al., 2018). The static environment allows us to model consumer substitution between peak and off-peak times, a factor abstracted in most existing works due to computational complexity, thus characterizing the core trade-offs for restaurants that choose how much demand to divert to off-peak periods in order to preserve capacity for walk-in customers.

We begin by documenting substantial price variation on the platform across times of day. For the average restaurant, the gap between the highest and lowest price within a day is approximately 20% of the menu price. We show that these intra-day price patterns are strongly correlated with restaurants' walk-in traffic and that platform customers are willing to shift their reservation times toward lower-priced hours within the same day. Taken together, these findings point to peak-load pricing as the primary driver of observed intra-day price variation. By contrast, we find only modest heterogeneity in price sensitivity among platform users, suggesting that intertemporal price discrimination plays a limited role.

We next construct and estimate a model of demand and supply to disentangle restaurants' pricing incentives. In the demand model, platform customers arrive first and choose a restaurant and dining time. Each consumer has an ideal dining time but may shift across times or restaurants in response to discounts. We allow consumers' price sensitivities to vary across time, generating scope for intertemporal price discrimination. Given platform reservations, walk-in customers arrive exogenously and fill remaining tables. Because restaurants must honor reservations, binding capacity constraints imply that platform bookings crowd out walk-in diners who would otherwise pay full price. This crowd-out creates an opportunity cost of accepting reservations during peak periods, generating incentives for peak-load pricing: restaurants increase (decrease) peak (off-peak) period prices to steer platform demand toward less capacity-constrained times.

We estimate the model using reservations, prices, and restaurant walk-in traffic. On the demand side, we estimate an average price elasticity of approximately -10, with little intra-day variation. Consumers are primarily willing to substitute across dining times within the same restaurant, rather than switch across restaurants. On the supply side, we find that, for approximately 39% of restaurants, capacity binds with at least 50% probability during peak periods, and prices account for the possible revenue loss due to this binding capacity. If we counterfactually eliminate capacity constraints, intertemporal price variation declines by 94%. Together, these findings indicate that peak-load pricing is the almost sole driver of observed intra-day price variation.

We evaluate the efficiency gains from peak-load pricing by simulating a counterfactual in which

we assume a uniform discount rate within a day. Under uniform pricing, reservations tend to concentrate in peak periods, thus increasing the crowd out of walk-in customers. In this counterfactual, a platform booking displaces 0.21 walk-in customers (or displaces one 21% of the time), compared to 0.14 under variable pricing. Peak-load pricing hence decreases the displacement of walk-in customers by 32%. The improved capacity allocation increases the total welfare generated on the platform by 9.7% and improves both consumer surplus and restaurant profits. Therefore, peak-load pricing provides substantial efficiency gains and increases welfare.

We also find that, despite the welfare gains, the platform may not have incentives to offer variable pricing to restaurants. This incentive misalignment is amplified under profit-sharing, which is typically viewed as a more efficient contract. Under the observed per-reservation fee contract, variable pricing increases the platform's revenue by 11.3%. Thus, the platform's incentive is aligned with the restaurants, rationalizing the platform's decision to allow variable pricing in the data. In contrast, under a profit-sharing contract—which eliminates double marginalization—the platform loses revenue from variable pricing among 69% of restaurants. This misalignment arises from the fact that platform and off-platform markets share the same capacity. Peak-load pricing shifts platform reservations toward lower-priced off-peak periods, while the freed-up peak capacity is filled by off-platform diners whose higher-priced payments accrue only to restaurants. Profit-sharing exacerbates this effect because the platform's compensation depends on the (lower) platform transaction price. We show that, if the platform does not allow variable pricing among restaurants with misaligned incentives, overall welfare gains from variable pricing drop by 44%. Therefore, the misaligned incentives can prevent efficient pricing technologies from being deployed, and this inefficiency depends on the contract form between firms and the platform.

Our study is among the first to present empirical evidence that time-varying pricing generates substantial welfare gains, even when adopted by profit-maximizing firms. These gains arise when variable pricing is used primarily to manage capacity constraints by reallocating demand away from congested peak periods. The efficiency gains from improved capacity allocation and increased overall supply outweigh potential welfare losses from intertemporal price discrimination.

Because costly capacity investments often serve as an entry barrier, markets such as airlines where pricing power is a welfare concern tend to exhibit firm capacity constraints (Evans and Kessides, 1993). Our findings thus contribute to policy debates on regulating forms of variable pricing in concentrated markets: the welfare consequences depend on the mechanism driving price variation.

Taking a step back, our findings highlight a previously unexplored source of inefficiency in platform contracts that arises when participating firms serve both on- and off-platform markets using a shared resource. Because surplus generated off the platform is largely non-contractible, platforms cannot capture the surplus created by their own policies—such as allowing variable pricing—when those policies benefit paticipating firms through off-platform channels. Unlike standard vertical relationships in which profit-sharing contracts resolve channel incentive misalignment, the conventionally efficient contracts align incentives only over on-platform transactions, neglecting the off-platform returns to the welfare-improving policy, and therefore can prevent such a policy from being implemented.

1.1 Related Literature

We build on a growing literature that studies the welfare implications of variable pricing by capacity-constrained firms. A core strand of this literature examines firms' dynamic pricing strategies under capacity constraints, including studies of airlines (Williams, 2022, Hortaçsu et al., 2024), hotels (Cho et al., 2018), concert tickets (Sweeting, 2012), fresh food (Sanders, 2024), ride hailing (Castillo, 2025, Ming et al., 2025, Buchholz et al., 2025), and cruises (Joo et al., 2020). These settings share a common feature: capacity or inventory is perishable, in the sense that unsold units—such as airline seats, hotel rooms, or, in our setting, restaurant tables—"expire" if not utilized at the time of consumption.

We add to this literature by focusing on the firm incentive to redistribute sales *across products* (i.e., time slots), a key component for peak-load pricing but not addressed in existing studies due to its computational complexity.[4] Unlike conventional dynamic pricing environments in which the

[4]The unpublished manuscript by Lazarev (2013) also accounts for consumer substitution in the airline industry, but

firm sells a single product (e.g., a flight) within a limited timeframe, many business establishments such as restaurants, sports gyms and theme parks must operate with a fixed capacity in both low- and high-demand periods. Thus, inducing consumer substitution across periods is the first-order consideration. We also document unique implications of firm revenue management in platform markets where the ability to flexibly adjust prices is provided by the platform: the presence of misaligned incentives highlights the importance of identifying the firm's pricing incentives for platform designs.

We also relate to the literature on intertemporal price discrimination, in which sellers use time to segment consumers with heterogeneous valuations (Nair, 2007, Hendel and Nevo, 2013). This practice is a form of second-degree price discrimination: in the absence of targeted pricing, firms design price schedules that induce consumers to self-select according to their willingness to pay. Theoretical work has extensively studied this screening mechanism (Mussa and Rosen, 1978, Maskin and Riley, 1984), and a growing empirical literature documents its relevance across settings (Anderson and Song, 2004, Ghili and Yoon, 2024, Gardete et al., 2025). We add to this literature by considering how intertemporal price discrimination interacts with perishable capacity constraints. In such environments, price discrimination can coexist with peak-load pricing, and the resulting consumer sorting need not reduce welfare.

Our paper also builds on the extensive literature on peak-load pricing in public infrastructure contexts. Foundational theory characterizes peak-load pricing as a mechanism for balancing demand and supply in capacity-constrained public services such as electricity and water (Williamson, 1966, Turvey, 1968). More recent work studies consumer responses and welfare consequences of peak-load pricing regimes in regulated markets for electricity and natural resources (Joskow, 2012, Borenstein, 2012, Brecko and Hartmann, 2023, Garcia-Osipenko et al., 2025). We extend these insights to private markets, where prices are set by profit-maximizing firms and flexible, time-varying pricing is enabled by a digital platform rather than mandated by a regulator.

Finally, we relate to a recent literature examining how digital technologies facilitate price ad-

do not explicitly separate firms' capacity constraints from their operating marginal costs.

justment (Aparicio and Misra, 2023). Several studies document technology-enabled time-varying pricing across a wide range of settings, including pharmaceuticals (Brown and MacKay, 2023), food retail (Aparicio et al., 2024), gasoline prices (Assad et al., 2024), ride hailing (Castillo and Mathur, 2023), short-term rentals (Huang, 2022), hotels (Garcia et al., 2022), and car rentals (Filippas et al., 2023). MacKay et al. (2023) document that time-varying food delivery fees smooth demand and reduce marginal costs for a large restaurant chain. Our work extends this literature by documenting possible technology underprovision in platform markets: the technology provider's incentives can be misaligned with those that benefit from the technology, thus even a socially desirable technology may not be implemented.[5] Our study thus also adds to the growing literature on platform design (Einav et al., 2016) by empirically studying environments in which participating firms also operate in external markets and the welfare consequences of such externality.[6]

2 Empirical Setting

2.1 A Dining Reservation Platform

Our empirical setting is a dining reservation platform operating in Germany and the U.K. Restaurants that join the platform list available reservation times and specify a percentage discount for each listed time. Discounts vary by day of the week and by 30-minute time slots within a day (e.g., Friday 8:30 PM). Restaurants choose discount rates between 0% and 50% in increments of 10 percentage points. Appendix A illustrates the user interface.

When a customer books a table through the platform, the discount applies to the total bill with no upper limit. For example, a reservation made with a 10% discount reduces the final bill—

[5]This underprovision of technology connects to the literature on platform leakage, which studies how firms can actively divert transactions to off-platform channels to avoid platform fees (Hagiu and Wright, 2024, Xie and Zhu, 2023). In our paper, firms cannot actively divert traffic, but can reallocate shared capacity to off-platform market using variable pricing.

[6]Despite a growing theory literature (Armstrong, 2006, Belleflamme and Peitz, 2019, Bakos and Halaburda, 2020), empirical works on platforms with multi-homing firms are sparse to date (Li and Zhu, 2021, Chung et al., 2024, Sullivan, 2025).

covering food and drinks[7]—by 10%. In this sense, the platform discount is equivalent to a uniform price reduction applied to the restaurant's regular menu. The platform earns revenue through a fixed per-reservation commission.

The platform is the primary provider of flexible pricing technology in this market, allowing restaurants to set time-varying discounts in a seamless manner and thereby substantially reducing menu costs (Hobijn et al., 2006). Restaurants can adjust prices without reprinting menus or investing in proprietary pricing technology.[8] This flexibility enables restaurants to attract demand during off-peak periods, when capacity would otherwise go unused, and to generate incremental revenue.

Because consumers primarily join the platform to access discounts rather than as a substitute for their usual dining choices, they tend to be relatively price-sensitive and distinct from restaurants' regular customer base.[9] Although regular customers could in principle substitute toward booking through the platform, we show below that awareness of the platform among walk-in diners is limited, resulting in minimal substitution between platform and off-platform demand. Accordingly, throughout we treat platform customers and walk-in customers as distinct segments that draw on a shared, capacity-constrained set of tables.

Restaurants may have various incentives to adopt flexible pricing on the platform. One motive arises from capacity management. Restaurants serve both platform users, who receive discounts, and walk-in customers, who pay full price, using a fixed and capacity-constrained set of tables. During peak periods, restaurants can typically fill capacity with walk-in demand at full price, so accepting a discounted platform reservation entails an opportunity cost in the form of foregone full-price sales. By contrast, during off-peak periods, capacity is more likely to be underutilized, making discounted reservations more attractive. Time-varying discounts allow restaurants to re-

[7]Some restaurants impose exclusions (e.g., premium liquor items). In the vast majority of cases, such exclusions apply to a small subset of menu items.

[8]To the best of our knowledge, no other platform offers pricing flexibility at this level of temporal granularity. While some restaurants offer coarse discounts such as happy hours, discounts on the platform are typically larger in magnitude and subject to fewer restrictions.

[9]A parallel example in the U.S. is Groupon, which provided daily deals across multiple categories (Li et al., 2018). Groupon-style restaurant deals typically take the form of fixed-value vouchers or set menus that are valid only during specific times and often involve substantial usage restrictions, making them a coarse and inflexible instrument for managing demand compared to the fine-grained, time-varying discounts studied in this paper. Moreover, customers typically pre-purchase vouchers on Groupon rather than making a dining reservation at the time of consumption.

flect these differences by setting higher prices during peak periods and lower prices during off-peak periods, thereby steering platform demand toward times when excess capacity is available. This pricing behavior corresponds to peak-load pricing and, by improving the allocation of capacity over time as well as total supply, is expected to increase both restaurant profits and total welfare.

Another motive for variable pricing is intertemporal price discrimination among platform users. If customers differ in their willingness to pay across dining times—for example, if less price-sensitive consumers are also less willing to shift away from peak periods—restaurants may set higher prices during peak times to induce consumer sorting and increase profits. This mechanism represents a form of second-degree price discrimination that does not rely on capacity constraints and whose welfare implications are, in general, ambiguous. Distinguishing between these two pricing motives—capacity management versus price discrimination—by examining how prices respond to intertemporal variation in off-platform traffic and in consumers' price sensitivity is a central identification objective of the paper.

The platform operates in multiple cities in Germany and the U.K., with its largest presence in Berlin. By March 2020 (prior to the COVID-19 outbreak), 231 restaurants in Berlin were active on the platform. Participating restaurants are predominantly mid-range, small-scale establishments not affiliated with a chain. The platform restricts participation to restaurants above a minimum quality threshold, requiring a Google review rating of approximately 4.1 or higher. By design, participating restaurants are not expected to be a random sample of restaurants in the city. We discuss the implications of this non-representative market environment for external validity later in the paper.

2.2 Data

The dataset covers the universe of reservations made through the platform for restaurants in Berlin between October 2018, when the platform launched, and November 2020. For each reservation, we observe anonymized consumer identifiers, as well as the date and time of the reservation. In addition, for each restaurant $\times$ date we observe the full set of offered time slots and the corre-

sponding discount rates, regardless of whether a reservation is made. This allows us to observe both realized transactions and the complete menu of prices faced by consumers.

We also observe a rich set of restaurant characteristics, including name, location, cuisine type (e.g., Mexican), regular price per person, and average Google rating. The regular price per person reflects the average amount a typical diner would pay in the absence of a discount. The restaurant provides this information to the platform upon participation and the measure is displayed to consumers when browsing restaurants on the platform.

Prices are not realized at the time of reservation, as the final bill depends on the specific items ordered at the restaurant. We therefore define the price consumers expect to pay for a given time slot as follows: $P = \text{Regular Price} \times (1 - \text{Discount Rate})$. We assume that consumers base their reservation decisions on this expected price.[10]

Additionally, we construct a measure of restaurant traffic using data from Google Maps. Each restaurant's Google Maps page reports a traffic index that captures the relative volume of visitors by day of the week and hour of the day. This index is based on aggregated and anonymized Google device location data and is scaled between 0 and 100, where 0 indicates either that the restaurant is closed or has no guests, and 100 corresponds to the busiest period in the week. A visual example is shown in Appendix B.1. Our dataset contains a snapshot of this traffic measure for each restaurant $\times$ hour-of-day $\times$ day-of-week tuple.

We use the traffic index as a measure for walk-in demand—that is, diners without a reservation or those booking through channels other than the platform.[11] Note that the walk-in traffic is not a direct measure of capacity constraints, as it does not account for a restaurant's seating capacity. Instead, we use the supply-side model to infer time-varying occupancy rates that rationalize observed pricing behavior while accounting for alternative explanations. While the absence of direct seating-capacity data is a limitation, the traffic measure provides the key variation in how capacity

[10] Upon our manual investigation, the regular price roughly corresponds to one main dish, up to two side dishes and a drink. We abstract from the possibility that the platform users systematically consume differently from regular diners, as the actual bill is not observed.

[11] Although this measure also includes platform customers, we show below that platform transactions represent a small share of total restaurant traffic. As a result, platform demand has a negligible effect on this measure.

constraints shift over time, leaving time-invariant seating capacity as the primary parameter to be inferred from the model. Additional details on the traffic measure are provided in Appendix B.1.

To consider periods with relatively stable demand and supply conditions, we remove the first 11 months of observations corresponding to the platform's launch period, as well as observations after March 8, 2020, when the COVID-19 pandemic affected dining behavior. We further restrict the sample to restaurants that received at least 10 reservations during the six-month analysis window and to customers who made at least three reservations.[12] Finally, we consider time slots between 3:00 PM and 10:00 PM—corresponding to at most 15 half-hour time slots per day—which capture peak dinner periods as well as surrounding off-peak times and account for approximately 76% of reservations. Additional details on the sample selection process are provided in Appendix B.2.

2.3 Summary Statistics

The data include 6,188 reservations made by 1,357 customers across 80 restaurants over 190 days, spanning September 1, 2019 to March 8, 2020. Panel A of Table 1 reports summary statistics at the restaurant level. On average, restaurants offer 13.2 half-hour time slots per day out of the 15 slots considered in the analysis, indicating that most restaurants are available for reservations throughout the day. Consistent with this, 73% of restaurant–date pairs list either 14 or all 15 time slots as available on the platform.

Restaurants set discount rates ranging from 0% to 50%, with an average discount of 26%, indicating that discounts offered are economically meaningful. The resulting average discounted price per person is €13.2, with prices ranging from €4.25 to €54.00. Approximately 78% of available options are priced below €15 per person, suggesting that participating restaurants are concentrated in the lower- to mid-price range of the market.

On average, a restaurant receives one reservation every two days through the platform, al-

[12] We focus on repeat users because we exclude each user's first reservation from the likelihood function. Including first reservations would yield a potentially misspecified likelihood function, as we do not observe whether a user was active on the platform prior to her first recorded reservation. Consumers with three or more reservations account for approximately 22% of users 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)	13.199	3.024	2	15
Discount Rate	.257	.101	0	.5
Price (per Person, euro)	13.180	7.589	4.25	54
Number of Reservations (per Day)	.522	.948	0	9
Walk-in Traffic	48.203	24.305	0	100
(B) Customer-level:				
Number of Reservations	4.560	3.350	1	41
Average Discount Rate Booked	.317	.052	.1	.5
Average Price Booked (per Person)	10.672	3.162	5.65	37.625
Number of Restaurants	80			
Number of Customers	1,357			
Number of Reservations	6,188			

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.

though there is substantial heterogeneity across restaurants. The most popular restaurant receives an average of 3.29 reservations per day, while the least popular receives just 0.06. The relatively low volume of platform-based reservations indicates that the platform accounts for a small share of restaurants' overall demand.[13] Finally, the walk-in traffic data show that restaurants operate at nearly half (48.2%) the volume relative to their highest level throughout a typical week.

Panel B of Table 1 reports summary statistics at the consumer level. On average, consumers make 4.6 reservations over the six-month period. The average discount applied to their bookings is 32%. Additional details on other variables used in the demand estimation are provided in Appendix B.3. Histograms of prices and reservation activity are available in Appendix Figure I1.

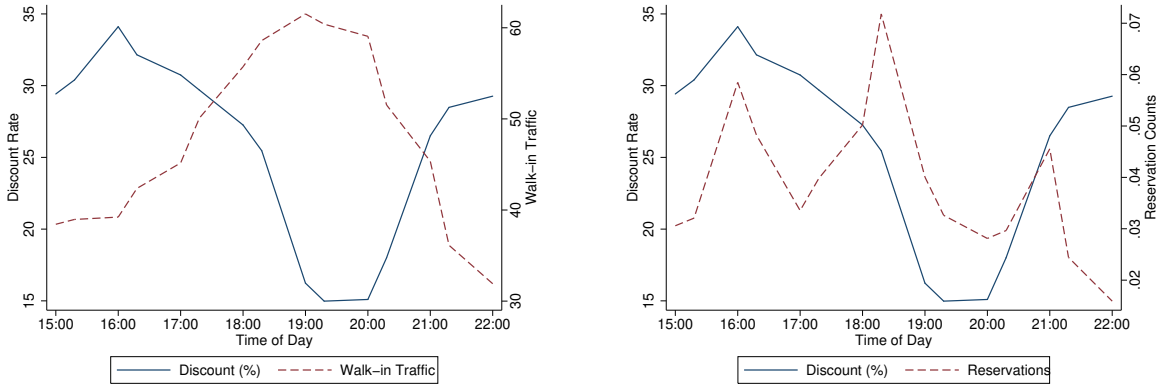
[13]The fact that restaurants adjust discounts across time, and raise prices during peak periods, suggests that even a small number of platform reservations can crowd out walk-in customers at the margin, allowing us to infer that many restaurants operate close to capacity.

3 Descriptive Evidence

We begin by documenting key empirical patterns in pricing and reservation behavior. First, we show that platform discounts closely track fluctuations in walk-in traffic, suggesting that variable pricing responds to expected capacity utilization. Second, platform users primarily respond to price variation by shifting their dining times within the same day, rather than across days. Finally, there is little heterogeneity in consumers' price elasticities, indicating a narrow scope for intertemporal price discrimination among platform users.

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 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 foot traffic: the negative correlation implies restaurants reduce platform discounts when they anticipate high walk-in traffic.



(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.

A simple regression of the discount rates (in percentage points, pp) on the traffic measure and restaurant fixed effects yields a coefficient of -0.12 (standard error of 0.001): a one-unit increase in

the measure of walk-in traffic is associated with a 0.12 pp lower discount rate for a given restaurant. This negative correlation is consistent with the hypothesis that pricing reflects capacity constraints. Because walk-in customers do not receive platform discounts, changes in discount rates cannot mechanically affect walk-in traffic, ruling out reverse causality from prices to traffic.[14]

Panel (b) depicts the average number of platform reservations per restaurant by time of day against the corresponding average discount rates. In contrast to walk-in traffic, platform reservations peak at 6:30pm, just before 7pm, when discount rates drop sharply. Two additional local peaks in reservations occur at 4pm, when discounts are highest, and at 9pm, immediately following a sharp increase in discount rates. These patterns are consistent with peak-load pricing: restaurants use intra-day price variation to steer platform reservations toward less congested time slots, while preserving peak dining periods for walk-in customers who pay full price.[15]

These figures also provide suggestive evidence that platform discounts do not induce substantial substitution from walk-in diners to the platform. Panel (b) shows that, despite average discounts of around 30% during off-peak periods, restaurants receive only about 0.4 platform reservations per day in those time slots, consistent with limited awareness of the platform among the broader dining population. Moreover, nearly all restaurants offer discounts even during peak periods—walk-in diners should be strictly better off by switching to platform reservation while still dining at their preferred time—yet receive, on average, only about 0.1 peak-time platform reservations. We therefore interpret discounts offered during peak periods as reflecting some price discrimination between walk-in diners and a small, price-sensitive segment of platform users, in a setting where substitution from walk-in to platform demand is limited.[16]

Restaurants also vary discount rates across days of the week, though the magnitude of this variation is substantially smaller than that of intra-day price variation. The largest difference in

[14]Panel (a) admits an alternative interpretation in which prices reflect systematic differences in platform demand between peak and off-peak times (which may be correlated with walk-in traffic) rather than capacity constraints. The reservations pattern shown in Panel (b) help mitigate this concern, as reservation volumes during peak times are among the lowest observed during the day. In the absence of capacity constraints, one would expect both prices and quantities to be higher during peak periods. We later adopt a flexible demand model to account for this alternative hypothesis.

[15]Platform users may have an intrinsic preference for off-peak dining times. We explicitly account for this possibility in our structural model.

[16]The model incorporates this form of third-degree price discrimination across markets as a fixed incentive.

discount rates across days—observed at 4pm between Sundays and Mondays—is approximately 7.5 percentage points, roughly half the size of typical intra-day price gaps. This more limited variation across days suggests that restaurants primarily use variable pricing to shift demand across time slots within a day rather than across dates. Appendix C.1 provides additional evidence on price variability across days of the week.

Interestingly, restaurants adjust their discount schedules infrequently. Among the 80 restaurants in the sample, 18 never change their price schedule, and among the remaining 62 restaurants, the average frequency of any price adjustment is once every 71 days. By contrast, the vast majority of reservations are made at the last minute: 88.7% of reservations occur on the day of dining, and fewer than 3% are made more than three days in advance. This indicates that, unlike in markets such as airlines where prices change as the departure date approaches, restaurants set discount schedules well in advance and do not adjust prices dynamically as tables are occupied. Moreover, nearly all time slots remain available on the platform until shortly before the dining time (e.g., 30 minutes prior), suggesting that platform reservations alone do not saturate restaurant capacity (but walk-in traffic does). Thus, we assume a platform reservation may crowd out a walk-in customer, but does not crowd out another 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:

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

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's core innovation is its technology that enables restaurants to adjust prices in response to demand fluctuations, observed prices may be endogenous to unobserved demand shifters. Identifying valid instruments is also challenging, as much of the observed price variation reflects consumer sorting across time slots.[17] To address this concern, we exploit the predetermined nature of restaurants' time-varying pricing schedules.

As documented above, restaurants rarely update discount rates on a day-to-day basis; instead, they assign fixed discount schedules to each day of the week and maintain them over extended periods. These infrequent adjustments are likely responses to gradual, long-run shifts in demand, and the exact timing of schedule changes is plausibly uncorrelated with idiosyncratic daily demand shocks. Following this identification strategy, we 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. Additionally, we include calendar day fixed effects, μ_t, to account for platform growth and demand seasonality that are common across restaurants.

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 -8.99. 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 2.47, whereas that to the same time slot of the adjacent day has an unexpected sign and is small. Combined

[17]The walk-in traffic measure is not a valid instrument because it may be correlated with the platform customers' demand.

Table 2: The Effect of Prices on Reservations

	(1)	(2)	(3)	(4)
$P_{j,t,\tau}$	-.821 (.027)	-8.745 (.413)	-10.183 (1.017)	-.038 (.004)
$\bar{P}_{j,t,\tau\pm1}$			2.467 (.839)	.007 (.003)
$\bar{P}_{j,t\pm1,\tau}$			-.023 (.980)	.003 (.004)
j-dow(t)-τ FE	No	Yes	Yes	Yes
Date FE	No	Yes	Yes	Yes
Observations	156,558	156,531	117,894	117,894

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.

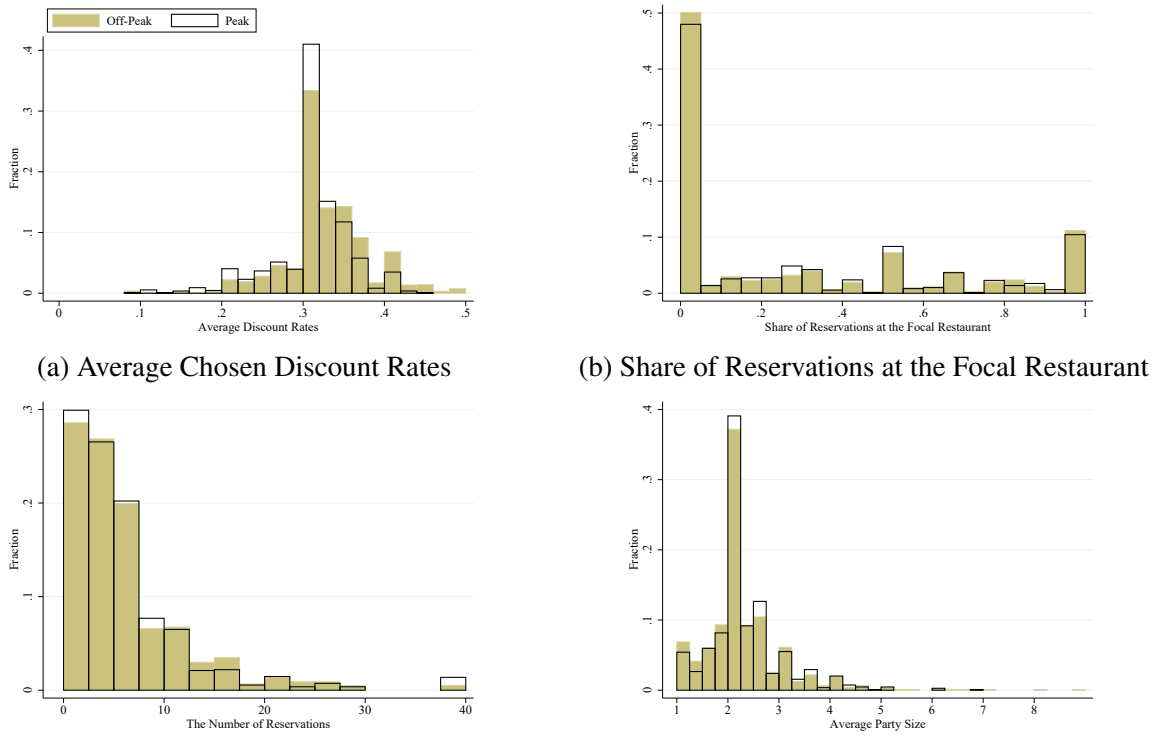
with the relatively limited variation in prices across dates, these results suggest that the primary channel through which peak-load pricing operates is intra-day substitution, rather than inter-day reallocation of demand.

A large share of observations report zero reservations, and thus the log price specification may not precisely represent price elasticity. As a robustness check, we estimate 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 -12.71, and the intra-day cross-price elasticity for adjacent time slots is 2.36.

These results collectively demonstrate that the observed price variation effectively induces platform customers to shift their reservations from peak to off-peak hours. This behavior is consistent with the patterns shown in Panel (b) of Figure 1. For restaurants facing binding peak-time capacity constraints, such intra-day demand shifting increases overall capacity utilization. Appendix C.2 reports similar elasticity estimates obtained using alternative sources of identification.

3.3 Possible Role of Intertemporal Price Discrimination

In addition to peak-load pricing, restaurants may use variable pricing for intertemporal price discrimination among platform users. For instance, when less price-sensitive customers have a stronger preference for peak hours, higher peak-time prices allow restaurants to extract additional surplus. In this case, observed price differences reflect heterogeneity in consumers' willingness to pay across time slots, in addition to incentives for capacity management.



(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 rate), 0.03 (share of reservations at the focal restaurant), 0.03 (the number of reservations), and 0.05 (average party size).

A direct implication of the price discrimination hypothesis is that consumers who reserve during peak versus off-peak times should differ in their price sensitivity. 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 (30.7%).[18]

The observed differences in price sensitivity arise 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 an off-peak-time booking and 25% for a peak one.

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 (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 the average chosen discount rate.

Taken together, these results provide suggestive evidence that variable pricing induces some consumer sorting across time slots, consistent with a potential role for intertemporal price discrimination. However, the economic magnitude of this effect is modest: the difference in average discount rates chosen by off-peak versus peak-time diners is only 1.7 percentage points. Moreover, platform users exhibit relatively similar observable characteristics, pointing to a limited scope for intertemporal price discrimination on the platform. In the structural model, we allow for heterogeneity in consumers' price sensitivities and time preferences to explicitly account for price

[18] 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 I2.

discrimination as an 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 $\beta_{i,\tau_{it}}^\tau$. 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 time may vary across days and follow a type-specific distribution $F_{it}(\tau_{it})$. The distribution varies 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 realized ideal dining time, allowing flexible correlations between a consumer's price sensitivity and her preferred dining time. This specification thus allows time-varying price sensitivity within a person — for instance, an early dinner is typically associated with a casual dining occasion with some scheduling flexibility, whereas peak-time reservations may be for business purposes.[19]

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 $q_{\tau_j t} = M \times Pr_t(j|X_t, p_t, \tau_t)$, 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:

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

[19] Time-varying $\beta^\tau_{i,\tau_{it}}$ substantially increases the model flexibility to capture any systematic shift in platform demand across time. Accounting for all the intertemporal shifts in platform demand is crucial to properly isolate out the restaurants' price-discrimination incentive and attribute the remaining intertemporal price variations to (partly unobservable) capacity constraints.

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 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:

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

where c_r is the marginal cost of serving a customer, c_r^p is the platform commission per reservation, $n_{\tau_j t}$ 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.[20]

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.[21] 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 restaurant can sell on the platform without crowding out walk-in customers. We assume that $x_{\tau_j t}$ is independent across τ_j and

[20] 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 D, we show that even when the average party size exceeds one and potentially differs between platform and walk-in customers, it only affects the interpretation of L_r and $n_{\tau_j t}$.

[21] 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}.

t, and follows a Poisson distribution with rate $\lambda_{\tau_j t}$. In our empirical application, we parameterize $\lambda_{\tau_j t}$ as follows:

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

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 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 time slots. 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. Solving the maximization problem yields the following first-order conditions:

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

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 $Pr(q_{\tau_{j'} t} > x_{\tau_{j'} t})$, 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 with most preferred dining time 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 allow the platform users' ideal dining time to differ from that of walk-in diners by characterizing $F_{it}(\tau_{it})$ as a mean-preserving spread of the distribution of walk-in diners. Specifically, we denote by $\tilde{F}_t(\bar{g}_{\tau_{it}t})$ the empirical distribution of our walk-in traffic measure at time τ_{it} on day t, $\bar{g}_{\tau_{it}t}$, aggregated across restaurants (we detail the construction of it in Appendix B.1). We parameterize $F_{it}(\tau_{it})$ as follows:

$$F_{it}(\tau_{it};\beta_i^F) = \bar{\tilde{F}}_t(\bar{g}_{\tau_{it}t}) + (\tilde{F}_t(\bar{g}_{\tau_{it}t}) - \bar{\tilde{F}}_t(\bar{g}_{\tau_{it}t}))\beta_i^F,$$

where $\bar{\tilde{F}}_t(\bar{g}_{\tau_{it}t})$ is an average value of $\tilde{F}_t(\bar{g}_{\tau_{it}t})$ within a day (which equals 0.0667, given there are 15 time slots). The segment-specific parameter β_i^F determines the difference in time preference between platform users and walk-in diners. $\beta_i^F < 1$ implies that platform users' ideal time is more dispersed to off-peak times, whereas $\beta_i^F > 1$ implies more concentration for peak times.[22] Our

[22]When the estimate of β_i^F produces $F_{it}(\tau_{it}) < 0$, we set it to zero and scale up all the other terms proportionally to ensure $\sum_{\tau_{it}} F_{it}(\tau_{it}) = 0$.

model thus allows for possible selection in which platform users sort into the platform because they have a higher baseline preference toward off-peak hours.

We identify and estimate β_i^F, along with the switching cost parameter, $\beta_{i,\tau_{it}}^{\tau}$, and the nest parameter, σ_i, by utilizing the panel structure of the data. The switching cost parameter is identified by how far across time a user is willing to substitute away due to a price change. The nest parameter is identified by asymmetric substitution between restaurants and within a restaurant across times. β_i^F is identified by user decisions that do not involve prices: for instance, among identically-priced time slots, users should select a slot closest to their ideal dining time. Thus, the parameters that govern the intertemporal substitution are identified by separate variations in the data.

To identify the price coefficients, we exploit restaurants' infrequent adjustments of their discount schedules, the same strategy we used in Section 3.2. We assume that the control variables (X_{jt}) accounts for all predictable demand shifters 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 (outside of the model) 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, district fixed effects, fixed effects for year-month pairs and for weekend days. The coefficients for these variables are segment-specific. In addition, we include district $\times$ cuisine type fixed effects, whose coefficients are common across segments (due to the dimensionality of these interactions). To measure 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.[23] In

[23]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

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_r^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 $g_{\tau_j+\tau',t}$.[24] 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.[25]

This identification strategy exploits our unique market structure where two pricing incentives arise from separate channels. On the one hand, walk-in demand affects the peak-load pricing incentive, but it is independent of intertemporal price discrimination because walk-in diners do not receive discounts. On the other hand, heterogeneous platform customers can create the intertemporal price discrimination incentive, but the small platform demand by itself does not require peak-load pricing in the absence of walk-in traffic. This structure substantially simplifies an otherwise nontrivial separation task: if such an external shift in capacity does not exist—suppose instead that no walk-in market exists and peak-load 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 discrimi-

it could be positive if higher restaurant density suggests a more active business landscape.

[24]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.

[25]This identification strategy requires any factor that affects platform demand that is correlated with $g_{\tau_j t}$ be in the demand function, or we misattribute the unobserved demand shifts as capacity-driven incentives. We establish this by allowing $g_{\tau_j t}$ to flexibly affect both the platform customers' preferred time (in $F(\tau_{it})$) and their utility (in X_t).

nation 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:

$$L(\theta) = \prod_{k} \left(\int_{\beta_i} \prod_{t=\underline{t}_k}^{189} \left(\int_{\tau_{it}} Pr_{it}(j_{kt}|X_t, p_t, \tau_t) dF_{it}(\tau_{it}) \right) dG(\beta_i) \right),$$

where k represents a consumer in the data and i represents the discrete type. For each consumer k, the likelihood corresponds to the joint probability that the consumer takes the observed action (j_{kt}) on each day during the sample period. We denote by $\underline{t}_k$ the first day that consumer k is observed in the data, which we define as the day following their first reservation.[26]

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.[27]

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 76.2% of the customer base. Both segments show low intercepts, driven by the fact that an average consumer makes only four reservations

[26] 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."

[27] 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 E.

over the six-month period. For each segment, we estimate four price coefficients and the cost of switching time slots, corresponding to the consumer's most preferred dining time being each 90-minute interval.

	Segment 1		Segment 2	
Intercept	-7.459	(.464)	-12.961	(1.484)
Price (τ_{it} ∈[3pm, 4pm])	-.345	(.014)	-.747	(.021)
Price (τ_{it} ∈[4:30pm, 5:30pm])	-.466	(.022)	-.782	(.029)
Price (τ_{it} ∈[6pm, 7pm])	-.374	(.010)	-.657	(.017)
Price (τ_{it} ∈[7:30pm, 8:30pm])	-.343	(.008)	-.641	(.017)
Price (τ_{it} ∈[9:00pm, 10:00pm])	-.330	(.014)	-.625	(.022)
Switching cost (τ_{it} ∈[3pm, 4pm])	-.014	(.006)	-.038	(.021)
Switching cost (τ_{it} ∈[4:30pm, 5:30pm])	-.000	(.002)	-.061	(.049)
Switching cost (τ_{it} ∈[6pm, 7pm])	-.195	(.037)	-1.083	(.318)
Switching cost (τ_{it} ∈[7:30pm, 8:30pm])	-.118	(.020)	-.252	(.074)
Switching cost (τ_{it} ∈[9:00pm, 10:00pm])	-.000	(.002)	-73.39	(99.99)
Review Ratings	-.291	(.086)	.394	(.310)
Popularity	.061	(.017)	.236	(.067)
Regular Price	.221	(.001)	.538	(.002)
Nest Parameter	.570	(.015)	.000	(.061)
Peak Concentration (β_i^F)	3.739	(.420)	1.816	(.256)
Fraction Segment 1	.762	(.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. Parameter estimates not reported in this table are available in Appendix Table I1.

Table 3: Demand Parameter Estimates

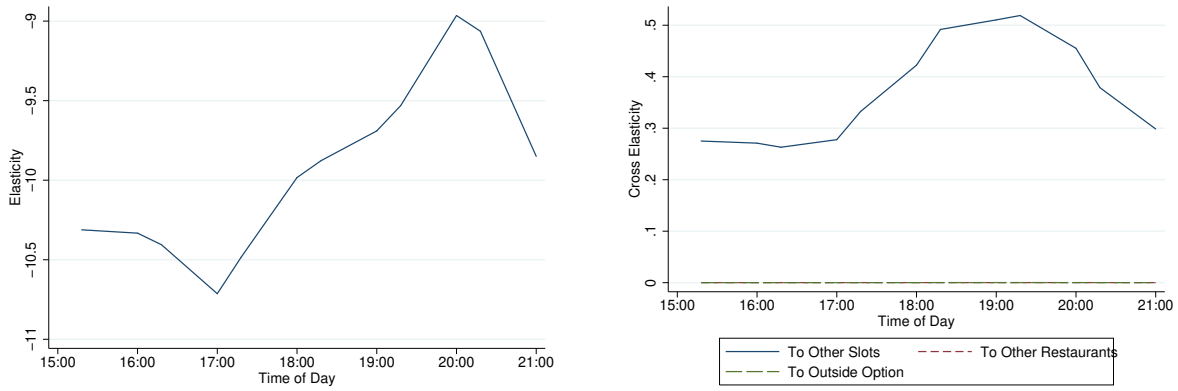
Among segment 1 consumers, we find that price coefficient is mostly flat during the day, except for the 4:30-5:30pm period in which the price sensitivity increases by about 20% to -0.47. The cost of switching between time slots is higher when the consumer's most preferred time is during the peak time, whereas it is close to zero elsewhere. The willingness to accept to switch to the adjacent time slot—given by the ratio between the switching cost parameter and the price coefficient—is about 0.3 to 0.5 euros when their preferred dining time is between 6pm and 8:30pm, and it is near

zero during other times. Taken together, majority consumers are willing to shift their dining time with very small compensation, though shifting the dining demand for the peak time requires a slightly higher compensation value.

Furthermore, segment 2 has a large negative intercept and larger (in magnitude) price coefficients than segment 1, indicating this segment rarely dines out. Similar to segment 1, this segment is easier to shift their dining time when they prefer early dining time. We find that when a consumer from this segment prefers a late evening time, she has little willingness to switch to other times, indicating the presence of late-night reservations that do not substitute.[28] Overall, both segments show willingness to shift their dining time in response to discounts, 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 restaurant quality. 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. The nest parameters indicate that segment 1 consumers mostly switch to other slots of the same restaurant, whereas segment 2 evenly substitute across restaurants and times. Finally, the relative preference of platform users toward the peak time (β_i^F) is larger than one for both segments, indicating that the ideal dining time for the platform users are more concentrated to the peak times than the observed distribution of walk-in customers. Thus, we do not find evidence of user selection that platform users have an innate preference for off-peak times. Appendix Figure I3 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 periods.

[28]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 3: 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.

In Panel (a) of Figure 3, we show the average own-price elasticity within the day and across time slots. We estimate an own-price elasticity between -9 and -10.5, indicating that consumers on the platform are quite price sensitive, likely more so than average diners. The model-implied price elasticities are within the range of the reduced-form estimates presented in Table 2 (-9 to -12.8 across specifications), 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 F. 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 smaller reservation volume around 5 pm.[29] As we show in Section 7.1, this shift in elasticity is not large enough to allow meaningful intertemporal price discrimination.

Panel (b) of Figure 3 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 restaurant (as per the

[29]The model mechanically generates lower elasticity at the start and the end of the day, because these slots lack one adjacent slot. Figure 3 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.

solid line), while substitution to other restaurants is negligible. Although the cross-elasticity 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. The within-restaurant substitution is larger during the peak time, because segment 1 consumers, who primarily substitute across times, have stronger preferences toward peak times (as represented by a larger value of β_i^F).

6.2 Supply

In Figure 4, we show the distribution of marginal cost estimates. The average marginal cost is €7.70, and the standard deviation is 6.61. The average profit margin is 59.0% of the regular price and 43.3% of the discounted price. Our marginal cost estimates are statistically significant for 76 out of the 80 restaurants.

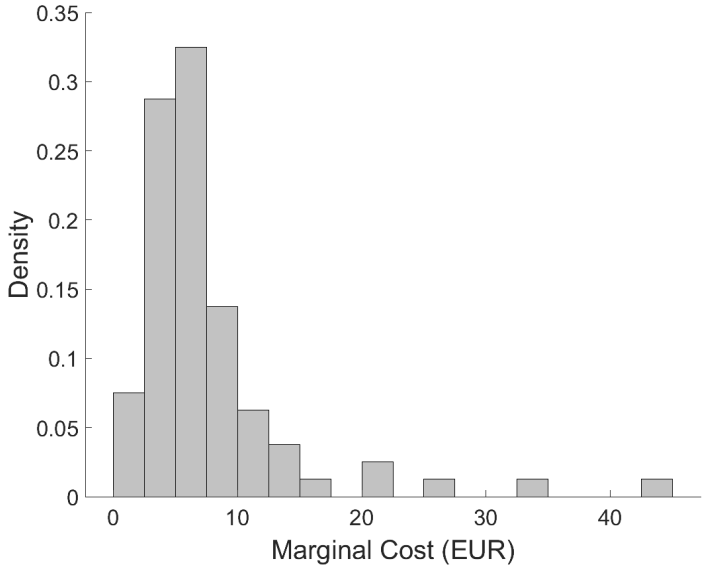
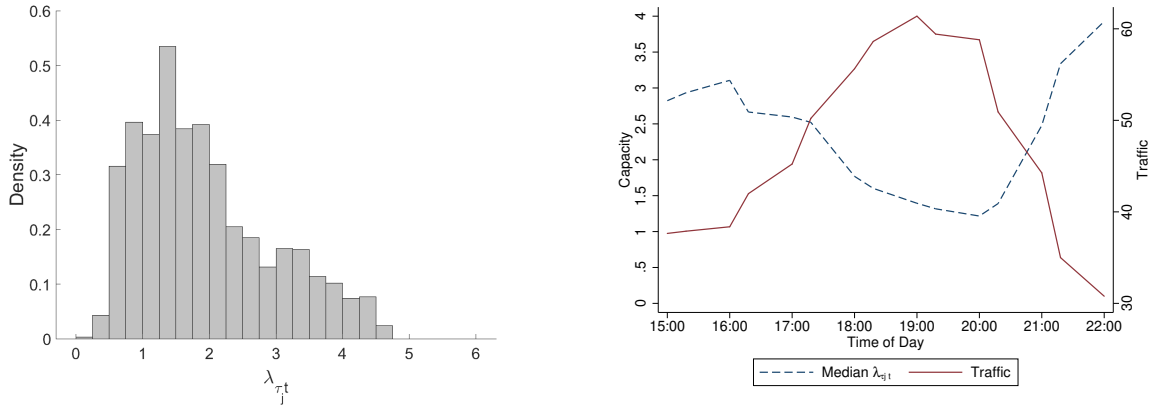


Figure 4: Marginal Cost

Note: The unit of observation is a restaurant.

Figure 5 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.[30] In our data, 80% of observations exhibit more than one percent of binding capacity, which translates to an informative value of $\lambda_{\tau_j t}$. The average and the median value of $\lambda_{\tau_j t}$ are 1.93 and 1.72, respectively, with the standard deviation of 1.00.



(a) Overall distribution

(b) Across-time Variations

Figure 5: Expected Available Capacity ($\lambda_{\tau_j t}$)

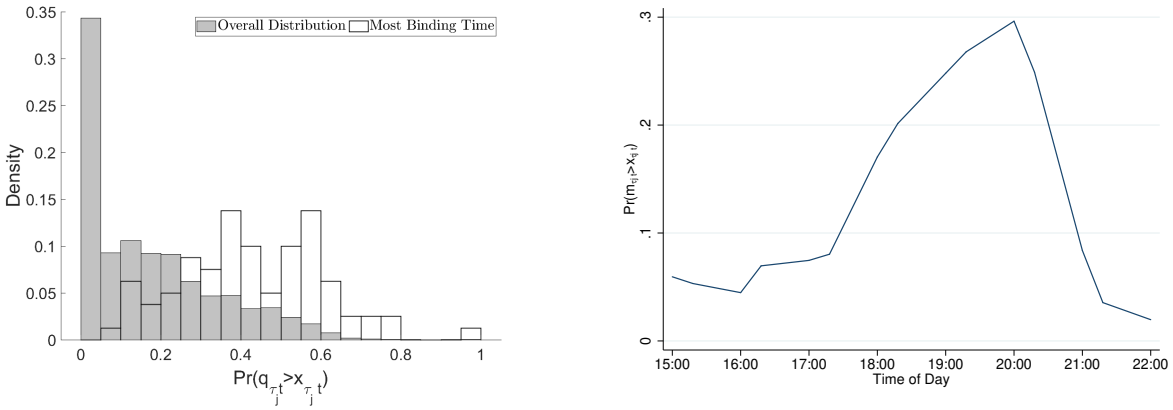
Note: In Panel (a), the unit of observation is restaurant-day-time tuple. $\lambda_{\tau_j t}$ estimates that imply less than one percent chance of binding capacity are not informative and hence are dropped. In Panel (b), the dashed line corresponds to the median value (across restaurants and dates) 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 6, 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 0.171, with a standard deviation of 0.166. Focusing on the

[30] 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 are obtained as pure parametric extrapolations and are irrelevant for our study.

busiest time for each restaurant, we find that 31 restaurants (39% of the sample) experience a sell-out probability greater than 50%. Panel (b) displays the within-day variation in sellout probability, evaluated at the sample median. During peak hours between 7pm and 8pm, a median restaurant faces a 25-30% 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. We present the model fit for the intra-day pricing patterns in Appendix Figure I4. Overall, the predicted prices closely track the observed variation throughout the day, although the model predicts the restaurants should start raising the prices earlier in the day.[31] The average predicted price is €13.79, compared to an average observed price of €13.18.



(a) Distribution

(b) Across-time Variations

Figure 6: 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.

These results indicate 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

[31] This is because the observed change in discount rates lags the surge of the traffic, as seen in panel (a) of Figure 1.

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, due to the selection of users and restaurants.

7 Welfare and Profit Implications of Peak-load Pricing

We now use the model to analyze underlying mechanisms and welfare implications of variable pricing. 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, thus demonstrating that peak-load pricing improves welfare even when the firm's objective is profit maximization. We also show that variable pricing benefits the platform under the observed per-unit-fee contract. In Section 7.3, we consider an alternative contractual scheme in which the platform and each restaurant share their joint profit, and show that variable pricing decreases the platform's profit among the majority restaurants. The misaligned incentives reduce the platform's incentive to enable variable pricing, thus partly offsetting the relative efficiency of profit-sharing contracts over per-unit-fee contracts that arises from eliminating double-marginalization.

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 observed price schedule on the platform. 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 five estimated coefficients for that segment.

This counterfactual eliminates the restaurants' incentive for intertemporal sorting between platform users, while keeping the peak-load pricing incentive to shift the (now mostly homogeneous) platform users to off-peak times.[32] 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 intertemporal price discrimination and capacity constraints. Figure 7 displays the intra-day price variability under each scenario, focusing on the set of restaurants that are available at all slots. The solid line is the average prices under the factual benchmark.

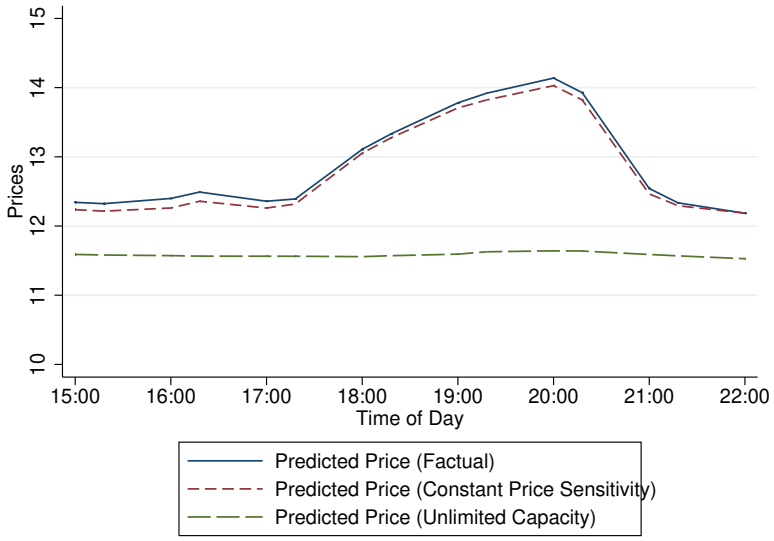


Figure 7: Predicted Prices

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

We find that capacity constraints have a substantial influence on the pricing schedule. Removing the capacity constraint (the long-dashed line in Figure 7) reduces the average price by 14%, from €13.79 to €11.92. 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 94% on average, representing a marked reduction in price dispersion during the day. In contrast, eliminating the intertemporal variation in consumer price sensitivity has a more limited impact. The short-dashed line in Figure 7 shows the price patterns

[32]This analysis keeps the heterogeneity between platform users and walk-in diners, thus not eliminating all the price discrimination opportunities in the market. We only eliminate the part of price discrimination to which intertemporal pricing is relevant, a more relevant object to compare with peak-load pricing incentives.

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 intertemporal 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 within the day (allowing prices to vary across days). We hold other market conditions fixed, such as the number of users and participating restaurants, and restaurants' regular prices and operating hours.[33] 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.[34] We assume the per-unit fee, c_r^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.[35]

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. Under variable pricing, the platform generates the welfare of approximately €205 per day, compared to

[33]The former can be restrictive if the platform differentiates itself by this variable-pricing ability. Because the platform offers restaurants access to new consumer segments and restaurants can always offer tables without a discount, restaurants are likely better off not quitting the platform 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 abandon the platform merely due to the lack of time variability in discounts.

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

[35]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 Pricing**	(2) **Uniform Pricing**
Panel (A): Gains Generated by the Platform		
Total Welfare Generated by the Platform (per Day, €)	205.1	186.9
Restaurant Profits from Participation	87.90	80.49
Platform Profit	28.03	25.18
Platform Users' Surplus	89.12	81.24
Panel (B): Effect on The Platform		
Daily Reservations	32.24	29.43
Average Price Booked (€)	13.20	13.98
Fraction of Peak-Time Reservations (7-8 pm)	.245	.350
Restaurant Profit per Reservation (€)	4.407	5.192
Surplus per Reservation (€)	6.910	7.690
Panel (C): Effect on Walk-in Customers		
Displaced Walk-in Customers per Reservation	.140	.205
Lost Profit due to Displacement per Reservation (€)	1.454	2.212

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 G.

Table 4: Impact of Variable Pricing

€187 under uniform pricing. This represents a lift of roughly 9.7%.

Nearly all 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 €87.9 per day under variable pricing, compared to €80.5 under uniform pricing, representing a 9.2% increase. Figure 8 displays the distribution of the percentage profit change across restaurants. About half of the restaurants see a modest profit increase of up to 10%, while three restaurants experience gains of over 100%. Five restaurants exhibit a slight profit decline, likely due to intensified price competition, but the magnitude of the loss is negligible (0.05% to 0.22%). These minor losses suggest that any competitive disadvantages from adopting variable pricing are economically small.

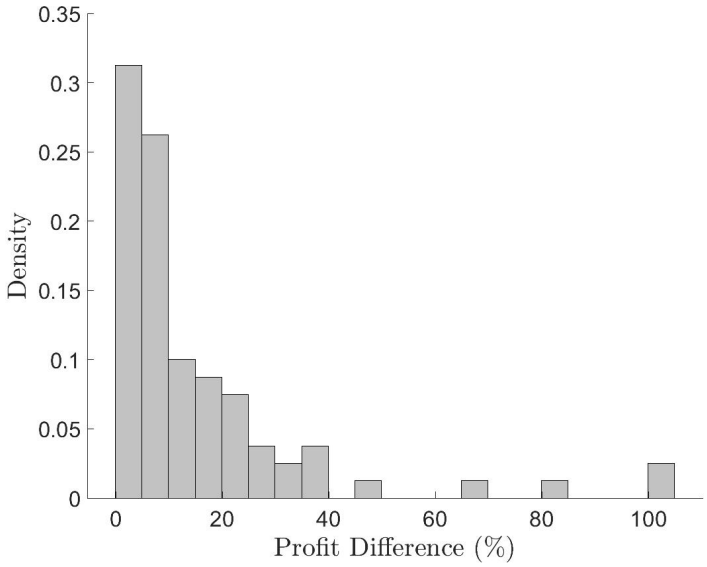


Figure 8: Profit Increase due to Variable Pricing (% of Uniform Pricing Profit)

Note: The unit of analysis is a restaurant - date pair. The rightmost bar denotes an increase above 100%.

The platform's total commission increases by 11.3% from €25.2 to €28.0 per day. The gain arises because a lower average transaction price substantially increases the number of reservations, as we show in Panel (B). The users' surplus also increases by 9.7% from €81.2 to €89.1. Overall, these findings indicate that variable pricing benefits both restaurants and the platform in a roughly equal manner under the observed per-unit-fee contract.

Panels (B) and (C) of Table 4 decompose the total welfare gain into on- and off-platform transactions, thereby illustrating how peak-load pricing reallocates capacity. Panel (B) focuses on outcomes within the platform. We find that variable pricing leads to a 9.5% net increase in reservation volume (from 29.4 to 32.2), with a 5.6% drop in the average transaction price (from €13.98 to €13.20). Variable pricing redistributes the demand away from congested time slots through lower off-peak prices: under uniform pricing, approximately 35.0% of all reservations occur during peak hours (7 to 8pm), whereas this share declines to 24.5% under variable pricing.

We find this shift to off-peak times reduces average restaurant profit per reservation from €5.19 to €4.41, reflecting the lower price point booked. Notably, for the majority restaurants, their *total* profit from platform customers also decreases even after accounting for the increased reservation

volume. The surplus a reservation generates on the platform (not accounting for displaced walk-ins) decreases from €7.69 to €6.91. These declines indicate that restaurants primarily benefit from preserving peak-time seats for full-price walk-in customers, albeit at the cost of lower platform profits. This sales diversion to off-platform markets, however, can reduce the platform's incentive to allow variable pricing, a question we turn to in the next section.

Panel (C) reports the welfare impact among walk-in customers. We find that variable pricing reduces the number of displaced walk-in customers from 0.21 to 0.14 per reservation, or a 32% decrease. The forgone profits from these displaced walk-in diners also decrease from €2.21 to €1.45 per reservation, or a decrease of 34%. Because the forgone profits under the uniform pricing account for 29% of the surplus a reservation generates on the platform, these improvements in allocative efficiency are economically meaningful. In particular, because the efficiency gains per reservation are quite large, variable pricing reduces the total volume of displaced walk-in traffic during the day by 11.1% despite the larger reservation volume. Once again, the measured gains do not include the increased surplus of walk-in customers, thus representing a conservative estimate of the actual gains 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. These results offer new empirical evidence that profit-maximizing behavior by firms can contribute to welfare improvements in markets characterized by capacity constraints.

7.3 Misaligned Incentives and Platform Compensation Scheme

We next show that misaligned incentives between restaurants and the platform can inhibit the use of variable-pricing technology despite its welfare benefits, and that such misalignment depends on platform compensation scheme. In practice, our focal platform charges restaurants a flat fee per reservation. As we show in Section 7.2, variable pricing increases the platform's profit by 11.3%

under this scheme, indicating that the platform's incentive is aligned with that of restaurants and rationalizing the platform's decision to allow variable pricing in the data. However, conventional wisdom posits that this scheme also typically creates a double-marginalization problem, in which the platform fee corresponds to an increase in marginal costs for the restaurants, leading to an inefficient outcome due to upward-distorted prices.

In this section, we consider a counterfactual environment in which the platform and each restaurant shares their joint profit, a scheme conventionally known to eliminate double marginalization. We show that while a profit-sharing contract enhances the efficacy of peak-load pricing, it substantially reduces the platform's gains, thus potentially leading to lower adoption of variable pricing in the market. We thus show that employing a profit-sharing contract—thought to be efficient in many settings—might not lead to an efficient outcome in the context of peak-load pricing in platform markets. We discuss mechanisms behind the misaligned incentives and general implications to platform contracts.

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

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

where ρ_r^p corresponds to the fraction of the restaurant's profit from platform customers that is shared to the platform. Correspondingly, the platform's profit is given as follows:

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

Note that only the profit on the platform is shared—not the walk-in profit—because contracting on off-platform outcomes is practically infeasible. This feature presents the core incentive misalignment: the off-peak discounts offered by restaurants are excessive from the platform's point of view as the platform does not benefit from the freed-up peak-time capacity. Although as we show below the profit-sharing contract remains more efficient than per-unit fees, this inability to contract

on off-platform outcomes prevents the market from achieving full efficiency.

	Profit Share			Per-Unit Fee	
	Variable	Partly Variable	Uniform	Variable	Uniform
	(1)	(2)	(3)	(4)	(5)
Welfare Generated by the Platform	270.4	257.2	240.3	205.1	186.9
Restaurants' Profit	102.80	97.44	91.07	87.90	80.49
Platform Profit	25.87	26.15	25.18	28.03	25.18

Note: The metrics shown in the three rows are identical to the first three rows of Table 4. All the metrics are per day, summed over restaurants. "Partly variable" outcome corresponds to the case in which the platform exclude restaurants with misaligned incentives.

Table 5: The Role of Platform Compensation

In what follows, we fix restaurant-specific ρ_r^p at the value such that the platform's profit under the baseline case of uniform pricing is identical to the factual per-unit-fee contract. In Table 5, we present the results. Columns 1 and 3 correspond to variable- and uniform-pricing regimes under the profit-sharing contract (we discuss column 2 below). Columns 4 and 5 correspond to the factual per-unit fee structure and are identical to Table 4, presented for comparison. The metrics shown in the three rows are identical to the first three rows of Table 4. Comparing columns 3 and 5, 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 the welfare of €240.3 per day, as compared to €186.9 under the per-unit fee, or an increase of 28.5%. The restaurants' profits also increase by approximately 13.1% from €80.5 to €91.1 per day. Thus, we find that a profit-sharing contract can lead to a substantial efficiency gain by eliminating double marginalization, in the baseline environment of uniform pricing.

We find that the profit-sharing scheme also substantially enhances the efficacy of peak-load pricing. Comparing columns 1 and 3, we find that variable pricing increases welfare by 12.5% from €240.3 to €270.4, accompanied with a 12.8% increase in restaurant profits from €91.1 to €102.8. Compared to the corresponding gains under the observed per-unit fee contracts, these numbers amount to 129% and 139% in percentage terms, or 165% and 158% in Euro values.

Eliminating double marginalization, by lowering the restaurants' effective marginal costs, allows restaurants to offer deeper off-peak discounts, thus further improving the restaurants' capacity utilization and providing complementary benefits to the use of peak-load pricing.

However, we also find that the platform's gain from variable pricing is limited to only 2.7% (from €25.2 to €25.9) under the profit-sharing scheme, which is only 24% of the gains under the per-unit-fee scheme (11.3%). The smaller gains are the representation of the aforementioned misaligned incentives; restaurants offer deep off-peak discounts to induce intertemporal substitution of platform customers from peak- to off-peak hours, lowering the average transaction price on the platform while increasing peak-time walk-in sales that only benefit the restaurants. This diversion of peak-time sales particularly hurts the platform under the profit-sharing scheme because, unlike per-unit fees, the platform's profit directly reflects the transaction prices that drop sharply.

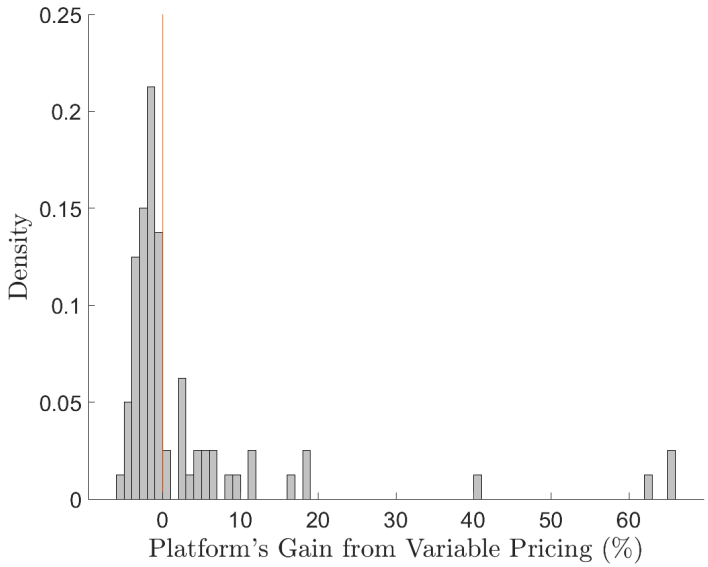


Figure 9: Platform's Gains from Allowing Variable Pricing

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 243%.

Figure 9 depicts the distribution of the platform's gains from allowing variable pricing across restaurants under the profit-sharing scheme. We find that the platform's profit decreases among 55 restaurants, or 69% of the sample, indicating that the platform potentially lacks incentive to provide

variable pricing to a large fraction of restaurants.[36] In Appendix H, we show that the restaurants' heterogeneous capacity constraints are the primary determinant of the magnitude of the incentive misalignment and discuss the underlying mechanisms.

This incentive misalignment has welfare implications: in column 2 of Table 5, we present the "partially variable" outcome, in which the platform does not allow these 55 restaurants with misaligned incentives to adopt variable pricing. If the platform has the ability to tailor the interface at the restaurant level, this outcome is a possible benchmark in which it satisfies all the parties' incentive compatibility. Comparing columns 2 and 3, we find a 7% overall welfare gain from variable pricing (from €240.3 to €257.2), which is 44% smaller than the infeasible first-best outcome in column 1. Notably, the welfare gain of €16.9 is comparable to the gain under the per-unit fee scheme (€18.2, comparing columns 4 and 5), indicating that misaligned incentives largely offset the increased efficacy of peak-load pricing that arises through eliminating double-marginalization.[37]

Overall, these results uncover a previously underexplored form of incentive misalignment between the platform and participating firms. When firms serve both the platform and off-platform markets using a shared capacity, the platform's policy can benefit firms through improved off-platform market outcomes, while causing negative externality to platform sales. Because the platform cannot extract such gains from off-platform markets, it may lose incentives to implement a welfare-improving policy. In particular, we show that this inefficiency can be magnified when a profit-sharing contract is used as the platform compensation scheme. Therefore, a conventionally efficient contract can still produce substantial inefficiency if adopted as a platform compensation scheme.

[36] This incentive misalignment is not a major issue under the observed per-unit-fee contract, 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 four restaurants, and the largest percentage decline is only 0.22%.

[37] Profit sharing still outperforms the per-unit-fee contract, because it achieves the outcome in column 2 of Table 5 as compared to one in column 4, producing the daily welfare gain of €52.1. However, compared to the true (infeasible) benefit of adopting the profit-sharing contract while still allowing variable pricing by all restaurants (i.e., €65.3, comparing column 1 and column 4), the welfare gain from adopting the profit-sharing contract is 20% smaller due to the incentive misalignment.

8 Conclusion

We study the welfare effects of peak-load pricing in a platform market with profit-maximizing firms. We find that time-varying prices—by day of the week and 30-minute time slots—increase total welfare by 9.7%, even when firms' pricing decisions are driven purely by profit maximization. These results show that capacity constraints can generate pro-social pricing outcomes among firms with market power, and that the efficiency gains from improved capacity allocation can outweigh incentives for surplus extraction. Our findings therefore contribute to ongoing policy debates on regulating firms' pricing power, particularly in markets where capacity constraints are prevalent.

We also document a novel source of incentive misalignment in platform markets that leads to underprovision of welfare-improving pricing technologies. Although peak-load pricing benefits firms and consumers, platforms may lack incentives to enable such pricing when surplus generated off the platform is non-contractible. The relevance of this mechanism extends well beyond restaurants. Many services—such as ride-hailing, hair and nail salons, fitness classes, massage studios, and coworking spaces—operate under similar conditions of fixed short-run capacity, fluctuating demand, a mix of reserved and walk-in access, and a reservation system owned by a third-party platform. In these settings, realizing the social benefits of peak-load pricing — improving product availability by smoothing demand over time without costly capacity investments — requires coordinated policy choices across parties, which may fail to emerge when gains are unevenly distributed.

Finally, while our small, lesser-known platform setup provides a strong identification strategy with "excluded" walk-in traffic, we acknowledge that this clean identification comes at the cost of generalizability of the quantitative findings: our platform is designed to match capacity-constrained restaurants with price-sensitive users, thus representing an environment in which peak-load pricing is most effective. Larger platforms (e.g., Opentable) face a more heterogeneous user base comparable to off-platform consumers, potentially limiting the scope of peak-load pricing while increasing firm incentives for price discrimination. Pricing strategies on large platforms also create substitution from off-platform markets, making the firm pricing incentives more complex (e.g., firms

may price discriminate between markets, a factor we held constant) and their identification less transparent. We argue that the core incentive misalignment persists in these environments: to fully harness the benefits of dynamic pricing technologies, platform contracts must account not only for on-platform outcomes but also for the broader off-platform externalities. We call for further studies on the platform incentives in these policy-relevant markets.

References

ANDERSON, E. T. AND I. SONG (2004): "Coordinating price reductions and coupon events," Journal of Marketing Research, 41, 411–422.

APARICIO, D., Z. METZMAN, AND R. RIGOBON (2024): "The pricing strategies of online grocery retailers," Quantitative Marketing and Economics, 22, 1–21.

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

ARMSTRONG, M. (2006): "Competition in two-sided markets," The RAND Journal of Economics, 37, 668–691.

ASSAD, S., R. CLARK, D. ERSHOV, AND L. XU (2024): "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market," Journal of Political Economy, 132, 723–771.

BAKOS, Y. AND H. HALABURDA (2020): "Platform Competition with Multihoming on Both Sides: Subsidize or Not?" Management Science, 66, 5599–5607.

BELLEFLAMME, P. AND M. PEITZ (2019): "Platform competition: Who benefits from multihoming?" International Journal of Industrial Organization, 64, 1–26, 2017 EARIE Proceedings.

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," .

BROWN, Z. Y. AND A. MACKAY (2023): "Competition in pricing algorithms," American Economic Journal: Microeconomics, 15, 109–156.

BUCHHOLZ, N., L. DOVAL, J. KASTL, F. MATEJKA, AND T. SALZ (2025): "Personalized pricing and the value of time: Evidence from auctioned cab rides," Econometrica, 93, 929–958.

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. (2025): "Who benefits from surge pricing?" Econometrica, 93, 1811–1854.

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," .

CHUNG, H. D., Y. M. ZHOU, AND S. ETHIRAJ (2024): "Platform Governance in the Presence of Within-Complementor Interdependencies: Evidence from the Rideshare Industry," Management Science, 70, 799–814.

EINAV, L., C. FARRONATO, AND J. LEVIN (2016): "Peer-to-peer markets," Annual Review of Economics, 8, 615–635.

EVANS, W. N. AND I. N. KESSIDES (1993): "Localized market power in the US airline industry," The Review of Economics and Statistics, 66–75.

FILIPPAS, A., S. JAGABATHULA, AND A. SUNDARARAJAN (2023): "The limits of centralized pricing in online marketplaces and the value of user control," Management Science, 69, 7202–7216.

GARCIA, D., J. TOLVANEN, AND A. K. WAGNER (2022): "Demand estimation using managerial responses to automated price recommendations," Management Science, 68, 7918–7939.

GARCIA-OSIPENKO, M., N. V. KUMINOFF, S. PERRY, AND N. VREUGDENHIL (2025): "Optimal Second-best Menu Design: Evidence from Residential Electricity Plans," Tech. rep., National Bureau of Economic Research.

GARDETE, P., D. M. SCHMITT, AND F. STAHL (2025): "Pricing and Consumption in Subscription Settings," SBE Working Paper Series 674.

GHILI, S. AND R. YOON (2024): "An Empirical Analysis of Optimal Nonlinear Pricing in Business-to-Business Markets," SSRN Working Paper 4367335.

HAGIU, A. AND J. WRIGHT (2024): "Marketplace leakage," Management Science, 70, 1529–1553.

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

HOBIJN, B., F. RAVENNA, AND A. TAMBALOTTI (2006): "Menu costs at work: restaurant prices and the introduction of the euro," The Quarterly Journal of Economics, 121, 1103–1131.

HORTAÇSU, A., O. R. NATAN, H. PARSLEY, T. SCHWIEG, AND K. R. WILLIAMS (2024): "Organizational structure and pricing: Evidence from a large us airline," The Quarterly Journal of Economics, 139, 1149–1199.

HUANG, Y. (2022): "Pricing frictions and platform remedies: the case of Airbnb," Ssrn working paper 3767103.

JOO, M., D. K. GAURI, AND K. C. WILBUR (2020): "Temporal distance and price responsiveness: Empirical investigation of the cruise industry," Management Science, 66, 5362–5388.

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

LAZAREV, J. (2013): "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets," Working paper.

LI, H., Q. SHEN, AND Y. BART (2018): "Local market characteristics and online-to-offline commerce: An empirical analysis of Groupon," Management Science, 64, 1860–1878.

LI, H. AND F. ZHU (2021): "Information Transparency, Multihoming, and Platform Competition: A Natural Experiment in the Daily Deals Market," Management Science, 67, 4384–4407.

MACKAY, A., D. SVARTBÄCK, AND A. G. EKHOLM (2023): "Dynamic Pricing, Intertemporal Spillovers, and Efficiency," Harvard business school strategy unit working paper.

MASKIN, E. AND J. RILEY (1984): "Monopoly with Incomplete Information," The RAND Journal of Economics, 15, 171–196.

MING, L., T. TUNCA, Y. XU, AND W. ZHU (2025): "Market Formation, Pricing, and Value Generation in Ride-Hailing Services," Manufacturing & Service Operations Management, 27, 1551–1570.

MUSSA, M. AND S. ROSEN (1978): "Monopoly and Product Quality," Journal of Economic Theory, 18, 301–317.

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

ROBINSON, J. (1969): The economics of imperfect competition, Springer.

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.

SULLIVAN, M. (2025): "Fee Optimality in a Two-Sided Market," Working paper.

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.

TALLURI, K. AND G. VAN RYZIN (2004): "Revenue Management Under a General Discrete Choice Model of Consumer Behavior," Management Science, 50, 15–33.

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

VARIAN, H. R. (1989): "Price discrimination," Handbook of industrial organization, 1, 597–654.

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.

XIE, Y. AND H. ZHU (2023): "Platform leakage: Incentive conflicts in two-sided markets," Working paper.

Web Appendix to "Peak-Load Pricing in Platform Markets"

A User Experience

To illustrate the look and feel of the platform from the customer's perspective, Figure A1 shows a screenshot of the mobile interface. Customers browse restaurants and observe available reservation times together with the corresponding discounts offered at each time slot. For a given restaurant, discounts typically vary across times within the same day, with higher discounts displayed for off-peak hours and lower or zero discounts for peak dining times. Indeed, we observe the possibility to make a reservation at a 0% discount. Discounts are transparent and fixed at the time of reservation, and customers do not observe or negotiate prices beyond the posted discount schedule.

*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). We are grateful for the insightful comments and suggestions by Juan Camilo Castillo, Sarah Moshary, and seminar and conference participants at Cornell-Rochester Marketing Workshop, IIOC 2025, SICS 2025, Marketing Science 2025, WEAI 2025. An alternative disclosure plan was submitted for this article. The raw data are not available; however, all the codes and disguised data are available in a repository to be determined.

These materials have been supplied by the authors to aid in the understanding of their paper. The AMA is sharing these materials at the request of the authors.

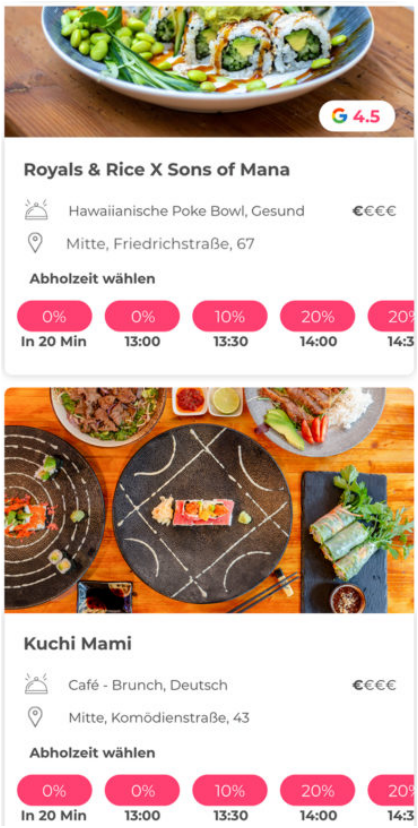


Figure A1: Phone Screenshot

B Further Details on Data

B.1 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 B1. The traffic measure represents the total volume of foot traffic at each hour of day, constructed from Google devices' location data.[1] 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

[1]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.

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, because historical data is not available on Google Map. We assume that the traffic pattern remains the same between the two data periods.

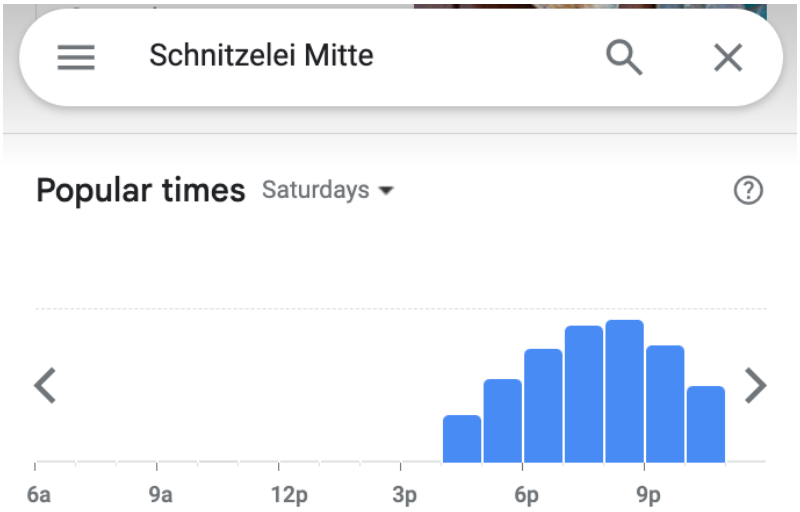


Figure B1: Histogram of Traffic on Google Maps

In various parts of the estimation procedure, we also use the market-level aggregate measure of traffic volume, which we denote by $\bar{g}_{\tau t}$. For instance, we use the empirical distribution of $\bar{g}_{\tau t}$ as an input to estimate each consumer's most preferred dining time, $F_{it}(\tau_{it})$. We also use the ratio $g_{\tau_j t}/\bar{g}_{\tau_j t}$ (i.e., the traffic volume of a given restaurant relative to the aggregate traffic volume at time τ_j) 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 is our proxy for restaurant capacity. By weighting $g_{\tau_j t}$ by a measure of restaurant capacity, we 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.2 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-10:00 pm window (5,103 obs.)
- Restaurants with no walk-in traffic measure available (3,107 obs.)

These criteria leave approximately 12,995 reservations by 6,217 consumers across 101 restaurants. However, 59% of these consumers only make one reservation throughout the sample period, and some restaurants receive very few reservations. Because we 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 6,188 reservations made by 1,357 customers across 80 restaurants during 189 days between September 1st, 2019 and March 8th, 2020.

B.3 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 D1, 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 17.3 Euro per person, with the standard deviation of 9.2. 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 to aggregate the restaurant traffic up to the market level. On average, restaurants offer 18 tables with the standard deviation of 8.7. Although in principle the number of available tables can be a restaurant's strategic decision, these restrictions are never binding (i.e., the reservations are strictly less than the offered tables) and thus we abstract this decision away from the model.[2]

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 D1: Summary statistics

	Mean	Std. Dev.	Min	Max
Regular Price	17.314	9.188	8.5	60
Review Rating	4.343	.243	3.6	4.8
Maximum tables offered	18.038	8.653	4	50

Note: Each unit of observation is a restaurant.

C Further Reduced-form Analyses

C.1 Price Schedule across Days

In Figure D1, 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

[2]Offering tables at 0% discount is strictly better for the restaurant than not offering tables, suggesting one reason why the number of tables is not utilized as a strategic tool in the focal environment.

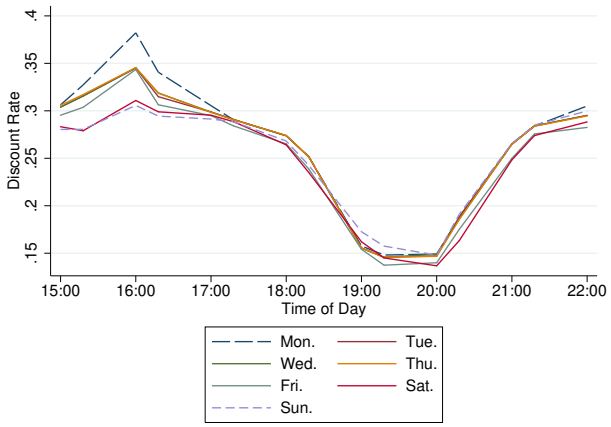


Figure D1: Price Schedule across Days

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

Sunday with the smallest variability. We find similar price schedules across days. The largest price gap 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.

C.2 Robustness of the Identification of the Price Coefficient

Our main identification strategy for the price coefficient exploits local variation in prices around the date of price adjustment. Because restaurants keep the same price schedule for weeks and months, we argue that these prices are not fine-tuned at each calendar-day level. Thus, when the restaurant does adjust the schedule, the adjustment is likely a response to a longer-term shift in the market conditions and thus is not correlated with the demand shock of that particular day. We study robustness of this strategy against 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 than our main specification. Although Expression (1) mainly exploits the aforementioned variations by exclusively focusing on variations across calendar days within restaurant - day of week - time of

day tuple, the regression still includes other sources of price variations: for instance, it does not exclude the variation in prices between two calendar days that are far apart, as long as they belong to the same tuple. Even when the local price adjustments do not reflect the demand shock on that focal day, price differences between two distant dates can well reflect long-term difference in demand, e.g., seasonality. To address this concern, we now exclusively use observations within two weeks of the date of price change, purely exploiting the variation in prices right around the adjustments. 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. We present the results in Table C1. 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. The cross-elasticity terms are extremely collinear and thus are dropped. We find the own elasticity of -11.1, which is approximately 9% larger than that of the main specification.

The other threat to identification is the possibility that restaurants do set prices at the calendar-day level, so that a 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 exploit a different source of price variation. Specifically, we include a set of fixed effects at restaurant - date - hour of day level. Because the observations are restaurant - date - 30 minute level, these 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 specification thus exploits any 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 own elasticity of 12.05, or approximately 18% higher than our main results. The difference is not statistically significant, partly due to the noisier estimate that arises from the smaller degree of freedom. The intra-day cross-elasticity term also remains within the same ballpark of the main specification. We find a larger across-day elasticity, but the estimate is quite noisy because of high collinearity between the regressors. Given that all three specifications produce similar own price elasticities despite the substantial differences in the source of identification, we argue that the true price coefficient is also located within the same ballpark of these estimates, which we also recover in the structural analysis (between 9 to 10.7: see panel (a) of Figure 3).

Table C1: Various Sources of Identification

	$\log(R_{j,t,\tau})$		
$\log(P_{j,t,\tau})$	-10.183 (1.017)	-11.056 (1.135)	-12.371 (2.129)
$\log(\bar{P}_{j,t\pm1,\tau})$	2.467 (.838)	-	2.050 (1.803)
$\log(\bar{P}_{j,t,\tau\pm1})$	-.023 (.980)	-	1.827 (2.329)
j-dow(t)-τ FE	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
j-t-hour(τ) FE	No	No	Yes
$\pm$ 2 wks of price change	No	Yes	No
Observations	117,894	2,875	103,938
Degree of Freedom	112,437	1,999	49,635

Note: 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. j, t and τ correspond to restaurant, date and time, respectively.

D 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 that each party includes N guests (common across platform and walk-in diners), then the objective function is given as follows:

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

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.

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

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 $\lambda_{\tau_j t}$. We parameterize $\lambda_{\tau_j t}$ 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, and 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:

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

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 $\rho_{r\tau}$ terms 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).[3] The reason we include the time-varying discount factors 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

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

$\bar{n}_{\tau_j+\tau',t}$ are both nonnegative, based on our microfoundation that these numbers correspond to current and future consumer arrivals.[4]

F The Use of Popularity Measure to Account for Price Endogeneity

We show in Appendix C.2 that our identification strategy produces a robust estimate for the reduced-form price coefficient. We next examine whether our structural specification can adequately mimic this reduced-form identification strategy. Recall that the set of interactive fixed effects we used for the reduced-form estimate is not included in the structural model due to computational complexity. We instead opt to use a variety of continuous variables to mimic the 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 the restaurant - date - time fixed effect. Whether or not a single-dimensional continuous variable, included as a linear function in the utility, can adequately replace the slew of fixed effects is an obvious concern. In this section, we include the restaurant popularity measure up to its third-order term in the utility rather than as a linear function, and show the estimates remain stable.

Table F1 shows the estimates of price coefficients and costs for switching time slots under this alternative 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. 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 F1, we show the within-day evolution in own- (left figure) and cross-elasticity (right figure) when we include the third-degree polynomial

[4]For 12 restaurants in the data, the specification above does not produce sufficient statistical power. We assume $\rho_{r\tau} = 1$ across all times for these restaurants.

of the popularity term. Compared to the results in Figure 3, we find that both the own- and cross-elasticities are within 5% of our main specification. These results show that the linear specification we employ in the main text adequately captures the within-restaurant, across-time variation in demand.

	Popularity in Cube			
	Segment 1		Segment 2	
Price (τ_{it} ∈[3pm, 4pm])	-.338	(.012)	-.776	(.016)
Price (τ_{it} ∈[4:30pm, 5:30pm])	-.475	(.022)	-.823	(.033)
Price (τ_{it} ∈[6pm, 7pm])	-.389	(.013)	-.685	(.014)
Price (τ_{it} ∈[7:30pm, 8:30pm])	-.355	(.012)	-.654	(.016)
Price (τ_{it} ∈[9:00pm, 10:00pm])	-.320	(.013)	-.647	(.014)
Switching cost (τ_{it} ∈[3pm, 4pm])	-.013	(.005)	-.064	(.030)
Switching cost (τ_{it} ∈[4:30pm, 5:30pm])	-.000	(.002)	-1.105	(.885)
Switching cost (τ_{it} ∈[6pm, 7pm])	-.221	(.047)	-1.207	(.271)
Switching cost (τ_{it} ∈[7:30pm, 8:30pm])	-.120	(.022)	-1.809	(.600)
Switching cost (τ_{it} ∈[9:00pm, 10:00pm])	-.000	(.001)	-73.29	(99.99)

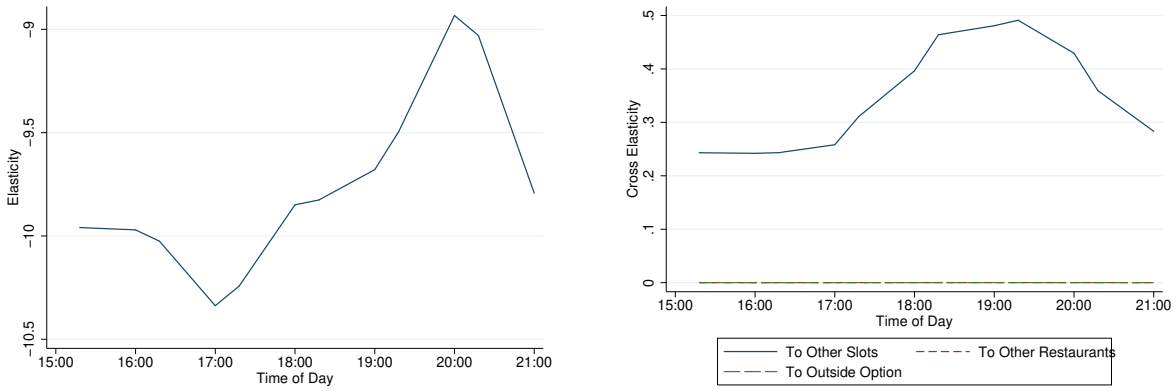
Note: Standard errors are reported in parentheses. "Early", "Peak" and "Late" stand for the most preferred dining time between 3pm and 4:30pm, 5pm and 6:30pm and 7pm and 9:30pm, respectively. The specification also includes fixed effects for year-month pairs, for weekend days, for restaurant locations (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 F1: Demand Parameter Estimates

G Definition of Welfare Metrics

In this section, we formally define the welfare metrics we use 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:



(a) Own Elasticity

(b) Cross Elasticity

Figure F1: 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.

Total welfare generated by the platform on day t is given as follows:

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

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 across restaurants 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 $\sum_{j \in J_{rt}} (q_{\tau_j t}(p_{jt} - c_r - c_r^p)) / \sum_{j \in J_{rt}} q_{\tau_j t}$. The numerator corresponds to the restaurant profit from platform users, excluding walk-in diners. We present its average across restaurants in Table 4.

- Platform Surplus per Reservation:

  Platform surplus per reservation is given by the sum of platform user's surplus per reservation, restaurant profits from platform users per reservation and the platform profit per reservation. The joint profit between the restaurant and the platform (per reservation per restaurant) is given by $\sum_{j \in J_{rt}} (q_{\tau_j t}(p_{jt} - c_r)) / \sum_{j \in J_{rt}} q_{\tau_j t}$. We take its average across restaurants.

  The user surplus per reservation corresponds to the expected utility generated from the reservation, conditional on the reservation made (i.e., the utility from reservation exceeds all the other alternatives). More formally, for each user with type β_i and τ_{it}, the consumer surplus from making a reservation at restaurant-time pair j is given by $\mathbb{E}[u_{ijt} | u_{ijt} > u_{ij't}]$. The closed-form expression exists, and the exact formula is available in Arcidiacono and Miller (2011), Lemma 3. To compute its market average, we first take its expectation over all consumer types. We thus have the average consumer surplus from each option j. We next take its weighted average across j, in which the weights are given by the frequency of reservation, $q_{\tau_j t}$. These operations produces a measure of the platform users' average surplus per reservation.

- 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 reser-

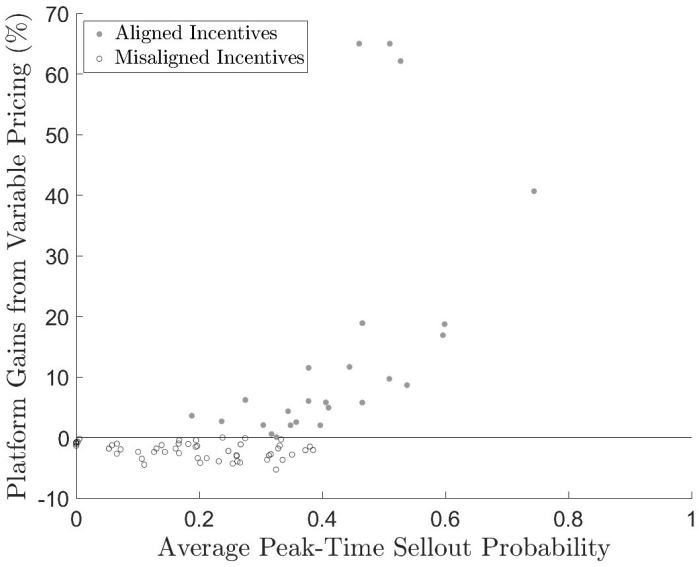


Figure H1: Platform's Gains Plotted against Restaurant Capacity Constraints

Note: Each dot corresponds to a restaurant, averaged across dates. The vertical axis is truncated at 65% for exposition. The two dots located exactly at 65% on the figure correspond to 141% and 234%, respectively.

vations to create a per-reservation metric.

- Lost Profit due to Displacement per Reservation:

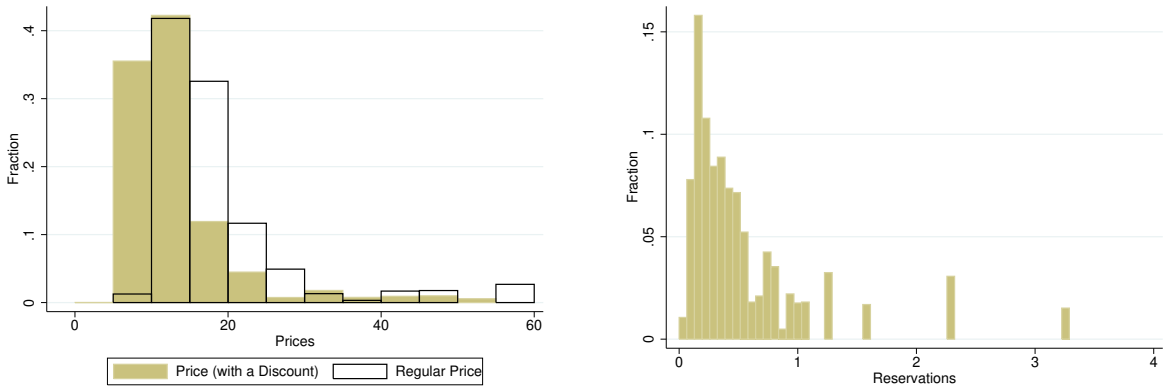
  Lost profit due to demand displacement for each restaurant in each day is given by $\sum_{j \in J_{rt}} (\mathbb{E}[\min\{n_{\tau_j t}, L_r\}] - \mathbb{E}[\min\{n_{\tau_j t}, L_r - q_{\tau_j t}\}])(p_r^{regular} - c_r)$. We divide it by the number of reservations to create a per-reservation metric.

H Mechanisms of Incentive Misalignment

We find whether or not the platform benefits from allowing a restaurant's variable pricing is primarily determined by the restaurant's peak-time capacity constraints. In Figure H1, we present the restaurant-level scatterplot, in which the vertical axis corresponds to the platform's gains from allowing the restaurant's variable pricing under the profit-sharing scheme (the metric in Figure 9), and the horizontal axis corresponds to the average peak-time (7 pm - 8 pm) sellout probability of the restaurant. We find a positive correlation, indicating that the platform incentive is more aligned

with restaurants that face a strict peak-time capacity constraint. The mechanism is as follows. When variable pricing is not allowed, platform users choose peak times, exacerbating congestion. Restaurants with tight peak constraints thus set the 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 baseline sales is close to zero), by offering variable pricing and generating incremental off-peak sales without cannibalizing the restaurants' peak-time sales to walk-in diners. On the other hand, for restaurants that have enough capacity to accommodate platform users during the peak times under the uniform-pricing regime, the mechanism in the main text dominates, that is, peak-load pricing lowers these users' average transaction prices, reducing the platform profit. Note, however, that this lack of misaligned incentives among restaurants with tight constraints is likely specific to our research context, in which platform users are systematically more price sensitive. If the platform users were comparable to walk-in diners, the platform sales would not "shut down" even when the restaurant uniformly offers close-to-zero discounts on the platform, and thus the mechanism in the main text would likely be the first order among these restaurants. Thus, the overall magnitude of incentive misalignment can be larger in a more representative platform.

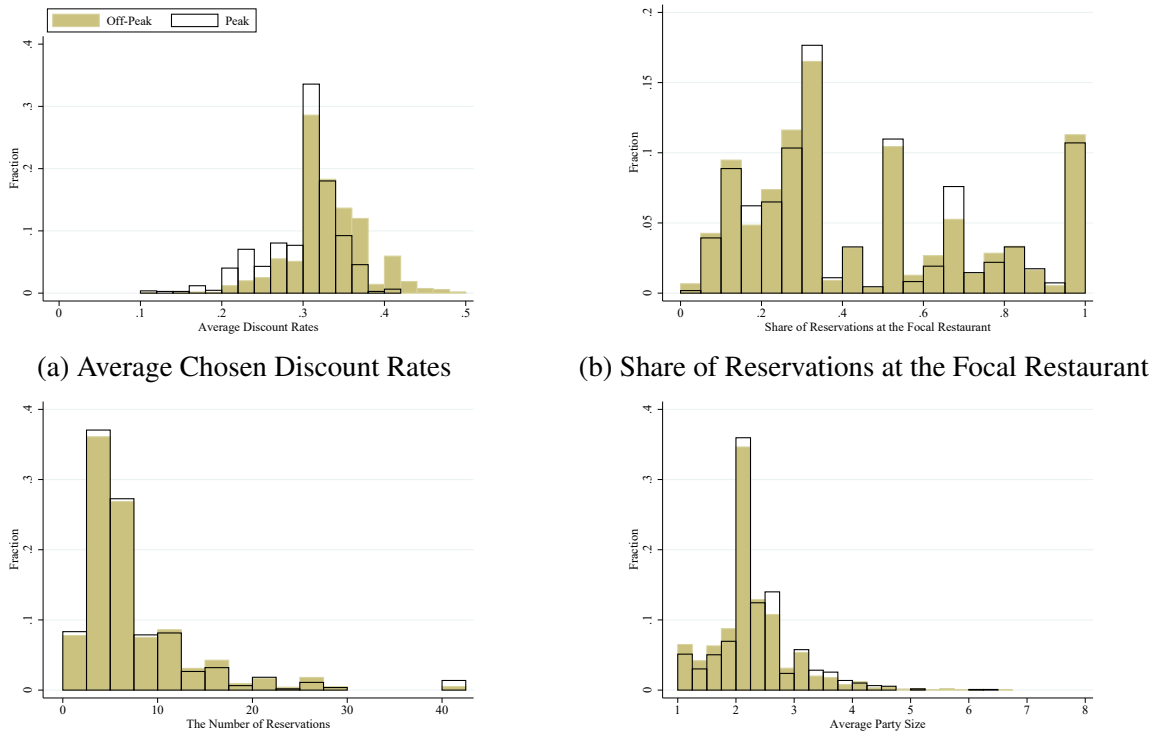
I Additional Figures and Tables



(a) Prices (b) Reservations

Figure I1: 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. The regular price corresponds to the prices that an walk-in guest will pay without the online discount. Panel (b) is the histogram of average daily reservations of restaurants.



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

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

Figure I2: Distribution of Average Characteristics Between Consumers Choosing Peak and Off-peak Times

Note: This figure 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.

		Segment 1		Segment 2	
Time					
	Oct 2019	-.085	(.054)	-.001	(.101)
	Nov 2019	-.132	(.053)	-.156	(.102)
	Dec 2019	-.508	(.057)	-.519	(.108)
	Jan 2020	-.341	(.055)	-.181	(.101)
	Feb 2020	-.447	(.056)	-.360	(.103)
	Mar 2020	-.549	(.094)	-.481	(.163)
	Weekend	-.227	(.033)	-.206	(.056)
District					
	Friedrichshain	-1.101	(.344)	-5.214	(.442)
	Kreuzberg	.968	(.290)	-.019	(.309)
	Mitte	.186	(.299)	-3.244	(.321)
	Neukölln	-1.549	(.465)	-4.441	(.728)
	Prenzlauer Berg	.145	(.312)	-4.684	(.421)
	Schöneberg	.153	(.317)	-.438	(.435)
Cuisine					
	Asian	1.711	(.223)	4.795	(.390)
	European	.700	(.216)	1.611	(.320)
	Middle Eastern	1.632	(.238)	3.476	(.438)
	Other	.567	(.236)	2.818	(.720)
Local Restaurant Density (within 1km)					
	24-43 restaurants	-.419	(.063)	-1.444	(.146)
	44-59 restaurants	.285	(.058)	.286	(.104)
	$>$59 restaurants	.099	(.065)	-.708	(.150)

Note: Standard errors are reported in parentheses. The baseline observation is Sep 2019 in Charlottenburg district with American cuisine and $<= 23$ restaurants around it. In addition to these controls, we include interactions between restaurant locations and cuisines whose coefficients are common across segments. Their estimates are available upon request.

Table I1: Demand Parameter Estimates

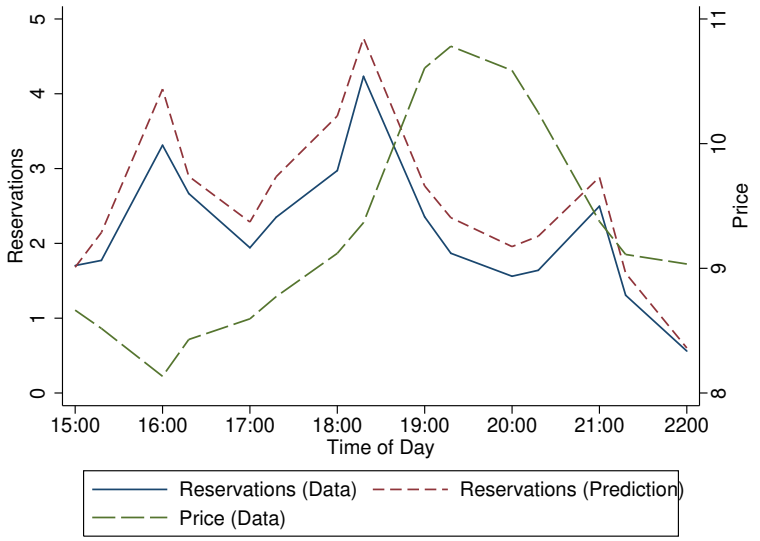


Figure I3: Demand Model Fit

Note: Reservations represent the total bookings on the platform (across all restaurant) at each slot, and Price corresponds to the average price across restaurants.

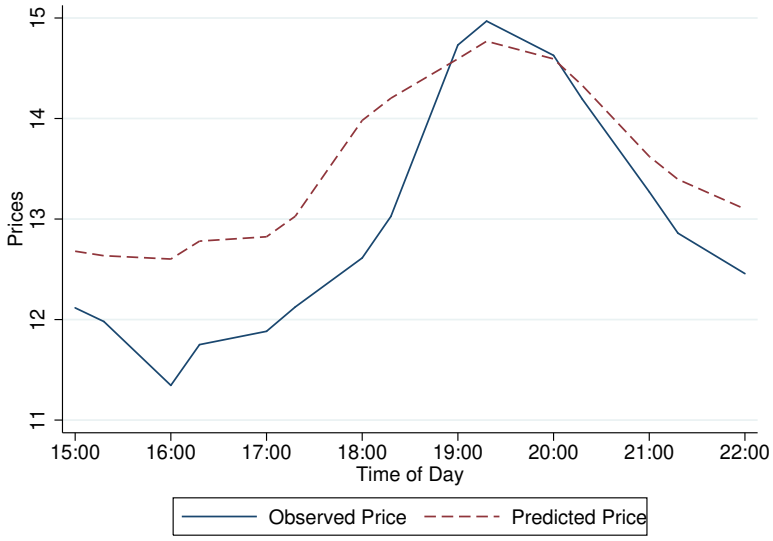


Figure I4: Model Fit for Prices

Note: The solid line represents the observed average prices (across restaurants and dates) 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.