European Journal of Operational Research 215 (2011) 257–267

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Dynamic lot-sizing in sequential online retail auctions

Xi Chen [a], Archis Ghate [a,*], Arvind Tripathi [b]

[a] Industrial and Systems Engineering, University of Washington, Seattle, USA
[b] Information Systems and Operations Management, University of Auckland Business School, Auckland, New Zealand

ARTICLE INFO

Article history:
Received 7 August 2010
Accepted 27 May 2011
Available online 6 June 2011

Keywords:
Auctions/bidding
Dynamic programming
e-Commerce

ABSTRACT

Retailers often conduct non-overlapping sequential online auctions as a revenue generation and inventory clearing tool. We build a stochastic dynamic programming model for the seller's lot-size decision problem in these auctions. The model incorporates a *random number* of participating bidders in each auction, allows for *any* bid distribution, and is not restricted to any specific price-determination mechanism. Using stochastic monotonicity/stochastic concavity and supermodularity arguments, we present a complete structural characterization of optimal lot-sizing policies under a second order condition on the single-auction expected revenue function. We show that a monotone staircase with unit jumps policy is optimal and provide a simple inequality to determine the locations of these staircase jumps. Our analytical examples demonstrate that the second order condition is met in common online auction mechanisms. We also present numerical experiments and sensitivity analyses using real online auction data.

1. Introduction

Online auctions of retail goods have become a significant component of modern internet commerce. Several large retailers such as Dell (www.dellauction.com) and Sam's Club (auctions.samsclub.com) increasingly use online auctions as a revenue generation mechanism (Bapna et al., 2008, 2010). In combination with scrapping excess inventories to firms like Overstock (www.overstock.com), large retailers also use online auctions as an inventory clearing tool. The auction giant eBay (www.ebay.com) and other similar firms such as Ubid (www.ubid.com) provide auction hosting services to retailers like IBM, Sharp and Fujitsu, and also to individual sellers. Companies like Truition (www.truition.com) and ChannelAdvisor (www.channeladvisor.com) specialize in helping businesses conduct online auctions (Odegaard and Puterman, 2006). Based on empirical data available in Vakrat and Seidmann (2000) and Pinker et al. (2010) have noted that most retail auction websites conduct a sequence of multi-unit auctions of identical items. These auctions were also observed to be the operational norm by Bapna et al. (2008), Pinker et al. (2003) and Tripathi et al. (2009). Pinker et al. (2003) have summarized various research issues in such auctions.

Lot-sizes, that is, the number of units to be auctioned in each auction, are one of the key decision variables in sequential auctions (Pinker et al., 2003, 2010; Segev et al., 2001; Tripathi et al., 2009; Vakrat and Seidmann, 2000). A small lot-size induces bidder competition thus increasing the clearing-price. The total revenue may still be lower than one would hope because the number of units sold is small. Uncertainty in the number of participating bidders (demand) in each auction and that in their bids increases decision complexity. For instance, an auction with too large a lot-size may fail due to insufficient demand. Inventory holding costs and the possibility of scrapping inventory to save and recover some of these costs introduce additional economic tradeoffs.

Two papers have investigated inventory scrapping and/or lot-sizing decisions in sequential online retail auctions (Pinker et al., 2010; Tripathi et al., 2009).

Pinker et al. (2010) studied these problems under the following restrictions: a fixed number of participating bidders in each auction, uniform bid distributions with support [0,1], and a truth revealing multi-unit Vickrey mechanism. These assumptions enabled them to formulate a deterministic dynamic program, wherein a closed-form lot-sizing policy was derived by equating derivatives of value functions to zero within a backward induction procedure. The optimal lot-size decreased at a constant rate from one auction to the next. This rate increased with inventory holding costs and decreased with the number of bidders per auction. In their model, it was optimal to scrap inventory only one time before beginning the entire sequence of auctions.

Tripathi et al. (2009) also assumed a fixed number of participating bidders in each auction, and employed a multi-unit Dutch mechanism. Using uniform bid distributions, they first optimized the lot-size over a sequence of auctions assuming that the lot-size

* Corresponding author. Address: Industrial and Systems Engineering, Box 352650, University of Washington, Seattle, WA 98195, USA. Tel.: +1 (206) 616 5968; fax: +1 (206) 685 3072.
E-mail address: archis@uw.edu (A. Ghate).

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.05.051

did not change over time. This led to a simple closed-form lot-size expression that resembled the well-known Economic Order Quantity (EOQ) formula in inventory management (Heyman and Sobel, 2003). They also devised a goal programming method to estimate bid distributions from online bid data.

Segev et al. (2001) focused on predicting auction clearing-prices using an orbit queue Markov chain model, and compared these predictions with data obtained from Onsale (www.onsale.com), a Sillicon Valley start-up company. They proposed a deterministic dynamic programming model for lot-size optimization under the restrictive assumption that all items on sale will be sold owing to a sufficiently large number of participating bidders but did not attempt to solve it.

Odegaard and Puterman (2006) considered an auctioneer with two identical items on hand, and determined an optimal time-point at which the second item should be "released" for an auction. They derived conditions to ensure an optimal control-limit release-time policy. This control-limit was decreasing in holding cost.

Vulcano et al. (2002) studied a problem motivated by airline ticket selling websites like Priceline (www.priceline.com). The seller first observed bids from potential travelers, and then chose how many and which bids to accept, as opposed to publicly pre-committing lot-sizes at the beginning of each auction before receiving bids as practiced in retail auctions (Odegaard and Puterman, 2006; Pinker et al., 2010; Segev et al., 2001; Tripathi et al., 2009). Consequently, they solved a variable supply allocation problem rather than a lot-size optimization problem to obtain a structural result similar to ours but utilized different mathematical analysis and sufficient conditions as developed by Myerson (1981) and Maskin and Riley (2000). This work was later extended to an infinite-horizon joint auctioning and pricing problem under holding and ordering costs (van Ryzin and Vulcano, 2004).

The basic setting in our paper is similar to Pinker et al. (2010) and Tripathi et al. (2009) in that we consider a seller who conducts a sequence of non-overlapping online auctions of retail goods. However, in contrast to their work, we incorporate uncertainty in the number of participating bidders (stochastic demand) in each auction; do not restrict our formulation to any specific clearing-price determination rule; and allow for any bid distribution (see Section 2 for details). To the best of our knowledge, this is the first paper that successfully overcomes all mathematical difficulties introduced by this generalization in the retail pre-committing setting to provide a complete structural characterization of optimal inventory scrapping and lot-sizing policies as in Theorem 2.1.

More specifically, under the second order condition (7) on the single-auction expected revenue function, we show that a threshold inventory-scrapping policy, and a monotone staircase with unit jumps lot-sizing policy are optimal. This condition roughly requires that the marginal single-auction expected revenue, normalized by the probability of sufficient number of bidders participating, be decreasing in lot-size. It is then optimal to scrap all inventory above a time-dependent threshold inventory level, and not to scrap any inventory below the threshold. This threshold equals the inventory level at which the scrap-value of a unit exceeds its marginal value over all remaining auctions. Moreover, if lot-size x is optimal in post-scrapping inventory i, then either lot-size x or lot-size $x+1$ is optimal in post-scrapping inventory $i+1$. This unit jump in optimal lot-size occurs when the normalized marginal single-auction expected revenue exceeds the discounted marginal value of saving the additional unit for future auctions. See Theorem 2.1 and its proof in Section 3 for precise detailed versions of these statements. Section 3.1 includes several examples where our second order condition is met. Numerical results and sensitivity analyses conducted using real online auction data are presented in Section 4. Limitations and potential extensions of our model are discussed in Section 5.

2. Problem description and mathematical formulation

Consider a seller with some initial inventory of identical units on hand. We assume that the seller conducts a sequence of $1 \leqslant T < \infty$ auctions indexed by $t = 1, 2, \ldots, T$. The seller uses a fixed auction mechanism in all auctions and this mechanism is disclosed to the bidders. Examples of auction mechanisms include multi-unit Vickrey as on eBay, multi-unit Dutch as on Sam's Club, and Yankee as on Ubid.

Under stochastic demand, one-shot scrapping as in Pinker et al. (2010) may not be optimal; in fact, it may lead to negative marginal values. It is essential to dynamically exploit the flexibility to scrap inventory *even if* the scrap-value is zero. Thus, at the beginning of auction t, the seller makes two decisions after observing inventory i on hand: (i) the number of units y to be scrapped for a value of $s \geqslant 0$ per unit, and (ii) of the $i - y$ remaining units, the lot-size x to be put up for the tth auction. This lot-size is disclosed to the potential bidders at the start of the tth auction.

A random number N of bidders who wish to buy one unit each then place their bids. The probability mass function (pmf) of N is denoted $g(\cdot)$, its support is denoted $\mathcal{N} \subseteq \mathbb{N}_+$, and its distribution function is denoted $G(\cdot)$. Consistent with the existing literature, we assume that bidders are independent across auctions, and identical both within an auction and across auctions. A detailed discussion of practical limitations introduced by this assumption is included, for instance, in Section 4 of Pinker et al. (2010), and we do not repeat it here (also see Section 5). More specifically, the final bids in all auctions are independent and identical (iid) random variables B with distribution $F(\cdot)$, finite expectation, and support $\mathcal{B} \subseteq \mathbb{R}_+$ whose smallest element is denoted L. The existence of a probability density function for B is not assumed, and in particular, B may be discrete. We remark that our setting is flexible and general enough to allow bid distributions that are statistically estimated using online data, those derived from game theoretic analyses of how bidders might behave in sequential online multi-unit auctions, and the ones obtained through a combination of these two approaches (see Bapna et al., 2002, 2003, 2008; Fatima, 2008; Jiang and Leyton-Brown, 2007; Pinker et al., 2010; Tripathi et al., 2009 for examples of such techniques).

If the actual number of bidders n in an auction is more than the lot-size x, the seller generates revenue through bidder competition by selling all x units. We denote this revenue by $\pi(x; n)$, and it is derived from the expected value of a mechanism-specific order-statistic of $F(\cdot)$. See the beginning of Section 3.1, and in particular, Eqs. (17)–(19) for specific examples of $\pi(x; n)$. If $n \leqslant x$, the auction fails due to a lack of bidder competition. Note that this scenario does not arise in Pinker et al. (2010) and Tripathi et al. (2009) owing to their assumption of deterministic demand. When an auction fails, the seller sells one unit to each of the n bidders for amount L (Pinker et al., 2003). Equivalently, in the language of Pinker et al. (2003), the "minimum bid" of the auction is set to L. This is better than selling for any price less than L. The seller may however benefit from using a minimum bid requirement of some $\Lambda > L$, and then selling each unit in a failed auction for Λ. This introduces challenging tradeoffs, which require dynamic optimization of Λ and are not the focus of this paper (also see Section 5). Another entirely different possibility for the seller is to cancel a failed auction, returning the x units originally intended for sale back to the inventory held. Canceling auctions disappoints bidders who did participate, leading to negative feedback from these unsatisfied bidders. This hurts the seller's reputation that is critical for success in e-commerce (Ba and Pavlou, 2002; Resnick et al., 2006). We therefore do not follow this alternative approach.

The holding cost of carrying each unit in inventory during auction t is denoted $h \geqslant 0$, and it is assumed to be incurred for the