Using Transaction Data for the Design of Sequential, Multi-unit, Online Auctions[1]

Edieal Pinker [*]

Abraham Seidmann [*]

Yaniv Vakrat [+]

Simon School Working Paper CIS00-03

Revised: October 25, 2003

[*] W. E. Simon Graduate School of Business Administration, University of Rochester, Rochester, NY, 14627.

[+] McKinsey and Company, Palo Alto, CA, 94304, and The Graduate School of Business, Stanford University, Stanford, CA., 94305.

---

[1] We acknowledge the helpful comments on earlier versions of this work from seminar participants at MIT, Stanford University, UCLA, 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, the Technion, WISE'99, ICIS '99, and HICSS2000.

Abstract

Internet auctions for consumers' goods are an increasingly popular selling venue. The Internet's computational ability makes possible the sale of multiple units of the same good in a single auction. Many sellers, instead of offering their entire inventory in a single auction, split it into sequential auctions of smaller lots, to reduce the negative market impact of large lots. Information technology also makes it possible to collect and analyze bid data from online auctions. In this paper we develop and implement a new model of sequential online auctions to explore the potential benefits of using real bid data from earlier auctions to improve the design of future auctions.

Assuming a truth-revealing ascending auction model, we quantify the effect of the auction lot size on the closing price. We then develop a model for allocating inventory across multiple auctions. We investigate how the available inventory should be split into multiple lots and how many sequential auctions should be run. Solving the dynamic programming formulation, we prove that the lot size drops from period to period. The rate of decline in the optimal lot size increases with the number of bidders, the holding costs, and the discount factor, while it decreases in the dispersion of consumers' valuations of the good. Finally, we extend this model to incorporate dynamically the results of previous auctions as feedback into the design of consecutive auctions, updating the lot size and number of auctions. We demonstrate that 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 show how the benefits of using detailed transaction data for the design of sequential, multi-unit, online auctions is influenced by the inventory holding costs, the number of bidders, and the dispersion of consumers' valuations.

Key words: Information Technology, Information Systems, Internet, Auctions

1. Introduction

Online auctions rank among the more successful uses of the Internet for commerce. An increasing number of companies see the web as an additional marketing channel complementing their stores and posted-price online catalogs. For example, the Sam's Club[1] auction site has many single and multi-unit auctions of electronics and computers, while on Dell's site[2] both Dell and other sellers auction numerous new and used computers. Dovebid[3] specializes in auctions liquidating the assets of businesses.

The reason for the growing popularity of online auctions is that the Internet has revolutionized auctions in several ways. The computational ability at the auction site for the first time enables the conduct of multi-unit auctions with sophisticated allocation rules based on price, quantity, or time of arrival. The Internet also allows the participation of geographically separated sellers and bidders, who can track progress in real time and search for products at a relatively low cost using software tools such as BidXS.com[4] and AuctionPatrol.com.[5] Using the multimedia capabilities of the Internet, sellers can provide rich descriptions of the items being auctioned and therefore are not very limited in the variety of products that they can sell. Online auctions have thus become a viable and active sales channel for many more 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 automatically collected. This information has tremendous value to the auctioneer, since it provides insight into bidders' valuations for the auctioned goods and thus can be used in the design 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, 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

---

[1] http://auctions.samsclub.com/.
[2] http://www.dellauctions.com/.
[3] http://www.dovebid.com.
[4] http://www.bidxs.com/.
[5] http://www.auctionpatrol.com/.

clearing price. This means that in the auction environment companies can determine the quantity offered, while the market determines the price in response to this quantity. The greater the lot size offered, the lower the price received for each unit. We call this phenomenon the *market impact* of the lot size. As a result of these marked differences, traditional inventory management approaches (Zipkin 2000), which assume a fixed price and an uncertain quantity of sales, address decisions that are not as relevant to the auction setting. Moreover, 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. The highest bids made by these individuals were used as proxies for their true reservation prices.

In this paper, we model a similar strategy of using the bid data from previous auctions to improve the design 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.[6]

At all major Internet auction websites, it is common to observe sequential auctions for the same good. 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. Sequential auctions and the market impact of the lot-size decision 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

---

[6] 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 Seidmann and Vakrat (1999a)).

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. By conducting sequential auctions for the same product and observing the leading bids file, the auctioneer can learn about customers' valuations for the product. This capability raises the following questions: how can one implement the learning process; how will it effect the operational decisions; and what are its economic benefits to the firm?

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 the 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 the operations of a firm through the integration of inventory management and multi-unit auctions.
- 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 design of future auctions by embedding the inventory management and auction planning decisions within a practical Bayesian analysis framework.

The paper proceeds as follows. In § 2, we briefly review some of the relevant literature and outline the theoretical background, which is necessary to understand the main results of the paper. In § 3, we initially model the multi-unit sequential online auctions, assuming that the bidder valuation distribution is known, 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 bidders' valuations and that setting the lot size optimally can have a significant impact on profits. In § 4, we extend our model to include the mechanism we have developed for analyzing auction data to improve future auctions. In this framework, each auction provides the auctioneer with *feedback* that can be used to improve the auction design. We investigate the value of using this feedback in conjunction with the dynamic optimization of the auction design. We conduct an extensive set of numerical experiments and determine under what circumstances this sophisticated approach to auction design yields significant benefits to the auctioneer. We offer our conclusions in § 5.

2. Literature review

The Internet has created new opportunities for the use of auctions. As a result, firms face questions that have not yet been addressed by the literature. For broad overviews of how the Internet has changed auctions, see Klein (1997), Turban (1997), and Pinker et al. (2001). **{This might be better as a footnote.}** In online auctions, the auctioneer decides on the quantities that are put on sale, while market forces determine the auction's closing prices. The quantity offered is a decision variable, while the price is determined by market forces.

In this paper, we consider a firm with a given initial inventory that it wishes to sell using online auctions. We therefore look at multi-period *profit* maximization with an unknown distribution of bidders' valuations, where the revenue function is derived from an auction model. In our model, the price is endogenous and depends on the lot size due to its negative market impact. We develop a Bayesian feedback mechanism for learning about the unknown distribution of bidders' valuations.

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, as it becomes available, into the decision model. Most of the results in 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 problem and consequently our results are very different.

There is an enormous body of theoretical literature on auctions. Most papers analyze agents' strategies when an auctioneer wants to sell a single item and tries to achieve the highest possible payment for it, while each bidder wants to acquire that item at the lowest possible price, although there are many papers on multi-unit auctions as well. The interested reader can refer to excellent surveys by Klemperer (1999, 2000). In this paper we assume that the seller conducts private value, multi-unit auctions in which each bidder only has demand for a single unit. Below, we briefly review the auction theory that is directly relevant.

The most common single-item auction mechanism found on the Internet today is the English auction (see Lucking-Reilly (2000) and Beam and Segev (1998)). In its traditional form, this is an oral, open, ascending-price auction. Each bidder is free to raise her bid, until there is only one bidder left. Obviously, this winner must be the one with the highest bid. It is important to observe, however, that the price she pays for the good is the *reservation price* (the upper limit valuation) of the bidder with the *second-highest* value (plus possibly ε, i.e., a tiny increment to go above the second-highest reservation price). This is exactly the intuition behind the well-known revenue equivalence result: in the *independent private value environment* the English auction is equivalent to the second-price sealed-bid auction, (Vickrey 1961). We assume, in this paper, that bidders have independent private values. In *private value* auctions, the value of the good depends only on the bidder's own preferences. The key is that the winning bidder may not resell the item, as in that case the value would depend on other bidders' valuations. Online auctions for consumer goods, such as those we analyze in this paper, fall into this category.[7] Each Internet shopper may have a different consumption value for the good being auctioned.

---

[7] In B2C and C2C online auctions, although some buyers/bidders are resellers, most are individual consumers (Vakrat and Seidmann 2000). Under such circumstances, the private value assumption seems reasonable. Realistically, one can think of each product being offered as having both a common and a private value component. For the purposes of this paper we are assuming that the private component dominates.

Vickrey (1961) proposed the second-price sealed-bid auction, and it is often known as the Vickrey auction. In the Vickrey auction, the highest bidder wins the auction but only pays the second highest bid, and each participant bids without knowing the bids of the others. Vickrey (1961) shows that this auction mechanism has the useful property that it is truth-revealing: it is optimal for a bidder to bid her true valuation for the product. If she bids higher, she may end up with a loss if she wins. If she bids lower, there is a smaller chance of winning, but the winning price is unaffected--it is determined by the second-highest valuation anyway. In general, an agent's strategy is defined as her bid as a function of her private value and prior beliefs about other agents' valuations.

The most common mechanism for online multi-unit auctions is frequently called the "Yankee auction." In a Yankee auction, one or more identical items are auctioned at the same time. Bidders specify a price and the quantity of items they wish to purchase. Bids are typically ranked by price first and then quantity. Winning bidders pay their bid price. When bidders only have demand for a single item, the Yankee auction is the multi-item generalization of the single-item English auction discussed above. The second-price sealed-bid model can also be extended to multi-unit auctions in which k units ($k \geq 1$) are offered in a single auction and bidders have single-unit demands, in what is called a sealed bid *$(k+1)$ - price* auction (Vickrey 1961). In this auction, the uniform price is set by the bid on the $k+1^{th}$ unit (since there are only k units on sale, it is not a winning bid). The dominant strategy is again to reveal fully one's private information (i.e., one's reservation price). As in the single-unit auction, the Vickrey mechanism is truth-revealing because the winning bidders do not determine the price they pay. At the same time, revealing the true value guarantees a non-negative consumer surplus. Although Rothkopf et al. (1990) explain why practitioners refrain from using Vickrey auctions, the Vickrey auction serves as a good approximation for determining the outcomes of auctions in the real world. Harris and Raviv (1981) show that when bidders have independent private values and are risk-neutral, the expected revenue from the Yankee auction is equivalent to that of the sealed bid *$(k+1)$ - price* auction. This means that when modeling multi-item auctions we can use the *$(k+1)$ - price* format to represent situations in which the Yankee mechanism is actually used and accurately capture the seller's revenue.

Although online auctions are receiving considerable attention from management scientists and economists (see Pinker et al. (2001) for a survey), there has been little analytical modeling of the optimal ways to operationalize online auctions. Here we review some exceptions. Bapna, Goes, and Gupta (2000 and 2001), who develop a discrete model to compare the current online auction practice with an ideal Vickrey mechanism. They classify online bidders into opportunists, participators, and evaluators, and they study the impact of the bidding increment. Other aspects of online auctions are addressed by Beam and Segev (1998), who conducted a field study of some of the early practices in online auction businesses. 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 design, so long as it can be shown that the design changes do not alter the predictions. Vakrat and Seidmann (1999a) have found empirical evidence of significant discounts at online auctions over online catalog prices for the same goods. These results suggest that online auction participants 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.

Vulcano et al. (2002) model a firm with a fixed initial inventory that receives bids 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 value threshold set dynamically by the firm. They propose a novel auction scheme in which the bidders reveal their bids and then the firm determines the value thresholds for that auction. They show numerically that this mechanism outperforms 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 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 auction scheme they propose has yet to be seen in practice. In our paper we specifically model auctions with committed quantities, as are found today. By making specific distributional assumptions about bidders' valuations, we are able to derive detailed results about optimal lot-sizing polices that are informative to managers. In addition, we are able to develop a Bayesian learning mechanism for analyzing auction data to improve future auctions.

3. Optimization of multi-unit sequential auctions with a known bidder valuation distribution

Suppose that a firm has a fixed number of units, x_0, to auction. It can offer the whole lot in one auction, or it can split the lot among multiple separate sequential auctions spread over a longer period of time. Running a number of sequential auctions over a longer period of time 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 tradeoffs 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, in this section, we develop a multi-period model of a firm auctioning off an initial inventory of a product. 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. We assume that the number of bidders (or demand) in each period is a constant, n, that is independent of the auction's lot size. We also assume the following: bidders have independent private values for the good being sold; bidders' valuations follow a uniform distribution; bidders in one auction are independent of bidders in other auctions; bidders have single-unit demand; and auctions are conducted using a Vickrey format (sealed bid ($k+1$) - price**{The italicization is inconsistent: should it be like this, or should "- price" also be italicized, as above?}**). Our modeling assumptions best suit a situation in which consumers are shopping for an item and come to the auction. If they do not succeed in making a purchase, they will go elsewhere. 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), and because the sellers have no interest in providing advance information regarding follow-up auctions. .

The assumptions of single-unit demand and independent private values are standard in the auction literature. At the end of this section we discuss the sensitivity of our results to our other assumptions.

3.1 Model formulation

We start with determining the price in a single multi-unit auction. Our method is based on a single-period model that was developed in Vakrat and Seidmann (1999b). For completeness, we briefly outline that single-period model in what follows. Suppose there are n independent bidders. Bidder i, $i=1,\ldots,n$, has a fixed reservation price, v_i, beyond which she will not bid. For simplicity, we assume that these values are independently and uniformly distributed on $[\mu - s, \mu + s]$, where μ is the distribution mean and $\boldsymbol{s}$ is its dispersion. Each bidder is interested in one unit, and if she wins the auction she receives a utility U($v_i - p$), where p is the price paid by bidder i if she wins the auction. The utility function is increasing in its argument. In addition, the bidders are risk-neutral.

Assuming a Vickrey auction format, bidder i's bid is given by the bidder's valuation (or reservation price) for a single unit of the product, v_i. Let $v_{(1)},\ldots,v_{(n)}$ be the order statistics defined on the actual reservation prices, where $v_{(1)}$ is the highest bid, $v_{(2)}$ is the next highest bid, and so on. From the analysis above, the bidders with values $v_{(1)},\ldots,v_{(k)}$ are the winning bidders, and they will pay a price equal to $v_{(k+1)}$. The expected value of the auction price, p_a, given n bidders in the auction and $k \le n$, can be written as

$$\begin{aligned} E\left[p_a \mid n,k\right] &= E\left[v_{(k+1)} \mid n\right] = \int_{x=\mu+s}^{\mu+s} x f_{v(k+1)|n}(x)\,dx \\ &= \mu - s + 2s\frac{n-k}{n+1}. \end{aligned} \tag{2}$$

We can see in Equation (2) 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 μ - s. If we define $p(k, n)$ as the expected auction closing price when the lot size is k and there are n bidders, then

$$p(k,n) = \begin{cases} \mu + s - 2s\, m\,(k+1) & \text{for } k < n \\ \mu - s & \text{otherwise,} \end{cases} \qquad (3)$$
$$where \;\; m = 1/(n+1).$$

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 market impact is inversely related to the number of bidders. We will henceforth refer to m as the marginal market impact.

Assuming that the seller plans T auctions, let $i = 1,2,\ldots,T$ be the index of the period, and let k_i be the auction lot size in period i. 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 $k_i \in [0, x_i]$. The firm must select T so that $J^*_{i,T}(x_i)$ is maximized. In each period, the auctioneer's revenue is determined by the number of units he decides to offer times the auction's closing price. Revenue in each period therefore is $k_i\, p(k_i,n)$.

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 $x_i ht$ on the current stock. Finally, we assume a fixed cost of C to run each auction, which includes, but is not limited to, advertising, 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 number of auctions T and initial inventory x_1, recursively as the following deterministic dynamic program, (P1):

$$\begin{aligned}
&\text{Problem P1:}\\
&J^*_{i,T}(x_i) = \underset{k_i \in [0,x_i]}{Max}\, J^*_{i+1,T}(x_i - k_i) + k_i p(k_i) - htx_i - C\\
&for\ \ i = 0,1,.....,T\\
&where: \qquad (1)\\
&x_i = x_{i-1} - \min(k_{i-1}, n),\\
&k_i \le n,\\
&k_i \le x_i \ \ and \ \ x_i, k_i \ \ \text{integer.}
\end{aligned}$$

The lot size may not exceed the current inventory level; the seller cannot auction items he does not possess. The firm can solve (P1) for different values of T and different values of $x_1 \le x_0$ and pick the pair that maximizes profits. If $x_1 < x_0$, the firm scraps $x_0 - x_1$ units before starting the auctions. For brevity, we have formulated the model without discounting because in a finite-horizon model discounting greatly complicates all the equations. In the Appendix we present our main results from this section with discounting. Numerical experiments have shown that discounting does not change the general nature of our results.

3.2 Analysis

Given the bounded closing price function in Equation (3) 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 x_i can be non-integer.

Assumption 2: We assume that $n > x_1/T$ so that there is sufficient demand to sell all the items in inventory. Otherwise, the seller must scrap the excess x_1-nT units to avoid unnecessary holding costs.

Assumption 3: $\frac{8s(nT-x_1)}{T(T-1)(n+1)} > ht$. This assumption guarantees that the lot size will always be less than n in our analytical results. We discuss this in more detail later.

We first show that we only need to consider lot sizes greater than zero.

Remark 1: *If the firm conducts T auctions, then $k_i>0$ for all $i=1,2, \ldots, T$.*

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 1-3, define $\boldsymbol{T_f}$ to be a set of T such that*

$$x_i \geq \frac{x_i}{T-i+1} + \frac{ht(T-i)(n+1)}{8s} \text{ for all } i \leq T.$$

The optimal number of auctions to conduct is

$$T^* = ArgMax\{J_{1,T}(x_1)\} \text{ for } T \in \boldsymbol{T_f}.$$

The optimal lot size in each period is given by

$$k_i = \frac{x_i}{T-i+1} + \frac{ht(T-i)(n+1)}{8s} \quad , \tag{4}$$

and the auctioneer's resulting profit from period 1 to period T is given by

$$J_{1,T}(x_1) = \frac{T(T^2-1)}{6}\frac{(ht)^2(n+1)}{16s} + (\mu + s - \frac{2s}{n+1})x_1 - \frac{2sx_1^{\ 2}}{T(n+1)} - \frac{(T+1)htx_1}{2} - TC\,. \tag{5}$$

Proof: We first consider an approximation that differs from the optimization problem (P1) in two ways: (a) $p(k_i, n) = \mu + s - 2sm(k_i+1)$ for all k_i, even $k_i > n$, and (b) $x_i = x_{i-1} - k_{i-1}$ for all $k_{i-1} \in [0, x_{i-1}]$. Approximation (a) violates the definition in Equation (3), and (b) violates the state transition equation in

Equation (1). For the approximate problem we can use Remarks 1 and 2 and backward induction (see the Appendix for details) to show that, in any period i,

$$k_i^* = \min\left(\frac{x_i}{T-i+1} + \frac{h(T-i)(n+1)}{8s}, x_i\right). \tag{6}$$

If $k_i^* = x_i$ for some $i < T$, then $T \notin \boldsymbol{T_f}$, and it is suboptimal to offer items in all T auctions, even if the auctioneer has already paid the fixed cost of running T auctions, TC. The optimal number of auctions to run, then, is $T^* < T$. Restricting ourselves to $T \in \boldsymbol{T_f}$ therefore will not result in a suboptimal solution.

Given $T \in \boldsymbol{T_f}$, $k_i^* = \dfrac{x_i}{T-i+1} + \dfrac{ht(T-i)(n+1)}{8s}$. By Assumptions 2 and 3, $k_1^* < n$. It is easy to see from Equation (6) that k_i^* is decreasing in i, so $k_i^* < n$ for all i. If $k_i^* < n$ for all i, we no longer deviate from the original problem (P1), so Equation (4) defines the optimal lot size for (P1). We can apply backward induction (see the Appendix) to get the total profit function in closed form:

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

Before making some observations about the results of Proposition 1 we discuss the implications of Assumptions 1-3. Assumption 1 introduces some 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 avoid the consideration of cases in which the firm would be forced to scrap units. This would be a concern if demand were stochastic, but with deterministic demand the firm would scrap the excess $x_1 - nT$ units to avoid holding costs, suggesting that x_1 was set too high. Since we have assumed that x_1 is determined exogenously to Problem (P1), we ignore these situations by only considering T sufficiently large that $nT \geq x_1$. Assumption 3 guarantees that lot sizes in each period, determined with Equation (4), are strictly smaller than n. When this condition is violated, e.g., when holding costs are very high, it is in the interest of the seller to initially "scrap" large quantities of the product to reduce inventory. As inventory decreases, the condition will eventually be satisfied in all future auctions, and then the results of Proposition 1 will hold. When the condition in Assumption 3 is violated it is impossible to derive closed-form analytical results, but we find in numerical experiments[8] 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 and Assumptions 1-3 will hold in general.

Given an optimal number of periods, T^* (obtained by maximizing Equation (5)), we can analyze the form of the optimal lot size in each period. The first term of Equation (4) represents a lot-size policy that is constant across all periods. Such a policy is achieved if at the beginning of each period the auctioneer splits the remaining inventory into the number of periods that are left until the end (e.g., in period 1 it is $\frac{1}{T} x_1$, 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 $\frac{1}{4}\frac{ht}{sm}$ from period to period.

Proof:

$$k_{i-1} - k_i = x_{i-1}\left(\frac{1}{T-i+2} - \frac{1}{T-i+1}\right) + \frac{k_{i-1}}{T-i+1} + \frac{ht(n+1)}{8s}.$$

Substitute

$$k_{i-1} = \frac{x_{i-1}}{T-i+2} + \frac{ht(n+1)(T-i+1)}{8s}$$

into the equation above and simplify.

$$\Rightarrow k_{i-1} - k_i = \frac{2ht(n+1)}{8s} = \frac{ht}{4sm}.$$

□.

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 Equation (2), more bidders means more competition and consequently a higher closing price. With larger

---

[8] We have conducted numerical experiments considering cases in which Assumption 3 does not hold. We find that once the inventory has been reduced enough the lot size drops below n and then decreases until it reaches 0,

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. More units can be shifted to early periods in order to save on the total inventory holding costs.

Given an optimal number of auctions *T*, the optimal lot size in each auction is independent of μ but dependent on *s*. Proposition 1 also shows that higher dispersion in consumers' valuations 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 the auctioneer will encounter 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 balance the lot size when the dispersion is high.

To demonstrate the benefit of optimizing the auction 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 Figures 1(a)-(d), we show the percentage improvement in profit of the optimal lot-sizing policy over the constant-lot-size policy as we vary the parameters *h*, *n*, *s*, and *C* respectively. In Figures 2(a)-(d), we plot the optimal number of auctions for the same cases used in Figures 1(a)-(d). For example, in Figure 1(a) 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 lot-sizing policy is 10.7% greater than the profit from the constant-lot-size policy. In Figure 1(b) we see that if we reduce the number of bidders in the base case to 7, the optimal lot-sizing policy provides a 75% improvement over the constant-lot-size policy.

To determine the optimal lot-sizing policy, we solve the problem formulated in Equation (1) for different values of $x_1 \leq 30$, choosing the initial inventory that maximizes profits. If the optimal starting

---

following Equation (4).

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 $i = [1, \mathbf{int}(x_1 / K)]$, and if $x_1 \bmod K > 0$, then $k_{\mathbf{int}(x_1 / K)+1} = x_1 \bmod K$, where $K = [1, n]$, x_1=30, and **int()** is the integer part function. **{In the figures, “Lot-size” shouldn’t have a hyphen unless you actually include “Policy”}**

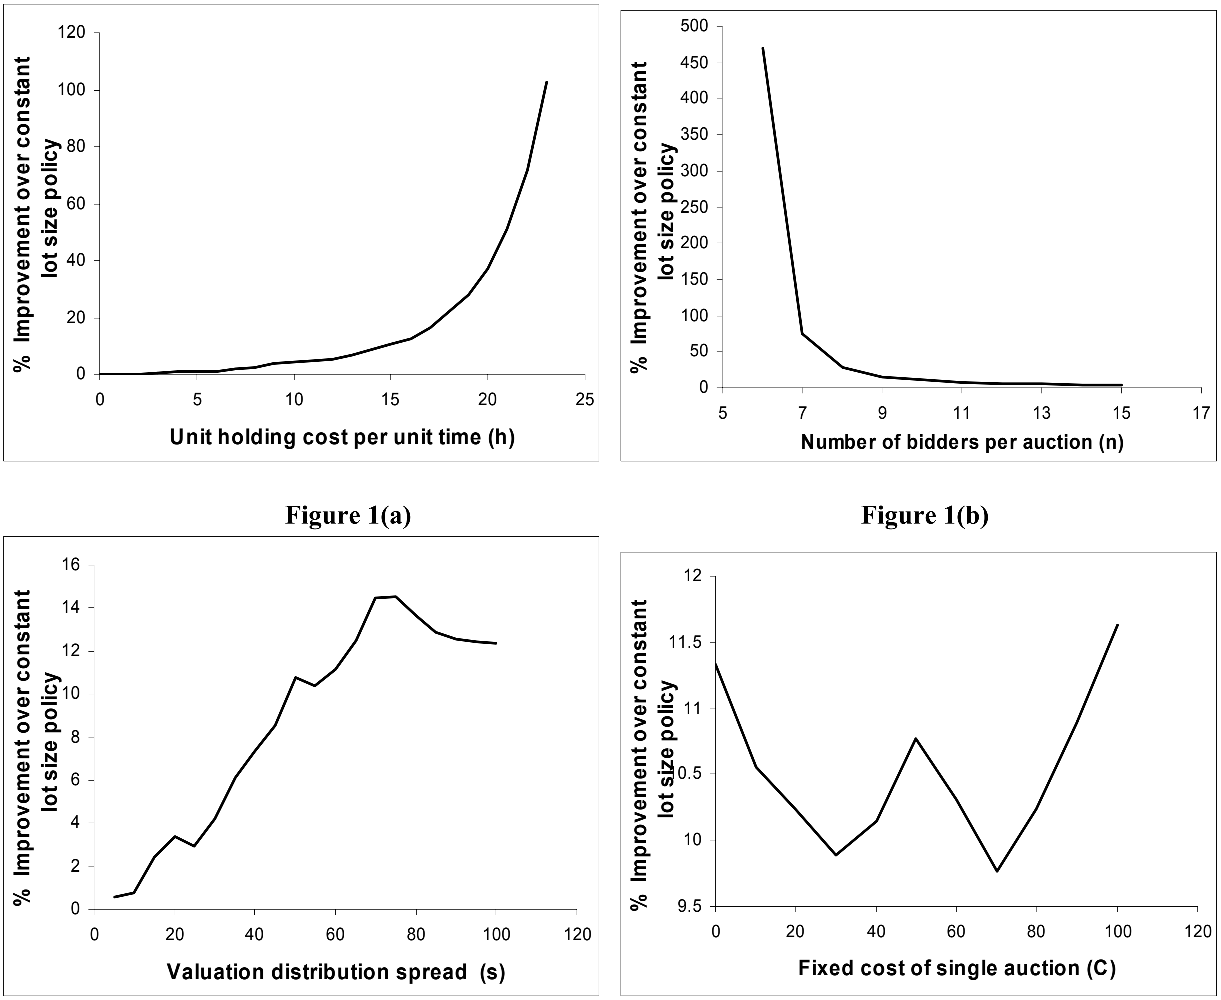

Figure 1(a) **Figure 1(b)**

Figure 1(c) **Figure 1(d)**

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.

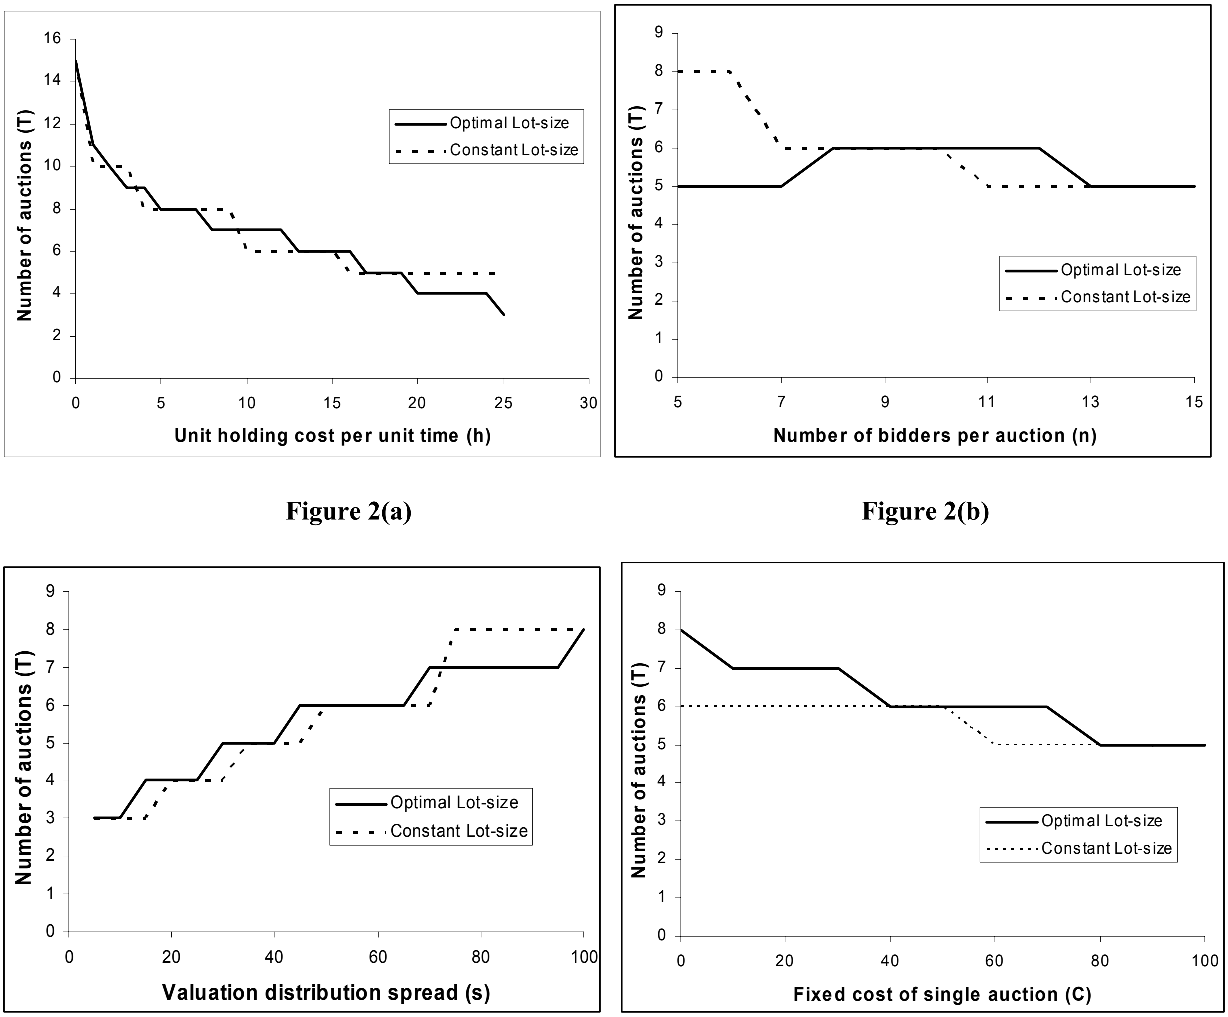

Figure 2(a)

Figure 2(b)

Figure 2(c)

Figure 2(d)

Optimal number of auctions for the constant-lot-size policy and 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.

Figures 1(a) and 2(a) show that when holding costs are low it is optimal to have many small auctions, and as a result there is little difference in the lot sizes used in both approaches. As holding costs increase, both policies require reducing the number of auctions, but the optimal lot-sizing 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. Since this is not an option in the constant-lot-size policy, the relative performance of the optimal lot-sizing policy increases dramatically.

Figures 1(b) and 2(b) show the effect of the number of bidders. When there are few bidders, using the optimal lot-sizing policy is more important, because it takes the market impact of the lot size into account. When the number of bidders is large, the two policies converge to a smaller number of auctions with large lot sizes. Interestingly, we see that the optimal number of auctions is not a monotone decreasing function of the number of bidders in the optimal lot-sizing policy. The reason is that when the number of bidders is small the optimal policy is to scrap some of the original inventory. As the number of bidders increases, it is optimal to scrap less of the original inventory and run more auctions to generate more revenue. When there are enough bidders that it is optimal to auction the entire original inventory, increasing the number of bidders further makes it possible to conduct fewer auctions and reduce total holding costs.

Figures 1(c) and 2(c) show 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)**{Italicize?}** - price auction mechanism, the risk of a low auction price is reduced, making larger lots and fewer auctions preferable. In this scenario there is little advantage from the optimal lot-sizing policy. When the valuation spread increases, it is optimal to conduct more auctions with smaller lots, and the benefits of the optimal policy increase. The non-monotonicity of Figure 1(c) 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 Equation (2), the expected price formula. The derivative of Equation (2) with respect to the spread is $2(n - k) / (n + 1) - 1$, and it is positive for $k < n / 2 - 1$ and negative for larger k.

Figures 1(d) and 2(d) show the effect of the fixed cost of running each auction, C. As expected, the optimal number of auctions decreases as the cost of running each increases. However, it does not appear that the relative performance of the two policies is very sensitive to C. The fluctuations observed in Figure 1(d) are due to integrality effects more than anything else.

To summarize, we have found that 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, there can be significant benefit to using a lot-sizing optimization such as ours over more naïve lot-sizing approaches.

3.3 Discussion of modeling assumptions

We 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. These bidders can be viewed as bargain hunters who repeatedly participate in auctions until they are in one that they can win at a low price. From the perspective of our model these are low-valuation bidders, i.e., those who fall in the left tail of the valuation distribution. It is important for us to know how far the left tail extends beyond the mean. In the next section, we develop a feedback mechanism for the seller to use to infer this distributional information from observed bids. The impact of the bargain-hunting strategic behavior described above should not be overestimated. A consumer who shades her bids in anticipation of future auctions by the same seller for the same good is taking a number of risks. In our model (and in practice), the seller does not announce the schedule for future auctions. The consumer therefore does not know whether there will be another auction, how much inventory is available to auction, or when another auction may take place. Given that auction participation has costs in terms of monitoring and delay in purchase, a consumer may be advised to bid her true valuation in the first auction. If a bidder shades her bids in general anticipation of other purchase opportunities, we can view this as incorporated into the overall valuation distribution.

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 continuously through time according to a random process (see Pinker et al. 2001). While 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.

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 over the constant-lot-size policy only increase.

Finally, we have assumed that bidders' valuations are drawn from a uniform distribution. We have conducted numerous simulations of the expected price (Equation (3)) using different sums of uniform distributions to test the robustness of Equation (3) 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 in conjunction with the tests with Poisson numbers of bidders and found little difference in the expected 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.

4. Multiple sequential auctions with information feedback

In the previous section, we showed how the optimal number of auctions to conduct and the optimal auction lot sizes depend on the distribution of bidders' valuations and other parameters. 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 investigate how the auctioneer can use observations of the bidders' reservation prices as feedback on his estimate of the dispersion of customers' valuations. We show that this new information can change the design of later auctions in the series.

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

1. Proposition 1 demonstrates 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).
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 for the mean of consumers' valuations, while the dispersion would be more difficult to estimate. Most items that are offered via Internet auctions can be typically 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 dispersion of consumers' valuations.

In what follows, we present a structured information-gathering method, built on Bayesian techniques, in which the auctioneer has a prior estimate of the dispersion of consumers' reservation values (denoted by s in the previous section). Running the auction and monitoring the bidding activity provide the auctioneer with a signal of the actual dispersion. This observation of the dispersion has its own distribution. The auctioneer then updates his belief and forms a posterior that will help him design the next period's auction better. We describe the Bayesian updating method used here and the key assumptions.

4.1. Learning the dispersion of consumers' valuations

Suppose the auctioneer is not sure about the value of s and therefore the spread of the consumers' reservation values. Assume that his *prior distribution* on s is $g_1(s)$, distributed uniformly between [*a, b*], where $a \geq 0$ (the dispersion spread can not be negative) and $b \leq \mu$ (because consumers' valuations [v_j's] must be non-negative). We assume that the auctioneer observes all bidders' reservation prices and uses this information to update his estimation of s or of a and b.

In practice, there is considerably more information available to the online auctioneer about reservation prices than theory predicts. In the traditional (offline) English auction setting, the auctioneer cannot observe all the bidders' reservation prices. Rather, the auctioneer observes the highest bid needed to win, which is in theory a small increment above the $k+1^{\text{th}}$ highest reservation price in an auction with k items. The true reservation price of the winner may be much higher than the winning bid. In the traditional (offline) auction, the auctioneer's observations are censored, and the estimation problem therefore becomes more complex. The computerization of auctions has made it easier to collect the reservation price data, as discussed in § 1 above. For example, by analyzing the bid history, the auctioneer can identify the bid at which each individual participant dropped out of the competition, and then, by adding one increment (one tick), the auctioneer gets a very good proxy for the reservation price of these individuals. In addition, the use of bidding agents in an online auction could also provide the auction house with reservation values directly. To explore the potential of learning from bids and to simplify the analysis, we assume the use of a sealed-bid Vickrey auction mechanism. It has been shown (Vickrey 1961) that under this mechanism, bidders do reveal their true reservation prices.

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:

$Y_i = \max\left\{\left|v_j - \mu\right| : j = 1, \ldots, n\right\} \quad for \quad i = 1, \ldots T^*.$

As mentioned above, in the first period the auctioneer's prior on s is

$$g_1(s) = \frac{1}{b-a} \qquad s \in [a, b].$$

During one auction, the auctioneer observes n different reservation prices. Conditioned on s, the cumulative distribution of Y_i for n random draws is

$$F\left(Y_1 \middle| n,s\right)=\begin{cases}0 & y_1 > s,\\ \left(\dfrac{y_1}{s}\right)^n & y_1 \le s.\end{cases} \tag{10}$$

The observed signal (the maximal deviation from the mean) from the first auction is denoted by y_1. We can now use Bayes' formula to calculate the auctioneer's posterior on s for the second period:

$$\begin{aligned} g_2(s) = g_1\left(s \middle| y_1, n\right) &= \frac{g_1(s) f\left(y_1 \middle| n,s\right)}{\int f\left(y_1 \middle| n,s\right) g_1(s)\, ds} \\ &= \frac{\dfrac{1}{b-a}\left(\dfrac{y_1}{s}\right)^{n-1}\dfrac{n}{s}}{\int_{y_1}^{b}\left(\dfrac{y_1}{s}\right)^{n-1}\dfrac{n}{s}\dfrac{1}{b-a}\, ds} \\ &= \frac{n-1}{s^n\left(\dfrac{1}{y_1^{n-1}}-\dfrac{1}{b^{n-1}}\right)} \qquad \mathit{for}\ s \in \left[y_1, b\right]. \end{aligned}$$

The limits of the integral in the likelihood function are y_1 to b. As the auctioneer observed a maximal deviation of y_1, he is positive that s is at least as large. 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,

$$l_i = \max\left\{l_{i-1}, y_i\right\} \qquad i = 1, \ldots., T^* \quad \text{where } l_1 = y_1 .$$

Using Bayesian updating, we can obtain the posterior distribution of s for each period in the auction sequence.

Proposition 2: After i periods, the auctioneer's posterior (which is period $i+1$'s prior) for the spread of consumers' valuations (s) is

$$g_{i+1}(s) = g_i\left(s \middle| l_i\right) = \frac{in-1}{s^{in}\left(\dfrac{1}{l_i^{in-1}}-\dfrac{1}{b^{in-1}}\right)} \qquad \mathit{for}\ s \in \left[l_i, b\right] . \tag{11}$$

Proof: Follows from the period 1 definition by induction.

Knowing the posterior distribution, the auctioneer can calculate his updated estimate for the dispersion of consumers' valuations. We can see from Proposition 2 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*. Period *i* tells us how many bids have been observed so far, and l_i gives us the lower bound on *s*.

We can incorporate the learning method described above into the optimization problem that was analyzed in § 3. The analysis above naturally implies an additional state variable, which is the *lower bound on s* (namely l_i). The auction design dynamic program therefore is now redefined in a two-dimensional state space, is now *stochastic*, and can be formulated as follows:

Problem (P2)

$$J_i^*(x_i, l_i) = \begin{cases} \max\limits_{k_i \in [0, x_i]} E_s \left[\mathrm{E}_{l_{i+1}} [J_{i+1}^*(x_i - k_i, l_{i+1}) \mid s, l_i] + k_i(\mu + s - 2sm(k_i + 1)) - x_i h - C \right] & \text{for } x_i > 0 \\ 0 \quad \text{when } x_i = 0 & \end{cases}, \quad (12)$$

$$i = 1, \ldots, x_i$$

where the distribution of *s* is determined by l_i and the number of auctions run so far. 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 $J_1^*(x_1, l_1)$ using the finite-horizon dynamic program Problem (P2), in which $J_i^*(0, l_i) = 0$. $J_i^*(0, l_i) = 0$ simply means that no more auctions are conducted once the inventory has been exhausted.

4.2. Numerical experiments

To illustrate the potential of using auction feedback in sequential auctions, we conduct a set of illustrative numerical experiments. We assume that the spread of customers' valuations (*s*) is unknown with a uniform distribution [*a, b*]. We then compare three approaches to the auction design:

- Full information
- Constant-lot-size policy

- Auction feedback

In the full-information scenario, we assume that the actual value of s ($a \leq s \leq b$) is known when the auction is designed. The full-information scenario therefore serves as an upper bound on the profit the auctioneer can receive. In the constant-lot-size policy, we assume that the auctioneer distributes the initial inventory equally across all the sequential auctions as in § 3, 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) for updating the distribution of s each period based on the bids observed so far. As in § 3, 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 base case scenario parameter values: $\mu = \$100$, $C = \$50$, $h = \$15$, $n = 10$, and $t = 1$, with s uniform on [$a = 1$, $b = 99$].

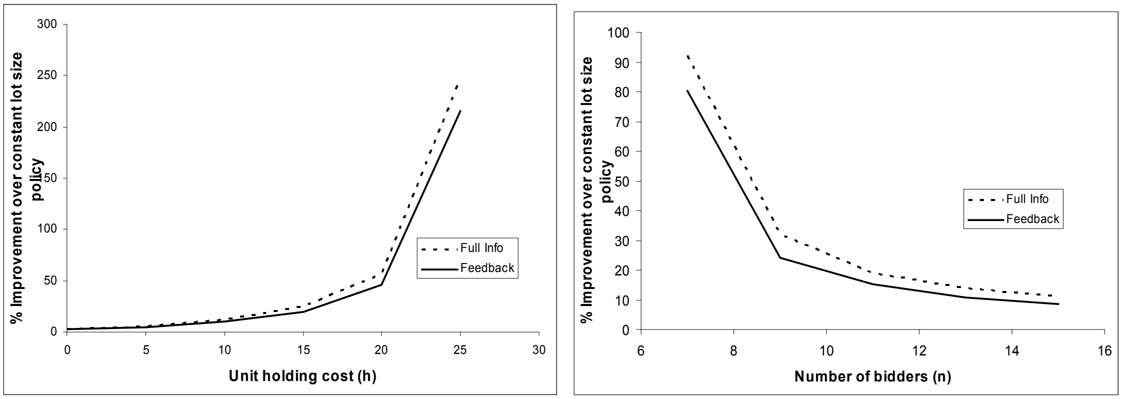

Figure 3(a) Figure 3(b)

Percentage improvement in profit over a constant-lot-size policy from using an auction feedback mechanism and when there is full information about the valuation dispersion parameter *s*, as a function of (a) holding costs and (b) the number of bidders per auction.

In Figure 3(a) we plot the percentage revenue improvement over the constant-lot-size policy with no feedback for the full-information and auction feedback approaches as a function of the unit holding costs. We see that the performance is similar to that in Figure 1(a). The benefit of the auction feedback mechanism, relative to the constant-lot-size policy, increases with the holding costs. This makes sense, as higher holding costs increase the sensitivity of the total cost function to the lot-size policy. As a result, the gap between the feedback mechanism and the full-information scenario increases with holding costs too.

Since the optimal lot size is dependent on *s*, any uncertainty in the value of *s* will reduce profits more when optimal lot sizing is more important. We see the same effects in Figure 3(b), where we look at how the performance improves as a function of the number of bidders. There is an additional phenomenon at work when the number of bidders varies. When the number of bidders is low, there are fewer observations of the bidders' valuation distribution, so it takes more auctions for the auction feedback mechanism to get a good estimate of the valuation spread, *s*. The auction feedback mechanism thus does not perform as well relative to the full-information scenario when the number of bidders is small.

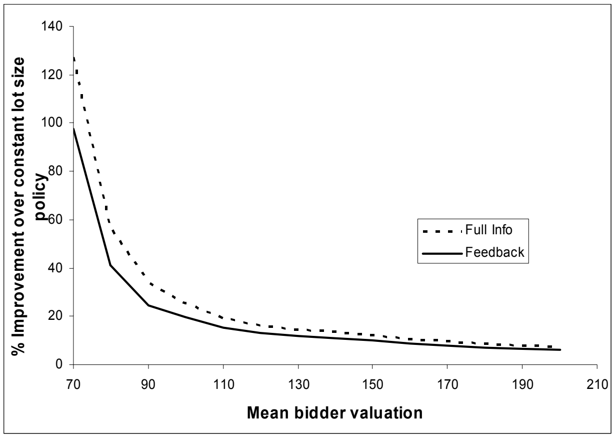

Figure 4: Percentage improvement in profit over the constant-lot-size policy from using an auction feedback mechanism and when there is full information about the valuation dispersion parameter *s*, as a function of the mean bidder valuation μ. Assume that *s* is Unif[1, μ].

In Figure 4 we see that the advantage of the auction feedback mechanism decreases as the mean bidder valuation increases. This is because when the mean valuation is low, prices will be low, and auction revenue therefore decreases. If prices decrease relative to the auction costs, the effect is similar to increasing costs, as we see in Figure 3(a).

In Figures 5(a) and 5(b), we fix the expected value of *s* and vary its coefficient of variation by adjusting the parameters *a* and *b*. In Figure 5(a), the distribution of *s* is centered at 25, and in Figure 5(b), it is centered at 50. In both cases we find that as the uncertainty in *s* increases the gap in performance between the auction feedback mechanism and the full-information cases increases. Three things are happening here. First, as the uncertainty in *s* increases it is harder for the feedback mechanism to learn *s*, increasing the distance from the full-information case. Second, greater uncertainty in *s* also increases the benefit of the auction feedback mechanism over the more naïve constant-lot-size policy that assumes *s* is at its mean value. Finally, we see that the auction feedback mechanism with optimal lot sizing does better when the spread in the bidders' valuations is lower because it can set larger lots in early auctions without as much price risk.

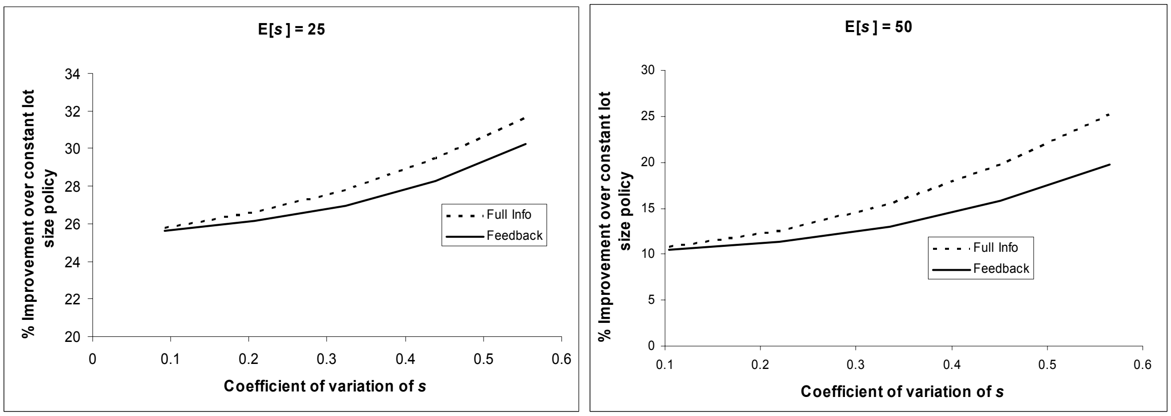

Figure 5(a) Figure 5(b)

Percentage improvement in profit over a constant-lot-size policy from using 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.

To summarize, we see in Figures 3-5 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 § 3, 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 2 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 Figures 1-5. For each period we list the starting inventory, the largest spread observed thus far, the optimal number of auctions to run given the current distribution of *s*, the optimal lot size, and the expected value of *s* given the observed bids.

Period (i)	Inventory at the start of period (x_i)	Inventory at the end of period (x_i-k_i)	Observed spread (l_i)	Optimal number of remaining auctions (T_i)	Optimal lot size (k_i)	Price $p(k_i)$	$E[s\|l_i]$
1	30	22	0	5	8	$68.18	50.50
2	22	12	19	4	10	$50.00	21.37
3	12	5	25	2	7	$77.30	26.39
4	5	0	26	1	5	$95.50	26.93

Table 2: An example of a sample path of the optimal lot size k_i and the optimal number of auctions T_i when using auction feedback.

We observe (in Column 5) that the optimal number of auctions dynamically changes. Initially (before the start of auction 1), five auctions are planned, but after two auctions are observed, the optimal total number of auctions is revised down to four (i.e., before the start of auction 3, it is optimal to conduct two more auctions). We also observe (in Column 6) that the optimal 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 50.50 to 21.37. A smaller spread reduces the risk of low prices and therefore makes a larger lot size more attractive. As a result, in Column 7 we see that the closing price decreases and then increases. In Column 8, we see how the auctioneer's belief about the expected value of *s* is converging 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 Equation (11) and shown to be less than 10%.

We can summarize the results of our experiments with the Bayesian learning methodology as follows:

- Using auction feedback provides increasing benefits when the number of bidders is smaller.

- The auction feedback scenario performs asymptotically like the full-information scenario as holding costs decrease.
- Auction feedback provides increasing benefits relative to a non-Bayesian policy with an increase in holding costs.
- Auction feedback can lead to lot-sizing polices that are very different from those developed when the valuation spread parameter is assumed to be fixed.

The performance results of the auction feedback mechanism we have demonstrated are, of course, dependent on our modeling assumptions and limited to the numerical experiments we have conducted. We have assumed that both the bidders' valuations and the prior distribution of the parameter *s* are uniform. In practice, these distributions may be quite disparate. Furthermore, a seller may need to use the bid data to estimate more than one parameter of the valuation distribution. We have also 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 Vickrey auctions and some bidders may not get the opportunity to bid up to their reservation values. An interesting 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.

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 design problem. Our results demonstrate 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 find 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. On the other hand, splitting the initial inventory across too many auctions results in excessive holding and administrative costs.

Our analysis and previous work on auctions show 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 auction design. Online auctions offer an effective mechanism for collecting data about the actual valuations of bidders that 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 design decisions for consecutive auctions of the same item. 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. Using the bidding information from previous auctions leads to dynamic adjustments in both the optimal number of remaining auctions and the optimal lot size. Our work serves as a form of proof of concept that incorporating bidding information into auction design has a significant potential to increase sellers' profits.

There are several interesting avenues for future research in this area. We have not included the minimum bid and reservation price as endogenous design parameters despite the fact that they may influence bidders' participation. It may also be interesting to extend our model to include production decisions as well as lot-sizing decisions. 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, we have no doubt that exciting research will continue to appear in the area of the optimal design of online auctions, on both the empirical and the theoretical frontier.

References

1. Azoury, K. 1985. Bayes Solution to Dynamic Inventory Models under Unknown Demand Distribution. *Management Science.* **31** pp. 1150-1160.
2. Bapna, R., Goes, P., and A. Gupta. 2000. A Theoretical and Empirical Investigation of Multi-item On-line Auctions. *Information Technology and Management Journal.* **1** pp. 1-23.
3. Bapna, R., Goes, P., and A. Gupta. 2001. On-Line Auctions: Insights and Analysis. *Communications of the ACM*. **44** (11), pp. 42-50.
4. Beam, C., and A., Segev, Arie. 1998. Auctions on the Internet: A Field Study. *Haas School of Business Working Paper, University of California Berkeley.* **WP-1032**.
5. Harris, M., and A. Raviv. 1981. Allocation mechanism and design of auctions. *Econometrica*. **49**(6).
6. Klein, S. 1997. Introduction to electronic auctions. *Electronic Markets.* 7(4) pp. 3-6.
7. Klemperer, P. 1999. Auction Theory: A Guide to the Literature. *Journal of Economic Surveys.* **13** (3) pp. 227-286.
8. Klemperer, P. (ed.). 2000. The Economic Theory of Auctions. Cheltenham, UK: Edward Elgar.
9. Lovejoy, W. 1990. Myopic Policies for Some Inventory Models with Uncertain Demand Distributions. *Management Science.* **36** pp. 724-737.
10. Lucking-Reilly, D. 2000. Auctions on the Internet: What's Being Auctioned, and How? *Journal of Industrial Economics*. **48** pp. 227-252.
11. Moon, Y., 1999. OnSale Inc. *Harvard Business Review Case Study-599091.*
12. Pinker, E., Seidmann, A. and Vakrat, Y. 2001. The Design of Online Auctions, Business Issues and Current Research. *University of Rochester, Simon School Working Paper.*
13. Rothkopf, M., Teisberg, T. and E. Kahn. 1990. Why are Vickrey Auctions Rare? *Journal of Political Economy.* **98** (1) pp. 94-109.
14. Scarf, H. 1959. Bayes Solutions of the Statistical Inventory Problem. *Annals of Mathematical Statistics* **30** pp. 490-508.
15. Segev, A., Shanthikumar, G. and C. Beam. 2001. Optimal Design of Internet Auctions. *Information Technology and Management Journal*, **2** pp. 121-163.
16. Turban, E. 1997. Auctions and bidding on the Internet: an assessment. *Electronic Markets* 7(4).
17. Vakrat, Y. 2000. Optimal Design of Online Auctions. *Unpublished Ph.D. Thesis, University of Rochester*.
18. Vakrat, Y. and A. Seidmann. 2000. Implications of the Bidder's Arrival Process on the Design of Online Auctions. *Proceedings of the Hawaii International Conference on System Sciences – 33.*
19. Vakrat Y. and A. Seidmann. 1999a. Can Online Auctions Beat Online Catalogs? *Proceedings of the International Conference on Information Systems 1999.*
20. Vakrat, Y. and A. Seidmann. 1999b. Analysis and Design Models of Online Auctions. *Proceedings of the 4th INFORMS Conference on Information Systems and Technology, Cincinnati*. pp. 64-78.
21. Vickrey, W. 1961. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance.* **16** pp. 8-37.
22. Vulcano, G., Van Ryzin, G. and C. Maglaras. 2002. Optimal Dynamic Auctions for Revenue Management. *Management Science.* **48** (11) pp. 1388-1407.
23. Wang, R. 1993. Auctions versus Posted-Price Selling. *American Economic Review*. **83** (4) pp. 838-851.
24. Zipkin, P. 2000. Foundations of Inventory Management. McGraw-Hill.

Appendix

Proposition 1 {See comments below for "Discounting"}

Given Assumptions 1-3, define $\boldsymbol{T_f}$ to be a set of T such that

if $k_i = \frac{x_i}{T-i+1} + \frac{ht(T-i)(n+1)}{8s}$, then $k_i \leq x_i$ for all $i \leq T$. The optimal number of auctions to conduct is $T^* = ArgMax\{J_{1,T}(x_1)\}$ for $T \in \boldsymbol{T_f}$, and the optimal lot size in each period is given by

$$k_i = \frac{x_i}{T-i+1} + \frac{ht(T-i)(n+1)}{8s}. \tag{1}$$

The auctioneer's resulting profit from period 1 to period T is given by

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

Details of the proof of Proposition 1:

To derive a closed-form solution to the dynamic program, we will use backward induction. We know, by Lemma 2, that the last period (period T) profits are

$$J^*_{T,T}(x_T) = x_T(\mu + s - 2sm(x_T + 1)) - x_T ht - C.$$

Let us consider period T-1. The problem that the auctioneer faces is

$$\begin{aligned} J^*_{T-1,T}(x_{T-1}) = Max \quad & (x_{T-1} - k_{T-1})(\mu + s - 2sm(x_{T-1} - k_{T-1} + 1)) + \\ & k_{T-1}(\mu + s - 2sm(k_{T-1} + 1)) - (x_{T-1} - k_{T-1})ht - x_{T-1}ht - 2C \end{aligned}.$$

Note that the problem is quadratic in k_{T-1}. It is also easy to show that the first coefficient of the polynomial is negative. As a result, the function is concave in k_{T-1}. Using the sufficient first-order conditions, we obtain the optimal number of units to offer in period T-1:

$$k_{T-1} = \frac{1}{2} x_{T-1} + \frac{1}{8} \frac{ht}{sm}.$$

The profit function thus can be rewritten as

$$J^*_{T-1,T}\left(x_{T-1}\right) = \frac{1}{16} \frac{(ht)^2}{sm} + (\mu + s - 2sm)\left(x_{T-1}\right) - smx^2_{T-1} - \frac{3}{2} htx_{T-1} - 2C.$$

Let us investigate the problem the auctioneer faces in period *T*-2:

$$J^*_{T-2,T}\left(x_{T-2}\right) = \underset{k_{T-w}}{Max} \frac{1}{16} \frac{(ht)^2}{sm} + (\mu + s - 2sm)\left(x_{T-2} - k_{T-2}\right) - sm\left(x_{T-2} - k_{T-2}\right)^2$$
$$- \frac{3}{2} ht\left(x_{T-2} - k_{T-2}\right) - 2C + k_{T-2}\left(\mu + s - 2sm\left(k_{T-2} + 1\right)\right) - x_{T-2}ht - C$$

It is again easy to see that the multi-period profit function is concave in k_{T-1}. Consequently, the optimal number of units to offer this period is

$$k^*_{T-2} = \frac{1}{3} x_{T-2} + \frac{2}{8} \frac{ht}{sm},$$

and the resulting profit function is

$$J^*_{T-2,T}\left(x_{T-2}\right) = \frac{4}{16} \frac{(ht)^2}{sm} + (\mu + s - 2sm)x_{T-2} - \frac{2}{3} smx_{T-2}{}^2 - \frac{4}{2} htx_{T-2} - 3C.$$

Note that both k^*_{T-1} and k^*_{T-2} are of the form specified in Equation (1). Also note that it is possible for either k^*_{T-1} or k^*_{T-2} (as defined above) to exceed the initial inventory in the respective periods. In such a case, because of the concavity of the profit function the optimal lot size in a period would be the initial inventory in that period. Yet if this happens it means that *T* is not in set T_f, and the case therefore is not relevant to our analysis.

We assume that in period *i* the profit function is given by

$$J_{i,T}^{*}(x_i)=\frac{1}{16}\frac{(ht)^2}{sm}\frac{(T-i)(T-i+1)(T-i+2)}{6}+(\mu+s-2sm)x_i \\ -\frac{2}{T-i+1}smx_i^{\ 2}-\frac{(T-i+2)}{2}htx_i-(T-i+1)C \tag{3}$$

The profit functions we have derived for periods *T*-1 and *T*-2 are of this form as well.

In period *i*-1 the profit, then, is

$$J_{i-1,T}^{*}(x_{i-1})=\underset{k_{i-1}}{Max}\ k_{i-1}(\mu+s-2sm(k_{i-1}+1))-x_{i-1}ht-C+J_{i,T}(x_{i-1}-k_{i-1})\,. \tag{4}$$

The first-order conditions give

$$k_{i-1}^{*}=\frac{1}{T-(i-1)+1}x_{i-1}+\frac{(T-(i-1))}{8}\frac{ht}{sm},$$

which is of the form specified by Equation (1). Substituting k_{i-1}^{*} into Equation (4) gives

$$J_{i-1,T}^{*}(x_{i-1})=\frac{1}{16}\frac{(ht)^2}{sm}\frac{(T-(i-1))(T-(i-1)+1)(T-(i-1)+2)}{6}+(\mu+s-2sm)x_{i-1}-\frac{2}{T-(i-1)+1}smx_{i-1}^2 \\ -\frac{(T-(i-1)+2)}{2}htx_{i-1}-(T-(i-1)+1)C \,.$$

This is of the form in Equation (3). When *i=1*, we get the following equation:

$$J_{1,T}^{*}(x_1)=\frac{1}{6}(T-1)T(T+1)\frac{1}{16}\frac{(ht)^2}{sm}+(\mu+s-2sm)x_1-\frac{2}{T}smx_1^2-\frac{1}{2}(T+1)htx_1-TC\,.$$

Analysis with discounting

With discounting, the firm's problem becomes

$$J^*_{i,T}(x_i) = \underset{k_i \in [0,x_i]}{Max}\, \beta\, J^*_{i+1,T}(x_i - k_i) + k_i p(k_i) - htx_i - C$$

$$for\ \ i = 0,1,.....,T$$

$$where:$$

$$x_i = x_{i-1} - k_{i-1},$$

$$k_i \le n,$$

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

where $\beta \in [0,1]$ is the discount factor.

To simplify notation we define the following:

$$G(i) = \sum_{l=0}^{T-i} \beta^{\,l}$$

$$H(i) = \sum_{l=0}^{T-i} l\beta^{\,l}$$

$$B(i) = \sum_{l=0}^{T-i} (l+1)\beta^{\,l} \qquad .$$

$$Q(i) = \sum_{l=1}^{T-i} \frac{(T-l+2)!}{6(T-l-1)!}\beta^{\,l-1} + \sum_{l=1}^{T-i-1} \frac{(T-i-l+2)!}{6(T-i-l-1)!}\beta^{\,T-i+l-1}$$

$$M = \mu + s - 2sm$$

Using induction as in Proposition 1, it can be shown that the optimal lot size in period *i* when $\beta < 1$ is given by

$$k_i^* = \frac{\beta^{\,T-i}}{G(i)} x_i + \frac{ht}{4sm}\frac{H(i)}{G(i)} + \frac{M}{4sm}\frac{(G(i) - (T-i+1)\beta^{\,T-i})}{G(i)}. \tag{5}$$

For $\beta = 1$, Equation (5) collapses into Equation (1).

The profit function in period *i* is given by

$$\begin{aligned} J^*_{i,T}(x_i) = {} & \frac{(T-i+1)M\beta^{\,T-i}}{G(i)} x_i - \frac{htB(i)}{G(i)} x_i - \frac{2sm\beta^{\,T-i}}{G(i)} x_i^2 + \frac{(ht\beta)^2 Q(i)}{8smG(i)} \\ & + \frac{M^2(1-\beta)^2 Q(i)}{8smG(i)} + \frac{2htM\beta(1-\beta)Q(i)}{G(i)} - CG(i) \end{aligned} \quad . \tag{6}$$

For $\beta = 1$, Equation (6) collapses into Equation (3). With discounting, the profit functions are too complex to analyze algebraically. However, numerical experiments show that with greater discounting, i.e., β decreasing, the earlier lots increase in size but lot sizes are still monotonically decreasing.