Pricing Frictions and Platform Remedies: The Case of Airbnb

Yufeng Huang*

October 3, 2021

Click here for the latest version

Abstract

Pricing in a complex environment is difficult for individual sellers. Whereas the platform tries to aid seller pricing, its different objectives might steer seller behavior towards the platform's goal. This paper empirically studies pricing frictions on Airbnb and explores the equilibrium consequence of different platform designs. I first show that pricing frictions are prevalent. Then, leveraging natural variation in the platform's interface design, I demonstrate that sellers' price-setting costs and cognitive constraints are plausible drivers of the frictions. I then estimate a structural equilibrium model and find that pricing frictions lead to a 14% consumer welfare loss and a 0–15% seller-profit loss. Finally, I ask: How to ameliorate these frictions? The platform's revenue-maximizing algorithm does not lead to market-clearing prices because it fails to internalize sellers' high opportunity costs of time. However, a simple platform design, where the platform sets price variation but gives sellers the final decision right to determine the price levels, will eliminate almost all frictions.

Keywords: Pricing frictions, Algorithmic pricing, Market design, Airbnb

*Univerity of Rochester Simon Business School. Email: yufeng.huang@simon.rochester.edu. This paper benefits from numerous comments and suggestions from Kristina Brecko, Zach Brown, Hana Choi, Andrey Fradkin, Tong Guo, Avery Haviv, Przemyslaw Jeziorski, Ilya Morozov, Tobias Klein, Tesary Lin, Bowen Luo, Olivia Natan, Adam Smith, Avner Strulov-Shlain, Takeaki Sunada, Kosuke Uetake, Caio Waisman, Yan Xu, and seminar and conference participants at AI and Machine Learning Conference, China Virtual IO Seminar, European Quantitative Marketing Seminar, IIOC, Marketing Science Conference, Northwestern University, Seoul National University, University of Pennsylvania, University of Rochester, and University of Toronto. I thank Miao Xi and Chen Cao for their excellent research assistance.

1 Introduction

> "[F]or many hosts, finding the right price for their space can be both time-consuming and challenging... Even many experienced hosts told us that they find pricing difficult, especially as seasons change, special events come to town, and more listings emerge in their neighborhood."
>
> – Janna Bray, Head of Research for Airbnb.[1]

There is rising attention to online platforms' adoption of new pricing technologies and how these technologies might steer the behavior of market participants. An important rationale for adopting pricing technologies is that sellers face significant frictions in setting prices—in that their prices fail to react to differences or changes in market conditions. Pricing frictions might be particularly pronounced for amateur sellers, who lack the managerial capabilities compared to professionals (Goldfarb and Xiao, 2011; Li et al., 2016). Further, whereas the platform has an incentive to assist seller-pricing, its objective does not necessarily align with sellers'. As a result of the misaligned incentives, platform-provided pricing technologies (pricing interfaces and algorithms) might distort prices away from sellers' (or consumers') desired outcomes and towards that of the platform. What is the equilibrium consequence of sellers' pricing frictions under the platform's design of pricing interfaces and algorithms? Can alternative market designs improve market outcomes?

In this paper, I study heterogeneous Airbnb sellers' pricing decisions under the existing market environment created by the platform, and I explore alternative market designs using counterfactual experiments. Airbnb sellers face a difficult pricing problem similar to airline or hotel pricing. On the one hand, each product (a night of stay) is capacity-constrained and perishable, which requires dynamic pricing. On the other hand, different nights face drastically different demand (e.g., tourist-season nights sell out quickly), which calls for flexible prices that respond to market conditions. This paper uses detailed data to document the extent of seller-pricing frictions, and exploits natural

[1] Source: https://airbnb.design/smart-pricing-how-we-used-host-feedback-to-build-personalized-tools/. Extracted in April 2021.

variation in the pricing-interface design to separate different mechanisms of the frictions. Then, estimating a tailored structural model, I quantify the consumer, seller, and platform loss due to these frictions. Lastly, I demonstrate that a simple market design, which balances the platform's and sellers' information sets and incentives, would improve the payoff of all parties.

I begin by presenting the distribution of pricing frictions and the extent to which these patterns reflect the platform's interface and algorithm design. The platform provides a feature-limited pricing interface and a revenue-maximizing pricing algorithm (Ye et al., 2018), and the latter is not well received among sellers. Consistent with these anecdotes, observed prices fall into two extremes: Most sellers use inflexible pricing strategies (consistent with the pricing interface), and a small set of sellers use highly flexible, algorithm-like prices.

Next, I investigate two plausible explanations for these pricing frictions. One explanation is that the platform's limited-feature interface creates the frictions, which I refer to as (interface-induced) price-adjustment costs. A different explanation is that cognitively-constrained sellers find the pricing problem inherently difficult, and as a result, they opt for simpler, heuristic prices. I exploit the platform's pricing-interface change to separate the two explanations. In 2019, Airbnb introduced "last-minute discounts" that allow sellers to set one *automated* price adjustment for each night. I find that some sellers respond to this feature by setting a sharp last-minute discount. Yet, most sellers do not adopt dynamic pricing, despite it being automated—consistent with the hypothesis that sellers face cognitive constraints and dynamic pricing is "too difficult" for them. This finding suggests that price-adjustment costs is *a* mechanism, but not the only mechanism, behind the pricing frictions.

To quantify pricing frictions and examine how these frictions impact equilibrium outcomes, I construct and estimate a structural model that characterizes consumer arrival and demand, and sellers' pricing and participation decisions. The demand model builds on the existing literature on capacity-constrained products (notably Pan, 2019) but extends this literature on two novel aspects. The first extension is to account for a large set of differentiated products, with arbitrary time-varying demand that involve a large set of fixed effects. I adopt a likelihood-based fixed-point algorithm (in the spirit of Chintagunta and Dubé, 2005), which makes computation scalable on large datasets. The second extension is to propose a new "uniform-pricing instrument" to address price endogeneity. The idea is that adjusting multiple nights' prices together is less costly than doing so night-by-night. The costly price adjustment makes it so that prices of far-apart nights (with uncorrelated demand shocks) co-move with each other.

Estimating the demand model on the San Francisco Airbnb market, I find rich heterogeneity on consumer arrival rates and price sensitivities. For example, the segmentation structure implies that consumers who arrive early are more price-sensitive compared to those who arrive later. National holidays see more early-arrival consumers, whereas March, the peak tourist season for San Francisco, sees a surge of last-minute bookings. Optimal pricing in this market *should* factor in the heterogeneity in consumer traffic across nights and over time. Besides, the average price elasticity aligns well with field experimental evidence presented by Jeziorski and Michelidaki (2019).

I then construct and estimate a supply-side model, which nests the two types of pricing frictions. Each seller can be of two types: The first type solves for optimal finite-horizon dynamic prices while facing price-adjustment costs (*a la* Calvo, 1983). The second type is cognitively constrained, sets non-dynamic prices (i.e., maximizing expected profit by setting prices that do not vary with time-to-check-in), and faces additional price-setting costs. This model exploits how different sellers respond to the interface change to separate the importance of cognitive constraints from price-adjustment costs. Finally, to characterize seller heterogeneity as flexibly as possible, I separately estimate, by granularly defined market segments, sellers' marginal costs, price-adjustment costs, and degrees of cognitive constraints (using a method proposed by Bonhomme et al., 2019).

Supply-side estimates reveal rich heterogeneity in sellers' pricing frictions and their underlying drivers. First, only about 15% of listings set frictionless prices—they price-in demand shifters for different nights and are able to update prices (for a given night) as time changes and new market information is revealed. Given the platform's price-setting interface, manually setting such flexible prices is unlikely to be feasible. As such, one can *interpret* these listings' prices as being controlled by some algorithm. Second, about 35% of listings *can* set dynamic prices but find it too costly to adjust them often (until the 2019 interface change makes discounting the price easier). Finally, about half the sellers are cognitively constrained and do not set dynamic prices. Many of these sellers also face high price-setting costs that result in inflexible prices across nights. As a result, prices of these listings often do not account for night-specific demand shocks or market conditions.

Also, many listings bear considerable marginal costs. These costs likely reflect the sellers' opportunity cost of time to host guests per night. The median marginal cost is $37 per night, and a quarter of listings have a cost above $64. These costs are nontrivial relative to the median price at $150 (and the 75 percentile at $235). Because the platform's payoff depends on a fixed share of sellers' total booking revenue, it likely wants seller-*revenue*-maximizing prices (throughout the paper, I will assume that the platform will not help sellers collude). In contrast, sellers' *profit*-maximzing prices are higher, so as to internalize their opportunity costs of time. The incentive incompatibility implies that platform-controlled prices might not be at the seller-optimal level.

How important are the pricing frictions? I compare the observed market outcome with the counterfactual where all frictions are absent (which I refer to as the "first best"). I find that consumers and some sellers would have gained significant surplus compared to the baseline. Consumer surplus would have been 14% higher without all frictions, and a quarter of sellers would have gained at least 6% profit. In contrast, the platform has limited incentives to ameliorate the pricing frictions because it (and some sellers) gains relatively little from doing so.

Given the sizable potential gains from ameliorating these frictions, I then ask: What *realistic* remedies can the platform provide to improve market outcomes? I consider three different approaches. The first is the platform's revenue-maximizing pricing algorithm, "Smart Pricing" (Ye et al., 2018). I simulate the market outcome if the platform enforces a seller-revenue-maximizing algorithm and fully controls pricing (but ruling out collusion). Prices are unsurprisingly lower because the algorithm does not internalize seller costs. Consequently, whereas consumers and the platform gain from lower prices, most sellers lose and prefer setting prices themselves (consistent with sellers' resistance to Smart Pricing).

The second approach is to improve the platform's price-setting interface. I simulate an "ideal" interface where the platform can eliminate all seller price-adjustment costs (but cannot change sellers' cognitive constraints). Although this scenario is overly optimistic and assumes that the platform can *directly* change some seller frictions, I find that consumer, seller, and platform payoffs are only modestly improved. This is because significant part of the friction is due to sellers' cognitive constraints. As such, complicating the pricing interface is unlikely a fruitful direction.

The third and final approach is to redesign the platform's pricing to (1) let the platform determine how prices vary, but (2) let the seller determine what the base price should be. This market design leverages the platform's informational (data to estimate demand) and technological (algorithms to automate pricing) advantages over individual sellers. Nevertheless, this design limits the platform's misaligned incentive by giving sellers the ability to set their own price levels. With this market design, sellers now make simpler pricing decisions by only setting one price, and the platform bears the main burden of deciding how prices should vary across markets and over time. Nevertheless, I demonstrate that prices are close to the first best, and so are consumer, seller, and platform surplus. Ameliorating the pricing frictions is feasible in practice.

Related literature. The paper's primary contribution is to the recent stream of literature on pricing frictions. Cho and Rust (2010), Pan (2019), Leisten (2020), and Hortaçsu et al. (2021) document the lack of price variation in capacity-constrained industries (rental cars, Airbnb, and airline), and attribute the frictions to managerial mistakes (Cho and Rust), menu costs (Pan), and organizational frictions (Leisten and Hortaçsu et al.). DellaVigna and Gentzkow (2019), Hitsch et al. (2019), Arcidiacono et al. (2020), Strulov-Shlain (2019), and Huang et al. (2020) document grocery prices' (lack of) response to demand features. Bloom and Van Reenen (2010), Bloom et al. (2019), Goldfarb and Xiao (2011), and Hortaçsu et al. (2019) study firm heterogeneity and show that firm size and manager's education play a role in the firm's decision quality. This paper contributes to this literature by separating different mechanisms driving the pricing frictions and by exploring alternative platform designs where such mechanisms (as well as the platform's and sellers' incentives) play an important role. Closely related to this paper, Pan (2019) presents evidence of sticky Airbnb prices and estimates listings' implied price-adjustment costs. This paper builts on Pan (2019) and further estimates sellers' cognitive constraints and opportunity costs of time (marginal costs), both of which impact platform design in fundamental ways. Also, in a contemporaneous paper, Filippas et al. (2021) present experimental evidence from a decentralized car-rental platform's transition from seller pricing to centralized pricing. They show that enforcing centralized pricing will decrease prices, increase revenue, and lead to a high seller exit (or non-participation) rate. This paper does not observe such a regime shift, but instead, seeks to understand market participants' objectives and constraints, and use this understanding to explore different platform designs. One such counterfactual is to move all sellers to centralized pricing, where my results are well in line with Filippas et al..

This paper's secondary contribution is to extend the previous literature and present a framework to study capacity-constrained markets with pricing frictions. On the demand side, the framework follows Williams (2021) and Pan (2019) but includes a fixed-point algorithm (in the nature of Berry et al., 1995) to accommodate a large set of fixed effects. The combination of likelihood-based estimation with a nested-fixed-point algorithm goes back to Goolsbee and Petrin (2004); Chintagunta and Dubé (2005) and is recently applied in Tuchman (2019). The demand estimation also uses a new "uniform pricing" instrument. Both the algorithm and instrument has wider applicability to other markets. On the supply side, the model extends Pan (2019) (who characterizes revenue-maximizing pricing with menu costs) to capture a mixture of rational and cognitively-constrained seller decisions. To my knowledge, the marketing and industrial organization literature has a limited set of tools to systematically characterize firms' heterogeneous, bounded-rational behavior. Although the specific model is tailored to the empirical setting, the idea of characterizing the seller mixture (but going beyond cognitive hierarchy (Goldfarb and Xiao, 2011; Hortaçsu et al., 2019), which solely focuses on firm beliefs about opponents' strategies) might be generalizable.

Finally, the paper is broadly related to the recent discussions on algorithmic pricing, with a particular focus on pricing algorithms as a technology (related to Brown and MacKay, 2019).[2]

[2]This paper is also broadly related to Dubé and Misra (2017) and Jin and Sun (2019), who show that a targeted-pricing model and provision of market information can help sellers (although Jin and Sun, 2019 also show that seller

The paper is also broadly related to the vast literature on Airbnb and sharing platforms.[3,4]

The remainder of the paper is organized as follows. Section 2 discusses the background and introduces the data. Section 3 presents the main evidence on pricing behavior and heterogeneity. Section 4 presents the structural model (with some model details discussed in Appendix D). Section 5 discusses estimation results. Section 6 discusses counterfactual simulation results and implications. Section 7 concludes the paper.

2 Background and data

General context. Airbnb is the dominant platform for the short-term rental market ("short-term" refers to the lease term shorter than 30 days). Sellers ("hosts") register their listings on the platform, upload pictures, enter descriptions, and set prices. Consumers ("guests") arrive through the platform's front page and search by the destination city and check-in/check-out dates.

I assume that booking is instant and cancellation is negligible. Most listings either support instant booking or respond quickly to requests.[5] Airbnb does not have a free cancellation policy during the sample period I study, and many listings have steep penalty for cancellations.[6]

pricing does not react to information). In addition, the pricing patterns across sellers (with some algorithm-like behavior) is related to Chen et al. (2016), who study Amazon sellers' algorithm-like pricing behavior.

[3]On pricing, Pavlov and Berman (2019) present a theoretical model to highlight the tradeoff between platform-centralized pricing (which internalizes the cannibalization effect between sellers) and pricing-in sellers' quality differences. My paper also speaks to the role of pricing algorithms but highlights a different tradeoff. I do not entertain the possibility that the pricing algorithm directly acts as monopoly pricing (which might violate anti-trust laws, although it is certainly possible in reality).

[4]On different topics, Zervas et al. (2017) estimates the impact of Airbnb listings on hotel revenue and demonstrates a sizable substitution effect, primarily on low-end hotels. Farronato and Fradkin (2018) and Li and Srinivasan (2019) structurally characterize Airbnb and hotels' demand and supply, emphasizing that Airbnb hosts' flexibility plays a crucial role because hotels are capacity-constrained. Barron et al. (2020) and Garcia-López et al. (2020) examine the effect of Airbnb listings on rental and housing prices. Fradkin et al. (2018); Zervas et al. (2020); Proserpio et al. (2018); Zhang et al. (2019) study reputation, reciprocity, and image quality on Airbnb.

[5]Guests can book the preferred listing if the listing supports instant booking (28% listings support instant booking in my sample). If not, guests can inquire about the listing, and 98% sellers respond to the inquiry within a day (60% sellers respond to requests within an hour).

[6]During my sample period, Airbnb's cancellation policy is typically much stricter than hotels. 25% listings employ a "flexible" cancellation policy, allowing cancellation 14 days before check-in (or 48 hours after booking if booked in less than 14 days). The 14% service fee is not refundable (see, e.g., https://www.bnbspecialist.com/airbnb-service-fee-when-refundable/, accessed in September 2021). Beyond the "flexible" cancellation policy, 32% listings employ a "moderate" policy and 43% a "strict" policy, further tightening the window in which a refund (net of service fee) can be issued and increasing the penalty outside of this window.

Upon booking, consumers pay the per-night price, a fixed cleaning fee (set by the seller), a percent service fee set by the platform, and taxes (if Airbnb collects logding taxes on behalf of the city). Airbnb keeps the payment until the stay is concluded. Airbnb also charges a percent seller fee. This paper focuses on San Francisco, in which the platform charges a 3% (ad valorem) fee to sellers and a 14% fee, plus a 14% transient lodging tax, to consumers.

Sellers might manage listings on their own or commission a management company. Management services are expensive: a typical management company charges 12-40% of the total revenue.[7] Interviews with practitioners suggest that the use of such services is uncommon. Few hotel or resort chains operate directly on Airbnb,[8] although some Airbnb hosts might be professional sellers.

Price-setting interface. To understanding the nature of pricing frictions on this market, I investigate the platform's standard price-setting interface. Figure 1 shows screenshots of the interface (captured in June 2020). The seller can set a base price for all nights. She can set a different price for weekends (Fridays and Saturdays). Beyond that, she can individually set prices for each night on the price calendar. To do so, she can select one night (or a range of *consecutive* nights), enter a "nightly price," and click "save" to confirm. In addition, once a night's price is set, the only way to change it is to manually adjust the price (changing the base price will not affect any existing night). Given this interface, manually setting and changing the nightly prices is a labor-intensive task. It is therefore not surprising to see many hosts complaining about the lack of a more convenient interface. For example, Joanna14, an experienced Airbnb host, describes her frustration: *"It is proving TOO time consuming to maintain the pricing using the Airbnb standard [interface]... It is just too basic and does not allow enough flexibility!"*[9]

Airbnb has implemented two changes that affects price flexibility. First, in November 2015, Airbnb launched a pricing algorithm, "Smart Pricing." Sellers can opt in to Smart Pricing (see

[7]Source: https://www.guestready.com/blog/airbnb-management-service-choose. Accessed in June 2020.

[8]I examine whether the host's name contains keywords such as "hotel", "resort", or other terms that indicate that the host is a company. I find these keywords only consist of 0.61% of all listings, suggesting that very few hosts brand themselves as a company (but a host can still be a professional seller). Full list of the keywords: Apartments, Corp, Guest, House, Hotel, Inc, Rental, Resort, Room.

[9]Source: https://community.withairbnb.com/t5/Hosting/Lack-of-seasonal-pricing-forcing-me-to-consider-leaving-Airbnb/td-p/328832/. Extracted in April 2021.

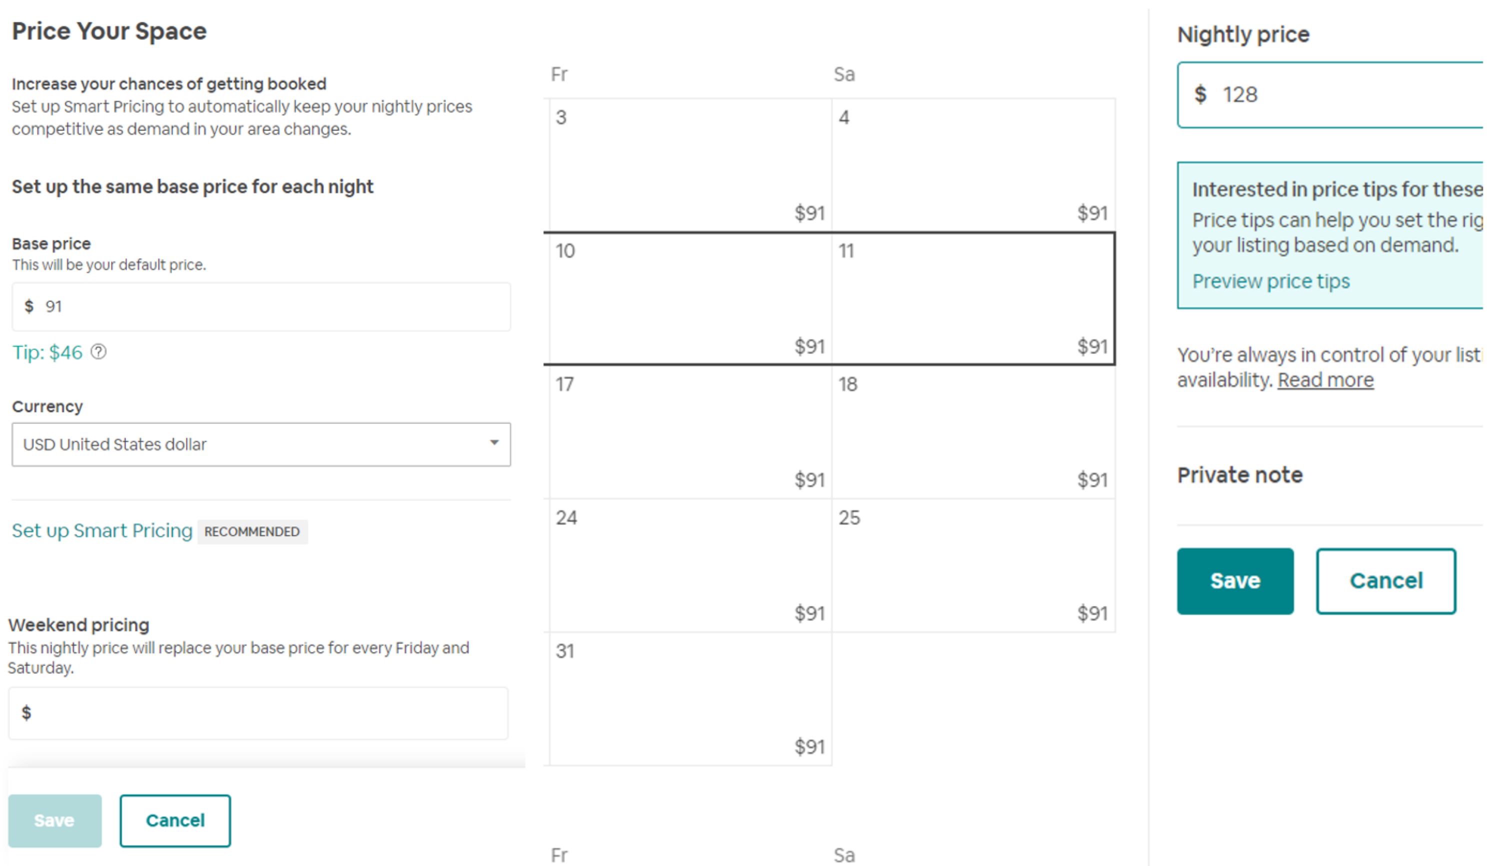

Figure 1: Standard price-setting interface on Airbnb

Notes: This figure presents three aspects of the pricing interface (at different part of the webpage). The top-left portion is the interface to set one base price for all dates. The bottom-left is the interface to set a weekend price (for Friday and Saturday). The right panel is the calendar view, which allows the user to manually set prices for each calendar date (overriding the price set elsewhere) or for a group of consecutive dates that are multi-selected.

top-left panel of Figure 1), in which case all prices are set by the platform's algorithm by default.[10] Whereas Airbnb does not officially discuss what the algorithm is, Ye et al. (2018), who are affiliated with Airbnb, explain that the algorithm first estimates a reduced-form consumer demand function using observed prices, and then solves for the revenue-maximizing prices for each listing. If sellers have positive marginal costs (as my results show), the revenue-maximzing algorithm will price lower than sellers. Consistent with this prediction, numerous articles and discussions argue that sellers would have wanted higher prices but cannot achieve so with Smart Pricing.[11,12]

Another platform change that affects price flexibility is the introduction of "last-minute discounts" in early 2019. This new feature allows the seller to pre-set a percent price discount and a threshold in the lead time (time before check-in). For a given night, when the number of days before check-in falls under the threshold, the percent discount is automatically applied to the price of the night – without manual intervention by the seller. This feature is a part of the "professional hosting tool," which includes other features (e.g., monthly discounts) that are less relevant for this paper.

Data and sample selection. The data come from Inside Airbnb (`insideairbnb.com`) and are publically available under CC0 1.0 Universal License. These data cover all listings from 96 cities or regions worldwide and are collected from Airbnb roughly at the monthly frequency since early 2015. At each data collection instance t, two datasets are relevant to this research. The first dataset includes characteristics of the listing observed on sampling date t, such as the seller's identity, the listing's available features or amenities, location measured by neighborhood, zipcode, and coordinates, and average rating and price.

The second dataset is the calendar data. On sampling date t, I observe the availability status

[10]The seller can still manually override each night's price in the calendar, in the same way as when she manually sets nightly prices. She can also set a price floor and ceiling, which simply bounds the algorithm's price.

[11]See, e.g., https://www.hostyapp.com/smart-pricing-sets-airbnb-rates-low/. Also see numerous forus discussions on reddit/r/airbnb.

[12]Beyond using the platform's algorithm, sellers can use paid third-party pricing software. Typically, using third-party interfaces incurs a fee (usually 1% of total revenue) and requires the sellers to set up the software through Airbnb's API. While I do not have direct measures of who uses a pricing software, I later demonstrate that price variations are low and display clear patterns consistent with the standard price-setting interface, suggesting that the majority of sellers still use the standard interface to set prices.

of each *night* τ—that is, 1 if the listing is available when searched on date t, and 0 otherwise. Booking can start one year before check-in, and thus, I typically have 12 monthly observations for each night τ. I interpret the listing-night as sold in-between $t-1$ and t, if the night is available on $t-1$ but unavailable on t. In case the listing-night is always unavailable for the entire 12-month duration, I interpret the seller has "blocked" night τ, i.e., made it unavailable from the start. The price of τ is shown if the night is available at t, and missing if the night is unavailable. I address missing prices below.

I focus on Airbnb listings in San Francisco and take a subsample based on three criteria. The first is to condition on listings who require a minimum stay of below three nights. The duration-of-stay requirement is to eliminate weekly and long-term rental listings, which belong to a different market.[13] The second criterion restricts attention to the most popular listing types: private rooms (38% of all listings in San Francisco), studios and single-bedroom apartments (31%), and two-bedroom apartments (17%). This screening step eliminates shared rooms and large houses, both of which are uncommon and do not appeal to mainstream customers. The last criterion is to condition on listing-years when the seller blocks no more than a quarter of nights. 25% of listings are dropped from the sample, which I interpret as part-time sellers (who might have very different incentives than full-time sellers).[14] After these three steps, I arrive at a sample of 18,054 listings, operated by 12,856 sellers, observed over 54 months. In total, the sample consists of 30,864,535 observations on the listing-night-sampling date level.

Missing data and interpolation. The monthly sampling rate creates a truncation problem: I do not see the exact date when the booking happens, and, in case when the last sampling date is still far from the night of stay, the night might be sold after the last observation. This truncation problem necessitates the interpolation of the booking outcomes.

I leverage the fact that some nights are close to the last sampling date, in which case the

[13]On the platform, 67% listings require a minimum stay of no more than 3 nights, 11% require a minimum stay of between 4 and 7 nights, 1% between 8 and 29 nights, and 21% above 30 nights. Requiring a minimum stay above 30 nights will put the listing in the long-term rental market and will exempt it from the hotel-lodging tax and other regulations.

[14]Of the remaining listings, 75% are available for at least 305 nights out of a year, and 50% are available all year.

truncation problem disappears. The data collection timing is fixed and thus should be orthogonal to unobserved market outcomes. Thus, conditional on observables such as day-of-the-week and seasonality, nights further away from the last sampling date should have the same *true* occupancy rate as nights closer to it, but the *measured* occupancy rate is lower due to truncation. I leverage this feature to interpolate the occupancy rate (and draw occupancy outcomes from the interpolated occupancy rate) for nights that are further away from the last sampling date. Appendix A provides additional details.

The low sampling frequency also creates a missing price problem. I do not see prices after a listing is booked, and thus do not see the exact price when the booking occurs. To interpolate prices, I leverage the observation that pricing strategies are highly simplistic for most listings. Section 3 documents that most listings either set uniform prices across all nights or follow simple pricing strategies that set prices conditional on weekend and month-of-the-year. For these listings, I interpolate the price of the booked night by estimating a pricing-policy function using the prices of other (not-yet-booked) nights. For example, if all observed prices are uniform, I assume that the unobserved price is the same as all observed prices. Appendix A provides details and shows that the majority of the interpolated nights are associated with simple policy functions. Still, for a small fraction of listings with flexible prices, the interpolated prices (coming from unsold nights) might be systematically biased. In this paper, I only use the observed (non-missing) prices in descriptives and supply-side estimation (and the interpolated prices are used in demand estimation).

3 Pricing strategies and frictions: empirical observations

This section presents summary statistics about seller characteristics and explores the extent to which these characteristics explain the heterogeneity in pricing strategies across sellers. I start by demonstrating that most sellers set unsophisticated prices that do not seem to respond to demand and opportunity cost differences, whereas a small fraction of them set sophisticated prices. Then, I show that this heterogeneity likely comes from inherent "seller type" differences. Although multi-listing sellers appear to be more sophisticated, this difference does not come from within-seller changes in scale (reflecting, e.g., management or thinking costs) or experience. I further examine plausible drivers of the lack of flexible pricing, and I demonstrate that price-adjustment costs explain some, but not all, lack of sophisticated pricing.

3.1 Heterogeneity in pricing strategies

Examples, and possible algorithmic pricing. To get a sense of what pricing strategies look like, I randomly draw 25 listing-sampling date level observations and plot their prices across nights. Figure 2 shows that prices follow strong uniformity patterns for most of these listings. 10 out of 25 listings have completely uniform prices, some listings have weekday-weekend patterns, and most others have large clusters of consecutive nights set at the same price. However, a few of them have significantly higher degrees of price variation.

Next, I examine listings whose prices are far from uniform. Specifically, I examine the set of listing-sampling dates, where the residual price variation conditional on weekend and month (detailed later), falls into the top-5%. I randomly draw 25 listing-sampling date observations out of this set. I find that the pricing patterns are complex for almost all of them and are difficult to be set by the standard interface. Meanwhile, prices share common, intuitive patterns: They are highly seasonal, with higher summer prices and lower winter prices, a polynomial-like baseline, and clear weekend patterns (but with heterogeneous weekend surcharges). One might speculate that these prices are set by algorithms or at least set with the algorithm's recommendation.

Measure of price variability. To formalize the above patterns, I construct three price-variability measures. First, to get a sense of overall price variation, I compute the standard deviation of log price across all nights τ for listing j holding fixed the sampling date t:

$$\text{std}\left(\log\left(\text{price}_{j\tau t}\right)\right)|_{j,t}. \tag{1}$$

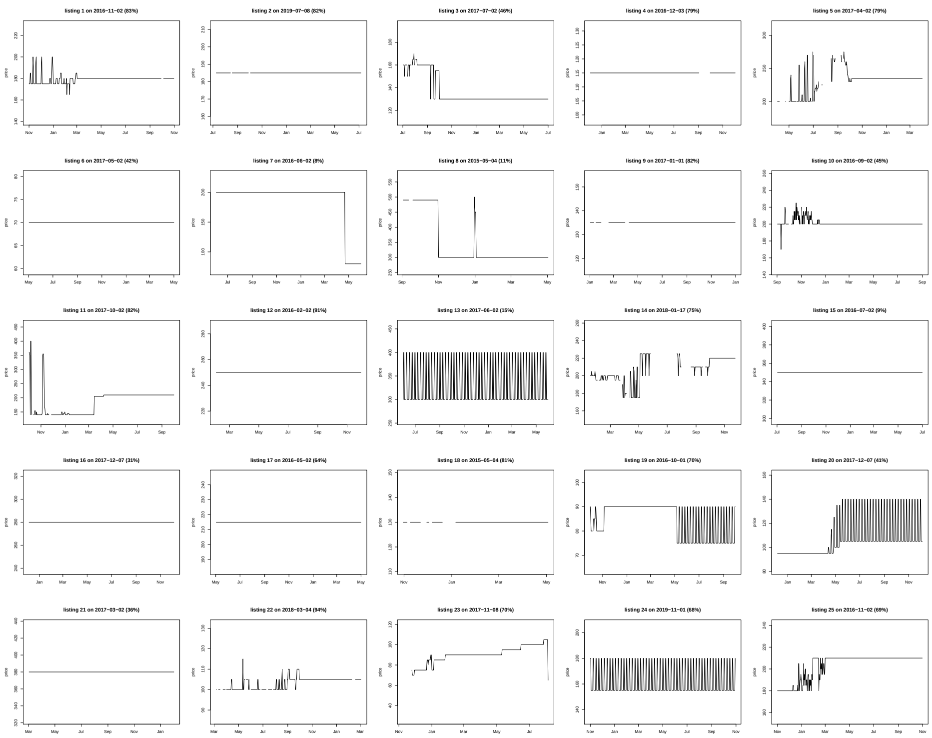

Panel A: all listings

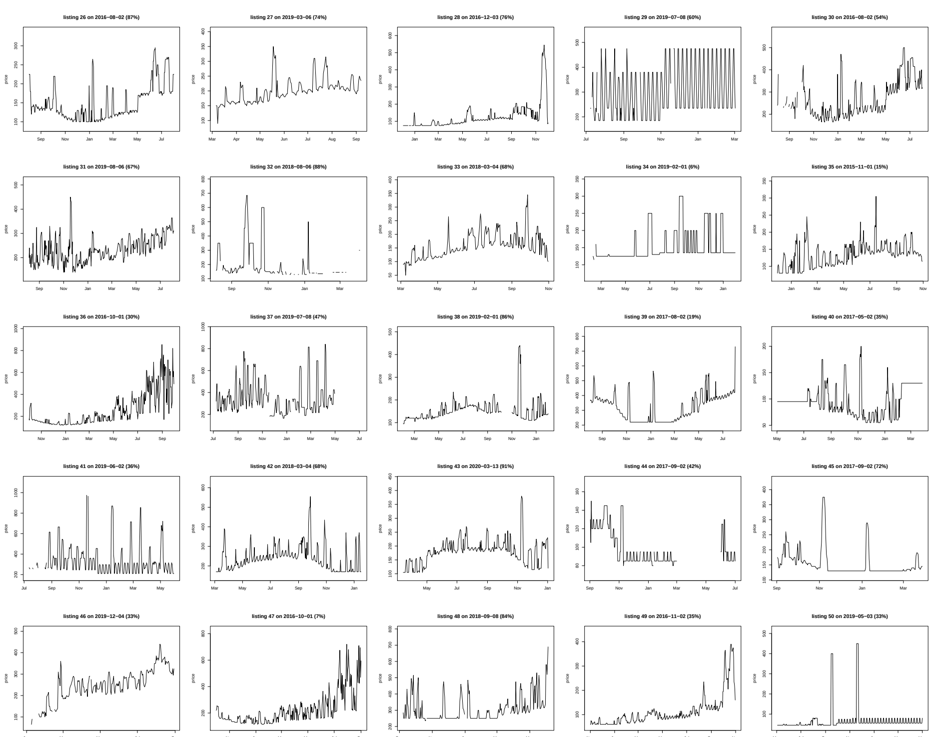

Panel B: listings with residual price variation *above* the 95th percentile (5% of sample)

Figure 2: Pricing patterns of randomly drawn listings

Notes: Prices for 50 randomly drawn listings. The X-axis is the night of the stay. The top 25 are drawn from all listings. The bottom 25 are drawn conditional on the residual price variation above the 95th percentile.

This standard deviation measures the percent-price variation across nights for a given listing: If the listing sets uniform prices across all nights and do not change prices over time-to-check-in, this standard deviation should be zero. If the listing's pricing policy function is simple, this standard deviation should be absorbed by state variables such as a weekend dummy (or in addition, month dummies).

The second measure focuses on the degree of price changes over lead time, i.e., time before check-in. Optimal Airbnb pricing *should* be dynamic: If the lead time is high, selling the listing today precludes selling it tomorrow (or the option value to sell in the future is high). In contrast, if check-in is imminent, the option value is low. As a result, prices should start high to capture the option to get high willingness-to-pay customers, and decrease (in expectation) over time as the option value dwindles. One way to measure the extent of dynamic pricing is to compute the average percent-difference between the last-month price and the initial price, or

$$\mathbb{E}\left[1-\frac{\text{price}_{j,\tau,12}}{\text{price}_{j,\tau,1}}|j,t\right]. \tag{2}$$

Third, to measure the degree of price uniformity across salient demand shifters, I compute the percent summer price premium (where summer is the quarter from July to September), weekend price premium, and holiday price premium. These "%premiums" capture the extent to which the host sets different prices across nights, of which the underlying demand is different. I mainly present the summer price premium in this section but show all measures in the appendix. This construction is closely related to Leisten (2020), who examines the extent to which chain and independent hotels use prices to capture salient and non-salient demand shifters.

Summary statistics. I summarize sellers' scale of operation and years of experience in the market, the distribution of price and quantity, and the three measures of price variability across nights and over time-to-check-in. Table 1 presents this summary. One finds sizable heterogeneity in the scale of operation and experience. For example, the median property is operated by single-listing sellers with 3 years of experience, whereas 25% properties are operated by sellers with at least 3

Table 1: Summary statistics across listings

	mean	5 pct	25 pct	median	75 pct	95 pct
total number of listings	4.002	1	1	1	3	11
years of experience on Airbnb	3.622	0	2	3.5	5	7
number of nights supplied (per 365 nights)	329	212	303	360	365	365
occupancy rate	0.612	0.045	0.370	0.693	0.872	0.997
price	188	72	108	150	235	410
std. of log price across nights	0.076	0.000	0.000	0.050	0.110	0.263
number of distinct prices (per 365 nights)	16	1	2	4	16	75
(negative of) %last-month discount	0.05	-0.11	-0.01	0.02	0.10	0.31
%summer price premium	0.024	-0.040	-0.003	0.004	0.035	0.158

Notes: All variables are measured at the listing-month level, and then averaged across months for each listing.

listings or at least 5 years of experience.

The table also presents considerable heterogeneity across the average occupancy rate and price levels of sellers. The top quartile of sellers are able to fill their properties at least 87% of the time, whereas the bottom quartile have their properties mostly empty. There is also a significant price dispersion: The 75th quantile of price, at $235, is more than two times of the 25th quantile at $108.

Further, the table demonstrates a systematic lack of price variation across nights and over time to check-in: The median standard deviation of price across nights is only about 5% of price, which is only about $7 at the median price. An alternative measure of price variability is to count the number of distinct price points (normalized by the number of nights supplied). The median listing has only 4 distinct price points per 365 nights, consistent with low price variation. In addition, summer prices are less than 0.4% higher for the median listing, and the last month (before check-in) prices are 2% lower than the initial price, despite high summer demand (shown later) and the decline in option value as the check-in date approaches. However, a small fraction of listings do display significant price variations. For example, a quarter of listings set at least a 11% standard deviation between nights (at least 16 distinct price points), charge at least 3.5% summer price premium, and provide 10% or more last-month discounts.

The lack of price variation over time, for the majority of listings, are consistent with price rigidity shown by Pan (2019). In addition, the lack of price flexibility across nights further suggest Airbnb hosts' inability or unwillingness to set prices that reflect demand differences across nights.

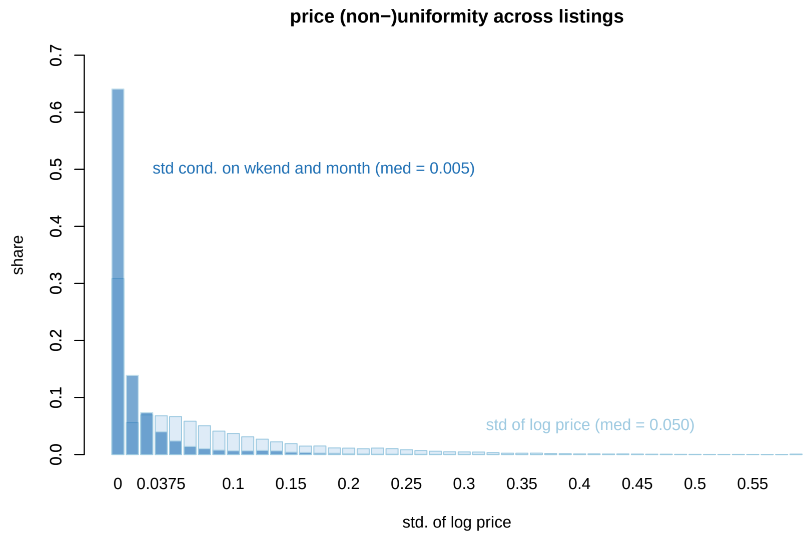

Figure 3: Unconditional and conditional (on weekend and month) price variation

Notes: Histogram of the unconditional standard deviation of log prices across nights, and the conditional one given weekend and month of the night.

This evidence is in line with uniform retail pricing across markets (DellaVigna and Gentzkow, 2019).

Decomposition of price variation. One conjecture of the lack of price variation is that many sellers use the standard price-setting interface provided by the platform, where all price changes and most flexible pricing policies (other than weekend pricing) have to be implemented manually. This interface's lack of feature implies that observed prices are functions of simple state variables, namely weekend dummies and month of the year (the latter is a proxy for "chunky" prices). To offer initial summary statistics related to this conjecture, Figure 3 shows residual (percent) price variation after controlling for weekend and month-of-the-year dummies. For the median listing, the residual price variation is down to 0.5% of price, one-tenth of the unconditional price variation.

3.2 Firm size and pricing strategy

I further examine the extent to which sellers' different pricing strategies can be explained by their observed characteristics. In particular, can "firm size"—the number of listings operated by a given seller—explain the degree of price variation for Airbnb sellers? The hypothesis is that multi-listing sellers (as a proxy for professional sellers) are more sophisticated in their pricing decisions. That is, we should see a larger degree of price variability (within listing) for sellers (*h*) with more listings, in the following regression,

$$\text{std}\left(\log\left(\text{price}\right)\right)_{jt} = \beta_{\#\text{listing}_{ht}} \mathbb{I}_{\#\text{listing}_{ht}} + \delta_{m(j)t} + X_{jt}\gamma + \varepsilon_{jt}. \tag{3}$$

In this regression, $\text{std}\left(\log\left(\text{price}\right)\right)_{jt}$ is the standard deviation of log price across nights τ within the same listing j-sampling date t. To interpret $\beta_{\#\text{listing}_{ht}}$ as the "firm size effect," one should compare similar listings operated by different sellers. I control for neighborhood-sampling time fixed effects $\delta_{m(j)t}$, and observed characteristics X_{jt} which includes (1) fully saturated neighborhood$\times$listing type$\times$number of rooms$\times$max number of guests fixed effects, and (2) amenity fixed effects.

Table 2 (A) shows that, conditional on observationally similar listings, multi-listing sellers adopt more sophisticated pricing strategies than single-listing sellers. The standard deviation of prices across nights is 1.1 to 6.4 percentage points (pp.) higher for multi-listing sellers. Likewise, multi-listing sellers offer a 1.4-3.6 pp. higher last-month discount and charge a 0.4-2.1 pp. higher summer price premium. All three measures are statistically significant and economically large (across different measures, 6+ listing sellers' degree of price variations is amost twice that of single-listing sellers). Appendix Table 4 show that this finding is robust to other measures of price variability.

Not explained by scale economy or learning. One plausible explanation of the difference in pricing frictions between seller types is that single-listing sellers have small operation scale and thus find it cumbersome to use sophisticated strategies. In other words, good pricing practices

Table 2: Pricing strategy differences and sellers' number of listings

(A) *Across*-seller differences (without seller FEs)

	Dependent variable:		
	std of %prices	%last-month discount	%summer premium
	(1)	(2)	(3)
2 listings	0.011***	−0.014**	0.004*
	(0.003)	(0.005)	(0.002)
3-5 listings	0.023***	−0.036***	0.016***
	(0.004)	(0.007)	(0.003)
6+ listings	0.064***	−0.021	0.021**
	(0.015)	(0.018)	(0.009)
baseline Y	0.067	-0.044	0.019
seller FE	no	no	no
loc.-time/type and amenities FE	yes	yes	yes
Observations	71,565	70,686	70,594
R^2	0.413	0.393	0.399
Note:			*p<0.1; **p<0.05; ***p<0.01

(B) *Within*-seller changes (with seller FEs)

	Dependent variable:		
	std of %prices	%last-month discount	%summer premium
	(1)	(2)	(3)
2 listings	−0.0001	0.004	−0.006**
	(0.003)	(0.005)	(0.003)
3-5 listings	0.001	0.005	0.001
	(0.004)	(0.007)	(0.004)
6+ listings	0.020**	−0.003	−0.0002
	(0.008)	(0.011)	(0.010)
seller FE	yes	yes	yes
loc.-time/type and amenities FE	yes	yes	yes
Observations	71,565	70,686	70,594
R^2	0.746	0.820	0.727
Note:			*p<0.1; **p<0.05; ***p<0.01

Notes: Baseline Y is the conditional mean of the dependent variable for single-listing sellers. Panel A focuses on differences between sellers with different #listings. Panel B controls for seller fixed effects (FEs), focusing on within-seller changes in #listings. Standard errors are clustered at the seller level.

incur higher fixed costs (e.g. hiring a manager) and thus is not uptaken by everyone. A related explanation is that good pricing practices come from learning by doing. That is, single-listing sellers are less experienced and find it costly to implement or "think through" a sophisticated strategy. Table 2 (B) shows that, as the number of listings increases *within* a seller, there is virtually no response in the degree of price flexibility. This finding suggests that sophisticated pricing is not a *choice* driven by fixed costs. Further, Appendix Table 3 demonstrates that learning does not explain differences in pricing strategies.

Main explanation: persistent seller heterogeneity in frictions. My main explanation of the finding is persistent heterogeneity between sellers in their pricing strategies, which is correlated with (but not driven by) the number of listings. In particular, single-listing sellers face a greater extent of pricing frictions, and thus, tend to set fixed prices across nights and/or over time.

In this context, one plausible driver of this friction is the time or effort in setting flexible prices under the standard interface. One would imagine that single-listing sellers face higher price-setting costs than multi-listing sellers (yet the above discussion suggests that these costs are not fixed costs). Another explanation is that some sellers face persistent cognitive constraints, such that they are bounded by using simple, heuristic-like, pricing strategies. For example, many surveys to professional managers (e.g., Hall and Hitch, 1939; Noble and Gruca, 1999) document that they use simple pricing heuristics such as "cost-plus" pricing. Many sellers on Airbnb (especially single-listing sellers) are not professional managers, suggesting that they might be more likely to follow simple behavioral heuristics. Section 3.3 and the structural model seek to distinguish between these two explanations.

3.3 Response to changes in the price-setting interface

Based on the above findings, this section further asks: Are the frictions driven by sellers' price-setting costs or sellers' cognitive constraints? Separating the two is important. If price-setting costs are the primary explanation, the platform can improve the interface to help with prices' flexibility.

But if sellers are constrained by their cognitive abilities, simply making the pricing interface more complex might not help them (and might even hurt them). When sellers face cognitive constraints, the platform might play a more active role in price-setting.

The ideal variation to separate the two mechanisms is to make pricing less costly for sellers and observe how many sellers (or which segment of sellers) react to this change. If all sellers are constrained by price-setting costs, they should set flexible prices when the price-setting costs are reduced. If all sellers are bounded by cognitive abilities, the flexible interface will not make prices more flexible.

From early 2019, Airbnb implemented a a new feature, "last-minute discounts," to make price adjustments easier.[15] This feature allows the seller to set one percentage markdown if the lead time falls below one threshold (set by the seller). Therefore, measuring the extent to which prices as a function of lead time changes before and after this interface update will help us understand nature of pricing frictions over time. If prices stay fixed because adjusting them is a labor-intensive process, sellers should set the optimal percent-discount after this feature is launched, and prices will be closer to the optimal dynamic-pricing path (but given that the feature only permits one percent change, prices will not coincide with the optimal path). Yet, if prices are fixed because sellers are cognitively constrained and do not think about dynamic pricing, prices will remain fixed under the new interface.

I estimate a regression of log price on a set of lead-time dummies (in weeks) by year, controlling for listing fixed effects δ_j.

$$\log\left(\text{price}_{j\tau t}\right) = \gamma_{\tau-t,y(t)} \cdot \mathbb{I}_{\tau-t} \times \mathbb{I}_{y(t)} + \delta_j + \varepsilon_{j\tau t}, \tag{4}$$

where $\gamma_{\tau-t,y(t)}$ captures the average price path as a function of lead time $\tau - t$ and separately by the year of the sampling date $y(t)$. The idea is to approximate the average pricing profile as a

[15] I use google archives to pinpoint that the change date is around January 2019, indicated by a surge of discussion about this feature. But there are sparse report that some hosts have received pilot trials of this system in 2018. Although the official website states that this feature is available for hosts with at least two listings, single-listing hosts also report having access to this feature.

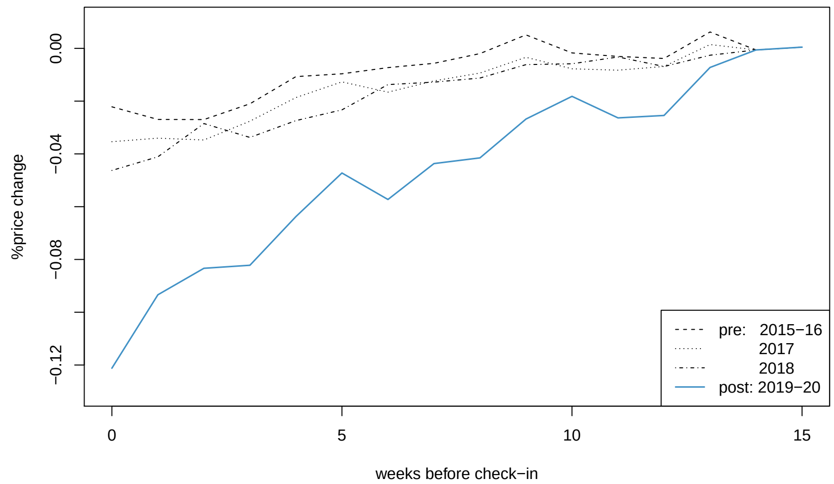

Figure 4: Price paths over lead time, by calendar year

Notes: Regression coefficients from Equation (4).

function of lead time and to see whether the shape of this profile changes in 2019 after the change of pricing interface (whereas other years's price paths should not differ much). Figure 4 shows the change of price profiles by year and demonstrates that the 2019 profile trends down more. This change is consistent with the conjecture that price-adjustment costs at least partly explain the price rigidities before 2019. After 2019, sellers can set a policy to lower prices as the night approaches (and different sellers might set different thresholds for the price drop to happen), instead of visiting the pricing page often to adjust prices.

However, a significant lack of dynamic pricing remains after the platform's change in interface. One way to present this observation is to look at the fraction of sellers using a dynamic pricing strategy and how this fraction changes before and after 2019.[16] The first row of Table 3 shows that 45% sellers use a dynamic pricing strategy, higher than the 36% before the introduction of the last-minute discount feature. However, 55% sellers –more than half of them– still do not set meaningful dynamic-pricing paths. The lack of dynamic pricing suggests that price-setting costs are not the only explanation.

[16]I define dynamic pricing as the average slope (over lead time) is steeper than 1% per month.

Table 3: Fraction of sellers using dynamic pricing: before and after 2019

	2015-2018	2019-2020
All sample	0.36	0.45
Sellers with six or more listings	0.38	0.46

Notes: Share of listings using dynamic pricing, defined as the average price decline per month of lead time is greater than 1% (in the last four months of lead time).

Not lack of awareness to the feature change. Another explanation for the lack of dynamic pricing after 2019 is that some sellers are not aware of the pricing interface change. Whereas many sellers have to activate the last-minute discount feature, the platform explicitly pushed this feature, as part of the "professional hosting tool," to sellers with six or more listings. The second row of Table 3 conditions on this subset of sellers and find low adoption of dynamic pricing, virtually identical to the first row. This finding suggests that the lack of awareness of the dynamic-pricing feature is not the main explanation.

4 Structural model and estimation

This section introduces the structural model, which contains three parts: (1) consumer demand for Airbnb listings, (2) sellers' dynamic pricing decisions, (3) and sellers' participation decisions on the platform. The purpose of constructing a model is to uncover demand primitives and sellers' marginal costs, establish the "first-best" market outcome (where all listings price optimally to maximize profits), and simulate counterfactual market outcomes under alternative market designs.

I introduce two incremental innovations to the demand model. First, I combine a nested fixed-point algorithm—typically used in random-coefficient logit models (Berry et al., 1995; Goolsbee and Petrin, 2004; Chintagunta and Dubé, 2005; Tuchman, 2019)—with sparse-demand models (Williams, 2021; Pan, 2019). This algorithm vastly improves the scalability of existing sparse-demand models and makes them viable on large datasets of differentiated products. Second, I present a new instrumental variable for price, leveraging the extent of price frictions in the context.

Due to computer-memory constraints, the structural model uses a random subsample consisting of 75% of the full sample. Appendix D discusses model details.

4.1 Consumer demand for Airbnb

Model setup. Two types of customers, denoted $k = 1,2$, arrive at the San Francisco Airbnb market. Consumer i of type k comes in month t and looks for a listing for the night τ in zipcode m. Her utility for booking listing j is

$$u^k_{ij\tau} = \delta_{jq(\tau)} + \alpha^k \log\left((1+r)\cdot p_{j\tau t(i)}\right) + \xi_{j\tau t(i)} + \varepsilon_{ij\tau}. \tag{5}$$

The consumer can choose not to book any listings, in which case she books a hotel and exits the market. Thus, the arrival time t is fixed given consumer identity i (hence the notation $t(i)$, later simplified as t). $\delta_{jq(\tau)}$ are fixed effects for listing j in the quarter of the night, $q(\tau)$, which absorbs quality, amenities, ratings and reviews, cleaning fee, and other unobserved features that vary infrequently. $p_{j\tau t(i)}$ is the price for night τ if booked in period t where the consumer i comes to the platform. r is a constant representing the percent service fees Airbnb charges on top of the list price. $\xi_{j\tau t(i)}$ captures unobserved demand shocks for night τ at time t – for example, a game occurs on day τ and is announced before time t. It will be clear later that, because I use a control function to address price endogeneity, I will parameterize $\xi_{j\tau t(i)} = \sigma\eta_{j\tau t(i)}$, where $\eta_{j\tau t(i)}$ is the error term of a first-stage price equation, and σ is an additional scale parameter. ε_{ijt} is a Type-1 extreme value error term.

I normalize $u_{i0\tau} = \varepsilon_{i0\tau}$ if the consumer does not book any listing. With this structure, we have a logit demand at the individual level:

$$s^k_{ij\tau} = \frac{\exp\left(\delta_{jq(\tau)} + \alpha^k \log\left((1+r)\cdot p_{j\tau t(i)}\right) + \xi_{j\tau t(i)}\right)}{1 + \sum_{j'\in J_{m\tau t}} \exp\left(\delta_{j'q(\tau)} + \alpha^k \log\left((1+r)\cdot p_{j'\tau t(i)}\right) + \xi_{j'\tau t(i)}\right)}, \tag{6}$$

where $J_{m\tau t}$ is the available set of listings at the time t for night τ in zipcode m.

Type k customers arrive at a Poisson rate $\lambda^k_{m\tau t}$, for zipcode m, night τ, and booking date t. I assume the arrival rate depends on the customer type, lead time (a linear function), whether τ is on a weekend or a national holiday, day of the week dw, and month of the quarter mo (the first month of each quarter is normalized to zero, given δ_{jq}'s). Specifically,

$$\lambda^k_{m\tau t} = \gamma^k_{0m} \exp\left(-\gamma^k_1 \cdot (\tau - t) + \gamma^k_2 \cdot \mathbb{I}_{\text{holiday}(\tau)} + \sum_{dw=1,\ldots,6} \gamma^k_{2+dw} \cdot \mathbb{I}_{\text{DOW}(\tau)=dw} + \sum_{mo\in\{2,3,5,6,8,9,11,12\}} \gamma^k_{8+mo} \cdot \mathbb{I}_{\text{MOY}(\tau)=mo} \right). \tag{7}$$

For the largest market $\bar{m}$ (zipcode 94110), I normalize $\gamma^1_{0\bar{m}} = 1000$. That is, in the last month before check-in, 1,000 segment 1 consumers arrive to book listings around Mission District. I normalize γ^1_{0m} for other markets to be proportional to the total number of rentals relative to $\bar{m}$. I can then estimate the number of consumers in segment 2, and how each segment's arrival changes with time-to-check in, day of the week and month of the year (see Appendix D.2 for detail).

Nested fixed point algorithm and estimation. There are 33,354 listing-quarter level intercepts, δ_{jq}'s, which are nonlinear parameters in the model. Jointly estimating them via maximum likelihood is infeasible, yet omitting them (for example, assuming that listings only differ in observable characteristics) makes demand too restrictive.[17] To allow for δ_{jq}'s in the model, I introduce a nested fixed-point algorithm to the model, bridging the gap between demand for capacity-constrained perishable products (Williams, 2021; Pan, 2019; Hortaçsu et al., 2021) and demand analysis in non-capacity-constrained products (Berry et al., 1995).

The main idea is to collapse observed (binary) occupancy *outcomes* on the listing-night level to compute (continuous) occupancy *rates* on the listing-quarter level. Built on Williams (2021) and Pan (2019), Appendix D.2 derives a close form expression for the *quarterly* occupancy rate,

$$\frac{1}{|q|} \sum_{\tau\in q} \text{occupancy}_{j\tau} = \bar{s}_{jq}\left(\delta_q\right) := \frac{1}{|q|} \sum_{\tau\in q} \left(\sum_t \left(1 - \exp\left(-s^1_{ij\tau} \cdot \lambda^1_{m\tau t} - s^2_{ij\tau} \cdot \lambda^2_{m\tau t}\right)\right) A_{j\tau t} \right). \tag{8}$$

Denote $A_{j\tau t}$ as the availability of listing j, night τ at the beginning of month t. The left-hand side of the first equality is the *observed* quarterly occupancy rate. The right-hand side is a closed-

[17]For example, a vast literature highlight the importance of reputation (e.g., Fradkin et al., 2018). The reputation effect is consistent with my descriptive evidence that prices tend to increase as the host stays longer on the platform. Also, Zhang et al. (2019) show that image quality plays a role in driving demand for Airbnb listings. Reputation (reviews) and pictures are two examples of many unobserved demand shifters.

form expression of the occupancy rate, a function of all listing fixed effects in the quarter, $\delta_q = \left(\delta_{1q}, ..., \delta_{Jq}\right)'$ and other demand parameters, α^k, γ^k and σ. Given these parameters, one can exactly solve for δ_{jq}'s by stacking J equations with J unknowns for each quarter.

Estimation follows the maximum likelihood estimator (with a nested fixed point) similar to Goolsbee and Petrin (2004), Chintagunta and Dubé (2005), and Tuchman (2019). Specifically, for each set of trial parameters $\left(\alpha^1, \alpha^2, \gamma^1, \gamma^2, \sigma\right)$, I compute δ_{jq} via a fixed point algorithm, which then allows me to evaluate the likelihood value.[18]

Uniform-pricing instrument (control function). Price $p_{j\tau t}$ might be endogeneous to unobserved demand shifters $\xi_{j\tau t(i)}$. These demand shifters potentially capture unobserved local events on night τ or other changes within a listing-quarter.

I propose a new price instrument that leverages the extent of pricing frictions in the market. The main idea is that prices often set uniformly across unrelated nights. In particular, far-apart nights, say τ', might have the same price with the focal night τ *only because* it is convenient to the seller to set one price for different night. The costly price adjustments in the standard interface further makes prices of consecutive nights either persistent or adjusting at the same time. Thus, pricing frictions create correlation in prices between dates τ and τ' even when the underlying demand shocks are uncorrelated (after controlling for listing-quarter fixed effects δ_{jq}). Based on this idea, for a focal night τ, I construct the average price for nights τ' in *different quarters* of the focal date τ. I further use the one-month lagged price for each τ' to guard against simultaneity. These "uniform pricing" instruments are, to my knowledge, new to the literature on demand estimation and might have broader applicability beyond Airbnb given the wide variety of markets with uniform prices (DellaVigna and Gentzkow, 2019; Adams and Williams, 2019; Hitsch et al., 2019).

I show in Appendix section B that this instrument strongly predicts price, with an excluded-variable F-statistic on the order of 10,000 in linear specifications with the same set of controls.

[18] Availability $A_{j\tau t}$ is known in the data during estimation. The fixed point (8) can be solved quickly because $A_{j\tau t}$ is known (different from Tuchman, 2019, who needs to simulate individual states). In the counterfactual, while I need to I forward-simulate availability $A_{j\tau t}$, the δ_{jq}'s are already solved, and one does not need to repeat the fixed point algorithm.

I also show that this "uniform pricing" instrument produces very different results from using the lagged price as an instrument, with the latter potentially endogenous to time-invariant unobservables on each night (e.g., events on given dates).

Given the nonlinear demand model, I adopt a control function approach (Petrin and Train, 2010) where I allow prices to be a function of observables $x_{jt\tau}$ and the uniform pricing instruments (or strictly speaking, excluded variables) $z_{jt\tau}$, or

$$p_{j\tau t} = \left(x_{j\tau t}, z_{j\tau t}\right) \cdot \phi + \eta_{jt\tau}. \tag{9}$$

I estimate this first stage and obtain the residual $\hat{\eta}_{j\tau t}$ as the proxy for demand shocks—that is, $\hat{\xi}_{j\tau t} = \sigma \hat{\eta}_{j\tau t}$. This approach implicitly assumes that prices fully capture unobserved demand factors $\xi_{j\tau t}$.

4.2 Optimal dynamic pricing decisions

Listings make two decisions. Each listing makes a participation decision for each quarter: Conditional on having entered the market, it decides whether to stay in the market. Then, the listing makes pricing decisions. I characterize observed prices as mixtures of different extent of cognitive constraints: (1) dynamic pricing for each night, with no cognitive constraint (but possibly facing price-adjustment costs), (2) cognitively constrained sellers making setting non-dynamic prices, which does not vary over time but *might* (or might not) vary across nights. In this section and the next, I characterize the two types of pricing decisions.

Optimal prices are determined by a finite-horizon dynamic pricing problem where the seller faces price-adjustment costs. I model these adjustment costs as sellers drawing a chance to set prices with probability μ_j (Calvo, 1983). The paramter μ_j is heterogeneous across sellers and captures the possibility that the seller does not make price adjustments even if it is profitable to do so.

Now, sellers' optimal dynamic pricing accounts for the possibility of future inaction. Denote the probability that *at least one* customer books listing j for night τ during month t as $q_{j\tau t}$. Denote the inclusive value for choosing all listings other than j as $\omega_{j\tau t}$ (Pan, 2019), which represents the "state" of the market governing the competitive set for listing j. Also define $\pi_{j\tau t}$ as the static flow profit of listing j if date τ can be rented out in month t:

$$\pi_{j\tau t}\left(p,\omega_{j\tau t}\right)=q_{j\tau t}\left(p,\omega_{j\tau t}\right)\cdot\left(p\cdot(1-f)-c_j\right) \tag{10}$$

where $f=0.03$ is the fixed fee platform sets for all sellers in San Francisco and c_j is the marginal cost for hosting for a night.

With these notations, one can characterize the *observed* price path set by an optimizing listing j, with inaction probability $1-\mu_j$, as

$$p_{j\tau t}^{\text{dynamic}}=\begin{cases}p_{j\tau t}^{*} & \text{probability } \mu_j\\ p_{j\tau,t-1}^{\text{dynamic}} & \text{probability } 1-\mu_j\end{cases}. \tag{11}$$

Note that $p_{j\tau t}^{*}$ is the optimal price when the seller sets prices (with probability μ_j), which is different from the observed prices $p_{j\tau t}^{\text{dynamic}}$. When setting the optimal price, the seller solves an optimal dynamic programming problem for the optimal price p^*,

$$\max_p \pi_{j\tau t}\left(p,\omega_{j\tau t}\right)+\left(1-q_{jt\tau}\left(p,\omega_{j\tau t}\right)\right)\mathbb{E}\left[V_{j\tau,t+1}\left(p,\omega_{j\tau,t+1}\right)|p,\omega_{j\tau t}\right], \tag{12}$$

i.e., the seller balances current profit $\pi_{j\tau t}\left(p,\omega_{j\tau t}\right)$, which depends on the probability that night τ can be rented out now, and future value $V_{j\tau,t+1}$, which is multiplied by the probability that night τ remains on the market. The value $V_{j\tau,t+1}$, and thus optimal dynamic price $p_{j\tau t}^{*}$ can be solved via backward induction. The details (including an illustrative example) are presented in Appendix D.3.

Discussions. Several properties of this structure warrant discussion. First, the value $V_{j\tau,t+1}$ gives an "option value" for the night τ, such that the seller has an incentive not to sell it immediately but keep prices high in search of a high-willingness-to-pay consumer. This term drives the optimal dynamic prices higher (in expectation) than the static optimal price until the lead time $\tau - t$ gets to zero (when pricing reduces to a static optimization problem). In other words, optimal dynamic prices decrease as the time-to-check-in approaches because of the declining option value in $V_{j\tau,t+1}$.

Second, if the seller draws an action and can adjust prices, she rationally expects a probability $1-\mu_j$ that her current prices will be carried into the next period (hence the value $V_{j\tau,t+1}$ also depends on p). This expectation changes her current price p^*, generally making p^* declining quicker than the optimal price without the inaction probability. See Appendix D.3 for a detailed discussion with an example.

Third, sellers form rational expectation on future states $\omega_{j\tau t'}$ for $t' > t$. I assume that sellers take the predicted $\omega_{j\tau t'}$ using information up to t, or $\mathbb{E}_t\left[\omega_{j\tau t'}\right]$.[19] In past versions, I also assumed perfect foresight and second-order Markov belief, and have found similar estimates between these assumptions.

Lastly, one feature I simplify away is multi-product pricing. In principle, owning multiple listings will raise the price of each listing due to cannibalization among them. However, given the large choice sets and that no seller holds a sizable portion of the market, substitution within a seller can be safely ignored, which significantly simplifies the pricing model.

4.3 Pricing with (non-dynamic) cognitive constraints

If price-adjustment costs are the only source of departure fully flexible optimal prices, dynamic price paths $p^{\text{dynamic}}_{j\tau t}$ should approximate the observed prices. However, section 3.3 has shown that many sellers do not use dynamic pricing when such strategies are made easier to implement.

[19]Specifically, for each seller segment l (introduced later), I estimate

$$\omega_{j\tau t'} = b_{0l} + b_{1l} \cdot t + b_{2l} \cdot t^2 + b_{3l} \cdot \omega_{j\tau t} + \Delta\omega_{j\tau t'}$$

and take the prediction, $\mathbb{E}_t\left[\omega_{j\tau t'}\right]$, to approximate seller expectation.

Consistent with this finding, I characterize sellers' pricing behavior subject to a non-dynamic-pricing constraint (interpreted as a cognitive constraint). These non-dynamic-pricing sellers are also subject to price-setting costs, limiting the degree of flexibility of their pricing policy.

Specifically, for listing j in quarter q, a non-dynamic pricing seller will divide all nights into $1 \leq K_{jq} \leq 90$ bins (where each bin includes consecutive nights), and manually set one price for each bin. A high K (e.g., $K \to 90$) reflects low price-setting costs, and as a result, flexible prices reflecting the different underlying demand across nights (although they are still time-invariant). A low K reflects the case where the seller finds it costly to manually set prices on the calendar, in which case her prices will show little variation despite the different underlying demand.

Specifically, listing j in quarter q draws the number of distinct "price bins" from a Poisson distribution, $K_{jq(\tau)} \sim Poiss\left(\rho_j\right)$, and within each bin, prices must be the same. $\rho_j \in (1,\infty)$ is the expected number of price points that the listing can set in a quarter. The higher the ρ_j, the more flexible the prices are. I further assume that the listing knows which consecutive nights (bins) to set the same price, conditional on the number of bins $K_{jq} = K$. I first search for K clusters of consecutive nights where similar consecutive $\bar{\lambda}_{j\tau}$ (average consumer arrival across the two segments) are clustered together. This results in a partition of nights such that nights in each partition $\tau \in \left(\bar{\tau}^{k-1}, \bar{\tau}^{k}\right]$ will be charged the same price. Therefore, given K, one can write the pricing problem as setting one price to optimize the total profit from a given partition of nights:

$$\bar{p}_j^k(K) = \max_p \sum_{\tau=\bar{\tau}^{k-1}+1}^{\bar{\tau}^k} \left(p - c_j\right) \cdot \mathbb{E}\left[\text{occupancy}_{j\tau}(p)\right]. \tag{13}$$

The notation $p_j^k(K)$ implicitly ackowledges that the partition-specific prices depends on the number of partitions, K. The exact partitions $(\bar{\tau}^{k-1}, \bar{\tau}^k]$ also changes with K.

Integrating over the realization K, one can summarize the listing's optimal non-dynamic prices as

$$\bar{p}_{j\tau}^{\text{non-dynm}} = \mathbb{E}\left[\bar{p}_j^{k(\tau)} | \rho_j\right]. \tag{14}$$

Modelling observed prices. I then characterize the observed prices of listing j as a weighted average between the optimal dynamic price $p_{j\tau t}^{\text{dynamic}}$ and the optimal non-dynamic prices $\bar{p}_{j\tau}^{\text{non-dynm}}$,

$$p_{j\tau t} = \theta_j p_{j\tau t}^{\text{dynamic}} + \left(1 - \theta_j\right) \bar{p}_{j\tau}^{\text{non-dynm}}. \tag{15}$$

Here, $\theta_j \to 1$ (0) will indicate that listing j is the type that sets dynamic (non-dynamic) prices. Whereas it is hard to interpret a θ_j in-between 0 and 1 as a fixed seller type, I will later show that many listings have θ_j's close to either 0 or 1.

4.4 Identification of the pricing model: An example

I now illustrate which moments in the data identify sellers' price-changing probability (μ_j), number of price points (ρ_j), and their tendency to set dynamic prices (θ_j). I will present one example and compute its observed prices over time and across different nights, given the estimated demand but under different supply-side parameters. All identification arguments work within a listing, although there is some pooling across listings in estimation.

I first set $c_j = 0$, assume that the listing is fully optimizing and is not subject to any frictions ($\theta_j = 1$, $\mu_j = 1$). The right-most panel in Figure 5 (A) presents the distribution of these "ideal" prices across nights and over lead time (months before check-in). The solid line represents median prices in a given month relative to check-in, and the dashed lines are the 5th and 95th percentiles. The corresponding (right-most) panel in Figure 5 (B) shows the last-month price for each night of the quarter. One finds that prices start at around \$110 in the median but range between below \$100 to over \$140 depending on the night-of-stay's popularity. Moreover, whereas the median price goes down the closer it is to the check-in date (due to the increasing risk of not selling by the check-in date), prices for high-demand nights increase to over \$200 at one month before check-in (due to the scarcity of those nights and the large number of late-arriving customers). This is the optimal price path.

I then add price-adjustment costs, setting the inaction probability $\mu_j = 0.25$. That is, the listing

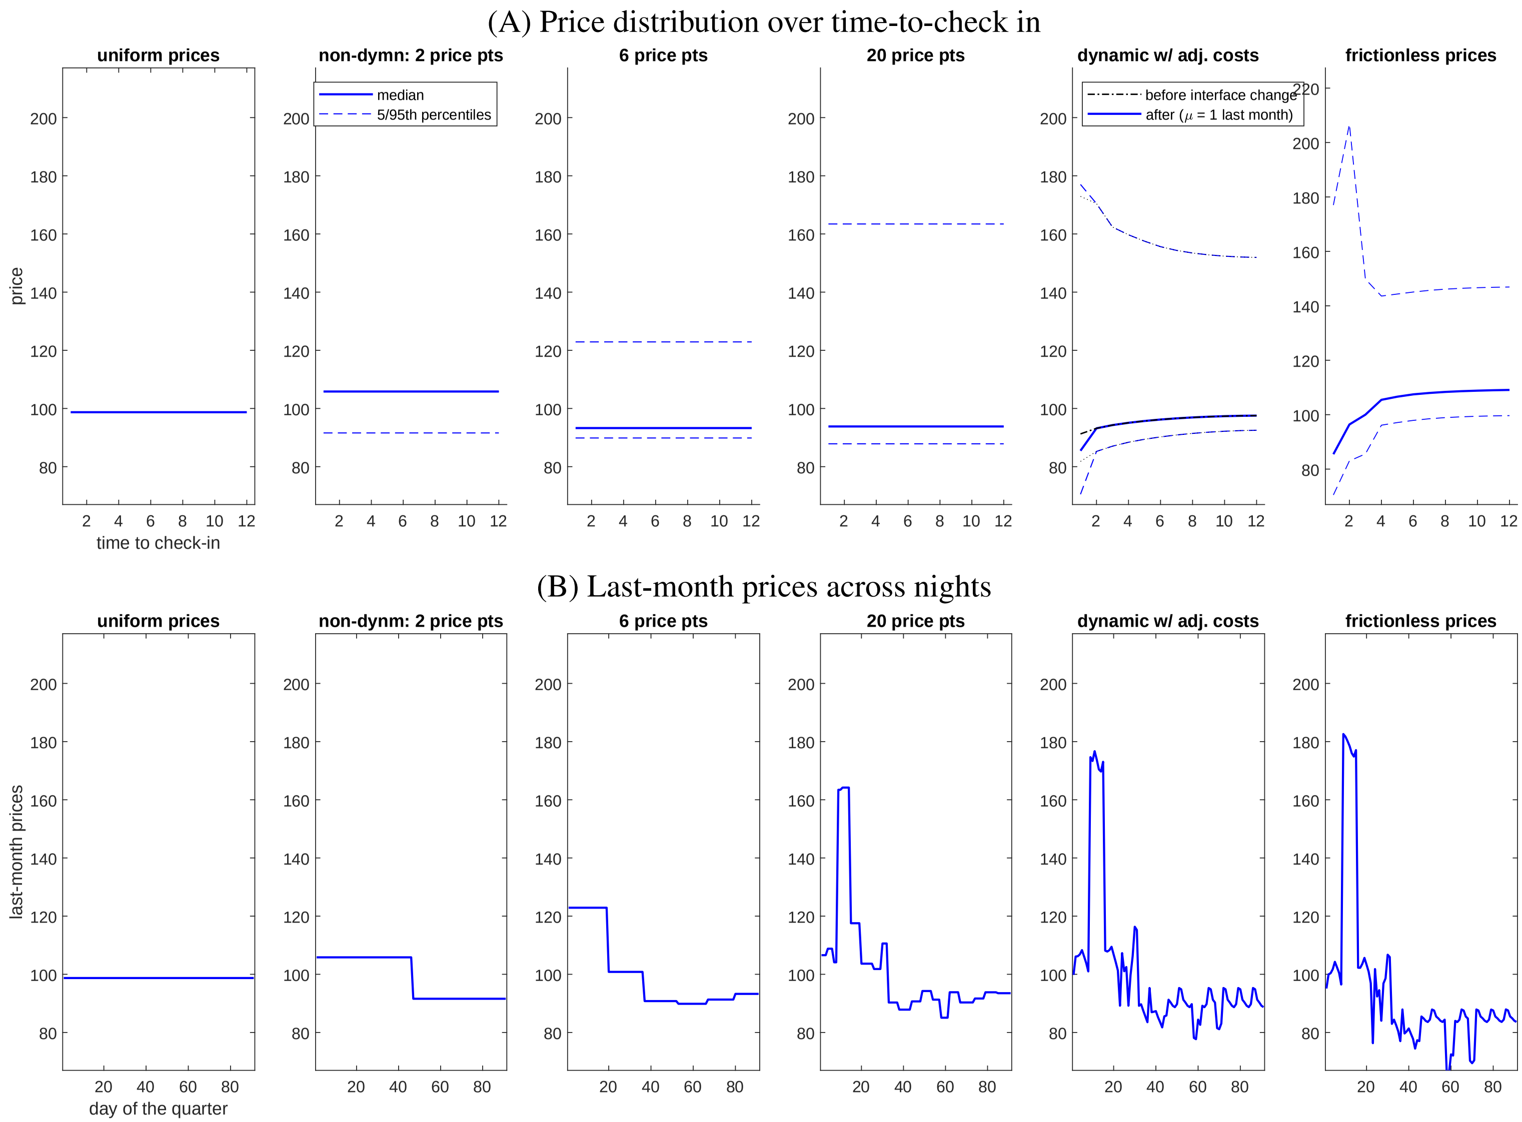

Figure 5: Price variation over time and across nights: A numerical example

Notes: A stylized example that demonstrates identification of supply-side parameters. All panels take demand for the same listing as given and assume zero marginal costs. The top panel focuses on price variations over time-to-check in, whereas the bottom presents last-month prices across different nights.

sets prices about 1/4 of the months and the forward-looking, rational seller knows about it. Shown in the second panel to the right of Panel A and B, the seller-set prices follows the dash-dot line (which, until the last period, overlaps with the solid line). Price-adjustment costs smooths out the optimal price variation over time: In each period, there is a possibility that the seller does not adjust prices when she should, and the seller adjusts prices preemptively in response to this possibility. Therefore, the degree of price variation over time identifies μ_j.

What happens when sellers are cognitively constrained and sets *non*-dynamic prices? The first four panels from the left of Figure 5 (A and B) show three examples of non-dynamic prices: A complex scheme with 20 different nightly prices, and less flexible pricing schemes with six, two, and one price. Sellers who can set many price points (high ρ_j) will be able to capture most of the nightly demand variation, despite not being able to adjust prices over time. Conversely, sellers who cannot set many price points (low ρ_j) will set inflexible nightly prices, with many contiguous nights "lumped together." The degree of flexibility across nights identify ρ_j.

Finally, what identifies the cognitive constraint from price-adjustment costs? Recall that the 2019 change in the price interface allows the seller to automatically adjust prices in the last period, resulting in a chance in the last-month price path (see the solid line, which departs from the dash-dot line).[20] Note that even if $\mu_j = 0$, the last-minute discount feature will still drive prices down in the last month. Therefore, the extent of last-month price change after 2019 distinguishes between low-μ_j sellers and sellers who are cognitively constrained from dynamic pricing ($\theta_j = 0$).

4.5 Estimation of the pricing model

I follow Bonhomme et al. (2019) and Pan (2019) to cluster all listings into segments based on their observed characteristics and observed actions. The underlying idea come from Bonhomme et al. (2019), who show that such *ex ante* clustering can approximate a model with continuously-distributed persistent heterogeneity with a flexible distribution. This method is desirable for my

[20]I assume sellers set last-minute discounts in the last month, which is a simplifying assumption but it agrees with the pattern shown in Figure 4.

paper because it does not impose strong shape restrictions on the joint distribution of seller parameters.

I cluster all listings into 150 groups by their observed prices and characteristics. Specifically, for each listing j, I first obtain its demand intercept, the number of listings operated by the owner, the listing's median price, price discount in the last two months before check-in, and the difference in its last-month discount after and before the 2019 interface change. Demand and the number of listings are important characteristics to control for. The vector of price moments closely resembles my identification arguments (which is important for the ex-ante split segments to resemble the underlying heterogeneity, see Bonhomme et al., 2019). I then use hierarchical clustering to group all listings into 150 clusters.

I give ρ_j an upper bound of 31 to ease computation burden. The above numerical example shows that 20 prices are already quite flexible to represent the shape of (a cross-section of) optimal prices.

Next, for each cluster l which j belongs to, I estimate $c_{j(l)}$, $\mu_{j(l)}$ and $\theta_{j(l)}$ using a generalized methods of moment (GMM) approach. Specifically, I match the following moments for each cluster l: (1) median price conditional on months-to-check-in (m_{1l}, a 12×1 vector), (2) median changes in last month's price before and after 2019 (m_{2l}, a scalar), and (3) median interquartile range of price across weekday nights within the last month (m_{3l}, a scalar). The choice of moments directly follow my identification strategy above. The (one-step) GMM objective is

$$\min_{c,\mu,\theta} (m_{1l}, m_{2l}, m_{3l})' I_{14} (m_{1l}, m_{2l}, m_{3l}) \tag{16}$$

with identity weights I_{14}. After estimation, I compute asymptotic standard errors from the variance-covariance matrix.[21]

[21]The variance-covariance matrix is given by $(\Gamma\Gamma')^{-1}$ where Γ is the Jacobian matrix of all moments on all parameters (Hansen, 1982).

4.6 Quarterly participation decisions

I construct a static entry/exit model to characterize listings' decision to participate in the Airbnb market. Recall that $q_{j\tau t}$ denotes the probability that listing j for date τ is rented by any customer who arrives in month t (conditional on being available before t). The expected total profit for quarter q is the sum of expected profit for each night τ in that quarter, which, in turn, integrates out the likelihood that date τ is rented out in each of the 12 months when it is available,

$$\Pi_{jq} = \mathbb{E}\left[\sum_{\tau \in q} \underbrace{\sum_{t=1}^{12}\left(\prod_{\iota=1}^{t-1}\left(1-q_{j\tau\iota}\right)\right) q_{j\tau t}}_{\text{prob. rented in any of the 12 months}} \cdot \underbrace{\left(p_{j\tau t}\cdot(1-f)-c_j\right)}_{\text{markup after platform fee}}\right]. \tag{17}$$

Given the quarterly profit,[22] I estimate the fixed costs by imposing that the listing decides to stay in the market (conditional on having entered the market) if it earns positive net expected profit:

$$\Pi_{jq} - F_{jq} > 0. \tag{18}$$

Here, fixed costs F_{jq} is a combination of the opportunity cost of participating in Airbnb (as opposed to renting on other platforms, or going to the long-term rental market, or using the apartment for oneself). I parameterize this cost as

$$F_{jq} = \bar{F}_{l(j)} + \psi_1 \mathbb{I}_{\text{post regulation}} + \psi_2 \text{dist}_j + \psi_3 \text{dist}_j^2 + \psi_4 \zeta_{jq} \tag{19}$$

where $\bar{F}_{l(j)}$ is segment l's fixed cost. $\mathbb{I}_{\text{post regulation}}$ is an indicator for the 2018 San Francisco regulation, which imposes a mandatory license requirement with annual fee (and an application

[22] To estimate the profit for listings that do not currently operate on the market, one needs to infer the demand intercept $\hat{\delta}_{jq}$ had it operated in the market. For observed listing-quarters, I estimate

$$\delta_{jq} = \delta_j^1 + \delta_q^2 + \Delta\delta_{jq}$$

and project $\hat{\delta}_{jq} = \hat{\delta}_j^1 + \hat{\delta}_q^2$. The R-squared of the above equation is 0.993. The projected δ's are only used to estimate fixed costs.

process). dist_j characterizes the listing's distance to Union Square (in miles) and proxies for the higher land cost in downtown San Francisco (or higher forgone rent). ζ_{jq} is a type-1 extreme value error term, which implies a binary-logit choice probability for the participation decision. I jointly estimate all parameters by maximum likelihood.

5 Estimation results

Demand estimates. Table 4 presents parameter estimates of consumer demand, omitting δ_{jq}'s, which are computed from the fixed point. I find that both segments are sensitive to price, with segment 2 (the segment that arrives later) being less price sensitive. The difference in price sensitivity between segments is consistent with the myopic consumer assumption; if consumers were to wait for discounts, late-adopters should be more price sensitive. I also find the residual from the control function, $\eta_{j\tau t}$, has a sizable coefficient, suggesting that prices are endogenous to demand shocks.

The average price elasticity in San Francisco is -2.51. This average elasticity is consistent with Jeziorski and Michelidaki (2019), who use experimental variation to estimate the price coefficient for Airbnb in San Francisco. This similarity gives face validity to the price instrument.

Further, the estimated day-of-week fixed effects suggest a clear weekend demand surge for both segments. In addition, the month-of-the-year fixed effects reflect a more nuanced seasonality structure. Spring is the high-demand season for San Francisco. Segment 1 demand peaks in late-Spring (June) whereas segment 2 demand peaks in early-Spring (March). This nuanced seasonal-demand structure is identified by the correlation between occupancy-rate level and booking timing: For example, June listings are likely booked early, reflected in the high segment 1 demand for that month (segment 1 comes earlier).

These arrival-process parameters can be graphically represented by Figure 6, where I plot the implied segment-specific arrival rate by time-to-check in, day-of-the-week, and month-of-the-year. I also find that the model predictions matches with the observed sales outcomes over all three dimensions. These rich demand variations, over time and across different nights, will be important

Table 4: Demand parameter estimates

	Segment 1	std err	Segment 2	std err
log(price)	-3.214	0.011	-2.695	0.008
#customers: baseline (last month)	1.000		2.549	0.027
%Δ by months to check-in	-0.383	0.001	-1.553	0.005
%Δ on holidays	0.196	0.012	-0.043	0.020
%Δ February	-0.204	0.007	0.508	0.007
%Δ March	0.091	0.007	0.784	0.008
%Δ May	0.203	0.006	0.364	0.009
%Δ June	0.534	0.006	0.230	0.010
%Δ August	0.310	0.006	-0.345	0.012
%Δ September	0.269	0.006	-0.123	0.011
%Δ November	-0.227	0.006	-0.358	0.008
%Δ December	-0.396	0.007	-0.359	0.008
%Δ Monday	0.002	0.005	-0.068	0.008
%Δ Tuesday	-0.004	0.005	-0.034	0.008
%Δ Wednesday	-0.004	0.005	0.036	0.008
%Δ Thursday	0.018	0.005	0.085	0.007
%Δ Friday	0.167	0.005	0.225	0.007
%Δ Saturday	0.165	0.005	0.247	0.007
scale of price residual (control fn)	2.714	0.009		

Notes: Nonlinear parameters from the demand model. Implied δ_{jq}'s are not reported in the table. Number of observations = 16,674,620 (33,354 listing-quarter $\times$ 91 check-in days $\times$ 12 months = 36,422,568, but only 16.7 million observations are when the listing is available). Log likelihood at convergence = -3,503,185. Asymptotic standard errors computed from the inverse Hessian matrix.

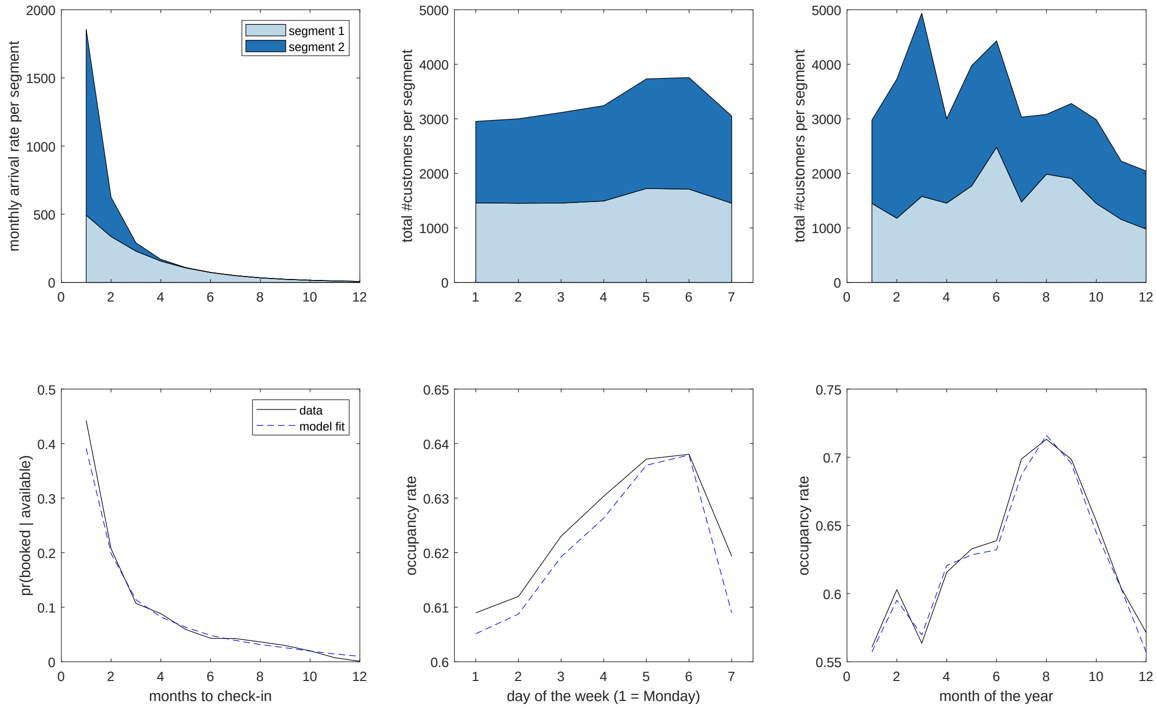

Figure 6: Customer arrival: model estimates and fit

Notes: Top panels: implied average *monthly* arrival rate for the two segments over time-until-check-in (left), and the *total* number of customers over day-of-the-week (middle) and month-of-the-year (right). Right: empirical and model-implied booking rate (left), which is defined as the probability of being booked in a given month conditional on availability, and occupancy rate (middle and right), which is defined as the probability of a given night-over-stay ever being booked.

drivers of *optimal* prices – ones that a rational, frictionless sellers should set.

Marginal costs and fixed costs. Figure 7 summarizes the distribution of marginal and fixed costs across listings. Recall that supply-side parameters are separately estimated for all 150 segments.

The average marginal cost is \$60 to host guest(s) for a night. I find significant heterogeneity in marginal costs across sellers, suggesting that different sellers might see the hassle of operating on Airbnb differently. The median cost is \$38, only 0.7% listings have negative costs, and the interquartile range is [\$24, \$68]. This interquartile range of marginal costs is equivalent to 1.5 to 4 hours of the minimum wage (\$16) in San Francisco (the city's average hourly wage is \$36 in 2019). These costs reflect the time to clean the property and communicate with guests, but might

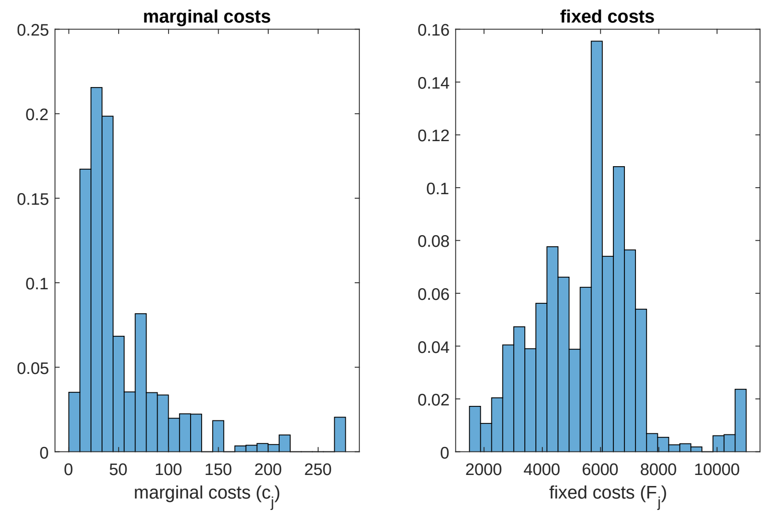

Figure 7: Distribution of marginal costs and fixed costs

Notes: Marginal distribution of nightly marginal costs and monthly fixed costs. Both distributions are winsorized at the 98th percentile.

also include depreciation of the property and other hassle costs. The marginal costs push prices above the revenue-maximizing level, typically considered as a benchmark for Airbnb and hotels.

I also find the monthly fixed cost at $5,540 on average, with an inter-quartile range at [$4,240, $6,510]. The 25th percentile is slightly higher than San Francisco's monthly rent in 2018-2019 (at about $3,500),[23] whereas the 75th percentile is much higher. Appendix Table 7 reports the estimates for ψ's, implying much lower fixed costs at locations further away from downtown: 3 miles away from Union Square, the average monthly fixed cost is about $3,340 lower. Accounting for this decline in fixed costs from downtown to outskirts of the city, these fixed costs estimates seem comparable to hosts' opportunity costs to listing their apartment on Airbnb.

Cognitive constraints (the use of non-dynamic pricing strategies). Figure 8 summarizes the distribution of listings' propensity to use non-dynamic pricing (θ) and price-setting costs (μ and ρ). First, the left panel presents the distribution of θ_j. 48% listings have $\theta_j > 0.5$ (and for most of them, $\theta_j \to 1$), meaning that they can set dynamic prices if price-setting costs are zero. The remaining half are cognitively constrained and set non-dynamic prices.

[23]Source: https://sf.curbed.com/2019/10/2/20895578/san-francisco-median-rents-market-census-september-2019. Extracted in May 2021.

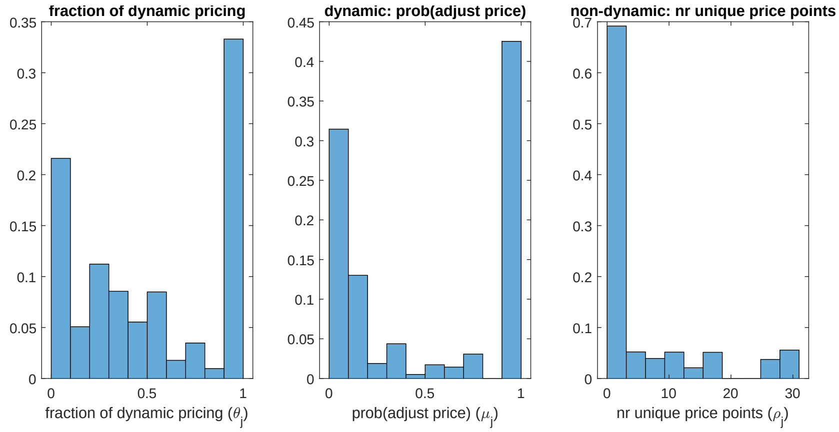

Figure 8: Distribution price-friction parameters

Notes: Marginal distribution of demand intercept (quality), marginal costs, probability of changing prices, fraction of sellers using dynamic pricing strategies, fraction using non-uniform pricing strategies, and quarterly fixed costs.

The middle panel presents the distribution of μ_j —the listing's probability of adjusting prices in each month— for the set of listings with $\theta_j > 0.5$. μ_j is clearly bi-modal. 43% of listings always adjust prices. 31% listings never adjust prices before the last-minute discount feature was available (and start using the feature after it was introduced). The remaining 26% have some price-adjustment costs such that they only get to change prices in some months (and most of them change prices rarely). Listings with price-adjustment costs will benefit from platform's effort that makes dynamic pricing less costly.

The right panel characterizes the distribution of ρ_j—the number of price points—for those who set non-dynamic prices. The average of ρ_j is 5.3. 74% of listings can set 5 or below price points in expectation. The low ρ_j suggests that the existing price-setting interface also lowers the degree of flexibility across nights. As a result, most of the listings using non-dynamic prices might also benefit from a more flexible price-setting interface.

Correlation with observed characteristics. I project each estimated supply-side parameters on observed characteristics to examine what explains differences in pricing strategies (and participation decisions). For example, I estimate a linear regression of marginal costs (c_j) on a vector of listing and seller characteristics, controlling for time and neighborhood fixed effects. Table 5

Table 5: Decomposition of supply-side primitives on listing and seller characteristics

	c_j (\$)	se	θ_j (%)	se	μ_j (%)	se	ρ_j (nr)	se
intercept	92.3	5.6	38.7	3.0	52.7	3.7	9.5	1.0
log(nr listing)	5.7	0.5	8.2	0.3	-0.5	0.3	1.2	0.1
host is superhost	-0.4	0.9	-2.0	0.5	2.6	0.6	0.5	0.2
respond in 1 day	-8.9	1.1	1.7	0.6	3.1	0.7	0.9	0.2
instant booking	-6.4	0.9	-1.6	0.5	-1.1	0.6	-0.1	0.2
flexible cancellation	10.3	1.0	-2.7	0.6	-5.1	0.7	-0.9	0.2

Notes: Regression results of estimated supply parameters on observed listing and seller characteristics. Additional controls not reported in the table are: fully saturated listing type (e.g. entire apartment) $\times$ max number of guests $\times$ property type (e.g. townhouse) fixed effects, amenity indicators (TV, internet, parking, washer/dryer, breakfast, allow pets), and length of the listing's descriptions.

presents these findings.

Controlling for listing characteristics, the biggest difference between single- and multi-listing sellers is on the extent of cognitive constraints, or whether the seller *can* set dynamic prices. One standard deviation in the number of listings explains 0.23 standard deviation of θ_j. In addition, one standard deviation in the number of listings explains 0.07 standard deviation of marginal costs, 0.11 standard deviation of ρ_j, and -0.01 standard deviation of μ_j. This result echoes the earlier finding that multi-listing sellers set sophisticated pricing strategies.

6 Counterfactual

I start by comparing equilibrium outcomes under the factual scenario (referred to as the "baseline"), where sellers set prices subject to frictions, to the frictionless market outcome where all frictions are eliminated ("first-best"). Whereas the first-best scenario is *unattainable* in practice, its simulation shows the upper bound of the potential gains from alleviating seller-side frictions. Then, I examine three platform remedies. The first is where the platform enforces a revenue-maximizing pricing algorithm. Such an algorithm approximates the "Smart Pricing" algorithm (Ye et al., 2018). The second scenario involves improving the pricing interface, which—in perhaps an overly optimistic scenario—eliminates sellers' price-adjustment costs but does not affect their cognitive constraints. The third scenario involves a more fundamental redesign of the platform: The platform provides a flexible price-adjustment function to sellers, and based on these functions, each seller sets the base price and leverages the platform's price-adjustment function. In this scenario, the platform leverages its informational and technological advantage, provides assistance to sellers, but does not take away sellers' rights to set higher-than-revenue-maximizing prices. I compare profits and consumer surplus across the five scenarios. All calculations are based on the period from May 2015 to December 2017, the sample period before the San Francisco regulation.

Counterfactual 0: First-best. The first two columns of Table 6 compare the baseline to the first-best. First-best prices are set where all pricing frictions are eliminated (i.e., $\mu_j = 1$, $\theta_j = 1$, and ρ_j is irrelevant in this case), and as a result, prices can target night-of-stay level demand differences and thus are more dispersed. One standard deviation of the price (within listing, across nights) increases from 4% in the baseline to 6% in the first-best for the median listing. Also, first-best prices generally decrease over time as the option value for waiting for additional customers dwindles. I show that the percent last-month discount increases to 30% from 2% (i.e., almost completely sticky prices) for the median listing. These two aspects lead to lower last-month prices and higher occupancy rates.

I further explore who gains and who loses in the first-best. The median seller gains as net profits increase from \$2,350 to \$2,440 per quarter, or by 3.8%. Table 7 further presents the distribution of within-seller profit changes. 1% sellers lose because they do not face significant frictions but their competitors can now price more flexibly (and often lower). Still, most sellers gain from the first best, and 5/95 percentiles of the profit gain is [0%, 15%]. Other than sellers, consumers also gain as their surplus (measured in utils) increases by 14%. The platform gains little—its profit increases by 2.5%. Almost all market participants gain from eliminating all frictions. The frictions significantly impact consumers and some sellers. Yet, around half the sellers and the platform are not affected much by the frictions (below 3% loss in surplus). The platform might have limited incentives to spend significant efforts to eliminate the frictions.

Table 6: Counterfactuals: median-seller outcomes, platform profit and consumer surplus

	baseline	first best	revenue-max.	ideal interface	platform-assist
last-month price	133.66	120.49	99.27	130.26	118.49
price dispersion across nights	0.04	0.06	0.09	0.05	0.06
last-month discount	-0.02	-0.30	-0.38	-0.10	-0.29
occupancy rate	0.73	0.77	0.90	0.74	0.79
seller quarterly profit ($k)	2.35	2.44	2.16	2.38	2.42
seller participation rate	1.00	1.00	0.92	1.00	1.00
total platform revenue ($m)	2.36	2.42	2.40	2.38	2.42
average consumer surplus (util)	4.11	4.69	6.03	4.23	4.84

Notes: This table summarizes counterfactual market outcomes in the baseline, first-best (i.e., sellers set prices without frictions), and under the two platform remedies. All seller-level outcome (last-month price, price dispersion, last-month discount, occupancy rate, and profits) are at the median. Price dispersion is the standard deviation of log(price) across nights. Last-month discount is the ratio between last-month price and first-month price, minus one. Profit is the net quarterly profit after substacting fixed costs.

Table 7: Counterfactuals: within-seller profit changes relative to the baseline

	first best	revenue-max.	ideal interface	platform-assist
fraction who earn negative profit	0.00	0.08	0.00	0.00
– who lose relative to baseline	0.01	0.56	0.00	0.09
– who gain relative to baseline	0.69	0.31	0.40	0.62
profit increase relative to baseline: 5 %	-0.00	-1.00	-0.00	-0.02
– 25%	0.00	-0.19	-0.00	-0.00
– median	0.03	-0.02	0.01	0.02
– 75%	0.05	0.02	0.02	0.05
– 95%	0.15	0.08	0.06	0.14

Notes: This table summarizes within-seller changes in profit in the counterfactual scenarios (relative to the baseline scenario). Fraction who lose (gain) profit is defined as the fraction of sellers who earn less than 99% profit (more than 101% profit) relative to the baseline, where I put a 1% buffer to filter out sellers who are close to indifferent.

Counterfactual 1: Revenue-maximizing algorithm. The platform has argued that "smart pricing," an algorithm that maximizes the revenue of each seller (Ye et al., 2018), will help eliminate pricing frictions. Maximizing seller revenue is in line with the platform's incentive because the platform takes an ad valorem fee (fixed percent on sellers' revenue). Yet, maximizing revenue might not be aligned with sellers' incentives, given that many of them have non-zero marginal costs. I now examine what happens if the platform enforces a seller-revenue-maximizing algorithm upon all sellers.[24]

Column 3 of Table 6 shows that, because revenue-maximizing prices do not factor in seller marginal costs, they are 26% below the baseline prices and 18% below the first-best prices. As a result, I find higher price dispersion across nights and heavier last-month discounts than the first-best. Occupancy rate jumps up to 90%. However, these changes are not aligned with sellers' interests. The median seller now earns only $2,160 per quarter, down by 11% from the first-best (and 8% from the baseline). Further, Table 7 shows that 8% sellers now earn negative profits and will exit the platform, 56% sellers are worse-off compared to the baseline, and 25% sellers' profits decrease by 19% or worse.

However, consumers and the platform gain from these changes. Platform profit is 1.7% above the baseline and 0.9% below the first-best, suggesting that, from the platform's perspective, a simple revenue-maximizing algorithm (that is easy to implement) could bring its outcome towards the first-best. Consumers surplus is up by 47% from the baseline and 29% from the first-best, because of the much lower prices and higher occupancy rate (note that these gains are achieved despite the lower variety, because of seller exit).

These equilibrium outcome predictions—decrease in prices, increase in occupancy rate, and the higher the seller exit rate—are highly consistent with a recent paper by Filippas et al. (2021). Filippas et al. work with a rental-car platform to experimentally impose a revenue-driven, centralized pricing algorithm on a random set of sellers (and in some cases, they give sellers a limited

[24]This scenario assumes that each seller solves for the frictionless revenue-maximzing prices. It is my interpretation of a (global) revenue-maximizing algorithm. But it does not necessarily use the same demand model as in Ye et al. (2018).

degree of control over prices). They find that the algorithm decreases prices, increases utilization rate, and creates a high amount of seller exit.[25] The platform in Filippas et al.'s later changed the pricing rule to centralized pricing (but giving sellers some room of price adjustments).

My counterfactual findings, as well as the connection to Filippas et al.'s observed market outcomes, highlight an important conflict of interest between the platform and its sellers. Whereas the platform would like to gain full control over pricing and eliminate pricing frictions through automation, doing so is detrimental to sellers (but will benefit consumers due to the low prices). In reality, Airbnb might not enforce such an algorithm in fear of seller backlash, and "Smart Pricing," introduced in 2015, might need further tuning. Nevertheless, the platform's preference for using this algorithm is consistent with its advocating for the algorithm over alternative pricing tools.

Counterfactual 2: A flexible price-setting interface. If sellers maintain control of pricing, can the platform do anything to alleviate pricing frictions? The 2019 interface change suggests some room for improvement from the pricing interface. I simulate the counterfactual impact if the platform introduces a fully-flexible dynamic pricing interface and a flexible interface across nights. In perhaps an overly optimistic scenario, this "ideal" interface eliminates all price-adjustment costs, setting $\mu_j = 1$ and $\rho_j = 31$, for all j. Meanwhile, this scenario does not affect sellers cognitive constraints, keeping θ_j's at their estimates.

The last columns of Table 6 and 7 show that the median last-month price discounts are much steeper than in the baseline, whereas price dispersions remain the same. The median seller profit is now \$2,380 per quarter, up by 1.3% from the baseline—only a third of the total potential gain from the first-best. Likewise, the platform gains a similar amount as the mandatory revenue-maximizing algorithm case, and consumer surplus is only up by 2.9%, far from the first-best.

The result shows that even in such an ideal scenario, the scope for improving pricing through reducing the price-adjustment costs is limited. This result is consistent with the estimates that an important source of frictions is sellers' cognitive constraints. One would imagine that in re-

[25]Exit rates are as high as 30% in their experiment. Some exits might be driven by sellers agonized by the platform's unilaterally changing the pricing rule. The authors acknowledge this possibility.

ality, such a complex price-setting interface might have backfired: sellers who are cognitively constrained might not react positively to the much more complicated pricing interface.

Counterfactual 3: Platfom-assisted pricing. In a different counterfactual scenario, I investigate whether the platform can leverage its advantage in information and technology to assist seller pricing. The platform has an information advantage over sellers because it has all the booking data and can estimate demand. The platform also has an advantage in pricing technology because it can implement any pricing rules or algorithms (and thus design how the market works). Although the current "Smart Pricing" algorithm tries to leverage these advantages, the platform's disadvantage is (1) not knowing sellers' marginal costs and (2) unable to commit to seller-incentive-aligned price levels (due to misaligned incentives). I now entertain a scenario where the platform commits to *not* set the base price level, but still leverages its informational and technological advantages to assist seller pricing.

Specifically, the platform redesigns pricing into two stages. In the first stage, the platform presents a price-adjustment function, $a_{ml\tau t}$, for market m, listing type $l(j) = 1, ..., 150$, night-of-the-quarter τ, and time t (with time-to-checkin $\tau - t$). The platform announces that the consumer price will be

$$p_{j\tau t} = \bar{p}_{jq(t)} \times \left(1 + a_{ml(j)\tau t}\right), \tag{20}$$

and that sellers only set quarterly price $\bar{p}_{jq(t)}$ — which is at the same level of uniform pricing in the model. In the second stage, sellers observe the now-committed $a_{ml\tau t}$ and set $\bar{p}_{jq(t)}$.[26]

How would the new equilibrium look like, and how far are we from the first best? The last column of Table 6 and 7 present the results. First, I find that the median price level is close to (but a bit lower than) the first-best, and that the degree of price adjustments over time, price dispersion, and occupancy rate are almost identical to the first-best. On the aggregate, the platform's new price

[26]To implement this counterfactual, for each j-τ-t, I take the ratio between the first-best prices over the counterfactual uniform (and time-invariant) prices to compute $\tilde{a}_{j\tau t}$. Then, I average these a's for each market m, time t, and for each type l to get $a_{ml\tau t}$. One might imagine that $a_{ml\tau t}$'s can be crude if the groups l are crude. One might also imagine that $a_{ml\tau t}$'s can be further optimized by the platform. For this counterfactual exercise, I use the crude (and potentially suboptimal) $a_{ml\tau t}$'s to illustrate that improvements can still be gained. Letting the platform strategically choose a's is beyond the paper, given the enormous computation burden.

function mimick's that of the first-best and forges the "shape" of equilibrium prices. As a result of this similarity, consumer surplus is now 3.2% above the first-best and is much higher than the baseline scenario (due to the slightly lower prices). Sellers, on the other hand, mostly gain (or at least not lose) from this change. The median seller profit is 3.0% higher than the baseline. 62% of sellers gain from this cahnge, and 9% sellers lose. Virtually no seller chooses to exit. All in all, the crude platform-assisted pricing scheme helps consumers and most sellers, and it also simplifies seller decision-making and loads most of the decision burden to the platform.

7 Summary

Pricing in a complex environment is difficult for individual sellers. While providing aid to seller pricing, the platform might have incentives to steer prices towards its objective. This paper shows that seller-pricing frictions —which I refer to as prices' lack of response to market conditions— are prevalent on Airbnb. I demonstrate that two mechanisms jointly drive the frictions: sellers' price-setting costs and cognitive constraints. I also estimate a tailored structural equilibrium model, which recovers rich consumer arrival processes and demand, sellers' pricing frictions, and their opportunity costs of time.

I demonstrate that the recovered pricing frictions lead to a 14% consumer welfare loss and heterogeneous profit losses for sellers—for example, the profit loss' 5–95 percentile range is [0%, 15%]. Given the loss, are there any ways that the platform can ameliorate the frictions? Based on the estimates, a flexible interface will unlikely help because significant frictions come from sellers' cognitive constraints (instead of interface-driven menu costs). Enforcing a revenue-maximizing algorithm will also see limited return because such an algorithm does not internalize sellers' opportunity costs of time. However, I show that a simple redesign will eliminate almost all frictions, where the platform sets price variation but commits to give sellers the final decision right to determine the price levels. Based on this analysis, ameliorating the pricing frictions is feasible in practice.

References

Adams, Brian and Kevin R Williams (2019), 'Zone pricing in retail oligopoly', *American Economic Journal: Microeconomics* **11**(1), 124–56.

Arcidiacono, Peter, Paul B Ellickson, Carl F Mela and John D Singleton (2020), 'The competitive effects of entry: Evidence from supercenter expansion', *American Economic Journal: Applied Economics* **12**(3), 175–206.

Barron, Kyle, Edward Kung and Davide Proserpio (2020), 'The effect of home-sharing on house prices and rents: Evidence from airbnb', *Marketing Science* .

Berry, Steven, James Levinsohn and Ariel Pakes (1995), 'Automobile prices in market equilibrium', *Econometrica: Journal of the Econometric Society* pp. 841–890.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten and John Van Reenen (2019), 'What drives differences in management practices?', *American Economic Review* **109**(5), 1648–83.

Bloom, Nicholas and John Van Reenen (2010), 'Why do management practices differ across firms and countries?', *Journal of economic perspectives* **24**(1), 203–24.

Bonhomme, Stéphane, Thibaut Lamadon and Elena Manresa (2019), 'Discretizing unobserved heterogeneity', (2019-16).

Brown, Zach and Alexander MacKay (2019), 'Competition in pricing algorithms', *Available at SSRN 3485024* .

Calvo, Guillermo A (1983), 'Staggered prices in a utility-maximizing framework', *Journal of monetary Economics* **12**(3), 383–398.

Chen, Le, Alan Mislove and Christo Wilson (2016), 'An empirical analysis of algorithmic pricing on amazon marketplace', pp. 1339–1349.

Chintagunta, Pradeep K and Jean-Pierre Dubé (2005), 'Estimating a stockkeeping-unit-level brand choice model that combines household panel data and store data', *Journal of Marketing Research* **42**(3), 368–379.

Cho, Sungjin and John Rust (2010), 'The flat rental puzzle', *The Review of Economic Studies* **77**(2), 560–594.

DellaVigna, Stefano and Matthew Gentzkow (2019), 'Uniform pricing in us retail chains', *The Quarterly Journal of Economics* **134**(4), 2011–2084.

Dubé, Jean-Pierre and Sanjog Misra (2017), 'Scalable price targeting'.

Farronato, Chiara and Andrey Fradkin (2018), 'The welfare effects of peer entry in the accommodation market: The case of airbnb'.

Filippas, Apostolos, Srikanth Jagabathula and Arun Sundararajan (2021), 'The limits of centralized pricing in online marketplaces and the value of user control'.

Fradkin, Andrey, Elena Grewal and David Holtz (2018), 'The determinants of online review informativeness: Evidence from field experiments on airbnb', *Available at SSRN 2939064* .

Garcia-López, Miquel-Àngel, Jordi Jofre-Monseny, Rodrigo Martínez-Mazza and Mariona Segú (2020), 'Do short-term rental platforms affect housing markets? evidence from airbnb in barcelona', *Journal of Urban Economics* **119**, 103278.

Goldfarb, Avi and Mo Xiao (2011), 'Who thinks about the competition? managerial ability and strategic entry in us local telephone markets', *American Economic Review* **101**(7), 3130–61.

Goolsbee, Austan and Amil Petrin (2004), 'The consumer gains from direct broadcast satellites and the competition with cable tv', *Econometrica* **72**(2), 351–381.

Hall, R. L. and C. J. Hitch (1939), 'Price theory and business behaviour', *Oxford Economic Papers, Mayo* .

Hansen, Lars Peter (1982), 'Large sample properties of generalized method of moments estimators', *Econometrica: Journal of the Econometric Society* pp. 1029–1054.

Hitsch, Günter J, Ali Hortacsu and Xiliang Lin (2019), 'Prices and promotions in us retail markets: Evidence from big data'.

Hollenbeck, Brett (2018), 'Online reputation mechanisms and the decreasing value of chain affiliation', *Journal of Marketing Research* **55**(5), 636–654.

Hortaçsu, Ali, Fernando Luco, Steven L Puller and Dongni Zhu (2019), 'Does strategic ability affect efficiency? evidence from electricity markets', *American Economic Review* **109**(12), 4302–42.

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg and Kevin R Williams (2021), 'Organizational structure and pricing: Evidence from a large us airline'.

Huang, Yufeng, Paul B Ellickson and Mitchell J Lovett (2020), 'Learning to set prices', *Available at SSRN 3267701* .

Jeziorski, Przemyslaw and Nadia Michelidaki (2019), 'Implications dynamic pricing: Experimental evidence', *http://sics.haas.berkeley.edu/pdf_2019/paper_j.pdf* .

Jin, Yizhou and Zhengyun Sun (2019), 'Information acquisition and the return to data'.

Leisten, Matthew (2020), 'Information, managerial incentives, and scale: Evidence from hotel pricing'.

Li, Hui and Kannan Srinivasan (2019), 'Competitive dynamics in the sharing economy: an analysis in the context of airbnb and hotels', *Marketing Science* **38**(3), 365–391.

Li, Jun, Antonio Moreno and Dennis Zhang (2016), 'Pros vs joes: Agent pricing behavior in the sharing economy', *Ross School of Business Paper* (1298).

Noble, Peter M and Thomas S Gruca (1999), 'Industrial pricing: Theory and managerial practice', *Marketing science* **18**(3), 435–454.

Pan, Qi; Wang, Wen (2019), 'Costly price adjustment and automated pricing: The case of airbnb'.

Pavlov, Vladimir and Ron Berman (2019), 'Price manipulation in peer-to-peer markets'.

Petrin, Amil and Kenneth Train (2010), 'A control function approach to endogeneity in consumer choice models', *Journal of marketing research* **47**(1), 3–13.

Proserpio, Davide, Wendy Xu and Georgios Zervas (2018), 'You get what you give: theory and evidence of reciprocity in the sharing economy', *Quantitative Marketing and Economics* **16**(4), 371–407.

Strulov-Shlain, Avner (2019), 'More than a penny's worth: Left-digit bias and firm pricing', *Available at SSRN 3413019* .

Tuchman, Anna E (2019), 'Advertising and demand for addictive goods: The effects of e-cigarette advertising', *Marketing Science* **38**(6), 994–1022.

Williams, Kevin (2021), 'The welfare effects of dynamic pricing: Evidence from airline markets'.

Ye, Peng, Julian Qian, Jieying Chen, Chen-hung Wu, Yitong Zhou, Spencer De Mars, Frank Yang and Li Zhang (2018), 'Customized regression model for airbnb dynamic pricing', pp. 932–940.

Zervas, Georgios, Davide Proserpio and John W Byers (2017), 'The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry', *Journal of marketing research* **54**(5), 687–705.

Zervas, Georgios, Davide Proserpio and John W Byers (2020), 'A first look at online reputation on airbnb, where every stay is above average', *Marketing Letters* pp. 1–16.

Zhang, Shunyuan, Nitin Mehta, Param Vir Singh and Kannan Srinivasan (2019), 'Can lower-quality images lead to greater demand on airbnb?'.

Appendix

A Interpolation of prices and occupancy

This section provides more details about the interpolation of price and occupancy outcomes.

Prices. I do not observe the price of the nights that are booked. Because listings are sampled once a month, the last price observation might be a few weeks before the actual booking date. However, fortunately, most prices do not vary much over time or are characterized by simple price-policy functions, making it feasible to interpolate the missing prices.

In the raw data (at the listing-night-sampling date level), 13.5% price observations are missing. I start with the set of listings which charge uniform prices, i.e. the same price across all nights for a given sampling date. For these listings (given sampling date), if I find that observed prices are uniform, I interpolate all missing prices using this uniform price.[27] This step fills in 3.9% observations, which is 29% of all missing prices and is consistent with Figure 3.

Next, I find listings-sampling date observations where all nights' prices can be characterized by one baseline plus one weekend surcharge. This is the case for 8% of missing prices, in which case I fill them in using the observed, stylized pricing policy. Similarly, I also fill in prices that vary by calendar month of the night plus a weekend surcharge (11% missing prices) and weekly prices plus a weekend surcharge (2% missing prices). At this point, half of the missing prices have been interpolated).

Further, I examine intertemporal variation in prices for the given night.[28] Specifically, I examine nights for which prices do not vary at all. I find that an additional 15% of missing prices fall into this category, in which case I interpolate the price of the night by the constant, observed prices.

By this step, 35% of the original missing-price data are still missing. I now take a stronger stance and assume that missing prices is equal to the last-observed price (i.e. price last observed before the check-in is occupied). This step fills in 19% missing prices, or 2.7% of all price data.

Occupancy. Recall that occupancy could be under-measured because some nights are last-observed weeks before the stay date, during which time it might be booked by a late-arrival customer. I leverage the fact that some nights are last-observed very close to the stay date –in which case there should be little truncation problem– and interpolate the occupancy event for dates that are (1) last observed far from the stay date, and (2) not booked when last observed. To do so, I assume that the expected occupancy rate is the same up to observed characteristics of the date, such as weekday or month of the year (seasonality). However, observed occupancy is different because of the difference in truncation. Given this assumption, I estimate a simple linear regression of binary (eventual) occupancy for listing j date τ, as a function of the degree of truncation (the number of days between the last observation and τ) interacted with month of the year and weekday of τ,

$$\text{occupancy}_{j\tau} = f\left(\text{truncation}_{j\tau}, \tau\right) + \varepsilon_{j\tau} \tag{21}$$

[27] In practice, I allow for a 0.5% standard deviation in unexplained price differences, to accommodate the possibility of a scraper error. This threshold is below $1 at the median price so can be safely ignored.

[28] I do not prioritize this dimension because, intertemporally, prices are not sampled at high frequency.

which I then parameterize by quadratic specifications of $\text{truncation}_{j\tau}$ interacted with fully saturated set of fixed effects.

The estimated $\hat{f}$ predicts, given the night τ, the expected occupancy if truncation becomes zero. Denote this difference $\Delta\text{occupancy}_{j\tau}$. I find that if the night is truncated by two weeks, occupancy rate is predicted to be 5 percentage points higher, or 9% relative to the observed (truncated) occupancy rate at 0.56. In the extreme, if the night is truncated by 30 days, occupancy rate is predicted to be 11 percentage points higher. Across all dates, occupancy rate would have been 0.69 if there were no truncation.

Lastly, for nights that I do not see occupancy, I sample binary outcomes from $\Delta\text{occupancy}_{j\tau}$, which is interpreted as "additional occupancy events" were there no truncation. I interpolate the cases where $\Delta\text{occupancy}_{j\tau}$ is one. The interpolation finds 765,996 occupancy events, 16% over a baseline of 4,680,931 cases.

B Uniform-pricing instruments: detail

To identify the price coefficient, I leverage the predominant uniform pricing and construct an exogenous price shifter based on prices of other, unrelated nights. Nights that are unrelated to the focal date have uncorrelated demand shocks, but are often set the same price. This section demonstrates the strength of the instrument and performs several robustness checks.

I estimate

$$\log\left(\text{price}_{j\tau t}\right) = \beta^1 \log\left(\bar{\text{price}}_{j,-q,t-1}\right) + \delta^1_{jq} + X_\tau \gamma^1 + u_{j\tau t} \tag{22}$$

where X_τ are characteristics of the night, including weekend, seasonality (quadratic specification of week-of-the-year), and holiday indicators, and δ^1_{jq} are listing-quarter fixed effects. This is exactly the same set of controls used in structural demand estimation. The excluded variable ("instrument") is log price of nights in other quarters observed at the previous month, $\log\left(\bar{\text{price}}_{j,-q,t-1}\right)$. Given the first stage, I estimate the second stage of the IV regression

$$\text{sale}_{j\tau t} = \alpha^2 \log\left(\hat{\text{price}}_{j\tau t}\right) + \delta^2_{jq} + X_\tau \gamma^2 + \varepsilon_{j\tau t} \tag{23}$$

where $\text{sale}_{j\tau t}$ is an indicator of whether one customer occupies night τ in month t.

I find that, given the set of controls, log price of other nights is still strongly correlated with the focal date's price. t-statistics are on the order of 100, implying that the F-statistic of the excluded variable is on the order of 10,000. I find that the second stage linear-log price coefficient is -0.467, implying an average price elasticity of -2.4 at the mean of the dependent variable at 0.217. This exercise confirms the source of identification in the structural model and that the driver of price variation is strong.

I further perform two robustness checks and one placebo test. First, one might be concerned that only some listings charge uniform prices and thus the IV recovers the *local* average price coefficient for these uniform-pricing listings, who might systematically differ from others. To address this concern, I estimate the same first and second stage regressions using a sub-sample of listings with high degrees of price variation. I take the top quartile of listings with the highest price variation (defined as the average inter-quartile range of prices across nights). Column 2 shows the results from this subset, I still find strong correlation in prices between nights across quarters, and find virtually the same price coefficient from the IV estimate. This result suggests that the across-check-in-date price variation leveraged by this IV is widely applied to most listings and thus recovers the average price effect of a representative listing.

Second, one might be concerned that the correlation across nights might come from demand rather than supply. I already focus on nights that are in different quarters of the focal date, potentially avoiding correlated demand from unobserved holidays and local events (such as a music festival) that introduce correlated demand shocks across consecutive nights. In addition, I perform a robustness check including nearby nights and show that the results are robust. Column 3 of the table shows the result using log average price of all nights (including the focal date). I find that whereas the first stage is a little stronger, potentially capturing the fact that prices of nearby dates have higher correlation, the second stage results are virtually unchanged. This result suggests that one should not worry about potential demand correlation between nearby nights. One plausible explanation is that these potential correlated demand shocks are weak and thus are averaged out in calculating the mean price across all nights.

Finally, one might further wonder how the uniform-pricing IV performs compared to using lagged prices as IV, which is often thought of as non-ideal but is still used in past and contemporary literature. In this case, a concern for using lagged price as IV is the presence of check-in-date-specific demand shocks. Column 4 shows the result of using lagged price of the focal night as IV for the current price. While this IV is very strong, the second stage estimate is only 1/3 of the preferred specification, implying a price elasticity of about -0.8. This placebo check showcases the importance of unobserved demand shifters on the night level, which is corrected for by the uniform-pricing IV but not the lagged price IV.

C Additional descriptive evidence

C.1 Multi-listing sellers are more responsive to demand shifters

Given that multi-listing sellers set more flexible prices, a natural follow-up question arises: Can their prices capture demand shifters better, or is it that they face different demand in the first place? To speak to this question, I collapse data to the level of listing (j) and nights (τ) and regress occupancy rate on *salient* characteristics of the night, such as the summer dummy (but also weekend and holidays), and the interaction between that and the log number of listings.

$$\text{occupancy}_{j\tau} = \alpha \log\left(\text{price}_{j\tau}\right) + \beta_1 \log\left(\#\text{listing}_{j\tau}\right) + \beta_2 \text{summer}_\tau + \beta_3 \log\left(\#\text{listing}_{j\tau}\right) \cdot \text{summer}_\tau + \delta_j + \lambda_{y(\tau)} + \varepsilon_{j\tau}. \tag{24}$$

I also control for the price, listing fixed effects, and calendar year fixed effects. Equation (24) examines whether the occupancy rate is higher during the summer and whether this relationship differs across sellers with different scale.[29] The main parameter of interest is whether multi-listing sellers listing face the same summer-demand shifter, i.e., to test against $\beta_3 = 0$. I also estimate a similar regression with log price on the dependent variable, to test whether multi-listing sellers set

[29]As an aside, the capacity-constrained nature of the market makes it that optimal prices depends on the level of the occupancy rate (as opposed to only depending on the price elasticity). Holding elasticity fixed, the higher the occupancy rate, the more likely the listing will be rented out early at a given price, and the higher the optimal price should be.

Appendix Table 1: Uniform-pricing instruments and alternative specifications

Panel A: first stage

	Dependent variable:			
	log(price)			
	(1)	(2)	(3)	(4)
log(avg lag price, diff quarter)	0.150***	0.114***		
	(0.001)	(0.003)		
log(avg lag price, all dates)			0.222***	
			(0.001)	
log(lag price, same date)				0.560***
				(0.0005)
weekend	0.028***	0.100***	0.028***	0.013***
	(0.0001)	(0.001)	(0.0001)	(0.0001)
holiday	−0.0004	−0.007***	−0.0005	−0.0004
	(0.0003)	(0.001)	(0.0003)	(0.0003)
days to checkin	0.001***	0.002***	0.001***	0.001***
	(0.00001)	(0.00004)	(0.00001)	(0.00001)
days to checkin squared	−0.00001***	−0.00001***	−0.00001***	−0.00001***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)
listing × checkin quarter FE	yes	yes	yes	yes
week, weekend, holiday FE	yes	yes	yes	yes
Observations	3,601,194	498,281	3,601,194	3,471,476
R^2	0.970	0.909	0.970	0.979
Note:	*p<0.1; **p<0.05; ***p<0.01			

Panel B: second stage

	Dependent variable:			
	occupancy			
	(1)	(2)	(3)	(4)
log(price)	−0.467***	−0.510***	−0.447***	−0.137***
	(0.017)	(0.048)	(0.012)	(0.003)
weekend	0.022***	0.061***	0.021***	0.013***
	(0.001)	(0.005)	(0.001)	(0.0004)
holiday	−0.005***	−0.012***	−0.005***	−0.005***
	(0.001)	(0.003)	(0.001)	(0.001)
days to checkin	−0.010***	−0.010***	−0.010***	−0.010***
	(0.00004)	(0.0001)	(0.00003)	(0.00003)
days to checkin squared	0.00005***	0.0001***	0.00005***	0.00005***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)
listing × checkin quarter FE	yes	yes	yes	yes
week, weekend, holiday FE	yes	yes	yes	yes
Observations	3,601,194	498,281	3,601,194	3,471,476
R^2	0.275	0.223	0.276	0.289
Note:	*p<0.1; **p<0.05; ***p<0.01			

Notes: first and second stage of IV estimates of sales (probability that one consumer rents the listing in a given month) on price, where the price is instrumented by the average lagged price of nights that are in different quarters of the focal date ("uniform-pricing IV"). Alternative IVs are compared.

Appendix Table 2: Price dispersion across weekends, summer, and holidays

	Dependent variable:							
	occupancy	log(price)	occupancy	log(price)	occupancy	log(price)	occupancy	log(price)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(listing)	0.010***	0.014***	0.009***	0.020***	0.009***	0.024***	0.010***	0.011***
	(0.001)	(0.0005)	(0.001)	(0.0005)	(0.001)	(0.0005)	(0.001)	(0.0005)
log(price)	−0.104***		−0.104***		−0.097***		−0.111***	
	(0.002)		(0.002)		(0.002)		(0.002)	
summer quarter	0.033***	0.027***					0.033***	0.027***
	(0.001)	(0.0003)					(0.001)	(0.0003)
summer × log(listing)	−0.008***	0.023***					−0.008***	0.023***
	(0.001)	(0.0003)					(0.001)	(0.0003)
weekend			0.021***	0.028***			0.021***	0.028***
			(0.001)	(0.0003)			(0.001)	(0.0003)
weekend × log(listing)			0.001	0.013***			0.001	0.013***
			(0.001)	(0.0002)			(0.001)	(0.0002)
holiday					−0.023***	−0.007***	−0.017***	−0.001*
					(0.001)	(0.001)	(0.001)	(0.001)
holiday × log(listing)					0.001	−0.013***	−0.0005	−0.010***
					(0.001)	(0.0005)	(0.001)	(0.0005)
listing FE	yes	yes	yes	yes	yes	yes	yes	yes
year FE	yes	yes	yes	yes	yes	yes	yes	yes
Observations	2,347,006	2,347,006	2,347,006	2,347,006	2,347,006	2,347,006	2,347,006	2,347,006
R^2	0.340	0.940	0.340	0.940	0.339	0.939	0.340	0.941
Note:								*p<0.1; **p<0.05; ***p<0.01

Notes: Estimation results of Equation (24).

different summer price premiums than single-listing sellers.[30]

Table 2 finds that, conditional on price, occupancy rate is 3 percentage points higher during the summer, 2 percentage points higher during the weekend, and 2 percentage points lower during a public holiday.[31] In addition, multi-listing sellers do not face different demand shifters, as β_3 is indistinguishable from zero (except for summer, in which case multi-listing hosts face a slightly smaller demand increase). Nevertheless, I reproduce the previous finding that multi-listing sellers set higher prices for the summer and weekends, and lower prices for holidays. These results suggest that multi-listing sellers face the same demand but can set different prices to capture demand. Hence, these results support the general picture that persistent seller differences create large degrees of heterogeneity in pricing strategies.

[30]This exercise is related to Leisten (2020), who examines hotels' ability to price in college football games (a "non-salient" demand shifter) relative to their ability to price in salient shifters, and Huang et al. (2020), who examine supermarket's ability to set prices that capture product-level idiosyncratic demand.

[31]The lower demand during holidays can potentially be explained by higher supply during holidays, reducing the residual demand for each listing.

Appendix Table 3: *Within*-seller effects of scale and experience

	Dependent variable:				
	log(end price)	std of %prices (checkin d.)	%last-month discount	%summer premium	occupancy
	(1)	(2)	(3)	(4)	(5)
1 year expr	0.012*	−0.001	0.006	0.016***	−0.022**
	(0.006)	(0.004)	(0.005)	(0.005)	(0.010)
2 years	0.016	−0.003	−0.003	0.008	−0.029**
	(0.011)	(0.005)	(0.009)	(0.006)	(0.015)
3 years	0.015	−0.002	0.001	0.010	−0.037*
	(0.015)	(0.006)	(0.012)	(0.007)	(0.019)
4+ years	0.018	−0.007	0.001	0.005	−0.047**
	(0.018)	(0.007)	(0.014)	(0.009)	(0.023)
2 listings	0.011*	0.0001	0.008*	−0.006**	−0.025***
	(0.006)	(0.003)	(0.005)	(0.003)	(0.008)
3-5 listings	0.009	0.001	0.005	0.001	−0.039***
	(0.010)	(0.004)	(0.008)	(0.004)	(0.012)
6+ listings	0.006	0.021***	−0.003	−0.0002	−0.015
	(0.020)	(0.008)	(0.015)	(0.009)	(0.024)
host FE	yes	yes	yes	yes	yes
loc-time FE	yes	yes	yes	yes	yes
loc-room type FE	yes	yes	yes	yes	yes
amenities FE	yes	yes	yes	yes	yes
Observations	70,612	71,491	70,612	70,521	69,854
R^2	0.975	0.747	0.751	0.728	0.851
Note:					*p<0.1; **p<0.05; ***p<0.01

Notes: Regression similar to Table 2, but with seller fixed effects and years of experience as covariates.

C.2 No learning

I present additional evidence, testing whether the main result in Section 3.2 is driven by within-seller variation in scale and experience. The hypothesis is that sellers might develop experience in setting prices and might increase price flexibility/sophistication over time. Appendix Table 3 confirms the earlier finding (Table 2) that almost most differences in price variability between sellers of different scale have vanished once I control for seller fixed effects. In addition, the table also shows that more experienced sellers price higher, do not price less uniformly across nights, and their prices lower less (not more) over time. These results is consistent with more experienced sellers face higher demand due to reputation effects (Hollenbeck, 2018; Fradkin et al., 2018), and do not explain the between-seller difference in the degree of sophistication in pricing.

D Model details

D.1 Demand model detail: aggregation and fixed points

Aggregation. A given day τ of a listing j is listed on the market for at most 12 months (t). In each month in expectation, $\lambda_{\tau t}^{k}$ customers of each type $k = 1,2$ will come to examine this listing and each customer has $s_{ij\tau}^{k}$ probability of booking it. Define $S_{j\tau t} = 1$ if *one of* the consumers books listing j, night τ in period t, and $A_{jt\tau} = 1$ if listing j night τ is available at the *beginning* of period t. If consumers choose independently,[32] Pan (2019) derives, under a model of homogeneous consumers, the probability that *no customers* from a given segment k books the listing if that night is available:

$$\Pr\left(S_{j\tau t} = 0 | A_{j\tau t} = 1, k\right) = \exp\left(-s_{ij\tau}^{k} \cdot \lambda_{\tau t}^{k}\right). \tag{25}$$

Based on this result, the probability that *no customer from either segment* books the listing is

$$\begin{aligned}\Pr\left(S_{j\tau t} = 0 | A_{jt\tau} = 1\right) &= \Pr\left(S_{j\tau t} = 0 | A_{j\tau t} = 1, k = 1\right) \cdot \Pr\left(S_{j\tau t} = 0 | A_{j\tau t} = 1, k = 2\right) \\ &= \exp\left(-s_{ij\tau}^{1} \cdot \lambda_{\tau t}^{1} - s_{ij\tau}^{2} \cdot \lambda_{\tau t}^{2}\right).\end{aligned} \tag{26}$$

Next, one can write down the expected occupancy rate (i.e., whether night τ of listing j is ever booked) as

$$\begin{aligned}\mathbb{E}\left[\text{occupancy}_{j\tau}\right] &= \Pr\left(S_{j\tau 1} = 1\right) + \Pr\left(S_{j\tau 2} = 1 | A_{j\tau 2} = 1\right) \cdot \Pr\left(S_{j\tau 1} = 0\right) + ... \\ &+ \Pr\left(S_{j\tau t} = 1\right) \cdot \prod_{\iota < t} \Pr\left(S_{j\tau \iota} = 0 | A_{j\tau \iota} = 1\right) + ... \\ &= \Pr\left(S_{j\tau 1} = 1\right) + \Pr\left(S_{j\tau 2} = 1 | A_{j\tau 2} = 1\right) \cdot \Pr\left(A_{j\tau 2} = 1\right) + \\ &\Pr\left(S_{j\tau 3} = 1 | A_{j\tau 3} = 1\right) \cdot \Pr\left(A_{j\tau 3} = 1\right) + ...\end{aligned} \tag{27}$$

that is, the occupancy rate is the sum of probability that the listing is booked in each period, which is in turn, the probability of being available by the start of a period and being booked in the same period.

We now discuss the invertibility of this demand system, so as to solve for δ_{jq} given realized quantity, following BLP and Berry, Ghandi and Haile (2013). Equation (27) outlines the expectation of occupancy rate conditional on a night being available at the beginning. One can write down the sample analog of this expectation,

$$\frac{1}{|q|} \sum_{\tau \in q} \text{occupancy}_{j\tau} = \frac{1}{|q|} \sum_{\tau \in q} \left(\Pr\left(S_{j\tau 1} = 1\right) + \Pr\left(S_{j\tau 2} = 1 | A_{j\tau 2} = 1\right) \cdot A_{j\tau 2} + ...\right) \tag{28}$$

which gives the fixed point equation (8) in Section 4.1.

Nested fixed point algorithm and estimation. I estimate demand using data each listing j and for each available night τ, over all the months t of observations during which the listing is available (up to 12 months). The probability of the listing being booked in month τ given availability at the

[32] I assume away variations in the choice set within the period t.

beginning of that month, is $\Pr\left(S_{j\tau t}=1|A_{j\tau t}=1\right)$ given by Equation (26). Likelihood function is simply

$$\text{likelihood}=\prod_{j,\tau,t}\Pr\left(S_{j\tau t}=1|A_{j\tau t}=1\right)^{S_{j\tau t}}\cdot\Pr\left(S_{j\tau t}=0|A_{j\tau t}=1\right)^{1-S_{j\tau t}}. \tag{29}$$

To compute the objective function at each set of trial parameters (α,σ,γ), we iterate Equation (8) and solve for all δ_{jq}'s as a function of these parameters, then use the δ's to compute the likelihood. The outer loop then finds the set of parameters that maximizes the log likelihood. The fixed point computation is costly but the compute time concentrates on many fixed-point iterations of Equations (5)-(8), with mostly the same data (but different parameter values). These computation tasks can be vastly accelerated in a graphical-processing unit (GPU).

D.2 Demand model detail: identification of arrival rates from the demand intercept

A challenge is to separate $\delta_{jq(\tau)}$ from $\lambda^k_{m\tau t}$, or to separate customer arrival rate from preferences in explaining a given that I only observe the occupancy rate. One needs normalizations. Specifically, I normalize $\gamma^1_{m0}=1000$ for zipcode 94110 (the largest market in terms of total Airbnb rentals). That is, for segment 1, I assume 1000 customers will arrive in the *last* month right before the stay date (if the day is not a holiday or weekend, and on January first when the week-of-the-year effect is zero) in this particular market. For every other zipcode, I assume segment 1's last-month arrival rate is proportional to that of 94110's, based on the total number of observed Airbnb rentals. Finally, I assume, for segment 2, $\lambda^2_{t,\tau}=0$ for $\tau\leq 4$. That is, there are no segment 2 customers in before eight months to check-in. This normalization should separate the baseline demand intercept from the customer arrival rate.

Further, recall that arrival is heterogeneous, and so are some utility parameters such as the price intercept. Separating the heterogeneous preferences and arrival rates further relies on the following arguments. First, the distribution of timing of when listings are booked identifies γ^k_1 given the basaeline γ^k_{0m}. Second, the size of segment 2 is identified by changes in the sensitivity to price, weekend, and holiday. If the price sensitivity changes very little, then one rationalizes the data mostly by segment 1. However, if the average price sensitivity changes by a lot, one would rationalize the fraction of segment 2 (i.e. $\gamma^2_{0m}/\gamma^1_{0m}$), together with the differences in the price sensitivity (i.e. $\alpha^2-\alpha^1$), by the empirical pattern of how price sensitivity changes over time.

D.3 Supply: Optimal dynamic pricing with probablistic inaction

Denote the probability of a sale for listing-stay $j\tau$ in period t (i.e. the probability that *one of the customers* book listing j for night τ, during month t) as

$$\begin{aligned} q_{j\tau t} &:= 1-\Pr\left(S_{j\tau t}=0|A_{j\tau t}=1\right) \\ &= 1-\exp\left(-s^0_{ijt}\cdot\lambda^0_{m\tau t}-s^1_{ijt}\cdot\lambda^1_{m\tau t}\right), \end{aligned} \tag{30}$$

and the individual choice probability for each segment $k = 1,2$ is

$$s_{j\tau t}^{k} = \frac{\exp\left(\delta_{jq} + \alpha^{k}\log\left((1+r)\cdot p_{j\tau t}\right) + \eta_{j\tau t}\right)}{1+\exp\left(\delta_{jq} + \alpha^{k}\log\left((1+r)\cdot p_{j\tau t}\right) + \eta_{j\tau t}\right) + \omega_{jt\tau}^{k}}$$
$$= \left(1 + \frac{1+\omega_{j\tau t}^{k}}{\exp\left(\delta_{jq} + \alpha^{k}\log\left((1+r)\cdot p_{j\tau t}\right) + \eta_{j\tau t}\right)}\right)^{-1}. \tag{31}$$

where, following the idea in Pan (2019), we denote the sum of exponential utility of other listings as the market-wide states $\omega_{j\tau t} = \left(\omega_{j\tau t}^{1}, \omega_{j\tau t}^{2}\right)$:[33]

$$\omega_{j\tau t}^{k} = \sum_{j'\neq j} \exp\left(\delta_{j'q} + \alpha^{k}\log\left((1+r)\cdot p_{j'\tau t}\right) + \eta_{j'\tau t}\right).$$

Also note that the arrival rate $\lambda_{j\tau t}$ contains time-invariant states about the night, such as whether it is a weekend, holiday, or the effect of seasonality. These states, as well as the lead time τ, are important states that drive the pricing decisions. Relatively, the effect of $\eta_{j\tau t}$ is small and thus I set all η to zero when computing optimal prices.

Now, denote the price path by an optimizing listing j with attention probability μ_j as $\tilde{p}_{j\tau t}$. These prices are set according to and one can solve the problem backwards.

Illustrating example. I start by illustrating the solution of the problem with two periods. In period $T = 12$, the problem is static because the continuation value is zero. Hence we have

$$\max_{p} \pi_{j\tau,12}\left(p, \omega_{j\tau,12}\right). \tag{32}$$

The first-order condition is

$$p_{j\tau,12}^{*} = \frac{c_j}{1-f} - \left(\frac{\partial q_{j\tau,12}}{\partial p_{j\tau,12}}\right)^{-1} q_{j\tau,12}. \tag{33}$$

In period $T-1 = 11$, the problem is different in two ways. First, if the room is not occupied in this period, it can still be listed in the next period, creating an option value that drives the prices higher. Second, if the manager gets a chance to act, she knows that she might not get another chance to act next period, making it so that her choices today is partially tied to her payoff tomorrow. The value function in this period reflects these two elements:

$$\begin{aligned} V_{j\tau,11} = \max_{p} q_{j\tau,11}\cdot\left(p\cdot(1-f)-c_j\right) + \left(1-q_{j\tau,11}\right)\cdot \\ \left(\mu_j\mathbb{E}\left[V_{j\tau,12}|\omega_{j\tau,11}\right] + \left(1-\mu_j\right)\mathbb{E}\left[q_{j\tau,12}|\omega_{j\tau,11}\right]\left(p\cdot(1-f)-c_j\right)\right) \end{aligned} \tag{34}$$

where, with one minus the probability of a sale, the manager gets her expected payoff renting the listing in month 12 (the option value). However, with probability $1-\mu_j$, she does not get a chance to change the price and would rent at the same price as she sets now (with μ_j she enters the optimal

[33] The two ω's are highly correlated given the variation comes from the set of available competitors.

decision problem in period 12). Collect terms and take the first-order condition and one gets

$$\left(\frac{\partial q_{j\tau,11}}{\partial p}+\frac{\partial\left(\left(1-q_{j\tau,11}\right)\cdot\left(1-\mu_j\right)\mathbb{E}\left[q_{j\tau,12}|\omega_{j\tau,11}\right]\right)}{\partial p}\right)\left(p\cdot(1-f)-c_j\right)+$$
$$q_{j\tau,11}\cdot(1-f)+\left(1-q_{j\tau,11}\right)\left(1-\mu_j\right)\mathbb{E}\left[q_{j\tau,12}|\omega_{j\tau,11}\right]\cdot(1-f)+\frac{\partial\left(\left(1-q_{j\tau,11}\right)\mu_j\right)}{\partial p}\mathbb{E}\left[V_{j\tau,12}|\omega_{j\tau,11}\right]=0, \tag{35}$$

which is closed-form once we solve for $V_{j\tau,12}$.

A note on the expectation operator. I approximate the expectation by drawing ω's from the empirical distribution.

General problem. In general, one can generally write down the value function as follows, suppressing j and τ subscripts:

$$V_t=\max_p(1-f)\left(q_t+(1-\mu)(1-q_t)\mathbb{E}\left[\sum_{\iota=t+1}^{T}(1-\mu)^{\iota-t-1}\left(\prod_{\iota'=t+1}^{\iota-1}(1-q_{\iota'})\right)q_\iota|\omega_t\right]\right)\left(p-\frac{c}{1-f}\right)+$$
$$(1-q_t)\mathbb{E}\left[\sum_{\iota=t+1}^{T}(1-\mu)^{\iota-t-1}\left(\prod_{\iota'=t+1}^{\iota-1}(1-q_{\iota'})\right)\mu V_\iota|\omega_\tau\right], \tag{36}$$

this value function can be solved in closed-form via backward induction. Denote

$$\Delta_t=(1-f)\left(q_t+(1-\mu)(1-q_t)\mathbb{E}\left[\sum_{\iota=t+1}^{T}(1-\mu)^{\iota-t-1}\left(\prod_{\iota'=t+1}^{\iota-1}(1-q_{\iota'})\right)q_\iota|\omega_t\right]\right) \tag{37}$$

and

$$\Omega_t=(1-q_t)\mathbb{E}\left[\sum_{\iota=t+1}^{T}(1-\mu)^{\iota-t-1}\left(\prod_{\iota'=t+1}^{\iota-1}(1-q_{\iota'})\right)\mu V_\iota|\omega_t\right] \tag{38}$$

and one can write the FOC in a simple term

$$p_t^*=\frac{c}{1-f}-\left(\frac{\partial\Delta_t}{\partial p}\right)^{-1}\left(\Delta_t+\frac{\partial\Omega_t}{\partial p}\right). \tag{39}$$

D.4 Counterfactual: Seller pricing with known price-adjustment factors

One counterfactual I entertain is what happens when the platform first proposes a pre-committed set of price-adjustment factors (say summer premium or last-minute discounts), and based on these adjustment factors, sellers set one price for each listing. This section outlines the seller optimal pricing under this counterfactual scenario.

Denote seller price as $\bar{p}_{jq}$ and platform's adjustment factor as $a_{ml\tau t}$ (for market m, listing-type l, night τ, and based on time before checkin $t-\tau$). The final price before tax and consumer fee is

$\left(1+a_{ml\tau t}\right)\bar{p}_{jq}$. We can write the seller objective as (suppress m and l)

$$\max_p \sum_\tau \Bigg(q_{j\tau,1}\left(\left(1-f\right)\left(1+a_{ml\tau,1}\right)p-c_j\right)+\left(1-q_{j\tau,1}\right)q_{j\tau,2}\left(\left(1-f\right)\left(1+a_{ml\tau,2}\right)p-c_j\right)+$$

$$\left(1-q_{j\tau,1}\right)\left(1-q_{j\tau,2}\right)q_{j\tau,3}\left(\left(1-f\right)\left(1+a_{ml\tau,3}\right)p-c_j\right)+\ldots+\prod_{\iota=1,\ldots,11}\left(1-q_{j\tau,\iota}\right)q_{j\tau,12}\left(\left(1-f\right)\left(1+a_{ml\tau,12}\right)p-c_j\right)\Bigg)$$

$$=\max_p \sum_\tau \left(\sum_{t=1}^{T}\prod_{\iota\leq t-1}\left(1-q_{j\tau,\iota}\right)q_{j\tau,t}\left(\left(1-f\right)\left(1+a_{ml\tau,t}\right)p-c_j\right)\right).$$

Take the first-order condition and we arrive at the "base price" of each listing at

$$p_j^{\text{counterfactual}}=-\frac{\sum_\tau\sum_t\left(\left(1-f\right)\left(1+a_{ml\tau,t}\right)\prod_{\iota\leq t-1}\left(1-q_{j\tau,\iota}\right)q_{j\tau,t}-\left(\frac{\partial q_{j\tau,t}}{\partial p}\prod_{\iota\leq t-1}\left(1-q_{j\tau,\iota}\right)-\sum_{\iota'\leq t-1}\frac{\partial q_{j\tau,\iota'}}{\partial p}\prod_{\iota\neq\iota'}\left(1-q_{j\tau,\iota}\right)q_{j\tau,t}\right)\times c_j\right)}{\sum_\tau\sum_t\left(\frac{\partial q_{j\tau,t}}{\partial p}\prod_{\iota\leq t-1}\left(1-q_{j\tau,\iota}\right)-\sum_{\iota'\leq t-1}\frac{\partial q_{j\tau,\iota'}}{\partial p}\prod_{\iota\neq\iota'}\left(1-q_{j\tau,\iota}\right)q_{j\tau,t}\right)\left(1-f\right)\left(1+a_{ml\tau,t}\right)}.$$

E Appendix: Additional tables and figures

Appendix Table 4: Sellers' number of listings on other measures of price variation

	Dependent variable:			
	cond. std of %prices	freq. of price change	%weekend premium	%holiday premium
	(1)	(2)	(3)	(4)
2 listings	0.003**	0.027***	0.007*	−0.001
	(0.001)	(0.008)	(0.004)	(0.001)
3-5 listings	0.010***	0.055***	0.007*	−0.006***
	(0.002)	(0.010)	(0.004)	(0.001)
6+ listings	0.017***	0.021	0.014	−0.014***
	(0.004)	(0.022)	(0.011)	(0.004)
baseline Y	0.016	0.192	0.032	-0.008
seller FE	no	no	no	no
loc-time FE	yes	yes	yes	yes
loc-room type FE	yes	yes	yes	yes
amenities FE	yes	yes	yes	yes
Observations	71,550	65,737	70,484	70,767
R^2	0.553	0.336	0.559	0.505

Note: *p<0.1; **p<0.05; ***p<0.01

Notes: See notes for Table 2.

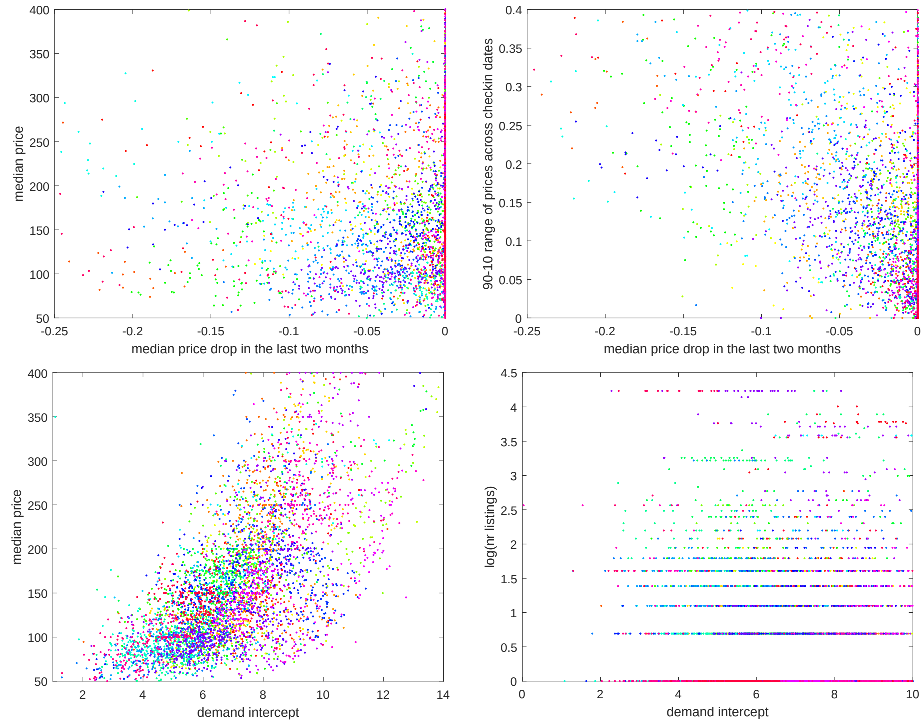

Appendix Figure 1: Hierarchical clustering of price level, price change, demand, and seller scale

Notes: These figures illustrate the clustering of listings based on the moments described in Section 4.5.

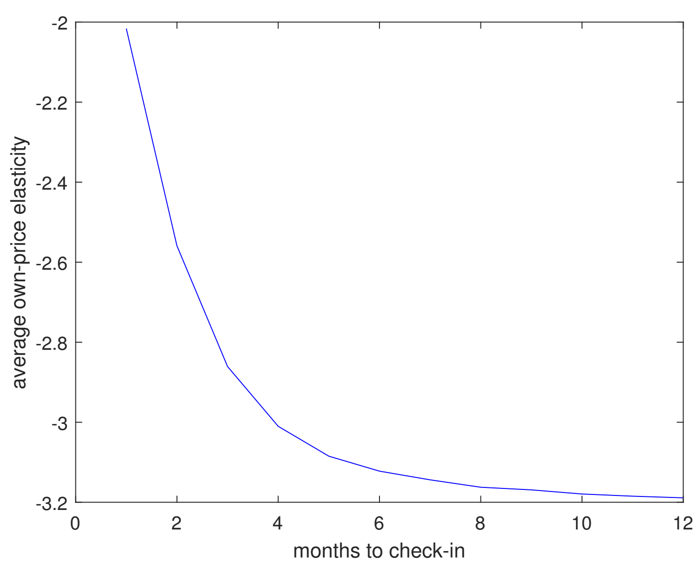

Appendix Figure 2: Average own-price elasticity over time-to-check-in

Notes: Implied elasticities over time-to-check-in.

Appendix Table 5: Supply-side estimates: top 75 segments

	% of sample	marginal cost ($)	std err	prob(adjust price) (μ, %)	std err	expected nr. unique prices (ρ)	std err	dynamic pricing (θ, %)	std err	fixed costs ($)	std err
group001	2.9	39.7	0.5			0.7	0.1	0.0	2.5	6473.6	367.3
group002	2.8	16.7	0.3	100.0	3.6			100.0	1.0	6870.9	391.1
group003	2.6	24.5	0.4	0.0	0.0	0.3	0.1	8.8	0.1	5960.9	376.5
group004	2.2	17.0	0.5	0.0	4.5	0.2	0.4	25.3	3.6	7455.9	440.9
group005	2.0	26.1	0.4	0.0	18.2	0.5	0.3	12.2	0.9	5894.6	399.2
group006	1.9	23.0	0.4	0.3	12.3	0.0	0.7	31.2	4.0	5771.6	422.3
group007	1.9	25.2	0.4	18.0	2.6	0.2	1.0	70.9	3.9	6792.2	470.8
group008	1.8	20.7	0.0			0.3	0.0	0.0	0.1	7096.7	499.5
group009	1.8	87.7	0.3	0.0	0.0	0.0	0.0	24.3	0.4	6686.4	437.1
group010	1.8	67.4	0.4	15.8	2.2	31.0	5.2	51.9	5.3	6066.6	389.2
group011	1.7	38.6	0.3	100.0	2.1			100.0	0.7	6011.8	413.5
group012	1.5	74.3	0.6	0.0	3.1	0.0	0.0	31.0	2.7	7536.9	508.1
group013	1.4	72.0	0.4	0.0	3.0	0.8	0.2	32.2	3.1	5449.0	411.7
group014	1.3	38.9	0.7	99.7	195.2	25.5	6.8	2.0	4.4	4839.0	421.3
group015	1.3	20.6	0.4	0.0	5.9	0.9	1.0	29.6	7.8	5043.1	424.9
group016	1.2	34.2	0.4	100.0	14.8	0.9	0.2	25.9	2.5	6505.6	480.1
group017	1.2	123.4	0.5	8.8	4.0	31.0	14.2	56.5	3.9	6491.5	440.1
group018	1.2	92.2	0.4	14.1	4.0	4.1	2.3	52.9	5.3	6422.9	455.1
group019	1.1	27.2	0.5	99.9	46.1	1.0	0.1	8.2	1.7	6182.1	438.8
group020	1.1	25.9	1.1	0.0	57.2	1.5	0.2	5.0	9.0	4195.7	440.9
group021	1.1	55.0	0.5	100.0	53.8	12.7	3.2	11.8	7.0	5765.2	492.6
group022	1.1	25.9	0.5	100.0	3.8	0.0	26.9	99.5	2.8	6082.9	506.5
group023	1.1	12.9	0.4	100.0	7.3			100.0	1.0	3229.4	406.4
group024	1.0	37.5	0.5	0.0	7.6	0.0	0.8	25.1	7.7	6049.1	503.9
group025	1.0	25.2	0.7	0.0	9.8	1.1	0.7	22.8	4.5	4631.2	458.0
group026	1.0	18.4	0.5	99.3	66.7	1.5	0.6	11.3	9.4	4701.1	488.1
group027	1.0	33.4	0.4	95.3	1151.5	1.3	0.5	0.8	4.4	5323.6	493.0
group028	1.0	84.8	0.0	0.0	0.0	31.0	0.0	34.3	0.0	4535.3	472.6
group029	1.0	218.7	0.4	1.7	3.2			100.0	3.6	5943.3	520.3
group030	1.0	41.9	0.3	100.0	80.6	6.6	1.1	4.6	0.4	3044.8	435.1
group031	1.0	42.5	0.3	100.0	5.5	0.1	0.2	47.5	2.8	2790.2	417.6
group032	0.9	25.6	0.4	81.3	49.2	0.7	0.7	21.1	5.2	4121.2	446.6
group033	0.9	67.1	0.4			7.4	0.5	0.0	9.1	4111.5	436.9
group034	0.9	12.9	0.5	100.0	3.0	26.5	51.2	96.0	2.6	4244.6	425.8
group035	0.9	35.8	0.4	100.0	7.2	21.2	6.9	55.0	3.0	4346.4	471.0
group036	0.9	2.3	0.7	55.1	22.5	2.1	0.4	15.2	3.5	3599.7	442.8
group037	0.9	21.8	0.5	100.0	14.8	4.5	1.5	40.7	3.7	3295.2	441.2
group038	0.8	12.6	0.6	100.0	5.6	1.5	0.3	50.5	2.3	5656.3	513.5
group039	0.8	23.4	0.5	0.7	7.6	0.0	1.4	60.8	11.5	6018.2	535.1
group040	0.8	53.4	0.4	100.0	4.7	10.4	1.6	44.1	1.8	4153.1	455.6
group041	0.8	37.8	0.5	100.0	63.6	31.0	4.2	5.9	4.5	3325.2	437.4
group042	0.8	93.8	0.5	85.6	324.1	9.5	1.3	2.4	5.4	6734.1	530.6
group043	0.8	68.4	0.0	0.0	0.0			100.0	0.0	4395.5	472.7
group044	0.8	14.1	0.3	37.2	4.2	0.0	0.5	80.9	2.4	3548.0	458.6
group045	0.8	7.9	0.4	100.0	2.4			100.0	0.4	3870.8	454.4
group046	0.8	29.5	0.5	1.6	3.1	16.3	7.3	38.6	3.8	6141.9	559.3
group047	0.7	39.9	0.5	0.0	8.7			100.0	4.0	7349.6	653.0
group048	0.7	106.8	0.4	6.0	1.3	0.1	0.3	79.2	2.2	4400.9	526.2
group049	0.7	150.9	0.3	100.0	1.7			100.0	0.4	6367.5	580.7
group050	0.7	101.1	0.6	71.2	3.3			100.0	2.7	15097.1	2110.4
group051	0.7	31.0	0.4	100.0	2.5	1.1	13.0	98.9	2.1	5982.6	532.2
group052	0.7	914.4	0.0	50.9	0.0	15.7	Inf	49.9	0.0	5162.4	1138.7
group053	0.7	335.0	0.3			0.0	0.0	0.0	0.0	5629.0	593.5
group054	0.7	124.3	0.3	100.0	4.9			100.0	2.3	3836.4	587.7
group055	0.7	21.3	0.4	0.0	547.0	0.5	0.5	0.8	10.9	7080.0	679.1
group056	0.7	20.1	0.3	30.7	5.8	2.1	1.1	44.4	6.0	6886.8	633.8
group057	0.7	120.0	0.4	6.0	1.9			100.0	4.5	7908.0	723.3
group058	0.7	-4.2	0.0	10.2	0.6	0.1	0.1	5.8	0.0	10952.6	987.3
group059	0.7	87.0	0.4	88.2	110.9	27.3	1.0	4.9	3.1	5411.0	573.5
group060	0.6	59.8	0.8			3.2	0.4	0.0	4.3	5305.9	587.3
group061	0.6	119.5	0.4			0.8	0.3	0.0	1.6	5765.0	589.8
group062	0.6	47.8	0.4	0.7	2.3	31.0	4.6	52.6	6.5	4621.7	562.6
group063	0.6	38.4	0.5	4.7	2.2	2.6	0.3	45.2	1.8	6079.6	553.9
group064	0.6	24.7	0.4	100.0	10.8			100.0	3.8	4674.3	557.4
group065	0.6	72.5	0.5	0.1	2.0			100.0	2.5	4265.6	486.8
group066	0.6	35.0	0.4	95.1	2.5	11.6	1.0	50.2	1.2	2368.1	540.4
group067	0.6	40.6	0.3	100.0	5.1	1.1	40.1	99.2	3.3	2025.2	541.9
group068	0.6	63.2	0.5	8.1	2.8	30.8	27.8	57.4	23.1	3683.6	537.2
group069	0.6	20.4	0.3	66.3	6.9	2.4	0.2	28.3	2.6	2690.6	509.0
group070	0.6	91.4	0.4	0.0	18.9			100.0	2.1	5817.6	683.9
group071	0.6	35.4	0.4	100.0	4.9			100.0	1.4	2702.8	553.9
group072	0.6	23.2	0.8	0.0	9.0			100.0	76.5	4378.0	600.5
group073	0.5	52.6	0.8	53.9	485.5	10.2	11.3	1.3	22.7	8130.3	812.1
group074	0.5	40.2	0.4	13.0	2.4	3.5	2.2	74.9	3.2	5887.8	708.2
group075	0.5	50.7	3.3	0.0	3.0	0.1	24.5	94.9	12.7	4942.2	588.2

Notes: Supply-side estimates for the top-75 segments. Estimates of ρ is not presented when $\theta \to 1$, and estimates of μ is not presented when $\theta \to 0$.

Appendix Table 6: Supply-side estimates: bottom 75 segments

	% of sample	marginal cost ($)	std err	prob(adjust price) (μ, %)	std err	expected nr. unique prices (ρ)	std err	dynamic pricing (θ, %)	std err	fixed costs ($)	std err
group076	0.5	64.7	0.4	100.0	0.6			100.0	0.3	5476.1	649.5
group077	0.5	36.5	0.4	3.2	5.0	31.0	17.5	31.6	10.6	4805.3	567.9
group078	0.5	22.3	1.4	0.0	7.6	0.5	3.3	65.7	24.7	6539.7	656.1
group079	0.5	90.5	0.6	99.9	2.5	30.8	6.3	39.7	5.2	3921.8	619.5
group080	0.5	113.5	0.3	0.1	1.1			100.0	1.3	4606.7	566.2
group081	0.5	34.4	0.4	99.9	91.0	18.1	5.6	7.8	6.2	2759.9	507.7
group082	0.5	28.4	0.4	15.2	2.5	9.7	45.5	98.0	1.3	1966.0	556.7
group083	0.5	150.6	0.3	100.0	4.0			100.0	1.5	4181.2	578.4
group084	0.5	62.7	0.4	100.0	2.6	0.0	2.6	94.7	2.2	3130.0	575.1
group085	0.5	46.1	0.7	2.5	59.0	11.8	30.1	21.9	13.2	4962.1	740.7
group086	0.5	30.7	0.4	80.2	4.4	3.5	0.3	30.2	1.5	2407.5	592.5
group087	0.5	34.6	0.0	0.0	0.0	31.0	0.0	98.3	0.0	6777.4	847.2
group088	0.4	41.6	0.5	7.0	1.9	0.0	5.2	95.6	3.7	3728.9	617.2
group089	0.4	207.1	0.4	38.5	2.4	31.0	65.6	98.7	1.9	7459.2	885.8
group090	0.4	72.7	0.4	98.8	1.1	6.1	5.9	93.2	1.0	2485.0	607.7
group091	0.4	39.9	0.4	57.7	7.5			100.0	7.9	2733.3	609.5
group092	0.4	30.0	0.3	100.0	1.4			100.0	0.2	1412.1	744.9
group093	0.4	20.7	0.3	100.0	4.9			100.0	0.6	5518.5	652.1
group094	0.4	188.7	0.7			3.3	2.1	0.1	39.0	7320.5	1133.2
group095	0.4	3.1	0.7	2.0	11.2	0.3	0.3	6.7	8.3	5970.4	735.4
group096	0.4	55.5	0.7	1.7	16.9			100.0	33.2	5198.1	1235.9
group097	0.4	49.1	0.6	100.0	112.8	17.7	21.8	8.0	11.9	10545.4	1221.4
group098	0.4	23.8	0.4	78.5	3.8	0.9	0.3	47.2	2.3	1620.0	665.2
group099	0.4	29.0	0.5	100.0	4.8			100.0	9.1	2844.7	787.4
group100	0.4	39.3	0.5	39.2	3.8	30.8	160.4	96.6	12.1	3611.0	658.8
group101	0.4	32.1	0.4	19.0	9.6	3.6	1.4	28.3	6.5	2899.3	655.5
group102	0.4	10.8	0.3	3.8	13.0	2.1	1.0	32.8	11.3	4710.2	659.7
group103	0.4	53.2	0.4	30.7	1.7			100.0	3.0	2580.4	628.6
group104	0.4	6.2	0.4	100.0	22.5	0.8	0.7	48.7	5.3	4683.6	732.7
group105	0.4	145.9	0.6	68.6	12.7			100.0	3.4	3833.4	756.9
group106	0.4	29.9	0.4	78.1	5.6	27.3	9.5	53.5	2.6	4317.5	699.4
group107	0.3	44.5	0.4	98.5	2.9			100.0	1.5	10006.7	1256.4
group108	0.3	54.3	0.4	29.9	2.5			100.0	3.1	1830.2	606.2
group109	0.3	54.6	0.5	0.0	5.3			100.0	10.0	1596.7	649.3
group110	0.3	40.0	0.4	100.0	2.7	31.0	12.1	93.8	1.9	3530.6	746.4
group111	0.3	127.6	0.3	0.0	0.8			100.0	1.7	6967.0	871.3
group112	0.3	166.7	0.7	2.3	10.7			100.0	7.9	4972.7	714.8
group113	0.3	62.2	0.5	21.4	2.5			100.0	4.5	8850.5	1267.9
group114	0.3	28.0	0.5	14.3	23.4	16.8	19.6	22.2	22.7	5722.3	878.2
group115	0.3	50.7	0.5	19.8	4.3			100.0	2.9	4566.6	713.2
group116	0.3	24.3	0.3	100.0	2.9			100.0	0.3	3460.1	825.5
group117	0.3	49.2	0.4	0.0	0.3			100.0	0.3	9986.6	832.5
group118	0.3	14.9	0.6	100.0	9.0	0.0	0.1	53.0	1.3	4100.7	762.4
group119	0.3	46.0	0.7	0.0	0.9	8.6	1814.7	99.7	2.4	6915.0	877.7
group120	0.3	153.3	0.6	0.0	3.2			100.0	2.1	8587.6	1254.9
group121	0.3	110.0	0.4	0.0	2.1			100.0	0.7	3823.7	715.8
group122	0.3	275.6	0.4	100.0	0.7			99.8	1.9	2831.0	798.9
group123	0.3	308.9	0.4	100.0	0.6			100.0	0.5	10434.3	1357.2
group124	0.2	59.9	0.3	49.1	2.1			100.0	3.0	3205.4	859.7
group125	0.2	196.6	0.6	0.0	0.3			100.0	1.3	6840.0	575.4
group126	0.2	195.1	0.6	0.0	2.5			100.0	2.1	12483.2	2117.0
group127	0.2	60.0	0.5	79.7	1.5	31.0	2.0	67.3	1.6	4459.2	819.4
group128	0.2	117.0	0.4	0.0	2.1			100.0	1.5	15113.0	2549.5
group129	0.2	115.8	0.4	0.0	1.7			100.0	2.2	11369.8	1680.2
group130	0.2	76.6	0.3	10.0	0.9			100.0	1.3	3465.6	987.6
group131	0.2	23.1	0.4	100.0	0.8	30.1	14.0	70.3	2.1	2507.3	876.4
group132	0.2	67.5	0.4	100.0	1.9			100.0	0.4	9307.9	1241.2
group133	0.2	53.6	0.6	66.4	1.4	7.6	10.0	89.9	2.3	5216.5	952.0
group134	0.2	36.0	0.5	59.0	2.5	28.8	29.9	79.2	4.3	509.3	1011.3
group135	0.2	21.0	0.4	21.6	5.2	31.0	15.3	64.9	6.6	7471.7	1438.0
group136	0.2	66.8	0.4	0.0	0.2			100.0	0.5	6952.2	1137.5
group137	0.1	35.2	0.4	68.4	0.4			100.0	0.5	3045.9	1145.8
group138	0.1	100.6	0.4	39.6	3.2			100.0	4.1	6366.8	1400.6
group139	0.1	89.0	0.7	98.1	1.0			100.0	7.8	6840.0	575.4
group140	0.1	68.0	0.0	53.3	0.0	6.5	Inf	46.4	0.0	15443.9	2556.6
group141	0.1	15.4	0.0	53.9	0.0	2.8	Inf	53.9	0.0	6840.0	575.4
group142	0.1	29.8	0.4	19.8	3.7			100.0	0.8	11587.9	2571.9
group143	0.1	13.3	0.4	79.9	1.5			99.9	1.1	5789.5	1058.6
group144	0.1	104.8	0.4	100.0	1.5			100.0	0.2	3378.5	1104.3
group145	0.1	57.2	0.4	22.4	1.1			100.0	1.8	440.9	1636.5
group146	0.1	93.4	0.4	100.0	0.4	31.0	91.5	98.1	3.0	6840.0	575.4
group147	0.1	302.2	0.7	50.7	2.5			100.0	5.0	6840.0	575.4
group148	0.1	44.2	0.4	91.1	3.7			100.0	1.8	6840.0	575.4
group149	0.0	176.8	0.5	100.0	0.4			100.0	0.8	6840.0	575.4
group150	0.0	-50.0	19.5	100.0	326.5	0.0	0.0	52.7	2.7	6840.0	575.4

Notes: Supply-side estimates for the bottom-75 segments.

Appendix Table 7: Fixed costs: auxiliary parameter estimates

	par est	std err
post 2018 regulation	-0.902	0.103
distance to union sq	-3.338	0.245
distance squared	0.790	0.065
scale of the fixed cost error	3.438	0.152

Notes: Reports fixed cost parameters except for segment-specific average fixed cost (summarized in Figure 8).