European Journal of Operational Research 245 (2015) 555–570

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Optimal minimum bids and inventory scrapping in sequential, single-unit, Vickrey auctions with demand learning

Archis Ghate*

Department of Industrial and Systems Engineering, University of Washington, BOX 352650, Seattle, WA 98195, USA

ARTICLE INFO

Article history:
Received 3 September 2014
Accepted 10 March 2015
Available online 16 March 2015

Keywords:
Sequential auctions
Bayesian learning
Dynamic programming

ABSTRACT

Retailers often conduct sequential, single-unit auctions and need to decide the minimum bid in each auction. To reduce inventory costs, it may be optimal to scrap some of the inventory rather than holding it until it is auctioned off. In some auctions, the seller may be uncertain about the market response and hence may want to dynamically learn the demand by observing the number of posted bids. We formulate a Markov decision process (MDP) to study this dynamic auction-design problem under the Vickrey mechanism.

We first develop a clairvoyant model where the seller knows the demand distribution. We prove that it is optimal to scrap all inventory above a certain threshold and then auction the remaining units. We derive a first order necessary condition whereby the bidders' virtual value at an optimal minimum bid equals the seller's marginal profit. This is a generalization of Riley and Samuelson's result from the one, single-unit auction case. When the virtual value is strictly increasing, this necessary condition is also sufficient and leads to a structured value iteration algorithm.

We then assume that the number of bidders is Poisson distributed but the seller does not know its mean. The seller uses a mixture-of-Gamma prior on this mean and updates this belief over several auctions. This results in a high-dimensional Bayesian MDP whose exact solution is intractable. We therefore propose and compare two approximation methods called certainty equivalent control (CEC) and Q-function approximation. Numerical experiments suggest that Q-function approximation can attain higher revenues than CEC.

1. Introduction

Private and public institutions, large retailers, and individual sellers use sequential, single-unit, online auctions of identical items as a revenue generation and inventory management tool (Chen, Ghate, & Tripathi, 2011; Pinker, Seidmann, & Vakrat, 2003, 2010). These auctions are conducted either on the seller's own website as for example at Dell (dellauction.com) and Sam's Club (auctions.samsclub.com), or through a facilitator such as eBay.com or Liquidation.com. Other examples of retail auction websites include eBid.net, uBid.com, bidz.com, CQout.com, QuiBids.com, bidcactus.com, and skoreit.com. Owing to the large number and the significant total monetary value of retail goods sold in online auctions, there has been a growing interest in the design and analysis of such auctions.[1]

The minimum bid, which is also often called the starting bid or the public reserve price, is the key design variable in single-unit, online auctions of retail goods (Pinker et al., 2003). The seller's minimum bid requirement is posted on the website along with other transaction rules.[2] The auction is set in motion when a bid above this minimum is posted. The minimum bid serves as a barrier to entry, and hence by requiring a low minimum bid, the seller hopes that she will be able to induce enough bidders to participate, thereby increasing competition

* Tel.: +1 2066165968.
E-mail address: archis@uw.edu

[1] The corporate website of eBay, http://www.ebayinc.com/who, accessed 15.02.12, stated that "with more than 97 million active users globally (as of Q2 2011), eBay is the world's largest online marketplace, where practically anyone can buy and sell practically anything. Founded in 1995, eBay connects a diverse and passionate community of individual buyers and sellers, as well as small businesses. Their collective impact on e-commerce is staggering: In 2010, the total value of goods sold on eBay was $62 billion – more than $2000 every second." Liquidation.com, website accessed on 15.02.12, stated that their "marketplaces provide nearly 1.2 million registered professional buyers access to a global, organized supply of wholesale, surplus and salvage assets in over 500 product categories. Since inception, it has conducted over 1.8 million online transactions generating over $1 billion in gross merchandise value."

[2] eBay, Liquidation.com, uBid.com, CQout.com, Sam's club, bidz.com, dellauction.com, and others offer this format. For instance, dellauction.com states (accessed 28.05.13), "Bid. Win. Enjoy... It's that simple! Auctions start at $.99." CQout.com states (accessed 28.05.13), "Worried that your item will sell for £10's when you wanted £100's? It won't. When you list an item for auction you can set a reserve price below which the item will not be sold. Only once bidding reaches or exceeds the reserve are you, as the seller, obliged to sell the item for that price to the highest bidder."

http://dx.doi.org/10.1016/j.ejor.2015.03.015
0377-2217/© 2015 Published by Elsevier B.V.

and hence the closing price (Bajari & Hortacsu, 2003; Kaiser & Kaiser, 1999; Ku, Galinsky, & Murnighan, 2006; Pinker et al., 2003). However, a low minimum bid requirement allows bidders to post small bids and if a sufficient number of bids are not received, it can lead to a low closing price. On the other hand, a high minimum bid requirement can turn a significant number of bidders away, but it ensures a good closing price if at least one bid is posted. These economic trade-offs were studied empirically by Vakrat (2000) in hundreds of online auctions, and they are discussed in detail in Pinker et al. (2003).

When a seller conducts a sequence of single-unit auctions of retail goods online, additional economic trade-offs in the minimum bid decisions become relevant. Bidders in one auction compete with other bidders in the same auction directly but they also compete with bidders in other auctions indirectly through the opportunity cost of finite inventory. Moreover, owing to the uncertainty in the number of bidders (this is a well-known feature of online retail auctions (Pinker et al., 2003)) and also in the bidders' valuations, inventory evolves stochastically and hence optimal minimum bids should vary dynamically with inventory on hand. Optimal minimum bids should also depend on the cost of holding inventory. In fact, Chen et al. (2011) and Pinker et al. (2010) have shown that in the presence of inventory costs it may be optimal to dynamically scrap some of the inventory rather than holding it until it is auctioned off.[3] This interplay between scrapping and minimum bid decisions further complicates the seller's problem.

Finally, in some auctions, and particularly in auctions of fashion goods, of seasonal products, of new products, and where the seller is new to the market, the seller may not know *a priori* the distribution of the number of potential bidders, that is, the demand distribution, in its entirety. This is because sellers of such items typically face significant uncertainty about market response, for example, due to varying fashion trends, and more generally, owing to the lack of sufficient past demand data about similar products (Araman & Caldentey, 2009; Aviv & Pazgal, 2005a, 2005b; Farias & Van Roy, 2010; Potoff & Beil, 2007). In this scenario, the seller can learn the demand distribution sequentially over multiple auctions by observing the number of posted bids (bidders that are excluded by the minimum bid requirement are not observed). This introduces an additional, exploration *versus* exploitation trade-off between demand learning and revenue maximization.

We present a Markov decision process (MDP) model to optimize these various trade-offs in minimum bid and inventory scrapping decisions in sequential, single-unit, online Vickrey auctions that a seller conducts while simultaneously learning demand. This paper is organized as follows. Our work is positioned in the context of the existing literature in the next section. In Section 3, we first present the clairvoyant MDP model, where the seller is assumed to know the demand distribution. In Section 3.1, we state the structure of an optimal policy for this clairvoyant MDP. Detailed economic intuition behind this structure is provided in Section 3.2. A streamlined version of the value iteration algorithm, which exploits this optimal policy structure, is proposed in Section 3.3. This algorithm is implemented on several numerical examples in Section 4. Specifically, in Section 4.1, our examples focus on a Poisson-distributed number of bidders with Beta-distributed valuations. In Section 4.2, we conduct sensitivity analyses for these examples: we explore the directional dependence of optimal minimum bids and of optimal scrapping thresholds on the mean demand, the holding cost, and the scrapping price. Our Bayesian model for learning the Poisson demand is developed in Section 5. The relevant Bayesian update formulas for a mixture-of-Gamma prior on the mean Poisson demand are introduced in Section 5.1. These formulas are utilized in the development of our Bayesian MDP formulation in Section 5.2. Exact solution of this MDP is intractable; we therefore propose two approximation methods in Sections 5.3 and 5.4. These two methods are compared via computational experiments in Section 6. Finally, we conclude by discussing some of our modeling assumptions, and outline directions for future research in Section 7. Proofs of all results are provided in four appendices as supplementary material.

2. Literature review and our contributions

Our work relates to two distinct streams of research in operations management: the first is about optimal design of online auctions, and the second is about demand learning in inventory control and dynamic pricing problems. We attempt to review only the most relevant papers from this vast literature here.

Yu, Wang, and Dang (2006) studied the problem of selling a single unit in a single eBay-like auction, and derived an implicit first order necessary condition for an optimal minimum bid. An extension of this problem, where the seller can reauction the unit if it did not sell in previous auctions, was also discussed. One contribution of Yu et al. was that they incorporated the effect of the so-called insertion fee that the seller incurs for conducting an online auction on eBay into the classic single-auction model of Riley and Samuelson (1981) for traditional auctions. In addition, Yu et al. also enhanced Riley and Samuelson's classic model, which assumed a fixed number of bidders, by allowing a Poisson distributed number of bidders with a known mean. Yu et al. did not consider sequential auctions and also did not consider inventory costs, scrapping, or demand learning.

Chapter 6 in the doctoral dissertation of Hawkins (2003) pursued a data-driven approach to optimal design of sequential, online auctions. Hawkins sought to make as few assumptions about bidder behavior as possible, and to simultaneously optimize multiple auction-design variables such as minimum bid, buy-it-now price, and auction-duration, with models built using publicly available online data from eBay. Owing to the complexity of the resulting models, analytical results were not possible. An experimental, approximate dynamic programming approach was instead employed to numerically obtain good values of design variables. Hawkins did not model demand learning.

Vulcano, Van Ryzin, and Maglaras (2002) presented a sequential allocation model where the seller determines the number of units to release from her inventory after observing submitted bids. Each unit has a hidden reserve price that depends on the inventory on hand. The seller is better-off by providing herself with this additional flexibility whereby she does not commit to allocation decisions until after seeing all bids in each auction. However, Hawkins (2003), Odegaard and Puterman (2006), and Pinker et al. (2010) have noted the somewhat peculiar nature of this non-committing mechanism; Pinker et al. (2010) observed that it had not been used in online retail auctions. This non-committing setting in Vulcano et al. was motivated by airline ticket allocation websites such as Priceline.com. Consistent with this motivating application from airline revenue management, Vulcano et al. did not consider inventory costs or scrapping. A special case of our clairvoyant model in Section 3 with no inventory holding costs and zero salvage value of scrapping is similar to a special case of the setting in Vulcano et al. with single-unit, Vickrey auctions. Note, however, the subtle point that the seller in Vulcano et al. needs to use a modified Vickrey mechanism to induce the bidders to bid their true valuations (see their Theorem 2 on p. 1396). The model in Vulcano et al. was generalized to include inventory holding and ordering over an infinite horizon in van Ryzin and Vulcano (2004). Again, neither of these two papers incorporated demand learning.

There is a parallel body of literature on sequential auctions, where the seller does not use a minimum bid but instead optimizes the

[3] For example, the store T. J. Maxx buys excess inventory from other retailers. Their website, http://www.tjmaxx.com/how-we-do-it/, accessed 30.06.12, stated that "when a designer overproduces and department stores overbuy, we swoop in, negotiate the lowest possible price, and pass the savings on." Another similar company, Overstock.com, also offers retailers an alternative inventory clearing channel.