

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

Rodney J. Garratt[†] Mark Walker[‡] John Wooders[§]

August 23, 2009

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 tended not to bid their values, suggesting that the non-optimal bidding of novice subjects is robust to substantial experience in non-experimental auctions. The data suggest that the *kind* of eBay experience a subject has – selling or buying – has an effect on bidding behavior.

*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, University of Arizona, Tucson, AZ 85721 (mwalker@arizona.edu).

[§]Department of Economics, University of Arizona, Tucson, AZ 85721 (jwooders@eller.arizona.edu).

1 Introduction

In a second-price private-value 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, so that people have time to reflect on how outcomes are affected by their decisions. Hence people with some experience in real-world auctions might be 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. This enables us to test the hypothesis that real-world experience in auctions leads to more nearly optimal bidding in SPAs and to shed light on the external validity of

¹Kagel 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. Guth and Ivanova-Stenzel (2003) find that overbidding is reduced if bidders do not know the distribution of their rivals' values.

previous laboratory results for SPAs.

eBay auctions are not exactly the same as SPAs, but they 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. 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.² 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 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 appears to be *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 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. A possible explanation is provided at the end of Section 3.

²However, during an eBay 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. Of course, in an eBay auction bidders don’t know how many rival bidders there will be.

³Roth 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 an SPA.

⁴After our experiment was conducted, eBay introduced “Second Chance Offers,” which allows a seller, after the close of his auction, to 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.

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.

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.⁵ 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.

THE AUCTION

Each of our experimental auctions had 5 bidders, whose values were randomly drawn from the uniform distribution on the interval $[\$25, \$125]$. In addition to their profits or losses

⁵As of May 2008 a seller can leave only positive feedback for the buyer.

⁶Moreover, 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.

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 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 eBay bidders. The mean feedback score for participants was 253, whereas the mean feedback score for invited subjects was 260, an insignificant difference. Indeed, the Mann-Whitney test of whether the feedback scores of invitees and participants are drawn from the same distribution yields a *p*-value of .86.

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. The rules were the same, but a new value was randomly drawn for each subject. In this second round of auctions, several months after the first round, 37 of the original 67 subjects submitted bids. One would not expect bidding behavior to differ across the two rounds: previous studies (Kagel and Levin, Harstad, Kagel and Levin, Harstad) have found that overbidding, value-bidding and underbidding frequencies remain roughly constant over as many as twenty rounds of bidding by the same subjects in second-price auctions. In fact, bidding behavior did not differ substantially across the two rounds (see Section 3).⁷

HOW THE AUCTIONS WORK

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.⁸ 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.⁹ 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.

⁷There is no evidence that the outcome in round 1 influenced the decision to participate in round 2. Probit regressions reported in columns (a) and (b) of Table 3 show no significant impact of round 1 earnings on the probability of participating in round 2.

⁸A sample webpage is available at www.econ.ucsb.edu/~garratt/auction/sample.html. The text of the invitation is here: www.econ.ucsb.edu/~garratt/auction/email.html.

⁹Each 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 the values and bids for each round of the experiment. There are 11 outliers – bids greater than \$1000 – shown as points along the top edge of the graph. Of these, five were for \$9,999,999,999, which was the largest bid that could be entered on the webpage. Such bids demonstrate 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.¹⁰

	Under bids	Value bids	Over bids	Total
eBay Round 1	27 (40.3%)	15 (22.4%)	25 (37.3%)	67 (100%)
eBay Round 2	16 (43.2%)	7 (18.9%)	14 (37.8%)	37 (100%)
eBay Combined	43 (41.3%)	22 (21.2%)	39 (37.5%)	104 (100%)

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

	Under bids	Value bids	Over bids	Total
K&L Round 1	5 (25.0%)	1 (5.0%)	14 (70.0%)	20 (100%)
K&L Round 2	0 (0.0%)	6 (30.0%)	14 (70.0%)	20 (100%)
K&L Rounds 1 and 2	5 (12.5%)	7 (17.5%)	28 (70.0%)	40 (100%)
K&L All Rounds	27 (5.7%)	127 (27.0%)	316 (67.2%)	470 (100%)

Table 2: Frequency of under, over, and value bids in K&L

¹⁰The 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 (d) 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.

Table 2 shows the number of under bids, value bids, and overbids by subjects in the first two rounds of the K&L experiment, along with the combined totals for the first two rounds and the combined totals for all rounds. The first conclusion we draw is that the extent of value bidding was not significantly greater in our auctions than in the K&L auctions: 21.2% of bids in our experiment were value bids and 17.5% of the K&L bids in the first two rounds were value bids.

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

Tables 1 and 2 nevertheless indicate 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.¹¹

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 (c) 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 the magnitude of his bidding error. 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.

¹¹K&L’s data comes from two sessions of 10 subjects each. In these experiments values were uniformly distributed between \$0 and \$28.60.

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.¹² Column (e) of Table 3 shows that subjects who had experience as a seller on eBay tended to bid less than subjects who had only bought, bidding \$14.19 less on average. Column (f) 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 the subject placed in that auction. (Columns (g) and (h) repeat these regressions with clustered standard errors.)

The regression results reported in Table 3 exclude the 11 outlier bids along the top edge of Figure 2. 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 for which we can distinguish buyer and seller experience (including the outlier bids).¹³ 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 from the same multinomial distribution.¹⁴

¹²eBay’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 makes it more difficult to identify a difference between sellers and buyers if such a difference exists.

¹³Four of our subjects, who made 7 bids in total, were no longer registered users when we identified subjects with experience selling.

¹⁴The 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 buyers and for sellers independence across rounds is not rejected at the 10% level.

	Under bids	Value bids	Over bids	Total
eBay – only buyer	13 (29.5%)	11 (25.9%)	20 (45.5%)	44
Round 1	9 (31.0%)	9 (31.0%)	11 (38.0%)	29
Round 2	4 (26.7%)	2 (13.3%)	9 (60.0%)	15
eBay – sometimes seller	27 (50.9%)	9 (17.0%)	17 (32.1%)	53
Round 1	15 (41.1%)	6 (17.6%)	13 (38.2%)	34
Round 2	12 (63.2%)	3 (15.8%)	4 (21.0%)	19

Table 4: Frequency of under, over, and value bids 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 bid-minus-value 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 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.

A possible explanation why the subjects with experience selling on eBay bid differently than subjects without such experience is found in the psychology literature on deductive reasoning processes.¹⁵ 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 about 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 may be traders, for example, buying items and then reselling them at higher prices on eBay, aiming to achieve a positive markup, and these subjects may have 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.

¹⁵We 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.

4 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 suggests that bidders who are experienced as sellers bid non-optimally but differently than bidders with no selling experience. Here we investigate the effect of non-optimal bidding on revenue from the auction, on bidder surplus, and on the auction’s efficiency.

Revenue

Table 5 shows how non-optimal bidding by the subjects in our experiment and in the K&L experiment affect the revenue the auction generates.¹⁶ The table’s first column shows the expected revenue the auction would generate if bidders bid their values, given the empirical distribution of the values actually drawn in the experiment. The next column shows the expected revenue when five bids are randomly drawn from the empirical distribution of all the subjects’ actual bids.¹⁷ The third column shows the expected “excess revenue” – the expected amount by which the revenue from actual bids exceeds the revenue from value bidding. The last column contains an estimate of the standard error of the expected excess revenue in each case.¹⁸

	Expected Revenue from Value Bidding	Expected Revenue from Actual Bids	Expected Excess Revenue	Standard Error of Expected Excess Revenue
eBay Round 1	\$95.89	\$90.45	-\$5.44	\$3.55
eBay Round 2	\$86.13	\$78.78	-\$7.35	\$5.05
K&L Round 1	\$19.48	\$20.30	\$0.82	\$0.63
K&L Round 2	\$22.18	\$22.82	\$0.64	\$0.12
K&L Final Round	\$14.26	\$15.12	\$0.86	\$0.18

Table 5: Effect of non-optimal bidding on auction revenue

In both rounds of our experiment the net effect of non-optimal bidding is to reduce expected revenue, despite the nearly equal incidence of underbids and overbids (*cf.* Table 2),

¹⁶The last row of Table 5 is based on round 23 of the SPA sessions in K&L .

¹⁷Whenever the second-highest of the five bids exceeded the winning bidder’s value by more than \$15, revenue was computed as the winning bidder’s value plus the \$15 participation fee forgone by the bidder.

¹⁸In order to obtain estimates of standard errors, the values in this column and the remaining two columns were computed by averaging the results of ten runs of the Mullin-Reiley recombinant estimation calculator. The calculator is available at <http://www.u.arizona.edu/~dreiley/papers/RecombinantEstimator.xls>.

but the effect is not statistically significant.¹⁹

In the K&L experiment the expected revenue, given the empirical distribution of bids, is slightly higher than would be expected had subjects bid their values. This increase in expected revenue was statistically insignificant in round 1, but in the other two K&L rounds the excess revenue was close to five times the standard error, statistically significant at any reasonable level.²⁰

Efficiency

Given the empirical distribution of bid-value combinations in round 1 of our experiment, the probability is only 0.53 that in a randomly selected group of 5 of our subjects the bidder with the highest value also placed the highest bid. This number drops to 0.33 in round 2 (standard errors are 0.065 and 0.091, respectively).²¹ In contrast, in K&L’s experiment, the probability that the outcome is efficient is 0.69 in round 1, 0.84 in round 2, and it increases to 0.96 in the final round (standard errors are 0.144, 0.085 and 0.025, respectively). The higher efficiency in K&L is perhaps surprising since both experiments saw similarly low amounts of value bidding. However, bidding errors in K&L’s experiment tend to be uniformly in the direction of overbidding, which does not impact efficiency as much as a mixture of over- and under-bidding.²² Efficiency of SPAs may therefore be less than one would predict from experiments with student subjects, who are perhaps more homogeneous than typical bidders in real-world auctions.

Table 6 shows how the surplus is divided between the seller and the bidders in our auctions and in those conducted by K&L. It reports the results of randomly drawing groups of 5 bid-value pairs from the empirical distribution of such pairs. The “Efficient Surplus” column shows the expected highest value when forming such groups; the “Seller Revenue” column shows the expected value of the second highest bid; and the “Bidder Profit” column shows the expected value of the difference between the value of the winning bidder and the price he pays (the second highest bid). Forgone surplus is obtained by subtracting the sum of seller revenue and bidder profit from the efficient surplus. This table shows that auction efficiency, as measured by the percentage of surplus captured, was substantially lower in our SPAs than in K&L’s auctions with student subjects.

¹⁹None of the ten excess revenue estimates generated by our recombinant calculations differed from zero at the 10% level of significance.

²⁰K&L assumes a \$10 limit on liability, however this was not binding in these rounds.

²¹If we exclude the outlier bids, these probabilities rise to 0.62 and 0.37 with similar standard errors.

²²If everyone overbids by the same fixed amount, for example, the outcome remains efficient.

	Seller Revenue	Bidder Profit	Forgone Surplus	Efficient Surplus	Captured Surplus
eBay Round 1	\$90.59	\$8.69	\$15.32	\$114.60	86.6%
eBay Round 2	\$78.87	-\$1.72	\$26.94	\$104.09	74.1%
K&L Round 1	\$20.32	\$1.05	\$2.04	\$23.41	91.3%
K&L Round 2	\$22.82	\$2.68	\$0.18	\$25.68	99.3%
K&L Final Round	\$15.09	\$4.83	\$0.06	\$19.98	99.7%

Table 6: Distribution of the surplus and auction efficiency

Bidders suffered larger losses in our experiment than in K&L, both in absolute and percentage terms. In round 1 of our experiment, the bidders captured only \$8.69 (or 46%) of the \$18.71 ($= \$114.6 - \95.89) of potential surplus they would have obtained under value bidding. In round 2 bidders on average suffered losses. In K&L, by contrast, bidders captured respectively, 26.7%, 76.6%, and 84.4% of the surplus achievable under value bidding in the first, second, and final round.

5 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 previous experiments, where subjects have been observed to typically overbid and almost never underbid, the subjects in our experiment (who were all highly experienced) 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 plausible explanation of the greater tendency of subjects in our experiment to underbid is that many of them had substantial experience *selling* in eBay auctions. Subjects with experience as sellers bid less, 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.

Our experiment 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 use data from subjects' field experience to explain their behavior in experiments.

In light of the overbidding exhibited by student subjects in laboratory experiments on SPAs, one might ask why SPAs are not used more commonly in practice in order to generate more revenue. Our results suggest an explanation. Frequent underbidding by bidders with auction experience may eliminate any potential gain in revenue from an SPA.

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] Guth, W. and R. Ivanova-Stenzel (2003): "Asymmetric auctions experiments with(out) commonly known beliefs," *Economic Letters* **80**, p. 195-199.
- [7] Harstad, R. (2000): "Dominant Strategy Adoption, Efficiency and Bidders' Experience with Pricing Rules." *Experimental Economics* **3**, pp. 261-280.
- [8] 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.

- [9] Johnson-Laird, P. N. (1999): “Deductive Reasoning,” *Annual Review of Psychology* **50**, 109-135.
- [10] 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.
- [11] 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.
- [12] List, J. (2003): “Does Market Experience Eliminate Market Anomalies?” *Quarterly Journal of Economics* **118**, pp. 41-71.
- [13] Mullin, C. and D. Reiley (2006): “Recombinant estimation for normal-form games, with applications to auctions and bargaining.” *Games and Economic Behavior* **54**, pp. 159-182.
- [14] Palacios-Huerta, I. and O. Volij (2006): “Experientia Docet: Professionals Play Minimax in Laboratory Experiments,” mimeo.
- [15] 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	Participate in Round 2		Absolute Value of Bid		Bid					
	(a) Probit	(b) Probit	(c) OLS	(d) OLS	(e) OLS	(f) OLS	(g) OLS, clustered SE	(h) OLS, clustered SE		
Specification										
Constant		0.978 (0.673)	21.863 (3.747)**	10.695 (10.573)	14.043 (9.560)	15.594 (14.272)	14.043 (8.915)	15.594 (13.607)		
Value				0.765 (0.131)**	0.838 (0.118)**	0.825 (0.118)**	0.838 (0.091)**	0.825 (0.096)**		
Round dummy: 0 if round 1, 1 if round 2				-10.092 (19.202)						
Round dummy x Value				0.193 (0.246)						
Feedback score			-0.004 (.010)							
Value in Round 1			-0.0003 (0.002)							
Value in Round 2			-0.005 (.002)**							
Round 1 Earnings			0.002 (0.007)							
Auction type: 1 if Morgan, 0 if Comic			0.124 (0.126)			13.682 (7.127)		13.682 (7.714)		
Log of highest bid subject placed in eBay auction						-3.365 (6.516)		-3.365 (7.022)		
Number of bids a subject placed in the eBay auction from which he/she was recruited						-1.568 (3.757)		-1.568 (3.786)		
Seller dummy: 1 if some seller feedback, 0 if only buyer feedback						-14.193 (7.067)**		-14.193 (7.975)*		
Observations	67	67	93	93	86	86	86	86		
R-squared		.067	.010	.383	.380	.412	.380	.412		

Standard errors are in parentheses: * significant at 10% level, ** significant at the 5% level. Column (a) reports estimated differences in the probability of participating in the second round, while column (b) reports estimates from the same probit regression. Regressions (c)-(h) exclude bids over \$1000 (11 obs.). Regressions (e) - (h) exclude bids by four subjects for whom we were unable to obtain their buyer/seller feedback (7 obs.); Regressions (g) and (h) include clustered standard errors to account for the fact that the 86 observations are based on the choices of 58 subjects.

Table 3: Regression results