Available online at www.sciencedirect.com
SCIENCE DIRECT
European Economic Review 50 (2006) 307–322
EUROPEAN ECONOMIC REVIEW
www.elsevier.com/locate/econbase

Bidding behavior in competing auctions: Evidence from eBay☆

Sajid Anwar[a,b], Robert McMillan[c], Mingli Zheng[d,*]

[a]*School of Business, James Cook University, P.O. Box 6811, Cairns, QLD 4870, Australia*
[b]*IGSM University of South Australia, Adelaide, SA 5000, Australia*
[c]*Department of Economics, University of Toronto, Toronto, Ontario M5S 3G7, Canada*
[d]*Faculty of Social Sciences and Humanities, University of Macau, Taipa, Macau, China*

Received 25 February 2003; accepted 4 October 2004
Available online 30 November 2004

Abstract

Much of the existing auction literature treats auctions as running independently of one another, with each bidder choosing to participate in only one auction. However, in many online auctions, a number of substitutable goods are auctioned concurrently and bidders can bid on several auctions at the same time. Recent theoretical research shows how bidders can gain from the existence of competing auctions, the current paper providing the first empirical evidence in support of competing auctions theory using online auctions data from eBay. Our results indicate that a significant proportion of bidders do bid across competing auctions and that bidders tend to submit bids on auctions with the lowest standing bid, as the theory predicts. The paper also shows that winning bidders who cross-bid pay lower prices on average than winning bidders who do not.

JEL classification: D44; D440

Keywords: Competing auction; Cross-bidding; Auction empirics

☆This paper is based on Chapter 1 of Mingli Zheng's 2002 University of Toronto PhD thesis.
*Corresponding author.
E-mail address: mlzheng@umac.mo (M. Zheng).

0014-2921/$ - see front matter
doi:10.1016/j.euroecorev.2004.10.007

1. Introduction

In standard auction theory, the typical assumption is that there is a single seller and several bidders, with the seller acting as a monopoly and capturing all the information rents. In practice, sellers often do not have monopoly power but rather compete against other sellers, giving buyers the opportunity to choose among many auctions. This is especially common in online auctions.

The best known online auction site, eBay, has evolved to act as a clearinghouse for the sale of a large number of homogeneous goods. At any given time, a number of virtually identical items are available for sale: on the afternoon of February 12, 2004, for example, the search item 'Windows XP Home Edition' turned up 196 active auctions for this software. Further, the cost associated with bidding is very low and bidders can easily monitor several online auctions concurrently. Thus it is possible for bidders to bid across several competing auctions at roughly the same time.

A handful of papers in the literature consider the case in which sellers compete against each other (see for instance McAfee, 1993; Peters and Severinov, 1997).[1] These papers assume that bidders can only choose to buy from one seller, and the only equilibrium involves bidders randomizing over available sellers. In this case, it is entirely possible that some auctions should have many bidders while other auctions have few or no bidders, and consequently that some profitable trades may not be realized. Several recent papers provide empirical analyses of the online auction market, surveyed by Bajari and Hortacsu (2004) in their comprehensive review of the empirical literature.[2] It is worth noting that all the papers in this rapidly growing empirical literature treat online auctions as independent of one another.

The empirical analysis in this paper is prompted by recent theoretical work by Peters and Severinov (2002) on bidding behavior in competing auctions, in which they consider market equilibrium involving competing auctions similar to those used in eBay. When there is no bidding cost and no fixed ending time for auctions, Peters and Severinov demonstrate that the strategy in which bidders always submit a bid on an auction with the lowest 'standing' bid and then bid with the minimum increment is a (weak) perfect Bayesian equilibrium. This implies that bidders should bid across competing auctions and if a bidder becomes the highest bidder, he/she should pause bidding until other bidders outbid him/her. It is interesting to note that bidding once and bidding one's true valuation is not an equilibrium in this environment; intuitively, if bidders bid their true valuations and bid only once, they may be trapped into paying a higher price than is necessary and therefore may not have an

[1] Recently, in a paper entitled "Competing Auctions," Ellison et al. (2004) consider competition among different online auction sites (markets), focusing on the coexistence of these markets rather than competition among auctions themselves.

[2] Topics addressed include late bidding (see Roth and Ockenfels, 2002; Ockenfels and Roth, 2004; Schindler, 2003), the winner's curse (Bajari and Hortacsu, 2003; Yin, 2003), reputation mechanisms (Ba and Pavlou, 2002; Cabral and Hortacsu, 2003; Resnick et al., 2003), and the choice of different trading mechanisms (see Lucking-Reiley, 1999; Ivanova-Stenzel and Salmon, 2003).

opportunity to switch to other less competitive auctions.[3] For competing auctions, the final price of one auction is affected by the existence of other auctions. Furthermore, prices will tend to be uniform across competing auctions.

In the current paper, we examine whether or not the behavior of bidders in competing auctions corresponds to the strategy prescribed by Peters and Severinov (2002) using data from competing auctions in eBay.[4] In particular, we examine the impact of auction substitutability on bidding behavior, assessing whether agents do indeed bid across competing auctions (''cross-bid''). We also examine whether cross-bidding results in bidders paying lower prices for the objects they bid for.

The results presented in this study are based on a unique dataset collected from competing auctions for CPUs that took place on eBay over several months. We consider auctions with almost the same ending time and involving the sale of identical or very similar items. To ensure homogeneity among these items, we restrict attention to groups of auctions that share the same seller, and which have the same description, the same starting price and delivery method, and ending at almost the same time. From the data, we generate three competing auction samples: competing auctions ending on the same day, competing auctions ending within an hour, and competing auctions with ending times within a minute of each other. This allows us to investigate the effect of increasing the degree of substitutability among auctions on the behavior of auction participants.

Our results provide first compelling empirical evidence in support of competing auctions theory. We find clear evidence that a high proportion of bidders do cross-bid. Further, bidders tend to bid on the auction with the lowest standing bid, in conformity with the strategy proposed by Peters and Severinov. In terms of substitutability, we find that as the difference in the ending time across auctions becomes smaller, so more bidders tend to bid across competing auctions. In order to assess the potential gains from following the Peters and Severinov bidding strategy, we compare the winning price paid by bidders who bid across competing auctions with those who do not. We find that on average, winning bidders who actively bid across competing auctions pay around 91 percent of the price paid by those winning bidders who do not cross-bid.[5]

[3]Consider a situation that involves two competing auctions and four bidders. Suppose that the sellers' valuations are $0, and bidders' valuations are $10, $9, $5, and $4, respectively. If all bidders choose only one auction and bid their true valuation then the two highest valuation bidders may end-up bidding at the same auction, the winner paying a price of $9. The bidder with valuation of $5 will win the other auction. If bidders bid across competing auctions and bid with minimum increment then the two highest valuation bidders would win these two auctions and pay lower prices.

[4]Other auction markets also sell many similar items but use mechanisms different from that in eBay. Sotheby's and Christie's sell many identical items, such as wine, in one auction. These identical items are sold sequentially, lot by lot. Telecommunication spectrum auctions make all lots available for bidding at the same time. The auctions take place over a number of rounds, until no further bidding takes place on all lots being offered. The bidding strategies appropriate for these kinds of auctions are completely different from those in eBay environment we study.

[5]At the time the data were collected, Intel Pentium III 933 MHz CPUs sold for around US$100. Thus cross-bidding could lead to a $9 discount on average, which is non-trivial.

The rest of the paper is organized as follows: Section 2 describes the institutional features of eBay, and the construction of the dataset we use is discussed in Section 3. Section 4 provides evidence on bidding behavior in competing eBay auctions. In Section 5, we compare the average price paid by bidders who cross-bid versus bidders who do not, and Section 6 concludes.

2. Institutional features of eBay auctions

As the most extensively studied online market, eBay provides a rich resource for the study of auctions in general and, because a large number of very similar (if not identical) items are auctioned at any given time, it also provides an excellent opportunity to study competing auctions in particular.

eBay is an e-commerce website which provides a central market for buyers and sellers to trade using auctions. Sellers pay a fee whereas bidders pay no fee, and sellers also choose an auction type for the sale of their goods. (In this paper, we only consider auctions that use eBay's standard online auction format.[6]) They set a starting bid, a minimum bid increment and specify the duration of the auction. Sellers provide a detailed description of the item, which usually includes the method of delivery and method of payment, and they also have the option to set a secret reserve price. If they do, during the auction process eBay indicates whether the reserve price has been met or not, and if the reserve price has not been met at the end of the auction, sellers have the option of not selling the item.

The mechanism used by eBay resembles a second-price auction. At any time, eBay shows the current standing bid of the auction, which is the current *second highest* bid; in the special cases where there is no bid or only one bid, the current standing bid is the starting bid. When a bidder submits a bid, he/she knows the current standing bid, the identity of the seller (along with the seller's "feedback value" discussed below), the starting and ending time of the auction and a description of the item. The bidder also knows how many bids have already been submitted, the identity of the bidders and the times each of the bids were made. However, the exact amount of each bid is not revealed until the end of the auction. The final price is the second highest bid plus a specified minimum increment.

eBay plays no role in the actual exchange of items at the end of the auction. The winning bidders and sellers contact each other directly to complete the transaction. Since bidders cannot inspect the goods, sellers have incentive to provide false information, and the winning bidders may not want to pay and complete the transaction. In order to facilitate the completion of transactions and to boost the confidence of potential market players, eBay uses a feedback system. After each transaction (whether successful or not), the seller and the winning bidder can send feedback about the other party to eBay, marked as positive, neutral and negative,

[6]eBay conducts a variety of other types of auction, including reserve price auctions, Buy It Now, Fixed Price, New eBay Store Buy It Now Listings, private auctions, Dutch auctions, and restricted-access auctions. See http://pages.ebay.com/help/sellerguide/selling-type.html.

with values of +1, 0, −1, respectively, along with brief comments. Each trader is allocated a feedback number, which is the sum of his or her feedback values. The feedback information is public and always associated with the trader, though this system is far from foolproof, as eBay cannot prevent traders, particularly those with a bad reputation, from changing identities. (See Bajari and Hortacsu (2004) for a survey of empirical analyses of the reputation mechanism in eBay.)

When an auction ends, eBay provides detailed information about the associated bid history.[7] An example of the history of one auction extracted from eBay's website is presented in Appendix A. The top half of the Appendix reports basic information about the auction. It was a three-day auction, which started on October 31, 2001 at 22:51:27 PST and ended on November 3, 2001 at 22:51:27 PST. The seller's feedback value was 11. The starting bid set by the seller was $10 and the minimum bid increment, $0.50. The auction received 10 bids from 4 different bidders. The cost of shipping was $5 and the optional shipping insurance cost, $5.

The lower half of the figure shows the detailed bidding history, sorted by the amount of each bid. The auction received its first bid of $17.50 by "planetorb" around 23 hours after the start of the auction. Then 4 hours before the end of the auction, bidder "raheem112" started to bid. This bidder would have realized that he/she was not the only bidder though he/she would not have known the exact amount of the other bid. Since there was only one bidder at that time, the standing bid was still $10. Bidder "raheem112" first placed a bid of $11 and found that the standing bid increased to $11, indicating that he/she was not the highest bidder. "raheem112" then increased his/her bid subsequently in 2 minutes to $13, $14, then $15, until at last he/she became the highest bidder with a bid of $20. The standing bid then became $17.50. In the last hour of the auction, the bidder "iteachcomputers" submitted a bid of $20, increasing the standing bid to $20. Bidder "raheem112" increased his/her bid to $21 followed by another bidder who bid $23.99. Bidder "planetorb" finally won the auction with an unknown bid. Since the minimum increment is $0.50, the final price was $24.49 (the second highest bid $23.99 plus the bid increment $0.50). There was no bid retraction or cancellation for this auction. Similar information on completed eBay auctions is available to the public for one month following the end of the auction.

3. Data

Our sample consists of CPU auctions that took place on eBay during the period January 26–May 18, 2002. As a first step towards data collection, we entered "CPU" as the key word on eBay's search page for completed auctions.[8] The search results showed CPU-related auctions completed for up to 15 days prior to the day the search was conducted. A Java program was then used to download links to the

[7] Before the middle of 2000, eBay provided only the highest bid of a bidder and his/her last bidding time. After that date, more detailed bidding histories are provided, as used in this study.

[8] The option to search for completed auctions is no longer available on eBay, as of February 2004.

history of each item, in turn used to download the bid history and other relevant information. This generated data on all kinds of CPUs and CPU-related items, from more than 71,000 auctions in total. The sample was further refined by selecting only those auctions that involved the sale of a single item.[9]

Our full sample consists of 41,849 auctions. Of these, 7333 auctions did not attract any bids and 2624 auctions had a secret reserve price, the secret price being met in only 1533 auctions. We also observed some auctions with non-serious sellers.[10] In the full sample, the average final price is \$60.91, with standard deviation \$90.93. The wide variation in price is largely due to the existence of many different models and items in different conditions. However, even for identical items sold at almost the same time (such as items in the same competing auctions group), we still observe significant price variation.[11] The mean of the starting bids was \$25.17 with standard deviation \$69.85. On average, each auction attracted 7.13 bids (with standard deviation 7.14). Variation in the number of bids can be explained in part by variation in the starting bid—the higher the starting bid, the lower the number of bidders and subsequent bids.

Competing auctions in our sample satisfy the following conditions: first, the items auctioned are reasonably homogeneous in quality (including warranty provisions) and have similar delivery methods and shipping costs; second, the auctions end at approximately the same time. Although the conditions of properly working CPUs and related products in a specific category (such as Pentium III 800 retail boxes) are largely the same, they may still differ in other respects. For example, some may be new and the box may not have been opened, while others may have been used for several months; some CPUs may have been under warranty, and the method of delivery, the shipping cost and the method of payment may also differ.

To overcome complications associated with this heterogeneity, we focus on a sub-sample consisting of groups of auctions with very similar features. In particular, we choose auctions run by the same sellers, selling items with the same product descriptions, the same delivery method and the same shipping cost, using the fact that some sellers on eBay sell a number of almost identical items. Except for different ending times, these auctions are completely indistinguishable to bidders. Because we focus on auctions with the same seller, we are likely to understate the true extent of cross-bidding. Even so, we will find clear evidence of cross-bidding using this more conservative sample in Section 4 below.[12]

The starting and ending times of auctions included in our sample are not the same. Auctions that end at approximately the same time compete more directly against each other than those with sizeable differences in ending times. From the CPU

[9]Items sold using Buy It Now are also excluded from the sample.

[10]For example, one seller set a starting bid of US\$999 for a single CPU. Unsurprisingly, no one bid on this item.

[11]This fact is documented and analyzed in Zheng (2002). It may be explained by random matching of bidders and sellers. Cross-bidding can reduce price variation to some extent.

[12]Drawing the sample in this way means that we cannot examine competition among competing sellers. That remains a topic for future research.

Table 1
Descriptive statistics on the three competing auction sub-samples

	Daily sample	Hourly sample	Minute sample
Number of groups of competing auctions	1943	1021	328
Average number of auctions in each group of competing auctions[a]	2.77	2.63	2.55
Average number of bidders in each group	10.74	9.30	7.72
Average number of bids submitted to each group	22.18	20.58	16.13

[a]As defined in the text, a 'group of competing auctions' consists of auctions involving the sale the same CPU, sold by the same seller, and all ending within a given interval of each other (one day, one hour or one minute).

auctions sample, we obtain three sub-samples of competing auctions.[13] Each observation in these samples is a group of auctions that consist of 2 or more competing auctions. The first sample (the daily sample) includes groups of competing auctions ending within the same day. The second sample (the hourly sample) includes groups of competing auctions ending within one hour of each other. The third sample (the minute sample) includes groups of competing auctions ending within one minute of each other. Auctions appearing in the minute sample also appear in the hourly sample and those appearing in the hourly sample also appear in the daily sample.

Table 1 provides a description of the three samples. We consider only those groups of competing auctions in which at least one auction receives a positive number of bids. The minute sample has the smallest size, consisting 328 groups of competing auctions. The hourly and daily samples consist of 1021 and 1943 groups of competing auctions, respectively. However, the average number of auctions per group for the three samples is not very different. A group in the minute sample consists of 2.55 auctions while a group in the daily sample consists of 2.77 auctions. While this may seem surprising, a careful look at the data shows that there are many groups with only two auctions, and sellers set the ending time many minutes or even hours apart. These auction groups appear in the daily (or hourly) sample, but not in the minute sample, which causes the number of auctions in a group in the daily sample to be similar in size to the number in the minute sample. Table 1 also shows the average number of bidders and the number of bids submitted in the different groups.

It should be noted that bidders may need to buy more than one unit for personal use or alternatively, they may be professional dealers. The optimal strategies are likely to be different depending on whether bidders have single or multiple demands,

[13]Since a bidder responding in the last minute of the auction is more likely to be using special software or to be a professional bidder, a one minute cutoff might isolate professional bidders. We have also tried samples including groups of competing auctions ending within 5 minutes, half an hour, and 12 hours. The results are similar to the results obtained from minute, hourly and daily samples reported here.

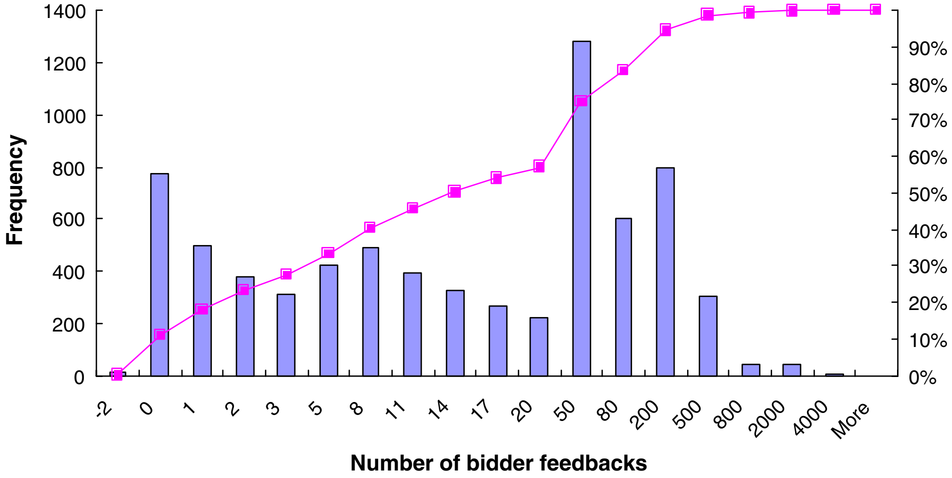

Fig. 1. Histogram of bidders' feedback values.

with bidders requiring multiple units bidding across more than one auction at the same time. The existence of bidders demanding multiple units in our sample would bias our claim about cross-bidding. Because we know bidder identities, we are able to address this problem, distinguishing genuine cross-bidders from bidders demanding multiple units in the next section.

One might be concerned that bidders, especially novices, do not fully understand the auction mechanism used on eBay and may therefore experiment with different strategies. In this paper, we use the feedback values as a rough indicator of traders' experience on eBay.[14] eBay also uses the feedback values in daily trading. For example, to use the ''Buy It Now'' feature, sellers must have a feedback value greater than 10, otherwise they have to go through a process of identity verification in eBay, at a cost of $5.[15] Fig. 1 shows the distribution of the bidders' feedback values. The daily sample includes 7149 different bidders: among them, only 8 bidders had a negative feedback value,[16] accounting for 0.1% of all bidders, and bidders with feedback value 0 account for 10% of all bidders. The feedback value of the remaining 90 percent of bidders was positive, while the feedback value of 60 percent of bidders was greater than 8, suggesting that they tend to have a history of successful transactions on eBay.

[14]Before March 2003, eBay displayed a feedback summary value given by the number of positive feedbacks received from other unique users, minus the number of negatives. We use this value as an indicator of experience because we did not collect data on neutral or negative feedback. In fact, Resnick et al. (2001) report that only 0.6% of the feedbacks left by buyers about sellers are negative or neutral. In March 2003, eBay began displaying the percentage of positive feedbacks along with the total score.

[15]See http://pages.ebay.com/help/sell/bin.html.

[16]If a bidder has a different feedback number during the period, we use the highest feedback number.

Table 2
Evidence on cross-bidding behavior
Panel A: Descriptive statistics

	Daily sample	Hourly sample	Minute sample
Average number of cross bidders including multi-unit bidders	2.32 (3.42)	2.85 (3.33)	2.46 (2.89)
Proportion of cross bidders including multi-unit bidders	0.19 (0.23)	0.31 (0.28)	0.32 (0.29)
Average number of cross bidders excluding multi-unit bidders	1.97 (2.93)	2.30 (2.88)	2.02 (2.66)
Proportion of cross bidders excluding multi-unit bidders	0.14 (0.17)	0.20 (0.20)	0.20 (0.21)

Panel B: Hypothesis tests for equality in the proportion of cross bidders comparing different samples

	Including multi-unit bidders	Excluding multi-unit bidders
Equality between daily sample and hourly sample	$t = -11.76$ $p = 0.000$	$t = -3.56$ $p = 0.000$
Equality between the hourly sample and the minute sample	$t = -0.55$ $p = 0.58$	$t = 0$ $p = 1$

Note: Panel A includes groups with positive bids only. The numbers in parentheses are standard deviations. In Panel B, the *t*-values correspond to a test of the null hypothesis that the proportion of bidders who cross bid is equal between the daily and the hourly sample, and between the hourly and the minute sample, respectively. *p*-values are based on a two-tailed test.

4. Bidding behavior in competing auctions

Using these data on detailed individual bidding history from eBay, in this section, we study detailed bidding behavior by examining the behavior of all bidders within each group of competing auctions.[17] In particular, we test whether bidders bid across competing auctions, and whether or not a bid is submitted on an auction with the lowest standing bid within the relevant auction group.

Table 2 provides evidence on the first of these issues—whether or not bidders bid across competing auctions. Panel A reports the average number of bidders bidding across competing auctions and the proportion of such bidders in the relevant competing auctions group. It does so on two separate bases: including and then excluding multiple unit demand bidders. If we include multiple unit demand bidders in the calculation (the more liberal measure), the mean proportion of bidders who bid across auctions is positive in each case—19 percent in the daily sample, 31 percent in the hourly sample and 32 percent in the minute sample. Further, the data

[17]Determining whether a bidder cross-bids and whether a bid is submitted to the auction with the lowest standing bid is not a trivial task. It involves the comparison of auction prices at any point of time and the tracking of the whole bidding histories for all auctions in a competing auction group. We used Java's TreeSet data structure for this purpose. See Bruce (2002, Chapter 9) for details.

clearly reject the null hypothesis that the proportion of bidders who cross-bid is zero, the t-statistics for the daily sample, hourly sample and minute sample being 36.4, 35.4 and 20.0, respectively, with the corresponding one-tailed p-values all less than 0.001.

It is clear that the proportion of bidders who cross-bid is very similar for the minute sample and hourly sample, and that both proportions are significantly larger than the proportion in the daily sample. In panel B of Table 2, we carry out a formal test to compare the proportion of bidders who cross-bid in the different samples. The null that the proportion of cross-bidders in the daily and hourly samples are equal is easily rejected in favor of the alternative hypothesis that the proportion in the hourly sample is greater. This finding is consistent with the notion that cross-bidding is more likely on auctions which are closer substitutes—in this case, closer substitutes across time.

One might wonder whether the cross-bidding observed here is due solely to the existence of bidders who want to acquire multiple units of the good. Bidder intentions are not observable, but it can be argued that the bidders who are attempting to acquire multiple units are likely to bid aggressively on several auctions thereby increasing their chances of winning these auctions. In contrast, a bidder seeking a single unit is likely to bid aggressively on only one auction at a time, thereby avoiding a situation in which he/she wins more than one auction. Accordingly, such bidders will hesitate to become the highest bidders in more than one auction, especially towards the end of the auction. Based on these considerations, we define 'true' cross-bidders to be those who (a) are never the winner of more than one auction and (b) are never the highest bidder in more than one auction on the last day of the auction.

Criterion (b) may cause a bias in specific cases. For example, if an auction ends before the last day of the auction, and a bidder bids aggressively to attempt to acquire multiple units, criterion (b) cannot be used to identify such a multiple unit bidder, and the bidder will be classified (incorrectly) as a 'true' cross-bidder. Fortunately, this case rarely occurs in eBay because of the prevalence of late bidding. In Fig. 2, we provide a histogram of the bid submission times for all bids in the hourly sample, for illustration. More than 40 percent of the bids are submitted during the final 10 percent of the remaining auction time, consistent with the late bidding noted by Roth and Ockenfels (2003).

Table 2 also reports cross-bidding results on a second basis, by excluding multiple unit demand bidders. Multiple unit demand bidders account for only a relatively small proportion of all cross-bidders—on average around 10–20 percent of all samples. After excluding bidders who are attempting to acquire multiple units, we still observe a high percentage of bidders who cross-bid. The mean proportion of bidders who cross-bid is 14 percent in the daily sample, 20 percent in the hourly sample and 20 percent in the minute sample. For the null hypothesis that the proportion of bidders who cross bid is 0, the t-statistics for the daily sample, the hourly sample and the minute sample are 36.3, 32.0 and 17.2, respectively, with the corresponding one-tailed p-values less than 0.001 in each case. For the comparison of the proportion of bidders who bid across competing auctions among three samples,

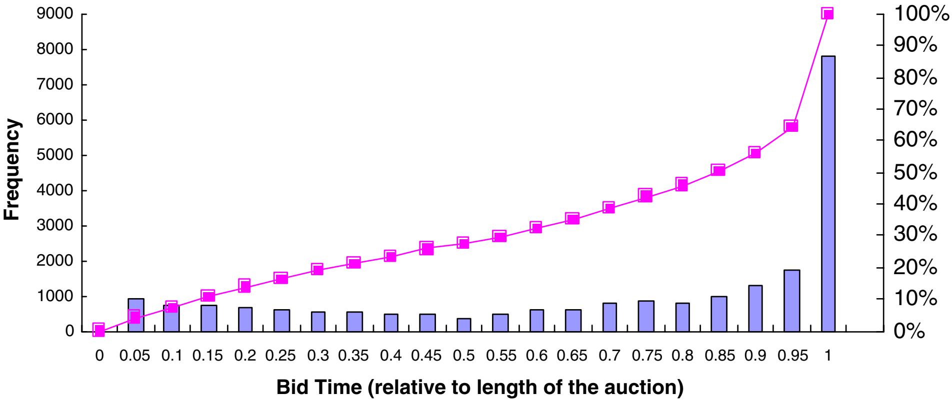

Fig. 2. Distribution of bid submission times.

the results are similar to the ones including the multiple unit demand bidders (see Panel B of Table 2): the proportion of bidders who bid across competing auctions is not significantly different between the hourly sample and the minute sample, while the proportion of cross-bidders is significantly smaller in the daily sample. Again, this is consistent with the notion that cross-bidding is more likely on auctions that are closer substitutes across time.

Another important implication of competing auctions theory is that bidders tend to bid on auctions where the standing bid is the lowest. This strategy ensures that a bidder who is attempting to buy only a single unit of a commodity never wins more than one auction. In addition, bidders avoid being trapped into paying excessively high prices.

In order to determine the proportion of bids submitted on auctions with the lowest standing bid, simply dividing the number of such bids by the total number of bids submitted is not appropriate. To see this, if all auctions in a group have ended save one, then there will be no more competition and bids submitted to this remaining auction might be considered as bidding on the auction with the lowest standing bid. However, this method will tend to over-estimate the proportion of bids submitted to the auction with the lowest standing bid, especially for groups of auctions where the ending times of auctions are further apart, as with the daily sample. In order to calculate the relevant proportion, we first isolate those bids submitted while at least one other competing auction was in progress. (The auction groups where the number of such bids is zero have been excluded from following calculations.) Among these bids, we calculate the proportion of bids that were submitted on auctions where the standing bid was the lowest. The results are shown in Table 3.

This table reports summary statistics on the number of bids submitted while at least one other competing auction was in progress, and the number of bids submitted on the auction with the lowest standing bid. The mean proportion of those bids

Table 3
Proportion of bids submitted to auctions with the lowest standing bid

	Minute sample	Hourly sample	Daily sample
Number of bids (while at least one other competing auction was in progress) received by each group	15.88 (17.75)	19.17 (22.15)	19.30 (24.41)
Number of bids submitted to the auction with lowest standing bid	9.45 (8.89)	10.65 (9.31)	9.16 (8.93)
Proportion of bids submitted to the auction with the lowest standing bid	0.76 (0.22)	0.72 (0.24)	0.62 (0.30)

Notes: Standard deviations are in parentheses. Auction groups on which no bids were submitted (while at least one competing auction was in progress) were excluded for the above calculations. The value of *t*-statistic concerning the null hypothesis that the proportion of bids submitted to auctions with the lowest standing bid is equal across the minute and the hourly sample is 1.33. The value of the *t*-statistic relating to the null hypothesis that the proportion of bids submitted to auctions with the lowest standing bid is equal across the hourly and the daily sample is 5.01.

submitted on the auction with the lowest standing bid is high—62 percent for the daily sample, 72 percent for the hourly sample, and 76 percent for the minute sample. We also observe significant variation in these proportions.

It is worth noting that the mean proportion, reported in the third row of the table, is much higher than the ratio of the mean number of bids submitted on auctions with the lowest standing bid (reported in the second row) to the mean of the total bids (in the first row). This difference can be accounted for by the existence of auction groups with a low total number of bids but a very high proportion of bids submitted on auctions with the lowest standing bid.

The high proportions in Table 3 need to be placed in perspective. If bidders were to choose an auction in a group randomly each time they placed a bid, our method of calculation would tend to generate proportions in the range 30–50 percent.[18] The results in Table 3 are much higher. For the null that the proportion is 50 percent, the *t*-statistics for the daily sample, hourly sample and minute sample are 17.6, 28.8 and 21.4, respectively, rejecting the hypothesis of random bidding in favor of bidding on the auction with the lowest standing bid.

We might expect an increased tendency to bid on a competing auction with the lowest standing bid when the auctions in an auction group compete more closely. We do observe a slight monotonic tendency, especially comparing the hourly data and daily data. However, this may simply be due to the fact that groups of auctions that compete more closely are smaller in size, so even when bidders submit bids randomly, a monotonic tendency may still be apparent. Further research into this lack of monotonicity may be fruitful.

[18]This is because there are at least two auctions in our calculation, and we know that the average number of auctions in a group is less than 3 in all three samples. Thanks to a referee for making this point.

5. Do cross-bidders pay lower prices?

Bidding across competing auctions affords bidders greater choice. More importantly, cross-bidding allows bidders to switch to auctions where the standing bid is low, and as the theory of competing auctions predicts, cross-bidding should result in bidders paying lower prices in competing auctions.

In order to test the above prediction, we divide each group of competing auctions into two subgroups: a group consisting of auctions in which the winner of the auction is a true cross-bidder, as defined in Section 4, and a group consisting of the remainder of the auctions. We calculate the mean price paid by the winners in each subgroup and the ratio of these two average prices. In some groups, all winners were 'true' cross bidder while in others, none of the winners were—both these types of groups are excluded from the analysis. We therefore consider only those groups in which there are a mixture of winners—those who are true cross-bidders and winners who are not, allowing us to focus on those auctions that involve both types of winner and so make a meaningful comparison of prices paid. The results are summarized in Table 4.

Within the minute sample, only 77 groups of competing auctions include winners who bid across competing auctions and winners who do not. The ratio of the average price paid by cross-bidders to the average price paid by non-cross bidders is 0.91, with a standard deviation of 0.17. This indicates that there is a 9 percent discount for those winning bidders who cross-bid. For an average price of around 60 dollars in the sample, the saving is not trivial.

The table also reports a test of the null hypothesis that the ratio of the two population means is equal to 1 against the alternative that the ratio is less than 1. Winners who bid across competing auctions on average pay a lower price as compared to those who do not, the difference being statistically significant—the relevant t-statistic has a value of -4.64.

The hourly sample is larger, including 244 groups of competing auctions with winners who bid across competing auctions and winners who do not. The ratio of the average price paid by cross bidders and non-cross bidders is 0.93, with a standard

Table 4
Hypothesis tests comparing the ratio of the price paid by cross-bidding winners and non-cross-bidding winners

	Mean	Standard deviation	t-statistic	p-value
Daily sample ($N = 394$)	0.94	0.18	-6.62	<0.000
Hourly sample ($N = 244$)	0.93	0.27	-4.05	<0.000
Minute sample ($N = 77$)	0.91	0.17	-4.64	<0.000

Note: 'Mean' and 'Standard deviation' denote the mean and standard deviation of the ratio of the price paid by cross-bidding winners and non-cross-bidding winners. We also report the corresponding one-tailed p-value. H_o: the ratio of price paid by cross-bidding winner and by non-cross-bidding winner is equal to 1, versus H_a: the ratio is less than 1.

deviation of 0.27. Within the daily sample, there are 394 groups of competing auctions with winners who bid across competing auctions and winners who do not. The ratio of the average price paid by cross bidders and by non-cross bidders is 0.94 with a standard deviation of 0.18. Hypothesis tests show the ratio is significantly different from 1 in both samples. For the null hypothesis that the ratio of the two population means is equal to 1 against the alternative that the ratio is less than 1, the t-statistic is -4.05 for hourly sample and -6.62 for daily sample.

This evidence indicates that bidders who cross-bid on competing auctions are strictly better off. It is also interesting to note that the reward from cross-bidding appears to increase as the ending times of the auctions in a competing group become closer, though a formal hypothesis test shows that the differences are not statistically significant. This increasing gain can be attributed to the fact that it is easier to coordinate bidding on competing auctions if these auctions end at approximately the same time.

6. Conclusion

Most existing studies assume that all auctions, including on-line auctions, are independent. This paper provides the first related pieces of evidence that lend support to competing auctions theory. Our analysis of eBay data shows that bidders' behavior is very different in the presence of competing auctions. Specifically, bidders tend to bid across competing auctions and bid on the auction with the lowest standing bid. Our investigation shows that a greater proportion of bidders cross-bid if competing auctions end at approximately the same time. We also find that winners who bid across competing auctions tend to pay lower price than those who do not.

The theory of competing auctions proposed by Peters and Severinov (2002) suggests that, as a result of cross bidding, the final price paid by winning bidders is likely to be uniform. This remains the subject of future research. In addition, the current paper takes the behavior of sellers as exogenous, focusing on bidding behavior in competing auctions. We plan to investigate the behavior of sellers in a separate paper.

Acknowledgements

We thank Paul Anglin, Simon Board, Michael Denny, Alex Maynard, Joanne Roberts and especially Michael Peters for helpful comments. We are also very grateful to Fang Xiao for her help with the Java programming. This paper has greatly benefited from constructive comments from the editor and two anonymous referees. All remaining errors are the authors' sole responsibility. Financial support from a University of Macau research Grant (No. RG305/02-03) is gratefully acknowledged.

Appendix A. A bidding History Page from eBay

EBay Bid History for
Intel Pentium II Xeon 450MHz 512k - Used (Item # 1292121814)

Currently	$24.49	First bid	$10.00
Quantity	1	# of bids	10
Time left	Auction has ended.		
Started	Oct-31-01 22:51:27 PST		
Ends	Nov-03-01 22:51:27 PST		
Seller (Rating)	nhahmad (11)*me		

View page with email addresses (Accessible by Seller only) Learn more.

Bidding History (Highest bids first)

User ID	Bid Amount	Date of Bid
planetorb (9)	$24.49	Nov-03-01 22:30:37 PST
ibgeek (4)	$23.99	Nov-03-01 22:18:01 PST
raheem112 (0)	$21.00	Nov-03-01 22:03:35 PST
iteachcomputers (1)	$20.00	Nov-03-01 22:06:49 PST
raheem112 (0)	$20.00	Nov-03-01 18:33:27 PST
planetorb (9)	$17.50	Nov-01-01 21:11:05 PST
raheem112 (0)	$15.00	Nov-03-01 18:33:16 PST
raheem112 (0)	$14.00	Nov-03-01 18:32:50 PST
raheem112 (0)	$13.00	Nov-03-01 18:32:39 PST
raheem112 (0)	$11.00	Nov-03-01 18:31:52 PST

Remember that earlier bids of the same amount take precedence.
Bid Retraction and Cancellation History
There are no bid retractions or cancellations.

References

Ba, S., Pavlou, P., 2002. Evidence of the effect of trust building technology in electronic markets: Price premium and buyer behavior. MIS Quarterly 26 (3), 243–268.

Bajari, P., Hortacsu, A., 2003. Winner's curse, reserve prices and endogenous entry: Empirical insights from eBay auctions. Rand Journal of Economics 3 (2), 329–335.

Bajari, P., Hortacsu, A., 2004. Economic insights from internet auctions. Journal of Economic Literature XLII, 457–486.

Bruce, E., 2002. Thinking in Java, third ed. Prentice-Hall, Englewood Cliffs, NJ.

Cabral, L., Hortacsu, A., 2003. Dynamics of seller reputation: Theory and evidence from eBay. Mimeo., University of Chicago.

Ellison, G., Fudenberg, D., Mobius, M., 2004. Competing auctions. Journal of the European Economic Association 2 (1), 30–66.

Ivanova-Stenzel, R., Salmon, T., 2003. Bidder preferences among auction institutions. Economic Inquiry.

Lucking-Reiley, D., 1999. Using field experiments to test equivalence between auction formats: Magic on the internet. American Economic Review 89 (5), 1063–1080.

McAfee, P., 1993. Mechanism design by competing sellers. Econometrica 61 (6), 1281–1312.

Ockenfels, A., Roth, A.E., 2004. Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. Games and Economic Behavior, forthcoming.

Peters, M., Severinov, S., 1997. Competition among sellers who offer auctions instead of prices. Journal of Economic Theory 75, 141–197.

Peters, M., Severinov, S., 2002. Internet auctions with many traders. Working paper, University of Toronto.

Resnick, P., Zeckhauser, R.Z., Swanson, J., Lockwood, K., 2003. The value of reputation on eBay: A controlled experiment. Working Paper, Harvard Kennedy School.

Roth, A., Ockenfels, A., 2002. Last minute bidding and the rules for ending second price auctions: Theory and evidence from a natural experiment in eBay. American Economic Review 92 (4), 1093–1103.

Schindler, J., 2003. Late bidding on the internet. Mimeo., Vienna University of Economics and Business Administration.

Yin, P.-L., 2003. Information dispersion and auction prices. Working paper, Harvard Business School.

Zheng, M., 2002. Price formation in eBay. Mimeo., University of Macau.