

Managing Market Thickness in Online B2B Markets

Kostas Bimpikis

Graduate School of Business, Stanford University, Stanford, CA 94305
kostasb@stanford.edu

Wedad J. Elmaghraby

Robert H. Smith School of Business, University of Maryland, College Park, MD 20742
welmaghr@rhsmith.umd.edu

Ken Moon

The Wharton School, University of Pennsylvania, Philadelphia, PA 19104
kenmoon@wharton.upenn.edu

Wenchang Zhang

Robert H. Smith School of Business, University of Maryland, College Park, MD 20742
wzhang@rhsmith.umd.edu

We explore marketplace design in the context of a B2B platform specializing in liquidation auctions. This setting enables us to illustrate the value a platform can generate by appropriately managing its *effective* availability of supply over time while demand endogenously determines *when* to participate in the marketplace. This value remains sizable even though the platform's *aggregate* levels of supply and potential demand are not affected by the design changes we consider. Specifically, we study two design levers: (i) the platform's *listing policy*, which sets the ending times of incoming auctions (hence, the frequency of market clearing). Exploiting a natural experiment, we illustrate that consolidating auctions' ending times to certain weekdays increases the platform's revenues by 7.3% mainly by inducing a higher level of bidder participation. (ii) The second lever is a *recommendation system* that can be used to reveal information about real-time market thickness to potential bidders. Both levers use the platform's availability of supply to incentivize participation from buyers who decide strategically when/how often to participate. Thus, their optimization highlights a novel trade-off. Namely, when the platform consolidates auctions' ending times, more bidders may participate in the marketplace (demand-side competition); but ultimately auctions for substitutable goods cannibalize one another (supply-side competition). To optimize these design decisions, we estimate a structural model that endogenizes bidders' dynamic behavior, i.e., their decisions on whether/how often to participate in the marketplace and how much to bid. We find that appropriately designing a recommendation system yields an additional revenue increase (on top of the benefits obtained by optimizing the platform's listing policy) by reducing supply-side cannibalization and altering the composition of participating bidders.

Key words: Online markets; Market clearing; Market thickness, Matching supply with demand; Natural experiment; Structural estimation.

1. Introduction

The emergence of Internet-enabled platforms, such as Airbnb and Lyft, has highlighted that online marketplaces greatly reduce frictions that prevent buyers and sellers in previously disparate spheres from connecting, thereby increasing the volume of trade in a number of markets. Typically, such platforms neither own nor directly control the goods involved in each transaction, but act as intermediaries. Thus, their success relies heavily on the design features of their respective marketplaces, e.g., the ways in which they organize and present information to the buyers and the timing with which they match and clear (portions of) the market.

The opportunity for online intermediaries to create value has manifested itself not only in the cases exemplified by Airbnb and Lyft, but has also reshaped retail operations, particularly with regard to the handling and resale of liquidation inventory. The present paper explores marketplace design in the context of an e-commerce platform specializing in liquidation inventory including merchandise that either remained unsold in its primary market (e.g., due to low demand levels) or was returned by customers. The secondary market within which the platform operates is of great economic significance: roughly 20% of inventory goes unsold in the fast fashion industry ([Ferdows et al. 2005](#)), and whereas brick-and-mortar retailers encounter a 9% return rate on products – for online retailers, the return rate is a staggering 30%.¹ Overall, it is estimated that in 2012, the size of this excess/return product market was \$424 billion or 2.9% of the entire US GDP. However, given the uncertainty surrounding the volume, quality, and composition of their excess and returned merchandise, retailers have come to expect mere cents-on-the-dollar recovery rates from traditional channels. Thus, they typically offload this inventory to business buyers further down the retail food chain, such as discount stores, or donate it to qualifying recipients for tax purposes.

Online business-to-business (B2B) auction platforms connect an increasing number of retailers to deeper pools of potential business buyers, both domestic and foreign. Given the diversity of potential bidders, which range from large wholesale liquidators to small mom-and-pop stores, online auctions crucially facilitate price discovery and constitute one of the major sale mechanisms in secondary markets. In 2016, Liquidity Service Inc., one of the fastest-growing online B2B auction platforms, sold merchandise worth more than \$600 million in aggregate retail value.² Today, many US chain retailers, including Best Buy, Walmart, Home–Depot, Amazon, Target, and Costco, utilize online B2B auction platforms to liquidate their products.

The amount of inventory sold in these online platforms is highly variable, owing to the uncertain and dynamic nature of when products are returned and when excess inventory is pulled from shelves

¹ <http://www.sdexec.com/article/12037309/statistics-reveal-8-to-9-percent-of-goods-purchased-at-stores-get-returned-and-25-to-30-percent-of-e-retail-orders-are-sent-back>

² <http://investors.liquidityservices.com/phoenix.zhtml?c=195189&p=irol-reportsannual>

and made available for resale. In turn, this results in a varying number of auctions being open on the platform at any point in time. The uncertainty in supply coupled with the uncertain valuation of potential buyers, who are downstream resellers with access to different resale channels, implies that liquidation platforms face a familiar operational challenge: how to tailor their design so as to profitably match supply with demand.

One relatively under-explored lever that an auction platform can employ to attain this match is its *listing policy* or, more specifically, the timing of auctions' closing dates. By aligning or, conversely, spreading out the closing dates of auctions, the platform can induce different levels of market thickness. In turn, the level of market thickness has first-order revenue implications for the platform as it determines participation in the marketplace. Specifically, buyers who face uncertainty with regard to actual supply levels on a given day may choose to participate only when they expect adequate availability, implying that even a *fixed* level of aggregate supply should be coordinated and allocated with this uncertainty in mind.

To this end, using a proprietary dataset collected from a leading online B2B platform, we investigate the role and efficacy of the platform's choice of listing policy in inducing different levels of market thickness and coordinating the behavior of market participants so as to influence market outcomes (i.e., the auctions' final prices). Notably, although, as we show, the listing policy alters neither the platform's underlying supply of arriving liquidation inventory nor its pool of potential bidders, its role in incentivizing bidder participation in auctions can be of first-order importance. Exploiting a natural experiment, we find that the platform's listing policy significantly impacts its revenues: implementing a listing policy to concentrate ("batch") auctions' ending times to certain days of the week increases sellers' revenues by 7.3%. Further evidence supports the hypothesis that the existence of market frictions, specifically *participation costs* associated with the bidders visiting the platform and determining their bidding strategies, drives this result.

Our finding underscores the economic significance of inducing the optimal thickness in the marketplace. In doing so, the platform faces the following trade-off. On the one hand, increasing supply availability on a given day (by having more auctions ending on that day) can profitably incentivize demand-side participation. On the other hand, having auctions end on the same day can induce them to cannibalize one another. Put differently, marketplace demand curves corresponding to a day are ultimately downward-sloping in the quantity supplied on that day.

Prescriptively, we study two relatively simple market design levers available to the platform to profitably calibrate its market thickness: the listing policy and (targeted) recommendations that aim to provide real-time information about the state of the platform to (a subset of) bidders. However, optimizing these levers involves complex demand-side behavior. For this purpose, we develop a structural model of bidders' behavior, i.e., whether and when they choose to participate

in the marketplace and in which auctions and how much they select to bid. These decisions are based on bidders' equilibrium beliefs about the supply-side availability of auctions on any given day and the demand-side competition from other bidders. The model also incorporates the dynamics of the platform's bidder pool and the heterogeneity of the bidders' participation costs, valuations, and potential demand. Out-of-sample revenue forecasts obtained from our estimated structural model align closely with the treatment effect as obtained from the natural experiment.

By simulating the model, we demonstrate how to design the platform's listing policy so as to induce the revenue-maximizing market thickness for any underlying level of expected supply. Our approach accounts for uncertainty in the actual *realization* of supply and bidders' equilibrium participation and bidding decisions. Further, in an effort to further reduce marketplace uncertainty and frictions, we also consider a recommendation system that notifies potential bidders of the supply conditions on the auction site. This is motivated by the fact that, although potential bidders may form accurate beliefs about the *expected* number of auctions on the platform on any given day, they typically do not know their *exact* number prior to visiting the platform itself. The recommendation system is built to inform a (randomly) sampled set of recipients about the realized supply on a day without them having to visit the platform, when the supply is higher than a given threshold. We find that appropriately designing such a recommendation system improves the platform's revenue by an additional 1.6% on the days when the platform sends out recommendations. This gain is achieved by reducing cannibalization during days with high supply and shifting the composition of participating bidders toward those with higher valuations.

1.1. Related Literature

Online marketplaces face a number of design challenges when seeking to match supply with demand so as to maximize revenue. A recent stream of papers explores how different aspects of marketplace design may be used to shape the incentives of market participants. For two-sided service platforms, [Cachon et al. \(2017\)](#) and [Kabra et al. \(2017\)](#) deliver novel pricing prescriptions based on how users respond to higher service levels. On the other hand, [Bimpikis et al. \(2017\)](#) consider pricing for spatially dispersed demand in a ride-sharing network. In an online auction setting, [Balseiro et al. \(2015\)](#) introduce the notion of a fluid mean-field equilibrium and illustrate its practical appeal in setting reserve prices. In addition, [Golrezaei et al. \(2017\)](#) show empirically that the platform's revenues increase by 3% when it boosts bids in a customized fashion based on bidders' past behaviors. Our work contributes to this literature by empirically demonstrating the impact of market thickness on bidders' participation and bidding decisions on a platform specializing in B2B

auctions.³ In addition, our focus is mainly on the use of non-price levers that affect the availability of supply-side inventory and, in turn, influence demand.⁴

Recent literature also connects participants' transaction costs and information frictions to outcomes in online markets. [Fradkin \(2017\)](#) studies the role of the search engine in reducing transaction costs and improving matches on Airbnb, while [Tadelis and Zettelmeyer \(2015\)](#) explore the disclosure of product information within online marketplaces. [Moon et al. \(2018\)](#) study the revenue and welfare implications of costs associated with customers monitoring a retailer's online channel for changes in the price and availability of inventory in which they are interested. [Horton \(2017\)](#) suggests that introducing a signaling feature that allows workers to indicate availability could increase surplus by as much as 6% in an online labor market. Closer in spirit to our research questions, [Cullen and Farronato \(2016\)](#) use data from TaskRabbit, a marketplace for domestic tasks, to empirically demonstrate that the growth of online peer-to-peer markets is largely affected by the thickness they induce. [Li and Netessine \(2018\)](#) find that higher market thickness actually leads to lower matching efficiency in an online peer-to-peer holiday rental platform.⁵ We contribute to this line of work by illustrating how design levers such as the platform's listing policy may lead to a sizable increase in the platform's revenues.

Relatedly, prior work explores both the effects of inventory availability on strategic demand and the related benefits of reducing buyers' uncertainty about availability. [Dana and Petruzzi \(2001\)](#), [Su and Zhang \(2009\)](#), and [Petruzzi et al. \(2009\)](#) consider settings where prospective buyers incur a search or opportunity cost when visiting physical stores with the intent to purchase a product, if it is available. The seller decides on inventory levels, potentially across stores, while customers form beliefs about the resulting availability. Moreover, [Su and Zhang \(2009\)](#), [Allon and Bassamboo \(2011\)](#), and [Alexandrov and Lariviere \(2012\)](#) examine how the seller may benefit from reducing consumers' uncertainty about its inventory availability. Empirically, [Gallino and Moreno \(2014\)](#) study the impact of sharing inventory information on consumer behavior through credible "buy online, pick up in store" offers, concluding that brick-and-mortar stores drew increased traffic by resolving availability risks. In contrast to these settings, the online platform that we focus on can neither set prices nor control the arrival of inventory to its marketplace. Nonetheless, we highlight that listing policies and state-contingent recommendations (communicated to a subset of bidders) can successfully complement the platform's efforts in boosting the bidders' participation rates.

³ Relatedly, [Mendelson and Tunca \(2007\)](#) and [Etzion and Pinker \(2008\)](#) explore strategic behavior in B2B spot markets while [Pinker et al. \(2010\)](#) focus on the revenue impact of the lot size in the context of a sequence of online auctions.

⁴ In recent work, [Arnosti et al. \(2014\)](#), [Halaburda et al. \(2017\)](#), [Kanoria and Saban \(2017\)](#), and [Akbarpour et al. \(2018\)](#) explore (non-price) interventions to improve efficiency in the context of matching platforms.

⁵ Relatedly, [Tunca \(2008\)](#) illustrates the potential for welfare losses in large markets.

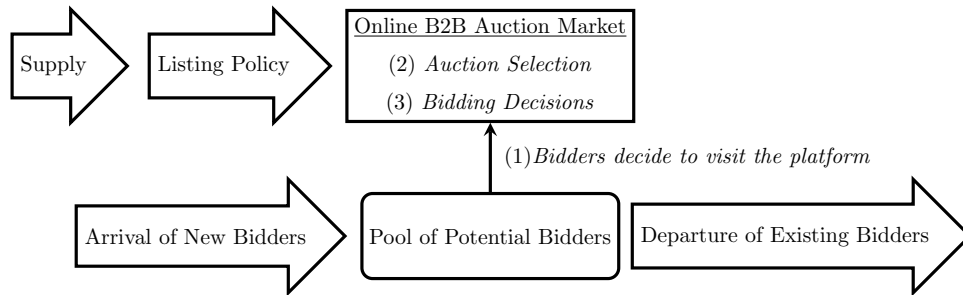
Finally, our paper falls within the growing body of literature employing structural estimation methods to study auction markets (e.g., [Jofre-Bonet and Pesendorfer \(2003\)](#), [Kim et al. \(2014\)](#), and [Backus and Lewis \(2016\)](#)) and to address questions of operational interest (e.g., [Olivares et al. \(2008\)](#), [Li et al. \(2014\)](#), and [Bray and Mendelson \(2015\)](#)).

2. Data and Background

Our dataset was obtained from a leading online platform managing private B2B auction markets for the liquidation inventory of more than thirty US big-box retailers, such as Costco, Walmart, Sears, and Home Depot (henceforth referred to as *sellers*). In contrast to traditional two-sided online markets, where a multitude of sellers enter and exit freely, each seller on our platform is associated with a private online auction market through which only its own merchandise (excess or returned inventory) can be sold. Under a long-term contract, the platform supports each seller’s online auction presence in exchange for a fixed commission taken out of the generated revenues. Sourced from customer returns, trade-ins, and unsold items, the supply of liquidation inventory reaching the platform can be highly stochastic and is beyond the platform’s control; rather, it is primarily driven by a seller’s reverse logistics. In other words, the supply of these B2B auction markets can be considered exogenous to the platform’s design.

Merchandise is sold through an ascending English auction with proxy-bidding (similar to the format used on eBay). Each auction lasts for 1 to 4 days and offers a bundle of similar products for sale in an all-or-nothing auction. While electronics account for most of the platform’s annual revenues, which are in the hundreds of millions of U.S. dollars, a broad range of product categories are auctioned on the site, including household appliances, furniture, and apparel. A typical auction lot contains a box of goods from the same product category and in roughly the same condition (e.g., unused, or used and in good condition). Bidders may access the auction’s *manifest*, which provides a brief description of the items included in the box. In addition, bidders are able to observe the current second-highest bid (the *standing bid*) and the time remaining in the auction. In contrast to the standing bid, the highest bid currently placed in an auction cannot be observed by bidders.

The demand side of the market consists of downstream resellers specializing in liquidation inventory (henceforth referred to as *bidders*). Reflecting bidders’ individual downstream resale channels, both bidders’ valuations and levels of demand (i.e., the number of auctions they intend to win) are substantially heterogeneous (see [Pilehvar et al. \(2016\)](#)). Bidders need to go through a separate registration process for each such private market in order to be able to view the market’s available inventory and submit bids. Finally, the pool of active bidders features the continuous arrival of new registrants and the continuous exit of some existing bidders, both of which are driven by the demand of downstream secondary markets worldwide. [Figure 1](#) summarizes the market dynamics and bidders’ major decisions within the platform.

Figure 1 Market Setting and Bidders' Decisions on a Given Day

2.1. Markets for iPhones

In what follows, we empirically examine whether the choice of listing policy has any impact on the platform's revenues. To this end, we use data on the iPhone auctions from the platform's two major mobile phone sellers. We restrict attention to iPhone auctions for two reasons. First, to identify the listing policy's revenue impact, we must carefully control for the characteristics of the particular products sold in each auction; iPhones are well-defined products with retail values that are derived in a straightforward way from their observable specifications, such as their model, carrier, condition, and time since the model's release date. This is contrary to other merchandise sold on the platform, such as furniture and household appliances, whose retail values depend on a large set of somewhat subjective features. Second, the revenue generated by iPhone sales alone accounts for approximately 73% of the revenue of the two major cell phone markets we study. We call these two major markets Market A and Market B.

Covering February 2013 to October 2015 (our *observation period*), our dataset tracks the entire bidding history (i.e., every bid's time of submission and dollar amount) in each of the 679 auctions in Market A and each of the 497 auctions in Market B. Moreover, we are able to track bidders' behavior across all auctions in the observation period.

Table 2 summarizes the auctions and bidding activity in the two markets. On average, 211 new bidders register per month in Market A, and 261 do so in Market B. Auction lot sizes vary substantially, both within markets (the corresponding coefficients of variation are 0.54 for Market A and 0.53 for Market B) and across markets (the average number of iPhones per auction is 150 in Market A versus 62 in Market B). Participation per auction, whether tallied in bidders or bids, tends to be higher in Market B's auctions. Nonetheless, Market A exhibits slightly higher average per-device revenues (\$116.10) than Market B (\$106.60).

2.2. Bidders

More than 2,200 bidders placed at least one bid in either Market A or Market B during the observation period. In line with prior work on B2B markets (Bapna et al. 2004, Pilehvar et al. 2016) the markets' bidder pools consist of experienced resellers – certified and registered in the market –

Table 2 Auctions in Markets A and B (means with standard errors in parentheses)

	Market A	Market B
Auction duration (in days)	2.60 (1.54)	2.90 (1.09)
Auction lot size (devices)	150.3 (81.0)	61.5 (32.7)
Avg. monthly registrations	211.0 (44.6)	260.7 (63.2)
Number of bidders per auction	4.96 (2.09)	8.20 (2.62)
Number of bids per auction	19.6 (12.1)	27.0 (15.9)
Avg. number of auctions per auction ending day	2.07	1.16
Total number of auctions	679	497
Avg. final per-device price (\$)	116.10 (49.97)	106.57 (51.47)

that are heterogeneous in both their demand profiles and valuations. Over 30% of bidders in both markets exhibit demand for multiple auctions. More specifically, these *multi-unit* (MU) bidders either submit winning bids for two or more concurrent auctions or submit a bid for a new auction shortly after winning an auction. The remaining bidders, whom we call *unit-demand* (UD) bidders, exhibit demand for winning only a single auction lot within our observation period. Though multi-unit bidders are fewer in number, their behavior affects the operations of the platform substantially, as they win over 80% of the auctions in both markets during our observation period.

In addition to their demand characteristics, bidders can also be classified by their registration status in the two markets. We refer to those who are registered in both markets as *cross-market* bidders, and those who are registered in only Market A or Market B during our study period as *Market A* or *Market B* bidders, respectively. Cross-market bidders can observe and participate in auctions from both markets. On the other hand, Market A and Market B bidders can observe auctions only from the single market in which they have registered.

In behavior similar to “sniping” in B2C auctions (Bajari and Hortacsu 2003), we find that bidders predominantly submit bids on an auction’s last day. Auctions typically close within the time window from 6pm to 8pm, and most bids are placed during that time. As we report in Table 3, the median last bid per bidder per auction comes quite late in both markets (after 99.0% of the auction’s total duration in Market A and 80.3% of the auction’s total duration in Market B). Moreover, the median winning bid in both markets arrives when 99.7% of the auction’s duration has elapsed (similar to 98.3% in Bajari and Hortacsu (2003)). Because an auction’s final price materializes toward its end, we focus on each auction’s ending day – as opposed, for example, to the entire time it is open – when studying the impact of the platform’s design on revenues.

2.3. Listing Policies

In both Markets A and B, as soon as the seller makes an inventory lot available to the platform (the inventory’s *arrival time*), a corresponding auction commences and is listed on the platform. When it lists the auction, the platform must decide and display the auction’s ending time. The policy governing how auctions’ ending times are determined is the platform’s *listing policy*.

Table 3 Bidder Types and Bidding Activity

	Market A	Market B
Percentage of MU bidders (%)	36.4	33.2
Percentage of auctions won by MU bidders (%)	91.9	73.8
Avg. number of auctions a MU bidder participates in	19.65	12.47
Avg. number of auctions won by a MU bidder ever	4.19	1.76
Avg. number of auctions a UD bidder participates in	2.45	1.69
Avg. number of auctions won by a UD bidder ever	0.29	0.12
Median of normalized time of first bid per bidder-auction (%)	97.2	60.0
Median of normalized time of last bid per bidder-auction (%)	99.0	80.3
Median of normalized time of winning bid per auction (%)	99.7	99.7

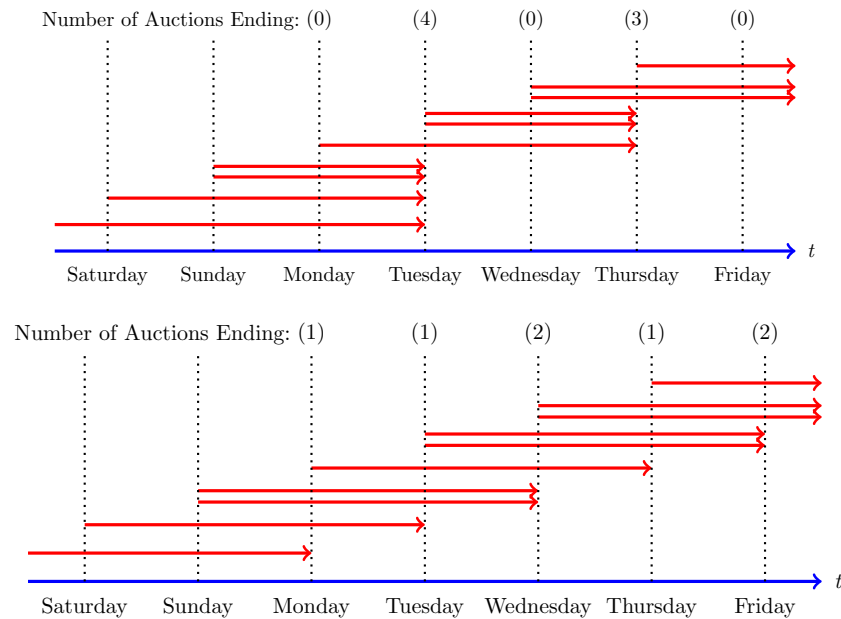
Table 4 Distribution of Inventory Arrivals (Auction Postings) by Day of Week

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Market A	16.7%	23.5%	20.4%	17.8%	16.3%	5.2%	0.0%
Market B	16.7%	21.4%	20.0%	17.3%	21.8%	3.0%	0.0%

Uniform listing. Market B’s auctions typically close 3 business days after the time they are listed on the platform (which is the time when the corresponding inventory is made available by the seller). Owing to the fact that the arrival times of inventory are approximately uniformly distributed across the weekdays in both markets (Table 4), the number of auctions expected to close is approximately constant across the weekdays. We refer to this listing policy, which sets auction ending times to enforce a fixed duration, as *uniform* listing. Market B uses uniform listing throughout our observation period, while Market A uses it from February 2013 to February 2014.

Batch listing. By contrast, in the second half of our observation period starting in November 2014 (see Section 3 for more details), Market A’s auctions closed only on Tuesdays and Thursdays. Consequently, the duration for which an auction remains open on the platform varies depending on the arrival time of its associated inventory lot. We use the term *batch* to refer to the listing policy that closes auctions only on Tuesdays and Thursdays.

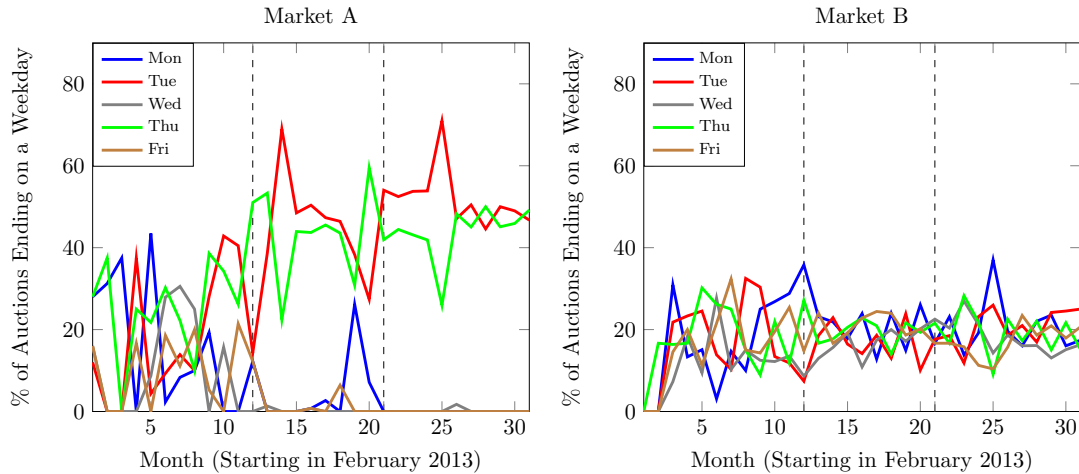
Effects of different listing policies. Even when receiving the same stream of inventory arrivals, different listing policies result in several salient divergences in outcomes. Figure 5 illustrates using the uniform and batch listing policies as examples. First, as shown in Figure 5, the number of days on which auctions close is different: two days of the week under batch listing, compared to five days under uniform listing. We call such days the platform’s “auction-clearing days.” Second, it naturally follows that the average number of auctions cleared on one of the platform’s auction-clearing days will depend on the platform’s choice of listing policy. Given that bidders predominantly bid in an auction on its last day (Table 3), different listing policies result in different levels of supply-side availability of auctions for the bidders (and they induce different levels of thickness in the market). As an example, Figure 5 illustrates that compared to uniform listing, batch listing maintains a higher level of supply-side availability on days when auctions are scheduled to close.

Figure 5 Batch (above) and Uniform (below) Listing Examples (each red arrow represents an auction)

3. Listing Policies, Market Thickness, and Platform’s Revenues

Prior work has studied both empirically and theoretically how to design an (online) marketplace so as to induce a better match between supply and demand; yet, the primary focus has mostly been on providing incentives for additional supply to join the marketplace (e.g., through surge pricing in ride-hailing platforms) or on smoothing out supply to match exogenous demand. By contrast, in our setting the aggregate levels of supply (incoming inventory) and potential demand (bidder pool) do not seem to be affected by the design decisions we consider. Instead, our analysis illustrates the potential merits of appropriately managing the *effective* availability of supply when the demand side endogenously determines *when* to actively participate in the marketplace.

By exploiting a natural experiment using Market B as a control, we find a substantial benefit attached to Market A’s change in listing policy. We estimate that implementing batch instead of uniform listing boosts Market A’s revenues by 7.3%, amounting to roughly \$3.8M annually. Our findings suggest that bidders face participation costs that cause them to strategically consider when to participate in the marketplace based on the supply-side availability they expect (i.e., number of active auctions). In other words, much like the platform offering service-level guarantees, assuring adequate market thickness attracts participation. More broadly, the fact that coordinating a fixed supply process can profitably incentivize marketplace participation, points to a novel operational trade-off. On the one hand, thickening the market increases demand-side participation, but, on the other hand, it induces supply-side cannibalization as substitutable auctions compete against one another. In the parlance of classic supply and demand theory, the downward-sloping demand curve

Figure 6 Percentage of Auctions Ending on Each Weekday Aggregated by Month in Market A (Left) and Market B (Right)

dictates that market-clearing unit prices ultimately fall as the quantity supplied by the market grows.

As with any field data, our natural experiment possesses considerable richness; that said, we would like to point out one potential limitation. While our data permits us to carefully account for individual auctions' characteristics and for depreciation in the auctioned phone models (the latter is shared in the form of trends by the two markets), our study is not a perfect natural experiment: a minority of bidders are cross-registrants who bid in both Markets A and B. However, we show that their presence *dilutes* the treatment effect of batch listing as compared to the case of completely independent markets; this is because these cross-registered bidders act as arbitrageurs who bring the treated and control prices closer together (Appendix A.3). Besides, we discuss how cross-registrants' participation is consistent with the hypothesis of participation costs in Section 3.4 and Appendix A.4.

3.1. Natural Experiment by Change in Listing Policy

Before February 2014, both markets enacted uniform listing and cleared an approximately equal number of auctions on each weekday (see Section 2.3). In February 2014, Market A had a change in personnel; management reports that this change was unrelated to the performance of the platform or its marketplaces. Nonetheless, the change triggered Market A to alter its listing policy twice. First, in February 2014, it constrained its auctions to close only on 3 days of the week (Mondays, Tuesdays, and Thursdays). Then, in November 2014, it adopted the batch policy described in Section 2.3 to close its auctions only on Tuesdays and Thursdays. Meanwhile, Market B's listing policy remained unaltered throughout the observation period.

Table 7 Supply and Demand Profiles across Markets within Pre- and Post-treatment Periods

	No. of weekly auctions		No. of weekly registrations	
	Market B	Market A	Market B	Market A
Pre-treatment Period	2.49	2.93	30	26
Post-treatment Period	5.11	6.34	58	47
Increase Rate (%)	105%	116%	93%	81%

Because we are interested in the effect of supply-side availability (and induced market thickness) on participation and revenues, we study a natural experiment: batching the listing policy serves as the relevant treatment of interest. We observe two clearly defined periods: the *pre-treatment* period (February 2013 to February 2014), during which both markets practiced uniform listing, and the *post-treatment* period (November 2014 to October 2015), during which Market A alone batch-listed its auctions. Figure 6 depicts the monthly percentages of auctions closing on each of the 5 weekdays, aggregated monthly for each of the Markets A and B: we observe the changes over the pre-treatment and post-treatment periods, as well as the intervening time period.

In addition to the market’s listing policy, an auction’s per-device revenue is affected by the attributes of the iPhones in the auction, market-level characteristics such as the supply process and the registration rate of new bidders, temporal effects that encompass new product releases, depreciation, and price fluctuations in the overall iPhone market. Observing both markets’ pre-treatment period characteristics allows us to account for unobserved differences between the two markets. In the data, we note that the differences in these market-level features between the two markets remain roughly constant over time.⁶ In Table 7, we summarize the numbers of auctions and the newly registered bidders per week for both markets across the pre- and post-treatment periods. Market A’s aggregate levels of supply grew at a slightly faster rate than that of Market B while its pool of potential bidders (demand) grew at a somewhat slower rate than that of Market B. According to classic supply and demand theory, we would expect an increase in Market B’s revenues relative to those of Market A. However, we observe the opposite. Using the average treatment effect (ATE) methodology reviewed in Imbens (2004), we estimate the effect of the platform’s listing policy on auctions’ final prices while controlling for product attributes, temporal effects, and static market-level differences. Within the observation window, we do not observe changes in any of the market design levers except for the listing policy.

3.2. Descriptive Comparison and Difference-in-Differences Analysis

As a primer to the analysis, Figure 8 compares the two markets’ average, per-device revenues in the pre-treatment and post-treatment periods, respectively. While iPhone prices fall over time in

⁶ Our average treatment effect analysis explicitly controls for auction-level characteristics such as auction size and product types, as detailed in Section 3.3.

Figure 8 Averages and Trends of Per-device Prices (Avg. selling prices in the pre- and post-treatment periods (left) and trends of final price per unit in log scale (right))

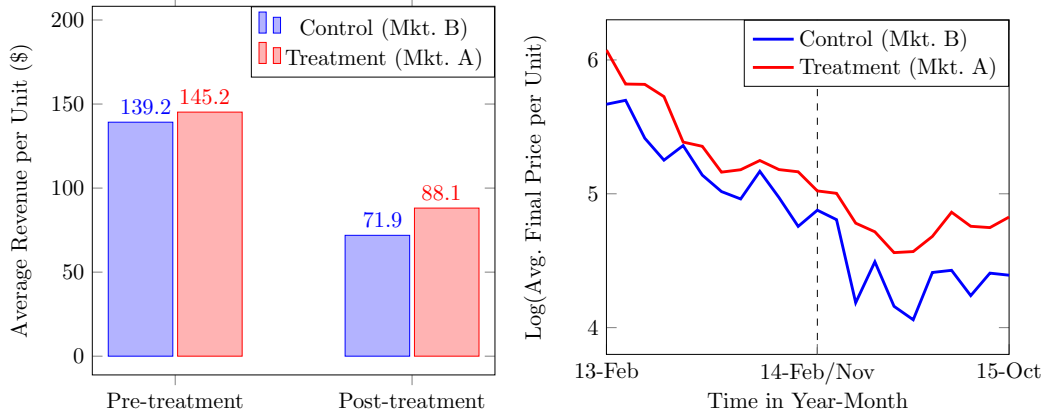


Table 9 Difference-in-Differences Estimate of Policy Switch on Final Price per Unit

<i>Dependent Variable</i>	
Log (Final Price per Unit)	
Treatment: Batch Policy	10.1%*** (0.027)

*p<0.1; **p<0.05; ***p<0.01
Confidence intervals constructed via bootstrap.

both markets (as the phone models depreciate), the left plot of Figure 8 highlights an emerging gap in the markets’ per-device revenues during the post-treatment period relative to the pre-treatment period, consistent with the batch listing policy affording Market A a post-treatment boost in per-device revenue. The right plot of Figure 8 shows that the pre-trends in both markets are relatively parallel. Notably, the gap between the price trends increases during the post-treatment period.

A difference-in-differences estimator transparently compares the post-treatment revenue difference between the two markets against any that existed in the pre-treatment period. As shown in Table 9, the markets’ revenue difference increases by 10.1% post-treatment, which we might attribute to the effect of Market A’s new listing policy. Because it remains conceivable that the markets’ inventory compositions could have changed across the periods, we now carry out an average treatment effect analysis that controls for granular differences in the attributes of individual auction lots and for weekly time trends affecting the phones’ (hence lots’) valuations.

3.3. Average Treatment Effect

First, our analysis controls explicitly for observables, such as phone model, carrier, and auction lot size. Additionally, we include a weekly time fixed effect to account for temporal price fluctuations in the broader iPhone market, leveraging that these effects are simultaneously present in both markets. Finally, we exploit the dataset’s pre-treatment period, during which both markets used

the uniform listing policy, to measure and account for the effect of unobservable differences between the two markets, including the information they provide about the quality of their products.⁷

Specifically for this analysis, we assign an auction lot j at time t to the treatment group if it is posted to Market A (indicator variable A_{jt}) and to the control group if it is instead posted to Market B. Based on its observed attributes, each auction lot exhibits a propensity for being assigned to the treatment group, known as its propensity score. As developed in the related literature (e.g., Rosenbaum and Rubin (1983), Hirano and Imbens (2001), Hirano et al. (2003)), an efficient estimate of the average treatment effect (ATE), τ , that accounts for such attributes can be obtained as the difference between post-treatment and pre-treatment outcomes appropriately weighted using their associated propensity scores. We estimate separate ATEs for each of the pre-treatment and post-treatment periods (τ_{Pre} and τ_{Post} , respectively) and are ultimately interested in the difference between the two. While both τ_{Pre} and τ_{Post} include the net revenue effects of the markets' unobserved differences (e.g., product quality categories, reputations, and bidder compositions), only τ_{Post} captures the additional revenue effect of batch listings under Market A's revised listing policy. Accordingly, the corresponding estimators, $\hat{\tau}_{\text{Pre}}$ and $\hat{\tau}_{\text{Post}}$, are each obtained by Expression (8), which can be found in Appendix A.1. As a (technical) aside, note that relative to employing propensity score methods alone, the mixed-methods approach, which combines propensity score weighting and a regression model for the outcome, may improve the precision of the estimation by directly addressing the correlation between the covariates and the outcome (Imbens 2004). One explicit benefit is "double robustness" (Robins et al. 1994, Wooldridge 2007, Graham et al. 2012), whereby the estimator is consistent if either the parametric propensity score model or the outcome regression model is correctly specified. In this study, we adopt an estimator from a class of augmented inverse probability weighting (AIPW) estimators introduced in Robins et al. (1994) (see Appendix A.1).

3.4. Revenue and Participation Effects of Batch Listing

As reported in Table 10, we find the per-device revenue effect of the batch listing policy, $\hat{\tau}_{\text{Post}} - \hat{\tau}_{\text{Pre}}$, to be positive and statistically significant. We find that the batch listing policy, in comparison with the uniform listing policy, yields a 7.3% average increase in Market A's per-device revenue (Table 10).⁸ This estimate translates to more than \$3.8M in additional revenues annually for Market A.

⁷ More than 90% of Market A iPhone auctions are classified into quality "A/B," and more than 99% of Market B iPhone auctions are classified into quality "Used/Functional." In our sample, we focus exclusively on auctions with these two quality types.

⁸ We conduct robustness checks to support the validity of our estimates (see Appendix A.2), including examining the distributional overlap of covariates for the treatment and control groups. In Appendix A.2, we carry out robustness checks on the only covariate that exhibits possibly insufficient overlap (i.e., auction lot size).

Table 10 Estimated ATE of Policy Switch on Final Price per Unit and No. of Bidders per Auction

	<i>Dependent Variables</i>	
	Log (Final Price per Unit)	No. of Observed Bidders per Auction
Treatment: Batch Policy	7.3%*** (0.009)	0.49** (0.22)

*p<0.1; **p<0.05; ***p<0.01

Confidence intervals constructed via bootstrap.

Setting a market’s listing policy may appear to be an innocuous choice born of convenience, chance, or custom: a priori it is not clear why the platform would expect anything beyond a marginal impact on its revenues, given that its listing policy does not have an impact on either its exogenous supply of auctions or its demand pool of certified bidders. We examine whether higher supply-side availability of inventory (auctions) on the platform actually results in more bidders participating per auction. To this end, we apply a similar average treatment effect analysis to the one in the previous section to estimate the effect of the policy switch on the number of observed bidders per auction. All explanatory variables remain the same as those in the ATE analysis of the final price. As shown in the last column of Table 10, under the batch policy, Market A attracts 0.49 more bidders per auction, on average. Besides, we find that the cross-market bidders’ participation rate in Market B on Tuesdays and Thursdays is substantially higher within the post-treatment period (i.e., 84%) than that within the pre-treatment period (i.e., 63%). The findings suggest that it is costly for bidders to participate in the auctions; thus, they strategically decide when to visit the platform (see Appendix A.4). Lastly, an additional difference-in-differences analysis confirms that the switch of the listing policy has little impact on the bidders’ bidding amounts (see Appendix D.3), which implies that the listing policy’s revenue gain is not driven by bidders’ bidding behavior.

In summary, the above discussion suggests the following mechanism to explain the revenue increase: an increase in market thickness boosts bidder participation both on aggregate and per auction. In turn, this additional traffic results in higher revenues for the platform. Underlying this mechanism is the fact that visiting the platform on a given day, monitoring the set of available auctions, and choosing whether and how much to bid entails a significant cost for bidders.

Despite of the positive revenue effect of increasing market thickness, there exists potential pitfalls to *thicken* the market. First, batching too many auctions together may cannibalize and decrease per-device revenues by increasing the number of immediately available substitutes (see Appendix D.2 for empirical evidence of the auction cannibalization). Second, some bidders may be disincen-tivized from participating on the platform due to the increased participation of competing bidders. In what follows, we develop a structural model to characterize the optimal market thickness.

4. Structural Model

We present a dynamic, structural model endogenizing the bidders' decisions on whether and when to visit the platform, which auction(s) to participate in, and how much to bid. Our discrete-time model captures the behavior of a dynamic pool of potential bidders who are heterogeneous in their valuations and demand profiles while facing an exogenously stochastic supply of liquidation inventory lots arriving to the platform to be auctioned. We assume that there are up to S_t auctions on Day t each of which is for a single unit whose valuation is appropriately normalized. Given that auction lots differ from one another, we assume that a bidder's valuation for a given auction has an idiosyncratic term. Motivated by our assertion that cross-market bidders only dampen the price impact resulting from batching auctions' ending times (Appendix A.3), we consider only a single market. To capture heterogeneity in the bidders' demand, we define the following two bidder types:

DEFINITION 1. (*Bidder Types*) Bidders are risk-neutral and have private values. Each bidder belongs to one of the following two types:

(i) Unit demand (UD): A UD bidder is interested in winning only one auction throughout her *lifetime* on the platform. Each UD bidder is permanently endowed with a private valuation for winning an auction lot that is drawn independently from distribution F^{UD} .

(ii) Multi-unit (MU): A MU bidder is interested in winning multiple auctions on the platform. In particular, we assume that a MU bidder is interested in winning up to K auctions in a single *day*, where $K > 1$ is exogenously specified, regardless of her prior win history. Each MU bidder's endowed private valuation for an auction lot is drawn independently from distribution F^{MU} .

In contrast to UD bidders, who operate on relatively low volumes, MU bidders represent repeat buyers that interact regularly with the platform. To ensure that the model is tractable, we assume that the bidders' private valuations for each type follow the Weibull distribution, which fits the data reasonably well. As mentioned in Section 2.2, most of the bidding activity takes place toward the ending time of an auction. In such a case, we assume that bidders on any given day submit bids only to auctions that are closing on that day. Also, more than 90% of the auctions end within the narrow time window between 6 and 8 PM. Hence, we find it reasonable to assume that if a bidder decides to visit the platform on a given day, the timing of her visit (and subsequent bidding) within the day is exogenous; i.e., it is not part of her decision-making process.⁹ On each day t , the following sequence of events transpires:

(1) *Supply*. First, S_t new auctions are listed on the platform, with S_t following the count distribution P_{Supply} . Each posted auction possesses characteristics affecting its idiosyncratic fit with a bidder's resale channels. Therefore, Bidder ℓ 's valuation for an auction lot j is the sum of Bidder

⁹ The exogenous timing assumption is essential for tractability purposes as well.

ℓ 's endowed product valuation, x_ℓ , and an idiosyncratic term, $\zeta_{\ell j}$. Each $\zeta_{\ell j}$ is independently drawn from mean-zero normal distributions F_ζ^{MU} and F_ζ^{UD} , with standard deviations ν^{MU} for MU bidders and ν^{UD} for UD bidders, respectively.

(2) *Platform Participation.* New bidder arrivals join the existing pool of potential bidders, under an exogenous arrival process. If t is an auction-ending day as defined in Sections 2.3 and 2.2, all potential bidders in the pool must decide simultaneously whether to visit the platform to bid in ending auctions. Notably, while each bidder ℓ knows her own valuation x_ℓ for a standard unit throughout, a bidder does not know either the current state of the platform (i.e., how many auctions are ending that day and/or the current standing bids in those auctions) or the idiosyncratic component of her lot valuation $\zeta_{\ell j}$ for any listed auction j . Instead, her decision on whether to visit the platform is based on her expected payoff from visiting, which in turn depends on her ex ante beliefs about the likely state of the platform. Formally, upon visiting the platform, the observable state of the market on Day t is $\omega_{\ell t} \triangleq (n_t, \mathbf{s}_t)$, where n_t is the market thickness (number of auctions ending) on Day t , with support $\mathcal{N} = \{0, 1, \dots, \bar{N}\}$, and \mathbf{s}_t denotes the vector of standing bids of the auctions that are ending on Day t . As discussed in Section 4.1, the bidder's ex ante beliefs about $\omega_{\ell t}$ anticipate the platform's equilibrium steady-state distribution. Against her expected payoff, bidder ℓ weighs her daily participation cost, $c_{\ell t}$, on Day t , which captures the potential cost (in time and effort) associated with participating on the platform. Conditional on a bidder's type, her participation cost is drawn independently each day from an exponential distribution with rate μ^{MU} for MU and μ^{UD} for UD bidders. Thus, across bidders and days, whether or not a bidder visits the platform depends both on her endowed valuation x_ℓ and on her realized participation cost $c_{\ell t}$.

(3) *Auction Selection and Bidding.* Upon visiting the platform, a bidder observes the realized state $\omega_{\ell t}$ and the idiosyncratic term of her valuation $\zeta_{\ell j}$ for each available auction j . First, she decides which auction(s) to participate in based on her private valuations and her beliefs about the currently highest rival bids given $\omega_{\ell t}$, and she then determines how much to bid in each auction so as to maximize her expected payoff.

(4) *Departure.* Given the nature of their demand, UD bidders depart with certainty upon winning an auction. In addition, we let α^{MU} and α^{UD} denote the daily retention probabilities of those MU and UD bidders who do not win an auction, respectively. In particular, at the end of Day t , each bidder departs the bidder pool with probability $(1 - \alpha^{TY})$, where $TY \in \{MU, UD\}$.

4.1. Bidder Strategies and Equilibrium

The size of the market (i.e., Market A averages about 160 active bidders daily) makes it impractical for an individual bidder to fully track her competitors and the history of their actions. Instead, we assume that bidders respond to their steady-state beliefs about their rivals, which are not meaningfully affected by their own actions (Backus and Lewis 2016). This assumption approximates well

a setting that involves a large group of anonymous bidders, with similarities to other assumptions that have been employed in related settings, such as the notions of oblivious equilibrium (e.g., Weintraub et al. (2008)), stationary competitive equilibrium (e.g., Bodoh-Creed et al. (2016)), and mean field equilibrium (e.g., Iyer et al. (2014), Yang et al. (2016)). Prior to defining the steady-state equilibrium, we first introduce a bidder's optimal actions on a given day, given her beliefs about the underlying market state Ψ and the highest rival bid G .

Platform Participation. A bidder determines whether or not to participate in the platform on a given day so as to maximize her expected payoff. In particular, for MU bidders, the corresponding maximization problem can be reduced to that of a single period, since their payoffs on any given day are independent of their actions on other days. She receives a payoff of zero if she does not visit. Specifically, a MU bidder's platform participation decision is

$$\sigma_{VST}^{MU}(x_\ell, c_{\ell t}) = \begin{cases} 1, & \text{i.e., visit, if } r^{MU}(x_\ell; G, \Psi) - c_{\ell t} \geq 0 \\ 0, & \text{i.e., wait, otherwise,} \end{cases} \quad (1)$$

where $r^{MU}(x_\ell; G, \Psi)$ denotes the expected payoff for a MU bidder who visits the platform on Day t . For UD bidders, since they have demand for only one auction throughout their lifetime, their actions on a given day affect their future payoffs. Thus, a UD bidder's platform visit decision is

$$\sigma_{VST}^{UD}(x_\ell, c_{\ell t}) = \begin{cases} 1, & \text{i.e., visit, if } v^{UD}(x_\ell; G, \Psi) - c_{\ell t} \geq \alpha^{UD} v_f(x_\ell; G, \Psi) \\ 0, & \text{i.e., wait, otherwise,} \end{cases} \quad (2)$$

where $v^{UD}(x_\ell; G, \Psi)$ denotes the aggregated payoff for a UD bidder who visits the platform on Day t , and $v_f(x_\ell; G, \Psi)$ denotes her continuation value if she does not exit the platform after Day t .¹⁰

Auction Selection and Bidding. After a bidder (of type $TY \in \{MU, UD\}$) participates in the platform, she selects which auctions to bid in, denoted by σ_{SLT}^{TY} , and how much to bid in the selected auctions, denoted by σ_{BID}^{TY} . Naturally, an auction's standing and highest rival bids are not independent. Thus, the bidder updates her beliefs over the highest rival bids after observing the corresponding standing bids and the realized market thickness. Under the updated beliefs, σ_{SLT}^{TY} and σ_{BID}^{TY} jointly maximize her expected payoff on the platform on a given day. For MU bidders, the corresponding maximization problem is

$$(\sigma_{SLT}^{MU}, \sigma_{BID}^{MU}) = \arg \max_{\sigma_{SLT}, \sigma_{BID}} \sum_{j \in \sigma_{SLT}} \int_{s_j}^{\sigma_{BID,j}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j, \quad (3)$$

where p_j is the highest rival bid in Auction j , and $g_j(p_j | \omega_{\ell t})$ denotes its conditional PDF. The term $\int_{s_j}^{\sigma_{BID,j}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j$ stands for the expected payoff from bidding in Auction j . If the bidder chooses Auction j , her optimal bidding strategy is to bid up to her valuation (i.e.,

¹⁰ The characterizations of $r^{MU}(x_\ell; G, \psi)$, $r^{UD}(x_\ell; G, \Psi)$, $v^{UD}(x_\ell; G, \Psi)$, and $v_f(x_\ell; G, \Psi)$ are detailed in Appendix B.1.

$\sigma_{BID,j}^{MU}(x_\ell; \zeta_{\ell j}, \omega_{\ell t}) = x_\ell + \zeta_{\ell j}$. Then, her optimal auction selection decision $\sigma_{SLT}(x_\ell; \zeta_{\ell 1}, \dots, \zeta_{\ell \bar{N}}, \omega_{\ell t})$ is to choose up to K auctions with the highest expected payoffs. For UD bidders, the corresponding maximization problem is

$$(\sigma_{SLT}^{UD}, \sigma_{BID}^{UD}) = \arg \max_{\sigma_{SLT}, \sigma_{BID}} \sum_{j \in \sigma_{SLT}} \int_{s_j}^{\sigma_{BID,j}^{UD}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j + \alpha^{UD} (1 - G_j(\sigma_{BID,j}^{UD} | \omega_{\ell t})) v_f(x_\ell; G, \Psi), \quad (4)$$

where $G_j(p_j | \omega_{\ell t})$ denotes the conditional CDF of the highest rival bid in Auction j . The first term in Expression (4) is equal to the instantaneous payoff if the bidder wins Auction j , where we establish that $g_j(p_j | \omega_{\ell t})$ is independent of the vector of optimal bids (Appendix B.2). The second term is equal to the bidder's (expected) payoff if she does not win the current auction (which occurs with probability $1 - G_j(b_{\ell j} | \omega_{\ell t})$). Then, we can show that her optimal bid in Auction j is

$$\sigma_{BID,j}^{UD}(x_\ell; \zeta_{\ell j}, \omega_{\ell t}) = x_\ell + \zeta_{\ell j} - \alpha^{UD} v_f(x_\ell; G, \Psi), \quad (5)$$

and her optimal auction selection $\sigma_{SLT}^{UD}(x_\ell; \zeta_{\ell \cdot}, \omega_{\ell t})$ is to choose the auction with the highest payoff.¹¹

Note that the bidders' decision whether to visit the platform on any given day largely depends on their beliefs about the steady state of the market (i.e., Ψ) on that day. On the other hand, once a bidder is already on the platform and has observed the number of auctions closing on that day, her decisions regarding which auction(s) to participate in and how much to bid are mainly driven by her beliefs about the highest rival bids (i.e., G). Thus, anticipated market thickness has a more direct impact on a bidder's decision whether to visit the platform compared to which bidding strategy to use. Formally, the notion of steady-state equilibrium we employ is defined as follows.

DEFINITION 2 (EQUILIBRIUM). A steady-state equilibrium is a tuple $(\{\sigma^{MU}, \sigma^{UD}\}, \{G, \Psi\})$ such that:

- (*Optimality*) For bidder ℓ with type $TY \in \{MU, UD\}$, her best response comprises three decisions on day t ; that is, it takes the form $\sigma^{TY}(x_\ell; c_{\ell t}, \zeta_{\ell \cdot}, \omega_{\ell t}) = [\sigma_{VST}^{TY}, \sigma_{SLT}^{TY}, \sigma_{BID}^{TY}]$, where σ_{VST}^{TY} , σ_{SLT}^{TY} , and σ_{BID}^{TY} are defined by Expressions (1), (2), (3), and (4), respectively, given the steady-state distributions for the market state Ψ and the highest rival bids G .

- (*Consistency*) Steady-state distributions $\{G, \Psi\}$ are induced by bidders following strategy $\sigma^{TY}(x_\ell; c_{\ell t}, \zeta_{\ell \cdot}, \omega_{\ell t})$, $TY \in \{MU, UD\}$.

In equilibrium, both Ψ and G are functions of the market design (e.g., listing policy). A market design lever impacts the platform's revenues by influencing bidders' beliefs about the market state and the competition level. With the notion of how bidders behave in equilibrium, we then establish that an equilibrium exists under a (mild) technical assumption.

¹¹ As expected, UD bidders shade their bids by their continuation values. As they are forward-looking and realize that their purchases today come at the expense of winning an auction in the future, they implicitly discount their willingness to pay for a present auction.

PROPOSITION 1. *Assume that $\nu^{MU} = \nu^{UD} = 0$ and $\bar{N} \leq K$. Then, an equilibrium exists.*

The proof of Proposition 1 is in Appendix B.3. As stated in the proposition, equilibrium existence is shown under the assumption that bidders have the same valuation for all auctions and the daily demand of MU bidders is sufficiently large. For the general case where the assumptions of the proposition are relaxed (i.e., $\nu^{MU} \geq 0$, $\nu^{UD} \geq 0$, and $\bar{N} > K$), we provide an algorithm that efficiently converges to the equilibrium (see Appendix D.4), which, in turn, generates a broad set of counterfactuals (see Section 6). It is also worthwhile to highlight that given Ψ and G , all of the bidders' decisions can be expressed analytically. This is crucial in enabling us to structurally estimate the model in the following section.

5. Structural Estimation

In this section, we outline our structural estimation approach and present our estimates. More specifically, we estimate our structural model on Market A's data exclusively from February 2013 to February 2014 (i.e., the pre-treatment period), throughout which the uniform listing policy was used. By doing so, we are able to: (i) validate our structural model by deriving out-of-sample projections for Market A's pre-treatment-to-post-treatment revenue improvement for comparison with the estimate we obtained from our ATE analysis in Section 3.3; and (ii) reserve the post-treatment period's data for use in our counterfactual market design analysis.

As is common in the auction literature (Backus and Lewis 2016), we first normalize the observed bids to adjust for heterogeneity in the products' features and condition for time fixed effects. As we argue in detail in Appendix C.1, among the iPhone models available on the platform (i.e., iPhone 4, iPhone 4s, and iPhone 5), bidders do not strongly favor one product over the others (after appropriately adjusting for differences in their valuations). Therefore, we can reasonably treat the iPhone models as substitutes (after appropriately normalizing the corresponding bids). Post-normalization, bidders' valuations are treated as drawn from a common distribution, which is assumed to be the Weibull distribution, throughout the observation period.¹²

The structural estimation follows two steps: we conduct (1) a nonparametric estimation of the platform's steady-state distribution of auctions and bids, followed by (2) a maximum simulated likelihood (MSL) estimation of our modeling primitives for bidders' valuation distributions, participation costs, and retention rates. The pre-treatment period lasts for a year, so it is a fair assumption that Market A's pre-treatment period data are generated from a steady-state equilibrium. The corresponding equilibrium distributions Ψ and G are directly estimated in step (1). Notably, Ψ and G are equilibrium outcomes rather than exogenously specified model primitives. Using the iterative algorithm detailed in Appendix D.4, they are re-computed numerically for each counterfactual scenario corresponding to a given market design (e.g., listing policy) in Section 6.

¹² We shift the normalized bids to ensure that they fall within the support of the Weibull distribution.

5.1. Estimating the Platform’s Steady State

The steady-state belief about the market consists of the distribution of the market thickness as well as the distribution of the corresponding standing bids. In particular, it takes the form of $\Psi(\omega_{\ell t}) = P_{MKT}(n_t)\psi(\mathbf{s}_{\ell t}|n)$, where P_{MKT} is estimated by the empirical distribution of auctions ending on a given day. On the other hand, the PDF of standing bids $\psi(\mathbf{s}_{\ell t}|n)$ conditional on n is estimated using kernel density estimator to minimize misspecification bias.

Furthermore, a bidder forms beliefs about the highest rival bids given a market state $\omega_{\ell t}$. Recall that $g_j(p_j|\omega_{\ell t})$ denotes the PDF of the highest rival bid in the j^{th} auction given the state $\omega_{\ell t}$. Again, we use a kernel density estimator for $g_j(y|\omega_{\ell t})$ to mitigate misspecification bias. The kernel density estimators for $\Psi(\omega_{\ell t})$ and $g_j(p_j|\omega_{\ell t})$ are detailed in Appendix C.2. Lastly, we let the upper bound K on the MU bidders’ demand be equal to 14 auctions (i.e., the maximum number of auctions in which a MU bidder submitted a bid on a single day in our observation period).¹³

5.2. Estimating the Bidders’ Primitives

Before describing in detail how we estimate them, we reiterate that the model primitives for a bidder of type $TY \in \{MU, UD\}$ include: (i) her endowed valuation of the product F^{TY} that is drawn from the Weibull distribution with scale parameter λ^{TY} and shape parameter γ^{TY} ; (ii) her daily participation cost that is exponentially distributed with rate μ^{TY} ; and (iii) her daily retention rate α^{TY} that represents her probability of remaining in the bidder pool for another period. As we argue in Appendix C.4, our dataset exhibits sufficient variation to identify the above primitives.

Given a bidder’s endowed valuation and her beliefs over the market state and the highest rival bids (which are estimated in Section 5.1), her decisions, including whether and when to visit the platform, which auction(s) to participate in, and how much to bid, can be analytically obtained as we described in Section 4. This, in principle, allows us to recover all model primitives in an efficient manner by iteratively obtaining the likelihood of a given equilibrium outcome as a function of the set of model primitives using simulation and updating the primitives accordingly.

Bidder ℓ ’s observed bidding history \mathbf{X}_ℓ consists of two parts. The first part is her participation sequence, $\mathbf{B}_\ell = [B_\ell^{t_\ell}, \dots, B_\ell^{t_\ell+l_\ell}]$, where $B_\ell^{t_\ell} \in \{0, 1\}$ indicates whether she placed a bid on Day t (here, t_ℓ and $t_\ell + l_\ell$ denote the first and last observed bidding days for the bidder in our sample). The second part comprises her bids on Day t , $\mathbf{b}_{\ell t}$, in the auctions she entered and the standing bids, $\mathbf{S}_{\ell t}$, in the auctions available but not entered within the period from Day t_ℓ to Day $t_\ell + l_\ell$.

Given her endowed valuation x_ℓ and her presence in the bidder pool, we specify the likelihood of each of her observed behaviors. The following formulas apply to both types of bidders;

¹³ During our observation period, no bidder submits bids to more than 14 auctions on any single day.

thus we omit the superscripts specifying the bidder type. On Day t , observing a bid by Bidder ℓ implies that she (i) visits the market (which occurs with probability $P^V(X_\ell|\theta)$) and (ii) places bid $\mathbf{b}_{\ell t}$ (which has likelihood $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \theta)$). Thus, the likelihood of her placing bid $\mathbf{b}_{\ell t}$ is $P^V(X_\ell|\theta)L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \theta)$. On the other hand, if no bid is observed on that day, there are two possibilities: (i) she chooses not to visit the platform (which occurs with probability $1 - P^V(x_\ell|\theta)$), or (ii) she visits the platform but finds it optimal to not place a bid (which occurs with probability $P^V(x_\ell, \theta)L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta)$). Thus, the corresponding likelihood of not observing a bid is $1 - P^V(x_\ell|\theta) + P^V(x_\ell|\theta)L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta)$.¹⁴ In addition, we specify the likelihood associated with her exit from the bidder pool. Given that the day when she exits from the bidder pool cannot be observed, we assume that she leaves the bidder pool on any day within E days¹⁵ of her last bid (we then take the expectation over the likelihoods of all possible exit days).

In summary, the overall likelihood of a bidder's entire bidding history conditional on her endowed valuation x_ℓ is given by $\mathcal{L}_\ell(\mathbf{X}_\ell|x_\ell, \theta) =$

$$\begin{aligned} & \left(\alpha^{l_\ell} \prod_{t=t_\ell}^{t_\ell+l_\ell} \underbrace{\left(P^V(x_\ell|\theta)L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \theta) \right)^{B_\ell^{t_\ell}}}_{\text{Placing the bid(s) } \mathbf{b}_{\ell t} \text{ on day } t} \underbrace{\left(1 - P^V(x_\ell|\theta) + P^V(x_\ell|\theta)L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta) \right)^{1-B_\ell^{t_\ell}}}_{\text{Not placing a bid on day } t} \right) \\ & \cdot \underbrace{\left(1 + \sum_{t'=1}^{E-1} \alpha^{t'} \prod_{t=t_\ell+l_\ell+1}^{t_\ell+l_\ell+t'} \left(1 - P^V(x_\ell|\theta) + P^V(x_\ell|\theta)L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta) \right) \right)}_{\text{Exiting bidder pool within } E = 14 \text{ days since the last bidding day}} (1 - \alpha). \end{aligned} \quad (6)$$

As the bidder's valuation x_ℓ is not observed, we need to consider the unconditional likelihood function, $\mathcal{L}_\ell(\mathbf{X}_\ell|\theta) = \int_{x_\ell} \mathcal{L}_\ell(\mathbf{X}_\ell|x_\ell, \theta) f(x_\ell|\lambda, \gamma) dx_\ell$, where $f(x_\ell|\lambda, \gamma)$ denotes the PDF of the Weibull valuation distribution. As $\mathcal{L}_\ell(\mathbf{X}_\ell|\theta)$ has no closed-form expression, the MLE approach is computationally intractable. To overcome this issue, we construct the simulated likelihood function $\hat{\mathcal{L}}_\ell(\mathbf{X}_\ell|\theta)$ by employing Monte Carlo integration.¹⁶ The MSL estimate $\hat{\theta}_{MSL}$ is obtained by maximizing the log of the simulated likelihood of all bidders, namely $\hat{\theta}_{MSL} = \arg \max_\theta \sum_\ell \log(\hat{\mathcal{L}}_\ell(\mathbf{X}_\ell|\theta))$.

5.3. Estimation Results

In Table 11, we report our estimates for MU and UD bidders, including their valuation distribution, average daily participation cost, and daily retention probability.

These estimates highlight several differences between MU and UD bidders. In particular, MU bidders possess on average higher endowed valuations than UD bidders, but exhibit substantially lower variability in idiosyncratic valuations within-type across individual bidders. On average, the

¹⁴ Expressions for $P^V(X_\ell|\theta)$, $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \theta)$, and $L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta)$ are detailed in Appendix C.4.

¹⁵ Our estimates are obtained under $E = 14$ days. We have tested other values (e.g., $E = 21$ days and 28 days) and found that the resulting estimates do not differ significantly.

¹⁶ For additional discussion on $\hat{\mathcal{L}}_\ell(\mathbf{X}_\ell|\theta)$, refer to Appendix C.4.

Table 11 Maximum Simulated Likelihood Estimates for the Structural Model

	<u>MU Bidders</u>	<u>UD Bidders</u>
<u>Valuation per auction</u> (x)		
Mean	\$18,594 (\$138)	\$15,459 (\$203)
Standard Deviation	\$876 (\$127)	\$897 (\$164)
Idiosyncratic Error (SD) (ν)	\$2,978 (\$72)	\$4,313 (\$244)
<u>Avg. Daily Participation Cost</u> (μ)	\$107.52 (\$15.38)	\$70.64 (\$10.06)
<u>Retention Rate</u> (α)	0.968 (0.001)	0.843 (0.014)
Number of Bidders	74	113

Note: Standard deviations of the estimates are in parentheses. Valuations are at normalized scale.

MU type's valuation is higher by \$20.86 per unit (the average auction lot size in Market A is 150.3 units), which is both statistically and economically significant. By comparison, UD bidders exhibit more variability in their idiosyncratic terms. Despite the non-dominance of valuation components across MU and UD bidders, UD bidders shade their bids (Expression (5)), which contributes to MU bidders winning the majority of auctions.

Per daily platform visit, MU bidders incur a substantially higher average participation cost of \$107.52 compared with \$70.64 for UD bidders. On the other hand, MU bidders tend to have a lower *per-auction* participation cost, as they typically participate in multiple auctions within a day. Evidence suggests that the MU bidders' higher cost to visit on a day is associated with having to review, compare, and match downstream channels for multiple auctions: MU bidders, on average, spend 87.8 minutes on the platform per day, compared with 75.9 minutes for UD bidders. Depending on the profit margin of the bidders' downstream resale channels, participation costs in the order of \$100 can be substantial, thus affecting their incentives to participate on the platform.¹⁷

To validate the model, we compare the predicted distributions of the number of bidders per auction and the final price in the post-treatment period with those observed in the dataset. Our model is fairly accurate in predicting both distributions (see Appendix C.5).

¹⁷ The estimates obtained for the bidders' participation costs may find use in designing platform's incentive schemes (e.g., discounts or free-shipping coupons) as a way to boost bidders' participation rates.

6. Implications for Platform Design

We explore how our findings from Section 5 lead to implications for platform design, with a focus on relating the performance of different listing policies to their induced levels of market thickness.

A platform’s listing policy influences revenues by manipulating the induced market thickness. Separate strands in the existing literature contemplate dual, but countervailing, effects from market thickness/availability on the behaviors of potential buyers and on revenues. For example, using a dataset of notebook auctions, [Chan et al. \(2007\)](#) suggest that increasing the number of concurrently available and substitutable products reduces bidders’ willingness to pay by up to 10.2%. On the other hand, the operations literature suggests that higher product availability may increase the seller’s revenue by stimulating demand (e.g., [Su and Zhang \(2009\)](#)). In a matching context, [Gan and Li \(2016\)](#) present evidence that thicker markets enhance efficiency by raising matching probabilities. In this subsection, we decompose and analyze the contending effects of supply-side cannibalization and demand-side participation in response to market thickness. Within this framework, we study how market thickness can be adjusted through the platform’s listing policy to balance these effects.

In simulating counterfactuals, we evaluate the performance of four listing policies that differ in how auction ending times are distributed throughout the week under six different levels of incoming supply (which together induce a market thickness level). The four listing policies are

- (i) Uniform: The number of auctions ending on each weekday is roughly the same.
- (ii) Three-Day Batch: Auctions end only on Mondays, Wednesdays, and Fridays.
- (iii) Batch: Auctions end only on Tuesdays and Thursdays.
- (iv) Single-Day Batch: Auctions end only on Wednesdays.

Given the same supply level, market thickness on auction-ending days increases by switching from policy (i) to policy (iv). We consider six supply levels, denoted by “ $\frac{1}{3}x$,” “ $\frac{1}{2}x$,” “ $1x$ ” (baseline), “ $2x$,” “ $3x$,” and “ $5x$.” In particular, we first derive our “baseline” supply case by fitting Market A’s supply data to a gamma distribution¹⁸ in order to appropriately account for variability observed in the data. To derive the remaining five supply levels, we simply scale the average supply level of the baseline case by the corresponding factor while keeping the coefficient of variation the same. For our entire counterfactual analysis, we keep the demand side (i.e., the bidders’ arrival process) fixed and defined by the non-parametrically estimated process for Market A’s post-treatment period.¹⁹ Table 12 summarizes our simulation results. For each combination pairing a supply level with a listing policy, we present the ratio of the revenue obtained relative to the revenue under the batch

¹⁸ The supply process during the post-treatment period is highly variable (as shown in the left plot of Figure 15), and we find that the gamma distribution fits the data reasonably well.

¹⁹ Note that for each listing policy (or recommendation system in Section 6.1), we simulate the equilibrium beliefs for the market state and the highest rival bids (i.e., Ψ and G), using the iterative algorithm of Appendix D.4. The resulting (equilibrium) steady-state distributions are consistent with the bidders’ strategies.

Table 12 Each Listing Policy’s Revenues as Percentage of Revenue under Batch Listing

Supply Level	Uniform	Three-Day Batch	Batch	Single-Day Batch
$\frac{1}{3}x$	90.1% [0.6]	96.5% [1.1]	100% [1.6]	106.1% [3.2]
$\frac{1}{2}x$	91.3% [0.9]	96.6% [1.5]	100% [2.2]	105.5% [4.4]
1x (baseline)	93.5% [2.1] (\$−299, 581)	97.2% [3.4] (\$−129, 050)	100% [5.2]	104.8% [10.3] (\$221, 229)
2x	97.8% [4.2]	101.0% [6.9]	100% [10.5]	95.0% [20.6]
3x	109.6% [6.3]	110.8% [10.3]	100% [15.7]	88.4% [31.0]
5x	122.1% [10.1]	117.8% [16.7]	100% [25.4]	81.9% [50.2]

Note. Values in bold correspond to the row-optimal listing policies with a significance level of 0.001. Market thickness averages are in brackets.


listing policy at that supply level. Therefore, each row of Table 12 reports the relative performance of the four listing policies at the row’s associated supply level. In the brackets next to the relative revenues, we display the average market thickness associated with each case.

At the baseline supply level, our counterfactual revenue improvement due to switching Market A from the uniform listing policy to the batch listing policy is roughly 7.0%, offering out-of-sample validation of the 7.3% increase estimated in our ATE analysis of Section 3.3 using time-wise separate data. Besides, the optimal policy in the baseline case turns out to be the single-day batch policy, which yields 4.8% more revenue than the batch policy. These relative differences in performance translate into substantial revenue gains. Over the 10 months of the post-treatment period, Market A’s revenue from iPhone 4, iPhone 4s, and iPhone 5 auctions amounted to \$4,608,941: the 11.3% relative difference between the uniform and the best-performing single-day policy translates to \$520,810 of additional revenue for the platform.²⁰

The relative performance of the four policies and consequently which listing policy is “optimal” hinges on the platform’s underlying supply level. For low levels of supply, “ $\frac{1}{3}x$ ” and “ $\frac{1}{2}x$,” the single-day batch policy performs best by inducing the thickest market. In this scenario, the participation cost exhibits a dominant effect on the platform’s revenues by deterring bidders’ visits; thus, the platform finds it optimal to increase the bidders’ expected payoff per visit by providing more options each auction-ending day. By contrast, for high levels of supply (e.g., “5x”), the uniform policy and the three-day batch policy perform relatively better than the rest of the policies by maintaining a thinner market. In this case, the cannibalization between auctions becomes the dominant effect; thus, the platform finds it optimal to spread out the auction ending times throughout the week.

The market thickness induced by the platform’s listing policy impacts revenues. Intuitively, platform revenues are maximized when the platform’s market thickness comes close to the level that effectively balances its trade-off of overcoming bidder participation costs (thus, inducing higher

²⁰ Perhaps reflecting our projections, as of January 2017, Market A enforced a single-day listing policy that ends all auctions exclusively on Tuesdays.

Figure 13 Effect of Market Thickness on Auctions' Bidder Participation and Final Prices (gray curve-95% CI)

participation rates) against competition as substitutes among auctions on the supply side. Using simulated data from the above counterfactual simulations, Figure 13 (the left plot) depicts an auction's average final per-device revenue as a function of the market's (log) average market thickness. Interestingly, low market thickness does not lead to a high average final price, due primarily to the participation costs and the hurdles to entry they present. As shown in the middle plot of Figure 13, the expected number of MU bidders per auction drops from 10 to 4 when market thickness decreases from 7 auctions to 1. On the other hand, cannibalization is also evident in the same plot: the number of MU bidders in an auction keeps decreasing when market thickness increases beyond 10 auctions. By comparison, UD bidders are much less sensitive to the shift in market thickness. As displayed in the right plot of Figure 13, there is no statistically significant increase in the expected number of UD bidders per auction associated with higher market thickness. Illustrating the trade-off between demand-side participation and supply-side cannibalization as market thickness increases, the average final price is maximized at moderate levels of market thickness.


Lastly, Table 14 summarizes the effects of the policy switch on the bidders' equilibrium behavior, including the rate at which they visit the platform and participate in auctions and the size of their bids. As we do not have any data on bidders' website visits that did not result in bids, we employ a simulation approach to estimate the change in the bidders' rate of platform visits associated with the policy switch. As is clearly shown in Table 14, the increase in the platform's revenues is primarily due to the increase in traffic to the platform (which, consequently, results in higher auction participation rates). On the other hand, additional analysis in Appendix D.3 shows that the listing policy has little effect on bidders' bidding decisions.

6.1. Targeted Recommendations

In this section we discuss the design of recommendation systems for a competitive online market (e.g., online auction markets). Complementing the listing policy, a recommendation system is intended to mitigate cannibalization on days of high market thickness by incentivizing buyers

Table 14 Revenue Funnel of Participation Effects from Listing Policies

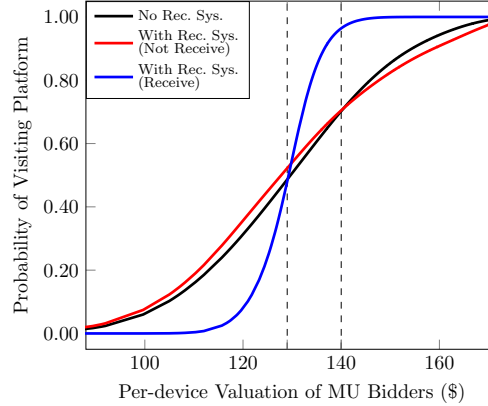
Mechanisms	Change in %	Evidence Type
Platform Visit	11.9%	Counterfactual Simulation
Auction Participation	9.7% (p -value = 0.03)	Observed
Bidding Decision	Not Significant	Observed

Figure 15 Variability of Daily Market Thickness and Demand Response in Market A

to visit the platform. As is evident in the left plot of Figure 15, daily market thickness (blue curve) exhibits large variations even though the listing policy is fixed. On the other hand, bidder participation does not scale up to match the realized supply; thus, the number of bidders per auction (red curve) drops on days where the realized supply is high (e.g., from Day 1 to Day 20 and on Days 55, 89, and 98).²¹ As shown in the right plot of Figure 15, the ratio of bidders who submit at least one bid to all active bidders in the bidder pool remains consistently below 60%. Therefore, the decline in the number of bidders per auction on these days is not the result of fluctuations in the overall bidder pool. Instead, we argue that the daily supply-demand mismatches are mainly due to the fact that it is costly for bidders to continuously monitor the market and obtain up-to-date information about the realized supply; thus, they decide whether to visit the platform based on their beliefs about the market in steady state. Namely, they cannot account for the inherent variability in the *realized* number of auctions available on the platform.

Given this, the platform can benefit from communicating credible market thickness information to bidders as a way to incentivize additional visits. Yet, sharing information about the daily supply to the entire bidder pool may actually backfire. Specifically, our counterfactual simulation of the *full information* policy, which discloses the realized market thickness daily to all bidders, indicates that in the resulting equilibrium the final price drops by 28.3% (on days when the realized supply is low, bidder participation drops significantly and, consequently, the final price per device drops).

²¹ The left plot of Figure 15 also implies that sellers do not strategize the weekly supply or the time of their listings to match the demand.

Figure 16 Effects of Recommendations on Marketplace Participation

To remedy the adverse effect of the low-supply information revealed from the communication, we consider recommendation systems that send recommendations to disclose the market state only when the market thickness is above threshold κ and only to η fraction of randomly-selected MU bidders. We focus exclusively on MU bidders because they win most auctions (Table 3) and UD bidders are much less responsive to market thickness information (right plot in Figure 13).

If a bidder receives a recommendation, she will update her belief about the realized supply of auctions on the platform, taking into account both the platform’s revealed supply state and how other bidders who received the same recommendation may respond. If she does not receive a recommendation, she will still update her belief by inferring that one of the following two cases must be true: (1) the market thickness on that day is lower than κ ; (2) The market thickness exceeds κ , but she was not selected to receive the information. Due to Case (1), the bidder will adjust her belief about market thickness downwards if she does not receive a recommendation. Our counterfactual will elucidate how this affects the participation decisions of various types of bidders.

We optimize over parameters κ and η in Market A, which operates under its optimal single-day batch policy. The optimal system sends recommendations to 10% of the MU bidders when the market thickness on a given day is above 22. Recommendations are sent out roughly on 11.0% of the days in the simulation horizon. Implementing this recommendation system further improves the platform’s overall revenue by 0.9%, including a 1.6% increase during the recommendation days.

Figure 16 plots MU bidders’ equilibrium rates of visiting the platform in the presence and absence of (optimal) recommendations. Bidders’ responses to receiving a recommendation (blue curve) vary drastically depending on their valuations. For bidders with lower valuations (below \$131), receiving a recommendation is actually bad news, and substantially decreases their probability of visiting. By contrast, bidders with higher valuations (above \$131) become significantly more likely to visit the platform after receiving a recommendation. Since the supply of auctions is higher on days when recommendations are sent out, high-valuation bidders’ expected payoff per visit is higher. Besides,

their increased platform visits intensify the competition and deters the low-valuation bidders. Thus, the optimal recommendation system increases revenues mainly by altering the composition of participants through self-selection on the market’s high-inventory days. Figure 16 also reveals how bidders response to not receiving a recommendation (red curve). First, a high-valuation bidder (above \$140) who did not receive a recommendation in the presence of a recommendation system is slightly less likely to visit the platform than in the absence of a recommendation system (black curve). Such a bidder infers that the supply of auctions may be low, which, in turn, weakens her incentive to participate. On the other hand, for low-valuation bidders (below \$140), not receiving the recommendation can be interpreted as facing less competition on the platform. Thus, the incentives for participation for low-valuation bidders are (slightly) higher.

7. Concluding Remarks

In this paper, we empirically illustrate the role of a platform’s design levels in inducing the optimal level of market thickness and, consequently, in generating higher revenues. In particular, we highlight that optimizing a seemingly innocuous market design choice (i.e., the listing policy) affects revenues significantly by inducing the appropriate level of market thickness on the marketplace. We also explore the design of a recommendation system that selectively informs bidders about the market state and establish that it can benefit the B2B auction platform by mitigating cannibalization among substitutable auctions.

Using a proprietary dataset, obtained from a leading online B2B auction platform, we estimate that inducing higher market thickness (by concentrating the auctions’ ending times on certain weekdays) leads to a 7.3% increase in the platform’s revenues. Additional analysis points to the presence of a significant participation cost associated with visiting the platform and bidding in open auctions that adversely affects bidders’ participation rates. Motivated by the descriptive results, we develop and estimate a structural model, which endogenizes bidders’ decision-making including whether and when to visit the platform, which auction(s) to participate in, and how much to bid. Notably, the revenue impact of inducing higher market thickness predicted by counterfactual simulation on the estimated model is consistent with the results from the reduced-form analysis.

Complementary to illustrating the revenue impact of the listing policy, we discuss the design of a recommendation system, which alters the composition of participants through self-selection. Appropriately designing the system yields an additional revenue increase by successively increasing the level of competition between bidders on days when the incoming supply is significantly higher than average.

More broadly, our work highlights that marketplace design can have significantly positive revenue implications for online two-sided platforms by mitigating frictions that impede participation. Given

their growing prominence, we believe that exploring the impact of platform design, especially focusing on non-price levers, on market thickness and, consequently, revenues and welfare in the context of other two-sided marketplaces is a very fruitful direction for future research.

References

- Akbarpour, Mohammad, Shengwu Li, Shayan Oveis Gharan. 2018. Thickness and information in dynamic matching markets. Working paper, Stanford University.
- Alexandrov, Alexei, Martin A. Lariviere. 2012. Are reservations recommended? *Manufacturing & Service Oper. Management* **14**(2) 218–230.
- Allon, Gad, Achal Bassamboo. 2011. Buying from the babbling retailer? The impact of availability information on customer behavior. *Management Sci.* **57**(4) 713–726.
- Arnosti, Nick, Ramesh Johari, Yashodhan Kanoria. 2014. Managing congestion in dynamic matching markets. *ACM Conf. Econom. Computation.* 451–451.
- Backus, Matthew, Gregory Lewis. 2016. Dynamic demand estimation in auction markets. Tech. rep., National Bureau of Economic Research.
- Bajari, Patrick, Ali Hortacsu. 2003. The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econom.* **34**(2) 329–355.
- Balseiro, Santiago R., Omar Besbes, Gabriel Y. Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* **61**(4) 864–884.
- Bapna, Ravi, Paulo Goes, Alok Gupta, Yiwei Jin. 2004. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS Quart.* **28**(1) 21–43.
- Bimpikis, Kostas, Ozan Candogan, Daniela Saban. 2017. Spatial pricing in ride-sharing networks. *Oper. Res., forthcoming* .
- Bodoh-Creed, Aaron, Joern Boehnke, Brent Richard Hickman. 2016. How efficient are decentralized auction platforms? Working paper, University of Chicago.
- Bray, Robert L, Haim Mendelson. 2015. Production smoothing and the bullwhip effect. *Manufacturing & Service Oper. Management* **17**(2) 208–220.
- Cachon, Gerard P, Kaitlin M Daniels, Ruben Lobel. 2017. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Oper. Management* **19**(3) 337–507.
- Chan, Tat Y., Vrinda Kadiyali, Young-Hoon Park. 2007. Willingness to pay and competition in online auctions. *J. Marketing Res.* **44**(2) 324–333.
- Cullen, Zoë, Chiara Farronato. 2016. Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. Working paper, Harvard Business School.
- Dana, James D., Nicholas C. Petruzzi. 2001. Note: The newsvendor model with endogenous demand. *Management Sci.* **47**(11) 1488–1497.

-
- Etzion, Hila, Edieal J. Pinker. 2008. Asymmetric competition in B2B spot markets. *Production and Operations Management* **17**(2) 150–161.
- Ferdows, Kasra, Michael A. Lewis, Jose A.D. Machuca. 2005. Zara’s secret for fast fashion. *Harvard Bus. Rev.* **82**(11) 98–111.
- Fradkin, Andrey. 2017. Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb. Working paper, Massachusetts Institute of Technology.
- Gallino, Santiago, Antonio Moreno. 2014. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Sci.* **60**(6) 1434–1451.
- Gan, Li, Qi Li. 2016. Efficiency of thin and thick markets. *J. Econometrics* **192**(1) 40–54.
- Golrezaei, Negin, Hamid Nazerzadeh, Vahab Mirrokni. 2017. Boosted second price auctions for heterogeneous bidders. Working paper, Massachusetts Institute of Technology.
- Graham, Bryan S., Cristine Campos de Xavier Pinto, Daniel Egel. 2012. Inverse probability tilting for moment condition models with missing data. *Rev. Econom. Stud.* **79**(3) 1053–1079.
- Halaburda, Hanna, Mikolaj Jan Piskorski, Pinar Yildirim. 2017. Competing by restricting choice: The case of matching platforms. *Management Sci.* **64**(8) 3574–3594.
- Hirano, Keisuke, Guido W. Imbens. 2001. Estimation of causal effects using propensity score weighting: An application to data on right heart catheterization. *Health Services and Outcomes Res. Methodology* **2**(3) 259–278.
- Hirano, Keisuke, Guido W. Imbens, Geert Ridder. 2003. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* **71**(4) 1161–1189.
- Horton, John. 2017. Buyer uncertainty about seller capacity: Causes, consequences, and partial solution. Working paper, New York University.
- Imbens, Guido W. 2004. Nonparametric estimation of average treatment effects under exogeneity: A review. *Rev. Econom. Stat.* **86**(1) 4–29.
- Iyer, Krishnamurthy, Ramesh Johari, Mukund Sundararajan. 2014. Mean field equilibria of dynamic auctions with learning. *Management Sci.* **60**(12) 2949–2970.
- Jofre-Bonet, Mireia, Martin Pesendorfer. 2003. Estimation of a dynamic auction game. *Econometrica* **71**(5) 1443–1489.
- Kabra, Ashish, Belavina Elena, Girotra Karan. 2017. The efficacy of incentives in scaling marketplaces. Working paper, University of Maryland.
- Kanoria, Yash, Daniela Saban. 2017. Facilitating the search for partners on matching platforms: Restricting agent actions. Working paper, Stanford University.
- Kim, Sang Won, Marcelo Olivares, Gabriel Y. Weintraub. 2014. Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Sci.* **60**(5) 1180–1201.

- Li, Jun, Nelson Granados, Serguei Netessine. 2014. Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* **60**(9) 2114–2137.
- Li, Jun, Serguei Netessine. 2018. Higher market thickness reduces matching rate in online platforms: Evidence from a quasi-experiment. *Management Sci.*, *forthcoming* .
- Mendelson, Haim, Tunay I. Tunca. 2007. Strategic spot trading in supply chains. *Management Sci.* **53**(5) 742–759.
- Moon, Ken, Kostas Bimpikis, Haim Mendelson. 2018. Randomized markdowns and online monitoring. *Management Sci.* **64**(3) 1271–1290.
- Olivares, Marcelo, Christian Terwiesch, Lydia Cassorla. 2008. Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Sci.* **54**(1) 41–55.
- Petruzzi, Nicholas C., Kwan E. Wee, Maqbool Dada. 2009. The newsvendor model with consumer search costs. *Prod. Oper. Management* **18**(6) 693–704.
- Pilehvar, Ali, Wedad J. Elmaghraby, Anandasivam Gopal. 2016. Market information and bidder heterogeneity in secondary market online B2B auctions. *Management Sci.* **63**(5) 1493–1518.
- Pinker, Edieal, Abraham Seidmann, Yaniv Vakrat. 2010. Using bid data for the management of sequential, multi-unit, online auctions with uniformly distributed bidder valuations. *European Journal of Operational Research* **202**(2) 574–583.
- Robins, James M., Andrea Rotnitzky, Lue Ping Zhao. 1994. Estimation of regression coefficients when some regressors are not always observed. *J. Amer. Stat. Assoc.* **89**(427) 846–866.
- Rosenbaum, Paul R., Donald B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* **70** 41–55.
- Su, Xuanming, Fuqiang Zhang. 2009. On the value of commitment and availability guarantees when selling to strategic consumers. *Management Sci.* **55**(5) 713–726.
- Tadelis, Steven, Florian Zettelmeyer. 2015. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *Amer. Econom. Rev.* **105**(2) 886–905.
- Tunca, Tunay I. 2008. Information precision and asymptotic efficiency of industrial markets. *Journal of Mathematical Economics* **44**(9-10) 964–996.
- Weintraub, Gabriel Y., C. Lanier Benkard, Benjamin Van Roy. 2008. Markov perfect industry dynamics with many firms. *Econometrica* **76**(6) 1375–1411.
- Wooldridge, Jeffrey M. 2007. Inverse probability weighted estimation for general missing data problems. *J. Econometrics* **141**(2) 1281–1301.
- Yang, Pu, Iyer Krishnamurthy, Peter Frazier. 2016. Mean field equilibria for competitive exploration in resource sharing settings. *Proc. 25th Internat. Conf. on World Wide Web*. International World Wide Web Conferences Steering Committee, 177–187.

Appendix A: Natural Experiment

A.1. AIPW Estimator of the Average Treatment Effect

The propensity score e_{jt} of an auction j that gets listed at time t is its probability of being assigned to Market A on the basis of its vector of observable attributes AUC_{jt} and the listing time t at which the auction lot becomes available. To estimate the propensity score, we specify

$$\text{Logit}(e_{jt}) = \tilde{\beta}_0 + \tilde{\delta}_t + \tilde{\beta}^T AUC_{jt}, \quad (7)$$

where $\tilde{\delta}_t$ and $\tilde{\beta}$ respectively denote the week- t fixed effect and the attribute-coefficient vector. Using the dependent variable A_{jt} , we first estimate model (7) by logistic regression to obtain predicted propensity scores \hat{e}_{jt} .

The estimator of the ATEs (i.e., τ_{Pre} and τ_{Post}) is characterized as follows:

$$\hat{\tau} = N_{obs}^{-1} \sum_{jt} \left(\frac{A_{jt} \cdot LFP_{jt}}{\hat{e}_{jt}} - \frac{A_{jt} - \hat{e}_{jt}}{\hat{e}_{jt}} \cdot \widehat{LFP}_{jt,1} \right) - N_{obs}^{-1} \sum_{jt} \left(\frac{1 - A_{jt} \cdot LFP_{jt}}{1 - \hat{e}_{jt}} + \frac{A_{jt} - \hat{e}_{jt}}{1 - \hat{e}_{jt}} \cdot \widehat{LFP}_{jt,0} \right), \quad (8)$$

where N_{obs} is the period's sample size, corresponding to the pre- and post-treatment periods, respectively, and \hat{e}_{jt} is the estimated propensity score for auction jt obtained by Expression (7). We restrict attention to auctions with estimated propensity scores between 0.2 and 0.8 in order to ensure that each unit will not have a weight that is more than 0.025, and that the requisite overlap assumption holds (Imbens 2004). While the classic ATE estimate is derived as the difference in observed outcomes weighted appropriately using the associated propensity score projections \hat{e}_{jt} , Robins' AIPW estimator adds terms involving the projections $\widehat{LFP}_{jt,1}$ and $\widehat{LFP}_{jt,0}$. In turn, these projections are derived from estimating the following linear models of log-revenue outcomes, as an alternative to the classic ATE:


$$LFP_{jt,A} = \beta_{0,A} + \delta_{t,A} + \beta_A^T AUC_{jt,A} + \epsilon_{jt,A}, \quad (9)$$

where A is the indicator of Market A, $\beta_{0,A}$ is the intercept, $\delta_{t,A}$ is a fixed effect for week t , and β_A denotes the attribute-coefficient vector. Note that Robins' AIPW is a doubly robust estimator; that is, it offers a consistent estimator of the ATE if *either* the classic ATE propensity score model (i.e., Expression (7)) or the linear outcome model (i.e, Expression (9)) is well specified.

A.2. Assessing Overlap

The purpose of this section is to assess the overlap in the covariates of Markets A and B and, subsequently, to argue that our estimation approach is indeed valid. To this end, we restrict attention to observations that have predicted propensity scores between 0.2 and 0.8 and, first, present the corresponding distributions of predicted propensity scores and summary statistics.²² In our propensity score model, we include variables to capture weekly fixed effects (i.e., the time at which an auction was held), the auction lot size, and fixed effects for product models and carriers. We then calculate the differences between the means of each of the covariates in the treatment and control groups, respectively. A covariate with a difference in means that is greater than 0.25 standard deviations is considered to be lacking overlap.

²² Given the sample size, the selection of the cutoff points follows the rule that no observation will have a weight that is more than 0.05 (Imbens 2004)

Figure 17 Histograms of Propensity Scores in Samples of the Pre-treatment Period (left) and the Post-treatment Period (right).**Table 18** Overlap of Covariates in the Selected Samples of the Pre- and Post-treatment Periods

	Pre-treatment			Post-treatment		
	Ctr. Group	Trt. Group	Diff./S.D.	Ctr. Group	Trt. Group	Diff./S.D.
Auction Lot Size	77.34 (37.58)	80.52 (37.32)	0.08	63.29 (37.65)	92.70 (26.53)	0.78
iPh. 4-AT&T	.023 (.152)	.012 (.112)	-0.07	.075 (.264)	.084 (.278)	0.03
iPh. 4-Sprint	.192 (.394)	.194 (.395)	0.01	.104 (.306)	.069 (.254)	-0.11
iPh. 4-Verizon	.291 (.454)	.256 (.436)	-0.07	.037 (.191)	.045 (.207)	0.03
iPh. 4s-AT&T	.175 (.380)	.161 (.368)	-0.03	.162 (.369)	.259 (.438)	0.26
iPh. 4s-Sprint	.103 (.304)	.129 (.336)	0.08	.223 (.416)	.196 (.397)	-0.06
iPh. 4s-Verizon	.133 (.339)	.183 (.386)	0.14	.163 (.369)	.165 (.371)	0.01
iPh. 5-Sprint	.036 (.188)	.031 (.175)	-0.02	.154 (.361)	.105 (.307)	-0.13
iPh. 5-Verizon	.020 (.143)	.014 (.119)	-0.04	.032 (.177)	.039 (.194)	0.04

During the pre-treatment period, the treatment effect relates to unobserved differences between Markets A and B. As deduced from the left plot of Figure 17, there is a sufficiently large overlap between the treatment group (Market A) and the control group (Market B). In the first three columns of Table 18, we present the means of the covariates (with standard deviations in parentheses) and the cross-market differences for each covariate. All differences in means lie within 0.14 standard deviations, which indicates that the two groups are well balanced.

During the post-treatment period, the treatment effect relates to both the unobserved differences between Markets A and B and the listing policy switch. In the right plot of Figure 17, we provide a histogram with the predicted propensity scores of the treatment group (Market A) and the control group (Market B). Although the overlap is not as large as that in the pre-treatment period, we believe that the level of overlap in the post-treatment period is sufficient. As both distributions spread out within $[0.2, 0.8]$, we have observations from both groups in every propensity score bin. In the last three columns of Table 18, we present the covariate means (with standard deviations in parentheses) and the cross-group differences for each covariate. The differences in means for all covariates (with the exception of the auction lot size) are

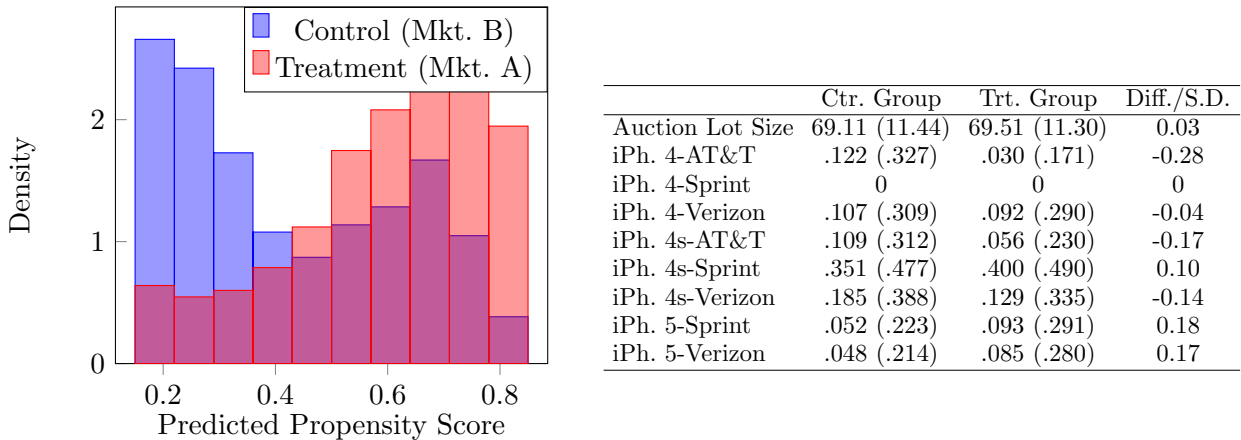
within 0.26 standard deviations. In general, an outcome regression model relies on control observations when predicting the outcome of treated units if they were *not* treated. A large difference in covariate distributions across groups implies that the predictions heavily rely on extrapolation; hence, they are sensitive to the specification of the outcome regression model.

To evaluate how our ATE estimation is affected by the lack of overlap in the auction lot size (in the post-treatment period), we conduct the following robustness checks:

- (i) First, we rerun the ATE analysis based on a subsample in which the treatment and control groups have a large overlap and there are sufficient observations to make statistically significant inferences. Although the assumption of overlap in the alternative approach is more credible compared to the original estimation, the estimated effect does not necessarily apply to all observations.
- (ii) Second, to test whether the estimation is sensitive to the specification of the outcome regression model, we conduct two additional robustness checks by replacing the auction lot covariate with its log scale value (to make it concave) and by adding a quadratic term for the auction lot (to make it convex).

For the subsample in (i), we select auctions with lot size between 50 and 90. This results in 49 observations for the control group and 58 observations for the treatment group. The subsample accounts for 27% of the original sample. To assess the overlap for the subsample, we present the histogram of predicted propensity scores (left plot of Table 19) and the differences in means of the covariates (right table of Table 19). Compared with the original sample, the overlap in the subsample is sufficiently higher.

Table 19 Histogram of the Predicted Propensity Scores in the Selected Subsample in the Post-treatment Period (left) and a Table Assessing the Overlap of Covariates in the Selected Subsample (right)



We perform the ATE analysis based on the subsample and then compare the result with that obtained from the original sample. The results can be found in Table 20. The estimated ATE of the subsample is significantly positive and on a similar scale to the ATE estimate corresponding to the original sample. Though the subsample estimate may suffer from selection bias, comparison of the two results suggests that the lack of overlap for the auction lot size covariate seems to have a very limited impact on the ATE estimation. In what follows, we conduct another analysis to check whether our ATE estimation is sensitive to the specification of the outcome regression model.

Table 20 ATEs of the Original Sample and the Selected Subsample

Auction lot size specification	Original Sample	Selected Subsample
Linear	7.3%*** (0.009)	8.1%*** (0.011)
<i>Note:</i>		*** p<0.01

Table 21 Robustness Check of Auction Lot Size Specification

Auction lot size specification	Original Sample
Linear	7.3%*** (0.009)
Concave (Log)	7.2%*** (0.009)
Convex (Quadratic)	7.3%*** (0.009)
<i>Note:</i>	*** p<0.01

The key issue corresponding to the lack of overlap between treated and controlled units is that the controlled outcome predictions of the treated units with outlying values heavily rely on extrapolation, as few control units are observed in this region when estimating the outcome regression model. In other words, the prediction precision may be sensitive to the specification of the outcome regression model. In our case, the predicted outcomes may be subject to the functional form of auction lot size. In our robustness check (ii), we consider concave (log) and convex (quadratic) specifications regarding the auction lot size in the outcome regression model. As shown in Table 21, the ATE estimates remain almost the same under different specifications. In particular, there is a strong linear relationship between the final price of an auction and its size. This is quite intuitive given that the cell phones sold in an auction are relatively homogeneous; thus, the price is likely to increase linearly with the lot size. Within a moderate lot size range, economies of scale and demand satiation effects are not significant; thus, the dependence between the final price and the lot size can be well described by a linear relationship. This implies that the estimated ATE is not significantly influenced by outliers, including those with auction lots larger than 90, whose predicted control outcomes heavily rely on extrapolation. In conclusion, our original ATE estimation remains reliable.

A.3. Cross-market Bidders

In this section, we establish that the presence of bidders that are cross-registered in both Markets A and B implies that our estimate on the revenue impact of the switch in the listing policy is conservative. In other words, if the markets were truly independent, then switching to a batch listing policy would result in a higher increase in the platform’s revenues than our estimated 7.3%. Our approach involves developing a model that features cross-market bidders and computing the platform’s revenues as a function of their share in the total population via simulation.

A.3.1. Model. First, we provide an outline of the model.

(a) **Model Setting and Assumptions.**

1. The platform has two markets, A and B. Each market has a single-market bidder pool of fixed size M . Both markets share a cross-market bidder pool of fixed size L . Each bidder has a private

valuation x drawn from a common distribution with cdf $F(\cdot)$, which has a positive support. The bidder has demand for K auctions upon entry into the pool.

2. The platform operates over an infinite horizon. All auctions are identical, second-price auctions and last for one day. Under the uniform policy, auctions are listed every day; under the batch policy, auctions are listed every other day.
3. Each day, a bidder first decides whether or not to participate on the platform. If she does, she incurs a participation cost c drawn from the exponential distribution with rate μ . If the bidder chooses to participate on a given day, she selects a number of auctions equal to her unsatisfied demand to participate at random. For each auction j she participates in, she places a bid equal to $x + \epsilon_j$, where ϵ_j is an idiosyncratic term drawn from $N(0, \sigma_\epsilon^2)$. At the end of the day, she exits the bidder pool with probability $1 - \rho$. As soon as a bidder exits, a new bidder joins so that the pool size remains constant.

(b) **Notation.** We use k and s_m to denote a bidder's demand that is still not satisfied and the number of auctions in Market m that the bidder has chosen to participate in, respectively. Moreover, we introduce the following notation:

1. $G_{mt}(y)$, $g_{mt}(y)$: the CDF and PDF of the highest rival bid in an auction in Market m on Day t .
2. $P_{mt}(w; x, k)$: the probability of a single-market bidder winning w auctions in Market m on Day t .
3. $P_t(w_A, w_B; x, s_A, s_B)$: the probability of a cross-market bidder winning w_A Market A auctions and w_B Market B auctions on Day t .
4. $u_{mt}(x, k)$: a single-market bidder's expected payoff in Market m on Day t .
5. $u_t(x, k)$: a cross-market bidder's expected payoff by visiting the platform on Day t .
6. $v_{mt}(x, k)$: a single-market bidder's aggregate payoff on Day t .
7. $v_t(x, k)$: a cross-market bidder's aggregate payoff on Day t .

(c) **Participation decisions when both markets implement the uniform policy.** Suppose that there are N_m auctions in Market m on each day. Since both markets use the uniform policy, we omit subscript t .

1. Single-market bidders: The expected payoff of a single-market bidder in Market m when she visits the platform is given by

$$u_m(x, k) = \min\{N_m, k\} \int_0^x (x - y)g_m(y)dy + \rho \sum_{w=0}^{\min\{N_m, k\}} P_m(w; x, \min\{N_m, k\})v_m(x, k - w). \quad (10)$$

The bidder visits the platform if $u_m(x, k) - c \geq \rho v_m(x, k)$, where $\rho v_m(x, k)$ is the payoff corresponding to waiting. Hence, her aggregate payoff is equal to

$$v_m(x, k) = \mathbb{E}[\max\{u_m(x, k) - c, \rho v_m(x, k)\}] = u_m(x, k) + \frac{1}{\mu} \exp\left(-\mu(u_m(x, k) - \rho v_m(x, k))\right) - \frac{1}{\mu}. \quad (11)$$

Lastly, combining (10) and (11), we can solve for $u_m(x, k)$ and $v_m(x, k)$ numerically for $k = 1, 2, \dots, K$, given g_m, P_m and x .

2. Cross-market bidders: The expected payoff of a cross-market bidder when she visits the platform is given by

$$u(x, k) = \frac{1}{|\mathcal{C}(k, N_A, N_B)|} \sum_{(k_A, k_B) \in \mathcal{C}(k, N_A, N_B)} \left(\sum_{m \in \{A, B\}} k_m \int_0^x (x-y) g_m(y) dy + \rho \sum_{w_A=0}^{k_A} \sum_{w_B=0}^{k_B} P(w_A, w_B; x, k_A, k_B) v(x, k - w_A - w_B) \right), \quad (12)$$

where $\mathcal{C}(k, N_A, N_B) = \{(k_A, k_B) : k_A + k_B = \min(k, N_A + N_B), k_A = 0, 1, \dots, N_A, k_B = 0, 1, \dots, N_B\}$. Then, she will visit the platform if $u(x, k) - c \geq \rho v(x, k)$, where $\rho v(x, k)$ is the payoff corresponding to waiting. Hence, her aggregate payoff is equal to

$$v(x, k) = \mathbb{E}[\max(u(x, k) - c, \rho v(x, k))] = u(x, k) + \frac{1}{\mu} \exp\left(-\mu(u(x, k) - \rho v(x, k))\right) - \frac{1}{\mu}. \quad (13)$$

Lastly, combining Expressions (12) and (13), we can solve for $u(x, k)$ and $v(x, k)$ numerically, for $k = 1, 2, \dots, K$, given g_A, g_B, P , and x .

- (d) **Participation decisions when Market A implements the uniform policy and Market B implements the batch policy.** In Market A, there are 0 auctions on odd days, denoted by $t = 1$ below, and $2N_A$ auctions on even days, denoted by $t = 2$ below. In Market B, there are N_B auctions every day.

1. Single-Market-A bidders: For a single-market bidder in Market A, she only visits the platform on even days. Then,

$$u_{A2}(x, k) = \min(2N_A, k) \int_0^x (x-y) g_{A2}(y) dy + \rho^2 \sum_{w=0}^{\min(2N_A, k)} P_{A2}(w; x, \min(2N_A, k)) v_{A2}(x, k - w). \quad (14)$$

She will visit the platform if $u_{A2}(x, k) - c \geq \rho^2 v_{A2}(x, k)$, where $\rho^2 v_{A2}(x, k)$ is the payoff corresponding to waiting. Hence, her aggregate payoff is equal to

$$v_{A2}(x, k) = \mathbb{E}[\max(u_{A2}(x, k) - c, \rho^2 v_{A2}(x, k))] = u_{A2}(x, k) + \frac{1}{\mu} \exp\left(-\mu(u_{A2}(x, k) - \rho^2 v_{A2}(x, k))\right) - \frac{1}{\mu}. \quad (15)$$

Lastly, combining (14) and (15), we can solve for $u_{A2}(x, k)$ and $v_{A2}(x, k)$ numerically for $k = 1, 2, \dots, K$, given g_{A2}, P_{A2} , and x .

2. Single-Market-B bidders: For a single-market bidder in Market B, her expected payoff of visiting the platform on day $t \in \{0, 1\}$ (let $t' \in \{0, 1\} \setminus \{t\}$) is

$$u_{Bt}(x, k) = \min(N_B, k) \int_0^x (x-y) g_{Bt}(y) dy + \rho \sum_{w=0}^{\min(N_B, k)} P_{Bt}(w; x, \min(N_B, k)) v_{Bt'}(x, k - w). \quad (16)$$

She will visit the platform if $u_{Bt}(x, k) - c \geq \rho v_{Bt'}(x, k)$, where $\rho v_{Bt'}(x, k)$ is the payoff corresponding to waiting. Hence, her aggregate payoff is equal to

$$v_{Bt} = \mathbb{E}[\max(u_{Bt}(x, k) - c, \rho v_{Bt'}(x, k))] = u_{Bt}(x, k) + \frac{1}{\mu} \left(-\mu(u_{Bt}(x, k) - \rho v_{Bt'}(x, k)) \right) - \frac{1}{\mu}. \quad (17)$$

Lastly, combining Expressions (16) and (17), we can solve for $u_{B1}(x, k), u_{B2}(x, k), v_{B1}(x, k)$, and $v_{B2}(x, k)$ numerically for $k = 1, 2, \dots, K$, given $g_{B1}, g_{B2}, P_{B1}, P_{B2}$, and x .

3. **Cross-market bidders:** For a cross-market bidder visiting the platform, her expected payoffs on odd and even days are given as follows:

$$u_1(x, k) = \min(k, N_B) \int_0^x (x-y)g_{B1}(y)dy + \rho \sum_{w_B=0}^{\min(k, N_B)} P_1(0, w_B; x, 0, \min(k, N_B))v_2(x, k-w_B), \quad (18)$$

$$\text{and } u_2(x, k) = \frac{1}{|\mathcal{C}(k, 2N_A, N_B)|} \sum_{(k_A, k_B) \in \mathcal{C}(k, 2N_A, N_B)} \left(\sum_{m \in \{A, B\}} k_m \int_0^x (x-y)g_{m2}(y)dy + \rho \sum_{w_A=0}^{k_A} \sum_{w_B=0}^{k_B} P_2(w_A, w_B; x, k_A, k_B)v_1(x, k-w_A-w_B) \right), \quad (19)$$

where $\mathcal{C}(k, 2N_A, N_B) = \{(k_A, k_B) : k_A + k_B = \min(k, 2N_A + N_B), k_A = 0, 1, \dots, N_A, k_B = 0, 1, \dots, N_B\}$.

On day $t \in \{0, 1\}$ (with $t' \in \{0, 1\} \setminus t$), she will visit the platform if $u_t(x, k) - c \geq \rho v_{t'}(x, k)$, where $\rho v_{t'}(x, k)$ is the payoff corresponding to waiting. Hence, her aggregate payoff on Day t is equal to

$$v_t(x, k) = \mathbb{E}[\max(u_t(x, k) - c, \rho v_{t'}(x, k))] = u_t(x, k) + \frac{1}{\mu} \exp\left(-\mu(u_t(x, k) - \rho v_{t'}(x, k))\right) - \frac{1}{\mu}. \quad (20)$$

Lastly, combining Expressions (18), (19), and (20), we can solve for $u_1(x, k)$, $u_2(x, k)$, $v_1(x, k)$, and $v_2(x, k)$ numerically for $k = 1, 2, \dots, K$, given $g_{A2}, g_{B1}, g_{B2}, P_1, P_2$ and x .

- (e) **Simulation.** We provide a brief summary of how we simulate bids on each day (through Day T):

- Day 0: We initialize $g_{mt}(y)$, $P_{mt}(w; x, k)$, and $P_t(w_A, w_B; x, s_A, s_B)$ for the first $2M$ single-market bidders and L cross-market bidders in the three bidder pools.
- Day t : First, we update the bidder pools by replacing the bidders who exit with new bidders. The new bidders form their beliefs about the state of the market, i.e., $g_{mt}(y)$, $P_{mt}(w; x, k)$, and $P_t(w_A, w_B; x, s_A, s_B)$ based on all bids between Day 1 and Day t . Then, all active bidders make platform-visit decisions and, if they choose to visit the platform, they place bids in a random subset of the auctions. When all bids are placed, auctions are allocated to their winners and we update the latter's demand that has not been yet satisfied.

A.3.2. Simulation Results.

- (a) **Parameters.** Our simulation scenarios are based on the following parameters:

$$\begin{array}{llll} N_A = N_B = 3 & 2M + L = 40 & F = N(50, 15) & F_\epsilon = N(0, 1) \\ K = 8 & \mu = 0.01 & \rho = 0.8 & T = 2,000. \end{array}$$

We chose the parameters above in order to make our simulation scenarios relatively comparable to what we observe in the dataset. We then simulate the model for 1,000 days and analyze the corresponding bids.

- (b) **Arbitrage effect of cross-market bidders.** We evaluate the effect of the policy switch under a number of scenarios corresponding to different ratios for the cross-market bidders, ranging from 0% to 80%. Under the parameters noted above, Market A and Market B are identical when they both implement the uniform policy. Thus, we only need to consider the cross-market price difference when Market A is under the batch policy, while Market B is under the uniform policy.

Table 22 Average Final Price per Market Corresponding to Different Simulation Scenarios

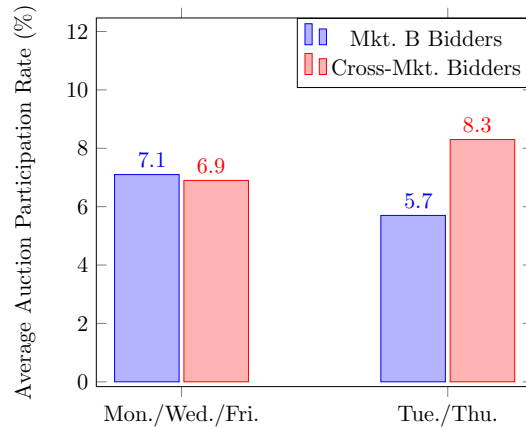
Ratio of cross-market bidders	0%	10%	20%	30%	60%	80%
Market A	54.64	55.63	56.81	57.66	59.49	60.67
Market B	49.37	50.61	51.92	53.02	55.67	57.25
Abs. Change	5.28	5.03	4.89	4.64	3.82	3.42
Change in %	10.7%	9.9%	9.4%	8.8%	6.9%	6.0%

As is evident in Table 22, the revenue increase corresponding to the policy switch is decreasing in the proportion of cross-market bidders in the population of bidders. This is due to the fact that cross-market bidders play the role of an arbitrageur and tend to shrink the price gap between the two markets. In particular, when Market A auctions are expected to have higher final prices than those in Market B, a cross-market bidder will certainly choose to bid in Market B, thus intensifying the competition in Market B and resulting in a decrease in the price gap between the two markets. In summary, this analysis implies that our estimate of the revenue impact of the policy switch in Section 3.3 is conservative compared to the case where the two markets are completely independent (i.e., when there are no cross-market bidders). Therefore, designating our setting as a “natural experiment” is well justified.

A.4. Additional Evidence of Participation Costs

We provide additional evidence for participation costs driving the revenue increase in Market A by examining the platform-visit decisions of cross-market bidders. To that end, we examine whether cross-market bidders exhibit different participation patterns in Market B across weekdays before and after the adoption of the batch policy in Market A. In particular, we compare bidders’ participation rates in Market B on Tuesdays and Thursdays to those on Mondays, Wednesdays, and Fridays by computing the ratio of the total number of participants on Tuesdays and Thursdays over the total number of participants on the remaining weekdays. Assuming that cross-market bidders are equally likely to participate in the market on any given weekday would imply a ratio of 2/3. During the pre-treatment period, we observe the cross-market bidders’ participation ratio to be 63% (i.e., close to 2/3) in Market B, consistent with the use of the uniform listing policy in both markets. However, while Market B persisted in employing the uniform listing period throughout the post-treatment period, its participation ratio for cross-registrants increased to 84%, suggesting that these bidders strongly preferred to participate in Market B on the auction-clearing days of the *other* market (i.e., Tuesdays and Thursdays in Market A). Within the same period, an average cross-market bidder bids in a higher fraction of available auctions in Market B on Tuesday/Thursday than on Monday/Wednesday/Friday (as shown in Figure 23).²³ This spillover effect in the cross-market bidders’ participation rates suggests that the costs involved in visiting the platform, carefully examining the inventory of open auctions, and placing a bid are substantial. Thus, bidders are strategic in their decision to visit the platform and actively participate in auctions. In other words, bidders choose to visit the platform only when their expected payoffs from doing so exceed some threshold.

²³ We formally analyze this spillover effect in the cross-market bidders’ participation rate using a difference-in-differences methodology in Appendix D.1.

Figure 23 Average Auction Participation Rates in Market B in the Post-treatment Period

In summary, the above discussion (including Section 3.4) suggests the following mechanism to explain the revenue increase: an increase in market thickness boosts bidder participation both on aggregate and per auction. In turn, this additional traffic results in higher revenues for the platform. Underlying this mechanism is the fact that visiting the platform on a given day, monitoring the set of available auctions, and choosing whether and how much to bid entails a significant cost for bidders. In other words, although one expects that the auction platform would eliminate frictions and an auction's ending time would not affect its final price, we demonstrate that this is not the case (at least in the liquidation auctions). Thus, optimizing over listing policies (i.e., auctions' ending times) brings significant benefits to the platform.

Appendix B: Proofs

B.1. Characterizations of $r^{MU}(x_\ell; G, \psi)$, $r^{UD}(x_\ell; G, \Psi)$, $v^{UD}(x_\ell; G, \Psi)$, and $v_f(x_\ell; G, \Psi)$

When a MU bidder ℓ participates in the platform and observes market state $\omega_{\ell t}$, her conditional payoff becomes

$$u^{MU}(x_\ell; \zeta_\ell, \omega_{\ell t}, G) = \max_{\sigma_{SLT}, \sigma_{BID}} \sum_{j \in \sigma_{SLT}} \int_{s_j}^{\sigma_{BID}, j} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j. \quad (21)$$

Therefore, her expected payoff for visiting the platform is

$$r^{MU}(x_\ell; G, \Psi) = \int_{\omega} \int_{\zeta_\ell} u^{MU}(x_\ell; \zeta_\ell, \omega_{\ell t}, G) d\mathbf{F}_\zeta(\zeta_\ell) d\Psi(\omega_{\ell t}). \quad (22)$$

For a UD bidder ℓ , her conditional payoff after visiting the platform is given by

$$u_j^{UD}(x_\ell; \zeta_{\ell j}, \omega_{\ell t}, G, \Psi) = \max_{\sigma_{BID}, j} \int_{s_j}^{\sigma_{BID}, j} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j + \alpha^{UD} (1 - G_j(b_{\ell j} | \omega_{\ell t})) v_f(x_\ell; G, \Psi). \quad (23)$$

Therefore, her aggregated payoff for participating in the platform is:

$$v^{UD}(x_\ell; G, \Psi) = \int_{\omega} \int_{\zeta_{\ell j^*}} u_{j^*}^{UD}(x_\ell; \zeta_{\ell j^*}, \omega_{\ell t}, G) dF_\zeta^{UD}(\zeta_{\ell j^*}) d\Psi(\omega_{\ell t}),$$

where $j^* = \sigma_{SLT}^{UD}(x_\ell; \zeta_\ell, \omega_{\ell t})$ denotes the auction in which it is optimal to bid. Lastly, a UD bidder's continuation payoff is characterized by the following Bellman equation:

$$v_f(x_\ell; G, \Psi) = \int_0^\infty \max \left\{ v^{UD}(x_\ell; G, \Psi) - c, \alpha^{UD} v_f(x_\ell; G, \Psi) \right\} \mu^{UD} e^{-\mu^{UD} c} dc. \quad (24)$$

B.2. Independence between \mathbf{b}_ℓ and $g_j(p_j | \omega_{\ell t})$

In this subsection, we show that Expressions (21) and (23) can account for the effect of \mathbf{b}_ℓ on the distribution of the highest rival bids.

Recall that the market state is denoted by $\omega_{\ell t} = (n_t, \mathbf{s}_{\ell t})$, and let w_j denote the current highest bid in auction j , which is unobservable. The conditional PDF of w_j is denoted by $h_j(w_j | n_t, \mathbf{s}_{\ell t})$. Let p_j denote the highest rival bid (i.e., the final price) in auction j , and let $k_j(p_j | n_t, \mathbf{s}_{\ell t})$ denote its conditional PDF. The expected payoff of a bidder that places a bid $b_{\ell j}$ in auction j is

$$\int_{s_j}^{b_{\ell j}} \left(\int_{w_j}^{b_{\ell j}} (x_\ell + \zeta_{\ell j} - p_j) k_j(p_j | n_t, w_j, \mathbf{s}_{\ell t, -j}) dp_j \right) h_j(w_j | n_t, \mathbf{s}_{\ell t}) dw_j,$$

where s_j denotes the current standing bid in auction j , and $\mathbf{s}_{\ell t, -j}$ is the vector of the standing bids excluding auction j . In turn, this is equal to:

$$\int_{s_j}^{b_{\ell j}} (x_\ell + \zeta_{\ell j} - p_j) \left(\int_{s_j}^{p_j} k_j(p_j | n_t, w_j, \mathbf{s}_{\ell t, -j}) h_j(w_j | n_t, \mathbf{s}_{\ell t}) dw_j \right) dp_j.$$

Let

$$g_j(p_j | n_t, \mathbf{s}_{\ell t}) = \int_{s_j}^{p_j} k_j(p_j | n_t, w_j, \mathbf{s}_{\ell t, -j}) h_j(w_j | n_t, \mathbf{s}_{\ell t}) dw_j.$$

We can verify that $g_j(p_j | n_t, \mathbf{s}_{\ell t})$ is a probability density function and that it is independent of \mathbf{b}_ℓ . ■

B.3. Proof of Proposition 1

To establish the existence of the equilibrium, we first establish that the mapping from beliefs about the highest rival bids to their actual distribution is continuous and compact. Then, we conclude that the mapping has a fixed point, using the Schauder fixed-point theorem.

Mapping Γ . We denote the unconditional CDF of the highest rival bids by \mathbf{G} . Its j^{th} component is the following $G_n^j(y_n^j) = \int_{\omega} G_n^{(j)}(y_n^j | \omega) d\Psi(\omega)$, which is the unconditional distribution of the highest rival bid y_n^j in the auction with the j^{th} -lowest standing bid. We then specify the mapping between \mathbf{G} and the resulting unconditional rival bids distribution $\hat{\mathbf{G}}$ as

$$\begin{aligned} \hat{\mathbf{G}}(\cdot) &= \Gamma(\cdot; \mathbf{G}) = P(\mathbf{f}_h(\mathbf{x}_{MU}, \mathbf{m}_{MU}, \mathbf{b}, \mathbf{m}_{UD}, o, n) \leq \cdot; \mathbf{G}) \\ &= \int \left(\sum_{o, n, \mathbf{m}_{MU}, \mathbf{m}_{UD}} \mathbb{1}(\mathbf{f}_h(\mathbf{x}_{MU}, \mathbf{m}_{MU}, \mathbf{b}(\mathbf{x}_{UD}; \mathbf{G}), \mathbf{m}_{UD}, o, n) \leq \cdot) \right. \\ &\quad \left. \cdot P(o, n) P^{MU}(\mathbf{m}_{MU}; \mathbf{x}_{MU}, \mathbf{G}) P^{UD}(\mathbf{m}_{UD}; \mathbf{b}(\mathbf{x}_{UD}; \mathbf{G}), \mathbf{G}) \right) P(d\mathbf{x}_{MU}, d\mathbf{x}_{UD}). \end{aligned}$$

The notation we use above can be summarized as follows:

- (i) o : the exogenous order of bidders in the pool visiting the platform.
- (ii) n : number of auctions (i.e., market thickness).
- (iii) \bar{N} : upper bound of n .
- (iv) $\mathbf{m}_{MU}, \mathbf{m}_{UD}$: sets of MU and UD bidders who visit the platform.
- (v) $\mathbf{x}_{MU}, \mathbf{x}_{UD}$: vectors of valuations of MU and UD bidders in the bidder pool.
- (vi) $\mathbf{b}(\mathbf{x}_{UD}; \mathbf{G})$: bidding function of UD bidders.
- (vii) $\mathbf{f}_h(\cdot)$: vector of functions that generate the final price in each auction.

We can ignore all discrete variables (i.e., $\mathbf{m}_{MU}, \mathbf{m}_{UD}, o$, and n) as they only take a finite number of values. It is sufficient to establish the continuity and compactness of the following mapping $\Gamma_{sub}: \mathcal{H} \rightarrow \mathcal{H}$, where \mathcal{H} is the space of probability CDFs, with fixed $\mathbf{m}_{MU}, \mathbf{m}_{UD}, o$, and n . Specifically,

$$\begin{aligned} \tilde{\mathbf{G}}(\cdot) &= \Gamma_{sub}(\cdot; \mathbf{G}) \\ &= \int \left(\mathbb{1}(\mathbf{f}_h(\mathbf{x}_{MU}, \mathbf{m}_{MU}, \mathbf{b}(\mathbf{x}_{UD}; \mathbf{G}), \mathbf{m}_{UD}, o, n) \leq \cdot) \right. \\ &\quad \left. \cdot P(o, n) P^{MU}(\mathbf{m}_{MU}; \mathbf{x}_{MU}, \mathbf{G}) P^{UD}(\mathbf{m}_{UD}; \mathbf{b}(\mathbf{x}_{UD}; \mathbf{G}), \mathbf{G}) \right) P(d\mathbf{x}_{MU}, d\mathbf{x}_{UD}). \end{aligned}$$

Note that $\Gamma = \sum \Gamma_{sub}$. In addition, we have the following specifications for the components of $\Gamma_{sub}(\mathbf{G})$:

$$\begin{aligned} P^{MU}(\mathbf{m}_{MU}; \mathbf{x}_{MU}, \mathbf{G}) &= \prod_{\ell^{MU}} P(c_{MU} \leq r^{MU}(x_{\ell^{MU}}; \mathbf{G}))^{\mathbb{1}(\ell^{MU} \in \mathbf{m}_{MU})} P(c_{MU} > r^{MU}(x_{\ell^{MU}}; \mathbf{G}))^{\mathbb{1}(\ell^{MU} \notin \mathbf{m}_{MU})}, \text{ and} \\ P^{UD}(\mathbf{m}_{UD}; \mathbf{b}, \mathbf{G}) &= \prod_{\ell^{UD}} P(c_{UD} \leq r^{UD}(b_{\ell^{UD}}; \mathbf{G}))^{\mathbb{1}(\ell^{UD} \in \mathbf{m}_{UD})} P(c_{UD} > r^{UD}(b_{\ell^{UD}}; \mathbf{G}))^{\mathbb{1}(\ell^{UD} \notin \mathbf{m}_{UD})}, \end{aligned}$$

where $r^{MU}(\cdot; \mathbf{G})$ and $r^{UD}(\cdot; \mathbf{G})$ denote the expected payoffs per market visit for MU and UD bidders, respectively. Furthermore, Proposition 2 (specified in Appendix B.4) implies that bidders simply choose to bid in the auction(s) with the lowest standing bid(s).²⁴ The expected payoffs per market visit for both types of bidders are:

$$r^{MU}(x_{\ell^{MU}}; \mathbf{G}) = \sum_{n=1}^{\bar{N}} p_n \sum_{j=1}^{\min(n, K)} \int_0^{x_{\ell^{MU}}} G_n^j(y) dy, \text{ and } r^{UD}(b_{\ell^{UD}}; \mathbf{G}) = \sum_{n=1}^{\bar{N}} p_n \int_0^{b_{\ell^{UD}}} G_n^1(y) dy,$$

²⁴ We provide proofs for the propositions in the remainder of the Appendix.

where ℓ^{MU} and ℓ^{UD} are indices corresponding to the MU and UD bidders. The probability that there are n auctions is denoted by p_n . The optimal bid $b_{\ell^{UD}}$ placed by UD bidder ℓ^{UD} is:

$$b_{\ell^{UD}} = x_{\ell^{UD}} - \frac{\alpha^{UD}}{1 - \alpha^{UD}} \left(r^{UD}(b_{\ell^{UD}}; \mathbf{G}) + \frac{1}{\mu^{UD}} \exp(-\mu^{UD} r^{UD}(b_{\ell^{UD}}; \mathbf{G})) - \frac{1}{\mu^{UD}} \right),$$

where μ^{UD} is the parameter of the participation cost distribution of the UD bidders.

By definition, the vector of functions that generate the final price in each auction can be denoted by $\mathbf{f}_h = (f_h^1, \dots, f_h^{\bar{N}})$, and the element f_h^j maps all bidders' auction selections and bidding decisions to the final price in auction j . For ease of exposition, we define f_h^j as a function of the ordered valuations of the participating bidders. In auction j , we use $x_{MU}^{j,(\ell)}, x_{UD}^{j,(\ell)}$ to denote the ℓ^{th} -highest valuation among MU and UD bidders, respectively, and we use \mathbf{x}_{MU}^j and \mathbf{x}_{UD}^j to denote the vectors of valuations of MU and UD bidders who choose auction j . We can write f_h^j as

$$f_h^j(\mathbf{x}_{MU}^j, \mathbf{x}_{UD}^j; \mathbf{G}) = \begin{cases} x_{MU}^{j,(2)}, & \text{if } x_{MU}^{j,(2)} > b(x_{UD}^{j,(1)}, \mathbf{G}) \\ b(x_{UD}^{j,(2)}, \mathbf{G}), & \text{if } b(x_{UD}^{j,(2)}, \mathbf{G}) > x_{MU}^{j,(1)} \\ x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}, \mathbf{G}), & \text{otherwise.} \end{cases}$$

So far, we have completed the specification of mapping Γ_{sub} . In addition, we note that the steady-state beliefs of the market state (i.e., $\Psi(\omega)$) and the highest rival bids (i.e., $G_j(y|\omega)$) affect bidders' behavior only through \mathbf{G} . Thus, to establish consistency, it is sufficient to show that \mathbf{G} is induced by bidders playing their optimal strategies. In what follows, we show the continuity and compactness of Γ_{sub} .

Continuity of Γ_{sub} . We prove the following three lemmas to establish the continuity of Γ_{sub} .

LEMMA 1. *Given \mathbf{m}_{MU} , mapping $\Gamma_1(\cdot; \mathbf{G}) \triangleq P^{MU}(\mathbf{m}_{MU}; \cdot, \mathbf{G})$ is uniformly bounded and Lipschitz continuous in \mathbf{G} .*

Proof: The mapping is uniformly bounded since $P_{MU}(\mathbf{m}_{MU}; \cdot, \mathbf{G})$ is a probability CDF. To show that Γ_1 is Lipschitz continuous, first note that for all $\mathbf{G}, \mathbf{G}' \in \mathcal{H}$ and for all $x_{\ell^{MU}} \in [0, B]$ (recall that B is the upper bound of the valuations), we have:

$$|r^{MU}(x_{\ell^{MU}}; \mathbf{G}) - r^{MU}(x_{\ell^{MU}}; \mathbf{G}')| \leq \sum_{n=1}^{\bar{N}} p_n \sum_{j=1}^{\min(n, K)} \int_0^{x_{\ell^{MU}}} |\mathbf{G}_n^j(y) - \mathbf{G}'_n^j(y)| dy \leq \bar{N}^2 B |\mathbf{G} - \mathbf{G}'|_{\infty}.$$

Therefore, $r^{MU}(\cdot; \mathbf{G})$ is uniformly bounded, and Lipschitz continuous in \mathbf{G} .

Second, using the fact that the participation costs follow the exponential distribution, we show that $\exp(-\mu^{MU} r^{MU}(\cdot; \mathbf{G}))$ is uniformly bounded and Lipschitz continuous in $r^{MU}(\cdot; \mathbf{G})$. Note that $0 \leq r^{MU}(x_{\ell^{MU}}; \mathbf{G}) \leq \bar{N}B$. Moreover, for any $u, u' \in [0, \bar{N}B]$, we have

$$|\exp(-\mu^{MU} u) - \exp(-\mu^{MU} u')| \leq \int_{u'}^u |\mu \exp(-\mu^{MU} t)| dt \leq \mu^{MU} |u - u'|,$$

and $\exp(-\mu^{MU} u) \in [0, 1]$. In turn, this implies that $\exp(-\mu^{MU} r^{MU}(\cdot; \mathbf{G}))$ is uniformly bounded and Lipschitz continuous in $r^{MU}(\cdot; \mathbf{G})$. From the definition of $P^{MU}(\mathbf{m}_{MU}; \cdot, \mathbf{G})$, which is equal to the product between $\exp(-\mu^{MU} r^{MU}(x_{\ell^{MU}}; \mathbf{G}))$ and $1 - \exp(-\mu^{MU} r^{MU}(x_{\ell^{MU}}; \mathbf{G}))$ for a finite number of bidders, we conclude that $P^{MU}(\mathbf{m}_{MU}; \cdot, \mathbf{G})$ is Lipschitz continuous in \mathbf{G} . ■

LEMMA 2. Given \mathbf{m}_{UD} , mapping $\Gamma_2(\cdot; \mathbf{G}) \triangleq P^{UD}(\mathbf{m}_{UD}; \mathbf{b}(\cdot; \mathbf{G}), \mathbf{G})$ is uniformly bounded and Lipschitz continuous in \mathbf{G} .

Proof: The mapping is uniformly bounded as $P^{UD}(\mathbf{m}_{UD}; \mathbf{b}(\cdot; \mathbf{G}), \mathbf{G})$ is a probability CDF. To establish Lipschitz continuity, we decompose Γ_2 into three parts:

- (i) Mapping $\Gamma_{2a}(\cdot; \mathbf{G}) \triangleq P^{UD}(\mathbf{m}_{UD}; \cdot, \mathbf{G})$;
- (ii) Function $f_{2b}(\mathbf{b}) \triangleq P^{UD}(\mathbf{m}_{UD}; \mathbf{b}, \mathbf{G})$ given \mathbf{G} ; and
- (iii) Mapping $\Gamma_{2c}(\cdot; \mathbf{G}) \triangleq \mathbf{b}(\cdot; \mathbf{G})$.

We show that each part is uniformly bounded and Lipschitz continuous separately. Note that \mathbf{m}_{UD} is fixed. The claim for $\Gamma_{2a}(\mathbf{G})$ follows using the same argument as in Lemma 1. For $f_{2b}(\mathbf{b})$, it is sufficient to show

$$P(c_{UD} \leq r^{UD}(b; \mathbf{G})) = 1 - \exp(-\mu^{UD} r^{UD}(b; \mathbf{G})),$$

is Lipschitz continuous in b given \mathbf{G} . This holds since

$$\left| \frac{\partial P(c_{UD} \leq r^{UD}(b; \mathbf{G}))}{\partial b} \right| = \left| \mu^{UD} \exp(-\mu^{UD} r^{UD}(b; \mathbf{G})) \sum_{n=1}^{\bar{N}} p_n \mathbf{G}_n^1(\mathbf{b}) \right| \leq \mu^{UD}.$$

Lastly, mapping $\Gamma_{2c}(\mathbf{G}) : \mathcal{H} \mapsto [0, B]^L$, where recall that L denotes the upper bound of the size of the bidder pool, is uniformly bounded. For a UD bidder with $x_{\ell UD} \in [0, B]$ and for all $\mathbf{G}, \mathbf{G}' \in \mathcal{H}$, according to the optimal bidding decision of UD bidders, we have

$$b = x_{\ell UD} - F(b, \mathbf{G}), \text{ and } b' = x_{\ell UD} - F(b', \mathbf{G}'),$$

where

$$F(b, \mathbf{G}) = \frac{\alpha^{UD}}{1 - \alpha^{UD}} \left(r^{UD}(b; \mathbf{G}) + \frac{1}{\mu^{UD}} \exp(-\mu^{UD} r^{UD}(b; \mathbf{G})) - \frac{1}{\mu^{UD}} \right).$$

Note that $b - b' + F(b, \mathbf{G}) - F(b', \mathbf{G}') = 0$, which further implies that

$$b - b' + F(b, \mathbf{G}) - F(b', \mathbf{G}') = F(b', \mathbf{G}') - F(b', \mathbf{G}).$$

The left-hand side of the above equation is equal to $\int_{b'}^b 1 + \frac{\partial F}{\partial b}(u, \mathbf{G}) du$. Thus,

$$\frac{\partial F}{\partial b}(u, \mathbf{G}) = \frac{\partial F}{\partial r^{UD}} \frac{\partial r^{UD}}{\partial b}(u, \mathbf{G}) = \frac{\alpha^{UD}}{1 - \alpha^{UD}} \left(1 - \exp(-\mu^{UD} r^{UD}(u; \mathbf{G})) \right) \left(\sum_{n=1}^{\bar{N}} p_n \mathbf{G}_n^1(u) \right) > 0,$$

which implies that $|b - b' + F(b, \mathbf{G}) - F(b', \mathbf{G}')| \geq |b - b'|$. On the other hand,

$$|F(b', \mathbf{G}') - F(b', \mathbf{G})| \leq \frac{2\alpha^{UD}}{1 - \alpha^{UD}} |r^{UD}(b'; \mathbf{G}') - r^{UD}(b'; \mathbf{G})| \leq \frac{2\alpha^{UD} \bar{N} B}{1 - \alpha} |\mathbf{G} - \mathbf{G}'|_{\infty}.$$

Thus,

$$|\mathbf{b} - \mathbf{b}'|_{\infty} \leq \frac{2\alpha^{UD} \bar{N} B}{1 - \alpha^{UD}} |\mathbf{G} - \mathbf{G}'|_{\infty}.$$

Finally, we conclude that $\Gamma_{2c}(\mathbf{G})$ is Lipschitz continuous in \mathbf{G} , using the triangular inequality. ■

LEMMA 3. Mapping Γ_3 corresponding to auction j and defined as follows:

$$\Gamma_3(\cdot; \mathbf{G}) \triangleq \int \mathbb{1}(f_h^j(\mathbf{x}_{MU}^j, \mathbf{x}_{UD}^j; \mathbf{G}) \leq \cdot) P^o(d\mathbf{x}),$$

is uniformly bounded and Lipschitz continuous in \mathbf{G} , given $o, n, \mathbf{m}_{MU}, \mathbf{m}_{UD}$.

Proof: For any $y^j \in [0, B]$, and for any $\mathbf{G}, \mathbf{G}' \in \mathcal{H}$, we have

$$\begin{aligned} \Gamma_3(y^j; \mathbf{G}) &= \tilde{G}(y^j) \\ &= \int \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(2)} \leq y^j) P^o(d\mathbf{x}) + \int \mathbb{1}(x_{MU}^{j,(1)} < b(x_{UD}^{j,(2)}; \mathbf{G}) \leq y^j) P^o(d\mathbf{x}) \\ &\quad + \int \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \geq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y^j) P^o(d\mathbf{x}). \end{aligned}$$

Then, $|\tilde{G}(y^j) - \tilde{G}'(y^j)| \leq I_1 + I_2 + I_3$, for I_1, I_2 , and I_3 , is defined as:

$$\begin{aligned} I_1 &= \int \left| \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(2)} \leq y^j) - \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}') < x_{MU}^{j,(2)} \leq y^j) \right| P^o(d\mathbf{x}) \\ I_2 &= \int \left| \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(2)} \leq y^j) - \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}') < x_{MU}^{j,(2)} \leq y^j) \right| P^o(d\mathbf{x}) \\ I_3 &= \int \left| \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \geq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y^j) \right. \\ &\quad \left. - \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}'), x_{MU}^{j,(1)} \geq b(x_{UD}^{j,(1)}; \mathbf{G}'), x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}') \leq y^j) \right| P^o(d\mathbf{x}). \end{aligned}$$

The claim follows by showing that I_1, I_2 , and I_3 are bounded by $\|\mathbf{G} - \mathbf{G}'\|_\infty$. For brevity, we establish that I_1 is bounded by $\|\mathbf{G} - \mathbf{G}'\|_\infty$ (the proofs for I_2 and I_3 follow a similar approach). Specifically,

$$\begin{aligned} &\int_{b(x_{UD}^{j,(1)}; \mathbf{G}) \geq x_{MU}^{j,(2)}, b(x_{UD}^{j,(1)}; \mathbf{G}') < x_{MU}^{j,(2)} \leq y^j} \mathbb{1} P^o(\mathbf{x}) \\ &\leq \int_{b(x_{UD}^{j,(1)}; \mathbf{G}') < x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G})} \mathbb{1} P^o(d\mathbf{x}) = \int F_{MU}^{j,(2)}(b(x_{UD}^{j,(1)}; \mathbf{G})) - F_{MU}^{j,(2)}(b(x_{UD}^{j,(1)}; \mathbf{G}')) P^o(dx_{UD}^{j,(1)}) \\ &\leq |f_{MU}^{j,(2)}|_\infty \int b(x_{UD}^{j,(1)}; \mathbf{G}) - b(x_{UD}^{j,(1)}; \mathbf{G}') P^o(dx_{UD}^{j,(1)}) \leq M_{MU}^{j,(2)} \cdot \frac{2\alpha \bar{N} B^2}{1 - \alpha_{UD}} \|\mathbf{G} - \mathbf{G}'\|_\infty, \end{aligned}$$

where $F_{MU}^{j,(2)}(x_{MU}^{j,(2)})$ and $f_{MU}^{j,(2)}(x_{MU}^{j,(2)})$ are the CDF and PDF of the second-highest valuation among MU bidders in auction j , respectively. The last inequality holds from Lemma 1, and $M_{MU}^{j,(2)} \triangleq \sup_{\mathbf{G}} |f_{MU}^{j,(2)}|_\infty < \infty$, which is independent of \mathbf{G} . ■

In summary, combining Lemmas 1, 2, and 3, we conclude that Γ_{sub} is uniformly bounded and Lipschitz continuous in \mathbf{G} . Therefore, Γ is a continuous mapping.

Compactness of Γ_{sub} . Here we show that Γ_{sub} is a compact mapping using the Arzela–Ascoli theorem.

To apply the theorem, Γ_{sub} needs to satisfy the following conditions:

1. The image $\Gamma_{sub}(\cdot; \mathcal{H})$ is uniformly bounded.
2. Sequence $\{\Gamma_{sub}(\mathbf{y}; \mathbf{G}_n)\}$ is equicontinuous in $\mathbf{y} \in [0, B]^{\bar{N}}$.

The fact that $\Gamma_{sub}(\mathcal{H})$ is uniformly bounded can be derived in a straightforward manner. To establish the equicontinuity of the mapping, we show that for any $\mathbf{G} \in \mathcal{H}$ and for any $\mathbf{y}, \mathbf{y}' \in [0, B]^{\bar{N}}$, we have

$$|\Gamma_{sub}(\mathbf{y}; \mathbf{G}) - \Gamma_{sub}(\mathbf{y}'; \mathbf{G})| \leq K_0 \cdot \|\mathbf{y} - \mathbf{y}'\|,$$

where K_0 is a constant, which is independent of \mathbf{G} . It is sufficient to show that the following inequality holds for the j^{th} auction given y_j, y'_j and fixed $\mathbf{m}_1, \mathbf{m}_2, n, o$:

$$|\Gamma_{sub,j}(y_j; \mathbf{G}) - \Gamma_{sub,j}(y'_j; \mathbf{G})| \leq J_1 + J_2 + J_3,$$

where J_1 , J_2 , and J_3 are defined as follows:

$$\begin{aligned} J_1 &= \int \left| \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(2)} < y_j) - \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(2)} < y'_j) \right| P^o(d\mathbf{x}) \\ J_2 &= \int \left| \mathbb{1}(x_{MU}^{j,(1)} < b(x_{UD}^{j,(2)}; \mathbf{G}) \leq y_j) - \mathbb{1}(x_{MU}^{j,(1)} < b(x_{UD}^{j,(2)}; \mathbf{G}) \leq y'_j) \right| P^o(d\mathbf{x}) \\ J_3 &= \int \left| \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y_j) \right. \\ &\quad \left. - \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j) \right| P^o(d\mathbf{x}). \end{aligned}$$

In each case, the difference within the absolute value takes value $-1, 0$, or 1 . For brevity, we only show that J_3 is bounded by $|y - y'|$ (Using similar arguments, one can show that J_1 and J_2 are bounded by $|y - y'|$ as well). We have

$$\begin{aligned} J_3 &= \int \mathbb{1}(x_{MU}^{j,(2)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), x_{MU}^{j,(1)} \leq b(x_{UD}^{j,(1)}; \mathbf{G}), y_j < x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j) P^o(d\mathbf{x}) \\ &\leq \int \mathbb{1}(y_j < x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j) P^o(d\mathbf{x}) \\ &= \int \left(\mathbb{1}(y_j < x_{MU}^{j,(1)} \leq y'_j) \cdot \mathbb{1}(x_{MU}^{j,(1)} \leq b(x_{UD}^{j,(1)}; \mathbf{G})) + \mathbb{1}(y_j < b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j) \cdot \mathbb{1}(b(x_{UD}^{j,(1)}; \mathbf{G}) < x_{MU}^{j,(1)}) \right) P^o(d\mathbf{x}) \\ &\leq \int_{y_j}^{y'_j} P^o(dx_{MU}^{j,(1)}) + \int_{y_j < b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j} P^o(dx_{UD}^{j,(1)}). \end{aligned} \quad (25)$$

For the first term in Expression (25), we have

$$\int_{y_j}^{y'_j} P^o(dx_{MU}^{j,(1)}) \leq \left| \frac{dF_{MU}^{j,(1)}}{dx_{MU}^{j,(1)}} \right|_{\infty} \cdot |y'_j - y_j| \leq M_{MU}^{j,(1)} \cdot |y'_j - y_j|,$$

where $M_{MU}^{j,(1)} \triangleq \sup_{\mathbf{G}} |f_{MU}^{j,(1)}|_{\infty} < \infty$. For the second term in Expression (25), we define vectors \mathbf{z}, \mathbf{z}' as the solutions to the following equations:

$$z_j = y_j + F(y_j, \mathbf{G}), \text{ and } z'_j = y'_j + F(y'_j, \mathbf{G}).$$

Recall that from Lemma 2:

$$F(b, \mathbf{G}) = \frac{\alpha^{UD}}{1 - \alpha^{UD}} \left(r^{UD}(b; \mathbf{G}) + \frac{1}{\mu^{UD}} \exp(-\mu^{UD} r^{UD}(b; \mathbf{G})) - \frac{1}{\mu^{UD}} \right).$$

Therefore,

$$\int_{y_j < b(x_{UD}^{j,(1)}; \mathbf{G}) \leq y'_j} P^o(dx_{UD}^{j,(1)}) = \int_{z_j < x_{UD}^{j,(1)} \leq z'_j} P^o(dx_{UD}^{j,(1)}) \leq M_{UD}^{j,(1)} \cdot |z'_j - z_j|,$$

where $M_{UD}^{j,(1)} \triangleq \sup_{\mathbf{G}} |f_{UD}^{j,(1)}|_{\infty} < \infty$. Finally,

$$|z'_j - z_j| \leq |y'_j - y_j| + |F(y'_j, \mathbf{G}) - F(y_j, \mathbf{G})| \leq \frac{1}{1 - \alpha^{UD}} |y'_j - y_j|.$$

Therefore, we establish that $\{\Gamma_{sub}(\mathbf{y}; \mathbf{G})\}$, $\mathbf{G} \in \mathcal{H}$, is equicontinuous in \mathbf{y} . Applying the Arzela–Ascoli theorem, we obtain that Γ_{sub} is a compact mapping. ■

Finally, employing the Schauder fixed-point theorem implies the existence of a fixed point such that $\Gamma(\mathbf{G}) = \mathbf{G}$. This completes the proof of the proposition. ■

B.4. UD Bidders' Auction Selection

PROPOSITION 2. *Assume that $\nu^{MU} = \nu^{UD} = 0$ and $K \geq \bar{N}$; the UD bidders' optimal auction selection decision is to choose the auction with the lowest standing bid.*

Proof. The assumptions of the proposition (i.e., $\nu^{MU} = \nu^{UD} = 0$ and $K \geq \bar{N}$) directly imply that MU bidders place a bid in all open auctions (provided that they place at least one bid). Given this observation, the proposition states that it is optimal for a UD bidder to bid in the auction with the lowest standing bid at the time she places her bid, if all other competing UD bidders also bid in the auction with the lowest standing bid at the time they place their bids. In other words, participating in the auction with the lowest standing bid (at the time she determines which auction to participate in) is an equilibrium strategy for a UD bidder.

Prior to proving the proposition, we first state and prove Lemmas 4, 5, and 6 that establish the following: if the *last* agent that places a bid is a UD bidder, it is optimal for her to bid in the auction with the lowest standing bid, assuming that all other UD bidders also place bids in the auctions with the lowest standing bids (at the time they decide which auction to participate in).

Suppose that there are n auctions on the platform, and they are ordered with respect to their standing bids (i.e., auction i 's standing bid is no greater than auction j 's if $i < j$). It is sufficient to show that the conditional CDFs of the current winning bids satisfy the following relationship:

$$G_1^0(w|\mathbf{s}) \geq G_2^0(w|\mathbf{s}) \geq \dots \geq G_n^0(w|\mathbf{s}), \quad (26)$$

where $G_j^0(w|\mathbf{s})$ is the conditional CDF of the current winning bid in the auction with the j^{th} -lowest standing bid. In what follows, we focus on the case where the number of UD bidders is known and denoted by N . One can extend the results to the case where N is unknown, by taking the expectation over the number of UD bidders on the platform.

We first start with the simplest case, where there are only two auctions and all bidders are UD bidders.

LEMMA 4. *Suppose that there are 2 identical auctions and $N + 1$ UD bidders placing bids in an exogenous sequence. It is optimal for the last UD bidder to bid in the auction with the lowest standing bid, assuming that the rest of the UD bidders also place bids in the auctions with the lowest standing bids (at the time they choose which auction to participate in).*

Proof. It is sufficient to show that, after the first N UD bidders place their bids, the conditional CDFs of the winning bids given the standing bids (s_1, s_2) , with $s_1 \leq s_2$, follow the relationship

$$G_1^0(w|s_1, s_2) \geq G_2^0(w|s_1, s_2),$$

where $G_j^0(w|s_1, s_2)$ is the conditional CDF of the winning bid in auction j .

We use X_1, X_2, \dots, X_N to denote the random variables corresponding to the bids of the first N bidders. The corresponding order statistics are denoted by $X_{(1)}, X_{(2)}, \dots, X_{(N)}$ (i.e., $X_{(\ell)}$ is the ℓ^{th} -largest bid among these N bids).

One important observation is that after all bids are submitted, the two standing bids (s_1, s_2) and the two (invisible) winning bids, denoted by (W_1, W_2) , are comprised of $X_{(1)}, X_{(2)}, X_{(3)}$, and $X_{(4)}$. We can verify this

observation by way of contradiction. In particular, suppose that the last four bids are $X_{(1)}, X_{(2)}, X_{(3)}$, and $X_{(\ell)}$ and $X_{(\ell)} < X_{(4)}$. Note that the 4th-highest bid can always be placed (since it is impossible to have both standing bids greater than $X_{(4)}$). Before $X_{(4)}$ is removed from these four remaining bids, it has to first serve as the standing bid of an auction. In this case, $X_{(4)}$ is the lowest standing bid among these two auctions, and it is outbid by a higher incoming bid (which can only be one of $X_{(1)}, X_{(2)}$, or $X_{(3)}$). In this case, bidding $X_{(\ell)}$ has no impact on the vector of standing bids, which yields the contradiction.

Given this observation, we note that there are the following three cases to consider:

Cases	(s_1, s_2)	(W_1, W_2)
1	$s_1 = X_{(4)}, s_2 = X_{(3)}$	$W_1 = X_{(1)}, W_2 = X_{(2)}$
2	$s_1 = X_{(4)}, s_2 = X_{(3)}$	$W_1 = X_{(2)}, W_2 = X_{(1)}$
3	$s_1 = X_{(4)}, s_2 = X_{(2)}$	$W_1 = X_{(3)}, W_2 = X_{(1)}$

In Case 1, we have

$$\Delta_1 \triangleq G_1^0(w|s_1, s_2, \text{Case 1}) - G_2^0(w|s_1, s_2, \text{Case 1}) = P(X_{(1)}|X_{(4)} = s_1, X_{(3)} = s_2) - P(X_{(2)}|X_{(4)} = s_1, X_{(3)} = s_2).$$

Similarly, in Case 2, we have:

$$\begin{aligned} \Delta_2 &\triangleq G_1^0(w|s_1, s_2, \text{Case 2}) - G_2^0(w|s_1, s_2, \text{Case 2}) = P(X_{(1)}|X_{(4)} = s_1, X_{(3)} = s_2) - P(X_{(2)}|X_{(4)} = s_1, X_{(3)} = s_2) \\ &= -\Delta_1. \end{aligned}$$

We also note that the conditional probabilities of a bidding sequence falling into Case 1 or Case 2 given (s_1, s_2) are identical (i.e., $P(\text{Case 1}|s_1, s_2) = P(\text{Case 2}|s_1, s_2)$). The reason is the following: for any bidding sequence that falls into Case 1, switching the order of bids $X_{(1)}$ and $X_{(2)}$ results in a bidding sequence that falls into Case 2, as (i) the decisions of these two bidders are the same, provided that their bids are larger than the standing bids, and (ii) both bids are invisible (thus, they do not alter the sequence of the following bids). In other words, the total number of bidding sequences that lead to Case 1 is the same as the corresponding number for Case 2. We also note that each bidding sequence has the same chance of occurring (as the order in which bidders place their bids is exogenous); thus, $P(\text{Case 1}|s_1, s_2) = P(\text{Case 2}|s_1, s_2)$ for any s_1, s_2 . Finally, in Case 3, we have:

$$\begin{aligned} \Delta_3 &\triangleq G_1^0(w|s_1, s_2, \text{Case 3}) - G_2^0(w|s_1, s_2, \text{Case 3}) \\ &= P(X_{(3)} \leq w|X_{(4)} = s_1, X_{(2)} = s_2) - P(X_{(1)} \leq w|X_{(4)} = s_1, X_{(2)} = s_2). \end{aligned}$$

Note that $\Delta_3 \geq 0$ since $\{X_{(1)} \leq w, X_{(4)} = s_1, X_{(2)} = s_2\} \subseteq \{X_{(3)} \leq w, X_{(4)} = s_1, X_{(2)} = s_2\}$. Summarizing the discussion above, we obtain that

$$G_1^0(w|s_1, s_2) - G_2^0(w|s_1, s_2) = P(\text{Case 1}|s_1, s_2)\Delta_1 + P(\text{Case 2}|s_1, s_2)\Delta_2 + P(\text{Case 3}|s_1, s_2)\Delta_3 \geq 0.$$

■

Next, we extend Lemma 4 to the case where there are n auctions, but still all bidders are UD.

LEMMA 5. *Suppose that there are n identical auctions and $N + 1$ UD bidders placing bids in an exogenous sequence. It is optimal for the last UD bidder to bid in the auction with the lowest standing bid, assuming that the rest of the UD bidders also place bids in the auctions with the lowest standing bids (at the time they choose which auction to participate in).*

Proof. It is sufficient to show that, after the first N bidders place their bids, the conditional CDFs of the winning bids given the standing bid vector \mathbf{s} follow the relationship specified in Expression (26). Equivalently, in the following proof, we show that for any $i < j$, we have

$$G_i^0(w|\mathbf{s}) \geq G_j^0(w|\mathbf{s}).$$

Observe that after the first N bids are submitted, the vector of standing bids \mathbf{s} and the vector of current winning bids \mathbf{W} are comprised of $X_{(\ell)}$, where $1 \leq \ell \leq 2n$. This follows using a similar argument as in the proof of Lemma 4.

Next, we consider the following three cases separately for auctions i and j .

Cases	(s_i, s_j)	(W_i, W_j)	
1	$s_i = X_{(l)}, s_j = X_{(k)}$	$W_i = X_{(L)}, W_j = X_{(K)}$	Given $K < L < k < l$ i.e., $X_{(l)} < X_{(k)} < X_{(L)} < X_{(K)}$ $\{\mathbf{X}_{-(l,k,L,K)} \in (\mathbf{s}_{-i,-j}, \mathbf{W}_{-i,-j})\}$
2	$s_i = X_{(l)}, s_j = X_{(k)}$	$W_i = X_{(K)}, W_j = X_{(L)}$	
3	$s_i = X_{(l)}, s_j = X_{(L)}$	$W_i = X_{(k)}, W_j = X_{(K)}$	

In the table, $\{\mathbf{X}_{-(l,k,L,K)} \in (\mathbf{s}_{-i,-j}, \mathbf{W}_{-i,-j})\}$ denotes the set of all valid assignments of the remaining $X_{(\ell)}$ ($1 \leq \ell \leq 2n$ and $\ell \neq l, k, L, K$) to the standing bids and the winning bids of all auctions excluding auctions i and j , which are denoted by $\mathbf{s}_{-i,-j}$ and $\mathbf{W}_{-i,-j}$, respectively.


Specifically, for a given assignment in Case 1, denoted by a_1 , we have:

$$s_i = X_{(l)}, s_j = X_{(k)}, W_i = X_{(L)}, W_j = X_{(K)} \text{ and } s_m = X_{m'}, W_m = X_{m''},$$

where $m \neq i, j, m' \neq l, k, m'' \neq L, K$, and $m' > m''$. Then, we have:

$$\begin{aligned} \Delta_1(a_1) &\triangleq G_i^0(w|s_i, s_j, a_1) - G_j^0(w|s_i, s_j, a_1) \\ &= P(X_{(L)} \leq w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \neq i, j, m' \neq l, k) \\ &\quad - P(X_{(K)} \leq w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \neq i, j, m' \neq l, k). \end{aligned}$$

Based on a given bidding sequence that results in an outcome in assignment a_1 , we can develop a new bidding sequence that leads to an outcome in the assignment described in Case 2, denoted by a_2 . Specifically, we can switch the order of two bidders who placed bids $X_{(L)}$ and $X_{(K)}$ in the bidding sequence of a_1 , to show that the new bidding sequence belongs to a_2 and leads to the same vector of standing bids \mathbf{s} as the original one in a_1 . In particular, the following is an illustration of a bidding sequence in a_1 :



Each circle stands for a bid, and the two red circles represent bids $X_{(L)}$ and $X_{(K)}$. At the first red circle, a bidder's auction choice is the same regardless of whether her bid is $X_{(L)}$ or $X_{(K)}$. Suppose that the corresponding auction is auction i , given the definition of a_1 . After the switch, at the first red circle, $X_{(K)}$ will be placed in auction i . Note that in a_1 , the highest losing bid in that auction at the first red circle is no greater than $X_{(l)}$; thus, when $X_{(K)}$ is placed, it does not get revealed as $X_{(K)} > X_{(l)}$. In other words, the subsequence of bidding until the second red circle will not be affected after the switch. Similarly, at the second red circle, the bidder with bid $X_{(L)}$ will choose to bid in auction j (the same decision at this circle in the sequence of a_1). In addition, bid $X_{(L)}$ does not get revealed as the highest losing bid in auction j at the second red circle since it is no greater than $X_{(k)}$, which, in turn, is lower than $X_{(l)}$. In summary, the bidding

sequence in a_1 is identical to the new sequence in a_2 except that we switched the bids corresponding to the two red circles. Moreover, these two sequences lead to the same vector of standing bids.

Given a_2 in Case 2, we have:

$$\begin{aligned}\Delta_2(a_2) &\triangleq G_i^0(w|s_i, s_j, a_2) - G_j^0(w|s_i, s_j, a_2) \\ &= P(X_{(K)} \leq w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \neq i, j, m' \neq l, k) \\ &\quad - P(X_{(L)} \leq w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \neq i, j, m' \neq l, k) \\ &= -\Delta_1(a_1).\end{aligned}$$

In addition, $P(a_1|\mathbf{s}) = P(a_2|\mathbf{s})$, as the numbers of assignments in Cases 1 and 2 are the same, and their elements can be matched one to one using the above argument.

Given an assignment in Case 3, we have:

$$s_i = X_{(l)}, s_j = X_{(L)}, W_i = X_{(k)}, W_j = X_{(K)} \text{ and } s_m = X_{(m')}, W_m = X_{(m'')},$$

where $m \neq i, j, m' \neq l, L, m'' \neq k, K$, and $m' > m''$. Then, we have:

$$\Delta_3(a_3) \triangleq P(X_{(k)} \leq w | s_i = X_{(l)}, s_j = X_{(L)}, a_3) - P(X_{(K)} \leq w | s_i = X_{(l)}, s_j = X_{(L)}, a_3).$$

Note that $\Delta_3(a_3) \geq 0$, we have

$$\{X_{(K)} \leq w, s_i = X_{(l)}, s_j = X_{(L)}, a_3\} \subseteq \{X_{(k)} \leq w, s_i = X_{(l)}, s_j = X_{(L)}, a_3\}.$$

Therefore, by taking the expectation of $\Delta_1(a_1)$, $\Delta_2(a_2)$ and $\Delta_3(a_3)$ over a_1 in Case 1, a_2 in Case 2, and a_3 in Case 3, we have:

$$G_i^0(w|\mathbf{s}) - G_j^0(w|\mathbf{s}) = \sum_{a_1 \in \text{Case 1}} P(a_1|\mathbf{s})\Delta_1(a_1) + \sum_{a_2 \in \text{Case 2}} P(a_2|\mathbf{s})\Delta_2(a_2) + \sum_{a_3 \in \text{Case 3}} P(a_3|\mathbf{s})\Delta_3(a_3) \geq 0.$$

■

Lastly, we consider the general case, which includes MU bidders. In particular, we establish the following lemma.

LEMMA 6. *Suppose that there are n identical auctions. Furthermore, $N+1$ UD bidders and N' MU bidders place their bids in an exogenous sequence. If the last one is a UD bidder, it is optimal for her to bid in the auction with the lowest standing bid, assuming that the rest of the UD bidders also place bids in the auctions with the lowest standing bids (at the time they choose which auction to participate in).*

Proof. It is sufficient to show that, after all MU bidders and the first N UD bidders place their bids, the conditional CDFs of the winning bids given the standing bid vector \mathbf{s} follow the relationship specified in Expression (26).

Since a MU bidder participates in all auctions (if she chooses to bid at all), there is at most one MU winner at any given time. Thus, there are two types of outcomes: either all winners are UD bidders (Case 1) or there is a single MU winner (Case 2). In what follows, we show that Expression (26) holds in both cases.

In Case 1, where the highest $2n$ bids are placed by UD bidders, the claim follows directly from Lemma 5.

In Case 2, where there is at least one auction being won by a MU bidder, we denote the winning MU bid by M . Given the vector of standing bids \mathbf{s} after the submissions of all bids, for any two auctions i and j with $i < j$, we aim to show that:

$$G_i^0(w|\mathbf{s}, \text{Case 2}) \geq G_j^0(w|\mathbf{s}, \text{Case 2}).$$

We then use W_i and W_j to denote the winning bids in auctions i and j , and consider the following three cases:

- (i) $W_i = M, W_j = M$, where the MU bidder is winning both auctions.
- (ii) $W_i = M, W_j = X_{(l)}$, where the MU bidder is winning auction i but not auction j .
- (iii) $W_i = X_{(l)}, W_j = X_{(l')}$, where the MU bidder is winning neither auction.

In sub-case (i), we have $G_i^0(w|\mathbf{s}, \text{Case 2}(i)) = G_j^0(w|\mathbf{s}, \text{Case 2}(i))$. In sub-case (ii), we have $X_{(l)} \geq M$; otherwise, the MU bidder is also winning auction j . Thus, $G_i^0(w|\mathbf{s}, \text{Case 2}(ii)) \geq G_j^0(w|\mathbf{s}, \text{Case 2}(ii))$, because $\{X_{(l)} \leq w, \mathbf{s}, \text{Case 2}(ii)\} \subseteq \{M \leq w, \mathbf{s}, \text{Case 2}(ii)\}$. In sub-case (iii), we can again apply the argument in the proof of Lemma 5 to show that $G_i^0(w|\mathbf{s}, \text{Case 2}(iii)) \geq G_j^0(w|\mathbf{s}, \text{Case 2}(iii))$. In summary, we have shown that with the presence of MU bidders, Expression (26) holds. ■

Proof of Proposition 3. The optimal bid, given that the bidder has already chosen an auction, can be directly derived using the first-order conditions for the payoff maximization problem. For the auction selection decision, we use induction to show that choosing to bid in the auction with the lowest standing bid is optimal, given that all other UD bidders follow the same strategy.

Step 1. Consider the bidder that places the last bid on a given day (Bidder A). Lemma 6 directly implies that it is optimal for Bidder A to choose to bid in the auction with the lowest standing bid.

Step 2 (Induction step). We then assume that Auction 1, which has the lowest standing bid, is the optimal auction choice for a bidder given \mathbf{s} and N future competing bidders. Specifically,

$$v_1^N(b, x; \mathbf{s}) \geq v_2^N(b, x; \mathbf{s}) \geq \dots \geq v_n^N(b, x; \mathbf{s}),$$

where $v_i^N(b, x; \mathbf{s})$ denotes the expected payoff for the bidder if she places a bid in auction i . The conditional expected payoff $v_i^N(b, x; \mathbf{s})$ can be expressed as

$$v_i^N(b, x; \mathbf{s}) = \int_{s_i}^b \left(\int_{w_i}^b (x - p_i) \hat{g}_i^N(p_i | w_i, \mathbf{s}_{-i}) + \alpha^{UD} (1 - \hat{G}_i^N(b | w_i, \mathbf{s}_{-i}) v_f(x)) \right) g_i^0(w_i | \mathbf{s}) dw_i.$$

In auction i , $\hat{g}_i^N(p_i | \mathbf{s})$ and $\hat{G}_i^N(p_i | \mathbf{s})$ respectively stand for the PDF and CDF of the highest rival bid against the *incumbent* winner given standing bids \mathbf{s} and N incoming bidders.

The optimal bid of the next UD bidder (Bidder B) is denoted by²⁵ b_1 . Without loss of generality, we show that for Bidder A, inequality $v_1^{N+1}(b; \mathbf{s}) \geq v_2^{N+1}(b; \mathbf{s})$ holds. We construct four information sets and establish that the inequality holds in each set.

(i) $b_1 \leq s_1$ (i.e., Bidder B does not submit a competing bid). In this case, it follows from the induction step that it is optimal for Bidder A to place a bid in Auction 1.

²⁵ For brevity, we omit the discussion of the case where the next bidder is a MU bidder, since the claim follows in a similar fashion when we classify the information set.

(ii) $s_1 < b_1 \leq \min(W_1, s_2)$ (i.e., whether Bidder B submits a bid depends on which auction Bidder A places her bid in). In this case, Bidder B will not bid in any auction if Bidder A bids in Auction 1 because bid b_1 is lower than both the standing bid in Auction 2 and the new standing bid in Auction 1 (i.e., the previously winning bid W_1). If Bidder A chooses to bid in auction 1, she will have only N future competitors and, as a result, the state would be updated from (s_1, s_2, \dots, s_n) with $N + 1$ incoming bids to (b_1, s_2, \dots, s_n) with N incoming bids. Therefore, her conditional expected payoff is

$$\begin{aligned} v_1^{N+1}(b, x; \mathbf{s}) &= \int_{b_1}^b \left(\int_{w_1}^b (x - p_1) \hat{g}_1^N(p_1 | w_1, \mathbf{s}_{-1}) dp_1 + \alpha^{UD} (1 - \hat{G}_1^N(b | w_1, \mathbf{s}_{-1})) v_f(x, G, \Psi) \right) g_1^0(w_1 | b_1, \mathbf{s}_{-1}) dw_1 \\ &= v_1^N(b, x; b_1, \mathbf{s}_{-1}). \end{aligned}$$

On the other hand, if Bidder A places a bid in Auction 2, Bidder B will bid in Auction 1. In this case, the standing bid in Auction 1 increases to b_1 . Her conditional expected payoff from bidding in Auction 2 is

$$\begin{aligned} v_2^{N+1}(b, x; \mathbf{s}) &= \int_{s_2}^b \left(\int_{w_2}^b (x - p_2) \hat{g}_2^N(p_2 | b_1, w_2, \mathbf{s}_{-1, -2}) dp_2 \right. \\ &\quad \left. + \alpha^{UD} (1 - \hat{G}_2^N(b | b_1, w_2, \mathbf{s}_{-1, -2})) v_f(x, G, \Psi) \right) g_2^0(w_2 | b_1, \mathbf{s}_{-1}) dw_2 \\ &= v_2^N(b, x; b_1, \mathbf{s}_{-1}). \end{aligned}$$

Note that $b_1 \leq s_2 \leq \dots \leq s_n$ and, by induction, we conclude that

$$v_1^N(b, x; b_1, \mathbf{s}_{-1}) \geq v_2^N(b, x; b_1, \mathbf{s}_{-1}).$$

(iii) $W_1 < b_1, W_1 < s_2$ (i.e., Bidder B will submit a bid regardless of where Bidder A places her bid). In this case, no matter which auction Bidder A chooses, Bidder B always places a bid in Auction 1. If Bidder A enters Auction 1, Bidder B will also bid in Auction 1 and outbid the new standing bid W_1 . Bidder B's bid b_1 therefore serves as the new winning bid in Auction 1. Bidder A's conditional expected payoff can thus be written as

$$\begin{aligned} v_1^{N+1}(b, x; \mathbf{s}) &= \int_{W_1}^b \left(\int_{b_1}^b (x - p_1) \hat{g}_1^N(p_1 | b_1, \mathbf{s}_{-1}) dp_1 + \alpha^{UD} (1 - \hat{G}_1^N(b | b_1, \mathbf{s}_{-1})) v_f(x, G, \Psi) \right) g_1^0(b_1 | W_1, \mathbf{s}_{-1}) db_1 \\ &= v_1^N(b, x; W_1, \mathbf{s}_{-1}). \end{aligned}$$

Similarly, if Bidder A places a bid in Auction 2, her conditional expected payoff can be written as

$$\begin{aligned} v_2^{N+1}(b, x; \mathbf{s}) &= \int_{s_2}^b \left(\int_{w_2}^b (x - p_2) \hat{g}_2^N(p_2 | W_1, w_2, \mathbf{s}_{-1, -2}) dp_2 \right. \\ &\quad \left. + \alpha^{UD} (1 - \hat{G}_2^N(b | W_1, w_2, \mathbf{s}_{-1, -2})) v_f(x, G, \Psi) \right) g_2^0(w_2 | W_1, \mathbf{s}_{-1}) dw_2 \\ &= v_2^N(b, x; W_1, \mathbf{s}_{-1}). \end{aligned}$$

Note that $W_1 \leq s_2 \leq \dots \leq s_n$ and, by induction, we conclude that

$$v_1^N(b, x; W_1, \mathbf{s}_{-1}) \geq v_2^N(b, x; W_1, \mathbf{s}_{-1}).$$

(iv) $s_2 < b_1, s_2 < W_1$ (i.e., Bidder B will also submit a bid regardless of Bidder A's decision). In this case, Bidder B will submit a bid in Auction 2 if Bidder A enters Auction 1, whereas Bidder B will place a bid in Auction 1 if Bidder A places a bid in Auction 2. When Bidder A places a bid in Auction 1 knowing

that $W_1 > s_2$, she will update her belief about W_1 accordingly (i.e., the conditional PDF of W_1 becomes $g_1^0(w_1|s_2, \mathbf{s}_{-1})$). Her conditional expected payoff then becomes

$$\begin{aligned} v_1^{N+1}(b, x; \mathbf{s}) &= \int_{s_2}^b \left(\int_{w_1}^b (x - p_1) \hat{g}_1^N(p_1|w_1, \min(b_1, W_2), \mathbf{s}_{-1, -2}) dp_1 \right. \\ &\quad \left. + \alpha^{UD} (1 - \hat{G}_1^N(b|w_1, \min(b_1, W_2), \mathbf{s}_{-1, -2})) v_f(x, G, \Psi) \right) g_1^0(w_1|s_2, \min(b_1, W_2), \mathbf{s}_{-1, -2}) dw_1 \\ &= v_1^N(b, x; s_2, \min(b_1, W_2), \mathbf{s}_{-1, -2}), \end{aligned}$$

where W_2 is the current winning bid in Auction 2. On the other hand, when Bidder A submits a bid in Auction 2, she expects Bidder B to place a bid in Auction 1. In addition, she only knows that the standing bid is the smaller one between b_1 and W_1 (i.e., $\min(b_1, W_1)$). Therefore, her conditional expected payoff of submitting a bid in Auction 2 is

$$\begin{aligned} v_2^{N+1}(b, x; \mathbf{s}) &= \int_{s_2}^b \left(\int_{w_2}^b (x - p_2) \hat{g}_2^N(p_2|\min(b_1, W_1), w_2, \mathbf{s}_{-1, -2}) dp_2 \right. \\ &\quad \left. + \alpha^{UD} (1 - \hat{G}_2^N(b|\min(b_1, W_1), w_2, \mathbf{s}_{-1, -2})) v_f(x, G, \Psi) \right) g_2^0(w_2|\min(b_1, W_1), s_2, \mathbf{s}_{-1, -2}) dw_2 \\ &= v_2^N(b, x; \min(b_1, W_1), s_2, \mathbf{s}_{-1, -2}). \end{aligned}$$

Note that both W_1 and W_2 have the same conditional PDF (i.e., $g_1^0(w_1|s_2, \mathbf{s}) = g_2^0(w_2|s_2, \mathbf{s})$) as Auctions 1 and 2 are indistinguishable given Bidder A's information set. Furthermore, expressions $v_1^N(b, x; s_2, \min(b_1, W_2), \mathbf{s}_{-1, -2})$ and $v_2^N(b, x; \min(b_1, W_1), s_2, \mathbf{s}_{-1, -2})$ are symmetric in terms of Auctions 1 and 2. Thus,

$$E[v_1^N(b, x; s_2, \min(b_1, W_2), \mathbf{s}_{-1, -2})] = E[v_2^N(b, x; \min(b_1, W_1), s_2, \mathbf{s}_{-1, -2})]$$

and for Bidder A, bidding in Auction 1 is a weakly dominant strategy.

In summary, given N incoming competing bids, it is a weakly dominant strategy for a UD bidder to bid in the auction with the lowest standing bid. This strategy depends on neither the bidders' valuations nor the market state. In the case of an unknown number of incoming bidders, one can derive the same conclusion by taking an expectation over N . ■

Table 24 Hypothesis Testing Results of Product Preference

(Model 1, Model 2)	Population Level Test		Individual Level Test
	\hat{R}	p -value	% of Bidders <i>without</i> a Strong Preference
(iPhone 4, iPhone 4s)	0.055	0.157	94.7%
(iPhone 4, iPhone 5)	0.000	0.987	92.6%
(iPhone 4s, iPhone 5)	-0.039	0.122	96.8%

Appendix C: Structural Model

C.1. Substitutability of iPhone 4, iPhone 4s, and iPhone 5 Models

In this section, we present descriptive evidence supporting the fact that bidders view auctions for iPhone 4, iPhone 4s, and iPhone 5 as substitutes. In particular, we test whether bidders significantly prefer one model to the others at both the population and individual bidder levels.

At the population level, given a model pair, i.e., (iPhone) Model 1 and (iPhone) Model 2, we use b_{1t} and b_{2t} to denote the number of bids placed in auctions for Model 1 and Model 2 phones that end on Day t , respectively. Also, we use n_{1t} and n_{2t} to denote the total number of auctions for Model 1 and Model 2 phones that end on Day t , respectively. Then, we propose to test the following null hypothesis:

HYPOTHESIS 1. *On a given day, the expected ratio of bids placed in Model 1 auctions to bids placed in Model 2 auctions is equal to the ratio of the total number of Model 1 auctions to the total number of Model 2 auctions.*

Specifically, we focus on the following statistic:

$$\hat{R} = \frac{1}{T} \sum_{t=1}^T \left(\frac{b_{1t}}{b_{1t} + b_{2t}} / \frac{n_{1t}}{n_{1t} + n_{2t}} - 1 \right),$$

where T is the number of days when at least one Model 1 auction ends. On a given day, a consistent mismatch between $b_{1t}/(b_{1t} + b_{2t})$ and $n_{1t}/(n_{1t} + n_{2t})$ indicates that bidders prefer one model to the other. For an iPhone model, the statistic \hat{R} manages to capture not only the daily mismatch between its shares of bids and auctions, but also how the trend of bid shares matches the trend of auction shares over time. If \hat{R} differs significantly from zero, we reject Hypothesis 1 and conclude that one model is preferred to the other at the population level; thus, the models cannot be considered as substitutes.

Our findings are presented in the first two columns of Table 24. For all three model pairs, we cannot reject Hypothesis 1. This implies that on a daily basis, the number of bids for each model is proportional to the number of auctions for that model. Therefore, each model is not more or less popular than other models at the population level. Thus, we can treat the three models as substitutes.

Furthermore, we test whether the models are substitutable at the individual bidder level. We find that more than 90% of the bidders do not show a strong preference for a specific model. In particular, given Bidder ℓ , we construct the following null hypothesis to test:

HYPOTHESIS 2. *On a given day, when Bidder ℓ participates in at least one Model 1 or Model 2 auction, the expected ratio of Model 1 auctions to Model 2 auctions that Bidder ℓ participates in is equal to the ratio of the total number of Model 1 auctions to the total number of Model 2 auctions.*

In other words, we focus on the following statistic for Bidder ℓ :

$$\hat{R}_\ell = \frac{1}{|\mathcal{T}_\ell|} \sum_{t \in \mathcal{T}_\ell} \left(\frac{p_{1t\ell}}{p_{1t\ell} + p_{2t\ell}} / \frac{n_{1t}}{n_{1t} + n_{2t}} - 1 \right),$$

where \mathcal{T}_ℓ is the set of days when Bidder ℓ participated in at least one Model 1 or Model 2 auction, and $p_{1t\ell}$ (resp., $p_{2t\ell}$) are the numbers of Model 1 (resp., Model 2) auctions which Bidder ℓ participated in on Day t . If \hat{R}_ℓ differs significantly from zero, we reject Hypothesis 2 and conclude that the bidder has a strong preference for either Model 1 or Model 2. Specifically, we say that the bidder does not exhibit a strong preference for either of the two models if the p -value of the test result is less than 0.01. We report our findings in the last column of Table 24. For each model pair, there are more than 90% of bidders without a strong preference. Therefore, the findings provide convincing evidence that supports the idea that auctions of the three models can be viewed as substitutes.

C.2. Kernel Density Estimators

In this section, we specify the kernel density estimators for $\Psi(\omega_{\ell t})$ and $g_j(p_j|\omega_{e \ll t})$. Assuming the within-day timing of a participating bidder's bids to be exogenous, the steady-state vector $\mathbf{s}_{\ell t}$ of standing bids encountered by bidders per platform visit can be viewed as identically and independently distributed. We use $\mathbf{S}_n = \{\mathbf{s}_{\ell t, n}\}$ and $\mathbf{y}_n = \{y_{\ell t, n}\}$ to denote the observations of the vectors of standing bids²⁶ and highest rival bids (final prices), respectively, when the market thickness is n auctions. The corresponding sample size is denoted by N_n . The kernel density estimator of the standing bids conditional on market thickness n can be written as $\hat{\psi}(\mathbf{s}_n|n) = \frac{1}{N_n} \sum_{\ell t} K_{h_s}(\mathbf{s}_n - \mathbf{s}_{\ell t, n})$, where $K_{h_s}(\mathbf{s})$ is the multivariate Gaussian kernel density with identical bandwidth h_s :

$$K_{h_s}(\mathbf{s}) = \left(\frac{1}{\sqrt{2\pi}h_s} \right)^n \prod_{j=1}^n \exp\left(-\frac{s_j^2}{2h_s^2} \right).$$

The selection of the optimal bandwidth h_s^* is conducted by minimizing the cross-validation (CV) estimator of the risk function. Finally, the kernel density estimator of the highest rival bid in auction j (i.e., $y^{(j)}$), given market state $\omega = (n, \mathbf{s})$, can be written as

$$\hat{g}_j(y^{(j)}|\omega) = \frac{\frac{1}{N_n} \sum_{\ell t} K_{h_{y,1}}(y^{(j)} - y_{\ell t, n}^{(j)}) K_{H_y}(\mathbf{s} - \mathbf{s}_{\ell t, n})}{\frac{1}{N_n} \sum_{\ell t} K_{H_y}(\mathbf{s} - \mathbf{s}_{\ell t, n})},$$

where $y_{\ell t, n}^{(j)}$ is the j^{th} element of $y_{\ell t, n}$, and $h_{y,1}$ is the bandwidth associated with the highest rival bid. For tractability, we assume that the bandwidth matrix of standing bids takes the form $H_y = h_{y,2}I_n$. The optimal bandwidths $h_{y,1}^*$ and $h_{y,2}^*$ are determined by minimizing the CV estimator of the risk function, $CV(h_{y,1}, h_{y,2})$, of the conditional density estimation.²⁷

²⁶ If a bidder places multiple bids within a day, we assume that the market state is observed upon her last bid of the day, which reflects the most recent market state for her decision.

²⁷ We provide additional details on the bandwidth selection procedure for both $\hat{\psi}_n(\mathbf{s}_n)$ and $\hat{g}_j(y|\omega)$ in Appendix C.3.

C.3. Bandwidth Selection in Kernel Density Estimation

In this subsection, we describe the bandwidth selection process we use when we non-parametrically estimate the distributions of the steady state of the market and the highest rival bids, respectively. The risk, or mean integrated squared error $R(f, \hat{f})$, which we define below, is a metric for the distance between a density estimator \hat{f} and the true density f :

$$R(f, \hat{f}) = E_X \left(\int (f(x) - \hat{f}(x|X))^2 dx \right),$$

where $X \sim f$. As f is unknown, the risk cannot be evaluated. Therefore, we consider the cross-validation (CV) estimator of risk, which is defined as follows:

$$CV(\mathbf{h}) = \int (\hat{f}(x))^2 dx - \frac{2}{n} \sum_{i=1}^n \hat{f}_{(-i)}(X_i),$$

where $\hat{f}_{(-i)}(X_i)$ is the kernel density estimator obtained by removing the i^{th} observation, X_i . In our case, \hat{f} is the kernel density estimator with bandwidth \mathbf{h} , which is the parameter to be optimized (i.e., the bandwidth \mathbf{h} is chosen so that the CV estimator of risk is minimized).

For the kernel density estimator of the standing bids, the CV estimator of risk can be further specified as:

$$CV(h_s) = \frac{1}{N^2 |H_s|} \sum_{\ell t} \sum_{\ell' t'} \bar{K}(H_s^{-1}(\mathbf{s}_{\ell t} - \mathbf{s}_{\ell' t'})) - \frac{2}{N(N-1) |H_s|} \sum_{\ell t} \sum_{\ell' t' \neq \ell t} K(H_s^{-1}(\mathbf{s}_{\ell t} - \mathbf{s}_{\ell' t'})),$$

where N denotes the total number of observations (here, we use ℓt and $\ell' t'$ as observation indices). Furthermore, the bandwidth matrix is given by $H_s = h_s I_n$ and K, \bar{K} denote the kernel and the corresponding convolution functions, respectively.

Note that when estimating the conditional distribution of highest rival bids, directly optimizing the cross-validation of the risk function may lead to very small bandwidths (due to the fact that multiple vectors of standing bids can be associated with the same vector of highest rival bids). To resolve this issue, we employ bootstrap sampling from the original dataset, drawing only one observation $(y_{\ell t, n}, \mathbf{s}_{\ell t, n})$ per auction (to ensure that we do not duplicate the vectors of highest rival bids within a bootstrap sample). Given a bootstrap subsample b , the subsample CV estimator of risk can be specified as

$$CV^{(b)}(h_{y,1}, h_{y,2}) = \frac{1}{N} \sum_{\ell t} \frac{\sum_{\ell' t' \neq \ell t} \sum_{\ell'' t'' \neq \ell t} K_{H_y}(\mathbf{s}_{\ell t}^{(b)} - \mathbf{s}_{\ell' t'}^{(b)}) K_{H_y}(\mathbf{s}_{\ell t}^{(b)} - \mathbf{s}_{\ell'' t''}^{(b)}) K_{\sqrt{2}h_{y,1}}(y_{\ell'' t''}^{(b)} - y_{\ell' t'}^{(b)})}{\left(\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}_{\ell t}^{(b)} - \mathbf{s}_{\ell' t'}^{(b)}) \right)^2} - \frac{2}{N} \sum_{\ell t} \frac{\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}_{\ell t}^{(b)} - \mathbf{s}_{\ell' t'}^{(b)}) K_{h_{y,1}}(y_{\ell t}^{(b)} - y_{\ell' t'}^{(b)})}{\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}_{\ell t}^{(b)} - \mathbf{s}_{\ell' t'}^{(b)})}.$$

Lastly, we consider the average of the subsample CV estimators of risk (we draw 50 bootstrap samples) as the objective function to optimize when selecting the bandwidth for the kernel density estimation.

C.4. Simulated Maximum Likelihood Estimation

In this subsection, we provide expressions for $P^V(x_\ell | \theta)$, $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t} | x_\ell, \theta)$, and $L_{\ell t}^{NB}(\mathbf{S}_{\ell t} | x_\ell, \theta)$, where $\theta = [\alpha, \mu, \nu, \lambda, \gamma]$, and describe how we derive the simulated likelihood function $\hat{L}_\ell(\mathbf{X}_\ell | \theta)$, using importance sampling. Before presenting the technical details, we preview how each primitive is estimated by exploiting the variations in the data.

The data contain sufficient variations for us to identify all the model primitives. First, in the data, a bidder (whose endowed valuation x_ℓ is given) may choose a subset of auctions on a given day and place bids in the selected auctions. The variations in her auction selections and in her bids can be mainly explained by the idiosyncratic valuation distribution (characterized by ν). In addition to the bid variations of a given bidder, we observe sizable bid variations across bidders (e.g., some bidders consistently submit much higher bids than others in the same auction). The cross-bidder bid variations are characterized by the bidders' endowed valuations (i.e., x_ℓ), whose distribution is characterized by λ^{TY} and γ^{TY} . Second, we observe that each bidder's platform-visit sequence follows a stochastic process with a certain frequency (e.g., some bidders tend to participate on the platform more often than others). In our model, a bidder's platform-visit pattern is determined by the interplay between her expected payoff, which depends on the endowed valuation (i.e., λ^{TY} and γ^{TY}) as well as the idiosyncratic valuation (i.e., ν^{TY}), and the distribution of participation cost (characterized by μ^{TY}). Note that λ^{TY} , γ^{TY} , and ν^{TY} are pinned down by the above bid variations; thus, the variation in bidders' platform visits can be employed to identify μ^{TY} . Lastly, in the data, each bidder exits the platform at some point, while her stay in the bidder pool may vary from a few days to a few months. The variations in the duration of bidders' stays in the bidder pool are then used to estimate the retention rate α . Therefore, our data have enough variations to recover all the modeling primitives.

We use $P^V(x_\ell|\theta)$ to denote the probability that Bidder ℓ visits the platform on Day t :

$$P^V(x_\ell|\theta) = P(c_{\ell t} \leq r(x_\ell; \theta, \hat{G}, \hat{\Psi})) = 1 - \exp(-\mu r(x_\ell; \theta, \hat{G}, \hat{\Psi})),$$

where recall that $r(x_\ell; \theta, G, \Psi)$ is the unconditional payoff per platform visit (Section 4). Given x_ℓ , θ^{MU} , \hat{G} , and $\hat{\Psi}$, a MU bidder's payoff per platform visit r^{MU} can be calculated using Expressions (21) and (22). For a UD bidder, given x_ℓ , θ^{UD} , \hat{G} , and $\hat{\Psi}$, we solve for r^{UD} and the continuation valuation v_f simultaneously from

$$\begin{aligned} & r^{UD}(x_\ell; \theta^{UD}, \hat{G}, \hat{\Psi}) \\ &= \int_{\omega} \int_{\zeta_\ell} \sum_j^{n(\omega)} \mathbb{1}\{\sigma_{SLT}^{UD}(x_\ell; \zeta_\ell, \omega, \theta^{UD}) = j\} \int_{s_j}^{x_\ell - \alpha v_f(x_\ell; \theta^{UD}, \hat{G}, \hat{\Psi}) + \zeta_{\ell j}} \hat{G}(p_j|\omega) dp_j d\mathbf{F}_\zeta(\zeta_\ell; \nu^{UD}) d\hat{\Psi}(\omega), \end{aligned} \quad (27)$$

$$\text{and } v_f(x_\ell; \theta^{UD}, \hat{G}, \hat{\Psi}) = \frac{1}{1 - \alpha^{UD}} \left(r^{UD}(x_\ell; \theta^{UD}, \hat{G}, \hat{\Psi}) + \frac{1}{\mu^{UD}} \exp(-\mu^{UD} r^{UD}(x_\ell; \theta^{UD}, \hat{G}, \hat{\Psi})) - \frac{1}{\mu^{UD}} \right), \quad (28)$$

where Expression (27) is derived from Expressions (5) and (23), and Expression (28) is derived from Expression (24). For computational tractability, we employ Monte Carlo integration to approximate all the integrations over ζ_ℓ . Any integration of \hat{G} or $\hat{\Psi}$ can be done quickly due to the selection of the Gaussian kernel.

Conditional on Bidder ℓ visiting the platform on Day t , we let $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \theta)$ denote the likelihood of her placing bid(s) $\mathbf{b}_{\ell t}$, and we let $L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \theta)$ denote the likelihood of her not placing any bid.

First, assume that ℓ is a MU bidder (who is interested in winning $K = 14$ auctions, which are *all* open auctions, on a given day).²⁸ Then, when she visits the platform on Day t , she will not bid in auction j if and only if her bid is lower than the current standing bid $s_{\ell j}$ (i.e., $b_{\ell j} = x_\ell + \zeta_{\ell j}^{MU} < s_{\ell j}$). Therefore,

²⁸ Recall that K is the observed maximum auction participation on a given day.

$$L_{\ell t}^{NB, MU}(\mathbf{S}_{\ell t}|x_{\ell}, \theta^{MU}) = \prod_{j' \in \mathcal{A}^t} \Phi(s_{\ell j'}|x_{\ell}, \nu^{MU}),$$

where $\Phi(\cdot|x_{\ell}, \nu^{MU})$ denotes the CDF of the normal distribution with mean x_{ℓ} and standard deviation ν^{MU} , and \mathcal{A}^t denotes the set of all auctions on Day t . On the other hand, if she places bids $\mathbf{b}_{\ell t} = \{b_{\ell j}\}$, and \mathcal{J}_{ℓ}^t denotes the set of auctions that Bidder ℓ chooses to participate in on Day t , the likelihood of placing these bids is given by $\prod_{j \in \mathcal{J}_{\ell}^t} \phi(b_{\ell j}|x_{\ell}, \nu^{MU})$. Therefore,

$$L_{\ell t}^{B, MU}(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_{\ell}, \theta^{MU}) = \prod_{j \in \mathcal{J}_{\ell}^t} \phi(b_{\ell j}|x_{\ell}, \nu^{MU}) \prod_{j' \in \mathcal{A}^t \setminus \mathcal{J}_{\ell}^t} \Phi(s_{\ell j'}|x_{\ell}, \nu^{MU}).$$

Next, assume that Bidder ℓ has unit demand. Then, given that she visits the platform on Day t , she will not place a bid if and only if all auctions on the day have standing bids that are higher than her potential bids (i.e., $b_{\ell j} = x_{\ell} - \alpha^{UD} v_f(x_{\ell}; \theta^{UD}, \hat{G}, \hat{\Psi}) + \zeta_{\ell j}^{UD} < s_{\ell j}, \forall j \in \mathcal{A}^t$). Therefore,

$$L_{\ell t}^{NB, UD}(\mathbf{S}_{\ell t}|x_{\ell}, \theta^{UD}) = \prod_{j' \in \mathcal{A}^t} \Phi(s_{\ell j'}|x_{\ell} - \alpha^{UD} v_f(x_{\ell}; \theta^{UD}, \hat{G}, \hat{\Psi}), \nu^{UD}).$$

On the other hand, if she places a bid $\mathbf{b}_{\ell t} = b_{\ell j}$ in auction j , this implies that participating in auction j yields the highest payoff among all open auctions on the same day (as was shown in Proposition 2). In addition, if she chooses auction j (i.e., $\sigma_{SLT}^{UD}(x_{\ell}; \zeta_{\ell}, \mathbf{S}_{\ell t}, \theta^{UD}) = j$), the likelihood of placing bid $b_{\ell j} = x_{\ell} - \alpha^{UD} v_f(x_{\ell}; \theta^{UD}, \hat{G}, \hat{\Psi}) + \zeta_{\ell j}^{UD}$ is equal to $\phi(b_{\ell j}|x_{\ell} - \alpha^{UD} v_f(x_{\ell}; \theta^{UD}, \hat{G}, \hat{\Psi}), \nu^{UD})$. Thus,

$$L_{\ell t}^{B, UD}(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_{\ell}, \theta^{UD}) = \phi(b_{\ell j}|x_{\ell} - \alpha^{UD} v_f(x_{\ell}; \theta^{UD}, \hat{G}, \hat{\Psi}), \nu^{UD}) \int_{\zeta_{\ell}} \mathbb{1}\{\sigma_{SLT}^{UD}(x_{\ell}; \zeta_{\ell}, \mathbf{S}_{\ell t}, \theta^{UD}) = j\} d\mathbf{F}_{\zeta}(\zeta_{\ell}; \nu^{UD}).$$

To obtain the simulated likelihood, $\hat{L}_{\ell}(\mathbf{X}_{\ell}|\theta)$, we first rewrite the unconditional likelihood function as

$$\mathcal{L}_{\ell}(\mathbf{X}_{\ell}|\theta) = \int_{x_{\ell}} \mathcal{L}_{\ell}(\mathbf{X}_{\ell}|x_{\ell}, \theta) \frac{f(x_{\ell}|\lambda, \gamma)}{f_I(x_{\ell}|\lambda_I, \gamma_I)} f_I(x_{\ell}|\lambda_I, \gamma_I) dx_{\ell} = E_{x_I} \left(\mathcal{L}_{\ell}(\mathbf{X}_{\ell}|x_I, \theta) \frac{f(x_I|\lambda, \gamma)}{f_I(x_I|\lambda_I, \gamma_I)} \right), \quad (29)$$

where $f_I(x_I|\lambda_I, \gamma_I)$ is PDF of the candidate Weibull distribution with parameters λ_I and γ_I . Then, we draw M samples for the bidder's valuation $x_{m, I}, m = 1, 2, \dots, M$ from the candidate Weibull distribution. Finally, we (approximately) compute integral (29) using Monte Carlo integration and obtain

$$\hat{\mathcal{L}}_{\ell}(\mathbf{X}_{\ell}|\theta) = \frac{1}{M} \sum_{m=1}^M \mathcal{L}_{\ell}(X_{\ell}|x_{m, I}, \theta) \frac{f(x_{m, I}|\lambda, \gamma)}{f_I(x_{m, I}|\lambda_I, \gamma_I)}.$$

In our estimation process, we draw $M = 200$ samples and set λ_I, γ_I such that they satisfy the first and second moment conditions of the empirical bid distribution.

C.5. Model Validation

To validate the model, we compare the predicted distributions of (i) the number of bidders participating in an auction and (ii) the final price in the post-treatment period, with those observed in the dataset. Given that these distributions concern outcomes in the post-treatment period, they can be viewed as out-of-sample predictions (recall that the model was estimated using pre-treatment data). Our findings provide support for the validity of the model. In particular:

Table 25 Distributions of Number of Bidders per Auction and Final Price per Unit in the Post-treatment Period

	No. of Bidders per Auction		Final Price per Unit (\$)	
	Mean	Standard Deviation	Mean	Standard Deviation
Observed	5.24	2.33	168.75	15.49
Simulated	5.79 (0.68)	2.01 (0.35)	167.97 (5.24)	18.30 (1.99)

(a) The mean and the standard deviation of the simulated distribution of participating bidders (who placed at least one bid) per auction in the post-treatment period are shown in the first two columns of Table 25 (with standard errors in parentheses). As is evident from the table, our counterfactual simulations provide a relatively accurate prediction of the corresponding distribution. In particular, both the mean and the standard deviation of the simulated distribution are not significantly different from those observed in the data.

(b) Similarly, we simulate the mean and the standard deviation of the final price distribution (on a normalized scale) in the post-treatment period. The results are shown in the last two columns of Table 25 (with standard errors in parentheses). As one can deduce from the table, the simulations provide an accurate out-of-sample prediction of the average final price in the post-treatment period.

Appendix D: Supporting Material

D.1. Spillover Effects on Participation Rates

This subsection focuses on auction participation behaviors in Market B and the subset of cross-market bidders (i.e., bidders that are registered in both Markets A and B). The objective is to establish that the change in Market A's market thickness has a strong positive effect on the cross-market bidders' participation rates in Market B. To this end, we specify the following linear model for the bidders' participation rates in Market B:

$$APR_{\ell tw} = \xi_{\ell} + \eta_w + \beta_1 CMB_{\ell} + \beta_2 TT_t + \beta_3 CMB_{\ell} \cdot TT_t + \epsilon_{\ell tw}, \quad (30)$$

where the dependent variable $APR_{\ell tw}$ is defined as

$$APR_{\ell tw} = \frac{\text{No. of Mkt. B auctions Bidder } \ell \text{ bids on Day } t \text{ in Week } w}{\text{No. of Mkt. B auctions ending on Day } t \text{ in Week } w}.$$

In Equation (30), ξ_{ℓ} denotes the fixed effect associated with Bidder ℓ 's idiosyncratic participation behavior, and η_w denotes a weekly fixed effect introduced to control for changes over time.

We include fixed effect ξ_{ℓ} to control for Bidder ℓ 's idiosyncratic participation behaviors. We also include fixed effect η_w to control for changes over time (e.g., releases of new products on participation rates) in Week w . Furthermore, CMB_{ℓ} and TT_t are binary variables denoting whether i is a cross-market bidder and whether the day of the Week t is Tuesday or Thursday, respectively. Thus, coefficient β_1 captures the participation pattern of cross-market bidders in Market B, and β_2 captures the baseline difference in participation rates between [Tuesday, Thursday] and [Monday, Wednesday, Friday]. Finally, the quantity of interest, β_3 , is meant to capture the potential spillover effect on the cross-market bidders' participation rates on Tuesdays and Thursdays.

For the results that follow, we focus only on the treatment period and restrict attention to iPhone 4 (thus, we do not need to control for product characteristics). In this sample, cross-market bidders account for 49% of the total Market B bidders. First, we plot in Figure 23 the average participation rates corresponding to cross-market bidders and exclusively to Market B bidders on Monday/Wednesday/Friday and Thursday/Thursday, respectively. Although on Monday/Wednesday/Friday, cross-market and Market B bidders have similar participation rates, on Tuesday/Thursday, the participation rate of cross-market bidders is substantially higher than the participation rate of Market B bidders.

Second, we employ difference in differences to estimate coefficient β_3 . As we report in Table 26, we estimate that the relative increase in the cross-market bidders' participation rates that can be attributed to the change in Market A's market thickness is roughly 39% (2.7% in absolute terms).

The size of the spillover effect provides evidence supporting the presence of participation frictions associated with visiting a market on any given day and actively bidding and monitoring the auctions listed on the platform.

D.2. Regression Analysis in Support of Cannibalization among Auctions

In this section, we establish that high market thickness may have a negative impact on the final price of the platform's auctions. To this end, we conduct a regression analysis using data on Samsung Galaxy S3 auctions in Market A. We address the potential issue of demand endogeneity by using an additional dataset tracking the retail price of Galaxy S3 phones that was obtained from Amazon.

Table 26 Spillover Participation Effect of the Batch Policy

	<i>Dependent Variable</i>
	Auction Participation Rate
Spillover Participation Effect $\hat{\beta}_3$	0.027*** (0.005)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Demand Endogeneity. To address the potential issue arising from demand endogeneity, we use a dataset tracking the price of Galaxy S3 phones that was obtained from Amazon. The prices on Amazon, which is a retail market, are largely exogenous to the prices on the platform, which is a secondary market. This allows us to control for market trends in the demand for Galaxy S3 phones; thus, we can attribute price fluctuations we observe on the platform to changes in the induced market thickness.

Data and Variables. The data we use for this analysis includes (1) detailed Galaxy S3 auction data from Market A, and (2) weekly retail price data from Amazon. Our observation window is from March 15, 2013 to December 7, 2015.²⁹ In the Galaxy S3 auction dataset, there are 376 observations within the observation window. For each auction, the dataset contains the start and ending times of the auction, its final price, lot size, and carrier information. Market thickness on a given day is measured by two metrics: the number of auctions ending on that day and the number of units in those auctions. The dataset we obtained from Amazon tracks the retail price of a Galaxy S3 phone within our observation window.

In Figure 27, we plot the price trends of Galaxy S3 phones in Market A (red curve) and on Amazon (green curve) as well as the weekly supply in Market A (blue curve). The supply shock in the local market (corresponding to the dashed line in the figure) is due to the release of a new Samsung Galaxy model.

Empirical Evidence. Here, we establish that high market thickness may have a negative impact on the final price of an auction. To this end, first, we compute the correlation between market thickness and the auction's final price without controlling for demand endogeneity. As already mentioned, market thickness on a given day is measured in two ways: (1) by the number of auctions ending on that day and (2) by the number of units in those auctions. We find that the correlation coefficients corresponding to the market thickness metrics are equal to -0.30 and -0.31, respectively, thus establishing the negative impact of market thickness on price.

Next, we address demand endogeneity by explicitly controlling for the market value of Galaxy S3 phones using the Amazon data. In the regression analysis, we select the final price per unit as the dependent variable. We include fixed effects for both the product carrier and new product releases.

The estimation results are displayed in Table 28. As is evident in the results for Models 1 and 2 (without including the fixed effect corresponding to a new product release), high market thickness has a significantly negative impact on the final price. Particularly, we note that adding one Galaxy S3 auction to the platform's

²⁹ We collect Amazon's retail prices for Samsung Galaxy S3 from CamelCamelCamel.com, a tool tracking prices for products sold on Amazon.

Figure 27 Weekly Supply and Prices of Samsung Galaxy S3 (the vertical dashed line highlights the release of Galaxy S5 on 04-11-2014)

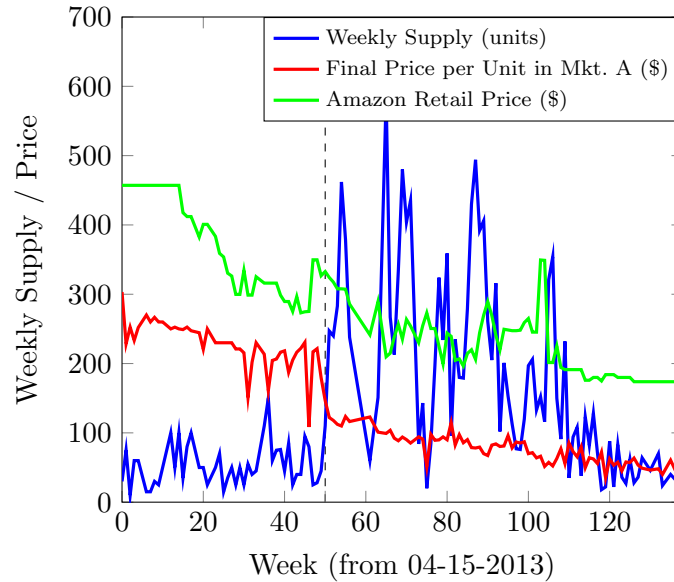


Table 28 Regression Results in Support of the Cannibalization Effect

	<i>Dependent Variable</i>			
	Final Price per Unit			
	Model 1	Model 2	Model 3	Model 4
No. of units per day	-0.11*** (0.02)		-0.043** (0.017)	
No. of auctions per day		-7.63*** (1.39)		-3.06*** (1.06)
Auction lot size		-0.15** (0.06)		-0.08* (0.048)
Amazon retail value	0.54*** (0.02)	0.53*** (0.02)	0.226*** (0.027)	0.225*** (0.027)
FEs of Rel. of S4, S5, S6	No	No	Yes	Yes
FEs of carriers	Yes	Yes	Yes	Yes
Constant	-54.42*** (7.39)	-43.42*** (8.02)	122.44*** (15.56)	128.20*** (15.61)

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

market thickness results in a decrease in the final price per unit by \$7.63. Similarly, adding 10 units of Galaxy S3 phones decreases the final price per unit by \$1.10.

Furthermore, we control for possible supply endogeneity by considering the exogenous shocks corresponding to new product releases (i.e., the releases of Samsung Galaxy S4, S5, and S6). When a new Samsung Galaxy model is released, we expect a jump in the number of Galaxy S3 phones that become available in the secondary market. As shown in Table 28 (Models 3 and 4), the cannibalization effect remains significant even after we account for the release of new products. In particular, our results indicate that adding an auction

Table 29 Difference-in-Differences Result of Effect on Bids

	<i>Dependent Variable</i>
	Log(Bid)
<i>IsPost</i>	−0.763*** (0.118)
<i>IsMktA</i>	1.206*** (0.070)
<i>IsMktA * IsPost</i>	−0.173 (0.146)
FEs of Products	Yes
Constant	3.359*** (0.765)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

to the platform’s market thickness results in a decrease in the final price per unit by \$3.06, whereas adding 10 Galaxy S3 phones results decreases their unit price by \$0.43.

D.3. Listing Policy Has Little Impact on Bids

In this section, we look into how bidders’ bidding strategies respond to the listing policy switch in Market A. We focus on a sample of bidders who placed a bid in either only Market A or only Market B in both the pre- and post-treatment periods (i.e., cross-market bidders are excluded). Therefore, in the post-treatment period, Market A bidders are directly subject to the policy switch, while Market B bidders are not. There are 14 Market A bidders, 21 Market B bidders, and 814 total bids observed in the sample. Then, we estimate the following difference-in-differences model to capture any effect on their bidding strategies:

$$\log(Bid_i) = \beta_0 + FEProd_i + \beta_1 IsMktA_i + \beta_2 IsPost_i + \beta_3 IsMktA_i * IsPost_i + \epsilon_i,$$

where for each observation i , Bid_i is the variable corresponding to the placed bid (normalized to a single unit), $FEProd_i$ captures fixed effects associated with product features, $IsMktA_i$ is the dummy variable of Market A auctions, and $IsPost_i$ is the dummy variable of the post-treatment period. Thus, the coefficient of interest is β_3 , which captures the effect of the policy switch on the size of the bids. As shown in Table 29, the policy switch does not seem to affect bidders’ bidding strategies, given that β_3 is not significantly different from zero. This finding rules out the alternative mechanism that could explain the increase in revenues observed after the listing policy switch, i.e., that the increase is the outcome of bidders changing their bidding strategies.

D.4. Algorithm for Estimating a Steady-State Equilibrium

In this subsection, we describe an iterative algorithm that converges to the steady-state equilibrium. Given the stochastic supply, a dynamic bidder pool, and a listing policy, a bidder’s decision whether to visit the platform, which auction(s) to participate in, and how to bid are all endogenously determined. Overall, the algorithm starts by initializing the beliefs for the highest rival bid G and the state distribution Ψ , based on which bidders decide whether to visit the platform and how to bid. Once the corresponding bidding data are

generated, they are used to re-estimate G and Ψ . The algorithm iterates between the data simulation step and the belief estimation step until they converge (Algorithm 1). The expected payoff $r(x; G, \Psi)$ is calculated numerically for MU and UD bidders separately. We employ kernel density estimation for the conditional distribution of the highest rival bid G and the distribution of standing bids Ψ , which are specified in Section 5.1.

Acknowledgments

The authors gratefully acknowledge the data science team of our corporate sponsor for their data and close collaboration. We thank Tunay Tunca and Dennis Zhang for their helpful comments and suggestions.

Algorithm 1 Estimating a steady-state equilibrium

-
- 1: Assign a convergence threshold ϵ .
 - 2: Assign a positive value to Δ , such that $\Delta > \epsilon$.
 - 3: Iteration $i = 0$.
 - 4: Initialization:
 - 5: The highest rival belief $G^{(0)}$.
 - 6: Market state distribution $\Psi^{(0)}$.
 - 7: $r^{TY}(x_\ell, G^{(0)}, \Psi^{(0)})$, $\forall \ell$ in the set of all bidders with type $TY \in \{MU, UD\}$.
 - 8: **while** $\Delta > \epsilon$ **do**
 - 9: Iteration $i \leftarrow i + 1$.
 - 10: Simulate bidding history $\Omega^{(i)}$ under beliefs $G^{(i-1)}$ and $\Psi^{(i-1)}$.
 - 11: **for** day $t = 1 \dots T$ **do**
 - 12: Update bidder pool on Day t :
 - 13: Existing bidder exits *w.p.* $1 - \alpha^{TY}$.
 - 14: Add new bidders $\sim P_{Arr}^{TY}$.
 - 15: Initialize market state ω :
 - 16: Add new auctions $\sim P_{Supply}$.
 - 17: Set standing bids.
 - 18: **for** Bidder ℓ in the pool with type $TY \in \{MU, UD\}$ **do**
 - 19: Participation cost $c_{\ell t} \sim Exp(\mu^{TY})$.
 - 20: **if** $\sigma_{PRT}^{TY}(x_\ell; c_{\ell t}) = \text{Visit}$ **then**
 - 21: Observe the market state ω and idiosyncratic terms ζ_ℓ .
 - 22: Select auction(s) according to $\sigma_{SLT}^{TY}(x_\ell; \zeta_\ell, \omega)$.
 - 23: Place bid(s) $\sigma_{BID}^{TY}(x_j; \zeta_\ell, \omega)$ in selected auction(s).
 - 24: Update standing bids in market state ω .
 - 25: **end if**
 - 26: **end for**
 - 27: **end for**
 - 28: Estimate $G^{(i)}$ and $\Psi^{(i)}$ non-parametrically from $\Omega^{(i)}$.
 - 29: Compute $r^{TY}(x_\ell, G^{(i)}, \Psi^{(i)})$, $\forall \ell$ in the set of all bidders with type $TY \in \{MU, UD\}$.
 - 30: $\Delta \leftarrow \|G^{(i)} - G^{(i-1)}\| + \|\Psi^{(i)} - \Psi^{(i-1)}\|$.
 - 31: **end while**
