M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 11, No. 1, Winter 2009, pp. 33–51
ISSN 1523-4614 | EISSN 1526-5498 | 09 | 1101 | 0033
informs
DOI 10.1287/msom.1070.0192

A Pooling Analysis of Two Simultaneous Online Auctions

Damian R. Beil
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109,
dbeil@umich.edu

Lawrence M. Wein
Graduate School of Business, Stanford University, Stanford, California 94305,
lwein@stanford.edu

Motivated by the ease with which online customers can bid simultaneously in multiple auctions, we analyze a system with two competing auctioneers and three types of bidders: those dedicated to either of the two auctions and those that participate simultaneously in both auctions. Bidding behavior is specified and proven to induce a Bayesian Nash equilibrium, and a closed-form expression for the expected revenue of each auctioneer is derived. For auctioneers selling a single item, partial pooling—i.e., the presence of some cross-auction bidders—is beneficial to both auctioneers as long as neither one dominates the market (e.g., possesses more than 60%–65% of the market share). For multi-item auctions, pooling is mutually beneficial only if both auctioneers have nearly identical ratios of bidders per items for sale; otherwise, only the auctioneer with the smaller ratio benefits from pooling. Pooling's impact on revenue decreases with the number of bidders, suggesting that popular auction sites need not be overly concerned with mitigating bidding across auctions.

Key words: auctions; bidding; pooling
History: Received: October 2, 2006; accepted: July 6, 2007. Published online in *Articles in Advance* January 4, 2008.

1. Introduction

In this paper we analyze partial pooling in business-to-consumer (B2C) auctions on the Internet, where partial pooling occurs because bidders may simultaneously participate in more than one auction. Partial pooling arises from a reduction in search costs enabled by the Internet (see Pinker et al. 2003 for a discussion of search costs and for a review of Internet auction research relevant to operations management). Anwar et al. (2006) provide empirical evidence of partial bidder pooling in eBay business-to-consumer CPU auctions, where approximately 20% of bidders facing competing auctions selling the same item participate in multiple auctions. In addition, consumers are increasingly using intelligent software, or *shopbots*, to scan the Internet for auctions selling the particular item they want; although in 2002 only about 6% of Internet shoppers used a shopbot, the percentage has steadily increased since 1997 (Montgomery et al. 2004). One such tool, BidXS, currently searches about 300 consumer auctions.

In an attempt to understand the impact of partial pooling, we analyze a simple setting in which two auctioneers use simultaneous open-ascending (English) auctions to sell the same commodity to three types of bidders: those devoted to each individual auction and those participating simultaneously in both auctions (what we call shared bidders). Our main objective is to examine the expected revenue impacts of partial pooling under a proposed bidding strategy: Bidders not currently winning an item bid the minimum amount necessary to take the lead in the lowest-priced auction they have access to, provided such a bid would not exceed their private valuation.

It is not obvious whether auctioneers benefit from partial pooling. On one hand, the selling price in traditional auctions increases with the number of bidders, which suggests that competing Internet auctioneers selling the same commodity at the same time might benefit from partial pooling. But pooling is a two-way street: These additional bidders have access to the items of competing auctioneers. A few simple examples illustrate the flavor of the auction pooling problem.

Ignoring reserve prices, suppose single-item auction 1 had bidders with valuations of \$10, \$6, and \$5, and single-item auction 2 had bidders with valuations of \$9, \$7, and \$3. If each auction was run in isolation, i.e., no pooling, the closing price in auction 1 would be \$6 and that of auction 2 would be \$7. If instead all six bidders are shared, i.e., full pooling, then because under the proposed bidding strategy the \$7 bidder participates in both auctions and drives up the price in auction 2, both auctions close at \$7. Here pooling is mutually beneficial: auctioneer 2 is indifferent to the pooling and auctioneer 1 clearly gains. But the pooling picture is not always so rosy.

Instead suppose that auction 2 had bidder valuations of \$4, \$3, and \$2. The closing prices for isolated auctions would be \$6 in auction 1 and \$3 in auction 2. Under full pooling both would close at \$5, because the \$10 bidder would transact in auction 1 and the \$6 bidder in auction 2 (or vice versa) and thus would escape competing against each other for a single item in auction 1. In this situation, full pooling makes auctioneer 2 better off, but auctioneer 1's price is lowered as a result of less bidding competition for his item, and he is worse off.

Finally, consider a partial-pooling case in which only the \$5 and the \$4 bidders are shared bidders. Unlike full pooling, this partial pooling situation is mutually beneficial: auction 1 closes at \$6 and auction 2 at \$4, because under our proposed bidding strategy the \$5 bidder drives up the price in auction 2, but the \$10 and \$6 bidders are still forced to compete for auctioneer 1's single item.

Thus, from a revenue-maximizing auctioneer's ex ante perspective (the perspective we adopt in this paper), we see that full or partial pooling could help or hurt, depending on the distribution of bidder valuations and the subset of total bidders who are shared bidders. We focus our analysis on the more general and realistic (Anwar et al. 2006 find that approximately 20% of bidders are shared) partial-pooling case, whose analysis subsumes the full pooling case. More generally, given the size of the three bidder subsets and a prior belief over bidder valuations, we apply our closed-form expected revenue results to characterize when partial pooling is mutually beneficial for the auctioneers. Although auctioneers have very limited discretion over the amount of partial pooling that occurs, they do have several options to influence it on the margin. To discourage pooling, auctioneers could change the timing of their auctions to avoid competing auctions. To encourage pooling, auctioneers could share their bidder lists with other auctioneers or advertise each others' auctions on their websites.

This paper examines how the amount of partial pooling between competing auctions affects expected auctioneer revenue. A prelude to our closed-form expected revenue results is a bidding equilibrium analysis reminiscent of Walrasian tatonnement, as in other multi-item auction studies (e.g., Demange et al. 1986, Gul and Stacchetti 2000, Milgrom 2000). However, unlike these studies, our equilibrium is built around a model of *decentralized* auctioneers who do not jointly give feedback to bidders; therefore, our equilibrium is an extension of Peters and Severinov (2006), who assume decentralized auctioneers but consider only full (not partial) pooling and focus on economic efficiency. More to the point, none of these papers (nor any others to our knowledge) examines the expected auctioneer revenue impact of various levels of partial pooling, which is our goal. Also, although these studies provide various characterizations for ex post revenue (as the optimal dual solution of an assignment LP, for instance), we study expected or ex ante revenue for different bidder valuation distributions, which requires detailed calculations using order statistics.

This paper is organized as follows. In §2 the model is described and our proposed bidder behavior is specified and shown to induce a Bayesian Nash equilibrium. A closed-form expression for the auctioneers' expected revenue is derived in §3. Sections 4 and 5 investigate the impact of pooling in single-item and multi-item settings, respectively. Concluding remarks are offered in §6.

2. The Model

In this section we describe our model of two simultaneous competing auctions. Section 2.1 describes notation and the simultaneous auction game we use to characterize bidders' equilibrium behavior—essentially that bidders bid in the lowest-priced auction they have access to and only bid the minimum amount required to take the lead. For readers interested in game theoretic details of the bidding strategy, §2.2 provides formalizations and illustrations, although such details are not required to understand the paper's subsequent sections. Two additional assumptions made for the price analyses are presented in §2.3.

2.1. Simultaneous Auction Game

For $i = 1, 2$, auctioneer i has m_i identical items for sale and has d_i (mnemonic for dedicated) bidders participating solely in auction i (see Figure 1); we let $d = d_1 + d_2$ denote the total number of dedicated bidders. In addition, there are s *shared*, or *pooled*, bidders who participate simultaneously in both auctions. To avoid introducing new notation, we sometimes refer to these groups of bidders simply as the "d_i bidders" or the "s bidders." We place no restrictions on the sizes of d_1, d_2, and s. The $d + s$ bidders each seek a single item, and each bidder's private valuation for the item is an independent draw from a common valuation distribution. We assume both auctions use an English format, as is typical on the Internet (Beam and Segev 1998, Lucking-Reiley 2000), and we simplify interauction dynamics by assuming in our model that both auctions start and end at precisely the same times.

Behavior in the competing auction game is sufficiently complex that a somewhat detailed description of the game is required. Bidding proceeds in turns as follows. We say that a bidder enters the market the moment she places her first bid in an auction. After every entry to the market, each previously entered bidder is given at least one turn, i.e., an opportunity to submit a new bid or to pass. Following this round of bidding, a new bidder enters the market, provided there is a bidder willing to do so, and the next round of bidding begins.

When auctioneer i receives a new bid, he (i) posts the current *standing bid* of his auction equal to the $(m_i + 1)$st highest submitted bid he has received or his reserve price if he has received fewer than $m_i + 1$ bids exceeding his reserve price, and (ii) announces the bidder's identity and the identities of the winning bidders (but not values of the winning bids). All bids exceeding the current standing bid are winning bids. If there are only $m_i - m < m_i$ such bids, the remaining winning bids are the earliest—up to the mth earliest—bids that equal the current standing bid. If only $k < m_i$ bids have been submitted in the auction, the auctioneer does not assign any winner to the remaining $m_i - k$ items. All new bids must exceed the current standing bid by at least the minimum bid increment δ. We call a bidder's bid *successful* if the bidder is a winner after submitting the bid; otherwise, we say the bid is *unsuccessful*. When submitting a new bid, the bidder can designate it as a replacement, or *revision*, for one of her current winning bids, if any exist. Bids can be revised upwards only, and a revision's time stamp supersedes that of the replaced bid. The standing bid is assumed to be the $(m_i + 1)$st highest bid submitted, ignoring any revised (i.e., replaced) bids; hence, if a bidder revises one of her winning bids, the standing bid does not change.

The auction proceeds as described above until there are no new bidders who wish to enter the market and two or more rounds of bidding pass with no new bids submitted. At this point, for each winning bid in auction i, the bidder who submitted it is awarded one item and pays auctioneer i a transaction price equal to the final standing bid in auction i.

We assume the public information available to bidders with access to auction i at some stage t (this can be thought of as the tth turn in the procedure described) to be the auction i history of standing bids, and winner and bidder identities and time stamps, up to and including stage t. This resembles the public information of eBay and Yahoo! auctions. All bidders have an identical, known payoff function structure equal to their private valuation for a single item minus the sum of their transaction prices, or zero if they do not transact. This implies that a bidder is

Figure 1 The Pooling Model

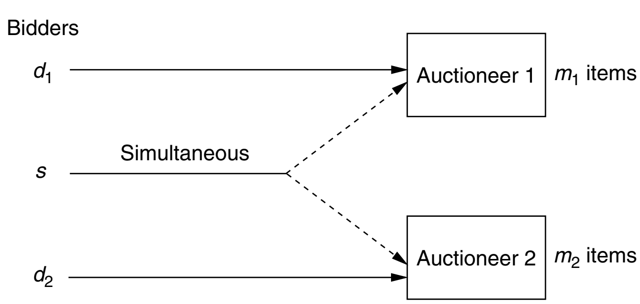

indifferent between either not transacting or winning an item and paying her valuation. Here the assumption that bidders seek a single item is taken to mean that a bidder's valuation for multiple items is the same as their valuation for a single item. It is known that an auctioneer's payoff for each item sold is the item's transaction price minus the auctioneer's cost, or zero for items not sold.

In reality, issues such as currency denominations lead auctions to incorporate the minimum bidding increment $\delta > 0$. To avoid complex bid-ordering issues and to allow bidders to bid up to their true valuations, we assume that valuations and bids are discrete and restricted to lie in the grid $\{0, \delta, 2\delta, \ldots\}$.

For the equilibrium analysis in the following subsection we treat the case $\delta > 0$. With the equilibrium result in hand, we then begin analyzing the auction prices that arise (§§3–5), albeit with two additional assumptions discussed in §2.3: We let $\delta \to 0$, assume continuously distributed valuations having a probability density function, and ignore reservation prices.

2.2. Bidder Behavior

It is a well-known (e.g., the survey paper by Klemperer 1999) equilibrium for bidders seeking a single item in a traditional English auction to bid up to their true valuation. This reasoning applies in our model but alone is insufficient; equilibrium bidding for shared bidders, who have a choice about which auction(s) to participate in as they bid up to their true valuation, makes things more complex. This subsection begins by formalizing σ^*, the equilibrium bidding strategy for the simultaneous auction game. Following this, we use examples to motivate the bidding strategy used by shared bidders, discuss behavioral assumptions, and conclude with a result stating that σ^* is indeed an equilibrium.

2.2.1. Strategy Formalization. We define the strategy σ^* as follows. If it is a dedicated bidder's turn to bid, then

(a) If the bidder is currently winning an item in an auction with a bid exceeding the auction's current standing bid by the minimum bid increment δ, then the bidder should pass.

(b) Otherwise, if the bidder is currently winning no item or is winning an item and her winning bid does not exceed the auction's current standing bid by at least δ, the bidder should submit a bid equal to the standing bid plus δ (designating this bid as a revision if the latter case holds), provided such a bid does not exceed the bidder's valuation. If such a bid would exceed the bidder's valuation, the bidder should pass. If it is a shared bidder's turn to bid, then

(c) If the bidder is currently winning an item in an auction, or if both the current standing bids equal or exceed the bidder's valuation, then the bidder should pass.

(d) Otherwise, if auction i has the unique lowest standing bid, submit in auction i a bid that exceeds this standing bid by δ.

(e) Otherwise, if both standing bids are equal and if in the time since the bidder's previous turn all bidders passed on their turns and no new bidder entered the market, if auction i is the unique auction in which the most recent m_i bids have been successful, the bidder should submit a bid according to (d) in auction $-i$. (Here $-i$ denotes the complement auction; e.g., if $i = 2$, $-i$ equals 1.) If the most recent m_i bids in both auctions $i = 1, 2$ have been successful, then the bidder should submit a bid according to (d) in either auction with equal probability. Otherwise, the bidder should pass.

Since under σ^* bidders bid incrementally by the minimum bid increment δ, the winning bids in an auction can assume only two values: the standing bid or δ above the standing bid. Furthermore, every unsuccessful bid raises the standing bid by δ. These observations pertain to part (e) of σ^*, which essentially requires knowledge of the winning bids, which are not announced by the auctioneer. The winning bids can be inferred using the history of successful and unsuccessful bids, which in turn can be inferred from the history of bidder and winning bidders' identities and time stamps. (The most recent bid was successful if and only if the current list of winners contains the most recent bidder.)

Consider auction i. Suppose its current standing bid is p, and suppose in auction i there have been exactly k successful bids since the most recent unsuccessful bid (or since time zero if all bids have been successful). Since under σ^* bidders bid incrementally by the minimum bid increment δ, k of auction i's winning bids equal $p + \delta$ and any remaining winning bids (if they exist) equal p.

2.2.2. Strategy Illustration. Figure 2 provides a graphic backdrop for explaining the shared bidders' equilibrium bidding strategy when facing single-item auctions. Consider a sole shared bidder with no winning bids and a valuation of \$4 facing standing bids of \$1 in auction 1 and \$2 in auction 2, Figure 2(a). Clearly the shared bidder could bid \$2, \$3, or \$4 in auction 1 and/or \$3 or \$4 in auction 2 (in this example $\delta = \$1$). Suppose that she submits a single successful bid of \$4 in auction 1; unknown to her, she is facing competing dedicated bidder valuations of \$3 and \$4 in auction 1 and \$2 and \$3 in auction 2, so eventually she will transact in auction 1 at a price of \$4 (under strategy σ^*(a)–(b), dedicated bidders bid up to their true valuation in whichever auction they have access to and place a bid even when indifferent between bidding and doing nothing). In contrast, the bidder could have transacted at \$3 in auction 2 had she instead bid *conservatively*, which we define as follows: If a shared bidder does not currently have a winning bid, she submits an *incremental* bid, i.e., one just large enough to take the lead in the auction with the lower standing bid, as long as this bid does not exceed her own valuation. Conservative bidding aims to avoid inadvertently aggressive participation in the more competitive of the two auctions, and is captured by σ^*(c)–(d). However, conservative bidding is alone not enough to prevent such ensnarement.

Figure 2 **Standing Bids and Bidder Valuations**

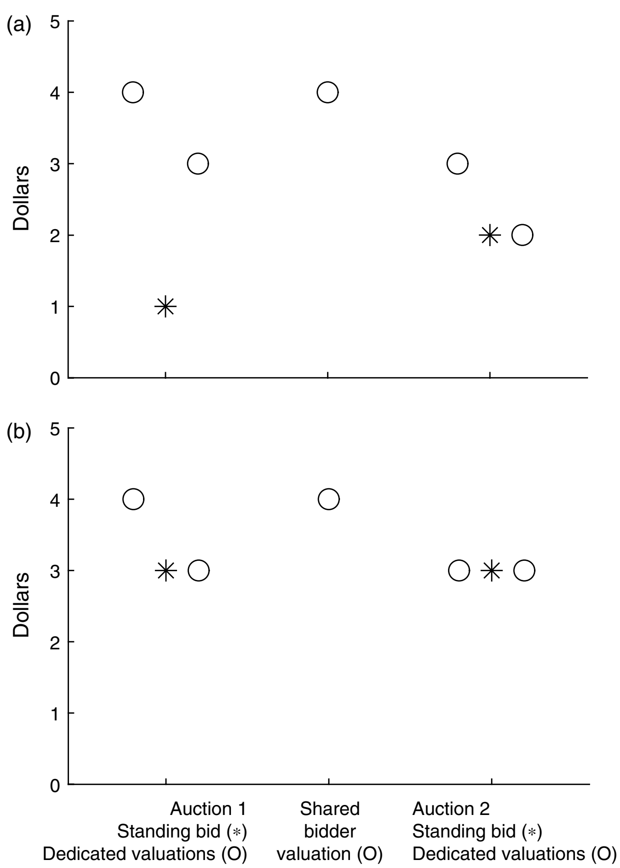

Now consider the shared bidder instead facing standing bids of \$3 in both auctions. Suppose the current winning bids in both auctions are also \$3. Clearly, the shared bidder could bid \$4 in either auction. If, unknown to her, she is competing against dedicated bidder valuations as graphed in Figure 2(b), a bid of \$4 would be victorious in either auction but would result in a payment of \$4 in auction 1 and only \$3 in auction 2. To avoid being misled to a premature bid (as a bid of \$4 in auction 1 would be), the equilibrium bidding strategy for shared bidders is *patient* (σ^*(e)): While standing bids in both auctions are equal, the shared bidder casts no bid unless both auctions' dedicated bidders have all completed their bidding (in a round of the game, no new bidders enter and all bidders pass on their turn to bid); if both auctions' standing bids are equal and all dedicated bidders have completed their bidding, the bidder submits a bid just high enough to take the lead in the auction with the lowest winning bid, or she places one such bid in a randomly chosen auction if the two auctions' winning bids match. With such patience, in our example, auction 1's winning bid would reach \$4 before the shared bidder casts her bid, and she would be sure to bid \$4 in auction 2 instead of auction 1.

To summarize, under the equilibrium bidding strategy σ^*, all bidders bid up to their true valuation incrementally, but shared bidders do so conservatively and patiently. It may seem that by bidding earlier the dedicated bidders enjoy a slight advantage over patient shared bidders, but Proposition 1 shows that in fact this is part of an equilibrium strategy for all bidders. (Although this result does not rule out other possible equilibria in which dedicated bidders do not bid earlier, we leave this direction to future work.) Thus, the equilibrium bidding strategy is more complex than the well-known "bid your valuation right off the bat" strategy that is dominant for isolated, second-price auctions. Shared bidders monitor both auctions and dynamically select the cheaper auction to bid in; this naturally relates the revenue of one auction to the other; how this plays out is the crux of this paper's main analyses, beginning in the next section.

2.2.3. Bidding Assumptions Discussion. Our model clearly is an abstraction meant to capture the salient features of simultaneous auctions; four key assumptions on information, behavior, and model structure deserve attention before moving on:

(i) The informational assumption that shared bidders can bid patiently by detecting when dedicated bidding stops proves useful for the equilibrium bidding analysis but is unlikely to hold in practice, where bidders do not necessarily bid in a round-by-round fashion and dedicated bidders might bid patiently. However, this informational assumption has little impact on the practical predictions of the model as long as shared bidders bid conservatively: If it is the only assumption that fails to hold in practice, shared bidders will overpay by at most δ and hence the model will misestimate the auctions' closing prices by at most δ. This is akin to Demange et al. (1986) and Milgrom (2000), who show that multiunit auction outcomes with myopic conservative bidding strategies approximate the competitive outcome.

(ii) Our model ignores bidding costs, and conservative bidding might be prohibitively expensive in practice if bidding is costly. Nonetheless, empirical evidence from eBay in Anwar et al. (2006) suggests that (what we call) shared bidders tend to submit bids in the auction with the lowest standing bid and to bid more frequently near the end of the auctions, behaviors consistent with conservative (and patient) bidding. More generally, a richer equilibrium model, which is beyond the scope of this paper, would require the incorporation of search costs (e.g., the cost of using intelligent search software) from both bidders' perspectives and the software's perspective (i.e., optimal pricing of these services); see Bakos (1997) and Kephart and Greenwald (2002) for research on when sellers post prices rather than operate auctions, and see Hann and Terwiesch (2003) for an empirical estimate of these costs in an Internet auction setting.

(iii) Our simplified model assumes d and s bidder valuations are identically distributed. In reality, shared bidders might value the item more than their dedicated counterparts and be willing to participate in more auctions to win an item when few are available for sale (e.g., Beanie Babies). In contrast, shared bidders might have lower valuations and participate in multiple auctions to cheapen their transaction price, especially if search costs are low and items are abundant (e.g., CPUs).

(iv) Another simplification used in our model is its structural assumption of coincident auctions. In their empirical analysis, Anwar et al. (2006) observe that shared bidding activity is diminished as auction ending times become dissimilar, but they do not find it eliminated even for auctions ending a day apart, suggesting that the shared bidder model is still relevant if our coincidence assumption fails to some degree in practice. For modeling, s could be decreased as a proxy to capture dissimilar auction ending times. In the literature, Lavi and Nisan (2005) study multi-item auctions in which items must be sold before dissimilar expiration times.

2.2.4. Equilibrium Result. Our paper analyzes the impact of the amount of partial pooling on auctioneers' expected revenues (or prices), for example, determining conditions under which expected prices increase if the s compartment grows at the expense of the d_i compartments (in such situations, auctioneers would want to encourage pooling, causing dedicated bidders to become shared bidders). As such, the equilibrium analysis of this section is merely a prelude to our expected price analyses. Nonetheless, we include it because existing equilibrium models cannot be directly applied to our setting, as we describe.

Our model is related to auction assignment mechanisms for heterogeneous goods (e.g., Demange et al. 1986). In Demange et al.'s auction, the auctioneer offers prices and bidders report their demand set. Eventually the mechanism terminates in minimal Walrasian prices, resulting in an ex post equilibrium (Gul and Stacchetti 2000). However, these models do not capture ascending auctions, where bidders make binding bids on individual items rather than report demand sets (Milgrom 2000), and *decentralized* auctioneers cannot centrally offer prices and process bidders' demand reports. Only one previous study (to our knowledge) (Peters and Severinov 2006) tackles both issues in an equilibrium analysis, and our equilibrium strategy is an extension of that (we add dedicated bidders and patient bidding by shared bidders). (Others, such as Milgrom 2000, study a nonstrategic bidding algorithm [essentially conservative bidding] that does not rely on the centralized feedback assumption.)

Our bidding model has three notable differences from that studied by Peters and Severinov (2006): Ours includes dedicated bidders, auctioneers can sell multiple items instead of just single items, and we consider only two simultaneous auctions instead of many. The first difference is the most important, as our inclusion of dedicated bidders creates asymmetry between the auctions' ending prices, which is absent in Peters and Severinov, where all bidders are shared. Unfortunately, this prevents us from greatly leveraging Peters and Severinov's analysis in our equilibrium proof. Because we avoid assuming symmetry, the effect of bidder strategy deviations intricately depends on the particular combinatorial mix of valuations across the shared and dedicated bidders. Even computing equilibrium expected prices is a complex task in our model (Proposition 3), in contrast to the setting in which all bidders are shared, where both prices would simply equal the $(m_1 + m_2 + 1)$st highest trader's valuation. More to the point, Peters and Severinov (2006) do not analyze partial pooling, while our paper focuses on analyzing how the amount of partial pooling impacts auctioneers' expected revenues (or prices). In Appendix A (appendixes online) we show that no bidder can benefit by unilaterally deviating from the bidding strategy outlined above.

Proposition 1. *Bidding strategy σ^* induces a Bayesian Nash equilibrium in the simultaneous English auction game.*

Although σ^* resembles the most common equilibrium assumption found in the literature on isolated English auctions, namely incrementally bidding by the minimum bid increment and dropping out of the auction at one's true value, additional equilibria might exist. For example, in the literature on isolated English auctions, papers have examined equilibria involving strategic jump bids for signalling purposes under costly bidding (Daniel and Hirschleifer 1998) or affiliated values (Avery 1998), or to take advantage of a large minimum bid increment (Isaac et al. 2007).

2.3. Two Additional Assumptions for Expected Price Analyses

Under the above bidder behavior, the resulting selling prices in the two auctions involve the ordering of order statistics, as seen in the next section. Before moving to analyses of expected auction prices, which comprise the remainder of the paper, we will make and carry forward two additional assumptions.

First, because continuous order statistics are easier to analyze than their discrete counterparts, we find it mathematically convenient to analyze the limiting case $\delta \to 0$. That is, in the sequel we assume that bidder valuations are continuously distributed and have a probability density function and that the bidding increment is arbitrarily small. For example, our analysis of an auction's expected price when assuming that bidder valuations are distributed according to cdf $F \sim U[0, 1]$ can be thought of as an approximation to the auction's expected price for a game featuring a discrete cdf F_δ that assigns equal probability to any valuation in the discrete mesh $\{0, \delta, 2\delta, \ldots, 1-\delta, 1\}$, where δ is small.

Second, for simplicity, in the sequel we assume that the auctioneers' reserve prices are set to zero. Clearly, in the simultaneous English auction an optimal reserve price in one auction must be predicated on the auctioneer's cost and the reserve price of the other auction. Under their model with multiple auctioneers, Peters and Severinov (2006) apply an asymptotic analysis and prove that auctioneers' optimal reserve prices must converge to their true costs as the number of bidders and auctioneers become large. Optimal reserve prices would likely exceed true costs in our two-auction setting, but we leave such analysis for future research. Here we focus on the revenue effects of interbidder competition, effects that change according to the extent to which bidders do or do not participate in both auctions. Note that using a reserve price in a single-item, independent-private-values, English auction generates less revenue than adding an additional bidder (Bulow and Klemperer 1996).

3. Expected Price in Partially Pooled Multi-Item Auctions

Because the selling price need not be identical for the two auctions, for convenience we adopt the viewpoint of auctioneer 1 and denote the price of auction 1 by $\Pi(d_1, d_2, s, m_1, m_2)$. The analysis of $\Pi(d_1, d_2, s, m_1, m_2)$ is not straightforward because the price in a partially pooled auction 1 depends not only on the valuations of its bidders, but—via shared bidders—can depend indirectly on the valuations of bidders in auction 2. Proposition 3 contains the paper's key building block, a closed-form expression for moments of the price in auction 1. First, we derive the probability distribution of the price in auction 1.

Proposition 2. *The price of auction* 1 *satisfies*

$$\begin{aligned} &P(\Pi(d_1, d_2, s, m_1, m_2) = \pi) \\ &\quad = P(X_{m_1+m_2+1:d+s} = \pi)[1 - P(X_{m_2+1:d_2+s} < X_{m_1+1:d_1}) \\ &\qquad\qquad - P(X_{m_1+1:d_1+s} < X_{m_2+1:d_2})] \\ &\qquad + P(X_{m_1+1:d_1} = \pi\,, X_{m_2+1:d_2+s} < \pi) \\ &\qquad + P(X_{m_1+1:d_1+s} = \pi, \pi < X_{m_2+1:d_2}), \end{aligned} \tag{1}$$

where $X_{b:a}$ denotes the bth largest order statistic out of draws from the valuation distribution F, with $X_{b:a} = 0$ if $a < b$.

Proof. The proof relies on a conditioning argument, which determines whether auctions 1 and 2 have the same price. The prices in the two auctions will differ when circumstances preclude the s bidders from setting the prices of both auctions simultaneously. Such circumstances arise when the valuations among one auction's dedicated bidders are sufficiently high relative to all other bidders' valuations. More specifically, if the $(m_1 + 1)$st highest valuation among the d_1 bidders is higher than the $(m_2 + 1)$st valuation among all the d_2 *and* s bidders, then the winning price in auction 1 will exceed that of auction 2—there is no rational incentive for the s bidders to bid up auction 2's price towards auction 1's higher price. Call this event U_1, for unequal prices with auction 1's price higher; i.e., $U_1 = \{X_{m_1+1:d_1} > X_{m_2+1:d_2+s}\}$. The price in event U_1 is captured by the second term on the right side of (1), while the third term captures this situation in reverse; we call this latter event $U_2 = \{X_{m_2+1:d_2} > X_{m_1+1:d_1+s}\}$.

In §1, the final example, where the d_1 bidders had valuations of \$10 and \$6, the d_2 bidders had valuations of \$3 and \$2, and the s bidders had valuations of \$5 and \$4, resulted in a selling price of \$6 in auction 1 and \$4 in auction 2. This is an illustration of the event U_1: The d_1 valuations were so high that the s bidders did not influence the price in auction 1.

Suppose instead that the d_1 bidders had valuations of \$5 and \$4, the d_2 bidders had valuations of \$3 and \$2, and the s bidders had valuations of \$10 and \$6. Then the s bidders both win an item and in so doing set the prices of both auctions to \$5. If we had replaced the \$5 d_1 bidder with a \$12 bidder, the \$6 bidder would walk away empty handed, but her final bids would set the price of both auctions to \$6. These latter two situations illustrate the final event we tackle in our proof.

The key to Equation (1)'s simplicity lies in the fact that when neither event U_1 nor U_2 occurs (call this event E), the resulting prices of the auctions are equal. Essentially, both auctions' dedicated bidders' valuations are competitive with each other, and the s compartment valuations are not so small that the higher-priced dedicated auction could overwhelm all the s and the other auction's dedicated bidders (as in U_1 and U_2). Consequently auctions 1 and 2 receive shared bids all the way through to the formation of closing prices, which are set to be equal by some shared bidder(s) winning, or price-setting dropout, bid(s). Note that the probability of event E is independent of the actual $(m_1 + m_2 + 1)$st highest value among the $d + s$ valuations (it is simply the fraction of bidder valuation orderings that yields E); hence, conditioning on E does not change the distribution of the $(m_1 + m_2 + 1)$st highest value. Since this is exactly the price π when both auctions share the same price (event E), the distribution of π under event E is consistent with $m_1 + m_2$ items being allocated to the highest of the $d + s$ bidders, which is captured by the first term on the right side of Equation (1). □

Proposition 2 elucidates the scenarios that can occur in a partially pooled system. The second term on the right side of (1) captures the event that auction 1's price is determined by the d_1 dedicated auction 1 bidders, without competitive influence from the s or d_2 bidders. In this case, auctioneer 1 does not benefit from pooling. In contrast, in the third term in (1), the price of auction 1 is determined by competitive bidding among the d_1 and s bidders, and auction 1 effectively co-opts all the s bidders. A middle ground is tread by the first term, in which both auctions experience the same price, equal to what would be realized if the auctions were formally merged.

Proposition 3. *Let* $\binom{a}{b} = 0$, $X_{b:a} = 0$ *if* $a < b$. *The moments of the price of auction* 1 *are*

$$\begin{aligned} &\mathrm{E}[\Pi^k(d_1, d_2, s, m_1, m_2)] \\ &\quad = \mathrm{E}[X^k_{m_1+1:d_1+s}] \\ &\qquad - \sum_{\substack{i=0 \\ \text{s.t. } d+s>m_1+i}}^{m_2-1} \frac{(d_1+s)\binom{d_1+s-1}{m_1}\binom{d_2}{i} - d_1\binom{d_1-1}{m_1}\binom{d_2+s}{i}}{(d+s)\binom{d+s-1}{m_1+i}} \\ &\qquad \cdot \mathrm{E}[X^k_{m_1+i+1:d+s} - X^k_{m_1+m_2+1:d+s}]. \end{aligned} \tag{2}$$

See the online appendix for a proof of this, and all remaining, propositions. The moments in Proposition 3 are expressed as deviations (because of competition from auction 2) from the idealized value $\mathrm{E}[X^k_{m_1+1:d_1+s}]$ that would be achieved if auctioneer 1 had sole access to the s shared bidders. As a check, note that when $s = 0$ (no pooling), this deviation equals zero, and $\mathrm{E}[X^k_{m_1+1:d_1+s}]$ equals the kth moment of auction 1's price in the absence of pooling.

Exact analysis of Equation (2) is difficult for several reasons: The price depends on five variables; the summed terms in (2) are fractions of binomial coefficients; and even the primitive order statistics in (2) are generally intractable except for a few specific distributions. To make further progress, we investigate various special cases in the next two sections.

4. Partial Pooling in Single-Item Auctions

In this section we assume that each auctioneer sells only one item, so price and revenue are interchangeable. To avoid trivial cases, we will assume that $d+s \geq m_1+m_2+1$. Plugging $m_1 = m_2 = 1$ into (2), applying

$$\mathrm{E}[X^k_{b:a}] = \sum_{i=a-b+1}^{a} \binom{a}{i}\binom{i-1}{a-b}(-1)^{i-(a-b)-1}\mathrm{E}[X^k_{1:i}]$$

(David 1970, p. 38),

and suppressing the arguments m_1 and m_2 from Π yield

$$\begin{aligned} &\mathrm{E}[\Pi^k(d_1, d_2, s)] \\ &\quad = \mathrm{E}[X^k_{2:d_1+s}] + s(2d_1+s-1) \\ &\qquad \cdot \left(\tfrac{1}{2}\mathrm{E}[X^k_{1:d+s}] - \mathrm{E}[X^k_{1:d+s-1}] + \tfrac{1}{2}\mathrm{E}[X^k_{1:d+s-2}]\right). \end{aligned} \tag{3}$$

4.1. The Symmetric Case

This subsection further specializes the single-item auction to the symmetric case, where $d_1 = d_2$. We express the expected price of auction 1 as $\Pi(d, s)$, where there are $d = d_1 + d_2$ dedicated bidders and s shared bidders. The total number of bidders, which we refer to as the *total market size*, is denoted by $n = d+s$. The assumption that $d+s \geq m_1+m_2+1$, or $n \geq 3$, avoids the trivial case in which both auctioneers receive zero revenue. We define the *pooling proportion* to be the fraction of all bidders who are shared; in this subsection, this quantity is given by s/n, where to ensure d is even we assume $s \in \{0, 2, \ldots, n\}$ when n is even, and otherwise $s \in \{1, 3, \ldots, n\}$. This set-up allows us to determine whether partial pooling is beneficial or detrimental to the auctioneers.

Proposition 4. *Let* $n-s$ *be even, and suppose that the valuation distribution* F *satisfies*

$$\mathrm{E}[X_{2:(n+s)/2} - X_{2:(n+s)/2-1}] \quad \textit{is nonincreasing for } 2 \leq s \leq n; \tag{4}$$

then the expected revenue in auction 1 *of the symmetric auction system satisfies*

$$\mathrm{E}[\Pi(n-s, s)] \leq \mathrm{E}[\Pi(n-s-2, s+2)] \quad \textit{for } 0 \leq s \leq n-4,$$

$$\mathrm{E}[\Pi(2, n-2)] = \mathrm{E}[\Pi(0, n)] = \mathrm{E}[X_{3:n}], \quad \textit{and} \tag{5}$$

$$\mathrm{E}[\Pi(n-s-2, s+2)] - \mathrm{E}[\Pi(n-s, s)] \quad \textit{is nonincreasing in } s \textit{ for } 0 \leq s \leq n-2. \tag{6}$$

In words, Proposition 4 states that if the expected revenue of a single-item auction has nonincreasing marginal returns in the number of bidders when the number of bidders is between $\lceil n/2 \rceil$ and n, then the expected revenue is nondecreasing in the pooling proportion (with strict equality from $(n-2)/n$ to 1) and has nonincreasing marginal returns. The main result is that when a lone auction's revenue is concave in the number of bidders, increasing the amount of pooling with a symmetric auctioneer boosts both auctioneers' expected revenues. The remainder of this subsection explores the implications of Proposition 4.

In the rest of the paper, $U[a, b]$ denotes the uniform distribution on $[a, b]$, $\exp(\lambda)$ is the exponential distribution with rate λ, and Pareto(α, k) is the Pareto distribution with cumulative distribution function $F(x) = 1-(k/x)^\alpha$, where $\alpha > 1$ (to ensure finite mean) and $x \geq k > 0$. The first two distributions are also analyzed in Ellison et al. (2004). For modeling purposes, fatter-tailed valuation distributions might be more appropriate for art or collectibles (where some bidders might have extremely high valuations for products that others value very little) rather than commodity-like products, for which thinner-tailed (e.g., the uniform) distributions might be more applicable.

4.1.1. Sensitivity of Pooling Benefits to Bidder Valuation Distribution. The complex nature of order statistics makes it difficult to rephrase condition (4) in Proposition 4 in terms of the underlying valuation distribution; however, we have Proposition 5.

Proposition 5. *Proposition 4 holds for $U[a,b]$, $\exp(\lambda)$, and Pareto(α,k) valuations.*

In addition to Proposition 5, numerical experiments suggest that Proposition 4 holds for the normal, lognormal, and chi-squared distributions, to list a few. However, the condition can be shown to fail: For example, the Weibull distribution can be shown (by computation) to fail condition (4) if the shape parameter is sufficiently small compared with n^{-1} (e.g., approximately 0.3 for $n \leq 20$). A heuristic approach to characterizing the valuation distributions that satisfy condition (4) suggests that, along the domain's tail, the log of the pdf must decrease no more than twice as fast as the hazard rate; see Appendix E (online) for details of this approximation.

4.1.2. Sensitivity of Pooling Benefits to Market Size. To assess the maximum possible benefit, we look at the relative price increase from full pooling relative to no pooling. For n even, $n/2 \geq 2$, this quantity is given by

$$\frac{\mathrm{E}[X_{3:n}-X_{2:n/2}]}{\mathrm{E}[X_{2:n/2}]} = \begin{cases} \dfrac{1}{(n+1)(a(n+2)/((b-a)(n-2))+1)} & \text{for } U[a,b], \\[2ex] \dfrac{\sum_{i=n/2+1}^{n} 1/i - 1/2}{\sum_{i=2}^{n/2} 1/i} \in \left[\dfrac{\ln(n+1)-\ln(n/2+1)-1/2}{\ln(n/2)}, \right. \\ \qquad\qquad \left. \dfrac{\ln 2 - 1/2}{\ln(n/2+1)-\ln 2}\right] & \text{for } \exp(\lambda), \\[2ex] \dfrac{2\alpha-1}{2\alpha} \displaystyle\prod_{j=n/2+1}^{n}\left(1+\frac{1}{\alpha j-1}\right)-1 & \text{for Pareto}(\alpha,k), \end{cases} \tag{7}$$

where we have used the distributions' closed-form order statistics per Appendix D (online) and for the exponential case applied the fact that, for $u \geq l$,

$$\ln\left(\frac{u+1}{l}\right) = \int_l^{u+1} \frac{dx}{x} \leq \sum_{j=l}^{u} \frac{1}{j} \leq \int_{l-1}^{u} \frac{dx}{x} = \ln\left(\frac{u}{l-1}\right). \tag{8}$$

Hence, the relative price increase from full pooling decreases as $O(1/n)$ in the $U[a,b]$ case and decreases no more slowly than $O(1/\ln(n/2))$ in the $\exp(\lambda)$ case. In contrast, returns from full pooling for the Pareto case increase with the number of bidders, converging to $(2\alpha - 1/2\alpha)2^{1/\alpha} - 1$ as $n \to \infty$ (see Appendix F). This limit is concave in α and maximized at 6.14% at $\alpha \approx 1.8$. For thinner-tailed valuation distributions, when the market is large, the highest valuations begin to converge near an upper limit (e.g., near b for the $U[a,b]$ case), making the pooled and unpooled auction prices similar. Intuitively, such convergence is less a concern with fatter-tailed distributions, where ever-larger, but rare, valuations are possible as the market size grows. For example, depending on the total market size being small, medium, or large, the best performer among the three valuation distributions (measured by relative price increase) is (respectively) the light-tailed $U[a=0,b]$, the exponential-tailed $\exp(\lambda)$, or the heavy-tailed Pareto; see Figure 3, in which we fix α at 2.2, yielding a "medium" range of $8 \leq n \leq 78$. The plots in Figure 3 hold for all positive values of λ, b, and k; see Equation (7). Thus Figure 3 is a fair comparison of pooling effects across the distributions; for example, it does not rely on any assumptions about the relative sizes of the distributions' mean values.

Figure 4(a) illustrates this subsection's main points: The percentage price increase is increasing and concave in the pooling proportion, and the percentage increase shrinks in the market size n for the thin-tailed uniform and exponential distributions and is larger for exponential than for uniform or Pareto at the given (medium) market sizes. Figure 4(b) computes the coefficient of variation (standard deviation divided by the mean) of the revenue and shows that this quantity is decreasing in the pooling proportion and decreasing in the market size for uniform and exponential distributions. Hence, there is no risk-reward trade-off with pooling: It simultaneously

Figure 3 Sensitivity to Market Size of Price Increase Under Full Pooling Relative to No Pooling for $U[0, b]$, $\exp(\lambda)$, and Pareto$(2.2, k)$ Valuations

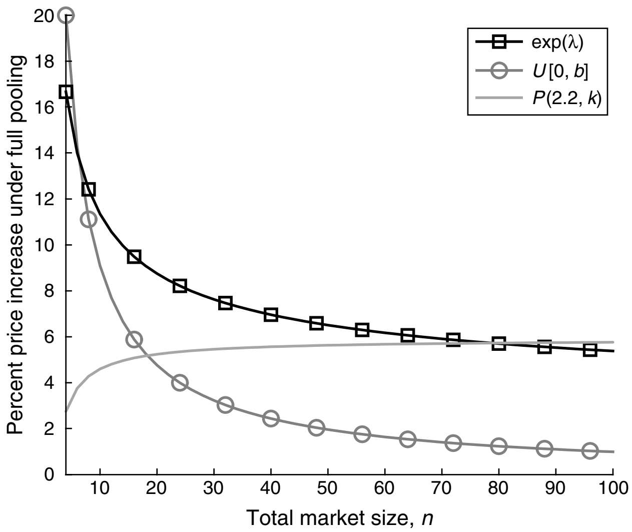

Figure 4 (a) The Percentage Price Increase Relative to No Pooling and (b) the Coefficient of Variation of Price, as a Function of the Pooling Proportion p, for $n = 24$ and 40 Total Bidders, and for $U[0, b]$, $\exp(\lambda)$, and Pareto$(2.2, k)$ Valuations

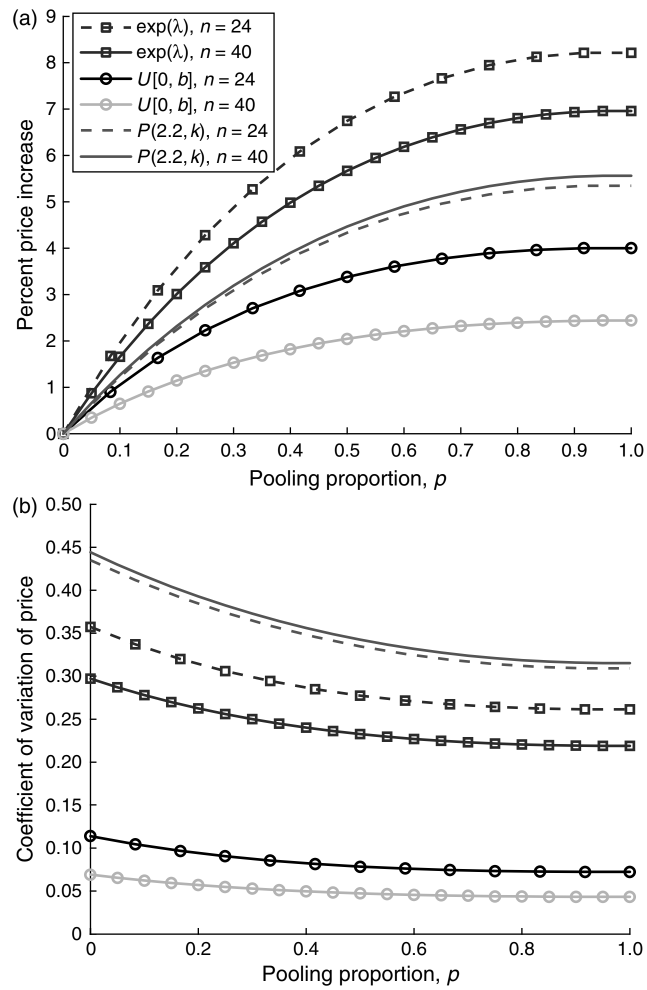

increases the mean and decreases the coefficient of variation of the revenue. The plots in Figure 4 hold for all positive λ, b, and k (as in Figure 3, the plots are of ratios for which the effects of these parameters can be shown to cancel out).

4.2. The Asymmetric Case

In §4.1 we found that partial pooling benefits both auctioneers when each has the same number of dedicated bidders. This subsection investigates whether partial pooling is mutually beneficial in the asymmetric case. More specifically, we consider partial pooling relative to the base case in which there is no pooling, which could be achieved (ignoring inter-auction temporal dynamics) by scheduling the two auctions at different times. In the absence of pooling, we assume that auction i has $n_i \geq 2$ dedicated bidders, so auctioneer i's revenue is $\mathrm{E}[X_{2:n_i}]$. In the partially pooled scenario, where simultaneous single-item auctions are held, we assume that a fixed proportion p of all $n = n_1 + n_2$ bidders participates in both auctions. More specifically, to account for integrality, we assume in the partially pooled scenario that $\lfloor pn_i \rfloor$ bidders from auction i participate in both auctions simultaneously and $n_i - \lfloor pn_i \rfloor$ are dedicated. Case $p = 1$ refers to full pooling and $0 < p < 1$ to partial pooling.

To assess the desirability of partial pooling for a given pooling proportion p, we define $n_2^*(p, n_1)$ as the minimum number of bidders auctioneer 2 must have for auctioneer 1 to find partial pooling with auctioneer 2 profitable (relative to the unpooled scenario). Given the pooling proportion p, $\lfloor pn_2^*(p, n_1) \rfloor$ is the minimum number of bidders that auctioneer 2 must share to make auctioneer 1 favor partial pooling with him. By symmetry, we can apply the same idea from the perspective of auctioneer 2. Taking $n_1^*(p, n_2)$ to be the analogous operator on a fraction and a positive integer, we call the set of integers (n_1, n_2) where $n_2 > n_2^*(p, n_1)$ and $n_1 > n_1^*(p, n_2)$ the *mutually feasible pooling region*. The following two propositions seek to characterize the mutually feasible pooling region for the $U[a, b]$ and $\exp(\lambda)$ distributions, respectively.

Proposition 6. *For* $U[a, b]$ *valuations, if* $n_1 p \in \{1, 2, \ldots, n_1\}$, *then*

$$\left\lfloor p\left(\frac{n_1}{2.3}+\frac{1}{2}\right)\right\rfloor \le \lfloor p n_2^*(p, n_1)\rfloor \le \left\lceil p\left(\frac{n_1}{2}+\frac{1}{2}\right)\right\rceil. \quad (9)$$

Proposition 7. *For* $\exp(\lambda)$ *valuations, if* $n_1 p \in \{1, 2, \ldots, n_1\}$, *then*

$$\lfloor n_1\theta_l(p)\rfloor \le \lfloor p n_2^*(p, n_1)\rfloor \le \lceil (n_1+1)\theta_u(p, n_1)\rceil, \quad \textit{where}$$
$$\theta_l(p)=e^{p/2}-1, \quad \textit{and}$$
$$\theta_u(p, n_1)=\exp\left(\frac{p(4n_1-n_1p-2)}{6n_1-4}\right)-1.$$

4.2.1. Minimum Market Share. Figure 5 graphs the exact mutually feasible full pooling (i.e., $p=1$) regions for uniform and exponential valuations. Note that the uniform region contains the exponential region; i.e., the requirements for full pooling in the exponentially distributed valuations case are more stringent than in the uniformly distributed valuations case. To explore this further, it is more intuitive to discuss these results in terms of the minimum market share, $n_2^*(p, n_1)/(n_1+n_2^*(p, n_1))$, that the smaller auctioneer (auctioneer 2) must possess for the larger auctioneer to want to pool with him (i.e., the slope of the bottom two lines in Figure 5). The minimum required market share is approximately 1/3 for the $U[a, b]$ case in Proposition 6. To assess the market share bounds

Figure 5 **Mutually Feasible Full Pooling Regions, for Uniform** (—) **and Exponential** (- -) **Valuations**

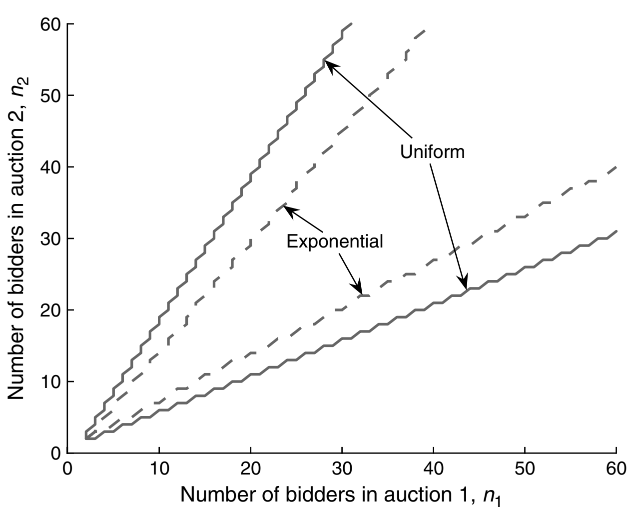

arising from Proposition 7, we look at $\theta_l(p)/(p+\theta_l(p))$ and $\theta_u(p, n_1)/(p+\theta_u(p, n_1))$

$$\frac{\theta_l(p)}{p+\theta_l(p)} \approx \frac{p/2+p^2/8+p^3/48}{3p/2+p^2/8+p^3/48},$$
$$\lim_{p\to 0}\frac{\theta_l(p)}{p+\theta_l(p)}=\frac{1}{3},$$
$$\lim_{p\to 1}\frac{\theta_l(p)}{p+\theta_l(p)}=1-e^{-1/2}\approx 0.393,$$
$$\lim_{p\to 0}\lim_{n_1\to\infty}\frac{\theta_u(p, n_1)}{p+\theta_u(p, n_1)}=\lim_{n_1\to\infty}\lim_{p\to 0}\frac{\theta_u(p, n_1)}{p+\theta_u(p, n_1)}=0.4,$$
$$\lim_{p\to 1}\frac{\theta_u(p, n_1)}{p+\theta_u(p, n_1)}=1-e^{-1/2}\approx 0.393,$$
$$\arg\max_{p\in[0,1]}\lim_{n_1\to\infty}\frac{\theta_u(p, n_1)}{p+\theta_u(p, n_1)}$$
$$=\lim_{n_1\to\infty}\arg\max_{p\in[0,1]}\frac{\theta_u(p, n_1)}{p+\theta_u(p, n_1)}=0.404.$$

Hence, the bounds on the market share in the $\exp(\lambda)$ case range from 1/3 to 0.404 when n_1 is large and the lower and upper bounds equal 0.393 under full pooling. That is, the hurdle for pooling appears to be somewhat higher in the $\exp(\lambda)$ case than in the $U[a, b]$ case. We turn now to the Pareto distribution, for which the following proposition shows that the pooling hurdle grows to the highest possible value—a 50–50 market split—as the tail fattens but shrinks to the hurdle for the exponential case as the tail becomes thinner.

Proposition 8. *For the Pareto*(α, k) *distribution*,

$$\lim_{\alpha\to 1^+} n^*_{2,\,Pareto(\alpha,k)}(p, n_1)=n_1, \quad \textit{and}$$
$$\lim_{\alpha\to\infty} n^*_{2,\,Pareto(\alpha,k)}(p, n_1)=n^*_{2,\,\exp(\lambda)}(p, n_1).$$

The first limit in Proposition 8 follows from algebra (see Appendix I), but the second limit follows from the fact that linear transformations of the valuations do not alter the $n_2^*(p, n_1)$ analysis, and, for $X\sim$ Pareto(α, k), $\alpha X-\alpha k$ converges in distribution to an exponential random variable with mean k as α becomes large.

4.2.2. Sensitivity of Minimum Market Share to Valuation Distribution. Collectively, Propositions 6–8 suggest the hurdle to pool is, relatively, lowest for uniform, highest for Pareto, and in the middle for exponential distributions. This ordering is essentially the flip side of the returns to symmetric pooling behavior illustrated in Figure 3. As the market size increases, returns to symmetric pooling are better sustained by fatter-tailed distributions, because bids do not pile up near an upper limit. As such, the expected price growth is more sustainable because of the possibility of ever larger, but rare, bidder valuations. But, because expected prices with fatter-tailed distributions are more driven by large, rare valuations, pooling with an auction that has a significantly smaller market share is less desirable because of the low odds that the additional bidders will have valuations large enough to improve on the expected price of the larger auction without pooling. In other words, the more prices are driven by rare events, the more equity in the number of bidders is needed to ensure that benefits of pooling are mutual. Equation (17) addresses this mathematically, showing that losing a very rare, high valuation (adding an item) is for fatter-tailed distributions less outweighed by the benefit of adding a few more bidders.

4.2.3. Sensitivity of Minimum Market Share to Pooling Proportion. Figure 6 plots the exact values of the minimum required market share, $n_2^*(p, n_1)/(n_1 + n_2^*(p, n_1))$, versus n_1 for $U[a, b]$ and $\exp(\lambda)$ valuations, for both $p = 0.2$ and $p = 1$ ($p = 0.2$ is an empirical estimate of eBay's proportion of pooling bidders; Anwar et al. 2006). These curves indicate that $n_2^*(0.2, n_1)$ is approximately equal to $n_2^*(1, n_1)$, suggesting that the mutually feasible pooling region is well described by analyzing the full-pooling case. That is, the plots in Figures 6(a) and 6(b) are roughly concentrated near values 1/3 and 0.393, respectively.

The main result from §4.2 is that, regardless of the pooling proportion, for the practical (nonheavy-tailed) valuation distributions studied here, partial pooling appears mutually beneficial as long as neither auctioneer dominates the market (has more than 60%–65% market share). This suggests that market share is the most important factor in the auctioneer's decision on whether to encourage or discourage pooling; that is good news, because the simplicity of counting bidders shown in bid histories (available on eBay and Yahoo! auctions) may make market share data easier to monitor and estimate than pooling proportions and bidder valuation distributions.

Figure 6 **The Minimum Market Share, $n_2^*(p, n_1)/(n_1 + n_2^*(p, n_1))$, of Auction 2 such that Auctioneer 1 will Want to Pool, as a Function of the Number of Bidders in Auction 1, for Pooling Proportions $p = 0.2$ and $p = 1$**

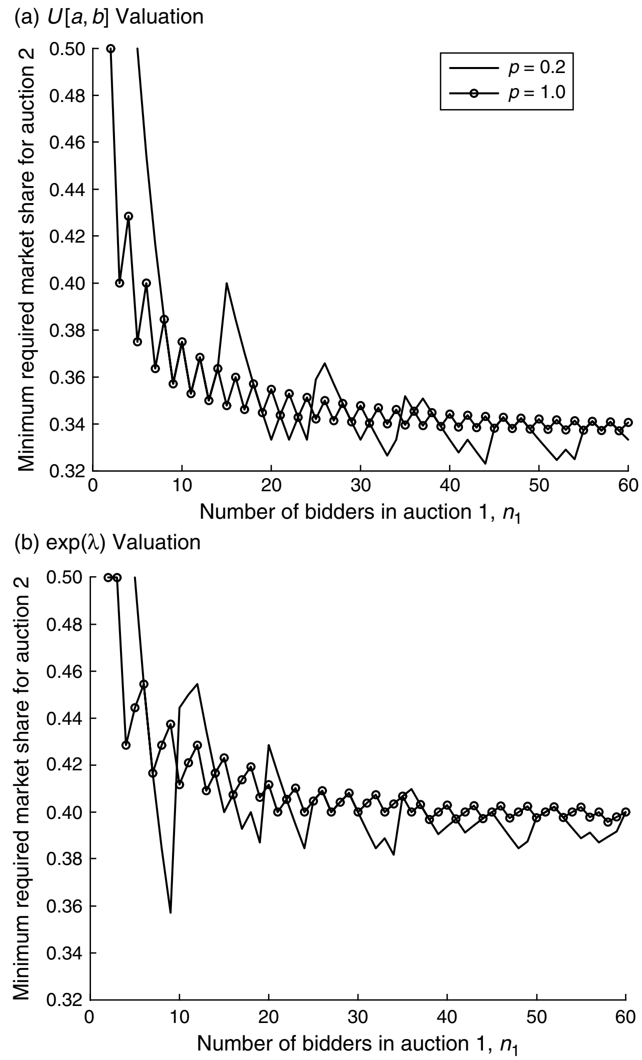

Note. Peaks and valleys in the graphs are caused by the integrality of n_1 and n_2.

5. Pooling in Multi-Item Auctions

Although the complex expression in (2) for the expected price in multi-item pooled auctions precludes direct extension of the single-item results of §4, in this section we compute the expected revenue increase of full symmetric pooling relative to no pooling and partially characterize the mutually feasible pooling region for multi-item auctions. Full pooling is the focus of §§5.1–5.2, and partial pooling is considered in §5.3.

5.1. Full-Pooling Analysis: Symmetric Case

To help assess the benefit of pooling multi-item auctions, we compute an auctioneer's relative revenue increase from symmetric full pooling relative to no pooling, the analogue of Equation (7) for single-item auctions. The expressions involve m, the number of items sold in each auction (n is again the total market size and is assumed even with $n/2 \geq m+1$):

$$\frac{\mathrm{E}[X_{2m+1:n} - X_{m+1:n/2}]}{\mathrm{E}[X_{m+1:n/2}]} = \begin{cases} \dfrac{1}{(n+1)(a(n+2)/((b-a)(n-2m))+1)} & \text{for } U[a,b], \\[2ex] \dfrac{\sum_{i=2m+1}^{n} 1/i - \sum_{i=m+1}^{n/2} 1/i}{\sum_{i=m+1}^{n/2} 1/i} & \text{for } \exp(\lambda), \\[2ex] \displaystyle\prod_{j=m+1}^{2m}\left(1-\frac{1}{\alpha j}\right)\prod_{j=n/2+1}^{n}\left(1+\frac{1}{\alpha j - 1}\right) - 1 & \text{for Pareto}(\alpha,k). \end{cases} \tag{10}$$

5.1.1. Sensitivity of Pooling Benefits to Market Size and Competition Ratio. Echoing the insights for single-item auctions gleaned from (7), we find once again that fatter tails accommodate more sustained improvement when pooling large markets: For fixed m, the relative revenue returns are decreasing in total market size when valuations are uniform or exponential, but relative revenue returns actually increase with total market size when valuations are Pareto distributed. (We omit the details that prove this, but for the exponential case one can show the result is true for $n = 2m+2$ and then use induction to establish the same for $n > 2m+2$.) However, for fixed market size, a similar analysis shows that relative returns to pooling for uniform, exponential, and Pareto distributions all weakly increase with the bidder-to-item ratio, a quantity that we refer to as the *competition ratio* (see Figure 7, which illustrates these points for two- and three-item auctions). As in the single item case (Figure 3), we see that expected price growth is more sustainable under larger market sizes with fatter-tailed valuations because of the possibility of ever larger—but rare—bidder valuations. Yet fatter tails imply more sensitivity to additional items, as more items can allow bidders with large, rare valuations to win an item without bidding the price up to their valuation. In contrast, returns to pooling are less sensitive for the uniform distribution for which large, rare valuations are less an issue (with $a = 0$ the returns are actually insensitive to the number of items; see Equation (10)).

Figure 7 **Sensitivity to Market Size of Price Increase Under Full Pooling Relative to No Pooling for** $U[0,b]$, $\exp(\lambda)$, **and** Pareto$(2.2,k)$ **Valuations, with** $m=2$ **and** $m=3$ **Items**

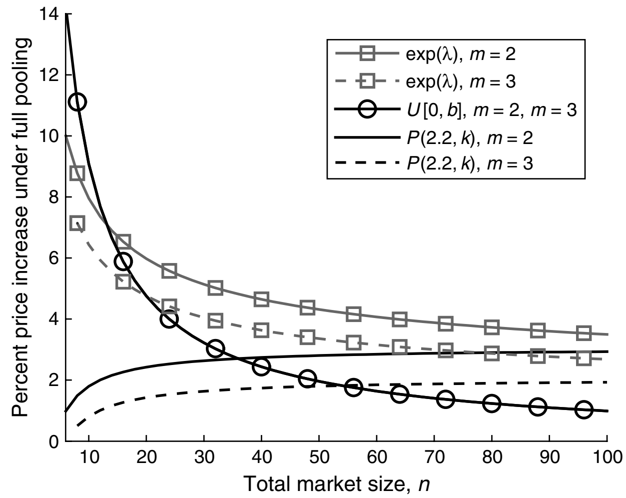

5.2. Full-Pooling Analysis: Asymmetric Case

Let us define $n_2^*(p, n_1, m_1, m_2)$ as the natural generalization of the quantity $n_2^*(p, n_1)$ introduced in §4.2: It is the minimum number of bidders auction 2 must possess for auctioneer 1 to favor pooling (with proportion p) over no pooling. In this subsection we develop insight into the behavior of n_2^* for general distributions, leading off with the uniform, exponential, and Pareto cases. We assume $n_1 \geq m_1 + 1$. For the $U[0,1]$ case, auction 1's price under full pooling ($p=1$) is $\mathrm{E}[X_{m_1+m_2+1:n_1+n_2}] = (n_1+n_2-m_1-m_2)/(n_1+n_2+1)$. Because this equals $\mathrm{E}[X_{m_1+1:n_1}] = (n_1 - m_1)/(n_1+1)$ if $n_2/m_2 = (n_1+1)/(m_1+1)$, we conclude that

$$n_2^*(1, n_1, m_1, m_2) = \left\lceil \frac{m_2}{m_1+1}(n_1+1) \right\rceil \quad \text{for } U[a,b], \tag{11}$$

where we have generalized to $U[a,b]$ since n_2^* is insensitive to linear transformations of the underling valuations. An analogous approach for $\exp(\lambda)$ valuations, using (8) to bound summations in the difference $E[X_{m_1+m_2+1:n_1+n_2}] - E[X_{m_1+1:n_1}]$, yields

$$\left\lfloor \frac{m_2}{m_1+1} n_1 \right\rfloor \le n_2^*(1, n_1, m_1, m_2) \le \left\lceil \frac{m_2}{m_1}(n_1+1) \right\rceil \quad \text{for } \exp(\lambda). \quad (12)$$

Similarly,

$$\lim_{\alpha \to 1} n^*_{2,\,\text{Pareto}(\alpha,k)}(1, n_1, m_1, m_2) = n_1 \frac{m_2}{m_1}, \quad \text{and} \quad (13)$$

$$\lim_{\alpha \to \infty} n^*_{2,\,\text{Pareto}(\alpha,k)}(1, n_1, m_1, m_2) = n^*_{2,\,\exp(\lambda)}(1, n_1, m_1, m_2); \quad (14)$$

see Section J for a derivation of Equation (13).

5.2.1. Competition Ratio Parity and Mutually Feasible Full Pooling. For m_i, n_i large, we can use the asymptotic approximation

$$E[X_{i:j}] \approx F^{-1}\left(\frac{j-i+1}{j}\right) \quad \text{for } j-i, j \text{ large, } 0<i<j, \quad (15)$$

where F^{-1} denotes the inverse cdf. (Although we will not need it here, the full result, including a characterization of the asymptotically normal distribution of $X_{i:j}$ as $j-i$, $j \to \infty$, is proved in Theorem 5.8 of Balkema and de Haan 1978.) Equation (15) allows us to conclude that when $n_2 \approx (m_2/m_1)n_1$, $E[X_{m_1+m_2+1:n_1+n_2}] \approx E[X_{m_1+1:n_1}]$; i.e., full pooling generates approximately the same expected price as no pooling when m_i and n_i are large and the auctions' competition ratios (number of bidders divided by number of items) are nearly equal. This asymptotic result, along with Equations (11)–(14), suggests that the approximation $n_2^*(1, n_1, m_1, m_2) \approx (m_2/m_1)n_1$ is in general valid if m_i and n_i are at least moderate in size. Moreover, the monotonicity of auction i's revenue in the pooling proportion p in Proposition 4 suggests that if $p = 1$ generates revenue no higher than $p = 0$, then neither will the cases $0 < p < 1$. If this is true, then (15) implies for large m_i and n_i, $n_2^*(p, n_1, m_1, m_2) \approx (m_2/m_1)n_1$. Furthermore, a similar analysis from auctioneer 2's viewpoint suggests that $n_1^*(p, n_2, m_1, m_2) \approx (m_1/m_2)n_2$; consequently, we predict that for large m_i and n_i, the mutually feasible pooling region consists of only a very small area in (n_1, n_2) space surrounding the line $n_2 = (m_2/m_1)n_1$. That is, both auctions need almost the same competition ratio for pooling to be mutually beneficial. Numerical experiments depicted in Figure 8 corroborate these claims, even for moderate m_i, n_i.

Figure 8 **The Mutually Feasible Pooling Region for Pooling Proportions** $p = 0.5$ (□) **and** $p = 0.75$ (·)

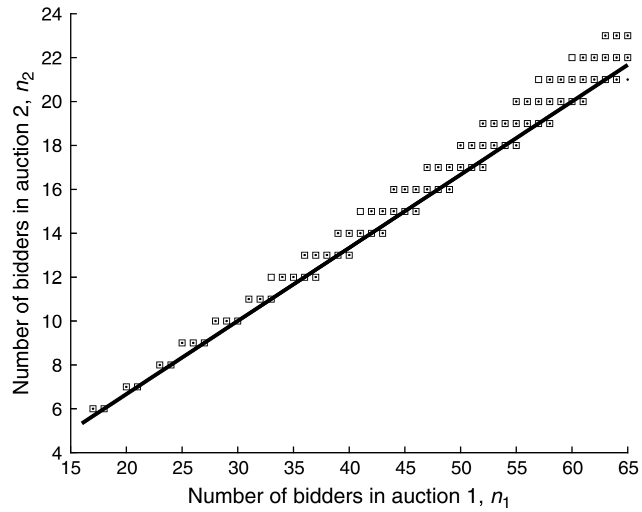

Notes. Also displayed is the line $n_2 = (m_2/m_1)n_1$. The valuations are lognormal with $\mu = 5$, $\sigma = 0.2$, and there are $m_1 = 15$ and $m_2 = 5$ items for sale.

5.3. Competition Ratio vs. Market Size

To further investigate the multi-item mutually feasible pooling region, this subsection studies the effects on partial pooling of each auction's competition ratio (i.e., bidder-to-item ratio) and market size (i.e., number of bidders). In this terminology, an auction with six bidders and two items has the same competition ratio but twice the market size as an auction with three bidders and one item. For $U[0, 1]$ valuations, the expression $E[X_{m+1:n}] = (n-m)/(n+1)$, $n \ge m+1$ reveals that, for an isolated auction, the expected price is increasing in both the competition ratio and market size.

Under partial pooling, however, the impacts of competition ratio and market size are more complex, as described in Proposition 9 (for $U[a, b]$ and $\exp(\lambda)$ valuations). Because Equation (2) becomes increasingly unwieldy as m_2 grows above 1, in Proposition 9 we create disparate competition ratios and market sizes by taking auction 1 to be a first-price auction (with slight abuse of notation, $m_1 = 0$) and auction 2 to have either one or two items. While auction 1 in this case is not interpretable as English, the qualitative nature of Proposition 9 is based on the mathematics of Equation (2) (an expression of order statistics) and is indicative of the competition ratio versus market size relationship when m_1, $m_2 \geq 1$, as discussed at the end of this subsection.

PROPOSITION 9. *Let*

$$\beta_j \triangleq \mathrm{E}[\Pi(n-j, 2(n-j), 3j, 0, 2)] - \mathrm{E}[\Pi(n-j, n-j, 2j, 0, 1)] \tag{16}$$

and suppose $n \geq 4$. *For both* $U[a,b]$ *and* $\exp(\lambda)$ *valuations there exists* $\bar{j} \in [1, \ldots, n]$ *such that* $\beta_j > 0$ *for* $j \in [1, \ldots, \bar{j}]$, $\beta_j \leq 0$ *for* $j \in [\bar{j}+1, \ldots, n]$. *That is, when* $n_1 = n_2 = n$, $m_1 = 0$, $m_2 = 1$, *doubling the market size of auction* 2 (*while maintaining auction* 2*'s competition ratio at* n, *which is less than auction* 1*'s competition ratio*) *increases the pooling price of auction* 1 *only if the pooling proportion* j/n *is below* $\bar{j}/n$. *More precisely,*

$$\text{for } U[a,b], \begin{cases} \dfrac{3}{5} \leq \dfrac{\bar{j}}{n} & \text{if } n \geq 4; \\[2ex] \dfrac{4}{5} \leq \dfrac{\bar{j}}{n} & \text{if } n \geq 9, \end{cases} \qquad \text{and}$$

$$\text{for } \exp(\lambda), \quad \frac{\bar{j}}{n} \leq \frac{1}{2} \quad \text{if } n \geq 4.$$

Proposition 9 shows that when auction 1 has a higher competition ratio than auction 2, auction 1's price is positively or negatively impacted by increased market size in auction 2, depending respectively on whether the pooling proportion is small or large.

To elaborate on the intuition behind Proposition 9, when the pooling proportion is large ($p \approx 1$), the event E (the prices of the two auctions are equal) in the proof of Proposition 2 occurs with high probability; i.e., the auctions behave as if they were fully pooled. In this case, the auction price is driven by the pooled competition ratio, $(n_1+n_2)/(m_1+m_2)$, which is the weighted average (weighted by items for sale, or equivalently, by market size if the competition ratios are held fixed) of the original auctions' competition ratios, $m_1/(m_1+m_2)(n_1/m_1) + m_2/(m_1+m_2)(n_2/m_2)$. Hence, the auctioneer with the higher competition ratio (auctioneer 1) will have his price reduced by partial pooling, and this effect is exacerbated if the auction with the smaller competition ratio (auction 2) has the higher market share. This explains the observation in §5.2 that, for $p \approx 1$, the mutually feasible pooling region is closely approximated by the line $n_2 = (m_2/m_1)n_1$. In contrast, when the pooling proportion is small ($p \ll 1$), the event U (unequal prices in the two auctions) in the proof of Proposition 2 is highly likely to prevail. In particular, shared bidders are apt to be co-opted. In this case, increasing the market size of the auction 2 increases the number of bidders that auctioneer 1 can co-opt and thus makes partial pooling more attractive to auctioneer 1.

Proposition 9 also suggests that, relative to the $\exp(\lambda)$ case, the impact of auction 2's market size under $U[a,b]$ valuations is positive in a larger interval of pooling proportions. While Proposition 9 does not address Pareto(α, k) valuations directly, numerical experiments indicate that, among the uniform, exponential, and Pareto cases, the interval $[0, \bar{j}/n]$ is shortest for the Pareto case, with $\bar{j}/n$ approaching zero as α approaches 1 (as with Proposition 8, the case $\alpha \to \infty$ is subsumed by $\exp(\lambda)$ analysis). This nesting is reminiscent of the mutually feasibly set nesting seen in §4.2; indeed, the §4.2 analysis rests explicitly on the trade-off between adding bidders and adding items, and the same trade-off essentially governs the sign of β_j: see Equation (A54) in the proof of Proposition 9, Appendix K (online). In particular, the impact of an additional bidder relative to an additional item is given by

$$\begin{aligned} &\frac{E_{U[a,b]}[X_{m+1:n+1} - X_{m+1:n}]}{E_{U[a,b]}[X_{m:n} - X_{m+1:n}]} \\ &\quad = \frac{m+1}{n+2} > \frac{m}{n+1} = \frac{E_{\exp(\lambda)}[X_{m+1:n+1} - X_{m+1:n}]}{E_{\exp(\lambda)}[X_{m:n} - X_{m+1:n}]} \\ &\quad > \frac{\alpha m - 1}{\alpha(n+1) - 1} = \frac{E_{\mathrm{Pareto}(\alpha,k)}[X_{m+1:n+1} - X_{m+1:n}]}{E_{\mathrm{Pareto}(\alpha,k)}[X_{m:n} - X_{m+1:n}]}, \end{aligned} \tag{17}$$

suggesting that, in the multi-item case, the number of bidders required to offset additional items is smaller for auctions having $U[a,b]$ valuations, larger for Pareto(α, k) valuations, and in the middle for $\exp(\lambda)$ valuations. These differences "wash out" as n and m become large.

To illustrate these ideas, we provide some numerical plots of mutually feasible pooling regions (Figure 9). These plots fix the market size of auction 1 at $n_1 = 20$, vary n_2 (either $n_2 = 20$ or 60), and express

Figure 9 **Mutually Feasible Pooling Region Under $n_1 = 20$, $n_2 = 20$ (□), and $n_2 = 60$ (·)**

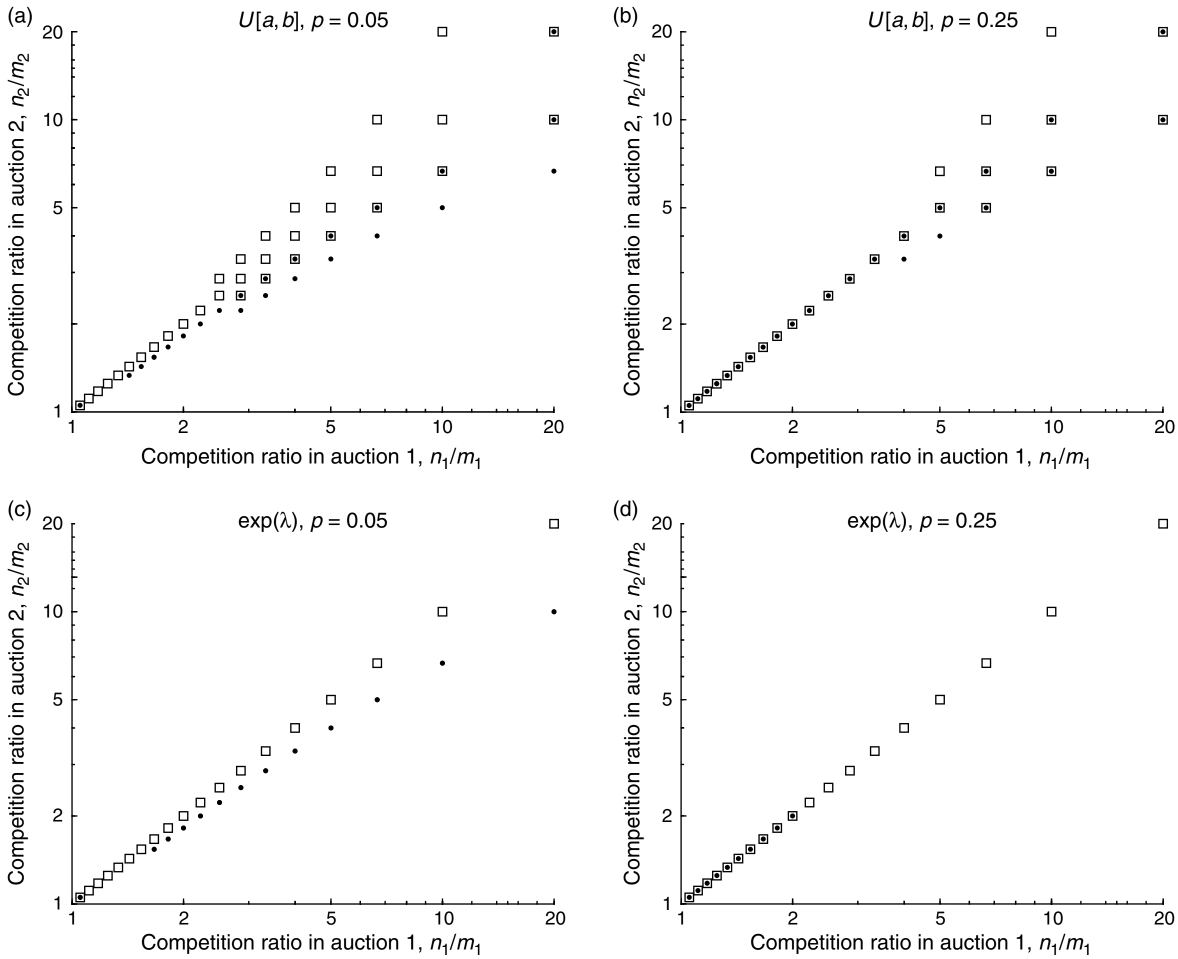

Notes. The four figures consider two pooling proportions ($p = 0.05$ and $p = 0.25$) and two valuation distributions ($U[a, b]$ and $\exp(\lambda)$). For readability, the figures are plotted on a log-log scale.

the mutually feasible pooling region in the two-dimensional competition ratio space $(n_1/m_1, n_2/m_2)$ on a log-log scale. Comparing the $n_2 = 20$ and $n_2 = 60$ regions in Figure 9(a), we see that increasing the market size of auction 2 (from $n_2 = 20$ to $n_2 = 60$) causes the mutually feasible pooling region to move slightly clockwise; i.e., auction 2 has smaller competition ratios. However, a comparison of Figures 9(a) and 9(b) reveals that this effect of increasing auction 2's market size is mitigated when the pooling proportion is increased from $p = 0.05$ to $p = 0.25$. Analogous observations can be made in the $\exp(\lambda)$ plots in Figures 9(c) and 9(d). Comparing Figures 9(a) and 9(c) (or Figures 9(b) and 9(d)) shows that the impact of increasing auction 2's market size (i.e., increasing n_2) is greater for $U[a, b]$ valuations than for $\exp(\lambda)$ valuations. In summary, Figure 9 corroborates the insights of Proposition 9: Market size can compensate for competition ratio under small pooling proportions, and, relative to $\exp(\lambda)$ valuations, impacts of market size are greater for $U[a, b]$. Finally, as in Figure 8, the mutually feasible pooling regions in Figure 9 are once again narrow areas near the equal competition ratio line $n_1/m_1 = n_2/m_2$.

6. Conclusions

The main goal of this analysis is to understand the impact of partial pooling—i.e., a subset of bidders participating simultaneously in more than one Internet auction—on the auctioneers' expected revenues.

We use a stylized model of two simultaneous English auctions with others dedicated to each auction and others participating simultaneously in both auctions, accompanied by bidder behavior shown to induce a Bayesian Nash equilibrium. We show that, for single-item auctions, auctioneers having the same number of dedicated bidders benefit (through both a higher expected revenue and a lower coefficient of variation of revenue) as the proportion of pooled bidders—i.e., those participating in both auctions—in the overall population grows, provided that the revenue from a second-price, single-item auction would be concave in the number of bidders. This latter condition relies only on the underlying distribution of bidder valuations. In the case where a single-item auction has more bidders than the other, the bigger auctioneer only finds partial pooling desirable if the market share of the smaller auction exceeds some threshold. This threshold is largely independent of the amount of pooling (i.e., the proportion of bidders participating simultaneously in both auctions), but depends on the valuation distribution: The threshold is approximately $1/3$ for $U[a,b]$ valuations and $1-e^{-1/2}\approx 0.393$ for $\exp(\lambda)$ valuations and is larger still for Pareto(α,k) distributions, nearing $1/2$ as the shape parameter $\alpha\to 1$. Although the threshold appears to become more stringent with the tail size of the valuation distribution, for the practical (nonheavy-tailed) valuation distributions studied, partial pooling appears mutually beneficial as long as neither auctioneer dominates the market (has more than 60%–65% market share), regardless of the pooling proportion.

For auctioneers selling many items to many bidders, we find that the key characteristic of each auction is its competition ratio, which we define as the number of bidders divided by the number of items. Relative to the case of single-item auctions, the mutually beneficial pooling region in multi-item auctions is much smaller and is restricted to cases where both auctions have nearly identical competition ratios (see Figures 8 and 9). From the viewpoint of the auction with the larger competition ratio, pooling with the auction with the smaller competition ratio is desirable only if this latter auctioneer has a large market share and the pooling proportion (i.e., fraction of bidders participating in both auctions) is small. Our analysis also shows that, for both single-item and multi-item auctions, the valuation distribution governs the degree to which larger markets mute or enhance the effects of pooling. For example, for symmetric single-item auctions with total market size n, the benefit of full pooling over no pooling is $O(1/n)$ for $U[a,b]$ valuations, approximately $O(1/\ln(n/2))$ for $\exp(\lambda)$ valuations, and increasing in n for Pareto(α,k) distributions. A practical result suggested by these findings is that pooling effects diminish with market size when we ignore the possibility of heavy-tailed valuations.

Because a transition to full pooling can be achieved by merging the two auctions in our model, our work is tangentially related to market-merging research in the finance literature, e.g., Pagano (1989), and more recently the economics literature, e.g., Ellison and Fudenberg (2003). These papers characterize equilibria in which two markets can coexist, where pressures to merge arise from liquidity (or market size) (Pagano 1989, Ellison and Fudenberg 2003), and pressures not to merge arise from externalities (buyers prefer markets with fewer buyers) (Ellison and Fudenberg 2003). Ellison et al. (2004) perform research along these lines in an auction context, in which buyers and sellers both make an ex ante decision to dedicate themselves to one of two auction houses. Roughly speaking, Ellison et al. (2004) find that two auction houses can coexist when their forces of liquidity and externalities are in balance. In our study, we find that a similar qualitative statement holds when pooling is mutually beneficial for auctioneers, albeit from a different modeling set-up (we allow buyers to participate across auctions, and there are only two sellers, each selling multiple items). Thus, well-balanced markets can assure coexistence if relying on the decisions of many sellers and buyers (e.g., competing double-auction environment) or can, as we find in this paper, lead to pooling if they rely on the mutual decisions of two sellers who dominate the sell side of the market (competing English auctions).

Our model is limited in scope: In addition to assuming English auctions with no reservation price, we ignore the fact that auction closings and openings are continuous-time events. Incorporating asynchronous auctions into our partial pooling model would be an interesting but daunting task. Nonetheless, taken together, our results suggest that competing single-item auctioneers can benefit from partial pooling as long as their market shares are not too lopsided, and consequently partial pooling should be mutually encouraged by the auctioneers in this case. For multi-item auctions, the prime beneficiaries of pooling are the auctioneers with lower competition ratios. However, for auctioneers with high competition ratios that operate large auctions, the negative impact of pooling is likely to be too small to engage in efforts to reduce pooling, such as product differentiation through policies, service or advertising which might be costly.

Electronic Companion

An electronic companion to this paper is available on the *Manufacturing & Service Operations Management* website (http://msom.pubs.informs.org/ecompanion.html).

Acknowledgments

The authors thank Jérémie Gallien for helpful comments on an earlier draft of this paper, as well as the referees and an associate editor whose valuable suggestions improved this paper considerably. The first author was partially supported by a graduate fellowship from the Massachusetts Institute of Technology Sloan School of Management.

References

Anwar, S., R. McMillan, M. L. Zheng. 2006. Bidding behavior in competing auctions: Evidence from eBay. *Eur. Econom. Rev.* **50** 307–322.

Avery, C. 1998. Strategic jump bidding in English auctions. *Rev. Econom. Stud.* **65** 185–210.

Bakos, J. Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Sci.* **43** 1676–1692.

Balakrishnan, N., A. C. Cohen. 1991. *Order Statistics and Inference: Estimation Methods*. Academic Press, San Diego, CA.

Balkema, A. A., L. de Haan. 1978. Limit distributions for order statistics II. *Theory Probab. Appl.* **23** 341–358.

Beam, C., A. Segev. 1998. Auctions on the Internet: A field study. Working Paper 1032, Haas School of Business, University of California, Berkeley.

Bulow, J., P. Klemperer. 1996. Auctions versus negotiations. *Amer. Econom. Rev.* **86** 180–194.

Daniel, K. D., D. Hirshleifer. 1998. A theory of costly sequential bidding. Working Paper 98028, University of Michigan Business School, Ann Arbor.

David, H. A. 1970. *Order Statistics*. John Wiley, New York.

Demange, G., D. Gale, M. Sotomayor. 1986. Multi-item auctions. *J. Political Econom.* **94** 863–872.

Ellison, G., D. Fudenberg. 2003. Knife-edge or plateau: When do market models tip? *Quart. J. Econom.* **118** 1249–1278.

Ellison, G., D. Fudenberg, M. Möbius. 2004. Competing auctions. *J. Eur. Econom. Assoc.* **2** 30–66.

Gul, F., E. Stacchetti. 2000. The English auction with differentiated commodities. *J. Econom. Theory* **92** 66–95.

Hann, I., C. Terwiesch. 2003. Measuring the frictional costs of online transactions: The case of name-your-own-price channel. *Management Sci.* **43** 1563–1579.

Isaac, R., T. C. Salmon, A. Zillante. 2007. A theory of jump bidding in ascending auctions. *J. Econom. Behav. Organ.* **62** 144–164.

Johnson, N. L., S. Kotz. 1970. *Continuous Univariate Distributions*, Vol. 1. John Wiley, New York.

Kephart, J. O., A. R. Greenwald. 2002. Shopbot economics. S. Parsons, P. T. Gmytrasiewicz, M. Woolridge, eds. *Game Theory and Decision Theory in Agent-Based Systems* (*Multiagent Systems, Artificial Societies, and Simulated Organizations*). Kluwer Academic Publishers, Norwell, MA, 119–158.

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

Lavi, R., N. Nisan. 2005. Online ascending auctions for gradually expiring items. *SODA '05*: *Proc. 16th Annual ACM-SIAM Sympos. Discrete Algorithms, Vancouver, British Columbia*. SIAM, Philadelphia, 1146–1155.

Lucking-Reiley, D. 2000. Auctions on the Internet: What's being auctioned and how? *J. Indust. Econom.* **48** 227–252.

Milgrom, P. R. 2000. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* **108** 245–272.

Montgomery, A. L., K. Hosanagar, R. Krishnan, K. B. Clay. 2004. Designing a better shopbot. *Management Sci.* **50** 189–206.

Pagano, M. 1989. Trading volume and asset liquidity. *Quart. J. Econom.* **104** 255–274.

Peters, M., S. Severinov. 2006. Internet auctions with many traders. *J. Econom. Theory* **130** 220–245.

Pinker, E., A. Seidmann, Y. Vakrat. 2003. Managing online auctions. *Management Sci.* **49** 1457–1484.