Revenue Equivalence Revisited*

Radosveta Ivanova-Stenzel† Humboldt University
Timothy C. Salmon‡ Florida State University

December 2007

Abstract

The conventional wisdom in the auction design literature is that first price sealed bid auctions tend to make more money while ascending auctions tend to be more efficient. We re-examine these issues in an environment in which bidders are allowed to endogenously choose in which auction format to participate. Our findings are that more bidders choose to enter the ascending auction than the first price sealed bid auction and this extra entry is enough to make up the revenue difference between the formats. Consequently, we find that both formats raise approximately the same amount of revenue. They also generate efficiency levels and bidder earnings that are roughly equivalent across mechanisms though the earnings in the ascending might be slightly higher.
JEL Codes: C91, D44
Key Words: bidder preferences, private values, sealed bid auctions, ascending auctions, endogenous entry

1 Introduction

Early work in the auction literature focused on the famous Revenue Equivalence Theorem (RET) for which the principles were established in Vickrey (1961) and then later proven to be more general in Myerson (1981) and Riley and Samuelson (1981). The claims of the base version of the theorem are that so long as the assumptions of the symmetric independent private values (SIPV) environment hold, then all standard auctions will yield the same expected profit to an auctioneer and the same expected surplus to bidders.[1] As important and fundamental result as this was and is to the auction

*We thank seminar participants at Indiana University, The Johns Hopkins University and Florida State University for helpful comments as well as attendees at the 2006 ASSA winter meeting and the 2006 INFORMS conference. We would also like to thank the Associate Editor and two anonymous referees for providing useful feedback that helped to improve the paper. We acknowledge the National Science Foundation for providing funding for these experiments.

†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.

[1] The key assumptions underlying this theorem, as described in Krishna (2002), are independence of bidder values, risk neutrality of bidders, lack of bidder budget constraints and that all bidder values are drawn from the same distribution. The formal requirements that two mechanisms in the SIPV environment must satisfy in order to apply the RET are that (i) a bidder with a value of v must have the same probability of winning in both and (ii) a bidder who draws the lowest value in the range, $\underline{v}$, must expect the same surplus in both institutions.

literature, it has to date received little empirical support. Since the publication of Cox, Roberson, and Smith (1982) it has been a long established and easily re-verified result that the RET fails to universally hold and in particular that first price sealed bid auctions generate significantly more revenue than ascending auctions in the SIPV environment.[2] For example, Kagel and Levin (1993) show that in auctions with five bidders first price auctions generate significantly higher average revenues than second price auctions. When applying the "bid function approach", i.e. each subject submits a complete vector of bids (a bid function), Güth, Ivanova-Stenzel, Königstein, and Strobel (2002) find that the switch from first price to second price rule in auctions with three bidders leads to substantial decrease in expected prices.

The reason the RET fails to hold is that bidders are found to bid above the risk neutral Nash equilibrium in the first price sealed bid auction which leads to higher than expected revenue. The standard equilibrium prediction works well for the ascending auction leading to revenue in-line with the prediction. While the reason for the "overbidding" observed in first price auctions has generated a substantial amount of controversy over the years (see the December 1992 issue of *American Economic Review* as well as Kagel (1995) for a sampling of the arguments), the fact of the overbidding and it's implication for which mechanism will generate more revenue have been generally well accepted. We seek to call this conventional wisdom into question.

This is related to an empirical puzzle regarding the auction formats actually used by auctioneers in certain domains, which is that despite all of the evidence collected on this issue suggesting that first price auctions are revenue dominant, the ascending auction dominates many field applications. In one survey of online auction sites, Lucking-Reiley (2000) finds that of 142 surveyed auction sites, 121 use ascending price auctions while only 21 used sealed bid auctions. It is difficult to explain this preponderance of ascending auctions based on the assumption that auctioneers are revenue maximizers and the assumption that the previous laboratory evidence that first price auctions raise more revenue translates into field applications. Our study will attempt to provide a solution for this contradiction based on forces one will observe in markets but are unaccounted for in previous experimental studies, i.e. endogenous entry choice by the bidders.

Previous investigations of the revenue generation abilities of first price sealed bid and ascending auctions have typically shared one key aspect in common. This is the assumption that the number of bidders, n, is fixed and constant across both mechanisms. Thus, the correct way to state the

[2] Other violations are also found. For example, the Dutch auction can generate lower revenues than the first price auction and the English auction may yield lower revenue than the second price auction as shown in Coppinger, Smith, and Titus (1980) and Kagel, Harstad, and Levin (1987).

established empirical result is that given a common number of bidders, the sealed bid first price auction will raise more money than the ascending auction. In naturally occurring auctions the number of bidders at an auction is not exogenously determined and fixed. Our question is what happens to the revenue generation of the two auction formats when n is not held constant across both formats but is allowed to vary endogenously based on the choices of the bidders?

This is an important auction design issue because auctioneers can not in practice enforce participation. Perhaps the key ability of a successful auctioneer is his or her ability to attract bidders to an auction. While there are many factors involved in attracting bidders to an auction (e.g., the entry or reserve price, the reputation of the seller, the quality of the listing), our view is that the format of the auction itself may be one of the important factors for the entry decision of a bidder.

There is existing theoretical and empirical literature that examines entry decisions in regard to auctions, but most of it deals with contexts quite different than the one we are concerned with here. For example, Engelbrecht-Wiggans (1993), Smith and Levin (1996), Smith and Levin (2002), Pevnitskaya (2004) and Palfrey and Pevnitskaya (2008) consider the decision of whether to enter an auction or not and study variously the nature of the entry choice itself and/or the effect of endogenous entry on bidding behavior. Bajari and Hortacsu (2003), Lucking-Reiley (1999a) and Levin and Smith (1994) look at a bidder's 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 and there are other papers that examine the issue of auctioneers competing in such an environment. Our point of concern is to look at an environment in which bidders are allowed to endogenously choose which auction to participate in when the main characteristic that distinguishes the auctions is the format being used and to re-examine the generally accepted opinion that first price auctions tend to make more money while ascending auctions tend to be more efficient.

In prior work, Ivanova-Stenzel and Salmon (2004a) and Ivanova-Stenzel and Salmon (2008), we established that bidders do prefer ascending (English) and Vickrey auctions, i.e. second price rule, to first price sealed bid auctions (will be abbreviated as just the "sealed bid" auction), on a *ceteris paribus* basis and that they are typically willing to pay a higher entry fee to participate in ascending auctions instead of sealed bid auctions on the basis of that preference.[3] This was determined by first allowing subjects to choose to enter into either a two-bidder ascending or sealed bid auction

[3]We note that this result requires symmetric bidders. Klemperer (2002) argues that ascending auctions can discourage the entry of any disadvantaged or weak bidders because they know their chances of winning are very small. This claim is supported by Goeree, Offerman, and Schram (). Our work is aimed at situations in which all bidders are *a priori* equivalent as should be the case in most on-line auction markets such as eBay, Yahoo etc.

at equal entry fees multiple times and then raising the entry price of the auction format they chose most often to determine how much they were willing to pay for this preference. We discovered that the bidders were not willing to pay an entry fee to get into the ascending auction up to a level that would make a risk neutral person indifferent between the two formats. Generally, our subjects would cease paying the higher entry fee for the ascending auction and instead choose the sealed bid auction at a price at which, even with the higher entry fee, their expected profit would have been greater by choosing the ascending auction. This led to our result discussed in Ivanova-Stenzel and Salmon (2004b) showing that auctioneers could not exploit this revealed preference for the ascending auction by using entry fees. Alternatively, we also found that if we assumed bidders were risk averse and estimated their CRRA parameter based on some initial training rounds, these parameters predicted the nature of the switch-over prices quite well except that on average subjects were willing to pay more than predicted. We then extrapolated from these revealed preferences to show that if our model were accurate then many of these bidders would have been willing to enter into ascending auctions with 3, 4 and perhaps even more bidders instead of a sealed bid auction with only 2. This result suggests very strong preferences for the ascending auction that auctioneers may be able to exploit profitably through a competitive effect by using ascending auctions instead of sealed bid auctions.

In this paper we will present the results of an experiment designed to test these propositions. That is, would bidders in fact be willing to enter larger ascending than sealed bid auctions and if so, which auction format ends up earning more revenue? Our results will show that more bidders choose to enter the ascending than the sealed bid auction and this extra entry is enough to make up the revenue difference between the formats. Consequently, we find that there is no statistically significant difference in the revenue raised by the two formats. Further, we find both to be equally efficient and to generate close to the same average earnings to bidders though the ascending generates slightly more. The fact that there are differences in the number of bidders between the formats is an important point to realize in properly understanding these results. Since more bidders enter into the ascending auction this means the total surplus to be divided between bidders and sellers should be (and is) larger than in the sealed bid auction. If bidders arbitrage between the two formats to equalize their earnings between formats, that is not a guarantee that the residual surplus that becomes revenue to the seller will also be equalized. Consequently it is quite an interesting finding that this seems to hold empirically.

There is a related paper, Engelbrecht-Wiggans and Katok (2005), that considers auction choice experiments similar to ours. The authors also allow endogenous entry between the two formats and one finds results embedded in their tables similar to those for which we will provide evidence. However, examining the revenue differences and such between the auction formats was not a focus of that paper. Furthermore, their results are based on small sample sizes (in some cases a single session) which means that their results can only be taken as a preliminary investigation of this phenomenon. Based on seeing some of their results, the authors were not convinced that their subjects were able to accurately understand expected profit differences between auction institutions. The thrust of that study involved using various experimental treatments to understand the difficulties subjects have in evaluating differences in expected profits derived from differences in auction formats. Their concern is that "the standard laboratory setting, where subjects see a randomly-drawn value each period makes the problem more difficult by impeding learning because subjects do not see outcomes of their actions in any systematic way." Our experiment was constructed specifically to avoid this problem and make sure that subjects would experience each auction format with a single value matched up against the same opponents who also possessed the same values across formats (i.e., values were independent across subjects but each subject bid with the same value in multiple mechanisms). By also doing this systematically multiple times in auctions with different numbers of bidders, we should have given the subjects their best chance at understanding profit differentials between auction formats as well as how profits might change as the number of competitors change. Thus, we believe that our results on these issues are more definitive.

The closest theoretical paper to our interest is Smith and Levin (1996), which contains an exploratory approach to the theory of auction revenue with endogenous mechanism choice in a similar environment in that they allow for risk averse bidders and try to solve for an entry equilibrium and revenue ranking across mechanisms. Based on the results of Matthews (1987) showing that bidders possessing IARA preferences will prefer the sealed bid instead of the ascending auctions and CARA preferences will lead to indifference, Smith and Levin (1996) are able to establish a clear revenue ranking in those cases with the sealed bid auction being clearly revenue dominant. DARA is perhaps the more empirically interesting case and Matthews (1987) shows that bidders with these preferences will prefer the ascending to the sealed bid auction. Determining a revenue ranking here is difficult because while in equilibrium the ascending auction will attract more bidders, the revenue ranking is highly dependent on how many more bidders the ascending auction attracts compared to the sealed bid auction and how far above the RN level those bidders bid in the sealed bid auction. Smith and Levin were able to derive no clear result in this case to suggest which of the two mechanisms should be revenue dominant. We do not point this out as a criticism of their paper but rather to make it clear that existing theory suggests an indeterminate result in regard to which format will raise more money when endogenous entry is allowed under the empirically plausible assumption of DARA risk preferences.

In addition to the results described above, our experiments will show that risk aversion (or degree of overbidding) does not significantly effect bidders' entry decisions. Furthermore, assuming the type of risk aversion bidders possessed satisfies DARA,[4] we find that the expected utility for choosing the ascending auction is not higher than the expected utility for choosing the sealed bid auction for any risk averse bidder.

The remainder of the paper is organized as follows: In section 2 we will provide a description of our experiment. In section 3 we present our results and in section 4 we provide a discussion regarding how we believe our results should be interpreted.

2 Design of Experiment

Our experiment consisted of two phases. The first phase we will refer to as the learning phase as in this phase we had the subjects participate in several rounds of ascending auctions[5] (will be abbreviated henceforth as the "A" auction) and sealed bid auctions (will be abbreviated henceforth as the "SB" auction) with differing but exogenously fixed n's. This was to allow subjects to develop some experience with both auction formats and to give them some idea of how profit levels might vary with n. In the second phase, subjects were allowed to repeatedly choose to participate in either a sealed bid auction or an ascending auction in which the number of competitors was endogenously determined.

We conducted four sessions of our experiment with 12 subjects per session. In the learning phase, subjects participated in twelve rounds of fixed n auctions. In each round they were informed of their value for winning the auction which was drawn independently for each subject from a uniform distribution over the integers in the range $[0, 100]$. All values were denoted in a fictitious currency termed ECU for Experimental Currency Unit. The subjects then participated in both a SB and A auction using the same value. They would first submit a bid in the SB auction and

[4] 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.

[5] The ascending auctions we used were clock auctions in which the price would start at 0 and increase by 1 ECU every second up to a maximum of 150 (which is above the maximum value of a subject). Each subject had a button on their screen to exit the auction and the auction ended when only 1 bidder remained with that bidder paying the price at which the last bidder dropped out.

then immediately participate in the A auction and only after the A auction was completed would they be informed of the results from both auctions. These twelve rounds consisted of two blocks of six period cycles in which the subjects participated in three rounds with $n = 2$ auctions and then three with $n = 4$. At the end of the learning phase the subjects received feedback telling them the session-wide average profit in each auction format and for each n. They also received information on their total profit up to this time.

For the main phase of the experiment, we split each session of twelve subjects into separate groups of six. These groupings were held constant throughout the rest of the session which keeps the two groups independent for statistical purposes. The choice of the size of the groups is a very important one for this experiment. Having six in the pool allows for auction sizes to vary in what is a key region for this issue. Auctioneers should be mostly concerned about choosing the right mechanism to maximize revenue in the small n range. At large n's one will find that the revenue differences between mechanisms, especially as a percentage of total revenue, get quite small but they are large for small n. Consequently, if an auctioneer knows he or she will get 8 or more bidders, the choice of an auction mechanism is much less important than if the number of prospective bidders is in the 2-6 range. It is also the case that the marginal increase in revenue from one additional bidder is largest for small n and, therefore, this is the range in which an auctioneer will find it most important to entice an extra bidder to their auction. This is why we are concentrating on relatively small pools. Further, having an even number in the prospective bidder pools allows for the possibility that bidders can split evenly between auctions should that be their natural preference. Were we to use an odd number and find that one format attracted more bidders than the other, that would be less convincing since the result is almost required by the design. Thus, having groups of size six allows us to concentrate exactly on the important range of the size distribution without rigging our results in favor of finding differences in the number of bidders choosing each auction.[6]

The main phase of our experiment consisted of 30 rounds in which subjects could choose in which auction to participate. To maximize the number of auction choice periods in the time allotted, we did not have the subjects actually play the auctions in each round. Instead, each set of auctions was only conducted with a 20% probability. This allowed us to get more observations on the auction choice behavior of the subjects which is important because this represents a non-trivial coordination problem that could take a while to equilibrate. Furthermore, having the possibility

[6] In online auctions, the observed average number of bidders is also rather small. For example, in Bajari and Hortacsu (2003) the observed average number of bidders is 3.05. In an extensive review on online auctions Ockenfels, Reiley, and Sadrieh (2006) reports average number of bidders not higher than 5.0.

of choosing to be the only bidder in an auction could lead to substantially increased variance in the behavior. In order to decrease the difficulty of the coordination problem we eliminated the possibility of choosing to be the only bidder in an auction. We did this by taking two of our six group members in each round and placing one in the SB and the other in the A auction by default, allowing the other four to choose. This meant that anyone choosing to enter into an auction knew that they would face at least one other bidder. Our view was that this design choice would help to sharpen the inference.[7]

Subjects were not told of their value before choosing which auction format to enter. It was only revealed after their entry choice and if the auctions were to be conducted that round. It was, however, made clear to them that they would receive the same value regardless of which format they chose. Whether the auction round was conducted or not, they were told the number of people choosing both formats in each round. If the auctions were being conducted that round, the subjects were told how many other bidders were in the auction format they were participating in that round before the bidding started so that they would know the relevant n while formulating their bids. At the end of each round they were told whether or not they won that round, the price paid by the winner and how much profit they made.

There are several considerations that led to the decision not to inform subjects about their valuations before choosing the auction format. First is the theoretical tractability of our design as well as the fact that it is most similar to the assumptions in theoretical treatments of endogenous entry. For example, in Levin and Smith (1994) the value is not known prior to entry choice (but number of bidders is known after entering and before bidding). Dealing with the potential for cut-off equilibria in which the entry choices of bidders would depend on the value as well as some subtle updating and strategic signaling issues would have added more complexity than seemed warranted. Further, our goal for this study was to add one element (endogenous entry) to standard IPV experimental auctions and determine the effect. This design seemed to be the most straightforward way of minimizing the differences compared to past experimental designs to allow us to focus on this key issue.

Another important detail to our design is the choice the subjects have between two different auctions. An alternative design that might be considered "cleaner" in some sense would have been to allow the endogenous entry choice in our experiments to be between participating in an auction

[7] We also note that real auctioneers could have the option of placing a reserve price or not holding the auction with a small number of bidders (e.g., one or two) to avoid the possibility of auctions with 0 revenue.

of a single format versus not participating. The problem with such a design for our interests is that a critical component to that design is the structure of the outside option given to subjects who choose not to participate in a given round. Whether the outside option is a fixed sum of money or a lottery could be quite important as would the amount of money available in both and the variance in the second. Our view was that the most field relevant way to design that outside option would be to have its value also be determined endogenously in the experiment by allowing them to choose another auction. A prospective eBay bidder wishing to procure a particular object could choose to enter an eBay auction for it while his alternatives could include choosing a different eBay auction, an Amazon auction, an auction on some specialty auction site or purchasing from a retail outlet. While we could not allow for this full range of options, we believe our design captures the key element for our interest because our interest is whether or not one auctioneer can pull bidders away from another auctioneer based solely on the rules of the auction format he or she chooses. Our design also matches with the description of the formation of early internet auction markets described in Lucking-Reiley (1999b) in which sellers would be selling certain types of trading cards and different sellers would be differentiated largely by the mechanism they selected for use and in fact many of the sellers gave potential bidders a choice almost identical to the one we give in the experiment. So while our design may not capture the exact nature of the choice faced by current eBay bidders, it does come closer to capturing the choice bidders faced when online auction markets were forming and is therefore quite empirically relevant for trying to understand why over time sellers appear to have chosen to use variations on the ascending format much more often than the sealed bid format.

All experiments were conducted at Technical University Berlin. The software for the experiments was programmed using z-Tree (Fischbacher (2007)). Earnings from the experiment were translated into Euros at the exchange rate of 1 ECU = € 0.10. Subjects' total earnings ranged from € 3.00 to € 39.50 with an average of € 22.22 in sessions that lasted approximately 2 hours.[8]

3 Results

The first prediction we wish to test is the one based upon our work in Ivanova-Stenzel and Salmon (2004a) stating that we should observe bidders willing to choose larger ascending rather than smaller sealed bid auctions. The outcome of this test leads to our first result.

[8]These numbers include a show up fee, i.e., starting capital, of €3.00.

	Auction	1-A	1-B	2-A	2-B	3-A	3-B	4-A	4-B	Average
Number of Bidders (Full)	**SB**	2.73	2.60	2.67	2.57	2.37	2.23	2.23	2.10	2.44
	A	3.27	3.40	3.33	3.43	3.63	3.77	3.77	3.90	3.56
Number of Bidders (Part)	**SB**	2.50	2.75	2.80	2.60	2.80	2.80	2.25	2.50	2.63
	A	3.50	3.25	3.20	3.40	3.20	3.20	3.75	3.50	3.38
Revenue	**SB**	47.08	61.50	41.20	59.40	57.00	59.20	41.25	49.00	51.95
	A	50.42	52.00	71.60	59.00	54.60	41.80	40.00	69.75	54.90
Winner's Surplus	**SB**	13.50	12.33	8.80	5.00	8.20	15.20	20.00	12.50	11.94
	A	20.08	23.08	19.60	21.40	19.80	35.40	21.00	5.50	20.73
Avg. Bidder Earnings	**SB**	8.15	5.26	3.07	1.87	3.42	5.93	14.81	5.19	5.96
	A	6.37	7.00	5.92	6.82	6.20	11.77	6.31	1.44	6.48
Efficiency	**SB**	0.974	0.998	0.935	0.994	1.00	1.00	1.00	1.00	0.988
	A	1.00	0.986	1.00	1.00	1.00	1.00	1.00	1.00	0.998

Table 1: Summary of key descriptive statistics.

Result 1: The ascending auction attracted more bidders than the sealed bid auction.

The first piece of evidence to support this result can be found in the top four rows of table 1 which contains summary statistics regarding the average size of each auction as well as the average revenue, surplus/earnings and efficiency for each subject group. We ran four sessions of the experiment and in each session we divided subjects into two independent groups of size six with each group participating in 30 rounds of the auction choice game (there are 240 total periods of observation with 30 for each of the eight groups). Thus, to use the most conservative approach to examining our results we can use the average auction size, revenue and so forth over a session from each of those groups as a single observation giving us eight observations. Table 1 presents the averages for each group for each measure and therefore it contains every observation of data necessary for tests of this nature. While this delivers independence between groups, the observations regarding the ascending and sealed bid auctions are not, however, independent from each other inside of a group. This can be dealt with by using paired tests to determine if there are statistically significant differences in the related observations.

Note that we observe revenue, surplus and efficiency only in those periods in which auctions were conducted while we observe the number of bidders choosing each auction format in all rounds. Table 1 shows data for the auction measures obviously only from the rounds in which auctions were conducted (52 out of 240) while we have shown two versions of the statistics concerning the number of bidders. The first and second rows of table 1 contain data based upon the full set of rounds (240) while the third and forth rows show results only from those rounds in which auctions were conducted (52). Both sets of results on the number of bidders in each auction show that for each group, the average size of the ascending auctions was larger than the average size of the sealed bid auctions. Thus, even with the relatively small sample size, the results of a Wilcoxon signed rank sum test are a foregone conclusion. The actual Wilcoxon tests lead to $p-$values for either set of data of around 0.01 while paired t-tests regarding the hypothesis that the means of the two distributions are equal deliver a result with the same interpretation yielding p-values of <0.01 for either set of data points.

Figure 1 shows an alternative way of looking at the number of bidders choosing the two different auction formats as it contains histograms of the number of periods with auctions of size 1-5. Recall that in each round, one bidder was forced to be in the SB and one in the A auction. Thus, the smallest auction possible was with $n = 1$ and this involves not a single bidder choosing that format on their own. There were no A auctions with only the single designated bidder while there were 34 SB auctions (14 %) that did not attract a single additional bidder. Overall there were 240 auction pairs resulting from entry choices. Of these, 30 pairs resulted in more bidders in the SB than in the A auction (12.5 %), 79 pairs resulted in the even 3/3 split (33%), and the other 131 pairs resulted in the A auction having more bidders than the SB auction (54.5%). These results quite clearly show that many of the subjects desired to be in the A auction even if it meant competing against a larger number of opponents.

It is possible, though, that most subjects chose the A auction early on but then began choosing the SB auction more often once they saw that there were typically more bidders in the A than the SB auction. To examine this issue, we look at whether the frequency with which subjects choose the A versus SB auction changed over time. We conducted a simple linear fixed effects panel regression using each group as the unit of observation with the size of the A auction as the dependent variable and period or the natural log of the period (to allow for a nonlinear effect) as the independent variable. In the first regression we get a coefficient of 0.006 on period with a standard deviation of 0.006 and a p-value of 0.358 while using the natural log of the period gives a coefficient of 0.002, standard deviation of 0.0672 and a p-value of 0.971. Alternatively, we could compare the distribution of the size of the A auctions in the first 10 periods to the last ten periods using the 8 groups as the units of observation again. The mean size of an A auction in the first 10 rounds is 3.525 and it is 3.6875 in the last 10. A paired t-test regarding the equality of these means delivers a test statistic of -1.926 and a p-value of 0.096 while a Wilcoxon signed rank test yields a Z-score of -1.625 and p-value of 0.104. These results show that the preference for the A auction is robust over time and does not degrade when subjects learn that choosing the A auction typically

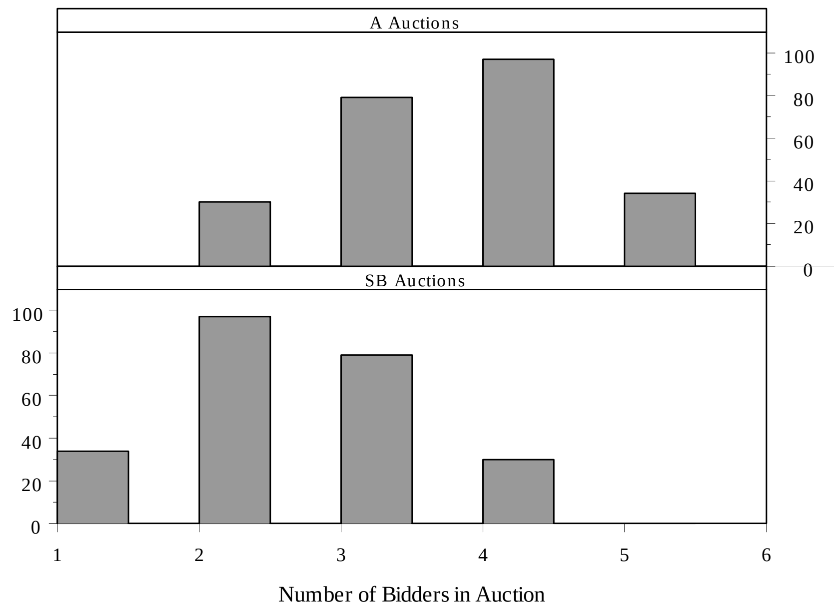

Figure 1: Histograms of number of potential A and SB auctions.

means choosing the auction with higher n.

With this result established we can now move to the result of key importance for auction designers.

Result 2: The revenue generated by both formats is not statistically different.

A visual demonstration of the similarity in the revenue generated by both formats can be found in the two panels on the left side of figure 2 which show density plots of the distribution of auction revenues from both auction formats and of the distribution of the difference between the paired observations. The two revenue distributions themselves appear quite similar and the distribution of paired differences appears approximately normal with a mean and median at approximately zero. The aggregated summaries of the underlying data can be found in table 1. The fifth and sixth rows contain the average revenue across all auctions in a session from each independent group. Notice that in about half of the cases the SB auction raises more while in the other half the A auction raises more. The overall average revenues are 51.95 for the SB and 54.90 for the A auction.[9] The most conservative tests regarding the similarity between these data involve using each of those groups

[9] For completeness we note that given the auctions conducted were bidders bidding according to the RN equilibrium, the overall averages would have been 38.40 for the SB and 53.66 for the A auction. This shows the standard result that bidders bid higher than predicted in the SB but as predicted in the A auction.

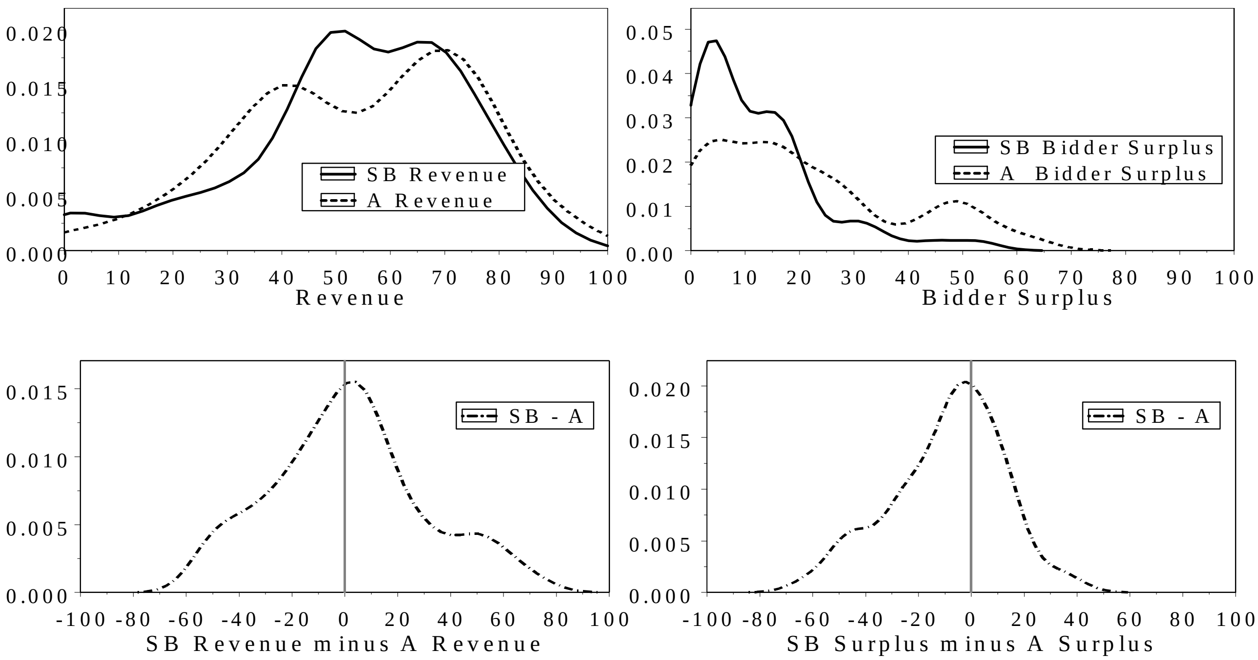

Figure 2: Density plots of the distributions of revenue, surplus, and the revenue / surplus difference between each set of paired observations.

level averages as the observations. Figure 2 provides evidence that the underlying distribution (at least at the individual period level) is close enough to normal so that standard tests are well specified. A paired t-test regarding the null hypothesis that the means of the two distributions are equivalent using the group averages as observations results in a p-value of 0.61 while a similar Wilcoxon signed-rank test results in a p-value of 0.95. These tests suggest strongly that there is no statistically significant difference between the revenue distributions.

Given the small number of observations at the group level, there is obviously a concern that the inability to reject the null hypothesis is due to low power of the test rather than a true equality of the distributions. That is certainly always a concern with experimental data. Our main point in regard to this is the fact that the conventional wisdom from prior literature is that we should expect the SB to raise more revenue than the A auction. Our results show the A auction generating more money than the SB auction and while that difference is not significant, the result is at least strong enough to call the conventional wisdom of SB revenue superiority into doubt.

Additional tests can be done to lend further support to the claim that the result is not just from a small sample size. There are at least three ways one might restructure the simple distribution tests as a means of decreasing the standard errors involved in the test to make rejecting the null hypothesis more likely. The first might be to conduct the same tests using each period as an observation rather than aggregating at the session level. This is the least conservative test of the hypothesis as it involves maximizing the relevant sample size at the cost of violating the independence assumption across observations. These tests generate a $p-$value of 0.808 for a paired $t-$test and $p-$value of 0.778 for a Wilcoxon signed rank-test. So this least conservative test also generates strong results suggesting that the differences between the means and medians of the distributions are not statistically significantly different. Intermediate approaches involve conducting OLS or fixed effect panel regressions of revenue on the type of auction and time measures defining the panel by the independent groups and clustering the standard errors by the groups. These approaches will take advantage of the individual observations in the data but will also take into account the lack of independence of the multiple observations per group. The results of such regressions are a foregone conclusion though from the results of the prior two tests and so we will refrain from including the full set of regression results on these specifications to conserve space. According to a broad range of specifications in this vein we used, the coefficient on the dummy variable for the type of auction is always small in magnitude and highly insignificant.[10] While all of these tests can not prove the lack of significance result is not the result of type 2 error, the fact that all tests conducted show the same unambiguous result should suggest that the result is quite robust and is not simply due to the small sample.

The simple distribution tests and uncomplicated regressions we use here are particularly appropriate given the issue we are addressing. We could, for example construct a more complex regression analysis with revenue as a function of the auction format, n, and other characteristics of bidders to correct for any composition differences we see in the auction formats. We are, however, specifically uninterested in testing the revenue difference between these mechanisms in such a *ceteris paribus* structure because we are interested in comparing them when these things are not held equal but rather allowed to vary endogenously. Further, these other issues regarding how auction revenue changes based upon n and other characteristics are well understood from previous literature.

We can also again check to see whether this relationship changes over time and the answer is no. Regressing the difference in revenue between the A and SB auctions and conducting fixed effect panel regressions with period and then $\ln(period)$ as explanatory variables again leads to results showing a lack of statistical significance; coefficient on period is -0.435 with standard deviation of 0.495 and p-value of 0.384 while the coefficient on $\ln(period)$ is -7.416 with a standard deviation of 5.904 and p-value of 0.216. We also find a lack of significance conducting distribution tests regarding

[10] Full results are available from the authors upon request.

the revenue differential between the first ten and last ten periods; paired $t-$test yields t-statistic of 0.511 and p-value of 0.631 while the Wilcoxon signed rank test yields a Z-score of 0.524 and a p-value of 0.600. These tests of differences over time are problematic though as the data is not very dense due to our not running auctions in all periods. Since the number of auctions conducted was not that large we should not expect any change of bidding behavior based on time. On the other hand, a selection bias could enter in if certain types of bidders gravitate towards separate auction formats over time. In any event, we observe no time trends in revenue.

While revenue equivalence between these institutions does seem to be obtained by allowing endogenous entry that does not mean that all of the claims of the Revenue Equivalence Theorem are satisfied. In our setting, n is not constant and therefore equivalence in revenue does not automatically imply equivalence in bidder surplus. Before comparing bidder surplus across mechanisms, we next establish that one of the classic assumptions in the RET turns out to be valid which is that both formats are approximately equally efficient.

Result 3: There is no statistically significant difference in the efficiency between the two formats

Once again, the group level data is presented in table 1. What we find is that virtually all of our auctions were 100% efficient, including the sealed bid auctions. Paired test results regarding the hypothesis that the means of both distributions are equal result in a t-statistic of 1.24 with p-value of 0.26 and a Wilcoxon signed rank Z-statistic of 1.46 with a p-value of 0.14. Thus, we can not reject the hypothesis that these two distributions are equivalent. This is contrary to the results found in other studies beginning with Cox, Roberson, and Smith (1982) showing that sealed bid auctions are typically less efficient than ascending. It is important to note, however, that our measure of efficiency is the classical measure of the ratio of the value of the bidder who won the auction to the value of the bidder with the highest value. Thus, both formats being equally efficient does not imply that they generate equivalent total surplus. Since there are more bidders in the ascending auctions, this means there must be more total surplus in them. This can be easily seen in table 1 because total surplus is equal to revenue plus the surplus of the winner and since the latter is much larger in the ascending auctions, total surplus is obviously larger as well. This result shows that the total surplus realized in both auctions is approximately the same percentage of the total surplus possible to be achieved.

Our next result attempts to verify the last component of the RET, that bidders earn equal surplus across mechanisms, but this is complicated by the nature of the experiment.

Result 4: Winning bidders earn on average more money in the ascending auctions than sealed bid auctions. Average surplus earned by all bidders is slightly higher in the ascending than sealed bid auctions but the difference is insignificant.

Table 1 again shows the group averages for bidder surplus in the two auction formats. In all but one pair the surplus earned by the winners is greater in the A than in the SB auction and paired statistical tests will show that this difference is clearly statistically significant. This is not a surprising result because we have previously established that prices are about the same between mechanisms but there are more bidders in the A auction. The latter fact means that the expected highest value in the groups participating in the A auction must be greater than the expected highest value of those participating in the SB auction; actual averages of the highest values in each format are 75.75 for the A auction and 64.80 for the SB auction. Because we have also found that the auctions are highly efficient, the average values of the winners must be very close to these levels and so it must also be the case that the winners in the A auction earn more surplus. This does not mean, however, that bidders earn more on average or in expectation. This is so because of the fact that the number of competitors in each format is different. We can take one step towards more accurately determining what is going on by looking at average profits of bidders in both mechanisms. Average profit is calculated by taking each auction and dividing the realized surplus of the winner by the number of bidders in that specific auction to yield a properly weighted average profit measure. Table 1 also contains the average bidder earnings for each of the eight groups as separate samples. The average earnings are 6.48 and 5.96 for the A and SB auctions. This suggests a slight advantage for the A auction but this difference is not statistically significant. A paired $t-$test on the group level data delivers a $p-$value of 0.772 while a Wilcoxon signed rank test yields a $p-$value of 0.742. The two panels on the right side of figure 2 show distribution plots of the earnings in each format as well as the difference between the paired observations.

We should note that this result is affected by an important choice in regard to which auction rounds of data to consider. In many respects it would be reasonable to remove from consideration the rounds in which only one bidder was in the SB auction as they count for 0 revenue and very high bidder surplus. These data points are certainly outliers and even though there are few of them in the data (only two of these rounds were actually played) they are strong enough to effect these bidder earning results (though they have no effect on any of the other results in the paper). If we remove these data points, the average earnings are 6.78 and 4.29. These differences are larger and large enough to lead to a borderline significant difference. A paired $t-$test delivers a

$p-$value of 0.045 while a Wilcoxon signed rank test yields a $p-$value of 0.109. Which of these two representations of the data is the more reasonable is a debatable issue. Our interpretation of the results is that in either case there is perhaps a small average earnings difference in favor of the A auction but it was likely small enough such that subjects may not have been able to detect it.

Were the numbers of bidders in both formats equivalent, then looking at either the surplus of the winners or the average surplus would be sufficient to allow us to conclude that the expected surplus is equivalent across institutions. While looking at these averages by themselves is interesting, ultimately what we want to do is obtain a more accurate measure of expected earnings which involves conditioning on the value of the bidder. Thus, we want to provide an answer to the question of which format should generate higher earnings conditioned on the value v a bidder possesses.

In the process of calculating a more careful expected utility value for our potential bidders, we need to account for the fact that bidders do not bid according to the risk neutral equilibrium which could mean that they might possess utility functions other than risk neutrality. One of the standard explanations for the "overbidding" observed in the first price auctions is risk aversion and in calculating the expected utility of participating in either auction we will make our calculations general enough to also allow for an investigation of how such preferences might impact the results. A key detail in such an analysis is whether or not there are selection effects due to risk preferences or more generally propensity to overbid in the SB auction driving the auction choice behavior. There are a number of reasons why risk aversion or propensity to overbid might lead to self-selection into auction formats. One argument is based on the notion that risk averse bidders should be more willing to enter into the SB auction. This is so because while they may expect roughly equal or maybe even lower profit in the SB than the A auction, the earnings in the sealed bid auction are more stable which is attractive to risk averse bidders. Additionally, with the smaller n in the SB auction, they will win more often than in the A auction making some surplus more likely. If this selection occurs then this could have an impact on how we model the expected opponent bidders might face in SB auctions. Based upon our prior experimental results in Ivanova-Stenzel and Salmon (2004a), we showed that bidders who are more risk averse than average should actually be willing to pay slightly more to enter an A auction than bidders who are closer to being risk neutral. This was due to the extra stability in earnings in the SB auction not making up for the substantially lower earnings in the SB relative to the A auction. However, the difference we found was not large. Therefore, our hypothesis is that if the overbidding is due to risk aversion then we should not see much correlation between risk aversion and auction choice.

Selection effects could also occur if subjects are risk neutral but are heterogeneous in their propensity to overbid for non-risk aversion related reasons. A bidder who regularly bids well above the RNNE level in the SB auction will find himself to be making little money in those auctions and the fact that his profits are more stable should not make up at all for the expected profit gap between the A and SB auction. Thus, it would be reasonable to expect that such bidders would avoid the SB auctions and instead participate in the more lucrative (especially for them) A auctions. A bidder who bids closer to RN might well have little reason to avoid the SB auctions as when she wins, her profit is quite large. Such behavior would lead to a positive correlation between degree of overbidding and likelihood of choosing the ascending auction. This hypothesis is complicated by the fact that if the bidder is risk neutral then there must be an alternative explanation for the overbidding. The exact nature of that explanation could lead to a different prediction. While examining this selection issue, we can simultaneously look more generally at whether or not we can identify the determinants of a bidder's choice of an auction mechanism.

Result 5: Neither risk aversion nor more generally the degree of overbidding has a statistically significant effect on auction choice. Most experience based variables also have no statistically significant effect on auction choice.

All we observe about the bidders prior to their choices of auctions is their results from the learning phase. Thus, we will be looking at whether or not experience in that phase impacted auction choice. The chief characteristic we can observe about an individual is their risk aversion or degree of overbidding. We will use in one of our specifications a subject's estimated risk aversion parameter assuming a CRRA utility function of $u(x) = x^{\alpha}$. This delivers a bid function of $b^{*}(v_i) = \frac{(n-1)v_i}{\alpha+(n-1)}$ which suggests a regression of $b_{it} = \beta_{i1} + \beta_{i2}v_{it} + \varepsilon_{it}$ to capture the degree of risk aversion of the bidder. In our regressions regarding auction choice we will include $\alpha_i = \frac{1}{\beta_{i2}}(n_{it} - 1) - n_{it} + 1$ which can be interpreted as a measure of risk aversion or more generally as a measure of overbidding.[11] In our second specification we use β_{i2} itself as a direct measure of overbidding that is independent of a specification of a utility function.[12] Both measures, α_i and the average of the β_{i2} values, should and do yield highly similar outcomes.

[11] To estimate the risk aversion parameter we used all of the sealed bid auctions from the learning phase which consists of 6 auctions with $n = 2$ and 6 with $n = 4$. We also ran the regressions using risk aversion parameters estimated for each size n separately as well as the average of those two estimates. There was not a substantial difference in the final results.

[12] The measure we actually used is the average of the values found for β_{i2} when using only the data from $n = 2$ auctions and then using only the data from $n = 4$ auctions.

Before presenting the regression results it is worth describing in broad terms the nature of the choice paths of the subjects. Subjects had 20 rounds each when they were allowed to choose between A and SB auctions (in the remaining 10 rounds they were used as the default bidder in either the A or SB auction). On average subjects chose the A auction 12.8 times out of 20 and the SB auction 7.2. Further, 33 out of the 48 subjects (69%) chose the A auction more often than the SB auction and 7 subjects (15%) chose the A auction exclusively. No subjects chose the SB auction exclusively but there was 1 subject who chose it 18 times and 2 who chose it 17 (14 subjects though chose the A auction 17 or more times). This is simply further confirmation of the fact that the A auction was the more popular of the two. The econometric implication of this fact is that almost one third of the subjects had choices with too little variability to be explainable as anything other than pure preference for the A auction. The other subjects were switching between formats in a manner that might suggest their preference for the A auction is weaker or at least that there were conditions that would lead to them choosing the SB auction. To determine what those conditions are, we can conduct both cross section and panel regressions to try to understand the factors that determined those choices.

Our first approach will be a simple cross section regression using each subject as a separate data point. The dependent variable is the fraction of times a subject chose the ascending auction.[13] The independent variables are various measures from the learning phase. To make the estimation more parsimonious and to focus on what is more likely to be a driving factor in a subject's decision making, all variables except the risk/overbidding parameter are differences between the relevant values in the A and in the SB auction. Note that we do have separate variables for the $n = 2$ and $n = 4$ auction periods. These variables include NxProfit which is the average profit the subject made in the A periods minus the profit he made in the SB periods in the $n = x$ periods, NxWin which is the difference in the win percentages for that subject in the $n = x$ periods and then the session level average profit differentials, NxSesProfit. We have also computed the regression with actual levels instead of differences and we end up with the same interpretation.

The results of both specifications of this regression are found in table 2. There is only one variable that is significant at the 5% level in both specifications, N2Win, and it has a positive coefficient. This suggests that if a subject wins in the A more often than in the SB auction in $n = 2$ periods, then he is more likely to choose the ascending later on. There are a few other

[13] We have conducted similar regressions using a count variable instead of a percentage variable and the results are no different. A Poisson regression using the count variable also leads to no difference in the conclusions.

	Specification 1				Specification 2			
Parameter	**Coeff**	***p*-Value**	**Coeff**	***p*-Value**	**Coeff**	***p*-Value**	**Coeff**	***p*-Value**
RA	0.105	0.487	0.09	0.475	-	-	-	-
Avg β_{i2}	-	-	-	-	0.037	0.887	0.31	0.195
N2Profit	-0.001	0.145	-	-	-0.002	0.052	-	-
N2Win	0.733	0.009	-	-	0.768	0.006	-	-
N4Profit	-0.004	0.059	-	-	-0.004	0.086	-	-
N4Win	-0.492	0.074	-	-	-0.406	0.141	-	-
N2SesProfit	-0.000	0.918	-	-	-0.000	0.781	-	-
N4SesProfit	0.003	0.143	-	-	0.002	0.160	-	-
Constant	0.542	0.125	0.61	<0.01	0.037	0.887	0.401	0.035
Obs(Groups)=464(48)	$\overline{R}^2 = 0.171$		$\overline{R}^2 = -0.01$		$\overline{R}^2 = 0.161$		$\overline{R}^2 = 0.02$	
	$Prob > F = 0.039$		$Prob > F = 0.475$		$Prob > F = 0.047$		$Prob > F = 0.195$	

Table 2: Cross section regressions with fraction of times a subject has chosen the ascending auction as the dependent variable.

variables that might be considered borderline significant, but they are small in magnitude. Further, the overall explanatory power of both regressions is quite low with $\overline{R}^2 = 0.171$ and $\overline{R}^2 = 0.161$ respectively. A legitimate concern with either of these specifications is that there may be a great deal of colinearity between several of these variables. We have conducted regressions with subsets of these variables and find no difference in the outcomes. Therefore, we chose the more compact presentation. Due to the importance of the results regarding risk aversion/overbidding we also included in table 2 the results of the regressions with those variables by themselves. The coefficients on risk preferences/degree of overbidding are still insignificant.

The results from this regression match with those in Ivanova-Stenzel and Salmon (2004a) showing that preferences between auction formats do not seem to be influenced much by outcomes in the learning stage. Further, neither measure of risk aversion nor degree of overbidding influences auction choice behavior. This suggests that not only is there no self-selection into auction formats according to risk aversion but also that there is no self-selection based upon some other characteristic associated with overbidding. This is consistent with our hypothesis that the overbidding may be derived from risk aversion and inconsistent with at least one naive hypothesis about how auction choice behavior should depend on overbidding if the behavior is not driven by risk aversion.

We can alternatively look at the subjects' choices between formats in a time series manner by examining their individual choices across time using panel techniques. However, given the design of our experiment (i.e., the fact that auctions were not conducted in all periods) it is difficult to conduct such regressions due to the sparseness of the data. In most of the 30 periods of the experiment, the only data we observe is the number of subjects choosing each auction format. If we conduct a panel logit regression with fixed effects to correct for any subject level differences using the choice of format as the dependent variable (1=A, 0=SB) then we get a coefficient of -0.070 on the number of people in the ascending auction last period with a standard error of 0.059 and a p-value of 0.233. If we include variables regarding their experience in the auctions conducted (i.e. whether they won in period $t-1$, profit in $t-1$ etc. . .) the data becomes very sparse due to all of the missing variables causing interpretation of the coefficients to be suspect. Given that, when conducting such regressions we find that all variables are insignificant except the previous period's choice. We conclude from this analysis that the subjects have some intrinsic preference for one format over another represented by the fixed effect or constant in the regression and even beyond that there is some amount of inertia in regard to their choices. This analysis does not reveal any other strong factors which figure into a subject's decision of which format to enter. This may well indicate that most subjects were following a mixed strategy for entering into each auction as might be reasonable to avoid situations in which everyone simultaneously switches from one crowded auction format to the other.[14] The probabilities for those mixed strategies were not 50/50 though and these results are strong enough to show that even the potential mixed strategy propensities used by individuals were not correlated with overbidding or risk aversion.

To calculate the expected utilities of the bidders we will again assume a CRRA utility function or $u(x) = x^\alpha$ with $\alpha \in (0, 1]$ where $\alpha = 1$ is risk neutrality and $\alpha < 1$ is risk aversion and calculate the expected utility for choosing each auction format for all $\alpha \in \{0.01, 0.02..., 0.99, 1\}$. The equations for performing these calculations can be found in the appendix. The results are presented in figure 3 under two different specifications of the underlying data. The left diagram presents the results in which we assumed the distribution of auction sizes corresponded to that derived from the full set of 240 choice periods. The diagram on the right draws the auction distribution from only those rounds in which an auction was conducted. Notice that for the diagram derived from all rounds of data, the expected utility line for the SB is above that of the A for all possible α while for the diagram using the data from the conducted auctions the expected utilities are virtually equivalent for all α. The difference between the two diagrams is derived from the fact that while 32 (14%) of the 240 total rounds had one bidder in the SB, by chance only 2 (4%) of the 52 rounds in

[14] Smith and Levin (2002), Ochs (1990) as well as Meyer, Van Huyck, Battalio, and Saving (1992), who examine coordination problems based on individual entry decisions in different locations taken by subjects operating in decentralized markets, show that observed variations in the aggregate number of entrants across locations are characterized quite well by the mixed strategy symmetric equilibrium.

which auctions were conducted had one bidder in the SB instead of the expected 6-7. Due to the very high bidder earnings in those situations, this small under representation of these periods in the conducted auctions has an outsized effect on the EU representation. This also explains the counterintuitive nature of the expected utility comparison in the panel using all of the data which shows that even for the risk neutral individuals, they would have an expected utility benefit from participating in the SB auction even though average earnings as presented before suggests the opposite. The reason is that the *potential* earnings in those periods in which a bidder is alone in the SB auction and receives his full value as earnings are very high. These rounds make the SB quite attractive in theory, but due to the small underrepresentation of those rounds in the realized auctions the actual earnings did not reflect this potential.[15]

It is not entirely clear which representation of earnings subjects are most likely to consider when making their entry decisions. They would of course calculate neither explicitly but they would have an intuitive knowledge of the likely distribution of auction sizes and if that were the basis for the decision then the left diagram is the better representation. If they were responding to past profits then perhaps the right diagram is the better representation. In either case, the implication is that the difference in the expected utilities is not large though one might argue for a slight advantage for the SB. The fact that the relationship between the expected utilities is the same for all α is consistent with the fact that we do not observe auction choice to be correlated with α. Finally, if one believes that the most appropriate representation is the one in the left panel then there might be a suggestion of "over-entry" into the A auction perhaps due to a preference for the less complicated auction format.

It is of course important to note that this comparison is only valid under the maintained assumptions that bidders are bidding as if they are risk averse, that they are deciding between auction formats under the same "as-if" risk aversion parameter and assuming that the CRRA utility function is a reasonable representation of their preferences. Different functional specifications for the risk aversion could lead to different conclusions and if bidders do not possess risk averse preferences then the entire analysis does not apply. We provide this analysis mainly as an illustration, but there are reasons to suspect that it is a reasonable approach to modeling the behavior. As mentioned before the CRRA utility function is indicated as reasonable because the general nature of the auction choice behavior matches with the predictions under the DARA assumption established by

[15] It is also important to point out that no subject could choose the SB and end up in the format by himself due to our design. So the existence of these periods with 1 bidder SB auctions should not have made it more attractive to choose the SB, but rather they made it more attractive to randomly be placed in it.

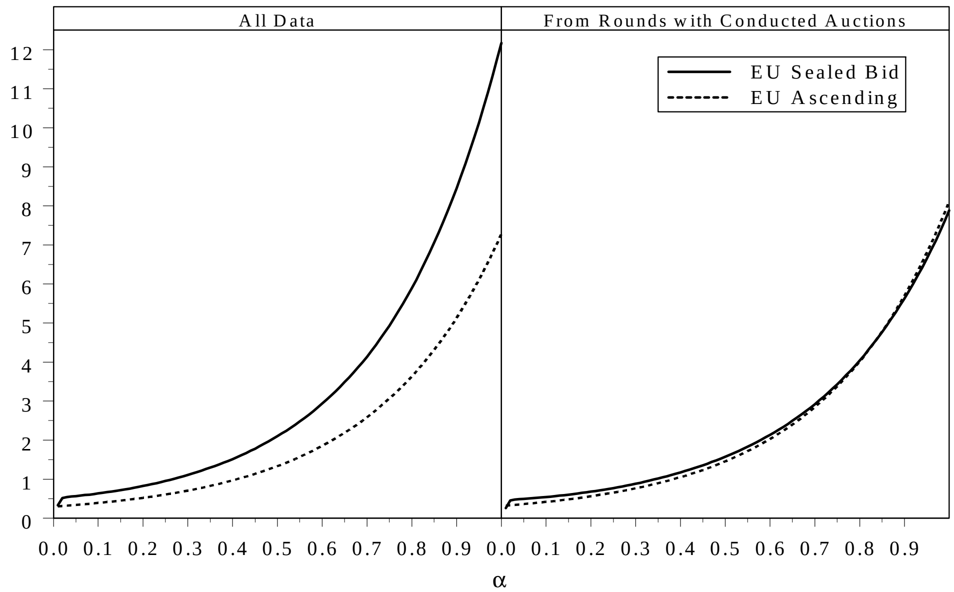

Figure 3: Bidders' expected utility for any risk preference.

Matthews (1987).[16] While there are other utility functions also satisfying DARA that one could choose, CRRA also seems the most appropriate to use given its common use in the literature as a standard benchmark for risk averse behavior.

There are several approaches one can take in rationalizing all of the results discussed here. One approach involves rejecting the idea that risk aversion accounts for the bidding behavior in the first price auctions. The auction choice behavior observed here could be explained by arguing that subjects are basically risk neutral and are arbitraging between the two formats. The slight advantage for the A auction may not have been large enough to be detectable in the relatively short time horizon for earnings to be completely equalized. The main deficiency here is that one is left without an explanation for the overbidding in the first price auction. Depending on how that issue is dealt with there may be a puzzle regarding why those bidders who overbid by more in the SB auction than others were not self-selecting out of it and we are left without an explanation for the behavior observed in Ivanova-Stenzel and Salmon (2004a) which is inconsistent with this arbitrage story.

If one is willing to accept that risk aversion can at least partially account for the bidding behavior in first price auctions and that people have a non-pecuniary preference for the ascending auction, we can provide an explanation for all of the observed behavior. The results presented above match

[16] The CRRA utility function, of course, satisfies DARA.

quite well with what was predicted in Ivanova-Stenzel and Salmon (2004a). We observe that most subjects are willing to choose an ascending auction instead of a sealed bid auction even if this might lead to losing some degree of expected utility. The reason for this choice is easily explainable based on the difference in mental costs of participating in sealed bid auctions relative to ascending or minimizing the possibility of regret.[17]

4 Conclusion

The results from this paper clearly show that the revenue ranking of sealed bid and ascending auctions for field applications is not necessarily as resolved an issue as one might think from reading most of the past experimental literature. Our results show that entry decisions are quite important to this issue and ignoring them can lead to a distorted view of which auction mechanism is likely to raise the most revenue. In our specific environment by allowing subjects to freely choose to enter either a sealed bid or an ascending auction, we found results showing that the two mechanisms result in approximately equal revenue to the sellers as well as equal efficiency, and a small advantage for the ascending auction in regard to bidder surplus. This provides at least a partial resolution to the question of why auctioneers continue to use ascending auctions in the field despite the prior evidence suggesting first price auctions generate more revenue.

The environment we chose to examine in our experiments is, of course, highly stylized and it is important to be clear on what our results suggest about what one should expect to find in naturally occurring markets. It is clear from our results that bidders were arbitraging to some degree between these two formats. That accounts for the approximate equality in average earnings / expected utility across mechanisms. This is behavior exactly in line with what economists expect to see in market settings and it seems quite reasonable to expect that this observed tendency in our data should generalize outside of the laboratory. The fact that this behavior also leads to equal revenue in both formats is more difficult to generalize due to the fact that we use only a single pool size and a single value distribution. In the exploratory results in Engelbrecht-Wiggans and Katok (2005), they have a few sessions with larger pools (one with 20 bidders) and find a comparable result. This suggests our finding is broader than just in our environment.

The key result in our study is not the approximate revenue equivalence but rather the fact that

[17] Regret aversion has been recently gaining traction in its own right as an explanation for overbidding. Engelbrecht-Wiggans and Katok (2007) present a fairly general theory of regret in auctions and show how it can organize a great deal of existing data. Regret aversion could also provide an explanation for the observed auction choice behavior – subjects prefer ascending auctions because there is never any regret in them.

we have shown that the revenue in the ascending auction can rise high enough when endogenous entry is allowed to call into question the assumed revenue superiority of the sealed bid auction. We argue that this result is driven by subjects arbitraging between mechanisms on the basis of expected utility comparisons but with a component of this calculation due to some disutility for participating in first price sealed bid auctions. The reason why this behavior seems to robustly lead to approximate revenue equivalence is still unclear. That result may well be an artifact of the uniform value distribution as that is used in both this study and Engelbrecht-Wiggans and Katok (2005). It is certainly possible that in other environments the ascending auction could raise more money while in others the first price may raise more. This indicates a clear need for a more general theory of endogenous auction choice expanding on Smith and Levin (1996) that can better characterize when we should expect this approximate revenue equivalence and when we should expect one format to raise more money than the other. What is clear though is that the first price sealed bid auction is not necessarily revenue dominant and that endogenous entry is a key component in revenue comparisons that should not be overlooked.

Another avenue to investigating this issue might involve examining a bidder's choice between entering an auction of a specific type versus accepting some outside option as this is an alternative characterization of the decision faced by participants in online auctions. We chose not to conduct the experiments in this manner mostly due to the difficulty in properly calibrating the outside option to generate a "fair" comparison. However, there is reason to expect that the effect we provide evidence for will be robust in this alternate setting. First, if bidders possess a preference for the ascending auction, then this preference should transfer leading to an edge for the ascending auction in attracting bidders away from any given outside option. Second, Pevnitskaya (2004) and Palfrey and Pevnitskaya (2008) show both theoretically and using experimental data that when faced with the choice of a certain outside option (e.g. buying the item in a store at a known price) versus a first price auction, the more risk averse bidders will choose the outside option with greater likelihood. Because these are the bidders driving the high revenues in first price auctions, the result is that the outside option blunts the only advantage of the first price auction.[18] Consequently, there is every reason to suspect that our result would be robust to allowing the choice to be participating

[18] There is another paper on outside options that finds the opposite result, Kirkchamp, Poen, and Reiss (2005), but the implementation of the outside option is quite different. In Kirkchamp, Poen, and Reiss (2005), the outside option is not chosen by the bidder as an alternative to participating in the auction but rather is implemented as an amount earned if the bidder loses the auction. This is found to increase overbidding in the first price auction. The issue that paper does not investigate is the important entry decision of the bidders which will lead to n being variable. Thus it is not clear that their findings will translate into an environment when entry in an auction is a choice.

in an auction versus choosing an outside option.

APPENDIX A: Computation of Expected Utilities

To calculate the expected utility in both formats we used equations 1 and 2. To calculate the expected utility in the first price auction, we have to sum over all possible sizes of auctions, $1-5$, adding the expected utility of participating in an auction of each size multiplied by the probability of the auction being of that size. This latter probability is obtained from the data as the empirical percentage of auctions conducted of each size and given by $g(n)$. For an auction of any specific size, we sum over all possible values the bidder might draw the surplus the bidder makes if he wins multiplied by his probability of winning with all of this multiplied by the probability of drawing that particular value. We are calculating the probability of any bid winning using the empirical bid distribution from all auctions conducted of each size which means we are taking no stand in regard to the process generating the competing bids. Thus, for any bid x we calculate the probability that the bidder draws $n-1$ opponent bids from the empirical distribution that are less than x. This is given by $B(x,n)^{n-1}$. We are assuming for simplicity that the bid function used by the risk averse bidder is the standard linear bid function for a risk averse bidder without the correction for the non-linear "hook" at high values derived from the heterogeneous risk preferences of the opponents. Ignoring the hook on the function will lead to our underestimating the expected utility of choosing a first price auction. As our results will show though, an error in this direction will not affect our conclusions so we opted for this simpler approach for calculating expected utilities. All of this is summarized in equation 1.

$$EU^{SB} = \sum_{n=1}^{5} \left(\sum_{v=0}^{100} (v - \frac{n-1}{n-1+\alpha} v)^{\alpha} B(\frac{n-1}{n-1+\alpha} v, n)^{n-1} \frac{1}{101} \right) g(n) \quad (1)$$

To calculate the expected utilities in the A auction, we must follow a similar approach. Again we sum the expected utility of participating in an auction of size n over all possible auction sizes and multiply by the empirical fraction of auctions of each size, $h(n)$. Since bidders in A auctions do bid for the most part as theory predicts, the rest of the calculation is perfectly straightforward and requires no questionable assumptions. We sum over all possible values and then over all possible prices the bidder might pay adding the utility in each case times the probability of that price being the price paid and finally multiplying by the probability that the bidder draws that particular value. This calculation is given by equation 2.

$$EU^{A} = \sum_{n=1}^{5} \left[\sum_{v=0}^{100} \left(\sum_{p=0}^{v} (v-p)^{\alpha} (n-1) \frac{1}{101} \left(\frac{p}{101} \right)^{n-2} \right) \frac{1}{101} \right] h(n) \quad (2)$$

References

Bajari, P. and A. Hortacsu (2003). Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions. *Rand Journal of Economics 3*(2), 329–355.

Coppinger, V., V. L. Smith, and J. Titus (1980). Incentives and Behavior in English, Dutch and Sealed-Bid Auctions. *Economic Inquiry 18*, 1–20.

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.

Engelbrecht-Wiggans, R. (1993). Optimal Auctions Revisited. *Games and Economic Behavior 5*, 227–239.

Engelbrecht-Wiggans, R. and E. Katok (2005). Experiments on Auction Valuation and Endogenous Entry. In J. Morgan (Ed.), *Experimental and Behavior Economics, Advances in Applied Microeconomics*, Volume 13, pp. 171–196. Elsevier.

Engelbrecht-Wiggans, R. and E. Katok (2007). Regret in Auctions: Theory and Evidence. *Economic Theory*.

Fischbacher, U. (2007). z-Tree: Zurich Toolbox For Readymade Economic Experiments. *Experimental Economics 10*(2), 171–178.

Goeree, J. K., T. Offerman, and A. Schram. Using First-Price Auctions to Sell Heterogeneous Licenses. *International Journal of Industrial Organization*.

Güth, W., R. Ivanova-Stenzel, M. Königstein, and M. Strobel (2002). Bid Functions in Auctions and Fair Division Games: Experimental Evidence. *German Economic Review 3*(4), 461–484.

Ivanova-Stenzel, R. and T. C. Salmon (2004a). Bidder Preferences Among Auction Institutions. *Economic Inquiry 42*, 223–236.

Ivanova-Stenzel, R. and T. C. Salmon (2004b). Entry Fees and Endogenous Entry in Electronic Auctions. *Electronic Markets 14*(3), 170–177.

Ivanova-Stenzel, R. and T. C. Salmon (2008). Robustness of Bidder Preferences Among Auction Institutions. *Economic Inquiry*. Forthcoming.

Kagel, J. H. (1995). Auctions: A Survey of Experimental Research. In J. H. Kagel and A. E. Roth (Eds.), *The Handbook of Experimental Economics*, pp. 501–586. Princeton, New Jersey: Princeton University Press.

Kagel, J. H., R. M. Harstad, and D. Levin (1987). Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study. *Econometrica 55*(6), 1275–1304.

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

Kirkchamp, O., E. Poen, and J. P. Reiss (2005). Outside Options: Another Reason to Choose the First-Price Auction. Working Paper, University of Madgebourg.

Klemperer, P. (2002). What Really Matters in Auction Design. *Journal of Economic Perspectives 16*(1), 169–190.

Krishna, V. (2002). *Auction Theory.* San Diego, CA: Academic Press.

Levin, D. and J. L. Smith (1994). Equilibrium in auctions with Entry. *American Economic Review 84*(3), 585–599.

Lucking-Reiley, D. (1999a). Experimental Evidence on the Endogenous Entry of Bidders in Internet Auctions. Working Paper, University of Arizona.

Lucking-Reiley, D. (1999b). Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet. *American Economic Review 89*(5), 1063–1080.

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.

Meyer, D. J., J. B. Van Huyck, R. C. Battalio, and T. R. Saving (1992). HistoryŠs Role in Coordinating Decentralized Allocation Decisions. *Journal of Political Economy 100*(2), 292–316.

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

Ochs, J. (1990). The Coordination Problem in Decentralized Markets: An Experiment. *Quarterly Journal of Economics 105*, 549–559.

Ockenfels, A., D. Reiley, and A. Sadrieh (2006). Online Auctions. NBER Working Paper 12785.

Palfrey, T. R. and S. Pevnitskaya (2008). Endogenous Entry and Self-selection in Private Value Auctions: An Experimental Study. *Journal of Economic Behavior and Organization*. Forthcoming.

Pevnitskaya, S. (2004). Endogenous Entry in First-Price Private Value Auctions: the Self-Selection Effect. Working Paper.

Riley, J. G. and W. F. Samuelson (1981). Optimal Auctions. *American Economic Review 71*, 381–392.

Smith, J. L. and D. Levin (1996). Ranking Auctions with Risk Averse Bidders. *Journal of Economic Theory 68*, 549–561.

Smith, J. L. and D. Levin (2002). Entry Coordination in Auctions and Social Welfare: An Experimental Investigation. *International Journal of Game Theory 30*(3), 321–350.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance 41*, 8–37.