

Bidder Preferences among Auction Institutions*

Radosveta Ivanova-Stenzel[†]
Humboldt University

Timothy C. Salmon[‡]
Florida State University

April 16, 2003

Abstract

This study examines bidder preferences between alternative auction institutions. In particular we seek to experimentally characterize the degree to which bidders prefer an ascending auction over a sealed bid auction. We find very strong *ceteris paribus* preferences for the ascending institution with bidders choosing it overwhelmingly often when entry prices for the two auctions are the same. When the entry prices of the two auctions differ, many subjects can be shown to be willing to pay far more to enter the ascending auction than is explainable by their risk attitudes when accounting for their expectations about the risk preferences of their opponents.

JEL Codes: C91, D44

Key Words: bidder preferences, private values, sealed bid auctions, ascending auctions

1 Introduction

Auctions have become a pervasive method of exchange in the on-line world as each day thousands of auctions take place online and this trade volume totals to billions of dollars worth of goods per year¹. This large volume of auction transactions implies the existence of a large number of sellers competing for buyers. The obvious implication is that the competition among the sellers for the pool of potential buyers can be fierce and any competitive edge a seller can find could be important. One such competitive edge a seller might exploit is using an auction design that attracts bidders away from their competitors.

When designing a real auction or modeling a theoretical one, this entry decision of prospective bidders is rarely considered. Most auction analysis is performed assuming that a certain number of bidders will participate for certain or perhaps that the number of bidders is unknown and randomly determined. It should be clear however, that the most crucial part of a successful auction is encouraging as many bidders as possible to participate. In general this should be expected to

*The authors would like to thank various participants in seminars at Florida State University, University College London, Max Planck Institute for Research into Economic Systems, and ESA Meeting 2002 for comments and advice. We thank Mark Isaac, Sabine Kröger and Frank Heiland in particular for many helpful conversations and suggestions. Financial Support from the *Deutsche Forschungsgemeinschaft*, SFB 373 (“Quantifikation und Simulation ökonomischer Prozesse”), Humboldt-University Berlin, is gratefully acknowledged. Complete sets of instructions for the experiments in this paper are available upon request to the authors.

[†]Humboldt-University of Berlin, Department of Economics, Spandauer Str. 1, D-10178 Berlin, Germany, e-mail: ivanova@wiwi.hu-berlin.de

[‡]Department of Economics, Florida State University, Tallahassee, FL, 32306-2180, tsalmon@garnet.acns.fsu.edu. Phone: 850-644-7207 Fax: 850-644-4535.

¹For a survey of the on-line auction activity see Lucking-Reiley (2000).

have a positive effect on revenue (at least in non-common value environments) and in certain types of auctions it may help to combat the possibility of bidder collusion. Since there are typically competing auctions available for similar goods or even outside options that bidders can pursue when auctions are for unique goods, it is important to understand how the aspects of an auction format can effect a person's decision to enter.

Consider a bidder who is faced with the choice of entering one of two auctions for similar or even identical objects. How does this bidder make his decision of which auction to enter? The obvious answer is that the bidder will enter into the auction that maximizes their expected utility so long as that expected utility is greater than some reservation value. The real question, then, is how are these expected utilities constructed? Profit from participating in the auction is an obvious argument. There are also a number of environmental considerations that might effect this decision that would be difficult to account for precisely such as the reputation and trustworthiness of the auctioneer, quality of the advertisement for the auction and things of this nature. It is also possible though that the format of the auction itself can have an impact on the preferences of the bidder. This latter point will be the issue of this study. The particular focus will be looking at bidder preferences between the two most common standard auction formats used in the field; the sealed bid first price (will be abbreviated as just the "sealed bid" auction) and the ascending or English auction. The other reason we are interested in comparing these two auction formats rather than between the ascending and second price or first price and descending is that it seems reasonable to expect bidders to have preferences between the two due to the strategic differences between them. Such differences lead to substantial differences in terms of the difficulty of deciding how to bid and also in the possibility that an outcome leads to a bidder experiencing some form of regret. Regret may occur in a first price auction when a bidder loses to a bid that is below their value as they may think that if only they had bid higher, they could have won. While such feelings of regret may be considered "irrational," that does not preclude their existence. Such a scenario should not reasonably occur in an ascending auction.

If one considers the situation of a bidder choosing between two auctions that differ only by whether the auction is being conducted according to an ascending or sealed bid format, it is not immediately obvious which would be the most preferred even assuming a standard symmetric independent private values environment. Were the bidders risk neutral, then of course revenue equivalence would hold and the bidders would be indifferent. If the bidders are risk averse, the situation is more complex. As shown in Milgrom and Weber (1982), RA bidders will bid higher in a sealed bid auction than risk neutral bidders and therefore expect to make a lower surplus than if they participated in an ascending auction where they will bid identically to risk neutral bidders. That would imply a preference for the ascending since the surplus if they win is higher. On the other hand, the surplus in the ascending auction is more variable than the surplus in the sealed bid and a RA decision maker dislikes a variable outcome causing the sealed bid to be more attractive. These conflicting attractions lead to a lack of a general conclusion about which format a RA bidder would prefer. Matthews (1987) presents a solution to this dilemma by showing that if a bidder possesses decreasing absolute risk averse (DARA) preferences, they will prefer the ascending, increasing absolute risk averse (IARA) the sealed bid and constant absolute risk averse (CARA) preferences will lead to indifference. These results will serve as a useful back-drop to the analysis below.

There are several other empirical and theoretical papers that look at the issue of endogenous entry choices in regard to auctions, Bajari and Hortacsu (2000), Lucking-Reiley (1999), Harstad (1990), Engelbrecht-Wiggans (1993), Levin and Smith (1994), McAfee (1993) and McAfee and McMillan (1987) to name a few. These papers are examining entry decisions in contexts quite different than we are concerned with here. In general, they are looking at either the decision of

whether to enter an auction or not, or at the choice of which auction to enter based upon the entry price or reserve price being the main or even only characteristic upon which the auctions differ. Our point of concern is to look at the entry decision when the main characteristic that distinguishes two auctions is the format being used to conduct the auction and examine the preferences that underlie those decisions. We note that in Klemperer (2002), the author discusses such preferences and proposes that ascending auctions can actually discourage disadvantaged or weak bidders from entering. We only consider symmetric bidders so the relative strengths of bidders is not an issue. Klemperer’s claim is investigated in Goeree and Offerman (2002) where they find that if weak bidders are allowed to choose sequentially among themselves between entering an auction or playing a fixed lottery, they tend to enter first price auctions more than ascending.

The most closely related prior study is Ivanova-Stenzel and Sonsino (2001). In this paper the authors conduct an experiment comparing the outcomes in a “one-bid” auction, i.e. a standard sealed bid first price auction, to a “two-bid” auction, a modified version where subjects may submit two bids: a high bid and a low one in which the winner pays his low bid if this was higher than all other bids or the high one if other bidders have bid above his low bid but below his high. One of the issues they examined was which auction format the bidders preferred. They accomplished this by allowing bidders in one part of the experiment to repeatedly choose between participating in a one-bid auction competing with one other bidder or a two-bid auction against one other player. The results showed a strong preference for the two-bid auction.

The one shortcoming in the methodology in Ivanova-Stenzel and Sonsino (2001) as a means of measuring bidder preferences is that it only allowed for comparing auctions on what might be called a *ceteris paribus* basis. That is, with all things being equal, which auction would the bidders prefer? If one wishes to use the results from a study of this sort to argue in favor of adopting a new design, it would be useful to take the investigation a step further and measure the intensity of this preference or find an answer to the question “how much are bidders willing to pay for their preferred auction format?” The current study will extend the methodology in Ivanova-Stenzel and Sonsino (2001) to explore this additional question in the context of bidder preferences between the one-bid version of the first price auction and the standard ascending or English auction.

The outcome of these experiments will show the existence of very strong preferences for the ascending auction. When subjects are given a choice between the first price and the ascending auction on a *ceteris paribus* basis, they overwhelmingly choose the ascending auction. As expected, the surplus achieved by the winners in the ascending auctions far exceeds the surplus obtained in the sealed bid auctions. When the subjects are asked to pay to get into an ascending auction, however, some are evidencing a willingness to pay that far exceeds what appear to be the most appropriate theoretical predictions.

Section 2 of the paper will explain the design and conduct of the experiments. Section 3 contains the analysis of results and section 4 will conclude. There is also an appendix to the paper which contains some technical details relating to a few of the computations made in the analysis.

2 Design of Experiments

The experiments for this study were divided into two distinct phases. In the first phase, the learning or training phase, the subjects played both a sealed bid and an ascending auction for 10 consecutive rounds with each auction being conducted with two bidders. Each round consisted of the bidder playing one of each auction type with the same value. The bidders possessed private values which were randomly drawn from the set $V = \{0, 1, 2, \dots, 99, 100\}$ with all values $v_i \in V$ being equally likely. After seeing their value draw, subjects were asked to submit their bid to be used in the

sealed-bid auction. Subjects could choose integer bids between 0 and 150, which did allow them to overbid their highest possible value. All values were denoted in a fictitious currency termed ECU for Experimental Currency Unit. Before bidders were informed about the results of the sealed bid auction, they participated in a Japanese or ascending clock auction². At the end of the round the bidders observed a feedback-window specifying the results from both auction formats indicating whether or not they won, the price paid by the winner in each, the private value of the buyer, their own profit in the auction and their total profit in the current round. They were not given cumulative profit numbers, only numbers from the current round. There were 10 participants in each experiment session and in each round, subjects were randomly re-paired to bid against a new opponent. In a given round, subjects competed against the same opponent in both the ascending and sealed bid auctions.

The idea for this phase was to allow subjects time to figure out how to bid in these auctions as well as to understand the formats well enough for them to form preferences between them. The reason for having subjects play both auctions with the same value was an attempt to minimize any negative impressions a bidder might receive about an auction format due to a random series of bad draws on one format while getting good draws in the other. In four out of the six sessions, at the end of the learning phase there was a summary screen detailing the average profit achieved by the winner across both auction types. This screen was eliminated in two of the sessions. The purpose of including this information screen was to aid subjects in learning about the average actual profitability for participating in the two mechanisms and it was removed in the two sessions to determine if it had any effect.

In the second phase of the experiment, the preference assessment phase, the participants played an extended auction-selection game for 30 rounds. In a single round of this phase, bidders were asked to choose to enter either an ascending or sealed bid auction, knowing that regardless of which they chose they would be competing against one other bidder³. In each round, both auction formats had an entry price attached to choosing them which the bidder had to pay regardless of whether or not they won. This choice of which auction to enter was made before observing their realized private value for the auction. In the first 10 rounds of this phase, the entry prices for both auction formats (sealed bid and ascending) were the same (1.40 ECU). The preferred auction design was identified for each individual as the one they chose in at least 5 rounds of these 10 rounds. In the remaining 20 rounds, the entry price for the preferred auction format was varied in each round using a grid consisting of entry prices ranging from .7 to 14 ECU with an increment of .7. This range was decided upon based on two pilot sessions to identify a reasonable range of values that yielded fine enough resolution to identify bidder preferences while still being wide enough to allow the observation of the maximum willingness to pay of most subjects. To avoid the possibility that the subjects would see the experiment as a simple grid exercise and become bored or disinterested, the grid was not presented in an ascending order, rather the order was randomized. To make the grid structure even less apparent, we added an ϵ to each element of the grid, where ϵ is a random variable normally distributed on the range $(-.05, .05)$.⁴ There are other ways of eliciting a subject's willingness to pay such as running a second-price sealed bid auction for the right to

²The price started at 0 and began increasing at the rate of 1 ECU every 2 seconds. The auction concluded when one of the bidders clicked on a button to indicate they were withdrawing from the auction with the remaining bidder winning the auction at the price the first bidder dropped out at.

³To guarantee that an even number of subjects participated in each mechanism, only 9 out of the 10 participants were able to choose an auction type in each round. The 10-th participant was automatically assigned to whichever auction type had an odd number of people selecting it. The identity of the "10-th" player was changed in each round, so that each subject played the balancing role once every ten rounds or three times among the 30 rounds.

⁴The actual entry price order all subjects saw was {8.39, 2.10, 0.70, 4.92, 12.61, 1.42, 6.27, 4.20, 9.79, 11.15, 13.27, 5.59, 11.90, 9.07, 2.80, 10.49, 7.01, 3.50, 13.98, 7.74}.

enter each format that some might be inclined to find more straightforward. There are, however, two main advantages of our approach. The first is that it allows us to conduct consistency checks on the elicited WTP through using this randomized grid. We are also able to use our results to specifically test the effect of entry prices on entry which is in itself an important issue for applied auction design.

After subjects made their choices concerning the auction type, the round was played with 20% probability. This was a session wide determination, not specific to any particular player. At the end of each auction that was actually conducted, subjects were informed whether or not they won the auction, the price paid by the winner, the entry price they paid, their private (reselling) value and their own payoff in the current round. Note that an entry price was only charged to a subject if the auction was conducted. In the rounds in which no auction was conducted, no entry prices were charged.

All experimental sessions were conducted with the use of a computer based software system, created with z-Tree (Fischbacher (1998)). All experiments were conducted at Humboldt-University, Berlin and most participants were students of economics or business administration. They had been invited by leaflets to participate in an experiment announced to last about two hours which turned out to be approximately accurate. The conversion rate of the ECU earned by each subject into cash was: 1 ECU = 0.04 EUR or about US\$0.035 (at the time the experiment was conducted). In addition, subjects were paid a fixed participation fee of 2.50 EUR or about \$2.20. Subjects' total earnings ranged between 7.35 EUR (\$6.47) to 26.60 EUR (\$23.41) with a mean of 16.83 EUR (\$14.81).

3 Results

There are three basic questions that arise from these experiments which are 1. What did people prefer?, 2. How much were they willing to pay for that preference? and 3. What can account for that willingness to pay? Each of these will be answered in order.

3.1 Which institution did subjects prefer?

When the entry prices for the two auction institutions were equal, subjects overwhelmingly preferred the ascending auction. There was only one subject choosing the sealed bid exclusively while 39 chose the ascending exclusively. In fact only 5 out of the 60 subjects chose the sealed bid more often than the ascending. There were 9 choices made by each subject in this phase of the auction and the average number of times the ascending was chosen was 7.87 with a median of 9 while the numbers were 1.13 and 0 for the sealed bid.

3.2 How much were they willing to pay?

There are several different ways to look at how much the subjects were willing to pay for their most preferred auction. Since only five subjects evidenced a preference for the sealed bid, we will be ignoring willingness to pay for it and will concentrate on analyzing willingness to pay for the ascending.

One characterization of this would be to construct a pseudo demand curve showing how many people are willing to pay each possible entry price for the ascending auction. Such a construct is shown in figure 1. This pseudo demand curve exhibits the standard characteristics of a normal demand curve. It shows that as the price rises, fewer people are willing to pay for the ascending auction. It is a bit jagged, however, indicating that there are some bidders making choices that




Figure 1: Pseudo demand curve indicating the number of subjects choosing the ascending auction format at each entry price.

are not purely monotonic. For example, a subject may have declined to pay a price of 2.1 for the ascending auction, but agreed to pay a price of 2.8 or 3.5. This is also partially an artifact of not allowing one person to choose at each price. For example, at one point we may observe 35 subjects willing to pay a price of 2.1 and 36 willing to pay 2.8 because the person held out at 2.1 was willing to pay that price and higher, then at the price of 2.8, he was allowed to choose and accepted that price while one of the people not willing to pay at 2.1 were now held out.

Theoretically, expected profit would have been the same between both institutions had the subjects bid as risk neutral expected utility maximizers. In reality, average profit to winners from the ascending auctions was 37.02 and 19.78 for sealed bid⁵. This implies a difference of 17.24. If subjects expected to win half the time, this implies an expected profit differential of $17.24/2 = 8.62$. The theoretical average profits, or the profits that would have been obtained had bidders bid according to risk neutral Nash equilibrium bidding strategies, were 37.40 and 33.63 (37.64 and 33.28 for phase 1)⁶.

If bidders were concerned only with expected profitability, they should have been willing to pay up to 8.62 on top of the entry price for the sealed bid auction, 1.4, or a total entry price of 10.02 to be in the ascending versus the sealed bid auction. The horizontal dotted line in the demand curve graph represents this cut-off price. At this line, expected profits from the two formats are equal. Below this line, the subjects will be making less on average from participating in the sealed bid auction. This indicates that fewer than 10 out of 60 people were willing to pay more than this

⁵These numbers are average profits to winners from all auctions excluding any entry price payments, if we just look at phase 1 numbers, the results are 37.02 and 19.32.

⁶The theoretically expected profits for the ascending auction are a bit higher than they should be. This was due to a degree of correlation somehow getting into the values for the two bidders. It is unlikely this was detectable by the subjects and should have had little impact on their choice behavior. The most likely impact would have been to increase the level of the price that leads to a switch-over from the ascending auction to the sealed bid, but results will show that experiential variables such as this have no effect on the switch-over price.

to participate in the ascending auction, most were only willing to pay far less.

To get a more precise picture of the willingness to pay of the subjects requires estimating their switch-over price or the price at which the subjects would switch from choosing the ascending auction to the sealed bid. The best way to understand this process is to visualize the price of the two auctions starting off the same, subjects choosing the ascending, and then the price of the ascending slowly rising. At some point, the subject will be choosing the sealed bid. We want to find the price that best describes the point at which each subject finds the ascending auction no longer worthwhile to choose. To do that we will propose that there exists some price differential, δ_i , that will lead to the subject to switch from choosing the ascending to the sealed bid. They will bid on the ascending so long as the actual price differential $p_t(A) - p_t(SB)$ is less than this threshold. If we make a reasonable specification of probabilistic, rather than deterministic choice, δ_i can be estimated for each subject by finding the δ_i that solves the following:

$$\max_{\delta_i} \sum_{t=1}^T [2\rho^*(t) - \sum_{j=1}^2 \rho_j^2] \quad (1)$$

$$\rho(c_t) = \begin{cases} \frac{e^{\delta_i - (p_t(A) - p_t(SB))}}{1 + e^{\delta_i - (p_t(A) - p_t(SB))}} & \text{if } c_t = A \\ 1 - \frac{e^{\delta_i - (p_t(A) - p_t(SB))}}{1 + e^{\delta_i - (p_t(A) - p_t(SB))}} & \text{if } c_t = SB \end{cases}$$

This specification is essentially minimizing the mean squared deviation of the predictions. Since many of the predicted probabilities will be close to 0 and 1, this specification should be expected to be superior to a standard log-likelihood specification⁷.

The δ_i for each subject represents the price differential that will make them prefer to choose the sealed bid institution. To obtain the actual price for the ascending auction at which this switch-over should be observed we must add 1.4 as this is the static price for the sealed bid auction. Figure 2 contains a histogram summarizing the estimated switchover prices for all subjects. Note that the few negative observations in the graph represent those bidders who preferred the sealed bid auction in periods 1-10. The average price that lead people to switch from the ascending to the sealed bid is 5.95 (6.61 considering only those evidencing a preference for the ascending auction in periods 1-10) while the average profit difference between the two institutions was 8.62 leading to a implied switch-over price of 10.02 if the subjects were only concerned about average profits. This leads to the same conclusion that was implied by the pseudo demand curve above which is that subjects were willing to pay significantly less than would be implied by the expected payoff differential alone. The expected payoff differential is of course a very crude measure of what a bidder should be willing to pay and will be improved on in the next section.

Based upon the demand curve seen in figure 1 it is clear that the subjects were not displaying purely monotonic preferences. It is therefore important to get some characterization of the degree to which the choices of the subjects were consistent and purposeful instead of random. If bidders were perfectly consistent in their choices, we would expect to see one of two patterns to their choices. One is maintaining a constant choice throughout the second phase, such as a choice path consisting of all A's. A second would be choosing the ascending auction up to some price and then switching once and for all to the sealed bid, which we might represent as an A-SB path. It might also be reasonable to find that preferences are somewhat probabilistic and when the two auctions are roughly equal in expected utility the subjects choose randomly. This would lead to

⁷See Selten (1998) and Friedman (1983) for a discussion of the problems of using a log-likelihood function for this sort of a problem.




Figure 2: Histogram of estimated prices at which subjects chose to switch from ascending to sealed bid.

a reasonable expectation that say for a price of 2.8 a bidder is observed choosing the ascending, 3.5 sealed bid, switching back to the ascending at the next price and sealed bid at the one above that and then staying constant for the rest. This would be three switches and can be described as an A-SB-A-SB path. In fact the median number of switches was 3 (average 2.92, mode 1). There are of course many other consistency measures one might be interested in examining. We have investigated several additional measures and all indicate that in most cases, the decisions of the subjects appear purposeful and consistent with a probabilistic choice model⁸.

3.3 What can account for the observed willingness to pay?

There are a couple of obvious things to check to see if they can account for the observed differences in willingness to pay. One might think that preferences over auction institutions would be formed by outcomes from the learning phase of the experiment. Subjects who ended up with higher profits in the ascending auction may have been more willing to pay for it or perhaps those who won more often in the sealed bid auction would be less willing to pay for the ascending auction and so forth. While the design of the experiment attempts to control for these effects, there is no guarantee that the attempt was successful. Another possibility intentionally not controlled for in the design is that the treatment of whether or not the subjects saw the summary statistics from the results in the training phase could lead to different behavior. While it is unclear how this might effect the outcome, it seems possible that it could.

Table 1 contains the results of an analysis of whether or not these experience variables based on the training phase have any influence on the observed switchover price, $\delta_i + 1.4$. It contains the results from two different approaches to looking at the issue. The first involves looking at the

⁸Further details of this analysis are contained in the working paper version of this study that is freely available from the authors upon request. These details have been omitted here to conserve space.

	Correlation Coefficients		Regression Coefficients	
	Full Sample	Partial	Full Sample	Partial
Treatment	-.28* (.03)	-.13 (.32)	-3.94* (.05)	-1.05 (.43)
Avg SB Profit	.06 (.68)	.07 (.59)	.11 (.71)	.08 (.67)
% SB wins	.08 (.55)	.09 (.51)	-.57 (.93)	2.03 (.63)
Avg A Profit	.14 (.30)	.05 (.73)	.15 (.55)	.01 (.93)
% A wins	.01 (.45)	.04 (.78)	-2.25 (.83)	-2.10 (.76)
Session SB Profit	-.15 (.25)	-.04 (.80)	-.82 (.35)	.29 (.61)
Session A Profit	-.04 (.78)	-.01 (.48)	1.70 (.21)	-.58 (.54)
Constant	-	-	-15.32 (.45)	13.67 (.34)
			$R^2 = .13, F = 1.15$ $R^2 = .04, F = .31$	

Table 1: This table contains correlation coefficients of the correlation between $\delta_i + 1.4$ and various experience variables from the learning phase. The table also contains standard OLS regression coefficients of regressions with these experience variables as the independent variables and $\delta_i + 1.4$ as the dependent. Results for both are presented from the full sample and with the outliers removed. Numbers in () are the p-values of the parameters.

raw correlation coefficients between $\delta_i + 1.4$ and variables such as average profit in the sealed bid auctions, percent of times the individual won a sealed bid auction and so forth. The second set of results are from a simple regression with $\delta_i + 1.4$ as the dependent variable and these variables as the independent ones. The p-values of the coefficients are listed next to them in (). Examining both the correlation coefficient results and the regression results in the Full Sample cases shows that the only variable that might be significantly related to the switchover price is the treatment variable. This is somewhat misleading, though, as by chance the one bidder shown in figure 2 to have a $\delta_i + 1.4 < -10$ and the four bidders found to have $\delta_i + 1.4 > 20$ were in separate treatments. These five bidders are certainly outliers from the rest of the population. If they are dropped from the sample and the correlation and regression coefficients recomputed, the significance of this variable disappears and all other variables remain insignificant as shown in the Partial Sample columns in the table.

We present both sets of results because while the regression results may be the most natural test of whether or not these experience variables are related to the switchover price, these variables are highly correlated with each other making the interpretation of the coefficients somewhat suspect. The analysis based on the correlation coefficients, however, does not suffer from this problem and it shows the exact same results.

These results show that the heterogeneity in observed willingness to pay does not seem to be derived from the fact that different bidders experienced different results during the training phase. That is actually an encouraging result as theoretically, these variables should not have any impact on these decisions and the experiment was designed to insure they would have no effect. This leaves us with the likelihood that the observed differences were based upon unobserved heterogeneity in the preferences of the bidders rather than upon the observed heterogeneity of their experiences. One possible source of this preference heterogeneity is in the risk preferences of the subjects.

A casual examination of the bids observed in the sealed bid auctions, reveals the standard pattern observed in most sealed bid auction experiments which is bids far in excess of those predicted by the risk neutral Nash equilibrium. This might suggest that bidders possessed some form of risk averse preferences and this risk aversion may have influenced their choice of auction formats. Since the predominant preference was for the ascending auction and Matthews (1987) shows that

bidders possessing DARA preferences will prefer the ascending auction, this would lend support to the assumption that the type of risk aversion bidders possessed satisfies DARA. We will make one further assumption and assume that the preferences of the subjects can be described by the constant relative risk averse (CRRA) utility function, $u(x) = x^\alpha$, as this satisfies DARA. Using this utility function, we can estimate the risk preferences of the bidders based upon their bids in the sealed bid auctions and generate predictions of the switch-over prices that bidders would have had, were their choices guided by the same risk attitudes they exhibited in their bidding behavior. This predicted switch-over price is computed by finding an entry price for the ascending auction that makes the expected utility of participating in the two different auction formats equal given that the entry price for the sealed bid auction was 1.4. Due to the length of the equations for performing these calculations, they have been placed in the appendix.

There are two important issues involved in performing these calculations. The most obvious involves the fact that when someone loses an auction yet still pays an entry price, their utility is $(-e)^\alpha$ where e is the entry price they paid. Since taking roots of negative numbers leads to problems, we must use some measure of wealth to add into the utility function to insure a positive argument. What to use for this measure of wealth is certain to be a controversial issue. We investigated several different measures of wealth and found that the one that worked best was the cash balance of the subject at the time of the decision divided by ten. This measure is constructed to be the smallest wealth measure we can use, the amounts typically correspond to about \$0.50-\$1.00, and still be able to evaluate the equations. We refer to this as the “no wealth” case for convenience. The other two wealth specifications we spent the most time examining were the actual cash balance in the experiment and the cash balance multiplied by 100. The latter was intended to be a proxy for the notion that subjects considered their external wealth as well. Both lead to predictions that did not match the data well at all. Reasons for why these specification did not and should not have been expected to fit well are explained below.

The second technical issue with the computation of a predicted switchover price deals with the expectations subjects have regarding the risk attitudes of their opponents in sealed bid auctions. Here too we investigated multiple specifications. We found that the best fitting and most reasonable specification was one in which subjects assume their opponent will possess the average degree of risk aversion in the population (the “average opponent” case). Since the auction choice behavior occurs after observing several auctions from the first phase of the experiment, it seems quite reasonable to assume that bidders have developed some intuition about the degree of risk aversion of likely opponents. An alternative, and more plausible explanation, is that subjects were able to form a belief about the average bid/value ratio the other subjects used in their bidding. Either way of motivating this belief is observationally equivalent. While we could use the actual empirical distribution as the beliefs for this case, that would be overly cumbersome and likely not obtain better results than modeling subjects as though they assume their opponent possesses the average level of risk aversion in the population. The derivation of the bid functions under this belief system can be found in the appendix⁹.

With these issues dealt with, risk aversion parameters can be estimated by running a standard OLS regression with the equation $b_i = \beta + \gamma v_i$ where $\gamma = \frac{1}{1+\alpha}$ as is developed in Cox, Roberson, and Smith (1982) and Cox, Smith, and Walker (1988). Doing so yields a distribution of risk aversion parameters with an average of 0.66¹⁰ which is what will be used as the average level of risk aversion

⁹The other specification considered was the “equivalent opponent” hypothesis in which the subjects were presumed to form a belief that their opponent would possess the same risk aversion parameter as themselves.

¹⁰As a technical note, there were only 8 out of the 60 bidders who possessed intercepts that were statistically significant at the 5% level. Excluding these 8 yields an average α of 0.60.

in the population in the calculations below.¹¹ One might suggest that since we have put wealth into the utility function for auction choice, we should also have wealth in the utility function for bidding in sealed bid auctions and thus for estimating risk aversion. We have done so, but will not present either the methods or details of the results here to conserve space. We have chosen to omit the results for two reasons. First is that the general nature of our conclusions do not appear to change. Second, the average degree of risk aversion is estimated to be -19.05 (requires changing the utility function to $u(x) = x^\alpha/\alpha$ to allow for negative values of α) assuming subjects consider their wealth to be their cash balance at the time of the bid plus show-up fee in the experiment. We believe this is simply an incomprehensibly large degree of risk aversion and indicates a specification bias in the utility function from the addition of wealth. This is related to the issues discussed in Rabin and Thaler (2001)¹² which suggest that when one assumes that decision makers consider their wealth when making decisions over small amounts of money, you find levels of risk aversion far greater than is consistent with behavior concerning larger amounts of money. Thus it is most appropriate to assume that decision makers do not consider their wealth position when making choices. This argument also explains part of the reason we find that the No Wealth model works best in the auction choice setting¹³.

Figure 3 shows a dual scatterplot of the predicted switchover prices with the No Wealth-Average Opponent model as well as the observed switchover prices against the risk aversion parameters for each bidder. The most salient feature of the predicted relationship between switchover prices and the risk aversion parameter is that it is quite flat. This is seemingly counterintuitive as bidders who are quite risk averse expect to make very little money in a sealed bid auction and should therefore be willing to pay more to enter into an ascending. The problem with this notion is that bidders who are quite a bit more risk averse than the average expect their probability of winning to be high in the sealed bid auction relative to the ascending. While they expect to make little surplus in the sealed bid, the expected probability of getting the surplus is very high and this causes the sealed bid auction to be very appealing relative to the more risky yet lucrative ascending auction. Thus even highly risk averse bidders would not be willing to pay much to enter into the ascending auction. Examining the pattern of observed switchover prices plotted against the risk aversion parameters, this predicted relationship is precisely what emerges for a good portion of the subjects. This pattern in the observed switchover prices can therefore be matched by the no wealth-average opponent model, as demonstrated in figure 3. The other model specifications investigated can

¹¹For the equivalent opponent model this is a correctly specified estimation. For the average opponent case, the estimated bid function is an approximation only, due to the existence of the flat region for bids above $\frac{1}{1+\alpha}(100) = \frac{1}{1+.66}(100) = 60.241$. A bidder will be predicted to bid 60.241 anytime their value is such that $\frac{1}{1+\alpha}(v) \geq 60.241$. For some extremely risk averse bidder, say $\alpha = .3$, this threshold value is, e.g., 78.313. For only about half of the subjects is this ever an issue, those more risk averse than average, and even for them, this flat portion of the bid function covers a relatively small range of the value space. Thus the bias introduced by ignoring this in the estimations should not be expected to be severe while keeping the same risk aversion values throughout the analysis aids in continuity of exposition.

¹²It is important to note that the specific problem reported in Rabin and Thaler (2001) does not actually apply to the risk aversion model as applied to bidding behavior. The easiest way to see this is that the loss aversion model proposed to “fix” the problem is the same as the risk aversion model we have used in this context as losses are not possible. More generally, both the results reported here and those in the Rabin and Thaler (2001) paper seem to suggest that the more reasonable interpretation of both is not that the risk aversion model is not applicable but rather that the problem is in the assumption that people always consider their external wealth position in any decision. Removing this assumption leaves us with reasonable levels of risk aversion here and would do the same for the examples in Rabin and Thaler (2001). Of course a complete discussion of this issue goes well beyond the scope of this paper and those unconvinced by this short note are directed to Cox and Sadiraj (2002) for a more in-depth examination of these issues.

¹³For a more complete explanation of why we argue the Wealth and equivalent bidder assumptions do not fit the data, please see the working draft of this paper.




Figure 3: Scatterplot of both observed and predicted switchover prices using the No Wealth-Average Opponent model against estimated risk aversion parameter for each individual.

not match this property of the data. The equivalent opponent model predicts a strong negatively sloped relationship between switchover price and risk aversion, as do other specifications of the wealth assumption.

The no wealth-average opponent model appears quite good at rationalizing the choice behavior of many of the subjects, but there is a large group of subjects who were observed to be willing to pay much more than would be predicted based upon their degree of risk aversion alone which accounts for the regression line between the observed points being pulled above the predicted line. A plausible explanation for the fact that many bidders appear to be willing to pay more for the ascending auction than they “should” is that for various reasons participating in the sealed bid auction incurs a certain amount of disutility. This disutility could come from the extra effort that is required to figure out how to determine a bidding strategy, the mental anguish over seeing someone else win with a bid lower than your value or any number of other factors.¹⁴ We can augment the framework developed above to obtain an estimate of this disutility. We will model this disutility using a parameter λ that will be considered to be the equivalent of an extra entry price for the sealed bid auction in the utility function. Thus when a bidder wins a sealed bid auction, their utility is $(W + v - b - e + \lambda)^\alpha$ and when the bidder loses, $(W - e + \lambda)^\alpha$. We could of course frame this willingness to pay as based upon extra utility from participating in the ascending auction, but the previous formulation seems more natural and identification problems preclude simultaneously identifying both. To find the value of λ for each bidder, we will use the No Wealth-Average opponent model and solve for the value of λ that makes the bidder indifferent between the sealed bid auction at an entry price of 1.4 and the ascending auction at their observed switchover price, $\delta_i + 1.4$ given their estimated degree of risk aversion. Details can again be found in the

¹⁴We conducted a short post-experimental questionnaire on this issue and subjects gave the following main reasons for preferring the ascending auction: (i) higher payoffs, (ii) easier decision-problem, (iii) decision independent from what others do, (iv) avoidance of the risk to lose by bidding too low, and (v) no uncertainty.




Figure 4: Histogram of the estimated values of λ , or value of participating in a sealed bid auction, found for the subjects.

appendix.

For some bidders, λ will be negative indicating a disutility for the sealed bid auction while for others λ could be positive. These will obviously correspond to bidders who were found to be willing to pay more/less than the theoretical prediction. A histogram summary of the λ 's found in the population can be seen in figure 4. Most of the mass in the population is between 2 and -2. There are only a few subjects who possess a $\lambda > 2$ (8) while quite a number possess a $\lambda < -2$ (15) and some significantly so. This matches with the results found in figure 3 showing relatively few and only minor over-predictions of switch-over prices yet a number of large under-predictions.

It is difficult to understand exactly what this λ means in its raw form. It is therefore useful to translate λ into a different measure of the degree to which subjects prefer the ascending auction to the sealed bid. The most natural such measure would be to find the n -bidder ascending auction that is utility equivalent to the 2-bidder sealed bid auction. To do that we solve for the \hat{N} that makes the expected utility of being in a sealed bid auction, assuming the computed value of λ for the bidder, equal to the expected value of being in an n -bidder ascending auction. Details on how this calculation is done are found in the appendix. There were 25 bidders found to have a $\lambda < 0$ out of the 46 who were found to be risk averse and not prefer the sealed bid auction. Of these 25 bidders, only 2 were found to possess an $\hat{N} = 2$, 14 had an $\hat{N} = 3$, 3 with $\hat{N} = 4$, 1 with $\hat{N} = 5$, 1 with $\hat{N} = 9$ and 4 bidders were found to have values of \hat{N} that were essentially arbitrarily large¹⁵. This measure of the degree to which a bidder prefers an ascending auction is quite strong due to the very large jump in expected value between a 2 and 3 bidder auction. Thus the fact that so many bidders would be willing to participate in ascending auctions with larger numbers of bidders even when they could participate in a 2 bidder sealed bid auction indicates quite strong preferences for the ascending auction.

¹⁵Possible values out to 40 were checked and the ascending still generated higher utility. Since expected values vary so little at this point and beyond it was considered not worthwhile to search at higher values.

4 Conclusion

In this study we attempt to do two things. First, we develop a methodology capable of measuring bidder preferences for various attributes of an auction design. Second, we apply this methodology to looking at bidder preferences between sealed bid and ascending auctions. The results show that when the entry prices for the two auction institutions are equal, subjects overwhelmingly prefer the ascending auction. However, the revealed willingness to pay for that preference in the second phase is much lower than the average realized profit differential between the two auction mechanisms. One hypothesis of subject behavior is that they are more or less risk neutral and bidding above the risk neutral Nash equilibrium level in the sealed bid auctions is a function of some phenomenon other than risk aversion. Were that true, it would have been reasonable to expect bidders to have paid up to the 10.02 level for the ascending auction since below that level the expected profit from the ascending is greater than the expected profit from the sealed bid. This did not happen and we were forced to search for an alternative explanation.

We first tried to explain the observed switch-over prices through experiential data concerning the subjects experiences in the learning phase. Data on the individual outcomes from the learning phase had no explanatory power over the observed switchover prices. We therefore tried modeling the subjects choosing as if they possessed risk averse preferences and utility or disutility from bidding in a sealed bid auction. The results from measuring this disutility indicated that many of the bidders would be willing to enter into ascending auctions with at least 3-4 bidders instead of a 2 bidder sealed bid auction. Considering the sizable expected value difference in going from a 2 bidder to a 3 or 4 bidder auction, this suggests very strong preferences for the ascending. Further, these preferences seem to be of the sort that an auction designer could exploit to increase their revenue. The degree to which an auctioneer can do so requires additional testing to be able to make a definitive conclusion on the matter which is the subject of ongoing research.

There is a way of looking at our results that would suggest much less significance to them. This is derived from the results found in Isaac and James (2000) as in that paper the authors attempt to estimate the risk aversion parameters possessed by subjects using two different procedures (sealed bid auctions and the Becker-DeGroot-Marschak procedure) and find the estimations obtained using both of these procedures to be quite different. The implication is that either risk preferences are not invariant between choice mechanisms or perhaps that the bidding behavior in auctions is truly not a function of risk aversion. Either would suggest that our attempt to measure risk preferences in the sealed bid auction behavior and use that to predict behavior in the auction choice setting is doomed to failure. These experiments do not allow us to adequately deal with such an argument, but in future work we will be subjecting the predictions on entry behavior this model suggests to detailed tests to determine if they are robust.

There is a final important insight that can be derived from the observed willingness to pay of most of the bidders in regard to entry prices for their preferred auction. Many optimal auction designs (see, e.g., Bulow and Roberts (1989), Riley and Samuelson (1981)) rely upon the use of properly chosen entry prices for full surplus extraction. The theories generally rely on the assumption of a fixed pool of bidders. Our results cast doubt on the true optimality of those designs for cases in which bidders have alternative options. If the bidders have an alternative option and if the entry price is too high, they may well take advantage of it. The results have not been included above to conserve space, but if we were to consider these experiments an attempt to determine if adding an entry fee to the ascending auction would increase revenue, the answer would be a clear “no”. Total revenue in the ascending auction is almost monotonically decreasing in the entry price. The reason for that is that while per auction revenue is increasing in the entry price, the auctioneer will be able to conduct fewer and fewer of these auctions as the entry price rises

since more bidders will be pursuing their outside options. The latter effect outweighs the former leading to total revenue decreasing as the entry price rises. These results show that entry prices do not have to be very high in order to scare off would-be bidders even when those bidders have a significant preference for the more expensive mechanism. These results support the theoretical predictions derived by Levin and Smith (1994) that the optimal design for the seller in private value auctions is one without an entry fee, though for different reasons.

APPENDIX A: Computation of Expected Utilities

Sealed Bid Auctions: Bid Functions and Expected Utility

As shown in Cox, Roberson, and Smith (1982), the optimal bid function for a risk averse bidder in the SIPV environment assuming that values are uniformly distributed on the range $[0, 100]$ can be defined by $b^*(v) = v \left(\frac{N-1}{N-1+\alpha} \right)$ so long as $b < 100 \left(\frac{N-1}{N-1+\alpha} \right)$. Above this point, the bid function becomes non-linear. This non-linearity is due to allowing that bidders may believe that their opponents have different risk aversion parameters than themselves. For the average opponent model, we are using the “hook” portion of the bid function turns into a flat upper bound. Assuming the average risk aversion parameter is $\bar{\alpha}$, the bid function is:

$$b^*(v) = \begin{cases} v \frac{N-1}{N-1+\alpha} & \text{for } v \in [0, \min(\frac{N-1+\alpha}{N-1+\bar{\alpha}} * 100, 100)] \\ 100 \frac{N-1}{N-1+\bar{\alpha}} & \text{for } v \in [\min(\frac{N-1+\alpha}{N-1+\bar{\alpha}} * 100, 100), 100] \end{cases} \quad (2)$$

Computing the expected utility for choosing to enter a sealed bid auction under this model requires first being careful to specify the expected probability of winning given any value v . Now bidder i bids b_i assuming that their opponent is bidding $\frac{N-1}{N-1+\bar{\alpha}} v_j$, then

$$\begin{aligned} \Pr(\text{win}) &= \Pr(b_i > b_j = \frac{N-1}{N-1+\bar{\alpha}} v_j) = \Pr\left(\frac{N-1+\bar{\alpha}}{N-1} b_i > v_j\right) \\ &= F\left(\frac{N-1+\bar{\alpha}}{N-1} b_i\right)^{N-1} = \left(\frac{\frac{N-1+\bar{\alpha}}{N-1} b_i + 1}{101}\right)^{N-1} \end{aligned}$$

If $b_i = v \frac{N-1}{N-1+\alpha}$ then this becomes

$$\left(\frac{\frac{N-1+\bar{\alpha}}{N-1} b_i + 1}{101}\right)^{N-1} = \left(\frac{\left(\frac{N-1+\bar{\alpha}}{N-1}\right) v \left(\frac{N-1}{N-1+\alpha}\right) + 1}{101}\right)^{N-1} = \left(\frac{v \left(\frac{N-1+\bar{\alpha}}{N-1+\alpha}\right) + 1}{101}\right)^{N-1}$$

At least this is i 's probability of winning along the sloped portion of the bid function. Along the flat portion of the bid function, i 's belief is that their opponent will never bid above $100 \frac{N-1}{N-1+\bar{\alpha}}$ thus if i bids that plus ϵ , he expects to win with a probability of 1.

If we add in the terms for wealth, W , entry price, e , and the hypothesized utility for bidding in a sealed bid auction, λ , the the expected utility of bidder i for participating in a sealed bid auction

assuming he is facing a bidder with risk aversion of $\bar{\alpha}$ is :

$$EU_S(W, e, \alpha, \lambda) = \sum_{v=0}^{\min(\frac{1+\alpha}{1+\bar{\alpha}}(100), 100)} \left(\begin{aligned} & \left(W + \frac{\alpha}{1+\alpha}v + \lambda - e \right)^\alpha * \frac{1+\bar{\alpha}}{101}v+1 * \frac{1}{101} + \\ & \left(W + \lambda - e \right)^\alpha * \left(1 - \frac{1+\bar{\alpha}}{101}v+1 \right) * \frac{1}{101} \end{aligned} \right) + \sum_{v=\min(\frac{1+\alpha}{1+\bar{\alpha}}(100), 100)}^{100} \left(\left(W + v - \frac{1}{1+\bar{\alpha}}(100) + \lambda - e \right)^\alpha * 1 * \frac{1}{101} \right)$$

Ascending Auctions: Bid Functions and Expected Utility

In an ascending auction the risk attitude of the opponent is irrelevant thus any model of expectations produces the same result. For a given value v , a bidder will pay the next highest value if they do win or their utility will be $(W + v - v_{(2)} - e)^\alpha$. The distribution of $v_{(2)}$ is given by the standard distribution of an order statistic for the k 'th highest draw, $f_{(k)}(x) = \frac{(n)!}{(n-k)!(k-1)!} F(x)^{n-k} [1 - F(x)]^{k-1} f(x)$, for the case in which x is drawn from the discrete uniform distribution on $[0, v]$, $k = 1$ and $n = N - 1$, N being the actual number of bidders in the auction. This simplifies to $(N - 1) \left(\frac{x+1}{v+1} \right)^{N-2} \frac{1}{v+1}$. To find the expected value of being in the ascending auction, we must first compute the expected surplus that occurs when the bidder wins multiplied by the probability of winning and then we add on a second term for the case in which the bidder loses. Since we will be allowing for the possibility of more than 2 bidders in the ascending auctions, we have to generalize to N bidders.

$$EU_A(W, \alpha, e, N) = \sum_{v=0}^{100} \left(\begin{aligned} & \left(\sum_{x=0}^v (W + v - x - e)^\alpha * (N - 1) \left(\frac{x+1}{v+1} \right)^{N-2} \frac{1}{v+1} \right) \\ & * \left(\frac{v+1}{101} \right)^{N-1} * \frac{1}{101} \\ & + (W - e)^\alpha \left(1 - \left(\frac{v+1}{101} \right)^{N-1} \right) * \frac{1}{101} \end{aligned} \right) \quad (3)$$

Finding Predicted Switchover Prices

In the experiments, it was observed that for equal entry prices most bidders preferred the ascending auction. The entry price for the sealed bid auction is then held static at 1.4 while the entry price of the ascending auction, e , is increased until we find a point that the bidder prefers to choose the sealed bid auction. Given a level of risk aversion, α , and a level of wealth, W , it is possible to compute predictions as to what this price should be by solving the following equation for e .

$$EU_S(W, 1.4, \alpha, 0) = EU_A(W, e, \alpha, 2) \quad (4)$$

Computing λ and \hat{N}

The disutility of participating in a sealed bid auction is the entry price equivalent of what would make the expected utility of participating in a sealed bid auction equal to the expected utility of being in the ascending at the entry price the subject was observed to pay, $\delta_i + 1.4$. This is done by finding the value of λ that solves the following:

$$EU_S^A(W, 1.4, \alpha, \lambda) = EU_A(W, \delta_i + 1.4, \alpha, 0, 2) \quad (5)$$

Finally, the \hat{N} -bidder ascending auction a bidder finds equivalent to participating in a two bidder sealed bid auction assuming entry prices of 0 for both is found by solving the following for \hat{N} :

$$EU_S^A(W, 0, \alpha, \lambda) = EU_A(W, 0, \alpha, 0, \hat{N}) \quad (6)$$

References

- Bajari, P. and A. Hortacsu (2000). Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions. Stanford University Working Paper.
- Bulow, J. I. and D. J. Roberts (1989). The Simple Economics of Optimal Auctions. *Journal of Political Economy* 97, 1060–1090.
- Cox, J. C., B. Roberson, and V. L. Smith (1982). Theory and Behavior of Single Object Auctions. In *Research in Experimental Economics*, Volume 2, pp. 1–43. JAI Press Inc.
- Cox, J. C. and V. Sadiraj (2002). Risk Aversion and Expected Utility: Coherence for Small- and Large-Stakes Gambles. University of Arizona Working Paper.
- Cox, J. C., V. L. Smith, and J. M. Walker (1988). Theory and Individual Behavior of First Price Auctions. *Journal of Risk and Uncertainty* 1, 61–99.
- Engelbrecht-Wiggans, R. (1993). Optimal Auctions Revisited. *Games and Economic Behavior* 5, 227–239.
- Fischbacher, U. (1998). z-Tree: Zurich Toolbox For Readymade Economic Experiments. Zurich University.
- Friedman, D. (1983). Effective Scoring Rules for Probabilistic Forecasts. *Management Science* 29(4), 447–454.
- Goeree, J. K. and T. Offerman (2002). First-Price Auctions to Sell Heterogeneous Licenses. CREED Working Paper.
- Harstad, R. M. (1990). Alternative Common Values Auction Procedures: Revenue Comparisons with Free Entry. *Journal of Political Economy* 98, 421–429.
- Isaac, R. M. and D. James (2000). Just Who Are You Calling Risk Averse. *Journal of Risk and Uncertainty* 20(2), 177–187.
- Ivanova-Stenzel, R. and D. Sonsino (2001). Comparative Study of One-Bid versus Two-Bid Auctions. Working Paper, Humboldt University.
- Klemperer, P. (2002). What Really Matters in Auction Design. *Journal of Economic Perspectives* 16(1), 169–190.
- Levin, D. and J. L. Smith (1994). Equilibrium in auctions with Entry. *American Economic Review* 84(3), 585–599.
- Lucking-Reiley, D. (1999). Experimental Evidence on the Endogenous Entry of Bidders in Internet Auctions. Working Paper, University of Arizona.
- Lucking-Reiley, D. (2000). Auctions on the Internet: What's being auctioned and How? *Journal of Industrial Economics* 48(3), 227–252.

- Matthews, S. (1987). Comparing Auctions For Risk Averse Buyers: A Buyer's Point of View. *Econometrica* 55(3), 633–646.
- McAfee, P. (1993). Mechanism Design by Competing Sellers. *Econometrica* 61(6), 1281–1312.
- McAfee, P. and J. McMillan (1987). Auctions with Entry. *Economics Letters* 23, 343–347.
- Milgrom, P. R. and R. J. Weber (1982). A Theory of Auctions and Competitive Bidding. *Econometrica* 50(5), 1089–1122.
- Rabin, M. and R. H. Thaler (2001). Risk Aversion. *Journal of Economic Perspectives* 15(1), 219–232.
- Riley, J. G. and W. F. Samuelson (1981). Optimal Auctions. *American Economic Review* 71, 381–392.
- Selten, R. (1998). Axiomatic Characterization of the Quadratic Scoring Rule. *Experimental Economics* 1, 43–62.