

The Effects of Proxy Bidding and Minimum Bid Increments within eBay Auctions

ALEX ROGERS, ESTHER DAVID, and NICHOLAS R. JENNINGS

University of Southampton

and

JEREMY SCHIFF

Bar-Ilan University

9

We present a mathematical model of the eBay auction protocol and perform a detailed analysis of the effects that the eBay proxy bidding system and the minimum bid increment have on the auction properties. We first consider the revenue of the auction, and we show analytically that when two bidders with independent private valuations use the eBay proxy bidding system there exists an optimal value for the minimum bid increment at which the auctioneer's revenue is maximized. We then consider the sequential way in which bids are placed within the auction, and we show analytically that independent of assumptions regarding the bidders' valuation distribution or bidding strategy the number of visible bids placed is related to the logarithm of the number of potential bidders. Thus, in many cases, it is only a minority of the potential bidders that are able to submit bids and are visible in the auction bid history (despite the fact that the other hidden bidders are still effectively competing for the item). Furthermore, we show through simulation that the minimum bid increment also introduces an inefficiency to the auction, whereby a bidder who enters the auction late may find that its valuation is insufficient to allow them to advance the current bid by the minimum bid increment despite them actually having the highest valuation for the item. Finally, we use these results to consider appropriate strategies for bidders within real world eBay auctions. We show that while last-minute bidding (sniping) is an effective strategy against bidders engaging in incremental bidding (and against those with common values), in general, delaying bidding is disadvantageous even if delayed bids are sure to be received before the auction closes. Thus, when several bidders submit last-minute bids, we show that rather than seeking to bid as late as possible, a bidder should try to be the first sniper to bid (i.e., it should "snipe before the snipers").

This research was funded by the ARGUS II DARP (Defence and Aerospace Research Partnership) and the DIF-DTC project (8.6) on Agent-Based Control. The ARGUS II DARP is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University, and Southampton University, and is funded by the industrial partners together with the EPSRC, Ministry of Defence (MoD), and Department of Trade and Industry (DTI).

Authors' addresses: A. Rogers, E. David, and N. R. Jennings, Department of Electronics and Computer Science, Southampton University, Southampton, SO17 1BJ, U.K; email: {acr, ed, nrj}@soton.ac.uk; J. Schiff, Department of Mathematics, Bar-Ilan University, Ramat Gan, 5290 Israel; email: Schiff@math.biu.ac.il.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permission@acm.org. © 2007 ACM 1559-1131/2007/08-ART9 \$5.00 DOI 10.1145/1255438.1255441 <http://doi.acm.org/10.1145/1255438.1255441>

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]:—Economics

General Terms: Design, Economics, Theory

Additional Key Words and Phrases: Online auctions, electronic commerce, sniping, proxy bidding, bid increment

ACM Reference Format:

Rogers, A., David, E., Jennings, N. R., and Schiff, J. 2007. The effects of proxy bidding and minimum bid increments within eBay auctions. *ACM Trans. Web.* 1, 2, Article 9 (August 2007), 28 pages. DOI = 10.1145/1255438.1255441 <http://doi.acm.org/10.1145/1255438.1255441>

1. INTRODUCTION

The growth of the Web has ensured that electronic commerce, once the preserve of large corporations, is now an everyday activity for millions of private individuals and small businesses. eBay, the world's most popular online auction house, represents a canonical example of this phenomenon. Through a simple and intuitive Web interface, it allows buyers and sellers to come together within a worldwide virtual market, and its popularity is evidence of both its successful design and the public demand for these services¹.

Described briefly, the eBay auction protocol is a variation of an ascending price auction with a minimum bid increment and a fixed closing time. Buyers interested in bidding within an auction do not submit bids directly. Rather, they must use a proxy bidding system that requires that they specify the maximum amount that they are willing bid (with the constraint that this amount must exceed the current auction price plus a minimum bid increment, d). The proxy bidding system then automatically submits bids on their behalf, and the eBay protocol guarantees that the bidder who has entered the highest amount wins the item, but pays no more than the amount entered by the second highest bidder plus the minimum bid increment. Indeed, eBay recommends that bidders simply enter the maximum amount that they are willing to pay for the item, and leave the proxy bidding system to perform all the resulting bidding (see eBay help: Bid increments), available online at <http://pages.ebay.com/help/buy/bid-increments.html>.

While eBay has been the subject of significant academic research, the majority of this work has been limited to phenomenology (observing what happens in actual auctions and supplying qualitative explanations of the observed effects). As such, it has often been stated that the eBay auction protocol behaves as a second price auction whereby the expected auction revenue is equal to the second highest bidder's valuation plus the minimum bid increment [Bapna 2003; Ockenfels and Roth 2006]. However, this assumption apparently contradicts previous results (including our recent work) that shows that, whenever the bids that may be submitted to the auctioneer are restricted to certain discrete levels, then the auction generates less revenue than the second highest bidders' valuation [Rothkopf and Harstad 1994; David et al. 2007]. To resolve this contradiction, in this article we construct a mathematical model of the eBay

¹In 2006, eBay listed more than 2.2 billion items and generated gross merchandise sales of \$52.2 billion [eBay 2006].

auction protocol. Using this model, we carry out, for the first time, a detailed analysis of how the minimum bid increment and the eBay proxy bidding system affect the properties of the eBay auction.

First, we consider the revenue of an eBay auction, and we calculate analytically the expected revenue in the case that two bidders with private valuation use the eBay proxy bidding system (and, as suggested by eBay, these bidders enter their valuations as the maximum amount that they are willing to pay for the item as soon as they become aware of the auction). We show that, in this case, the revenue of the auction is dependent on the value of the minimum bid increment, and that there exists an optimal value at which this revenue is maximized. At this optimal value, the revenue does indeed exceed the second highest bidder's valuation (but by an amount that is less than the minimum bid increment). To fully understand this effect, we compare it to the case where the two bidders employ a pedestrian bidding strategy (i.e., each time they are outbid, they increase the current bid by the minimum bid increment). We again analytically calculate the expected revenue of the auction, and we show that, in this case, the auction generates a revenue that is less than the valuation of the second highest bidder. Comparing these two results indicates that it is the interaction of the eBay proxy bidding system and the minimum bid increment, and not the minimum bid increment alone, that is responsible for the increased auction revenue. We confirm this result for larger numbers of bidders through simulation.

Second, we consider the fact that, unlike conventional auctions, an eBay auction does not commence with all the bidders being present. Instead, the bidders arrive in a random order that is dependent on the time at which they first become aware of the auction's existence (typically through searching the eBay Web site). We show through simulation that the proxy bids of the early bidders cause the current price of the auction to increase rapidly, and hence bidders who enter the auction late may find that their valuation is insufficient to allow them to advance the current bid by the minimum bid increment. Thus the number of bids that are observed in an eBay auction (and recorded in the auction history) can be substantially smaller than the number of bidders who would have liked to place a bid. We calculate an analytic expression for this relationship in the limiting case that the minimum bid increment $d = 0$ and we show that the number of bids received is approximately proportional to the logarithm of the number of bidders who would have liked to place a bid. This analysis is attractive as it is independent of any assumptions regarding the distribution of bids placed within the auction (or the valuations of the bidders). It also represents the maximum number of bids that may be observed; increasing either the minimum bid increment or the starting bid only reduces the chance that a bidder will have a valuation that is sufficient to allow them to place a bid. This result suggests that care must be taken when attempting to infer the behavior of bidders through observation of the bid history, and thus empirical studies that neglect this effect may be incorporating systematic errors. In addition, it shows that the minimum bid increment introduces an inefficiency in to the auction, in that the bidder with the highest valuation does not always win. We show through simulation that, the earlier bidders submit their valuation to the eBay proxy bidding system, the greater their chance of winning the auction.

This last observation appears to contradict the extensive literature that suggests that last minute bidding or sniping is an effective strategy within an eBay auction [Roth and Ockenfels 2003]. Thus we consider how these observations determine the bidding strategy that a bidder should adopt within a real eBay auction. Our results show that, while late bidding is effective in the case of common values or when other bidders engage in incremental bidding (i.e., rather than using the proxy bidding system as eBay intended, they repeatedly increase the amount that they are willing to pay whenever they are outbid²), in general, due to the inefficiency discussed above, bidding late is disadvantageous. Thus, when several bidders snipe, we show that, rather than seeking to place the last bid, bidders should try to be the first sniper to submit their bid (i.e., they should “snipe before the snipers”).

These results significantly advance the understanding of the eBay auction protocol. They can be used by institutions attempting to construct similar online auctions, since our analysis indicates how the value of the minimum bid increment should be chosen to maximize the revenue that the auction generates. In addition, our observation that the number of bids observed within an eBay auction is significantly smaller than the number of bidders who attempted to place a bid (and the related analytical calculation) is extremely important for researchers performing empirical studies of eBay auctions using the auction bid history. Finally, our insights into the timing of snipe bids can be used by bidders in real eBay auctions to maximize their chance of winning the item.

The remainder of the article is organized as follows. Section 2 presents related work. Section 3 describes the eBay protocol and the proxy bidding system in more detail and Section 4 presents our analysis of the revenue of the eBay auction. Section 5 considers how the sequential update of the current bid affects the auction properties and Section 6 considers the resulting bidding behavior. Finally, we conclude and discuss future work in Section 7.

2. RELATED WORK

The growth of Web-based electronic commerce has initiated much research into the design of novel mechanisms for online auctions [Fontoura et al. 2002], and also effective bidding strategies for automated bidding agents [Guo 2002; Dumas et al. 2002; Anthony and Jennings 2003]. However, much less research has addressed the specific implementation details of auction mechanisms that have proved to be popular and effective in real online settings (such as eBay). What work does exist in this area has tended to focus on observing bidding within eBay auctions (typically through an analysis of the bid history that eBay makes available on its Web site) and then attempting to form qualitative explanations of the observed effects (see Bajari and Hortacsu (2003) and Ockenfels et al. (2007) for reviews of this work). Typical work within this area has considered how the closing price of the auction is affected by the various auction settings, such as the starting bid, the reserve price, the shipping costs, and the

²Pedestrian bidding, where the bid is raised by the minimum bid increment each time, is an extreme form of incremental bidding.

auction duration [Lucking-Reiley 2000; Bajari and Hortacsu 2003; Hossain and Morgan 2006; Gerding et al. 2007].

In addition, other research has considered the behavior of bidders within these auctions [Shah et al. 2003; Bapna 2003; Roth and Ockenfels 2003]. A key observation in this respect has been the fact that in many eBay auctions a significant number of bids are placed in the last few minutes before the auction closes. Many explanations have been proposed for this last-minute bidding and it is commonly believed that sniping is the best response to naïve bidders engaging in incremental bidding [Roth and Ockenfels 2003]. It has also been suggested that, in common-value auctions, such last-minute bidding prevents the disclosure of information, and thus prevents other bidders from responding with an updated bid [Bajari and Hortacsu 2003].³ Furthermore, a number of researchers have suggested that last-minute bidding may also protect a bidder from a form of shill bidding, commonly known as maximum-bid fishing, whereby the seller submits small incremental shill bids to expose the value of the highest bid, and thus forces the highest bidder to pay their full bid [Barbaro and Bracht 2004; Engelberg and Williams 2005].

However, despite this extensive body of research, little work has considered the specific details of the eBay auction protocol, such as the operation of the eBay proxy bidding system, the use of a minimum bid increment,⁴ and the fact that bidders generally arrive within the eBay auction after bidding has commenced. Instead, the eBay auction protocol is commonly assumed to behave as a second-price auction, and it is often stated that the expected auction revenue is equal to the second highest bidders' valuation plus the minimum bid increment [Bapna 2003; Ockenfels and Roth 2006]. In the work that we present here, we show that these assumptions are too simplistic. By constructing a mathematical model of the eBay auction protocol and then performing a detailed analysis of its properties, we are able to gain significant and novel insights into the operation of the eBay auction protocol that can be used by designers of online auctions, by other researchers performing empirical studies of real eBay auctions and finally, by bidders within these auctions.

3. THE EBAY AUCTION

We consider the single-item eBay auction protocol⁵ and, to explain its operation, we describe the sequence of events as the bidding proceeds. The auction process

³Last-minute bidding has also been shown to be an equilibrium strategy in private-valuation auctions where there is some small probability that last-minute bids are dropped [Ockenfels and Roth 2006].

⁴Other work has considered the effect of the minimum bid increment within ascending-price discrete-bid English auctions [Rothkopf and Harstad 1994; David et al 2005, 2007]. However, the additional details of the eBay auction make it difficult to apply these general results in this specific case, and this motivates the more detailed study that we present here.

⁵eBay also offers a multiple-item auction format that they often refer to as a *Dutch auction* (see eBay help: Multiple item auction, available online at <http://pages.ebay.com/help/buy/buyer-multiple.html>). This auction format is actually quite different from the descending price auctions used in the Dutch flower markets that are commonly referred to as *Dutch auctions* within the academic literature of auctions. In the multiple-item eBay auction, bidders specify the quantity of items that they wish to purchase and the maximum price that they are willing to pay per item.

Table I. Minimum Bid Increment Implemented by eBay (see ebay help:Bid Increments, available online at <http://pages.ebay.com/help/buy/bid-increments.html>)

| Current auction price | Bid increment |
|-----------------------|---------------|
| \$0.01–\$0.99 | \$0.05 |
| \$1.00–\$4.99 | \$0.25 |
| \$5.00–\$24.99 | \$0.50 |
| \$25.00–\$99.99 | \$1.00 |
| \$100.00–\$249.99 | \$2.50 |
| \$250.00–\$499.99 | \$5.00 |
| \$500.00–\$999.99 | \$10.00 |
| \$1000.00–\$2499.99 | \$25.00 |
| \$2500.00–\$4999.99 | \$50.00 |
| \$ 5000.00 and up | \$100.00 |

commences with the seller determining the auction duration (i.e., number of days that the auction will run) and the starting bid of the auction, s . This starting bid represents the minimum bid that is acceptable to the seller, and corresponds to the reserve price described in the standard auction literature⁶. Thus, the seller is not committed to selling the item until a bidder submits a bid with a value greater than or equal to this starting bid. However, once a suitable bid has been received, the current bid price of the auction is set to this starting bid and normal bidding can commence. From this point onward, all bids must exceed the current bid by a minimum bid increment, d . This minimum bid increment takes a predetermined value depending on the current bid price (see Table I). For example, when the bid current auction price is between \$1.00 and \$4.99, the minimum bid increment is \$0.25 (see eBay help: Bid increments, available online at <http://pages.ebay.com/help/buy/bid-increments.html>).

Bidders who wish to bid within the eBay auction cannot do so directly, but must use a proxy bidding system that submits bids on their behalf. They must enter the maximum price that they are prepared to bid, and, to be accepted, this amount must be greater than, or equal to, the current bid price plus the minimum bid increment. In practice, eBay recommends that bidders simply enter the maximum price that they are willing to pay for the item (i.e., their valuation), and leave the proxy bidding system to perform all the resulting bidding. While it is commonly assumed that the eBay proxy bidding system engages in pedestrian bidding (i.e., increasing the bid price by the minimum bid increment each time it is outbid), a detailed reading of the eBay bidding instructions and some experimental bidding in real auctions reveals that this is not the case. Rather, the eBay protocol maintains a record of the current

At the close of the auction, all the winning bidders pay the same price, and this price is equal to the lowest successful bid. These auctions are less frequently used by sellers, and when they are used, they typically close at the starting price since supply outstrips demand. Thus they are of less interest and we do not consider them here.

⁶In addition to this starting bid, eBay also allows the seller to set a secret reserve price. If the bidding fails to reach this secret reserve price, then the item remains unsold and can be relisted at a later date. Since this feature incurs an additional charge, and is much less common, we do not consider it here.


```

01  highest_bid = 0
02  current_bid = 0
03  starting_bid = s
04  starting_bid_met = false

05  function HandleBid( bid_price )
06    if starting_bid_met
07      if bid_price ≥ current_bid + d
08        if bid_price ≥ highest_bid
09          current_bid = min(highest_bid+d,bid_price)
10          highest_bid = bid_price
11        else
12          current_bid = min(bid_price+d,highest_bid)
13      else
14        RejectBid
15    else
16      if bid_price ≥ starting_bid
17        starting_bid_met = true
18        current_bid = starting_bid
19        highest_bid = bid_price
20      else
21        RejectBid

```

Fig. 1. Pseudocode for the rule by which the eBay protocol updates the current bid price.

bid price and the highest price entered into the proxy bidding system so far. Then, whenever a new bidder submits a bid (or more accurately, the new bidder informs the proxy bidding system of the maximum price it is prepared to bid), the current bid price immediately advances to the minimum of the highest price entered so far and the second highest price plus the minimum bid increment (see eBay help: Bid increments at the web address given above; see also Figure 1). The difference between this update rule and pedestrian bidding is significant. It means that the current bid price of the auction may take any value and is not restricted to integer multiples of the minimum bid increment, and, as we describe in the next section, it gives rise to an optimal value for the minimum bid increment.⁷ Finally, at the fixed closing time of the auction, the item is allocated to the bidder who has entered the highest amount into the proxy bidding system, and that bidder pays the current bid price.

3.1 Example

We illustrate the eBay protocol with a simple example. Consider an auction for a small item with a starting bid of \$1.99. Some time after the auction starts, the first bidder arrives. This bidder has a valuation for the item of \$3.50, and thus enters this amount into the proxy bidding system. In response, the current bid price of the auction is updated to \$1.99, and the minimum bid increment is set to \$0.25 (see lines 16–19 in Figure 1).

⁷Note that other online auctions, such as the popular Israeli site www.olsale.com, do not provide automatic bidding at all, but instead enforce manually pedestrian bidding by providing a button that advances the current bid price by a fixed bid increment each time it is clicked.

Soon afterward, another bidder notices the auction while browsing the eBay Web site. This bidder has a valuation for the item of \$3.40 and thus enters this amount in the proxy bidding system. Now, since the first bidder has entered a greater amount, a proxy bid is submitted on its behalf. The current bid price of the auction is thus updated to \$3.50 (i.e., the minimum of \$3.50 and $\$3.40 + \0.25 according to line 12 in Figure 1), and the second bidder is immediately informed that it has been outbid by another bidder. Note that if the auction were to close at this point, the bidder with the highest valuation would have won, but it would have paid less than the second valuation plus the minimum bid increment.

Next, a third bidder arrives with a valuation of \$3.60. Now the minimum amount that the proxy bidding system will accept is \$3.75 (i.e., $\$3.50 + \0.25), and thus despite having the highest valuation so far, this bidder is unable to submit a bid (see line 14 in Figure 1). In this case, the minimum bid increment has introduced an inefficiency into the auction.

Finally, a fourth bidder arrives with a valuation of \$4.00 and enters this amount in the proxy bidding system. Since this amount is greater than that entered by the first bidder, the current bid price jumps to \$3.75 (i.e., the minimum of \$4.00 and $\$3.50 + \0.25 according to line 9 in Figure 1). At the closing time of the auction, the item is allocated to the fourth bidder, who pays \$3.75 for it (plus shipping), and an inspection of the auction bid history will reveal just three bids (despite the fact that four bidders would have liked to bid).

Note that in our example, if the third bidder had arrived before the second, then the fourth bidder would still have won the auction, but the closing price would have been \$3.85. If the fourth bidder, with valuation \$4.00, had arrived first then the closing price of the auction could be \$3.65, \$3.75, or \$3.85, depending on which bidder arrived second, and only two bids would be recorded in the bid history.

4. ANALYSIS OF AUCTION REVENUE

Given the preceding description of the eBay auction protocol, in this section we consider in detail the effects that the minimum bid increment and the eBay proxy bidding system have on the expected revenue of the auction. In order to do so, we consider two cases that represent different simple strategies adopted by the bidders. Our purpose at this point is to understand the properties of the eBay auction mechanism, rather than searching for equilibrium bidding strategies for participants. In Section 6, we shall use the results and understanding of the eBay auction mechanism developed here to inform the bidding strategy adopted by bidders within real eBay auctions.

Thus, to this end, we first assume that the bidders use the proxy bidding system, and, as suggested by eBay, they simply enter their valuation for the item as the maximum amount they are willing to bid. They then allow the proxy bidding system to perform all the necessary bidding on their behalf. We then compare this proxy bidding case, to one where bidders perform pedestrian bidding. That is, they monitor the auction and whenever they are outbid by another bidder,

they increase their bid by the minimum possible amount (assuming that this amount is less than their valuation of the item).⁸

From this comparison, we show that, when bidders use the eBay proxy bidding system, there is an optimal value for the minimum bid increment, and, at this optimum, the revenue exceeds the second highest bidder's valuation (but by an amount that is less than the minimum bid increment). However, with pedestrian bidding, we show that the expected auction revenue decreases as the minimum bid increment is increased, and, in fact, the expected auction revenue is always less than that of the second highest bidder's valuation. These results show that it is the interaction of the eBay proxy bidding system and the minimum bid increment, not the minimum bid increment alone, that results in the expected auction revenue exceeding the second highest valuation.

In order to perform this analysis, we assume that N bidders, each with a private valuation that is drawn from a common valuation distribution, with density $f(v)$, in the range $[\underline{v}, \bar{v}]$, participate in an auction with a starting price s and a fixed minimum bid increment d .⁹ We will assume that, as in the real eBay auction, these bidders become aware of the auction at different times, and thus they enter the auction in a random sequence.

We now consider the two cases described above, and, in the first instance, we derive analytical results when $N = 2$ (in Section 4.4 we present simulation results for greater numbers of bidders).

4.1 Using the eBay Proxy Bidding System

To calculate the revenue when two bidders use the eBay proxy bidding system, we must consider the order in which the eBay protocol updates the current auction price. We assume that the first bidder has entered a bid with value v_1 , and thus the closing price of the auction depends on this value, the value of the bid of the second bidder, v_2 , and on s and d .

We illustrate this with an example, considering the case that v_1 is greater than $s + 2d$ but less than $\bar{v} - d$. Since v_1 is greater than the starting bid, s , the bid is accepted, and the current bid price of the auction is updated to s . The

⁸This pedestrian bidding can in principle be achieved by all bidders always entering the minimum possible amount to the existing eBay proxy bidding system, thus always increasing the current bid by the minimum bid increment. However, this would necessitate bidders watching the auction for its entire duration after they submitted an initial bid, so in practice we are thinking about an alternative possible implementation of a proxy bidding system.

⁹We note that the assumption of a fixed minimum bid increment, as opposed to one that increases incrementally with the auction price (as occurs within the real eBay auction and is shown in Table I), makes no difference to our analysis of the case when $N = 2$. In this case, the minimum bid increment is simply determined by the price range into which the auction starting price falls (and not by the value of submitted bids). In Section 4.4 where we present simulation results for a greater number of bidders, we maintain this assumption in order to allow comparison with these earlier results. However, we note that, since the eBay proxy bidding system increases the auction price rapidly, and since the range of the auction prices over which the minimum bid increment is large (e.g., over the range \$100.00 to \$249.99 the minimum bid increment has a value of \$2.00), the majority of the bids will be submitted under a constant minimum bid increment, and thus this is a reasonable assumption.




Fig. 2. Digram showing the expected revenue of the auction, E , when two bidders submit bids of v_1 and v_2 .

expected revenue of the auction, E , will now depend on the bid of the second bidder, v_2 , and there are five different cases to consider. These are shown in Figure 2 by the five regions labeled (1) to (5). If v_2 is less than $s + d$, this second bid will not be accepted, and thus the first bidder will be allocated the item and will pay the current bid of s (region 1 in Figure 2). If, however, v_2 is greater than or equal to $s + d$, but less than v_1 , then the second bid will be accepted, but the item will be allocated to the first bidder, who must pay the minimum of $v_2 + d$ and v_1 (regions 2 and 3). Finally, if v_2 is greater than v_1 , then this time the item will be allocated to the second bidder, and this bidder must pay the minimum of $v_1 + d$ and v_2 (regions 4 and 5).

To calculate the revenue in the general case, we must consider all the possible values of s and v_1 . In doing so, to simplify the analysis somewhat we make the explicit assumption that $s + 2d \leq \bar{v}$. In the case of a bidders' valuation distribution with infinite support (such as an exponential distribution), this assumption is always satisfied. However, in the case of distributions with finite support (such as a uniform distribution), it ensures that it is always possible to raise the current auction price by the minimum bid increment at least twice before the upper limit of the distribution is exceeded (see Appendix A for results when this condition is relaxed). We present these cases, and the resulting auction revenue, below:

$\underline{v} \leq v_1 < s$:

$$E = \begin{cases} 0, & \underline{v} \leq v_2 < s, \\ s, & s \leq v_2 < \bar{v}; \end{cases} \quad (1)$$

$s \leq v_1 < s + d$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_2, & s + d \leq v_2 < v_1 + d, \\ v_1 + d, & v_1 + d \leq v_2 \leq \bar{v}; \end{cases} \quad (2)$$

$s + d \leq v_1 < s + 2d$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_1, & s + d \leq v_2 < v_1, \\ v_2, & v_1 \leq v_2 < v_1 + d, \\ v_1 + d, & v_1 + d \leq v_2 \leq \bar{v}; \end{cases} \quad (3)$$

$s + 2d \leq v_1 < \bar{v} - d$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_2 + d, & s + d \leq v_2 < v_1 - d, \\ v_1, & v_1 - d \leq v_2 < v_1, \\ v_2, & v_1 \leq v_2 < v_1 + d, \\ v_1 + d, & v_1 + d \leq v_2 \leq \bar{v}; \end{cases} \quad (4)$$

$\bar{v} - d \leq v_1 \leq \bar{v}$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_2 + d, & s + d \leq v_2 < v_1 - d, \\ v_1, & v_1 - d \leq v_2 < v_1, \\ v_2, & v_1 \leq v_2 \leq \bar{v}; \end{cases} \quad (5)$$

Now, to calculate the expected auction revenue we must consider each individual case and weight the revenue that it generates by the probability that it occurs. We do so by integrating over the bidders' valuation distribution, $f(v)$, between the limits \underline{v} and \bar{v} , and thus we derive the result expression shown below:

$$\begin{aligned} E &= \int_{\underline{v}}^s f(v_1) \int_s^{\bar{v}} s f(v_2) dv_2 dv_1 \\ &+ \int_s^{s+d} f(v_1) \left[\int_{\underline{v}}^{s+d} s f(v_2) dv_2 \right. \\ &+ \left. \int_{s+d}^{v_1+d} v_2 f(v_2) dv_2 + \int_{v_1+d}^{\bar{v}} (v_1 + d) f(v_2) dv_2 \right] dv_1 \\ &+ \int_{s+d}^{s+2d} f(v_1) \left[\int_{\underline{v}}^{s+d} s f(v_2) dv_2 + \int_{s+d}^{v_1} v_1 f(v_2) dv_2 \right. \\ &+ \left. \int_{v_1}^{v_1+d} v_2 f(v_2) dv_2 + \int_{v_1+d}^{\bar{v}} (v_1 + d) f(v_2) dv_2 \right] dv_1 \\ &+ \int_{s+2d}^{\bar{v}-d} f(v_1) \left[\int_{\underline{v}}^{s+d} s f(v_2) dv_2 \right. \\ &+ \left. \int_{s+d}^{v_1-d} (v_2 + d) f(v_2) dv_2 + \int_{s+d}^{v_1} v_1 f(v_2) dv_2 \right. \\ &+ \left. \int_{v_1}^{v_1+d} v_2 f(v_2) dv_2 + \int_{v_1+d}^{\bar{v}} (v_1 + d) f(v_2) dv_2 \right] dv_1 \\ &+ \int_{\bar{v}-d}^{\bar{v}} f(v_1) \left[\int_{\underline{v}}^{s+d} s f(v_2) dv_2 + \int_{s+d}^{v_1-d} (v_2 + d) f(v_2) dv_2 \right. \\ &+ \left. \int_{v_1-d}^{v_1} v_1 f(v_2) dv_2 + \int_{v_1}^{\bar{v}} v_2 f(v_2) dv_2 \right] dv_1. \end{aligned} \quad (6)$$

At this point, the results are valid for any bidders' valuation distribution; however, since we will compare this analytical calculation to simulation

results, we must choose a particular distribution to consider. Thus, we make the common assumption that the bidders' valuations are drawn from a uniform distribution with support $[0, 1]$. Hence, $f(v) = 1$, $\underline{v} = 0$ and $\bar{v} = 1$, and evaluating the integrals is straightforward. The final result is an expression, in terms of the minimum bid increment, d , and the auction starting price, s , for the expected revenue of the auction:

$$E = \left(\frac{1}{3} + s^2 - \frac{4s^3}{3} \right) + (1 - 2s + s^2)d - \frac{5(1-s)}{2}d^2 + \frac{4d^3}{3}. \quad (7)$$

This expression shows that the expected auction revenue is dependent on both the starting bid of the auction and the minimum bid increment. It has a constant term that is equal to the revenue of a conventional English auction, and it has a maximum value (over the region $s \geq 0$, $d \geq 0$ and $s + 2d \leq 1$) when $s = 0$ and $d = \frac{1}{4}$. Thus the auction generates the maximum revenue when the starting bid of the auction is zero, and an optimal bid increment is selected. For the case that we consider here, at this optimal configuration the expected revenue of the eBay auction is $\frac{43}{96}$.

Interestingly, the revenue at this maximum is greater than that achieved by a conventional English auction using an optimal starting price and infinitesimal bid increments. In this case, the optimal starting price, s^* , can be calculated from the standard auction literature (i.e., the result that $s^* = [1 - F(s^*)]/f(s^*)$ where $F(\cdot)$ is the cumulative distribution function of the bidders' valuation distribution [Riley and Samuelson 1981]). For the uniform distribution that we consider here, $s^* = \frac{1}{2}$ and the revenue at this optimum is $\frac{5}{12}$. We note that the introduction of the minimum bid increment has produced a substantial lower optimal starting price, and, indeed, phenomenological studies have indicated that auctions with lower starting prices produce higher revenue, though there are many possible justifications for this [Bajari and Hortacsu 2003]. If we consider the case where $s = 0$, the expression in shown Equation (7) simplifies to

$$E = \frac{1}{3} + d - \frac{5d^2}{2} + \frac{4d^3}{3}. \quad (8)$$

Here it is clear that, as d is reduced to zero, the expected revenue of the auction falls to the expected valuation of the second highest bidder (i.e., $\frac{N-1}{N+1}$ or in this case $\frac{1}{3}$).

4.2 Using Pedestrian Bidding

In order to calculate the revenue when bidders engage in pedestrian bidding, we use previously derived results for the revenue of an ascending price English auction with discrete bid levels. In these auctions, the bids that may be submitted to the auctioneer are restricted to certain discrete bid levels, and, in general, this restriction leads the auction to generate less revenue than the equivalent continuous bid auction [David et al. 2007]. In our previous work, we

have shown that the expected revenue of these auctions is given by

$$E = \sum_{i=0}^{m-1} \frac{F(l_{i+1})^N - F(l_i)^N}{F(l_{i+1}) - F(l_i)} [l_i(1 - F(l_i)) - l_{i+1}(1 - F(l_{i+1}))] + l_m [1 - F(l_m)^N], \quad (9)$$

where N is the number of bidders, $F(\cdot)$ is again the cumulative distribution function of the bidders' valuation distribution, and there are $m + 1$ discrete bid levels with values $l_0 \cdots l_m$.

Now, the case in which two bidders engage in pedestrian bidding within an eBay auction is exactly analogous to this discrete bid auction. We can set $l_i = s + id$ and $m = \lfloor \frac{1-s}{d} \rfloor$ (where i is the integer index of the i th bid level and $\lfloor \cdot \rfloor$ denotes the floor function), and thus, perform the summation in Equation (9) analytically, to express the expected revenue as:

$$E = \left(\frac{1}{3} + s^2 - \frac{4s^3}{3} \right) - \frac{1-s}{3} d^2 + \frac{\mathcal{F}(1 - \mathcal{F}^2)}{3} d^3, \quad (10)$$

where $\mathcal{F} = \text{frac}(\frac{1-s}{d})$ — the fractional part of $\frac{1-s}{d}$. Note that, in this expression, the constant term is identical to that for the case of proxy bidding shown in equation (7). Thus both bidding strategies converge to the same expected revenue when $d = 0$ (and this revenue is equal to that generated by a conventional English auction). However, unlike in the case of proxy bidding, this expression does not contain a term proportional to d . Thus, given the condition that it is possible to make at least one bid (i.e., $\frac{1-s}{d} \geq 1$), then it is easy to show that $\frac{1-s}{3} d^2 - \frac{\mathcal{F}(1-\mathcal{F}^2)}{3} d^3 \geq 0$ (since $\mathcal{F} < 1$), and, hence, the revenue in the case of pedestrian bidding is always less than the equivalent continuous bid auction.¹⁰

4.3 Revenue Comparison

Having calculated the expected auction revenue when two bidders use the eBay proxy bidding system or engage in pedestrian bidding, we compare these analytical results to simulation results. Thus, we simulate the behavior of the bidders, the auction protocol, and the proxy bidding system, and we then hold repeated auctions with randomly drawn bidders' valuations. In each auction, we record the closing price, and thus, we calculate the mean auction revenue.

¹⁰Rothkopf and Harstad [1994] considered a similar case in which two bidders, with valuation drawn from a uniform distribution with support $[0, 1]$, engage in pedestrian bidding within an oral ascending price English auction. They showed that if there are $m + 1$ evenly spaced bid levels, with $l_0 = 0$ and $l_m = 1$, then the expected revenue of the auction is given by

$$E = \frac{1}{3} - \frac{1}{3m^2}. \quad (11)$$

When $s = 0$ and $d = 1/m$, the expression in Equation (10) simplifies to give the same result, and thus clearly shows that increasing the number of bid levels (effectively decreasing the minimum bid increment) causes the revenue to more closely approximate the continuous bid case.

Furthermore, they showed that show raising the bid by the minimum bid increment, and not jump bidding (i.e., raising the bid by greater than this minimum increment), is an equilibrium bidding strategy in the case of nondecreasing valuation distributions such as the uniform distribution considered here.

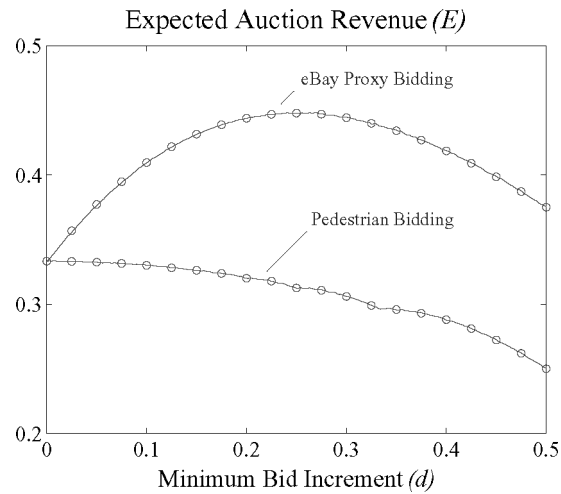


Fig. 3. Comparison of simulation (*solid lines*) and calculated (*circles*) results showing the dependence of the expected auction revenue on the minimum bid increment for both the eBay proxy bidding system and pedestrian bidding. Results are for 2 bidders with valuations drawn uniformly on $[0,1]$ and $s = 0$. Simulation results are averaged over 500,000 auctions.

Figure 3 shows simulation results (for proxy bidding and pedestrian bidding) as solid lines for $s = 0$ and a range of minimum bid increments, from 0 to 0.5. In this plot, and all subsequent ones, the simulations are repeated a sufficient number of times (in this case 500,000 times) that the standard error in the mean auction revenue approaches the thickness of plotted line, and thus error bars are not shown. In this figure, we also show (as circles) the analytical results calculated in the previous section. Clearly, the simulation and calculated results match exactly.

As previously described, the eBay proxy bidding system shows an optimal value for the minimum bid increment, at which the expected revenue reaches a maximum that is greater than the expected second highest valuation (but less than the expected second highest valuation plus the minimum bid increment). The cause of the optimal bid increment is clear. When the bid increment is small, increasing its value directly increases the closing price of the auction, and thus the expected auction revenue increases with increasing bid increment. However, increasing the bid increment also increases the probability that one of the bidders is excluded from the auction (this bidder will be unable to raise the current auction price by the minimum bid increment). Since the closing price is dependent on the bid of the second highest bidder, increasing the bid increment beyond a certain value results in a decrease in the expected auction revenue.

However, in the case of pedestrian bidding, the expected revenue decreases as the minimum bid increment increases, and our previous analysis has shown that this is a general phenomenon that occurs whenever the bidding within an auction is restricted to discrete bid levels (see David et al. [2007] for a more detailed explanation of this effect). Thus these results show that it is the

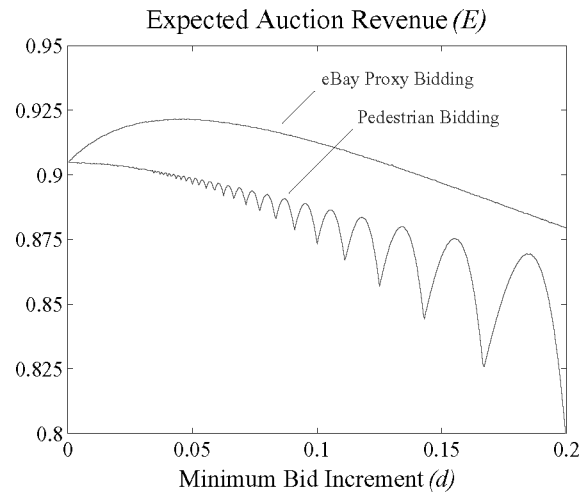


Fig. 4. Simulation results showing the dependence of the expected auction revenue on the minimum bid increment for both the eBay proxy bidding system and pedestrian bidding. Results are for 20 bidders with valuations drawn uniformly on $[0,1]$ and $s = 0$. Results are averaged over 500,000 auctions.

interaction of the eBay proxy bidding system and the minimum bid increment, rather than the minimum bid increment alone, that results in the increase in revenue seen in the eBay auction protocol.

4.4 More Than Two Bidders

When there are more than two bidders, the analysis of both proxy bidding and pedestrian bidding becomes more complex. It is possible to obtain analytic expressions for the revenue for any fixed N , but these are increasingly long and difficult to work with. However, we can easily perform simulations, and thus Figure 4 shows comparisons of simulation results for the two bidding behaviors when $N = 20$. As before, we see that the eBay bidding proxy shows an optimal bid increment (at a smaller value of d than before), at which the expected revenue has a maximum value. However, pedestrian bidding results in a general decrease in auction revenue as the bid increment increases. As expected, both curves meet at the point where $d = 0$, and the revenue at this point corresponds to the expected valuation of the second highest bidder.¹¹

Figure 5 shows more simulation results for the eBay bidding proxy when there are 10, 20, and 40 bidders participating. These results show that, as the number of bidders increases, the size of the optimal bid increment decreases. In the case of the uniform bidders' valuation distribution, the optimal bid increment is approximately equal to $\frac{1}{N+1}$. While we have not been able to fully

¹¹Note that the pedestrian bidding plot shows a number of minima, or “corners,” that correspond to points where an integer multiple of the minimum bid increment falls exactly on the upper limit of the bidders' valuation distribution. This effect also occurs in Figure 3, and is captured by the analytical results (by the term dependent on \mathcal{F}), but is much less apparent due to the smaller number of bidders.

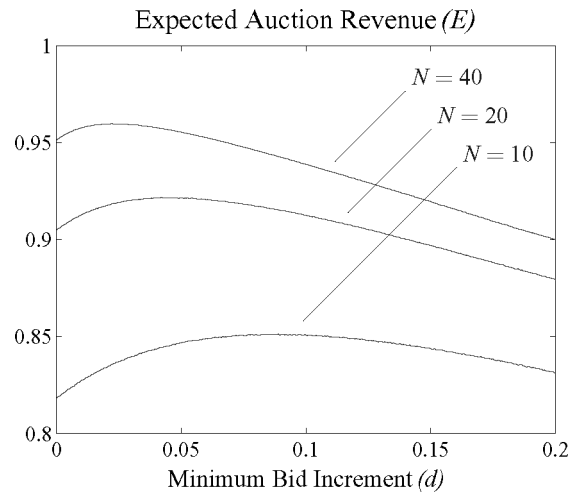


Fig. 5. Simulation results showing the dependence of the expected auction revenue on the number of bidders using proxy bidding. Results are shown for 10, 20, and 40 bidders with valuations drawn from a uniform distribution on $[0,1]$. The starting price $s = 0$ and results are averaged over 500,000 auctions.

characterize this result analytically, through simulation we have observed that the optimal bid increment is closely related to the expected difference in value between the two highest bids received by the auctioneer (in the case of a uniform distribution on $[0,1]$ the expected difference between any two bids is equal to $\frac{1}{N+1}$). We can understand this intuitively since increasing the bid increment further increases the probability that the bidders with the highest valuations will be excluded from the auction (i.e., their valuation will not be sufficient to allow them to raise the current auction price by the minimum bid increment), while decreasing the bid increment will reduce the amount that the winning bidder pays, and in doing so will decrease the revenue.

We have confirmed this result for a number of different valuation distributions. For example, Figure 6 shows similar results when the bidders' valuations are drawn from a normal distribution with mean 100 and standard deviation 5. In this case, the optimal bid increment is again closely related to the expected difference in value between the highest bids, but, due to the tail of the distribution, there is much less dependence on the number of bidders. It is interesting to note that, in this example, the values of the mean and standard deviation were chosen to match data from a study of real eBay auctions for GameBoy Advance consoles [Shah et al. 2003]. For this price range (\$100.00 to \$249.99), the actual eBay bid increment is \$2.50, and is thus close to the optimum shown in our simulations.¹²

¹²One criticism of this correspondence may be that the number of bidders shown in our simulation is much greater than that generally observed in real eBay auctions. However, in the next section, we show that this is not necessarily the case, and in general the number of bids observed in the auction is much smaller than the number of bidders who attempted to place a bid (and are denoted by N here).

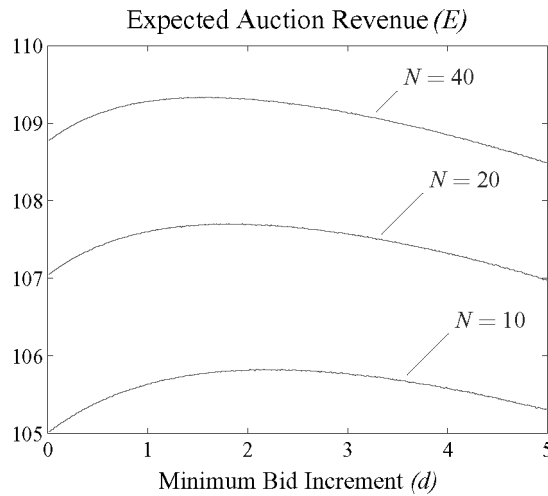


Fig. 6. Simulation results showing the dependence of the expected auction revenue on the number of bidders using proxy bidding. Results are shown for 10, 20, and 40 bidders with valuations drawn from a normal distribution $N[100, 5]$. The starting price $s = 0$ and results are averaged over 500,000 auctions.

These results indicate that in the context of eBay auctions the minimum bid increment is a *design parameter*, and thus it is a variable which the auctioneer may wish to adjust for different auctions in order to maximize the expected revenue. This is in contrast to other auction protocols, in which the minimum bid increment (or choice of discrete bid levels) is a forced necessity, which helps control the duration of the auction, but reduces revenue [Rothkopf and Harstad 1994; David et al. 2007].

5. ANALYSIS OF THE BID HISTORY

In the previous section, we considered the effects that the eBay proxy bidding system and the minimum bid increment have on the revenue of the auction. In this section, we continue our analysis of the proxy bidding system, and consider the fact that, unlike conventional auctions, an eBay auction does not commence with all the bidders being present. Instead, the bidders arrive in a random order that is defined by the time at which they first become aware of the auction's existence. Since the eBay bidding proxy protocol immediately updates the current bid price whenever a new bid is submitted, many of the bidders who enter the auction will find that the current price of the auction has already advanced to a price such that they are unable submit an acceptable bid (i.e., they can not advance the auction price by the minimum bid increment).

This effect leads to two important observations, and we present an analysis of both here. Firstly, we show that the number of bids that are observed in an eBay auction (and recorded in the auction history) can be substantially smaller than the number of bidders who would have liked to place a bid. Second, we show that the minimum bid increment introduces an inefficiency into the auction, and thus the probability of any bidder winning the auction is dependent on when

the bidder attempts to place its bid. Bidders who submit their valuations to the eBay bidding proxy earlier in the auction have a greater chance of winning the auction.

5.1 Expected Number of Bids Observed

In our simulations of the eBay auction protocol, we have assumed that N bidders participate within the auction and attempt to place a bid. However, as described above, many of these bidders find that their valuation is insufficient to allow them to advance the current bid price by the minimum bid increment, and thus their bid is not accepted by the proxy bidding system. Thus, despite there being N potential bidders, the number of bids observed in the bid history, which we will denote n , can be less than N .

In the general case, the number of bids observed will depend on the the starting price, the bid increment, the bidders' valuation distributions, and the order in which the bidders attend the auction. However, we are able to calculate analytically the expected number of bids observed by considering the special case when $d = 0$ and $s = 0$. This case is attractive as it represents the maximum number of bids that are observed; increasing either the minimum bid increment or the starting bid only reduces the chance that a bidder will have a valuation sufficient to submit a bid. Most importantly, the result is completely independent of any assumption regarding the distribution of the bids submitted and is thus independent of assumptions the bidders' valuation distributions (including the assumption that they all share a common valuation distribution) and their bidding strategies.

Thus, in order to calculate this upper bound, we consider that N bidders come to the auction in a random order and each submits a single bid into the eBay proxy bidding system. We make no assumptions regarding the specific timing of these bids, but only that order in which the bids are submitted is independent of their valuation. Thus our model covers the case that bidders behave as instructed by eBay and submit their valuation to the proxy bidding system as soon as they become aware of the auction.¹³

In this case, with both $d = 0$ and $s = 0$, the bids submitted by the first two bidders will always be accepted. Subsequent bids will only be accepted if the bid amount is greater than the second highest bid received so far. Thus, if we denote by \mathcal{H}_N the expected number of times that the second highest bid received changes as we sequentially consider the N bids, then, the expected number of bids observed, n , is given by

$$n = 1 + \mathcal{H}_N. \quad (12)$$

¹³Although it is not our principal interest here, it is important to note that this model also applies to cases where bidders may also attempt to snipe, and thus submit a last minute bid. In addition, we are not restricted to considering truthful bidding so long as we assume that these decision are independent of the bidder's valuation. This may not actually be the case in real eBay auctions since such strategic decisions are likely to be influenced in some way by the bidder's valuation and the current price of the auction when they first discover it. Indeed, there may also be correlations between the bidders' valuation and the time at which the auction was first noticed (i.e., bidders with higher valuations may seek out auctions more proactively than bidders with low valuations)

We can derive a recursive expression for \mathcal{H}_N by considering the effect of adding one more bid to the original sequence of N bids. If this additional bid has the highest, or the second highest value, out of all $N + 1$ bids, then it will be accepted. The probability of this occurring is identical to the probability that the value of the last bid in a random permutation of $N + 1$ bids is the highest or second highest out of all $N + 1$ bids. It is simply given by $\frac{2}{N+1}$, and thus, we have

$$\mathcal{H}_{N+1} = \mathcal{H}_N + \frac{2}{N+1}. \quad (13)$$

Now, when there are just two bids, the second highest bid can change once and once only, and thus $\mathcal{H}_2 = 1$. Using this result and the recursive expression above gives the following simple result:

$$\mathcal{H}_N = \sum_{j=2}^N \frac{2}{j}. \quad (14)$$

We can approximate this expression for large N using the standard asymptotic approximation for the harmonic series:

$$\sum_{j=1}^N \frac{1}{j} = \ln(N) + \gamma + \frac{1}{2N} + O(N^{-2}), \quad (15)$$

where γ is Euler's constant, which has a value of 0.577215665. Thus, for large N , the expected number of bids observed is given approximately by

$$n \approx 2 \ln(N) + (2\gamma - 1) + \frac{1}{N}. \quad (16)$$

Figure 7 shows a comparison of this analytical expression and simulation results, and the agreement between the two is extremely good, even for relatively small N . Most striking about this result is the fact that the number of bids observed is significantly less than the number of bidders who attempt to place bids. For example, when 40 bidders attempt to place bids, on average less than eight of them are successful. The logarithmic relationship between the number of potential bidders and the bids observed is particularly striking since it contrasts with the more familiar result that the longest increasing subsequence in a random permutation follows a \sqrt{N} relationship as N becomes large [Johansson 1998].

Note again that this is a general result that is independent of any assumptions regarding how the value of the bids was derived or the timing of their submittal. It is dependent solely on the sequential manner in which the eBay auction protocol updates the current bid price (and only assumes that the order that they are submitted is independent of their valuation). In addition, it represents an upper bound in that increasing the starting price, s , or the bid increment, d , generally results in fewer bids being observed. We illustrate this in Figure 7 by also showing the number of expected bids received for several values of d when bidders' valuations are uniformly drawn on $[0,1]$ and, as before, bidders directly enter their valuations into the proxy bidding system. As

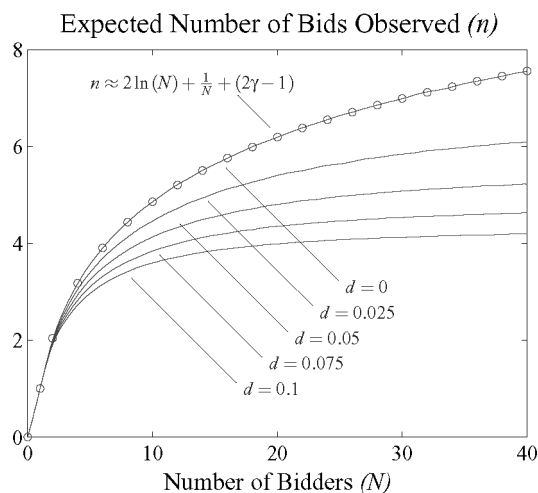


Fig. 7. Comparison of simulation (*solid lines*) and calculated (*circles*) results showing the number of bids observed (n) compared to the number of bidders who attempted to place a bid (N). The starting price $s = 0$ and for results where $d > 0$, bidders' valuations are drawn uniformly on $[0,1]$. Results are averaged over 500,000 auctions.

expected, as the minimum bid increment increases, the number of successfully submitted bids decreases.

This result is significant as it indicates that the bid history (i.e., the record of submitted bids that eBay records) is far from complete and thus empirical studies of eBay auctions that use this bid history in a naïve manner will incorporate systematic errors into their results. A number of researchers have noted similar findings, and have presented approaches to attempt to deal with this issue. For example, Harstad and Rothkopf (2000) considered an oral English auction and noted that many of the bidders who are present and willing to bid may not actually submit bids since they are unable to attract the attention of the auctioneer before the auction price exceeds their private valuation. They proposed an “alternating recognition” model of oral English auctions as an alternative to the model proposed by Milgrom and Weber (1982) in which the prices at which all bidders depart the auction are directly observed. This “alternating recognition” model results in a lower expected revenue for the auctioneer. Similarly, Haile and Tamer (2003) recognized that the bids submitted by bidders within an English auction with private valuations may not necessarily represent the maximum amount that they are willing to pay since again the auction price may advance past their valuation before they can attract the auctioneer’s attention (although they still assumed that the total number of potential bidders is known). They used this observation to calculate a nonparametric bound on the bidders’ valuation distribution. Finally, Jiang and Leyton-Brown (2007) described the problem of missing bids in eBay auctions, and used an expectation-maximisation algorithm to estimate distributions of their number and valuation (although they ignored the minimum bid increment and also assumed that the maximum bid received from any bidder represents its private valuation).




Fig. 8. Simulation results showing the probability of each bidder winning the auction depending on the order in which they bid. There are 20 bidders with valuations drawn uniformly on $[0,1]$ and $s = 0$. Results are averaged over 10^7 auctions.

However, none of these researchers presented a quantitative analysis of this effect, and in particular the logarithmic dependence of n on N has not been previously exposed. This dependence indicates that, when N is large, the majority of bidders are hidden from the auctioneer, and thus visible bidders are a small sample of a larger population.

5.2 Effect of Bid Timing

As well as affecting the number of bids that are observed in the bid history, the fact that bidders may find that they are unable to advance the current bid price by the minimum bid increment also has a key impact on whether or not the bidder with the highest valuation does in fact win the auction. When $d = 0$, the bidder with the highest valuation is always able to submit its bid, and thus will always win the auction. However, when $d > 0$ this is not the case. It is possible for the current price of the auction to have increased to the point that the bidder who has the highest valuation is actually unable to bid, since its valuation is less than the current price plus the minimum bid increment. This introduces a source of inefficiency into the eBay protocol, and thus, when $d > 0$, the order in which the bids are submitted has an effect on determining the winner of the auction (an example of this was presented in Section 3.1).

Figure 8 shows simulation results for an auction in which 20 bidders attempt to place bids. In this plot, the x -axis denotes the order in which the bidders attempt to place their bids within the auction (i.e., 1 is the first bidder and 20 is the last bidder) and the y -axis is the probability of this bidder winning the auction. Each bidder's valuation is drawn uniformly in the range $[0, 1]$, and the bidders submit this valuation to the eBay proxy bidding system. Since the bidders are symmetrical, when $d = 0$ all the bidders have an equal 1 in 20 chance of winning the auction (i.e., the figure shows a straight line at 0.05).

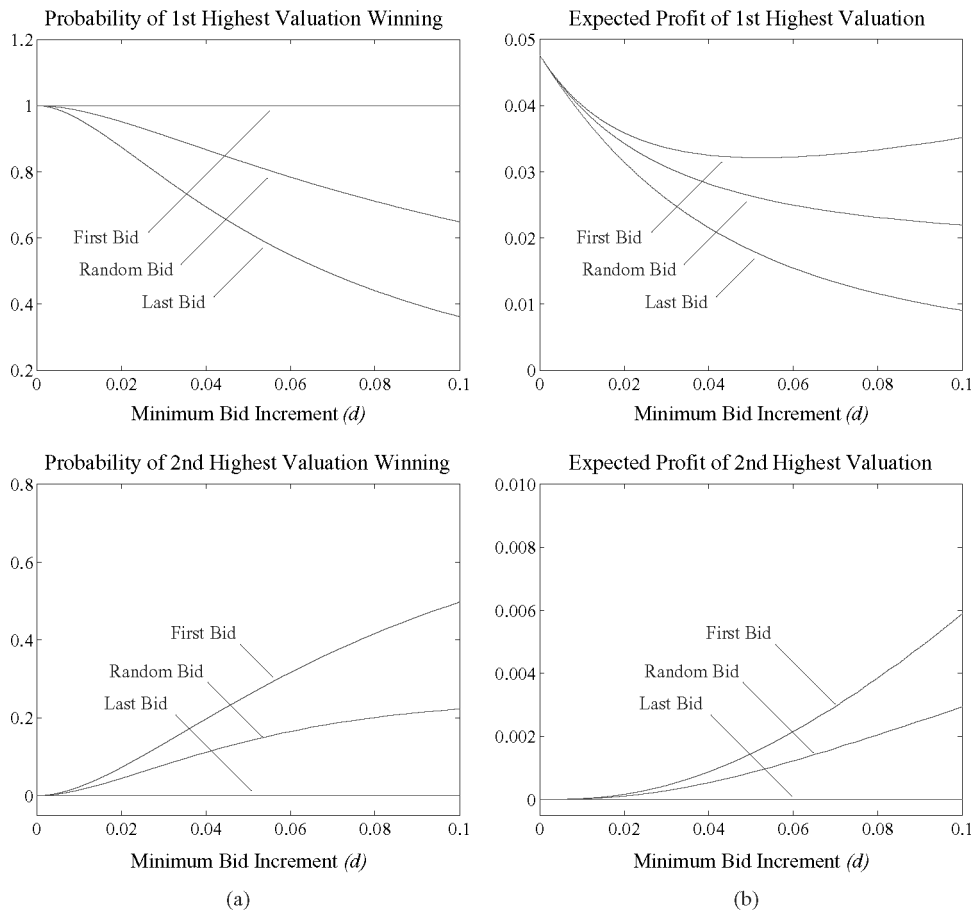


Fig. 9. Simulation results showing (a) the probability of winning, and (b) the expected profit, when the bidders with the highest and second highest valuations bid first, at a random time, and last. There are 20 bidders with valuations drawn uniformly on $[0,1]$ and $s = 0$. Results are averaged over 10^7 auctions.

However, as the value of the bid increment is increased, those bidders who bid early have a greater chance of winning. This is because the bidders who arrive later, while possibly having greater valuations, will find that the current price of the auction has increased to a point that leaves them unable to submit their bid.

Figure 9 shows similar results, but this time we consider just the bidders with the highest and second highest valuations. Specifically, in Figure 9(a), we show the probability of these bidders winning the auction, when they bid first, last, and at a random time. In Figure 9(b), we consider the same three cases, but we show the expected profit of the bidders with the highest and second highest valuations. We first note that the bidder with the highest valuation will always win the auction when either $d = 0$ or it bids first. Conversely, the bidder with the second highest valuation will never win if either $d = 0$ or it bids last.

However, as d increases, both bidders can increase their probability of winning by bidding earlier. This effect is mirrored in Figure 9(b). As expected, when the bidder with the highest valuation bids first (and thus wins the auction), its expected profit is greatest. However, when this bidder does not bid first, its expected profit decreases. Likewise, by bidding early, the bidder with the second highest valuation can increase its expected profit (since it now has a finite chance of winning the auction). Most notable in this plot is the fact that, while the effect of the bid increment on the auctioneer's revenue is relatively small, it has a most significant effect on the expected profit of the winning bidder. Thus the impact of the minimum bid increment is not felt by the auctioneer, but rather by bidders who attempt to submit their bids late and then find that they are not accepted by the proxy bidding system since they do not exceed the current auction price by the minimum bid increment. Indeed, this may occur despite them actually having the highest valuation for the item (and thus they would have won the auction had they bid earlier).

These results suggest that it is always better to submit a bid as early as possible in the auction. However, this observation would appear to contradict the literature that suggests that sniping is an effective and prevalent strategy within an eBay auction [Roth and Ockenfels 2003]. Thus, having completed our analysis of the properties of the eBay protocol, we consider how these properties determine the strategy that a bidder should adopt.

6. BIDDING BEHAVIOR

The analysis of the eBay auction protocol that we presented in the previous sections indicates that the timing of bids has a significant effect on the outcome of the auction. In particular, the analysis shows that bidders with private valuations, who enter these valuations as the maximum amounts that they are prepared to bid, should submit their bids as soon as they become aware of the auction. Indeed, to do so is an equilibrium strategy since delaying bidding always increases the probability that the current auction price will advance to a point such that it is impossible to submit a bid at all.

However, it seems this bidding behavior is not commonly observed in real eBay auctions. Rather, many bidders engage in incremental bidding, while others refrain from bidding until the last few minutes of the auction [Bajari and Hortacsu 2004]. A number of reasons have been proposed for this behavior. First, incremental bidding has been ascribed to the presence of inexperienced or naïve bidders within the auction. These bidders either do not understand the proxy bidding system, or are simply not prepared to enter their full valuations. Sniping is an effective response to these bidders. By submitting a bid in the last minutes of the auction, the snipers leave the incremental bidders insufficient time to respond, and thus do not have to outbid the full valuation of these bidders [Roth and Ockenfels 2003]. Second, in a common-value auction, where incremental bidding may be a sign of bidders actually updating their valuation of the item in response to other bids, sniping is again favored [Bajari and Hortacsu 2003]. By bidding early, a bidder may signal information about its valuation to other bidders. This leaking of information is avoided by

submitting a bid so close to the end of the auction that other bidders do not have time to update their valuations and respond with a revised bid. Finally, sniping may also protect a bidder from a form of shill bidding, commonly known as *maximum bid fishing*, whereby the seller submits small incremental shill bids to expose the value of the highest bid [Barbaro and Bracht 2004; Engelberg and Williams 2005].

Thus, in real eBay auctions there are clearly very good reasons for sniping, and a number of software tools are available to automate this behavior. These tools range from software programs that may be run on the bidders' own computers (e.g., JBidwatcher), to commercial online sniping services that either charge a fixed subscription fee or variable fee depending on the final auction closing price (e.g., eSnipe or AuctionStealer). They typically allow bids to be submitted in the last 20 s of the auction¹⁴ and measure Internet server delays to ensure that bids reach the eBay server at the correct time. Thus, given the ease of sniping, it is expected that within competitive auctions a number of bidders will use this strategy, and that this trend will increase in the future [Bapna 2003].

Given that by definition sniping does not allow other bidders time to respond with an updated bid (and this is almost certainly true in the last 20 s of the auction), in the final stages of an auction we should see exactly the type of bidding that we have modeled and analyzed here. That is, bidders must submit a single bid to the proxy bidding system that represents their final valuations for the item. This is even true in the case of common-value auctions, since from this point onward no more information can be revealed to the other bidders [Song 2004]. Thus, given the reasons described above, bidders in real eBay auctions should indeed snipe, but the analysis presented within this article indicates that they should refrain from bidding in the very last seconds of the auction. They should attempt to submit their bids late enough so that other bidders cannot respond with updated bids, but no later, to prevent themselves being excluded by similar bids submitted by their competitors.¹⁵ In other words, they should aim to "snipe before the snipers."

7. CONCLUSIONS AND FUTURE WORK

In this article, we have performed a detailed mathematical analysis of the eBay auction protocol. Our aim has been to understand how the particular features of the eBay auction, and specifically the minimum bid increment and proxy bidding system, affect the properties of the auction and the resulting behavior of bidders. We initially considered the revenue of the auction, and challenged the common assumption that the eBay protocol acts as a second price auction in which the expected revenue is equal to the second highest bidder's valuation

¹⁴Some online sniping services fix this time, others allow it to be specified by the bidder. In addition, others provide a tiered service depending upon payment. For example, AuctionStealer provides a free service to bid 10 seconds before the end of the auction, and a premium service to bid 3 seconds before the end.

¹⁵Note that this approach also minimizes the possibility that a snipe bid accidentally arrives after the auction has closed. This is a real risk when attempting to bid in the very last seconds of the auction, and is one that we do not explicitly consider here.

plus the minimum bid increment. We showed that this assumption is too simplistic, and in fact, the proxy bidding system and the minimum bid increment interact with one another to yield an optimal bid increment that depends on the number of participants.

We then considered the fact that bidders within an eBay auction are not all present at the start of the auction, but instead, arrive in sequence depending on when they first become aware of the auction's existence. When these bidders use the proxy bidding system, the result is that the current bid price of the auction increases rapidly, and thus the number of bids observed within the auction can be significantly smaller than the number of bidders who would have liked to place a bid. We showed that this leads to an inefficiency within the auction whereby the bidder with the highest valuation is no longer guaranteed to win.

Finally, we considered the bidders' behavior in real eBay auctions. We showed that the bidding behavior that we modeled and analyzed here also describes the bidding in the last seconds of an auction in which several bidders submit snipe bids (and this is true even in the case of common values). As such, we showed that our analysis informs the bidding strategy that should be adopted in this case, and we showed that bidders should "snipe before the snipers." This result contrasts with the common assumption (supported by commercial sniping software that offer premium sniping services in the very last seconds of the auction) that bidding as late as possible is most successful.

Our future work in this area concerns three main issues. First, we would like to characterize through experimentation the optimal time at which to submit a snipe bid. This time will depend on the responsiveness of incremental bidders within typical real-world auctions and is thus impossible to quantify analytically. This task is made more difficult by the fact that unsuccessful snipe bids are hidden from the bid history, and are only observed by the individual sniping service. Second, we would like to extend the model of the eBay auction that we consider. Our analysis has addressed the effect of the starting bid since this most closely parallels the reserve price considered in the standard economic auction literature. However, eBay also allows sellers to declare an invisible reserve price, and also charge a fixed shipping cost. Thus, in practice, all three factors contribute to the revenue of the eBay auction, and we would like to perform a detailed analysis of this. Finally, we would like to extend our analysis of the bidding behavior within the eBay auction to consider fully the effect that naïve, irrational, and common-value bidders may have on the auction revenue. Such considerations are particularly important in the case of online auctions where the experience of bidders varies greatly, and hence no equilibrium bidding strategy is apparent. In extending our analysis in this way, we hope to gain further insights into the behavior of bidders within real online auctions.

APPENDIX

EXPECTED EBAY AUCTION REVENUE IN TO ADDITIONAL CASES

In Section 4.1 we calculated the expected revenue of the eBay auction when two bidders used the proxy bidding system. In doing so we explicitly made the assumption that it was possible to raise the current auction price by the

minimum bid increment at least twice (i.e., $s + 2d \leq \bar{v}$). In this section we relax this constraint and consider the three additional cases that occur as the starting price and minimum bid increment are increased beyond this limit. Specifically, we consider the cases where $s + d \leq \bar{v} < s + 2d$, $\bar{v} < s + d$, and $\bar{v} < s$, and we note that in each case, the expected revenue of the auction decreases.

—*Case 1.* $s + d \leq \bar{v} < s + 2d$. This case is similar to the one presented in earlier in the article; however, it is never possible for the auction to close at a price equal to the bid of the first bidder plus the minimum bid increment. The restricted number of auction outcomes, given all possible combinations of s, d, v_1 , and v_2 , are now given by

$\underline{v} \leq v_1 < s$:

$$E = \begin{cases} 0, & \underline{v} \leq v_2 < s, \\ s, & s \leq v_2 < \bar{v}; \end{cases} \quad (17)$$

$s \leq v_1 < \bar{v} - d$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_2, & s + d \leq v_2 < v_1 + d, \\ v_1 + d, & v_1 + d \leq v_2 \leq \bar{v}; \end{cases} \quad (18)$$

$\bar{v} - d \leq v_1 < s + d$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_2, & s + d \leq v_2 < \bar{v}; \end{cases} \quad (19)$$

$s + d \leq v_1 < \bar{v}$:

$$E = \begin{cases} s, & \underline{v} \leq v_2 < s + d, \\ v_1, & s + d \leq v_2 < v_1, \\ v_2, & v_1 \leq v_2 < \bar{v}. \end{cases} \quad (20)$$

Integrating over the bidders' valuation distribution and again assuming a uniform valuation distribution with support $[0, 1]$ gives the expected revenue of the auction:

$$E = \frac{1-s}{2}[(1+3s^2) - d^2]. \quad (21)$$

As expected, at the boundary condition where $s + 2d = 1$ this expression yields the same expected revenue as that shown in Equation (7). However, the maximum revenue achievable in this case (over the region $s \geq 0, d \geq 0$, and $s + d \leq 1 < s + 2d$) occurs when $s = \frac{1}{11}(3 + \frac{4}{\sqrt{3}})$ and $d = \frac{1}{11}(4 - \frac{2}{\sqrt{3}})$. Note that the maximum revenue is less than that generated earlier in the article where $s = 0$ and $d = \frac{1}{4}$ were found to be optimum parameter values.

—*Case 2.* $\bar{v} < s + d$. In this case, once an initial bid has been received, it is not possible for any other bids to be accepted (no bidder can have a valuation greater than $s + d$ —the minimum acceptable bid). Thus in this case we have the following conditions:

$\underline{v} \leq v_1 < s$:

$$E = \begin{cases} 0, & \underline{v} \leq v_2 < s, \\ s, & s \leq v_2 < \bar{v}; \end{cases} \quad (22)$$

$s \leq v_1 < \bar{v}$:

$$E = s, \quad (23)$$

and the expected revenue of the auction (again assuming a uniform bidders' valuation distribution with support $[0, 1]$) is given by

$$E = s(1 - s^2). \quad (24)$$

Clearly, the auction closes at a price of s , unless both bidders have valuations less than s . The expected revenue is thus solely dependent on the starting price of the auction, and this actually corresponds to both bidders being given sequential take-it-or-leave-it offers at a price s . The maximum revenue in this case occurs when $s = \sqrt{\frac{1}{3}}$ and has value $\frac{2}{3}\sqrt{\frac{1}{3}}$. Note that this is less than the previous case, and the general case described in Section 4.1.

—*Case 3.* $\bar{v} < s$. In this case it is impossible for any bidder to submit an initial bid and thus $E = 0$.

ACKNOWLEDGMENTS

We would like to thank Axel Ockenfels for providing feedback and comment on an earlier version of this article.

REFERENCES

- ANTHONY, P. AND JENNINGS, N. R. 2003. Developing a bidding agent for multiple heterogeneous auctions. *ACM Trans. Internet Tech.* 3, 3, 185–217.
- BAJARI, P. AND HORTACSU, A. 2003. The winners curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econ.* 34, 2, 329–355.
- BAJARI, P. AND HORTACSU, A. 2004. Economic insights from internet auctions. *J. Econ. Lit.* 42, 457–486.
- BAPNA, R. 2003. When snipers become predators: Can mechanism design save online auctions? *Commun. ACM* 46, 12, 152–158.
- BARBARO, S. AND BRACHT, B. 2004. Shilling, squeezing, sniping: Explaining late bidding in online second-price auctions. Working article. University of Mainz, Mainz, Germany.
- DAVID, E., ROGERS, A., SCHIFF, J., KRAUS, S., AND JENNINGS, N. R. 2005. Optimal design of English auctions with discrete bid levels. In *Proceedings of the ACM Conference on Electronic Commerce*. 98–107.
- DAVID, E., ROGERS, A., JENNINGS, N. R., SCHIFF, J., KRAUS, S., AND ROTHKOPF, M. H. 2007. Optimal design of English auctions with discrete bid levels. *ACM Trans. Inte. Tech.* 7, 2, Article 12 (May 2007), 34 pages. DOI = 10.1145/1239971.1239976 <http://doi.acm.org/10.1145/1239971.1239976> Also available online at <http://eprints.ecs.soton.ac.uk/11548/>.
- DUMAS, M., ALDRED, L., GOVERNATORI, G., HOFSTEDE, A., AND RUSSEL, N. 2002. A probabilistic approach to automated bidding in alternative auctions. In *Proceedings of the 11th International World Wide Web Conference*. 99–108.
- E-BAY. 2006. Financial releases: Fourth quarter and full year 2006 financial results. Available online at <http://investor.ebay.com/>.
- ENGELBERG, J. AND WILLIAMS, J. 2005. A license to shill: Incremental bidding and eBay's proxy system. Working article. Kellogg School of Management, Northwestern University, Evanston, IL.
- FONTOURA, M., IONESCU, M., AND MINSKY, N. 2002. Law-governed peer-to-peer auctions. In *Proceedings of the 11th International World Wide Web Conference*. 109–117.
- GERDING, E., ROGERS, A., DASH, R. K., AND JENNINGS, N. R. 2007. Sellers competing for buyers in online markets: Reserve prices, shill bids, and auction fees. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*. 1287–1293.

- GUO, X. 2002. An optimal strategy for sellers in an online auction. *ACM Trans. Internet Tech.* 2, 1, 1–13.
- HAILE, P. A. AND TAMER, E. 2003. Inference with an incomplete model of English auctions. *J. Pol. Econ.* 11, 1, 1–51.
- HARSTAD, R. M. AND ROTHKOPF, M. H. 2000. An ‘alternating recognition’ model of English auctions. *Manag. Sci.* 46, 1–12.
- HOSSAIN, T. AND MORGAN, J. 2006. ...Plus shipping and handling: Revenue (non)equivalence in field experiments on eBay. *Advan. Econ. Analys. Policy* 6, 2, article 3.
- JIANG, A. X. AND LEYTON-BROWN, K. 2007. Bidding agents for online auctions with hidden bids. *Mach. Learn. J.* 67, 1/2(May), 117–143.
- JOHANSSON, K. 1998. The longest increasing subsequence in a random permutation and a unitary random matrix model. *Math. Res. Lett.* 5, 63–82.
- LUCKING-REILEY, D. H. 2000. Auctions on the internet: What’s being auctioned, and how? *J. Indust. Econ.* 48, 3, 227–252.
- MILGROM, P. R. AND WEBER, R. J. 1982. A theory of auctions and competitive bidding. *Econometrica* 50, 5, 1089–1122.
- OCKENFELS, A., REILEY, D., AND SADRIEH, A. 2007. Online auctions. In *Economics and Information Systems*, T. Hendershott, Ed. Elsevier, Amsterdam, The Netherlands, 571–628.
- OCKENFELS, A. AND ROTH, A. E. 2006. Late and multiple bidding in second-price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games Econ. Behav.* 55, 297–320.
- RILEY, J. G. AND SAMUELSON, W. F. 1981. Optimal auctions. *Amer. Econ. Rev.* 71, 381–392.
- ROTH, A. E. AND OCKENFELS, A. 2003. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econ. Rev.* 92, 4, 1093–1103.
- ROTHKOPF, M. H. AND HARSTAD, R. M. 1994. On the role of discrete bid levels in oral auctions. *Eur. J. Operat. Res.* 74, 572–581.
- SHAH, H. S., JOSHI, N. R., SUREKA, A., AND WURMAN, P. R. 2003. Mining ebay: Bidding strategies and skill detection. In *Proceedings of the Conference on Mining Web Data for Discovering Usage Patterns and Profiles (WEBKDD 2002)*. 17–34.
- SONG, U. 2004. Nonparametric estimation of an eBay auction model. Working article. Department of Economics, University of British Columbia, Vancouver, B.C. Canada.

Received April 2006; revised January 2007; accepted February 2007