Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Dynamic Decision Making in Sequential Business-to-Business Auctions: A Structural Econometric Approach

Yixin Lu, Alok Gupta, Wolfgang Ketter, Eric van Heck

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–24
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Dynamic Decision Making in Sequential Business-to-Business Auctions: A Structural Econometric Approach

Yixin Lu,[a] Alok Gupta,[b] Wolfgang Ketter,[c,d] Eric van Heck[d]

[a] School of Business, George Washington University, Washington, District of Columbia 20052; [b] Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455; [c] Faculty of Management, Economics and Social Sciences, University of Cologne, 50931 Köln, Germany; [d] Rotterdam School of Management, Erasmus University, 3062 PA Rotterdam, Netherlands

Contact: yixinlu@gwu.edu, http://orcid.org/0000-0001-6515-1940 (YL); alok@umn.edu, http://orcid.org/0000-0002-2097-1643 (AG); ketter@wiso.uni-koeln.de, http://orcid.org/0000-0001-9008-142X (WK); evanheck@rsm.nl (EvH)

Received: March 13, 2014
Revised: July 30, 2015; November 12, 2016; September 17, 2017
Accepted: April 20, 2018
Published Online in Articles in Advance: April 22, 2019

https://doi.org/10.1287/mnsc.2018.3118

Abstract. We develop a dynamic structural model of competitive bidding in multiunit sequential business-to-business auctions. Our model accounts for two notable characteristics of these auctions: (i) bidders have multiple purchase opportunities for the same product, and (ii) winning bidders in each round can acquire multiple units of the same product. We apply the model to bidding data from the world's largest flower wholesale market at which trades are facilitated through fast-paced, sequential, Dutch auctions. Using a two-step estimation approach, we are able to recover the structural parameters effectively and efficiently. We then conduct policy counterfactuals to evaluate the performance of alternative design choices. The results suggest that the current auction practice still has ample room for improvement. In light of this, we propose an optimization framework that can facilitate auctioneers' decisions in making the trade-off between revenue maximization and operational efficiency.

History: Accepted by Lorin Hitt, information systems.

Keywords: auction design • dynamic discrete games • sequential auction • structural modeling

1. Introduction

Auctions have long been used as effective mechanisms for price discovery and resource allocation. Since the seminal work of Vickrey (1961), a large body of literature has investigated various informational and strategic factors in auction design using the game-theoretic framework. Despite the elegance of the results, such game-theoretic work has largely focused on highly stylized settings. The gaps between those stylized settings and the real-world operating environment seriously limit the direct usefulness of the theoretical insights to practitioners (Rothkopf and Harstad 1994).

The proliferation of online auctions has fueled a wide stream of empirical auction research. Depending on the approach taken to characterize bidding behavior, this body of literature can be divided into two broad categories. The first category of work takes the reduced-form approach and seeks to identify the empirical regularities from auction data. For example, Roth and Ockenfels (2002) found that, in eBay auctions, a significant amount of bids were submitted in the last minute. Lu et al. (2016) identified five different bidding strategies in multichannel business-to-business (B2B) auctions and showed that bidders' choice of strategies is associated with their demand, budget constraint, and transaction cost. Despite its usefulness in understanding real-life bidding behavior, the reduced-form approach, in general, cannot be used to evaluate policy interventions as they often constitute changes in bidders' bidding strategies (Lucas 1976). The second category of work takes the structural approach, assuming that bidders are profit-maximizing agents and they make the best response to the current situation. The structural approach attempts to recover the primitives—for example, bidders' value distribution—of an auction model from the observed data. The estimates of the primitives allow us to simulate results under alternative auction designs[1] and perform policy counterfactuals (Hickman et al. 2012). For example, using the bidding data from collectible coin auctions at eBay, Bajari and Hortacsu (2003) estimated a parametric structural model and showed that, for any given reserve price, a secret reserve price strategy yields higher revenue for sellers.

Although the value of structural models is well acknowledged, so far, the bulk of the structural empirical literature has focused on the identification and characterization of static (isolated) auctions. It is noteworthy that in most real-world auctions, both sellers and buyers are likely to adapt their behavior to the outcomes of previous transactions (Rothkopf and Harstad 1994). Sellers may adjust the reserve price and the supply based on the observations of previous sales. Similarly, buyers will update their willingness to pay based on their demand and participation experience as well as the outcomes in previous auctions (Goes et al. 2010). In particular, when buyers realize that there are multiple opportunities to purchase their desired products, they will shade the bids in the current auction to account for future purchasing opportunities (Zeithammer 2006). Unfortunately, there has been little research that systematically investigates such dynamic features and their policy implications.[2]

Our study attempts to fill this research gap by examining competitive bidding in a dynamic B2B market: the Dutch flower auction (DFA) market. The Dutch flower auctions play a critical role in the global flower network. In 2017, Royal FloraHolland, the market leader, reported a turnover of €4.7 billion.[3] The DFAs use the multiunit sequential Dutch auction mechanism. A unique feature of these auctions as opposed to other sequential auctions is that winning bidders in each round can purchase multiple units of goods (i.e., a homogeneous bundle of flowers). This poses considerable challenges to modeling the dynamic decision making in these auctions.

Drawing upon the growing literature on dynamic oligopoly games (Aguirregabiria and Mira 2007, Bajari et al. 2007), we develop a dynamic structural model to characterize the bidding behavior in the DFAs. Compared with existing structural auction models, our model accounts for both bidders' forward-looking behavior and their multiunit demand in sequential rounds. To address the econometric and computational challenges associated with the estimation of the model, we adapt the two-step estimation methods in Bajari et al. (2007) and Jofre-Bonet and Pesendorfer (2003) and recover the structural parameters governing the model effectively and efficiently. Using the estimated structural parameters, we conduct policy counterfactuals to evaluate and quantify the impact of different design choices on revenue and turnaround. The results from counterfactual experiments show that the current intuition- and experience-based auction practice has ample room for improvement. In light of this, we propose a novel optimization framework that can leverage the structural properties of bidders' strategic bidding behavior and guide auctioneers' decision making under different market conditions.

Our paper makes three important contributions. First, we extend the prior literature on structural auction models by characterizing bidding dynamics in sequential auctions in which winning bidders can purchase multiple units in each round. Currently, most of the structural models are concerned with single-unit auctions (Hickman et al. 2012). Of the few papers that investigated the dynamics of multiunit auctions, bidders are either assumed to have single unit demand throughout an auction or they can acquire at most one unit in each round. In contrast to the existing works, we relax the unit-demand (unit-sale) assumption and study the strategic bidding behavior in a more general setting. Second, our work complements the existing research that concerns the optimization of lot sizes in sequential auctions (e.g., Pinker et al. 2010 and Chen et al. 2011). Although significant effort has been devoted to quantify the effect of lot size on auction process and outcome, previous studies implicitly assume that bidders' bidding strategies are invariant to the changes in lot-sizing policies. We, on the other hand, explicitly account for bidders' potential responses to policy changes while evaluating the performance of different policies. Third, we contribute to the nascent literature on smart markets in which high-performance computational tools can assist human decision makers in complex trading environments (Bichler et al. 2010). In particular, our structurally based optimization framework provides the basis of the design and development of effective decision support systems for auctioneers in the DFAs and thereby offers the first step in reengineering this complex B2B market toward a smart market.

The remainder of this paper is organized as follows. Section 2 provides a review of relevant literature. Section 3 introduces the empirical context and presents results from preliminary analysis. Sections 4 and 5 describe the structural modeling framework and estimation methods. The empirical results are presented in Section 6. In Sections 7 and 8, we demonstrate how to use the structural estimation results to perform policy counterfactuals and optimize key auction design parameters in sequential auctions. Finally, in Section 9, we draw conclusions and outline the future research directions.

2. Literature Review

Our research draws upon two streams of literature: (i) structural econometric analysis of auction data and (ii) dynamic oligopoly games.

2.1. Structural Analysis of Auction Data

The structural econometric analysis of auction data, pioneered by Paarsch (1992), has emerged as one of the most successful areas of empirical auction research. By assuming that all observed bids are the equilibrium bids, the goal is to recover the economic primitives (e.g., distribution of bidders' valuations) of the underlying auction model. Over the past decades, there has been significant progress in structural econometric analysis of common auctions (Hickman et al. 2012). The estimation methods can be broadly divided into two categories: parametric and nonparametric. The parametric estimation—for example, Donald and Paarsch (1996)—imposes explicit distributional assumptions on bidders' value distribution, which inevitably bears misspecification risks. In contrast, the nonparametric methods initiated by Guerre et al. (2000) do not require any a priori assumptions regarding the value distribution and in many cases offer computational advantages. Unfortunately, only a few standard auction models are nonparametrically identifiable (Athey and Haile 2002).

So far, most of the structural modeling research has restricted attention to static auctions (Paarsch 1992, Donald and Paarsch 1996, Guerre et al. 2000, Flambard and Perrigne 2006). Comparatively, dynamic auctions (e.g., sequential, Dutch auctions) are much less well understood. Our paper aims to fill this gap in the literature. In particular, drawing upon the empirical findings from Zeithammer (2006), we explicitly model bidders' strategic forward-looking behavior in sequential auctions. In this regard, our work is related to the recent work of Backus and Lewis (2016), in which the authors propose a demand system for a dynamic market with heterogeneous goods and directed search. However, their model restricts attention to unit-demand bidders (thus, winning bidders would exit at the end of each period with certainty) whereas we allow for acquisition of multiple units in each round.[4] It is noteworthy that in multiunit settings, bidders have the incentive to shade bids differently across units (Ausubel et al. 2014), and the equilibrium can involve mixed strategies (McAfee and Vincent 1993). Such complication poses significant challenges to the identification and estimation of the underlying models. On the methodology side, we build on Jofre-Bonet and Pesendorfer (2003) although the latter examines first-price, sealed-bid procurement auctions with a given number of bidders whereas we are dealing with fast-paced Dutch auctions in which only winning bids are revealed. An important innovation of our work is that we combine the estimated structural parameters with the institutional characteristics and build a dynamic optimization framework to facilitate auctioneers' decision making.

2.2. Dynamic Oligopoly Games

Given that in sequential auctions bidders' decisions not only affect their own payoffs in the current and future periods, but also the payoffs of their competitors, the bidding competition studied in our paper is conceptually similar to the dynamic oligopoly games in industrial organization literature. In their pioneering work, Ericson and Pakes (1995) (EP) modeled the dynamics of an oligopoly market with a set of incumbent firms and a large number of potential entrants and introduced the Markov perfect equilibrium (MPE) as the solution concept. The EP model and its extensions have been widely adopted to study strategic interactions of heterogeneous agents in many different fields, including marketing (Dubé et al. 2010, Yao and Mela 2011) and information systems (Huang et al. 2015).

Despite the theoretical attractiveness, one of the main challenges when applying the EP-type models to real-world settings is the dimensionality problem associated with the computation of MPE.[5] Specifically, because dynamic oligopoly models are analytically intractable, the "curse of dimensionality" is present even in a simple model with a small number of players (Bajari et al. 2007). As soon as we start to account for the potential heterogeneity among different players (e.g., differentiated products or heterogeneous costs), the computational burden quickly becomes prohibitive. In our case, given that more than 100 bidders may participate in the bidding competition in each auction, we draw upon the asymptotic environmental similarity (AES) property of large auctions (Swinkels 2001) to approximate bidders' decisions in sequential rounds of the auctions.

3. Empirical Context and Data

A good understanding of the institutional environment and the behavior of the economic agents (individuals or firms) involved in the environment is both helpful and necessary to any structural modeling work. In this section, we first introduce the empirical context, focusing on the auction mechanism and bidders' decision process, and then discuss the data and preliminary analyses that serve as important precursors to our structural modeling framework.

3.1. The Dutch Flower Auctions

The Dutch flower auctions play a vital role in the global floriculture sector, both as a marketplace and as a distribution hub (Kambil and van Heck 2002). The auctions attract growers not only from the Netherlands, but also from other countries—particularly, Belgium, Denmark, Ecuador, Ethiopia, Germany, Israel, Italy, Kenya, Spain, and Zimbabwe—who ship their flowers to the Netherlands for sale. Because flowers are highly perishable goods, realizing short lead times is of utmost importance. Most flowers make it from the growers to the retail shops within 24 to 48 hours.

On weekday mornings, flowers are brought to the auction halls before 4:30 a.m., and their quality is assessed before the auctions start at 6:30 a.m. It is worth mentioning that, currently, 99% of the cut flowers sold through the auctions are rated as A1, the highest quality level. Flowers are auctioned in separate lots, which are bundles of homogeneous products (i.e., from the same grower and with the same characteristics). The size of a lot can vary from a few units to more than a hundred units. Depending on the type and quality of flower, each unit consists of 10 to 50 stems. Currently, different product categories (e.g., roses, chrysanthemums) are auctioned based on a predetermined order throughout the year. However, within the same product category, the order of auctions is determined by random draws from the grower pool to ensure

Figure 1. Overview of the Two Buying Scenarios

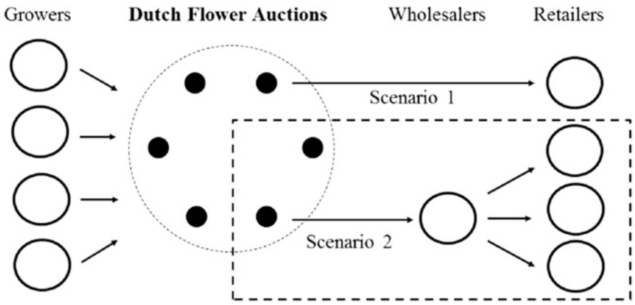

Notes. In the first scenario, retailers buy for themselves. In the second scenario, wholesalers buy for their retail clients.

Figure 2. Factors Impacting the Auctioneer's Decision

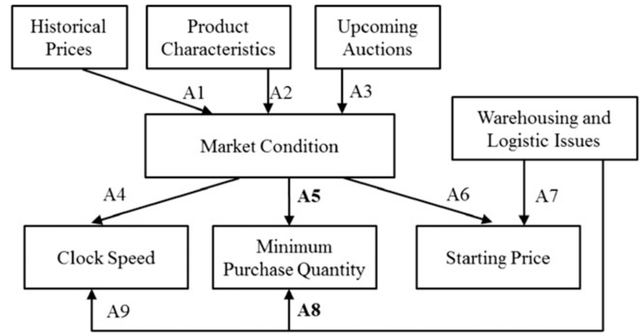

fairness. All the bidders and auctioneers are notified about the detailed auction schedule at the beginning of the day.

There are two buying scenarios through the DFAs (see Figure 1). In the first scenario, the DFA directly sells flowers to the retailers (e.g., florists) whereas, in the second scenario, the DFA sells to wholesalers who buy on behalf of their clients (e.g., retailers located in different countries). Around two thirds of the auctioned flowers are sold to wholesalers. In this research, we choose to focus on the bidding and auctioning decisions in the wholesale scenario.

3.1.1. The Mechanism. The DFAs use an open-cry, descending auction mechanism, which is also known as a *Dutch* auction. It is implemented using a single-handed clock that initially points to a high price and then quickly ticks down in a counterclockwise direction. As the price falls, each bidder can bid by pressing a button indicating that the bidder is willing to accept the current price. The first bidder who makes a bid wins. The winning bidder can select the purchase quantity (which must exceed the minimum quantity set by the auctioneer). If the winning bidder does not purchase all the available units in the current lot, the clock restarts at a high price and the auction continues. This process repeats until the entire lot is sold or until the price falls below the seller's reserve price, in which case any unsold goods in that lot are destroyed.[6] On average, each transaction takes three to five seconds.

3.1.2. Auctioning. Each auction clock is closely monitored by an auctioneer, who can influence the bidding competition by controlling the speed of the clock, the starting price, and the minimum purchase quantity. The auctioneers' objective is twofold: to realize high selling prices and to achieve a quick throughput. During the actual auction process, auctioneers need to set the auction parameters, taking into account internal conditions (e.g., warehousing) and market conditions, which, in turn, depend on the following: (i) historical prices, (ii) product characteristics (e.g., quality measures), and (iii) upcoming auctions. Because of the high complexity and extreme time pressure, even though auctioneers have many years of experience, their decisions are neither optimal nor standardized (Bichler et al. 2010). In particular, the clock speed is kept constant (36 milliseconds per tick), and the starting price in subsequent rounds is set by adding a constant incremental on top of the winning price in the current round (constant swingback). The minimum purchase quantity is the only parameter used by auctioneers to influence the real-time bidding dynamics. Figure 2 depicts factors impacting the auctioneer's decision. The links A5 and A8 are the focal processes we address in this paper.

3.1.3. Bidding. As a rule of thumb, no bidder in the DFAs should ever bid more than the bidder's *value* of the product (flowers) less any associated costs. A bidder's utility upon winning increases as a function of how far below the bidder's value the winning price is. Therefore, the key question a bidder faces concerns the right time to hit the buy button (i.e., at what price) and make a purchase. The answer to this question depends on the price at which the bidder expects the auction to clear—in other words, the bidder's predicted winning price.

These auctions constitute games of incomplete information. In particular, each auction comprises a set of players (the bidders), each with a set of actions (prices and quantities) and private information (resale values), and the outcome of an auction depends on the choices of all players rather than the choices of a single player as in a monopolistic setting or the choices of none as in a model of perfect competition. Figure 3 depicts factors impacting the bidder's decision. The links denoted by B1–B5 are the processes we model in the current paper.

3.2. Data

Our data consists of transaction details of large roses from a major auction site[7] between September 1 and October 10, 2014 (six weeks with 30 auction days).

Figure 3. Factors Impacting the Bidder's Decision

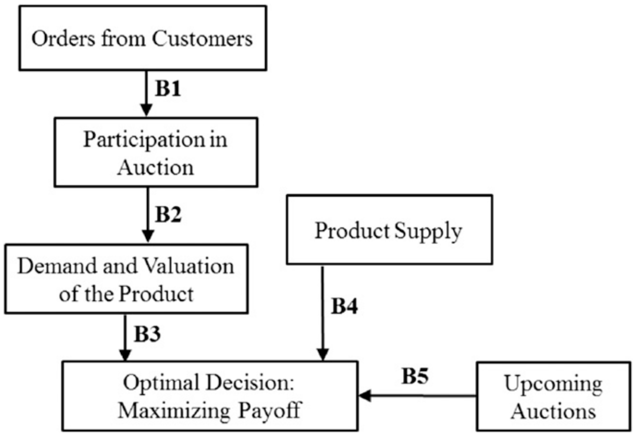

Table 1 gives a stylized yet representative example of a sequence of transactions that can be found in our data set. Here, we do not include all the attributes but only those relevant to our structural analysis: (i) transaction timing; (ii) product characteristics (e.g., product type and quality); (iii) supply-side information, which includes lot size and minimum purchase quantity; (iv) the precise market actors (grower identity and bidder identity); (v) demand-side information, that is, the number of bidders who registered at an auction; and (vi) bidders' real-time decision variables of price and quantity.

Further, unlike the sequential auctions studied in Brendstrup (2007) and Donald et al. (2006) in which only one unit is on sale in each round, in our case, multiple units can be purchased in each round. Thus, neither bidders nor the auctioneer knows in advance how many rounds an auction will take. Table 1 also shows that the winning prices are not monotonically decreasing or increasing,[8] and the same buyer (ID: 439) may end up winning multiple sublots at different prices. These observations are indicative of the complexity of the DFAs, and such complexity arises because of the flexibility of the auction mechanism.

The specific product subcategory we chose is *Red Naomi*, which is the most important Dutch-grown red rose in the high-end segment. Although all the products within this subcategory are rated as of the highest quality level, there remains considerable heterogeneity with respect to product characteristics (e.g., stem length, blooming stage). Because our study focuses on bidding dynamics in *multiunit* auctions rather than *multiobject* auctions,[9] ideally, we would like to have a homogeneous sample. However, this would leave us with a very small data set. In light of this, we created a sample with products that are close substitutes. In particular, we combined the transactions of Red Naomi with stem length of 50 cm, 55 cm, and 60 cm as reduced-form analyses showed no statistical difference between these products.[10] In total, we have 5,536 transactions from 1,012 auctions with 388 unique winning bidders. Table 2 provides summary statistics of the sample.

At the day level, the number of auctions varied from a minimum of 15 to a maximum of 45 with an average of 33.7, and the number of transactions (subauctions) varied from a minimum of 58 to a maximum of 326 with an average of 184.5. The average number of winning bidders also varied substantially from a minimum of 45 to a maximum of 154. Clearly, there were some peak days during the 30-day period under consideration. At the transaction level, we examined the number of registered bidders and the purchase quantity as well as the winning price. Again, we found considerable variation. Specifically, the minimum size of a purchase was as small as one unit whereas the maximum was as large as 162 units. This implies that the distribution of the purchase amount at the transaction level is highly skewed. As for the winning price, the minimum was as low as five cents, and the maximum was 67 cents.

It is well known that the prices of cut flowers are extremely volatile, partly because of their perishability. This is consistent with our observation from Figure 4, in which prices exhibit considerable variability from day to day. Specifically, in the first and third week of September, prices exhibited a downward trend whereas in the second week, there was a mild upward trend. Toward the end of September, prices were relatively stable, followed by a surge starting from October 3. Figure 4 also shows that, even within the same day, there was significant price variation.

Further, in light of the declining price anomaly documented in prior literature (Van den Berg et al. 2001), we also examined the price trend in sequential

Table 1. A Sample Entry in a Logbook

Transaction time	Grower ID	Product ID	Product quality	Available units	Minimum purchase quantity	Number of registered bidders	Buyer ID	Purchase quantity	Price (cents)
07:10:54	689	16207	A1	18	1	170	439	1	50
07:10:56	689	16207	A1	17	2	173	395	5	49
07:10:59	689	16207	A1	12	2	175	439	3	46
07:11:01	689	16207	A1	8	4	168	563	4	48

Note. The price is for each stem, not each unit.

Table 2. Summary Statistics

	Day-level variables			Transaction-level variables		
	Number of auctions	*Number of transactions*	*Number of winning bidders*	*Number of registered bidders*	*Purchase quantity*	*Price (cents)*
Mean	33.7	184.5	108.2	144.4	11.0	30.6
Standard deviation	8.7	67.2	31.8	19.4	12.3	9.1
Minimum	15.0	58.0	45.0	93.0	1.0	5.0
Maximum	45.0	326.0	154.0	217.0	162.0	67.0

Figure 4. Overview of Price Trend from Day to Day

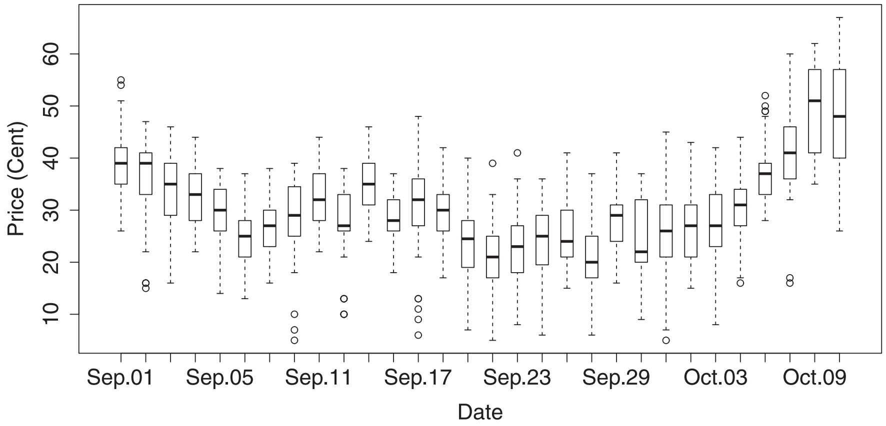

rounds. Overall, we found inconclusive evidence for the downward trend; in some cases, the winning price continued to decrease in sequential rounds, and in other cases, the winning price maintained the same level over time or even increased from an earlier round to a later round. Figure 5 illustrates the variability in the price trend. Here, we plotted the sequence of winning prices in each auction on September 1. The x-axis corresponds to the round number of a transaction in a given auction. A round number of two indicates the second sale in an auction.

A variety of theories have been developed to explain the price trend in sequential auctions (see Ashenfelter and Graddy 2003 for a review). However, almost all the theoretical work adopts a stylized setup with unit-demand bidders who are competing for two units. In this regard, the dynamic structural model developed in this paper serves as a useful starting point for understanding to what extent the predictions of prior literature carry over to more general and realistic settings.

Figure 5. Price Dynamics in Sequential Rounds (September 1)

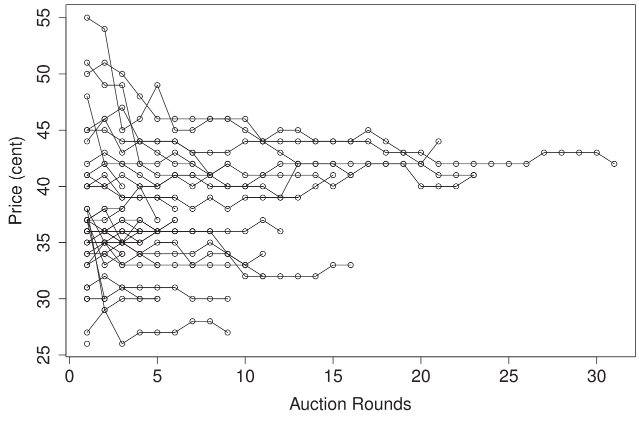

As we mentioned before, an important feature that differentiates the DFAs from other sequential auctions is that bidders can purchase multiple units of goods in each round. Panel (a) of Figure 6 shows the distribution of purchase quantities. We can see that (winning) bidders typically purchased no more than 30 units. We also plot the distribution of the *excessive* purchase quantities, that is, purchase quantities less the corresponding minimum purchase quantities (see panel (b) of Figure 6). The enormous amount of zeros indicates that many transactions take place for the minimum purchase quantity, suggesting that bidders might use the minimum purchase quantity as the reference point to determine their purchase quantity. Auctioneers must choose the minimum purchase quantity carefully to achieve desirable outcomes. In Section 8, we discuss in detail the optimization of the minimum purchase quantity in sequential rounds.

4. Modeling Framework

Consider a simplified market of the DFAs with a single auction clock. Each day, a homogeneous good is

Figure 6. Distribution of Purchase Quantities vs. Distribution of Excessive Purchase Quantities

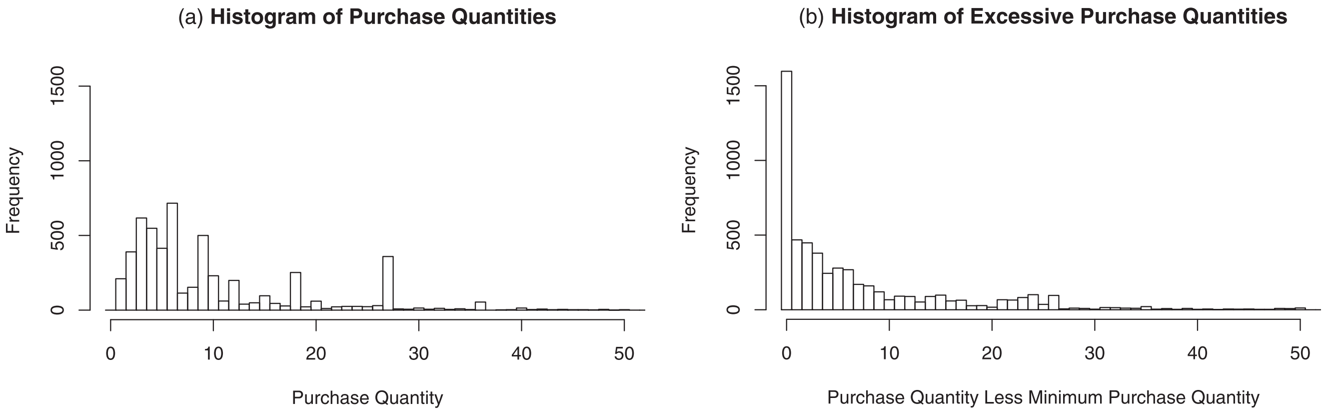

pushed to the market and auctioned using the mechanism described in Section 3.1. At the beginning of the day, an auction schedule is announced that specifies, for each lot $l = 1, 2, \ldots$, the quantity for sale $\bar{Q}_l$. There are N risk-neutral bidders indexed by i. The dynamic game proceeds as follows:

- Stage 1: Auctioneer announces the available units and the minimum purchase quantity.
- Stage 2: Bidders simultaneously and privately submit[11] their bids.
- Stage 3: Auctioneer announces the bidding outcome (i.e., the winning price and winner's identity), and the winning bidder selects the purchase quantity.

Because each lot is auctioned in several rounds, that is, there are multiple subauctions from the same lot, to simplify the notations, we index the subauctions $1, \ldots, t, \ldots$ using lexicographical order.[12] In the following, we use t to refer to a subauction as well as the time period corresponding to the subauction. Additionally, when there is no confusion, we use $\bar{q}_0^t$ to denote the available quantity for the lot under auction at time t.

For any bidder, prior to submitting a bid, the vector of commonly observable endogenous state variables at time t are $\mathbf{s}^t := (\bar{q}_0^t, \bar{q}_1^t, \ldots, \bar{q}_i^t, \ldots, \bar{q}_N^t)$, where $\bar{q}_i^t$ denotes the quantity that has been purchased by bidder i at the beginning of period t. We can also write $\mathbf{s}^t$ from bidder i's perspective as $(s_0^t, s_i^t, \mathbf{s}_{-i}^t)$, where $s_0^t = \bar{q}_0^t$, $s_i^t = \bar{q}_i^t$, and $\mathbf{s}_{-i}^t = \{\bar{q}_j^t : j \neq i\}$. Bidder i's action in period t, denoted by a_i^t, is a bid–quantity pair; that is, $a_i^t := (b_i^t, q_i^t)$. The collection of all bidders' actions in period t is denoted by $\mathbf{a}^t := (a_1^t, \ldots, a_i^t, \ldots, a_N^t)$.

4.1. Key Assumptions

As Chintagunta et al. (2006) point out, theory per se is rarely sufficient to enable a complete specification of a structural model. In fact, most structural models make strong assumptions about the form of the utility function, the distribution of unobservable components, and the nature of equilibrium in a given market. Our model is no exception. In what follows, we discuss the main assumptions used in our modeling framework.

To start with, because bidders are buying on behalf of their customers, their participation decision is primarily driven by customer orders. Specifically, as flowers are highly perishable, bidders' procurement decisions must been *synchronized* with customer orders.

Assumption 1 (Market Entry). *Bidders' participation in the auction market is order driven. In particular, at the beginning of the day, each potential bidder receives a private signal $\zeta_i = 0, 1, 2, \ldots$ regarding customer orders; ζ_i is independently and randomly drawn from a commonly known Poisson distribution with intensity λ; that is,*

$$\Pr(\zeta_i = m) = \frac{\lambda^m \exp(-\lambda)}{m!}, \quad m = 0, 1, 2, \ldots \tag{1}$$

Bidder i will enter the auctions on a given day if and only if bidder i receives a nonzero signal. In other words, bidders' entry follows a Bernoulli process. Thus, the number of bidders participating in the auctions (N) is distributed binomially with two parameters: the total number of bidders registered at the market, denoted by $\mathcal{N}$, and the participation probability, that is, $1 - \exp(-\lambda)$. Its probability mass function is given by

$$\Pr(N = n) = \binom{\mathcal{N}}{n} [1 - \exp(-\lambda)]^n \exp(-\lambda)^{\mathcal{N}-n}, \quad n = 0, 1, 2, \ldots \tag{2}$$

Note that our market entry assumption admits differential participation at the day level, which is consistent with the observation from the empirical data (see Table 2). The Poisson assumption is widely used in modeling customer behavior in supply chain environments (Caldentey and Vulcano 2007, Shen and Su 2007). In our case, given that customer orders are from different segments (countries or regions), it is reasonable to assume that they follow a Poisson distribution.

As Donald et al. (2006) has pointed out, *nonparticipation* in first-price, sealed-bid auctions does not introduce a problem in deriving the data-generating process of the winning bids. This is because the bidders' decision rule depends on the number of potential bidders instead of the number of actual bidders. Given the strategic equivalence between first-price, sealed-bid auctions and Dutch auctions, the decision rule and the data-generating process on a given day in the DFAs depends on the total number of bidders who have received a nonzero signal and decided to enter the market rather than the actual number of bidders who logged into the system in each round. Thus, we do not make further assumptions about bidders' entry or exit decisions at the auction level.

Next, drawing upon prior works on structrual analysis of auction data (see Hickman et al. (2012) for more details), we make the following assumption about the bidders' information structure.

Assumption 2 (Bidders' Valuations). *Bidders have independent, private valuations. In particular, bidder i's unit valuation for any homogeneous bundle in period t, denoted by* v_i^t*, is privately known and drawn independently from a twice differentiable, common knowledge conditional distribution function* $F(\cdot \mid s_0^t, s_i^t, \mathbf{s}_{-i}^t)$ *with compact support* $[\underline{v}, \bar{v}]$.

The independent private valuation (IPV) assumption is widely used in both theoretical and empirical auction literature.[13] In our case, the IPV assumption can be justified by the market structure: bidders are typically serving distinct market segments, and they come to the auctions with the willingness to pay of their customers. Note that within-bidder volatility in valuations across different subauctions may arise from unexpected demand shocks from the customer side and the change of market condition as well as the temporary budget constraint for distribution and delivery.

Finally, we assume winning bidders' purchase quantities in sequential rounds are conditional on customer orders and strictly exogenous. Such a strict exogeneity assumption can be justified by the B2B nature of the trades; that is, bidders are not purchasing the auctioned products for personal consumption, but for resale.[14] In particular, with the increasing use of the online channel in these auctions, bidders can easily communicate with their customers about the purchase quantities in real time. Based on empirical observation from Figure 6, we make the following assumption about bidders' purchase quantities to account for idiosyncratic shocks associated with customer orders.

Assumption 3 (Bidders' Purchase Quantities). *In period t, bidder i's purchase quantity is given by* $q_i^t = \underline{q}^t + r_i^t$*, where* $\underline{q}^t$ *is the minimum purchase quantity and* r_i^t *is the idiosyncratic shock, drawn from a commonly known, zero-inflated negative binomial distribution with probability density function* f_{zinb}.

According to Hickman et al. (2012), when bidders are allowed to purchase multiple units in each round, the dynamic system would be intractable even for a small number of bidders. The strict exogeneity assumption of purchase quantities is indispensable for our empirical estimation given the large number of active bidders in the market. Despite its restrictiveness, the parametric specification in Assumption 3 is well supported by empirical data.[15]

4.2. A Single Bidder's Decision Problem

For any bidder i, the bidding problem can be represented as a Markov decision process (MDP). Figure 7 provides an illustrative example of two sequential auctions, each with two units. Suppose the minimum purchase quantity is set to one throughout the auction. Bidder i starts in the upper-left node: it is the first round of the auction, there are two units for sale, and the bidder has not yet purchased any units. Bidder i will choose a bid–quantity pair at this state. There are four possible transitions from this state: either a competitor wins, which reduces the available quantity by one or two without increasing bidder i's purchased quantity, or bidder i wins, which increases bidder i's purchased quantity and reduces the available quantity by one or two. The probability that bidder i transitions to the winning state depends on bidder i's own bid and bidder i's competitors' bids. Once all the units are sold out in the first auction, bidder i transitions to the second auction.

Formally, the bidding problem for bidder i can be defined by the following MDP:

Figure 7. Bidder i's Decision Problem as an MDP: An Illustrative Example of Two Auctions, Each with Two Units

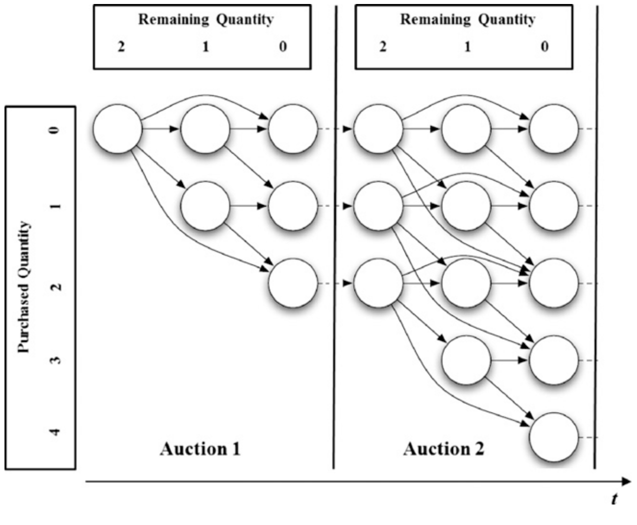

- **States:** $\{\mathbf{s}^t\}$, where $\mathbf{s}^t := (\bar{q}_0^t, \bar{q}_1^t, \ldots, \bar{q}_i^t, \ldots, \bar{q}_N^t)$.
- **Actions:** Given state $\mathbf{s}^t$, the set of actions for bidder i is $\{a_i^t := (b_i^t, q_i^t),\ \underline{b} \le b_i^t \le \bar{b},\ q_i^t = q^t + r_i^t,\ q_i^t \le \bar{q}_0^t\}$, where $\underline{b}$ and $\bar{b}$ refer to the lower and upper bounds of bids. Note that for states in which no item remains at time t ($\bar{q}_0^t = 0$), there only exists a single, *dummy* action; that is, bidder i moves to the next auction with certainty.
- **Transitions:** The transition function, denoted by $\omega : \mathbf{A} \times \mathbf{S} \rightarrow \mathbf{S}$, is a deterministic function of the states and actions. In particular, we have

$$\omega(\mathbf{a}^t, \mathbf{s}^t) = \begin{cases} s_0^{t+1} = \bar{q}_0^t - q_i^t,\ s_i^{t+1} = \bar{q}_i^t + q_i^t,\ \mathbf{s}_{-i}^{t+1} = \mathbf{s}_{-i}^t \\ \qquad\qquad \text{if } b_i \ge \max_{j \ne i} b_j \\ s_0^{t+1} = \bar{q}_0^t - q_{j^*}^t,\ s_i^{t+1} = s_i^t, \\ \mathbf{s}_{-i}^{t+1} = \{\{\bar{q}_j^t : j \ne i, j^*\}, \bar{q}_{j^*}^t + q_{j^*}^t\} \\ \qquad\qquad \text{if } b_i < b_{j^*} = \max_{j \ne i} b_j \end{cases} \tag{3}$$

- **Rewards:** Given the private valuation v_i^t and the state $\mathbf{s}^t$, bidder i's expected unit payoff in the current period is given by

$$\pi_i(\mathbf{s}^t, b_i^t, v_i^t) = (v_i^t - b_i^t) P(i \text{ wins} \mid b_i^t, s_0^t, s_i^t, \mathbf{s}_{-i}^t). \tag{4}$$

In addition to the current-period payoff, bidder i also takes into account bidder i's expected future payoffs when making the bidding decision. Thus, given the state variables in period t, bidder i's expected unit payoff, evaluated prior to the realization of the private valuations, is given by

$$E\left[\sum_{\tau=t} \beta^{\tau-t} \pi_i(\mathbf{s}^\tau, b_i^\tau, v_i^\tau)\right]. \tag{5}$$

The expectation is taken over the realization of bidder i's private valuation and all bidders' actions in the current period as well as the future realizations of state variables, actions, and private valuations. The parameter $\beta \in (0,1)$ is a common discount factor. It accounts for the impact of participation costs (e.g., the time and effort incurred monitoring the dynamics and deliberating the bids) on bidding decisions in sequential rounds. The endogenous state variables $\mathbf{s}^{t+1}$ are drawn from the probability distribution $P(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t)$. Note that because bidders' purchase quantities are strictly exogenous (see Assumption 3), maximizing the expected total payoff reduces to maximizing the expected unit payoff. This property is very useful as it allows us to adapt the two-step estimation frameworks in Bajari et al. (2007) and Jofre-Bonet and Pesendorfer (2003) to estimate our model.

4.3. The Equilibrium Concept

To analyze bidders' equilibrium behavior, we focus on pure-strategy Markov perfect equilibrium.[16] The MPE implies that bidders' bidding strategies only depend on the payoff-relevant information. Thus, we can describe the equilibrium bidding strategy of bidder i as a function $\sigma_i(\mathbf{s}^t, v_i^t) = b_i^t$. Because the strategy profile $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ is time invariant, we can drop the period index t. Given state $\mathbf{s}$, we can express bidder i's expected unit payoff under the equilibrium strategy profile $\boldsymbol{\sigma}$ recursively as a Bellman equation:

$$V_i(\mathbf{s}; \boldsymbol{\sigma}) = E\left[\pi_i(\boldsymbol{\sigma}(\mathbf{s}, \mathbf{v}), \mathbf{s}, v_i) + \beta \int_{\mathbf{s}'} V_i(\mathbf{s}'; \boldsymbol{\sigma}) \cdot dP(\mathbf{s}' \mid \boldsymbol{\sigma}(\mathbf{s}, \mathbf{v}), \mathbf{s}) \mid \mathbf{s}\right]. \tag{6}$$

Here, V_i is bidder i's ex ante value function in the sense that it reflects the expected unit payoff at the beginning of a period before the realization of private valuations. Following the literature of dynamic oligopoly games, a strategy profile $\boldsymbol{\sigma}$ is an MPE if, for any bidder i, given competitors' strategy profile $\boldsymbol{\sigma}_{-i}$, states $\mathbf{s}$, and Markov strategies σ_i',

$$V_i(\mathbf{s}; \boldsymbol{\sigma}) \ge V_i(\mathbf{s}; \sigma_i', \boldsymbol{\sigma}_{-i}). \tag{7}$$

Table 3 provides a summary of the notations used in our modeling framework.

Before moving to the details of the estimation strategy, we would like to discuss under which conditions our model can be identified. It is worth mentioning that although Guerre et al. (2000) shows that the value distribution is nonparametrically identifiable in the static first price auction, Rust (1994) demonstrates that the model primitives in a dynamic decision problem are not identified when the per-period payoff function is unknown. In our case, because the reward function and the discount factor β are both known, it follows from proposition 2 of Jofre-Bonet and Pesendorfer (2003) that the value distribution function is identifiable on the interval $[\underline{v}, \bar{v}]$.

5. Estimation Strategy

A standard approach to estimate MPE models is to employ dynamic programming (Rust 1994). However, this requires repetitive evaluation of the value function (Equation (6)) and is computationally demanding. Given the high dimensionality of our state space and action space, such an approach would be infeasible. In light of recent works on structural modeling (Yao and Mela 2011, Huang et al. 2015), we adapt the two-step estimation frameworks in Bajari et al. (2007) and Jofre-Bonet and Pesendorfer (2003): in the first step, we estimated the transition function, the bid distribution function, and the value function from the empirical data; in the second step, the distribution of bidders' private valuations is estimated based on the first-order condition. In what follows, we explain the estimation method in detail and demonstrate its effectiveness using Monte Carlo simulation.

Table 3. Summary of Notations

Notation	Explanation
$i, j, k; l; t$	Indices of bidders, index of auctions, index of time period
ζ_i	The private signal about customer orders for bidder i, which is randomly drawn from a Poisson distribution with intensity λ
$\mathcal{N}$, N	Number of potential bidders, number of active bidders
$\bar{Q}_l$	The initial available units for auction in lot l
$\bar{q}_0^t$	The remaining units at the beginning of period t
$\underline{q}^t$	The required minimum purchase quantity in period t
$\mathbf{s}^t$	The commonly observable state variables in period t
$(s_0^t, s_i^t, \mathbf{s}_{-i}^t)$	The state variables evaluated from bidder i's perspective in period t
v_i^t	Bidder i's unit valuation for the product in period t
$\mathbf{v}^t$	The collection of all bidders' unit valuations for the product in period t
$F(\cdot\|s_0^t, s_i^t, \mathbf{s}_{-i}^t)$	The cumulative distribution function of bidder i's private valuation defined in $[\underline{v}, \bar{v}]$
$f(\cdot\|s_0^t, s_i^t, \mathbf{s}_{-i}^t)$	The probability density function of bidder i's private valuation
b_i^t	Bidder i's bid in period t, which is bounded by $\underline{b}$ and $\bar{b}$
$\mathbf{b}_{-i}^t$	Bidder i's competitors' bids in period t
$G(\cdot\|s_0^t, s_i^t, \mathbf{s}_{-i}^t)$	The cumulative distribution function of equilibrium bids
$g(\cdot\|s_0^t, s_i^t, \mathbf{s}_{-i}^t)$	The probability density function of equilibrium bids
$\bar{q}_i^t$	The quantity purchased by bidder i at the beginning of period t
q_i^t	Bidder i's purchase quantity in period t
r_i^t	The difference between bidder i's purchase quantity and the minimum purchase quantity in period t, which is drawn randomly from a zero-inflated negative binomial distribution with probability density function f_{zinb}
a_i^t	Bidder i's action in period t, $a_i^t := (b_i^t, q_i^t)$
$\mathbf{a}^t$	The collection of all bidders' actions in period t
β	The fixed discount factor associated with future payoff
σ_i	Bidder i's equilibrium bidding strategy
π_i	Bidder i's period payoff
ω	Transition function of state variables
V_i	Bidder i's value function
$\phi(\cdot)$	The explicit function expressing bidder's private valuation in terms of the equilibrium bids, the distribution of equilibrium bids, and the value function
$\psi(\cdot\|s_0, s_i, \mathbf{s}_{-i})$	Auxiliary function, which is defined as the ratio of bid distribution function and probability density function of equilibrium bids at state $(s_i, \mathbf{s}_{-i})$

5.1. Estimation Method

Our estimation method is based on the first-order condition of optimal bids.[17] The general idea is to express the privately known valuations as an explicit function of the state, the observed bids (which are assumed to be equilibrium bids), the bid distribution function, and the value function.

Let $\phi(\cdot)$ denote bidder i's private valuation associated with a bid; $G(\cdot|s_0, s_i, \mathbf{s}_{-i})$ and $g(\cdot|s_0, s_i, \mathbf{s}_{-i})$ denote the distribution function and probability density function of equilibrium bids of bidder i at state $(s_0, s_i, \mathbf{s}_{-i})$, respectively. In addition, we define an auxiliary function ψ as follows:

$$\psi(\cdot|s_0, s_i, \mathbf{s}_{-i}) = \frac{g(\cdot|s_0, s_i, \mathbf{s}_{-i})}{G(\cdot|s_0, s_i, \mathbf{s}_{-i})}. \tag{8}$$

The first-order condition yields (see Appendix A.1 for details of the derivation):

$$\begin{aligned}\phi(b|s_0, s_i, \mathbf{s}_{-i}, \psi, V_i) = b &+ \frac{1}{\sum_{j\neq i}\psi(b|s_0, s_j, \mathbf{s}_{-j})} \\ &- \beta\sum_{j\neq i}\frac{\psi(b|s_0, s_j, \mathbf{s}_{-j})}{\sum_{k\neq i}\psi(b|s_0, s_k, \mathbf{s}_{-k})} \\ &\cdot (V_i(\omega_i(\mathbf{a}, (s_0, s_i, \mathbf{s}_{-i}))) \\ &- V_i(\omega_j(\mathbf{a}, (s_0, s_i, \mathbf{s}_{-i})))),\end{aligned} \tag{9}$$

where $\omega_i(\cdot)$ denotes the state transition given bidder i is the winner and $\omega_j(\cdot)$ the state transition given bidder $j\,(j \neq i)$ is the winner. For simplicity of notation, we write $\omega_j(\mathbf{a}, (s_0, s_i, \mathbf{s}_{-i}))$ as ω_j. Equation (9) states that the private valuation equals the bid plus a markup. The markup consists of two parts: the first part accounts for the level of competition in the current period, that is, the more intense the competition is, the more the bidder must pay to win in the current period; the second part accounts for the effect on the future discounted payoff if bidder i wins the auction instead of losing it to bidder j. When making the bidding decision, a strategic (forward-looking) bidder is making the trade-off of winning in the current round (but perhaps paying too much) and winning in a future round (but perhaps the product is sold out).

To estimate the distribution of bidders' private valuations, we need estimators for the transition function ω, the auxiliary function ψ, the discount factor β, and the value function V_i on the right-hand side of Equation (9). As described previously, the transition function is a given function (see Equation (3)); the bid distribution function as well as the auxiliary function can be directly estimated from the bidding data. The discount factor is typically nonidentifiable without exogenous variation in the bidding environment (Rust 1994). In light of this, we follow prior works (e.g., Yao and Mela 2011 and Huang et al. 2015) and set the discount factor β to 0.9 for our estimation.[18] Thus, the main challenge lies in the estimation of the value function: because the expression of the value function in Equation (6) involves the bidders' valuations that are unobserved and endogenous bidding decisions of multiple bidders, we cannot directly estimate it from the bidding data. In what follows, we explain how to approximate the value function.

The key idea underlying the approximation method is that the distribution of equilibrium bids determines the discounted sum of expected future payoffs. As such, we can represent the value function using the distribution of bids only. This is formalized in Proposition 1.

Proposition 1. *The value function can be represented as a recursive equation involving the bid distribution function. In particular,*

$$V_i(s_0, s_i, \mathbf{s}_{-i}) = \int_{\underline{b}}^{\bar{b}} \frac{1}{\sum_{j\neq i} \psi(b|s_0, s_j, \mathbf{s}_{-j})} dG^{(i)}(b|s_0, s_i, \mathbf{s}_{-i}) + \beta \sum_{j\neq i} \left(\int_{\underline{b}}^{\bar{b}} \left[1 + \frac{\psi(b|s_0, s_i, \mathbf{s}_{-i})}{\sum_{k\neq i} \psi(b|s_0, s_k, \mathbf{s}_{-k})} \right] \cdot dG^{(j)}(b|s_0, s_i, \mathbf{s}_{-i}) \right) V_i(\omega_j), \quad (10)$$

where $G^{(i)}(b|s_0, s_i, \mathbf{s}_{-i})$ denotes the probability that bidder i wins with a bid less than or equal to b given state variables $(s_0, s_i, \mathbf{s}_{-i})$.

The representation of the value function in Proposition 1 consists of two parts: the first part accounts for bidder i's current-period expected unit payoff; the second part accounts for bidder i's expected future unit payoffs. The proof of Proposition 1 is based on two observations. First, the winning probability can be written as a function of the distribution of bids of competitors. Second, the first-order condition of optimal bids provides an explicit expression of a bidder's private valuation in terms of the bidder's equilibrium bids and their distribution (see Equation (9)). The detailed proof of Proposition 1 is provided in Appendix A.2.

Based on Proposition 1, we can use a numerical method to approximate the value function. To start, we select a grid of state vectors $\hat{\mathbf{S}} = (\mathbf{s}_1, \ldots, \mathbf{s}_T)$ from the distribution of observed states and solve Equation (10) for each bidder on this grid. We restrict the range of the transition function ω to $\hat{\mathbf{S}}$ by defining a pseudo-transition function $\hat{\omega}(\mathbf{a}, \mathbf{s}) = \{\mathbf{s} \in \hat{\mathbf{S}} \mid \mathbf{s} \text{ is closest to } \omega(\mathbf{a}, \mathbf{s})\}$. For each $\mathbf{s} \in \hat{\mathbf{S}}$, we calculate the expected current-period unit payoff

$$U_i(\mathbf{s}) = \int_{\underline{b}}^{\bar{b}} \frac{1}{\sum_{j\neq i} \psi(b|s_0, s_j, \mathbf{s}_{-j})} dG^{(i)}(b|s_0, s_i, \mathbf{s}_{-i}) \quad (11)$$

and the $1 \times T$ vector of transition probabilities of the events that the states $(\mathbf{s}_1, \ldots, \mathbf{s}_T)$ are reached when bidder j wins,

$$W_{ij}(\mathbf{s}) = \int_{\underline{b}}^{\bar{b}} \left[1 + \frac{\psi(b|s_0, s_i, \mathbf{s}_{-i})}{\sum_{k\neq i} \psi(b|s_0, s_k, \mathbf{s}_{-k})} \right] dG^{(j)}(b|s_0, s_j, \mathbf{s}_{-j}) \times (\mathbf{1}_{\hat{\omega}(\mathbf{a},\mathbf{s})=\mathbf{s}_1}, \ldots, \mathbf{1}_{\hat{\omega}(\mathbf{a},\mathbf{s})=\mathbf{s}_T}). \quad (12)$$

Given this definition of U_i and W_{ij}, we can rewrite the value function in Equation (10) as

$$V_i(\mathbf{s}) = U_i(\mathbf{s}) + \beta \sum_{j\neq i} W_{ij}(\mathbf{s})(V_i(\mathbf{s}_1), \ldots, V_i(\mathbf{s}_T))^{tr}, \quad (13)$$

where $(\cdot)^{tr}$ denotes the transpose operator. We can rewrite Equation (13) in matrix form

$$[I - \beta W_i]V_i = U_i, \quad (14)$$

where I denotes the T-dimensional identity matrix, W_i denotes the $T \times T$ transition matrix given by $[\sum_{j\neq i} W_{ij}(\mathbf{s}_1), \ldots, \sum_{j\neq i} W_{ij}(\mathbf{s}_T)]^{tr}$, V_i denotes the vector $(V_i(\mathbf{s}_1), \ldots, V_i(\mathbf{s}_T))^{tr}$, and U_i denotes the vector $(U_i(\mathbf{s}_1), \ldots, U_i(\mathbf{s}_T))^{tr}$. The value function can be calculated by

$$V_i = [I - \beta W_i]^{-1} U_i. \quad (15)$$

For points outside the grid $\hat{\mathbf{S}}$, we can approximate the value function using a quadratic polynomial (see Judd 1998 for more details about the approximation).

Despite the theoretical attractiveness, however, the two-step estimation procedure described herein is computationally expensive and even infeasible because the state space will grow exponentially as the number of bidders increases. Further, because there are typically more than 100 bidders competing in each period and bidders have only a few seconds to make their bidding decision, it is unlikely a bidder can monitor and track each competitor's state in each period. Given these considerations, we draw upon the *asymptotic environmental similarity* property[19] of large markets (Swinkels 2001) and assume that any bidder i in these auctions makes bidding decisions only based on bidder i's own state s_i (i.e., the quantity purchased so far) and the market state s_0 (i.e., the current supply). In other words, we can remove $\mathbf{s}_{-i}$ from Equations (9) and (10). This helps to significantly reduce the computational cost. It is worth mentioning that such approximation shares the same spirit as the one described in Yao and Mela (2008).

5.2. Monte Carlo Simulation

To demonstrate the effectiveness of our estimation method, we conduct Monte Carlo experiments on a simulated auction market. In what follows, we first describe the data-generating process and then present the estimation results on the simulated data.

5.2.1. Data-Generating Process. Consider a simulated auction market with 100 risk-neutral, ex ante symmetric bidders competing for a homogeneous product. We focus on two scenarios that correspond to two different market conditions:[20] scarcity and balanced. The scarcity condition refers to the situation when demand exceeds supply substantially in the market and the balanced condition when demand is roughly equal to supply.

For illustration, we simplify the data-generating process as follows. The minimum purchase quantity is set to one unit for each round. Each bidder demands a maximum of five units in each auction. Further, we assume the total supply of the product is 100 units under the scarcity condition and 400 units under the balanced condition. For both conditions, we consider the grid of state vectors $\{s_{00}, s_{01}, s_{02}, s_{03}, s_{04}, s_{10}, s_{11}, s_{12}, s_{13}, s_{14}\}$, where s_{00} denotes the state in which the supply is low (remaining units less than half of the initial supply) and bidder i has not won anything yet, s_{0k} denotes the state in which the supply is low but bidder i has purchased k unit(s), s_{10} denotes the state with high supply (remaining units larger than or equal to half of the initial supply) and bidder i has not won anything, and s_{1k} denotes the state in which the supply is high and bidder i has purchased k unit(s) in previous rounds.

Given these states, we assume bidders' private valuations v are random draws from truncated normal distributions defined at $[0, 100]$. Specifically, we assume $F(v|s_{00}) = \mathbf{N}(60,5)$, $F(v|s_{01}) = \mathbf{N}(57.5,5)$, $F(v|s_{02}) = \mathbf{N}(55,5)$, $F(v|s_{03}) = \mathbf{N}(52.5,5)$, $F(v|s_{04}) = \mathbf{N}(50,5)$, $F(v|s_{10}) = \mathbf{N}(50,5)$, $F(v|s_{11}) = \mathbf{N}(47.5,5)$, $F(v|s_{12}) = \mathbf{N}(45,5)$, $F(v|s_{13}) = \mathbf{N}(42.5,5)$, and $F(v|s_{14}) = \mathbf{N}(40,5)$. Bidders' equilibrium bids are calculated by solving the corresponding MDP defined on the aforementioned grid. The winning bid and purchase quantity in each round are recorded. We repeat the data-generating process and generate bidding data for 100 auctions under each of the two scenarios.

5.2.2. Estimation Results. Using the simulated bidding data, we can recover the structural parameters (i.e., mean and standard deviation for each of the truncated normal distributions) by following the two-step estimation approach described in Section 5.1. Tables 4 and 5 summarize the statistics of the estimation under the scarcity and balanced market conditions, respectively. The numbers in parentheses are standard errors of the estimates. For state s_{14}, as there are not enough observations, it is not possible to estimate the parameters (see the "NA" in the last row of Table 4).

The results from Tables 4 and 5 demonstrate that our approach yields good estimates of the model primitives under both the scarcity market condition (i.e., 100 units for sale) and the balanced market condition (i.e., 400 units for sale).

Before moving to the discussion of empirical estimation, it is worth noting that the choice of the grid is critical for the performance of our estimation approach: an overly coarse grid cannot adequately capture the conditional distributions for bidders' private valuations whereas an overly fine grid might lead to instability of the estimation if we do not have enough observations. In the next section, we discuss the choice of an appropriate grid for our empirical estimation.

6. Empirical Estimation

This section describes the empirical estimation of our dynamic structural model. Section 6.1 begins with a discussion of discretization of the state space, the estimation of number of active bidders, and the distribution of purchase quantities, which are prerequisites for the application of the two-step estimation approach. Section 6.2 reports the estimates of the bid distribution functions and the value distribution functions. Given the observed price volatility in these auctions (see Figure 4), all the estimations are done on a day-to-day basis.

6.1. Preliminaries

Our data pose several challenges for empirical estimation. To start with, although the original high-dimensional state space is reduced to a two-dimensional space, it is still very sparse in the sense that we only have observations from a limited number of states, which would lead to instability of the estimation. In light of this, we discretized the state space into a two-dimensional grid based on the distribution of bidders' purchase quantities and the lot sizes. Specifically, we discretized s_0 (the remaining quantity in the current period) into five intervals: $[0, 9]$, $[10, 19]$, $[20, 49]$, $[50, 99]$, and $[100, +\infty)$. Similarly, for any bidder i, we discretized s_i (recall that $s_i := \bar{q}_i$) into five intervals: $[0, 5], [6, 10], [11, 20], [21, 50]$, and $(50, +\infty)$. Table 6 provides an overview of the discretization and the notations for the states in the grid.

Next, unlike prior works on English auctions or first-price, sealed-bid auctions (e.g., Jofre-Bonet and Pesendorfer 2003), in the DFAs, only winning bids are revealed and recorded. In other words, we need to estimate the distribution of bidders' equilibrium bids from the empirical distribution of winning bids. This can be done by using the order statistics (Paarsch et al. 2006); that is, $G^{\text{win}}(b^{\text{win}}|\mathbf{s}) = G(b|\mathbf{s})^{n(\mathbf{s})}$. Here, G^{win} denotes the cumulative distribution function of the winning bids b^{win}, which can be directly estimated from empirical data; $n(\mathbf{s})$ denotes the number of bidders in state $\mathbf{s}$, which requires the knowledge of the total number of active bidders in the market on

Table 4. Estimation Results of Monte Carlo Simulation Under Scarcity Condition

	True values		Estimated values	
	Mean	Standard deviation	Mean (Standard error)	Standard deviation (Standard error)
$F(v\|s_{00})$	60	5	59.11 (0.13)	4.83 (0.12)
$F(v\|s_{01})$	57.5	5	57.04 (0.12)	5.47 (0.22)
$F(v\|s_{02})$	55	5	54.77 (0.34)	6.03 (0.66)
$F(v\|s_{03})$	52.5	5	52.08 (0.78)	5.72 (1.22)
$F(v\|s_{04})$	50	5	51.34 (1.19)	5.91 (1.71)
$F(v\|s_{10})$	50	5	49.32 (0.09)	5.06 (0.13)
$F(v\|s_{11})$	47.5	5	47.64 (0.26)	5.80 (0.63)
$F(v\|s_{12})$	45	5	45.75 (1.02)	5.41 (1.43)
$F(v\|s_{13})$	42.5	5	44.21 (1.05)	6.12 (1.18)
$F(v\|s_{14})$	40	5	NA	NA

Note. When there are insufficient observations for a given state, it is not possible to estimate the corresponding parameters; thus we set the values to NA (i.e., not applicable).

Table 5. Estimation Results of Monte Carlo Simulation Under Balanced Condition

	True values		Estimated values	
	Mean	Standard deviation	Mean (Standard error)	Standard deviation (Standard error)
$F(v\|s_{00})$	60	5	58.66 (1.53)	3.83 (1.11)
$F(v\|s_{01})$	57.5	5	56.51 (0.16)	4.66 (0.17)
$F(v\|s_{02})$	55	5	54.32 (0.07)	5.03 (0.13)
$F(v\|s_{03})$	52.5	5	52.05 (0.07)	5.44 (0.16)
$F(v\|s_{04})$	50	5	49.73 (0.10)	5.96 (0.27)
$F(v\|s_{10})$	50	5	49.28 (0.07)	4.95 (0.12)
$F(v\|s_{11})$	47.5	5	47.03 (0.07)	5.48 (0.20)
$F(v\|s_{12})$	45	5	44.71 (0.09)	5.86 (0.2)
$F(v\|s_{13})$	42.5	5	42.44 (0.31)	6.46 (0.69)
$F(v\|s_{14})$	40	5	41.27 (1.52)	6.67 (1.57)

a daily basis. Following Assumption 1 in Section 4.1, the number of active bidders in the market is distributed binomially with two parameters: the total number of potential bidders registered at the market and the participation probability. Although the former number is not directly available, we can use the number of unique winner IDs (388) observed from the transaction data as the proxy given that the data covers an extended period.[21] Using the proxy of total number of potential bidders and the daily log-in data, we can estimate the participation probability (44.7%) and thereby the number of active bidders each day.[22]

Further, to characterize the state transitions, we also need to know the distribution of bidders' purchase quantities upon winning. Given Assumption 3 in Section 4.1, we only need to estimate the parameters characterizing the zero-inflated negative binomial distribution. Using maximum likelihood, we obtain the estimates of the parameters, and the zero-inflated density can be written as $f_{\text{ZINB}}(y) = 0.1 \cdot \mathbf{I}_0(y) + 0.9 \cdot f_{\text{NB}}(y; \theta_1, \theta_2), \theta_1 = 4, \theta_2 = 5$, where $\mathbf{I}_0$ is the indicator function and $f_{NB}(y; \theta_1, \theta_2)$ is the density function of negative binomial distribution; that is,

$$P(Y = y) = \frac{\Gamma(y + \theta_1)}{y!\Gamma(\theta_1)} \left(\frac{\theta_1}{\theta_2 + \theta_1}\right)^{\theta_1} \left(\frac{\theta_2}{\theta_2 + \theta_1}\right)^{y}, \quad y = 0, 1, \ldots$$

and $\Gamma(\cdot)$ is the gamma function.

6.2. Main Estimation Results

Following the work of Athey et al. (2011) and Jofre-Bonet and Pesendorfer (2003), we model the bid

Table 6. Overview of the Discretization of State Space

	A bidder's cumulative purchase quantity				
Remaining quantity	[0, 5]	[6, 10]	[11, 20]	[21, 50]	> 50
[0, 9]	s_{00}	s_{01}	s_{02}	s_{03}	s_{04}
[10, 19]	s_{10}	s_{11}	s_{12}	s_{13}	s_{14}
[20, 49]	s_{20}	s_{21}	s_{22}	s_{23}	s_{24}
[50, 99]	s_{30}	s_{31}	s_{32}	s_{33}	s_{34}
≥100	s_{40}	s_{41}	s_{42}	s_{43}	s_{44}

distribution as a Weibull distribution because of its versatility in modeling uncertain quantities with nonnegative values (Rinne 2008). Table 7 summarizes the estimation results of the Weibull parameters (the shape parameters placed above the scale parameters) for the first five days. If the number of winning bids corresponding to a given state are insufficient (e.g., we rarely observe winning bids from bidders with a cumulative purchase quantity of more than 50 when the remaining units were larger than 100), we cannot obtain any estimates of the parameters, and thereby, we set the parameter values to "NA" in the table. The estimation results for day 6 to day 30 can be found in Table B.1.

The empirical estimation results from Table 7 indicate high variability of winning bids across different states. This is not unexpected given that both bidders' private values and the competitive environment vary across different states. Besides, we can see that the distribution of winning bids changes from day to day. This finding is consistent with the observation from Figure 4 and reinforces the necessity to estimate the parameters in our structural model on a day-to-day basis.

With the estimated bid distribution, next, we can apply the two-step estimation approach to recover bidders' private valuations. Because we have already discretized the state space (see Table 6), we can first use Equation (14) to approximate the value functions numerically and then apply the first-order condition in Equation (9) to recover the private valuations. Because bidders' true valuations are not observable, we cannot directly evaluate the accuracy of the estimates. In light of this, we used the estimated bidders' valuations to compute bidders' best-response bids (estimated winning bids) by solving a series of MDPs defined in Equation (6) at each state in the grid and then compared the cumulative distribution of the estimated winning bids with the cumulative distribution of the observed winning bids. As an example, we depicted the comparison results on September 1 in Figure 8. We can see that the estimated distribution curve tends to be slightly more flat than the empirical distribution curve in the range of [10, 60], which results in a minor increase of the chances of observing a winning bid in the two end (i.e., [10, 20] and [40, 60]) under the estimated cumulative distribution. Nevertheless, overall, the two empirical distributions are very similar to each other.

Besides qualitative comparisons, such as Figure 8, we also performed the Kolmogorov–Smirnov two-sample test (Pratt and Gibbons 1981) for the estimation on each day to determine whether the underlying probability distributions differ. The results show that we cannot reject the null hypothesis that the estimated winning bids and the observed winning bids have the same distribution. This result provides strong support for the effectiveness of our estimation.

Table 7. Estimation Results of Weibull Parameters from Day 1 to Day 5

	s_{40}	s_{30}	s_{20}	s_{10}	s_{00}	s_{41}	s_{31}	s_{21}	s_{11}	s_{01}	s_{42}	s_{32}	s_{22}
D_1	2.1	2.3	2	NA	NA	2.9	2.9	2	NA	3.1	2.6	1.9	1.9
	23.5	22.7	21.5	NA	NA	25.6	27.1	21	NA	26.9	23.6	20.3	19.1
D_2	6.3	4.1	3.2	1.5	1.6	12.6	1.5	1.4	1.5	3.9	13.3	1.4	1.6
	36.1	31.1	29.1	19.5	18.9	38.6	20.2	13.4	18.4	22.5	39.2	15.5	16.6
D_3	5.9	1.5	4	NA	1.5	11.8	1.7	1.7	1.1	1.3	3	2.3	1.9
	32	15.9	25.9	NA	14.7	35.6	21.2	15.4	13.9	11.1	29	24.9	20
D_4	8.1	1.4	2.8	2.6	1.5	7.8	6.4	6.2	2.1	1.7	8.2	27.4	1.9
	36.3	17.1	20.9	23.8	16	35.5	33.6	26.3	19.8	14	35.1	37.1	18.1
D_5	4.2	6.3	1.9	12.8	3.6	3.2	3.5	2.2	3.2	4.9	2	2.8	1.9
	27.4	29.8	16.4	5.5	26.4	25.2	22.6	15.6	14.2	17.7	18.6	22.1	15.2

	s_{12}	s_{02}	s_{43}	s_{33}	s_{23}	s_{13}	s_{03}	s_{44}	s_{34}	s_{24}	s_{14}	s_{04}
D_1	2.3	3.6	2.4	1.9	1.9	1.5	1.9	4.2	3.2	2	NA	1.9
	24.2	27.6	24.6	18.7	18.1	19.8	18.3	35.5	27.8	20.7	NA	17.8
D_2	1.6	1.2	21.5	1.7	1.5	3.8	1.2	NA	NA	1.8	NA	1.7
	15.8	9.3	40.2	16.9	13.7	28.1	8.4	NA	NA	19.5	NA	15.6
D_3	3.1	1.3	3.6	2	1.9	3	3.4	NA	2.6	2	1.8	1.8
	24.4	9.8	29.3	20.4	17.9	23.6	17.8	NA	24.6	22.4	20.5	15.4
D_4	3.4	1.7	NA	NA	1.6	NA	2	NA	NA	2.1	NA	7.7
	25.3	14.1	NA	NA	15.6	NA	15.1	NA	NA	18.8	NA	16.6
D_5	1.2	1.3	2.7	3.6	17.5	NA	1.5	NA	NA	4.4	NA	3.3
	10.5	8.4	23	24.4	7.6	NA	9.3	NA	NA	24.5	NA	4.7

Notes. The estimates of the shape parameters are placed above the scale parameters. When there are insufficient observations for a given state, it is not possible to estimate the corresponding parameters; thus we set the values to NA (i.e., not applicable).

Figure 8. Comparison of Cumulative Distributions of the Estimated and Observed Winning Bids (September 1)

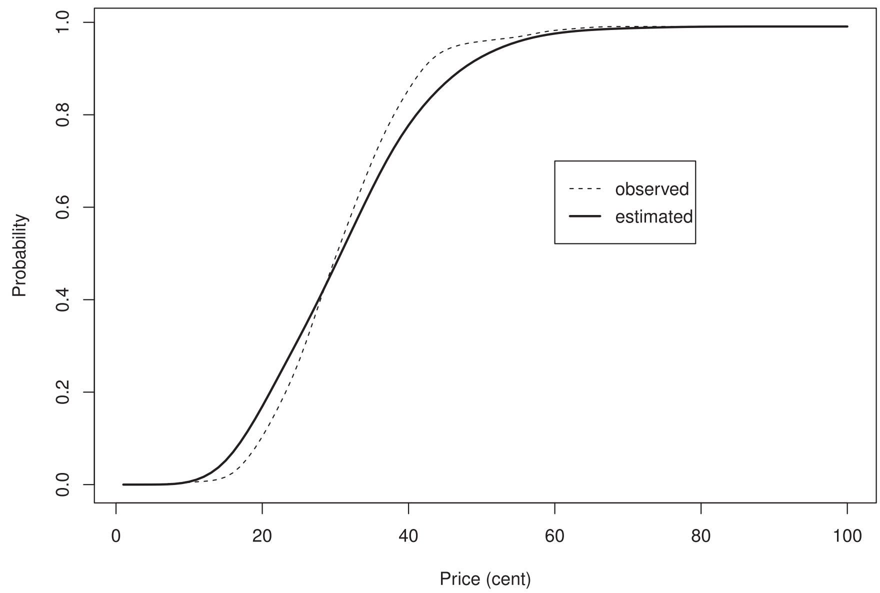

7. Policy Experiment

Although the parameter estimates from the previous section provide useful insights about bidder behavior, they offer limited guidance to auctioneers' decision making. In this section, we focus on the unique design parameter, namely, minimum purchase quantity, in the DFAs and evaluate the effects of different policy interventions related to this parameter.

7.1. Minimum Purchase Quantity

Minimum purchase quantity has a strong impact on the bidding dynamics and outcomes. On one hand, increasing minimum purchase quantity can speed up the auction process by forcing bidders to buy more in each round. On the other hand, given that bidders have nonincreasing marginal values, a large minimum purchase quantity may lead to less competition and a lower price. It is noteworthy that the conflicting effects of minimum purchase quantity are qualitatively analogous to the effects of lot size in sequential auctions. Specifically, prior research has shown that a large lot size often has a negative impact on the closing price whereas a small lot size leads to excess holding and administrative costs (Pinker et al. 2010). In this regard, the decision of minimum purchase quantity can also be considered as a generalized lot-sizing problem.

Currently, auctioneers mainly rely on their intuition and experience to decide the minimum purchase quantity in each round. Typically, they set a relatively low minimum purchase quantity at the beginning and gradually increase it as the auction proceeds. It is important to empirically determine whether this *rule of thumb* yields desirable outcomes.

7.2. Policy Simulation

We consider five alternative auction designs in which minimum purchase quantities are set as follows: (1) minimum purchase quantity is fixed to one, (2) minimum purchase quantity is fixed to three, (3) minimum purchase quantity is fixed to five, (4) minimum purchase quantity is fixed to seven, and (5) minimum purchase quantity is fixed to nine.[23] For each of the alternative designs, we simulated bidders' private valuations using the estimated conditional distributions and then calculated their optimal bids. Winners' purchase quantities were drawn from the estimated zero-inflated negative binomial distribution.

We repeated the simulation for 100 iterations based on which we compared the performance of the five alternative designs with the observed design. Note that, during the simulation process, we matched the supply in each auction to the supply in the observed data. Similarly, the auction parameters, such as reserve price and starting price, were matched to the observed data. The only thing we manipulated was the minimum purchase quantity in each round.

Given the nature of the market under consideration, we choose two performance indicators, revenue and turnaround, with which the latter is measured by the total number of rounds taken to finish the auctions. Table 8 provides an overview of the performance of each design (the numbers in the parentheses are the standard errors of the estimates). Here, the observed design refers to the way the minimum purchase quantities were set in those auctions, and we can directly calculate the total revenue and number of rounds from the transaction data. For each alternative design, we simulated bidders'

Table 8. Performance Evaluation of Different Auction Designs

	Total revenue (in euros)	Number of rounds
Observed design	920,116	5,536
Alternative design 1 (minimum purchase quantity = 1)	1,053,874 (10,518)	7,682 (33)
Alternative design 2 (minimum purchase quantity = 3)	1,004,604 (9,526)	6,296 (34)
Alternative design 3 (minimum purchase quantity = 5)	945,941 (8,925)	5,362 (25)
Alternative design 4 (minimum purchase quantity = 7)	886,203 (11,057)	4,667 (16)
Alternative design 5 (minimum purchase quantity = 9)	820,157 (11,759)	4,158 (16)

Note. The numbers displayed in the parentheses from row 2 to row 6 (Alternative design 1–5) are standard errors of the estimates.

valuations and purchase quantities, and we estimated the winning price; thus, the total revenue and number of rounds were subject to variability.

To start with, we can see alternative design 1 (i.e., minimum purchase quantity always set to one) generates the highest expected revenue: if we compare it with the observed design, the revenue is expected to increase by approximately 14.5%. However, such increase comes at a high cost of turnaround rate: compared with the observed design, the total number of transactions increased by 38.7%. Given the high perishability of the products and the tight auction schedule, such alternative design would not be a good option. Similarly, alternative design 2 is expected to increase the revenue (9.2%) at a significant cost of turnaround (expected number of rounds increases by 13.7%), which again makes it undesirable.

Alternative design 3 seems to be a good option: it generates 3% higher revenue while maintaining a high turnaround rate. In fact, the expected number of rounds to finish the auctions even decreases by 3%. Such observation suggests that (i) a moderate, constant minimum purchase quantity can yield comparable outcomes as the rule of thumb, and (ii) auctioneers' current practice in setting the minimum purchase quantity has ample room for improvement. Finally, if we compare the performance of alternative designs 4 and 5 with the observed design, we can see the expected revenue drops by approximately 3.7% and 11% whereas the turnaround speed increases by 15.7% and 25%. Note that the potential improvement in turnaround not only saves the operational and administrative costs, but also offers time slots for auctioning extra products that can be similar or different. At the end, if the partial loss of revenue from the previous auctions can be compensated or even surpassed by the revenue generated from auctioning extra products, such alternative design could be a viable option.

Based on these observations, a natural question is that of how we can derive the *optimal* values of minimum purchase quantities in such a dynamic environment. In the next section, we explore this optimization problem in detail.

8. Dynamic Optimization of Auction Design Parameters

The auctioneers in the DFAs represent the growers, and their main objective is to realize high revenue. Besides, it is important to achieve a quick turnaround because flowers are perishable goods. By controlling key auction parameters, such as starting price, minimum purchase quantity, and reserve price, auctioneers can influence the bidding dynamics as well as the outcomes. However, currently these parameters are not optimized because auctioneers cannot adequately process all the market information under the extreme time pressure. Instead, they rely on their intuition and experience to decide how to set these key auction parameters.

A promising way to address these limitations is to augment auctioneers' capabilities with high-performance decision support systems (Bichler et al. 2010). To be useful, these systems must be able to (i) make good predictions of future states (e.g., the number of bidders, the winning prices, and purchase quantities in the upcoming auctions) and (ii) optimize the key auction parameters based on the predictions. In this section, we discuss how to apply our dynamic structural model to optimize the key auction parameters. As with the policy evaluation in Section 7, we focus on the choice of the minimum purchase quantity in sequential rounds.[24]

8.1. Structurally Based Prediction

A number of methods have been used for price prediction in online auctions (Bapna et al. 2008, Wang et al. 2008). Despite the differences in the functional specifications, it is widely recognized that prediction models must be grounded, in one way or another, on market information (Muth 1961).

Based on the theoretical and empirical analysis of the DFAs earlier in this paper, we propose a novel prediction method that incorporates the structural properties of the underlying auction model while exploiting historical market data. Specifically, given the structural estimates of bidders' valuations and an auction policy (i.e., the set of key auction parameters), we can predict future auction outcomes by repeatedly solving the MDP

Figure 9. An Illustration of the Structurally Based Prediction Method

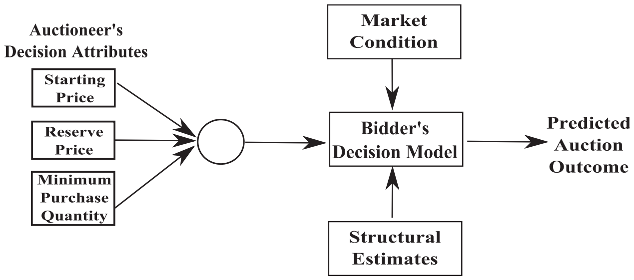

defined in Section 4.2 for each bidder. An important characteristic of this method is that it accounts for bidders' potential strategic responses to any policy changes (e.g., dynamic changes in minimum purchase quantity) and, thus, is immune from the Lucas critique (Lucas 1976, Chintagunta et al. 2006).

Figure 9 illustrates the general idea of our prediction method. Here, the market condition refers to the collection of market characteristics such as the supply of the product and the number of potential bidders that will affect the competition.

To evaluate the performance of such a prediction method, we split the total number of auctions on a given day into two parts evenly. The first half serves as the training set and is used to estimate the distribution of bidders' private valuations. The second half serves as the test set in which we compare the cumulative probability distribution of the predicted winning prices and the observed prices using Kullback–Leibler (KL) divergence (Kullback and Leibler 1951). The KL divergence measure can be interpreted as how much additional information is needed to achieve optimal prediction. A KL divergence less than or close to one suggests the predicted distribution matches well with the observed distribution (Ketter et al. 2012). In our case, the KL divergence measures are consistently less than two (ranging from 0.22 to 1.72) on 24 days, and for the remaining six days, the highest KL divergence values are still less than three. This indicates a good fit between the predicted and the actual winning prices.

8.2. Optimal Choices of Minimum Purchase Quantities

Auctioneers have to choose the minimum purchase quantity in each round. Note that the choice of minimum purchase quantity in the current round not only affects the purchase quantity in the current round, but also the competition in future rounds. In light of this, we can formulate the auctioneer's sequential decision problem as the following MDP:

- **States:** $\{s_0^t\}$, where $s_0^t := \bar{q}^t$.
- **Actions:** Given state s_0^t, the set of actions for the auctioneer is $\{a_0^t := \underline{q}^t | \underline{q}^t = 1, 2, \ldots, \bar{q}^t\}$.
- **Transitions:** The transition function, denoted by $\omega_0 : \mathbf{A}_0 \times \mathbf{S}_0 \rightarrow \mathbf{S}_0$, is a deterministic function of the states and actions. In particular, we have $\omega_0(a_0^t, s_0^t) = \bar{q}^t - q^t$, where q^t is the winners' purchase quantity in the current round.
- **Rewards:** Given state s_0^t, the auctioneer's expected payoff is given by

$$E\left[\sum_{\tau=t} \beta^{\tau-t} \max_i(b_i^\tau) \cdot q_i^\tau | s_0^t, \underline{q}^\tau; f_{\text{ZINB}}\right]. \tag{16}$$

The expectation in Equation (16) is taken over the realization of the winning price and purchase quantity in the current period as well as the future realizations of state variables, winning prices, and purchase quantities. Because we do not know the exact operational or administration costs associated with extra transaction rounds, we experiment with different values of the discount factor between 0.5 and 1.0 with a step size of 0.1; that is, $\beta = 0.5, 0.6, \ldots, 1$. Using a similar simulation procedure as in the policy evaluation, we compared the performance of the optimized designs with the observed design on the test data.

Figure 10 displays the 6×5 matrix of subplots of the auction outcomes under the optimized designs and observed designs over the 30 days. The y-axis shows the total revenue generated from the auctions in the test set, and the x-axis shows the number of rounds taken to finish the auctions. As Figure 10 demonstrates, for 24 out of the 30 days, the optimized designs generate higher expected revenues than the observed designs while maintaining a comparable turnaround. The observed performance gap between the optimized choices and the ones observed in the data suggests that auctioneers' current practice in setting minimum purchase quantities has ample room for improvement; that is, the rule of thumb—starting with a low minimum purchase quantity and gradually increasing it—does not necessarily produce the most desirable outcomes. Specifically, even if the general idea of such a heuristic is applicable in some cases, the auctioneers do not choose the right time to adjust the minimum purchase quantity. Our two-step optimization framework provides clear reference points for auctioneers in making the trade-off between revenue and turnaround. For example, for September 1, the auctioneer can achieve approximately 15% higher revenue (from 23,543 to 27,020) by following the optimized sequence of minimum purchase quantities should an extra six rounds of transactions (5%) be acceptable (when the discount factor β is set to 0.9). Further, the determination of the minimum purchase quantities is very flexible: auctioneers can leverage their experience to tailor the choice of the discount function in Equation (16) to the specific market conditions (Ketter et al. 2012). In this regard, our optimization framework shares the same

Figure 10. (Color online) Optimization Results on the Test Data

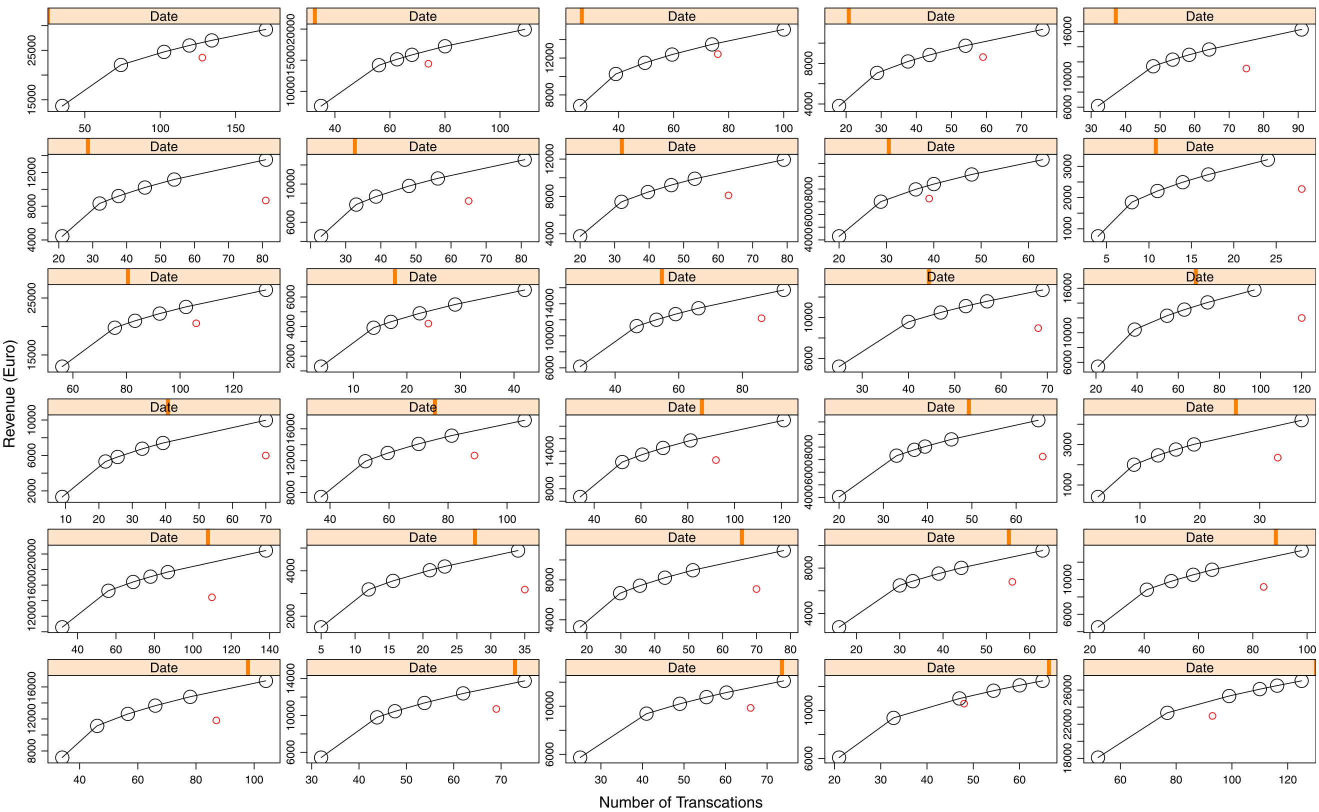

Notes. The six rows stand for the six transaction weeks, and the columns indicate the weekdays. The small (red online) circle denotes the observed outcome (in terms of total revenue and the number of rounds) whereas the big (black) circles on the curves stand for the optimized outcomes under different values of the discount factor; the rightmost point denotes the outcome with discount factor set to one (no discount).

spirit of the adaptive design proposed by Pardoe et al. (2010), which incorporates prior knowledge of bidder behavior to enhance the setting of the auction parameters.

It is worth mentioning that our current optimization framework has not taken into account bidders' learning effect; that is, bidders may also form and update their beliefs about auctioneer's dynamic decisions over time. Although a full treatment of designing and implementing effective decision support tools in the complex, dynamic market of DFAs is beyond the scope of the current study, the optimization framework described herein provides a useful starting point.

9. Conclusion

In this paper, we develop a dynamic structural model for multiunit sequential Dutch auctions in a B2B market. To the best of our knowledge, this is the first paper that adopts a structural econometric approach and systematically examines the dynamic decision-making problem in these complex auctions.

Our work makes important contributions to both the theory and practice of auction design. From the theoretical perspective, we extend the current structural empirical auction literature by modeling competitive bidding in sequential auctions in which bidders have multiple purchase opportunities over time and can fulfill multiunit demand in each round. The multiunit, sequential nature of these auctions poses both econometric and computational challenges for the structural estimation. To address these challenges, we adapt the well-known two-step estimation methods by exploiting the characteristics of the market processes and market participants (particularly bidders). This allows us to recover the structural parameters in a computationally efficient manner. To this point, it is worth noting that although individual bidders may deviate from the *optimizing* behavior quite often because of cognitive or computational constraints in this competitive environment, the empirical results indicate that at the aggregate level, our dynamic structural model well captures the essence of how bidders interact with each other under different market conditions.

From the practical point of view, this research provides valuable insights into auctioneers' decision making concerning the key auction parameters. As Klemperer (2002, p.184) points out, "auction design is *not*, 'one size fits all.'" Although extensive progress has been made in practical designs of single, isolated auctions, there is limited understanding of the bidding dynamics in sequential auctions, and the literature concerning the design issues of these auctions is still in its infancy. In the current study, we conduct a series of policy simulations based on the structural estimation of bidders' values, and the results indicate that the decisions of auctioneers in setting the key auction parameters (e.g., minimum purchase quantity) are often suboptimal. We then propose a novel, structurally based optimization framework that can account for bidders' strategic reactions toward policy interventions and guide auctioneers to choose minimum purchase quantities appropriately under different market conditions. In this regard, our paper shares the same spirit of Adomavicius et al. (2009) and Ketter et al. (2012) by providing an effective computational tool to facilitate decision making in a complex environment. However, unlike previous studies that mainly rely on simulated data to evaluate the performance of proposed tools, we demonstrate the effectiveness of our optimization framework using real transaction data, which makes it very promising for practical use. It is worth mentioning that, although the current study is specifically geared toward the DFAs, our optimization framework can be applied to other time-critical B2B markets, particularly the auction markets for perishable goods, such as fish, fruits, and vegetables. Some of the famous markets that have adopted the multiunit sequential auction mechanism include Ota Market (Japan's largest flower and vegetable wholesale market), Kunming International Flora Auction Trading Centre (China's largest flower-trading market), and Pan European Fish Auctions (the first European "virtual" B2B market in the fishing industry). The results and findings from the current study provide useful implications to both academics and practitioners in reengineering these complex markets (Kambil and van Heck 1998).

The current study bears several limitations that, nevertheless, provide possible directions for future research. First, our current model does not account for the potential complementarity of products. In reality, many bidders, especially small retailers, may be interested in purchasing an assortment of products on behalf of their customers. This means that bidders' valuation of certain products may be conditional on the purchase of other products. Such a package bidding feature would pose considerable challenges to the modeling and estimation of these auctions (Boutilier et al. 1999). Unfortunately, as we do not have access to bidders' order books, it is very difficult to conduct empirical studies in this aspect. In addition, given the focus of this study on multiunit sequential auctions, we have chosen a homogeneous sample and performed nonparametric estimation. It would be nice to have a comprehensive, parametric model that accounts for product heterogeneity across different auctions although such a model would inevitably require even stronger assumptions than our current model. Second, following prior works in structural modeling of sequential auctions, we have assumed that conditional on their states, bidders' valuations are randomly drawn from the same distribution. Although this assumption ensures our model is identifiable, it might be violated in reality. Specifically, when there exists significant heterogeneity in the downstream market, the estimates from our current model could be biased. To better understand the downstream market, we have tried very hard to approach wholesalers as well as retailers in the DFAs to obtain their order books. Once we have the information about their order books, we can work on extensions of our current model to account for observed heterogeneity. Third, as bidders are allowed to purchase multiple units in each round, we assume bidders' purchase quantities are strictly exogenous to ensure the tractability of the dynamic system. An exciting and challenging avenue for future research would be incorporating purchase quantity as an endogenous decision variable in the dynamic structural model. Finally, as already mentioned in Section 8.2, our current structurally based optimization framework does not account for bidders' learning regarding the dynamic setting of minimum purchase quantities. Future research may consider using computational platforms (Bichler et al. 2010) to explore the impact of this factor.

Acknowledgments

The authors thank the department editor, associate editor, and reviewers for insightful comments and suggestions. This research would not have been possible without the collaboration of Royal FloraHolland via the Veilen met Advies project. The first author also acknowledges the support of Erasmus Research Institute of Management (ERIM) via the Talent Placement Program.

Appendix A. Mathematical Proofs

Appendix A.1. First-Order Condition

The probability that bidder i wins the subauction at state $(s_0, s_i, \mathbf{s}_{-i})$ with bid b, denoted by $P(i\,\text{wins}|b, s_0, s_i, \mathbf{s}_{-i})$, can be written as

$$\prod_{j\neq i} G(b|s_0, s_j, \mathbf{s}_{-j}), \tag{A.1}$$

and the corresponding probability density function is

$$\sum_{j\neq i}\prod_{k\neq i,j} G(b|s_0, s_k, \mathbf{s}_{-k})g(b|s_0, s_j, \mathbf{s}_{-j}). \tag{A.2}$$

The probability that bidder $j, j\neq i$ wins when bidder i bids b, $P(j\,\text{wins}|b, s_0, s_i, \mathbf{s}_{-i})$, is given by

$$\int_{\underline{b}}^{b} -g(x|s_0, s_j, \mathbf{s}_{-j})\prod_{k\neq i,j} G(x|s_0, s_k, \mathbf{s}_{-k})dx. \tag{A.3}$$

The first-order condition of equilibrium bids yields

$$\begin{aligned}&(v-b)(P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i}))'-P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})\\&+\beta(P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i}))'V_i(\omega_i)\\&+\beta\sum_{j\neq i}(P(j\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i}))'V_i(\omega_j)=0.\end{aligned}\tag{A.4}$$

Substituting Equations (A.1)–(A.3) into Equation (A.4), we have

$$\begin{aligned}&(v-b)\sum_{j\neq i}\prod_{k\neq i,j}G(b|s_0,s_k,\mathbf{s}_{-k})g(b|s_0,s_j,\mathbf{s}_{-j})-\prod_{j\neq i}G(b|s_0,s_j,\mathbf{s}_{-j})\\&+\beta V_i(\omega_i)\sum_{j\neq i}\prod_{k\neq i,j}G(b|s_0,s_k,\mathbf{s}_{-k})g(b|s_0,s_j,\mathbf{s}_{-j})\\&+\beta\sum_{j\neq i}(-g(b|s_0,s_j,\mathbf{s}_{-j})\prod_{k\neq i,j}G(b|s_0,s_k,\mathbf{s}_{-k})V_i(\omega_j))=0.\end{aligned}\tag{A.5}$$

Dividing both sides of Equation (A.5) by $\prod_{k\neq i}G(b|s_0,s_k,\mathbf{s}_{-k})$ and rearranging it yields

$$\begin{aligned}&(v-b)\sum_{j\neq i}\frac{g(b|s_0,s_j,\mathbf{s}_{-j})}{G(b|s_0,s_j,\mathbf{s}_{-j})}-1+\beta V_i(\omega_i)\sum_{j\neq i}\frac{g(b|s_0,s_j,\mathbf{s}_{-j})}{G(b|s_0,s_j,\mathbf{s}_{-j})}\\&+\beta\sum_{j\neq i}\frac{-g(b|s_0,s_j,\mathbf{s}_{-j})}{G(b|s_0,s_j,\mathbf{s}_{-j})}V_i(\omega_j)=0.\end{aligned}\tag{A.6}$$

We can obtain Equation (9) by substituting the auxiliary function (e.g., Equation (8)) into Equation (A.6).

Appendix A.2. Proof of Proposition 1

To prove Proposition 1, we first derive the expectation of the value function. Substituting Equation (4) into Equation (6) and expressing the sum of discounted future payoffs with respect to all possible state transitions, we can obtain

$$\begin{aligned}V_i(s_0,s_i,\mathbf{s}_{-i})=&\int_{\underline{v}}^{\bar{v}}\Big[(v-b)P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})\\&+\beta\sum_{j=1}^{N}P(j\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})V_i(\omega_j)\Big]\\&\cdot f(v|s_0,s_i,\mathbf{s}_{-i})dv.\end{aligned}\tag{A.7}$$

Next, we substitute Equation (9), that is, the first-order condition for optimal bids, into Equation (A.7). After rearrangement, we can obtain

$$\begin{aligned}V_i(s_0,s_i,\mathbf{s}_{-i})=&\int_{\underline{v}}^{\bar{v}}\frac{1}{\sum_{j\neq i}\psi(b|s_0,s_j,\mathbf{s}_{-j})}P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})\\&\cdot f(v|s_0,s_i,\mathbf{s}_{-i})dv\\&+\beta\int_{\underline{v}}^{\bar{v}}\sum_{j\neq i}\frac{\psi(b|s_0,s_j,\mathbf{s}_{-j})}{\sum_{l\neq i}\psi(b|s_0,s_l,\mathbf{s}_{-l})}V_i(\omega_j)\\&\cdot P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})f(v|s_0,s_i,\mathbf{s}_{-i})dv\\&+\beta\sum_{j\neq i}P(j\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})V_i(\omega_j).\end{aligned}\tag{A.8}$$

Consider a change of variable of integration from v to b. We have

$$db=\frac{\partial b(v)}{\partial v}dv.\tag{A.9}$$

Let $\mathscr{b}^{-1}$ denote the inverse bidding function. We have

$$F(\mathscr{b}^{-1}(b)|s_0,s_i,\mathbf{s}_{-i})=G(b|s_0,s_i,\mathbf{s}_{-i}).\tag{A.10}$$

Taking the derivative on both sides of Equation (A.10) yields the following relationship between the probability density function of bidders' private valuations and their equilibrium bids:

$$f(\mathscr{b}^{-1}(b)|s_0,s_i,\mathbf{s}_{-i})\cdot\frac{\partial\mathscr{b}^{-1}(b)}{\partial b}=g(b|s_0,s_i,\mathbf{s}_{-i}),\tag{A.11}$$

where

$$\frac{\partial\mathscr{b}^{-1}(b)}{\partial b}=\frac{1}{\partial b(v)/\partial v}.\tag{A.12}$$

Combining Equations (A.9), (A.11), and (A.12), we have

$$f(v|s_0,s_i,\mathbf{s}_{-i})dv=g(b|s_0,s_i,\mathbf{s}_{-i})db.\tag{A.13}$$

Note that, given the state $(s_i,\mathbf{s}_{-i})$, the probability that bidder i wins with bid b is the probability that b is the highest bid; that is,

$$P(i\text{ wins}|b,s_0,s_i,\mathbf{s}_{-i})=\prod_{k\neq i}G(b|s_0,s_k,\mathbf{s}_{-k}).\tag{A.14}$$

Substituting Equations (A.13) and (A.14) into Equation (A.8) yields

$$\begin{aligned}V_i(s_0,s_i,\mathbf{s}_{-i})=&\int_{\underline{b}}^{\bar{b}}\frac{1}{\sum_{j\neq i}\psi(b|s_0,s_j,\mathbf{s}_{-j})}\prod_{k\neq i}G(b|s_0,s_k,\mathbf{s}_{-k})\\&\cdot g(b|s_0,s_i,\mathbf{s}_{-i})db\\&+\beta\int_{\underline{b}}^{\bar{b}}\left[\sum_{j\neq i}\frac{\psi(b|s_0,s_j,\mathbf{s}_{-j})V_i(\omega_j)}{\sum_{k\neq i}\psi(b|s_0,s_k,\mathbf{s}_{-k})}\right]\\&\cdot\prod_{k\neq i}G(b|s_0,s_k,\mathbf{s}_{-k})g(b|s_0,s_i,\mathbf{s}_{-i})db\\&+\beta\sum_{j\neq i}P(j\text{ wins}|s_0,s_i,\mathbf{s}_{-i})V_i(\omega_j).\end{aligned}\tag{A.15}$$

Finally, taking $\sum_{j\neq i}$ out from the second term and rearranging the second and third terms on the right-hand side of Equation (A.15), we have

$$\begin{aligned}V_i(s_0,s_i,\mathbf{s}_{-i})=&\int_{\underline{b}}^{\bar{b}}\frac{1}{\sum_{j\neq i}\psi(b|s_0,s_j,\mathbf{s}_{-j})}\\&\cdot\prod_{k\neq i}G(b|s_0,s_k,\mathbf{s}_{-k})g(b|s_0,s_i,\mathbf{s}_{-i})db\\&+\beta\sum_{j\neq i}\Bigg[\int_{\underline{b}}^{\bar{b}}\frac{\psi(b|s_0,s_i,\mathbf{s}_{-i})}{\sum_{k\neq i}\psi(b|s_0,s_k,\mathbf{s}_{-k})}\\&\cdot\prod_{k\neq j}G(b|s_0,s_k,\mathbf{s}_{-k})g(b|s_0,s_j,\mathbf{s}_{-j})db\\&+\int_{\underline{b}}^{\bar{b}}dG^{(j)}((b|s_0,s_i,\mathbf{s}_{-i})\Bigg]V_i(\omega_j).\end{aligned}\tag{A.16}$$

Recall that $G^{(i)}(b|s_0, s_i, \mathbf{s}_{-i})$ denotes the probability that bidder i wins with a bid less than or equal to b. Thus, we have

$$dG^{(i)}(b|s_0, s_i, \mathbf{s}_{-i}) = \prod_{k \neq i} G(b|s_0, s_k, \mathbf{s}_{-k}) \times g(b|s_0, s_i, \mathbf{s}_{-i})db. \tag{A.17}$$

Substituting Equation (A.17) into Equation (A.16) yields Proposition 1.

Appendix B. Estimation Results

Table B.1 summarizes the estimation results for Weibull parameters from day 6 to day 30.

Endnotes

[1] Here, by auction design, we are not only referring to different auction formats (e.g., English auction or first-price auction), but also the choices of key auction parameters, such as reserve prices, starting prices, and lot sizes.

[2] Backus and Lewis (2016) is a notable exception.

[3] More institutional details can be found at https://www.royalfloraholland.com/en.

[4] In this sense, it can be considered as a special case of bundle bidding in which the element of the bundle does not have an intrinsic value (Boutilier et al. 1999).

[5] With the estimation methods proposed by Aguirregabiria and Mira (2007) and Bajari et al. (2007), it is no longer necessary to solve for MPE when estimating the model primitives. However, solving for MPE is still needed for policy counterfactuals.

[6] In practice, the reserve price is set low enough, and we rarely observe auctioned goods being destroyed.

[7] Technically speaking, with the introduction of the online channel, bidders can choose to purchase from any one of the four auction sites within the country. However, in practice, wholesalers typically buy products from the auction site that is in close proximity to the location where they need to ship the goods.

[8] Van den Berg et al. (2001) show empirical evidence for declining price anomaly in the flower auctions; however, if we look at individual auctions, price trends are inconclusive in these multiunit sequential auctions.

[9] In multiobject auctions, bidders' information structure can be very complex in the presence of complementarity or substitutability between different objects. In particular, when L distinctive objects are auctioned, each bidder may have preferences over up to $2^L - 1$ bundles. Although a structural model for such multiobject sequential auctions would offer great insights, it is beyond the scope of this paper.

[10] We performed reduced-form analysis by regressing the winning prices on the aforementioned product characteristics; for stem length, we created two dummy variables corresponding to the stem length of 55 and 60 cm. The coefficients corresponding to the two dummy variables are statistically insignificant (p-value > 0.1).

[11] Strictly speaking, in a Dutch auction, only the winning bidder *explicitly* submits a bid. However, from the strategic perspective, it is equivalent to assume that each bidder submits a secret bid and such bid only gets revealed upon winning (Paarsch et al. 2006).

[12] Suppose there are two lots for auction on a given day. The first lot gets sold out in j rounds. The first subauction of the second lot has an index of $j + 1$.

[13] According to Milgrom and Weber (1982), the IPV assumption is appropriate for the study of nondurable consumer goods, such as flowers.

[14] Note that, as bidders are buying on order, the commonly used assumption of decreasing marginal utility does not hold in the DFA

Table B.1. Estimation Results for Weibull Parameters

	s_{40}	s_{30}	s_{20}	s_{10}	s_{00}	s_{41}	s_{31}	s_{21}	s_{11}	s_{01}	s_{42}	s_{32}	s_{22}	s_{12}	s_{02}	s_{43}	s_{33}	s_{23}	s_{13}	s_{03}	s_{44}	s_{34}	s_{24}	s_{14}	s_{04}
D_6	6.1	6.6	3.1	NA	5.8	7.6	1.8	1.6	1.7	2.5	2.6	1.6	1.7	1.6	1.4	3.5	1.8	3.3	1.1	1.7	NA	2.4	2	2.7	2.1
	13.8	23.1	15	NA	11	19.8	13.3	11.3	12.8	6.5	16.7	11.3	9.9	16.5	9.4	19.9	12.2	6.1	4.5	9	NA	16.3	12	15.2	1.6
D_7	11.9	2.1	2.2	2.1	2.6	3	2.5	5.5	3.6	1.5	NA	5	1.3	1.4	1.7	8.7	1.6	4.2	2.1	2	NA	NA	2.8	NA	1.9
	14.1	17.6	7.4	15.4	17.3	21.6	18.5	8.2	10.3	11.4	NA	8.2	3	12.2	11.6	18.6	11.9	12.2	16.1	12.2	NA	NA	9.1	NA	11.6
D_8	7.7	1.7	1.5	1.5	6.4	NA	3.1	6.5	1.9	0.8	NA	NA	1.8	9.1	1.5	NA	3.3	1.4	10.2	2.9	NA	NA	NA	NA	1.8
	27.4	19.2	11.4	14	15.7	NA	25	6.7	15.3	6.4	NA	NA	18.2	15.8	10.5	NA	25.9	11.7	14	4.3	NA	NA	NA	NA	14.1
D_9	NA	5.9	2.7	3.4	3.2	NA	31.6	1.3	2.1	3	NA	4.4	NA	NA	1.8	NA	2.2	NA	NA	1.7	NA	NA	NA	NA	4.7
	NA	33.9	25.4	23.3	20.7	NA	36.7	14.8	19.3	17.4	NA	32.3	NA	NA	16.9	NA	25.2	NA	NA	14.4	NA	NA	NA	NA	14
D_{10}	NA	NA	16.9	NA	7.2	NA	NA	7.9	NA	1.3	NA	NA	NA	27.7	3.2	NA	NA	NA	NA	4.1	NA	NA	NA	NA	NA
	NA	NA	8.1	NA	20.6	NA	NA	19.4	NA	5.4	NA	NA	NA	12.3	4.9	NA	NA	NA	NA	1.6	NA	NA	NA	NA	NA
D_{11}	3.1	2	3.6	NA	3.4	8.8	2.7	2.5	10.9	3	3.7	2.9	3.7	NA	2	25.3	2.1	2.1	NA	2.8	4.6	2.7	NA	NA	2.8
	27.8	20.6	25.4	NA	27.6	32.6	22.8	19.5	21.1	22.7	29.2	23.2	22.9	NA	18	34.8	22.1	18.9	NA	21.5	28.7	23.2	NA	NA	21.8
D_{12}	NA	9.4	4.2	NA	NA	NA	NA	NA	NA	1.7	NA	NA	0.9	NA	3.8	NA	NA	2.4	NA	3.2	NA	NA	NA	NA	NA
	NA	19.8	20.8	NA	NA	NA	NA	NA	NA	13.1	NA	NA	3.8	NA	21	NA	NA	18	NA	5.3	NA	NA	NA	NA	NA
D_{13}	4	2.1	1.7	1.4	2.3	NA	1.8	1.5	1	2.1	9	4.3	1.8	NA	1.2	8.6	2.3	1.6	NA	1.6	NA	NA	2.4	NA	9.2
	29.5	17.9	14.8	14.7	16.9	NA	16.6	13.2	7	17	34.7	25.4	17.1	NA	8.1	34.2	21.7	14.5	NA	3.8	NA	NA	23.1	NA	16
D_{14}	7	2.4	1.8	1.8	1.7	NA	6	1.7	9.7	1.3	NA	NA	1.9	2.1	3.6	NA	2.7	NA	NA	0.9	NA	NA	NA	NA	5.1
	28.9	22	15.9	18.5	14.3	NA	29	15.7	11.1	9.8	NA	NA	17.7	16.6	15.6	NA	22.8	NA	NA	34.3	NA	NA	NA	NA	7.2

Table B.1. (Continued)

	s_{40}	s_{30}	s_{20}	s_{10}	s_{00}	s_{41}	s_{31}	s_{21}	s_{11}	s_{01}	s_{42}	s_{32}	s_{22}	s_{12}	s_{02}	s_{43}	s_{33}	s_{23}	s_{13}	s_{03}	s_{44}	s_{34}	s_{24}	s_{14}	s_{04}
D_{15}	1.8	1.5	3.2	NA	1.4	1.7	1.9	1.5	1.4	1.6	1.7	5.3	1.3	1.5	1.5	3	1.5	4.5	2.7	1.6	2.1	1	13.9	1.2	2.1
	14	11.3	6.8	NA	10.6	2.3	3.1	5.4	11.8	10.2	3.8	10	5.8	12.5	4.9	12	10.2	13.8	13.3	5.2	21.8	9	2.3	10.6	8.8
D_{16}	1.8	2	8.3	2.3	8.7	2.5	4.2	4.4	1.7	1.5	NA	5.6	2.7	11.1	9.5	5.5	2.3	1.5	9.1	1.5	NA	NA	8.3	NA	9.1
	14	12.8	39.4	13.5	6.9	19	7.8	9.1	9.5	3.6	NA	16.3	13.1	5.3	39.8	6.1	16.1	8.1	39.7	8.4	NA	NA	39.4	NA	39.7
D_{17}	3.5	1.5	4	1	1.5	1.8	5.7	1.4	5.7	1.8	3	2.6	1	2	4.4	9.7	2.4	5.1	4.8	1.5	NA	3	1.1	NA	2.2
	22.6	9.6	12.3	10.3	8.6	14.5	9.5	8.1	15.3	12.4	12	7	8.8	10.4	3.1	24.8	14.6	6.1	12.3	6.1	NA	19.9	7.1	NA	6.5
D_{18}	5.2	1.5	5.1	5.4	6.8	9.6	2.3	3.8	1.3	1.3	12.2	2.6	1.6	1.8	1.3	0.9	1.6	5	1.2	1.8	NA	1.5	1.8	NA	1.3
	23.1	12.2	5.3	10.6	3.3	20.6	15.9	13.5	5.8	5.1	27.4	18	3.5	13.7	6	8.5	13	1.7	10.7	11.1	NA	15.7	15.1	NA	4.3
D_{19}	0.5	1.1	6.2	1.6	1.3	NA	2.8	5.5	5.4	2	NA	2.4	5.1	6.6	1.4	NA	NA	2.2	2.1	1.8	NA	NA	NA	NA	4.9
	31	10.9	6.4	12.7	9.6	NA	19.4	12.4	20.7	11.6	NA	16.4	5.8	4.9	11.1	NA	NA	15.1	14.1	11.4	NA	NA	NA	NA	10.8
D_{20}	8.1	1.7	8.3	NA	NA	NA	1.6	7.2	8.2	14.1	NA	2.1	1.4	NA	1.3	NA	1.2	1.6	NA	1.5	NA	NA	NA	NA	NA
	30.2	10.3	39.4	NA	NA	NA	12.5	8.1	39.4	6.3	NA	20.5	10.4	NA	4.2	NA	10.3	10.9	NA	3.9	NA	NA	NA	NA	NA
D_{21}	4.3	1.5	1.6	NA	NA	8.7	4.4	1.7	1.3	6	3.9	2.2	1.7	1.7	1.4	4.6	1.9	1.5	NA	5.3	NA	NA	2.8	NA	5.6
	25	13.9	12.1	NA	NA	18.9	17.7	14.2	11.9	19.9	14	15	12.6	14.2	12.4	13	14.2	10.7	NA	5.4	NA	NA	18.2	NA	9.7
D_{22}	8.3	2.9	8.7	NA	NA	14.6	NA	1.4	NA	8.9	NA	7.7	1.7	9.4	5	NA	NA	8.6	NA	1.4	NA	NA	NA	NA	7.6
	31.6	17	39.5	NA	NA	32.3	NA	11	NA	39.6	NA	14.1	14	39.8	12.5	NA	NA	39.5	NA	6.6	NA	NA	NA	NA	12.1
D_{23}	3.3	1.9	9.4	1.3	8.3	1.1	1.5	8.9	1.3	5.1	0.9	1.7	2	2	1.5	15.8	1	2.1	1.3	1.5	NA	1.3	NA	NA	4.2
	22.4	16.8	4.3	12.3	17	12.7	15	13.5	9	11.1	10.3	14	2.3	14.5	3.5	31.2	10.8	13.4	9.2	9.4	NA	11.8	NA	NA	14.5
D_{24}	NA	2.8	1.4	NA	1.4	NA	1.6	15.3	NA	1.8	NA	3.9	2.9	0.8	1.5	NA	1.5	1.4	NA	1.5	NA	NA	NA	NA	9.1
	NA	22.6	9.9	NA	9.5	NA	14	8.3	NA	12.9	NA	12.6	15.2	36.7	8.6	NA	16	11.5	NA	10.5	NA	NA	NA	NA	39.7
D_{25}	2.3	1.5	1.3	1.6	1.9	6.4	1.5	1.5	1.4	1.7	NA	4.9	4.8	1.4	1.4	10	3.6	9.4	NA	2.3	NA	NA	1.7	2.4	6.8
	20.5	14	10.8	10.6	12.4	27.8	16.1	11.9	11.7	12.6	NA	28.2	19.1	10.7	5.2	30.9	22.5	11.9	NA	1.6	NA	NA	15.4	20.1	16.2
D_{26}	2.4	1.4	11.7	5.5	3.1	2.4	1.9	1.8	2	8.4	6.3	3.3	6.7	NA	1.4	8.9	10.1	1.8	2.2	1.3	NA	2	2.8	NA	7.3
	20.8	14.1	10.6	22.2	19.4	21.4	15.6	17.4	17.1	15.6	27.5	24.6	22.8	NA	10.3	30	16.9	15.3	19.8	10.2	NA	17.5	22.5	NA	20.8
D_{27}	NA	10.8	1.9	NA	1.6	43.6	2.2	1.7	2.3	5.1	NA	4	3.5	2.5	2.2	NA	3	2.9	5.4	2.2	NA	NA	1.8	NA	3.3
	NA	36.2	19.5	NA	17.6	38.9	23.1	17.8	22.4	30.2	NA	30.3	25.6	23.5	20.2	NA	25.2	24.6	23.8	19.7	NA	NA	25.1	NA	25
D_{28}	NA	2.8	2.1	2	2.6	NA	2.5	1.6	2.7	2	NA	NA	1.9	2.1	7.2	NA	4.9	4.9	1.9	2.5	NA	NA	NA	NA	1.6
	NA	29.9	24	25.3	24.4	NA	30.1	28.4	25.6	23.4	NA	NA	22.3	23.5	31.4	NA	35.5	30.6	26.1	23.3	NA	NA	NA	NA	3.1
D_{29}	NA	1.5	1.8	NA	9.7	NA	NA	1.8	1.6	2.8	NA	NA	10	10	1.6	NA	1	2.6	NA	1.9	NA	NA	NA	NA	3
	NA	24	29.3	NA	36.3	NA	NA	23.2	25.9	29.1	NA	NA	37.1	37.2	20.3	NA	18.2	30.9	NA	24.5	NA	NA	NA	NA	30.4
D_{30}	NA	1.7	1.6	9.8	3.5	NA	7.4	9.4	9.3	8.4	NA	3.1	4.2	2.4	4.2	10	1.5	3.7	2.2	2	NA	NA	1.9	8.6	8.8
	NA	29.3	23.5	36.4	32.8	NA	39.6	36.8	35.6	35	NA	36.3	30.4	31.5	30.4	40	25.1	30.5	30.3	18.9	NA	NA	30.4	34.1	33.2

Notes. The estimates of the shape parameters are placed above the scale parameters. When there are insufficient observations for a given state, it is not possible to estimate the corresponding parameters; thus we set the values to NA (i.e., not applicable).

case; that is, the unit valuation of a five-unit bundle may be higher than the unit valuation of a two-unit bundle for a given bidder.

[15] We have also considered alternative specifications such as zero-inflated Poisson distribution. Comparatively, the zero-inflated negative binomial distribution provides the best fit.

[16] Drawing upon prior studies (Aguirregabiria and Mira 2007, Bajari et al. 2007), we focus on pure-strategy equilibria. Nevertheless, according to the "purification" interpretation of Harsanyi (1973), they are observationally equivalent to mixed-strategy equilibria (i.e., the probability distribution of players' actions is the same under the two equilibria). For a detailed discussion about purification of MPE in dynamic stochastic games, see Doraszelski and Escobar (2010).

[17] The first-order condition has been used in structural estimations of both static and dynamic games (Guerre et al. 2000, Jofre-Bonet and Pesendorfer 2003).

[18] We also experiment with $\beta = 0.85$ and $\beta = 0.95$ but observe minimal differences in the results.

[19] The general idea of AES is that, although each bidder's value may vary in different periods and different bidders' values may also be quite different in any given period, the market environment faced by bidders and, thus, their optimal decisions with any given value may be very similar.

[20] Although market conditions can be broadly classified as scarcity, balanced, and oversupply (Ketter et al. 2012), we do not consider the oversupply condition in the simulation experiments as it rarely happen in the DFA market.

[21] Ideally, we should use the number of unique log-in IDs as the proxy for the total number of potential bidders. However, such data are not available under the current bidding system. Nevertheless, given the B2B nature of this market, we believe that it is highly unlikely that a bidder who was interested in the products (i.e., had received customers' orders) never won any auction during the six-week period, and thereby the number of unique winner IDs serves as a good proxy for the total number of potential bidders in the market.

[22] Because we do not have sufficient information about bidders' daily demand, it is not possible to recover the Poisson parameter in Assumption 1. However, once such data becomes available, we can easily adapt the estimation process and determine the participation probability.

[23] The maximum of the minimum purchase quantity observed in our data set is nine.

[24] It is worth mentioning that, although we have chosen to focus on the optimal decisions regarding the dynamic setting of minimum purchase quantity, the methodology developed in this section can be adapted to optimize other auction design parameters, such as starting price, reserve price, or lot size.

References

Adomavicius G, Gupta A, Zhdanov D (2009) Designing intelligent software agents for auctions with limited information feedback. *Inform. Systems Res.* 20(4):507–526.

Aguirregabiria V, Mira P (2007) Sequential estimation of dynamic discrete games. *Econometrica* 75(1):1–53.

Ashenfelter O, Graddy K (2003) Auctions and the price of art. *J. Econom. Literature* 41(3):763–787.

Athey S, Haile P (2002) Identification of standard auction models. *Econometrica* 70(6):2107–2140.

Athey S, Levin J, Seira E (2011) Comparing open and sealed bid auctions: Evidence from timber auctions. *Quart. J. Econom.* 126(1):207–257.

Ausubel L, Cramton P, Pycia M, Rostek M, Weretka M (2014) Demand reduction and inefficiency in multi-unit auctions. *Rev. Econom. Stud*. 81(4):1366–1400.

Backus M, Lewis G (2016) Dynamic demand estimation in auction markets. NBER Working Paper No. 22375, National Bureau of Economic Research, Cambridge, MA.

Bajari P, Hortacsu A (2003) The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econom*. 34(2):329–355.

Bajari P, Benkard L, Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75(5):1331–1370.

Bapna R, Goes P, Gupta A, Karuga G (2008) Predicting bidders' willingness to pay in online multi-unit ascending auctions: Analytical and empirical insights. *INFORMS J. Comput*. 20(3): 345–355.

Bichler M, Gupta A, Ketter W (2010) Designing smart markets. *Inform. Systems Res*. 21(4):688–699.

Boutilier C, Goldszmidt M, Sabata B (1999) Sequential auctions for the allocation of resources with complementarities. *Proc. 16th Internat. Joint Conf. Artificial Intelligence (IJCAI-99), Stockholm*, 527–534.

Brendstrup B (2007) Non-parametric estimation of sequential English auctions. *J. Econometrics* 141(2):460–481.

Caldentey R, Vulcano G (2007) Online auction and list price revenue management. *Management Sci*. 53(5):795–813.

Chen X, Ghate A, Tripathi A (2011) Dynamic lot-sizing in sequential online retail auctions. *Eur. J. Oper. Res*. 215(1):257–267.

Chintagunta P, Erdem T, Rossi P, Wedel M (2006) Structural modeling in marketing: Review and assessment. *Marketing Sci*. 25(6): 604–616.

Donald SG, Paarsch HJ (1996) Identification, estimation, and testing in parametric empirical models of auctions within the independent private values paradigm. *Econometric Theory* 12(3): 517–567.

Donald SG, Paarsch HJ, Robert J (2006) An empirical model of the multi-unit, sequential, clock auction. *J. Appl. Econometrics* 21(8): 1221–1247.

Doraszelski U, Escobar J (2010) A theory of regular Markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoret. Econom*. 5(3):369–402.

Dubé JP, Hitsch GJ, Chintagunta PK (2010) Tipping and concentration in markets with indirect network effects. *Marketing Sci*. 29(2): 216–249.

Ericson R, Pakes A (1995) Markov-perfect industry dynamics: A framework for empirical work. *Rev. Econom. Stud*. 62(1): 53–82.

Flambard V, Perrigne I (2006) Asymmetry in procurement auctions: Evidence from snow removal contracts. *Econom. J*. 116(514): 1014–1036.

Goes P, Karuga G, Tripathi A (2010) Understanding willingness to pay formation of repeat bidders in online sequential auction. *Inform. Systems Res*. 21(4):907–924.

Guerre E, Perrigne I, Vuong Q (2000) Optimal nonparametric estimation of first-price auctions. *Econometrica* 68(3):525–574.

Harsanyi J (1973) Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *Internat. J. Game Theory* 2(1):1–23.

Hickman BR, Hubbard TP, Saglam Y (2012) Structural econometric methods in auctions: A guide to the literature. *J. Econom. Methods* 1(1):67–106.

Huang Y, Singh P, Ghose A (2015) A structural model of employee behavioral dynamics in enterprise social media. *Management Sci*. 61(12):2825–2844.

Jofre-Bonet M, Pesendorfer M (2003) Estimation of a dynamic auction game. *Econometrica* 71(5):1443–1489.

Judd KL (1998) *Numerical Methods in Economics* (MIT Press, Cambridge, MA).

Kambil A, van Heck E (1998) Reengineering the Dutch flower auctions: A framework for analyzing exchange organizations. *Inform. Systems Res*. 9(1):1–19.

Kambil A, van Heck E (2002) *Making Markets: How Firms Can Design and Profit from Online Auctions and Exchanges* (Harvard Business School Press, Boston).

Ketter W, Collins J, Gini M, Gupta A, Schrater P (2012) Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Inform. Systems Res*. 23(4): 1263–1283.

Klemperer P (2002) What really matters in auction design. *J. Econom. Perspect*. 16(1):169–189.

Kullback S, Leibler RA (1951) On information and sufficiency. *Ann. Math. Statist*. 22(1):79–86.

Lu Y, Gupta A, Ketter W, van Heck E (2016) Exploring bidder heterogeneity in multichannel sequential B2B auctions. *MIS Quart*. 40(3):645–662.

Lucas RE Jr (1976) Econometric policy evaluation: A critique. Brunner K, Meltzer AH, eds. *Carnegie-Rochester Conference Series on Public Policy*, vol. 1 (North Holland, Amsterdam), 19–46.

McAfee R, Vincent D (1993) The declining price anomaly. *J. Econom. Theory* 60(1):191–212.

Milgrom P, Weber RJ (1982) A theory of auctions and competitive bidding. *Econometrica* 50(5):1089–1122.

Muth JF (1961) Rational expectations and the theory of price movements. *Econometrica* 29(3):315–335.

Paarsch HJ (1992) Deciding between the common and private value paradigms in empirical models of auctions. *J. Econometrics* 51(1/2):191–215.

Paarsch HJ, Hong H, Haley R (2006) *An Introduction to the Structural Econometrics of Auction Data* (MIT Press, Cambridge, MA).

Pardoe D, Stone P, Saar-Tsechansky M, Keskin T, Tomak K (2010) Adaptive auction mechanism design and the incorporation of prior knowledge. *INFORMS J. Comput*. 22(3):353–370.

Pinker E, Seidmann A, Vakrat Y (2010) Using bid data for the management of sequential, multi-unit, online auctions with uniformly distributed bidder valuations. *Eur. J. Oper. Res*. 202(2): 574–583.

Pratt JW, Gibbons JD (1981) *Concepts of Nonparametric Theory* (Springer-Verlag, New York).

Rinne H (2008) *The Weibull Distribution: A Handbook* (Chapman and Hall/CRC, New York).

Roth A, Ockenfels A (2002) Last minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econom. Rev*. 92(4):1093–1103.

Rothkopf MH, Harstad RM (1994) Modeling competitive bidding: A critical essay. *Management Sci*. 40(3):364–384.

Rust J (1994) Structural estimation of Markov decision processes. Engle RF, McFadden DL, eds. *Handbook of Econometrics*, vol. 4 (Elsevier, Amsterdam), 3081–3143.

Shen Z, Su X (2007) Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production Oper. Management* 16(6):713–728.

Swinkels JM (2001) Efficiency of large private value auctions. *Econometrica* 69(1):37–68.

Van den Berg GJ, Van Ours J, Pradhan M (2001) The declining price anomaly in Dutch Dutch rose auctions. *Amer. Econom. Rev*. 91(4): 1055–1062.

Vickrey W (1961) Counterspeculation, auctions and competitive sealed tenders. *J. Finance* 16(1):8–37.

Wang S, Jank W, Shmueli G (2008) Explaining and forecasting online auction prices and their dynamics using functional data analysis. *J. Bus. Econom. Statist*. 26(2):144–160.

Yao S, Mela CF (2008) Online auction demand. *Marketing Sci*. 27(5): 861–885.

Yao S, Mela CF (2011) A dynamic model of sponsored search advertising. *Marketing Sci*. 30(3):447–468.

Zeithammer R (2006) Forward-looking bidding in online auctions. *J. Marketing Res*. 43(3):462–76.