Reserve Price Choices of Sellers in Laboratory First Price Auctions: The Role of Experience

Priyodorshi Banerjee

Indian Statistical Institute Calcutta

Shashwat Khare

Indian Statistical Institute Calcutta

P Srikant

Indian Institute of Management Calcutta

September 2015

Abstract

This essay analyzes choices of sellers, each setting a reserve price in a laboratory first price auction with automated equilibrium bidding. Subjects are allowed to practice and gain experience prior to making a single payoff-relevant choice. We found behavior of more experienced sellers to be consistent with benchmark theory: average reserve price for these sellers was independent of the number of bidders in the market and equaled the predicted level. Less experienced sellers however deviated from the theoretical benchmark: on average, they tended to shade reserve price below the predicted level, and positively relate it to the number of bidders.

Keywords: reserve price, first-price auction, experience, seller choice, probability weighting

JEL classification codes: C91, D44

We are grateful to IMT, Ghaziabad for allowing us to conduct sessions, and to Archishman Chakraborty, Sujoy Chakravarty, Ananish Chaudhuri, Sean Crockett, Sanmitra Ghosh, Nityananda Sarkar and Shubhro Sarkar for many helpful comments.

1 Introduction

The variety, volume and value of goods sold through auctions are enormous. Economic agents of many different types regularly participate in auctions, to buy as well as to sell. A large experimental literature has emerged investigating buyer behavior in auctions, and substantial progress has been made in understanding to what extent such behavior comports to theoretical predictions. Little corresponding attention however has been paid to sellers, with few studies documenting or analyzing seller choices.

It is important to fill this lacuna for at least two reasons. First, sellers regularly make critical choices in auctions, such as which format to use, or what reserve price to set. Predictions of price or other variables which ignore this endogeneity therefore stand the risk of being erroneous. And second, the economic theory of auctions makes sharp predictions regarding the choices of payoff maximizing sellers. For example, Myerson (1981) and Riley and Samuelson (1981) show that with independent, private values a seller using a standard form such as a first price auction may set a reserve price exceeding her valuation, opening a possibility of inefficiency, as such a choice can cause some mutually gainful trades to not occur. They further show that this optimal reserve price is independent of the number of bidders in the market. Empirical examination of these predictions is thus required to assess the strength of the theory.

In this paper, we study reserve price choices of subjects in the role of sellers in a first price auction facing automated equilibrium bidding. Our sellers each made a single payoff-relevant choice, but were allowed to gather experience through prior practice for a fixed amount of time on a graphical interface which simulated auctions and kept track of past outcomes and average revenue. We used a between subjects design, and two treatment variables: the first was the number of bidders (2 or 4), and the second the length of the practice phase (5 or 15 minutes), as a proxy for the degree of experience.

We found that benchmark theory was capable of explaining behavior of more experienced sellers, whose average reserve prices were independent of the number of bidders, and indistinguishable from the predicted level. Less experienced sellers however chose reserve prices below the predicted level. Their choices also increased with the number of bidders. We then investigated whether augmenting the benchmark model with probability weighting (see Tversky and Kahneman

1992) was capable of reconciling the latter findings. We found such an approach could rationalize why the reserve price increased with the number of bidders, but not why it was below the predicted level.

To our knowledge, only three prior papers have experimentally analyzed the problem of reserve price choice in auctions. Of these, only Chen, Katuščák and Ozdenoren (2010; *CKO* henceforth) study first-price auctions: Greenleaf (2004; *Gl* henceforth) and Davis, Katok and Kwasnica (2011; *DKK* henceforth) examine English and second-price auctions respectively. *Gl* does not report any conformity of observed reserve prices with benchmark predictions. Support for benchmark predictions in *DKK* is also very weak. *CKO* finds that reserve prices are in line with benchmark predictions, but this result is not directly comparable with ours, as bidding in their environment was not simulated, but generated by volunteer subjects. In terms of results, *DKK* is arguably the closest to us, as they also find that reserve prices chosen increase in the number of bidders and are typically below the benchmark level (we find this tendency too, but only among less experienced bidders). Further, they find weak evidence that experience can reduce the discrepancy between predicted and actual choices. Because there are many differences in design separating these articles from ours, we defer a detailed discussion and comparison till Section 6, after results have been presented.

The rest of the paper is organized as follows. Section 2 summarizes the theoretical benchmark and Section 3 details experimental design and procedure. Sections 4 and 5 present our main results. Section 6 discusses the literature. Section 7 reconsiders the behavior of less experienced sellers, while Section 8 concludes.

2 Theoretical benchmark

A rich and sophisticated theory has emerged for sealed bid, single unit, winner pay auctions when agents are risk-neutral, bidders are symmetric, and have independent private values, through the work of Vickrey (1961), Myerson (1981), Riley and Samuelson (1981) etc: see Krishna (2010) for a textbook presentation. Here we state the key results relevant for our experimental implementation, using the notation of Krishna (2010).

There is one seller and N (potential) bidders. The seller has 0 value for the object. Bidder i has a valuation x_i for the object which is the realization of a random variable X_i. These random variables $X_i, i = 1, \ldots, N$are independently and identically distributed over $[0, w]$ according to the increasing and continuously differentiable distribution function F, with corresponding density function f. The realization of X_i is observed only by bidder i; all else is common knowledge.

We look at a first price auction where the seller is allowed to set a publicly known reserve price. In such an auction bidders independently and simultaneously submit non-negative bids, and the highest bidder wins and pays the seller his bid in exchange for the object, provided this highest bid is no less than the reserve price. A reserve price $r \geq 0$ is a minimum level below which bids are not accepted: if all bids fall short of the reserve price, the object is not sold and the seller earns 0 revenue.

Let Y_1 be the highest order statistic of $N - 1$ of the random variables $(X_i)_{i=1}^N$, and let G and g respectively be the distribution and density functions of Y_1. Given r, the first price auction has a unique symmetric increasing perfect Bayesian Nash equilibrium bid function, and the equilibrium bidding strategy for any bidder with $x \geq r$ is

$$\beta(x) = E[\max(Y_1, r)|Y_1 < x] = r\frac{G(r)}{G(x)} + \frac{1}{G(x)}\int_r^x yg(y)dy \cdots (1)$$

Using (1), the seller's expected revenue can be computed as

$$\Pi(r) = N[r\{1 - F(r)\}G(r) + \int_r^w y\{1 - F(y)\}g(y)dy] \cdots (2)$$

The optimal or revenue maximizing reserve price can be determined using (2). The first order condition implies the optimal reserve price r^* must satisfy

$$r^* = \frac{1}{\lambda(r^*)} \cdots (3)$$

where $\lambda(x) = \frac{f(x)}{1-F(x)}$ is the hazard rate function associated with the distribution function F. When $\lambda(.)$ is increasing, the above condition is also sufficient, and the optimal reserve price is given by (3). A remarkable feature of the result is that the optimal reserve price is independent of the number of bidders. We test this key proposition in this paper.

For our experimental implementation, we assume $w = 100$ and F is uniform (for which λ is increasing). Further we look at two cases, one with $N = 2$ and the other with $N = 4$. This yields $r^* = 50$, using (3), and, using (1),

$$\beta(v) = \begin{cases} \begin{cases} \frac{v^2+r^2}{2v} & if\, N = 2 \\ \frac{3v^4+r^4}{4v^3} & if\, N = 4 \end{cases} & \text{for } v \geqslant r \\ 0 & \text{for } v < r \end{cases}$$

3 Procedure and treatments

We first describe our experimental procedure in detail. In any treatment, a subject assumed the role of a seller selling an object facing automated bidders. Subjects were informed of this fact and also that (i) a seller had 0 value for the object, (ii) each bidder had a privately known valuation for the object which was equally likely to be any integer in $\{0, .., 100\}$, (iii) a seller could set any integer in $\{0, .., 100\}$ as a reserve price, and (iv) every bidder used a particular formula which gave the bid as a function of private valuation and reserve price. They were explained the rules of the first price auction with reserve price, told how many bidders were present (2 or 4), and were given the relevant formula, which happened to be the risk-neutral symmetric increasing perfect Bayesian Nash equilibrium strategy of the auction game, given above.

A subject had to decide what reserve price to set. Once the choice had been entered on a computer interface (programmed using z-tree, Fischbacher 2007), the program would simulate an auction (generate bidder valuations, produce bids and determine outcome) and report revenue earned. The subject would then be paid a show-up fee and twice revenue earned in private.

We chose to inform sellers about the specific formula being used to generate bids, insuring they had all information required to form the expected payoff function. This allowed us to focus on the optimization problem, i.e., the determination of the optimal reserve price, given the payoff function. In principle of course sellers face a much more complex problem, where they first conjecture bidder behavior, then form a payoff function on its basis, and then optimize

Figure 1: Screenshot

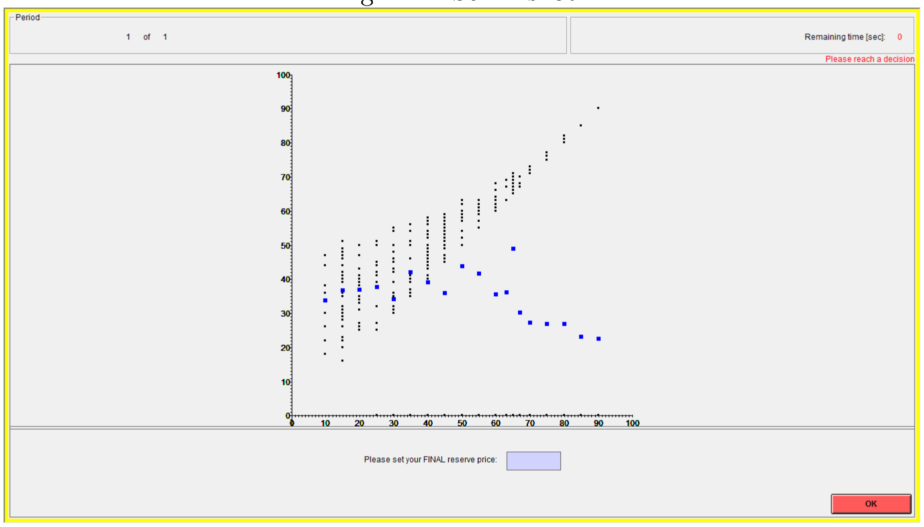

The horizontal axis is reserve price and the vertical axis is revenue. The black and blue dots respectively represent actual and average revenue for a given reserve price.

it. Our procedure, by eliminating the need to conjecture and thereby simplifying the formation aspect of the problem, allows a significant reduction in this complexity.[1]

Before making the final and payoff-relevant reserve price choice, subjects could practice on the same interface. Every time they entered a practice reserve price, the program would simulate an auction, display valuations, bids and outcome, and report revenue earned. It would also display this information graphically for every simulation, along with the history of past outcomes and the average (over all practice choices) revenue function (see Figure 1).

We had four treatments, based on a 2x2 factorial design, where the treatment variables were the number of bidders, a regular proxy for the degree of competition, and the length of the practice session, as a proxy for the degree of experience, or the level of learning. There were either 2 or 4 bidders, and the practice phases were for either 5 or 15 minutes. We denote corresponding

[1] Our procedure thus renders a potentially complex problem, rife with conjectural and strategic considerations, into a comparatively simpler one, and hence allows close focus on one aspect of the laboratory decision problem. We follow *Gl* and *DKK* in adopting this procedure. It can also be found outside the auction literature, such as in the Investment treatments in Oprea (2014), where subjects in the role of investors decide which of two firms to invest in after being made aware of the underlying stochastic processes as well as the strategies of the (automated) firms.

treatments by 2L, 4L, 2H, and 4H, where 2 and 4 stand respectively for two bidders and four bidders, and L and H stand respectively for low experience (5 minutes practice) and high experience (15 minutes practice). The practice phases were of fixed length: subjects could not terminate practice and move to the final choice before the full duration was over.

The experiment was conducted at the Institute for Management Technology, Ghaziabad, a business school near Delhi. There was one session per treatment. Sessions lasted for between 30 and 45 minutes from the time subjects assembled to the time they were paid. Subjects were students pursuing MBA and were recruited through email solicitations and flyers. There were 37, 40, 35, and 39 subjects in treatment 2L, 4L, 2H, and 4H respectively.[2]

After subjects had assembled, each in front of a terminal, instructions (see Appendix) had been handed and read out, numerical exercises done, and questions answered, a demonstration was given of the interface, to familiarize subjects with it. The practice session commenced at the end of the demonstration, after the cessation of which final and payoff-relevant reserve price choice had to be made. Subjects earned ₹ 300 on average.[3]

4 Final reserve prices

In this section we investigate whether sellers' choices of reserve price (final choice) concord with the theoretical prediction described earlier. Our hypotheses are that reserve prices chosen by subjects in every treatment is 50 on average, and that choices do not differ across treatments.

$$Hypotheses: r_{2L} = r_{4L} = r_{2H} = r_{4H} = 50$$

where r is average observed reserve price and subscripts denote treatments.

The following table gives mean and median reserve price for every treatment. It also presents results from the tests of the hypotheses that these means and medians are each equal to 50.

[2]40 subjects had registered for each session, but there were last minute no-shows in all sessions except 4L, leaving us with 151 subjects in total.

[3]The purchasing power parity exchange rate between the Indian Rupee and the US Dollar for 2010 was 16.8 rupees to a dollar according to the Penn World Tables (Heston, Summers and Aten 2012).

Table 1: Means and medians, with tests against 50

	2L	4L	2H	4H
mean	31.73	40.4	45.89	49.23
t-test	<0.0001***	0.0028***	0.1466	0.8142
median	35	42	47	53
Wilcoxon sign rank test	<0.0001***	0.0082***	0.1294	0.9055

The second and fourth rows give two-tailed p-values. *** denotes significance at the 1% level.

We find our hypotheses are upheld in treatments 2H and 4H. Average reserve price in these treatments are indistinguishable from 50. Our hypotheses are rejected however for treatments 2L and 4L, in each of which both averages are observed to be less than 50, with all differences strongly significant.

We now investigate whether average reserve price is the same across treatments. The following table gives the results of tests comparing average reserve prices.

Table 2: Treatment comparison tests

	2L/4L	2L/2H	2L/4H	4L/2H	4L/4H	2H/4H
t-test	0.037 **	0.0005 ***	0.0001 ***	0.1843	0.0499 **	0.436
Mann-Whitney rank sum test	0.0321 **	0.0005 ***	0.0002 ***	0.2599	0.0442 **	0.2334

Entries are two-tailed p-values. ** and *** denote significance respectively at the 5% and 1% levels. A blank indicates insignificant difference.

Table 2 thus shows that our hypothesis is upheld in the comparison between treatments 2H and 4H, i.e., reserve price chosen by high experience sellers does not depend on the number of bidders in the market, no matter whether average is represented by the mean or the median. The hypothesis is also upheld in the comparison between treatments 4L and 2H. However, our hypotheses fail in all other treatment comparisons, with 2L being significantly different from 4L, 2H and 4H, and 4L being significantly different from 4H. In particular, reserve price choices with two bidders differ from those with four bidders when sellers are inexperienced.

Collecting together results from tables 1 and 2, we find thus that our benchmark hypotheses are satisfied for treatments 2H and 4H, i.e., for more experienced subjects. Behavior of these subjects is in accordance with the theoretical predictions presented above. The reserve prices they choose are not dependent on average on the degree of competition as proxied by the number of bidders in the market, and moreover are statistically equal to 50, the predicted level.[4]

Behavior of less experienced bidders however fail to conform to the benchmark theoretical predictions. Their reserve price choices do depend on the number of bidders they face, and further are significantly different from 50. In particular, we find that no matter whether there are 2 or 4 bidders, reserve price chosen is less than 50, i.e., reserve price is shaded below the benchmark level, and also that the chosen reserve price is increasing in the number of bidders.

5 Improving choices with practice

The results presented above are consistent with benchmark theory being accurate in describing behavior of sellers when they face theoretical bidders, provided they have experience. The conjecture that experience causes adherence to the benchmark implies that final choice is based on information acquired during the practice phase, and hence should be forecasteable using data from the phase. It further implies that forecast errors should decrease with more practice on average, and so should be lower when 15 rather than 5 minutes are allowed for practice. Additionally, since there is no reason for the number of bidders present in the market to affect the experience-gathering process in our environment, forecast errors, given length of practice phase, should be independent of the number of bidders. We test the latter two as hypotheses in this section.

We used autoregressive distributed lag (ARDL) models for this purpose (see e.g. Box and Jenkins 1976 and Hamilton 1994). The ARDL model allows the forecast of the value of a variable in some period as a function of values of the same variable, as well as those of other variables, from prior periods.

[4] Kolmogorov-Smirnov and Kruskal-Wallis tests also indicated the samples generated by these two treatments come from the same distribution and population respectively. The Kolmogorov-Smirnov gave an exact p-value of 0.121, and the Kruskal-Wallis chi-squared test with ties gave a p-value of 0.2334.

Consider a subject who recorded a total of T practice choices. We index these from the first to the last by $\{0, 1, ..., T\}$ respectively, denote $\mathtt{T} = \{0, 1, ..., T\}$, and call any $t \in \mathtt{T}$ a period. Let the practice choice recorded in period t, for any $t \in \mathtt{T}$, be denoted p_t. Further, let y_t be the revenue obtained in period t corresponding to p_t. Then a general ARDL specification for forecasting p_{T+1}, the practice choice the subject would have recorded in period $T+1$ had such a period been available, is

$$p_{T+1} = \alpha + \sum_{\substack{t \in \mathtt{T}' \subset \mathtt{T} \\ \mathtt{T}' \neq \emptyset}} \beta_t p_t + \sum_{\substack{t \in \mathtt{T}^* \subset \mathtt{T} \\ \mathtt{T}^* \neq \emptyset}} \gamma_t y_t + u_t \cdots (4)$$

The error in any period is assumed to be an independent draw from some stationary distribution with 0 mean and constant variance.[5] If all the T observations are taken for the purpose of forecast ($\mathtt{T} = \mathtt{T}^{'} = \mathtt{T}^*$), then the model in (4) reduces to

$$p_{T+1} = \alpha + \sum_{t=0}^{T} \beta_t p_t + \sum_{t=0}^{T} \gamma_t y_t + u_t \cdots (5)$$

If we identify the time at which final choice is made with $T + 1$, and assume putative practice choice at this time is indistinguishable from final choice, absolute forecast error is $e = |r - \hat{p}_{T+1}|$, where r denotes the observed final choice, and $\hat{p}_{T+1}$ is the forecast. Denoting average error by e_i, where i indicates treatment, our hypotheses are therefore

$$Hypotheses: \begin{cases} A: \; e_{2L} = e_{4L}, \; e_{2H} = e_{4H} \\ B: \; e_{2L} > e_{2H}, \; e_{4L} > e_{4H} \end{cases}$$

Mean and median errors using the full sample forecasts from (5) are given in the table below. As can be seen, a counterintuitive feature is that errors in the H-treatments are higher than those in the L-treatments.

[5] We do not incorporate bid or valuations information from any period in the model, assuming revenue is sufficient.

Table 3: Means and medians with full sample forecasts

	2L	4L	2H	4H
mean	8.54	11.8	14.4	18.84
median	4.74	6.4	12.62	14.31

The table below presents the results of comparison tests. As indicated by the values in Table 3, our hypotheses are only partly upheld.

Table 4: Mean and median comparison tests

	2L/4L	2L/2H	4L/4H	2H/4H
t-test	0.2459	0.0236 **	0.0405 **	0.1615
Mann-Whitney	0.241	0.002 ***	0.0079 ***	0.4075

Entries are two-tailed p-values. ** and *** denote significance respectively at the 5% and 1% levels. A blank indicates insignificant difference.

Tables 3 and 4 shows the number of bidders indeed makes no difference to average errors (Hypotheses A). However, errors are found to increase with experience, contrary to our hypotheses of a negative relationship.

We probed the data at the level of the subject to understand what may have caused the rejection of Hypotheses B. We found that many subjects in all our treatments chose as final reserve price a number already entered during the practice phase several times, usually in separate blocks of consecutively repeated trials. Additionally, we found in several cases subjects apparently reaching a decision regarding final choice, and then practicing somewhat randomly afterward. This pattern was particularly prevalent in the high experience treatments (see Figure 2, which depicts data for 4 representative subjects, one from each treatment, as illustration). It is possible these subjects would have terminated practice and moved to the final choice at some point prior to the end of the session had that option been available.

Given these features of the data, it is probable that using the full sample of observed practice choices to determine forecast may yield poor results, particularly in the high experience treatments, since the model would then effectively

Figure 2: Practice and final reserve price choices of representative subjects

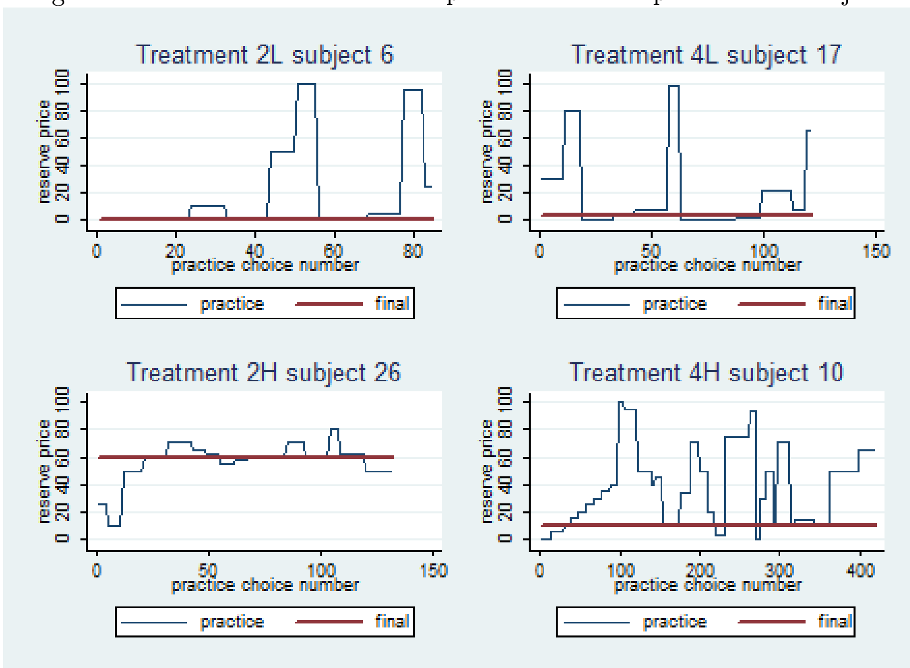

penalize agents who complete learning early. In particular, substantially improved forecasts may be likely for a significant number of subjects if some of the *final* observations were to be dropped to create a truncated sample on the basis of which to make the forecast.

We adopted the following procedure to account for this difficulty. We generated several forecasts for every subject, each using a sub-sample of observations (with the full sample included in the set of sub-samples). The same sub-sample was always used for the p and the y series. We then picked that sub-sample for which error was least. This became the set of observations used to make the forecast for that subject. We then compared average errors across treatments as before.

We picked the set of sub-samples over which to search and optimize in two different ways, since no natural standard was available. In either protocol, (i) all observations from the beginning were always contained in every member of the set, so if (p_t, t_t) was an element of any member for $t > 0$, so was (p_{t-1}, y_{t-1}), and hence any member was of the form $\{(p_0, y_0), ..., (p_t, y_t)\}$, and (ii) if all observations upto t constituted a member with $t < T$, then so did all observations upto $t+1$, i.e., if the set $\{(p_0, y_0), ..., (p_t, y_t)\}$ was a member and $t < T$, then $\{(p_0, y_0), ..., (p_t, y_t), (p_{t+1}, y_{t+1})\}$ was also a member. The protocols differed in the method of choosing the minimum t for which $\{(p_0, y_0), ..., (p_t, y_t)\}$ became a member. In protocol A this minimum was 100 for every subject who had recorded at least 101 practice choices, and was T for the remainder of subjects (T_j for each subject j).[6] For the latter group of subjects therefore, the set of sub-samples was a singleton and only contained the full sample.[7] In protocol P this minimum was the 70^{th} percentile for every subject.

We present the mean and median errors using both protocols in the following table.

Table 5: Mean and median errors

protocol	average	2L	4L	2H	4H
A	mean	7.64	8.32	3.39	1.42
	median	3.03	6.32	0.67	0.01
P	mean	3.03	2.33	1.91	1.37
	median	0.77	0.72	0.26	0.11

Comparison with Table 3 reveals that using either protocol results in a substantial decrease in errors in the H-treatments. The errors decrease in the L-treatments as well, mainly with protocol P. The new magnitudes are entirely in line with Hypotheses B, as surmized. The table below, which presents the results of comparison tests, shows that the differences are mostly statistically significant.

[6]The median numbers of observations in 2L, 4L, 2H and 4H were 85, 62.5, 131, 223 respectively.

[7]The numbers of subjects with 100 or less practice choices in treatments 2L, 4L, 2H and 4H were 19, 25, 2 and 1 respectively.

Table 6: Mean and median comparison tests using both protocols

protocol	test	2L/4L	2L/2H	4L/4H	2H/4H
A	t-test	0.7216	0.0263 **	0.0001 ***	0.0708 *
	Mann-Whitney	0.241	0.0261 **	<0.0001 ***	0.0014 ***
P	t-test	0.4719	0.1354	0.0819 **	0.4746
	Mann-Whitney	0.7136	0.0817 *	0.0052 ***	0.1484

Mann-Whitney entries are two-tailed p-values. t-test entries are two-tailed p-values for the comparisons 2L/4L and 2H/4H, and one-tailed p-values for the comparisons 2L/2H and 4L/4H. *, **, *** denote significance respectively at the 10%, 5% and 1% levels. A blank indicates insignificant difference.

Hypothesis B is now satisfied for the treatments with 4 bidders as errors are lower in 4H than in 4L, no matter which protocol is followed. The support for the hypothesis is weaker in the treatments with 2 bidders, as errors decline significantly with experience unambiguously under protocol A, but ambiguously under protocol P. In the latter case, while the t-test shows no decline, the Mann-Whitney test shows one, albeit weakly. Overall the optimization procedure thus yields substantial support for Hypotheses B.

Its cost manifests itself in the form of weaker support for Hypotheses A. While this hypothesis remains satisfied in the L-treatments, it receives ambiguous support in the H-treatments. Protocol P lends support to the hypothesis and indicates that errors are independent of the number of bidders in the high experience treatments. Protocol A however does not confirm the hypothesis. The Mann-Whitney test rejects the hypothesis strongly, while the t-test does so weakly.

6 Related literature

The approach to whether the degree of experience affects seller choices differentiates us from previous papers. This is because we investigate the impact of the degree of learning or experience through its use as a treatment variable, while others have considered the question within treatment only.

Another key aspect of design which distinguishes the current article from prior essays is the method by which experience is gathered. Here, subjects make a single payoff-relevant choice, and learn or gather experience through practice prior to making that relevant choice. Since in reality an experienced seller is one who is already experienced at the time of taking the critical decision, our approach affords an approximation by enabling the gathering of experience *before* taking payoff-relevant decision. Previous designs however have allowed subjects to learn only through making payoff-relevant choices, i.e., subjects have played a sequence of paying auctions, and gathered experience by playing. This approach by contrast thus has permitted the gathering of experience *while* taking payoff-relevant decisions.

Subjects in *Gl* participated in a sequence of 36 English auctions and set a reserve price in each. Bidding was simulated using the appropriate equilibrium strategy. The number of bidders was either 3 or 6, but this was varied within subjects, and not between subjects, as in our case, and so each subject faced 18 auctions with 3 bidders, and 18 with 6. He does not report average chosen reserve price, but his focus on a learning model with regret and rejoicing suggests conformity with benchmark predictions was poor.

30 first-price auctions were faced in sequence by subjects in *CKO*. Each auction additionally had two randomly matched bidders, who were also subjects. They find some evidence that bidding was in conformity with predictions (see their Table 3). Average reserve price was 39.8, while the risk-neutral prediction was 41.7.[8] Additionally, they find some weak learning effects. To the extent the two designs are comparable, our results on reserve price choices of experienced bidders thus confirm theirs.

As in *Gl*, where the number of bidders was varied within the treatment, the underlying set of experimental parameters was not stable in *CKO*. Specifically, there were two different distribution functions in accordance with which bidder valuations were drawn. One was faced 21 times on average, and the other 9 times. The stability of parameters within treatment, or its absence, is another design feature that separates our paper from others, as *DKK* also allowed vari-

[8]They also had a treatment with second-price auctions and found reserve prices there to be indistinguishable from those in the first-price auction treatment, as predicted. Bidding in the second-price auction however differed sharply from predictions.

ation. Specifically, they varied the number of bidders, like *Gl*, with n either 1, 2, 3, or 4 in some treatments, and either 1, 4, 7, or 10 in others.

A sequence of 60 second-price auctions (explained as English auctions) were faced by subjects in *DKK*, who had automated equilibrium bidding, like *Gl* and the current article.[9] They found little support overall for the hypothesis that benchmark theory adequately describes behavior. They did find some learning effects though, with sellers displaying a slight tendency to set reserve price more in coformity with the benchmark after experiencing more auctions. A key finding in their paper is that reserve price chosen increases with the number of bidders. This holds for the comparison between 2 and 4 bidders as well. They also find that reserve prices are usually below the benchmark value, and can exceed the predicted level for large number of bidders (7 or more).

Since we also find, for less experienced bidders, that reserve prices are below the benchmark level and an increase in the number of bidders leads to an increase in reserve price, our results therefore lend support to theirs. Of course a key difference remains in that we find conformity with the benchmark for more experienced bidders.

A possible reason for why seller behavior deviated from the benchmark in *DKK*, while it did not for our more experienced sellers, could be that sellers in *DKK* faced only 60 auctions, which was not far from the average of the number of practice auctions recorded in our L-treatments. By contrast, our more experienced sellers recorded a mean of over 200 practice auctions, and hence had considerably larger opportunities for learning.

7 Behavior of less experienced sellers

Our results indicate that experience may be a key differentiating factor in determining whether behavior can be expected to conform to the benchmark. The role of experience has been looked at in explaining deviations from benchmark bidding, in experiments with automated selling. There is considerable evidence that more experience is associated with bids deviating less from predicted levels (see surveys of results in Kagel 1995 and Kagel and Levin 2008). Our findings overall therefore suggest that the degree of experience can impact behavior not

[9] Like us, *Gl* and *DKK* also informed their subjects about the formula to be used to generate bids.

just on the buyer side in auction settings, but on the seller side as well.

A question remains as to how to address the behavioral patterns exhibited by less experienced sellers. With respect to the approaches commonly used to understand anomalies emerging from experimental bidding data, explanations based on aspects of preference such as risk (see Waehrer, Harstad and Rothkopf 1998 and Hu, Matthews and Zou 2010) or regret (see e.g. Loomes and Sugden 1982 and Filiz-Osbay and Ozbay 2007) are unlikely to be effective here as long as preferences are invariant to experience. Thus it seems likely that the benefits of increased practice are expressed mainly as lower errors in probability assessment (see e.g. Kahneman and Tversky 1979 and Goeree, Holt and Palfrey 2002) and hence reserve price computation.

We therefore follow *DKK* and use a probability weighting approach to investigate whether incorporating such an aspect into the model used to describe behavior of a less experienced seller can help explain our two key findings: reserve prices chosen by less experienced sellers are (i) below the benchmark level of 50, and (ii) increase with the number of bidders (see Table 1). The general idea of probability weighting is that agents behave as if the probabilities of certain events are different from their given values. Experimental evidence (see e.g. Kahneman and Tversky 1979) suggests subjects often behave as if they were inflating the probabilities of some low probability events and deflating the probabilities of some high probability events.

As in the literature (see Wu and Gonzalez 1996 and *DKK*), we assume there is a weighting function $w(p)$, mapping actual probability of occurence to $[0, 1]$, with (i) $w(0) = 0$ and $w(1) = 1$, (ii) $w'(p) > 0$ for all $p \in [0, 1]$, and (iii) $\exists \hat{p} \in (0, 1)$ such that $w(p) > p$ and $w''(p) < 0$ for $p \in (0, \hat{p})$ and $w(p) < p$ and $w''(p) > 0$ for $p \in (\hat{p}, 1)$.[10] (i) implies that probabilities of events with probability 0 or 1 will not get distorted through weighting. (ii) implies that events with higher probabilities will receive higher weights. (iii) implies low probability events will get overweighted, while high probability events will get underweighted.

Further, sellers are assumed to distort the actual cumulative probabilities p of some events and represent them as $w(p)$ instead when performing expected revenue calculations. We rewrite (2) as

[10] $w(.)$ can also be interpreted as a function representing errors in calculations of probability.

$$\Pi(r) = r[Y_2^{(N)}(r) - Y_1^{(N)}(r)] + \int_r^w z dY_2^{(N)}(z) \cdots (6)$$

where $Y_1^{(N)}(.)$ and $Y_2^{(N)}(.)$ are respectively the distribution functions of the highest and second-highest order statistics of the N random variables $(X_i)_{i=1}^N$. Denoting the derivatives of these two functions by $y_1^{(N)}(.)$ and $y_2^{(N)}(.)$ respectively, we can then rewrite condition (3), governing the optimal reserve price, as

$$r^* = \frac{Y_2^{(N)}(r^*) - Y_1^{(N)}(r^*)}{y_2^{(N)}(r^*) - [y_2^{(N)}(r^*) - y_1^{(N)}(r^*)]} = \frac{1 - F(r^*)}{f(r^*)} \cdots (7)$$

We now incorporate probability weighting into the seller's expected revenue function. (6) and (7) then respectively get modified as follows.

$$\Pi(r) = rw(Y_2^{(N)}(r) - Y_1^{(N)}(r)) + \int_r^w z dw(Y_2^{(N)}(z)) \cdots (8)$$

$$r^* = \frac{w(Y_2^{(N)}(r^*) - Y_1^{(N)}(r^*))}{y_2^{(N)}(r^*)w'(Y_2^{(N)}(r^*)) - [y_2^{(N)}(r^*) - y_1^{(N)}(r^*)]w'(Y_2^{(N)}(r^*) - Y_1^{(N)}(r^*))} \cdots (9)$$

We can now apply (9) to predict optimal reserve price choices, provided we assume a specific form for $w(.)$. We look at two alternative and well-discussed single-parameter families: the Tversky-Kahneman function (Tversky and Kahneman 1992) and the Prelec function (Prelec 1998), which are given sequentially below.

$$w^{TK}(p) = \frac{p^\alpha}{[p^\alpha + (1-p)^\alpha]^{\frac{1}{\alpha}}}, \alpha \in (0,1) \cdots (10)$$

$$w^P(p) = e^{-(-log_e p)^\alpha}, \alpha \in (0,1) \cdots (11)$$

The derivatives of these two functions are as follows.

$$\frac{dw^{TK}(p)}{dp} = \frac{\alpha p^{\alpha-1}}{[p^\alpha + (1-p)^\alpha]^{\frac{1}{\alpha}}} - \frac{\alpha p^\alpha [p^{\alpha-1} - (1-p)^{\alpha-1}]}{[p^\alpha + (1-p)^\alpha]^{(1+\frac{1}{\alpha})}} \cdots (12)$$

$$\frac{dw^P(p)}{dp} = \frac{\alpha(-log_e p)^{\alpha-1}}{p} e^{-(-log_e p)^{\alpha}} \cdots (13)$$

We calculated optimal reserve prices (rounded to the nearest integer) using (9) through (13) for $N = 2$ as well as $N = 4$, for all values of parameter α in $[0.34, 0.94]$ with a grid increment of 0.01. We chose this range because it is within most range estimates found in prior literature. Further, we got a unique solution for all values of α in this range, while the use of the TK function led to non-uniqueness for $\alpha < 1/3$.[11]

DKK performed a similar exercise, with two differences. Firstly, they only studied the TK function, and secondly, they fixed the value of the parameter at 0.65. They found that the predicted reserve price increased relative to the benchmark with the introduction of weighting. They also obtained a monotone increasing relationship between the predicted reserve and the number of bidders, but only for one of the two value distributions they study.

Our results are summarized in Figure 3. We too found that the probability weighted optimal reserve prices are usually higher than the unweighted benchmark. This is true for the TK function, for $N = 2$ as well as $N = 4$. This is also true for the P function for $N = 4$. For $N = 2$, the P function predicts a reserve price which exceeds the benchmark for $\alpha < 0.73$, and equals it for $\alpha \geq 0.73$. The probability weighting approach is thus unable to concur with our finding (i) above, that less experienced sellers tend to choose lower than benchmark reserve prices.

We further found that the Prelec function yields the prediction that sellers should choose higher reserve prices when the number of bidders is greater. This is in consonance with our finding (ii) above. The TK function yields the same prediction, but only for $\alpha < 0.57$. For $\alpha > 0.59$, the predicted reserve price for $N = 2$ exceeds that for $N = 4$.

Overall therefore, the probability weighting framework proved somewhat useful, as it helped understand why sellers may choose higher reserve prices when confronted with more bidders. However, it performed poorly in explaining why sellers might reduce reserve prices relative to the benchmark.

[11] See Ingersoll (2008) on problems of non-monotonicity associated with the TK function for low values of α.

Figure 3: Probability weighted optimal reserve prices

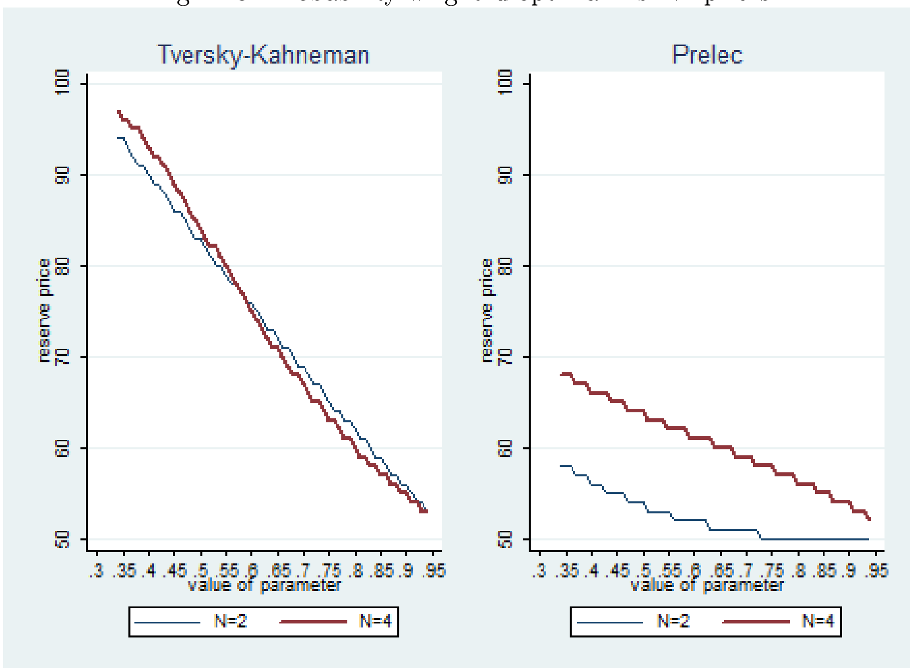

8 Conclusion

This paper has analyzed reserve price choices of laboratory sellers in a first price auction facing automated equilibrium bidding. Subjects were allowed to gather experience through practice prior to making a single payoff-relevant choice. We found close consonance between predicted and observed outcomes for sellers wth more experience, i.e., those who had faced a longer practice phase. Specifically, the reserve price choices of these sellers were at the predicted level, and did not vary with the number of bidders in the market.

The behavior of less experienced sellers deviated markedly from prediction. They tended to shade reserve price below the benchmark level and increase it in the presence of more bidders. A similar pattern of choices has also been reported by *DKK* in the context of a second price auction. Further research may be required here to establish whether these tendencies are of a general nature. A probability weighted variant of the standard model was determined to be capable of explaining why the reserve price may increase with the number of bidders, but not why it may fall below the predicted level.

It is well-known that experience can reduce discrepancy between observed and predicted bidding behavior in the laboratory. Our findings indicate that the degree of experience can perform a similar role for seller choices as well. In aiming to contribute to a better understanding of behavior in auction environments, this article has approached the problem from the side of the seller, or the designer of the mechanism. While auctions and similar mechanisms fill roles of great importance in the economy, there is a paucity of experimental research on seller behavior which needs to be removed to achieve improved analysis and prediction of price formation, trade flows etc. In this regard, questions of which kind of mechanism sellers choose are of at least as much centrality as questions of how sellers choose parameters defining particular mechanisms, which has been the focus of this paper. Questions of choice between mechanisms are left for future research.

References

Box, G. and G. Jenkins (1976): *Time Series Analysis: Forecasting and Control.* San Francisco: Holden-Day.

Chen, Y., P. Katuščák and E. Ozdenoren (2010): "Sealed-Bid Auctions with Human Auctioneers," working paper, University of Michigan.

Davis, A., E. Katok and A. Kwasnica (2011): "Do Auctioneers Pick Optimal Reserve Prices?" *Management Science* , 57, 177 – 92.

Filiz-Osbay, E. and E. Ozbay (2007): "Auctions with Anticipated Regret: Theory and Experiment," *American Economic Review*, 97, 1407 – 18.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10, 171 – 8.

Goeree, J., C. Holt and T. Palfrey (2002): "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104, 247 – 72.

Greenleaf, E. (2004): "Reserves, Regret, and Rejoicing in Open English Auctions," *Journal of Consumer Research*, 31, 264 – 73.

Hamilton, J. (1994): *Time Series Analysis.* Princeton: Princeton University Press.

Heston, A., R. Summers and B. Aten (2012): "Penn World Table Version 7.1," Center for International Comparisons of Production, Income and Prices, University of Pennsylvania.

Hu, A., S. Matthews and L. Zou (2010): "Risk-Aversion and Optimal Reserve Prices in First- and Second-Price Auctions," *Journal of Economic Theory*, 145, 1188 – 1202.

Ingersoll, J. (2008): "Non-Monotonicity of the Tversky-Kahneman Probability-Weighting Function: A Cautionary Note," *European Financial Management*, 14, 385 – 90.

Kagel, J. (1995): "Auctions: Survey of Experimental Research," in A. Roth and J. Kagel (eds.) *Handbook of Experimental Economics.* Princeton: Princeton University Press.

Kagel, J. and D. Levin (2008): "Auctions: A Survey of Experimental Research 1995 - 2008," wp, Ohio State University.

Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263 – 291.

Krishna, V. (2010): *Auction Theory, 2^{nd} edition.* Oxford: Academic Press.

Loomes, G. and R. Sugden (1982): "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal*, 92, 805 – 24.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58 – 73.

Oprea, R. (2014): "Survial versus Profit Maximization in a Dynamic Stochastic Experiment," *Econometrica*, 82, 2225 – 55.

Prelec, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497 – 527.

Riley, J. and W. Samuelson (1981): "Optimal Auctions," *American Economic Review*, 71, 381 – 92.

Tversky, A. and D. Kahneman (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297 – 323.

Vickrey, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8 – 37.

Waehrer, K., R. Harstad and M. Rothkopf (1998): "Auction Form Preferences of Risk-Averse Bid Takers," *RAND Journal of Economics*, 29, 179 – 92.

Wu, G. and R. Gonzalez (1996): "Curvature of the Probability Weighting Function," *Management Science* , 42, 1676 – 90.

Appendix: Instructions

Instructions for 2L and 2H:

You have an object which you have decided to sell using a *first-price auction*. In a first-price auction, all interested buyers submit a bid each. If all bids fall short of the *reserve price* (a minimum integer sale price that you set beforehand), the object remains unsold and the revenue from the auction is 0 and you earn 0. Otherwise, if at least one bid exceeds the reserve price, the highest bidder obtains the object and pays whatever he or she bid. This is the revenue from the auction and you earn twice this amount.

You have some information available: (1) There are two buyers interested in obtaining the object. (2) Each of them has a secret valuation for the object which is equally likely to be any integer between (and including) 0 and 100. (3) Bidders use the following rule to determine the bid:

If r be the reserve price set by you (which could be anything between and including 0 and 100), then a bidder whose valuation is v (something between and including 0 and 100) bids

$$b(v; r) = \begin{cases} 0 : & v < r \\ \frac{v^2+r^2}{2v} : & v \geqslant r \end{cases}$$

For example, suppose the two bidders have valuations 25 and 65. Then, if you set $r = 0$, revenue is 33. If you set $r = 25$, revenue is 37. If you set $r = 50$, revenue is 52. If you set $r = 75$, revenue is 0, as the good is not sold.

You can set any reserve price you want, including 0 (no reserve price) and 100 (maximal reserve). The reserve price has to be an integer.

You will shortly face a decision screen. There you will have to set the reserve price. Once you submit your selection, the computer will simulate bidding. It will draw a valuation for each bidder (a number between and including 0 and 100, randomly drawn), and then generate bids according to the rule given above. If both bidders' valuations fall short of the reserve price, you will earn 0. Otherwise you will earn twice the amount of the highest bid. What reserve price do you wish to set?

To help you decide, before you have to make the final choice of reserve price, you can practice for up to **[5/15]** minutes. In the practice examples on the computer screen, you can set a reserve price. Once you hit submit, the program draws buyer valuations randomly and runs an auction for you. You find out what valuations bidders had, what they bid, whether the object was sold or not, and the price obtained if sold.

Instructions for 4L and 4H:

You have an object which you have decided to sell using a *first-price auction*. In a first-price auction, all interested buyers submit a bid each. If all bids fall short of the *reserve price* (a minimum integer sale price that you set beforehand), the object remains unsold and the revenue from the auction is 0 and you earn 0. Otherwise, if at least one bid exceeds the reserve price, the highest bidder obtains the object and pays whatever he or she bid. This is the revenue from the auction and you earn twice this amount.

You have some information available: (1) There are four buyers interested in obtaining the object. (2) Each of them has a secret valuation for the object which is equally likely to be any integer between (and including) 0 and 100. (3) Bidders use the following rule to determine the bid:

If r be the reserve price set by you (which could be anything between and including 0 and 100), then a bidder whose valuation is v (something between and including 0 and 100) bids

$$b(v;r) = \begin{cases} 0 : & v < r \\ \frac{3v^4+r^4}{4v^3} : & v \geqslant r \end{cases}$$

For example, suppose the four bidders have valuations 5, 25, 45, and 65. Then, if you set $r = 0$, revenue is 49. If you set $r = 25$, revenue is 49. If you set $r = 50$, revenue is 54. If you set $r = 75$, revenue is 0, as the good is not sold.

You can set any reserve price you want, including 0 (no reserve price) and 100 (maximal reserve). The reserve price has to be an integer.

You will shortly face a decision screen. There you will have to set the reserve price. Once you submit your selection, the computer will simulate bidding. It will draw a valuation for each bidder (a number between and including 0 and 100, randomly drawn), and then generate bids according to the rule given above. If all bidders' valuations fall short of the reserve price, you will earn 0. Otherwise you will earn twice the amount of the highest bid. What reserve price do you wish to set?

To help you decide, before you have to make the final choice of reserve price, you can practice for up to **[5/15]** minutes. In the practice examples on the computer screen, you can set a reserve price. Once you hit submit, the program draws buyer valuations randomly and runs an auction for you. You find out what valuations bidders had, what they bid, whether the object was sold or not, and the price obtained if sold.