MANAGEMENT SCIENCE
Vol. 53, No. 5, May 2007, pp. 795–813
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5305 | 0795

DOI 10.1287/mnsc.1060.0646

Online Auction and List Price Revenue Management

René Caldentey, Gustavo Vulcano
Stern School of Business, New York University, 44 West Fourth Street, New York, New York 10012
{rcaldent@stern.nyu.edu, gvulcano@stern.nyu.edu}

We analyze a revenue management problem in which a seller facing a Poisson arrival stream of consumers operates an online multiunit auction. Consumers can get the product from an alternative list price channel. We consider two variants of this problem: In the first variant, the list price is an external channel run by another firm. In the second one, the seller manages both the auction and the list price channels.

Each consumer, trying to maximize his own surplus, must decide either to buy at the posted price and get the item at no risk, or to join the auction and wait until its end, when the winners are revealed and the auction price is disclosed.

Our approach consists of two parts. First, we study structural properties of the problem, and show that the equilibrium strategy for both versions of this game is of the threshold type, meaning that a consumer will join the auction only if his arrival time is above a function of his own valuation. This consumer's strategy can be computed using an iterative algorithm in a function space, provably convergent under some conditions. Unfortunately, this procedure is computationally intensive.

Second, and to overcome this limitation, we formulate an asymptotic version of the problem, in which the demand rate and the initial number of units grow proportionally large. We obtain a simple closed-form expression for the equilibrium strategy in this regime, which is then used as an approximate solution to the original problem. Numerical computations show that this heuristic is very accurate. The asymptotic solution culminates in simple and precise recipes of how bidders should behave, as well as how the seller should structure the auction, and price the product in the dual-channel case.

Key words: revenue management; online auction; dual channel; strategic behavior; asymptotic analysis
History: Accepted by Paul H. Zipkin, operations and supply chain management; received February 19, 2004. This paper was with the authors $9\frac{1}{2}$ months for 2 revisions.

1. Introduction

In the last few years, Revenue Management (RM) has widened its focus from capacity control and dynamic pricing to alternative selling mechanisms proposed by electronic commerce, such as group purchasing, online negotiations, and auctions. (See Talluri and van Ryzin 2004 for a reference on RM methods and applications, or the survey by Bitran and Caldentey 2003 for an overview of dynamic pricing models.) Although list pricing is probably still the most familiar and used pricing mechanism, online auctions are certainly an increasing phenomenon.

Nowadays, a huge variety of products is sold simultaneously through online posted price and auction channels, allowing consumers to compare prices and bid states easily across different channels in real time. This boost in market information and the corresponding reduction in search costs have a significant impact on consumers' purchasing behavior and should be considered by a seller when designing online sales mechanisms.

In this paper, we address the problem of an online seller who is endowed with a fixed initial inventory and faces a stochastic arriving stream of strategic consumers. To capture different e-business environments, we consider two alternative formulations. First, we consider the case in which the seller controls an online auction exclusively and competes with a third party that manages a list price channel. We refer to this case as the *single-auction channel* model. In the second formulation, the *dual-channel* model, the seller is a monopolist and controls both the auction and list price channels.

For these two scenarios, we are interested in answering the following questions: How should strategic consumers behave, that is, which channel should each consumer choose? What should the bidding strategy be of those consumers that enter the auction? Given consumers' behavior, how should the seller manage an online auction to maximize revenue? What should the length of the auction and the reservation price (i.e., minimum acceptable bid) be? How should the seller manage parallel online auction and list price channels to maximize revenue? In particular, the dual-channel case motivates an important managerial question: How should the seller design both channels to segment the population of consumers to extract as much revenue as possible from each segment? A seller does not want to offer a business model that cannibalizes itself, that is, if she offers a high posted price, she narrows the list price channel, and middle-to-high valuation consumers are tempted to join the auction, which could eventually close at a low price. On the other hand, if she posts a low list price, she widens the list price channel, pooling together low- and high-valuation consumers. In both cases, the seller runs the risk of decreasing revenues. Our purpose is to shed some light on the trade-offs that are inherent to this business environment.

More formally, we analyze a single-period model in which a seller operates a multiunit, uniform price, online auction, offering multiple units of a homogeneous good. The seller announces the inventory put up for sale Q_0, the auction duration T, and the auction reservation price v^R. Consumers with single-unit demand arrive according to a Poisson process. They have a private value for the product, independently drawn from a continuous distribution. They must decide whether to bid and wait for the auction outcome at time T, or buy at the posted price $\hat{P}$, and get the unit instantaneously. As mentioned above, we consider two variants of this problem:

- In the single-auction channel case, the fixed-price channel is external and run by another firm, which we assume has unlimited inventory. Hence, if a bidder is not among the winners of the auction, he can always buy the item at the posted price at a later time T, although his utility is discounted.
- In the dual-auction and list price channel case, the seller is a monopolist who manages both the auction and the list price channels. In this case, an auction takes place at time T only if there are units left unsold by then. Hence, supply is limited, and bidders that lose in the auction have no alternative market in which to buy the product.

The consumers' strategy consists of two decisions: (i) whether or not to join the auction, and (ii) in the case of joining the auction, what bidding strategy to use. The supply size under both scenarios produces different bidding behaviors. Regarding the first decision, we prove that a symmetric equilibrium strategy exists in both variations of the problem, and it is characterized by a threshold function in the space (valuation, time): For a consumer arriving at time t with valuation v, there is a threshold $H(v)$ such that if $H(v) \leq t$, then he will participate in the auction. Otherwise, he will buy at the posted price $\hat{P}$. This participation strategy can be computed using an iterative algorithm (in an appropriate function space) provably convergent under some special conditions. Unfortunately, this procedure is computationally intensive and does not lead to simple managerial insights. To overcome this limitation, we formulate an asymptotic version of the problem, in which the demand rate and the initial number of units grow proportionally large. We get a simple closed-form expression for the equilibrium strategy in this limiting regime, which is then used as an approximated solution for the original problem. Numerical computations show that this heuristic is very accurate.

Finally, we analyze the seller's optimization problem in both the single- and dual-channel settings, and plug the consumer's asymptotic participation strategy into them to compute the optimal values of the parameters Q_0, T, v^R (and $\hat{P}$ in the dual channel case), that the seller must announce to maximize revenues. We can then assert that the asymptotic solution culminates in precise and simple guidelines for how bidders should behave and how the seller should design the auction and list price channels.

The main insights that we obtained are the following. For the single-auction channel case, we find that the optimal number of units to offer is a nonmonotonic function of the external fixed price, $\hat{P}$, and that it is bounded above by 80% of the average demand. In addition, the optimal duration of the auction is an increasing function of $\hat{P}$. In the dual auction and list price channel case, we find that if the seller's initial inventory endowment is small or if her discount factor is large, then she does not have enough economic incentives to run a terminal auction; a single fixed-price channel is the optimal selling mechanism. On the other hand, if the endowment is large, or the seller discount factor is small, or buyers are impatient, then running both channels in parallel is optimal. In any of these cases, we show that a dual-channel operation can have a significant impact on revenues compared to a single fixed-price channel. The magnitude of the increase in revenues can be as large as 33% for the case of uniformly distributed valuations.

1.1. Literature Review

Auctions have been extensively studied in the economic literature (e.g., see the survey by Klemperer 1999 or the recent book by Krishna 2002). Price discrimination has been argued as one of the main reasons for using them (see Bulow and Roberts 1989). Maskin and Riley (1989) proved the optimality of the uniform price mechanism for the single-period multiunit auction.

Few papers have put auctions in an operational perspective. Specifically, regarding its connection with RM, Pinker et al. (2000) study how to run a sequence of standard multiunit auctions, using bidding information to learn about the consumer's valuation distribution. Vulcano et al. (2002) characterize an optimal dynamic auction for a firm selling a fixed capacity over a finite horizon.

The firm's choice between auctions and posted prices for the single-channel case has also been addressed (e.g., see Vany 1987, Wang 1993, Harstad 1990).

The problem of jointly managing auction and list price channels has not received much attention in the literature. New features like the *buy now* prices have been addressed by Budish and Takeyama (2001), although their model is limited to two bidders and two valuation types. Within the business-to-consumer (B2C) framework, the empirical study of Vakrat and Seidmann (1999) compares prices paid through online auctions and catalogs for the same product. They observe that auctions result in average prices 25% below the catalog ones. They build a simple model of single-unit auctions with a deterministic number of bidders, but ignoring consumer choice behavior. In the infinite-horizon model of van Ryzin and Vulcano (2004, §3.3), the seller operates auctions and posted prices simultaneously, and replenishes her stock in every period. However, the streams of consumers for both channels are independent, and the seller decides how many units to allocate to each of the channels separately.

Our research is mainly motivated by the work of Etzion et al. (2006). They analyze simultaneous online auctions and list price channels in a multiperiod B2C framework, where a seller with infinite supply maximizes her average expected revenue. Consumers arrive according to a Poisson process, and decide which channel to join. They found two optimal auction design strategies: short single-unit auctions and long multiunit auctions.

Our work differs from theirs in the way we model the supply side, because in our dual channel case scarcity plays a critical role, as is usually the case in RM: Given the risk that potentially no item could remain available for the auction by time T (which occurs when all the inventory is depleted through the list price channel), what should the consumer's participation strategy be? In the case of going for the auction, scarcity induces the standard dominant "bid your own value" strategy for multiunit uniform price auctions (see Krishna 2002, §2.2 for a comprehensive study of bidding behavior). The situation is different in the single-auction channel case, where the infinite supply (as is the case in Etzion et al. 2006) induces no bidder to bid higher than the posted price: In case he loses, he always has the chance to pay that price at time T. Now, given both settings, how should the seller structure the business, accounting for consumers' strategic behavior?

Etzion et al. (2006) work with additive consumer utility functions, and characterize the consumer equilibrium bidding strategy with a single value $\bar{t}$, such that all consumers with valuation below the posted price, and those arriving later than the threshold $\bar{t}$ with valuation above the posted price, will join the auction. All other consumers will go to the online catalog. Our equilibrium participation strategy turns out to be more complex because it is based on a multiplicative utility function, and as we said above, is defined by a continuous threshold function in the space (valuation, time). Furthermore, when computing the participation strategy in their paper, Etzion et al. assume that the total number of competing consumers is deterministic. We instead embed the random nature of the arrival process in the computation of the consumer's participation strategy.[1]

A distinguishing characteristic of our research is the asymptotic analysis of the game. We show that the complex threshold that describes the strategic behavior of the consumers can be easily computed in the limiting regime where the consumers' arrival rate and the number of units offered grow proportionally large, without missing the predictive power of the model.

Overall, we believe that the two models share some features, but contrast in these important dimensions, which are worth exploring.

Finally, the problem of analyzing the equilibrium of a system where consumers arrive during a time window has been addressed by few papers, but they are oriented to the characterization of the arrival pattern (e.g., Glazer and Hassin 1983 or Lariviere and Van Mieghem 2004). In our setting, the arrival process is exogenous, and we concentrate on characterizing the Nash equilibrium (in pure strategies) of the participation behavior of the consumers.

The remainder of this paper is organized as follows. We introduce the model for both variants of the problem in §2. In §§3 and 4, we study the consumers' problem of selecting an optimal participation and bidding strategy for both the single-auction and dual-channel models, respectively. We prove the existence of a symmetric equilibrium within a large class of participation strategies and use asymptotic analysis to characterize this equilibrium. The analysis in §§3 and 4 assumes that the initial inventory Q_0, the auction bidding period T, the reservation price v^R, and the posted price $\hat{P}$ are fixed. We turn to the seller's revenue maximization problem of optimally choosing Q_0, v^R, T, and $\hat{P}$ in §5. Finally, §6 summarizes our concluding remarks.

The paper includes two online appendices (provided in the e-companion).[2] Appendix A contains all

[1] In the final version of their paper, however, which was later than our first draft, they included a section in the appendix where they discussed the stochastic number of bidders assumption.

[2] An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

the proofs. Appendix B provides a detailed mathematical description of the dual-channel model.

2. Model Description

We study the problem faced by a firm (seller) endowed with an initial inventory $\bar{Q}$ of a homogeneous product. We take an RM point of view and assume that the seller cannot replenish her inventory throughout the selling season. We do, however, allow the seller to ration by choosing the quantity $Q_0 \leq \bar{Q}$ to put up for sale. The remaining quantity $\bar{Q} - Q_0$ is discarded at no extra cost or salvage value.

We discuss two variants of this problem. In the first variant, the seller manages a single-auction channel, and there is an external market with infinite supply where the same product is available at a fixed price $\hat{P}$. In the second one, the seller is a monopolist managing both the auction and the list price channels simultaneously. During the selling process, she allocates units through the list price channel. Therefore, for the seller at time $t = 0$ and for any arriving consumer, the effective number of units to be auctioned at time T is described by a random variable Q_T, with support $\{0, \ldots, Q_0\}$.

We also assume that the seller announces on her website the time left for the auction completion, the inventory $Q_0 \leq \bar{Q}$ that she decided to put up for sale, the reservation price v^{R}, and in the second variant, the posted price $\hat{P}$ as well. In particular, we assume that there is no update of the inventory level or state of the bidding process over time. For both settings, we model the auction as a sealed-bid multiunit uniform-price (also referred to as $(Q_T + 1)$-price) auction. In a $(Q_T + 1)$-price auction, the winners are the Q_T highest bidders in excess of the reservation price v^{R}, and all of them pay the value of the maximum between the $(Q_T + 1)$th highest bid and v^{R}.

On the demand side, consumers have single-unit requests, and visit the seller's website following a Poisson process with constant intensity λ. They are characterized by two quantities: (i) their arrival time, and (ii) their private valuation for the product. For notational convenience, we denote the private valuation of a consumer arriving at time t by v_t. Observe that this notation is well defined because, with probability one, the Poisson process has at most one arrival at any given time.[3] We also assume that the cumulative probability distribution F of the random variable v_t is a time-homogeneous and differentiable function with support $\mathcal{V} \triangleq [0, \bar{v}]$ that admits a density function $f(v)$.[4] Both λ and F are common knowledge. Without loss of generality, we assume from now on that $\bar{v} = 1$, that is, we scale all prices in this economy by $\bar{v}$.

When visiting the website, consumers must choose either to bid or to buy the product at a posted price $\hat{P}$ to maximize their own surplus. We assume that they are sensitive to delay, and denote by $u(t, \tau, v - p)$ the quasilinear discounted utility function of a consumer arriving at time t with valuation v who eventually gets a unit of product at time τ at a price p (paid at the moment of getting it). If the consumer never gets the object, we use the convention $\tau = \infty$. In particular, we consider an exponentially discounted utility function of the form:

$$u(t, \tau, v - p) = (v - p)\exp(-w(\tau - t)), \qquad (1)$$

where w is a fixed constant shared by all consumers that captures the consumers' disutility for waiting. As a side remark, we note that our main theoretical results are not especially tight to the functional form of the utility in (1) as long as it remains increasing in $v - p$ and decreasing in delay $\tau - t$.[5]

We assume that a consumer arriving at $t \in [0, T)$ chooses at this time whether to enter the auction or to buy at the fixed price. He bases this decision on his private valuation, his knowledge of the arrival rate λ and the distribution of valuations F, the initial inventory Q_0, the auction remaining time $T - t$, the reservation price v^{R}, and the posted price $\hat{P}$. This participation strategy can be characterized by a threshold function $H(\cdot)$ such that a consumer with valuation v_t enters the auction only if his arrival time t exceeds $H(v_t)$. Pictorially, $H(v)$ divides consumers' type space (valuation, arrival time) into two regions—as shown in Figure 1—one corresponding to the posted-price buyers and the other to the auction bidders.

The seller's problem is to design a single-auction channel in the first model (by setting a value for T, Q_0, and v^{R}), and a dual auction and list price channel in the second model (by setting also a value $\hat{P}$) to maximize her expected revenue, which is also exponentially discounted over time.

2.1. Discussion of the Model

On a theoretical level, our model is in many ways a variation of the classical, single-period, private-value

[3] Note that the subscript in v_t only refers to a customer index. By assumption, in our model the distribution of valuations is time independent.

[4] Note that setting the lower bound of the valuation support (say $\underline{v}$) to zero is w.l.o.g. In fact, suppose $\underline{v} < 0$. Then, because a revenue-maximizing seller will always choose $v^{\mathrm{R}} \geq 0$ and $\hat{P} \geq 0$, we can simply truncate the distribution of valuations to start at zero without affecting the outcome of the game. On the other hand, if $\underline{v} > 0$, then at optimality $v^{\mathrm{R}} > \underline{v}$ and $\hat{P} > \underline{v}$, and we can scale down all prices in the economy (subtracting $\underline{v}$) so that after scaling $\underline{v} = 0$.

[5] Note that the consumer's utility function is of the intertemporal type (e.g., see Mas-Colell et al. 1995, Chapter 20).

Figure 1 An Example of a Participation Strategy $H(v)$ for the Case $T = 1$ and $v^R = 0$

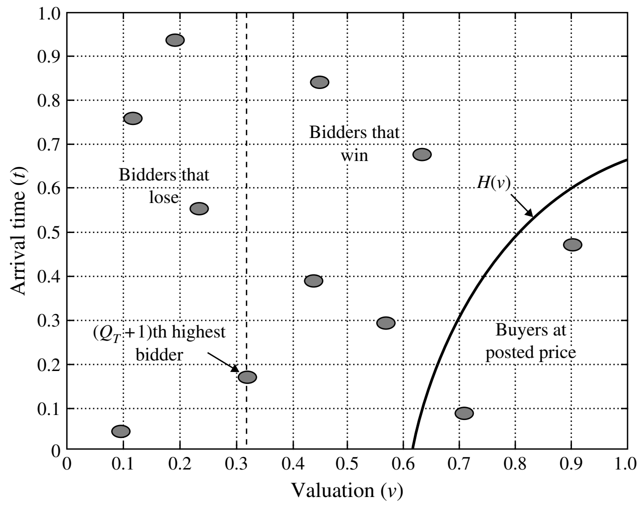

Note. The dots represent a sample path of the arrival process. In this case, $Q_0 = 6$. Two consumers buy at the posted price; and nine submit bids. There are 4 winning bids, and 5 losing bids. The vertical dashed line represents the auction price for this realization of demand.

auction model, but with two main distinctions: a stochastic arrival process embedded in it, and the coexistence of the alternative list price channel.

On a practical level, our model shares some of the limitations of the auction literature, such as the assumptions that bidders are symmetric or that the distribution F and the arrival rate λ are common knowledge (see Pinker et al. 2003 for further discussion). On the other hand, our model does capture important features of the real game. For instance, our model is able to explain the surge in bids close to the end of an auction with a sharp deadline, a phenomenon commonly observed in practice (e.g., Roth and Ockenfels 2002). Our result in Proposition 3 supports this fact (see also Figure 1), showing that the proportion of consumers that go to the auction is increasing in the elapsed time t.

The fact that when consumers arrive at the website they are only informed about the initial number of units Q_0 is reasonable in the single-auction channel, but it is certainly a limitation of the dual-channel model. This contrasts with the business practice of posting *buy now* prices while running an auction (e.g., eBay.com): Consumers are informed about the number of remaining units Q_t when visiting the website at time t. This assumption is made for mathematical tractability, because including this factor in the consumer choice behavior would add an extra dimension to the current space (valuation, time) for the participation decision.[6] Nevertheless, we think that the kind of arguments and technical tools used in this paper can be applied to this extension as well. In particular, we show that in an asymptotic sense (that we make precise in §§3 and 4), our proposed solution under partial information is also an equilibrium for the game in which the seller provides full information about the inventory level and bidding process over time. Hence, although restrictive, our model can be viewed as an asymptotically exact first-step approximation of the more complex game with full information.

Because there is no information update on the remaining number of units, consumers have no incentive to wait for a time window between arrival at the website and the final participation decision; otherwise, they would just be incurring a discount for waiting.

We model the auction as a multiunit uniform-price sealed-bid auction. In standard single-period auction theory, this auction format turns out to be strategically equivalent to the English auction[7] (e.g., see Krishna 2002). Actually, the English auction is by far the most common online auction format (see Pinker et al. 2003). Some empirical research supports the theoretical equivalence between sealed-bid uniform-price and open English auctions in the online practice (e.g., Lucking-Reiley 1999).

3. Single-Auction Channel

In this section, we study consumers' optimal purchasing strategy (i.e., channel selection and bidding) for the single-auction model. We assume that the external list price $\hat{P}$ has been posted and the seller (i.e., the auctioneer in this case) has already announced the parameters of the auction (Q_0, T, and v^R). The auctioneer's problem of optimally designing the auction is postponed to §5.

Let us start discussing the optimal bidding strategy. In our setting, where the sealed $(Q_T + 1)$-price auction[8] operates in parallel to an infinite supply fixed-price market, it is a weakly dominant strategy for bidders with valuations below the posted price to bid their true values, and it is dominant for bidders above the posted price to bid the posted price. In other words, if a consumer with valuation v decides to enter the auction, then he will bid $b(v)$, where

$$b(v) = \begin{cases} v & \text{if } v < \hat{P} \\ \hat{P} & \text{if } v \geq \hat{P} \end{cases} = \min\{v, \hat{P}\}. \tag{2}$$

[6] Gupta and Gallien (2007) provide an online auction model where real-time information of Q_t is revealed, but they need to restrict the analysis to single-unit auctions.

[7] In a multiunit English auction, bids are open for all to see. As in the multiunit uniform-price sealed-bid auction, if Q units are auctioned, the Q highest bidders win, and all the winners pay the highest losing bid.

[8] In this single-auction channel, $Q_T = Q_0$ because there are no items sold during the time window $[0, T)$.

The optimality of $b(v) = v$ for a low-valuation bidder with $v < \hat{P}$ follows from noticing that for this consumer the auction is the only profitable channel from which he can get the object. In addition, it is well known that for the $(Q_T + 1)$-price auction mechanism, bidding the true valuation is a dominant strategy, i.e., the strategy $b(v) = v$ maximizes the low-valuation bidder's expected utility independently of other bidders' strategies. For high-valuation consumers with $v \geq \hat{P}$, the posted-price channel is also an alternative to consider. Furthermore, and contrary to the auction channel, the fixed-price channel has unlimited supply, and bidders with $v \geq \hat{P}$ know that they can always get a unit at this fixed price $\hat{P}$. Therefore, these high-valuation consumers will never bid above the posted price, that is, $b(v) \leq \hat{P}$. In addition, a bidding strategy with $b(v) < \hat{P}$ is also suboptimal because under the $(Q_T + 1)$-price auction mechanism this strategy reduces the bidder probability of winning the auction (with respect to the strategy $b(v) = \hat{P}$) and at the same time does not affect the auction price in the case the bidder actually wins the auction. Thus, the high-valuation bidder is better off choosing $b(v) = \hat{P}$.

Interestingly, if we compare the bidding strategy $b(v) = \min\{v, \hat{P}\}$ of this model and the traditional bidding strategy $b(v) = v$ of the standard $(Q_T + 1)$-price auction, we conclude that, from the auctioneer's point of view, the presence of an uncapacitated fixed-price channel is equivalent to collapsing the range of valuations $[\hat{P}, \bar{v}]$ into a single value $\hat{P}$. For further details about the optimality of $b(v) = \min\{v, \hat{P}\}$, we refer the reader to Etzion et al. (2006, Lemma 1).

Equation (2) characterizes the dominant bidding strategy for all bidders. The probability distribution of the number of bidders and their valuations is the only piece of information that we need to fully characterize the output of the auction. We address this issue in the remainder of this section. Because v^{R} is kept constant for the rest of this section, we set it equal to zero ($v^{\mathrm{R}} = 0$), and rescale the bidder's valuations (and the corresponding probability distribution) accordingly. That is,

$$v \leftarrow \frac{v - v^{\mathrm{R}}}{1 - v^{\mathrm{R}}}, \qquad \hat{P} \leftarrow \frac{\hat{P} - v^{\mathrm{R}}}{1 - v^{\mathrm{R}}},$$
$$F(v) \leftarrow \frac{F(v) - F(v^{\mathrm{R}} + v(1 - v^{\mathrm{R}}))}{1 - F(v^{\mathrm{R}})}, \qquad (3)$$
$$\lambda \leftarrow \lambda(1 - F(v^{\mathrm{R}})).$$

This transformation achieves two objectives. First, only those consumers with valuation greater than or equal to the auction reservation price are considered, the others are discarded. This is true without loss of generality because discarded consumers have no real impact on the auction output. The second objective is that under this scaling, the range of valuations of the (nondiscarded) consumers remains $[0, 1]$. The corresponding scaling of the posted price $\hat{P}$, distribution F, and arrival rate λ follows from these two conditions.

3.1. Participation Strategy and Auction Price Probability Distribution

Because information about the number of units Q_0 to auction, as well as the probability distribution of consumer valuations F, are common knowledge, we can characterize the decision of a consumer arriving at time t, with private valuation v, by a threshold function $H(v)$ such that the consumer will place a bid if and only if $t \geq H(v)$. The fact that we can represent the participation strategy for all v-consumers (i.e., those consumers with valuation v) by a single threshold $H(v)$ is a consequence of the monotonicity of the utility function in the waiting time. In other words, if it is optimal for a v-consumer arriving at time t to wait $(T - t)$ time units for the auction, then it is also optimal for any other v-consumer arriving after t.

Two assumptions are used in this representation of the participation strategy. First, this characterization is based on the notion of a symmetric equilibrium in which all consumers use the same threshold function $H(v)$. In addition, we assume that a consumer arriving at time t is incapable of observing the number of bids already in the system. That is, we assume that the only information that a consumer uses to decide whether or not to enter the auction—besides λ, T, Q_0, and F—are his arrival time and private valuation.

We will denote by $\mathcal{H}$ the set of participation thresholds from which consumers choose their strategies. To keep our formulation reasonably simple, avoiding measure theory technicalities, we restrict consumers to the use of "well-behaved" participation strategies. Specifically, we assume that $\mathcal{H} \subseteq \mathcal{D}$, the set of piecewise continuous functions with right and left limits. Although a restriction, we believe this set $\mathcal{D}$ is large enough to include most strategies that are reasonable from a practical standpoint.[9] In Proposition 1, we will show that the set $\mathcal{D}$ is larger than necessary in the sense that in equilibrium any symmetric participation strategy $\mathcal{H} \in \mathcal{D}$ is actually continuous. Note that by our definition, the elements of $\mathcal{H}$ are functions taking values in $[0, T]$. Furthermore, for any $H \in \mathcal{H}$ and any valuation $v \in [0, \hat{P}]$, we must have $H(v) = 0$ and $b(v) = v$. This reflects the fact that any consumer with valuation in this range cannot afford to buy the product in the external market; the auction is his only

[9] For instance, it includes the set of càdlàg functions; those that are right continuous and have left limit. The problem of characterizing other types of equilibria that are not in $\mathcal{D}$ is certainly interesting from a theoretical perspective, but goes beyond the scope of this paper.

potentially profitable channel, no matter his arrival time. In summary, we define the set of potential bidding strategies as the set of functions

$$\mathscr{H} = \{H \in \mathscr{D},\ H\colon [0,1] \to [0,T],\ \text{such that } H(v) = 0 \text{ for all } v \in [0, \hat{P}]\}.$$

We still need to characterize what the equilibrium $H(v) \in \mathscr{H}$ looks like for $v \in (\hat{P}, 1]$. Consumers in this range will participate in the auction (and bid $b(v) = \hat{P}$) only if the expected auction price is small enough to compensate for the disutility associated with waiting for the auction closing time T.

For any $H \in \mathscr{H}$, let us define a useful random variable.

• $P_H(v)$: the auction price given that (i) there is a v-consumer that has joined the auction, and (ii) all other consumers use the participation strategy $H \in \mathscr{H}$. To compute the probability distribution of this price we need to estimate the number of bidders and their corresponding valuations. For this, consider $H \in \mathscr{H}$ and let us define for all $v \in \mathscr{V}$

$$\Lambda_H(v) \triangleq \lambda \int_v^1 (T - H(x))\,\mathrm{d}F(x) \triangleq \lambda T \eta_H(v),$$
$$\text{where } \eta_H(v) \triangleq \int_v^1 \left(1 - \frac{H(x)}{T}\right) \mathrm{d}F(x). \quad (4)$$

By the definition of the bidding function $H(v)$, the function $\Lambda_H(v)$ represents the average number of bidders with valuation in $[v, 1]$. Similarly, $\eta_H(v)$ represents the average fraction of arrivals with valuation in this range who choose the auction channel. We also note that the restriction $\mathscr{H} \subset \mathscr{D}$ ensures that $\Lambda_H(v)$ and $\eta_H(v)$ are well-defined functions continuous in v.

Let us denote by $B(\Lambda_H(v))$ the random number of bidders with valuation greater than or equal to v. Because consumers arrive according to a Poisson process with rate λ, $B(\Lambda_H(v))$ has a Poisson distribution with mean $\Lambda_H(v)$. We can now compute the probability distribution of $P_H(v)$ given a symmetric participation strategy $H \in \mathscr{H}$. Under a (Q_0+1)-price auction, $P_H(v) < x$ if and only if (i) $x > \hat{P}$, or (ii) the number of bidders with valuation greater than or equal to x is less than or equal to the number of objects in the auction minus $\mathbb{1}(x \le v)$. That is,

$$\mathbb{P}(P_H(v) < x) = \begin{cases} 1 & \text{if } x > \hat{P} \\ \mathbb{P}(B(\Lambda_H(x)) + \mathbb{1}(x \le v) \le Q_0) \\ \quad = \displaystyle\sum_{k=0}^{Q_0 - \mathbb{1}(x \le v)} \frac{(\Lambda_H(x))^k \exp(-\Lambda_H(x))}{k!} & \text{if } x \le \hat{P}, \end{cases} \quad (5)$$

where $\mathbb{1}(E)$ is the indicator function of event E. This also follows from the continuity of $\Lambda_H(x)$ that $\mathbb{P}(P_H(v) < x)$ is continuous in $x \in [0,1] - \{\hat{P}, v \wedge \hat{P}\}$.

3.2. Characterization of a Symmetric Participation Equilibrium $H(v)$

To characterize a symmetric participation equilibrium (SPE) $H \in \mathscr{H}$, we use the following two-step approach. First, we look at a consumer's best-response participation strategy assuming that other consumers use a fixed strategy $H \in \mathscr{H}$. We denote by $\mathscr{R}(H) \in \mathscr{H}$ this best-response participation strategy and refer to $\mathscr{R}$ as the best-response mapping on $\mathscr{H}$. Second, we impose the equilibrium condition $\mathscr{R}(H^*) = H^*$. Before moving into this analysis, we recall that the optimal strategy for consumers with valuation $v \le \hat{P}$ (independent of H) is to enter the auction, and so we must have $\mathscr{R}(H)(v) = 0$ for all $v \in [0, \hat{P}]$.

Suppose a consumer—referred to as consumer τ—arrives at time τ with private valuation $v_\tau > \hat{P}$, and suppose that every other consumer is using the participation strategy H. If consumer τ decides not to bid and buy a unit through the external fixed-price channel, then his expected utility would be $u(\tau, \tau, v_\tau - \hat{P})$. On the other hand, if he decides to bid, then his profit would be $u(\tau, T, v_\tau - \hat{P})$ if he does not get the object (because he can always buy the product in the external market at time T), and $u(\tau, T, v_\tau - P_H(v_\tau))$ if he indeed does get the object. Thus, a rational consumer τ enters the auction only if

$$\begin{aligned} &u(\tau, T, v_\tau - \hat{P})(1 - \mathbb{P}(P_H(v_\tau) < \hat{P})) \\ &\quad + \mathbb{E}[u(\tau, T, v_\tau - P_H(v_\tau)) \mid P_H(v_\tau) < \hat{P}]\mathbb{P}(P_H(v_\tau) < \hat{P}) \\ &\quad \ge u(\tau, \tau, v_\tau - \hat{P}), \end{aligned} \quad (6)$$

where $\mathbb{P}(P_H(v_\tau) < \hat{P})$ is the probability that bidder τ gets one of the auctioned objects at a price strictly less than the posted price. We still need to explicitly characterize this participation constraint in terms of the function $H(\cdot)$, but at this stage note that $H(\cdot)$ is embedded in Condition (6) through the random variable $P_H(v_\tau)$ (see Equation (5)).

Without loss of generality, we assume that $\hat{P} > 0$ (otherwise, the auction is meaningless because consumers will go directly to the external fixed price channel and get a unit at no risk). This is also consistent with the scaling in (3) where we set $v^{\mathrm{R}} = 0$, because all the bidders with valuations smaller than v^{R} will have zero utility, and will quit without purchasing in any of the channels.

We compute the best-response strategy $\mathscr{R}(H)$ for consumer τ by looking at the threshold function that is consistent with (6). First, note that in our setting, where consumers have the exponentially discounted utility function defined in Equation (1), Condition (6) is equivalent to

$$\frac{\hat{P} - \mathbb{E}[P_H(v_\tau) \mid P_H(v_\tau) < \hat{P}]}{v_\tau - \hat{P}} \mathbb{P}(P_H(v_\tau) < \hat{P}) \ge \exp(w(T - \tau)) - 1. \quad (7)$$

From this condition, we conclude that, in equilibrium $\mathbb{E}[P_H(v_\tau) \mid P_H(v_\tau) < \hat{P}] < \hat{P}$. That is, no consumer would have an incentive to bid if the expected auction price in case of winning is at least what he would have paid in the external market, at no risk. The following proposition characterizes $\mathcal{R}(H)$.

Proposition 1. *For the exponential utility function* (1), *Condition* (7) *is equivalent to*

$$\tau \geq T - \frac{1}{w}\ln\left(1 + \int_0^{\hat{P}} \frac{\mathbb{P}(P_H(v_\tau) < x)}{v_\tau - \hat{P}}\,dx\right). \qquad (8)$$

Thus, a consumer arriving at time τ with valuation v_τ enters the auction if and only if $\tau \geq \mathcal{R}(H)(v_\tau)$, where

$$\mathcal{R}(H)(v_\tau) \triangleq \begin{cases} 0 & \text{if } v_\tau \in [0, \hat{P}] \\ \left[T - \dfrac{1}{w}\ln\left(1 + \displaystyle\int_0^{\hat{P}} \frac{\mathbb{P}(P_H(v_\tau) < x)}{v_\tau - \hat{P}}\,dx\right)\right]^+ & \\ \quad \text{if } v_\tau \in (\hat{P}, 1]. \end{cases} \qquad (9)$$

This best-response mapping $\mathcal{R}(H)(v_\tau)$ is continuous in v_τ.

Because the best-response strategy $\mathcal{R}(H)(v)$ is continuous in $[0, 1]$, it follows that $\mathcal{R}$ effectively maps $\mathcal{H}$ into $\mathcal{H}$. Furthermore, because an SPE is characterized by the fixed-point condition $\mathcal{R}(H) = H$, we conclude that a symmetric participation equilibrium of this game $H^* \in \mathcal{H}$ is in fact continuous. Our next result extends this conclusion and shows that the best-response strategies are K-Lipschitz continuous[10] functions in $[0, 1]$, for an appropriate constant $K > 0$. This additional property of the bidding strategies becomes relevant in our proof of existence of an equilibrium. However, before we formally address this issue we need the following lemma.

Lemma 1. *For all $H \in \mathcal{H}$ there is a valuation $v_H > \hat{P}$ such that $\mathcal{R}(H)(v) = 0$ for all $v \in [0, v_H]$. The infimum of v_H over H satisfies*

$$\begin{aligned}\tilde{v} &= \inf_{H\in\mathcal{H}}\{v_H\} \\ &\geq \min\left\{1;\ \hat{P} + \exp(-wT) \right. \\ &\qquad \left. \cdot\left(1 + \int_0^{\hat{P}} \mathbb{P}(B(\lambda T[1 - F(x)]) \leq Q_0 - 1)\,dx\right)\right\}.\end{aligned}$$

In particular, this means that in equilibrium, all consumers with valuations slightly above the posted price (i.e., with valuation $v \in [\hat{P}, v_{H^*}])$ will join the auction regardless of their arrival time. Also, note that for some instances of the problem (e.g., when wT is small) it could be that $\tilde{v} = 1$. In these cases, we must have $\mathcal{R}(H)(v) = 0$ for all $v \in (\hat{P}, 1]$ and so $H(v) = 0$ is the unique SPE.

Proposition 2. *For the exponential utility function* (1) *and for all $H \in \mathcal{H}$, there is a positive constant K* (*independent of H*) *such that the best-response strategy $\mathcal{R}(H)(v)$ is a K-Lipschitz continuous function that satisfies $\mathcal{R}(H)(v) = 0$ for all $v \in [0, \tilde{v}]$. In addition, if $\tilde{v} = 1$, then it is optimal for every consumer, independent of his arrival time and private valuation, to enter the auction. That is, if $\tilde{v} = 1$, then $\{H(v) = 0$ for all $v \in [0, 1]\}$ is the unique* (*symmetric*) *participation strategy equilibrium.*

Proposition 2 characterizes an SPE for those special cases in which every consumer enters the auction. For the general case, finding a symmetric equilibrium $H(v)$ or even proving its existence is not an easy task. A standard way to approach the existence problem is to prove that the set of bidding strategies $\mathcal{H}$ has the fixed-point property (see Cheney 2001, §7.1, for details.)[11] and that the best-response mapping $\mathcal{R}$ is continuous in $\mathcal{H}$. We will take this approach here, although we first need to slightly modify our set of strategies $\mathcal{H}$. From our previous discussion and the result in Proposition 2, we can restrict the search of a symmetric equilibrium to those strategies H in $\mathcal{H}$ that are K-Lipschitz continuous and satisfies $H(v) = 0$ in $[0, \tilde{v}]$. For this reason, we redefine $\mathcal{H}$ to be this set:

$$\mathcal{H} \triangleq \{H\colon [0, 1] \to [0, T] \text{ s.t. } H \text{ is } K\text{-Lipschitz continuous and } H(v) = 0 \text{ in } v \in [0, \tilde{v}]\}.$$

Note that by Proposition 2, $\mathcal{R}$ is a well-defined mapping from $\mathcal{H}$ to $\mathcal{H}$.

Theorem 1. *The set of strategies $\mathcal{H}$ equipped with the uniform norm $\|X\| = \sup_{0\leq v\leq 1}\{|X(v)|\}$ in $[0, 1]$ exhibits the fixed-point property. In addition, for all $H, \widetilde{H} \in \mathcal{H}$, the mapping $\mathcal{R}$ satisfies*:

$$\|\mathcal{R}(H) - \mathcal{R}(\widetilde{H})\| \leq \frac{\lambda\mathbb{P}(B(Q_0 - 1) = Q_0 - 1)}{w} \cdot\left(\int_0^{\hat{P}} \frac{(1 - F(x))}{\tilde{v} - \hat{P}}\,dx\right)\|H - \widetilde{H}\|.$$

Therefore, $\mathcal{R}$ is a continuous mapping and there always exists an SPE. In addition, if

$$\frac{\lambda\mathbb{P}(B(Q_0 - 1) = Q_0 - 1)}{w}\left(\int_0^{\hat{P}} \frac{(1 - F(x))}{\tilde{v} - \hat{P}}\,dx\right) < 1,$$

then $\mathcal{R}$ is a contraction. In this case, the fixed point $\mathcal{R}(H^*) = H^*$ is guaranteed to be unique in $\mathcal{H}$ and can be

[10] We say that a real function $f(x)$ is K-Lipschitz continuous in $\mathcal{A}$ if for all $x, y \in \mathcal{A}$, there exists a constant K independent of x and y such that $|f(x) - f(y)| \leq K|x - y|$.

[11] A set $\mathcal{H}$ has the fixed-point property if every continuous mapping $\mathcal{R}\colon \mathcal{H} \to \mathcal{H}$ has a fixed point.

found through the iteration $H^{n+1} = \mathcal{R}(H^n)$ *starting at an arbitrary* $H^1 \in \mathcal{H}$.

We conclude our characterization of an SPE with the following proposition.

Proposition 3. *For the exponential utility function* (1), *an SPE* $H^*(v)$ *is an increasing and concave function of* $v \in [v_{H^*}, 1]$.

The monotonicity of $H^*(v)$ implies that in equilibrium consumers with large private valuation are less likely to enter the auction. High-value consumers lose more by waiting and thus are less likely to choose the auction; this leads to adverse selection and a poorer distribution of values among auction bidders.[12]

Another interesting property of $H^*(v)$ has to do with the resulting bidding pattern it induces. Let us define $\lambda_{H^*}(t)$ to be the equilibrium bidding rate at time t. Because of the concavity of $H^*(v)$, it follows that $\lambda_{H^*}(t)$ is increasing and convex in the range $t \in [0, H^*(1)]$ and remains constant afterwards at its maximum value λ. This bid-speeding-up feature (sometimes referred to as *late bidding* or *sniping*) has been empirically observed in online auctions with a rigid deadline, like the ones conducted through eBay (see, e.g., Roth and Ockenfels 2002, Shmueli et al. 2004).[13] In this respect, our model provides a simple description of this phenomenon based on consumers' sensitivity to delay.

Note that extending our analysis to include randomized strategies could extend the scope of our results. However, in our setting there is no need for this additional degree of complexity.

Proposition 4. *Without loss of generality, we can restrict our equilibrium analysis to pure participation strategies.*

To summarize, in this section we have characterized consumers' best-response participation strategy through a continuous mapping $\mathcal{R}$. We have shown that there always exists an SPE $H^*(v)$ satisfying $H^* = \mathcal{R}(H^*)$, which is identically zero in the range $[0, v_{H^*}]$, and K-Lipschitz continuous, increasing, and concave in the range $[v_{H^*}, 1]$. Moreover, we have been able to characterize those cases in which either $H^*(v) = 0$ or the algorithm $H^{n+1} = \mathcal{R}(H^n)$ effectively converges to H^*. Unfortunately, this analysis is not exhaustive in the sense that there are instances of the problem for which we do not have a guaranteed method for computing the equilibrium.[14] For this reason, in the next section we consider an asymptotic regime for which explicit solutions are obtained. As we will see, numerical experiments reveal that the behavior of the asymptotic approximation is satisfactorily accurate.

3.3. Asymptotic Analysis

In this section, we characterize the outcome of the auction using asymptotic analysis. In particular, we consider the limiting case in which both the number of units Q_0 and the average arrival rate λ grow proportionally large. In this regime, we show that characterizing the consumers' strategy (i.e., the threshold function $H(v)$) is equivalent to solving a deterministic problem, which we can do efficiently (Theorem 3).

Consider a sequence of instances of the problem indexed by n and let Q_0^n and λ^n be the corresponding number of units to auction and demand rate for instance n, respectively. All other parameters are kept independent of n. The asymptotic regime that we consider is defined by

$$\lim_{n\to\infty} \frac{Q_0^n}{n} = Q_0 \quad \text{and} \quad \lim_{n\to\infty} \frac{\lambda^n}{n} = \lambda. \tag{10}$$

For each instance n of the problem, we let $\rho^n \triangleq Q_0^n/(\lambda^n T)$. Then, $\lim_{n\to\infty} \rho^n = \rho$, for $\rho \triangleq Q_0/(\lambda T)$. For future reference, we refer to ρ as the *excess-supply ratio*, which can be viewed as a proxy for the average number of units available per arriving consumer. In this respect, the case $\rho < 1$ reflects the most interesting situation in which, on average, there are fewer units than consumers.

We denote by $P_H^n(v)$ the auction price for instance n given that a v-consumer enters the auction. The distribution of this price is given by Equation (5) replacing Q_0 by Q_0^n and λ by λ^n. We are now ready to characterize the asymptotic regime.

Theorem 2. *Suppose the participation strategy* $H(v)$ *is fixed. Then, in the limit as* $n \to \infty$, *the auction price* $P_H^n(v)$ *converges weakly*[15] *to the constant* $P_H^\infty(v) = \min\{\hat{P}, \eta_H^{-1}(\rho)\}$, *where* $\eta_H^{-1}(\rho) \triangleq \min\{v \in [0,1]\colon \eta_H(v) \le \rho\}$.

[12] We thank one of the referees for pointing out this observation.

[13] In a different auction-ending rule, bids are accepted after the original deadline, as long as there is some recent bidding activity, say within the last 10 minutes. This is the type of closing rule followed at Amazon.

[14] In practice, we have been able to find an SPE for all instances that we have tested using the following small step-size version of the iteration in Theorem 1 (see Bertsekas and Tsitsiklis 1996, Chapter 4):

$$H^{n+1} = \alpha H^n + (1-\alpha)\mathcal{R}(H^n),$$

where $\alpha \in [0, 1)$ is empirically selected.

[15] A sequence of distribution functions is said to *converge weakly* to a limit F (written $F_n \Rightarrow F$) if $F_n(y) \to F(y)$ for all y that are continuity points of F. A sequence of random variables X_n is said to *converge weakly* or *converge in distribution* to a limit X_∞ (written $X_n \Rightarrow X_\infty$) if their distribution functions $F_n(x) \triangleq P(X_n \le x)$ converge weakly (see Durrett 1996, §2.2).

It follows from Theorem 2 that the probability distribution of the limiting auction price is

$$\mathbb{P}(P_H^\infty(v) < x) = \begin{cases} 1 & \text{if } x > \hat{P}, \text{ or } x \le \hat{P} \text{ and } \rho > \eta_H(x) \\ 0 & \text{if } \rho < \eta_H(x) \text{ and } x \le \hat{P}. \end{cases}$$

Suppose $x \le \hat{P}$ and $\rho < \eta_H(x)$, that is, there are on average more bidders with valuation at least x than units to auction, $Q_0 < \Lambda_H(x)$. In this situation, there is a scarcity of units for bidders with valuation greater than or equal to x. Note that the monotonicity of $\eta_H(x)$ implies $x \le \min\{\hat{P}, \eta_H^{-1}(\rho)\}$, and it follows from Theorem 2 that the auction price will be higher than x. On the other hand, when there are more units to auction than bidders with valuation of at least x, that is, $\rho > \eta_H(x)$, the final auction price will be lower than x with certainty.

Note that the limiting auction price $P_H^\infty(v)$ does not depend on v because as the number of units and consumers goes to infinity, the auction price is unaffected by the decision of one particular bidder. In other words, we can think of this asymptotic regime as one in which we have a continuum of marginal consumers, each one having no impact on the overall outcome of the auction. In this respect, the asymptotic regime under consideration is of the fluid type. For notational convenience, let us define $P_H \triangleq P_H^\infty(v)$.

Suppose consumer τ with valuation v_τ arrives at time τ and suppose that every other consumer is using the participation strategy H.

- If $v_\tau \in [0, \hat{P}]$, the consumer enters the auction independently of the auction price because the posted price exceeds his valuation. Thus, we have that $\mathcal{R}(H)(v_\tau) = 0$ for all $v_\tau \in [0, \hat{P}]$.
- If $v_\tau \in (\hat{P}, 1]$, then it is optimal for consumer τ to participate in the auction only if

$$(v_\tau - P_H)\exp(-w(T-\tau)) \ge v_\tau - \hat{P}, \quad \text{or equivalently } \tau \ge h(v_\tau) \quad (11)$$

for the auxiliary threshold function $h(v_\tau) = T - 1/w \cdot \ln((v_\tau - P_H)/(v_\tau - \hat{P}))$. Because the logarithm in $h(v_\tau)$ goes to ∞ as $v_\tau \downarrow \hat{P}$, consumers with valuation greater, but close to $\hat{P}$, will enter the auction. In fact, let us define $v_H \ge \hat{P}$ to be the root of the equation

$$\left[T - \frac{1}{w}\ln\left(\frac{v - P_H}{v - \hat{P}}\right)\right]^- = 0 \quad \text{which is equivalent to}$$

$$v_H \triangleq \min\left\{\frac{\hat{P}\exp(wT) - P_H}{\exp(wT) - 1}, 1\right\}. \quad (12)$$

Then, by Condition (11) consumer τ with valuation $v_\tau \in (\hat{P}, v_H]$ will always enter the auction independently of his arrival time τ. On the other hand, consumer τ with valuation v_τ greater than v_H will enter the auction only if his arrival time is greater than $\mathcal{R}(H)(v_\tau)$, where $\mathcal{R}(H)(v_\tau) = T - (1/w)\ln((v_\tau - P_H)/(v_\tau - \hat{P}))$, for $v_\tau \in [v_H, 1]$. Clearly, if $v_H = 1$, then every consumer will enter the auction independently of his arrival time, in which case $H^*(v) = 0$ for all $v \in [0, 1]$.

In summary, given $H \in \mathcal{H}$ and the associated auction price P_H, the best-response participation strategy $\mathcal{R}(H)$ satisfies

$$\mathcal{R}(H)(v) = \begin{cases} 0 & \text{if } v \in [0, v_H] \\ T - \dfrac{1}{w}\ln\left(\dfrac{v - P_H}{v - \hat{P}}\right) & \text{if } v \in [v_H, 1]. \end{cases} \quad (13)$$

We note that $\mathcal{R}(H)(v)$ is increasing and concave in v, for all $v \ge v_H$.

To determine the equilibrium value of the auction price P_{H^*} and the corresponding participation strategy $H^*(v)$, we have to impose the equilibrium condition $\mathcal{R}(H^*) = H^*$. In this asymptotic regime, we can solve this fixed-point condition efficiently using Equation (13) and Theorem 2.

Theorem 3. *In the asymptotic regime under consideration, the auction price P_{H^*} is the unique solution in $[0, \hat{P}]$ to the equation*

$$F(v_{H^*}) - F(P_{H^*}) + \frac{1}{wT}\int_{v_{H^*}}^1 \ln\left(\frac{v - P_{H^*}}{v - \hat{P}}\right) dF(v) = \min\{\rho, \eta_{H^*}(0)\}, \quad (14)$$

where $v_{H^*} = \min\{(\hat{P}\exp(wT) - P_{H^*})/(\exp(wT) - 1), 1\}$ and the SPE strategy $H^*(v)$ is given by

$$H^*(v) = \begin{cases} 0 & \textit{if } v \in [0, v_{H^*}] \\ T - \dfrac{1}{w}\ln\left(\dfrac{v - P_{H^*}}{v - \hat{P}}\right) & \textit{if } v \in [v_{H^*}, 1]. \end{cases}$$

The next result characterizes two extreme outputs of the game and follows directly from Theorem 3.

Corollary 1. *In the asymptotic regime, it follows that*

(a) *the auction price equals the reservation price, that is, $P_{H^*} = 0$, if*

$$F(v_{H^*}) + \frac{1}{wT}\int_{v_{H^*}}^1 \ln\left(\frac{v}{v - \hat{P}}\right) dF(v) \le \rho, \quad \textit{where } v_{H^*} = \min\left\{\frac{\hat{P}\exp(wT)}{\exp(wT) - 1}, 1\right\} \quad \textit{and}$$

(b) *all arriving consumers enter the auction, that is, $H^*(v) = 0$ for all $v \in [0, 1]$, if*

$$F^{-1}(1-\rho) \le 1 - \exp(wT)(1 - \hat{P}), \quad \textit{in which case } P_{H^*} = F^{-1}(1-\rho).$$

Figure 2 Single-Channel Asymptotic Approximation for the Case with Valuations Uniformly Distributed in $[0, 1]$, and $T = w = 1$, $\hat{P} = 0.5$, and $\rho = 0.6$

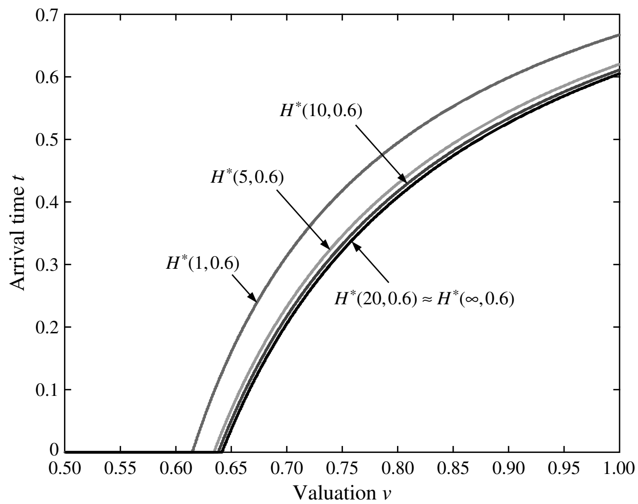

Q_0	P_H	P_H^{Approx}	Error (%)
1	0.1249	0.1346	7.20
5	0.1930	0.1979	2.43
10	0.2144	0.2168	1.12
20	0.2312	0.2319	0.30

Note. In this case, the asymptotic price is $P_{H^*}^{\infty} = 0.2578$.

Using a slight abuse of notation, let $H^*(Q_0, \rho)$ be the optimal participation strategy if the auctioneer has Q_0 units to auction and the excess supply ratio $(Q_0/\lambda T)$ is equal to ρ. In Figure 2, we compare the optimal asymptotic participation strategy $H^*(\infty, 0.6)$ (computed using Theorem 3) to four optimal bidding strategies: $H^*(1, 0.6)$, $H^*(5, 0.6)$, $H^*(10, 0.6)$, and $H^*(20, 0.6)$ (computed numerically using the iteration in Theorem 1). From the graph on the top of Figure 2, we can see that the asymptotic approximation mimics quite closely consumers' participation strategy even for small values of Q_0. As a matter of fact, for values of Q_0 greater than 10 units, the bidding strategies $H^*(Q_0, 0.6)$ and $H^*(\infty, 0.6)$ are almost indistinguishable.

The table on the bottom of Figure 2 compares the expected price of the auction P_H with the approximated value P_H^{Approx} obtained using the asymptotic participation strategy. In other words, P_H^{Approx} is the expected auction price if every consumer uses the participation strategy $H^*(\infty, \rho)$. As we can see, for a single-unit auction the error on the estimate is about 7%. However, for moderate multiunit auctions (with five or more items) the quality of the approximation improves considerably quickly. With 20 items the approximation is almost exact. As we expect, both the auction price P_H and the approximated auction price P_H^{Approx} converge to the asymptotic price P_H^{∞} as Q_0 goes to infinity. We note that the quality of results reported in this example were systematically replicated in all instances of the problem that we considered.

We conclude this section by applying our results to a particular instance of the problem with uniform valuations, a distribution widely considered in the auction literature. In §5, we will use this example to study the seller's optimization problem.

Example (Uniform Distribution Case). Suppose the valuations are uniformly distributed in $[0, 1]$, that is, $F(v) = v$. In this case, we can apply Theorem 3 to get the following cases:

- If $\hat{P} \le \hat{P}_1(\rho, wT)$ where $\hat{P}_1(\rho, wT)$ solves

$$\frac{-1}{wT}\left[(1-\hat{P}_1)\ln(1-\hat{P}_1) + \hat{P}_1 \ln\left(\frac{\hat{P}_1}{\exp(wT)-1}\right)\right] = \min\{1, \rho\},$$

then $P_{H^*} = 0$ and $H^*(v) = [T - (1/w)\ln(v/(v-\hat{P}))]^+$.

- If $\hat{P}_1(\rho, wT) < \hat{P} < \hat{P}_2(\rho, wT) \triangleq 1 - \exp(-wT) \cdot (1-(1-\rho)^+)$, then the auction price $P_{H^*} \in (0, \hat{P})$ solves $\eta_{H^*}(P_{H^*}) = \min\{1, \rho\}$, which in this case is the same as

$$\frac{1}{wT}\left[(1-P_{H^*})\ln(1-P_{H^*}) - (1-\hat{P})\ln(1-\hat{P}) - (\hat{P}-P_{H^*})\ln\left(\frac{\hat{P}-P_{H^*}}{\exp(wT)-1}\right)\right] = \min\{1, \rho\}.$$

- If $\hat{P} \ge \hat{P}_2(\rho, wT)$, then $H^* = 0$ and $P_{H^*} = (1-\rho)^+$. Note that in this case $\eta_{H^*}(P_{H^*}) = \min\{1, \rho\}$.

In other words, $\hat{P}_1(\rho, wT)$ is a lower bound on the posted price over which there will be enough people in the auction such that the resulting auction price will be positive. $\hat{P}_2(\rho, wT)$ is the minimum posted price, such that all consumers will participate in the auction. Figure 3 plots the auction price as a function of the posted price. As we can see, for low values of the posted price $\hat{P}$, the auction price coincides with the reservation price at zero. The intuition in this case is that a low posted price will induce a large fraction of high-valuation consumers to purchase at the posted price-resulting in fewer bidders left on the auction; bidders that, moreover, have a low valuation. As we can see from Figure 3, $\hat{P}_1(\rho, wT)$ increases, with ρ reflecting the fact that the higher the number of units in the auction, the lower the value of the bid of the last winning bidder. On the other extreme, $\hat{P}_2(\rho, wT)$ decreases with ρ; that is, the higher the ρ, the lower the posted price needed to induce all the players to participate in the auction. In summary, the auction price P_{H^*} is a nondecreasing function of $\hat{P}$ given by

$$P_{H^*} = \mathcal{P}(\hat{P}, \rho, wT) \triangleq \begin{cases} 0 & \text{if } \hat{P} \in [0, \hat{P}_1(\rho, wT)] \\ \text{solves } \eta_{H^*}(P_{H^*}) = \min\{1, \rho\} & \text{if } \hat{P} \ge \hat{P}_1(\rho, wT). \end{cases} \quad (15)$$

Figure 3 Auction Price P_H as a Function of the Posted Price $\hat{P}$ for Four Different Values of the Excess-Supply Ratio ρ

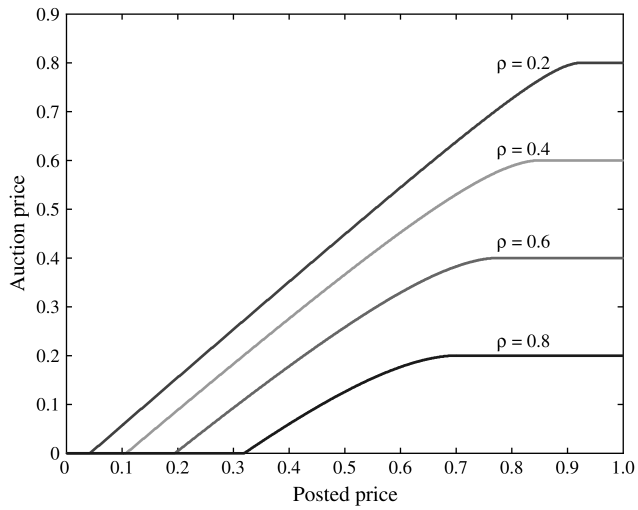

Note. In this numerical example $wT = 1$.

The corresponding number of units sold in the auction Q_{H^*} satisfies

$$Q_{H^*} = \lambda T \mathcal{Q}(\hat{P}, \rho, wT),$$

$$\text{where } \mathcal{Q}(\hat{P}, \rho, wT) \triangleq \begin{cases} \eta_{H^*}(0) & \text{if } \hat{P} \in [0, \hat{P}_1(\rho, wT)] \\ \min\{1, \rho\} & \text{if } \hat{P} \in [\hat{P}_1(\rho, wT), 1], \end{cases} \tag{16}$$

and

$$\eta_{H^*}(0) = \begin{cases} \dfrac{1}{wT}\left[\hat{P}\ln\left(\dfrac{(1-\hat{P})(\exp(wT)-1)}{\hat{P}}\right) - \ln(1-\hat{P})\right] \\ \qquad \text{if } \hat{P}\exp(wT) < \exp(wT) - 1 \\ 1 \quad \text{if } \hat{P}\exp(wT) \geq \exp(wT) - 1. \end{cases}$$

4. Dual Channel with Static List Price

In this section, we study the consumers' optimal buying strategy in the face of a monopolistic seller who has Q_0 units to sell, using two parallel channels: a list price channel in which she sets a constant list price $\hat{P}$ for the entire selling horizon, and the auction that will take place at time T with the remaining Q_T units. Most of the analysis of this dual channel mimics the steps we used in the previous section. For this reason, we summarize only the main results here, skipping many of the mathematical details. A complete analysis of this dual channel model can be found in online Appendix B.

One of the important differences of this model with respect to the one discussed in the previous section is that the fixed-price channel now has a limited supply of Q_0 units. Therefore, bidders that lose in the auction have no alternative market in which to buy the product. Hence, to decide whether to enter the auction or buy at the fixed price, an arriving consumer must estimate the joint probability distribution of the auction price (P_H) and the number of units that will be left unsold for the auction (Q_T). In general, this could be a rather complicated (path-dependent) task because the state of the system changes continuously. However, under our assumption that the seller does not reveal any information about the inventory position or bidding process over time (see §2 for more details), the actual computation of P_H and Q_T simplifies considerably. This assumption is certainly restrictive, and getting rid of it would be an important extension to our analysis (again, see Gupta and Gallien 2007 for some preliminary steps in this direction). We will argue at the end of this section, however, that in an asymptotic sense our equilibrium under partial information is also an equilibrium for the general case with full information. This is a positive conclusion that validates in part our (simplifying) informational assumptions.

Another important difference of this model is on the bidding side. In this new setting the optimal strategy for bidders is $b(v) = v$ as opposed to the strategy $b(v) = \min\{\hat{P}, v\}$ of the previous section. To see this, note that in this case a high-valuation consumer (i.e., with $v > \hat{P}$) who enters the auction and loses will not get the object at all, because all the units are cleared at T among buyers with valuations above v^{R}. Hence, once this high-valuation consumer decides to enter the auction, the subgame he plays is exactly equivalent to a $(Q_T + 1)$-price auction for which we know the optimal strategy is to bid the true valuation, $b(v) = v$.[16]

As in the previous section, we analyze the consumers' participation problem under the scaling in (3), which we revert to in §5 when we study the seller's optimization problem. For simplicity, we will keep the same notation that we use in the single-channel case, with the understanding that the value of some quantities (such as $\tilde{v}$, v_H, or P_H) will be slightly different.

Suppose $\hat{P}$ has been selected. As in the previous section, we restrict consumers' participation strategies to the set

$$\mathcal{H} = \{H \in \mathcal{D}\colon [0,1] \to [0,T] \text{ such that } H(v) = 0, \quad \text{for all } v \in [0, \hat{P}]\},$$

[16] Note that in this case Q_T is not a fixed quantity, but a random variable that depends on the initial number of units Q_0 and the number of consumers that select the fixed-price channel during $[0, T)$. The strategy is dominant even for this case, where the number of units to auction Q_T is uncertain. The argument is similar to the one used for the standard uniform price auction.

and we characterize the decision of a consumer by means of a threshold function $H \in \mathscr{H}$, such that a consumer arriving at time t with valuation v_t will join the auction only if $H(v_t) \le t$. The following is our main result from the point of view of the consumers' game.

Theorem 4. *For the dual-channel model with exponential utility function* (1), *a K-Lipschitz continuous symmetric equilibrium always exists in* $\mathscr{H}$. *In addition, for any symmetric equilibrium* H^* *there is a valuation* $v_{H^*} > \hat{P}$ *such that* $H^*(v) = 0$ *for all* $v \le v_{H^*}$.

The existence of v_{H^*} implies that there is a range of valuations above the list price $\hat{P}$, such that consumers with valuations in this range join the auction regardless of their arrival time.

Theorem 4 is important from a theoretical standpoint, but it does not lead to a simple way of characterizing a symmetric participation equilibrium H^*. As in §3.3, we can tackle the problem of determining H^* using asymptotic analysis when both the initial number of units Q_0 and the arrival rate λ grow proportionally large. Specifically, we consider a sequence of instances of the dual-channel problem indexed by n and let Q_0^n and λ^n be the corresponding initial number of units and demand rate for instance n. The asymptotic regime that we consider is again given by Equation (10).

For a given participation strategy $H \in \mathscr{H}$, we define the auxiliary function

$$\eta_{H-}(x) \triangleq \int_x^1 \frac{H(v)}{T}\, \mathrm{d}F(v) = \bar{F}(x) - \eta_H(x),$$

where $\bar{F}(x)$ stands for the tail distribution of the valuations. We also define $P_H^n(v)$ to be the auction price for instance n, given that a bidder with valuation v enters the auction, and Q_T^n to be the final random number of units to auction. The following result characterizes the asymptotic regime.

Theorem 5. *Suppose the participation strategy* $H(v)$ *and the static price* $\hat{P}$ *are given. Then, in the limit as* $n \to \infty$:

(i) *The rescaled number of units* Q_T^n/n *to sell through the auction converges weakly to a constant* $Q_T \triangleq (Q_0 - \lambda T \eta_{H-}(0))^+$.

(ii) *If a final auction takes place* (*i.e.*, $Q_T > 0$), *its price* $P_H^n(v)$ *converges weakly to a constant* $P_H^\infty \triangleq \min\{v \in [0,1]\colon \eta_H(v) \le \rho - \eta_{H-}(0)\}$, *where* $\rho = Q_0/(\lambda T)$.

We can use this result to characterize a symmetric participation equilibrium, H^*, in this asymptotic regime. Of course, we already know that $H^*(v) = 0$ in $v \in [0, \hat{P})$, so we only need to find the behavior of $H^*(v)$ in $v \in [\hat{P}, 1]$. We distinguish two cases.

Case 1: Limited-Supply. Suppose that the initial supply of units is *limited* in the sense that $\rho \le 1 - F(\hat{P})$. In this situation, consumers with valuation greater than $\hat{P}$ have no incentive to enter the auction because the auction price is guaranteed to be greater than or equal to $\hat{P}$. Therefore, the resulting participation strategy is $H^*(v) = T$ for all $v \ge \hat{P}$, and the auction never takes place because all the units will be bought at the posted price (i.e., $Q_T = 0$).

We note that $H^*(v) = T\mathbb{1}(v \ge \hat{P})$ is not the only SPE in this case. In fact, let us define $\tau^* \triangleq T\rho(1 - F(\hat{P}))^{-1}$. Then, any H of the form $H(v) = (\tau^* + h(v))\mathbb{1}(v \ge \hat{P})$, for an arbitrary nonnegative and bounded function $h(v) \le T - \tau^*$, is an SPE. In fact, for such an H, the initial Q_0 units will be depleted by time τ^* (i.e., $Q_{\tau^*} = 0$, because $\tau^* \lambda(1 - F(\hat{P})) = Q_0$). Therefore, any consumer arriving after τ^* will never get a unit, and so he becomes indifferent between the two channels.

Case 2: Abundant-Supply. Suppose that initial supply is *abundant* in the sense that $\rho > 1 - F(\hat{P})$. In this case, $Q_T > 0$, and some consumers with valuation smaller than $\hat{P}$ get units through the auction. From part (ii) in Theorem 5, it is not hard to see that in this abundant-supply case, the auction price satisfies $P_{H^*} = F^{-1}(1-\rho)$. Therefore, a consumer arriving at time τ with valuation $v_\tau \ge \hat{P}$ enters the auction only if $v_\tau - \hat{P} \le \exp(-w(T-\tau))(v_\tau - P_{H^*})$. We conclude that in this *abundant* case the unique SPE $H^*(v)$ is given by

$$H^*(v) = \begin{cases} 0 & \text{if } v \in [0, v_{H^*}] \\ T - \dfrac{1}{w}\ln\left(\dfrac{v - P_{H^*}}{v - \hat{P}}\right) & \text{if } v \in [v_{H^*}, 1] \end{cases}$$

$$\text{where } v_{H^*} = \min\left\{\frac{\hat{P}\exp(wT) - P_{H^*}}{\exp(wT) - 1}, 1\right\}.$$

Figure 4 compares the optimal asymptotic strategy $H^*(\infty, \rho)$ to four optimal participation strategies $H^*(Q_0, \rho)$ ($Q_0 = 1, 5, 10, 20$) for the case of $\rho = 0.5$. As in the single-auction channel case, the asymptotic approximation is very accurate even for small values of Q_0, and almost identical to the optimal strategy for values of Q_0 greater than 10 units. The table on the bottom compares the expected value of the auction price P_H and the approximation P_H^{Approx} obtained using the asymptotic participation strategy $H^*(\infty, 0.5)$.

Figure 4 **Dual-Channel Asymptotic Approximation for the Case with Valuations Uniformly Distributed in** $[0, 1]$, **and** $T = w = 1$, $\hat{P} = 0.6$, **and** $\rho = 0.5$

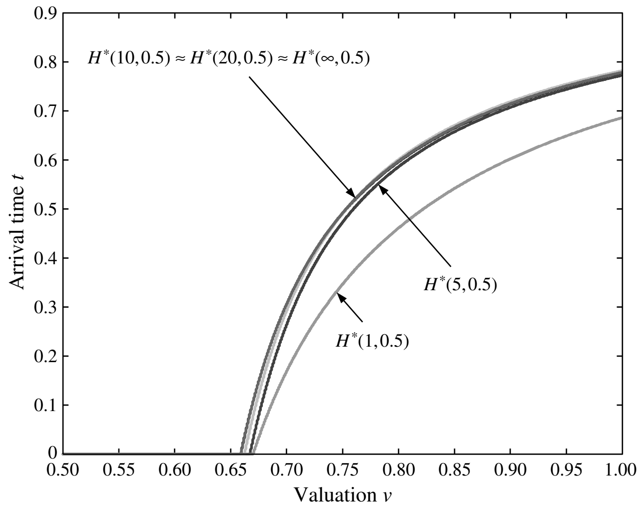

Q_0	P_H	P_H^{Approx}	Error (%)
1	0.2793	0.2834	1.428
5	0.4175	0.4184	0.213
10	0.4539	0.4540	0.034
20	0.4760	0.4760	0.002

Note. In this case, the asymptotic price is $P_{H^*}^{\infty} = 1 - \rho = 0.5$.

We conclude this section with an important remark. In the asymptotic model that we have considered, the number of units left for the auction Q_T and the auction price P_H converge weakly to a deterministic solution. Hence, in this limiting regime (and for any participation strategy $H \in \mathscr{H}$), consumers have no real uncertainty about the outcome of the game (P_H, Q_T). This certainty implies that a consumer arriving at time t does not need to know the state of the system at this time (in terms of remaining inventory or bidding process) to decide his participation strategy. He can perfectly infer this information based on the initial inventory, Q_0, the fixed price $\hat{P}$, and the participation strategy H. Therefore, in this asymptotic regime our equilibrium strategy H^* computed assuming imperfect information is also an equilibrium for the full-information case. In other words, the fact that the seller does not report the evolution of Q_t and the bidding process over time is immaterial in the asymptotic limit.

5. Seller's Optimization Problem

In the previous two sections we have studied the optimal buying strategy from the consumers side, assuming that the parameters of the auction and list price channels are fixed. In this section, we turn to the seller's problem and derive some managerial guidelines to support an optimal design of the auction and the list price channels. For simplicity, and given the accuracy of the asymptotic approximation, we will work under this limiting regime. Furthermore, we will assume that the original distribution of valuations is a standard uniform, that is, $F \triangleq \text{Unif}[0, 1]$. We use this uniform distribution to make the exposition cleaner, but we note that the general case uses exactly the same line of arguments.

As with most asymptotic analysis, the results in this section have to be interpreted with caution. This is because our fluid-type scaling in (10) has the property of "washing away" the stochasticity of the problem, capturing only first-order effects. Nevertheless, our numerical, experiments (together with the weak convergence results) suggest that this deterministic approximation is indeed robust even for small inventory levels.

The first step in formulating the seller's problem is to revert to the scaling in (3) and write the optimization in terms of the original parameters. The unscaled arrival rate, excess-supply ratio, and posted price are

$$\lambda \to \lambda(1 - v^{\text{R}}), \qquad \rho \to \frac{Q_0}{\lambda T(1 - v^{\text{R}})} = \frac{\rho}{1 - v^{\text{R}}}, \qquad \hat{P} \to \frac{\hat{P} - v^{\text{R}}}{1 - v^{\text{R}}}. \tag{17}$$

In the asymptotic regime under consideration, however, it turns out that there is no loss of generality in taking $v^{\text{R}} = 0$. This follows from: (i) the fact that in this asymptotic limit the final auction price (P_{H^*}) is a deterministic function of Q_0 and v^{R}; and (ii) the seller's ability to ration the number of units to auction. Therefore, for those instances in which $0 < v^{\text{R}} = P_{H^*}$,[17] there is a revenue-equivalent solution with $v^{\text{R}} = 0$ and the same auction price P_{H^*}, but with (possibly) fewer units being auctioned.[18]

We will take advantage of this property to simplify the analysis (and notation) of the seller's optimization problem.

5.1. Single-Auction Channel: Auctioneer's Optimization Problem

In the single-auction channel, the auctioneer has control over the duration of the auction (T) and the number of units to allocate in the auction Q_0. We assume that the auctioneer is endowed with $\bar{Q}$ units and

[17] These are the only instances in which the choice of the reservation price affects the seller's revenue.

[18] From a practical standpoint, we can implement the asymptotic solution by setting the reservation price v^{R} equal to the resulting auction price minus a safety factor $SF > 0$. That is, $v^{\text{R}} = P_{H^*} - SF$ where P_{H^*} is the auction price computed using asymptotic analysis. In this way, the safety factor will somehow recapture the stochasticity of the arrival process, acting as a hedge in case the true number of arrivals is below its expectation.

uses a discount factor α to penalize future payoffs. Hence, the auctioneer is interested in solving (possibly numerically) the following problem:

$$V_A(\bar{Q}) = \max_{T,\rho}\left\{ e^{-\alpha T} P_{H^*}(\rho, T, \hat{P}) Q_{H^*}(\rho, T, \hat{P}) \quad \text{subject to } 0 \le \rho \le \frac{\bar{Q}}{\lambda T} \right\}, \tag{18}$$

where $P_{H^*}(\rho, T, \hat{P})$ and $Q_{H^*}(\rho, T, \hat{P})$ are the auction price and number of units sold in the auction as a function of $\rho = Q_0/(\lambda T)$, T, and the fixed price $\hat{P}$. These values are computed in Equations (15) and (16), respectively.

As we can see, the optimal choice of T trades off the time value of money (captured by the discount factor $\exp(-\alpha T)$) with the corresponding volume of demand and auction payoff. On the other hand, the optimal choice of ρ (or equivalently Q_0) balances the auction price and demand. This rationing decision is due to the fact that (under a multiunit uniform price auction) expected revenues are not guaranteed to be increasing in the number of units. We will solve the optimization problem in (18) in two steps. First, we discuss this rationing decision for a given T and then determine the optimal value of T.

In Proposition 5 below, we show that the normalized revenue $P_{H^*}(\rho, T, \hat{P}) \times Q_{H^*}(\rho, T, \hat{P})$ is unimodal in ρ for every T and $\hat{P}$. This allows us to drop the supply constraint $\rho \le \bar{Q}/(\lambda T)$ with the understanding that if the unconstrained solution ρ^* exceeds this upper bound, then we will have to truncate the solution to $\rho^* = \bar{Q}/(\lambda T)$.

In the asymptotic regime under consideration, condition (15) implies that the auctioneer must restrict the choice of ρ so that $\hat{P} \ge \hat{P}_1(\rho, wT)$. Hence, we can simplify the search of an optimal solution imposing the condition $\eta_{H^*}(P_{H^*}) = \min\{1, \rho\}$. In addition, condition (16) implies that $Q_{H^*} = \lambda T \min\{1, \rho\}$. Finally, we note that if $\rho \ge 1$, then $P_{H^*} = v^{\mathrm{R}} = 0$, and so it is in the auctioneer's interest to do some rationing, selecting $\rho < 1$. This implies that at optimality, $Q_{H^*} = \lambda T \rho$.

The following proposition characterizes an optimal solution (ρ^*, P_{H^*}) for a fixed T. Before we state the result, we define $\tilde{P}(\hat{P})$ as the unique root in $[\tilde{P}_0, \hat{P}]$ of the equation

$$\left[\frac{(e^{wT}-1)(1-\tilde{P})}{\hat{P}-\tilde{P}}\right]^{1-2\tilde{P}} = \left[\frac{(e^{wT}-1)(1-\hat{P})}{\hat{P}-\tilde{P}}\right]^{1-\hat{P}} \quad \text{where } \tilde{P}_0 \triangleq \frac{(\hat{P}+1-\exp(wT))^+}{2-\exp(wT)}.$$

Proposition 5. *The revenue $\lambda T \rho P_{H^*}(\rho, T, \hat{P})$ is unimodal in ρ for every T and $\hat{P}$. In addition, the optimal (unconstrained) excess-supply ratio ρ^* and corresponding auction price P_{H^*} satisfy:*

Case 1. If $\hat{P} < 1 - e^{-wT}/2$, then

$$\rho^* = \frac{1}{2wT}\left[\ln\left(\frac{1-\tilde{P}(\hat{P})}{1-\hat{P}}\right) + \hat{P}\ln\left(\frac{(e^{wT}-1)(1-\hat{P})}{\hat{P}-\tilde{P}(\hat{P})}\right)\right]$$

and $P_{H^*} = \tilde{P}(\hat{P})$.

Case 2. If $\hat{P} \ge 1 - e^{-wT}/2$, then $\rho^* = P_{H^*} = 1/2$.

Note that in Case 2 the price $\hat{P}$ is sufficiently large that no arriving consumer will buy in this channel. Hence, in equilibrium, the resulting game looks exactly like a standard multiunit uniform-price auction. The corresponding ρ^* and P_{H^*} are the optimal design parameters for this auction in our large-capacity and sales volume setting.[19]

The following corollary is immediate.

Corollary 2. *It follows from Proposition 5 that in the limit as consumers become increasingly time sensitive or the auction time grows large, $\lim_{wT\to\infty} \rho^* = \lim_{wT\to\infty} P_{H^*} = \hat{P}/2$. On the other hand, in the limit as consumers become increasingly patient or the auction time is short, then $\lim_{wT\downarrow 0} \rho^* = 1 - \hat{P}$ and $\lim_{wT\downarrow 0} P_{H^*} = \hat{P}$.*

Figure 5 shows the behavior of ρ^*, P_{H^*}, and the auctioneer's revenue as a function of $\hat{P}$. As expected, both the auction price and revenues are nondecreasing in $\hat{P}$. On the other hand, the optimal ρ (or equivalently, the optimal number of units Q_0 to offer) is nonmonotonic in $\hat{P}$. Although we cannot get a closed-form solution for the maximum value of ρ^* as a function of $\hat{P}$ and wT, we can numerically show that ρ^* is always bounded above by 0.8. Hence, independent of $\hat{P}$, w, and T, the optimal number of units to auction should never exceed 80% of the average demand λT. We also note that, consistent with Case 2 in the proposition, when the fixed price is sufficiently large (i.e., $\hat{P} \ge 1 - \exp(-wT)/2$) the fixed-price channel has no effect on the choice of ρ^* and the auction output.

We conclude this subsection by discussing the optimal choice of T. Figure 6 depicts the auctioneer's optimal revenue as a function of T for three values of $\hat{P}$. Naturally, revenues increase with $\hat{P}$. Less obvious is the fact that the optimal T^* is bounded above by α^{-1} (in Figure 6, this value is represented by the dashed line). To see this, note that (i) the optimal normalized revenue $\rho^* P_{H^*}(\rho^*, T, \hat{P})$ is nonincreasing in T for every $\hat{P}$, and (ii) the function $\lambda T \exp(-\alpha T)$ is maximized at $T = \alpha^{-1}$. Hence, the auctioneer's revenue function $\lambda T \exp(-\alpha T)\rho^* P_{H^*}(\rho^*, T, \hat{P})$ attains its maximum in $[0, \alpha^{-1}]$. Furthermore, Figure 6 also reveals that T^* is increasing in $\hat{P}$. In other words, the higher the fixed price, the longer the duration of the auction.

[19] This outcome is also equivalent to the optimal design parameters of a list price channel (see Talluri and van Ryzin 2004, §6.2.6.2).

Figure 5 **Optimal Excess-Supply Ratio ρ, Auction Price P_{H^*} and Revenues as a Function of Posted Price $\hat{P}$**

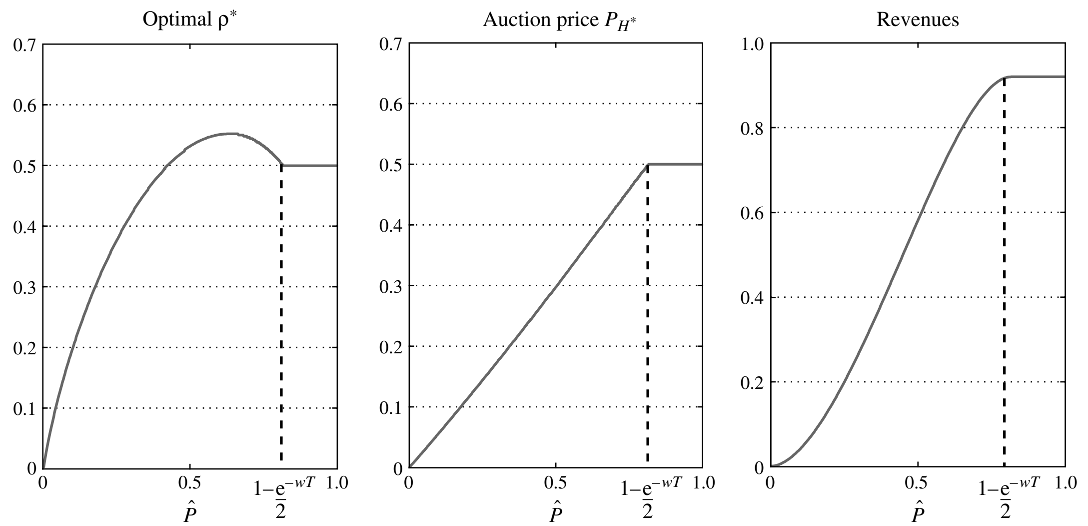

Note. Data: $w = T = \alpha = 1$ and $\lambda = 10$.

From Case 2 in Proposition 5, it follows that the upper bound $T^* = \alpha^{-1}$ is attained if $\hat{P} \geq 1 - \exp(-w/\alpha)/2$.

5.2. Dual Channel: Seller's Optimization Problem

In the dual-channel case, the seller must solve (possibly numerically) a more complex problem than the one in §5.1, as she also controls the list price channel $\hat{P}$. To write down the seller's optimization problem, we introduce one more piece of notation. We denote by $\lambda_{F^*}(t, \rho, T, \hat{P})$ the demand intensity for the fixed-price channel at time t given $(\rho, T, \hat{P})$. Combining the results in §4 and the uniform distribution assumption, it is not hard to show that (see also the online Appendix B)

$$\lambda_{F^*}(t, \rho, T, \hat{P}) = \begin{cases} \lambda(1-\hat{P})\mathbb{1}(t(1-\hat{P}) \leq \rho T) \\ \quad \text{if } \rho \leq 1 - \hat{P} \\ \lambda\left[1 - \dfrac{\hat{P}\exp(w(T-t)) - (1-\rho)^+}{\exp(w(T-t)) - 1}\right]^+ \\ \quad \text{if } \rho \geq 1 - \hat{P}. \end{cases}$$

The seller's optimization problem is given by

$$V_D(\bar{Q}) = \max_{T,\rho,\hat{P}} \left\{ \int_0^T e^{-\alpha t} \lambda_{F^*}(t, \rho, T, \hat{P}) \hat{P}\, dt + e^{-\alpha T} P_{H^*}(\rho, T, \hat{P}) Q_{H^*}(\rho, T, \hat{P}) \quad \text{subject to } 0 \leq \rho \leq \frac{\bar{Q}}{\lambda T} \right\}. \tag{19}$$

Unfortunately, the optimization problem above does not admit a simple analytical solution and we must rely on numerical computations to derive an optimal solution $(\rho^*, \hat{P}^*, T^*)$. Nevertheless, we have been able to identify some useful properties of an optimal solution that we summarize in the following proposition.

Proposition 6. (a) *Suppose that T is fixed and define $\delta = \alpha T \exp(-\alpha T)/(1 - \exp(-\alpha T))$. Then, the optimal solution $(\rho^*, \hat{P}^*)$ satisfies $1 - \hat{P}^* \leq \rho^* \leq (1 - \hat{P}^*)\exp(wT)$. Furthermore,*

Case 1. *If $\bar{Q}/(\lambda T) \leq (1-\delta)/(2-\delta)$, then*

$$V_D = (\lambda T \exp(-\alpha T)/\delta)\rho^* \hat{P}^*$$

where $\rho^* = 1 - \hat{P}^* = \bar{Q}/(\lambda T)$.

Figure 6 **Auctioneer's Revenues as a Function of T for Three Values of $\hat{P}$**

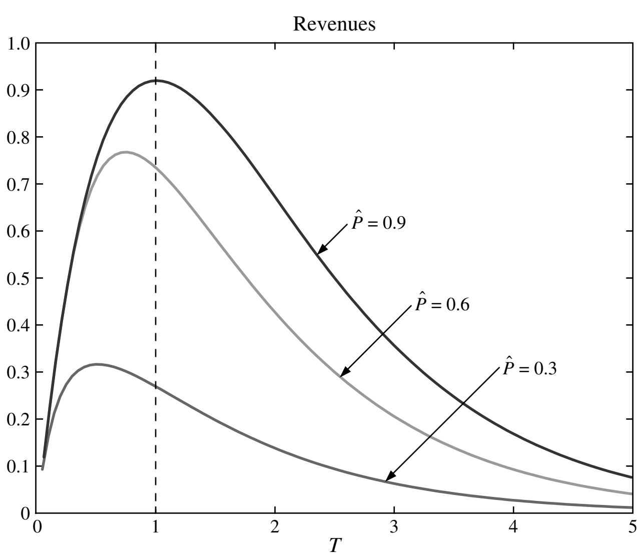

Note. Data: $w = \alpha = 1$ and $\lambda = 10$.

Case 2. *If* $\bar{Q}/(\lambda T) \geq (1-\delta)/(2-\delta)$, *then* V_D *is bounded by*

$$V_1 \triangleq \frac{\lambda T \exp(-\alpha T)}{\delta}\rho_1 P_1 \leq V_D$$
$$\leq \frac{\lambda T \exp(-\alpha T)}{\delta}(P_2^2 - (1-\rho_2)^2\delta) \triangleq V_2,$$

where $\rho_1 = 1 - P_1 = \min\{1/2;\ \bar{Q}/(\lambda T)\}$, $\rho_2 = \min\{\bar{Q}/(\lambda T);\ (3-\delta)/(4-\delta)\}$ *and* $P_2 = (1/2)[1+(1-\rho_2)\delta]$.

The lower and upper bounds V_1 *and* V_2, *respectively, are asymptotically tight in the following sense*

$$\lim_{w\downarrow 0} V_D = V_1, \quad \lim_{w\downarrow 0}\rho^* = \lim_{w\downarrow 0}(1-\hat{P}^*) = \rho_1 \quad \text{and}$$
$$\lim_{w\to\infty} V_D = V_2, \quad \lim_{w\to\infty}\rho^* = \rho_2, \quad \lim_{w\to\infty}\hat{P}^* = P_2.$$

(b) *The optimal time* T^* *is bounded above by the unique nonnegative root of*

$$\frac{\bar{Q}}{\lambda T} = \frac{1-\exp(-\alpha T) - \alpha T\exp(-\alpha T)}{2 - 2\exp(-\alpha T) - \alpha T\exp(-\alpha T)}.$$

The optimality condition $\rho^* = 1 - \hat{P}^*$ in Case 1 implies that the auction price equals the fixed price, and so all units are sold through the fixed-price channel. Note that δ is a decreasing function of αT, and so the condition $\bar{Q}/(\lambda T) \leq (1-\delta)/(2-\delta)$ tends to be satisfied for small values of $\bar{Q}$ and large values of αT. If $\bar{Q}$ is small, then the seller can target this limited inventory to high-valuation buyers, and there is no need for an auction. On the other hand, if αT is large, then any revenue collected in the auction will be penalized by a small discount factor $\exp(-\alpha T)$, and so the seller has no economic incentive to run an auction.

In Case 2, when $\bar{Q}/(\lambda T) \geq (1-\delta)/(2-\delta)$, it is in the seller's best interest to operate both channels. In this case, we cannot get a closed-form solution, but we derive upper and lower bounds for the seller's revenue. The lower bound V_1 is derived, assuming that the seller chooses (suboptimally) to operate only the fixed-price channel (i.e, $\rho_1 = 1 - P_1$). The upper bound V_2 is obtained, assuming that all buyers with valuation greater than the posted price will buy in this channel, independently of the expected auction price.

Interestingly, these bounds are asymptotically optimal. Intuitively, the first asymptotic result implies that as consumers become increasingly patient ($w \downarrow 0$), the seller is unable to segment the population using the two channels. Consumers are willing to wait for the auction if the auction price is below the fixed price. Hence, in the limit $w \downarrow 0$, the auction price must equal the fixed price, and all units are sold through the fixed-price channel. This is the worst possible scenario for the seller because the revenue reaches the lower bound V_1. From a practical standpoint, we can interpret this result as suggesting that the seller should choose a fixed price and initial inventory so that most of this inventory is sold through the fixed-price channel.

On the other hand, if consumers become increasingly impatient ($w \to \infty$), then the seller can perfectly segment these consumers among those that can afford to pay the fixed price and those that cannot. This gives the seller the ability to separate consumers exclusively based on their valuation and to optimally design the fixed price and the auction to achieve the maximum possible revenue V_2. The auction in this case is an effective selling mechanism that complements the fixed-price operation well.

The previous discussion raises an important question regarding the value of a dual-channel operation versus a traditional fixed-price channel. To address this question, we will compare the expected revenue collected using both channels, V_D, and the one collected using the fixed-price channel only. Note that if the seller uses only the fixed-price channel, she will choose a fixed price equal to $\max\{1/2;\ 1-\bar{Q}/(\lambda T)\}$ and collect a revenue equal to V_1 in Proposition 6. Figure 7 plots the ratio V_D/V_1 as a function of $\bar{Q}/(\lambda T)$ (left panel) and w (right panel) for three different values of α. Consistent with the results in Proposition 6, when $\bar{Q}/(\lambda T)$ or w are small or α is large, the ratio V_D/V_1 is close to one. Hence, in these cases there is little advantage in running an auction. On the other hand, for small values of α or large values of $\bar{Q}/(\lambda T)$ or w, the seller can significantly increase her discounted revenue by operating both channels (close to 22% in the examples in Figure 7). We can get an upper bound on this relative revenue increase by replacing V_D by the upper bound V_2. After some straightforward manipulations, combining Case 2 in Proposition 6 and the fact that $\delta \leq 1$, we get that

$$1 \leq \frac{V_D}{V_1} \leq \frac{V_2}{V_1} \leq \frac{4}{4-\delta} \leq \frac{4}{3},$$

that is, the seller can increase her revenues by as much as 33% by adding a terminal auction to a fixed-price operation. This upper bound is reached in the limit as $w \to \infty$ and $\bar{Q} \to \lambda T$.

6. Concluding Remarks

In this paper, we have proposed a model to analyze the problem faced by a seller when designing a single-channel online auction, or when managing a dual online auction and list price channel. The key to building this model is to understand consumers' strategic behavior, provided they choose to either join the auction or buy the product at the posted price.

Figure 7 **Optimal Dual-Channel Revenue (V_D) over Optimal Single Fixed-Price Channel Revenue (V_1)**

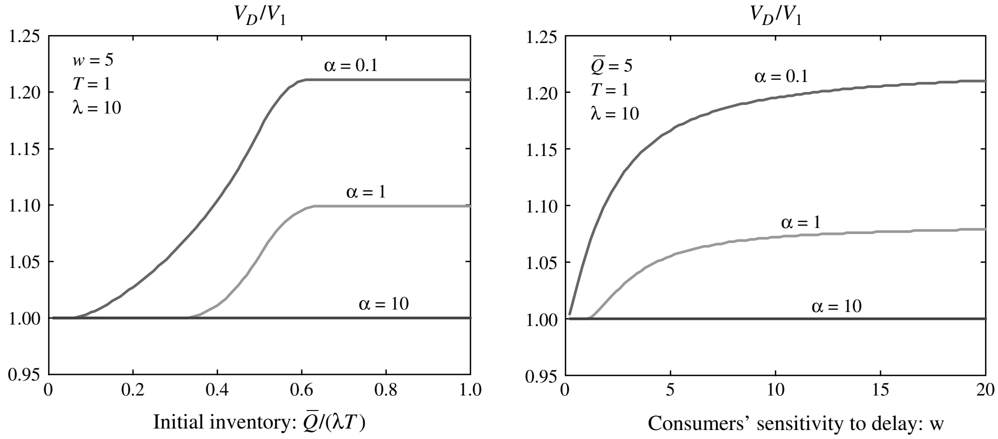

The private information of the consumers has two dimensions: the arrival time and the private value for one of the units being offered. Using a time-sensitive utility function, we showed that their participation equilibrium strategy is of the threshold type, that is, a consumer will join the auction if and only if his arrival time is higher than a function of his own valuation. Of course, for consumers with values below the list price, the optimal strategy is always to participate in the auction. For consumers with higher values, the threshold is nondecreasing in their own valuation. Interestingly, we found that there is always a range of values greater than the fixed price for which it is also optimal to enter the auction, regardless of the arrival time. We also found that, in equilibrium, the auction bidding rate is an increasing function of time (sniping); a phenomenon that is frequently observed in practice. In this respect, our model provides a simple explanation for this late bidding behavior based on consumers' sensitivity to delay.

At a theoretical level, we proved that a symmetric participation equilibrium always exists. We also proposed a contraction algorithm in a function space to find this threshold and proved its convergence under some conditions. When these conditions are not satisfied, we managed to find the fixed point of the algorithm by slightly perturbing the values of each iteration.

In general, the exact algorithm is computationally intensive and does not lead to clean managerial insights. To overcome these limitations, we proposed an asymptotic analysis for both settings. In this asymptotic regime the initial number of units and the demand rate grow proportionally large. In the limit, we showed that there is weak convergence to a unique equilibrium that we were able to characterize in closed form. Using numerical experiments, we also showed that this solution is indeed a good approximation even for moderate values of the inventory and arrival rate. Because of the simplicity and accuracy of the asymptotic analysis, we proposed to solve the seller's optimization problem using this limiting regime. In the case where the seller controls only the auction, we found that the optimal number of units to offer is a nonmonotonic function of the fixed price $\hat{P}$. Although we could not characterize this optimal supply in closed form, we derived a simple analytical expression for it and showed that it is always bounded above by 80% of the average demand. In terms of the optimal duration of the auction, we found that this optimal time is an increasing function of the fixed price, and it is bounded above by α^{-1} (the inverse of the seller discount factor). This upper bound is attained if $\hat{P} \geq 1 - \exp(-w/\alpha)/2$.

In the case in which the seller is a monopolist controlling both channels, we found that the auction is only useful if the excess-supply ratio $\bar{Q}/(\lambda T)$ exceeds the threshold $(1-\delta)/(2-\delta)$ (see Proposition 6). Hence, if the initial endowment $\bar{Q}$ is small or the discount factor α is large, then the seller does not have enough economic incentives to run a terminal auction; a single fixed-price channel is the optimal selling mechanism. On the other hand, if the excess-supply ratio is sufficiently large, then running both channels in parallel is optimal. In the latter case, we showed that a dual-channel operation can have a significant impact on revenues compared to a single fixed-price channel. The magnitude of the increase in revenues can be as large as 33% (for the case of uniformly distributed valuations), an upper bound that is reached in the limit as consumers' sensitivity to delay w grows large.

We believe that the techniques used to derive the equilibrium participation strategy here, in particular the asymptotic analysis, can also be used when extending our model to incorporate the number of remaining units as part of the information structure of the bidders. That would mean adding one coordinate to the threshold, leading to a three-dimensional surface, but we have not explored this direction in detail.

Other possible extensions are related to the seller's optimization problem. In our formulation, we have included the capital cost. However, for example, in the dual-channel case, one can easily add some holding cost for keeping the units until the end of the time horizon, such that the seller has an incentive to give more units through the list price channel. One could also add a penalty cost for keeping units at the end of the horizon, or a shortage cost for not being able to serve a consumer if the units are depleted before the time scheduled for the auction.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

The authors thank Victor Araman (New York University), Dorothée Honhon (University of Texas at Austin), Costis Maglaras (Columbia University), and Anshul Sheopuri (New York University) for helpful discussions.

References

Bertsekas, D., J. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5** 203–229.

Budish, E., L. Takeyama. 2001. Buy price in online auctions: Irrationality on the Internet? *Econom. Lett.* **72**(2) 325–333.

Bulow, J., J. Roberts. 1989. The simple economics of optimal auctions. *J. Political Econom.* **97**(3) 1060–1090.

Cheney, W. 2001. *Analysis for Applied Mathematics*. Springer-Verlag, New York.

Durrett, R. 1996. *Probability: Theory and Examples*. Duxbury Press, Belmont, CA.

Etzion, H., E. Pinker, A. Seidmann. 2006. Analyzing the simultaneous use of auctions and posted prices for online selling. *Manufacturing Service Oper. Management* **8**(2) 68–91.

Gallien, J., S. Gupta. 2007. Temporary and permanent buyout prices in online auctions. *Management Sci.* **53**(5) 814–833.

Glazer, A., R. Hassin. 1983. ?/m/1: On the equilibrium distribution of customer arrivals. *Eur. J. Oper. Res.* **13** 146–150. .

Harstad, R. 1990. Alternative common-value auction procedures: Revenue comparisons with free entry. *J. Political Econom.* **98** 421–429.

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

Krishna, V. 2002. *Auction Theory*. Academic Press, New York.

Lariviere, M., J. Van Mieghem. 2004. Strategically seeking service: How competition can generate Poisson arrivals. *Manufacturing Service Oper. Management* **6** 23–40.

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *Amer. Econom. Rev.* **89** 1063–1080.

Mas-Colell, A., M. Whinston, J. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.

Maskin, E., J. Riley. 1989. Optimal multi-unit auctions. F. Hahn, ed. *The Economics of Missing Markets: Information and Games*. Clarendon, Oxford, UK, 312–336.

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.

Pinker, E., A. Seidmann, Y. Vakrat. 2003. Managing online auctions: Current business and research issues. *Management Sci.* **49** 1457–1484.

Roth, A., A. Ockenfels. 2002. Last minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econom. Rev.* **92** 1093–1103.

Shmueli, G., R. Russo, W. Jank. 2004. Modeling bid arrivals in online auctions. Working paper, Smith School of Business, University of Maryland, College Park, MD.

Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston, MA.

Vakrat, Y., A. Seidmann. 1999. Can online auctions beat online catalogs? *Proc. 20th Internat. Conf. Inform. Systems*. Association for Information Systems, Atlanta, GA, 132–143.

van Ryzin, G. J., G. Vulcano. 2004. Optimal auctioning and ordering in an infinite horizon inventory-pricing system. *Oper. Res.* **52** 346–367.

Vany, A. De. 1987. Institutions for stochastic markets. *J. Institutional Theoret. Econom.* **143** 91–103.

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

Wang, W. 1993. Auctions versus posted-price selling. *Amer. Econom. Rev.* **83** 838–851.