OPERATIONS RESEARCH
Vol. 52, No. 3, May–June 2004, pp. 346–367
ISSN 0030-364X | EISSN 1526-5463 | 04 | 5203 | 0346

DOI 10.1287/opre.1040.0105

Optimal Auctioning and Ordering in an Infinite Horizon Inventory-Pricing System

Garrett van Ryzin
Graduate School of Business, Columbia University, 412 Uris Hall, New York, New York 10027, gjv1@columbia.edu

Gustavo Vulcano
Stern School of Business, New York University, New York, New York 10012, gvulcano@stern.nyu.edu

We consider a joint inventory-pricing problem in which buyers act strategically and bid for units of a firm's product over an infinite horizon. The number of bidders in each period as well as the individual bidders' valuations are random but stationary over time. There is a holding cost for inventory and a unit cost for ordering more stock from an outside supplier. Backordering is not allowed. The firm must decide how to conduct its auctions and how to replenish its stock over time to maximize its profits. We show that the optimal auction and replenishment policy for this problem is quite simple, consisting of running a standard first-price or second-price auction with a fixed reserve price in each period and following an order-up-to (basestock) policy for replenishing inventory at the end of each period. Moreover, the optimal basestock level can be easily computed. We then compare this optimal basestock, reserve-price-auction policy to a traditional basestock, list-price policy. We prove that in the limiting case of one buyer per period and in the limiting case of a large number of buyers per period and linear holding cost, list pricing is optimal. List pricing also becomes optimal as the holding cost tends to zero. Numerical comparisons confirm these theoretical results and show that auctions provide significant benefits when: (1) the number of buyers is moderate, (2) holding costs are high, or (3) there is high variability in the number of buyers per period.

Subject classifications: optimal auction; pricing; dynamic programming; infinite horizon; inventory control; basestock policy.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received March 2002; revision received November 2002; accepted March 2003.

Introduction

With the capabilities of Internet commerce, auctions have gained renewed popularity in both consumer and industrial markets (van Ryzin 2000). The new potential to use online auctions as an alternative to traditional list-price mechanisms raises some important theoretical and practical questions. In particular, which pricing mechanisms are optimal for sellers in any given context? How should these mechanisms be designed and implemented? How much benefit can alternative mechanisms provide over list pricing? And under what conditions are they most beneficial?

In this paper, we provide answers to these questions for a firm that orders, stores, and sells a homogeneous good over an infinite horizon. Our model is a stylized representation of a retailer, distributor, or producer who uses an auction mechanism for selling a replenishable product. The firm purchases its good at a constant unit cost from an outside supplier and incurs an increasing, convex holding cost on its inventory. There is zero leadtime for replenishment. Demand in each period is characterized as a random number of buyers, each of whom has his own, private value for the firm's good. The statistics of demand are assumed stationary over time and are known to the seller and all buyers. This demand model follows the assumptions of classical auction theory as described in the seminal work of Vickrey (1961), the influential paper of Milgrom and Weber (1982), the recent survey by Klemperer (1999), and earlier survey articles: McAfee and McMillan (1987a), Milgrom (1989), Rothkopf and Harstad (1994), Matthews (1995), and Wolfstetter (1996). As in this auction literature, we assume buyers act strategically to maximize their utility (i.e., their value minus the price they pay). As a result, the buyers' behavior depends on the auction and inventory policy of the firm. The firm must decide on an auction mechanism—that is, a set of rules for allocating goods to buyers and collecting payments from them—and a strategy for replenishing its stock that maximize its profits over an infinite horizon. We consider both the discounted and average profit criteria.

We analyze this problem using results from Maskin and Riley (2000), who show that the expected revenue for a seller in an auction depends only on the allocation—that is, which buyers receive the goods and which do not. By formulating a dynamic program in these allocation variables, we are able to characterize the optimal allocation and replenishment strategy for the firm. We then show that this optimal allocation can be achieved by conducting a first-price or second-price auction with a fixed reserve price in every period. The reserve price is related only to the replenishment cost of the good. The optimal replenishment policy is to order up to a fixed basestock level at the end of each period. Thus, the optimal policy is quite simple and familiar; namely, use a traditional auction with a reserve price as the selling mechanism and use a traditional basestock policy for replenishing inventory. We call this policy a *basestock, reserve-price-auction* policy. Moreover, the policy is easy to compute, and in the average-cost case reduces to a single parameter search over a closed-form profit function. We also extend these results to variations of the model, including the case where the firm sells in two markets—one fixed-price market and one auction market.

We then compare the basestock, reserve-price-auction policy to a list-price policy, which uses a fixed posted price in each period together with a basestock policy for replenishment. (The price and basestock level are jointly optimized.) This is the policy derived by Federgruen and Heching (1999) and shown to be optimal among all dynamic pricing and ordering policies under a model that is quite similar to ours. We show that this basestock, list-price policy is optimal for our problem as well in several limiting cases, including the case where there is only one buyer per period, the case where the number of buyers per period tends to infinity and the holding cost is linear, and the case where the holding cost is zero. A numerical study shows how the optimal basestock, reserve-price-auction policy compares to list pricing more generally. The results indicate that the auction policy is significantly better than list pricing under relatively specialized conditions, namely when the number of buyers per period is moderate (e.g., 5 to 10), the holding cost is large (e.g., holding cost rates of 1% of the value of the goods per period or higher), and when the variation in the number of buyers in a period is high. One can argue that many consumer and industrial markets do not match these conditions; consumer markets typically have high-volume demand and holding costs per period are less than 1%, though for specialty, low-volume products or big-ticket, high-tech products like personal computers, these conditions may hold. In some industrial markets—the sale of capital equipment for example—one encounters a modest volume of buyers and high holding costs, in which case our results suggest that the optimal auction policy can offer significant improvements in profit. Still, our model shows that list pricing is near optimal in many cases, which perhaps provides one explanation for its continued popularity, despite the promise of Internet-based auctions.[1]

Literature Review

While there is wide variety of work on auctions (see the survey articles mentioned above), analyzing joint auction and inventory decisions is a relatively new topic. In a finite horizon setting without replenishment, Segev et al. (2001) analyze a problem in which an auctioneer tries to sell multiple units of a product using a multiperiod auction; however, they assume customer bidding behavior is modeled exogenously by a Markov chain. Pinker et al. (2001) study how to run a sequence of standard k-unit auctions, using bidding information to learn about the customer valuation distribution, and determining the lot size k for each auction, the number of auctions to run, and the duration of each of them. Our earlier work, Vulcano et al. (2002), analyzes an optimal auction for a firm selling a fixed inventory over a finite horizon, and the approach we use here for the infinite horizon problem with replenishment closely follows it.

Our work is most closely related to research on stochastic inventory-pricing problems; see, for example, Federgruen and Heching (1999), Li (1998), Amihud and Mendelson (1983), Thomas (1974), Thowsen (1975), and Zabel (1972). Indeed, our problem is in many ways an auction version of the one studied by Federgruen and Heching (1999). They considered a problem in which a firm may choose a state-dependent list price and make replenishment decisions in each period. They showed that in the infinite horizon, stationary case, the optimal policy is a so-called *basestock, list-price* policy defined by two critical values p^* and z^*; if the inventory is above z^*, the firm orders nothing and selects an inventory-dependent price below p^*, which is decreasing in the inventory on hand; if the inventory is below z^*, the optimal policy is to order up to z^* and price at p^*. Thus, once the inventory level drops below z^*, the optimal policy is to use a fixed price and a fixed basestock level in each period. These results are essentially the fixed-price analogs of our results in Theorems 1 and 2 for our auction case. We also compare our results to a list-price policy of this type.

Overview

The remainder of this paper is organized as follows. In §1, we review the results we use on optimal auction design, formulate our inventory-pricing problem as a dynamic program, and present the main structural results on the basestock, reserve-price-auction policy. Section 2 provides the proofs of these main theorems. In §3, we analyze various extensions of the model, including the case where the firm has demand from both a fixed-price and an auction market and the case of the long-run-average profit criterion. In §4, we compare the basestock, reserve-price-auction policy to a basestock, list-price policy. Both theoretical and numerical comparisons are given. Finally, conclusions are given in §5. Several proofs are presented in the appendix.

Notation

We use the following notation: All vectors are assumed to be in $\mathbf{R}_+^n$ unless otherwise specified. v_j denotes the jth component of vector v, and $v_{-j} \equiv (v_1, \ldots, v_{j-1}, v_{j+1}, \ldots, v_n)$ is the vector of components other than j. Subscripts between parentheses stand for reverse order statistics; that is, for any vector v, $v_{(1)} \geqslant v_{(2)} \geqslant \cdots \geqslant v_{(n)}$. $\mathbf{Z}_+$ denotes the nonnegative integers. The positive part of a number a is $a^+ \equiv \max\{a, 0\}$. Analogously, $a^- \equiv -\min\{a, 0\}$.

The shorthand a.s. stands for almost surely; i.i.d. is shorthand for independent and identically distributed; and p.m.f. for probability mass function. A function is said to be increasing (decreasing) when it is nondecreasing (nonincreasing).

For a discrete-valued function $G(x)$, we define the difference $\Delta G(x) \equiv G(x) - G(x-1)$, and say that $G(\cdot)$ is concave (convex) when $\Delta G(x)$ is decreasing (increasing) in x.

1. Optimal Auctions, Model Formulation, and Statement of Main Results

In this section, we first review some results from the theory of optimal auctions that are required for our analysis. Readers familiar with auction theory may skip or only skim this section. We then formulate an inventory-pricing problem using this auction theory and state our main theoretical results on the optimal, dynamic auction-ordering policy. The proofs of these results are provided in §2.

1.1. Review of Results from the Theory of Optimal Auctions

The basic results on optimal auctions that we require are from Myerson (1981), Riley and Samuelson (1981), and Maskin and Riley (2000). The first two papers give the mathematical formulation of optimal auction design for a single good, and the third one extends these results to the multiunit setting.

Consider an auction in which we are selling one or more homogeneous objects to n buyers. Each buyer i wants at most one of the objects, which he values at v_i; and pretends to maximize his own expected surplus, defined as his valuation minus the amount paid to the auctioneer. The values v_i are private information, but it is common knowledge that they are i.i.d. with distribution F on a support $[0, \bar{v}]$. Buyers are assumed to be risk neutral, an assumption we discuss in more detail below.

An *auction mechanism* is a description of the auction, which specifies both allocation and payment rules. It is chosen by the seller, and is common knowledge. For example, in a k-unit first-price auction mechanism, buyers submit bids; the k highest bids win (the allocation rule), and all winners pay the amount offered (the payment rule). In a second-price auction, buyers submit bids; the k highest bids win (the allocation rule), but all winners pay the first losing bid, i.e., the $(k+1)$th bid (the payment rule). The buyers' behavior depends on the auction mechanism. Each buyer i seeks to maximize their expected surplus, which is the probability of winning times the difference between their value v_i and the amount they pay under the seller's mechanism. Assuming that buyers choose their strategies without collusion, they play a noncooperative game of incomplete information. The solution concept used in this context is that of a Bayesian equilibrium of Harsanyi (1967, 1968), an extension of the ordinary Nash equilibrium (1951).

Extending Myerson's (1981) results from single-unit auctions, Maskin and Riley (2000) showed the rather remarkable fact that the expected seller's revenue can be expressed *only* as a function of the allocation rule. Specifically, the allocation functions can be expressed as

$$q_i(v_i, v_{-i}) = \begin{cases} 1 & \text{if bidder } i \text{ is awarded a unit,} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

If the functions $q_i(\cdot, v_{-i})$ are increasing in v_i and buyers with value zero have zero expected surplus in equilibrium, then the expected revenue to the seller is given by

$$E_{v_i, v_{-i}}\left[\sum_{i=1}^{n} J(v_i) q_i(v_i, v_{-i})\right], \tag{2}$$

where

$$J(v) = v - \frac{1}{\rho(v)}, \tag{3}$$

and $\rho(v) = f(v)/[1-F(v)]$ is the hazard rate function associated with the distribution F. The function $J(v)$ is what Myerson (1981) calls the bidder's *virtual value*. From (2), it follows that all mechanisms that result in the same allocation q for each realization of v yield the same expected revenue. This is the so-called *Revenue Equivalence Theorem*.

For example, in a standard k-unit auction, one can show that both the first-price and second-price auctions award the k goods to the buyers with the k highest valuations. Thus, the allocation $q(v)$ is the same for each v and hence these two auctions generate the same expected revenue for the seller. This is true despite the fact that the bidding strategies and payments in each case are quite different.

Moreover, expression (2) can be used to design an optimal mechanism. This is achieved by simply choosing the allocation rule $q^*(v)$ that maximizes $\sum_{i=1}^{n} J(v_i) q_i(v_i, v_{-i})$ subject to any constraints one might have on the allocation (e.g., we have at most k items to allocate so we may require that the allocation q satisfies $\sum q_i \leqslant k$).

The following monotonicity assumption helps to simplify the analysis:

Assumption 1. *$J(v)$ is strictly increasing in v.*

This assumption simply ensures that higher-value bidders contribute higher expected revenues in (2). It holds when the hazard rate $\rho(v)$ either increases or does not decline too fast with v (formally speaking, we require $\rho'(v) > -\rho(v)^2$ for all $v \in [0, \bar{v}]$), and is satisfied by most standard distributions.[2] To illustrate, define

$$v_* = \max\{v\colon J(v) = 0\} \tag{4}$$

(and by convention, $v_* = \infty$ if $J(v) < 0\ \forall v$). Then, from (2) it follows that it is never optimal to allocate a unit to a buyer with valuation $v_i < v_*$. This simple observation is the basis for determining optimal reserve prices. Indeed, consider a standard k-unit auction with a second-price mechanism and reserve price v_*. One can show in this case that bidders bid their true values v_i and the items are awarded to the k highest bidders with valuations above v_*, which in fact produces the allocation $q^*(v)$ that maximizes (2) subject to the constraint that at most k items can be awarded.

Thus, the analysis of optimal auctions proceeds in two steps: (1) First, find an optimal allocation $q^*(v)$ that maximizes the revenue (or revenue net of costs) subject to any constraints one might have on the allocation; and then, (2) find an auction mechanism that achieves the allocation $q^*(v)$ for each realization v. Each of these steps requires a separate analysis. In the next section, we apply this two-step approach to analyze an inventory-pricing problem.

Finally, we note that the Revenue Equivalence Theorem and the optimal auction results can be extended to the case where the number of bidders n is random. In this case, all buyers and the seller know the distribution of n but buyers do not know n exactly when they formulate their bidding strategies. McAfee and McMillan (1987b, §4) analyzed this extension and showed that for risk-neutral buyers with symmetric priors on n, it is still optimal for the seller to allocate according to (2) and, moreover, that the first and second price auctions with reserve price v_* remain optimal. (The situation is different if buyers are risk averse or if they have different priors on n. See McAfee and McMillan 1987b.) Harstad et al. (1990) derive explicit equilibrium bidding functions for this random-number-of-bidders case.

1.2. An Inventory-Pricing Model

We are now ready to define and discuss our model. We first lay out the basic assumptions and problem statement. Once the basic model is defined, we then discuss the assumptions and implications in more depth.

1.2.1. Model and Assumptions. A firm stores, sells, and reorders units of an homogeneous good over an infinite time horizon, split in a discrete number of periods indexed by $t \geqslant 1$. The time index runs forward, so larger values of t represent later points in time. The firm starts a particular period with an initial (integral) inventory, denoted x, and sells these units through an auction. The problem is assumed to be stationary, so the statistics of demand are the same for all periods t.

In each period, N risk-neutral buyers participate in the auction. N is a nonnegative, discrete-valued random variable, distributed according to a probability mass function $g(\cdot)$ with support $[0, M]$ for some $M > 0$, and strictly positive first moment. We assume that the numbers of buyers N in each period are independent from one period to the next.

Each buyer requires one unit and has a reservation value v_i, $1 \leqslant i \leqslant N$, which represents the maximum amount buyer i is willing to pay for one unit of the good. Reservation values are private information, i.i.d. draws from a distribution $F(\cdot)$, which is strictly increasing with a continuous density function $f(\cdot)$ on the support $[\underline{v}, \bar{v}]$, with $F(\underline{v}) = 0$ and $F(\bar{v}) = 1$. Without loss of generality, we assume $\underline{v} = 0$ throughout. We assume that the virtual value $J(\cdot)$ derived from $F(\cdot)$ satisfies Assumption 1. We will use v both for the random vector of valuations (from the seller's perspective), and for a particular realization. Like the number of buyers, the valuations v are assumed to be independent from one period to the next. Thus, each period is an independent draw of N and v.

At the end of each period, the firm can reorder from a supplier at a unit cost c. Replenishment orders arrive instantly and backlogging is not allowed. If the firm decides to award k units through the auction in the current period and reorder y units from the supplier, then the initial stock of the next period will be $x - k + y$. The firm incurs a holding cost $h(x-k+y)$ on this inventory, which is paid at the end of the current period. The function $h(\cdot)$ is assumed convex and strictly increasing.

In terms of information structure, the distribution functions g and F are constant through time t, and are assumed common knowledge to the firm and all potential buyers. In terms of the buyer valuations, only buyer i knows his own (private) valuation v_i. Also, buyers cannot observe the number of other buyers prior to bidding, so they are uncertain about the number of competitors that they face. The selling firm also does not know the exact number of buyers when announcing the mechanism, but they observe the number of buyers that submit bids, which is not necessarily the total number of buyers N (e.g., buyers with values below a reserve price may simply choose not to bid and may therefore not be observed). Buyers do not have explicit information about the inventory position of the firm or its costs. However, they do have full information about the mechanism the firm selects, which in terms of their strategic behavior is all that matters to them. But effectively, as shown below, announcing the inventory position becomes a part of the optimal mechanism, because the firm has an incentive to sell all its stock if it receives a sufficient number of high bids, and the mechanism must (at least implicitly) reveal this fact.

The firm's problem is to design an auction mechanism and find a replenishment policy that maximizes its expected total discounted profit. As above, the auction mechanism is a set of rules for allocations and payments according to which the auction will be conducted. Each buyer, based on his private valuation, his knowledge of the distribution functions F, g, the inventory level x, and set of rules announced by the auctioneer—e.g., type of auction, reserve price—chooses his bid (or strategy) to maximize his expected utility. Then, the firm observes the set of submitted bids and applies the rules specified earlier to decide the number of units to award in the current period, whom to award the units to, and the payments to be made by the bidders (typically only the winners pay).

1.2.2. Discussion of the Model. On a theoretical level, our model is in many ways a natural combination of the classical, private-value auction model and dynamic inventory models. However, some of the assumptions are restrictive from a practical standpoint and their implications are worth examining in greater detail.

The first concerns the information structure. The assumption that buyers have the same priors on F and g is not unreasonable; it simply says that they are equally informed about the market. However, the fact that buyers have the same priors as the seller is less realistic. In particular, one might well imagine that the seller—who is conducting many auctions over an infinite horizon—would tend to learn over time, and as a result have much better information about the number of likely buyers and their reservation values than would an individual buyer, who may only occasionally participate in the auctions.

However, this assumption can be relaxed for the second-price auction mechanism discussed in §2.2.1, because under this mechanism a buyer's dominant strategy is to bid his value v_i. Thus, buyers do not need any information on the number of other buyers or their valuations to bid optimally in the second-price auction. For the first-price mechanism, in contrast, the more restrictive assumption that buyers know g and F is essential.

A second assumption is that the seller and buyers are risk neutral. That a large selling firm is risk neutral is quite reasonable, as typically each auction outcome is a small proportion of their wealth and they are making a very large number of gambles over an infinite horizon. So the fact that the firm is maximizing average profits is a quite natural assumption. In contrast, it is more reasonable to assume that individual buyers (e.g., consumers) are perhaps risk averse, because they may only participate in one auction and the values at risk may be a larger proportion of their wealth. Unfortunately, however, risk neutrality is a central assumption in the optimal auction theory of Myerson (1981), Riley and Samuelson (1981), and Maskin and Riley (2000). In this sense, our results share the limitations of this work.

If buyers are risk averse, then their preferences for the different types of auctions change, which affects both their bidding behavior and the seller's revenues. For example, in the traditional single-unit auction, risk-averse buyers prefer a first-price auction to a second-price auction because the amount they pay if they win in a first-price auction (their bid) is constant, whereas the amount they pay if they win in a second-price auction is uncertain (i.e., equals the second-highest bid). Hence, risk-averse buyers are willing to bid more in the first-price auction, which means the seller generates more revenue using a first-price auction. (See Klemperer 1999.) It is quite likely that a similar effect would occur in our context if buyers were risk averse, though we have not investigated this issue in detail.

However, risk neutrality is likely a better assumption if the buyers are other firms—perhaps procuring inputs from a supply auction. If one applies the model to industrial trade, the risk-neutrality assumption is therefore more realistic.

Another important assumption in our model is that the selling firm can wait until all bids are received before they decide on the number of units to allocate. It might arguably be more familiar to require the selling firm to announce the number of units they are putting up for auction prior to the bidding process.

However, the assumption that the number of units awarded can be varied based on the bid values is not as unrealistic as it first seems. For example, Lengwiler (1999) studies a variable-supply auction motivated by the problem of corporations that issue new securities to finance their operations. In this setting, the total number of securities issued is varied based on both the volume and value of the bids they receive. More to the point, in any auction in which a seller uses a reserve price, the quantity awarded is *implicitly* varied depending on how many bids (if any) exceed the reserve price. That is, by posting a reserve price the seller is effectively saying she will not necessarily sell all the units she has. This situation is quite close to our assumption. Indeed, we show that the optimal mechanism in our model in fact reduces to a standard multiunit auction with fixed reserve price. Hence, the "variable-supply" feature of the auction results in a quite familiar, k-unit, fixed-reserve-price auction, and moreover, our results show that this familiar auction is optimal among all possible variable-supply auctions.

Our assumption that the number of buyers N and their values v are i.i.d. also has some important implications. For one, it largely precludes situations where buyers are strategically attempting to time their purchases. For example, if a buyer anticipates that there may be a higher number of units available in the next period, then they might have an incentive to wait for the next auction. This would create dependencies between the inventory position and the number of customers N that arrive. However, we show below that the optimal policy eventually becomes one of running a sequence of identical auctions (same starting inventory, allocation rules, payments, etc.) over time, so the incentive on the part of customers to strategize over timing gradually disappears under our optimal policy.[3]

The independence assumption also precludes the case where buyers might rebid in later auctions, because in this case the number of unsuccessful bidders in the past may influence the distribution of N and v in the future. One possibility here is that unsuccessful bidders "drop out" of the market. For example, buyers might be impatient and buy elsewhere rather than waiting for the next auction. But this is a somewhat delicate explanation, because it implies that the periods are short enough that buyers are willing to wait for the auction result within a period, but the periods are long enough that they will not wait until the next period. Another possibility is that N and v are independent over time simply because buyers are not permitted to rebid.

For example, this strategy is used by Priceline.com; consumers who bid and fail are not allowed to rebid for seven days.

Also, the independence of N and v over time would again not be valid if the firm could learn over time about the valuations of customers through its repeated observations of bidding behavior, as is the case in the finite horizon model studied by Pinker et al. (2000). However, in our infinite horizon, stationary setting, it really does not make much sense to talk about "learning" because, implicitly, our model assumes that the firm has already been observing an infinite history of stationary bidding data. Indeed, in this sense F and g already reflect the firm's accumulated experience and learning over infinitely many past auctions, and thus repeated draws from F and g provide no new information. Of course, stationarity over an infinite horizon is indeed a strong assumption. As a practical matter, most systems are not stationary to this degree, and hence exploiting the information value of bids is a very important issue in practice. Again, see Pinker et al. (2000) for an analysis and discussion of learning in a finite horizon, dynamic auction.

Finally, in our model the auction intervals and the reorder intervals are the same. This is again a limitation, because the factors that drive the frequency of auctions (providing convenience to buyers, administrative costs, etc.) are likely to be different than those that drive the frequency of reorders (production cycles, delivery schedules, fixed ordering costs, etc.). Also, our model assumes zero leadtimes, while in reality there may be several periods of delay before orders arrive. Allowing the auction and reorder periods to be different and allowing for positive leadtimes would be worthwhile extensions, but would result in a more complicated analysis. Hence, we retain these assumptions as a starting point.

1.3. Dynamic Programming Formulation

We analyze this problem using a dynamic programming formulation in terms of the allocation variables $q(v)$ defined by (1). Define the value function $V(x)$ as the maximum expected discounted profit given an initial inventory $x = 0, 1, \ldots,$ which satisfies the Bellman equation

$$V(x) = E\left[\max_{\substack{q\in\{0,1\}^N\\ y\in\mathbf{Z}_+}}\left\{\sum_{i=1}^{N} J(v_i)\,q_i + \alpha V(x-k+y) - h(x-k+y) - cy:\ k = \sum_{i=1}^{N} q_i,\ k \leqslant x\right\}\right], \tag{5}$$

where $0 < \alpha < 1$ is the discount factor, k is the total number of units awarded, and y is the replenishment order for the next period. Note from first principles the state space can be bounded by M, because at most M buyers will arrive in any period, and because we can reorder at the end of every period, there is no need to stock more than M. Our objective is finding an optimal stationary policy, denoted $u^*(x)$, consisting of an allocation $q(\cdot)$ and a replenishment order $y(\cdot)$, that achieves $V(x)$.

We can reformulate our dynamic program using the variables q_i. Using Assumption 1, we can take advantage of the monotonicity of $J(\cdot)$. In this case, when the firm decides to award k units, it is optimal to assign them to the highest $J(v_i)$s (i.e., to the highest v_is). Using reverse order statistics, define

$$R(k) = \begin{cases} 0 & \text{if } k = 0,\\ \displaystyle\sum_{i=1}^{\min\{k,N\}} J(v_{(i)}) & \text{if } k > 0. \end{cases} \tag{6}$$

Note that $R(k)$ is a random function and that

$$\max\left\{\sum_{i=1}^{N} J(v_i) q_i:\ 0 \leqslant q_i \leqslant 1,\ \ \sum_i q_i = \min\{k, N\}\right\} = R(k),$$

so we can rewrite (5) in terms of k as follows:

$$V(x) = E\left[\max_{\substack{0\leqslant k\leqslant x\\ y\in\mathbf{Z}_+}}\{R(k) + \alpha V(x-k+y) - h(x-k+y) - cy\}\right], \quad x = 0, 1, \ldots. \tag{7}$$

Note that above we are assuming free disposal when $N < k \leqslant x$. This assumption is not essential for our analysis, but it helps to simplify the notation.

1.4. Statement of Main Theorems

We next state our main theorems, which characterize the optimal auction and replenishment policy for our problem. The first statement is presented in algorithmic form and the proof is provided in the next section.

THEOREM 1. *Consider the inventory-pricing problem described in* (7). *Define the optimal basestock level by*

$$z^* = \max\{z \in \mathbf{Z}_+:\ \alpha\Delta V(z) - \Delta h(z) - c > 0\}.$$

Then, the optimal stationary policy $u^*(x)$ is to allocate units to buyers and replenish stock according to the following procedure:

Step 1. *Allocate Units*

For $k = 1, 2, \ldots, \min\{x, N\}$, allocate the kth unit if either:

(i) $x - k \geqslant z^*$ *and* $J(v_{(k)}) > \alpha\Delta V(x-k+1) - \Delta h(x-k+1)$,

(ii) $x - k < z^*$ *and* $J(v_{(k)}) > c$,

else, do not award the kth unit and goto Step 2.

Step 2. *Replenish Stock*

If $x - k < z^*$, then order up to z^*, i.e., $y = z^* + k - x$; else order nothing ($y = 0$).

Figure 1. Illustration of optimal policy.

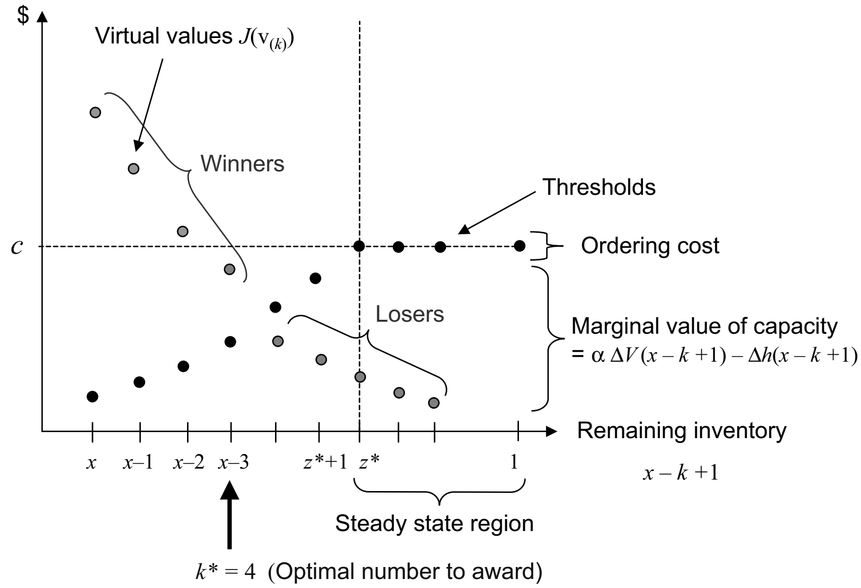

The policy says that while the current inventory is above the optimal basestock level z^* (Case (i)), then we will award the kth unit if the benefit from accepting the kth bid (its virtual value $J(v_{(k)})$) exceeds the profit of keeping the kth unit for the next period less the marginal holding cost for keeping it. The kth unit is not replenished in this case. Once the inventory reaches the optimal level z^* (Case (ii)), the firm awards a unit as long as the benefit from accepting a bid exceeds the cost of replacing the unit awarded; each such unit is replenished.

The optimal allocation policy is illustrated in Figure 1. There are x units to be auctioned in the current period. The black dots represent the threshold prices for the units, given by the marginal value of capacity for the units above the optimal basestock level z^*, and by the ordering cost c for the units up to z^*. The seller sorts the virtual value of the bidders (grey dots) in descending order and compares them with the threshold prices. In Figure 1, there are four winners in the auction: one unit is allocated to each of the top four value bidders, and the process restarts in the next period with $x-4$ units.

An interesting result of this allocation policy is that when the inventory is no more than the optimal basestock level z^*, the seller can achieve the optimal allocation by simply running a standard first-price or second-price auction in each period with a fixed reserve price

$$\hat{c} \equiv J^{-1}(c). \tag{8}$$

Indeed, we have the following characterization of the optimal policy in this case:

THEOREM 2. *Once the inventory reaches z^* units, the optimal policy in all subsequent periods is to use the following basestock, reserve-price-auction policy*: (1) *run a standard first-price or second-price, z^*-unit auction with fixed reserve price $\hat{c}$*; *and then* (2) *at the end of each period, order up to the optimal basestock level z^*.*

Because the problem is over an infinite horizon and the optimal policy only calls for ordering when the inventory drops below z^*, the firm eventually reaches a point where the above basestock, reserve-price-auction policy is optimal for all remaining time.

This result is significant on several levels. First, it shows that the classical first-price and second-price mechanisms remain optimal in the dynamic inventory setting. These are both familiar auction mechanisms, which are easy for buyers to understand and easy for sellers to implement. The inventory replenishment policy is also a familiar and simple basestock policy. This combination makes the optimal policy quite practical. On a theoretical level, the result is as simple as one could hope for in this setting. Finally, it is convenient as well from a computational perspective, because it reduces the optimal policy to a simple search over the single parameter z^*, as we show below in §3.4.

2. Analysis of the Optimal Policy

As mentioned above, the analysis proceeds in two steps. We first analyze the theoretical properties of the dynamic program (5) to characterize the optimal allocation of Theorem 1. We then use the structure of the optimal policy to define two auction mechanisms that achieve this allocation. These mechanisms reduce to the standard first-price and second-price auctions when the inventory is no more than z^*, which is the statement of Theorem 2.

2.1. Proof of Theorem 1

We analyze the infinite horizon dynamic programming formulation (7) as the limit of its corresponding finite horizon version. Defining $V_t(x)$ as the cumulative profit up to period t, we have

$$V_t(x) = E_{N,v}\left[\max_{\substack{0\leqslant k\leqslant x\\ y\in\mathbf{Z}_+}}\left\{R(k)+\alpha V_{t-1}(x-k+y) - h(x-k+y)-cy\right\}\right], \tag{9}$$

with boundary conditions

$$V_t(0)=0,\ t\geqslant 1 \quad\text{and}\quad V_0(x)=0,\ x\geqslant 0.$$

We require the following lemma, characterizing the inner optimization in (9):

Lemma 1. *Suppose $G(x)$ is concave and bounded above, and consider the problem*

$$\max_{\substack{0\leqslant k\leqslant x\\ y\in\mathbf{Z}_+}}\left\{R(k)+\alpha G(x-k+y)-h(x-k+y)-cy\right\}. \tag{10}$$

Let z^* be such that

$$z^* = \begin{cases} \max\{z\geqslant 1\colon \alpha\Delta G(z)-\Delta h(z)-c>0\} & \text{if } \alpha\Delta G(1)-\Delta h(1)-c>0,\\ 0 & \text{otherwise.}\end{cases} \tag{11}$$

Thus, the optimal solution (k^*, y^*) satisfies

(i)

$$y^* = \begin{cases} z^*-x+k^* & \text{if } z^*>x-k^*,\\ 0 & \text{otherwise.}\end{cases}$$

(ii)

$$k^* = \begin{cases} 0 & \text{if } R(1)+\Delta h(x)\leqslant \alpha\Delta G(x),\\ \max\{1\leqslant k\leqslant x-z^*\colon \Delta R(k)+\Delta h(x-k+1) > \alpha\Delta G(x-k+1)\} & \text{if } R(1)+\Delta h(x)>\alpha\Delta G(x) \text{ and } \Delta R(x-z^*+1)\leqslant c,\\ \max\{x-z^*+1\leqslant k\leqslant x\colon \Delta R(k)>c\} & \text{otherwise.}\end{cases}$$

Proof. To prove part (i), take (10) and fix a value of k, the number of units to award. We are then facing a problem only in the number of units to order from the supplier, $y(k)$. Define the inventory position $z\equiv z(k)=x-k+y(k)$. We can then express (10) as a problem in z:

$$\max_{z\in\mathbf{Z}_+}\left\{R(k)+\alpha G(z)-h(z)-cz+cx-ck\right\}.$$

By the concavity of $G(\cdot)$ and the convexity of $h(\cdot)$, the z^* in (11) is the optimal solution of this reformulated problem. Then,

$$y^*(k) = \begin{cases} z^*-x+k & \text{if } z^*>x-k,\\ 0 & \text{otherwise,}\end{cases}$$

and in particular, $y^*\equiv y^*(k^*)$ for some optimal k^* to be determined.

For part (ii), note that $y^*(k)=\max\{z^*, x-k\}-x+k$, turning (10) into a problem just in decision variable k:

$$\max_{0\leqslant k\leqslant x}\left\{R(k)+\alpha G(\max\{z^*,x-k\})-h(\max\{z^*,x-k\}) - c\max\{z^*-x+k,0\}\right\}. \tag{12}$$

For any $0\leqslant k\leqslant x$, we consider two cases according to its value:

(a) If $k\leqslant x-z^*$, then we can rewrite (12) as

$$\max_{0\leqslant k\leqslant x-z^*}\left\{R(k)+\alpha G(x-k)-h(x-k)\right\},$$

which is equivalent to

$$\max_{0\leqslant k\leqslant x-z^*}\left\{\sum_{i=1}^{k}[\Delta R(i)-\alpha\Delta G(x-i+1)+\Delta h(x-i+1)]\right\} + \alpha G(x)-h(x),$$

where the sum is defined to be 0 if $k=0$.

(b) If $k>x-z^*$, then problem (12) turns out to be

$$\max_{x-z^*+1\leqslant k\leqslant x}\left\{R(k)+\alpha G(z^*)-h(z^*)-cz^*+cx-ck\right\},$$

which is equivalent to

$$\max_{x-z^*+1\leqslant k\leqslant x}\left\{\sum_{i=x-z^*+1}^{k}[\Delta R(i)-c]\right\}+\alpha G(z^*)-h(z^*).$$

Essentially, the optimality of the proposed k^* is based on proving that the expression to maximize in (12) has decreasing increments in k. According to both observations above, we split the analysis in two cases. For case (a), note that $\Delta R(k)-\alpha\Delta G(x-k+1)+\Delta h(x-k+1)$ is decreasing in k ($\Delta R(k)$ is decreasing by Assumption 1, $\Delta G(\cdot)$ is increasing in k by its concavity, and $\Delta h(\cdot)$ is decreasing in k by its convexity). For case (b), observe that by (6),

$$\Delta R(k) = \begin{cases} J(v_{(k)}) & \text{if } 1\leqslant k\leqslant N,\\ 0 & \text{otherwise.}\end{cases}$$

Then, $\Delta R(k)-c$ is also decreasing in k.

To complete the proof, we have to check what happens at the transition point $k=x-z^*$. That is, we need to check if the last increment to its left is greater or equal than the first increment to its right, or in symbols, if

$$\Delta R(x-z^*)-\alpha\Delta G(z^*+1)+\Delta h(z^*+1) \geqslant \Delta R(x-z^*+1)-c. \tag{13}$$

By optimality of z^* (see formula (11)),

$$\alpha\Delta G(z^*+1)-\Delta h(z^*+1)-c\leqslant 0.$$

Because we also know that $\Delta R(\cdot)$ is decreasing, then

$$\begin{aligned}&\alpha\Delta G(z^*+1)-\Delta h(z^*+1)-c\\ &\quad\leqslant 0\leqslant \Delta R(x-z^*)-\Delta R(x-z^*+1)\end{aligned}$$

and Equation (13) is verified. Hence, the expression between the large brackets in (12) has decreasing increments in k, and k^* is the largest k for which this increment remains positive. □

To apply Lemma 1 to the finite horizon problem (9), we must verify that $V_{t-1}(\cdot)$ is bounded and concave. Indeed, take a realization (n, v) for problem (9), and assume that $V_{t-1}(\cdot)$ is concave and bounded. By letting $G(\cdot)=V_{t-1}(\cdot)$, Lemma 1 gives closed-form expressions for the optimal inventory level $z^*_{t-1}\equiv z^*$, the optimal number of items to award $k^*(x)\equiv k^*$, and the optimal number of units to replenish, $y^*(x)\equiv y^*$.

The next lemma establishes the boundedness of the value function. The proof is in the appendix.

Lemma 2. *For all $t\geqslant 0$, there exists $K>0$ such that $V_t(x)<K$ $\forall x$.*

We will also require the following lemma, which states that under the concavity condition if we have one more unit available to sell, we allocate at most one more unit to the buyers. It also relates the optimal allocation number to the optimal replenishment number. These properties are helpful both theoretically and computationally. The proof is in the appendix.

Lemma 3. *If $\Delta V_{t-1}(x)$ is decreasing in x, then for any realization (n, v), $k^*(x)\leqslant k^*(x+1)\leqslant k^*(x)+1$ for all $x\geqslant 0$. Moreover, if $k^*(x+1)=k^*(x)$, then*

$$y^*(x+1)=(y^*(x)-1)^+,$$

while if $k^*(x+1)=k^*(x)+1$, then

$$y^*(x+1)=y^*(x).$$

The following lemma establishes that $V_t(\cdot)$ is indeed concave; that is, the marginal value of capacity, $\Delta V_t(x)$, is decreasing in the remaining inventory. The proof is in the appendix.

Lemma 4. *$\Delta V_t(x)$ is decreasing in x.*

Proceeding with the finite horizon version in (9), we next show that we can constrain the feasible set for y so that the per-period profit is bounded both above and below. The proof is in the appendix.

Lemma 5. *There exist $\bar{y}\in\mathbf{Z}_+$ and $L>0$ such that $y^*\leqslant\bar{y}$ and $|R(k)-h(x-k+y)-cy|\leqslant L$ $\forall k, x, y$: $0\leqslant k\leqslant x$, $0\leqslant y\leqslant\bar{y}$. In particular, we can consider $\bar{y}=\bar{z}$, where $\bar{z}\equiv\max\{z\in\mathbf{Z}_+ : \alpha\bar{v}-c>\Delta h(z)\}$ is an upper bound for any optimal per-period inventory level.*[4]

Because both the per-period profit and initial function $V_0(x)$ are bounded, from Bertsekas (1995, §1.2, Assumption D and Proposition 2.1), we have that

$$V(x)=\lim_{t\to\infty}V_t(x)\quad \forall x\geqslant 0. \tag{14}$$

Furthermore, from Proposition 2.2 in Bertsekas (1995, §1.2), the limiting function $V(x)$ is the unique solution to Bellman's Equation (7). This limit allows us to extend the concavity to the infinite horizon profit function.

Lemma 6. *$\Delta V(x)$ is decreasing in x.*

Proof. Because $\Delta V_t(x)\geqslant\Delta V_t(x+1)$, taking the limit of both sides as $t\to\infty$, and using the property described by (14), $\Delta V(x)\geqslant\Delta V(x+1)$ as well. □

From Lemma 2, $V(\cdot)$ is bounded above. Because it is also concave, Lemma 1 gives a complete characterization of the minimizer for the right-hand side in that formulation, by taking the function $G(\cdot)\equiv V(\cdot)$. Following Bertsekas (1995, §1.2, Proposition 2.3), that minimizer is an optimal stationary policy. Indeed, we get the following technical description of the optimal policy, which translates algorithmically into our main Theorem 1:

$$k^*=\begin{cases}0 & \text{if } R(1)+\Delta h(x)\leqslant\alpha\Delta V(x),\\ \max\big\{1\leqslant k\leqslant x-z^*:\ \Delta R(k)+\Delta h(x-k+1) & \\ \qquad >\alpha\Delta V(x-k+1)\big\} & \\ & \text{if } R(1)+\Delta h(x)>\alpha\Delta V(x) \text{ and}\\ & \quad \Delta R(x-z^*+1)\leqslant c,\\ \max\{x-z^*+1\leqslant k\leqslant x:\ \Delta R(k)>c\} & \\ & \text{otherwise,}\end{cases}$$

where

$$z^*=\max\{z\in\mathbf{Z}_+:\ \alpha\Delta V(z)-\Delta h(z)-c>0\}.$$

Afterwards, order y^* units for replenishment, with

$$y^*=\begin{cases}z^*-x+k^* & \text{if } z^*\geqslant x-k^*,\\ 0 & \text{otherwise.}\end{cases}$$

Finally, observe that our system can be viewed as a finite state Markov chain, with states $\{0,1,\ldots,z^*,\ldots,\bar{z}\}$. The dynamics of the system are driven by the random variables (N, v), which induce a change in state through the decision variables k^* and y^*. Because of the structure of the optimal policy, it can be shown that the unique recurrent state is z^* (i.e., z^* is an absorbing state).

2.2. Analysis of the Optimal Auction Mechanism

The next step in our analysis of the problem is to construct auction mechanisms that implement the optimal allocation policy derived above. We will follow ideas introduced in Vulcano et al. (2002) to demonstrate that modified versions of two standard procedures—the first-price and second-price auctions—achieve the optimal allocation. We only outline the basic result for each mechanism in turn, and the reader is referred to Vulcano et al. (2002) for more details.

2.2.1. Second-Price Auction. In a traditional open, k-unit, ascending price auction or in the sealed-bid, second-price-Vickrey-auction, where all winners pay the maximum between the $(k+1)$th highest bid and the fixed reserve price (4), the dominant strategy for a buyer is to bid his true value. However, if one uses a straightforward application of the second-price mechanism in our setting, this is no longer true.

The following modified second-price mechanism avoids this pitfall: For $i \geqslant 1$, let

$$\hat{v}_i = \begin{cases} J^{-1}(\alpha\Delta V(x-i+1) - \Delta h(x-i+1)) \\ \qquad\qquad\qquad \text{if } 1 \leqslant i \leqslant x - z^*, \\ J^{-1}(c) \qquad\qquad \text{if } x - z^* + 1 \leqslant i \leqslant x. \end{cases} \tag{15}$$

The thresholds $\hat{v}_i$ are directly computable from the solution of (7), which uses common knowledge information, and is in principle known to all buyers and the seller. Following the argument in Vulcano et al. (2002, §3.3.1), suppose the firm acts as if a customer's bid is equal to his value. Then, given the vector of submitted bids b, the seller will award k items, where

$$k = \max\{i \geqslant 1\colon b_{(i)} > \hat{v}_i\}$$

and $k = 0$ if $b_{(1)} \leqslant \hat{v}_1$. All winners will pay

$$b_{(k+1)}^{(2\text{nd})} = \max\{b_{(k+1)}, \hat{v}_k\},$$

where $b_{(k+1)}$ is the $(k+1)$th highest bid and $\hat{v}_k$ is the threshold to award the kth unit. Ties between bids are broken by randomization.

Under this modified second-price mechanism, one can show that it is a dominant strategy for buyers to bid their own values. (See Vulcano et al. 2002, §3.3.1 for a detailed argument.) Moreover, because bids are equal to values, this mechanism achieves the optimal allocation of Theorem 1. Note that this fact makes it feasible to relax the assumption that bidders know the distributions F and g, because the dominant strategy of bidding one's own value holds regardless of the number of other buyers or their valuations. Also, note that when $x \leqslant z^*$, this mechanism is equivalent to a standard second-price auction with fixed reserve price $J^{-1}(c)$, because bids are awarded to the x highest value customers with virtual values in excess of c, which proves first part of Theorem 2.

Yet despite the many desirable properties of a second-price auction, Rothkopf et al. (1990) point out that they are somewhat uncommon in practice. Two possible explanations are: (1) bidders may fear truthful revelation of information to third parties with whom they will interact after the auction finishes, and (2) bidders may fear the auctioneer cheating, in the sense that the auctioneer could introduce "artificial" bids to raise the price paid. (Again, see Rothkopf et al. 1990 for a discussion of these issues.) In contrast, Lucking-Reiley (2000) argues that second-price auctions are indeed used, for example, for selling stamps or through the form of proxy bidding used in online auctions such as eBay.

2.2.2. First-Price Auction. In a first-price auction, items are awarded to the highest bidders and winners pay their bids. Note that if we can show that there exists a symmetric equilibrium bidding strategy $B(\cdot)$ that is strictly increasing in the bidders' values, then the firm can invert this bid function to infer each bidder's value, which it can then use to optimally award items.

Regarding the auction setup, the bidders are informed of the current inventory x, and of the following allocation rule: Given a vector of bids b, the seller will award k items, where $k = \max\{i \geqslant 1\colon B^{-1}(b_{(i)}) > \hat{v}_i\}$ and $k = 0$ if $B^{-1}(b_{(1)}) \leqslant \hat{v}_1$, where $\hat{v}_i$ is defined by (15). The items are awarded to the highest bidders, and winners pay their bids. Our first result is the following (see Vulcano et al. 2002, §3.3.2 and the appendix for a proof):

Proposition 1. *The first-price auction has a symmetric equilibrium, strictly increasing, bidding strategy $b_i = B(v_i)$. The strategy B depends on the current value of x as given by*

$$\hat{B}(v_i) = v_i - \frac{\int_{\hat{v}_1}^{v_i} P(v)\,dv}{P(v_i)} \quad \text{and}$$

$$B(v_i) \equiv \lim_{\varepsilon \to 0^+} \hat{B}(v_i - \varepsilon), \tag{16}$$

where $P(v)$ is the probability that a bidder with value v is among the winners,

$$P(v) = \begin{cases} 0 \qquad \text{if } k^*(v) = 0, \\ \displaystyle\sum_{n=1}^{M} \left\{ \sum_{k=0}^{k^*(v)-1} \binom{n-1}{k} [1-F(v)]^k [F(v)]^{n-1-k} \right\} g(n) \\ \qquad \text{if } k^*(v) \geqslant 1, \end{cases}$$

and $k^*(v) = \max\{0 \leqslant i \leqslant \min\{x, N\}\colon v > \hat{v}_i\}$, and by convention, $\hat{v}_0 < 0$.

Again, given this strictly increasing bidding function, the seller can invert the bids to determine a buyer's value. This information can then be used to implement the optimal allocation by checking for $k = 1, 2, \ldots, \min\{x, N\}$ whether $B^{-1}(b_{(k)}) > \hat{v}_k$ and stopping once this condition is violated. This proves the remaining part of Theorem 2.

Note that (16) shows, as one would expect, that under our first-price mechanism, because winners pay what they bid, buyers shade their values to make some positive surplus.

3. Some Extensions to the Basic Model

In this section, we consider some natural extensions to our auction model. We look at these in increasing order of difficulty.

3.1. Charging Holding Cost on the Ending Inventory

Suppose now that we charge the holding cost on the final inventory of each period, rather than on the starting level of the next period as in our original formulation. The dynamic program in this case is the same as (7), but the term $h(x-k+y)$ is replaced by $h(x-k)$. A basestock policy remains optimal for replenishment, but now the optimal basestock level is given by

$$z^* = \begin{cases} \max\{z \geqslant 1\colon \alpha\Delta G(z) - c > 0\} & \text{if } \alpha\Delta G(1) - c > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Regarding the number of units to award, we follow Lemma 1, part (ii). For case (a) in its proof, the allocation rule is the same, but it changes for case (b) by introducing the marginal holding cost. The optimal policy in Theorem 1 becomes:

Step 1. Allocate Units

For $k = 1, 2, \ldots, \min\{x, N\}$, allocate the kth unit if either:

(i) $x - k \geqslant z^*$ and $J(v_{(k)}) > \alpha\Delta V(x-k+1) - \Delta h(x-k+1)$,

(ii) $x - k < z^*$ and $J(v_{(k)}) > c - \Delta h(x-k+1)$,

else, do not award the kth unit and goto Step 2.

Step 2. Replenish Stock

If $x - k < z^*$, then order up to z^*, i.e., $y = z^* + k - x$; else order nothing ($y = 0$).

The case $x - k \leqslant z^*$, corresponding to the steady state of the system, can lead to more complicated auction mechanisms (see §2.2) than the ones presented in Theorem 2. However, if the holding cost is linear, so that $h(z) = a + hz$, then $\Delta h(x-k+1) = h$, a positive constant, and it is again optimal to run a first-price or second-price auction with a fixed reserve price, though the optimal reserve price is now $J^{-1}(c-h)$. This lowered reserve price (with respect to $\hat{c} = J^{-1}(c)$) reflects the fact that now the seller is willing to accept lower bids to avoid one period of holding cost.

3.2. Backorders

Consider the infinite horizon problem of §§1.2 and 1.3, but suppose the firm could award units beyond the current inventory level by backordering, incurring a penalty cost of $b(k)$ when k is the number of buyers backlogged. We assume that the function $b(\cdot)$ is convex increasing with $b(0) = 0$.

Following formulation (7), the dynamic programming formulation for this case is

$$V(x) = E_{N,v}\left[\max_{k,y \in \mathbf{Z}_+}\left\{R(k) + \alpha V(x-k+y) - h(x-k+y) - cy - b\big((k-x)^+\big)\right\}\right].$$

All the analyses developed can be extended to this setting, and we get similar formulas for z^* and y^* to the ones found in Lemma 1:

$$z^* = \begin{cases} \max\{z \geqslant 1\colon \alpha\Delta V(z) - \Delta h(z) - c > 0\} \\ \qquad \text{if } \alpha\Delta V(1) - \Delta h(1) - c > 0, \\ 0 \qquad \text{otherwise,} \end{cases}$$

and

$$y^* = \begin{cases} z^* - x + k^* & \text{if } z^* > x - k^*, \\ 0 & \text{otherwise.} \end{cases}$$

The main change is that the calculation of k^* involves the backorder cost once the seller goes beyond the stock on hand. Following the outline in the proof of Lemma 1, in this case we have cases (a) and (b) as before, plus a new case (c) corresponding to the situation $k > x$. It can be checked that when N bidders show up in a particular period, the optimal k is then

$$k^* = \begin{cases} 0 \qquad \text{if } R(1) + \Delta h(x) \leqslant \alpha\Delta V(x), \\ \max\{1 \leqslant k \leqslant x - z^*\colon \Delta R(k) + \Delta h(x-k+1) \\ \qquad > \alpha\Delta V(x-k+1)\} \\ \qquad \text{if } R(1) + \Delta h(x) > \alpha\Delta V(x) \text{ and} \\ \qquad\quad \Delta R(x - z^* + 1) \leqslant c, \\ \max\{x - z^* + 1 \leqslant k \leqslant x\colon \Delta R(k) > c\} \\ \qquad \text{if } \Delta R(x - z^*) + \Delta h(z^*+1) > \alpha\Delta V(z^*+1) \\ \qquad\quad \text{and } \Delta R(x+1) \leqslant c + \Delta b(1), \\ \max\{x + 1 \leqslant k \leqslant N\colon \Delta R(k) > c + \Delta b(k-x)\} \\ \qquad \text{otherwise.} \end{cases}$$

Regarding the mechanism design for this case, we should modify the definition of $\hat{v}_i$ in (15) to account for the backorder cost. So, suppose that the backorder cost is linear, of the form $b(w) = bw$, with $b > c$. Then,

$$\hat{v}_i = \begin{cases} J^{-1}\big(\alpha\Delta V_{t-1}(x-i+1) - \Delta h(x-i+1)\big) \\ \qquad\qquad \text{if } 1 \leqslant i \leqslant x - z^*, \\ J^{-1}(c) \qquad \text{if } x - z^* + 1 \leqslant i \leqslant x, \\ J^{-1}(b) \qquad \text{if } x + 1 \leqslant i \leqslant N. \end{cases}$$

That means that once the inventory drops below the optimal stationary inventory z^*, the firm essentially sets two reserve prices: one for the available on-hand units, and a higher one for the backlogged units. Both the first-price and second-price mechanisms can be extended to work in this case as well.

3.3. Combined Auction and List-Price Model

Often, firms that sell with an auction mechanism also use a regular, fixed-price mechanism in parallel. In the retail setting, this is often achieved by using each mechanism in a different channel (e.g., catalogue and web channels).[5] In industrial settings, a firm may have fixed-price demand as a result of long-term contracts, while at the same time participates in auctions from spot-purchase customers.[6]

We model this situation as follows: Consider the infinite horizon problem described in §§1.2 and 1.3, but assume that demand comes from two independent streams of customers—one stream for the auction channel and one for the list-price channel. In each period, there are N_A potential buyers that participate in the auction market, with valuations v drawn from a strictly increasing distribution F as before, and N_L buyers that participate in the list-price market. We assume that N_A follows a $g_A(\cdot)$ and that N_L follows a p.m.f. $g_L(\cdot)$. Distribution functions F, g_A, and g_L are common knowledge.

In the list-price market, we assume all N_L buyers are willing to purchase at a list price $r > c$, but the firm may ration its supply and only sell to a subset of these customers (or none at all). Let $V(x)$ denotes the seller's expected discounted profit given a starting inventory of x as before. Let q_i be the allocation binary variables for the auction and k_L be the number of units to award to the list-price buyers. The value function satisfies the Bellman equation

$$V(x) = E\Bigg[\max_{\substack{q\in\{0,1\}^{N_A}\\ y,k_L\in\mathbf{Z}_+}}\Bigg\{\sum_{i=1}^{N_A} J(v_i)q_i + rk_L + \alpha V(x-k+y) - h(x-k+y) - cy\colon k_A = \sum_{i=1}^{N_A} q_i,\, k_L \leqslant N_L,\, k = k_A + k_L,\, k\leqslant x\Bigg\}\Bigg], \quad (17)$$

where $0<\alpha<1$ is the discount factor.

In essence, the idea in analyzing this problem is simply to treat the N_L buyers in the list-price market as if they all had virtual values of r and combine them with the buyers in the auction market. Then, one finds the optimal allocation as before.

Following our previous arguments, relax the integrality of the variables q_i and redefine the function $R(k)$ as in (6). We rank the bidders' virtual values together with the price r, and if k is the total number of units to award through both channels, then $R(k)$ will represent the profit obtainable by optimally awarding those k units (not accounting for costs). Let

$$N_A(r) \equiv \big|\{v_i\colon J(v_i) > r, \forall\, 1\leqslant i\leqslant N_A\}\big|$$

be the random variable representing the number of buyers in the auction market with virtual values above the list price r. If $v_{(i)}$ is the ith reverse order statistic in the auction market, define

$$R(k) = \begin{cases} 0 & \text{if } k=0,\\ \sum_{i=1}^{k} J(v_{(i)}) & \text{if } 0<k\leqslant N_A(r),\\ \sum_{i=1}^{N_A(r)} J(v_{(i)}) + r(k-N_A(r)), & N_A(r) < k \leqslant N_A(r)+N_L,\\ \sum_{i=1}^{k-N_L} J(v_{(i)}) + rN_L, & N_A(r)+N_L < k \leqslant N_A+N_L. \end{cases}$$

Note that

$$R(k) = \max_{q_i,\,k_L}\Bigg\{\sum_{i=1}^{N_A} J(v_i)q_i + rk_L\colon 0\leqslant q_i\leqslant 1,\; k_A = \sum_{i=1}^{N_A} q_i,\, k_L\leqslant N_L,\, k = k_A+k_L\Bigg\},$$

so we can rewrite (17) in terms of k as follows:

$$V(x) = E\Bigg[\max_{\substack{0\leqslant k\leqslant x\\ y\in\mathbf{Z}_+}} \{R(k) + \alpha V(x-k+y) - h(x-k+y) - cy\}\Bigg].$$

We have thus reduced this case to a form essentially identical to our formula (7) in §1.3, observing that $\Delta R(k)$ is by construction decreasing in k. The allocation policy is therefore the same as before, again assuming we treat the fixed-price market buyers as if they are simply buyers with virtual valuations of r. For example, suppose $r > J(\bar{v})$ so that the list price is greater than any virtual value observed in the auction market (e.g., the auction is a "deep discount" market). Then, the optimal policy (assuming the starting inventory is at its steady state value of z^*) is to award as many units to the fixed-price market as possible. If there is any excess, it is sold in the auction market to the highest bidders with virtual values in excess of c. Afterwards, the firm replenishes its stock to bring its inventory up to z^* and the process repeats.

3.4. Average Profit Criterion

We next examine the long-run average profit per stage objective. This criterion leads to further simplification in computing the optimal policy and is useful in its own right.

Let k_t and y_t be the optimal number of units to award and replenish in period t, respectively, with $t\geqslant 1$. We denote by x_t the inventory level at the end of period t (i.e., initial inventory of period $t+1$). Let $V(x)$ be the maximum expected average profit when starting with $x_0 = x$ units of inventory in period $t=1$. This version of the problem can then be formulated as finding (nonanticipating) values k_t, y_t that maximize

$$\lim_{T\to\infty}\frac{1}{T}E\left[\left\{\sum_{t=1}^{T}[R(k_t)-h(x_t)-cy_t]\right\}: \\ x_t=x_{t-1}-k_t+y_t, k_t\leqslant x_{t-1}, \forall 1\leqslant t\leqslant T\right], \qquad (18)$$

where the function $R(\cdot)$ is defined in (6).

One can show that the optimal policy for our α-discounted problem is in fact Blackwell optimal (see Bertsekas 1995, §4.2, Definition 1.1); that is, it is simultaneously optimal for all discounted problems with discount factors $\alpha\in(\bar{\alpha}, 1)$ for some $0<\bar{\alpha}<1$. Therefore, using Proposition 2.2 in Bertsekas (1995, §4.2), $u^*(x)$ is optimal for the average profit problem within the class of all stationary policies. Moreover, because we have seen that $u^*(x)$ involves just one recurrent state, represented by the optimal inventory level z^*, then it is *unichain*.[7] Thus, by Proposition 2.6, condition 1, in Bertsekas (1995, §4.2), $u^*(x)$ is optimal within the class of all admissible policies. Furthermore, the corresponding average profit in Equation (18) is independent of the initial state x. As a result, the optimal average profit policy will again be a basestock, reserve-price-auction policy with reserve price $J^{-1}(c)$ and basestock level z^*.

Indeed, as a result from this fact, we can develop a quite simple procedure for finding the optimal basestock level z^* in the average profit case. Let $\Pi(z)$, described by

$$\Pi(z)=E\left[\max_{0\leqslant k\leqslant \min\{z,N\}}\{R(k)-ck\}\right]-h(z), \qquad (19)$$

be the average profit when following a policy of reordering up to a fixed basestock level z. We know that such a policy will be optimal for some z^*, so we simply need to search for a value z that maximizes $\Pi(z)$. In fact, the search is quite simple as the profit function is concave. The proof is in the appendix.

Proposition 2. *The profit function $\Pi(z)$ is concave in z.*

We know that $0\leqslant z^*\leqslant \bar{z}$, where $\bar{z}\equiv\max\{z\in\mathbf{Z}_+: \bar{v}-c>\Delta h(z)\}$ (see Lemma 5) or $\bar{z}=M$ if $h(\cdot)$ is not strictly convex. For a fixed z, we could use Monte Carlo simulation to calculate $\Pi(z)$: just sample instances for (N, v) and take the average. In some cases, however, this can be avoided. Specifically, we can rewrite (19) conditioning on N and on the number of values above the $\hat{c}$ given in (8). For a given realization n of N, and assuming by convention $v_{(n+1)}=0$, we have

$$\Pi(z)=\sum_{k=1}^{n}E[R(\min\{k,z\})-c\min\{k,z\}] \\ \cdot P(v_{(k)}>\hat{c}, v_{(k+1)}\leqslant\hat{c})-h(z),$$

where it can be easily checked that

$$P(v_{(k)}>\hat{c}, v_{(k+1)}\leqslant\hat{c})=\binom{n}{k}[1-F(\hat{c})]^k F(\hat{c})^{n-k}.$$

This closed form for the expected profit reduces to a simple expression when the buyer's values are uniformly distributed. Also, while computing the sum from term k to $k+1$, we can use the fact that

$$\binom{n}{k+1}=\frac{n-k}{k+1}\binom{n}{k},$$

to reduce the complexity introduced by the combinatorial numbers. Finally, taking advantage of the concavity of $\Pi(z)$, a binary search over the range for z^* gives an overall algorithm complexity of $O(n\log\bar{z})$. Henceforth, we will denote the optimal objective value $\Pi^*\equiv\Pi(z^*)$.

4. Comparisons to a List-Price Policy

We next consider how the optimal basestock, reserve-price-auction policy compares to a traditional, basestock, list-price policy. Specifically, we consider the case where the seller sets a fixed list price p in each period, and then replenish by ordering up to a fixed basestock level z. To be consistent with §3.4, the holding cost is charged on the initial inventory of the next period and we assume buyers who are interested in acquiring one unit at the posted list price submit "acceptances." If the number of acceptances exceeds the current inventory of the seller, the units are randomly rationed to the buyers. It is easy to see that under this pricing mechanism, a dominant strategy for the buyers is to submit an "acceptance" if and only if their own values are higher than the list price.

We will compare the profits earned under our optimal auction policy with those under the list-price policy for an optimal choice of p and z. Theoretical comparisons are provided first, followed by a numerical study of the two policies. The comparisons give some insight into when an auction-based policy has advantages over list pricing, and how much more beneficial it is.

Before proceeding, we note as mentioned in the Introduction that the dynamic list-price and inventory problem was studied extensively by Federgruen and Heching (1999, §4). They showed that in the infinite horizon, stationary case, the optimal policy is in fact a basestock, list-price policy of this form. However, there are differences between the problem Federgruen and Heching (1999) analyzed and our setting. The first, which is minor, is that the authors assume that the holding cost is incurred at the end of the period, resulting in an expected inventory-backorder cost function $G(y, p)$, where y is the inventory level at the beginning of the period, and p is the unit price. This can be mapped into our case by defining $G(y, p)=h(y)-E[b([y-D(p)]^-)]$, where $b(\cdot)$ is a backlog penalty function, and $D(p)$ is the random demand faced at price p. Second, and more fundamentally, ours is a lost sales model while Federgruen and

Heching allow backorders, and they do not provide a proof of the optimality of the basestock, list-price policy in the lost sales case. Still, this policy provides a useful benchmark for comparison.

4.1. Theoretical Comparisons

We first look at some theoretical comparisons of the optimal auction and list-price policies. We will restrict ourselves to the average-cost case, where the optimal profit is given by optimizing (19) over z, though similar results can be developed for the discounted case. Our analysis shows that in several important cases, the list-price policy is provably either optimal or asymptotically optimal, so there may be no benefit to the seller in using an auction policy in these (not unrealistic) settings.

To begin, let Π^*_{LP} denote the seller's average per-period expected profit under this list-price setting. Let s denote the reserve price set by the seller, and let

$$N(s) \equiv \left|\{v_i\colon\, v_i > s,\ \forall\, 1 \leqslant i \leqslant N\}\right|$$

be the random variable representing the number of buyers with valuations exceeding the reserve price s. The seller solves

$$\Pi^*_{\mathrm{LP}} = \max_{\substack{s\geqslant 0\\ z\in \mathbf{Z}_+}} E\big[(s-c)\min\{N(s), z\}\big] - h(z).$$

Fix a value of z, and define the corresponding function $\Pi_{\mathrm{LP}}(z)$ as

$$\begin{aligned}\Pi_{\mathrm{LP}}(z) &= \max_{s\geqslant 0} E\big[(s-c)\min\{N(s), z\}\big] - h(z)\\ &= \max_{s\geqslant 0}\left\{(s-c)\sum_{n=0}^{M} g(n)\left[\sum_{j=0}^{n}\min\{j,z\}\binom{n}{j}\right.\right.\\ &\qquad\qquad \left.\left.\cdot\,[1-F(s)]^j F(s)^{n-j}\right]\right\} - h(z), \qquad (20)\end{aligned}$$

where the last equality follows by conditioning on the outcome of $N(s)$. Note that the objective in (20) is continuous in s, positive for all $s \in (c, \bar{v})$, and zero when $s = c$ or $s = \bar{v}$. So, the maximum is guaranteed to exist in the interval $(c, \bar{v})$ and we can find it through standard line search methods.

In the next step, the seller must solve

$$\Pi^*_{\mathrm{LP}} = \max_{z\in\mathbf{Z}_+} \Pi_{\mathrm{LP}}(z).$$

The search space for z^*_{LP} is clearly bounded between 0 and M: the seller will not stock units beyond the maximum number of bidders that can show up in a particular period. We will keep the notation $\Pi^*_{\mathrm{LP}} = \Pi_{\mathrm{LP}}(z^*_{\mathrm{LP}})$.

4.1.1. Small Number of Buyers per Period. The first case where list pricing is optimal is when there is at most one buyer per period; we have the following proposition. The proof is in the appendix.

PROPOSITION 3. *If $N \leqslant 1$, then $\Pi^* = \Pi^*_{\mathrm{LP}}$.*

This shows that if the firm is receiving isolated bids (for example, as in Priceline.com's pricing mechanism), there is no inherent advantage to using auctions over list pricing—some aggregation of buyers is needed to gain a strict advantage through an auction policy. Intuitively, this is due to the fact that one needs to generate some bidding competition among buyers to realize a benefit from an auction. With at most one buyer bidding, no competition is created.

4.1.2. Large Number of Buyers per Period and Linear Holding Cost. While an auction is not beneficial if there are too few buyers per period, we next show that it is not beneficial if there are too many buyers either. Specifically, we show that if the holding cost is linear, list pricing becomes optimal for the limiting problem where the number of buyers per period tends to infinity. An analogous result was shown for the finite horizon optimal auction design problem in Vulcano et al. (2002, §5.1). Broadly speaking, as the number of buyers increases, the bid realizations become an accurate and dense sample of the value distribution F. As a result, the order statistics, which determine the winning set and price to be paid, will converge to appropriate fractiles of this distribution. The scaled auction profit then approaches a deterministic function of the basestock level z and of the minimum reserve price $\hat{c}$, making it asymptotically equivalent to the list-price profit.

To proceed, we consider a sequence of problems indexed by n. Without loss of generality, we analyze the profits from using a second-price auction with fixed reserve price $\hat{c}$ and an order-up-to policy with basestock level z^n, because we know from §3.4 that such a policy optimizes the long-run average profit for a suitable choice of z^n. The number of buyers in a period is denoted by N^n. For some $z \geqslant 0$ and $N \geqslant 0$, we assume

$$\frac{z^n}{n} \to z \quad \text{and} \quad \frac{N^n}{n} \to N.$$

If N^n is random, we assume that N^n/n converges to a deterministic limit almost surely. The next proposition states the result. A detailed proof is in the appendix.

PROPOSITION 4. *Consider a sequence of auction problems as described above with linear holding cost $h(z) = a + hz$. Then, the list-price policy is optimal for the limiting problem as n becomes large.*

4.1.3. Zero Holding Cost. Finally, the list-price policy is also optimal in the limit as the holding cost tends to zero. Specifically, we have (see the appendix for a proof):

PROPOSITION 5. *If the holding cost $h(z) = 0$, then $\Pi^*_{\mathrm{LP}} = \Pi^*$.*

Table 1. Profits for different numbers of buyers per period.

Buyers per Period	Auction			List Price			Profit Gap (%)
	Profit	z^*	Fillrate (%)	Profit	z^*_{LP}	Fillrate (%)	
1	0.021	1	100.00	0.021	1	100.00	0.00
5	0.128	2	90.39	0.124	3	98.74	3.20
10	0.268	4	95.93	0.261	4	96.62	2.50
50	1.404	14	95.03	1.381	16	98.88	1.62
100	2.835	26	94.90	2.798	30	99.32	1.31
1,000	28.723	242	95.86	28.544	259	99.80	0.62

The intuitive reason for this result is that with no holding cost, the firm will stock the maximum inventory M at the start of each period under both the optimal auction and list-price policies. As a result, there is no rationing of product, and thus buyers do not face any bidding competition. Without bidding competition, the auction produces the same profits as the list-price policy.

4.2. Numerical Comparisons

We next look at the results of some numerical examples that illustrate cases in which an auction policy is more profitable than a list-price policy. We restrict attention to the average profit versions of the problems as discussed in §3.4 and the beginning of §4.1, because the average profit criterion is quite natural, easier to compare than a discounted criterion, and the computations are quite straightforward. In every experiment that follows, we have solved the corresponding formulations (19) and (20) in closed form.

To study the impact of different parameters on the profits earned under both pricing mechanisms, we perturbed the parameters of the following base case: The ordering cost is normalized at $c = 1$; buyers' values are assumed uniform of width $\Delta = 0.5$ centered at c (i.e., buyers' values are centered at the cost, with Δ representing the dispersion in valuations); there are a constant $N = 50$ buyers per period; and the holding cost is linear of the form $h(z) = hcz$, where $h = 1\%$ is the one-period holding cost rate.

We then varied individual parameters of this base case to see the effect on the absolute and relative performance of each policy. Along with expected profit, we computed a "fillrate" for each policy, defined as the expected number of buyers who get an item awarded divided by the expected number who attempt to purchase (e.g., those with values above the reserve price in the auction, or those with values above the fixed price in the list-price case). Formally, the fillrate is the ratio $E[\min\{N(\hat{c}), z^*\}]/E[N(\hat{c})]$ in the auction case, and $E[\min\{N(s^*), z^*_{\mathrm{LP}}\}]/E[N(s^*)]$ in the list-price case, where s^* is the optimal price calculated by the algorithm. The fillrate gives a measure of the scarcity of inventory relative to demand and is a traditional service measure in inventory problems.

4.2.1. The Effect of the Number of Buyers per Period. In our first experiment, we studied how the profit is affected by the number of buyers in each period. The number of buyers N was assumed constant, but N was varied from 1 to 1,000. All other parameters are the same as in the base case. The results are summarized in Table 1.

As one would expect, the profits and inventory levels increase in both policies as the number of buyers increases. Also, as shown theoretically in Proposition 3, the list-price policy is optimal in the limiting case of just one buyer per period. In the other extreme as n gets large, again the list-price profit approaches the optimal auction profit, as predicted by the asymptotic result of Proposition 4. In particular, for our parameters, the limiting dynamic control problem (see formula (26) in the appendix) gives an estimate of the optimal auction basestock level of $z^* = 0.24N$, which is quite close to the values in the third column of Table 1 for the cases $N = 100$ and $N = 1{,}000$. The biggest benefit from the auction occurs at a moderate value of 5 buyers per period, where it achieves a 3.2% increase in profits over list pricing.

Note the fillrate and inventory level is higher in the list-price case. The intuition here is that the auction policy deliberately introduces some scarcity in the available goods to create more bidding competition among the buyers. This is consistent with the findings of Vakrat and Seidmann (2000), who study the impact of the number of units offered in online auctions on the total firm profits and find that profits are a unimodal function of the quantity available (our inventory level).[8]

4.2.2. The Effect of Variability in the Number of Buyers per Period. In our second experiment, we assumed that the number of bidders is uniformly distributed with mean 50. The variance of this distribution was then varied by changing the range of this discrete uniform random variable. All other parameters are the same as in the base case. Results are shown in Table 2.

The main observation here is that as the variance in the number of buyers increases, the seller's profit decreases under both the auction and the list-price policy and the auction becomes relatively more profitable. Thus, high levels of uncertainty about the number of bidders appears to favor the use of the auction policy. As in the previous experiment, we observe higher inventory levels and fillrates for the list-price case.

4.2.3. The Effect of Different Holding Costs. We next consider varying the holding cost through a change in the holding cost rate coefficient h, namely $h(z) = hcz$.

Table 2. Profits for different variances in the number of buyers per period.

Range of Buyers		Auction			List Price			
Min	Max	Profit	z^*	Fillrate (%)	Profit	z^*_{LP}	Fillrate (%)	Profit Gap (%)
50	50	1.404	14	95.03	1.381	16	98.88	1.62
40	60	1.398	15	96.10	1.374	17	98.99	1.74
30	70	1.386	15	93.69	1.358	18	98.69	2.03
20	80	1.371	17	94.59	1.339	20	98.73	2.33
10	90	1.354	18	93.31	1.319	21	98.29	2.59

Typically, this holding cost rate represents a cost of capital plus a rate of depreciation in the product's value over time. Table 3 shows the results.

The small difference in the expected profits for the lowest holding cost confirms the result of Proposition 5; low holding cost leads to high inventory levels, which reduces the bidding competition and hence the benefit of the auction. As the holding cost rises, the auction performs relatively better, achieving a large 21.67% improvement when the holding cost rate reaches 10%. This is simply the reverse effect: A high holding cost means that the firm is unwilling to stock much inventory. Because the number of buyers per period is unchanged, the number of buyers per unit of inventory increases; more competition among buyers is created and hence the auction policy performs relatively better.

It is worth pointing out, however, that there are few practical situations where holding cost rates of over 1% per period are observed, especially if one is considering auctions that are held relatively frequently (e.g., weekly). Rates this high are observed for products such as personal computers which become obsolete quickly, but for most goods weekly rates of less than 1% are the norm. Thus, the experiment suggests that either the product has to suffer fairly rapid depreciation or selling events have to be relatively infrequent (e.g., monthly or semiannual periods, not weekly) for the firm to realize a significant benefit from using auctions over list pricing.

4.2.4. The Effect of Different Levels of Variability in Buyers' Valuations. Finally, we looked at the effect of different levels of variability in the buyers' valuations for items. Recall that these values are assumed to be uniformly distributed with mean one (i.e., centered around the ordering cost $c = 1$) with range Δ. Thus, the values are $U(c-\Delta/2, c+\Delta/2)$. We then varied Δ. All other parameters are the same as in the base case. Table 4 shows the results.

The main observation is that the seller benefits, both in the auction and in the list-price setting, from increased variability in the buyer's valuations. This is to be expected, because the firm can extract more consumer surplus from high-value buyers as the variation increases.

The inventory level is also increasing with the variance, and the relative benefit from the auction policy is decreasing. Intuitively, there are two effects at work here. The first one is that with higher variation in valuations, there is more potential for profit gain through using the auction because it can potentially capture more consumer surplus. However, as the variability of the valuations increases, the level of inventory increases as well, reducing the bidding competition among buyers. In this example, this latter effect dominates the former. One can construct other cases where more variation in valuations increases the relative benefit of the auction, for example, when valuations are strictly higher than the cost c (rather than being centered at c as in this experiment).

5. Conclusions

With the rise of Internet commerce, auctions are increasingly viewed as a viable mechanism for pricing goods in retail and distribution businesses. Our results show that the optimal auction and replenishment policy in a stylized model of such systems is relatively simple, consisting of running a series of standard first-price or second-price auctions with fixed reserve price and following a simple order-up-to (basestock) policy for replenishment. Moreover, especially under the average profit criterion, this optimal basestock, reserve-price-auction policy is very easy to

Table 3. Profits for different levels of holding cost.

Holding Cost Rate (%)	Auction			List Price			
	Profit	z^*	Fillrate (%)	Profit	z^*_{LP}	Fillrate (%)	Profit Gap (%)
0.01	1.560	21	99.97	1.560	23	100.00	0.01
0.10	1.543	18	99.58	1.541	20	99.92	0.14
1.00	1.404	14	95.03	1.381	16	98.88	1.62
5.00	0.932	10	77.36	0.845	11	93.10	9.37
10.00	0.502	7	55.77	0.393	7	82.41	21.67

Table 4. Profits for different variances in buyers' valuations.

Values' Range	Auction			List Price			Profit
Δ	Profit	z^*	Fillrate (%)	Profit	z^*_{LP}	Fillrate (%)	Gap (%)
0.1	0.186	10	77.36	0.168	11	93.09	9.39
0.5	1.404	14	95.03	1.381	16	98.88	1.62
1.0	2.955	15	97.04	2.933	18	99.64	0.76
1.5	4.512	16	98.35	4.489	18	99.64	0.51
2.0	6.070	17	99.13	6.048	19	99.82	0.36

compute. Thus, the structure and computation of an optimal policy is surprisingly simple, familiar, and practical.

In addition, our comparisons with list pricing provide some insight into when an auction is likely to be beneficial. The somewhat notable finding here is that there are relatively few cases in which auctions provide large benefits. If the number of buyers per auction is either very small or very large, list pricing is provably optimal; if holding costs tend to zero, list pricing again becomes optimal. Auctions provided large benefits in our experiments only when (1) the number of buyers is moderate, (2) the holding cost per period is large, and (3) there is significant variability in the number of bidders. Thus, while auctions are quite fashionable in e-commerce, perhaps there are sound theoretical reasons why list pricing remains popular. But when the conditions are right, auctions can indeed provide significant improvements in profits over list pricing.

Appendix

Proof of Lemma 2. Note that it is enough to prove it for a particular instance (n, v). Actually, we will first prove that for all $t \geqslant 0$, there exist a $K_t > 0$ such that $V_t(x) < K_t$. By induction, the result is true for $t = 0$, because $V_0(x) = 0$. Suppose it is true for $V_{t-1}(\cdot)$, and define

$$H_t(x,k,y) = R(k) + \alpha V_{t-1}(x-k+y) - h(x-k+y) - cy.$$

Observe that $H_t(x, k, y) < M\bar{v} + \alpha K_{t-1}$, because $R(k) \leqslant M\bar{v}\ \ \forall\, 0 \leqslant k \leqslant x$. So, define

$$K_t \equiv M\bar{v} + \alpha K_{t-1},$$

with $K_0 = 0$, to have a time-dependent bound for the profit function. This sequence of bounds is increasing, with $\limsup_{t\to\infty} K_t = M\bar{v}/(1-\alpha)$, leading to $K \equiv M\bar{v} \cdot (1-\alpha)^{-1}$. □

Proof of Lemma 3. Recall that when referring to Lemma 1, we will be replacing the generic function $G(\cdot)$ by $V_{t-1}(\cdot)$; and accordingly, we will be talking about z^*_{t-1} instead of z^*.

We start with the left-hand side inequality. If $k^*(x) = 0$, the statement immediately holds true. If $k^*(x) > 0$, the next lines show that in fact $k^*(x)$ is a feasible allocation quantity when the inventory is $x+1$. Following cases (a) and (b) of the proof of Lemma 1, the transition value for k here will be $k = x + 1 - z^*_{t-1}$. Now, if $k^*(x) \leqslant x - z^*_{t-1} - 1$, then using the decreasing increments of $R(\cdot)$ and $V_{t-1}(\cdot)$, and the convexity of $h(\cdot)$, we get

$$\begin{aligned}
&\Delta R(k^*(x)) - \alpha\Delta V_{t-1}(x+1-k^*(x)+1)\\
&\quad + \Delta h(x+1-k^*(x)+1)\\
&\qquad \geqslant \Delta R(k^*(x)) - \alpha\Delta V_{t-1}(x-k^*(x)+1)\\
&\qquad\quad + \Delta h(x-k^*(x)+1) > 0.
\end{aligned}$$

If $k^*(x) = x - z^*_{t-1}$, using Lemma 1, part (ii) for the $x+1$-inventory case, we want to see that

$$\Delta R(x - z^*_{t-1}) - \alpha\Delta V_{t-1}(z^*_{t-1}+2) + \Delta h(z^*_{t-1}+2) > 0.$$

But this holds because $\Delta R(x - z^*_{t-1}) > c$ by optimality of $k^*(x)$ in the x-inventory case; and because of the optimality of z^*_{t-1} in (11), which gives

$$c - \alpha\Delta V_{t-1}(z^*_{t-1}+2) + \Delta h(z^*_{t-1}+2) \geqslant 0.$$

The case $k^*(x) \geqslant x - z^*_{t-1} + 1$ is easy: because of the optimality of $k^*(x)$ when $k \geqslant x - z^*_{t-1}$, $\Delta R(k^*(x)) > c$.

The optimal allocation $k^*(x+1)$ may be higher, however, due to the additional unit in stock. Hence, $k^*(x+1) \geqslant k^*(x)$.

For the right-hand side inequality, suppose by contradiction that $k^*(x+1) > k^*(x) + 1$.

First, take the case $k^*(x) \leqslant x - z^*_{t-1} - 2$. From Lemma 1, part (ii),

$$\begin{aligned}
&\Delta R(k^*(x)+2) + \Delta h(x+1-[k^*(x)+2]+1)\\
&\quad > \alpha\Delta V_{t-1}(x+1-[k^*(x)+2]+1).
\end{aligned}$$

But because $\Delta R(\cdot)$ is decreasing,

$$\begin{aligned}
&\Delta R(k^*(x)+1) + \Delta h(x-[k^*(x)+1]+1)\\
&\quad \geqslant \Delta R(k^*(x)+2) + \Delta h(x-k^*(x))\\
&\quad > \alpha\Delta V_{t-1}(x-[k^*(x)+1]+1),
\end{aligned}$$

contradicting the optimality of $k^*(x)$.

On the other hand, when $k^*(x) \geqslant x - z^*_{t-1} - 1$, then from our supposition, $k^*(x+1) \geqslant x + 1 - z^*_{t-1}$. So, we are in the last case of Lemma 1, part (ii) for the $x+1$-inventory case, and so $\Delta R(k^*(x)+2) > c$, contradicting again the optimality of $k^*(x)$. Furthermore, this tells us that in most cases where $k^*(x) \geqslant x - z^*_{t-1} - 1$, then $k^*(x+1) = k^*(x)$, except possibly when $k^*(x)$ is binding (i.e., $k^*(x) = x$).

Regarding the relation between $y^*(x+1)$ and $y^*(x)$, it follows by inspection from Lemma 1, part (i). □

Proof of Lemma 4. Again, recall that when referring to Lemma 1, we will be replacing $G(\cdot)$ by $V_{t-1}(\cdot)$; and accordingly, $z^*_{t-1} \equiv z^*$.

We proceed by induction on t. For $t = 0$, the theorem trivially holds because $V_0(x) = 0$ for all x. For period $t-1$, the inductive hypothesis (IH) is that $\Delta V_{t-1}(x) \geqslant \Delta V_{t-1} \cdot (x+1)$. Moreover, from (11), this assures the existence of the optimal inventory position z^*_{t-1}. We will then show that if the IH holds, then $\Delta V_t(x)$ is decreasing as well.

To do so, fix the number of bidders n and consider a given realization $v = (v_1, \ldots, v_n)$ of bidders's valuations. Define the maximized value in (12) as

$$\begin{aligned} &H_t(x,n,v)\\ &= \max_{0\leqslant k\leqslant x}\big\{R(k)+\alpha V_{t-1}(\max\{z^*_{t-1},x-k\})\\ &\qquad -h(\max\{z^*_{t-1},x-k\})-c\max\{z^*_{t-1}-x+k,0\}\big\}, \end{aligned} \tag{21}$$

and take the difference function

$$\Delta H_t(x,n,v) = H_t(x,n,v) - H_t(x-1,n,v).$$

Note that for random N and v,

$$\Delta V_t(x) = E_{N,v}[\Delta H_t(x,N,v)].$$

Thus, it suffices to establish that $\Delta H_t(x,n,v)$ is decreasing in x to prove that $\Delta V_t(x)$ is decreasing in x. For notational simplicity, we henceforth suppress the arguments n, v in $\Delta H_t(x,n,v)$ and simply use $\Delta H_t(x)$.

Using (21), cases (a) and (b) in the proof of Lemma 1 and Lemma 3, we make the following observations:

Observation 1. *If $k^*(x+1) = k^*(x) \equiv k^*$ and $k^* \leqslant x - z^*_{t-1}$, then*

$$\Delta H_t(x+1) = \alpha\, \Delta V_{t-1}(x+1-k^*) - \Delta h(x+1-k^*).$$

Observation 2. *If $k^*(x+1) = k^*(x) \equiv k^*$ and $k^* \geqslant x + 1 - z^*_{t-1}$, then*

$$\Delta H_t(x+1) = c.$$

Observation 3. *If $k^*(x+1) = k^*(x)+1$, then*

$$\Delta H_t(x+1) = \Delta R(k^*(x)+1).$$

Consider now $\Delta H_t(x+1)$ and $\Delta H_t(x)$. Given the different combinations of values that $k^*(x-1)$, $k^*(x)$, and $k^*(x+1)$ can take by Lemma 3, there are several cases to analyze:

Case 1. $k^*(x+1) = k^*(x) = k^*(x-1) \equiv k^*$ and $k^* \leqslant x - 1 - z^*_{t-1}$.

In this scenario,

$$\begin{aligned} \Delta H_t(x) &= \alpha\Delta V_{t-1}(x-k^*) - \Delta h(x-k^*)\\ &\qquad \text{(by observation 1)}\\ &\geqslant \alpha\Delta V_{t-1}(x+1-k^*) - \Delta h(x+1-k^*)\\ &\qquad \text{(by the IH and convexity of } h(\cdot))\\ &= \Delta H_t(x+1) \quad \text{(by observation 1)} \end{aligned}$$

Then, $\Delta H_t(x) \geqslant \Delta H_t(x+1)$.

Case 2. $k^*(x+1) = k^*(x) = k^*(x-1) \equiv k^*$ and $k^* = x - z^*_{t-1}$.

Here,

$$\begin{aligned} \Delta H_t(x) &= c \quad \text{(by observation 2)}\\ &\geqslant \alpha\Delta V_{t-1}(z^*_{t-1}+1) - \Delta h(z^*_{t-1}+1)\\ &\qquad \text{(by Equation (11))}\\ &= \Delta H_t(x+1) \quad \text{(by observation 1).} \end{aligned}$$

Then, $\Delta H_t(x) \geqslant \Delta H_t(x+1)$.

Case 3. $k^*(x+1) = k^*(x) = k^*(x-1) \equiv k^*$ and $k^* \geqslant x + 1 - z^*_{t-1}$.

By Observation 2, $\Delta H_t(x) = c = \Delta H_t(x+1)$.

Case 4. $k^*(x+1) = k^*(x) \equiv k^*, k^* > k^*(x-1)$, and $k^* \leqslant x - z^*_{t-1}$.

From Lemma 3, $k^* = k^*(x-1)+1$. Thus,

$$\begin{aligned} \Delta H_t(x) &= \Delta R(k^*(x-1)+1) \quad \text{(by observation 3)}\\ &= \Delta R(k^*)\\ &> \alpha\, \Delta V_{t-1}(x+1-k^*) - \Delta h(x+1-k^*)\\ &\qquad \text{(by optimality of } k^*)\\ &= \Delta H_t(x+1) \quad \text{(by observation 1).} \end{aligned}$$

Then, $\Delta H_t(x) > \Delta H_t(x+1)$.

Case 5. $k^*(x+1) = k^*(x) \equiv k^*, k^* > k^*(x-1)$, and $k^* \geqslant x + 1 - z^*_{t-1}$.

From Lemma 3, $k^* = k^*(x-1)+1$. Thus,

$$\begin{aligned} \Delta H_t(x) &= \Delta R(k^*(x-1)+1) \quad \text{(by observation 3)}\\ &= \Delta R(k^*)\\ &> c \quad \text{(by case (b) in the proof of Lemma 1)}\\ &= \Delta H_t(x+1) \quad \text{(by observation 2).} \end{aligned}$$

Then, $\Delta H_t(x) > \Delta H_t(x+1)$.

Case 6. $k^*(x+1) > k^*(x) > k^*(x-1)$.

From Lemma 3, $k^*(x+1) = k^*(x)+1$ and $k^*(x) = k^*(x-1)+1$. Then,

$$\begin{aligned} \Delta H_t(x) &= \Delta R(k^*(x-1)+1) \quad \text{(by observation 3)}\\ &\geqslant \Delta R(k^*(x-1)+2)\\ &= \Delta R(k^*(x)+1)\\ &= \Delta H_t(x+1) \quad \text{(by observation 3).} \end{aligned}$$

Therefore, $\Delta H_t(x) \geqslant \Delta H_t(x+1)$.

Case 7. $k^*(x+1) > k^*(x) = k^*(x-1) \equiv k^*$ and $k^* \leqslant x-1-z^*_{t-1}$.

Note that $k^*(x+1) = k^*+1$. So,

$$\begin{aligned}\Delta H_t(x) &= \alpha\Delta V_{t-1}(x-k^*) - \Delta h(x-k^*)\\ &\qquad \text{(by observation 1)}\\ &= \alpha\Delta V_{t-1}(x+1-[k^*+1]) - \Delta h(x+1-[k^*(x)+1])\\ &\geqslant \Delta R(k^*+1) \quad \text{(by Lemma 1, part (ii))}\\ &= \Delta H_t(x+1) \quad \text{(by observation 3).}\end{aligned}$$

Then, $\Delta H_t(x) \geqslant \Delta H_t(x+1)$.

Case 8. $k^*(x+1) > k^*(x) = k^*(x-1) \equiv k^*$ and $k^* \geqslant x - z^*_{t-1}$.

$$\begin{aligned}\Delta H_t(x) &= c \quad \text{(by observation 2)}\\ &\geqslant \Delta R(k^*+1) \quad \text{(by Lemma 1, part (ii))}\\ &= \Delta H_t(x+1) \quad \text{(by observation 3).}\end{aligned}$$

Then, $\Delta H_t(x) \geqslant \Delta H_t(x+1)$.

Thus, $\Delta H_t(x)$ is decreasing. Taking expected value over all possible realizations of (n, v) preserves this monotone property, and hence $\Delta V_t(x)$ is decreasing as well, which completes the induction. □

Proof of Lemma 5. We will work on a sample path argument, taking a realization (n, v). Start by noting that $\Delta V_t(x) < \bar{v}\ \forall\, t, x$: To this end, consider observations 1–3 of the proof of Lemma 4. From Observation 1 (for our purposes here, when $k^* = k^*(x-1) = k^*(x)$), we know that

$$\begin{aligned}\Delta V_t(x) &= \alpha\Delta V_{t-1}(x-k^*(x-1)) - \Delta h(x-k^*(x-1))\\ &< \Delta R(k^*(x-1)) \quad \text{(by Lemma 1, part (ii))}\\ &\leqslant J(\bar{v}) \quad \text{(by Equation (6))}\\ &< \bar{v} \quad \text{(by Equation (3)).}\end{aligned}$$

Together with the other two observations, we get

$$\Delta V_t(x) \leqslant \max\{c, \Delta R(k^*)\} < \bar{v}.$$

This is intuitive: the profit from an extra unit in inventory cannot exceed the maximum possible bid.

Now, take (11), and define

$$\bar{z} = \max\{z \in \mathbf{Z}_+\colon \alpha\bar{v} - c > \Delta h(z)\}.$$

Clearly, $\bar{z} \geqslant z^*_t\ \forall\, t$.

Take $\bar{y} \equiv \bar{z}$, and consider the following bounds for the per-period profit:

$$\min\{-h(x), -h(\bar{y}) - c\bar{y}\} \leqslant R(k) - h(x-k+y) - cy \leqslant M\bar{v}$$
$$\forall\, x, k, y\colon 0 \leqslant k \leqslant x, 0 \leqslant y \leqslant \bar{y}.$$

Hence, its modulus is bounded.

Moreover, because $y^*_t \leqslant z^*_t$, then $y^*_t \leqslant \bar{y}$. □

Proof of Proposition 2. Once more, we will argue on a sample path base. We start by defining

$$N(s) \equiv \big|\{v_i\colon v_i > s, \forall\, 1 \leqslant i \leqslant N\}\big|,$$

and let

$$Q(k) = \begin{cases} 0 & \text{if } k = 0,\\ \displaystyle\sum_{i=1}^{\min\{k,N\}} [J(v_{(i)}) - c] & \text{if } k > 0.\end{cases}$$

Note that $Q(k)$ is concave in k, because the increments $\Delta Q(k)$ are decreasing because of the ordering of the virtual values $J(v_{(i)})$.

By letting $\hat{c} = J^{-1}(c)$, we can rewrite (19) as

$$\Pi(z) = E_{N,v}[Q(\min\{z, N(\hat{c})\})] - h(z). \tag{22}$$

The composition $Q(\min\{z, N(\hat{c})\})$ turns out to be concave in z, and jointly with the convexity of $h(\cdot)$, we get the result. □

Proof of Proposition 3. The proof follows the guidelines of the one for an analogous result in the finite horizon problem described in Vulcano et al. (2002, Proposition 2, part (a)). We will show the result by proving the equality $\Pi(z) = \Pi_{\text{LP}}(z)\ \forall\, 0 \leqslant z \leqslant \bar{z}$. For $z = 0$, $\Pi(0) = \Pi_{\text{LP}}(0) = -h(0)$. For $z \geqslant 1$, from (22),

$$\Pi(z) = E_v[(J(v) - c)\mathbf{I}\{v > \hat{c}\}]P(N=1) - h(z) \tag{23}$$

and from (20),

$$\Pi_{\text{LP}}(z) = \max_{s\geqslant 0}\big\{E_v[(s-c)\mathbf{I}\{v > s\}]P(N=1) - h(z)\big\}. \tag{24}$$

For the latter, note that we can express the maximization problem as

$$\max_{s\geqslant 0}\{(s-c)(1-F(s))\}P(N=1) - h(z).$$

The optimal s^* is the solution to the first-order condition

$$s = c + \frac{1-F(s)}{f(s)}.$$

Using the definition of $J(\cdot)$ (see formula (3)), we can alternatively write

$$s^* = J^{-1}(c) = \hat{c}.$$

The list-price problem (24) becomes

$$\Pi_{\text{LP}}(z) = E_v[(\hat{c} - c)\mathbf{I}\{v > \hat{c}\}]P(N=1) - h(z).$$

Now, it will be enough to show that $E_v[J(v)\mathbf{I}\{v>\hat{c}\}]=E_v[\hat{c}\mathbf{I}\{v>\hat{c}\}]$:

$$\begin{aligned} E[J(v)\mathbf{I}\{v>\hat{c}\}] &= \int_{\hat{c}}^{\bar{v}} J(v)f(v)\,dv \\ &= \int_{\hat{c}}^{\bar{v}} vf(v)\,dv - \int_{\hat{c}}^{\bar{v}} (1-F(v))\,dv \\ &= E[v\mathbf{I}\{v>\hat{c}\}] - E[(v-\hat{c})\mathbf{I}\{v>\hat{c}\}] \\ &= E[\hat{c}\mathbf{I}\{v>\hat{c}\}]. \end{aligned} \tag{25}$$

Because both expressions (23) and (24) are equivalent for every z, the optimal values are the same. □

Proof of Proposition 4. Recall that in a second-price auction, the bidders' dominant strategy is to bid their own values independently of the number of items z^n and the number of bidders N^n. Consider the profit as a function of z^n and N^n, and note that for any realization v, the seller will award $x^{z^n,N^n}(v)$ units given by

$$x^{z^n,N^n}(v) \equiv \max\{i \leqslant \min\{z^n,\, N^n\}\colon v_{(i)} > \hat{c}\}$$

and each winner will pay

$$\pi^{z^n,N^n}(v) \equiv \max\{\hat{c},\, v_{(x^{z^n,N^n}(v)+1)}\},$$

for a total single-period profit of

$$X^{z^n,N^n}(v) = x^{z^n,N^n}(v)(\pi^{z^n,N^n}(v)-c) - hz^n - a.$$

Asymptotically, both $(1/n)x^{z^n,N^n}(v)$ and $\pi^{z^n,N^n}(v)$ converge to deterministic limits that depend on z and N. This leads to the following lemma characterizing the limiting, per-period profit (see Vulcano et al. (2002, Proposition 3 in the Appendix) for a proof of a similar result).

Lemma 7. *Let $\bar{F}(x)=1-F(x)$, and assume that the holding cost is linear of the form $h(z)=a+hz$. If $z\leqslant N\bar{F}(\hat{c})$, then as $n\to\infty$,*

$$\frac{1}{n}X^{z^n,N^n}(v) \to z[F^{-1}(1-z/N)-c]-hz \quad a.s.,$$

else if $z>N\bar{F}(\hat{c})$,

$$\frac{1}{n}X^{z^n,N^n}(v) \to N\bar{F}(\hat{c})[\hat{c}-c]-hz \quad a.s.$$

When the number of units z^n and the number of bidders N^n are large and $z^n\leqslant N^n\bar{F}(\hat{c})$, the firm's profit will be given by

$$X^{z^n,N^n}(v) = z^n[F^{-1}(1-z^n/N^n)-c]-hz^n+o(n),$$

which is independent of the particular instance of v. Dividing by n and letting $n\to\infty$, the limiting dynamic control problem can be written as

$$\Pi^{(a)} = \max_{0\leqslant z\leqslant N\bar{F}(\hat{c})}\{z[F^{-1}(1-z/N)-c]-hz\}. \tag{26}$$

Note that because the number of units to award is $\min\{z, N\bar{F}(\hat{c})\}$ but the holding cost is proportional to z, it is never optimal to use $z>N\bar{F}(\hat{c})$.

Solving the limiting problem (26) provides a simple approximation of the optimal basestock level. For example, when bidders' valuations are uniformly distributed, this average profit function reduces to a concave quadratic maximization problem that is readily solvable. Numerical results show that the approximate basestock level is quite accurate when N and z are large (see §4.2.1).

To see that this scaled auction problem is equivalent to the list-price policy, consider a similar sequence of problems for the list-price case. Let $x_{\mathrm{LP}}^{z^n,N^n}(v,s)$ be the number of units that the firm awards when facing an instance v for a given list price s. Then,

$$x_{\mathrm{LP}}^{z^n,N^n}(v,s) \equiv \max\{i\leqslant\min\{z^n,N^n\}\colon v_{(i)}>s\}$$

for a total single-period profit of

$$X_{\mathrm{LP}}^{z^n,N^n}(v,s) = (s-c)x_{\mathrm{LP}}^{z^n,N^n}(v,s)-hz^n-a.$$

Asymptotically, $(1/n)x_{\mathrm{LP}}^{z^n,N^n}(v,s)$ reaches a deterministic limit that depends on z, N, and s. An identical argument to Lemma 7 shows that for $z\leqslant N\bar{F}(s)$, as $n\to\infty$,

$$\frac{1}{n}X_{\mathrm{LP}}^{z^n,N^n}(v,s)\to z(s-c)-hz \quad \text{a.s.};$$

and for $z>N\bar{F}(s)$,

$$\frac{1}{n}X_{\mathrm{LP}}^{z^n,N^n}(v,s)\to N\bar{F}(s)(s-c)-hz \quad \text{a.s.}$$

In other words, when the number of units z^n and the number of buyers N^n become large, and $z^n\leqslant N^n\bar{F}(s)$ (note that it is enough to focus on this case), the seller's profit becomes

$$X_{\mathrm{LP}}^{z^n,N^n}(v,s)=z^n(s-c)-hz^n+o(n).$$

So, take a particular z^n in that range, and observe that

$$z^n\leqslant N^n\bar{F}(s)\Leftrightarrow s\leqslant F^{-1}\left(1-\frac{z^n}{N^n}\right).$$

Hence, maximizing over s yields

$$s^* = F^{-1}\left(1-\frac{z^n}{N^n}\right).$$

The limiting dynamic control problem in this case becomes

$$\Pi_{\mathrm{LP}}^{(a)} = \max_{z\geqslant 0}\{z[F^{-1}(1-z/N)-c]-hz\}.$$

It is again sufficient to focus on maximizing over $0\leqslant z\leqslant N\bar{F}(\hat{c})$, which yields the same limiting expression as (26). Thus, we have

$$\Pi_{\mathrm{LP}}^{(a)} = \Pi^{(a)},$$

which concludes the proof. □

Proof of Proposition 5. We will prove that if $h(z) = 0$, then $\Pi^*_{\text{LP}} = \Pi^*$. For the auction case, consider the benefit for a particular inventory level z:

$$\Pi(z) = E_{N,v}\left[\max_{0\leqslant k\leqslant \min\{z,N\}}\{R(k) - ck\}\right].$$

We can then take $z^* = M$, the maximum possible value for the random variable N. Then,

$$\begin{aligned}
\Pi^* &= E_{N,v}\left[\max_{0\leqslant k\leqslant N}\{R(k) - ck\}\right] \\
&= E_{N,v}\left[\max_{0\leqslant k\leqslant N}\sum_{i=1}^{k}[J(v_{(i)}) - c]\right] \\
&= E_{N,v}\left[\sum_{i=1}^{N}(J(v_i) - c)\mathbf{I}\{v_i > \hat{c}\}\right] \\
&= E[N](E_v[J(v)\mathbf{I}[v > \hat{c}]] - cE_v[\mathbf{I}\{v > \hat{c}\}]) \\
&\qquad \text{(because } N \text{ and } v \text{ are independent)} \\
&= E[N](E_v[\hat{c}\mathbf{I}\{v > \hat{c}\}] - cE_v[\mathbf{I}\{v > \hat{c}\}]) \\
&\qquad \text{(by (25) in the proof of Proposition 3)} \\
&= E[N](\hat{c} - c)E_v[\mathbf{I}\{v > \hat{c}\}]. \qquad (27)
\end{aligned}$$

For the list-price case,

$$\Pi_{\text{LP}}(z) = \max_{s\geqslant 0}\left[(s - c)E_{N,v}[\min\{N(s), z\}]\right].$$

As in the auction case, we can argue that $z^*_{\text{LP}} = M$. Thus, we get

$$\begin{aligned}
\Pi^*_{\text{LP}} &= \max_{s\geqslant 0}\{(s - c)E_{N,v}[\min\{N(s), z^*_{\text{LP}}\}]\} \\
&= \max_{s\geqslant 0}\left\{(s - c)E_{N,v}[N(s)]\right\} \\
&= \max_{s\geqslant 0}\left\{(s - c)E_{N,v}\left[\sum_{i=1}^{N}\mathbf{I}\{v_i > s\}\right]\right\} \\
&= \max_{s\geqslant 0}\{(s - c)E[N]E_v[\mathbf{I}\{v_i > s\}]\} \\
&\qquad \text{(because } N \text{ and } v \text{ are independent).}
\end{aligned}$$

Now, because Π^* dominates Π^*_{LP}, then from the expression in (27), it turns out that $s^* = \hat{c}$, and hence, $\Pi^* = \Pi^*_{\text{LP}}$. □

Endnotes

1. Of course, list pricing has many other advantages, including familiarity, ease of use, lack of delays, reduced risk for buyers, etc.
2. In particular, distributions that have increasing hazard rate include the uniform, normal, logistic, exponential, and extreme value (double exponential) distributions, etc. (See Bagnoli and Bergstrom 1989.)
3. Though this reasoning is admittedly circular, if buyers did indeed strategize over timing, the optimal policy might attempt to exploit this fact rather than mitigate it as our policy does.
4. This upper bound $\bar{z}$ only makes sense when the function $h(\cdot)$ is strictly convex. Otherwise, we can set $\bar{z} = M$: it could not be optimal to stock units beyond the maximum number of bidders that could eventually show up in one particular period.
5. See Vakrat and Seidmann (1999) for an experimental study comparing prices paid through online auctions and catalogs.
6. There is also another interesting case, in which a single stream of customers chooses which of the two channels to use based on maximizing their own utility. This question raises more complex behavioral issues that are beyond the scope of this paper, though worth further research.
7. A stationary policy is unichain if its associated Markov chain has a single recurrent class and a possibly empty set of transient states.
8. Vakrat and Seidmann (2000) build a theoretical model based on transaction data from 324 business-to-consumer online auctions. In their Figures 6.a and 6.b, the optimal number of units put through the auction is roughly speaking half of the mean demand.

Acknowledgment

The authors thank the Deming Center at the Columbia Business School for providing partial funding of this research.

References

Amihud, Y., H. Mendelson. 1983. Price smoothing and inventory. *Rev. Econom. Stud.* **50** 87–98.

Bagnoli, M., T. Bergstrom. 1989. Log-concave probability and its applications. Working paper 89-23, University of Michigan, Ann Arbor, MI.

Bertsekas, D. 1995. *Dynamic Programming and Optimal Control*, Vol. 2, 1st ed. Athena Scientific, Belmont, MA.

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3) 454–475.

Harsanyi, J. 1967. Games with incomplete information played by Bayesian players, I. *Management Sci.* **14** 159–182.

Harsanyi, J. 1968a. Games with incomplete information played by Bayesian players, II. *Management Sci.* **14** 320–334.

Harsanyi, J. 1968b. Games with incomplete information played by Bayesian players, III. *Management Sci.* **14** 486–503.

Harstad, R., J. Kagel, D. Levin. 1990. Equilibrium bid functions for auctions with an uncertain number of bidders. *Econom. Lett.* **33** 35–40.

Klemperer, P. 1999. Auction theory: A guide to the literature. *J. Econom. Surveys* **13**(3) 227–286.

Lengwiler, Y. 1999. The multiple unit auction with variable supply. *Econom. Theory* **14** 373–392.

Li, L. 1998. A stochastic theory of the firm. *Math. Oper. Res.* **13** 447–466.

Lucking-Reiley, D. 2000. Vickrey auctions in practice: From nineteenth century philately to twenty-first century e-commerce. *J. Econom. Perspectives* **14**(3) 183–192.

Maskin, E., J. Riley. 2000. Optimal multi-unit auctions. P. Klemperer, ed. *The Economic Theory of Auctions*, Vol. 2. E. Elgar Publishing, U.K., 312–336.

Matthews, S. 1995. A technical primer on auction theory I: Independent private values. Discussion paper no. 1096, CMS-EMS, Northwestern University, Evanston, IL.

McAfee, R., J. McMillan. 1987a. Auctions and bidding. *J. Econom. Literature* **25**(2) 699–738.

McAfee, R., J. McMillan. 1987b. Auctions with a stochastic number of bidders. *J. Econom. Theory* **43** 1–19.

Milgrom, P. 1989. Auctions and bidding: A primer. *J. Econom. Perspectives* **3**(3) 3–22.

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

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1) 58–73.

Nash, J. 1951. Non-cooperative games. *Ann. Math.* **54** 286–295.

Pinker, E., A. Seidmann, Y. Vakrat. 2000. Using transaction data for the design of sequential, multi-unit, online auctions. Working paper CIS-00-03, William E. Simon Graduate School of Business Administration, University of Rochester, Rochester, NY.

Riley, J., W. Samuelson. 1981. Optimal auctions. *Amer. Econom. Rev.* **71**(3) 381–392.

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

Rothkopf, M., T. Teisberg, E. Kahn. 1990. Why are Vikrey auctions rare? *J. Political Econom.* **98**(1) 94–109.

Segev, A., C. Beam, J. Shanthikumar. 2001. Optimal design of Internet-based auctions. *Inform. Tech. Management* **2**(2) 121–163.

Thomas, L. 1974. Price and production decisions with random demand. *Oper. Res.* **22** 513–518.

Thowsen, G. T. 1975. A dynamic, nonstationary inventory problem for a price/quantity setting firm. *Naval Res. Logist. Quart.* **22** 461–476.

Vakrat, Y., A. Seidmann. 1999. Can online auctions beat online catalogs? *Proc. 20th Internat. Conf. Information Systems*, Charlotte, NC.

Vakrat, Y., A. Seidmann. 2000. Implications of the bidder's arrival process on the design of online auctions. *Proc. 33rd Hawaii Internat. Conf. System Sci.*, Maui, HI.

van Ryzin, G. 2000. The brave new world of pricing. In Survey–Mastering Management. *Financial Times* (October 16) 6.

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

Vulcano, G., G. J. van Ryzin, C. Maglaras. 2002. Optimal dynamic auctions for revenue management. *Management Sci.* **48**(11) 1388–1407.

Wolfstetter, E. 1996. Auctions: An introduction. *J. Econom. Surveys* **10**(4) 367–420.

Zabel, E. 1972. Multiperiod monopoly under uncertainty. *J. Econom. Theory* **5** 524–536.