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

Platforms can obtain sizable returns by operationally managing their market thickness, i.e., the availability of supply-side inventory. Using data from a natural experiment on a major B2B auction platform specializing in the $424 billion secondary market for liquidating retail merchandise, we find that thickening the platform's market by consolidating the ending times of auctions to certain weekdays substantially increases its revenue by roughly 6.5%, due primarily to the bidders' participation frictions. We study two complementary design levers to calibrate and control the platform's market thickness in the face of complex demand-side decision making: (i) its listing policy, which determines the ending times of auctions, and (ii) a recommendation system. To optimize these design decisions, we first develop a structural model to characterize how bidders form expectations and respond to the imminent availability of auctions in equilibrium, including how frequently they visit the platform, in which auctions they choose to participate, and their bidding strategies. In calibrating its market thickness, the platform trades off increasing bidder participation in each auction by appropriately thickening the market (demand-side competition) against limiting the extent to which auctions for substitutable goods ultimately cannibalize one another under thicker market conditions (supply-side competition). Using our structural estimates, we illustrate how the platform can optimize its listing policy as a function of the incoming liquidation inventory and its bidder pool so as to achieve a supply-demand *sweet spot*, thereby increasing its revenue significantly relative to having auctions end after a fixed time. Furthermore, we find that real-time recommendations sent on the market's thickest days would add 3% revenue on such days (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 thickness, Matching supply with demand; Listing policy; Natural experiment; Structural estimation.

1. Introduction

The emergence of Internet-enabled services, such as Airbnb and Lyft, has highlighted that online marketplaces greatly reduce frictions to connect buyers and sellers in previously disparate spheres, thereby increasing the volume of trade in a number of markets. Typically, such platforms act as intermediaries between individuals and/or businesses; in particular, they neither own nor directly control the goods involved in each transaction. As such, 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 not only manifested itself in the more avant-garde businesses, exemplified by Lyft and Airbnb, but has also reshaped retail operations, including 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, i.e., the retailers' merchandise that either remained unsold in its primary market (e.g., due to decreasing 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 fast fashion (Ferdows et al. (2005)), 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: offloading this inventory to business buyers further down the retail food chain, such as discount stores, or donating to qualifying recipients for tax purposes.

Online business-to-business (B2B) auction platforms, such as our industry sponsor, connect an increasing number of retailers to deeper pools of potential business buyers, both domestic and foreign. Given the diversity of potential bidders, who 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 of more than $600 million in aggregate

[1] http://www.sdcexec.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

retail value.[2] Today, many national chain retailers, including Best Buy, Walmart, Home-Depot, Amazon, Target, and Costco, employ online B2B auctions to liquidate their products.

Supply in these online auctions is highly variable, owing to the uncertain and dynamic nature of product returns and retailer decisions as to when excess inventory is pulled from shelves. The manifestation of this variability on the platform is varying levels of auctions open at any point in time. The uncertainty in supply coupled with the stochastic demand for these products and the uncertain valuation of potential buyers implies that liquidation platforms face a familiar operational challenge: how to design online auctions so as to profitably match supply with demand. One lever auction platforms have in attaining this match is their *listing policy*, and specifically, the selection of auctions' closing dates. By aligning, or conversely, spreading out, the closing dates of auctions, the platform can manipulate its market thickness (number of auctions scheduled to close on any given day) and, in turn, influence the trade-off between bidder participation in any given auction (demand-side competition) with the degree of cannibalization across competing auctions (supply-side competition) in favor of increasing its own revenues.

Using a proprietary dataset collected from a leading online B2B platform, we investigate the role and efficacy of market thickness in coordinating the behavior of market participants so as to influence market outcomes, namely the auctions' final prices. Notably, while a platform's choice of listing policy alters neither its underlying supply of arriving liquidation inventory nor its pool of potential bidders, its role in incentivizing bidder behavior can be of operational (and revenue-enhancing) importance. Exploiting a natural experiment, we find that the platform's listing policy significantly impacts its revenue: implementing a listing policy to concentrate ("batch") auctions' ending times to certain days of the week increases sellers' revenues by 6.5%. Related evidence supports that the existence of market frictions, specifically the costs associated with the bidders visiting the platform and determining their bidding strategies, drives this result.

Prescriptively, we then study two complementary market design levers available to the platform to profitably calibrate its market thickness: its listing policy and a recommendation system. For this purpose, we first develop a structural model to characterize the bidders' equilibrium decisions, including how frequently they visit the platform, the auctions they choose to participate in, and their bidding strategies. Our model and its predicted equilibrium incorporate the dynamic nature of the platform's bidder pool, bidders' monitoring costs, valuation and demand heterogeneity, as well as the uncertainty surrounding the supply-side inventory of auctions available on any given day. The estimates we obtain from the structural model, using data from a separate time period, corroborate out-of-sample our findings from the natural experiment.

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

Our counterfactual results can be summarized as follows. Using our framework, we systematically explore listing policies for the auction platform to optimize an average market thickness *sweet spot* that depends on the underlying levels of supply against demand, while accounting for stochastic uncertainty on the supply side and potential bidders' complex and endogenous decisions about participating and bidding on the demand side. On one hand, the optimal *sweet spot* prescribes that the platform should restrict its auctions' ending dates to a single day per week when the rate of supply is sufficiently low, so as to incentivize bidder participation. Conversely, the platform should spread auctions' end days throughout the week to thin the market's daily closing inventory when the level of supply is high and, thus, cannibalization concerns loom large.

Furthering the spirit of reducing marketplace uncertainty and frictions, we next consider a recommendation system to notify potential bidders of supply conditions on the auction site (since while potential bidders may know the platform's listing policy, they do not typically know the exact number of auctions listed prior to visiting the platform itself). On days featuring especially ample supply, the recommendation system informs a (randomly) sampled set of recipients, allowing them to gain that information without having to visit the site. We find that decreasing supply uncertainty via such solicitations improves revenues by an additional 2.9% on these days. This gain is achieved by reducing cannibalization and shifting the composition of participating bidders towards those with higher valuations. Interestingly, designing such a recommendation system involves appropriately choosing the number of bidders to receive such information given that the value of the recommendation for bidders and, consequently, the likelihood that it will induce an increase in their participation is endogenously determined by the ensuing competition on the platform.

1.1. Related Literature

Online marketplaces face a number of design challenges in seeking to match supply with demand so as to maximize revenue. Relevant to our study, recent papers explore aspects of marketplace design that involve shaping or even tailoring 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 increased service levels whereas, Bimpikis et al. (2017) consider pricing for spatially dispersed demand in a ride-sharing network. In an online auctions setting, Balseiro et al. (2015) introduce the notion of a fluid mean-field equilibrium and illustrate its practical appeal in setting reserve prices. Golrezaei et al. (2017) show empirically that the platform's revenue increases by 3% when it boosts bids in a customized fashion based on the bidders' past behavior. Our work contributes to this literature by empirically demonstrating the impact of market thickness on how users participate and submit bids on a platform specializing in B2B auctions. 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.[3]

Recent literature connects participants' transaction costs and information frictions to outcomes in online markets. Fradkin (2017) studies the role of search to reduce transaction costs and improve matches on Airbnb, while Tadelis and Zettelmeyer (2015) explore the disclosure of product information in the online marketplace. Moon et al. (2017) study the revenue and welfare implications of customers' costs to monitor a retailer's online channel for changes in the price and availability of inventory they are interested in. 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) using data from TaskRabbit, a marketplace for domestic tasks, empirically demonstrate that the growth of online peer-to-peer markets is largely affected by the thickness they induce, whereas Li and Netessine (2017) find that higher market thickness actually leads to lower matching efficiency in an online peer-to-peer holiday rental platform. 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.

Also relatedly, prior work explores both the effects of inventory availability on strategic demand and the related benefits of reducing buyers' uncertainty around 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 to visit a physical store with the intent to purchase the product, if 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 availabilities. 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, our platform can neither set prices nor control the stochastic arrival of inventory to its marketplace. Nonetheless, we highlight that listing policies and state-contigent recommendations (communicated to a subset of bidders) can successfully complement the platform's efforts in boosting the bidders' participation rates.

Furthermore, our paper falls within the growing body of literature employing structural estimation methods to study auction markets, e.g., Jofre-Bonet and Pesendorfer (2003), Sailer (2006), Kim et al. (2014), and Backus and Lewis (2016). Although our focus is on the impact of market thickness on revenues for the platform (and not on the auction mechanism itself), our analysis

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

shares similarities with this line of work. Finally, our work is related to the burgeoning literature that uses structural estimation 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 national big-box retailers, such as Costco, Walmart, Sears, and Home Depot (henceforth, referred to as *sellers*). The platform's primary selling mechanism is the second-price auction. An auction lasts for two to three days on average and offers for sale an *auction lot*, which is a bundle of similar products typically exceeding $10,000 in its aggregate retail value. While electronics account for most of the platform's annual revenues, which are in the hundreds of millions of US dollars, a broad range of product categories are auctioned, including household appliances, furniture, and apparel.

Sourced from customer returns, trade-ins, and unsold items, the supply of liquidation inventory reaching the platform is highly stochastic and beyond the platform's control, driven primarily by the sellers' reverse logistics. 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 a good condition). Bidders may access the auction's *manifest*, which provides a brief description of the items included in the box. In addition, bidders 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 also Pilehvar et al. (2016)). Finally, a dynamic pool of potential bidders features the continuous arrival of new registrants and the ongoing exit of some existing bidders. Figure 1 summarizes the market dynamics and the bidders' major decisions.

2.1. Markets for iPhones

For our empirical analysis, we use data on the iPhone auctions held on the platform's two major mobile phone markets, Markets A and B, each managed by the platform for a distinct big-box seller. We restrict attention to iPhone auctions for two reasons. First, iPhones are well-defined products with retail values that are derived in a straightforward way from their observable specifications, i.e., model, carrier, condition, and time (i.e., depreciation), unlike other merchandise, such as furniture and household appliances, sold on the platform. Second, iPhone sales alone account for approximately 73% of these markets' revenues.

Figure 1 **Market setting and bidders' decisions on a given day**

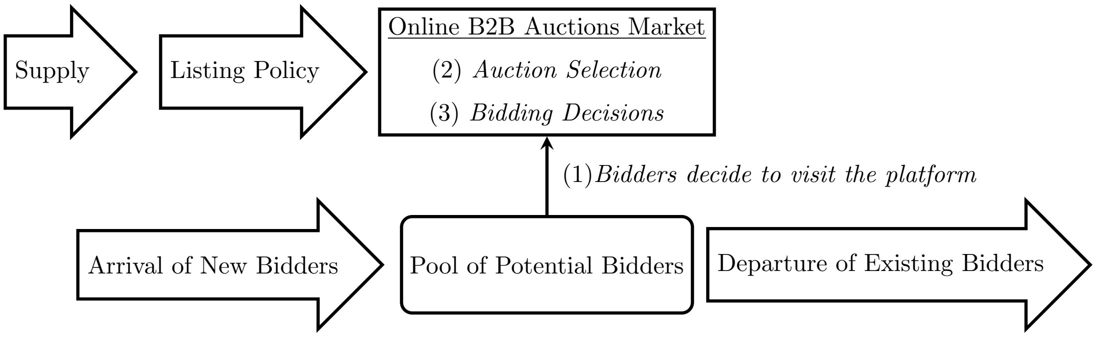

Table 1 **Summary statistics—mean values—for Markets A and B (standard errors in parentheses)**

	Market A	Market B
Auction duration (in days)	2.61 (1.55)	2.88 (1.09)
Auction lot size (units)	117.2 (60.6)	65.2 (41.1)
Avg monthly registrations	211.0 (44.6)	260.7 (63.2)
Number of bidders per auction	4.9 (2.1)	8.2 (2.6)
Number of bids per auction	19.6 (12.0)	27.2 (15.9)
Number of auctions per auction ending day	5.79 (4.71)	2.33 (1.58)
Total number of auctions	1,125	658
Avg final per-unit price ($)	118.60 (55.81)	109.02 (59.10)

Covering February 2013 to October 2015 (our *observation period*), our dataset tracks the entire bidding history, i.e., time of submission and dollar amount of every bid, in each of the 1,125 auctions in Market A and 658 auctions in Market B. Moreover, we are able to track bidders' behavior across multiple auctions throughout the observation period. Table 1 provides summary statistics comparing the two markets. On average, there are 211 newly-registered bidders in Market A and 261 in Market B per month. Auction lot sizes vary substantially, both within markets (the corresponding coefficients of variation are 0.51 for Market A and 0.63 for Market B) and across (the average number of iPhones per auction is 117 in Market A versus 65 in Market B). Participation per auction, whether tallied in bidders or bids, tends higher in Market B's auctions. Nonetheless, Market A exhibits slightly higher average per-unit revenues ($118.6) than Market B ($109.0).

In the remainder of the paper, we argue that this disparity cannot be explained by the markets' differences in product and bidder characteristics alone; instead, it also reflects the operational design of the two markets, specifically their respective listing policies.

2.2. Listing Policies

In this subsection, we discuss the *listing policies* employed by Markets A and B. In both markets, as soon as the seller makes available an inventory lot to the auction platform, a corresponding auction commences and is listed onto the platform (its *listing time*). However, the two markets differ in the length of their auctions, i.e., the duration over which an auction remains open for bids. Market

B's auctions typically close three days following their listing times. Because inventory arrives roughly uniformly throughout the week, the number of auctions ending on any given weekday is approximately the same. Used in Market B throughout our observation period, this listing policy is referred to as "uniform."

In contrast, over the second half of our observation period starting in November 2014 (more details on this in Section 3), Market A's auctions close only on Tuesdays and Thursdays. Consequently, the duration over which an auction remains open is variable, depending on its initial listing time. We use the term "batch" to refer to this listing policy. Figure 2 illustrates these two listing policies under the same supply process.

We define *market thickness* as the number of closing auctions on an auction-ending day under a listing policy. Bidders bid in an auction nearly exclusively on its closing day (Table 2), and thus the market thickness constitutes the platform's supply of concurrently available auctions. In particular, Tuesday and Thursday constitute the auction-ending days under the batch policy, while all weekdays qualify under the uniform policy. Figure 2 illustrates that listing policies lead to different distributions of market thickness: the batch policy maintains a "thicker" market than the uniform policy.

In the subsequent sections, we argue that the difference in the revenues generated by Markets A and B can be explained to a large extent by this seemingly innocuous divergence in their listing policies. The primary tradeoff can be summarized intuitively as follows. Batching auctions' closing times together increases the number of options available to a bidder visiting on an auction-ending day. On one hand, batching encourages bidders to enter auctions by coordinating and concentrating their costly participation onto a few pre-specified days: increased participation boosts the auctions' prices. On the other hand, closing multiple auctions on the same day implies that they compete as substitutes, potentially deflating their revenues. To balance these considerations, the platform's listing policy serves as an operational lever to modulate and match the supply of available, open auctions against the demand of participating bidders.

2.3. Bidders

More than 2,200 bidders placed at least one bid in Markets A and B during the observation period. In line with prior work on B2B markets, e.g., Bapna et al. (2004), Pilehvar et al. (2016), the markets' bidder pools consist of experienced resellers—certified and registered in the market—that are heterogeneous in both their demand profiles and valuations. Interestingly, at least 33% of bidders in both markets exhibit demand for multiple auctions. In particular, such *multi-unit* (MU) bidders either submit winning bids in two or more concurrent auctions or submit a bid in a new auction shortly after winning an auction. The remaining bidders, whom we call *unit-demand*

Figure 2 Examples of the batch (above) and uniform (below) listing policies—each red arrow represents an auction

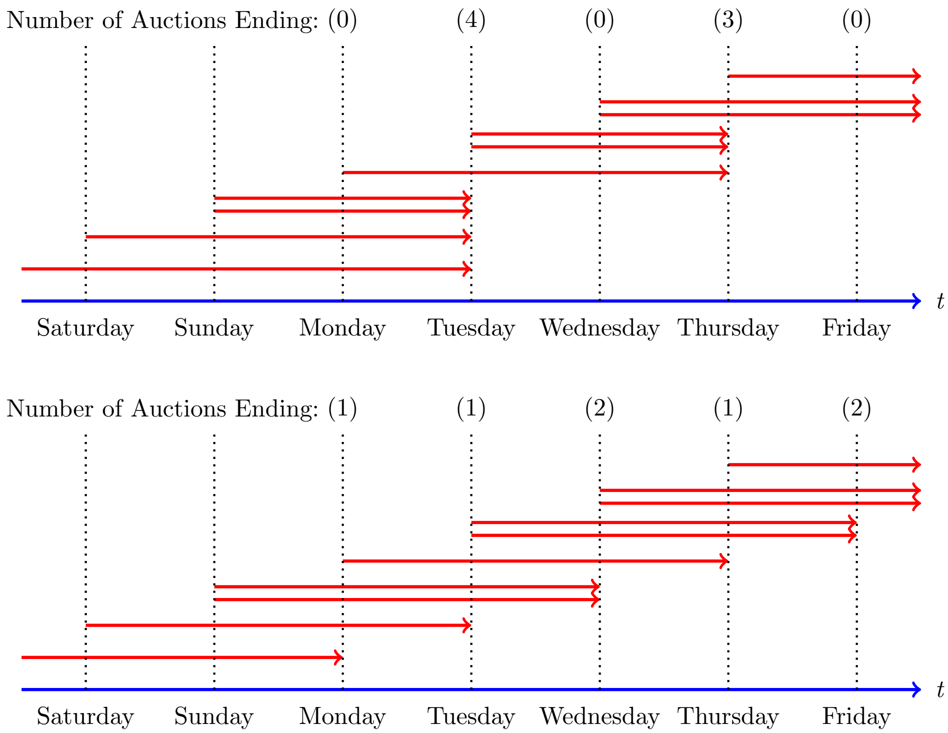

(UD) bidders, exhibit demand for winning only a single auction lot within our observation period. Though fewer in number, the behavior of multi-unit bidders affects the operations of the platform substantially, as they ended up winning more than 80% of the auctions in both markets during our observation period.

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

Akin to observations in B2C auctions (Bajari and Hortacsu (2003)), we find evidence of "sniping" behavior, in which bidders submit bids in an auction's last few minutes. As we report in Table 2, the median last bid per bidder per auction comes quite late in both markets (after 98.8% of the auction's total duration in Market A and 80.0% of the auction's total duration in Market B has elapsed). Moreover, the median winning bid arrives when 99.7% of the auction's duration has elapsed (similar to 98.3% in Bajari and Hortacsu (2003)); thus, an auction's final price materializes towards its end. Consequently, focusing on when auctions end—as opposed, for example, to the

Table 2 Summary statistics for bidders

	Market A	Market B
Percentage of multi-unit bidders (%)	36.4	33.2
Percentage of auctions won by multi-unit bidders (%)	88.7	83.8
Avg number of auctions a multi-unit bidder participates in	41.64	21.74
Avg number of auctions won by a multi-unit bidder	8.98	3.14
Avg number of auctions a unit-demand bidder participates in	5.31	3.09
Avg number of auctions won by a unit-demand bidder	0.65	0.30
Median of normalized time of last bid per bidder-auction (%)	98.8	80.0
Median of normalized time of winning bid per auction (%)	99.7	99.7

entire time they are open—is a reasonable assumption when studying the impact of the platform's listing policy on its revenues.

3. Natural Experiment

In this section, we estimate the revenue impact of the platform's listing policy by exploiting a natural experiment arising from a switch in Market A's listing policy.

Previewing our findings, we estimate that implementing batch instead of uniform listing boosts Market A's revenue by 6.5%, amounting to roughly $3.4M annually. Moreover, we examine the participation rates of the subset of cross-market bidders: we observe that when Market A batches its listings, cross-market bidders substantially increase their rates of participation in the *other* market, Market B, by 39% solely on Market A's auction-ending days, i.e., Tuesdays and Thursdays, but not on other days. This spillover pattern is consistent with bidders facing a daily cost to actively participate on the online platform. Our finding suggests that bidders strategically determine when to visit the platform and participate in auctions based on how thick they expect the underlying markets to be. Together, these results support using the platform's listing policy as an operational lever to match supply and demand in online markets.

3.1. Setting

As described in Section 2.2, both markets deployed the "uniform" listing policy before February 2014, under which their auctions closed roughly uniformly throughout the week. Triggered by a personnel change cited by management as being entirely unrelated to the performance of the platform and the markets, Market A altered its listing policy: first, in February 2014, constraining its auctions to close only on three days of the week (Monday, Tuesday, and Thursday); then in November 2014, adopting the "batch" policy described in Subsection 2.2 to close its auctions only on Tuesdays and Thursdays. Meanwhile, Market B's listing policy remained unaltered throughout the observation period.

Thus, we consider "batching" the listing policy as the treatment of interest. We observe two clearly defined periods: the *pre-treatment* period (November 2013 to February 2014), during which

Figure 3 **Percentage of auctions ending on each weekday aggregated by month in Market A (Left) and Market B (Right)**

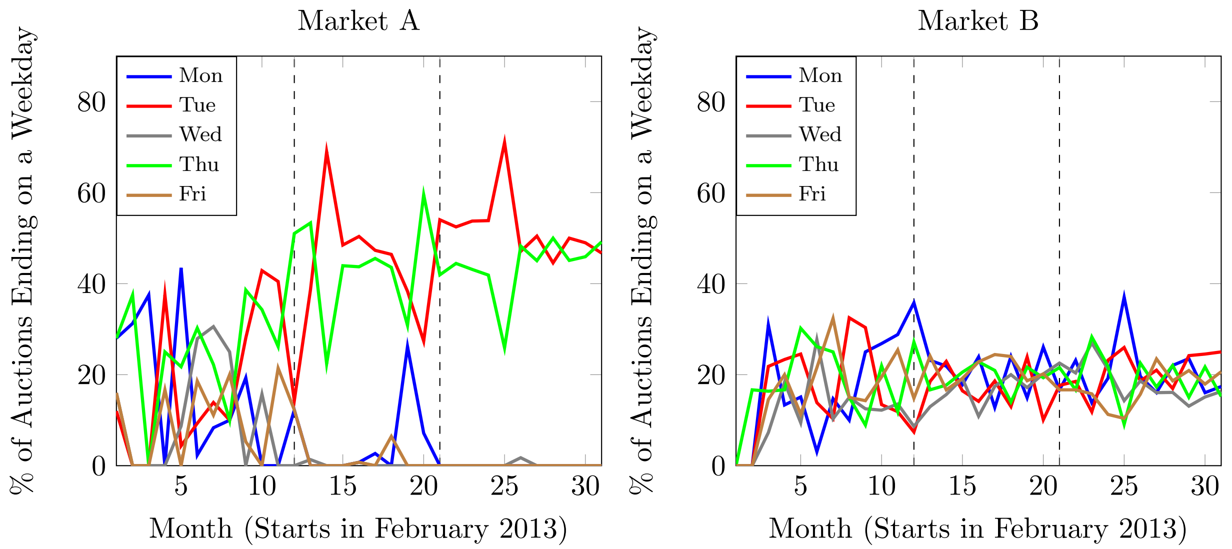

Table 3 **Number of auctions and average market thickness (in the brackets) per experimental subsample**

	Control (Market B)	Treatment (Market A)
Pre-Treatment Period	94 [1.25]	114 [1.52]
Post-Treatment Period	229 [1.12]	317 [3.87]

both markets practiced uniform listing, and the *post-treatment* period (November 2014 to October 2015), during which Market A alone batch listed its auctions. To illustrate the pre-treatment, intervening (from February to November 2014), and post-treatment periods, Figure 3 depicts the percentage of auctions ending on each of the five weekdays, aggregated monthly for each of Markets A and B. Importantly, observing both markets' pre-treatment allows us to account for unobserved differences between the two markets. The numbers of auctions and the average market thickness in the four relevant subsamples are shown in Table 3. Relying on this natural experiment, our goal in the remainder of this section is to credibly estimate the causal effect of a market's listing policy on its revenue.

As a primer to the analysis that follows, Figure 4 compares the two markets' average, per-unit revenues for the pre-treatment and post-treatment periods, respectively. While iPhone prices fall over time in both markets as the considered phone models depreciate, the plot highlights an emerging gap in the markets' per-unit 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-unit revenue.

Figure 4 **Average per-unit revenue in the pre-treatment and post-treatment periods in Markets A and B**

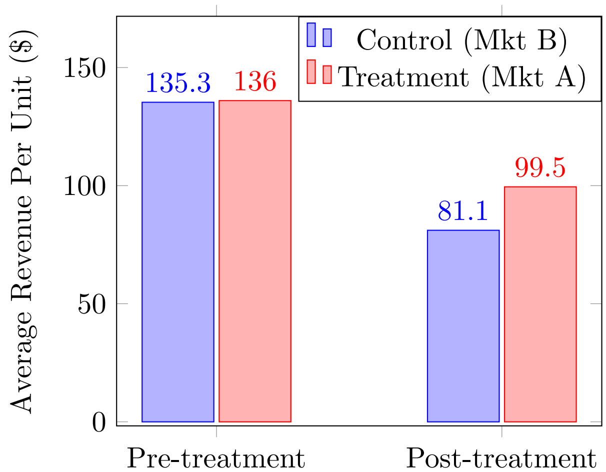

3.2. Average Treatment Effect

Beyond the posting market's listing policy, an auction's per-unit revenue is affected by the attributes of the iPhones composing its auction lot, the posting market's additional characteristics and operational practices, and temporal effects that encompass new product releases, depreciation, and price fluctuations in the overall iPhone market. Using the average treatment effect (ATE) methodology reviewed in Imbens (2004), we estimate the effect of the platform's listing policy on an auction's final price while controlling for these factors.

First, our analysis controls explicitly for observables, such as the relevant iPhone model, condition, 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 non-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 in their reputations and in the compositions of their respective bidder pools.

Specifically for this analysis, we designate an auction lot j at time t as treated if posted to Market A (indicator variable A_{jt}) and as not treated if instead posted to Market B. Based on its observed attributes, each auction lot exhibits a propensity for receiving the treatment, 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, τ, that accounts for such attributes can be obtained as the difference of post-treatment and pre-treatment outcomes appropriately weighted using their associated propensity scores. We estimate separate average treatment effects for each of the pre-treatment and post-treatment periods (τ_{Pre} and τ_{Post}, respectively) and are ultimately interested in their difference. 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 batching listings under Market A's revised listing policy.

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:

$$Logit(e_{jt}) = \tilde{\delta}_t + \tilde{\boldsymbol{\beta}}^T AUC_{jt}, \tag{1}$$

where $\tilde{\delta}_t$ and $\tilde{\boldsymbol{\beta}}$ respectively denote the week-t fixed effect and attribute-coefficient vector. Using the dependent variable A_{jt}, we first estimate model (1) by logistic regression to obtain predicted propensity scores, $\hat{e}_{jt}$.

The observed outcome of interest is the log-per-unit revenue obtained in auction jt, LFP_{jt}. We estimate the treatment effect, τ, by its augmented inverse probability weighted (AIPW) estimator (for details, refer to Bang and Robins (2005) and He et al. (2016)):

$$\hat{\tau} = N_{obs}^{-1} \sum_{jt} \Big(\frac{A_{jt} \cdot LFP_{jt}}{\hat{e}_{jt}} - \frac{A_{jt} - \hat{e}_{jt}}{\hat{e}_{jt}} \cdot \widehat{LFP}_{jt,1} \Big) - N_{obs}^{-1} \sum_{jt} \Big(\frac{1 - A_{jt} \cdot LFP_{jt}}{1 - \hat{e}_{jt}} + \frac{A_{jt} - \hat{e}_{jt}}{1 - \hat{e}_{jt}} \cdot \widehat{LFP}_{jt,0} \Big), \tag{2}$$

where N_{obs} is the period's sample size, i.e., corresponding to the pre- and the post-treatment periods respectively, and $\hat{e}_{jt}$ is the estimated propensity score for auction jt obtained by Expression (1). We restrict attention to auctions with estimated propensity scores between 0.1 and 0.9 in order to ensure that the requisite overlap assumption holds (Imbens (2004)). While the classical ATE estimate is derived as the difference in observed outcomes weighted appropriately using the associated propensity score projections, $\hat{e}_{jt}$, the 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 specification of log-revenue outcomes, as an alternative to the classical ATE:[4]

$$LFP_{jt} = \beta_A A_{jt} + \delta_t + \boldsymbol{\beta}^T AUC_{jt} + \epsilon_{jt}, \tag{3}$$

where β_A is the coefficient corresponding to the indicator of Market A, δ_t is a fixed effect for week t, and $\boldsymbol{\beta}$ denotes the attribute-coefficient vector. The AIPW estimator consistently estimates the average treatment effect if *either* the classical ATE's propensity score model, i.e., Expression (1), or the linear outcomes model, i.e, Expression (3), is well-specified.

Accordingly, $\hat{\tau}_{\text{Pre}}$ and $\hat{\tau}_{\text{Post}}$ are each estimated by Expression (2). As reported in Table 4, we find the per-unit revenue effect of the batch listing policy, $\hat{\tau}_{\text{Post}} - \hat{\tau}_{\text{Pre}}$, to be positive and statistically

[4] That is, $\widehat{LFP}_{jt,0} = \hat{\delta}_t + \hat{\boldsymbol{\beta}}^T AUC_{jt}$ and $\widehat{LFP}_{jt,1} = \hat{\beta}_A + \hat{\delta}_t + \hat{\boldsymbol{\beta}}^T AUC_{jt}$.

Table 4 **Price effect of the batch listing policy**

Dependent Variable: Log (Final Price Per Unit)	Treatment: Batched Policy
ATE	6.5%***
	(0.009)
	*p<0.05; **p<0.01; ***p<0.001

significant, with our estimated confidence interval constructed via bootstrap. Further validating our estimate, a difference-in-difference approach yields a significant, positive effect of similar magnitude.[5]

3.3. Revenue Effect of Batch Listing Policy

We find that the batch listing policy yields a 6.5% average increase in Market A's per-unit revenue, statistically significant at a p-value of 0.001, over the uniform listing policy. This estimate translates to more than $3.4M in additional revenues annually for Market A.

We discuss the intuition behind our finding. In particular, setting a market's listing policy may appear an innocuous choice borne of convenience, chance, or custom: why would the platform expect beyond a marginal impact on its revenue, given that its listing policy does not alter its exogenous supply of auctions nor its demand pool of certified bidders? Instead, we argue that the sizable revenue impact of the listing policy stems from a subtle interplay between the platform's market thickness and bidders' incentives to visit and actively engage with the platform.

We provide further evidence by examining the participation decisions of the subset of cross-market bidders. Before and after the adoption of the batch policy in Market A, we want to examine whether cross-market bidders exhibit different participation patterns across weekdays in Market B. In particular, we are interested in comparing their average participation rates in Market B on Tuesdays and Thursdays to those on Mondays, Wednesdays, and Fridays. To this end, we study the ratio of the total number of participants on Tuesdays and Thursdays over the total number of participants on the remaining weekdays. Then, assuming that cross-market bidders are equally likely to participate in the market on any given weekday would imply a ratio of 2/3.

As shown in Table 5, 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 rocketed to 84%, suggesting that these bidders heavily preferred to participate in Market B on the auction-ending days of the *other* market, i.e., Market A (Tuesdays and Thursdays). We formally

[5] Because our presented ATE estimate more effectively accounts for granular differences in auctions' attributes, we omit this additional analysis.

Table 5 Participation rates of cross-market bidders

	Pre-Treatment	Post-Treatment
Market B:		
Ratio of Cross-Market Bidders on TuTh over MWF	63%	84%
Market A:		
Average Number of Bidders per Auction	4.57	5.12

analyze this spillover effect in the cross-market bidders' participation using a difference-in-difference methodology in Appendix C.1.

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 incentivized to be selective in timing when to visit the platform and actively participate in auctions. In other words, bidders choose to visit only when their expected payoff from doing so exceeds some threshold. While our empirical finding of a positive revenue effect indicates that batch-closing multiple auctions can profitably coordinate bidders' participation and demand, there exist potential pitfalls to "thickening" the market. First, batching too many auctions together may cannibalize and decrease per-unit revenues by increasing the number of immediately available substitutes. Second, some bidders may be disincentivized from participating on the platform due to the increased participation of competing bidders—higher participation per auction under batched listing is shown in Table 5 for Market A.

Accordingly, the remainder of our paper addresses the appropriate design of a platform's listing policy to achieve its optimal market thickness. In Section 4, we develop a structural model endogenizing bidders' decisions. Section 5 covers the model's estimation to recover bidders' valuations and participation costs. Finally, Section 6 presents and discusses our counterfactual analysis, including an extension to a recommendation system that aims to further facilitate the process of matching supply with demand.

4. Structural Model

We present a dynamic, structural model endogenizing the bidders' decisions on whether and when to visit the platform, in which auction(s) to participate, and how much to bid. Our discrete-time model captures the behavior of a dynamic pool of potential bidders heterogeneous in their valuations and demand profiles, while facing an exogenously stochastic supply of liquidation inventory lots arriving to the platform to be auctioned.

Model Setup: We capture heterogeneity in the bidders' demand profiles in a tractable way by having two types of bidders as described below.

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 unit-demand 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 multi-unit bidder is interested in winning multiple auctions. 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 history. Each MU bidder's endowed private valuation for an auction lot is drawn independently from distribution F^{MU}.

In contrast to UD bidders that 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.

On each day t, the following sequence of events transpires.

(1) *Supply.* First, S_t new auctions are listed onto the platform, with S_t following the count distribution P_{Supply}. Each posted auction lot possesses characteristics affecting its idiosyncratic fit for a bidder's resale channels. Therefore, bidder ℓ's valuation for an auction lot j is the sum of bidder 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.* Next, new arrivals join the existing pool of potential bidders, under an exogenous arrival process. If t is an auction-ending day as defined in Sections 2.2 and 2.3, all potential bidders in the pool must decide simultaneously whether to visit the platform to bid on open auctions. Notably, while each bidder ℓ knows her own valuation x_ℓ for the product throughout, prior to visiting she observes neither the current state of the platform defined below nor 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, the observable (upon visiting the platform) state of the market on day t is:

$$\omega_{\ell t} := (n_t, \mathbf{s}_{\ell t}),$$

where n_t is the market thickness (number of auctions ending) on day t, with support $\mathcal{N} = \{0, 1, 2, ..., \bar{N}\}$, whereas $\mathbf{s}_t$ 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 evaluating the auctions

Figure 5 Timeline of events for bidder ℓ on an auction ending day

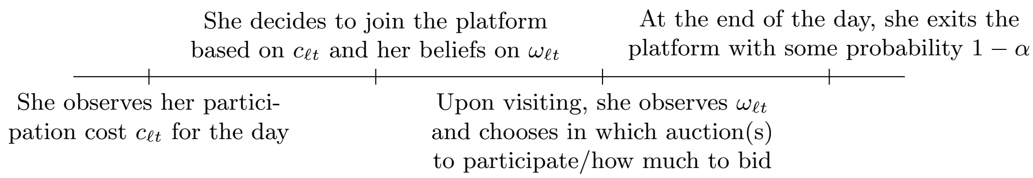

listed on the platform and determining a bidding strategy. Conditional on a bidder's type, her participation cost is drawn independently each day from an exponential distribution with rate μ^{MU} for MU and μ^{UD} for UD bidders. Thus, across bidders and days, whether a bidder visits the platform varies both on her endowed valuation and on her realized daily participation cost, which includes the value of a bidder's outside option.

(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 in which auction(s) to participate based on her private valuations and her beliefs about the currently highest rival bids given $\omega_{\ell t}$ and, then, she determines how much to bid in each auction so as to maximize her expected payoff.

(4) *Departure.* Finally, we let α^{MU} and α^{UD} denote the daily retention probabilities of MU and UD bidders respectively. In particular, at the end of day t, each bidder departs the bidder pool with probability $(1-\alpha^{TY})$, where $TY \in \{MU, UD\}$. Note that UD bidders also depart with certainty upon winning an auction.

4.1. Equilibrium

Given the size of the market (the number of active bidders in Market A is in the hundreds), it is 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), Hendricks and Sorensen (2016)). This assumption approximates well a setting that involves a large group of anonymous bidders, with similarities to other approaches 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 some notation pertaining to a bidder's actions on a given day (as specified in Figure 5) under the belief about the underlying market state Ψ and the highest rival bid G:

(i) A bidder's platform visit decision, which we denote by σ_{VST}^{MU} and σ_{VST}^{UD} depending on the bidder's type, is a mapping from her endowed valuation and her participation cost on that day to 1 (visit) or 0 (wait), i.e., for bidder ℓ we have:

$$\sigma_{VST}^{TY}(x_\ell; c_{\ell t}) : [0,\infty) \times [0,\infty) \to \{0,1\} \text{ , where } TY \in \{MU, UD\}.$$

Bidders determine when to visit so as to maximize their 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. Specifically,

$$\sigma_{VST}^{MU}(x_\ell, c_{\ell t}) = \begin{cases} 1, & \text{i.e., visit, } \quad \text{if } r^{MU}(x_\ell; G, \Psi) \geq c_{\ell t} \\ 0, & \text{i.e., wait, } \quad \text{otherwise} \end{cases}, \tag{4}$$

where $r^{MU}(x_\ell; G, \Psi)$ denotes the expected payoff for a MU bidder who visits the platform on day t. Similarly, for UD bidders, we have:

$$\sigma_{VST}^{UD}(x_\ell, c_{\ell t}) = \begin{cases} 1, & \text{i.e., visit, } \quad \text{if } r^{UD}(x_\ell; G, \Psi) \triangleq v^{UD}(x_\ell; G, \Psi) - \alpha^{UD} v_f(x_\ell; G, \Psi) \geq c_{\ell t} \\ 0, & \text{i.e., wait, } \quad \text{otherwise} \end{cases}, \tag{5}$$

where $v^{UD}(x_\ell; G, \Psi)$ denotes the expected 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 t. We provide additional details on $r^{MU}(x_\ell; G, \Psi)$, $v^{UD}(x_\ell; G, \Psi)$, and $v_f(x_\ell; G, \Psi)$ in Section 4.2.

(ii) A bidder's auction selection decision, which we denote by σ_{SLT}^{MU} and σ_{SLT}^{UD} depending on the bidder's type, is a mapping from her endowed valuation, the idiosyncratic valuation terms corresponding to each of the open auctions, and the realized state of the platform, to a binary vector that specifies which auction(s) the bidder chooses to participate in, i.e., for bidder ℓ we have:

$$\sigma_{SLT}^{TY}(x_\ell; \zeta_{\ell 1}, \cdots, \zeta_{\ell \bar{N}}, \omega_{\ell t}) : [0,\infty) \times \mathbb{R}^{\bar{N}} \times (\mathcal{N} \times \mathbb{R}^{\bar{N}}) \to \{0,1\}^{\bar{N}} \text{ , where } TY \in \{MU, UD\}.$$

Note that the j^{th} element of σ_{SLT}^{TY} denotes whether the bidder decides to participate in auction j.

(iii) A bidder's bidding decision, which we denote by σ_{BID}^{MU} and σ_{BID}^{UD} depending on the bidder's type, is a mapping from her endowed valuation, the idiosyncratic valuation terms corresponding to each of the open auctions, and the realized state of the platform to a vector of bids, i.e.,

$$\sigma_{BID}^{TY}(x_\ell; \zeta_{\ell 1}, \cdots, \zeta_{\ell \bar{N}}, \omega) : [0,\infty) \times \mathbb{R}^{\bar{N}} \times (\mathcal{N} \times \mathbb{R}^{\bar{N}}) \to \mathbb{R}^{\bar{N}} \text{ , where } TY \in \{MU, UD\}.$$

Here, the j^{th} element of σ_{BID}^{TY} denotes how much the bidder decides to bid in auction j.

Note that σ_{SLT}^{TY} and σ_{BID}^{TY} jointly maximize the expected payoff of a bidder who chooses to visit the platform on a given day. In particular, for MU bidders, the corresponding maximization problem can be written as:

$$(\sigma_{SLT}^{MU}, \sigma_{BID}^{MU}) = \underset{\sigma_{SLT}, \sigma_{BID}}{\arg\max} \sum_{\sigma_{SLT,j}=1} \int_{s_j}^{\sigma_{BID,j}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j | \omega_{\ell t}) dp_j, \tag{6}$$

where $g_j(p_j|\omega_{\ell t})$ denotes the conditional PDF of the highest rival bid in auction j. On the other hand, for UD bidders, we have:

$$(\sigma^{UD}_{SLT}, \sigma^{UD}_{BID}) = \underset{\sigma_{SLT}, \sigma_{BID}}{\arg\max} \sum_{\sigma_{SLT,j}=1} \int_{s_j}^{\sigma_{BID,j}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j|\omega_{\ell t}) dp_j \\ + \alpha^{UD}\big(1 - G_j(\sigma_{BID}|\omega_{\ell t})\big) v_f(x_\ell; G, \Psi), \quad (7)$$

where $G_j(p_j|\omega_{\ell t})$ denotes the conditional CDF of the highest rival bid in auction j (we revisit Expressions (6) and (7) in Section 4.2).

Formally, the notion of a 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 of three decisions on day t, i.e., it takes the form $\sigma^{TY}(x_\ell; c_{\ell t}, \zeta_{\ell \cdot}, \omega_{\ell t}) = [\sigma^{TY}_{VST}, \sigma^{TY}_{SLT}, \sigma^{TY}_{BID}]$, where $\sigma^{TY}_{VST}, \sigma^{TY}_{SLT}$, and σ^{TY}_{BID} are defined by Expressions (4), (5), (6), and (7), 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 $\sigma^{TY}(x_\ell; c_{\ell t}, \zeta_{\ell \cdot}, \omega_{\ell t})$, $TY \in \{MU, UD\}$.

Naturally, in the context of our model, an auction's standing and highest rival bids are not independent. Thus, bidders update their beliefs about the highest rival bids after observing the corresponding standing bids. Before describing in detail how bidders behave at equilibrium depending on their type, we first establish that an equilibrium exists under a (mild) technical assumption.

Proposition 1. *Assume that $\nu^{MU} = \nu^{UD} = 0$, i.e., the bidder-level idiosyncratic valuation terms are equal to zero. In addition, assume that the maximum number of auctions on any given day $\bar{N}$ is smaller than K, i.e., the upper bound on the number of auctions in which multi-unit bidders may consider to participate on a day. Then, an equilibrium exists.*

The proof of Proposition 1 can be found in Appendix A.2. Additionally, in Appendix C.2 we provide an algorithm that efficiently converges to the equilibrium and does not rely on the assumptions of the proposition, i.e., it applies to a setting where $\nu^{MU} > 0$, $\nu^{UD} > 0$, and the number of open auctions on a day is higher than K (thus, the algorithm enables a broad set of counterfactuals that we present in Section 6).

4.2. Bidders' Equilibrium Strategies

We set forth equilibrium decision-making for the multi-unit and unit-demand bidder types.

Multi-Unit Bidders: On any given day t, a multi-unit bidder can participate in up to K auctions. Her bidding decision does not affect her future payoffs as her future demand is independent of her current actions. Therefore, the conditional payoff of a multi-unit bidder ℓ with valuation x_ℓ given that she has already entered the market can be specified as follows:

$$u^{MU}(x_\ell;\zeta_{\ell\cdot},\omega_{\ell t},G) = \max_{\sigma^{MU}_{SLT},\sigma^{MU}_{BID}} \sum_{\sigma^{MU}_{SLT,j}=1} \int_{s_j}^{\sigma^{MU}_{BID,j}} (x_\ell + \zeta_{\ell j} - p_j) g_j(p_j|\omega_{\ell t}) dp_j. \tag{8}$$

We establish that $g_j(p_j|\omega_{\ell t})$ is independent of the vector of optimal bids (Appendix A.1). The intuition is that as long as the highest rival bid p_j is less than bidder ℓ's optimal bid for auction j, i.e., $b^*_{\ell j}$, which is the only situation considered in Expression (8), $b^*_{\ell j}$ will remain invisible thus will have no impact on the distribution of p_j. As a result, a multi-unit bidder would only choose to participate in the K auctions that lead to the highest expected payoffs for her. In particular, the optimal bids $b^*_{\ell j}$ can be derived directly from Expression (8). The following proposition summarizes the discussion above regarding equilibrium bidding from multi-unit bidders.

Proposition 2. *For any auction j, the optimal bid of a multi-unit bidder under beliefs G and Ψ is given by:*

$$\sigma^{MU}_{BID,j}(x_\ell;\zeta_{\ell j},\omega_{\ell t}) = x_\ell + \zeta_{\ell j}. \tag{9}$$

For the auction selection decision $\sigma^{MU}_{SLT}(x_\ell;\zeta_{\ell\cdot},\omega_{\ell t})$, the bidder would participate in the K auctions that lead to the highest expected payoffs.

In short, as one would expect, it is a dominant strategy for multi-unit bidders to submit bids equal to their valuations, as winning an auction does not affect their future payoffs (we omit the proof as it follows directly from this observation). Finally, when determining whether to visit the platform, bidder ℓ compares the daily cost of participating $c_{\ell t}$ with the unconditional payoff $r^{MU}(x_\ell;G,\Psi)$ given as follows

$$r^{MU}(x_\ell;G,\Psi) = \int_\omega \int_{\zeta_{\ell\cdot}} u^{MU}(x_\ell;\zeta_{\ell\cdot},\omega_{\ell t},G) d\mathbf{F}_\zeta(\zeta_{\ell\cdot}) d\Psi(\omega_{\ell t}). \tag{10}$$

Thus, the decision of a multi-unit bidder ℓ to visit the platform on a given day t takes the following simple form:

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

Unit-Demand Bidders: Unit-demand bidders are only interested in winning a single auction over their lifetime. Thus, their opportunity cost of winning today is the expected payoff of winning an auction in the future. In other words, unit-demand bidder ℓ has to take into account how her future payoff may be affected by her present bid. Thus, the conditional payoff of an UD bidder with valuation x_ℓ, given that she has already visited the platform and decided to submit a bid in auction j can be written as follows when the observed state of the market is $\omega_{\ell t}$:

$$u_j^{UD}(x_\ell;\zeta_{\ell j},\omega_{\ell t},G,\Psi)=\max_{b_{\ell j}}\int_{s_j}^{b_{\ell j}}(x_\ell+\zeta_{\ell j}-p_j)g_j(p_j|\omega_{\ell t})dp_j+\alpha^{UD}\big(1-G_j(b_{\ell j}|\omega_{\ell t})\big)v_f(x_\ell;G,\Psi). \tag{11}$$

The first term in Expression (11) is equal to the instantaneous payoff if the bidder wins auction j. The second term is equal to the bidder's (expected) payoff if she does not win the current auction (which occurs with probability $1-G_j(b_{\ell j}|\omega_{\ell t})$). Here, $v_f(x_\ell;G,\Psi)$ denotes the agent's continuation payoff, which is given by the following Bellman equation:

$$v_f(x_\ell;G,\Psi)=\int_0^\infty \max\Big\{v^{UD}(x_\ell;G,\Psi)-c,\alpha^{UD}v_f(x_\ell;G,\Psi)\Big\}\mu^{UD}e^{-\mu^{UD}c}dc, \tag{12}$$

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

and $j^*=\sigma_{SLT}^{UD}(x_\ell;\zeta_{\ell\cdot},\omega_{\ell t})$ denoting the auction it is optimal to bid in. Then, we can write:

$$\sigma_{SLT}^{UD}(x_\ell;\zeta_{\ell\cdot},\omega_{\ell t})=\arg\max_j\big\{u_j^{UD}(x_\ell;\zeta_{\ell j},\omega_{\ell t},G,\Psi)\big\}. \tag{13}$$

The following proposition summarizes the optimal bidding behavior for unit-demand bidders.

PROPOSITION 3. *The optimal bidding for a unit-demand bidder in auction j under beliefs G and Ψ is given by*

$$\sigma_{BID,j}^{UD}(x_\ell;\zeta_{\ell j},\omega)=x_\ell+\zeta_{\ell j}-\alpha^{UD}v_f(x_\ell;G,\Psi). \tag{14}$$

The bidder's auction selection decision $\sigma_{SLT}^{UD}(x_\ell;\zeta_\ell,\omega)$ is given by Expression (13).

The proof can be found in Appendix A.3. As expected, unit-demand bidders shade their bids by their continuation values. As they are forward-looking and 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.

Given that the payoff of entering the market on a day t is $v^{UD}(x_\ell;G,\Psi)-c_{\ell t}$ and the payoff of waiting is $\alpha^{UD}v_f(x_\ell;G,\Psi)$, we have

$$\sigma_{PRT}^{UD}(x_\ell;c_{\ell t})=\begin{cases}1, & \text{i.e., visit,} \quad \text{if } r^{UD}(x_\ell;G,\Psi)\triangleq v^{UD}(x_\ell;G,\Psi)-\alpha^{UD}v_f(x_\ell;G,\Psi)\ge c_{\ell t}\\ 0, & \text{i.e., wait,} \quad \text{otherwise}\end{cases}, \tag{15}$$

where recall that $c_{\ell t}$ is drawn from the exponential distribution with rate μ^{UD}.

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 to November 2014, i.e., subsequent to the market's pre-treatment period but prior to its post-treatment period. 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 that we then compare against the estimate we obtained from our average treatment effect analysis in Section 3.2; and *(ii)* reserve the post-treatment period's data for use in our counterfactual analysis regarding market design.

As common in the auctions literature (Hendricks and Sorensen (2016), Backus and Lewis (2016)), we first normalize the observed bids to adjust for heterogeneity in the products' features and condition for time fixed effects. We further shift the normalized bids to ensure that they fall within the support of the Weibull distribution. Post-normalization, bidders' valuations are treated as drawn from a common distribution throughout the observation period.

The structural estimation follows two steps: (1) first, a nonparametric estimation of the platform's steady-state distribution of auctions and bids, followed by (2) maximum simulated likelihood (MSL) estimation of our modeling primitives for bidders' valuation distributions, participation costs, and retention rates.

5.1. Platform's Steady State

To decide whether to visit the platform on a given day, a potential bidder compares her participation cost against her expected payoff. Given that the bidder cannot observe the realized state of the platform, i.e., the number of auctions and their standing bids, before visiting and incurring the participation cost, her decision is based on the beliefs she forms about the state. Accordingly, first we estimate the platform's equilibrium steady state as we describe below.

The steady-state belief of 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 following form:

$$\Psi(\omega_{\ell t}) = P_{MKT}(n_t)\Psi_n(\mathbf{s}_{\ell t}),$$

where P_{MKT} is directly estimated by the empirical distribution of auctions ending on a given day in the data. On the other hand, the PDF of standing bids $\Psi_n(\mathbf{s}_{\ell t})$ conditional on n is estimated using a non-parametric approach (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 probability density function 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 the misspecification bias. Lastly, we let the MU bidders' demand upper bound K to 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.

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 bids[6] and highest rival bids (final prices) respectively, when the market thickness is n auctions. The corresponding sample size is denoted as N_n. The kernel density estimator of the standing bids conditional on market thickness n can be written as

$$\hat{\psi}_n(\mathbf{s}_n) = \frac{1}{N_n}\sum_{\ell t} K_{h_s}(\mathbf{s}_n - \mathbf{s}_{\ell t,n}),$$

where $K_{h_s}(\mathbf{s})$ is the multivariate Gaussian Kernel density with identical bandwidth h_s, i.e.,

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

The selection of the optimal bandwidth h_s^* is conducted by minimizing the cross-validation (CV) estimator of the risk function (for more details, refer to Wasserman (2010), Li and Racine (2007)). 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

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

where $y^{(j)}_{\ell t,n}$ 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 (similar to the approach used in Hansen (2004)).[7]

5.2. Bidders' Primitives

Next, we obtain two-step maximum simulated likelihood (MSL) estimates of our modeling primitives for the bidders' valuation distributions, participation costs, and retention rates. To re-iterate, a bidder's endowed valuation for the product is drawn from her type's Weibull distribution, either F^{MU} or F^{UD}, characterized by the scale parameters, λ^{MU} and λ^{UD}, and the shape parameters, γ^{MU} and γ^{UD} respectively. A bidder's daily participation cost is exponentially distributed with rate

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

[7] 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 B.1.

μ^{MU} or μ^{UD} depending on the bidder's type. Similarly, the daily retention rates, α^{MU} and α^{UD}, represent the bidders' type-specific probabilities of remaining in the bidder pool for another day.

Bidder ℓ's bidding history $\mathbf{X}_\ell$ consists of two parts. The first part is her bidding sequence, $\mathbf{B}_\ell = [B_\ell^{t_\ell}, ..., B_\ell^{t_\ell + l_\ell}]$, 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 that she stays in the bidder pool, we specify the likelihood of each of her observed behaviors. The following formulas apply to both types of bidders, 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 $P^V(X_\ell|\mu)$) and *(ii)* places a bid (which has likelihood $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \nu)$). Thus, the overall probability of her placing a bid is given by:

$$P^V(X_\ell|\mu) L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \nu).$$

On the other hand, if no bid is observed on that day, there are two possibilities: *(i)* she doesn't visit the platform at all (which occurs with probability $1 - P^V(x_\ell|\mu)$), or *(ii)* she visits the platform but she finds it optimal not to place a bid (which occurs with probability $P^V(x_\ell, \mu) L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \nu)$). Thus, the corresponding likelihood of not observing a bid is given as

$$1 - P^V(x_\ell|\mu) + P^V(x_\ell, \mu) L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \nu).$$

Expressions for $P^V(X_\ell|\mu)$, $L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \nu)$, and $L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \nu)$ for both bidder types can be found in Appendix B.2.

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 days[8] after her last bid (and then we take the average over the likelihood of all possible exit days).

In summary, the likelihood of a bidder's entire bidding history conditional on her endowed valuation x_ℓ is given as:

$$\mathcal{L}_\ell(\mathbf{X}_\ell | x_\ell, \alpha, \mu, \nu) =$$

$$\Big(\alpha^{l_\ell} \prod_{t=t_\ell}^{t_\ell + l_\ell} \big(\underbrace{P^V(x_\ell, \mu) L_{\ell t}^B(\mathbf{b}_{\ell t}, \mathbf{S}_{\ell t}|x_\ell, \nu)}_{\text{Placing the bid(s) } \mathbf{b}_{\ell t} \text{ on day } t}\big)^{B_\ell^{t_\ell}} \big(\underbrace{1 - P^V(x_\ell, \mu) + P^V(x_\ell, \mu) L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell, \nu)}_{\text{Not placing a bid on day } t}\big)^{1 - B_\ell^{t_\ell}}\Big)$$

[8] Our estimates are obtained under $E = 14$ days. We have tested other values $E = 21$ and 28 days and found that the resulting estimates do not differ significantly.

$$\cdot \underbrace{\Big(1+\sum_{t'=1}^{E-1}\alpha^{t'}\prod_{t=t_\ell+l_\ell+1}^{t_\ell+l_\ell+t'}\big(1-P^V(x_\ell,\mu)+P^V(x_\ell,\mu)L_{\ell t}^{NB}(\mathbf{S}_{\ell t}|x_\ell,\nu)\big)\Big)(1-\alpha)}_{\text{Exiting bidder pool within } E=14 \text{ days since the last bidding day}}. \quad (16)$$

As the bidder's valuation x_ℓ is not observed, we need to consider the unconditional likelihood function,

$$\mathcal{L}_\ell(\mathbf{X}_\ell|\theta) = \int_{x_\ell} \mathcal{L}_\ell(\mathbf{X}_\ell|x_\ell,\alpha,\mu,\nu) f(x_\ell|\lambda,\gamma)dx_\ell,$$

where $f(x_\ell|\lambda,\gamma)$ denotes the PDF of the Weibull valuation distribution and $\theta=[\alpha,\mu,\nu,\lambda,\gamma]$ represents the vector of primitives associated with a bidder. 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 $\hat{\mathcal{L}}(\mathbf{X}_\ell|\theta)$ by employing Monte Carlo integration.[9] Then, the MSL estimate $\hat{\theta}_{MSL}$ is obtained by maximizing the log of the simulated likelihood of all bidders, i.e.,

$$\hat{\theta}_{MSL} = \arg\max_\theta \sum_\ell log\big(\hat{\mathcal{L}}_\ell(\mathbf{X}_\ell|\theta)\big).$$

5.3. Estimation Results

In Table 6, we report our estimates for each bidder type, including its valuation distribution, its average daily cost to participate, and its daily retention probability.[10]

These estimates highlight a number of differences between MU and UD bidders. In particular, UD bidders possess on average higher valuations than MU bidders but exhibit substantially higher variability in valuations within-type across individual bidders. On average, the UD type's valuation is higher by \$14.2 per unit, or \$1,667.1 per auction (the average auction lot size in Market A is 117.2 units), which is both statistically and economically significant. In contrast, the idiosyncratic component of a bidder's valuation for an auction lot, capturing how the lot's characteristics affect its fit to a bidder's resale channels, is statistically indistinguishable across the types. Despite their higher valuations on average, UD bidders shade their bids (as noted in Proposition 3), which contributes to MU bidders winning the majority of auctions.

Per daily platform visit, a MU bidder incurs a substantially higher average participation cost of \$98.28 compared with \$24.57 for a UD bidder. Nonetheless, she is still willing to visit the platform because she typically places bids for multiple auction lots. Evidence suggests that the MU bidder's higher cost to visit on a day is associated with having to review, compare, and match downstream

[9] For additional discussion on $\hat{\mathcal{L}}_\ell(\mathbf{X}_\ell|\theta)$, refer to Appendix B.2.

[10] To test the goodness of fit, we simulate bids from the estimated valuation distribution and compare them to the observed bids in the training data (both at a normalized scale). We simulate 6,000 bids from the estimated model and training data, separately. The average of the simulated bids is \$155.77 and their standard deviation is \$18.56 whereas the average of the observed bids is \$155.49 and their standard deviation of \$19.78. The difference between the two averages is not statically significant.

Table 6 Maximum Simulated Likelihood Estimates for the Structural Model

	Multi-Unit Bidders	Unit-Demand Bidders
Valuation *per-auction*		
Mean	\$14,757.6	\$16,424.7
	(\$220.6)	(\$367.4)
Standard Deviation	\$1,925.3	\$2,794.1
	(\$299.3)	(\$344.0)
Idiosyncratic Error (SD)	\$3,240.3	\$3,205.1
	(\$68.0)	(\$130.3)
Avg. Daily Participation Cost	\$98.28	\$24.57
	(\$2.93)	(\$1.52)
Retention Rate	0.989	0.926
	(0.000)	(0.005)
Number of Bidders	113	174

Note: standard deviations of the estimates are in the parenthesis. Valuation distributions are at normalized scale.

channels for multiple auctions: on average, a MU bidder spends 34.2 minutes on the platform per day, compared with 7.2 minutes for a UD bidder.[11] As a side remark, note that the estimates we obtain for the bidders' participation costs may serve as a guideline for the platform to design incentive schemes, e.g., in the form of discount or free-shipping coupons, as a way to boost the bidders' participation rates.

Lastly, MU bidders are more loyal to the platform than UD bidders. An average MU bidder stays much longer (90.9 days) in the bidder pool compared with an average UD bidder who has not yet won an auction (13.5 days). This indicates that MU bidders tend to visit the platform more regularly compared with UD bidders who tend to have a short interaction with the platform, despite potentially not being able to satisfy their demand.

6. Implications for Platform Design

By connecting a platform's listing policy to its bidders' participation rates and bidding behavior, our findings bear a number of implications for platform design. Using our estimates derived in Section 5, we evaluate several listing policies and relate their performance to the induced market thickness. As a by-product of this analysis, we confirm that the revenue increase of switching from

[11] Re-estimating our model on the pre-treatment data featuring thinner markets indicates that these costs are relatively stable across the range of market thickness we consider. The estimates are \$100.69 for MU bidders and \$24.32 for UD bidders, which are in line with the results in Table 6.

the uniform to the batch listing policy is of the order of 6–7%, validating our structural counterfactual projections out-of-sample against the average treatment effect estimate, 6.5%, derived in Section 3.2.

Furthermore, motivated by the recent efforts of our corporate sponsor, we explore the design and performance of a *recommendation system*, whose main purpose is to mitigate auction cannibalization on days when the stochastic supply of incoming auctions is significantly higher than average, leading to high market thickness. In particular, the system is designed to notify (a subset of) randomly selected MU bidders and inform them of the large number of auctions listed on the platform with the goal of increasing their participation rates. A challenge for the design of such a system is the possible unintended effect a recommendation has on bidders' endogenous entry behavior: informing a bidder of high market thickness may not only make participation more attractive for herself, but also for other informed bidders, thereby (perhaps) dissipate the attractiveness of participation. Thus, assessing the revenue impact of recommendations on the platform involves incorporating them in the bidders' equilibrium belief-formation process.

6.1. Supply Levels and Listing Policies

A platform's listing policy influences revenues by manipulating its market thickness. Separate strands in the existing literature contemplate dual but countervailing effects from market thickness/availability on the behavior 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 (Dana and Petruzzi (2001), 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 conducting counterfactuals, we evaluate the performance of four listing policies that differ in how auction end times are distributed throughout the week under seven different levels of incoming supply (which together induce a market thickness level). In particular, we simulate the performance of the following four listing policies:

(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 increases by switching from policy (i) through (iv). We consider seven supply levels, which we denote by "$\frac{1}{3}$x", "$\frac{1}{2}$x", "1x (baseline)", "2x", "3x", "4x", and "5x". In particular, we first derive our "baseline" supply case by fitting Market A's supply data to a gamma distribution, in order to appropriately account for the large 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 holding the coefficient of variation the same. For our entire counterfactual analysis, we hold 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.

Table 7 summarizes our simulation results. For each combination pairing a supply level with a listing policy, we present the average platform revenue generated as a percentage share of the average revenue generated by the batch listing policy at that supply level. Therefore, each row of Table 7 communicates 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.

Notably, as shown at the baseline supply level, our counterfactual revenue improvement of switching Market A from uniform to batch listing is roughly 6.4%, offering out-of-sample validation of the 6.5% increase estimated in our average treatment effect analysis of Section 3.2 using time-wise separate data. At the baseline supply, the optimal policy turns out to be the single-day batch policy, which yields 2.8% more revenue than the batch policy. These relative differences in performance translate into substantial revenue gains. Over the ten months of the post-treatment period, Market A's revenue from iPhone 4, 4s, and 5 auctions amounted to $\$4,608,941$: the 8.9% relative difference between the uniform and the best-performing single-day policy translates to $\$410,196$ of additional revenue for the platform. Perhaps reflecting our projections, as of January 2017, Market A enforced a single-day listing policy that ends all its auctions exclusively on Tuesdays.

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 revenue 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. In contrast, for high levels of supply, e.g., "4x" and "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.

Table 7 Revenue corresponding to each listing policy as a percentage of the revenue under the batch policy

Supply Level	Uniform	Three-Day Batch	Batch	Single-Day Batch
$\frac{1}{3}$x	93.0% [0.6]	96.1% [1.1]	100% [1.6]	**105.7%** [3.2]
$\frac{1}{2}$x	93.5% [0.9]	95.9% [1.5]	100% [2.2]	**104.4%** [4.4]
1x (baseline)	93.9% [2.1]	96.4% [3.4]	100% [5.2]	**102.8%** [10.3]
	($\$-283,241$)	($\$-166,440$)		($\$127,622$)
2x	94.8% [4.2]	97.7% [6.9]	**100%** [10.5]	**99.9%** [20.6]
3x	98.0% [6.3]	**100.4%** [10.3]	**100%** [15.7]	96.0% [31.0]
4x	102.0% [8.3]	**103.1%** [13.7]	100% [20.8]	93.8% [41.6]
5x	**126.8%** [10.1]	121.5% [16.7]	100% [25.4]	76.6% [50.2]

Note. Values in bolded correspond to the row optimal listing policies with significant level 0.001*. Market thickness averages are displayed in the brackets.*

Figure 6 Final price as a function of market thickness

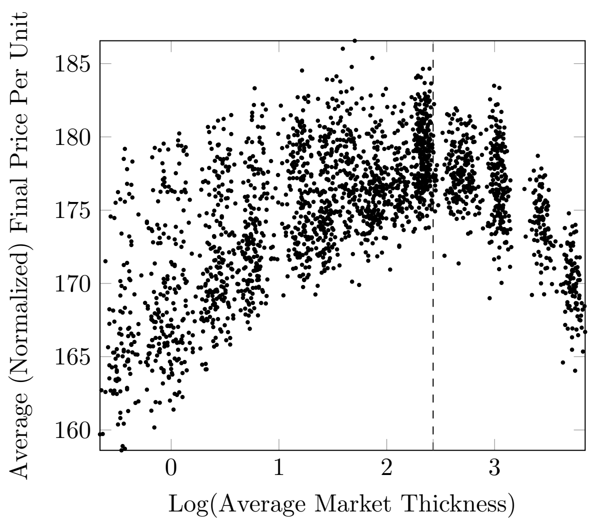

Lastly, we relate the platform's revenue outcomes to the market thickness induced by its listing policy. Intuitively, platform revenues are maximized when the platform's market thickness hews close to a "sweet spot" that effectively balances its trade-off of overcoming bidder participation costs (thus, inducing higher participation rates) against competition as substitutes among auctions on the supply side. Using simulated data from our counterfactual simulations described above, Figure 6 depicts an auction's average final per-unit revenue as a function of the market's (log) average market thickness. Interestingly, low market thickness does not lead to high average final price primarily as the result of participation costs. 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.

To further demonstrate that a market's per-unit revenue meaningfully relates to its induced market thickness, we estimate the following regression model which directly controls for the revenue effects of the supply level and the listing policy:

$$AvgFinalPrice = \beta_0 + \beta_1 * LogAvgMkt + \beta_2 * LogAvgMkt^2 + FE_S + FE_L + \epsilon, \tag{17}$$

where we include fixed effects for the supply level FE_S and listing policy FE_L. The dependent variable $AvgFinalPrice$ is the average per-unit final price, and $LogAvgMkt$ denotes the log of the average market thickness. Table 8's regression results illustrate the substantial quadratic relationship between market thickness and auction revenue, even after controlling for the fixed effects of supply and listing policy.

Table 8 Regression of average final price on average market thickness

	Dependent Variable:	
	$AvgFinalPrice$	
	Model (1)	Model (2)
$LogAvgMkt(\hat{\beta}_1)$	7.784^{***}	9.452^{***}
	(0.177)	(1.037)
$LogAvgMkt^2(\hat{\beta}_2)$	-1.874^{***}	-1.947^{***}
	(0.051)	(0.059)
FE_S		Y
FE_L		Y
Constant	169.556^{***}	169.793^{***}
	(0.141)	(0.520)
Observations	2,400	2,400
R^2	0.471	0.503
Adjusted R^2	0.470	0.501
Residual Std. Error	3.368 (df = 2397)	3.269 (df = 2389)
F Statistic	$1{,}065.059^{***}$ (df = 2; 2397)	241.666^{***} (df = 10; 2389)
Note:		*p<0.1; **p<0.05; ***p<0.01

6.2. Implementing Recommendations

In this section, we discuss the design of recommendation systems for a competitive online marketplace like the one studied in the paper. As a complement to the listing policy, recommendations intend to mitigate cannibalization on days of high market thickness by incentivizing buyers to visit the platform.

Normally, bidders obtain information about the daily supply of ending auctions only by visiting the platform, thereby incurring sizable participation costs to obtain this information. While a listing policy can be used to appropriately adjust the overall market thickness and level of bidder participation expected in steady state, the supply process remains highly variable. Thus, on some days, the market's inventory of available auctions may be significantly higher than expected. Indeed, the platform may then benefit from communicating credible recommendations to (a subset of) bidders to incentivize visits to the platform and, this way, maintain a "sweet spot" in market thickness. Interestingly and in contrast to a typical recommendation system, the value of such solicitations for bidders is endogenously determined by how they affect the bidders' equilibrium behavior, which, in turn, depends on when the platform sends them and what fraction of the bidder pool receives them. Given the fact that most of the auctions are won by multi-unit bidders (Table 2), we study how they respond to messages indicating high market thickness.

In particular, we consider recommendation systems that are defined by two parameters: a threshold κ on the number of ending auctions above which the platform sends a "recommendation" to a randomly selected fraction η of the multi-unit bidder pool. We assume that κ and η are known to all bidders. The recommendation is sent out at the beginning of those qualified days. The recommendation message provides information regarding the system's high market thickness state, i.e., the message indicates that the number of auctions ending on the day is above κ; thus, the bidders who receive it update their (steady-state) beliefs about the market state and determine whether to visit the platform. Furthermore, bidders who do not receive the message also update their steady-state beliefs and behave accordingly.

Considering Market A's baseline supply process and its identified optimal single-day listing policy, we simulate counterfactuals to derive the platform's optimal recommendation system across the equilibrium outcomes under different combinations of the parameters κ and η. In particular, we find that the optimal system sends recommendations to 10% of the MU bidders when the number of ending auctions on a given day is higher than 21. Recommendations are sent out roughly 9.97% of the days in the entire simulation horizon. Implementing this simple recommendation system further improves the platform's revenues by 2.9% on these days.

For additional intuition, Figure 7 plots multi-unit bidders' equilibrium rates of visiting the platform in the presence and absence of (optimal) recommendations. Several related observations follow. First, a bidder who did not receive a recommendation in the presence of a recommendation system (red curve) is slightly less likely to visit the platform than in the absence of a recommendation system entirely (black curve). Such a bidder infers that the supply of auctions may be low, which, in turn, weakens her incentive to participate. To guard against the strength of this effect, an optimal recommendation system appropriately calibrates its threshold κ, i.e., the frequency at which it sends out recommendations, and the share of recipients η.

Figure 7 **Impact of recommendations on the equilibrium rate of visiting the platform under the single-day batch policy**

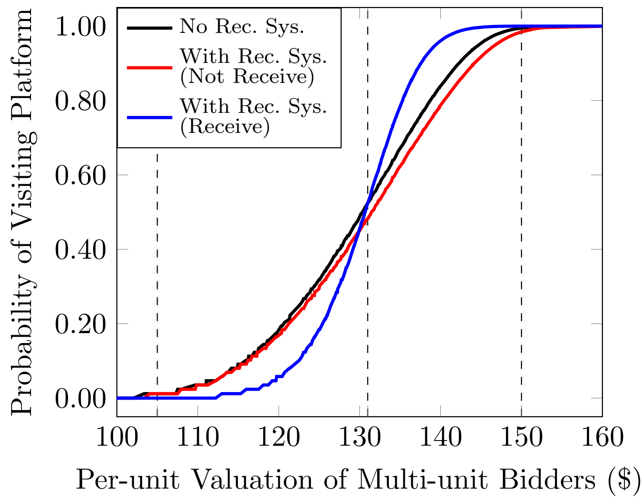

More interestingly, bidders' responses to receiving a recommendation (blue curve) vary depending on their valuations. For lower valuation bidders (valuation less than $131), receiving a recommendation is actually bad news, which substantially decreases their probability of visiting. In contrast, bidders with higher valuations (valuation higher than $131) become significantly more likely to visit the platform after receiving a recommendation. While the supply of auctions is higher on days on which recommendations are sent out, high valuation bidders become more likely to visit; thus, they intensify the competition by driving prices up and deter low valuation bidders. Lastly, the recommendation has a relatively smaller impact on bidders with very high valuations (higher than $150), who almost always visit the platform. Thus, the optimal recommendation system increases revenues mainly by altering the composition of bidders who actively participate in auctions in the market's high-inventory states. Thus, the optimal delivery of information is designed to target the participation of bidders through self-selection.

7. Concluding Remarks

In this paper, we empirically illustrate the impact of a platform's market thickness on the participation rates of its potential participants and consequently its revenues. In particular, we highlight that a seemingly innocuous market design choice, i.e., the platform's listing policy, affects revenues significantly by managing the market thickness induced on the marketplace. In addition, we explore the design of a recommendation system that selectively informs participants regarding the market state and establish that it can benefit the B2B auction platform by mitigating cannibalization among substitutable auctions.

Using a proprietary dataset we obtained from a leading online B2B auction platform, we estimate that inducing higher market thickness (by concentrating the auctions' ending times on certain days in the week) leads to a 6.5% increase in the platform's revenues. Additional analysis points to the presence of a significant participation cost associated with visiting the platform and bidding on open auctions that adversely affect bidders' participation rates. Motivated by the results of our descriptive analysis, we develop and estimate a structural model, which endogenizes bidders' decision-making including whether and when to visit the platform, which auction to participate, and how 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 (non-parametric) analysis.

Complementary to illustrating the revenue impact of the platform's listing policy and optimizing its design as a function of the incoming supply level, we discuss the implementation of a recommendation system that selectively informs market participants of the platform's state. Appropriately designing the system yields an additional revenue increase by successfully increasing the level of competition among bidders on days that the incoming supply is higher than average.

Furthermore, in addition to the counterfactual analysis, our structural estimates may find use in assessing a bidder's lifetime value for the platform, i.e., the financial value of the entire relationship between the bidder and the platform (Jain and Singh (2002)), and accordingly allow the platform to optimize the allocation of its marketing resources as a means to retain and attract bidders.

More broadly, our work highlights that systematic 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 addressing similar questions relating the design of an online platform with the market thickness it induces (and, consequently, its impact on revenues and welfare) in the context of other two-sided markets is a very fruitful direction for future research.

References

Alexandrov, Alexei, Martin A Lariviere. 2012. Are reservations recommended? *Manufacturing & Service Operations Management* **14**(2) 218–230.

Allon, Gad, Achal Bassamboo. 2011. Buying from the babbling retailer? The impact of availability information on customer behavior. *Management Science* **57**(4) 713–726.

Arnosti, Nick, Ramesh Johari, Yashodhan Kanoria. 2014. Managing congestion in dynamic matching markets. *ACM Conference on Economics and 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 Journal of Economics* 329–355.

Balseiro, Santiago R, Omar Besbes, Gabriel Y Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* **61**(4) 864–884.

Bang, Heejung, James M Robins. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**(4) 962–973.

Bapna, Ravi, Paulo Goes, Alok Gupta, Yiwei Jin. 2004. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS Quarterly* 21–43.

Bimpikis, Kostas, Ozan Candogan, Daniela Saban. 2017. Spatial pricing in ride-sharing networks. *Working paper* .

Bodoh-Creed, Aaron, Joern Boehnke, Brent Richard Hickman. 2016. How efficient are decentralized auction platforms? *Working paper* .

Bray, Robert L, Haim Mendelson. 2015. Production smoothing and the bullwhip effect. *Manufacturing & Service Operations 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. *Forthcoming in M&SOM* .

Chan, Tat Y, Vrinda Kadiyali, Young-Hoon Park. 2007. Willingness to pay and competition in online auctions. *Journal of Marketing Research* **44**(2) 324–333.

Cullen, Zoë, Chiara Farronato. 2016. Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. *Working paper* .

Dana, James D, Nicholas C Petruzzi. 2001. Note: The newsvendor model with endogenous demand. *Management Science* **47**(11) 1488–1497.

Ferdows, Kasra, Michael A Lewis, Jose AD Machuca. 2005. Zaras secret for fast fashion. *Harvard Business Review* **82**(11) 98–111.

Fradkin, Andrey. 2017. Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb. *Working paper* .

Gallino, Santiago, Antonio Moreno. 2014. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science* **60**(6) 1434–1451.

Gan, Li, Qi Li. 2016. Efficiency of thin and thick markets. *Journal of Econometrics* **192**(1) 40–54.

Golrezaei, Negin, Hamid Nazerzadeh, Vahab Mirrokni. 2017. Boosted second price auctions for heterogeneous bidders. *Working paper* .

Halaburda, Hanna, Mikołaj Jan Piskorski, Pınar Yıldırım. 2017. Competing by restricting choice: The case of matching platforms. *Management Science* .

Hansen, Bruce E. 2004. Nonparametric conditional density estimation. *Unpublished manuscript* .

He, Hua, Pan Wu, Ding-Geng Din Chen. 2016. Statistical causal inferences and their applications in public health research. *ICSA Book Series in Statistics* .

Hendricks, Kenneth, Alan Sorensen. 2016. The value of intermediaries in dynamic auction markets. *Working paper* .

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 research 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* .

Imbens, Guido W. 2004. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics* **86**(1) 4–29.

Iyer, Krishnamurthy, Ramesh Johari, Mukund Sundararajan. 2014. Mean field equilibria of dynamic auctions with learning. *Management Science* **60**(12) 2949–2970.

Jain, Dipak, Siddhartha S Singh. 2002. Customer lifetime value research in marketing: A review and future directions. *Journal of interactive marketing* **16**(2) 34–46.

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* .

Kanoria, Yash, Daniela Saban. 2017. Facilitating the search for partners on matching platforms: Restricting agent actions. *Working paper* .

Kim, Sang Won, Marcelo Olivares, Gabriel Y Weintraub. 2014. Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science* **60**(5) 1180–1201.

Li, Jun, Nelson Granados, Serguei Netessine. 2014. Are consumers strategic? structural estimation from the air-travel industry. *Management Science* **60**(9) 2114–2137.

Li, Jun, Serguei Netessine. 2017. Market thickness and matching (in)efficiency: Evidence from a quasi-experiment. *Working Paper* .

Li, Qi, Jeffrey Scott Racine. 2007. *Nonparametric econometrics: theory and practice*. Princeton University Press.

Moon, Ken, Kostas Bimpikis, Haim Mendelson. 2017. Randomized markdowns and online monitoring. *Forthcoming in Management Science* .

Olivares, Marcelo, Christian Terwiesch, Lydia Cassorla. 2008. Structural estimation of the newsvendor model: an application to reserving operating room time. *Management Science* **54**(1) 41–55.

Petruzzi, Nicholas C, Kwan E Wee, Maqbool Dada. 2009. The newsvendor model with consumer search costs. *Production and Operations 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 Science* **63**(5) 1493–1518.

Rosenbaum, Paul R, Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 41–55.

Sailer, Katharina. 2006. Searching the ebay marketplace. *Working paper* .

Su, Xuanming, Fuqiang Zhang. 2009. On the value of commitment and availability guarantees when selling to strategic consumers. *Management Science* **55**(5) 713–726.

Tadelis, Steven, Florian Zettelmeyer. 2015. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *The American Economic Review* **105**(2) 886–905.

Wasserman, Larry. 2010. *All of nonparametric statistics*. Springer Publishing Company, Incorporated.

Weintraub, Gabriel Y, C Lanier Benkard, Benjamin Van Roy. 2008. Markov perfect industry dynamics with many firms. *Econometrica* **76**(6) 1375–1411.

Yang, Pu, Iyer Krishnamurthy, Peter Frazier. 2016. Mean field equilibria for competitive exploration in resource sharing settings. *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 177–187.

Appendix A: Proofs

A.1. Independence between $\mathbf{b}_\ell$ and $g_j(p_j|\omega_{\ell t})$

In this subsection, we show that Expressions (8) and (11) have accounted for the effects of $\mathbf{b}_\ell$ 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 we use w_j to denote the current highest bid in auction j, which is unobservable. The conditional PDF of w_j is denoted by $h_j(w_j|n_t, \mathbf{s}_{\ell t})$. We use p_j to denote the highest rival bid, i.e., the final price, in auction j and its conditional PDF is denoted by $k_j(p_j|n_t, \mathbf{s}_{\ell t})$. The expected payoff of a bidder that places bid $b_{\ell j}$ in auction j is given as follows:

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

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

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

Let

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

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 $\mathbf{b}_\ell$. ■

A.2. Proof of Proposition 1

We use the Schauder fixed point theorem to establish equilibrium existence. First, we establish that the mapping from the 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:

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

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 multi-unit and unit-demand bidders who visit the platform

(v) $\mathbf{x}_{MU}, \mathbf{x}_{UD}$: Vectors of valuations of multi-unit and unit-demand bidders in the bidder pool;

(vi) $\mathbf{b}(\mathbf{x}_{UD}; \mathbf{G})$: Bidding function of unit-demand 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 $\Gamma_{sub}: \mathcal{H} \to \mathcal{H}$, where $\mathcal{H}$ is the space of probability CDFs, with fixed $\mathbf{m}_{MU}$,$\mathbf{m}_{UD}$, o, and n. Specifically,

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

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

$$P^{MU}(\mathbf{m}_{MU};\mathbf{x}_{MU},\mathbf{G}) = \prod_{\ell^{MU}} P\big(c_{MU}\le r^{MU}(x_{\ell^{MU}};\mathbf{G})\big)^{\mathbb{1}(\ell^{MU}\in\mathbf{m}_{MU})} P\big(c_{MU} > r^{MU}(x_{\ell^{MU}};\mathbf{G})\big)^{\mathbb{1}(\ell^{MU}\notin\mathbf{m}_{MU})}, \text{and}$$

$$P^{UD}(\mathbf{m}_{UD};\mathbf{b},\mathbf{G}) = \prod_{\ell^{UD}} P\big(c_{UD}\le r^{UD}(b_{\ell^{UD}};\mathbf{G})\big)^{\mathbb{1}(\ell^{UD}\in\mathbf{m}_{UD})} P\big(c_{UD} > r^{UD}(b_{\ell^{UD}};\mathbf{G})\big)^{\mathbb{1}(\ell^{UD}\notin\mathbf{m}_{UD})},$$

where $r^{MU}(\cdot;\mathbf{G})$ and $r^{UD}(\cdot;\mathbf{G})$ denote the expected payoffs per market visit for multi-unit and unit-demand bidders, respectively. Furthermore, Propositions 2 and 3 imply that bidders simply choose to bid in the auction(s) with the lowest standing bids.[12]The expected payoffs per market visit for both types of bidders can be expressed as:

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

$$r^{UD}(b_{\ell^{UD}};\mathbf{G}) = \sum_{n=1}^{\bar{N}} p_n \int_0^{b_{\ell^{UD}}} G_n^1(y)dy,$$

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

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

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

By definition, the vector of functions that generate the final price in each auction can be denoted by $\mathbf{f}_h = (f_h^1,...,f_h^{\bar{N}})$ and the element f_h^j maps all bidders' auction selection 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 multi-unit and unit-demand bidders respectively, and we use $\mathbf{x}_{MU}^j$ and $\mathbf{x}_{UD}^j$ to denote the vectors of valuations of multi-unit and unit-demand bidders who choose auction j. We can write f_h^j as:

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

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' behaviors 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 will show the continuity and compactness of Γ_{sub}.

[12] We provide proofs for the propositions in the remainder of the appendix.

Continuity of Γ_{sub}. We show 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 $x_{\ell^{MU}} \in [0,B]$ (where recall that B is the upper bound of the valuations), we have:

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

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

Second, using the fact that the participation costs follow the exponential distribution, we show that $exp\big(-\mu_{MU} r^{MU}(\cdot;\mathbf{G})\big)$ is uniformly bounded and Lipschitz continuous in $r^{MU}(\cdot;\mathbf{G})$. Note that $0 \leq r^{MU}(x_{\ell^{MU}};\mathbf{G}) \leq \bar{N}B$. Besides, 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 $\exp\big(-\mu^{MU} r^{MU}(\cdot;\mathbf{G})\big)$ is uniformly bounded and Lipschitz continuous in $r^{MU}(\cdot;\mathbf{G})$. By the definition of $P^{MU}(\mathbf{m}_{MU};\cdot,\mathbf{G})$, which is equal to the product between $\exp\big(-\mu^{MU} r^{MU}(x_{\ell^{MU}};\mathbf{G})\big)$ and $1-\exp\big(-\mu^{MU} r^{MU}(x_{\ell^{MU}};\mathbf{G})\big)$ for a finite number of bidders, we thus 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 $\mathbf{G}$;

(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 that

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

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

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

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 unit-demand bidder with $x_{\ell^{UD}} \in [0,B]$ and for all $\mathbf{G},\mathbf{G}' \in \mathcal{H}$, according to Proposition 3 we have

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

where

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

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,

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

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:*

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

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

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

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

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

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

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,

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

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 multi-unit 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}. In this part, we show that Γ_{sub} is a compact mapping using the Arzela-Ascoli theorem. To apply the theorem, Γ_{sub} needs to satisfy:

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 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:

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

and 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 $\mathbf{m}_1, \mathbf{m}_2, n, o$:

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

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

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

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

$$\begin{aligned}
J_3 &= \int \mathbb{1}\Big(x_{MU}^{j,(2)} \le b(x_{UD}^{j,(1)},\mathbf{G}), x_{MU}^{j,(1)} \le b(x_{UD}^{j,(1)};\mathbf{G}), y_j < x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)};\mathbf{G}) \le y_j'\Big) P^o(d\mathbf{x}) \\
&\le \int \mathbb{1}\big(y_j < x_{MU}^{j,(1)} \wedge b(x_{UD}^{j,(1)};\mathbf{G}) \le y_j'\big) P^o(d\mathbf{x}) \\
&= \int \Big(\mathbb{1}(y_j < x_{MU}^{j,(1)} \le y_j') \cdot \mathbb{1}\big(x_{MU}^{j,(1)} \le b(x_{UD}^{j,(1)};\mathbf{G})\big) + \mathbb{1}\big(y_j < b(x_{UD}^{j,(1)};\mathbf{G}) \le y_j'\big) \cdot \mathbb{1}\big(b(x_{UD}^{j,(1)};\mathbf{G}) < x_{MU}^{j,(1)}\big) \Big) P^o(d\mathbf{x}) \\
&\le \int_{y_j}^{y_j'} P^o(dx_{MU}^{j,(1)}) + \int_{y_j < b(x_{UD}^{j,(1)};\mathbf{G}) \le y_j'} P^o(dx_{UD}^{j,(1)}).
\end{aligned} \tag{18}$$

For the first term in (18), we have:

$$\int_{y_j}^{y_j'} P^o(dx_{MU}^{j,(1)}) \le \Big| \frac{dF_{MU}^{j,(1)}}{dx_{MU}^{j,(1)}} \Big|_\infty \cdot |y_j' - y_j| \le M_{MU}^{j,(1)} \cdot |y_j' - y_j|,$$

where $M_{MU}^{j,(1)} \triangleq \sup_{\mathbf{G}} |f_{MU}^{j,(1)}|_\infty < \infty$. For the second term in (18), we define vectors $\mathbf{z}, \mathbf{z}'$ as the solutions to the following equations:

$$z_j = y_j + F(y_j, \mathbf{G}),$$
$$z_j' = y_j' + F(y_j', \mathbf{G}).$$

Recall that from Lemma 2:

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

Therefore,

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

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

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

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. ■

A.3. Proof of Proposition 3

The assumptions of the proposition, i.e., $\nu^{MU} = \nu^{UD} = 0$ and $K \ge \bar{N}$, 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 in which auction to participate) 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 in which auction to participate).

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:

$$G_1^0(w|\mathbf{s}) \ge G_2^0(w|\mathbf{s}) \ge ... \ge G_n^0(w|\mathbf{s}), \tag{19}$$

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 that the number of UD bidders is known and denoted by N. One can extend the results to the case when 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 2 auctions and all bidders are UD bidders.

LEMMA 4. *Suppose 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 in which auction to participate).*

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:

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

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

We use $X_1, X_2, ..., X_N$ to denote the random variables corresponding to the bids of the first N bidders. The corresponding order statistics are denoted by $X_{(1)}, X_{(2)}, ..., X_{(N)}$, 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 using 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 4^{th} highest bid can always be placed (since it is impossible to have both standing bids being 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:

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

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 which 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

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

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

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

■

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

LEMMA 5. *Suppose 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 in which auction to participate).*

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 Expression (19). Equivalently, in the following proof, we show that for any $i < j$, we have

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

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

In the table, $\left\{\mathbf{X}_{-(l,k,L,K)} \in (\mathbf{s}_{-i,-j}, \mathbf{W}_{-i,-j})\right\}$ denotes the set of all valid assignments of the rest $X_{(\ell)}$ $(1 \le \ell \le 2n$ and $\ell \ne l,k,L,K)$ 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 $\mathbf{W}_{-i,-j}$.

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

$$\Delta_1(a_1) \triangleq G_i^0(w|s_i,s_j,a_1) - G_j^0(w|s_i,s_j,a_1)$$
$$= P(X_{(L)} \le w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \ne i,j, m' \ne l,k)$$
$$- P(X_{(K)} \le w | X_{(l)} = s_i, X_{(k)} = s_j, X_{(m')} = s_m, m \ne i,j, m' \ne l,k).$$

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, and, thus, show that the new bidding sequence belongs to a_2 and leads to the same standing vector $\mathbf{s}$ as the original one in a_1. In particular, the following is an illustration of a bidding sequence in a_1:

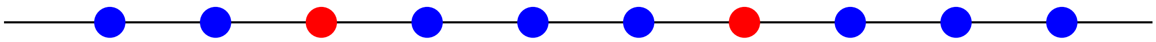

Each point stands for a bid and the two red points represent bids $X_{(L)}$ and $X_{(K)}$. At the first red point, a bidder's auction choice is the same regardless whether her bid is $X_{(L)}$ or $X_{(K)}$. Suppose the corresponding auction is auction i, given the definition of a_1. After the switch, at the first red point, $X_{(K)}$ will be placed in auction i. Note that in a_1, the highest losing bid in that auction at the first red point 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 point will not be affected after the switch. Similarly, at the second red point, the bidder with bid $X_{(L)}$ will choose to bid in auction j (the same decision at this point 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 point 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 points. Moreover, these two sequences lead to the same vector of standing bids.

Given a_2 in Case 2, we have:

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

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 $\Delta_3(a_3) \geq 0$ as conditioning on a_3 we have:

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

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 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 in which auction to participate).*

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 (19).

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 (19) 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 as 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 $G_i^0\big(w|\mathbf{s}, \text{Case } 2(i)\big) = G_j^0\big(w|\mathbf{s}, \text{Case } 2(i)\big)$. In sub-case *(ii)*, we have $X_{(l)} \geq M$, otherwise, the MU bidder is also winning auction j. Thus, $G_i^0\big(w|\mathbf{s}, \text{Case } 2(ii)\big) \geq G_j^0\big(w|\mathbf{s}, \text{Case } 2(ii)\big)$, because $\big\{X_{(l)} \leq w, \mathbf{s}, \text{Case } 2(ii)\big\} \subseteq \big\{M \leq w, \mathbf{s}, \text{Case } 2(ii)\big\}$. In sub-case *(iii)*, we can again apply the argument in the proof of the Lemma 5 to show that $G_i^0\big(w|\mathbf{s}, \text{Case } 2(iii)\big) \geq G_j^0\big(w|\mathbf{s}, \text{Case } 2(iii)\big)$. In summary, we have shown that with the presence of MU bidders, Expression (19) 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 her 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 (who we call bidder A) that places the last bid on a given day. 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,

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

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

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

In auction i, $\hat{g}_i^N(p_i|\mathbf{s})$ and $\hat{G}_i^N(p_i|\mathbf{s})$ 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, say bidder B, is denoted by b_1.[13] 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.

(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 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 only have N future competitors and, thus, the state would be updated from $(s_1, s_2, ..., s_n)$ with $N+1$ incoming bids to $(b_1, s_2, ..., s_n)$ with N incoming bids. Her conditional expected payoff is then:

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

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

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

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

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

(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 then be written as

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

[13] For brevity, we omit the discussion of the case that the next bidder is a MU bidder, since the claim follows in a similarly fashion of classification of the information set.

Similarly, if bidder A places a bid in auction 2, her conditional expected payoff can be written as:

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

Note that $W_1 \leq s_2 \leq ... \leq s_n$ 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:

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

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., $min(b_1, W_1)$. Her conditional expected payoff of submitting a bid in auction 2 is then:

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

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 $v_1^N\big(b,x;s_2,\min(b_1,W_2),\mathbf{s}_{-1,-2}\big)$ and $v_2^N\big(b,x;\min(b_1,W_1),s_2,\mathbf{s}_{-1,-2}\big)$ are symmetric in terms of auctions 1 and 2. Thus,

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

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. ■

Appendix B: Structural Estimation

B.1. Bandwidth Selection in Kernel Density Estimation

In this subsection, we describe the bandwidth selection process we use when we estimate non-parametrically 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:

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

where $X \sim f$. As f is unknown, the risk cannot be evaluated. Therefore, we consider the cross-validation (CV) estimator of risk (Wasserman (2010)), which is defined as follows:

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

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, following Li and Racine (2007), the CV estimator of risk can be further specified as

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

where N denotes the total number of observations (here, we use ℓt and $\ell' t'$ as observation indices). Furthermore, the bandwidth matrix is given as $H_s = h_s I_n$ 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:

$$\begin{aligned} CV^{(b)}(h_{y,1}, h_{y,2}) = \frac{1}{N}\sum_{\ell t} & \frac{\sum_{\ell' t' \neq \ell t}\sum_{\ell'' t'' \neq \ell t} K_{H_y}(\mathbf{s}^{(b)}_{\ell t} - \mathbf{s}^{(b)}_{\ell' t'}) K_{H_y}(\mathbf{s}^{(b)}_{\ell t} - \mathbf{s}^{(b)}_{\ell'' t''}) K_{\sqrt{2}h_{y,1}}(y^{(b)}_{\ell'' t''} - y^{(b)}_{\ell' t'})}{\big(\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}^{(b)}_{\ell t} - \mathbf{s}^{(b)}_{\ell' t'})\big)^2} \\ & - \frac{2}{N}\sum_{\ell t} \frac{\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}^{(b)}_{\ell t} - \mathbf{s}^{(b)}_{\ell' t'}) K_{h_{y,1}}(y^{(b)}_{\ell t} - y^{(b)}_{\ell' t'})}{\sum_{\ell' t' \neq \ell t} K_{H_y}(\mathbf{s}^{(b)}_{\ell t} - \mathbf{s}^{(b)}_{\ell' t'})}. \end{aligned}$$

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.

B.2. Simulated Maximum Likelihood Estimation

In this subsection, we provide expressions for $P^V(x_\ell,\mu)$, $L^B_{\ell t}(\mathbf{b}_{\ell t},\mathbf{S}_{\ell t}|x_\ell,\nu)$, and $L^{NB}_{\ell t}(\mathbf{S}_{\ell t}|x_\ell,\nu)$, and describe how we derive the simulated likelihood function $\hat{L}_\ell(\mathbf{X}_\ell|\theta)$ using importance sampling.

We use $P^V(x_\ell|\mu)$ to denote the probability that bidder ℓ visits the platform on day t, i.e.,

$$P^V(x_\ell,\mu) = P\big(c_{\ell t} \le r(x_\ell,G,\Psi)\big) = 1 - exp\big(-\mu r(x_\ell,G,\Psi)\big),$$

where recall that $r(x_\ell,G,\Psi)$ is the unconditional payoff per platform visit (see also Section 4).

Conditional on bidder ℓ visiting the platform on day t, we let $L^B_{\ell t}(\mathbf{b}_{\ell t},\mathbf{S}_{\ell t}|x_\ell,\nu)$ denote the likelihood of her placing bid(s) $\mathbf{b}_{\ell t}$ and we let $L^{NB}_{\ell t}(\mathbf{S}_{\ell t}|x_\ell,\nu)$ denote the likelihood of not placing any bid.

First, assume that ℓ is a multi-unit bidder (who could be interested in winning all open auctions on a day). 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., $b_{\ell j} = x_\ell + \zeta^{MU}_{\ell j} < s_{\ell j}$. Therefore,

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

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 $\mathcal{J}^t_\ell$ denotes the set of auctions that bidder ℓ chose to participate in on day t, the likelihood of placing these bids is given by $\prod_{j\in\mathcal{J}^t_\ell}\phi(b_{\ell j}|x_\ell,\nu^{MU})$. Therefore,

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

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., $b_{\ell j} = \sigma^{UD}(x_\ell;\omega,G,\Psi) + \zeta^{UD}_{\ell j} < s_{\ell j}, \forall j\in\mathcal{A}^t$. Therefore,

$$L^{NB,UD}_{\ell t}\big(\mathbf{S}_{\ell t}|\sigma^{UD}(x_\ell;\omega,G,\Psi),\nu^{UD}\big) = \prod_{j'\in\mathcal{A}^t}\Phi\big(s_{\ell j'}|\sigma^{UD}(x_\ell;\omega,G,\Psi),\nu^{UD}\big).$$

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 3). In addition, the likelihood of placing bid $b_{\ell j} = \sigma^{UD}(x_\ell;\omega,G,\Psi) + \zeta^{UD}_{\ell j}$ is equal to $\phi(b_{\ell j}|\sigma^{UD}(x_\ell;\omega,G,\Psi),\nu^{UD})$. Thus,

$$L^{B,UD}_{\ell t}\big(\mathbf{b}_{\ell t},\mathbf{S}_{\ell t}|\phi(b_{\ell j}|\sigma^{UD}(x_\ell;\omega,G,\Psi),\nu^{UD}\big) = \phi\big(b_{\ell j}|\sigma^{UD}(x_\ell;\omega,G,\Psi),\nu^{UD}\big)P(\text{choose auction } j),$$

where

$$P(\text{choose auction } j) = \int_{\zeta_\ell} \mathbb{1}\Big(j \in \arg\max_{j'}\{u^{UD}_{j'}(x_\ell;\zeta_{j'},\omega,G,\Psi)\}\Big) d\mathbf{F}^{UD}_\zeta(\zeta_\ell),$$

is calculated numerically via importance sampling.

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

$$\begin{aligned}\mathcal{L}_\ell(\mathbf{X}_\ell|\theta) &= \int_{x_\ell}\mathcal{L}_\ell(\mathbf{X}_\ell|x_\ell,\alpha,\mu,\nu)\frac{f(x_\ell|\lambda,\gamma)}{f_I(x_\ell|\lambda_I,\gamma_I)} f_I(x_\ell|\lambda_I,\gamma_I)dx_\ell \\ &= E_{x_I}\Big(\mathcal{L}_\ell(\mathbf{X}_\ell|x_I,\alpha,\mu,\nu)\frac{f(x_I|\lambda,\gamma)}{f_I(x_I|\lambda_I,\gamma_I)}\Big),\end{aligned} \tag{20}$$

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 $x_{m,I}, m=1,2,...,M$ from the candidate Weibull distribution. Finally, we (approximately) compute integral (20) using Monte-Carlo integration as

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

In our estimation process, we draw $M = 200$ samples and set λ_I, γ_I to be equal to the first and second moments of the distribution of bids.

Appendix C: Supporting Material

C.1. Spillover Effects on Participation Rates

This subsection focuses on auction participation behavior 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 rate in Market B:

$$APR_{\ell tw} = \xi_\ell + \eta_w + \beta_1 CMB_\ell + \beta_2 TT_t + \beta_3 CMB_\ell \cdot TT_t + \epsilon_{\ell tw}, \tag{21}$$

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 (21), ξ_ℓ denotes the fixed effect associated with bidder ℓ's idiosyncratic participation behavior, η_w denotes a week fixed effect introduced to control for changes over time.

We include fixed effect ξ_ℓ to control for bidder ℓ's idiosyncratic participation behavior. We also include fixed effect η_w to control for changes over time, e.g. release of new product on participations, 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, and 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 cellphones (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 8 the average participation rates corresponding to cross-market bidders and exclusively Market B bidders for 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 a difference-in-difference to estimate coefficient β_3. As we report in Table 9, we estimate that the relative increase in the cross-market bidders' participation rate 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.

Figure 8 **Average auction participation rates**

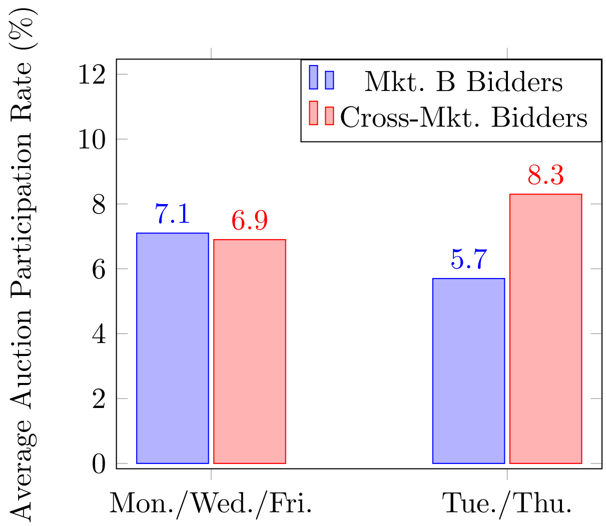

Table 9 **Spillover participation effect of the batch policy**

	Dependent Variable:
	Auction Participation Rate
Spillover Participation Effect $\hat{\beta}_3$	0.027^{***}
	(0.005)
Note:	*p<0.05; **p<0.01; ***p<0.001

C.2. 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 with 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 it converges (Algorithm 1). The expected payoff $r(x; G, \Psi)$ is calculated numerically for multi-unit and unit-demand 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

Assign a convergence threshold ϵ.
Assign a positive value to Δ, such that $\Delta > \epsilon$.
Iteration $i = 0$.
Initialization:
  The highest rival belief $G^{(0)}$.
  Market state distribution $\Psi^{(0)}$.
  $r^{TY}(x_\ell, G^{(0)}, \Psi^{(0)}), \forall \ell$ in the set of all bidders with type $TY \in \{MU, UD\}$.
while $\Delta > \epsilon$ **do**
  Iteration $i \leftarrow i + 1$.
  Simulate bidding history $\Omega^{(i)}$ under beliefs $G^{(i-1)}$ and $\Psi^{(i-1)}$.
  for day $t = 1...T$ **do**
    Update bidder pool on day t:
      1. Existing bidder exits $w.p.$ $1 - \alpha^{TY}$.
      2. Add new bidders $\sim P_{Arr}^{TY}$.
    Initialize market state ω:
      1. Add new auctions $\sim P_{Supply}$.
      2. Set standing bids.
    for bidder ℓ in the pool with type $TY \in \{MU, UD\}$ **do**
      Participation cost $c_{\ell t} \sim Exp(\mu^{TY})$.
      if $\sigma_{PRT}^{TY}(x_\ell; c_{\ell t})$ = Visit **then**
        Observe the market state ω and idiosyncratic terms $\zeta_{\ell \cdot}$.
        Select auction(s) according to $\sigma_{SLT}^{TY}(x_\ell; \zeta_{\ell \cdot}, \omega)$.
        Place bid(s) $\sigma_{BID}^{TY}(x_j; \zeta_{\ell \cdot}, \omega)$ in selected auction(s).
        Update standing bids in market state ω.
      end if
    end for
  end for
  Estimate $G^{(i)}$ and $\Psi^{(i)}$ nonparametrically from $\Omega^{(i)}$.
  Compute $r^{TY}(x_\ell, G^{(i)}, \Psi^{(i)}), \forall \ell$ in the set of all bidders with type $TY \in \{MU, UD\}$.
  $\Delta \leftarrow ||G^{(i)} - G^{(i-1)}|| + ||\Psi^{(i)} - \Psi^{(i-1)}||$.
end while