Behavior in Second-Price Auctions by Highly Experienced eBay Buyers and Sellers*

Rod Garratt† Mark Walker‡ John Wooders§

February 27, 2008

Abstract

When second-price auctions have been conducted in the laboratory, most of the observed bids have been "overbids" (bids that exceed the bidder's value) and there are very few underbids. Few if any of the subjects in those experiments had any prior experience bidding in auctions. We report on sealed-bid second-price auctions that we conducted on the Internet using subjects with substantial prior experience: they were highly experienced participants in eBay auctions. Unlike the novice bidders in previous (laboratory) experiments, the experienced bidders exhibited no greater tendency to overbid than to underbid. However, even subjects with substantial prior experience tend not to bid their values, suggesting that the non-optimal bidding of novice subjects is robust to substantial experience in non-experimental auctions. A key determinant of bidding behavior was whether a subject had ever been a seller on eBay.

*We thank Greg Crawford, Ron Harstad, John Kagel, and Dan Levin for helpful comments. We are grateful to John Kagel for providing us with data from experiments reported in Kagel and Levin (1993). This research was partially supported by Grant No. SBR-0099353 from the National Science Foundation.

†Department of Economics, University of California, Santa Barbara, CA 93106 (garratt@econ.ucsb.edu).

‡Department of Economics, Eller College of Management, University of Arizona, Tucson, AZ 85721 (mwalker@arizona.edu).

§Department of Economics, Eller College of Management, University of Arizona, Tucson, AZ 85721 (jwooders@eller.arizona.edu).

1 Introduction

In a second-price auction, bidding one's value is always a dominant strategy. However, when second-price auctions (SPAs, for short) have been conducted in the laboratory, roughly two-thirds of the subjects overbid (i.e., they submit bids that exceed their values). There are very few underbids.[1] Why do so many subjects fail to choose the dominant-strategy value bid? And given that they don't bid their values, why is overbidding so much more prevalent than underbidding?

Kagel, Harstad, and Levin (1987) suggest that overbidding in SPAs is due to subjects' "illusion that [bidding in excess of value] improves the probability of winning with no real cost to the bidder, as the second-high-bid price is paid." (p. 1299) Moreover, they argue that the reason this behavior does not go away with repeated play is that "punishment probabilities are weak, given that bidders start with the illusion that bids in excess of [value] increase the probability of winning without materially affecting the price paid, and the majority of the time the auction supports this supposition" (p. 1299). In short, subjects who overbid in SPAs are rarely confronted with the consequences of their "mistake." Hence, the learning that often occurs when laboratory subjects participate repeatedly in the same institution does not eliminate overbidding in second-price auctions.

The subjects in the laboratory experiments were students, who typically had little if any prior experience bidding in auctions. Moreover, the laboratory SPA experiments were of limited duration (typically about two hours), which might not provide sufficient opportunity for a subject to learn that value-bidding is a good strategy. In contrast, real-world experience is typically obtained on separate occasions over extended periods of time. Thus people have time to reflect on how outcomes are affected by their decisions. In addition, people who participate in large numbers of real-world auctions have strong incentives to learn since the stakes in cumulative terms are large. Hence people with some experience in real-world auctions might be

[1] Kagel and Levin (1993) (henceforth K&L) conducted one of the first laboratory experiments with SPAs. The subjects in their experiment were assigned independent private values for the item being auctioned. In SPAs with five bidders, 67% of the bids exceeded the bidder's value and fewer than 6% of the bids were less than the bidder's value. K&L obtained similar results in SPAs with ten bidders. Additionally, Kagel, Harstad, and Levin (1987) and Harstad (2000) report evidence of overbidding in SPAs with affiliated private values. Coppinger, Smith, and Titus (1980) and Cox, Roberson, and Smith (1982) report underbidding, but in these experiments subjects were not permitted to bid above their private values.

expected to bid in a way that conforms more closely to the theory.

We report on an experiment in which the subjects had a great deal of real-world experience: each subject had participated in at least fifty eBay auctions. The subjects were recruited directly from eBay, and the experiment was conducted on the Internet instead of in the laboratory. Our objective was to determine the impact of extensive experience in real-world auctions on behavior in SPAs.

eBay experience is relevant because eBay auctions share important features of an SPA. Auctions of a single item on eBay are conducted as ascending-price auctions in which bidders submit "proxy" bids, which represent their maximum willingness to pay.[2] Just as in an SPA, at the close of an eBay auction the bidder with the highest proxy bid wins the auction and pays the second-highest proxy bid. The proxy bidding system is designed to allow a bidder to ignore the dynamic aspects of an auction, and instead simply submit a bid equal to his value. Indeed, eBay advises bidders to value bid, telling them "Decide the maximum you're willing to pay and enter this amount." Elsewhere on its website, eBay advises bidders to think of their proxy bid as the amount they would tell a friend to bid for them if they were unable to attend the auction in person.[3,4] In effect, eBay advises bidders to follow their weakly dominant strategy in an SPA of bidding their value.

The behavior of our eBay subjects was similar in one respect to behavior observed in previous SPA experiments: just as in the laboratory experiments with inexperienced subjects, the subjects in our online experiment did not generally bid their values. In addition, the more highly-experienced of our subjects did not tend to bid closer to their values than the less experienced ones. In short, the failure of student subjects to bid their values in laboratory SPAs is *robust*: bidders with substantial real-world experience in field auctions also fail to value bid.

Our eBay bidders did, however, bid differently than student subjects. In contrast to the tendency to overbid but rarely underbid that the inexperienced bidders

[2] In particular, throughout an auction bidders can observe who is the current high bidder and the amount of the current high bid (given the proxy bids made so far), and bidders may increase their proxy bids.

[3] Roth and Ockenfels (2002) point out that eBay bidders commonly bid in the last minutes or seconds of the auction, a practice known as late bidding or sniping. A bidder who snipes cannot revise his bid and hence for such a bidder an eBay auction is a SPA.

[4] After our experiment was conducted, eBay introduced "Second Chance Offers," which allows a seller, after the close of his auction, to a make take-it-or-leave-it price offer to a non-winning bidder equal to the bidder's final bid. Clearly value bidding is no longer a dominant strategy for a bidder who anticipates the possibility of receiving such an offer.

in laboratory experiments displayed, the experienced bidders in our online auctions exhibited no greater tendency to overbid than to underbid. The number of subjects who underbid (41% of subjects) was almost exactly the same as the number who overbid (38%). We also find that it was subjects who had prior experience *selling* in eBay auctions who tended on average to underbid. Subjects who had never sold anything in an eBay auction (i.e., they had only been bidders) tended on average to overbid. These findings suggest that prior bidding experience does affect behavior, and that the *kind* of experience one has makes a difference.

An explanation for our finding that subjects with experience selling on eBay bid differently than subjects without such experience is found in the psychology literature on deductive reasoning processes.[5] This literature argues that humans use "case-based" reasoning when undertaking new tasks (see Johnson-Laird, 1999). They do not apply logical deductive reasoning to make inferences on how to behave, but rather they apply behavior from previous similar situations. Our subjects with experience selling were typically highly experienced as sellers, with a median of 57 feedbacks *as sellers*. These subjects are likely traders, buying items and then reselling them at higher price on eBay, while aiming to achieve a positive markup covering at least the opportunity cost of their time. It appears that these subjects brought that learned behavior with them to our experimental auction, continuing to bid less for an item than its value. Subjects with no experience selling, on the other hand, did not exhibit this bias, and tended to behave much like the novice subjects in prior experiments.

Other experiments have found that real-world experience can sometimes influence experimental subjects' behavior. This typically requires that the experimental context is very similar to a subject's prior experience. Dyer, Kagel, and Levin (1989) and Cooper, Kagel, Lo, and Gu (1999) contain discussions of this phenomenon. When experience does have an effect, subjects with relevant experience typically behave more in accordance with the predictions of economic and game theory than inexperienced subjects do; examples can be found in List (2003), Hannan, Kagel, and Moser (2002), and Palacios Huerta and Volij (2006). Our results are somewhat unusual in this regard. Only one kind of experience – selling – has an effect on subjects' behavior in SPAs, but instead of producing behavior that is more consistent with theoretical predictions, experience reverses the direction in which behavior is suboptimal.

[5] We thank John Kagel for this observation. Cooper and Kagel (2006) advance this argument in explaining the differing effects of meaningful versus generic context on cross-game learning in signaling games.

2 Experimental Procedures

Our goal was to recruit subjects who were highly experienced auction participants. eBay is an excellent venue for this purpose: eBay's publicly available feedback scores make it easy to identify people who have participated in a large number of eBay auctions. Every eBay user has a feedback profile: after the close of an eBay auction the winning bidder (and only the winning bidder) can leave feedback about the seller in the seller's feedback profile, and the seller can leave feedback in the winning bidder's profile. An eBay user's feedback score at any time is the number of positive entries in his profile minus the number of negative entries. We recruited subjects from eBay whose "feedback profiles" indicated that they had participated in at least fifty eBay auctions. These potential subjects were sent an invitation, by email, to a second-price sealed-bid auction to be conducted on the Internet. The invitation provided a link to a personalized auction web page which described the rules of the auction and provided the subject with his or her private value for the item being auctioned, as well as a form for submitting a bid.

Feedback scores typically understate a user's experience because (i) users often fail to leave feedback after a transaction, (ii) bidders who do not win the auction cannot receive feedback, and (iii) feedback cannot be reported for an auction in which the item does not sell.[6] Thus, an eBay user is likely to have participated in many more auctions than the number given by his feedback score. We recruited subjects from eBay, rather than Amazon or Yahoo auctions, as eBay is the dominant auction site and therefore provides a much larger pool of highly experienced subjects.

The Auction

Each of our experimental auctions had 5 bidders, whose values where randomly drawn from the uniform distribution on the interval [$25,$125]. In addition to their profits or losses from bidding in the auction, subjects received a $15 reward for participating. Hence, for the highest bidder, his total earnings were $15 plus his value minus the second highest bid. (If this total was negative, the loss was forgiven and the subject was paid nothing.) The other four bidders earned just the $15 participation reward. Subjects were fully informed of how their earnings would be determined. It is easy to verify that value bidding remains a weakly dominant strategy in a second-price

[6] Moreover, obtaining a negative feedback entry reduces the bidder's feedback score even though his experience has increased. Negative feedback, however, is generally a very small fraction of all feedback.

auction with a $15 limit on losses.

THE SUBJECTS

We recruited subjects by first downloading eBay Web pages for auctions in a specific category (Morgan silver dollars) that were listed as "Ending On The Current Day." The following day, after these auctions had closed, we examined the bid history of each auction. For every bidder in the bid history with a feedback score of 50 or higher we recorded (i) the bidder's eBay ID, (ii) his maximum bid, (iii) the number of times he had bid, and (iv) his feedback score. We continued this process until 50 unique IDs had been obtained. The process was then repeated to obtain 50 additional IDs from auctions of "Golden Age" collectable comics. We recruited subjects from these two auction categories because these auctions typically had many bidders, thereby reducing the difficulty of obtaining eBay IDs, and because bids in these auctions were in approximately the same range as the subjects' values would be in our experimental auction. Had we instead recruited subjects from eBay auctions where, for example, most bids were below $15, we might have introduced a significant bidding bias in our own auction.

A first set of auctions was conducted in a series of six sessions separated by a few days. In each session an invitation to participate in our experimental auction was sent via email to each one of the 100 eBay IDs we had collected for that session, as described above. A total of 67 people (out of 600) accepted our invitation and submitted bids.

With the highest possible value equal to five times the lowest possible value, those who were assigned lower values might have been less likely to participate in our auction. However, Figure 1 indicates that this did not happen. Figure 1 depicts the empirical *c.d.f.* of the values used when inviting subjects in each session of round 1 (drawn from the uniform distribution on [$25,$125]) and also the empirical *c.d.f.* of the values of the subjects who actually submitted bids at round 1. The two *c.d.f.*'s are virtually identical: the value assigned to an invitee apparently did not, on average, influence his decision whether to participate.

Figure 1 goes here.

Nor is there evidence that our experiment tended to attract the relatively unsophisticated subjects. The mean feedback score for participants was 253 (with a median score of 140), whereas the mean feedback score for invited subjects was 260 (with a median of 147), an insignificant difference.

Those who received our invitations might have been skeptical that they would actually get paid for their participation, and this might in turn have affected the bids they placed. In order to address this issue we subsequently invited our participants to a second auction, after they had actually received their first-round earnings. In this second round of auctions the rules were the same, but a new value was randomly drawn for each subject from the same uniform distribution as before. In this second round of auctions, several months after the first round, 37 of the original 67 subjects submitted bids. We did not expect bidding behavior to differ across rounds. We already knew from previous studies (Kagel and Levin, Harstad, Kagel and Levin, Harstad) that overbidding, value-bidding and under-bidding frequencies remain roughly constant over as many as twenty rounds of bidding experience in second-price auctions, and here we are only looking at 1 round of experience. In fact, bidding behavior did not differ substantially across the two rounds (see Section 3).

How the Auctions Worked

Each emailed invitation specified the deadline for submitting a bid, then directed the recipient to a Web page personalized uniquely to that invitee. The Web page described the rules of the auction, then asked three "quiz questions" about the auction rules, and then provided the subject with his value for the auctioned item.[7] A subject had no direct monetary incentive to answer the questions correctly, but was required to give answers to all three questions before he was allowed to submit a bid. The answers to the questions provide some indication of whether a subject understood the rules of the auction; 70% of the subjects answered all three questions correctly and another 24% answered two of the three questions correctly.

At the end of each session the bids were placed into groups of five in the order in which they arrived.[8] The subjects' earnings were calculated, and each subject was sent an email describing the bid and value of each bidder in his auction, as well as his own earnings. Each subject was then mailed a money order containing his earnings.

[7] A sample webpage is available at www.econ.ucsb.edu/∼garratt/auction/sample.html. The text of the invitation email is available here: www.econ.ucsb.edu/∼garratt/auction/email.html.

[8] Each auction had five bidders, but the number of bids received in a session was typically not a multiple of five. The "remainder group" in each session was filled out with bids randomly selected from the other groups. For example, if seven bids were received, then bids 1-5 formed one group which determined the payoffs of bidders 1-5. Bids 3-7, say, formed a second group, which determined the payoffs of only bidders 6 and 7.

3 Analysis of Bidder Behavior

Figure 2 shows bids and "high bids" (i.e., bids above $150) for each round of the experiment. The high bids were frequently for $9,999,999,999, which was the largest bid that could be entered on the webpage. Such bids illustrate the illusion that bidding in excess of value increases the probability of winning without any cost since the winner pays the second highest bid. It is immediate from Figure 2 that, despite having substantial experience with auctions in the field, eBay subjects typically do not value bid.

Figure 2 goes here.

Table 1 reports, by round, the number of under bids, value bids, and overbids by our subjects. We have retained the K&L definition of a value bid – i.e., any bid within five cents of the subject's value – but Table 1 would be almost unchanged if a value bid were defined as any bid differing from value by less than a dollar: only one of the 104 bids differed from value by more than five cents but less than a dollar. Using Pearson's chi-square goodness-of-fit test we cannot reject the hypothesis that the distributions over the three types of bids shown in Table 1 are the same at round 1 and round 2 of the experiment, and hence hereafter we pool the data at both rounds.[9]

	Under bids	Value bids	Over bids	Total
Round 1	27 (40.3%)	15 (22.4%)	25 (37.3%)	67 (100%)
Round 2	16 (43.2%)	7 (18.9%)	14 (37.8%)	37 (100%)

Table 1: Frequency by round of under, over, and value bids

Table 2 reports the number of underbids, value bids, and overbids by our subjects for both rounds combined. The table also shows the comparable numbers for the subjects in the K&L experiment. The first conclusion we draw is that value bidding was observed no more often in our auctions than in the K&L auctions: 21% of bids

[9] The value of the Pearson Q is .187 and the 5% critical value, with two degrees of freedom, is 5.99. The regression results in column (b) of Table 3 show that we cannot reject the null hypothesis that the linear bidding model is the same in both rounds. The joint test of the significance of the round dummy and the round dummy interacted with value has a p-value of .7621.

in our experiment were value bids and 27% of the K&L bids were value bids.

Experiment	Under bids	Value bids	Over bids	Total
eBay	43 (41.3%)	22 (21.2%)	39 (37.5%)	104 (100%)
K&L	27 (5.7%)	127 (27.0%)	316 (67.2%)	470 (100%)

Table 2: Frequency of under, over, and value bids

Conclusion 1: *The frequency of value bidding in our experiment is indistinguishable from its frequency in the K&L experiment. Only about one quarter of the bids are value bids in each case.*

Table 2 nevertheless indicates that bidding behavior in our auctions was dramatically different than in the K&L auctions. Subjects in the K&L auctions submitted more than ten times as many overbids as underbids (67% of bids were overbids and only 6% were underbids), which led K&L to conclude that overbidding is pervasive in SPAs and that underbidding is relatively unusual. In our auctions, however, only 38% of the bids were overbids and 41% were underbids.

Conclusion 2: *In our experiment the frequencies of overbidding and underbidding are indistinguishable from one another. There is no more tendency to overbid than to underbid.*

Our conjecture at the outset was that in SPAs bidding by people with significant experience in real-world auctions might conform more closely to the theory than bidding by inexperienced laboratory subjects. We have already seen in Conclusion 1 that the frequency of value (i.e., "correct") bidding in our data does not support this conjecture. Table 3 provides further evidence that the amount of a bidder's experience in real-world auctions does not affect bidding behavior. In Column (a) of Table 3 the magnitude (i.e., the absolute value) of subjects' bidding errors (the difference between value and bid) is regressed against the amount of a subject's experience, where experience is measured by subjects' feedback scores. The regression coefficient for feedback score does not differ significantly from zero: the amount of experience does not seem to affect bidding behavior. This is perhaps not surprising, since all of our subjects were highly experienced.

Conclusion 3: *Among our subjects, who were all highly experienced, variations in the amount of their experience have no systematic effect on their bidding behavior.*

Table 3 goes here.

While the *amount* of a subject's experience seems to have no effect on his behavior, the *kind* of experience a subject has appears to make a difference in bidding behavior. About half of our subjects had sold items on eBay, and the other half had only been bidders, never selling anything.[10,11] Furthermore, those who had been sellers typically had a great deal of experience as sellers: the median feedback count for them *as sellers* was 57. Column (c) of Table 3 shows that subjects who have experience as a seller on eBay tend to bid less than subjects who have only bought, bidding $14.19 less on average. Column (d) shows that experience as a seller remains significant even when controlling for the type of auction a subject was recruited from (Morgan dollar or comic), and controlling for the highest bid and the number of bids he placed in that auction.

The regression results reported in Table 3 exclude 11 bids of more than $1000. The majority of these bids (7 of 11) were made by subjects with experience as sellers, and hence to exclude them potentially biases our conclusions regarding seller experience. To address this possibility, we examine the effect of seller experience on the frequency of under, value, or overbidding, using all 97 bids (and, in particular, including bids over $1000) for which we can distinguish buyer and seller experience.[12] Table 4 shows that subjects who had been sellers behaved quite differently than those who had not. Subjects who had been sellers submitted 51% underbids and 32% overbids; those who had never sold submitted 30% underbids and 46% overbids, almost the exact reverse of the frequencies for those who had been sellers. We can reject at the 10% significance level (the p-value is 0.07) the hypothesis that the sellers' and the buyers' frequencies are realizations of independent draws from the same multinomial distribution.[13]

[10] eBay's feedback system did not indicate whether a user was a buyer or a seller in a transaction until one month after we conducted the first round of our experiment. We classified a subject as a "seller" if he received feedback as a seller in the following year. We only imperfectly observe whether a subject had experience as a seller. This works against us, making it more difficult to identify a difference between sellers and buyers, even when such a difference exists.

[11] The number of subjects who had sold and the number who had only bid are 34 and 29, respectively.

[12] Four of our subjects, who made 7 bids in total, were no longer registered users when we identified subjects with experience selling.

[13] The test of the equality of two multinomial distributions assumes independent draws, which may not be valid since the type of bid a bidder places may not be independent between the two rounds. However, for both buyers and sellers we cannot reject at the 10% level the hypothesis that the type of bid placed is independent between rounds.

Experiment	Under bids	Value bids	Over bids	Total
eBay – only buyer	13 (29.5%)	11 (25.0%)	20 (45.5%)	44 (100%)
eBay – sometime seller	27 (50.9%)	9 (17.0%)	17 (32.1%)	53 (100%)

Table 4: Frequency of under, over, and value bidding by type of experience.

As an added check whether the bidding behavior of buyers was different from that of sellers, we apply the Mann-Whitney rank-sum test to the entire bid distributions of sometime-sellers and only-buyers. The p-value of the Mann-Whitney test statistic is .093, a rejection at the 10% level that the differences between bids and values for the two groups of subjects were drawn from the same distribution. The two distributions of bid-minus-value are depicted in Figure 3, for differences of less than $100.

Conclusion 4: *Subjects who had sold on eBay exhibited a lower frequency of overbidding and a higher frequency of underbidding than subjects who had bought on eBay but never sold. The two groups' distributions of bids-minus-values do not appear to be the same: those who had sold typically bid less relative to their values than those who had not sold.*

Figure 3 goes here.

Perhaps subjects with selling experience view the values we provide them with differently from subjects who have only bought. Subjects who have sold on eBay may take the view that to earn a profit on an item it's necessary to buy it at a price below what they expect to sell it for. That is, they aim to pay strictly less than their induced value with the objective of realizing a positive markup. Subjects who have only bought may correctly view the induced value as how much it is worth to them to win the auction.

Revenue and Efficiency

Our Conclusion 1, above, reinforces results obtained previously in second-price-auction experiments: bidders (even experienced eBay bidders) generally do not bid their value, as the theory suggests they should. Moreover, our Conclusion 4 indicates that bidders who are experienced as sellers bid non-optimally but differently than bidders with no selling experience.

But how great is the effect of non-optimal bidding? What is its effect on the revenue the seller will obtain in the auction? On the surplus a bidder can expect? Or on the auction's efficiency?

Seller's Revenue

In our auction and in the K&L auction, the left-most column of Table 5 shows the *ex-ante* expected revenue that would accrue to the seller in the auction – i.e., the revenue the seller can expect before bidders' values are drawn – assuming the bidders all submit bids equal to their values. The center column again shows the seller's expected revenue if bidders bid their values, but now this expectation is determined from the empirical distribution of the values actually drawn in the experiment. The right-most column shows the expected revenue to the seller when five bids are randomly drawn from the empirical distribution of the subjects' actual bids.

	Uniform Draws	*Values Actually Drawn*	
	Value Bidding	Value Bidding	Actual Bids
eBay	\$91.67	\$92.09	\$90.97
K&L	\$18.87	\$18.28	\$19.26

Table 5: Effect of suboptimal bidding on seller revenue.

While there were nearly equal numbers of underbids and overbids by the eBay subjects (Table 2), the net effect on seller revenue is slightly negative: average revenue is only \$90.97, a reduction of \$1.12, or about one percent, due to non-optimal bidding. (The average revenue calculation is possibly biased downward, however, by the bidders' \$15.00 limited liability: whenever the second-highest bid exceeds the value of the winning bidder by more than \$15, we state seller revenue as the winning bidder's value plus \$15.) In the K&L experiment the seller's average revenue, given the empirical distribution of bids, is 98 cents higher than would be expected had subjects bid their value – an increase in revenue to the seller of about five percent.[14]

Subjects' non-optimal bidding in SPAs thus appears to have only a modest effect (in expectation) on seller revenue. Note, however, that the percentage effect on the seller's *profit*, or *surplus*, would generally be larger (possibly *much* larger), since that depends also on his reservation value, or cost, for the item at auction. For example, if the seller's reservation value in our auctions were \$75 (the midpoint of the bidders' distribution of values), then the \$1.12 expected reduction in his revenue due to non-optimal bidding yields about a 6 1/2% reduction in his expected profit. The K&L 5% increase in seller revenue would similarly be larger as a percentage of his profit.

[14] This expected revenue calculation for K&L assumes a \$10 limit on liability.

Bidders' Surplus

The effect on the bidders' expected surplus (i.e., their net earnings) from failing to bid their values is larger than the effect on seller revenue. Table 6 reports the average expected surplus (averaged across all 104 values that were actually drawn) from value bidding and from the bids the subjects actually placed. The left column reports the average expected surplus if each subject had bid his value, given the values that were actually drawn. The remaining two columns report a bidder's expected surplus if the opposing four bids are drawn at random from the actual bids placed: the center column is the average expected surplus (averaged across the 104 actual values) from value bidding; and the right column reports the average expected surplus from the bids actually placed.

	vs Value Bids	*vs Actual Bids*	
	Value Bidding	Value Bidding	Actual Bids
All bidders	\$3.84	\$2.77	\$1.05
eBay – only buyer	\$3.05	\$2.28	\$0.75
eBay – sometime seller	\$4.46	\$3.14	\$1.03

Table 6: Effect of nonoptimal bidding on bidders' expected surplus.

The left column is the benchmark against which to evaluate the cost to bidders of failing to bid their values: given the values actually drawn, it tells us the expected surplus had everyone placed the bid that was optimal for him, given his own value. Thus, as a group our subjects sacrificed \$2.79 in expected surplus per bidder, nearly 75% of the maximum achievable expectation, \$3.84. Given that other subjects were bidding non-optimally, the typical subject's non-optimal bid still reduced his expected surplus by more than 60%, to \$1.05 from the \$2.77 in expected surplus he could have achieved.

Efficiency

The effect of suboptimal bidding on auction efficiency is more striking. Given the empirical distribution of bid-value combinations in our experiment, the probability is only 0.46 that in a randomly selected group of 5 of our subjects the bidder with the highest value would also be the bidder who placed the highest bid. If we exclude from the empirical distribution the bid-value combinations in which the bid is above \$1000, this probability rises only to 0.56. In the first paid round of K&L's experiment, following several practice rounds, the subjects achieved an efficiency level of 0.69.

4 Concluding Remarks

eBay's auction institution operates much like a second-price auction. However, we find that even when highly experienced eBay participants bid in an actual second-price sealed-bid auction, they do not typically bid their values, as the theory suggests they should. Significant experience in a similar setting does not seem to help bidders learn how to bid optimally in second-price sealed-bid auctions. Thus, the non-optimal bidding previously discovered with student subjects in the laboratory appears to be a phenomenon that is robust even to substantial experience in non-experimental auctions.

In contrast with experiments in the laboratory, where subjects have been observed to typically overbid and almost never underbid, subjects in our experiment were just as likely to underbid as to overbid. Of course, this difference might be in part a result of differences in the experimental design. Since our experiment was conducted over the Internet, it necessarily used different instructions than the K&L experiment. For example, a quiz question in our instructions illustrated that a subject could lose part or all of his participation reward if he won the auction and the second highest bid was higher than his own value. This possibility was not explicitly mentioned in the K&L instructions.

However, a more likely explanation of the significantly larger propensity of subjects in our experiment to underbid is the presence of many subjects with substantial experience *selling* in eBay auctions. Subjects with experience as sellers bid less in a second-price auction, on average, than subjects who have only bought on eBay. This suggests that the revenue and efficiency properties of auctions might depend in a systematic way on the backgrounds and prior experiences of the bidders.

This highlights an advantage of using eBay participants as experimental subjects: eBay provides a rich set of publicly available data on each of its auctions, and this allows the experimenter to apply data from subjects' field experience to try to explain their behavior in experiments.

In light of the overbidding exhibited by student subjects in laboratory experiments on SPAs, and the increase in seller revenue that results, one wonders why SPAs are not used more commonly in practice. Our results suggest an explanation. The combination of frequent underbidding by bidders with auction experience actually leads to a reduction of revenue in SPAs relative to the theoretical prediction.

References

[1] Coppinger, V., Smith, V., and J. Titus (1980): "Incentives and Behavior in English, Dutch, and Sealed-Bid Auctions." *Economic Inquiry* **18**, pp. 1-22.

[2] Cooper, D. and J. Kagel (2006): "The Role of Context and Team Play in Cross-Game Learning," mimeo.

[3] Cooper, D., Kagel, J., Lo, W., and Q. Gu (1999): "Gaming against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers," *The American Economic Review* **89**, pp. 781-804.

[4] Cox, J., Roberson, B., and V. Smith (1982): "Theory and Behavior of Single Object Auctions," in V. L. Smith, ed., Research in Experimental Economics, vol. 2, Greenwich, Conn.: JAI press, pp. 1-43.

[5] Dyer, D., Kagel, J. and D. Levin (1989): "The Behavior of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis." *Economic Journal* **99**, pp. 108-115.

[6] Harstad, R. (2000): "Dominant Strategy Adoption, Efficiency and Bidders' Experience with Pricing Rules." *Experimental Economics* **3**, pp. 261-280.

[7] Hannan, L., Kagel, J. and D. Moser (2002): "Partial Gift Exchange in an Experimental Labor Market: Impact of Subject Population Differences, Productivity Differences, and Effort Requests on Behavior," *Journal of Labor Economics* **20**, pp. 923-951.

[8] Johnson-Laird, P. N. (1999): "Deductive Reasoning," *Annual Review of Psychology* **50**, 109-135.

[9] Kagel, J.. Harstad, R. and D. Levin (1987): "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica* **55**, pp. 1275-1304.

[10] Kagel, J. and D. Levin (1993): "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third Price Auctions with Varying Numbers of Bidders." *Economic Journal* **103**, pp. 868-879.

[11] List, J. (2003): "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics* **118**, pp. 41-71.

[12] Palacios-Huerta, I. and O. Volij (2006): "Experientia Docet: Professionals Play Minimax in Laboratory Experiments," mimeo.

[13] Roth, A. and A. Ockenfels (2002): "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review* **92**, pp. 1093-1103.

Figure 1

CDF of Invitee/Participant Values

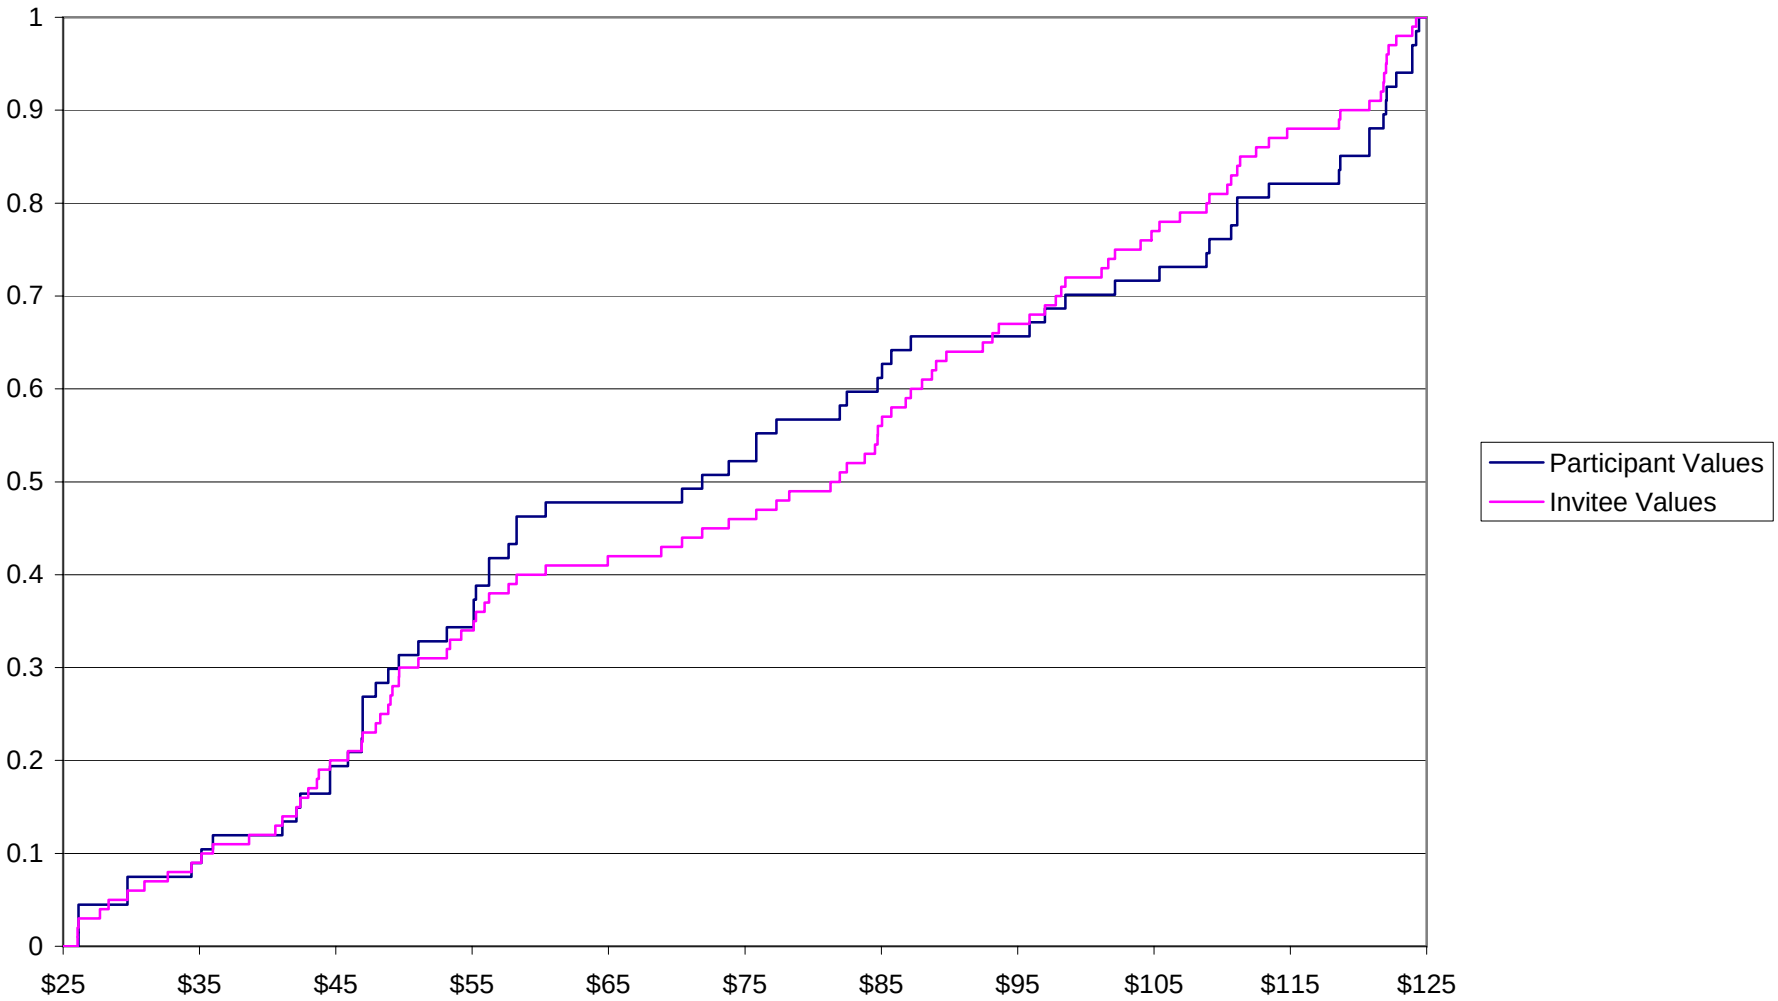

Figure 2

Bid vs. Value - Both Rounds

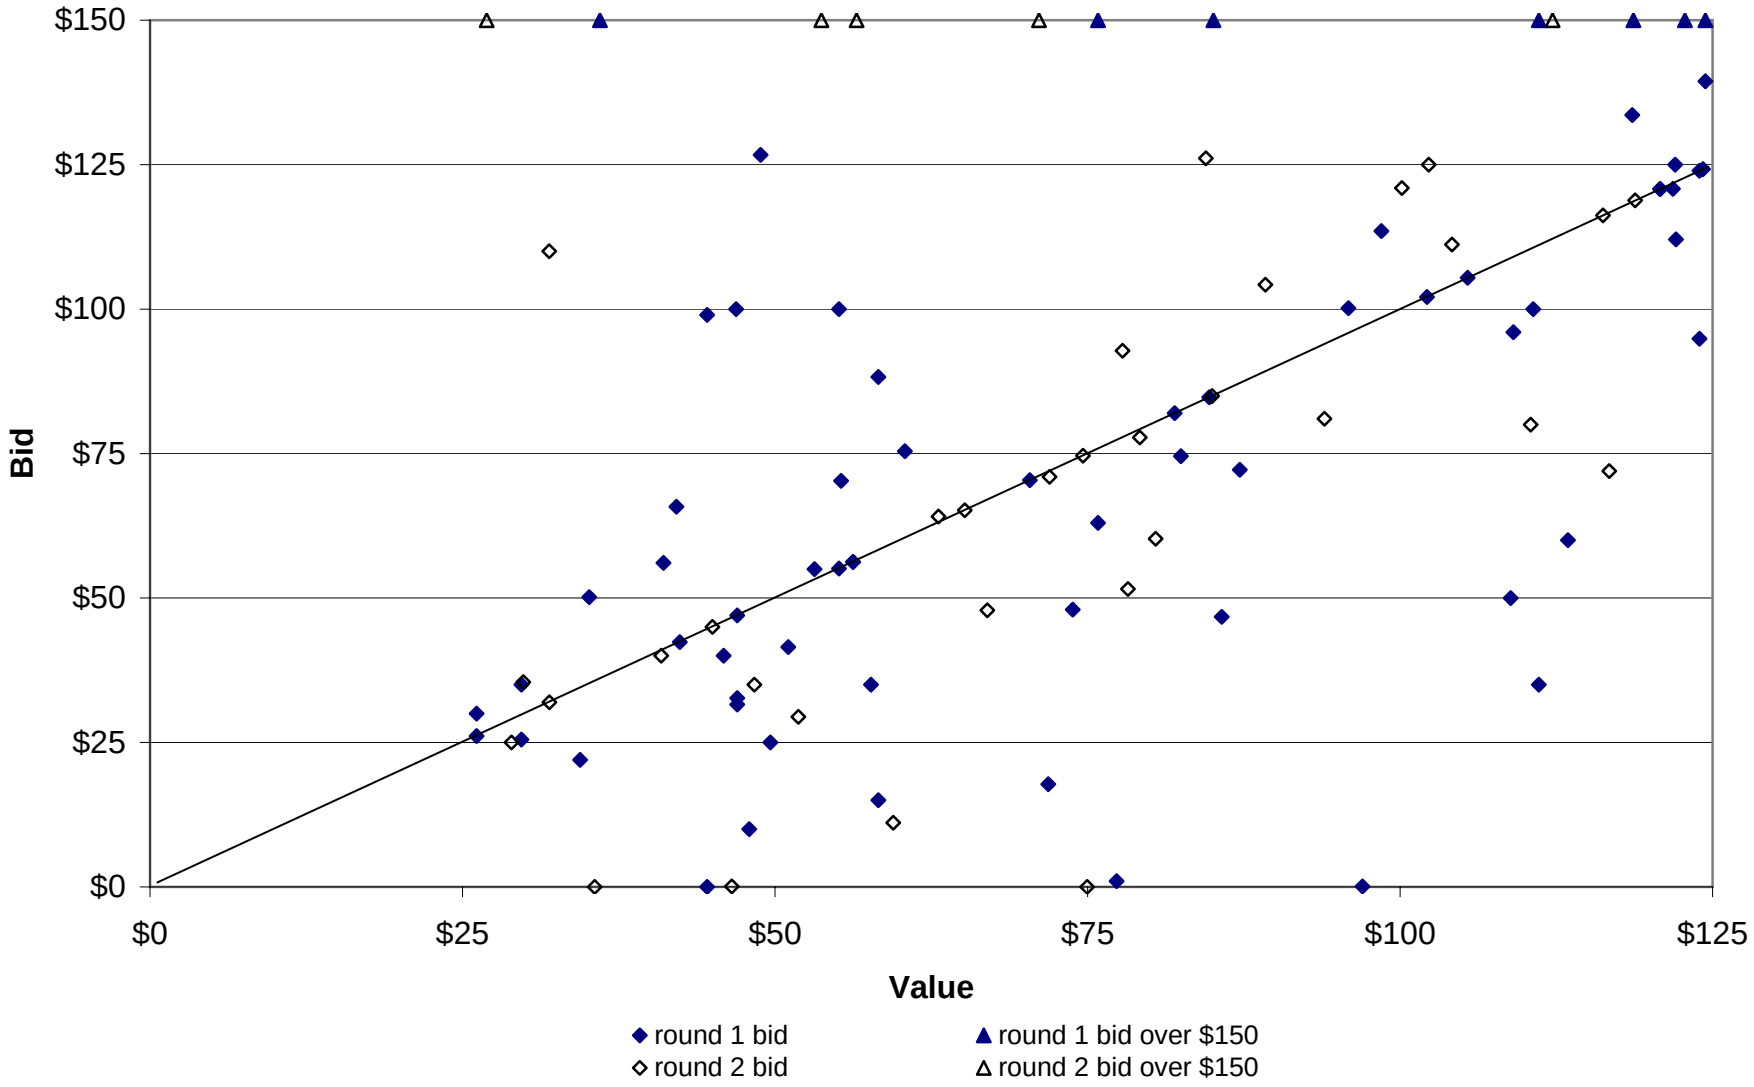

Figure 3

CDF of bid minus value

(for differences of less than $100)

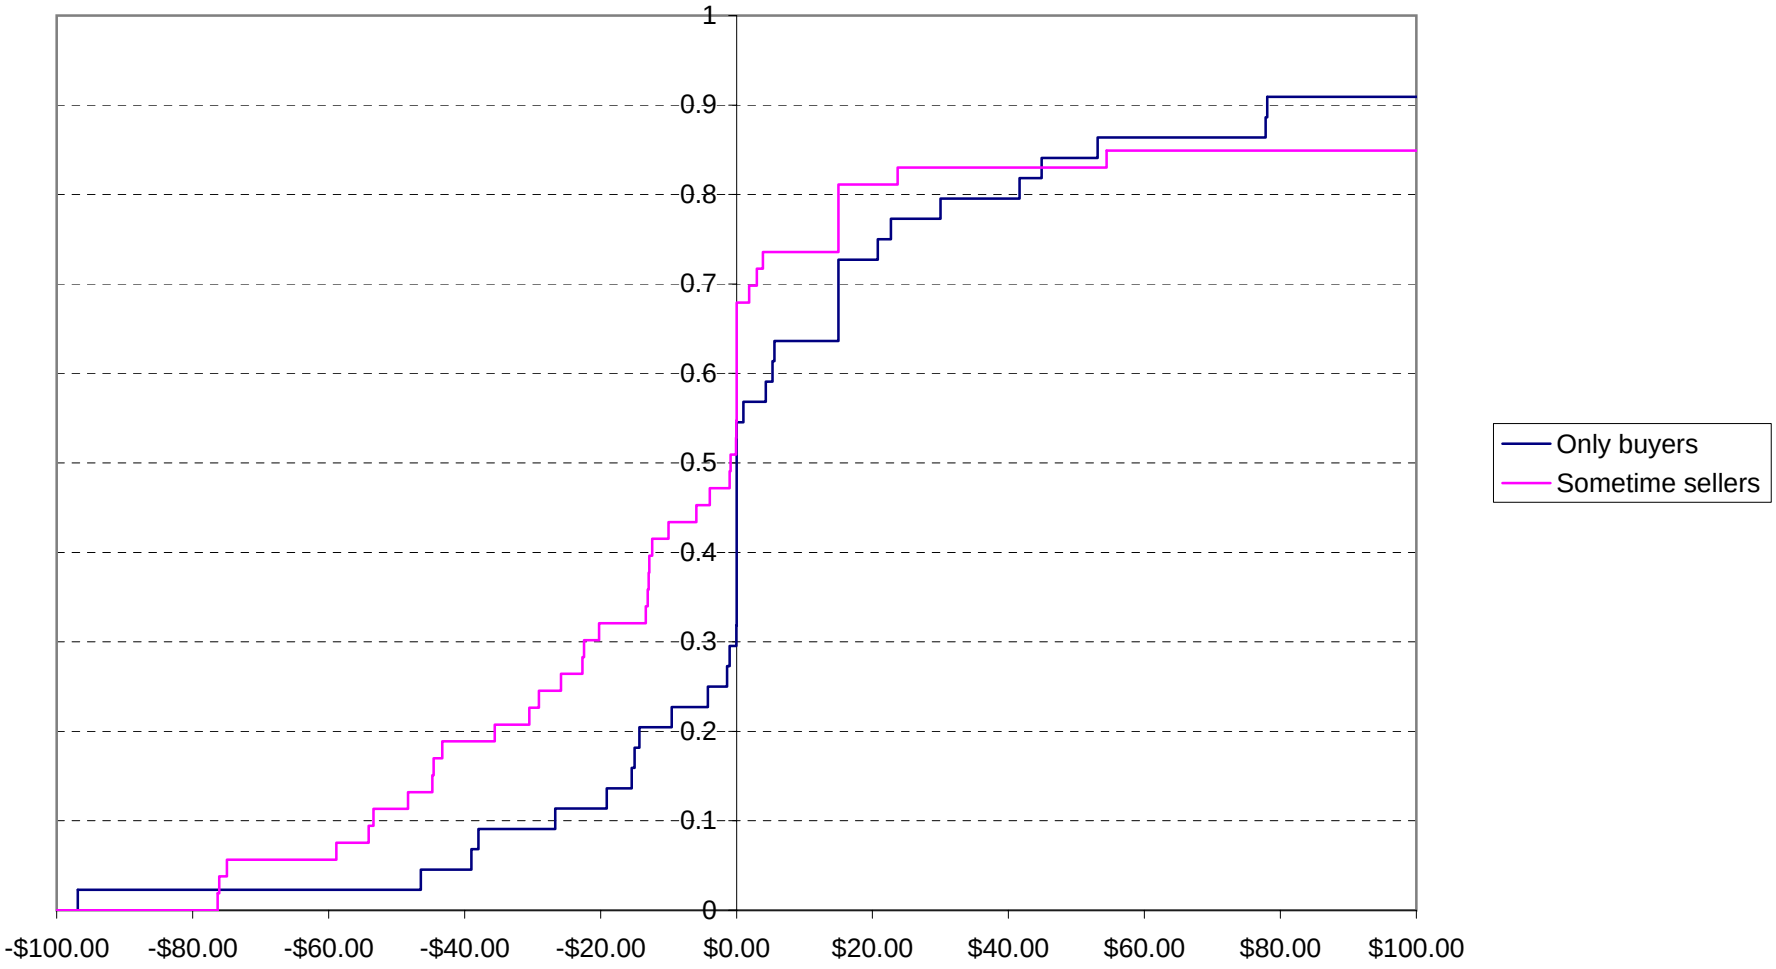

Dependent variable	Abs. Value of Overbid	Bid		
	(a)	(b)	(c)	(d)
Constant	21.916 (3.754)**	11.153 (11.352)	14.043 (9.560)	15.594 (14.272)
Value		0.760 (0.140)**	.838 (0.118)**	0.825 (0.118)**
Round dummy: 0 if round 1, 1 if round 2		-9.195 (21.027)		
Round dummy * Value		0.171 (0.275)		
Feedback score	-.004 (.010)			
Log of highest bid subject placed in eBay auction				-1.462 (2.830)
Auction type: 1 if Morgan, 0 if Comic				13.682 (7.127)
Number of bids a subject placed in the eBay auction from which he/she was recruited				-1.568 (3.757)
Seller dummy: 1 if some seller feedback, 0 if only buyer feedback			-14.193 (7.067)*	-15.204 (7.458)*
Observations	93	86	86	86
R-squared (adj.)	-.010	.331	.365	.375

Standard errors are in parentheses: * significant at 5% level, ** significant at the 1% level. All regressions exclude bids over $1000 (11 obs.). (b)-(d) exclude bids by four subjects for whom we were unable to obtain their buyer/seller feedback (7 obs.).

Table 3: Regression results