

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Interfaces with Other Disciplines

Using bid data for the management of sequential, multi-unit, online auctions with uniformly distributed bidder valuations

Edieal Pinker^{a,*}, Abraham Seidmann^a, Yaniv Vakrat^b^aSimon Graduate School of Business Administration, University of Rochester, Rochester, NY 14627, United States^bAdobe Systems Inc., San Jose, CA, United States

ARTICLE INFO

Article history:

Received 17 July 2008

Accepted 19 May 2009

Available online 27 May 2009

Keywords:

Information technology

Information systems

Internet

Multi-unit auctions

Sequential auctions

Dynamic programming

ABSTRACT

Internet auctions for consumers' goods are an increasingly popular selling venue. We have observed that many sellers, instead of offering their entire inventory in a single auction, split it into sequential auctions of smaller lots, thereby reducing the negative market impact of larger lots. Information technology also makes it possible to collect and analyze detailed bid data from online auctions. In this paper, we develop and test a new model of sequential online auctions to explore the potential benefits of using real bid data from earlier auctions to improve the management of future auctions. Assuming a typical truth-revealing auction model, we quantify the effect of the lot size on the closing price and derive a closed-form solution for the problem of allocating inventory across multiple auctions when bidder valuation distributions are known. We also develop a decision methodology for allocating inventory across multiple auctions that dynamically incorporates the results of previous auctions as feedback into the management of subsequent auctions, and updating the lot size and number of auctions. We demonstrate how information signals from previous auctions can be used to update the auctioneer's beliefs about the customers' valuation distribution, and then to significantly increase the seller's profit potential. We use several examples to reveal the benefits of using detailed transaction data for the management of sequential, multi-unit, online auctions and we demonstrate how these benefits are influenced by the inventory holding costs, the number of bidders, and the dispersion of consumers' valuations.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Online auctions rank among the more successful uses of the Internet for commerce. Increasing numbers of companies see the web as an additional marketing channel complementing their stores and posted-price online catalogs. For example, the Ubid and Sam's Club¹ auction sites have many single and multi-unit auctions of electronics and computers, while CompUSA maintains a companion auction site² offering numerous new and used computers. Dovebid³ specializes in auctions liquidating the assets of businesses (see Klein, 1997; Turban, 1997; Pinker et al., 2003). Online auctions have thus become a viable and active sales channel for many firms in both business-to-business and business-to-consumer markets. Furthermore, the use of computers to conduct the auctions means that all transactions are logged, and detailed information about all bids placed is collected automatically. This infor-

mation has tremendous value to the auctioneer, since it provides insight into bidders' valuations for the auctioned goods and thus can be used in the determination of future auction offerings.

The online auction is a fundamentally different way of selling goods and managing inventory than the posted-price mechanism. In the traditional posted-price setting, and in most of the classical work on inventory theory (Zipkin, 2000), the seller faces demand uncertainty, but the selling price is fixed. On the other hand, when using multi-unit auctions, the seller expects the entire lot to be sold but is uncertain about the clearing price. This means that in the auction environment companies can determine the quantity offered, and the market determines the price in response to this quantity. The greater the lot size offered, the lower the price established for each unit. We call this phenomenon the *market impact* of the lot size.

The computing power available today makes it quite practical to incorporate auction feedback into decision-making. For example, Moon (1999) describes how one online auction site leveraged the reservation price information that it possessed to generate an outstanding response rate to unsolicited emails; it created personalized email offers based on each consumer's willingness to pay. These offers were directed at individuals who participated in various auctions but did not bid high enough to win; their highest bids

* Corresponding author. Tel.: +1 585 275 2131; fax: +1 585 273 1140.

E-mail addresses: ed.pinker@simon.rochester.edu (E. Pinker), avi.seidmann@simon.rochester.edu (A. Seidmann).¹ <http://Ubid.com> and <http://auctions.samsclub.com/>.² <http://www.compusauctions.com/>.³ <http://www.dovebid.com>.

were used as proxies for their true reservation prices. We model a similar strategy of using the bid data from previous auctions to improve the management of future auctions. In general, there are three ways to gather this kind of information. First, if a firm runs its own auctions, it can collect the data directly from bid activity log files. Second, if a third-party hosts the auction, a firm can request this information as a condition of participation or as an additional service. Third, when online auctions are open to the public, one can make use of software agents such as AuctionScout.⁴

At all major Internet auction websites, it is common to observe sequential auctions for the same good.⁵ In fact, in Vakrat (2000), the results of extensive field data collection of online auctions show that about 85% of the goods auctioned at a popular business-to-consumer auction site were offered repeatedly. Nair et al. (2004) also give empirical evidence of both single and multi-unit sequential auctions for the same item on the Sam's Club auction site. Sequential auctions and the market impact of the lot size introduce a set of interesting tradeoffs for the seller. Smaller lot sizes will increase the revenue per unit, but, given an initial inventory, smaller lot sizes will increase the number of auctions necessary to dispose of it. There is a fixed cost to operating an online auction that is largely independent of the lot size offered, and there is a cost to holding items in inventory until they are auctioned off. Increasing the number of auctions and thereby spreading the sale of the inventory over multiple time periods will increase the auctioneer's costs.

In this paper, we consider a firm seeking to sell off an initial inventory of a product through a series of multi-unit auctions. Before the first auction, the firm has only incomplete information about the statistical distribution of its customers' valuations for the product. The profit-maximizing firm must decide how many units to offer in each auction and how many auctions to run. To address these decisions, we develop a dynamic optimization model, using a Bayesian learning framework, for analyzing auction data as the Internet auction site collects it. We do not intend our model to serve as a decision tool that provides precise calculation of the results of auctions, but rather, we seek to demonstrate the following managerial insights:

- How auctions can be integrated into a firm's operations.
- How sequential auctions can be used effectively to reduce the negative market impact of large lot sizes, and to identify key quantitative tradeoffs associated with that policy.
- How an auctioneer can use information from losing bidders in earlier auctions to learn about the market and to improve the management of future auctions by embedding the inventory management and auction planning decisions within a practical Bayesian analysis framework.

Our paper proceeds as follows. In Section 2, we briefly review some of the relevant literature and outline the theoretical background that is necessary to understand the main results of the paper. In Section 3, we formulate a general model of multi-unit sequential online auctions with seller learning in each auction. We then analyze two special cases. In the first, we assume that the bidder valuation distribution is uniform with known parameters, using a deterministic dynamic program. We show that the optimal lot size (the number of units offered) in each auction is strongly dependent on the spread (or second moment) of the bid-

ders' valuations, and that setting the lot size optimally can have a significant impact on profits. We then consider the case of a uniform valuation distribution in which the seller does not know the spread. In this case, each auction provides the auctioneer with *feedback data* that can be used to improve the management of subsequent auctions. We investigate the value of using this feedback in conjunction with dynamic optimization of the auction lot sizes. We conduct an extensive set of numerical experiments and determine under what circumstances this approach yields significant benefits to the auctioneer. Section 4 contains a critical review of our modeling assumptions. We offer our conclusions in Section 5.

2. Literature review

Much of the inventory management literature assumes an exogenous demand, independent of the lot size and of the price. There have been several studies in the area of Bayesian dynamic inventory models with some unknown demand parameter, for example Scarf (1959), Azoury (1985), and Lovejoy (1990). This inventory problem has been modeled as a dynamic program with a multi-dimensional state space, where Bayes' rule provides a well-defined procedure to incorporate new information into the decision model as it becomes available. Most of these papers are concerned with deriving multi-period minimum cost equations and characterizing the optimal reordering policy. Although we use a similar methodology (i.e., dynamic programming and Bayesian learning), our business context and consequently our results are very different. In our setting, with prices determined by auction, the uncertainty we focus on is in price and the valuation distributions of bidders and not demand.

Although online auctions are receiving considerable attention from management scientists and economists (see Pinker et al. (2003), Milgrom (2004), and Ockenfels et al. (2006) for surveys), there has been little analytical modeling of the optimal ways to operationalize online auctions. Here we review some exceptions. Bapna et al. (2000, 2001) develop a discrete model to compare the current online auction practice with an ideal Vickrey mechanism. They classify online bidders into opportunists, participants, and evaluators, and they study the impact of the bidding increment. Segev et al. (2001) use an orbit queue Markov chain to model online auctions and predict the final auction prices. In those cases in which their methodology can be proven to provide accurate predictions, one could also use it for auction management; so long as it can be shown that the changes to the auction conduct, proposed by the authors, do not alter the expected predictions. Vakrat and Seidmann (1999a,b) find empirical evidence of significant discounts at online auctions over online catalog prices for the same goods. These results suggest that participants in online auctions represent a separate customer segment that is willing to exchange the costs in monitoring and time of participating in an auction for lower prices than are found in posted-price settings.

Gallego and Van Ryzin (1994) study the related problem, to ours, of dynamic pricing of an inventory in continuous time setting when demand is a deterministic function of price. Vulcano et al. (2002) model a firm with a fixed initial inventory that receives bids for single-units over time from consumers. They develop a finite-horizon dynamic programming model in which the firm makes decisions to accept or reject bids in each period. For the i th highest bid to be accepted, it must exceed the i th highest threshold set dynamically by the firm. They propose a novel auction scheme in which the bidders reveal their bids before the firm even commits to the lot size. They show numerically that this mechanism may outperform multi-period auctions in which firms commit to a quantity in each auction before observing bids. Their result makes sense, because the firm is clearly better off making auction

⁴ Auction Scout is a data collection and monitoring software agent we have developed for unwrapping the HTML information displayed through the TCP/IP socket of third-party auction houses (see Vakrat and Seidmann (1999a,b)).

⁵ There even exist firms like Marketworks (<http://www.marketworks.com>) that specialize in setting up sequential auctions on public auction sites for firms trying to dispose of inventory.

quantity decisions after observing bids than before. This mechanism seems to fit best within the single-leg revenue management context that motivates the paper, yet the unique auction scheme they propose has yet to be seen in practice. In our paper we model auctions with a publicly committed lot size, as are found on most auction websites today. By making specific distributional assumptions about bidders' valuations, we are able to derive detailed results about optimal lot-sizing policies that are informative to managers. In addition, we are able to develop a Bayesian learning mechanism for analyzing auction data to improve the design parameters of future auctions.

3. Model formulation and analysis

Suppose that a firm has an inventory, x_0 , of units it would like to sell via auction. It can offer the entire lot in one auction, or it can split the lot among multiple separate sequential auctions spread over a longer period. Running several auctions potentially exposes the inventory to more demand, thus increasing prices and revenue. On the other hand, each additional auction incurs both fixed costs and per unit costs for the units held in inventory and carried over from period to period. A firm must balance these opposing effects when deciding how many units to offer in an auction. In making the lot-sizing decision, the firm also implicitly determines the number of auctions it conducts.

To address this question, we develop a multi-period model of a firm auctioning off an initial inventory of a product. For the reader's convenience notation is listed in Table 1. Each period is of duration t , and at the end of each period there is an auction of duration t_a that is included within t . For example, every 10 days the firm conducts a one-day auction that starts at the end of the ninth day. Following Etzion et al. (2006) we assume the following: bidders have independent private values for the good being sold and are risk neutral; bidders' valuations follow a uniform distribution; bidders in one auction are independent of bidders in other auctions; they have single-unit demand; and the number of bidders (or demand) in each period is a constant, n .

We also assume that auctions are conducted using a sealed-bid $(k + 1)$ - price format. In the $(k + 1)$ - price sealed-bid auction the dominant strategy is for the bidder to bid his true valuation for the good. Further, Harris and Raviv (1981) have shown that with independent private values and risk neutral bidders, expected revenue is equivalent to that of the Yankee auction which is another name for multi-unit, discriminatory, English auctions. In a Yankee auction each winning bidder pays their own bid price and bidding is open. This is a common format for multi-unit online auctions (Easley and Tenorio, 2004). The implication of revenue equivalence is that, from the perspective of the seller's revenue, modeling the

online auction as a $k + 1$ price sealed-bid auction will yield accurate results. However, we note that the actual bids will be different in these two formats, with implications for the Bayesian learning procedure we develop later. We discuss this further in Section 4. Our modeling assumptions best suit a situation in which consumers are shopping for an item, come to the auction, and if unsuccessful go elsewhere to make the purchase. They do not know whether an auction will be repeated, or when. They do not linger from one auction to the next because they need to make their purchase (imagine an office manager trying to equip a new office with a computer or fax machine). The assumptions of single-unit demand, independent private values, and a constant number of bidders are standard in the auction literature Fatima et al. (2008). In Section 4 we discuss the sensitivity of our results to these key modeling assumptions.

Assuming that the seller plans T auctions, let $i = 1, 2, \dots, T$ be the index of the period, and let k_i be the auction lot size in period i . In each period, the auctioneer's revenue is determined by the number of units he decides to offer and the auction's closing price. (A formulation of the model with discounting appears in Appendix 2.) In any period i in which the seller auctions k_i items to n bidders, we define the expected auction closing price, $p(k_i, n)$, and so the revenue in each period is given by $k_i p(k_i, n)$. Let $J_{i,T}(x_i, k_i)$ be the total profit of the best overall policy from the beginning of period i to the end of the planning horizon, if the inventory level entering period i is x_i and k_i units are chosen to be auctioned off. We denote by $J_{i,T}^*(x_i)$ the corresponding maximum value of $J_{i,T}(x_i, k_i)$ subject to $k_i \in \{0, \dots, x_i\}$. The firm must select T so that $J_{1,T}^*(x_0)$ is maximized.

Let h be the inventory holding costs incurred by the auctioneer per unit per time. In each period, the auctioneer incurs a holding cost of $x_i h t$ on the current stock. Finally, assume a fixed cost of C to set up and run each auction, which includes, but is not limited to, promotion, website hosting fees, and auction monitoring. Combining the revenue and the cost for each period, we can formulate the problem of finding the optimal lot size policy, for a given initial inventory x_0 , recursively as the following deterministic dynamic program:

Problem P1 :

$$J_{i,T}^*(x_i) = \underset{k_i \in \{0, \dots, x_i\}, T, x_1}{\text{Max}} J_{i+1,T}^*(x_i - k_i) + k_i p(k_i, n) - h t x_i - C$$

for $i = 1, \dots, T$
 where :
 $x_1 \leq x_0$
 $x_i = x_{i-1} - \min(k_{i-1}, n), i = 2, \dots, T$
 $k_i \leq x_i$ and x_i, k_i integer

The lot size k_i may not exceed the current inventory level x_i ; the seller cannot auction items he does not possess. If $x_1 < x_0$, the firm scraps $x_0 - x_1$ units before starting the auctions. In practice the seller will not know the bidder valuation distribution and thus will have considerable uncertainty about the auction closing price $p(k_i, n)$. One of the advantages of an online auction is that it is easy to collect information about bidders' behavior for the purpose of analysis. This information gives the seller feedback on the accuracy of his understanding of the market. We now build upon the formulation of Problem (P1) to formulate a general model of how a seller can incorporate bid observations into the management of sequential auctions.

Define $H_i \in \mathfrak{R}^{h_i}$ as the information set from the bid history through period $i - 1$ used by the seller in updating his beliefs about the bidder valuation distributions, where h_i could be any integer from 1 to $i - n$. Given a new set of final bids B_i from period i , we can define the process of incorporating it into H_i as a function $I_i(H_i, B_i) : \mathfrak{R}^{h_i} \times \mathfrak{R}^n \rightarrow \mathfrak{R}^{h_{i+1}}$, with $H_{i+1} = I_i(H_i, B_i)$.

Table 1
Notation.

k_i	Number of items to auction in auction i
T	Number of sequential auctions to run
x_i	Number of items available for auction i . For $i = 1$, a decision; for $i > 1$ (derived from k_i)
n	Number of bidders in each auction
h	Unit holding cost per unit time
t	Interval between the start of each auction
$p(k_i, n)$	Expected auction closing price given n bidders and k_i units auctioned
C	Fixed cost of running each auction
x_0	Initial inventory before any auction occurs
H_i	Information set from bid history used by seller
θ	Parameters of the bidder valuation distribution
$g_i(\theta)$	Joint pdf of the parameters of the bidder valuation distribution after $i - 1$ auctions

We define $g_i(\theta)$ as the joint pdf of the parameters θ of the bidder valuation distribution after $i - 1$ auctions. The functions $g_i(\theta)$ are determined by Bayesian updating after observing a new set of bids as follows:

$$g_{i+1}(\theta) = g_i(\theta|H_{i+1}). \tag{2}$$

We formulate the seller's optimization problem as follows:

$$\text{Problem P2 : } J_i^*(x_i, H_i) = \text{Max}_{k_i \in [0, x_i], x_1} E_\theta [E_{H_{i+1}} [\int_{i+1}^* (x_i - k_i) | \theta, H_i] + k_i p(k_i, n, \theta)] - X_i$$

$$J_{x_0}^*(x_{x_0}, H_{x_0}) = \text{Max}_{k_{x_0} \in [0, x_{x_0}]} E_\theta [k_{x_0} p(k_{x_0}, n, \theta)] - X_{x_0}$$

for $i = 1, \dots, x_0$ (3)

where :

$$x_1 \leq x_0$$

$$x_i = x_{i-1} - \min(k_{i-1}, n), \quad i = 2, \dots, x_0$$

$$k_i \leq x_i \text{ and } x_i, k_i \text{ integer}$$

$$\text{Here : } X_{x_0} = \begin{cases} hx_i + C & \text{when } x_i > 0 \text{ and } k_i > 0 \\ 0 & \text{when } x_i = 0 \\ hx_i & \text{when } k_i = 0 \text{ and } x_i > 0 \end{cases}$$

The problem is formulated as an x_0 period dynamic program because the number of periods in which the seller conducts auctions is unknown at the outset but will not exceed the number of units in the initial inventory, as the seller must offer at least one item to conduct an auction. Note that since it is possible to scrap excess inventory at the start (by choosing $x_1 < x_0$), it is suboptimal to not run an auction in a period if inventory is available. Problem (P2) is a general formulation of how a seller can use bid information to update his knowledge about bidders and allocate inventory to auctions.

In order to make it possible to derive specific results that will grant us insight into the potential benefits of using bid information in this way, we make the following assumptions about the seller's learning from the bidder valuation distribution. First, we assume that the valuation distribution is uniform on $[\mu - s, \mu + s]$. Second, we assume that the mean, μ , is known but the range (or dispersion), s , of the distribution is unknown. Third, we assume that the prior distribution on the dispersion is uniform. Finally, we assume that the information set used, H_i , is the largest deviation from the mean observed thus far.

We focus on learning the dispersion of consumers' valuations for the following reasons:

1. We show later (in Proposition 1) the importance of this dispersion in determining the optimal lot-sizing policy.
2. The dispersion significantly influences the seller's price risk. High dispersion, small numbers of bidders, and large lot sizes all increase the probability that the auction price will be determined by bidders drawn from the left tail of the valuation distribution (Vakrat and Seidmann, 1999a,b).
3. Previous research on auctions (Wang, 1993) has shown the impact of the variance in bidders' valuations on the efficacy of auctions.
4. We expect that the auctioneer would have a good sense of the mean of consumers' valuations, but the dispersion would be more difficult to estimate. Most items offered via Internet auctions can be related to publicly available prices (e.g., alternative posted-price channels), so the seller should be able to form a reasonable estimate of the mean. Selling items via a posted-price mechanism, however, does not provide the seller with explicit reservation price information and therefore does not guide him with respect to the underlying dispersion of consumers' valuations.

Assuming that the prior on the dispersion parameter is uniform creates a good test of the learning procedure. Using the largest deviation observed so far as the information set for the learning procedure keeps the state space small and simplifies the computations required.

Given the above assumptions we can precisely define the learning process as follows. Each time the online auctioneer runs an auction, he observes a random draw of the population's reservation prices (v_j 's). Let Y_i be the observed maximum distance from the mean:

$$Y_i = \max\{|v_j - \mu| : j = 1, \dots, n\} \text{ for } i = 1, \dots, T^*.$$

In the first period, the auctioneer's prior on s is $g_1(s) = 1/(b - a), s \in [a, b]$. During one auction, the auctioneer observes n different reservation prices. Conditioned on s , the cumulative distribution of Y_1 for n random draws is

$$F(y_1|n, s) = \begin{cases} 0 & y_1 > s \\ (y_1/s)^n & y_1 \leq s. \end{cases} \tag{4}$$

The observed signal (the maximal deviation from the mean) from each auction is denoted by y_i . We can now use Bayes' formula to calculate the auctioneer's posterior on s for the second period. To simplify the expressions, we define $a_1 = \max(a, y_1)$:

$$g_2(s) = g_1(s|y_1, n) = \frac{g_1(s)f(y_1|n, s)}{\int f(y_1|n, s)g_1(s)ds} = \frac{\frac{1}{b-a} \left(\frac{y_1}{s}\right)^{n-1} \frac{n}{s}}{\int_{a_1}^b \left(\frac{y_1}{s}\right)^{n-1} \frac{n}{s} \frac{1}{b-a} ds}$$

$$= s^{-n} \left(\frac{1}{a_1^{n-1}} - \frac{1}{b^{n-1}} \right)^{-1} (n - 1) \text{ for } s \in [a_1, b].$$

The limits of the integral in the likelihood function are a_1 to b . As the auctioneer observed a maximal deviation of y_1 , he is positive that s is at least as large as the maximum of that observation or the prior lower bound a . The next period's auction thus adds new information about the spread of the consumers' valuations only if the auctioneer observes a reservation price that is farther away from the mean than any he has seen so far (in all preceding periods). In general, for each period we define a lower bound on $s, l_i = \max\{l_{i-1}, y_i\}, i = 1, \dots, T^*$, where $l_1 = y_1$. Using Bayesian updating, we can obtain the posterior distribution of s for each period in the auction sequence. After i periods, the auctioneer's posterior (which is period $i + 1$'s prior) for the spread of consumers' valuations (s) is

$$g_{i+1}(s) = g_i(s|l_i) = s^{-in} \left(\frac{1}{\max(a, l_i)^{in-1}} - \frac{1}{b^{in-1}} \right)^{-1} (in - 1)$$

for $s \in [\max(a, l_i), b]$. (5)

Knowing the posterior distribution, the auctioneer can calculate his updated estimate for the dispersion of consumers' valuations. We can see from Eq. (5) that the distribution of s is completely specified by l_i and i because the number of bids in each period is known to be n . In this case, the number of auctions will not be determined beforehand. Period i determines how many bids have been observed so far, and l_i gives a lower bound on s . The solution to the entire planning problem is given by $J_1^*(x_1, l_1)$. The maximum number of auctions is x_1 , since at least one unit must be auctioned in each auction. Thus, even though the number of auctions is stochastic, we can calculate $J_i^*(x_1, l_1)$ using the finite-horizon dynamic program Problem (P2), in which $J_i^*(0, l_i) = 0, J_1^*(0, l_1) = 0$ simply means that no more auctions are conducted once the inventory has been exhausted.

As a baseline, we next consider the case in which the bidder valuation distribution is uniform with known parameters, Problem (P1). It is typically more difficult to derive structural results for finite-horizon dynamic programs than for those with an infinite-horizon. The analysis of P1 will show that the problem we are

considering has the added complexity that the horizon is endogenous to the problem. Adding the Bayesian learning in Problem (P2) forces us to resort to numerical analysis. Hence, when we consider the case in which the bidder valuation distribution is uniform with known mean but unknown dispersion, we present a representative set of numerical results to illustrate the potential of using Bayesian learning in this setting.

3.1. The bidder valuation distribution is uniform with known parameters

We start with determining the price in a single multi-unit auction. Suppose there are n bidders with values independently and uniformly distributed on $[\mu - s, \mu + s]$, where μ is the mean of the distribution and s is its dispersion. We assume that each bidder has demand for only one unit, and the auction follows a sealed-bid $k + 1$ price format. In such an auction, for $k < n$ items, the k highest bidders win an item, and the auction's closing price is given by the $k + 1$ highest bid. Under these assumptions, it can be shown (see Appendix 1) that, for lot size k and n bidders, the expected closing price, $p(k, n)$, is given by the following:

$$p(k, n) = \begin{cases} \mu + s - 2sm(k + 1) & \text{for } k < n \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where $m = 1/(n + 1)$. We can see in (6) that the expected auction price p_a increases with μ and n and decreases with k . If $k > n$, the products are sold at the auction's reserve price, which we assume to be $\mu - s$. Note that m is a measure of the competition in the auction; more bidders and correspondingly stronger competitive forces result in a smaller m ($0 < m < 1$). When the auctioneer offers an additional unit for sale, there is a negative market impact (i.e., a reduction in the closing price) of $2sm$. This impact is inversely related to the number of bidders. We henceforth refer to m as the marginal market impact. Given the bounded closing price function in Eq. (6) above, the integrality constraint on the lot size, and the fixed demand of n , it is not possible to derive an analytical expression for the optimal lot size policy. Yet by making some simplifying assumptions, we can derive analytical characterizations of how the optimal lot size in each auction depends on the problem's parameters. We later show that these assumptions do not limit the generality of our results.

Assumption 1. The lot size in each auction k_i can be non-integer.

Assumption 2. We assume parameters are such that it is never optimal to sell $k_i \geq n$ in any auction.

Assumption 1 introduces a rounding-off error into our solution but does not change the qualitative nature of the results. Assumption 2 is not very restrictive and serves only to eliminate cases in which the firm would be forced to scrap units. When the condition in Assumption 2 is violated it is impossible to derive closed-form analytical results, but we find in numerical experiments⁶ that the model's behavior is the same as when the condition holds. We therefore believe that the observations we make about the model's behavior based on Proposition 1 (below) and Assumptions 1 and 2 will hold in general.

Remark 1. If the firm conducts T auctions, then $k_i > 0$ for all $i = 1, 2, \dots, T$. Because of holding costs, k_i would never be zero if $k_j > 0$ for some $j > i$.

Remark 2. In the optimal policy the firm sells the entire remaining inventory in the last period; i.e., $k_T = x_T$.

If the firm does not auction the entire remaining inventory at the start of period T , then it must be optimal for the firm to scrap some of its inventory. If this were the case, it would have been preferable to scrap it before the first auction and avoid the holding costs of carrying it through to period T .

The following proposition defines the optimal profit of the auctioneer.

Proposition 1. Given Assumptions 1 and 2, define T^* to be the optimal number of auctions to conduct. Then the optimal lot size in each period $i = 1, \dots, T^*$, is given by

$$k_i = \frac{x_i}{T^* - i + 1} + \frac{ht(T^* - i)(n + 1)}{8s}, \quad (7)$$

We also have that $T^* = \text{ArgMax}\{J_{1,T}(x_i)\}$ for $T \in T_f$, where T_f is the set of T such that $x_i \geq \frac{x_i}{T-i+1} + \frac{ht(T-i)(n+1)}{8s}$ for all $i \leq T$. The auctioneer's resulting profit from period 1 to period T is given by

$$J_{1,T}(x_1) = \frac{T(T^2 - 1)}{6} \frac{(ht)^2(n + 1)}{16s} + \left(\mu + s - \frac{2s}{n + 1} \right) x_1 - \frac{2sx_1^2}{T(n + 1)} - \frac{(T + 1)htx_1}{2} - TC. \quad (8)$$

Proof. See Appendix 1. □

Given an optimal number of periods, T^* (obtained by maximizing Eq. (8)), we can analyze the form of the optimal lot size in each period. The first term of Eq. (7) represents a lot size policy that is constant across all periods. Such a policy is achieved if the auctioneer splits the inventory remaining at the beginning of each period across the number of periods that are left until the end (e.g., in period 1 it is $1/Tx_1$, in period 2 it is $\frac{1}{T-1}x_2$, etc.). The lot size decision affects the entire profit from that period on; the additional term corrects for this. The corrections are monotonically decreasing in the period's index; i.e., early periods will have larger lot sizes. Taking $k_{i-1} - k_i$, we can show the following.

Corollary 1. The optimal lot sizes determined in Proposition 1 decrease monotonically by a constant amount $(ht)/(4sm)$ from period to period.

Proof. See Appendix 1. □

We observe that higher inventory costs (h) result in a stronger correction; more units are moved toward early periods. Recall that $0 < m < 1$ and that a lower m implies more bidders. Also, according to Eq. (6), more bidders imply more competition and consequently a higher closing price. With larger n , the gap between the lot sizes of an early auction and a later auction increases, and more units can be offered in each auction without significantly affecting the price. Corollary 1 shows that as the number of bidders per auction increases, more units can be shifted to early periods in order to save on total inventory holding costs. According to Proposition 1, given an optimal number of auctions, T , the optimal lot size in each auction is independent of μ but dependent on s . Higher dispersion in consumers' valuations (s) results in fewer units shifted toward early periods. Greater dispersion creates an opportunity for the auctioneer to obtain higher closing prices on average, because if he is able to attract enough bidders, there is a higher probability that he will get bidders with high valuations. As a result, high dispersion creates stronger incentives to balance the lot sizes across periods, since deviating from the single-period optimum means greater loss of revenue. At the same time, large lot sizes increase the risk of the price being determined by a low bid. This is a further incentive for the auctioneer to try to balance the lot size when the dispersion is high.

⁶ Our numerical experiments show that once inventory has been reduced enough, the lot size drops below n and then decreases until it reaches 0, following Eq. (7).

3.2. Numerical experiments when the bidder valuation distribution known

To demonstrate the benefit of optimizing the lot sizes, we conduct a series of numerical experiments in which we compare the profit achieved using our optimization model and the profit achieved using a naïve constant lot size policy. We use a base case scenario in which the initial inventory is $x_1 = 30$, $\mu = \$100$, $s = \$50$, $C = \$50$, $h = \$15$, $n = 10$, and $t = 1$. In Table 2 we compare the results with a constant lot size versus those for the optimal lot size to illustrate what a solution looks like (quantities have been rounded for easier display). We see that it is optimal to dump two units before implementing a constant lot size policy and conduct five auctions. With optimal lot-sizing we see that it is optimal to dump one unit and then run progressively six smaller auctions. In Fig. 1a–d, we show the percentage improvement in profit of the optimal policy over the constant lot size policy as we vary the parameters h , n , s , and C , respectively. For example, in Fig. 1a we see that if the holding cost per unit time is $h = \$15$, as in the base case, then the profit from using the optimal policy is 5.4% greater than the profit from the constant lot size policy. Fig. 1b reveals that if we reduce the number of bidders in the base case to 6, the optimal policy provides a 5.5% improvement over the constant lot size policy.

To determine the optimal policy, we solve the problem formulated in Eq. (1) for different values of $x_1 \leq 30$, choosing the initial inventory that maximizes profits. If the optimal starting inventory is less than 30, the firm scraps some of the original inventory before starting the auctions. To determine the constant lot size policy, we use the profit maximizer of all policies of the form $k_i = K$ for $i = [1, \text{int}(x_1/K)]$, and if $x_1 \bmod K > 0$, then $k_{\text{int}(x_1/K) + 1} = x_1 \bmod K$, where $K = [1, n]$, and $\text{int}()$ is the integer part function. This means that in some cases, in the last period, the constant lot size policy

Table 2
Base case solutions, optimal lot-sizing versus constant lot-sizing.

Period	Constant lot size			Optimal			Cumulative profit	
	Inventory	Lot size	Price	Inventory	Lot size	Price	Constant policy	Optimal
1	28	6	\$86	29	7	\$77	\$46	\$54
2	22	6	\$86	22	6	\$86	\$182	\$190
3	16	6	\$86	16	5	\$95	\$408	\$375
4	10	6	\$86	11	4	\$105	\$724	\$580
5	4	4	\$105	7	4	\$105	\$1,034	\$845
6	0	0	0	3	3	\$114	\$1,034	\$1,092

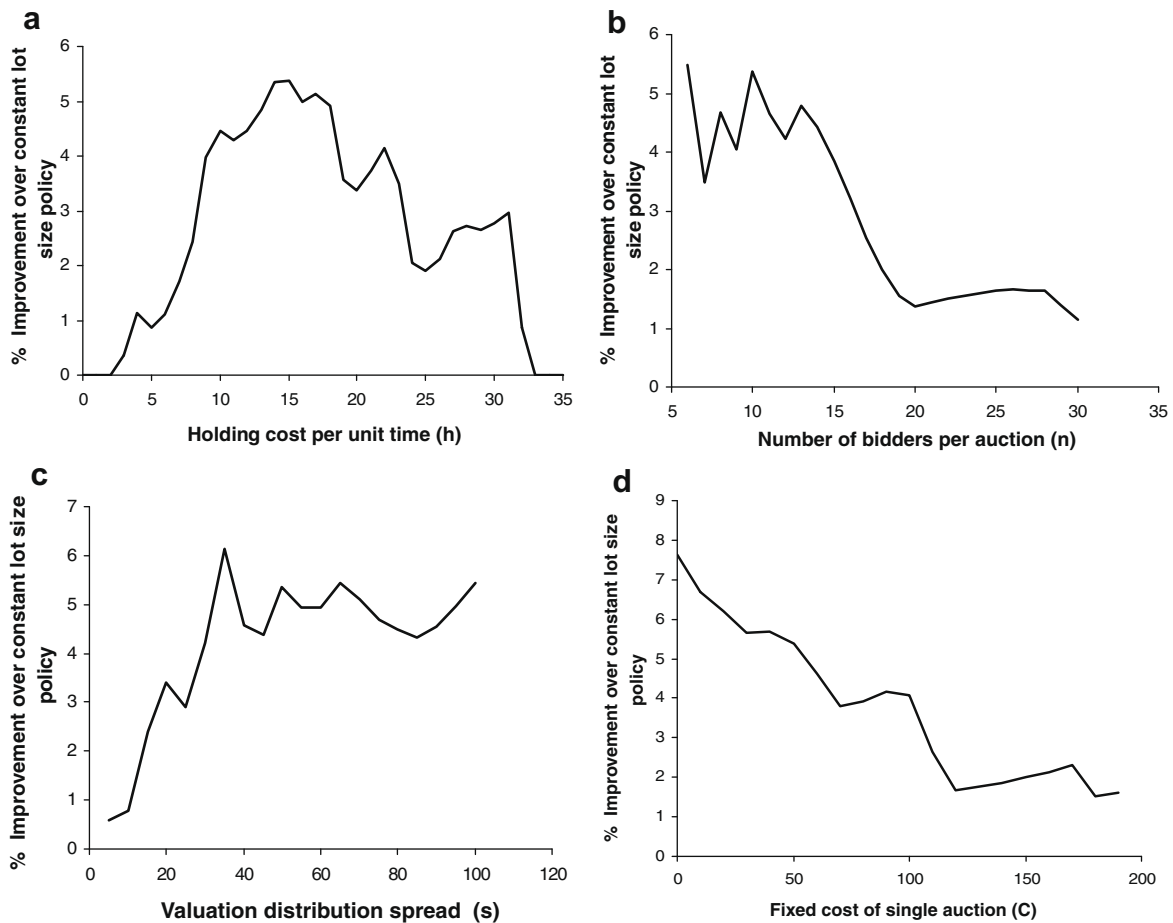


Fig. 1. Percentage improvement in profit over the constant lot size policy from using the optimal lot-sizing policy as a function of (a) holding costs, (b) the number of bidders per auction, (c) the spread of bidders' valuations, and (d) the fixed cost of a single auction.

has different and smaller lot sizes than in all earlier periods. Each constant lot size policy is checked for all possible values of $x_1 \leq 30$ to allow for the initial scrapping of some inventory.

Fig. 1a shows that when holding costs are low there is little difference in the lot sizes used in both approaches. This is because it is optimal to have many small auctions, and there is not much incentive to vary the lot size from period to period. As holding costs increase, both policies require reducing the number of auctions, but the optimal policy is to have larger lots early on, to reduce holding costs. The benefit of optimal lot-sizing thus increases. When holding costs increase, it becomes optimal to scrap some of the inventory rather than hold it for sale in future auctions, but more scrapping occurs using the constant lot size policy. When holding costs are very large, the optimal number of auctions to hold is low for both policies, and the results converge.

Fig. 1b illustrates the effect of the number of bidders. When there are few bidders, using the optimal policy is more important, because it takes the market impact of the lot size into account. As the number of bidders increases, the two policies converge to a smaller number of auctions with large lot sizes. The oscillations in Fig. 1b are caused by integrality and diminish as n grows.

In Fig. 1c we see the effect of changing the spread of the bidders' valuations. When the spread is small, the low end of the valuation distribution is close to the mean valuation. Since we are using a $(k + 1)$ - price auction mechanism, the risk of a low auction price is reduced, making larger lots and fewer auctions preferable. In this scenario there is minimal uncertainty, and therefore little advantage from the optimal lot-sizing policy. When the valuation spread increases, it is better to conduct more auctions with smaller lots, and the benefits of the optimal policy increase. The non-monotonicity of Fig. 1c is due to integrality effects and the countervailing effect of the possibility of higher auction prices when the spread increases. To understand this, it is helpful to look at Eq. (6), the expected price formula. The derivative of Eq. (6) with respect to the spread (s) is $2(n - k)/(n + 1) - 1$. It is positive for $k < (n - 1)/2$ and negative for larger k .

Fig. 1d shows the effect of the fixed cost of running each auction, C . There is a steady decline in relative performance differences as C increases because higher auction costs lead to fewer auctions and more upfront dumping. This leaves little room for significant differences between the two auction policies.

To summarize, when the number of bidders is low relative to the inventory, when the bidder valuation distribution has high variance, or when inventory carrying costs are high, using a lot-sizing optimization such as ours can yield significantly greater benefits than a more naïve constant lot size policy does.

3.3. Numerical investigation when the bidder valuation distribution is uniform with known mean and unknown range

To illustrate the potential of using auction feedback in sequential auctions, we conduct a set of numerical experiments. We assume that in each auction the spread of customers' valuation distribution is an unknown constant, s , with a prior distribution that is uniform on $[a, b]$, where a and b are known. We then compare three scenarios: full-information (this ideal serves as a benchmark), auction feedback (using the Bayesian updating model), and constant lot sizes. We show that using auction feedback can provide increasing benefits when the number of bidders is smaller and holding costs increase. Auction feedback can also lead to lot-sizing policies that are very different from those developed when the valuation spread parameter is assumed to be fixed.

In the full-information scenario, we assume that the actual value of s ($a \leq s \leq b$) is known when the auction is conducted. The full-information scenario therefore serves as an upper bound on the profit the auctioneer can receive. With the constant lot size policy, the auctioneer distributes the initial inventory equally across all the sequential auctions as in Section 3.2, and the number of auctions run is selected so that the expected profit is maximized over all possible values of s . In the auction feedback scenario, we solve Problem (P2) to update the distribution of s each period based on the bids observed so far. As in Section 3.2, we solve Problem (P2) for all possible values of x_0 and choose the profit-maximizing x_0 to allow for inventory scrapping. For these experiments, we use the following set of parameter values: $\mu = \$100$, $C = \$50$, $h = \$15$, $n = 10$, and $t = 1$, with s uniform on $[a = 1, b = 99]$.

In Fig. 2a and b and Fig. 3 we plot the percentage revenue improvement over the constant lot size policy with no feedback, for the full-information and auction feedback scenarios, as a function of the unit holding costs, number of bidders, and the mean bidder valuation, respectively. We observe the advantages of optimal lot-sizing and information, and how the various parameters affect the accuracy of the feedback mechanism. Getting the lot size "right" depends on knowing the valuation distribution or, s , and has the largest relative impact (what is being graphed) when profit margins are narrow, i.e., mean bidder valuations are low. We also see in the figures that the percent improvement in overall profits over the constant lot-sizing policy increases with holding costs, decreases with the number of bidders and decreases with the mean bidder valuation. Conversely, when the holding costs are low, there are many bidders, or bidders have high valuations, the seller makes a large profit regardless of the approach he uses. In these cases then the relative benefit of using the learning algorithm and there-

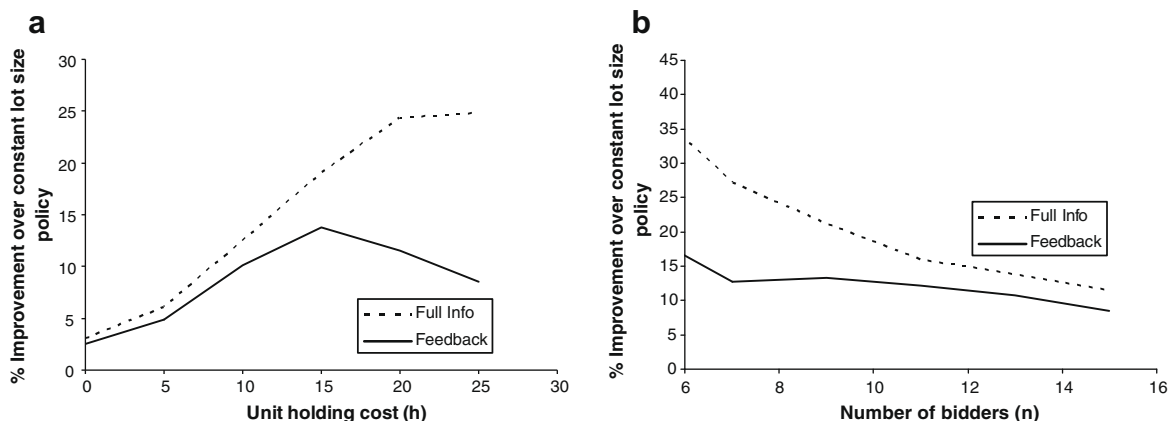


Fig. 2. Percentage improvement in profit over a constant lot size policy when the auction feedback mechanism is used and when there is full-information about the valuation dispersion parameter s , as a function of: (a) holding costs (h), and (b) the number of bidders per auction (n).

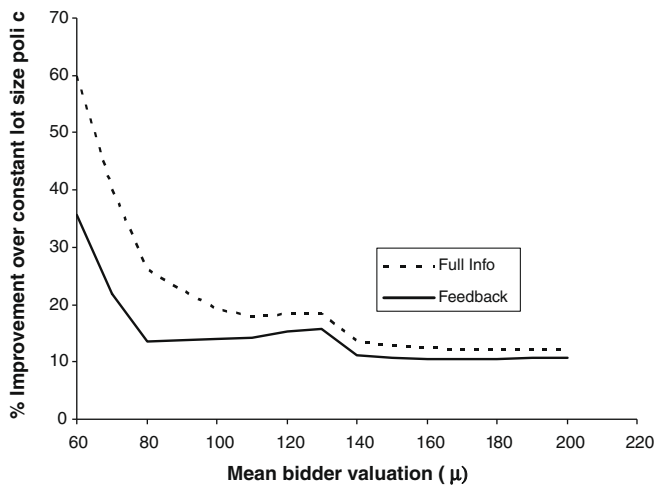


Fig. 3. Percentage improvement in profit over the constant lot size policy when the auction feedback mechanism is used or when there is full information about the valuation dispersion parameter s , as a function of the mean bidder valuation μ . s is $\text{Unif}[1, \mu]$.

fore of knowing s , is less important. As expected, the full information case provides higher profits relative to the feedback learning model and the feedback case converges to the full-information benchmark as the number of bidders, or the mean bidders' valuation increase.

The gap between the performance of the feedback mechanism and the full information case indicates the accuracy of the learning. The accuracy is driven by the sample size available to the learning procedure. When the holding costs are high or and when bidder valuations are low it will be optimal to dump some of the inventory up front and conduct fewer auctions. As a result a larger proportion of the units will be sold in the first auction before any learning can occur. When the number of bidders is small the seller

simply will learn more slowly and thus have a less accurate perception of the dispersion s .

In Fig. 4a and b, we fix the expected value of s and vary its coefficient of variation by adjusting the parameters a and b . In Fig. 4a, the distribution of s centers on 25, and in Fig. 4b, it centers on 50. Three things are happening here. First, as the uncertainty in s increases, it is harder for the feedback mechanism to learn s , and the gap in performance between the auction feedback and full-information cases increases. Second, greater uncertainty in s also increases the benefit of the auction feedback mechanism over the constant lot size policy, which does not use learning. Finally, the auction feedback mechanism with optimal lot sizing does better when the spread in the bidders' valuations is lower, i.e., $E[s]$ smaller, because it can set larger lots in early auctions without as much price risk.

To summarize, we see in Figs. 2–4 that the auction feedback mechanism works quite well in many settings. Furthermore, even when it has difficulty achieving results close to those in the full-information case, it still significantly outperforms the constant lot size policy.

In Section 3.2, when we assumed that the valuation spread parameter was known, we found that the optimal lot sizes decreased with each auction. When the auction feedback mechanism is used, this is not necessarily the case. In Table 3 we display a sample path of the solution of the auction lot-sizing problem with feedback using the same parameters as in the results displayed in Figs. 1–4 ($\mu = \$100$, $C = \$50$, $h = \$15$, $n = 10$, and $t = 1$) and with $s = 26$. For each period, we list the starting inventory, largest spread observed thus far, optimal number of auctions given the current distribution of s , optimal lot size, expected bid price, and expected value of s given the observed bids.

We observe (in Column 4) that the optimal number of auctions dynamically changes. Initially (before the start of auction 1), six auctions are planned, but after one auction is observed, the optimal remaining number of auctions is revised down to three (i.e., ending with auction four). The number of auctions changes repeatedly. We

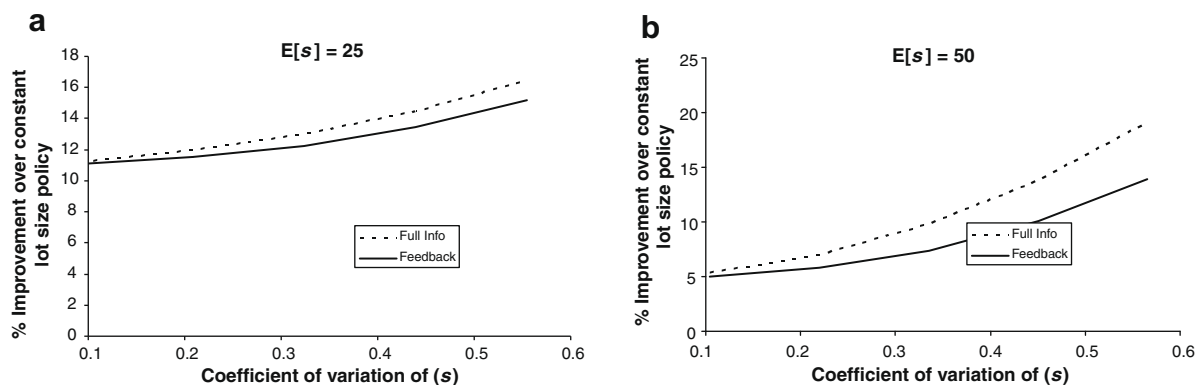


Fig. 4. Percentage improvement in profit over a constant lot size policy when the seller uses an auction feedback mechanism and when there is full-information about the valuation dispersion parameter s , as a function of the coefficient of variation of s , with (a) s uniform with mean 25, and (b) s uniform with mean 50.

Table 3

An example of a sample path of the optimal lot size and the optimal number of auctions when using auction feedback.

Period	Starting inventory	Largest observed spread l_i	Optimal number of remaining auctions	Optimal lot size	Price	$E[s l_i]$
1	30	0	6	7	77.05	50.0
2	23	19.8	3	9	83.12	22.28
3	14	25.3	3	6	96.4	26.72
4	8	25.3	2	5	114.12	26.7
5	3	25.7	2	2	116.97	26.46
6	1	25.7	1	1	103.64	25.8

also observe (in Column 5) that the optimal lot size does not change monotonically. The lot size increases after the first auction and then decreases. It increases because after the first auction, the expected value of the valuation spread, s , is revised downward from 50.0 to 22.28. A smaller spread reduces the risk of low prices and therefore makes a larger lot size more attractive. In Column 7, we see how the auctioneer's belief about the expected value of s rapidly converges with more observations. For example, after observing a maximum valuation spread of \$25 in three successive auctions (30 observations), the probability of seeing a spread of more than \$26 can be calculated using Eq. (5) and shown to be less than 10%.

4. Discussion of modeling assumptions

The performance results of the auction feedback mechanism we have used are, of course, dependent on our modeling assumptions and limited to the numerical experiments we have conducted. We have assumed that the seller observes all the bidders' true reservation values. As we discussed earlier, in practice, there is some censoring of reservation value data because most auctions are not sealed-bid $k + 1$ price auctions and some bidders may not get the opportunity to bid up to their reservation values. An avenue for further research would be to relax some of these assumptions and to develop tools for applying the auction feedback approach in practice. Our work here has demonstrated some of the potential for such tools.

We also have assumed that the bids made in one auction are independent of bids made in earlier auctions. We know that in practice this is not strictly correct, as some losing bidders may try again in later auctions. The main concern for us, regarding these repeat bidders, is not that they exist but rather do they alter the bidder valuation distribution from auction to auction in some significant way? We do not believe this to be a significant effect. To understand why, it is useful to contrast the online setting with the traditional physical world sequential auction. In a traditional auction a fixed number of people are present at the same location for a relatively short time to participate in a sequence of auctions. In such an environment there are a variety of strategic actions bidders might take. They could shade their bids below their reservation values in anticipation of the future auctions with less demand, as some participants may be satisfied already. They might also bid to signal other participants about their willingness to pay for certain items. In the online setting the auctions occur across much longer time frames, there is a constant flow of new participants, and there is uncertainty about if and when additional auctions will ever take place at all. This means that the cost of waiting for future auctions is greater and it is unclear who you are signaling with your bid and the value of the signal. Thus we believe the incentive to engage in such behaviors is diluted.

One could argue that for many goods sold in online auctions the bidders know there will be future and even concurrent auctions for the same or very similar items regardless of what one individual seller does. These outside options should lead to a reduction in the reservation value of those bidders who are willing to wait for a "bargain". From our perspective these behaviors should manifest themselves in the valuation distribution, but are stationary – whenever the market is in an equilibrium.

For the reasons given above, we do not believe that bidders waiting for bargains in future auctions will be a major factor. However, if they were a major factor, we conjecture that they would have two main effects. First, this phenomenon would make it optimal to have even more units in the earlier auctions. This will discourage bidders from "waiting around" and will reduce holding

costs. Second, the Bayesian learning procedure will have to be significantly changed because each auction will represent a sample from a different valuation distribution. All in all, the behavior of consumers at sequential auctions remains rather poorly understood. We see few robust findings, and little ability to discriminate between the different proposed theoretical explanations of this relatively new economic phenomenon (Houser and Wooders, 2005; Seifert, 2006).

We have also assumed that the number of bidders in each auction is a constant. It is a characteristic of online auctions that bidders arrive according to a random process (see Pinker et al., 2003). Although it would complicate our lot-sizing model and turn it into a stochastic dynamic program, we do not believe that explicitly modeling the number of bidders in each auction as a random variable will significantly alter our results. In fact, we have tested our equation for the expected auction closing price with n replaced by a Poisson random variable with mean n . We found that the resulting function of lot size k did not differ much from the deterministic version (see Appendix 3).

It is also possible that in practice the lot size influences the number of bidders. Large lots may attract more bidders than small lots. If that is the case, we believe that it will only reinforce the importance of using an optimal lot-sizing policy derived from an optimization model such as Problem (P1). If there is a strong relationship between n and k , choosing k determines not only the supply in an auction but the demand as well. We have conducted experiments with several functional forms for the relationship between n and k and have found that the marginal benefits of the optimal lot-sizing policy only increase.

Finally, we have assumed that bidders' valuations are drawn from a uniform distribution. We have conducted numerous simulations (see Appendix 3) of the expected price to test the robustness of Eq. (4) with different distribution shapes having the same mean, min, and max as the uniform distribution we used in our numerical experiments. We did these tests also using Poisson numbers of bidders and found little difference in the expected optimal auction closing price as a function of k . We thus have substantial evidence that the qualitative nature of our results is robust to relaxations of the modeling assumptions and in fact may be stronger than demonstrated here.

5. Conclusions

When conducting online auctions, a firm must determine the lot size for each auction and the number of auctions to run. To address these important issues, we develop a multi-period dynamic optimization model of the multi-unit auction management problem. We find that when auctioning multiple units of the same product, simultaneously optimizing lot sizes and the number of auctions can result in significant economic benefits for the seller. We also show that the optimal lot size is monotonically non-increasing across sequential auctions when the distribution of the bidders' valuations is known. Another interesting observation is the sensitivity of the total profit to the total number of auctions regardless of whether the lot sizes are chosen optimally or are identical. Having too few auctions, leads to a larger lot size per auction with a negative impact on the closing price. Splitting the initial inventory across too many auctions, results in excessive holding and administrative costs.

Our analysis and previous work on auctions indicate the importance of the variance in bidders' valuations on the outcomes of auctions. This explains why having the correct estimate for the dispersion of bidders' valuations is important in determining the optimal lot size. Online auctions offer an effective mechanism for collecting extensive detailed data about bidders' actual valuations,

which then can be used to estimate the dispersion in customers' valuations. To exploit this information, we extend our model to include a Bayesian framework for incorporating the results of previous auctions as feedback into management decisions for consecutive auctions of the same item. The number of future auctions and the corresponding lot sizes are updated after each auction. Numerical experimentation shows that incorporating Bayesian analysis into the multi-period, multi-item optimization model results in substantial benefits to the seller, primarily when the number of bidders is relatively small, holding costs are high, and there is a relatively large spread in bidders' valuations. Our work serves as a proof of concept that incorporating bidding information into auction management has a significant potential to increase sellers' profits.

There are several interesting avenues for future research in this area. We have not discussed other potential ways of mining detailed bidding information during or after an auction. In addition, we have not included the minimum bid and reservation price as endogenous design parameters, despite the fact that they may influence bidders' participation. Yet another important and related question involves finding ways that auctioneers can increase the number of bidders through advertising and price promotions. Given the way in which Internet technologies allow for unprecedented innovations in business practices, the optimal management of online auctions will continue to be a fruitful area of research.

Acknowledgement

We acknowledge the helpful comments on earlier versions of this work from seminar participants at MIT Stanford University, UCLA, Harvard Business School, the University of Chicago, the Kellogg School, the University of Rochester, NYU, the University of Pennsylvania, the University of Maryland, the University of Washington, Carnegie Mellon University, Tel-Aviv University, and the Technion.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at [doi:10.1016/j.ejor.2009.05.029](https://doi.org/10.1016/j.ejor.2009.05.029).

References

- Azoury, K., 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science* 31, 1150–1160.
- Bapna, R., Goes, P., Gupta, A., 2000. A theoretical and empirical investigation of multi-item online auctions. *Information Technology and Management Journal* 1, 1–23.
- Bapna, R., Goes, P., Gupta, A., 2001. Online auctions: Insights and analysis. *Communications of the ACM* 44 (11), 42–50.
- Easley, R., Tenorio, R., 2004. Jump bidding strategies in internet auctions. *Management Science* 50 (10), 1407–1419.
- Etzion, H., Pinker, E., Seidmann, A., 2006. Analyzing the simultaneous use of auctions and posted-prices for online selling. *Manufacturing and Service Operations Management* 8 (1), 68–91.
- Fatima, S.S., 2008. Bidding Strategies for Multi-object Auctions. In: Gimpel, H. et al. (Eds.), *Negotiation Auctions and Market Engineering LNBI*, vol. 2. Springer-Verlag, Berlin, Heidelberg, pp. 200–212.
- Gallego, G., Van Ryzin, G., 1994. Optimal dynamic pricing of inventories with stochastic demand over finite-horizons. *Management Science* 40 (8), 999–1020.
- Harris, M., Raviv, A., 1981. Allocation mechanism and design of auctions. *Econometrica* 49 (6), 1477–1499.
- Houser, D., Wooders, J., 2005. Hard and Soft Closes: A Field Experiment on Auction Closing Rules. In: Zwick, R., Rapoport, A. (Eds.), *Experimental Business Research*, vol. II. Kluwer Academic Publishers, pp. 279–287.
- Klein, S., 1997. Introduction to electronic auctions. *Electronic Markets* 7 (4), 3–6.
- Lovejoy, W., 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Science* 36, 724–737.
- Milgrom, P., 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- Moon, Y., 1999. OnSale Inc. Harvard Business Review Case Study-599091.
- Nair, S.A., Karuga, G., Tripathi, 2004. Effects of Sequential Auction Design on Bidding Behavior. University of Connecticut, Working Paper.
- Axel Ockenfels, A., Reiley, D., Sadrieh, A., 2006. Online Auctions, NBER Working Paper 12785.
- Pinker, E., Seidmann, A., Vakrat, Y., 2003. The design of online auctions, business issues and current research. *Management Science* 49 (11), 1457–1484.
- Scarf, H., 1959. Bayes solutions of the statistical inventory problem. *Annals of Mathematical Statistics* 30, 490–508.
- Seifert, S., 2006. Posted-price offers in internet auction markets. *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin.
- Segev, A., Shanthikumar, G., Beam, C., 2001. Optimal design of internet auctions. *Information Technology and Management Journal* 2, 121–163.
- Turban, E., 1997. Auctions and bidding on the internet: an assessment. *Electronic Markets* 7 (4), 7–11.
- Vakrat, Y., 2000. *Optimal Design of Online Auctions*. Unpublished Ph.D. Thesis, University of Rochester.
- Vakrat, Y., Seidmann, A., 1999. Can online auctions beat online catalogs? In: *Proceedings of the International Conference on Information Systems*.
- Vakrat, Y., Seidmann, A., 1999. Analysis and design models of online auctions. In: *Proceedings of the Fourth INFORMS Conference on Information Systems and Technology*, Cincinnati, pp. 64–78.
- Vulcano, G., Van Ryzin, G., Maglaras, C., 2002. Optimal dynamic auctions for revenue management. *Management Science* 48 (11), 1388–1407.
- Wang, R., 1993. Auctions versus posted-price selling. *American Economic Review* 83 (4), 838–851.
- Zipkin, P., 2000. *Foundations of Inventory Management*. McGraw-Hill.