Comments Welcome

Auctions: A Survey of Experimental Research, 1995 – 2008*

John H. Kagel and Dan Levin
Department of Economics
The Ohio State University
410 Arps Hall
1945 North High Street
Columbus, OH 43210-1172

kagel.4@osu.edu levin.36@osu.edu
614-292-4812 (phone) 614-688-4239
614-292-4192 (fax)

http://www.econ.ohio-state.edu/kagel/ (Web) http://www.econ.ohio-state.edu/levin/

7/22/2008

* Research support from the National Science Foundation for our experimental auction work is gratefully acknowledged. Any opinions, findings, and conclusions or recommendations in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Auctions: A Survey of Experimental Research, 1995 – 2008

John H. Kagel and Dan Levin
Department of Economics
The Ohio State University

Introduction

The first question we faced writing this updated survey of auction experiments is how to organize it. There have been hundreds of published and/or working papers describing experimental work on auctions since the first auction survey reported in the first Handbook of Experimental Economics (Kagel, 1995) so that it is quite impossible, and not even very useful, to cover them all. The early theoretical and experimental research on auctions was restricted to simple environments with a fixed and commonly known number of bidders, each demanding a single unit.[1] Accordingly, the 1995 survey focused on the Revenue Equivalence Theorem with respect to independent-private-value (IPV) auctions, with research on common value auctions largely restricted to demonstrating the overwhelming presence of a *winner's curse*.

Section I of this survey reviews the work since then on IPV auctions. Much of this research continues to be concerned with bidding above the risk neutral Nash equilibrium (RNNE) in first-price sealed bid auctions. Section 1.1 examines experiments looking at the consistency of risk aversion in explaining this overbidding across environments and the extent to which bidders accurately "calculate" the underlying tradeoffs between probability of winning and amount earned conditional on winning. Section 1.2 reviews recent experiments designed to explain overbidding in terms of regret theory (as opposed to expected utility). Empirical economists have developed techniques for analyzing field data on auctions that are designed to uncover the underlying distribution of bidder values. Section 1.3 looks at an econometric analysis designed to investigate the validity of these techniques using experimental data, where the underlying distribution of bidder values is known, and can be compared to the implied probability distribution. Some recent work on second-price private value auctions is reported in section 1.4. Section 1.5 looks at work on auctions with asymmetric valuation structures, where *weak* and *strong* bidders compete against each other. Section 1.6 ties up some loose ends looking at the role of cash balance effects on bidding in private value auctions and outcomes in some novel bidding environments.

Section II looks at single-unit common value auctions. Sections 2.1-2.3 look at some of the comparative static predictions of the theory, including the ability of English

[1] This would involve sellers in the case of procurement auctions.

auctions to raise revenue compared to first-price sealed bid auctions, bidding in auctions with *insiders*, and behavior in *almost common-value* auctions. Section 2.4 looks at results from the closely related "takeover" game, with a focus on the implications of the results for recent theories designed to explain the winner's curse. Section 2.5 ties up some loose ends: Examining the behavior of super experienced bidders (e.g., is the persistent bidding above the RNNE a best response to a rivals who are bidding more aggressively?), looking at auctions with both common and private value elements for all bidders, the role of selection bias, demographic and ability effects on the presence of a winner's curse, (e.g.; do "smarter" subjects bid closer to Nash prediction and/or make more money?), and whether susceptibility to the winners curse is minimal for sports card traders in field settings.

Section III takes up multi-unit demand auctions – auctions in which bidders demand multiple units that may be substitute goods or complements. Much of the work here has been spurred by the Federal Communications Commission's sale of spectrum (air wave) rights, beginning in the early 1990s, and the explosion of theoretical and applied research that followed (as well as the widespread application of auctions for the sale of government owned property rights that followed). Section 3.1 looks at bidding in uniform price and Vickrey auctions for substitute goods. The experiments here are concerned with the issue of *demand reduction* in the uniform price auctions, and the ability of the Vickrey mechanism to correct for this. Section 3.2 extends the study of multi-unit demand Vickrey auctions to different ways of implementing the Vickrey auction – dynamic versus static mechanisms. Multi-unit demand auctions with synergies are covered in Section 3.3, with sequential multi-unit demand auctions covered in Section 3.4. Mechanism design studies that deal primarily with the thorny issues associated with package bidding are covered in Chapter xx.

Section IV deals with several issues that do not fit in neatly elsewhere: collusion, an ever present concern in auctions (Section 4.1), creating money out of "thin air" via selling multiple units simultaneously to bidders who demand only a single unit (Section 4.2), practices in Internet auctions (Section 4.3), and accounting for entry in auctions (Section 4.4).

The literature is much more extensive and less focused this time around than in the 1995 survey. The good news is that it covers a lot of new ground. The bad news is that we cannot hope to cover all of the good papers out there. Our hope is that we have surveyed enough of the more important developments in enough detail for both the novice and experienced reader to benefit from the survey, and that we have established synthesis in some areas, while not leaving out too much of importance.

I. Single-Unit Private Value Auctions

Initial experimental research on auctions focused on the independent private values (IPV) model, with particular focus on the Revenue Equivalence Theorem. In the IPV model each bidder privately observes their own valuation (and knows it with certainty), bidders' valuations are drawn independently from the same commonly known distribution function, and the number of bidders is known. Under the revenue equivalence theorem (Myerson, 1981, Riley and Samuelson, 1981) the four main auction formats (as part of a much richer class of auctions) – first- and second-price sealed-bid auctions, English and Dutch auctions – yield the same average revenue assuming the same number of risk-neutral bidders and the same reserve price.[2] Further, first-price sealed-bid and Dutch auctions, as well as second-price sealed-bid and English auctions, are theoretically isomorphic yielding not just the same ex-ante expected revenue but also the same revenue (price) in any realization of bidders' signals. These two isomorphisms are particularly attractive as they do not depend on risk neutrality (as does the more general RET), which makes for more robust tests of the theory's predictions.

An experimental session typically consists of several auction periods under a given auction institution. Subjects' valuations are determined randomly prior to each auction with valuations typically being independent and identical draws (iid) from a uniform distribution. In each period the high bidder earns a profit equal to his value less the auction price; other bidders earn zero profit. Bids are commonly restricted to be nonnegative and rounded to the nearest penny. Theory does not specify what information feedback bidders ought to get after each auction, which usually differs between different experimenters (and

[2] The Dutch auction starts with a high price which is lowered until a bidder accepts at that price. In English auctions price starts low and increases until only one bidder is left standing and pays the price where the next to last bidder dropped out. In a first- (second-) price sealed-bid auction the high bidder wins the item and pays the highest (second-highest) bid.

is sometimes taken as a treatment variable), with the auction price usually announced along with bidders own earnings.

At the time of the 1995 survey it was clear that both the revenue equivalence theorem as well as the strategic equivalence between each of the two pair of auction formats failed. Further, there were persistent reports of significant bidding above the risk neutral Nash equilibrium (RNNE) benchmark in first-price sealed bid auctions, initial explanations of which focused on risk aversion, generating considerable controversy among experimenters (see the December 1992 issue of the *American Economic Review).* Sorting out between explanations of this overbidding relative to the RNNE has preoccupied a number of later papers as well, several of which are reviewed first.

1.1 Bidding Above the RNNE in First-Price Private Value Auctions:

Isaac and James (2000a) compare estimates of risk preferences from first-price IPV auctions to the Becker-DeGroot-Marshak (BDM) procedure for comparably risky choices.[3] The Spearman rank order correlation coefficient between individual subject estimates of risk preferences under the two institutions is significantly *negatively* correlated, as subjects who bid as if they are relatively more risk averse in first-price auctions are relatively more risk loving under the BDM measure (see Figure 1). The net result is that *aggregate* measures of risk preferences imply that bidders are risk averse in the first-price auction but risk neutral, or moderately risk loving, under the BDM procedure. Although it is well known from the psychology literature that different elicitation procedures will yield somewhat different quantitative predictions (see Camerer, 1995, pp. 657-61; Mellers and Cooke, 1996), a negative correlation between measures seems rather astonishing.

[Insert Fig 1 here]

Dorsey and Razzolini (2003) look at IPV auctions in which a single human bidder competes in a series of first-price sealed bid auctions with three simulated buyers who bid according to the RNNE. They compare bids in this setting to an equivalent lottery procedure in which the same subjects essentially pick their preferred probability of winning against their computerized rivals, with expected profits conditional on winning

[3] These experiments use computerized rivals who bid according to the RNNE bidding strategy in the first-price sealed bid auctions. This permits isolating the risk preferences of individual human bidders in each auction market.

being computed for them for each probability level chosen. Mean lottery-equivalent bids are compared to mean auction bids over the relevant range of valuations. As shown in Figure 2, mean bids are essentially the same between the two procedures over the interval [0, 750], the first three quarters of the uniform distribution from which valuations were drawn.[4] In the remaining interval the lottery equivalent bids are consistently *lower* than the auction bids, suggesting that probability miscalculations (how close rivals valuations are to your own) play some rle in bidding above the RNNE at such higher valuations. They also compare bids in first-price sealed-bid auctions where subjects are told the probability of winning the auction for each possible valuation with the lottery equivalent procedure. In this case the bids overlap over the entire range of valuations, which supports their probability miscalculation hypothesis. Finally, note in Figure 2 the humped back nature of the deviations from the RNNE over the range of possible valuations, with mean bids essentially equal to the RNNE over lower valuations, above the RNNE (with the difference growing) for middle valuations, with these differences decreasing over the upper 25% of valuations. We will return to this point later. A more recent experiment looking at much the same issues with similar manipulations reaches even stronger conclusions that biased probabilistic beliefs are the primary driving force behind overbidding, with risk aversion playing a lesser role than previously believed (Armantier and Treich, 2007).

[Insert Figure 2 here]

Neugebauer and Selten (2006) (NS) compare different information feedback treatments in a series of IPV, first-price sealed bid actions against computerized rivals. They focus on three types of information feedback: (i) no information about bids of computerized rivals, just telling bidders if they won the auction or not, (ii) information about the bid of the highest computerized rival when they did not win (i.e., the market price - the feedback usually employed in experiments) but not when they won, and (iii) the market price and the highest computer's bid in case of winning the auction. They look at differences between actual bids and the RNNE bid in the first auction period and averaged over the entire set of 100 auctions, and do this with different numbers of computerized rivals. The number of subjects bidding above the RNNE in the *first*

[4] Values are in pennies.

auction period is reasonably small under all three treatments – 22% - with minimal differences between the three treatments. However, averaged over all auction periods, there was significant movement towards overbidding in all three treatments with the largest increase in treatment (ii), where 75% of all subjects bid above the RNNE, with an average estimated risk tolerance parameter (r_i) of 0.78 (where 1-r_i is the Arrow-Pratt measure of constant relative risk aversion; CRRA). In contrast, under treatments (i) and (iii) 41% and 48% of subjects bid above the RNNE, with an average estimated r of 1.25 and 1.17 respectively (i.e., on average subjects act as if they are risk loving). NS use "learning direction theory" to explain the changes in bidding over time under the different feedback conditions.[5]

Goeree, Holt and Palfrey (2002) (GHP) report a series of first-price sealed bid auctions with two bidders with a limited number (6) of discrete values (requiring discrete bids as well). There employ a low and high values treatment with the same RNNE bid in both treatments, but with the cost of bidding above (below) the RNNE being higher in the low (high) values treatment. They employ discrete values in order to estimate a *quantal response equilibrium* (QRE). They find overbidding relative to the RNNE in both treatments with an estimated Arrow-Pratt measure of CRRA under the QRE of approximately 0.50 in both cases. They compare their QRE model with risk aversion to (i) a non-linear probability weighting model and (ii) a joy of winning model. The non-linear probability weighting model fits the data as well as the QRE with risk aversion but has one additional parameter and does not have the inverted S shape weighting function one might expect (overweighting of small probabilities and underweighting of large probabilities). Joy of winning adds nothing to the QRE estimates with risk aversion, while a pure joy of winning model fits the pooled data quite well, although not as well as the QRE with risk aversion.

GHP take on the Rabin (2000) critique that estimates of risk aversion from laboratory experiments do not plausibly scale up to larger gambles so that given the levels of risk aversion reported, subjects would (implausibly) avoid very attractive large

[5] There is considerable variation in the extent of overbidding relative to the number of computerized rivals under the different treatments as well, with more than 50% overbidding under all three treatments in auctions with 3 and 4 computer rivals, and less than 33% doing so with 9 computer rivals in treatments (i) and (iii) (67% in treatment ii).

gambles. Their response to this critique is that the relevant argument in subjects' utility function is gains and losses from particular gambles and/or is defined over a smaller time interval (e.g., within the experimental session itself) as opposed to changes in wealth. On this point, also see Cox et al. (2007).

1.2 *Overbidding and Regret Theory*

Rabin (2000) points out that alternatives to expected utility theory would seem to provide a more plausible account of modest-scale risk attitudes, allowing for both substantial risk aversion over modest stakes and nonrediculous risk aversion over large stakes. Filiz and Ozbay (2007) (FO) explore the implications of one such model, regret theory (Loomes and Sugden, 1982; Bell, 1982) in an experiment looking at bidding in first-price sealed bid auctions. In their analysis of first-price auctions they note that the information bidders receive at the end of the auction may generate one of two types of regret: (1) "Loser's regret" if a losing bidder could have won the item with a higher bid and earned positive profit and (2) "Winner's regret" if a winning bidder could have earned more by bidding less (money left on the table). They first demonstrate that loser's regret by itself will generate bidding above the RNNE and that winner's regret, by itself, will generate bidding below the RNNE. To isolate the effect of these two factors, and to judge their relative strength, they conduct a series of one-shot, first-price sealed bid auctions in which following completion of the auction (i) losers learn the winning bid but the winner learns nothing, (ii) winners learn the second highest bid but losers learn nothing, and (iii) a control treatment in which bidders learn nothing about others bids.

They run one-shot as opposed to repeated auctions on the grounds that their theory relies on bidders anticipating future regret in terms of their current decisions, while repeated auctions might result in regret from previous rounds impacting on current round decisions. However, to gather sufficient data they solicit bids for 10 possible valuations from each bidder with different lists for each of the four bidders competing against each other. They then *average* responses across subjects with the same lists who are in different auctions.[6] Table 1 gives the slopes of the linear bid functions estimated,

[6] There are two sessions per treatment, with the minimum number of subjects in each treatment apparently eight. Exact numbers are not provided.

along with the 95% confidence intervals for same.[7] The estimated slope of the control treatment (no information provided) is 0.79, just within the 95% confidence interval for the RNNE value of 0.75. The slope estimated from the winner's regret treatment is just below this (0.77), but is not significantly different from the no information treatment. However, the slope of the loser's regret treatment is 0.87, which is significantly higher than the no information treatment. Although averaging bids across subjects with the same valuations does not bias the estimated slope coefficients, it no doubt biases the standard errors downward as it removes any *between-subject* variation in bids. Thus, the statistical significance of their results is suspect.[8]

[Insert Table 1 here]

Engelbrecht-Wiggans and Katok (2007) (EK) report a similar experiment. In their study human bidders compete against two computerized rivals, bidding 20 times in a row at each of five different valuations. In addition to the three information feedback conditions employed in FO, they have a treatment with both types of regret present (winners know the second highest bid and losers know the winning bid). Regret theory fails to explain overbidding here as the average ratio of *bids/valuations* is essentially the same between the losers' and no regret treatments (0.760 vs. 0.764). While in contrast to FO winner's regret results in a significant reduction in the average ratio of bids to valuations compared to the no regret treatment (0.736 versus 0.764; p = 0.03, Wilcoxin non-parametric test).[9]

The NS experiment also has implications for regret theory as they report data from the first auction and averaged over a number of auctions in what amounts to a no regret treatment (NS treatment i) and a loser's regret treatment (NS treatment ii). For a bidder who is bidding against three and four computer rivals (the treatments that come

[7] Linear bid functions were estimated with no intercept.

[8] They also report the results of a survey in which subjects were asked to rate the intensity of emotions they would feel after they got the relevant information. They find that losers' regret is substantially more intense than the regret in the other two treatments. However, under the loser's regret treatment the question simply states "… you are not the winner and you learn the highest bid" as opposed to losing and possibly winning with a positive profit. FO go on to show that winners' regret has no role to play in second-price, English and Dutch auctions, but that loser's regret can impact Dutch auctions.

[9] They also have an interesting test for CRRA in which they report earnings back to bidders aggregated over 10 auctions, as opposed to auction by auction as is typically done. Under CRRA aggregated earnings should lead to less overbidding relative to the RNNE compared to knowing earnings following each auction. This test provides *no* support for CRRA.

closest to FO and EWK) NS find 50% of their subjects (11 out of 22) bidding above the RNNE in the first auction of their no regret treatment versus 9.1% (2 out of 22) in the loser's regret, which does *not* match FO's results. Averaged over all auctions they find essentially the same number of subjects bidding above the RNNE in the no regret and loser regret treatments (15 out of 22 versus 16 out of 22 subjects) which is qualitatively consistent with EK. But NS's average estimated risk preference parameter (r_i) is somewhat greater with loser's regret than with no regret, with the largest difference occurring with 4 computer rivals (0.87 versus 1.01).

One important methodological point these experiments emphasize is that results from earlier experiments can be, and often are, reinterpreted in light of new and different theoretical perspectives. This in turn calls for new experiments to see if the insights from the new perspective are satisfied in the data. On this score there is still more work to be done on anticipated regret if it is to explain bidding in private value auctions as (i) there are a number of inconsistencies between the results reported from different experiments and (ii) NS's work indicates that outcomes are sensitive to the number of bidders in the auction.

1.3 *Using Experimental Data to Corroborate Maintained Hypotheses Empirical Applications to Field Data:*

Bajari and Hortescu (2005) (BH) use experimental data from first-price IPV sealed bid auctions with three and six bidders to non-parametrically estimate bid functions. The primary purpose of their paper is to determine whether structural models of first-price auctions as applied to field data can generate reasonable estimates of bidders' private information.[10] The latter is an essential element of what econometricians hope to recover in examining field data. In using experimental data to do what econometricians typically do with field data, they have at their disposal bidders' actual valuations against which to judge the accuracy of the recovery process, data that is not available in field applications. Further, unlike with field data there is no question that in this case the econometricians are dealing with an IPV auctions as opposed to a common value or affiliated private value auction, or an auction with significant private and common value elements.

[10] For a similar exercise with respect to common value auctions see Armantier (2002).

The results are also of interest to experimenters as BH test between four competing models: (i) the RNNE, (ii) Nash equilibrium bidding but with (homogenous) risk averse (CRRA) bidders, (iii) an adaptive learning model in which bidders maximize their expected utility based on beliefs about the distribution of bids (with the latter formed on the basis of previous auction round outcomes), and (iv) QRE with risk averse (CRRA) bidders. Their results show that Nash bidding with risk aversion provides the best overall fit to the data.[11] Further, they are unable to reject a null hypothesis that the actual and estimated distribution of bidder valuations is the same under this specification. In evaluating the fit of the CRRA model they need to trim the upper bound of the support from which valuations were drawn (the top 5% of all bids) as there is a negative correlation between bids and values over this part of the support. This is consistent with Dorsey and Razollini's results (recall section 1.1 above) that at the highest private valuations the overbidding relative to the RNNE is decreasing.

QRE with risk aversion provides results similar to those of the Nash model with risk aversion, but does not correctly pin down the lower end of the support from which valuations are drawn. More generally, after cautioning readers that "…experimental environments may differ significantly form 'real' economic environments", BH conclude (p. 707) "… our finding that the Bayesian-Nash equilibrium bidding model with risk aversion performs quite well, even in this experimental setting, is encouraging for present and future users of structural econometric tools."

1.4 Recent Developments in Second-Price Private Value Auctions

The 1995 survey covered research showing a breakdown in strategic equivalence between second-price and English clock auctions, primarily as a result of bidding above value in the second-price auction as opposed to sincere bidding in the clock auctions. Since then there have been several experiments designed to better understand why subjects overbid in second-price auctions, as well as why subjects do so much better in English clock auctions. We review these below.

Shogren, Parkhurst, and McIntosh (2006) (SPM) report bids from second-price auctions conducted under a tournament structure so that bidder earnings depend on the

[11] The RNNE model provides a reasonably good fit to the data in auctions with six bidders, but risk aversion is necessary to explain bidding in auctions with three bidders.

total points earned over 20 second-price auction trials with a tournament type payoff structure: the player with the most total points earned \$120, the second most earned \$80, and so on, with the three lowest earning \$5 each. They compare bidding in the tournament to bidding in a series of 20 standard second-price auctions. Each auction had a total of 10 bidders who were repeatedly matched with each other. Valuations were iid from a uniform distribution with support [0, 20].[12]

Deviations from sincere bidding were much smaller in the tournament than in the standard second-price auctions with the difference between bids and values (bid – value) averaging 6.28 (63.51) in the standard auctions versus 0.96 (4.14) in the tournament (standard deviations are in parentheses). However, there are relatively small differences in the frequency with which the highest value bidder won the auction averaging 55.0% in the tournament versus 42.5% in the standard auctions, with similar results for the frequency with which the highest and second highest value bidders won (72.5% in the tournament versus 70.0% in the standard auctions).[13] SPM conjecture that the superior performance in the tournaments results from the fact that the typical mistake of bidding above value has a much greater adverse effect on outcomes given the tournament pay structure. These results are particularly surprising since under a tournament structure there is clear motivation for losing bidders to bid above value as this reduces the winner's, which may help in terms of winning the tournament.[14] This factor is, of course, not present in standard second-price auctions.

Garratt, Walker and Wooders (2004) (GWW) conduct a second-price auction using subjects who regularly participate in eBay auctions for Morgan ("Golden Age") silver dollars. Arguably these subjects have considerable field experience given the similarity between eBay and second-price auctions. (But there are significant differences

[12] In each auction bidders got to see if they won the auction and how much they got, with no one seeing what anyone else bid or earned.

[13] Defining *strict efficiency* as the frequency with which the highest value bidder wins and pays the bid of the second-highest value bidder, this occurred in 45% of the tournaments versus 22.5% in the standard auctions.

[14] Consider the last round of a two round tournament with valuations from the interval [0,10] and where bidder 1 has a lead of four points and gets a signal of 5. Bidding 8 dominates bidding sincerely as it assures winning the tournament. The experiment had 20 rounds and it is not clear what information bidders had, so that without any information as to standing, things would be considerably more complicated. But the example is still insightful.

between eBay and second-price sealed bid auctions; see Roth and Ockenfels, 2002, reported in section 4.3 below) GWW invited these bidders to participate in a standard sealed bid second-price auction with induced valuations from a support comparable to the range of values that Morgan silver dollars sell for. There were five bidders in each auction. After bidding once in a presumably one-shot auction, subjects were invited back for a second-round of bids, conducted as a control against possible skepticism that payoffs in round one were not for "real."

[Insert Fig 3 here]

Figure 3 shows bids and valuations from their experiment. Looking at these they conclude that "… despite having substantial experience with auctions in the field, eBay subjects do not value bid." (p. 7). GWW compare the frequency of sincere bidding to Kagel and Levin's (1993) (KL93) experiment, employing the same criteria that any bid within five cents of a subject's value is counted as sincere. They find essentially the same frequency of sincere bidding, 21.2% versus 27.0% in KL.[15] However, there is substantially more under bidding than over bidding compared to KL – 41.3% (37.5%) underbidding (overbidding) in GWW versus 5.7% (67.2%) in KL. They are able, at least qualitatively, to resolve this discrepancy in the pattern of deviations from sincere bidding after they break their data down into eBay only buyers versus eBay sometime sellers as sellers tend to underbid much more often than buyers do, 50.9% of all bids versus 29.5%. There is a corresponding discrepancy in the frequency of overbidding, 45.5% for buyers versus 32.1% for sellers, with both sets of differences statistically significant at the 10% level using a non-parametric Mann-Whitney test.[16]

Thus, they conclude that, "… those who had sold typically bid less relative to their values than those who had not sold." This result is reminiscent of Burns (1985) study comparing professional wool buyers to students in a continuous double auction market. In that experiment she reports that students performed much better than the wool

[15] Bajari and Hortacsu (2004) in surveying results from internet auctions claim that GWW "…found that bidders experienced on eBay do not overbid in second-price private-value auctions …Hence, field-experiments conducted on online auction sites may indeed provide a more qualified subject pool for auction experiments." It is hard to see how they could reach this conclusion given the data reported.

[16] The difference in overbidding between buyers and sellers is significant at the 5% level in a regression analysis they conduct. Subjects were told that they could lose part, or all of, their $15 participation fee in case they won the auction with the second-highest bid above their value.

buyers (earning more money with more efficient outcomes), in large measure because the wool buyers ignored subtle differences between the laboratory experiment and the wool market. The connection here is that people who sell on eBay will typically only buy if the price is below their value, as otherwise they cannot profit from resale, and one cannot expect them to ignore these habits when put into a new situation. This is consistent with the psychology literature which suggests that in deductive reasoning processes people typically employ short-cuts, developing mental models of situations and reasoning about them in the context of the model (Johnson-Laird, 1999). Thus, it is easy to see how Burns' wool buyers might behave in ways that are more appropriate to their customary environment which was similar to, but not exactly the same, as the laboratory environment. Similarly, it is easy to see how e-Bay sellers, who make a living by buying low and selling high, might deviate from sincere bidding by bidding less than their induced values, while buyers, as is typical of standard laboratory subjects, bid above their induced values. In short, there is no particular reason to think that experienced professionals will perform much better than student subjects when placed in a laboratory type setting, unless there are *strong and relevant* similarities between the field where they practice and laboratory environments.[17]

Andreoni, Che, and Kim (2007) (ACK) report the highest rate of sincere bidding in second-price auctions we are aware of – 77.3% overall (85.5% in the last 10 periods) – in auctions with four bidders and a uniform distribution of valuations. They find that sincere bidding drops substantially, largely replaced by overbidding, when subjects know their rivals resale values. They attribute this result to spite.[18] While spite might explain overbidding when rivals valuations are known, this does not provide a credible explanation for overbidding absent this information, as there is minimal overbidding in

[17] On this score also see Dyer, Kagel, and Levin (1989) (DKL) along with Dyer and Kagel (1996) (DK). Also see Frechette (2007) for a comprehensive survey of differences between professionals and student subjects in experiments.

[18] ACK employ an (augmented) dual market technique with subjects bidding in each of three markets with the same valuations. In the first market bidders only have information about the common distribution from which values were drawn. In the second market precise information about one other bidder's value is provided and in the third market information about all other bidders' values is provided. ACK also explore the impact of information about rivals' values in first-price auctions where the theory makes clear predictions, which are largely satisfied, at least qualitatively.

English clock auctions, which are strategically equivalent, and in which spite (as well as joy of winning) should play just as strong a role.

Cooper and Fang (in press) (CF) look at bidding in a series of two player second-price auctions with bidders valuations drawn from an approximate normal distribution. Their primary treatment variable consists of noisy information about rival's valuations, which in some cases is provided exogenously and in other cases can be purchased. In the control treatment, with no information about rival's valuations, just under 40% of all bids are sincere, with overbidding accounting for most of the deviations. Unlike ACK, with exogenously provided information about rival's valuations the rate of sincere bidding increases, especially with less noisy information. The probability of overbidding is reduced in response to costly mistakes (overbidding that causes subjects to lose money), with the apparent stability of bidding above value resulting from the infrequency of costly mistakes.

CF also find that subjects tend to buy costly information about rivals valuations (although the game has a dominant strategy so at least from a game theoretic perspective this is involves throwing money away), with these purchases diminishing over time. There is considerable heterogeneity in the subject population with respect to buying information: subjects who overbid the most tend to buy information more often. This suggests a split in the population between more "rational" types who neither overbid nor pay to buy essentially worthless information and less rational types who commit both types of mistakes.

1.5 Asymmetric Private Value Auctions

While much of the auction literature has focused on bidders that are *ex-ante* symmetric, in many real auctions it is commonly known that one or more bidders (the *strong* bidders) are likely to have higher valuations for the auctioned item than the other (*weak*) bidders. This extension of the private values model raises interesting theoretical questions (see Maskin and Riley, 2000) that have been explored in a handful of experimental studies which we review below.

Pezanis-Christou (2002) conducts an experiment based on a model with two risk-neutral bidders (i = 1, 2) each demanding a single unit. Bidders values are independent draws from a uniform distribution, with support [0, 100] for the strong bidder and support of either [-100, 100] or [-300, 100] for the weak bidder. The underlying support for the strong bidder first-order stochastically dominates (FOSD) that of the weak bidder. Negative bids are not allowed, with the weak bidder not allowed to bid when receiving a negative value. Each session consisted of either 60 or 72 auctions in which subjects' type (strong or weak) varied between auctions, as well as the two distributions for weak bidders' values.

The experiment contrasts differences in bidding between strong and weak bidders between first- and second-price sealed bid auctions as well as the impact of the two different supports for weak bidders on bidding in first-price auctions. Key comparative static predictions investigated are: (i) In the first-price auctions the strong bidder bids less aggressively than the weak bidder, ($b_S(v) < b_W(v)$), (ii) Efficiency is greater in second-price compared to first-price auctions, and (iii) Expected revenue is higher in the second-price auctions. The intuition underlying (iii) here is that since there is a positive probability that weak bidder will not bid (as a result of a negative value), strong bidders in the first-price auctions can maximize their expected earnings by placing very low bids ("low balling") when they get low values. In contrast, sincere bidding (bidding one's value) remains a dominant strategy in the second-price auction, resulting in higher revenue under the second-price institution. Further, these low-balling and revenue differences should be greater when the weak bidder has a greater likelihood of drawing a negative value (with support [-300, 100]).

As predicted, the second-price auctions have higher efficiency averaging 97% versus 95% in the first-price auctions with weak bidders support [-100, 100] and 99% versus 96% with [-300, 100].[19] The second-price auctions also generated higher efficiency in 9 out of 10 paired group comparisons ($p < 0.0$ x using a simple binomial test). Average revenue in the second-price auctions was approximately equal to its predicted expected value but, contrary to the theory, average revenue in the first-price auctions was greater than in the second-price auctions in both cases. Pezanis-Christou attributes this failure of the theory to bidders' difficulty in recognizing the profitable opportunities from low-balling, so that average revenue in the first-price auctions was greater than predicted. He rejects a general argument based on risk aversion as (i) with

[19] Efficiency is measured by the ratio of the [winner's value]/[highest value] * 100.

the weak bidders draws from the interval [-300, 100] the revenue ranking is not affected by risk aversion, (ii) simulations assuming both bidders are extremely risk averse cannot account for the reversal of the revenue ranking when the weak bidders values are drawn form the interval [-100, 100], and (iii) the size of the deviations in bids from the RNNE were decreasing over time, suggesting that subjects were employing an adaptive bidding strategy as opposed to a static, fully-optimizing one.[20] He goes on to show that the observed bid distribution for strong bidders is best approximated by the theoretical distribution with strong bidders assuming they face two other competitors rather than the single competitor they actually faced.

As predicted, strong bidders shave their bids more than weak bidders do in the first-price auctions under both treatments; where bid shaving is defined as the ratio $\varepsilon_i = (v_i - b_i)/v_i$ where v_i is bidder i's value. And they shave more when weak bidders draw values from [-300, 100]. About 46% of all second-price bids were equal to subjects' valuations, which is substantially larger than in previous experiments, with 40% of the bids above value. However, the fact that second-price revenues were close to their predicted levels indicates that whatever overbidding there was had to be relatively small.

Methodological Remark: Each of Pezanis-Christou's sessions had 12 subjects, who were told they would be randomly matched with another participant. In practice, the set of 12 subjects was divided into 3 groups of 4 subjects each with rotation within each group. This was done in an effort to obtain "three independent sets of observations per session instead of only one" as the unit of observation employed in the analysis is primarily session level data. The idea behind "only one" independent observation per session if randomly rotating among all 12 bidders in the session is given the repeated interactions between subjects this generates session level effects that will dominate the data. In this regard he is among a growing number of experimenters who believe this, and who break up their sessions into smaller subgroups in an effort to obtain more "independent" observations per session. This practice ignores the role of appropriate panel data techniques to correct for dependencies across and between subjects within a given experimental session.[21]

There are several important and unresolved issues in choosing between these two procedures. In both cases experimenters are trying to squeeze as much data as they can from a limited subject-payment budget. As experimenters who have consistently employed random rematching between all subjects recruited for a given session, and applied panel data analysis to appropriately account for the standard errors of the

[20] Note, however, that the revenue reversal remains even after looking at bidding in the last 30 auctions, where learning should have tended to stabilize.

[21] It should be clear that these remarks are *not* directed at this particular experiment but rather at a more general issue within the experimental community.

estimates, we are far from unbiased with respect to this issue. With this in mind we point out several things: First, advocates of repeated matching of the same small subset of subjects within an experimental sessions to generate more "independent" observations ignore the fact that there is no free lunch as: (i) they are implicitly lying to/deceiving subjects by not reporting the rotation rule employed and (ii) if subjects are as sensitive to repeated matching effects as they seem to assume under random matching between all subjects in a given experimental session, it seems plausible that repeated play within a small subset might generate super-game effects that will contaminate the data. Second, and more importantly, there have been a few experiments which have devoted treatments to determine the severity of possible session level effects from random rematching for the group as a whole. More often than not these studies find no differences, e.g. Cooper et al. (1993; footnote 13, p. 1308), Duffy and Ochs (2006). Also see Walker et al. (1987) and Brosig and Reiβ (2007) who find no differences when comparing bids in auctions with all human bidders against humans bidding against computers who follow the RNNE bidding strategy. For more on the econometrics of this issue see Frechette (2007).

Güth, Ivanova-Stenzelb and Wolfstetter (2005; GISW) investigate a simplified version of the asymmetric auction model in Plum (1992). Two risk-neutral bidders compete for the purchase of a single item in either a first- or a second-price sealed-bid auction. Bidders' valuations are independently drawn from a *uniform* distribution with support [50, 150] for the weak bidder versus [50, 200] for the strong bidder. As predicted, efficiency is consistently higher in the second-price auctions than in the first-price auctions averaging 98%, 99%, and 99% versus 97%, 97%, and 98% over the three time phases of the experiment.[22] Although the theory predicts that weak bidders' payoffs will be higher in a first-price auction, and strong bidders' payoffs higher in a second-price auction, both weak and strong bidders' average payoffs are significantly higher in second-price auctions. Bids are close to their predicted level in the second-price auctions (sincere bidding), but as is commonly the case, are substantially higher than predicted under the RNNE in the first-price auctions. This overbidding in the first-price auctions accounts for the failure of weak bidders' payoffs to be higher in the first-price auctions.

A closer look at bid patterns shows that strong bidders in the first-price auctions generally obey first-order rationality, as there are few bids above 150, the maximum possible valuation for weak bidders. Further, weak bidders shave their bids less than strong bidders at higher valuations ($v \approx 100$). Although, this satisfies a key qualitative prediction of the theory, the pattern differs from the predicted one as the differences in

[22] No statistical tests are reported for this.

bid shaving between weak and strong bidders does not increase regularly over higher valuations, and the differences are not nearly as large as the theory predicts at higher valuations. When given a choice, both weak and strong bidders overwhelming chose the second-price auction, consistent with the significantly higher payoffs both types get under this format.

Chernomaz (2006) looks at asymmetries resulting from two otherwise symmetric auction participants who merge to submit a single bid based on the highest of their private valuations.[23] This s*trong* bidder competes against a single w*eak* bidder. Each bidder draws a private value from a common uniform distribution, but by virtue of using the higher of their two private values, the value distribution for the strong bidder FOSD the weak bidder. Subjects participate in a series of first-price, sealed bid auctions under each of three treatments: (i) They bid as separate entities based on their private values in an auction with three bidders. (ii) "Merged" firms let each subject bid separately, with no communication, based on the higher of their two valuations. (iii) Merged firms submit a single agreed upon bid after they have the opportunity to communicate via an instant messaging system. Subjects' roles as weak or strong bidders remain fixed throughout a session, as do the pairings for the "merged" firm. The dual market technique is employed so that in each auction bids under all three treatments are based on the same valuations with the market to be paid off on determined randomly after all bids have been submitted. In equilibrium, the strong bidder bids less than the weak bidder with the same valuation, which in turn results in inefficient allocations compared to the symmetric first-price auctions. This design also permits comparing bids in the symmetric auction with those in the asymmetric auctions, with the latter predicted to be uniformly lower for both strong and weak bidders under the RNNE. As a result, revenue is predicted to decline and bidders' profits are predicted to increase, with the weak bidder getting a larger absolute increase in profits than the ("merged") strong bidders get after splitting their earnings.

[23] There are a couple of different ways to think about this. The firms have merged so that the bidder with the higher private value is the firm's value. Alternatively, there is a consortium of bidders who bid jointly and agree to allocate the item to the bidder with the highest value, along with some agreed upon device for splitting the profits. All of this is in the background, as in order to simplify the experimental design, who bids jointly is determined exogenously by the experimenter, with an exogenously determined rule for splitting the profits equally.

This last result has implications for the incentive to bid jointly in a model where joint bidding is determined endogenously.[24]

The experimental results show overbidding relative to the RNNE in the symmetric auctions and on the part of both weak and strong bidders in the asymmetric auctions. However, strong bidders bid less aggressively than in the symmetric auctions, although the difference is not as large as the theory predicts. Weak bidders tend to bid the same, or slightly higher, than in the symmetric auctions. Chernomaz shows that these differences in behavior can be partly accounted for by the different incentives weak and strong bidders face. Although subjects are not best responding (if they were indeed risk neutral) in the symmetric auctions, the best response function for the weak types is not much different between the symmetric and asymmetric auctions, indicating that they had very little *additional* incentive to change their bids given how the strong types were bidding. In contrast, there are relatively strong incentives for the strong bidders to reduce their bids between the symmetric and asymmetric auctions.

Contrary to the RNNE predictions, efficiency is *higher* with joint bidding than in the symmetric benchmark case. This can be explained by the reduction from three to two bidders, so that any inherent noise in bids is less disruptive to efficiency in the two bidder case as bidders' valuations are further apart, on average, than in three bidder case.[25] In addition, strong bidders benefit from joint bidding at least as much as the weak bidders (even after accounting for the fact that they split their profits), indicating that the incentives to bid jointly are stronger than predicted. Finally, there are essentially no differences in bids when members of the "merged" firm bid individually versus bidding jointly. But for some unknown reason, weak bidders tend to submit higher bids when the "merged" firm bids jointly.

Summing Up: Tests of revenue predictions in asymmetric private value auctions are confounded by the fact that subjects tend to bid well above the RNNE in first-price sealed bid auctions but bid close to the dominant strategy in second-price sealed bid auctions with only a couple of bidders. However, bid functions tend to move, at least qualitatively, in the right direction as in all three experiments strong bidders tend to bid

[24] This result is similar to results from horizontal-mergers in a Cournot oligopoly (Levin, 1990).

[25] This highlights the problem with comparing efficiency between different auction structures. One solution here is to normalize efficiency measures by the difference from random bidding.

less than weak bidders for comparable valuations. Efficiency tends to be lower in first-price compared to second-price auctions which is the same result reported for symmetric first- and second-price auctions (reviewed in Kagel, 1995). One secondary result of these experiments is that they show closer conformity to sincere bidding in second-price auctions with two bidders than typically found in experiments with larger numbers of bidders. However, rather large and persistent bidding above value remerges in Vickrey auctions with two bidders with flat demand for two units with a supply of two units (see Section 3.1 below).

1.6. Cash Balance Effects and the Role of Outside Earnings: In the typical auction experiment subjects bid in a series of auctions with payoffs following each auction period. As a consequence bidders' cash balances will vary over the course of the auction which, for a variety of reasons, may impact bidding; e.g., if subjects are risk averse with other than constant absolute risk averse preferences or for behavioral reasons such as having earnings targets or earning aspirations that they bring to the experiment. Further, since these cash balances are endogenous, absent proper instruments they cannot simply be included as a right hand side variable in estimating bid functions since this will result in biased estimates of coefficient for variables of interest.

Ham, Kagel and Lehrer (2005) (HKL) investigate cash balance effects in the context of an affiliated private value auction, employing the experimental design first used in Kagel, Harstad and Levin (1987).[26] HKL introduce exogenous variation in bidders' cash balances by simultaneously enrolling them in a lottery which has both positive and negative payoffs (but positive expected value), providing a more powerful test of the null hypothesis that cash balances do not affect bidding. In addition, they use instrumental variable estimators of the cash balance effect, with instruments based on the lottery earnings along with other exogenous variables (e.g., ranking of signal values) produced during the experiment. They also varied the number of bidders in each auction with either 4 or 6 bidders competing in an experimental session with 30 auction periods with a between groups design.

[26] The advantage of affiliated private values is that except for end point effects, bidders do not know if they have a high or low signal value. This is valuable since in IPV auctions bidders with low valuations know they have little chance of winning the auction which tends to result in a number of "throw away" bids (less serious bids) of one sort or another. The KHL affiliated private values model tends to eliminate such bids.

Their results show a small, but statistically significant, negative cash balance effect on bids; i.e., the larger cash balances are the lower subjects bid, other things equal. The quantitative effect of cash balances on the bid factor (the difference between a bidders value and what they bid) is to increase it from \$1.76 to \$2.36 in auctions with 4 bidders, and from \$1.27 to \$1.70 with 6 bidders (evaluated at the mean value for cash balances). KHL also estimate a time trend variable ($1/t$ where t is the number of auction periods) to capture any learning/adjustments on the part of bidders, which shows that bid factors decrease over time. This is consistent with NS's results (section 1.1 above) that bids tend to increase with experience and feedback regarding auction outcomes, as subjects were provided all bids and values following each auction.

HKL estimate the impact of not including cash balances on the estimated bid function. It biases the time trend coefficient downward – so that there is less of an increase in bids over time. In addition, since in auctions with larger numbers of bidders, subjects have lower earnings and smaller increases in their cash balances (holding the support from which values are drawn remains constant), part of the increase in bids found with increased numbers of bidders can be attributed to smaller cash balances.

HKL attribute the negative cash balance effect to target income earnings and/or income aspirations on bidding. They conjecture that the mechanism underlying this effect is that subjects, who are recruited for the experiment with the promise of cash earnings, enter the auction with some target income earnings in mind and quickly recognize that they must win an auction to realize these earnings, which promotes higher bids at first. However, as cash balances accumulate, bidders come closer to their target earnings which motivates them to take a chance on a bigger score by lowering their bids, even though this reduces their chances of winning. This effect is partly offset by the feedback regarding lost profit opportunities, which indices more aggressive bidding over time. This conjecture concerning the mechanism behind the cash balance effect remains to be tested directly. However, it does receive indirect support from at least one independent study.

Turocy and Watson (2007) (TW) report an experiment comparing bids in the typical first-price auction with resale values as opposed to auctions in which bidders have opportunities to obtain the same item at "outside prices." With outside prices each bidder has the same value for the item they are bidding on, with the winning bidder paying what they bid and all other bidders also getting the item but at randomly determined “outside prices.”[27] The outside price treatment does two things: (i) it frames the problem very differently as winning now involves an opportunity cost (what you would have earned by not winning given your outside price) and (ii) everyone is assured of earning a profit in each auction period. They find that bids in the auction are consistently lower with the outside price treatment than in the standard resale value treatment. TW attribute this difference exclusively to the first factor. However, the fact that everyone is assured of earning a profit in each auction should help eliminate anxieties about achieving a reasonable level of positive earnings regardless of whether they win or lose in the auction, the driving force that HKL conjecture underlies their cash balance effect.[28] [29] More direct tests, or more indirect evidence, is needed to confirm this conjecture or to overturn it in favor of an alternative explanation.

II Single Unit Common Value Auctions

In common value auctions (CVA) the value of the item is the same to all bidders. What makes common value auctions interesting is that although bidders don’t know the true common-value they receive signals (estimates**)** that are correlated (affiliated) with that value. Mineral rights auctions (e.g., outer continental shelf - OCS - oil lease auctions), are usually modeled as a common value auction. There is a common value element to most auctions. Bidders for a painting may purchase it for their own pleasure, a private value element, but also for investment and eventual resale, the common value element.

Experimental research on CVAs has focused on the “winner's curse.” Although all bidders obtain unbiased estimates of the item's value, they typically win in cases where they have (one of) the highest signal value(s). Unless this adverse selection problem is accounted for, it will result in winning bids that are systematically too high, earning below

[27] With risk neutral bidders the symmetric Bayes-Nash equilibrium bid functions are identical under the two treatments. There are some qualitative differences in the equilibrium bid functions if subjects are risk averse.

[28] To control for the higher earnings that the outside option treatment would produce TW payoff on a subset of randomly chosen auctions. Nevertheless, subjects are assured of positive earnings in all cases, unlike the resale value treatment. There are also response mode effects to consider given the interface differences between the resale value and outside price auctions.

[29] TW also report results for Dutch (descending price) auctions under both treatments; also see Turocy et al. (2007) for Dutch auction results compared to first price auctions. These results replicate those first reported in Coppinger et al. (1980) – lower prices in the Dutch compared to first-price auctions.

normal or negative profits - a disequilibrium phenomenon. Oil companies claim they fell prey to the winner's curse in early OCS lease sales, with similar claims made in a variety of other settings (e.g., free agency markets for professional athletes and corporate takeovers). Economists are naturally skeptical of such claims as they involve out-of-equilibrium play. Experiments reviewed in Kagel's (1995) survey clearly show the presence of a winner's curse for inexperienced bidders under a variety of circumstances and with different experimental subjects: average undergraduate or MBA students (Bazerman and Samuelson, 1983; Kagel and Levin, 1986), extremely bright (Cal Tech) undergraduates (Lind and Plott, 1991), experienced professionals in a laboratory setting (Dyer et al, 1989), and auctions in which it is common knowledge that one bidder knows, with certainty, the value of the item (Kagel and Levin, 1999). Papers reviewed there also dealt with several alternative explanations for the winner's curse – limited liability for losses (Hanson and Lott, 1991, Kagel and Levin, 1991, Lind and Plott, 1991) and joy of winning (Holt and Sherman, 1994).

We pick up the story here with experiments investigating the ability of English auctions to raise revenue compared to first-price sealed bid auctions and the effects of an insider who has better information than rival bidders. We look at behavior of very experienced bidders (how close do they get to the RNNE), bidding in "almost" common value auctions where one bidder values the item slightly more than others (and this is common knowledge), bidding in auctions with both private and common value elements for *all* bidders, new results on the closely related "takeover" game, as well as demographic and ability effects on the probability of falling prey to the winner's curse.

The auctions reported on here, unless otherwise noted, use the following experimental design: The common value, x_o is the same for all bidders and is chosen randomly from a *uniform* distribution with support [$\underline{v}$, v]. Each bidder i receives a private information signal, x_i, about the unknown value of the item based on an iid from a uniform distribution with support $[x_o - \varepsilon, \ x_o + \varepsilon]$.

2.1 *English Auctions*

Levin, Kagel, and Richard (1996) (LKR) implement an irrevocable exit, ascending-price (English clock) auction. Prices start at $\underline{v}$, the lowest possible value for x_o, and increase continuously. Bidders are counted as actively bidding until they drop out of the auction and are not permitted to reenter after that.[30] The last bidder earns a profit equal to x_o less the price at which the last bidder dropped out. Bidders observe the prices at which their rivals drop out of the bidding. Auctions of this sort have been run in Japan (Milgrom and Weber, 1982, Cassady, 1967). The irrevocable exit procedure, in conjunction with the public posting of drop-out prices, insures that in equilibrium bidders can infer their rivals' signal values from the drop-out prices.

The analysis focuses on signals in the interval $\underline{v} + \varepsilon \le x \le v - \varepsilon$. In a symmetric RNNE the bidder who holds the lowest signal value (x_L) drops out of the auction once the price reaches x_L.[31] This drop out price reveals x_L to the remaining bidders. Given the *uniform* distribution of signal values around x_o, and the fact that in a symmetric equilibrium any remaining bidder j wins only when she holds the highest signal, each bidder j ought to use $(x_L + x_j)/2$ (which provides a sufficient statistic for x_o) as their drop out price in the symmetric RNNE. Drop out prices other than x_L contain no additional information and should be ignored. Expected profits in the English auction are reduced by about 50% compared to a first-price sealed bid auction with private information with more than two bidders competing. As such, in equilibrium, the English auction is predicted to significantly raise average revenue compared to a first-price auction.

Earlier experimental results from first-price auctions with x_L *publicly announced* (KL, 1986) showed that when bidders suffered from a winner's curse, announcing x_L *lowered* revenue (contrary to the theory's prediction) as bidders with higher signal values upon observing x_L recognized that they were overestimating the common value. However, once bidders had adjusted to the winner's curse and were making reasonable profits relative to the RNNE benchmark, revealing x_L increased revenue via the linkage principle, as the theory predicted. The key difference between LKR's English clock

---

[30] Prices started at $\underline{v}$ as any other price rule would reveal information about x_o. Initially, the price increased every second with increments of $1.00. Once the first bidder dropped out there was a brief pause after which prices increased with smaller price increments.

[31] As in the second-price auction, but not the first-price, it is an equilibrium for the bidders with x_L to bid (drop out at) x_L. The intuition is roughly as follows: Given symmetry, the low signal holder knows that those remaining in the auction have higher signal values so that his estimate of x_o is higher than x_L. But the low signal holder can't profit from this additional information since not dropping at x_L pushes the price up, so that winning at a higher price, when others drop at equilibrium prices, assures the low bidder negative expected profit.

auctions and these earlier first-price sealed bid auctions is that information dissemination is endogenous in the clock auctions rather than exogenous as when x_L is publicly announced. As such higher signal holders must be able to recognize and process the relevant information, and low signal holders must recognize the futility of remaining active once the price exceeds their signal value in order for the results to generalize.

[Insert Table 2 here]

Table 2 shows averages of predicted and actual changes in revenue between English and first-price auctions for inexperienced bidders with the results classified by the number of bidders.[32] Average revenue is predicted to be higher in the English auctions in all cases for the set of signal values actually drawn, with significantly higher average revenue predicted for $n = 4$ for all values of ε and for n = 7 with $\varepsilon = \$12$.[33] However, for these inexperienced bidders, with the exception of $n = 4$ and $\varepsilon = \$24$, actual revenue is *lower* in the English auctions, with significantly lower average revenue for $n = 4$ and 7 with $\varepsilon = \$6$ ($p < 0.05$) and for $n = 7$ and $\varepsilon = \$12$ ($p < 0.10$).

These perverse revenue effects in terms of Nash equilibrium bidding theory are associated with negative average profits in both the first-price and English auctions (see Table 2). The negative average profits indicate that inexperienced bidders suffered from a winner's curse in both first-price and English auctions, but that the curse was relatively stronger in the first-price auctions. These results serve to generalize those reported for first-price auctions with x_L publicly announced (KL, 1986). However, the generalization is not complete as average profits in the English auctions were negative compared to positive average profits in the first-price auctions with x_L publicly announced. This suggests that information dissemination in the English auction is noisier than with x_L publicly announced. This probably results from two factors: (i) in the English auction bidders may not completely recognize the relationship between the first drop-out price

[32] Common-value auctions involve pure surplus transfers so that revenue differences are calculated as: $[\pi_E - \pi_F]$ where π_E and π_F correspond to profits in English and first-price auctions, respectively. This effectively normalizes for sampling variability in x_o by subtracting it from the price.

[33] T-tests are conducted for predicted revenue increases to measure the reliability of the prediction for the LKR sample data with one-tailed t-tests used since the symmetric RNNE makes unambiguous predictions regarding revenue increases. Two-tailed t-tests are used for determining statistical significance of actual revenue changes since the presence of a winner's curse promotes lower revenue in English auctions.

and x_L and (ii) there is some out of equilibrium play with low signal holders dropping out sometimes above and sometimes below their signal value. [34]

For more experienced bidders, English auctions raised average revenue with *n* = 4, with a statistically significant increase for ε = \$18 (see Table 3). However, for *n* = 7, there was essentially no difference in revenue between the first-price and English auctions. The significant increase in average revenue in English auctions with *n* = 4 was associated with the elimination of the worst effects of the winner's curse in the first-price auctions, as bidders earned a substantial share (more than 50%) of predicted profit. In contrast, with *n* = 7 bidders earned a relatively low share (21%) of predicted profits in the first-price auctions, indicating substantially stronger residual traces of a winner's curse, and highlighting the importance of largely eliminating the winner's curse for the revenue raising prediction of the theory to hold.

[Insert Table 3 here]

LKR develop an econometric model to characterize how bidders process information in the English auctions. As noted, the Nash bidding model predicts that bidders with higher signal values will average their own signal value with the first drop-out price observed, ignoring all intermediate drop-out prices. What they found, however, is that bidders placed weight on their own signal and the immediate past drop out price, ostensibly ignoring x_L and any earlier drop out prices. Further, as more bidders dropped out, subjects placed less and less weight on their own signal and more weight on the last drop-out price. This pattern, although inconsistent with the Nash model, is consistent with bidders acting "as if" they were averaging their own signal with the signal values underlying the drop out prices of *all* earlier bidders. They attribute the adoption of this signal averaging rule to the fact that (i) it is easy and quite natural to use and (ii) it yields results quite similar to the Nash rule without requiring that bidders explicitly recognize the adverse selection effect of winning the auction and/or knowing anything about sufficient statistics. One unanswered question raised by this analysis is if the signal-

[34] To further investigate this question Kagel and Levin (unpublished data) conducted some additional sessions with inexperienced bidders in which x_L was publicly announced prior to bidding in the English auctions. In auctions with 6 bidders and ε = \$12, average profits in the standard English auction were -\$1.55, with average profits in auctions with x_L announced of \$1.56 (t = 1.46, d. f. = 30, p < .10, 1-tailed test).

averaging rule would still be used with distribution functions where it leads to markedly different outcomes from the Nash equilibrium, as in this case bidders would have more opportunity to recognize and respond to the profit opportunities inherent in abandoning the signal-averaging rule.

2.2 Auctions with Insider Information

The standard common-value auction model assumes that all bidders are *ex-ante* symmetric with respect to the quality of their signals (estimates) regarding the common-value. It is quite natural to ask how robust the equilibrium is to the insertion of one bidder (or a subset of bidders), an *insider*, who is better informed than the other bidders, *outsiders*. The easiest place to start this analysis is to assume that it is commonly-known that there is a single insider with a better, more precise, estimate of the true value (at the extreme, an insider who knows the true value).

Evaluating auction performance with an insider compared to the symmetric baseline depends quite critically on the baseline chosen. Wilson (1967) employed a symmetric baseline in which all bidders have *only* public information so the seller extracts the entire surplus (bidders earn zero profits) in the resulting mixed strategy equilibrium (also see Hausch, 1987 and Hendricks et al., 1994 for similar models). Introducing an insider into this environment who knows the true value reduces the seller's revenue, as the insider can, and does, bid below the true value, earning positive profits. Since *ex-post* efficiency is not an issue in a pure common-value model, the insider's gains must be the seller's loss. We are unaware of any experiments that have tested these predictions of the Wilson-type model.

In contrast Kagel and Levin (1999; KL) introduced an insider into a symmetric baseline with bidders having affiliated private information signals along the lines characterized in section 2.1 above. They did this in the context of an experiment with the goal of determining if the presence of an insider who knows the true value would help bidders recognize the adverse selection effect conditional on winning against an informed insider thus mitigating, and possibly even eliminating, the winner's curse. Although this hypothesis failed (inexperienced outsiders suffered from as strong a winner's curse as inexperienced bidders with symmetric information), the experiment led to a very surprising and significant discovery: With more experienced subjects who had learned to overcome the worst effects of the winner's curse, earning substantial positive profits in first-price sealed bid auctions with symmetric information, the introduction of an insider actually increased seller's revenues! Table 4 reports these results.

[Insert Table 4 here]

This surprising outcome, particularly given the theoretical results from the Wilson-type models, might have led some skeptics to dismiss this new finding, arguing that "in laboratory experiments anything can happen." However, a further examination of the model revealed that the data are consistent with the model's predictions and to the discovery of the mechanism underlying the possibility of an insider raising seller's revenue. Unlike in Wilson-type models, in KL's symmetric information benchmark model all bidders have private information, so in equilibrium bidders make positive profits. As a result the introduction of a perfectly informed insider eliminates those baseline cases where the winning bidder makes very large profits, as the insider bids closer to the true value, on average, than bidders in the symmetric information baseline. (Evidence for this can be seen in the much smaller variance in average profits between asymmetric versus symmetric information setups reported in Table 4.) Further, unlike the early Wilson-type models, both insiders and outsiders earn positive average profits in equilibrium, as both agents have *private* information (insiders do not know the outsiders' private information signals). What the two types of models do have in common is that conditional on winning, insiders make much larger average profits compared to outsiders as they have superior information.[35]

KL (1999) argue that many "real world" cases are more realistically modeled with outsiders having some proprietary information and not just public information. In these circumstances, it may well be the case that the introduction of a single well-informed insider increases average sellers' revenue, and that both insiders and outsiders earn economic rents. This potential for an insider to raise average revenue in a common value auction had not been recognized in the literature prior to this.

[35] There is no analytic solution, or even readily calculated numerical solutions, to the system of differential equations that characterizes the Nash equilibrium in KL (1999). However, see Campbell and Levin (2006) for a model which solves for the Nash equilibrium analytically, in which the introduction of an insider may raise seller's revenues.

2.3 Almost Common Value Auctions

The standard model of pure common-value auctions assumes that all bidders have exactly the same value for the item. But how robust, theoretically and in practice, are the properties and performance of auctions to slight departures from the pure common value assumption? In many common value auctions it is common knowledge that one (or more) bidder(s) (*advantaged* bidders) get an extra payoff relative to the other (*regular*) bidders; e.g., in the FCC regional air wave rights auctions Pacific Telephone was widely believed to place a higher value on the West Coast regional area than their potential rivals because of their familiarity with the region and their existing customer base (Klemperer, 1998). Economic theory suggests that with two bidders and a second-price or English auction even the tiniest private value advantage can have an "explosive" effect on auction outcomes with the advantaged bidder always winning and earning very high profits (sharply reduced revenue) (Bikhchandani, 1988; Klemperer, 1998). However, the question of whether or not these predictions will emerge depends critically on bidder behavior. It is here where experiments can help sort out when, where and why we ought to be concerned about such explosive effects.

Avery and Kagel (1997) (AK) investigated the explosive effect of a small private value advantage in a second-price "wallet-game" auction. In this game there are two bidders who bid in a *second-price* auction for the value of a wallet while each of them privately observes the content of only one cell of the wallet's two cells. Let x_1 and x_2, represent the privately observed signals by the first and second bidder respectively. The value of winning the wallet for these bidders is: $V_1 = x_1 + x_2 = V_2$. Bidding twice their observed signal, $b(x_i) = 2x_i$, $i = 1, 2$ is both a unique symmetric equilibrium as well as an *ex-post* equilibrium: There is no regret, it is distribution free, and is not affected by risk attitudes.[36]

With a private value advantage, the valuation of the advantaged bidder (say bidder 1) becomes: $V_1 = x_1 + x_2 + \Delta$ (or $V_1 = x_1 + x_2 + \Delta x_1$, in the multiplicative form) where $\Delta > 0$, is presumed small. Essentially what the private-value advantage does is to destroy the symmetric equilibrium of the second-price auction. In the resulting

[36] The winner earns $x_{High} + x_{Low}$ and pays $2x_{Low}$ for a net gains of $x_{High} - x_{Low} > 0$, while if upon deviation the loser wins she earns $x_{High} + x_{Low}$ and pays $2x_{High}$ for a net gains of $x_{Low} - x_{High} < 0$. Thus, even ex-post, upon finding both signals no one regrets the outcome (given the strategy of the other bidder.)

asymmetric equilibrium the private-value advantage has a "snowball" effect resulting in the advantaged bidder winning all the time, bidding too high for the disadvantaged bidder to try to unseat him. This does not happen in a first-price auction but does in a second-price auction as the high bidder does not have to pay what he bids. In the experiment, the effect of the Δ value advantage on bids and prices was proportional rather than explosive as the difference in bids between advantaged and disadvantaged bidders for comparable signal values was closer to the private value advantage of $1 than to the $3 difference predicted under the explosive effect. In effect, both advantaged and disadvantaged bidders were bidding closer to the naïve expected value of the item conditional on their estimates of the common value, with the advantaged bidder simply adding their private value advantage to their estimate of the common value. AK explore a number of alternative explanations for this outcome. None fit as well as the naïve behavioral model in which advantaged bidders simply adding their private value advantage to their estimate of the common value.

Rose and Levin (in press) (RL) investigate the effect of a private value advantage in the two-person wallet game, this time using an ascending-price English clock auction. The key motivation for this experiment is that in virtually all experimental work behavior is much closer to equilibrium predictions in English type auctions compared to sealed-bid auctions (KHL, 1987; LKR; KL, 1999). As such there is a clear need to explore the potential explosive effect of a private value advantage in English auctions before concluding that small asymmetries do not matter very much in practice, particularly since English auctions are far more common than second-price sealed-bid auctions. RL do not find any evidence of the explosive effect either, with players clearly suffering from the winner's curse in both the symmetric and asymmetric auctions, as evidenced by the frequency with which they lost money. When tested against the data, the Nash equilibrium model and the expected value hypothesis (naïve expectations) are both rejected, although the expected value hypothesis provides a better fit than the Nash model.

Rose and Kagel (in press) (RK) explore the effects of a private value advantage in ascending-price clock auctions having the same structure as the English clock auctions in LKR. They employ twice-experienced subjects who have come to earn relatively large positive profits in first-price sealed bid auctions. This is important since virtually all research in common value auctions shows that when bidders suffer from an obvious winner's curse, as they did in the control treatments in AK and RL, the comparative static predictions of the Nash equilibrium bidding model fail to hold. However, once bidders have learned to overcome the worst effects of the winner's curse, the theory's comparative static predictions tend to hold.[37] Here too bidders responses to the private value advantage are closer to proportional than explosive as advantaged bidders won only 27.0% of the auctions, little more than one would expect based on chance factors alone (25.0%). Further, there are no significant differences in seller revenue between the pure common value and almost common value auctions. RK show that the same behavioral model employed in AK, with advantaged bidders simply adding their private value advantage to their estimate of the common value, better organizes the data than does the Nash model with its explosive outcome. From a broader perspective these results demonstrate that adjustment to equilibrium under a trial and error learning process (which seems to be how subjects learn to overcome the winner's curse; see KL, 1986) provides no assurance that after near equilibrium behavior is achieved, the comparative static predictions of the theory will be satisfied. That is, once the underlying economic structure is changed, it may still take a whole new learning process to approach the new equilibrium outcome.

Takeover battles for control of a company when bidders already have stakes/shares (toeholds) in the target company are quite similar to almost common value auctions.[38] Consider the wallet game when each of the two bidders for the firm has an i.i.d. signal and where the common value of the firm is the sum of the two signals. In addition every bidder already owns a share Q_j, j = 1,2, of the target firm. In this setup the (relative) ratio of the shares, $Q_j/(Q_1+Q_2)$ has a dramatic impact on the predicted outcome as (i) the probability of winning the auction by bidder j is $Q_j/(Q_1+Q_2)$, (ii) increasing a bidder's share, Q, makes that bidder more aggressive and (iii) increasing a bidder's toehold increases her profits regardless of her signal. Thus, although the equilibrium does not lead to the explosive outcome when both bidders have positive toeholds,

[37] In employing an English clock auction with four bidders RK are better able to address Klemperer's concerns regarding the use of ascending price auctions on bidders' reluctance to enter the bidding process.
[38] The discussion here, as well as the experiment that follows is based on Bulow et al. (1999).

behavior is quite sensitive to even small disparities in the relative size of bidders' shares. This is quite surprising as it predicts the same equilibrium outcome regardless of the absolute size of the toeholds as long as the relative ratio is the same; e. g., the same outcome is predicted when bidders 1 and 2 hold 20% and 10% shares as when they hold 0.2% and 0.1% shares!

Georganas and Nagel (2007) explore the predictions of the toehold model using an English clock auction. They find that larger toeholds raise the probability of winning and the profits of their owners as the theory predicts, and that revenue tends to fall the larger the discrepancy between the shares of the two players' toeholds. However, as in the almost common value auctions r, these results are not nearly as dramatic as the theory predicts. The paper concludes that laboratory subjects do not respond to small toeholds or to small disparities in toeholds to the extent that the theory predicts.

To sum up: The results of all four experiments reported on here agree that contrary to what theory predicts, a private value advantage leads to proportionate as opposed to explosive effects in almost common value second-price and English auctions. This is true even with experienced bidders who earn a respectable share of RNNE profits in pure common value auctions. The apparent reason for these failures is that bidders do not fully appreciate the adverse selection effect conditional on winning, which is exacerbated for regular bidders when facing an advantaged rival. As such, the behavioral mechanism underlying the explosive effect is not present, and there are no forces at work to replace it.

This leaves us quite skeptical of finding similar effects outside the lab *under the conditions the theory specifies.* Indeed, it would seem to require very sophisticated bidders for the explosive effect to be realized under these conditions. As such we would expect that bidders outside the laboratory would employ alternative strategies available to them in the less structured environment they operate in to "press" their private value advantage. PacTel appears to have done something like this in the FCC major trading area (MTA) sale of broadband personal communications service licenses for Los Angeles and San Francisco. PacTel, which held a substantial private value advantage, publicly announced their intentions to top their opponents bids, while obviously having the resources and a sufficiently large private value advantage to make such an announcement credible (Cramton, 1997), a strategy that lies outside the formal theory. As a side note, PacTel got the licenses in question but they were only partially successful in obtaining rock bottom prices, as there was rivalrous bidding based the personalities of the leading bidders, another element left out of the formal theory.[39]

2.4 New Results in the Takeover Game: Theory and Experiments:

The systematic overbidding resulting in a winner's curse for inexperienced bidders has attracted the attention of theorists in efforts to explain this behavior within a generalized Nash equilibrium that allows for a more relaxed belief system. Eyster and Rabin (2005) (ER) generalize the Nash equilibrium model by introducing the notion of a "cursed equilibrium," in which bidders correctly predict, and best respond, to the *distribution* of others' bids, but do not correctly perceive how these other bids depend on signals (i.e., they relax the sophisticated set of beliefs underlying the Nash equilibrium in common value auctions). This model rationalizes deviations from the standard Nash equilibrium depending on the degree of "cursedness".

Crawford and Iriberri (2007) (CI) rationalize the winner's curse by also allowing more flexible beliefs. Roughly, they allow different levels of "sophistication" ("level-*k*" reasoning) where they define a level-0 player as a bidder who picks randomly from the allowable domain of actions and level-k players' best respond to all other players being one level less sophisticated than they are. The remarkable thing about this approach is that (i) a combination of level-1 and level-2 players explains most inexperienced subject behavior in first-price sealed bid auctions (i.e., the high prevalence of a winner's curse) and (ii) the estimated frequencies of the two player types closely matches the frequencies reported in a variety of other, unrelated, experiments; having (i) without (ii) would simply be an exercise in data fitting.[40] Both the ER and CI models apply to common-value auctions and the closely related "take over" game.

[39] As Cramton (1997) notes there may also be some incentive under these circumstances for predatory bidding on the part of rivals that would work against the revenue reducing forces implied by the explosive effect and indeed seem to have been at play in the MTA broadband sales.

[40] As CI note, their model can*not* explain the winner's curse in second-price common value auctions or the persistent overbidding in first-price private value auctions. Nevertheless, this paper is important because it shows a totally unanticipated result for first-price common-value auctions. The failure to explain overbidding in second-price common value auctions can be rationalized by the fact that subjects simply do not understand second-price auctions very well, whether private or common value.

Nash equilibrium bidding in common value auctions requires complicated calculations of one's best response, involving both beliefs about rivals' rationality and strategic uncertainty. To circumvent these complications Charness and Levin (in press) (CL) employ a modified version of the takeover game creating a much simpler individual decision making environment where rational bidding does not depend on beliefs about rivals actions. In the takeover game (first studied in Samuelson and Bazerman, 1985) there are two players, a buyer (the acquiring firm) and a seller (the target firm). The buyer knows that the target's value, V_S, is a random variable uniformly distributed in the interval [\$0, \$100]. The value of the target firm to the buyer, V_B, is $V_B = 1.5V_S$. A buyer does not know V_S at the time he bids, but the seller employs the dominant strategy of only accepting offers that are greater than, or equal to, her (known) value $V_S = v_S$. In spite of the simplicity of this game, which abstracts from many of the complications embodied in a multi-player auction context, subjects still suffer from not recognizing the adverse selection effect and succumb to a winner's curse, bidding somewhere between the unconditional expected value to the seller of 50 and the unconditional expected value to the acquirer of 75 (see Kagel, 1995; and KL, 2002 for summaries of results from these experiments).

CL transform the game into an individual-choice task where subjects make a bid and then choose one of 100 'cards' numbered {0, 1, 2,…, 99}, that are displayed face-down on the computer screen. The same rules apply as in the takeover game in that if the card chosen is less than or equal to the bid, players receive 150% of the current value of the card less their bid, and zero otherwise. However, there are no other human agents whose behavior bidders need to establish beliefs for, either in the sense of ER or CI.[41]

CL find average bids to be 38.9, which is lower than the 50-75 average typically reported when the game is framed in terms of sellers rejecting bids that are below their value. Further, there are a higher proportion of very low bids (around 25%) in the 0-9 range than is typically reported. To further simplify the task, CL employ a version of the game where there only two possible values for the firm: 0 or 99 (50 cards with 0's and 50 cards with 99, face down and shuffled). This treatment circumvents the need to use

[41] Also see Tor and Bazerman (2003) who argue that the reason subjects succumb to the winner's curse in the takeover game is that they ignore the cognitive process of the sellers.

Baye's law to construct posterior beliefs, as well as the need to recognize the implications of the firm's values being drawn from a uniform distribution. Now, without any calculations, it should be clear that 0 dominates all other bids with the exception of 99, and in choosing between 0 and 99 the choice of 99 involves a rather unattractive gamble between a positive profit of 49.5 or a negative profit of 99, with both outcomes equally likely, for an expected profit of *negative* 24.25.

[Insert Figure 4 here]

Figure 4 reports the results from this last treatment under two sets of instructions, with one providing more detail than the other. While there are relatively few bids other than 0 and 99, 47% of all bids are still at 99. Since the latter may reflect a need for some "action" as opposed to always bidding zero, CL further modify the game so that the card values are either 20 or 119. Now, bidding 20 yields a profit of 10 half the time and zero otherwise. This results in an even further reduction in the frequency of non-optimal bids to 30%.[42] Finally, CL have subjects choose between lotteries whose payoffs are equivalent to the 0-99 and 20-119 treatments to rule out risk loving as a possible explanation for non-optimal bids.

CL note that taken literally, converting the takeover game to an individual choice game rules out both the ER and CI models as an explanation for the winner's curse, since there is no other player whose actions must be taken into account. However, assuming that subjects still frame the situation as a two-player game, with the computer as the second player, results from the two card treatments are inconsistent with both models as they both predict that all will be equal to 0 or 20.

To see this, recall that a cursed bidder gets the distribution of acceptances right, but doesn't take into account the correlation between acceptance and the seller's value. In the two-value case subjects know that bids less than 99 are accepted with probability 1/2 and bids of 99 or higher are accepted with probability 1. However, being cursed the bidder does not take into account what this implies for value, calculating it to be $\frac{3}{2}\times\frac{(0+99)}{2}=74.25$. Now the payoff from bidding b < 99 is $\frac{1}{2}\times(74.25-b)$ and from bidding b ≥ 99 is $1\times(74.25-b)$, so that the optimal bid is 0 by a simple dominance argument, no

[42] There is little change over time in the frequency of 0 bids or bids of 20 in the 0-99 and 20-199 treatments respectively. Further, it does not appear that many subjects consistently bid 0 or 20 in these treatments as well.

matter the level of cursedness (χ) in the ER model. Parallel reasoning holds when values are shifted to 20 and 119.

For the level-*k* model to have meaning in the context of the CL experiment one must assume that the computer plays the role of the random, level-0 player. But then the computer accepts *any* bid randomly, in which case the best response is to bid the smallest possible value, since this will maximize profits with no impact on the chance the bid will be accepted.

CL conclude that the fundamental problem underlying the winner's curse is the failure to fully account for payoffs contingent on winning the auction.

2.5 Additional Common Value Auction Results

2.5.1 Super Experienced Bidders: Kagel and Richard (2001) (KR) investigate bidding for super-experienced bidders - subjects who had participated in at least two, and up to four, prior auction sessions. These super-experienced bidders had learned to overcome the worst effects of the winner's curse in first price common-value auctions, rarely bidding above the expected value conditional on winning. However, they still earned less than 50% of the Nash equilibrium profits (at a cost of between \$2.50 - \$3.50 per auction, conditional on winning). KR examine a number of elements that might be responsible for the continued shortfall relative to the RNNE benchmark.

They first look at the bid function itself, which is quite complicated over the full support from which signals are drawn, to see if rules of thumb that boundedly rational bidders might employ were at fault. They find that subjects employ sensible piecewise linear bid functions that differ systematically from the RNNE benchmark (they are far too flat to begin with and completely overlook the nonlinear elements of the Nash bid function). However, simulations using these piecewise linear bid functions identify a symmetric rule of thumb equilibria (RTE) (an equilibrium in which bidders are restricted to using piecewise linear bid functions of the sort estimated) in which profits are equal to or *greater* than the RNNE benchmark. As such, bidders' inability to account for the complexities of the Nash bid function cannot account for the marked reduction in their earnings. KR also show that subjects are *not* best responding to overly aggressive bidding by rivals, as large sample estimates of best response intercepts are remarkably robust averaging around -\$18.00 with , = \$18; i.e., best responding to rivals overly aggressive bidding typically requires that intercepts of the piecewise linear bid functions be close to the symmetric RTE benchmark. Losses relative to best responding averaged 20% and 44% in auctions with 4 and 7 bidders, respectively. Thus, very experienced bidders still suffer from a winner's curse, albeit one that is much less pronounced and more subtle than the negative average profits inexperienced bidders suffer from.

KR suggest two primary reasons for these continuing losses relative to the RNNE and the RTE benchmarks. First, best responses are highly variable in small samples of the sort that bidders would have seen, sometimes pointing in the wrong direction (bid higher) and sometimes implying overly passive bidding (bid below x - ,). This makes best responding far more problematic than the large sample estimates suggest, and it could lead bidders to simply ignore any feedback once consistently positive profit levels were achieved. Second, large sample best responses require winning half as many auctions as were actually won. This involves a rather dramatic change in bidding, assuming that bidders are able to identify this fact. As such given that bidders are earning relatively large positive profits compared to their inexperienced selves, they may be reluctant to deviate from a rule of thumb that has proved capable of generating acceptable profit levels in such a high variance environment.

2.5.2 Auctions with Both Common and Private Value Elements: One of the simplifications in standard auction theory is that bidders are dealing with either a pure common value *or* a pure private value environment. However, most real world auctions have both private and common value elements. For example, bidders for oil leases may have a single estimate for the common-value of the oil in the ground that is affiliated with other bidders' estimates and, in addition, have an idiosyncratic cost of extracting the oil and delivering it for refining. The theoretical (and maybe behavioral) difficulty with multiple signals for all bidders is how to combine them into a single statistic that can be mapped into a bid. Goeree and Offerman (2003) develop such a model, and test it experimentally (Goeree and Offerman, 2002; GO).[43]

GO investigate a series of first-price sealed bid auctions, the main objectives of which are to evaluate those factors theory predicts will raise efficiency and revenue. These

[43] In their formulation the common value component depends on the *average* of bidders common value signals.

consist of (i) reducing the variance of the common value signals which ought to make the auction more efficient as it moves the environment closer to a pure private value auction, (ii) increasing the number of bidders which, in their design, reduces the weight bidders place on their common value signals thereby increasing efficiency, and (iii) releasing public information that reduces the importance of the common value element, thereby increasing efficiency and raising revenue as in a pure common value auction. In all cases, both the "rational" bid function in which bidders fully account for the adverse selection effect conditional on winning, and a naive bid function in which bidders fail to do so, predict the *same* winner as they are both functions of the same summary statistic. Hence, there is no efficiency loss in their design due to the winner's curse.

[Insert Figure 5 here]

GO report that the winner's curse is alive and well in their experiment as bids lie in between the naïve and Nash benchmarks (see Figure 5) even for experienced bidders, and lie closer to the naïve benchmark the less important a bidder's private value is relative to the common value.[44] Realized efficiency is roughly at the level predicted under the RNNE, with the winner's curse serving to raise seller revenue and reduce bidders' profits. This occurs because (i) almost all bidders suffer from a winner's curse and (ii) the degree of suffering is roughly the same across bidders, so that the size of the private value element serves to dictate who wins the auction. As predicted, efficiency increases the smaller the variance in the support the common value signals are drawn from, and with increases in the number of bidders. Ignoring the low variance treatment with its minimal scope for a winner's curse, public information regarding the common value increases bidders' profits in four out of five treatments, consistent with the comparative static prediction of the naive bidding model and the results reported in KL (1986).

2.5.3 Selection Bias and Demographic and Ability Effects: The transition from inexperienced bidders suffering persistent losses to experienced bidders earning respectable profits in common value auction experiments is characterized by large numbers of bidders going bankrupt, with these bankrupt bidders much less likely to return as experienced subjects. Further, the winner's curse involves a judgmental error –

[44] Bankrupt bidders from inexperienced subject sessions were *not* invited back for experienced sessions. This generates potential selection effects discussed in the next subsection.

the failure to account for the adverse selection effect conditional on winning – so that it joins a growing literature suggesting that limited cognitive abilities might help to explain many of the observed deviations from full rationality reported in experiments. Casari, Ham and Kagel (2007) (CHK) conduct an experiment designed to better understand the process whereby experienced bidders learn to avoid the winner's curse as well as the impact of ability and demographic effects on their ability to avoid or overcome it. Ability effects were measured by Scholastic Aptitude and American College Test (SAT/ACT) scores collected from university records. Unversity records also provided information regarding college major, grade point average (GPA), and gender.

Subjects participated in two sessions approximately one week apart. Starting cash balances were randomly varied across bidders with additional, random shocks to these balances through a lottery with positive expected value (recall the HKL experiment reported in section 1.6 above). Further, some sessions followed standard experimental procedures inviting all subjects back for additional sessions without any special inducements to return, while others recruited subjects who were committed to returning and were provided strong incentives to do so in the form of show-up fees and half of session one's earnings withheld until completion of session two. In this way CHK hoped to distinguish between learning via market selection effects (less able bidders going bankrupt, exiting the market and not returning for subsequent experimental sessions) versus individual bidders learning to avoid the winner's curse.

CHK report a number of substantive as well as methodological insights: First, not surprisingly, ability as measured by SAT/ACT scores matter in terms of avoiding the winner's curse. However, the nature of these ability effects are different from what one might expect as (i) *composite* SAT/ACT scores were consistently more significant than either math or verbal scores alone and (ii) the biggest and most consistent impact was that subjects with *below* median scores were more susceptible to the winner's curse, as opposed to those with very high scores doing exceptionally well. Second, there were clear demographic effects as inexperienced women were much more susceptible to the winner's curse than men, even after controlling for ability and college major, factors that are not typically controlled for in investigating gender effects in experimental economics.[45] However, women learned faster than men so that this difference disappeared with experienced bidders. Economics and business majors were much more susceptible to the winner's curse than other majors, and continued to earn lower profits even as experienced bidders controlling for SAT/ACT scores and gender. Controlling for selection effects, bidders are capable of substantial individual learning, even those subjects who start out being most susceptible to the winner's curse. However, more able bidders were more likely to return as experienced subjects, with this factor dominating learning between weeks 1 and 2 for those sessions that did not control for the winner's curse. As such previous studies that have not controlled for selection effects are likely to have substantially *overestimated* the amount of individual subject learning that occurs when moving from inexperienced to experienced bidders.

CHK also find that standard econometric estimators for dealing with selection effects in field data do not identify the selection effects present in their data, in spite of having a relatively large sample by experimental standards and well identified econometric models. However, the different experimental treatments built into the experimental design serve to identify, measure, and verify these effects. The latter is not surprising, since at least as far back as Fisher (1935) statisticians have understood that experimental design could permit the identification of casual effects.

As to why economics and business majors were more susceptible to the winner's curse, CHK suggest that this is more than likely a personality effect – business and economics students are by nature aggressive in business-type transactions – as the data are inconsistent with the hypothesis that a 'little knowledge is a dangerous thing.' The gender effect is much more difficult to explain. Two known factors that immediately come to mind, that women tend to be more risk averse than men and that men tend to be overrepresented in the upper tail of mathematical reasoning, fail as (i) risk aversion cannot explain succumbing to the winner's curse since the latter involves earning *negative* expected profits and (ii) the estimated bid functions show that mathematical ability does *not* play a critical role in succumbing to the winner's curse (and was controlled for in the statistical analysis). CHK conjecture that the greater susceptibility of

[45] Similar gender effects are identified by Charness and Levin (in press) in the closely related takeover game.

women to the winner's curse may reflect a relative lack of experience with strategic interactions compared to men, perhaps as a result of women shying away from competition more than men (Niederle and Vesterlund, 2007; see Chapter xx as well). This relative lack of familiarity might induce more aggressive bidding as a consequence of the failure to fully think through its implications.

Remark: CHK also compared their sample population to the university population from which their sample was drawn. The most interesting result here is that 20.2% of their sample were in the top 5% (of the national average) with respect to composite SAT/ACT scores (versus 4.9% for the university), with less than 8.9% scoring below the median (versus 20.9% for the university), indicating that much brighter students (as measured by SAT/ACT scores) enrolled in their experiment. Since there were no special elements associated with recruiting subjects in this case, these results suggest that subjects who voluntarily enroll in economics experiments more than likely over represent high ability students.

2.5.4 Is the Winner's Curse Confined to College Sophomores?: One inevitable question raised by laboratory experiments is whether the behavior reported is confined to the typical population of convenience, undergraduate students, as opposed to "real people" in field settings. Kagel's (1995) survey addressed this question in two ways: First, it reported a number of striking similarities between anomalous field data and the experimental outcomes that could be directly attributed to the winner's curse; e.g., public information lowered revenue in both cases.[46] Second, results from a laboratory experiment comparing experienced bidders from the construction industry with student subjects were reported that showed essentially no difference in the intensity of the winner's curse between the two subject populations; i.e., both suffered from a strong winner's curse (Dyer, Kagel and Levin, 1989). Follow up research suggested two key factors, which are not mutually exclusive, behind the executives performance in the lab and their apparent success in the field (Dyer and Kagel, 1996; DK): One is that the executives had learned a set of situation-specific rules of thumb which permit them to avoid the winner's curse in the field, but which could not be applied in the laboratory, such as their specialized experience with a given branch of the construction industry or familiarity with the architect responsible for supervising the work. Second, the bidding environment created in the experiment, which is based on theoretical work, is not fully

[46] To be sure there are alternative explanations for the field data (see KL, 1986), but the winner's curse is a much more straightforward explanation than the alternatives offered.

representative of the environment encountered in the construction industry; e.g., repeated play elements present in the field typically permit bidders to pull winning bids that are clearly too low relative to the expected cost of the project, and to do so without penalty.

Harrison and List (in press) (HL) report results that appear to be at odds with the contactor results. In their experiment they compare bidding by sports card dealers with non-dealers in a laboratory type setting under the symmetric information structure employed in KL (1986) (each bidder gets a random signal from the interval $[x_o - \varepsilon, x_o + \varepsilon]$) as well as the asymmetric information structure employed in KL (1999) (one bidder knows the true value, x_o, with certainty while all other bidders get a signal drawn from the interval $[x_o - \varepsilon, x_o + \varepsilon]$). Subjects bid in a single auction after having participated in 10 practice auctions. Treatments included two different values of ε ($6 and $12) and two different levels of competition – auctions with 4 and 7 bidders.[47]

Their results show that with symmetric information dealers rarely suffer from a winner's curse as they typically bid below the expected value conditional on winning. In contrast, non-dealers tend to bid above the expected value conditional on winning, with these differences statistically significant at conventional levels. Further, there are significant differences in the estimated bid function between dealers and non-dealers, with much of the difference resulting from the sharper discounting of bids relative to value with ε = $6 for dealers.[48] In the asymmetric information laboratory treatment HL are unable to reject a null hypothesis that dealer "outsiders" bid differently from non-dealer outsiders (outsiders being the only bidders susceptible to the winner's curse in the asymmetric information treatment; see section 2.2 above), with dealers suffering from a non-negligible frequency of the winner's curse (between 25%-30%).

HL interpret these results as follows: The absence of a winner's curse with symmetric information is consistent with the notion that dealers have experience in comparable settings. Further, since this experience is generated in the field and not in the lab, it provides support for the notion that "… context-specific experience does appear to

[47] Subjects are not provided with any starting capital balances or participation fees to cover potential losses in HL's experiment, bidding in isolation and returning at a later time to determine their earnings. Unfortunately HL do not report what was done in the case of losses or how the dry runs were conducted, information that is important for fully interpreting their results.
[48] The differential shading of bids relative to signal value averaged 40% (82%) for non-dealers versus 93% (88%) with ε = $6 ($12).

carry over to comparable settings, at least with these types of auctions." However, once dealers are taken out of their comfort zone, bidding as outsiders in the asymmetric information auctions, a role HL argue dealers rarely occupy in field settings, they look very much like the student subjects.

These results for dealers are at odds with DKL's results for construction contractors as well as experiments showing that student subjects have difficulty readily generalizing from one common value auction environment to another; for example in going from auctions with 4 to 7 bidders (KL 1986) or in going from a pure common value auction to an almost common value auction (see Section 2.3 above). Their results also represent a rather remarkable counter example to the psychology literature on learning generalizability which indicates that learning transfer, unless specifically taught for, does *not* generalize easily across different domains.[49] This is even more remarkable since HL do not specify any activities that traders routinely engage in that would establish experience related to common value auctions that could be generalized to their laboratory experiment, while in previous studies List and his co-authors assume that *the trading card market is best approximated by a private value auction* (List and Lucking-Reiley, 2000 and Engelbrecht-Wiggans, List and Reiley, 2006; reviewed in section 3.1 below).

As such we look for an alternative, *artifactual* basis for the dealers' superior performance in HL's symmetric information treatment. One need not look very far: Dealers in buying trading cards must purchase them at low enough prices to be able to sell them at a profit and would most certainly be in the habit of doing so; e.g., the LLR experiment shows that dealers bid just under \$50 for cards with a retail value of \$70 in a Vickrey auction. (Also see GWW discussed in section 1.4 above, who report similar results for eBay sellers in a second-price sealed bid auction with induced values.) So applying such large discount factors, which non-dealers would not be in the habit of

[49] For a good primer from the psychology literature on learning generalizability see Salomon and Perkins (1989). Or as the Noble laureate Richard Feynman (2005, p. 39) put it: "I don't know what's the matter with people: they don't learn by understanding; they learn by some other way - by rote, or something. Their knowledge is so fragile."

doing, could very well protect them from a winner's curse in the form of losses, but *not for the reasons that HL suggest*.[50]

Remark: HL also report a treatment in which subjects bid to purchase an unopened package of L*eaf* sports cards, packages containing 10 cards of unknown value, and an *established* retail price of between $9-$10. They argue that this represents a common value auction, which we agree with. However, it is *not* one in which there is any scope for a winner's curse since the cards have a well known market value; i.e., there is no scope for an adverse selection effect based on different estimates of value that anyone but a very poorly informed buyer might have. As such, this exercise is comparable to auctioning off a $10 bill. Plots of bid distributions bear this out as there is not a single bid above $10 for dealers and only a handful of bids above $10 for non-dealers.

HL also make a number of broader and related claims that deserve some discussion. First, they claim that their approach to "…*undertake(ing) experiments in naturally occurring settings in which the factors that are at the heart of the theory are identifiable and arise endogenously, and then to impose the remaining controls needed to implement a clean experiment*" (i.e., the *Leaf* trading card experiment using dealers; italics in the original) is superior to imposing controls exogenously on "a convenient sample of college students." As noted the outcome from this part of their experiment is totally irrelevant to identifying the presence or absence of a winner's curse since the factors at the heart of the theory are simply not present. Consequently, the lesson *we* take away from this treatment is *the need to establish a clear correspondence between a theory and its implementation to be able to draw valid conclusions from an experiment.* Second, they claim that the absence of a winner's curse among dealers in both the *Leaf* trading card treatment and in the more abstract laboratory treatment with symmetric information is "… consistent with the conclusion that *dealers in the field do not fall prey to the winner's curse providing tentative support for the hypothesis that naturally occurring markets are not in disequilibrium because of the WC* (winner's curse)" (italics in the original). Here, we would remind the reader that the term "winner's curse" was initially coined by three petroleum geologists (Capen, Clapp and Campbell, 1971) reporting on results from early outer continental shelf (OCS) oil lease auctions in an effort to explain low (or below normal) returns on these leases. The debate that this assertion set off for OCS leases, as well as similar claims in other settings, is what motivated experimental work investigating the winner's curse.[51] The fact that these experiments showed that the winner's curse is alive and well, persistent and robust, suggests that it is likely to exist *at least in the start up phase* of auction markets with a strong common value element. Finally, let us assume, that HL are correct that in relatively settled markets with very experienced bidders survivors no longer fall prey to the winner's curse. To us this is similar to arguing that in an environment ravaged by an infectious disease the disease no longer exists since the survivors have developed

[50] Why doesn't this same process help the construction contractors reported in DKL? The answer is simple. General contractors do not buy and sell in anything approaching the same way that card dealers or eBay sellers do. Rather they solicit bids from large numbers of subcontractors who are responsible for fulfilling their commitments, and then add in their own estimated general contactor costs.

[51] For example, auctions for book publication rights (Dessauer, 1981), professional baseball's free agency market (Cassing and Douglas, 1980; Bleckerman and Camerer, 1998), corporate takeover battles (Roll, 1986), and real estate auctions (Ashenfelter and Genesore, 1992).

immunity to the disease. It does not however imply that should a new disease strike the community that the survivors will be able to do any better than those who were never impacted by the original disease. The analogy to auction outcomes here is that they will not necessarily exhibit equilibrium responses to changes in the auction environment, such as increased numbers of rivals or the introduction of an insider with better information about the common value.

III. Multi-Unit Demand Auctions

Theoretical and experimental research up to 1995 focused almost entirely on auctions where each bidder demands a single unit of a homogenous commodity. Not much changes in the theory if the number of units for sale is greater than one as long as individual bidders continue to demand a single unit. However, in auctions where bidders demand multiple units, outcomes can change rather dramatically. The FCC spectrum auctions in the 90's provided the main incentive to better understand auctions where bidders demand multiple units, raising a host of new issues, many of which are of public policy importance. (The extensive use of Internet auctions also has played a major role in stimulating auction research as well.) Where and how can one design efficient multi-unit demand auctions? Are efficient multi-unit demand auctions very different from optimal (revenue maximizing) auctions? Multi-unit demand auctions also call attention to a much richer strategic environment where bidders may exercise demand reduction, bidding "passively" on some units in order to obtain other units at low prices. They also call attention to the difficult case of complements, with strong synergies generated as a consequence of winning multiple units, and the potential role of package bidding to help achieve more efficient outcomes.

In looking at multi-unit demand auctions we need to distinguish between small scale, traditional laboratory experiments designed to investigate some of the new theoretical/behavioral issues identified in the literature as opposed to mechanism design issues. In the latter the laboratory serves as a "wind tunnel" for comparing different mechanisms for specific public policy purposes, and there are virtually no comparable *series* of experiments against which to evaluate results (and often not much emphasis on the behavioral mechanisms behind the results reported). We will review the more traditional small scale experiments that focus on behavioral issues here, with mechanism design issues covered in Chapter xx.

3.1 Auctions with Homogeneous Goods - Uniform-Price and Vickrey Auctions: In multiple-unit, uniform-price auctions items are allocated to the high bidders at a price equal to the highest rejected bid. With bidders demanding multiple units, if the goods are substitutes, bidders have an incentive to reduce demand in an effort to obtain more favorable prices on the items actually won (Ausubel and Cramton, 1996; Englebrecht-Wiggans and Kahn, 1998). The argument for demand reduction is essentially the same as a monopsonist who takes account of the fact that with increased demand, the price they pay will increase as well. Cramton (1997) argues that the first nationwide FCC spectrum auctions could be best modeled as uniform price auctions of this sort.

Kagel and Levin (2001) experimentally investigate the sensitivity of bidders to these demand reduction possibilities, comparing behavior under a sealed-bid uniform-price auction with an ascending price, English clock auction in which bidders receive information regarding rivals' drop-out prices as the auction progresses. They study behavior in the simplest possible setting while still preserving the essential strategic elements of more complicated auctions: A human subject with flat demand for two units of a homogeneous commodity competes against different numbers of rivals demanding a single unit of the commodity, with the role of single unit buyers played by computers whose bids are equal to their private value (a dominant strategy for single unit buyers).

With independent private values drawn from a uniform distribution and with supply of two units, the equilibrium prediction for the "large" (human) bidder is to bid her value on unit 1 and to bid sufficiently low on unit 2 so as to not affect the market price. This holds irrespective of the value of the item, the number of computer rivals, or whether a sealed-bid or English auction format is used. For the sealed-bid auctions this requires bidding zero on the second unit, and is far from transparent. In contrast, the optimal bidding strategy in the clock auctions requires dropping out on the second unit at a price $p \in [0, v_2]$ where v_2 is the drop-out price of the second highest computer rival. This has exactly the same consequences as dropping out at 0, but the feedback information provided by rivals dropping out, and the flexibility in the dropping rule, makes the optimal bidding strategy substantially more transparent.[52]

[52] For example, assume a support for values of between [0, 100] with the values for both units for the human bidder, v_h, of 90. Suppose that *h* has no formal understanding of the optimal bidding strategy and

Results from this experiment showed clear evidence of demand reduction in the uniform-price auctions, but with substantially more demand reduction in the English auctions: 30.8% of all unit 2 bids were pivotal (higher than v_2, thereby setting the market price) in the sealed-bid auctions compared to 11.4% in the clock auctions.[53] However, there were even more striking differences between the two auction formats as: (i) There was a much higher frequency of bidding above value on the first (and even the second) unit in the sealed-bid auctions (comparable to the results for single-unit demand Vickrey auctions) and (ii) there were relatively few bids at 0 in the sealed-bid auctions, where they would have to be to insure not being pivotal. Figures 6 and 7 illustrate these differences between the two uniform-price auction formats with five computer rivals.

[Insert Figs 6 and 7 here]

KL show that the primary basis for the superior performance of the uniform-price clock auction over the sealed-bid version results from the feedback information regarding the computer's drop out prices. They did this in two ways. First, they conducted a clock version of the uniform-price auction in which there was no feedback, with the auction ending when the last bidder dropped out. In this case the clock was of no help to bidders as there was massive overbidding on both units quite similar to what was found in the sealed-bid auctions. They also conducted a sealed-bid version of the uniform-price auction in which v_2 was posted in a prominent place on bidders' computer screens. Subjects were not told how to use the information, just what it was and that it had been suggested that the information might prove helpful in determining how to bid. This treatment went a long way to moving the sealed-bid outcomes closer to the clock results as (i) it essentially eliminated the overbidding in the sealed-bid auctions and (ii) resulted in a level of demand reduction closer to the one reported for the clock auctions.

---

decides to remain active as long as $p \leq v_h$. Suppose that v_2 drops out at 50. Now h has two options, drop at 50 and earn an instant profit of 40 or remain active in an effort to win both units. In the latter case there are two events to consider (i) the highest computer rival (v_1) drops prior to $p = v_h$ in which case h's expected profit is 40 (as 70 is the expected drop price for v_1) or (ii) $v_1 \geq v_h \geq 90$ in which case h's expected profit is zero. Thus, dropping at $p = v_2$ dominates waiting and trying to win two units. This is not to say that these calculations are trivial but they are far more transparent than the ex ante calculations underlying the optimal bidding strategy in the sealed-bid auctions. Further, if h remains active once $p > v_2$ it should be increasingly transparent that she is competing against herself, which should lead to dropping out before the price is equal to v_h, which might help promote learning over time.

[53] Results are pooled over auctions with 3 and 5 computer rivals. All data are for the last 12 auctions in a session. Subjects were never told that the computers were following a dominant strategy, just that they would drop out at their randomly drawn values.

KL also compared outcomes in the uniform-price auctions to a dynamic Vickrey/Ausuble auction (Ausbel, 2004). This dynamic version of the Vickrey auction with drop-out information provided employs a "clinching" metaphor from sports leagues to characterize prices paid.[54] It generates sincere bidding in iterated deletion of dominated strategies and, under the demand structure employed in KL, is predicted to raise more revenue than the uniform-price auctions. Results from the Ausubel auctions are shown in Figure 8, where outcomes are reasonably close to sincere bidding.

[Insert Figure 8 here]

Comparing bidders' actual earnings relative to predicted earnings as a measure of how close bidders were to optimal outcomes, KL establish a clear ranking for the three auction institutions studied: the uniform-price sealed-bid auctions are furthest from the maximum predicted (only 13.6% of all subjects averaging within 5% of maximum possible profits), the uniform-price clock auction with feedback is next (46.5% of all subjects averaging within 5% of maximum possible profits), with the Ausubel auctions closest to the maximum (85.2% of all subjects averaging within 5% of maximum possible profits).[55] They conclude that like the uniform-price clock auction with feedback, the Ausubel auction benefits from the clock procedure with feedback to prevent overbidding. However, unlike the uniform price clock auction, the Ausubel auction encourages non-strategic bidding (full demand revelation), something that bidders are inclined to do even in the uniform-price auctions. Thus, the closer to optimal performance observed in the Ausubel auction partly results from an institution that accommodates itself to bidders' natural tendencies rather than any adjustments on bidders' part to the strategic requirements of the institution.

[54] Clinching works as follows: With 2 objects for sale, suppose at a given price, *p*, the human bidder (*h*) still demands two units, but the aggregate demand of all *other* bidders has dropped from 2 to1. Then, in the language of team sports, bidder *h* has clinched a unit no matter how the auction proceeds. As such, at that moment, *h* is awarded one unit at the clinching price, *p*. The auction continues with the supply reduced from 2 to 1 and *h*'s demand reduced to one unit. This process repeats itself until all units are allocated. In this way the auction sequentially implements the Vickrey rule that each bidder pays the amount of the $k^{\mathbf{th}}$ highest rejected bid, other than his own, for the k^{th}t unit won.

[55] Recall that in this experimental design closeness to equilibrium and closeness to maximum payoffs are one and the same since computer rivals all play their Nash strategies, so that profits provide a suitable metric for comparing outcomes across institutions. Z statistics using individual subjects as the unit of observation show all three of these differences to be statistically significant at better than the 1% level.

Remark: KL's uniform-price sealed-bid auction instructions included explicit advice against subjects bidding above their values along with examples as to how this could lead to losses. Motivation for this advice was to speed up equilibrium outcomes on unit 1 bids, a "nuisance" factor in terms of KL's primary interest of investigating demand reduction with respect to unit 2 bids. These procedures were criticized as biasing the sealed-bid outcomes too strongly in favor of equilibrium outcomes, in response to which additional sessions were run dropping the advice. As anticipated the primary impact was to reduce the frequency unit 1 bids above value, with essentially no impact on the overall frequency of demand reduction. The point of this remark is not to show that demand reduction in the sealed-bid auctions is robust to these procedural differences, but that by the turn of the century, with experiments in economics firmly entrenched in the economists tool kit, and behavioral economics making its way onto the stage, the referees and editors of a major journal were concerned with biasing procedures in *favor* of the theory. From our personal experience this reflects a significant (and welcome) shift from earlier referees (and journal) biases in favor of experimental outcomes *supporting* the theory under investigation with little regard, in some cases, to procedural biases that favored the theory.

There have been a number of subsequent experiments using all human bidders investigating these issues. List and Lucking-Reiley (2000) (LLR) look at demand reduction in a field experiment with subjects bidding for sports cards in a uniform-price sealed-bid auction. Each auction had two bidders who could bid on two identical units with supply of two units with subjects participating in a single auction. Since LLR do not know bidders' value for the sports cards, they employed a parallel series of sealed-bid Vickrey auctions (in which sincere bidding is a dominant strategy) as the reference point against which to evaluate the presence and extent of demand reduction in the uniform-price auctions.[56] They employed two types of sports cards – one with low ($3) and one with high ($70) book value – and conducted separate auctions for dealers and non-dealers.

LLR find that unit 2 bids are systematically lower in the uniform-price compared to the Vickrey auctions, with these differences statistically significant for the high valued cards for both dealers and non-dealers: $41.77 versus $30.60 for dealers, $28.82 versus $16.62 for non-dealers. They also found that unit 1 bids were consistently higher in the uniform-price auctions, with these differences statistically significant for the higher

[56] LLR recognize that their design also has an equilibrium in which subjects bid their value on both units (Ausubel and Cramton, 1996), but argue against this equilibrium on the grounds that it is a knife edge case, so that any small reduction in value for the second unit, which is likely to be present, would eliminate it.

valued sports cards. This stands in marked contrast to the sincere bidding predicted for unit 1bids in both cases.

Porter and Vragov (2006) replicate the LLR experiment, only with induced values and adding a uniform-price clock auction treatment to the mix. Sessions consisted of 30 auctions, with new randomly drawn valuations and auction partners in each auction. With supply of two units and two bidders each demanding two units with the same value they have the same multiple (symmetric) Nash equilibrium problem as LLR.

Their results largely replicate those reported so far. First, for the uniform-price sealed-bid auctions there is rather massive overbidding with respect to unit 1 bids and relatively large scale demand reduction with respect to unit 2 bids (see Figure 9). For the clock auctions unit 2 prices are close to their starting price and well below prices in the sealed-bid auctions, consistent with strong demand reduction. Their sealed-bid Vickrey auctions exhibited substantial bidding above value for both units (see Figure 10), consistent with the results reported for single-unit Vickrey auctions. Finally, they report more overbidding on unit 1 bids in the sealed-bid than in the Vickrey auctions, similar to the anomaly reported in LLR.

[Insert Figures 9 and 10 here]

Englemann and Grimm (2004) (EG) also investigate bidding for two homogenous items in auctions with two bidders each with flat demand for both units. They test for whether demand reduction occurs in both the clock and sealed-bid uniform-price auctions, and compare outcomes with a dynamic Vickrey (Ausubel) auction and a static Vickrey auction.[57] Their experimental design is hampered by the fact that subject pairings remain fixed over the full set of 10 auctions in each treatment, which leads to scattered efforts to promote collusive outcomes. However, after factoring out these collusive efforts, they conclude that their primary results are well in line with those reported in KL: (1) there is more demand reduction in the uniform-price clock auctions than in the uniform-price sealed-bid auctions, (2) there is close to sincere bidding in the Ausubel auction, and (3) there is a higher frequency of bidding above value on unit 1 bids in the uniform-price sealed-bid auction than in the uniform-price clock auction or the

[57] EG explicitly recognize that in their experimental design for the uniform-price auctions represents a knife edge case with multiple equilibria, both with and without any demand reduction, but argue that the demand reduction equilibrium is more plausible.

Ausubel auction. Based on this last result, they note that contrary to the theory, the uniform-price sealed-bid auction generates higher revenue to the auctioneer but lower efficiency than the Ausubel auction so that there might be a tradeoff there for the seller, a point that KL noted as well. Like LLR they find more overbidding on unit 1 in the uniform-price sealed-bid auctions than in the static Vickrey auction. The one inconsistency with earlier results is that they find surprisingly little learning within and across auction formats, in contrast to the modest learning reported in KL.

Remark: The one anomalous finding in these experiments is that unit 1 bids in the uniform-price sealed-bid auctions exceed those in the Vickrey auction. This cannot be attributed to the auction format, as both are sealed-bid. Levin (2005) argues that with two bidders each demanding two units with supply of two units there is a very appealing low revenue (implicitly collusive) Nash equilibrium that is also an *ex-post* equilibrium in which both bidders bid above their private value on their unit 1 bids, and zero on their unit 2 bids. Although the strategy of "bidding above value on unit 1" involves weakly dominated strategies, unlike the equilibrium LLR focused on, the alternative NE is particularly attractive given their experimental design as: 1. Being an *ex-post* equilibrium it is distribution free, an attractive equilibrium both in controlled laboratory experiments and in field studies where distributions are not induced; 2. The equilibrium has *no-regret,* so there is less incentive to correct actions than the one focused on in LR; 3. It allows a wide range of bidding (any high bid on the first unit will work) so it is easier to coordinate on; 4. It has a more equitable distribution of payoffs so is more robust against *other-regarding* preferences (e.g., envy or spite). Englebrecht-Wiggans, List and Reiley (2005) (EWLR), in their response to Levin's comment, argue that (i) the data in LLR are inconsistent with the beliefs underlying this alternative equilibrium, as it depends on both agents bidding zero on unit 2, while nearly three-fourths of all unit 2 bids were strictly positive in their experiment[58], and (ii) there is a similar equilibrium in weakly dominated strategies for their sealed-bid Vickrey auctions resulting in incentives to bid above value there as well. This alternative Nash equilibrium collapses once there are three or more bidders (with two units supplied) and/or with a positive reserve price. So that from a design point of view a simple change of parameters serves to test the predictions of Levin's alternative equilibrium. EWLR (2006) report such a test in an experiment with supply of two units and more than two bidders each demanding two units. The results of that experiment show that with three or five bidders unit 1 bids in the uniform-price auction are statistically indistinguishable from the Vickrey auction.

Summing Up: Both uniform-price sealed-bid and clock auctions with homogenous goods generate demand reduction as the theory predicts. But there is substantially more demand reduction and closer to equilibrium bidding in clock auctions which, although

[58] However, bidding zero on unit 2 is also possible in LLR's equilibrium since the underlying distribution of values is unknown and bidding errors can, at least partially, account for positive bids on unit 2, since deviations from Levin's proposed equilibrium can only be positive.

inconsistent with the theory, is consistent with results for single-unit demand auctions comparing second-price sealed bid and English clock auctions. The key mechanism behind these differences in the two uniform-price auctions formats appears to be the feedback provided by other bidders' drop-out prices in the ascending price version of the auction, which simplifies identifying better responses that tend to be closer to the equilibrium outcome. Ausubel's version of the dynamic Vickrey auction eliminates much of this demand reduction, with close to equilibrium outcomes (sincere bidding) as well. The static Vickrey auction generates overbidding relative to induced values, as do unit 1 bids in the uniform-price sealed-bid auctions, consistent with the results typically reported for single-unit second-price auctions. All of these results hold up both with simulated (computer) bidders and with all human bidders.

3.2 More on Multi-Unit Demand Vickrey Auctions

The 1995 survey summarized research showing that sincere bidding emerges quickly for most bidders in single-unit English clock auctions in contrast to the persistent overbidding in second-price sealed bid auctions. In auctions where bidders demand a single unit, the English clock auction and second-price sealed bid auction are strategically equivalent, with both yielding sincere bidding in weakly dominated strategies. In multi-unit demand auctions where bidders have weakly diminishing marginal valuations, the static Vickrey auction and the dynamic Ausubel auction with drop-out prices reported are no longer strategically equivalent. Rather the static Vickrey auction generates sincere bidding in weakly dominated strategies whereas the Ausubel auction generates sincere bidding through iterated deletion of dominated strategies, a weaker solution concept.[59] Nevertheless, the research summarized in this section shows that the Ausubel auction with drop-out information generates outcomes much closer to sincere bidding than either the static Vickrey auction or an Ausubel auction with no drop-out information provided. (The static Vickrey auction and the Ausubel auction without drop-out information are strategically equivalent.) While this may not be surprising from a behavioral perspective, it is surprising from a mechanism design perspective, which typically calls for employing

[59] While the first solution requires agents' rationality alone the later must add the requirement of common-knowledge of rationality, a far from trivial addition. Thus, the first solution concept is much more robust and desirable from a mechanism design standpoint.

a stronger rather than a weaker solution concept. This suggests a possible tradeoff between the simplicity and transparency of a mechanism and the strength of its solution concept when agents are not fully rational or are still learning. This has important implications for the mechanism design literature. In addition to summarizing these results, we report results from three studies that have looked at generalized versions of the Vickrey auction suitable for dealing with synergies between items.

Kagel, Kinross and Levin (2001; KKL) investigate different versions of the Vickrey auction in which a human bidder with flat demand for two units competes against computer rivals each demanding a single unit. They compared outcomes in a sealed-bid Vickrey auction with bidding in the dynamic Vickrey/Ausubel auction with drop-out information provided (hereafter referred to as the Ausubel auction). As anticipated, in the static Vickrey auction there was a high frequency of bidding above value for both units, with bidding above value on unit 1 more severe than unit 2 (see Figure 11). Comparing these results with those for the Ausubel auction (recall Figure 8 in section 3.1), it's clear that the Ausubel auction comes closer to sincere bidding than the Vickrey auction, resulting in significant improvements in efficiency (but lower revenue) than the Vickrey auctions than in the Ausubel auction.[60]

[Insert Figures 11 and 12 here]

In an effort to better identify the basis for the superior performance of the Ausubel auction, KKL compare bidding to an Ausubel auction without feedback (referred to as the Ausubel* auction). In this case prices increase continuously until all bidders have dropped out, or the clock price reaches the maximum valuation, with winners and prices only announced after the auction ends. Here too there is significantly less sincere bidding than in the Ausubel auction, but largely as a result of substantially more bids *below* value (see Figure 12). This suggests that framing winning and payouts in terms of clinching, in conjunction with the dynamic auction format, are largely responsible for eliminating the overbidding in the Ausubel auction, with the drop-out information providing information about the number of competitors left, encouraging bidders to remain active until the point that prices reach their value.

[60] Average efficiencies were 97.5% and 97.9% in the sealed-bid auctions with 3 and 5 computer rivals, compared to 99.1% and 99.3% in the Ausubel auctions. Although these differences are small, they are statistically significant in both cases.

One weakness with the KKL experiment is that with computers bidding their value, it only takes a single round of deletion of dominated strategies by the human bidder to achieve sincere bidding. With all human bidders it requires several more rounds of deletion of dominated strategies to arrive at sincere bidding. As such it is quite natural to ask whether the results with computerized rivals will extend to auction environment with all human bidders. Kagel and Levin (in press) address this question, looking at all three auction formats with four (human) bidders, each demanding two units and with supply of 2 or 3 units. Bidders' valuations were iid from a common uniform distribution. The results essentially replicate those reported in KKL as (i) there is substantially more sincere bidding in the Ausubel auction than in the other two auction formats, (ii) there is massive overbidding relative to valuations in the sealed-bid Vickrey auctions resulting in a relatively high frequency of negative profits conditional on winning and (iii) deviations from sincere bidding in the Ausubel* auction primarily consisting of bids below value, at least to begin with. Finally, note that it is the ascending prices in the Ausubel auction in conjunction with the provision of dropout information that underlies *both* the greater transparency of the auction rules and the weakening of the solution concept.[61]

Kagel, Pevnitskaya, and Ye (2007) compare the Ausubel auction to the strategically equivalent *survivor auction*. In the survivor auction, sealed bids are submitted in each round with the lowest bid announced, with bids on that unit no longer permitted and all subsequent bids required to be the same or higher than the low bid in the previous round. Winning items are announced and priced following the Vickrey rules as the auction proceeds. In spite of the similarity in structure and information feedback between the two auction formats, the Ausubel auction achieves significantly higher levels of sincere bidding and efficiency than the survivor auctions to begin with. It is only with experience that the survivor auctions come close to the performance of the Ausubel auction.[62] These results extend the breakdown between the theoretically isomorphic English and second-price auction formats reported in the case of single-unit demands to

[61] The drop-out information enriches the strategy space relative to the Vickrey or Ausubel* auctions by allowing bidders to have strategies that are contingent upon other players' drop-out prices. This enrichment also opens the door to different kinds of "misbehavior."

[62] These results are obtained in an experiment with all human bidders under the same design structure as in KL (in press).

the multi-unit case. In addition, they confirm the importance of feedback in conjunction with the repeated nature of bidders addressing the question of "am I in or out" in the clock auctions (as opposed to the difficulty of computing sensible bids in the static case) as responsible for the rapid emergence of sincere bidding in the Ausubel auctions.

Manelli, Sefton, and Wilner (2006) also compare the static Vickrey auction with the Ausubel auction for the private values case, reporting overbidding in the Vickrey auctions and closer to sincere bidding in the Ausubel auctions.[63] Englemann and Grimm (2004) also compare the sealed-bid Vickrey auction with the Ausubel auction, reporting little difference between the two auction formats, including rather limited bidding above value on unit 1 bids in the sealed-bid auctions. The latter is at odds with results reported for the other multi-unit sealed-bid Vickrey auctions reported in this section, as well as Porter and Vragov's (2006) results reported in the previous section. The best explanation we have to offer for this difference at this time is sampling variability as, for example, KL (2001) report one sealed-bid uniform price session with three computer rivals where there was very limited bidding above unit 1 values, even though this was far from the norm for the other sealed-bid sessions.[64]

In concluding this section, we briefly review results from three studies that have looked at generalized versions of the Vickrey auction for dealing with complementarities between items, the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961, Clarke, 1971, and Groves, 1973). These are package auctions that permit XOR bids, with bidders bidding for as many packages as they wish, but only winning on one of their bids; e.g., in the simple case of two items, A and B, with values V_A, V_B and V_{AB} (where V_{AB} is the value of getting *both* A and B with $V_{AB} > (V_A + V_B)$) agents are permitted to bid for A alone, for B alone, and for the package containing both A and B, but can win only *one* of the packages bid on. The VCG mechanism is designed to produce sincere bidding and maximum efficiency, using suitably generalized Vickrey pricing rules to allocate items. There are a number of technical issues associated with implementing the VCG mechanism, as well as potential tradeoffs between efficiency and seller revenue that are

[63] The primary focus of this paper is on comparing the advantages of the Ausubel auction to the static Vickrey auction in the presence of a significant *common-value* component.
[64] Note, this is not a session level effect since in KL subjects were bidding against computerized rivals and only saw results from their own auction.

of concern in using it, discussion of which goes well beyond the scope of the present review (see Ausubel and Milgrom 2006). Rather, our primary interest is to report the results of experiments applying the static (sealed-bid) version of the VCG mechanism. In short, these show (i) significant deviations from sincere bidding in the form of bidding *below* package values and (ii) failure to bid on all items, a necessary requirement for achieving the efficient outcome.

The experiment with the most complicated demand structure investigating the static VCG mechanism is Chen and Takeuchi (2005) (CT). In each auction bidders compete for four items, resulting in a total of 15 possible packages to bid on. Human subjects compete against two computer bidders who bid sincerely in one treatment and randomly in another, under two different information conditions - with and without information on how the computers were bidding. Sincere bidding is a weakly dominant strategy regardless of what the computers do or the information provided about their bidding strategy. Subjects participated in 10 auctions under each treatment condition.[65] The auction interface automatically computed the value of each of the 15 possible packages so that bidding on all packages was relatively easy.

In equilibrium (optimal bidding in this case) subjects must bid their value on *all* 15 packages. Subjects consistently failed to do so, with the average frequency of bidding on possible packages going from a low of 65%-66% for single item packages to a high of 83%-86% for combinations of items (83% for the package with all four items). This confirms one of the potential concerns with the VCG mechanism (and package bidding mechanisms in general), the complexity associated with formulating bids for all possible packages of interest. When bidding on items, subjects generally underbid rather than overbid, with 57% of subjects classified as under-bidders, 32% as sincere bidders, and 12% over-bidders. Losing bidders were significantly more likely to increase the number of packages they bid on as well as their bid to value ratio in the next auction, with winning bidders decreasing their bid to value ratio (albeit, to a smaller degree than losing bidders). These bid changes indicate that the dominant strategy is not transparent, with subjects adjusting their behavior according to a trial and error learning process.

[65] CT also compare the static VCG mechanism to an ascending price (iBEA) mechanism.

Isaac and James (2000b) studied a VCG mechanism with two items and synergies with similar results in that close to 50% of all bids were within 25 cents of the true value, and with underbidding being more prominent than overbidding. Morgan (2002) in an auction with three items and synergies reports sincere bidding 39% of the time, with underbidding substantially more prevalent than overbidding. Both sets of results are in line with CT; bidding below value in static multi-unit demand auctions with synergies employing the VCG mechanism. This is in contrast to bidding above value in both single-unit second-price auctions and multi-unit Vickrey auctions with homogenous goods. Exactly why this should be the case remains an open question.

Remark: The studies reported here have been more concerned with behavioral issues than with mechanism design issues. In a mechanism design context, it is perfectly reasonable for the instructions and working examples to point out the benefits of different bidding strategies (e. g., sincere bidding and bidding on all items in the case of the VCG mechanism) in the instructions describing how the mechanism works, and what it's supposed to do, and should be considered part of the mechanism. Whether or not this would completely clear up the problems with the static Vickrey mechanisms identified here is another matter as there is evidence to the contrary; e.g., KL report bidding above value on unit 1 in uniform-price sealed-bid auctions even with instructions intended to dissuade subjects from doing so. (Also see the limited bidding on multiple packages reported in Kagel, Lien, and Milgrom (2008) reported in section 3.3 below.) Thus, there is still scope for identifying mechanisms that are more in line with subjects' natural tendencies that also achieve (or come closer to achieving) a desired outcome.

3.3 Auctions with Synergies

Most of the work in this area has been concerned with mechanism design issues, particularly with respect to the issues raised in the FCC spectrum auctions. Here we cover several small-scale experiments concerned with underlying behavioral issues in the presence of synergies.

Kagel and Levin (2005) look at a simple model of auctions with synergies comparing uniform-price clock auctions to sealed-bid auctions. Their experimental design is similar to KL (2001) for the case of substitute goods - humans demanding two units of a commodity compete against computer rivals each demanding a single (and bidding their value – the dominant strategy for single-unit buyers). The standalone values for *h*, the human bidder, are the same (v_h), but winning both units generates synergies equal $3v_h$. The uniform-price auction creates incentives for demand reduction for the "large" bidder similar to those discussed in section 3.1. However, there is an opposing force to bid aggressively in order to capture the synergy bonus. The net effect of these two competing forces is an equilibrium with the following properties: (1) at lower valuations, the demand reduction force dominates so that *h* bids zero on her second unit (drastic demand reduction), (2) at the highest valuations the synergy force dominates so that *h* "goes for it," bidding high enough to insure winning both items, and (3) at middle valuations the two forces are at peak tension, counterbalancing each resulting in *h* bidding above her stand-alone value (but short of "going for it") in the sealed-bid auctions and "going for it," conditional on rivals' observed drop-out prices, in the clock auctions. In both auctions *h* faces an *exposure* problem for these middle valuations; the possibility of winning a single unit at a price above its standalone value thereby earning negative profits. Depending on the size of the potential loss, and risk preferences, bidders may refrain from suitably aggressive bidding in order to avoid these potential losses, resulting in inefficient outcomes and relatively low revenue.

KL look at bidding in auctions with 3 and 5 compute rivals. Given the complexity of the auction environment, they employed four values for the human bidder designed to span the strategy space and to induce maximum differences in behavior between the sealed-bid and clock auctions, while providing bidders with considerable experience at each value and multiple observations against which to evaluate behavior.[66] The lowest v_h, \$3.00, calls for complete demand reduction in both sealed-bid and clock auctions. The highest $v_{h,}$ \$5.10, requires "going for it," and insures a *secure* (minimum) profit in each auction. The two middle values make different predictions between sealed-bid and clock auctions: With v_h = \$4.00, in the sealed-bid auction *h* should bid the same on both units at prices modestly above their standalone values (\$4.34 with $n = 3$ and \$4.16 with $n = 5$). The clock auction also requires bidding above value on both units, but with a cutoff value, **P***, such that if v_2 # **P*** = \$4.50 (where v_2 is the second-highest computer value), *h* goes for it, as winning both units has positive expected value greater than the value of stopping the auction at p = v_2 and winning a single unit; otherwise *h* drops out of bidding on both units at the cutoff point **P*.** At the other middle value, v_h =

[66] Since single-unit bidders have a dominant strategy independent of *h*'s valuation, KL could employ a limited number of values for *h* without distorting equilibrium predictions.

\$4.40, h "goes for it" ($b_1 = b_2 \exists \$7.50$) in the sealed-bid auctions, regardless of the number of computer rivals. In the clock auction h continues to employ the same strategy as with v_h = \$4.00, only now the cutoff value **P*** = \$5.70.

The experimental results show that bidding is substantially closer to optimal play in the clock auctions (see Table 5), consistent with the evidence from virtually all other auction environments. Further, in most cases bidders behave sensibly, though not optimally: The highest valuation, where optimal play is relatively transparent, generates by far the highest level of optimal play, comparable to the highest levels reported in any experimental auction environment. Demand functions estimated for the sealed-bid auctions are monotonically increasing in bidders' valuations. And in the clock auctions, there is a higher frequency of "going for it" at the two middle valuations when the optimal play calls for "going for it" than at the lowest valuation.

[Insert Table 5 here]

Nevertheless, there is much out-of-equilibrium play in both the sealed-bid and clock auctions, with the most interesting and dramatic differences for the two middle valuations where bidders are exposed to possible losses. In the clock auctions the primary deviation from optimal play consists of demand reduction as opposed to "going for it" as the theory predicts. In contrast, in the sealed-bid auctions bidders consistently bid above value on both units (often well above what they should bid and with different bids on each item). This suggests that the clock auction format with feedback on rivals' drop-out prices makes it much more transparent to bidders that they are liable to lose money as a consequence of bidding above value. This heightened awareness of the perils of bidding above ones' value helps to improve bidder profits and to move bidding closer to equilibrium in single-unit private value auctions and in multi-unit demand auctions without synergies. However, with synergies it holds bidders back from achieving maximum profit and generates deviations from the equilibrium outcome.

Katok and Roth (2004; KR) look at synergies between homogenous goods that might result from economies of scale in a production process or in transporting products to market. They compare a descending price (Dutch) auction with an ascending, uniform-price auction. Each auction has three bidders with supply of two homogenous units; one "big" bidder who has a high value for both items and two small bidders who each want one unit. The Dutch auction is, in effect, a package auction since the winner gets to choose how many units to purchase. Under the Dutch auction the small bidders face a threshold problem. This threshold/free riding problem (first identified in Ledyard, Porter, and Rangel, 1997) results from the fact that with each small bidder demanding a single item, items A and B respectively, they are competing with the big bidder who is bidding on the package AB. For the small bidders to win, the sum of what they are willing to pay for each individual item must be larger than the package bid, so that the small bidders must *coordinate* their bids to reach the *threshold* needed to beat the package bid. However, each small bidder has an incentive to let the other one be more aggressive, as this will aid them getting the item they are interested in but with larger profits. The big bidder faces an exposure problem in the uniform price auction as he may not get both items. There is no threshold/free riding problem in the uniform price auction for the small bidders since given the uniform price rule no small bidder can obtain a unit at a lower price than the other small bidder. Thus, there is no incentive to free ride. KR's treatments are somewhat complicated but reduce to (i) an environment intended to create an exposure problem for the uniform-price auction, (ii) a threshold environment in which there is no danger of an exposure problem since a big bidder who is outbid on one unit can be expected to be outbid on both units and (iii) a super free riding environment which magnifies the threshold problem in (ii).

For the ascending price auctions (the exposure environment) big bidders suffer losses twice as often as predicted in equilibrium (33.0% of the time versus 16.5%) earning a single unit at prices above the unit's stand-alone value. In the free riding and super free riding environments where in the uniform-price auction the big bidder should never suffer losses, they do 17% of the time. The average price the large bidder pays in the exposure environment is similar in the ascending and descending treatments, but the large bidder wins two units more often in the Dutch auction (53.5% of the time versus 28.5%) so that it achieves higher efficiency under the same demand structure than in the uniform-price auction. Further, when the large bidder wins only one unit in the Dutch auction, it never loses money as the large bidder is able to successfully compete with the second small bidder. KR conclude that the Dutch auction performs better than the uniform-price auction as it does better in the exposure environment as well as in the free riding environment, with the uniform-price auction only performing better in the super free riding environment.

Chernomaz and Levin (2007; CL) investigate bidding in a first-price sealed-bid multi-unit demand auction with and without package bidding. Despite the general preference for iterative auctions, first-price sealed-bid auctions have been used in practice (Cantillon and Pesendorfer 2007, Epstein et al. 2002), having a number of attractive features such as their resistance to collusive behavior. With strong complementarities present package bidding solves the exposure problem. However, when complementarities are not present (or are relatively small) package bidding may be "abused" to gain a strategic advantage, as bidders demanding multiple units have an incentive to place a bid for the package that is higher than the sum of what they would bid for each item alone (Cantillon and Pesendorfer, 2007). In addition, bidders for individual items (and/or smaller combinations of items) face a threshold problem, which is exacerbated in a sealed-bid auction as there is no opportunity (as in an iterative auction) for small bidders to determine if prices are such that they should raise their bids in order to win their preferred item. The threshold problem lowers bids of the local bidders, which in turn induces lower bids by the global bidder, thereby lowering revenue.

CL's experiment operationalizes this environment, with two local bidders each demanding a single (non-overlapping) item competing against a global bidder who demands both items. They employ a two-by-two experimental design, varying the auction rules (with and without package bidding) and the synergy level (0% and 50%). The two local bidders have the same iid value from a uniform distribution. The global bidder draws a single value from the same uniform distribution, with the value to the global bidder for obtaining both items $v_g = 2\beta s_g$ where β represents the synergy value and s_g is the global bidder's value. This highly structured environment (in conjunction with some additional restrictions on subjects' bids) permits solving for equilibrium outcomes without compromising the essential tradeoffs inherent in the underlying structure of the auction.[67]

[67] These other simplifying restrictions require the global bidder to place the same bid on both items (since they have the same value) when bidding in separate markets and to only place a bid on the package when package bidding is permitted. This last restriction is predicted in equilibrium behavior when local bidders are symmetric, as in their design.

Absent synergies, auctioning each item separately is predicted to achieve 100% allocative efficiency (the frequency with which items are allocated to the bidders who value them the most), compared to 90.7% efficiency with package bidding. Introducing the 50% synergy value without package bidding reduces allocative efficiency to 93.3%, compared to 93.5% if package bidding is permitted. Under the RNNE, the bid function for the global bidder lies above that of the local bidders in all cases except when there is no package bidding and no synergies, so that with package bidding the global bidder is more aggressive than the local bidders even when the synergy value is zero, and is naturally more aggressive, with or without package bidding with synergies present. Finally, revenue is predicted to be lower with package bidding with or without synergies present, but other things equal the 50% synergy case is predicted to raise more revenue.

As is typical of private value auctions, bids are above the RNNE reference point for both local and global bidders under all four treatments. The only case in which bids are even near the RNNE reference point is when global bidders are not permitted to submit package bids in the 50% synergy treatment. This is *not* a result of the exposure problem dampening their bids, as in equilibrium the global bidder does not bid above her stand alone values in the 50% synergy treatment. Rather, it more than likely results from the fact that the RNNE itself requires higher bids so that the predicted outcome is closer to the more aggressive bidding subjects typically engage in. In all four treatments bids of global bidders are above those of the local bidders, even when there are no synergies present in the no package bidding treatment when they should be the same. Permitting package bidding has the strongest effect on local bidders, inducing less aggressive bidding as the theory predicts, consistent with the threshold problem and the theory's prediction.

Regarding market level outcomes: Efficiency, both allocative and the more typical fraction of expected surplus that bidders capture, are systematically below the levels predicted under the RNNE. There are no significant differences in allocative efficiency between any of the treatments. However, efficiency as measured by the fraction of the surplus bidders capture, decreases significantly when going from 0% to 50% synergies in the absence of opportunities for package bidding. This reflects the increased asymmetry between global and local bidders resulting from synergies in the absence of package bidding. Changes in seller revenue are directionally consistent with the theory, with synergies bringing in higher revenue. But permitting package bidding has a substantial negative effect on revenue, raising the least revenue regardless of whether or not synergies are present. This, in conjunction with the negligible positive effect of package bids on efficiency when synergies are present, lead CL to sound a cautionary note efficacy of first-price sealed-bid auctions.

Summary: To date there have been very limited small-scale experimental studies focusing on multi-unit demand auctions with synergies. The results of the experiments reported here confirm the existence of an exposure problem in the presences of synergies which results in less aggressive bidding and less of the synergies being realized. The introduction of package bidding in CLs sealed-bid auctions introduces significant threshold problems for local bidders. Results from the few VCG package auction experiments reported at the end of the previous section suggest that this is not a viable alternative to dealing with multiunit demand auctions with synergies, as the frequency of sincere bidding is relatively low and subjects only bid on a small percentage of the packages available to them (even in auctions with very few items), both of which can severely compromise the promised efficiency gains. And the VCG mechanism can result in very low revenues as well.

Kagel, Lien and Milgrom (2008; KLM) report results comparing a combinatorial clock auction (CCA) mechanism which permits package bidding with a simultaneous ascending clock auction (SAA) mechanism. This experiment is covered in more detail in Chapter xx dealing with mechanism design issues. What is relevant here is that: (1) KLM identify a clear threshold problem in the CCA auctions, although the magnitude of the effect is relatively small and (2) similar to the results reported in CT for the VGM mechanism with package bidding, subjects bid on only a tiny fraction of the packages available to them even though they were explicitly encouraged to do so and had a computer interface that made placing bids very easy. What bids are made in the CCA are typically directed at the most profitable packages, which generates high revenue and efficiency as long as these packages constitute the "relevant" ones for achieving an efficient allocation. However, in cases where the efficient outcome requires that all bidders obtain one or more items, there are marked reductions in efficiency under the

CCA (and relative to the SAA) as prices fail to direct bidders to the relevant packages. In particular the package containing all items is almost always the most profitable package for the global bidder so prices fail to direct global bidder to bid on the package needed to achieve an efficient allocation. Brunner, Goeree, Holt and Ledyard (2007) report similar results in CCA auctions, with global bidders directing too much attention to large packages in cases where the efficient allocation requires them to obtain smaller ones. The failure of prices to direct global bidders to bid on smaller packages constituting their share of an efficient allocation, in conjunction with the limited number of packages they tend to bid on with XOR bids, is an important, yet unresolved, issue in package bidding.

3.4 Sequential Auctions

There have been a number of experiments looking at sequential auctions. Much of the work has been devoted to exploring the declining-price anomaly whereby prices of homogenous goods decrease systematically over the course of selling multiple items (Ashenfelter, 1989, Ashenfelter and Genesove, 1992). Declining prices are an anomaly because economic intuition suggests that prices of identical items sold in a sequence at the same time and place should be the same when each bidder demands a single item. Weber (1983) proves this to be the case for risk neutral bidders. Further, although intuition suggests that risk aversion may cause prices to decline, McAfee and Vincent (1993) demonstrate that this can only be guaranteed if buyers' have strictly increasing absolute risk aversion, a questionable assumption. As reported on below, in almost all cases, experimental work confirms the declining price anomaly and when it does not, it reports the "right" results but not for the same reasons the theory predicts. We also review several experiments in which bidders have value for more than one unit, which establishes an interesting set of new issues.

Keser and Olson (1996; KO) report the first sequential auction experiment with unit demands with paid subjects.[68] Each auction market consisted of eight bidders with known supply of four units bidding in a sequence of first-price sealed-bid auctions. Each bidder made a bid for the first unit, with the highest bidder receiving that unit at the price

[68] In an earlier paper, Burns (1985) compared bidding in a sequential auction between wool buyers and students in which both groups were motivated "by a desire to succeed in their chosen field." Both groups started out with declining average prices, with the students eventually converged on constant average prices. However, the wool buyers continued to have declining average prices. Burns attributed the latter to rules of thumb relevant to field settings but not the more austere conditions of her experimental markets.

bid. The winning bidder was no longer permitted to bid, with the auction continuing with new bids solicited for a second unit, with this process repeating for all four units. Prices of units sold were announced following the sale of each unit. Values were iid from a uniform distribution. Four sessions with 20 auction periods each were conducted with subjects not permitted to bid above their values.

The symmetric RNNE bid function for unit *l* with iid values from a uniform distribution with support [0, 1] is

$$b^l(v)=\frac{n-k}{n-l-1}(v),$$

where v is the bidder's valuation, n is the number of bidders and k the number of units sold, so that bids on later units are substantially closer to bidders' values than earlier units. However, expected prices remain constant as bidders with higher values get units earlier than those with lower values.

Table 6 reports realized and predicted prices for each of the four units. There is some variation in predicted prices as a consequence of the random draws used in the experiment. Both average and median realized prices decline for later units consistent with the declining price anomaly. Further, prices were higher than the RNNE for all four units, only more so for early units. Overall, average efficiency was 98% compared to close to 85% based on randomly allocating units among the six highest value bidders. KO classify individual bidders as risk averse if the number of bids above the RNNE benchmark is greater than the number of bids below it. By this measure, 84% and 81% of all subjects were risk averse with respect to unit 1 and 2 bids, with these percentages decreasing to 72% and 53% for bids on units 3 and 4, suggesting that bidding above the RNNE cannot be attributed to some universal and stable characteristic of bidders preferences.[69]

[Insert Table 6 here]

We are aware of two replications of the KO experiment. One is by Salmon and Wilson (2008) involving the sale of two units with four bidders for up to 20 periods using an English clock procedure. This was used as a control treatment for the second-chance

[69] KO report two other treatments designed to represent the impact of agents bidding on behalf of principals, with agents penalized for failure to obtain items.

offer auctions discussed in Section 4.3 below. They report average prices of 335 for the first unit sold versus 273 for the second, compared to the equilibrium prediction of 270 for both units. Thus, here to prices are decreasing.

Neugenbauer and Pezanis-Christou (2007; NPC) report a series of first-price sequential auctions with eight bidders and four units supplied. Values were iid from a uniform distribution with support [0, 100], with winning prices announced following each unit sold. One major difference from the KO experiment is the number of auctions in a session – 100 here as opposed to 20 in KO.[70] NPC measure efficiency in terms of the proportion of allocations to bidders whose value ranking was lower than the order in which units were offered – yielding an average aggregate efficiency of 72%.[71] Misallocations were greater for units 1 and 2 than 3 and 4, indicating that the highest value bidders tend to "wait and see" regarding sales of early units, giving bidders with lower values a chance to win these early units. Average prices were approximately constant across units in this experiment – ranging from a high of 51.7 on the first unit to a low of 49.5 on the fourth unit. However, there are still a number of important deviations from the theory as (i) average prices are significantly *above* the RNNE prediction for the first three units sold, with bidding above the RNNE more pronounced for low compared to high value bidders and (ii) there are systematic deviations from the predicted order in which units will be sold with lower valued bidders tending to buy early units. Finally, like KO average prices were decreasing across units in the first 20 auctions in NPC, the total number of auctions in KO's sessions, so in this respect NPC replicate the declining price anomaly.[72]

Robert and Montmarquette (1999; RM) extend this single-unit demand design to multiple units: each of eight bidders had positive demand for m^i units, where m^i was iid from a *Poisson* distribution with a maximum m^i of 15, with total supply of 15 units. Once m^i was determined, the value for each of the m^i units was iid from a *uniform* distribution

[70] NPC also have treatments with uncertain supply, where the RNNE is predicted to result in decreasing prices.
[71] It is not clear if or how NPC avoid double counting here in the sense that the highest value bidder having not won the first unit sold is still very likely to win one of the other units. In using an unusual efficiency measure of this sort it is helpful to also include more standard measures and to compare efficiency to some sort of random allocation process. However, as shown below, this efficiency measure is helpful in explaining behavior in their auctions.
[72] As noted, Burns (1985) reports a similar result for her (unpaid) student subjects.

on [0, 100] and ranked in decreasing order to form a downward sloping demand curve for each bidder. RM compared bidding in three types of sequential auctions: Dutch (descending-price), English (ascending-price) and mixed Dutch and English. A round of Dutch auctions was conducted as follows. The first unit was offered at the highest possible price of 100 with the price lowered by one ECU every two seconds until a unit was purchased. The second unit was then offered at an initial selling price of 5 ECUs above the winning price for the first item, with this process repeated until all 15 units were sold. Bidders knew when a unit was purchased and the price at which it was purchased. English auctions followed similar rules beginning with a starting price of 40. In the Dutch-English auctions the first unit was sold following Dutch auction rules, with the second unit sold using English auction rules with a starting price equal to the winning Dutch auction price, with this process repeated until all units were sold.

RM characterize the properties of a symmetric RNNE yielding an efficient allocation for each of the three auctions to serve as a reference point against which to evaluate behavior. Unfortunately, there is no assurance that the equilibrium identified is unique. However, their model does demonstrate that there are sufficiently rich strategies to induce an efficient allocation in these complicated, multi-unit sequential auctions, with the equilibrium outcomes generating the same expected revenue (assuming risk-neutral bidders) across the three auction formats. In addition, the model offers sharp predictions about bidding: In each auction, the winner is the individual with the highest (re-indexed) valuation for the unit supplied in that stage, with the price paid for each unit approximated by

$$B^{Dutch}(v, s) \approx \max[0, v - 2.04s] \text{ and } B^{English}(v, s) \approx \max[0, v - 2.04(s-1)],$$

where v and s denote the item valuation and the number of unsold units remaining, so that bidding should be more aggressive (relative to valuations) as fewer items remain to be sold.

Efficiency is measured in the usual way as the sum of the valuations of the 15 units allocated divided by the sum of the 15 highest valuations. Losses, relative to full efficiency, averaged 0.84%, 0.94% and 0.45% for the Dutch, English and mixed auctions respectively. This compares to average efficiency losses of 14.3% in each auction under completely random bidding and 2.3% for budget constrained random bidders (whose bids are restricted to be between 0 and their valuation), suggesting that bidders were at least in part following the equilibrium bidding strategy. All three auctions demonstrate decreasing average winning prices. Using simulations based on structurally estimated bid functions, they note that at the start of each auction the standard deviation associated with the distribution of winning bids is quite large, which initiates a bias pushing winning bids higher than predicted at the beginning of an auction sequence, forcing adjustments later on that are responsible for the declining prices. They suggest that this is the result of the complexity associated with bidding on early units with so many units available to bid on as opposed to risk aversion arguments or absentee bidders (which could explain declining prices in field settings; Ginsberg, 1998).

Brosig and Reiß (2007; BR) look at the effects of capacity constraints on bidding in sequential auctions.[73] They argue that although many real life auctions run independently of each other, from the point of view of bidders, they form sequences of auctions once capacity considerations are taken into account in procurement auctions, or credit constraints are accounted for in ascending price auctions.[74] BR's experimental work focuses on isolating the role of opportunity costs/option values on bidding.

They study an IPV auction with two bidders and two consecutive first-price sealed-bid procurement auctions, where bidders have the capacity to undertake only a single project. Bidders learn their own project completion costs, with both their costs and their competitor's costs randomly drawn from the same uniform distribution with support [20, 100]. Bids greater than 100 were not accepted. BR employ a 2 × 2 design, varying the nature of the opponent (human or computer rival) and information feedback (no feedback or feedback regarding winners and prices).

[Insert Figure 13 here]

[73] Pitchik and Schotter (1988) have an earlier paper on budget-constrained bidders in sequential auctions. Their subjects have full information about each others values and budget constraints as their experiment focuses on testing between different equilibrium refinements.

[74] BR refer to two empirical studies as providing support for their design: Jofre-Bonet and Pesendorfer (2000, 2003) found that firms that did not win a highway paving contract earlier in a sequence of auctions were more likely to enter a subsequent auction than firms that had already won a contract. De Silva et al. (2002) found that in auctions held by the Oklahoma Department of Transportation firms that lost in morning auctions bid more aggressively in the afternoon auctions compared to firms that had won in the morning. Also see DK who report that the overhead rate attached to bids by general contractors are positively related to the number of jobs already won.

Figure 13 shows the relative frequency of bids in auction A conditional on $C^A > C^B$ versus $C^A < C^B$ for the different treatments. In all four treatments, as predicted, the bid distribution for $C^A > C^B$ stochastically dominates $C^A < C^B$. Further, the theoretical benchmark for bids with $C^A > C^B$ predicts that all bids will cluster at 100 in auction A, with their data showing that 77.9% of all bids were exactly equal to 100 and an additional 12.6% in the interval [99,100), suggesting that subjects understood the implications of the opportunity costs associated with winning auction A. Further, pooling across treatments, 87.8% of all auction A bids were higher, less aggressive, than the RNNE bid in the single auction control treatment. Since actual bids were below the RNNE in the control treatment (which corresponds to bidding above the RNNE high price auctions), the data support the prediction of higher bids in auction A than in the control treatment. Entry decisions were also affected by the opportunity cost of early bids with 78.2% correct entry decisions as measured against a benchmark of "*always* enter the first auction no matter what your first or second period cost is.@ However, this high percentage may overstate the support for the theory as always entering yields 57% correct entry decisions regardless of a bidder's costs, and overlooks more sensible rules of thumb. When both bidders meet in auction B, they tend to bid the same as predicted under the RNNE in the single unit control treatment, indicating that bidders failed to correctly account for the selection effect associated with their rival choosing not to enter auction A. This, however, is hardly surprising since subjects tend to ignore far more obvious adverse selection effects in common value auctions and the takeover game. Finally, there were minimal differences in results when competing against a human or computerized rival and with and without feedback.

Leufkens, Peeters, and Vorsatz (2006; LPV) consider the impact of positive synergies between items when bidding in a sequential private value auction. There are two stochastically equivalent objects for sale using a second-price sealed-bid auction with values iid from a uniform distribution on [0,100], with the same four bidders participating in both auctions. Valuations for the second auction were unknown when bidding in the first auction, but winning the first auction increased the winner's value in the second auction by $s > 1$. LPV investigate three treatments: a baseline with no synergies, one with mild synergies (s = 1.5) and one with strong synergies (s = 2.0). Subjects participated in 50 rounds of two auctions each.

Their model predicts that with s > 1 all bidders will increase their bids on the first unit *above* their valuation. However, since bidding is symmetric in round one regardless of the presence of synergies, round-one efficiency should be unaffected. Nevertheless, LPV find that positive synergies reduce round-one efficiency: Efficiency rates were, 99.6%, 97.7% and 97.3% using the share of surplus measure for s = 1, 1.5 and 2 respectively, with the percentage of auctions won by the high value holder decreasing as well from 92.4% to 84.4%, to 77.8%. These efficiency reductions are not very surprising given the presence of the exposure problem, in conjunction with heterogeneity in bidder risk preferences, with s > 1. As predicted the larger the synergy factor, the higher bids are above value in the first auction, with average and median overbids of 4.23/8.12/12.16 and 0.00/4.30/7.0 in the baseline, mild and strong synergy conditions. As predicted synergies led to higher prices in the first auction, but price did not increase significantly between s = 1.5 and s = 2.0: Average and median prices in the first auction were 65.2/71.1/75.4 and 68.0/73.0/78.0 with s = 1, 1.5 and 2 respectively. They found no statistical support for the prediction that prices would decrease between rounds one and two in the presence of positive synergies.

LPV's experiment is notable for the introduction of synergies into a sequential auction framework where predicted outcomes could be solved analytically. The main weakness in their design is that there is huge uncertainty as bidders have no idea what their round two values are when bidding in round one, which creates an unrealistically severe exposure problem for bidders in stage one. This, in turn, may have a strong impact on actual behavior that is not captured by the assumption of risk neutral bidders.

Summary: Multiple unit sequential auctions where bidders have private values and single unit demands exhibit the decreasing price anomaly observed in field settings, at least for moderately long series of auctions. This is not all that surprising since in single-unit auctions subjects tend to bid above the RNNE, with some heterogeneity in the degree of overbidding across subjects, so that both factors are likely to be exaggerated with multiple units sold in sequential auctions. Observing decreasing prices in this restricted environment suggests that its presence in field settings cannot be solely attributed to supply or value uncertainty, the presence of buyers' agents in the bidding pool, or other factors that may contribute to the phenomena in less structured field settings. Thus, the results establish an important connection between laboratory and field settings.

What's missing with respect to this line of research are direct comparisons of single unit sequential auctions with, for example, simultaneous or uniform price auctions in terms of the relative impact on revenue and efficiency in order to get some idea of which auction mechanism is likely to perform best in field settings. The popularity of sequential auctions in business to business auctions makes this an interesting topic to study. The BR and LVP papers begin to scratch the surface with respect to more complicated issues in sequential auctions. This leaves a number of unexplored questions that remain to be investigated.[75]

IV Additional Topics

4. 1. Collusion in Auctions: An issue of enduring concern in auctions is the possibility of collusion. This is not just an intellectual/theoretical exercise as collusion in auctions has been identified in a number of cases: Krishna (2002) reports that in the 1980s 75% of the cartel cases in the United States involving collusion were related to auctions. Klemperer (2002) argues that the issues of primary importance in practical auction design have to do with discouraging collusion, entry deterrence and predatory behavior. Collusion is a difficult topic to study in the laboratory since it is almost impossible to effectively introduce side payments, and experimental sessions have a natural end point which is likely to induce some unraveling at the end of an experimental session. Research reviewed in the 1995 survey involved providing subjects with explicit opportunities (even encouragement) to discuss and coordinate bidding strategies. We take up work since then, much of which still focuses on opportunities for bidders to discuss collusive strategies with impunity, and nearly all of which involve auctions with multiple units for sale.[76]

[75] See Chen-Ritzo et al., (in press) and Engelbrecht-Wiggans, Haruvy and Katok (in press) for papers starting to address issues in these reverse auctions.

[76] See Hu, Offerman, and Onderstal (2007) for the sole single-unit auction study we have identified since the 1995 survey.

Sherstyuck (2002) looks at tacit collusion in ascending-price auctions. She gives tacit collusion its "best chance" by incorporating a number of factors facilitating collusion: small numbers of bidders (three bidders with supply of two units) with identical resale values and repeated play between the same set of bidders. She compares a weak bid improvement rule on fostering collusion (rivals need only match a bid in order to have standing as the high bidder with tied bids settled randomly) versus a strict bid improvement rule. Bidders private values are drawn from either a wide support (iid from [50, 100]) or a narrow support ([90, 100]). Bidder profits at the competitive equilibrium are substantially higher in the former than in the latter, so that the relative return for collusion is higher with narrow support. A known reserve price of 10 was in effect throughout but explicit discussions were prohibited.

Her results show that with a weak improvement rule and the narrower support, [90,100], prices are quite low, although it takes some time for them to converge to the reserve price (see Figure 14a).[77] Collusion largely results from bid matching. In contrast, with a strict bid improvement rule and the narrow support, prices are higher averaging 60% of the competitive equilibrium prediction (compare prices between the weak and strong bid improvement rules in Figure 14a and b). However, unlike the weak bid improvement rule, with strict bid improvements, collusion takes the form of a bid rotation rule (adopted in two out of the four sessions).

[Insert Figure 14 here]

With a weak bid improvement rule and private values drawn from the wider support, [50, 100], prices are close to the competitive equilibrium level throughout, and remain so under the strict bid improvement rule. The contrasting effects of the weak bid improvement rule as a function of the support for bidder values can be explained by the greater return from collusion in the case of the narrow support, as the potential payoff when bidders compete is much smaller than with the wider support. However, there is a strong subject population effect under the weak bid improvement rule as under both narrow and wide supports, prices in sessions with Cal Tech students, who were veteran experimental subjects, converge to the collusive outcome. (All of the other experimental

[77] For similar results but almost immediate convergence to the reserve price see Sherstyuck (1999). In these auctions bidders' values were common knowledge.

sessions used relatively inexperienced University of Melbourne students.) It is unfortunate that Sherstyuck did not conduct sessions to determine if the Cal Tech subject population effect holds up with a strong bid improvement rule, and did not report experience levels in detail for the Cal Tech sessions.

Goswami, Noe and Rebello (1996) (GNR) look at collusion in multi-unit share auctions designed to resemble Treasury bill auctions. They compare the effect of nonbinding pre-play communication between bidders in uniform-price versus discriminatory auctions. In each auction there were 100 units for sale, with a value of 20 for all bidders. There were 11 bidders in each auction, with bidders specifying the number of units they were willing to purchase at each of three possible prices: 10, 15 and 20, with each bidder able to bid for up to 100 units. In the class of symmetric, pure strategy, Nash equilibria for the uniform-price auction, there exists both a competitive equilibrium and a collusive equilibrium in which bidders extract all the surplus from the auction. (There are also a variety of other collusive equilibria with prices that are less than 20.) These collusive equilibria do, however, require a great deal of delicate bid coordination. For example, in the most collusive equilibrium each of the 11 bidders demands 9 units at a price of 20, with all other bids at 10. This profile of strategies results in each bidder getting 9 units at the lowest possible equilibrium price of 10 (with one additional unit assigned randomly). It is easy to see that any unilateral deviation to get a larger share results in raising the price to 20, significantly lowering profits. In the discriminatory auction there is a unique symmetric Nash equilibrium in undominated strategies with all bids at 15.[78] The authors predict that the uniform-price auctions will clear at price levels of 10, 15, or 20 with the discriminatory auctions clearing at 15.

All sessions had at least 12 auctions, with a single set of bidders in each session. In the communication treatments bidders were allowed to speak to each other inbetween every other round. Following each auction bidders were told the actual market clearing price, their own allocation and their own payoff.

There were essentially no differences in clearing prices between the discriminatory auctions with and without communication: none cleared at the lowest

[78] There is also a Nash equilibrium with all bids at 20, but since bidders earn zero profits with discriminatory pricing, bidding 20 is dominated.

price of 10, with 65% clearing at 15 without communication versus 69% with communication (with the remainder clearing at the price of 20). In contrast, in the uniform-price auctions with communication 36% cleared at 10 versus 0% without, another 30% clearing at 15 with communication versus 16% without. As such average prices were substantially lower in the uniform-price auctions with communication than without, with no real difference as a result of communication for the discriminatory auctions. Naïve collusive outcomes predominated; e.g., all bidders agreeing to place all bids at 10, as opposed to the rather elaborate self enforcing collusive Nash equilibrium.[79] The authors conclude that uniform-price share auctions are more susceptible to collusion than discriminatory share auctions. The driving force behind this result might come from the fact that even if bids do not converge to one of these delicate Nash equilibria, collusive (Nash) equilibria are supported in the discriminatory format.

Sade, Schnitzlein, and Zender (2005) (SSZ) conduct an experiment similar to GNRs but with different results, as average revenue is quite similar between the discriminatory and uniform price formats: 462.4 in the discriminatory auctions versus 477.6 in the uniform price auctions. Although this difference is statistically significant at the 1% level comparing average revenue auction by auction (average revenue is consistently lower in the discriminatory auctions), it is not significantly different in a regression using session level data ($p > .25$), and in either case the differences are not very meaningful economically. One key difference between SSZ and GNR is that in SSZ there were four possible prices of 17, 18, 19, and 20 versus three possible prices of 10, 15, and 20 in GNR. Thus, there were fewer alternatives to coordinate on in GNR and the potential profits from collusion were substantially higher (at least in nominal terms) in GNR, both of which would tend to promote collusion. These differences remain to be resolved.

One interesting sidelight of the SSZ experiment is their use of both students and finance industry professionals, with the professionals generating higher average revenue (under both mechanisms) than the students, even though they had the same opportunities to collude. Regulations precluded cash payments to the professionals, so they were

[79] Unfortunately there is no direct accounting for the number of auctions that actually achieved the self enforcing collusive equilibria. But the impression one gets is that none of them did.

rewarded with prizes bearing the logos of the sponsoring universities. As such "winning" might have been more salient for the professionals.

Phillips, Menkhaus, and Coatney (2003) (PMC) study collusion in a series of sequential English auctions designed to mimic livestock auctions.[80] Several facilitating practices were employed: the same set of bidders over a series of seven auctions, knowledge about the number of units for sale, and communication via an online chat program. They investigate auctions with six and two bidders and between 19 and 30 (homogenous) units for sale in any given auction. Bidders had identical negatively sloped demand curves, with a reservation price set 20 points below the average price had all units been sold to the highest value bidders at their induced values. Collusion increased with bidder experience so that we focus on bidding in the last auction in each session.

The six bidder control treatment yields average prices at 77% of a norm in which each unit is sold at its valuation going from the highest to lowest value. Communication with or without bidder identification reduced average prices to between 50% and 52% of this norm, so that average prices were 65% to 67% less than when bidders could not communicate. These lower prices were accomplished primarily through bid rotation rules that communication facilitated. Further, while there was some cheating over the last several units in each set of auctions, it did not destroy effective rotation in subsequent auctions and/or lead to substantial unraveling in the last auction (see Figure 15). Information about quantity for sale had no impact compared to the control treatment.

[Insert Figure 15 here]

The baseline treatment with two bidders had average prices at 75% of the norm in which each unit is sold at its valuation going from the highest to lowest value. Communication reduced average prices to 58% of this norm. Unlike the six bidder auctions, information about quantity for sale without any opportunity for communication had almost the same effect as communication with average prices at 61% of the norm. PMC, using the chat records for support, suggest that the somewhat smaller effect of communication in the two buyer auctions resulted from disputes as bidders compared

[80] In auctions with two buyers, individual bidder demands were augmented in order to keep aggregate demand constant compared to the six bidder treatment. See Menkhaus, Phillips, and Coatney (2003) for a related experiment.

their relative gains, whereas it was too difficult to go beyond a simple bid rotation rule in the six buyer case. Note that collusion might have been even more effective in this study had there been no reserve price in place.

Kwasnica and Sherstyuk (in press) (KS) look at collusion in simultaneous ascending multi-unit demand (SAA) auctions. The experiment is inspired by Brusco and Lompomo (2002) (BL) who show that there exist collusive Nash equilibria in SAA auctions whereby bidders start bidding on the item of primary interest to themselves and, if there are no competing bids, they stop bidding, with each bidder obtaining their highest valued item at a very low price. This equilibrium is supported by the threat of competition and higher prices if rivals do not cooperate.[81] Although this equilibrium does not require repeated interactions with the same set of bidders, KS look for it in a repeated play setting as (i) this adds collusive opportunities via bid rotation to the strategy set and is more relevant to many auction settings outside the lab and (ii) it is no doubt substantially more difficult to achieve BL style collusion in one-shot games.[82] Their experimental design involved two objects for sale, with complementarities between items in some sessions. Each session consisted of between 6 and 25 SAA auctions with the end point not announced, with either two or five bidders in each auction and no opportunity for discussions between bidders. They define collusion, in the case of no complements, when prices are below 50% of the competitive equilibrium norm. With compliments they consider collusion to be present when both items are awarded to the bidder with the highest value for the package at a price equal to the second highest valuation.[83]

KS's strongest results are in auctions with no compliments and two bidders where 10% of the auctions with inexperienced bidders, and 55% of the auctions with experienced bidders, are classified as collusive. In contrast, none of the five bidder auctions were collusive, regardless of bidder experience. KS identify a number of the

[81] See BL for a full characterization of the Nash equilibrium which also holds for strong synergies between items and when two or more bidders have higher values for the same unit. Cramton and Swartz (2002) provide evidence for BL type collusion in the FCC's auctions for spectrum licenses. EG (section 3.1) report attempts at collusion in their SAA auctions with repeated matching.

[82] On this last point KLM report two clear instances of such a collusive outcome their CCA auctions involving three bidders with random rematching in each auction. These auctions had a more complicated structure than those reported here. The facilitating practice in KLM consisted of announcing provisional winners and involved experienced bidders.

[83] This standard for complements is more problematic than without them since bidders face an exposure problem which in and of itself may prevent an efficient allocation.

auctions that follow the BL mechanism for tacit collusion. Collusion was reasonably frequent in markets with two bidders and modest complements, averaging 31% of all auctions with inexperienced subjects. But was much less common with larger complementarities: 0 out of 16 auctions with inexperienced subjects and 2 out of 11 auctions with experienced subjects. Collusion in these cases was achieved through bid rotation rules.[84]

Li and Plott (2005) (LP), and Brown, Kamp, and Plott (2007) (BKP), study collusion in multi-unit demand SAA auctions with eight bidders and eight items. Their strategy is to induce collusion by using an 'incubator" technology and then study factors capable of mitigating the collusion. Their incubator technology involves: (i) Bidders valuations being "aligned" and "folded" so that each pair of bidders has a unique item they value the most, with bidder i's second highest valued item very close to bidder j's highest valued item and vice versa (in this way it's easy for a bidder to retaliate should her closest rival compete for her highest valued item), (ii) There is complete information about *all* bidders' valuations, and (iii) The same set of bidders compete over several auctions with an unknown end point and there are no opportunities for discussion between rivals. Under these conditions there exists a collusive Nash equilibrium of the sort specified in BL, as well as a Nash equilibrium with competitive prices.

Once collusion is established, LP explore several remedies including (i) dropping bidder identification so that competitors can not be directly identified, (ii) removing information about rivals values, (iii) using a fixed end time for the auction as opposed to a soft ending (bidding continues until no new bids are entered for 30 seconds), (iv) removing several items for sale (thus increasing competitive pressure) and (v) changing bidders expectations by having some pairs of bidders with the highest value for the *same* item. These remedies were implemented sequentially so that remedies (iv) and (v) were only implemented on top of remedies (i)-(iii). Remedies (i) – (iii) had minimal impact. Treatment (v) reliably broke up collusion, with competitive outcomes continuing after the aligned and folded preference structure was reinstated (but with bidders not knowing this as valuations continued to remain private).

[84] Interestingly, BL style tacit collusion can be achieved in the case of large compliments but not with moderate compliments.

BKP introduce a simultaneous *descending* price auction (SDA), in which case there no longer exists a collusive Nash equilibrium within the incubator structure. BKP argue that the SDA will break up the collusive equilibrium since with all bids being final, rivals can no longer profitably retaliate against each other. The SDA largely achieves the desired outcome with prices averaging 611 experimental currency units (ECUs) versus an average price of just under 800 if all items were sold at their full value. BKP liken the SDA to a sealed-bid auction. In this respect the results are similar to those reported in Sherstyuck (1999), and are consistent with the notion that it is more difficult to collude in sealed-bid than ascending-price auctions.

Offerman and Potters (2006) (OP) look at the question of whether auctioning of entry licenses induces collusion in the product market. Standard economic arguments hold that entry fees constitute a sunk cost so that they will not affect pricing in the product market. However, many companies claim that they will have to charge higher prices in the product market in order to recoup entry fees. In addition, OP note that if entry rights are auctioned off, this will result in selecting bidders with the highest profit expectations in the product market, which might foster tacit collusion as this is one way of achieving these higher profits.

OP employ a product market with price-setting duopolists with differentiated products, with a unique stage-game Nash equilibrium in which each duopolist charges a price of 60 ECUs and earns a profit of 5000 in each period. This compares with the joint profit maximizing collusive outcome with both firms charging 150 and earning profits of 9000.[85] Subjects received feedback following each period about their own and their opponent's price, quantity, revenue, cost and profit but were not allowed to discuss strategies. There were three treatments: (1) an auction treatment in which four subjects bid for entry rights, with the two highest bidders paying their bids, (2) a fixed cost treatment where the entry rights were randomly assigned at an exogenously determined entry fee comparable to the average fees in the auction treatment, and (3) a baseline treatment in which the entry rights were assigned randomly with no entry fees. In all three treatments, subjects first played the duopoly game for 10 periods against the same opponent. After that each subject was randomly assigned to a group of four which

[85] Demand was simulated in the product market with price taking consumers.

included their rival(s) from the first 10 periods. These groups remained fixed until the end of the experiment (20 more periods) with entry licenses, valid for five periods, auctioned off at the start of each block of five periods. The two remaining subjects received a fixed payment of 1000 per period, compared to expected earnings of 5000 in the (competitive) stage game Nash equilibrium.

[Insert Figure 16 here]

Figure 16 shows average prices in periods 1-10, which were approximately the same, before the three different entry treatments were introduced. In the first two five period blocks with entry (periods 11-20) average prices were significantly higher in both the auction and fixed cost treatments compared to the baseline treatment ($p < .10$), but not significantly different from each other. These differences from the baseline treatment were much less pronounced in the last ten entry periods, and were no longer statistically significant ($p > .10$). Average winning bids were close to 20,000, the net expected profit from the stage game Nash equilibrium, so that auction winners earned excess profits. The role of entry fees in fostering collusion is supported by Spearman rank order correlations between entry fees and average prices, which were positive and statistically significant at the 10% level in periods 11-20 for the auction treatment ($p = .14$ for the fixed cost treatment) and significant at the 5% level or better for both treatments in periods 21-30. Finally, the data show that collusion is "clustered" so that some groups had prices close to the stage game Nash equilibrium while others set prices at higher levels. As such it would be more accurate to say that entry fees increased the *probability* of collusion than that they increased the *degree* of collusion.[86] Further, the similarity in outcomes between the auction treatment and the fixed cost treatment would support industry arguments that entry fees by themselves will lead to higher prices (via tacit collusion) in concentrated industries.

Summary: All of the auctions considered here involved the same set of subjects competing in a series of auctions, usually with an unannounced end point. Repetition with the same cohort appears to be a key facilitating factor, a factor likely to be at play in field settings as well. Communication between bidders reliably facilitates collusion,

[86] There is considerably more to this rich experiment than reported on here including a monopoly treatment in which monopoly rights are bid for or simply awarded.

which seems hardly surprising. But as Whinston (2006) notes there is little in formal economic theory about the way in which prohibitions on (nonbinding) price agreements prevent anticompetitive prices, with the published empirical work offering surprisingly little evidence that preventing oligopolists from talking has a substantial effect on the prices they charge. Sealed-bid, pay what you bid type, auctions are more collusion proof than ascending price auctions which provide easier opportunities to detect and punish non-cooperators. Competitive pressures seem to play a role as well as suggested by the role played in breaking up collusion in LP and the role played (in the form of the support from which values were drawn) in Sherstyuk (2002).

4.2. Bidder's Choice Auctions: Creating Competition Out of Thin Air: The National Association of Realtors defines a bidders' choice auction as:[87]

> "A method of sale whereby the successful high bidder wins the right to choose a property (or properties) from a grouping of similar or like-kind properties. After the high bidder's selection, the property is deleted from the group, and the second round of bidding commences, with the high bidder in round two choosing a property which is then deleted from the group and so on, until all properties are sold."

This type of bidding is very popular when selling time-shares, condominiums, and building lots. Goree, Plott and Wooders (2004) (GPW) were the first to study these auctions experimentally, noting that they can create competition between bidders who are interested in *different* items, which they illustrate with the following example: Consider the case of two bidders and two items, with each bidder interested in a different item. When a standard SAA is conducted the seller's revenue is zero when bidders prefer different items, which occurs with probability one half. In contrast, in the first stage of an ascending price right to choose (A-RTC) auction there is always competition since bidders, not knowing their rival's preferences, run the risk that the stage-one winner will take their preferred item. GPW show that if bidders are risk neutral, the A-RTC will raise the same average revenue as the SAA, but if they are risk averse, the A-RTC will raise more revenue, which may account for its popularity.

[87] See http://www.aaauctionservice.com/glossery_files/glossery.htm

In GPW there are four bidders in each auction, with two items for sale, A and B. Each bidder had a 50% chance that either item A or B (but not both) would be their "preferred item," with value iid from a uniform distribution with support [20, 920].[88] The value for their non-preferred item was effectively set to zero, so that each bidder had positive value for only one of the two items in any given auction (i.e., there were zero substitution possibilities between items). In the A-RTC, after the first item was sold, bidders observed the item chosen, with the remaining item sold in an ascending price auction. In the SAA, items were sold simultaneously through two ascending price auctions, with bidders restricted to bidding in only one of the two auctions at any given time. Observed revenues were 19.3% higher in the A-RTC than in the SAA, with 100% efficiency in the SAA versus 98.4% efficiency in the A-RTC. The estimated coefficient of relative risk aversion, 1-*r*, is 0.39, consistent with the higher revenues found in the A-RTC.

Eliaz, Offerman, and Schotter (2008) (EOS) studied an RTC auction with four items for sale, with two bidders each having a randomly drawn value for one of the goods, and zero value for the others (once again, zero substitution possibility between goods). Goods were sold using a series of second-price sealed-bid auctions, with the high bidder in each phase choosing her preferred item. Bidders who did not win their preferred item were not permitted to bid in subsequent phases. Bidders other than the winning bidder received no information other than that one of the items they had zero value for had been sold. The control treatment involved four separate good-by-good (GBG) second-price sealed-bid auctions, with each subject only permitted to bid for her preferred item. Two GBG treatments were employed, one with no minimum bid requirement and one with a revenue maximizing minimum bid requirement. With risk neutral bidders, expected revenue is the same between the RTC and GBG (with no minimum bid requirement), with risk aversion generating higher expected revenue in the RTC auction.

The RTC auctions raised significantly more revenue than either the unrestricted or optimal GBG formats (40.4% and 13.9% higher revenue respectively). Average

[88] There were no restrictions on the probabilities so that it was possible to have less than two bidders whose preferred item was A or B.

efficiency was comparable between the RTC and unrestricted GBG auctions (98.2% versus 98.3%) and higher than in the optimal GBG* auctions (87.9%). To test whether risk aversion was the source of the higher revenues in the RTC auctions, EOS employed a no information RTC (NIRTC) auction where, after each phase, bidders were not informed as to which item was sold and no (losing) bidder was eliminated from the auctions, which proceeded as usual. When a winning bidder does not find the good she values, she still pays the second-highest bid and chooses a good at random, a good for which she has zero value. As a result, following the first item sold, risk averse bidders are essentially bidding in a second-price auction in which the high bidder wins a lottery that awards that bidder her most preferred good with some probability and nothing with the complementary probability. Risk neutral bidders will bid the expected value but risk avers bidders will bid strictly less than their expected value for the item. They then compare bidding assuming homogeneous risk averse bidders in both the RTC and the NIRTC auctions to the data. The RTC auction with risk aversion provides a better fit to the data than with risk neutrality. However, the NIRTC auction with risk aversion provides a worse fit to the data than assuming risk neutrality. Thus, EOS reject risk aversion as the explanation for the RTC auctions raising more revenue than the GBG auctions. Rather, their preferred explanation is that subjects act as if they are facing fiercer competition than they actually face, with simulations showing that in *both* the NIRTC and full information RTC subjects act as if they are effectively competing in an auction with six bidders for their preferred item (as opposed to the two actually competing for their preferred item). Finally, note that the higher revenue and efficiency achieved in the RTC auctions compared to the optimal GBG auctions belies a recurring theme in the mechanism design literature that tradeoffs must be made between efficiency and maximizing seller revenue.

Remark: One side note here is that subjects in the GBG auctions basically bid their value in accordance with the theory, consistent with findings reported earlier that the dominant bidding strategy organizes data well in single-unit second-price sealed bid auctions with only two bidders (see section 1.5 above).

Salmon and Iachini (2007; SI) examine a “pooled” RTC auction, with a number of similarities to the NIRTC auctions in EOS. They conduct a sealed-bid auction with multiple units for sale with all bidders submitting a *single* bid at the same time. Bidders’ values are perfectly correlated across items, so that each bidder has the exact same ordinal ranking across items. (Think of selling several condominiums in a given building, each of which is ranked from highest to lowest based on its scenic view. But because of the location of the building relative to where bidders work, bidder *i* ranks each apartment uniformly higher than bidder *j*.) Thus, unlike the other RTC auctions reported on, there are some substitution possibilities between items, albeit with common ordinal preferences over the goods.

Bids are ranked from highest to lowest with the high bidder getting first choice, the second highest bidder second choice, and so on, with all winners paying what they bid. Following Menezes and Monteiro (1998), assuming symmetric bidding strategies, SI provide the first-order differential equation that defines the bid function, solving it numerically for both risk neutral and loss averse preferences.[89] Loss aversion is relevant here since bidders can lose money when bidding according to the RNNE, as a bid designed to get a higher valued unit may end up securing a unit with a lower value, but the bidder must still pay the price bid. They compare outcomes in the pooled RTC auctions to an SAA in which subjects are restricted to holding the high bid on one item at a time.[90]

Their results show that seller revenue is uniformly, and substantially, higher in the pooled RTC auctions than in the SAA (41.8% higher), well above the revenue predicted under the RNNE for the RTC auctions. In fact, bidders suffer persistent losses, with bidder profits well below those in the SAA. Revenue and prices in the SAA are very close to those predicted under the efficient allocation. Efficiency is essentially the same between the two auction formats, averaging around 95% in both cases. Looking at individual bids in the RTC auction, the shape is essentially the same as the theory predicts, but bid functions are displaced upward relative to where they would be under the RNNE. SI explore a number of alternative explanations for this upward displacement of the bid function, with their preferred explanation consisting of "attentional" bias

[89] The loss averse specification uses the utility function and parameter values reported in Tversky and Kahneman (1992). SI note that there is little difference between risk neutral and risk averse bidding given their parameter values.
[90] Their SAA follows the format employed in the FTC spectrum auctions, with a countdown clock that resets every time a new bid is submitted. The auction ends when no new bids are submitted for any items, with winning bidders paying what they bid.

whereby bidders focus most of their attention on winning the best 2 or 3 items in the auction, largely ignoring the possibility of being "stuck" with lower valued items.[91] Finally, SI note that assuming their results translate outside the laboratory, the kind of pooled auction format they employ would have trouble sustaining itself, as persistent losses would reduce incentives to bid in such auctions, as well as generating defaults on bids.

Summary: The three RTC experiments reported on provide strong evidence for their revenue raising ability compared to either an SAA or a GBG format. The results reported in EOS and SI lie totally outside what theory predicts. The losses associated with the pooled RTC auctions in SI would seem to limit their use in field settings. The total lack of substitutability between commodities in the other two experiments seems somewhat unrealistic for the situations these auctions are intended to represent. Thus, there is scope to explore either an ascending, or sequential, RTC auction in which bidders demand a single unit but the items are substitutable to some degree, and to compare outcomes in these auctions with either an SAA or GBG auction. Nevertheless, the results of these three experiments are an exciting new application of experimental methodology designed to better understand the basis for RTC auctions found in field settings.

4.3 Internet Auctions: Internet auctions provide new opportunities to conduct experiments to study old and new puzzles. Lucking-Reiley (1999; LR) used the Internet to sell collectable trading cards under the four standard auction formats (Dutch, English, first- and second-price sealed-bid auctions), investigating the revenue equivalence theorem. He finds that Dutch auctions produce 30% higher revenue than first-price auctions, a reversal of previous laboratory results, and that English and second-price auctions produce roughly equivalent revenue. These results are interesting but lack the controls present in more standard laboratory experiments; i.e, there may well be a common value element to the trading cards and Dutch auctions provide an opportunity to use the game cards immediately, which cannot be done until the fixed closing time in the first-price auctions.

eBay auctions have a fixed closing time with many bidders submitting bids just seconds before the closing time (sniping), while others increase their bids over time in

[91] EOS note that attentional bias can also explain bidding above the RNNE in their NIRTC auctions, but not in their RTC auctions, so do not pursue it.

response to higher bids. In contrast, Amazon auctions automatically extend the closing time in response to late bids (as "soft" closing), with much less last minute bidding than in comparable eBay auctions. These differences raise two questions addressed by Roth and Ockenfels (2002; RO): (1) Why the sharp differences in last minute bidding between the two auction designs and (2) since eBay has a number of characteristics similar to a standard second-price auction, why the increased bidding by the same bidder over time?

RO suggest several (rational) reasons for sniping in (essentially) private value eBay auctions with their fixed deadline: (i) implicit collusion on the part of snipers in an effort to get the item at rock bottom prices since congestion will result in some of the last minute bids not being recorded at the web site and/or (ii) a best response to incremental bidding on the part of less sophisticated bidders in an effort to avert a bidding war. They also note that motivation for sniping for items with a significant common value component could result from (i) better informed bidders' efforts to conceal their superior information on high valued collectables and/or (ii) bidders updating their valuation of items as bids come in. Because there are a number of other differences between eBay and Amazon than their ending rules, as well as the difficulty of clearly distinguishing between private value and common value auctions in field settings, Ariely, Ockenfels, and Roth (2005) go beyond the field observations reported in RO to conduct a laboratory experiment in which the only difference between auction institutions is the ending rule for private value goods – a dynamic eBay auction with either a .8 or 1.0 probability that a late bid will be accepted and an Amazon style auction with a .8 probability that a late bid will be accepted, in which case the auction is automatically extended. Their results show quite clearly that there is more late bidding in both eBay auctions compared to the Amazon auction. Further, there is significantly more late bidding in the eBay treatment where last minute bids would be recorded with probability 1 than with probability .8, which rules out the tacit collusion hypothesis for sniping and is consistent with the hypothesis that it represents best responding to incremental bidding on the part of less sophisticated bidders.

Salmon and Wilson (2008) investigate the practice of second-chance offers to non-winning bidders in Internet auctions when selling multiple (identical) items. They compare a two-stage game with a second-price auction followed by an ultimatum game between the seller and the second-highest bidder versus selling the two items in a sequential English auction. As predicted the auction-ultimatum game mechanism generates more revenue than the sequential English auction, providing a potential explanation for the practice of second-chance offers to losing bidders.

Shahria and Wooders (2007; SW) study the practice of a "buy-now" option popular in eBay and Yahoo and other Internet auctions.[92] For a private value auction when bidders are risk averse a suitably chosen buy-now price will raise revenue as it extracts a risk premium from bidders wishing to avoid the uncertainty over whether they will win and the price paid in a first-price sealed-bid auction (Reynolds and Wooders, in press).[93] In contrast, for common value auctions, if bidders are sophisticated and do not suffer from a winner's curse, there is no buy-now price that raises revenue for risk neutral or risk averse bidders.

SW's results support the risk aversion predictions for the private values case as first-price sealed-bid auctions with a buy-now price raised average revenue by 6.8% compared to the control treatment with no buy-now price, and by 11.9% conditional on the buy-now price being accepted ($p < .01$ in both cases), with buy-now prices accepted in 45% of the auctions. Introducing a buy-now price that is a little above the (unconditional) expected value of the item in an ascending-price common value auction raises revenue by 4.2% ($p > .10$), but consistent with a winner's curse is accepted in 78.9% of the auctions. Further, bidders tend to drop out earlier when the buy-now price was not accepted compared to controls with no buy-now price, even though rejection of the buy-now price is completely uninformative. SW explain these anomalous results through an extension of the naïve bidding model developed in KL (1986) in which bidders make no adjustment to the adverse selection effect conditional on winning the item and fall prey to the winner's curse.

4.4 Auctions with Entry: Most of the theoretical literature on auctions assumes that the number of bidders, N, is fixed. The fixed N paradigm simplifies the analysis and allows for easy comparisons of revenue and efficiency between different auctions mechanisms, and is an essential assumption underlying the revenue equivalence theorem. The key motivation for looking at endogenous entry is that it's both costly and time consuming to

[92] In a buy-now auction bidders have a chance to get the item at a fixed price before any bids are placed.
[93] There are other motivations for the buy-now option not captured in the experiment. Matthews (2003) shows that in eBay auctions a buy now price increases revenue when bidders are impatient. Matthews and Katzaman (2006) show that risk averse sellers benefit as it reduces the variance in revenue.

prepare bids so that it is part and parcel of the auction process. As such it should not be swept under the rug by assuming an exogenously determined number of entrants. Further, casual observation shows that at times the number of actual bidders in similar situations varies a lot, leaving the impression that it is governed by a stochastic, rather than deterministic, process. Thus, a natural question, both theoretically and experimentally, is how sensitive are the typical auction results to dropping the fixed N assumption, extending the analysis to allow for endogenous entry.[94]

There have been two main approaches to modeling auctions with entry. Both start by assuming N potential entrants and an entry cost, c, (e.g. bid preparation costs) since otherwise all potential bidders enter and we are back in the fixed N setup. The first approach assumes that *ex-ante* bidders are symmetrically informed so that any information bidders' have before they enter the auctions is public information, so that they find out their private information signals only after incurring the entry cost. In this case the theory has focused on two types of equilibria: A *deterministic, asymmetric equilibrium* in which bidders use pure entry strategies with exactly n* bidders, the number of bidders that can enter profitably, enter the auction. The remaining (N– n*) bidders remain out and have no further impact on the auction (see Smith, 1982, 1984, Engelbrecht-Wiggans, 1987, 1993, and McAfee and McMillan, 1987). The second model Levin and Smith (1994; LS) has a unique *symmetric mixed strategy equilibrium* that determines a probability of entry, q*, which leaves all bidders just indifferent between entering and staying out.[95] This results in entry being a random variable that is governed by a binomial distribution with N and q* as the two parameters, with q* depending on the expected rewards from entry relative to its cost, q* = Q*(c, N).[96] We refer to the first equilibrium as ''deterministic" and to the second as ''stochastic." A second modeling approach is to assume that the N potential bidders obtain their private information regarding their type *before* they decide to enter. This approach generates a unique pure strategy equilibrium characterized by a cutoff value for a bidders' type, which determines

---

[94] Endoginizing entry decisions also forces one to take account of bidders' preferences over auctions (Mathews, 1987, McAfee and McMillan, 1987).

[95] There are an enormous number of asymmetric equilibria involving subsets of bidders that enter, or stay out, deterministically while the rest enter with the same probability.

[96] This approach also allows having symmetric risk-averse bidders. For example if $u(x) = x^{\rho}$ where ρ being the CRRA parameter than in equilibrium $q^* = Q^*(c, N| \rho)$.

who will enter and who will stay out (see, for example, Palfrey and Pevnitskaya, in press). Here, the realized number of entrants is a random variable governed by a binomial distribution with N and q*, where q* in this case represents he probability of a player's type exceeding the cutoff level.

Smith and Levin (2001; SL) conduct an experiment to examine whether their stochastic bidding model predicts better than the deterministic model. The experiment focuses on entry so that payoffs for entry matched the expected gains a bidder would get in the symmetric RNNE after entering (i.e., there is no bidding phase to the experiment after entry). This provides a rich set of comparative static predictions to use in discriminating between the stochastic and deterministic models.[97] Each experimental session consisted of a series of market periods with subjects electing whether to enter the ''market,'' with those electing to stay out paid a fixed sum of money designed to represent the opportunity cost of entry. Before each period, the number of potential entrants, N, the cost of entering, c, and the schedule of payoffs conditional on entering (which were deceasing in *n*) were publicly announced. After all entry decisions were made the total number of actual entrants was announced (without any information about bidder identities). There were two main treatments, one with a small number of potential entrants (N = 4) and one with a larger number of potential entrants (N = 8), with four different costs of entry within each treatment.[98] Payoffs were such that at each cost level there was room for "profitable" entry by at least one bidder but not more than three bidders.

The aggregate data strongly support the stochastic model. First, the deterministic model predicts that the number of bidders actually entering the auction is, other things equal, independent of the potential number of entrants (N) as opposed to the stochastic model's prediction that average entry will increase with larger N. For each of the four cost levels employed, this prediction is satisfied, with these increases significantly

[97] There is a large, closely related, earlier experimental literature on coordination games (see Ochs, 1995, and Rapoport et al, 1998, and references therein). The key difference between the Smith and Levin experiment and these earlier ones is linking the payoff structure to what would have been earned in the RNNE of a well defined auction market. These predictions are sensitive to bidders' risk preferences. However, the apparent risk premium demanded by subjects was close to zero, so that predictions for the risk-neutral case are used throughout.

[98] Entry costs of {$0.50, $1.00, $1.50, $4.00}.

different from zero for three of the four cost levels. Further, although the average increase in entry rates was somewhat greater than predicted in the stochastic model, none were significantly greater than predicted. Both models predict that the average number of entrants will fall as entry costs increase, but by varying degrees. At each of the cost levels with sufficient numbers of observations to perform exact tests, the observed reduction in entrants is significantly smaller than the deterministic prediction. In contrast, although there is a tendency for entry rates to decrease more than predicted in the stochastic model, in no case were the differences large enough to reject the stochastic model's prediction. The stochastic model predicts that subjects will enter with sufficient frequency to reduce expected profits to zero, since in equilibrium bidders are indifferent between entering and staying out. Profits for entrants averaged -\$0.02 per subject, per period, over all auctions, very close to the zero-profit prediction and substantially below the \$0.34 profit level of the deterministic model.[99] Finally, the stochastic model predicts that the total surplus generated in the auction will *increase* when the number of potential entrants *decreases*, whereas the deterministic model predicts no change. Simulating seller revenue and adding it to bidders' actual profits to compute total surplus, for all cost levels reductions in N caused social surplus to increase. This provides strong empirical support for one of the most intriguing policy implications of the stochastic model: Other things equal thicker markets are less efficient due to increased costs of entry, so that society may benefit from measures designed to limit the number of potential bidders.

While the preceding shows that the stochastic model organizes the aggregate data rather well, and substantially better than the deterministic model, there were also some significant deviations from the stochastic model at the individual subject level. The stochastic model assumes that bidders are symmetric, which implies that for each treatment they all employ the same (symmetric) entry probability. The data soundly reject this. Among inexperienced subjects, this hypothesis is rejected (at the 5% level) 26% of the time versus an expected rejection rate of 5%, and for experienced subjects it is rejected 33% of the time. Thus, it does not appear that all subjects rely on the common entry probability, q*, underlying the stochastic model. The failure to find a uniform

[99] These are "pure economic profits" above and beyond the return for staying out.

probability of entry across all subjects invites further research to identify a more accurate, stochastic *asymmetric* entry model (also see Ochs, 1995).

The main weakness of the LS model is that its symmetric equilibrium uses a mixed strategy. When the N potential bidders are risk-neutral, or symmetrically risk averse, a mixed strategy equilibrium is unavoidable. Palfrey and Pevnitskaya (2008; PP) purify the mixed strategy equilibrium by assuming that the number of potential entrants, N, are drawn from a population with heterogeneous (homegrown) risk preferences. As such there is a critical level of risk-aversion (the cut off level) for which bidders who are more risk-averse than the cutoff level select to stay out in order to avoid the costs associated with entry, while the less risk averse enter and bid.

PP explore an environment with either 4 or 6 potential bidders in an IPV first-price sealed bid auction. Sessions with no entry costs and fixed numbers of bidders served as the control treatment against which to evaluate bidding in the auctions with endogenous entry. Entry costs, ω, were represented by a fixed payoff for staying out. Varying N, the upper bound of the support from which valuations were drawn, $\bar{u}$, and the value of the outside option, they employed treatments where the RNNE entry probability, q^*, was either .5 or .35, representing "high" and "low" anticipated entry rates. Bidders' types, needed to purify the mixed strategy equilibrium, are determined by their risk preferences, with bidders who are more tolerant of risk entering the auction, after which they learn their value for the item.

Comparing auctions with endogenous versus exogenous entry, as predicted the estimated slopes of the bid functions are smaller with endogenous entry in 11 out of 12 cases, consistent with the prediction that with endogenous entry the more risk averse subjects choose to stay out of the auction. Further, comparing estimated slopes of bid functions for different realized values of n, slopes are larger with $q^* = .5$ than .35 in all cases, again consistent with the prediction that with higher entry rates more risk averse bidders enter the auctions, resulting in more aggressive bidding. PP conclude that subjects who enter the auction are, on average, less risk averse than those who stay out. However, actual entry rates are consistently higher than predicted under risk neutrality – averaging .61 with $q^* = .50$ and .45 with $q^* = .35$. In contrast, their model predicts that entry rates will be lower than the risk neutral model's prediction, as they assume a population of heterogeneous risk averse bidders (which is necessary to reconcile their model with the usual result that nearly all subjects bid above the RNNE in first-price auctions with exogenous entry).

Given this excess entry, average profits from entering were substantially and consistently lower than the outside option (approximately 50% less). One potential source of this excess entry is that sitting out is boring, with entering providing some entertainment value. To test this, PP employ a treatment in which after selecting to sit out subjects have the opportunity to play a simple computer game. While the average entry rate declined significantly from .61 to .54 (with q* = .50), it was still significantly above the predicted upper bound for entry. One important difference between this experiment and SL who got close to risk neutral bidding is that in SL subjects did not bid following "entry," but were instead given payoff information conditional on the number of bidders actually entering. Hence, potential entrants would have been better informed about the expected value of entering versus staying out in SL, with the natural variation in earnings conditional on entering inherent in a real action making it even more difficult to effectively asses the expected value of entering. The importance of clear information concerning expected profits will be shown in the next experiment reported on here.

Finally, the most direct test of the cut-off model of entry would be to examine individual subject behavior: Do the same, least risk-averse, subjects almost always enter under the same treatment conditions? Unfortunately, the authors do not provide this data. However, individual subject data from SL shows strong variability in entry rates for individual bidders, although not as much as predicted under the mixed strategy equilibrium. One side note here: The PP data is consistent with earlier results reported (HKL; section 1.6 above) that when bidders have an alternative source of income, so that winning the auction is not their sole source of earnings, bids in private value auctions are consistently closer to the RNNE than absent outside earning opportunities. There are a number of possible explanations for this.

Ivanova-Stenzel and Salmon (2007) (ISS) compare bidding in a first-price sealed-bid auction with an English clock auction when bidders have a choice as to which auction to enter. That is, they take endogenous entry to its logical conclusion by having bidders choose which of two different auction formats to bid in. In this case cost of entry cost is the opportunity cost of participating in the other auction. The key question posed in this experiment is whether revenue equivalence can be restored through competition between auction formats.

Bidders first participated in a "learning phase," where they bid in both an English clock auction and a first-price sealed-bid auction with exogenously determined numbers of bidders, in order to provide bidders with experience with both auction formats. Further, to insure that bidders knew the likely payoffs from the two formats, at the end of the learning phase subjects received feedback regarding the session-wide average profit for each format for all values of n. In the second phase subjects split into two groups of six bidders each and proceeded to bid in 30 rounds of auctions, choosing which auction to participate in.[100] The main result of the experiment is that the English clock auctions attracted more bidders than the sealed-bid auctions, with the net effect that average revenue was essentially the same in both formats, as was average efficiency. However, winning bidders earned slightly more money on average in the English auctions.

ISS conclude that the key result in their study is not the approximate revenue equivalence, but rather that revenue in the English auctions increases sufficiently with endogenous entry to call into question the assumed revenue superiority of the sealed-bid auction, as bidders' arbitrage between mechanisms. This may well be true in a laboratory study where subjects are well informed regarding expected profits between the two mechanisms, as well as in field settings where bidders have extensive experience with both mechanisms. However, looking at the results of LR's field experiment where he investigated revenue equivalence between auction formats, absent these conditions this conclusion would appear to be premature as his first-price sealed-bid auctions consistently raised more revenue than e-Bay's "English" auction format: Actual revenue as a percentage of the list prices for cards averaged 80.8% and 107.8% in the first-price auctions versus 73.7% and 74.6% in the English auctions.

Two other entry related studies are worth mentioning. Goeree, Offerman and Sloof (2006) study multi-unit demand auctions when two incumbent firms face a potential negative externality in the form of a new entrant who, if winning items, will

[100] To assure competitiveness one bidder was assigned to each format without any choice so that each of the remaining four bidders could not enter and find herself the only bidder in that market. See ISS for a number of other important details regarding the innovative procedures employed.

compete in the resulting product market.[101] The paper compares the performance of a uniform-price clock auction with that of a sealed-bid discriminatory auction, both with and without externalities imposed by new entrants. Their results show that both auction formats induced similar high levels of entry. However, the basis for the high levels of entry reported differ between the two auctions: Entry is high in the ascending price auction because of demand reduction, while in the discriminatory auction entry is encouraged because of incumbents' failure to coordinate their bids to entry.

Kagel, Pevnitskaya, and Ye (2008; KPY) look at entry in markets with indicative bidding. Indicative bidding is a two-stage auction process sometimes used in the sale of business assets with very high values. In the first stage the auctioneer solicits a large group of interested buyers to submit non-binding bids, with the highest of these non-binding bids used to establish a short list of final (second-stage) bidders. These short-listed bidders then engage in extensive studies to acquire more information about the asset for sale, after which they submit firm and final bids (typically in a standard first-price sealed-bid auction). Ye (2004) shows that there does not exist a symmetric increasing equilibrium with indicative bidding so that the most qualified bidders may not be selected to be on the short-list to compete in stage two, which may result in substantial efficiency losses. In contrast, there are a number of alternative two-stage bidding procedures that, in theory at least, guarantee efficient bidding in the sense that the short-list consists of those bidders with the highest preliminary (first-stage) valuations, while preserving the best properties of indicative bidding; namely, avoidance of the costly (thorough) asset valuation process for all but the short-list of final stage bidders. Most prominent among these is a uniform-price first-stage auction in which first-stage bids are binding and establish an entry fee (the highest rejected first-stage bid) for those bidding in stage two (Ye, 2004). This is the type of situation tailor made for an experiment since there is no guarantee that the alternatives to indicative bidding will produce fully efficient outcomes, nor any other way to seriously evaluate the efficiency losses associated with indicative bidding.

KPY compare the uniform-price, two-stage bid process with indicative bidding. Their results show that indicative bidding performs as well as the uniform-price process

[101] The experiment was inspired by developments in spectrum auctions in both the US and Germany.

in terms of efficiency. This is a result of (1) sufficient heterogeneity in first-stage bids under the uniform-price process to destroy 100% entry efficiency and (2) first-stage bids under indicative bidding being highly correlated with first-stage value which results in fairly high entry efficiency. The latter is reflective of the fact that bidders with low first-stage valuations tending to lose money as a result of entry, while those with higher valuations earning consistent positive profits. Further, indicative bidding does better on other dimensions as it yields higher average profits and fewer bankruptcies in the initial auction periods due to systematic overbidding under uniform prices in stage one. Although these higher revenues are good for sellers in the short-run, they indicate the greater difficulty bidders have early on, with the more complex uniform-price two-stage process, which would more than likely destroy its long run viability. KPY report similar problems with a discriminatory first-stage auction. These results suggest a trade-off between types of mechanisms: One with clear equilibrium predictions insuring efficiency in theory, but involving relatively complex rules and calculations for bidders, the other with no clear equilibrium prediction but with relatively simple rules. As such the results are similar to those reported in KL (in press) for multi-unit demand Vickrey type auctions (see section 3.2 above).

V Summary and Conclusions

Experimental research in auctions has continued apace along with the extensive theoretical work on auctions since the appearance of Kagel's survey in *The Handbook of Experimental Economics* in 1995. Results reported in the original, 1995 survey focused on the Revenue Equivalence Theorem and initial investigations of the winner's curse, so that they could be easily summarized. In contrast, it is essentially impossible to summarize the work reported in the current survey in a few sentences given the much broader scope of the issues covered in experimental auction research since 1995. We anticipate a continued flowering of auction experiments given the mainstream applications of auctions for privatization of government assets, the continued growth of online auctions and business-to-business procurement auctions and the attention of theorists to better understand the many different variations in auction design in practice and in efforts to design more efficient auction institutions and/or ones that raise more revenue.

References

Andreoni, J., Y. K. Che, and J. Kim. 2007. Asymmetric information about rivals' types in standard auctions: An experiment. *Games and Economic Behavior* 59:240-259.

Ariely, D., A. Ockenfels, and A. E. Roth. 2005. An experimental analysis of ending rules in internet auctions. *RAND Journal of Economics* 36:890-907.

Armantier, O. 2002. Deciding between the common and private values paradigm: An application to experimental data. *International Economic Review* 43:783-801.

Armantier, O. and N. Treich. 2007. Subjective probabilities in games: An application to the overbidding puzzle. Working Paper. New York Federal Reserve.

Ashenfelter, O. 1989. How auctions work for wine and art. *Journal of Economic Perspectives* 3:33-36.

Ashenfelter, O. and D. Genesore. 1992. Testing for price anomalies in real estate auctions. *American Economic Review: Papers and Proceedings* 82:501-505.

Ausubel, L. M. 2004. An efficient ascending-bid auction for multiple objects. *The American Economic Review* 94:1452-1475.

Ausubel, L. M. and P. C. Cramton. 1996. Demand revelation and inefficiency in multi-unit auctions. Mimeograph. University of Maryland.

Ausubel, L and Milgrom, P. 2006. The lovely but lonely Vickrey auction. In *Combinatorial Auctions*, P. Cramton, Y. Shohani and R. Steinberg (eds). MIT Press.

Avery, C. and J. H. Kagel. 1997. Second-price auctions with asymmetric payoffs: An experimental investigation. *Journal of Economics and Management Strategy* 46:573-604.

Bajari, P. and A. Hortacsu. 2005. Are structural estimates of auction models reasonable? Evidence from experimental data. *Journal of Political Economy* 113:703-741.

Bajari, P. and A. Hortacsu. 2004. Economic insights from internet auctions. *Journal of Economic Literature* 42:457-486.

Bazerman, M. H. and W. F. Samuelson. 1983. I won the auction but don't want the prize. *Journal of Conflict Resolution* 27:618-634.

Bell, D. E. 1982. Regret in decision making under uncertainty. *Operations Research* 30:961-981.

Bikchandani, S. 1988. Reputation in repeated second-price auctions. *Journal of Economic Theory* 46:97-119.

Blecherman, B. and C. F. Camerer. 1998. Is there a winner's curse in the market for Baseball Players? Mimeograph. Brooklyn Polytechnic University**.**

Brosig, J. and J. P. Reib. 2007. Entry decisions and bidding behavior in sequential first-price procurement auctions: An experimental study. *Games and Economic Behavior* 58:50-74.

Brown, A. L., H. Kamp, and C. R. Plott. 2007. The nature of collusion facilitating and collusion breaking power of simultaneous ascending price and simultaneous descending price auctions. Forthcoming in Eco*nomic Inquiry*.

Brusco, S. and G. Lopomo. 2002. Collusion via signaling in simultaneous ascending bid auctions with heterogeneous objects, with and without and complementarities. *Review of Economic Studies* 69:407-436.

Brunner, C., J. K. Goeree, C. A. Holt, and J. O. Ledyard. 2007. An experimental test of flexible combinatorial spectrum auction formats. Working paper. California Institute of Technology.

Bulow, J. M. Huang and P. Klemperer. 1999. Toeholds and takeovers. *Journal of Political Economy* 107:427-454.

Burns, P. 1985. Experience and decision making: A comparison of students and businessmen in a simulated progressive auction. In Research in Experimental Economics, V. L. Smith (ed.) Jai Press, Greenwich, CT.

Campbell, C. and D. Levin. 2006. When and why not to auction. *Economic Theory* 27:583-596.

Cantillon, E. and M. Pesendorfer. 2007. Combination bidding in multi-unit auctions. CEPR discussion paper No. 6083.

Capen, E. C., R. V. Clapp, and W. M. Campbell. 1971. Competitive bidding in high-risk situations. *Journal of Petroleum Technology* 23: 641-653.

Casari, M., J. C. Ham, and J. H. Kagel. 2007. Selection bias, demographic effects and ability effects in common value auction experiments. *American Economic Review* 97:1278-1304.

Cassady, R. 1967. *Auctions and auctioneering*. Berkeley and Los Angeles: University of California Press.

Cassing, J. and R.W. Douglas. 1980. Implications of the auction mechanism in baseball's free agent draft. *Southern Economic Journal* 47:110-121.

Charness, G. and D. Levin. In press. The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics.*

Chen-Ritzo, C. H., T. P. Harrison, A. M. Kwasnica, and D. J. Thomas. In press. Better, faster, cheaper: An experimental analysis of a multi-attribute reverse auction mechanism with restricted information feedback. *Management Science*.

Chen, Y. and K. Takeuch. 2005. Multi-object auctions with package bidding: An experimental comparison of *i*BEA and Vickrey. Working paper. University of Michigan**.**

Chernomaz, K. 2006. On the effects of joint bidding in independent private value auctions: An experimental study. Unpublished manuscript. Ohio State University.

Chernomaz, K. and D. Levin. 2007. Efficiency and synergy in multi-unit auction with package bidding: An experimental study. Unpublished manuscript. Ohio State University.

Clarke, E. 1971. Multi-part pricing of public goods. *Public Choice* 11:19-33.

Cooper, D. J. and H. Fang. (in press). Understanding overbidding in second price auctions: An experimental study. *Economic Journal.*

Cooper, R. W., D. DeJong, R. Forsythe, and T. Ross. 1993. Forward induction in the battle of the sexes games. *American Economic Review* 83:1303-1316.

Coppinger, V. M., V. L. Smith, and J. A. Titus. 1980. Incentives and behavior in English, Dutch and sealed-bid auctions. *Economic Inquiry* 43:1-22.

Cox, J. C. and V. Sadiraj. 2008. Risky decisions in the large and in the small: Theory and experiment. In *Risk aversion in experiments*, J. C. Cox and G. W. Harrison eds., Research in Experimental Economics, Bingley, UK:Emerald.

Cramton, P. C. 1997. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy 6:431-495.*

Cramton, P. and J. A. Schwartz. 2002. Collusive bidding in the FCC spectrum auctions. *Contributions to Economic Analysis and Policy* 1:1-20.

Crawford, V. and N. Iriberri. 2007. Level-k auctions: Can boundedly rational strategic thinking explain the winner's curse? *Econometrica* 75:1721-70.

De Silva, D. G., T. Dunne, and G. Kosmopoulou. 2002. Sequential bidding in auctions of construction contracts. *Economic Letters* 76:239-244.

Dessauer, J. P. 1981. *Book publishing*. New York: Bowker.

Dorsey, R. and L. Razzolini. 2003. Explaining overbidding in first price auctions using controlled lotteries. *Experimental Economics* 6:123-140.

Duffy, J. and Ochs, J. 2006. Cooperative behavior and the frequency of social interactions. Working Paper. University of Pittsburgh.

Dyer, D., J. H. Kagel, and D. Levin. 1989. A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *Economic Journal* 99:108-115.

Dyer, D. and J. H. Kagel. 1996. Bidding in common value auctions: How the commercial construction industry corrects for the winner's curse. *Management Science* 42:1463-1475.

Eliaz, K., T. Offerman, and A. Schotter. 2008. Creating competition out of thin air: An experimental study of right-to-choose auctions. *Games and Economic Behavior* 62:383-416.

Engelbrecht-Wiggans, R. 1987. On optimal reservation prices in auctions. *Management Science* 33:763-770.

Engelbrecht-Wiggans, R. 1993. Optimal auctions revisited. *Games and Economic Behavior* 5:227-239.

Engelbrecht-Wiggans, R. and C. M. Kahn. 1998. Multi-unit auctions with uniform prices. *Economic Theory* 12:227-258.

Engelbrecht-Wiggans, R., J. A. List, and D. H. Reiley. 2005. Demand reduction in Multi-unit auctions: Evidence from a sportscard field experiment: Reply. *American Economic Review* 95:472-476.

Engelbrecht-Wiggans, R., J. A. List, and D. H. Reiley. 2006. Demand reduction in multi-unit auctions with varying numbers of bidders: Theory and evidence from a field experiment. *International Economic Review,* 47:203-231.

Engelbrecht-Wiggans, R. and E. Katok. 2007. Regret in auctions: Theory and Evidence. *Economic Theory* 33:81-101.

Engelbrecht-Wiggans, R., E. Haruvy, and E. Katok. In press. A comparison of buyer-determined and price-based multi-attribute mechanisms. *Management Science.*

Engelmann, D. and Grimm, V. 2004. Bidding behavior in multi-unit auctions – An experimental investigation and some theoretical insights. Working paper, Department of Economics, University of London.

Epstein, R., L. Henríquez, J. Catalán, G. Y. Weintraub, and C. Martínez. 2002. A combinational auction improves school meals in Chile. *Interfaces* 32:1-14.

Eyster, E. and M. Rabin. 2005. Cursed equilibrium. *Econometrica* 73:1623-1672.

Feynman, R. P. 2005. *Classic Feynman: All the Adventures of a Curious Character.* New York: W. W. Norton.

Filiz, E. and E. Y. Ozbay. 2007. Auctions with anticipated regret. *American Economic Review* 97:1407-1418.

Fisher R. A. 1935. *The Design of Experiments* Edinburgh: Oliver and Boyd.

Frechette, G. 2007. Session effects in the laboratory. Unpublished manuscript. New York University.

Garratt, R., M. Walker, and J. Wooders. 2004. Behavior in second-price auctions by highly experienced eBay buyers and sellers. Working paper, No. 1181., Department of Economics, University of California Santa Barbara.

Georganas, S. and R. Nagel. 2007. Auctions with Toeholds: An experimental study. Mimeograph. University of Bonn.

Ginsburgh, V. 1998. Absentee bidders and the declining price anomaly in wine auctions. *Journal of Political Economy* 106:1302-1319.

Goeree, J. K., C. A. Holt, and T. R. Palfrey. 2002. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104:247-272.

Goeree, J. K. and T. Offerman. 2002. Efficiency in auctions with private and common values: An experimental study. *American Economic Review* 92:625-643.

Goeree, J. K. and T. Offerman. 2003. Competition bidding in auctions with private and common values. *Economic Journal* 113:598-614.

Goeree, J. K., C. Plott, and J. Wooders. 2004. Bidders' choice auctions: Raising revenues through the right to choose. *Journal of the European Economics Association* 2:504-515.

Goeree, J. K., T. Offerman, and R. Sloof. 2006. Demand reduction and preemptive bidding in multi-unit license auctions. Mimeograph**.** California Institute of Technology.

Goswami, G., T. H. Noe, and M. J. Rebello. 1996. Collusion in uniform-price auction: Experimental evidence and implications for treasury auctions. *The Review of Financial Studies* 9:757-785.

Groves, T. 1973. Incentives in teams. *Econometrica* 41:617-631.

Güth, W., R. Ivanova-Stenzel, and E. Wolfstetter. 2005. Bidding behavior in asymmetric auctions: An experimental study. *European Economic Review* 49:1891-1913.

Ham, J. C., J. H. Kagel, and S. F. Lehrer. 2005. Randomization, endogeneity and laboratory experiments: The role of cash balances in private value auctions. *Journal of Econometrics* 125:175-205.

Hansen, R. G. and J. R. Lott, Jr. 1991. The winner's curse and public information in common value auctions: Comment. *American Economic Review* 81:347-61.

Harrison, G. W. and J. A. List. In press. Naturally occurring markets and exogenous laboratory experiments: A case study of the winner's curse," *Economic Journal*.

Hausch, D. B. 1987. An asymmetric common value auction model. *RAND Journal of Economics* 18:611-621.

Hendricks, K., R. H. Porter, and C. A. Wilson. 1994. Auctions for oil and gas leases with an informed bidder and a random reservation price. *Econometrica*. 62:1415-1444.

Holt, C. A. Jr. and R. Sherman. 1994. The loser's curse and bidder's bias. *American Economic Review* 84:642-652.

Hu, A., T. Offerman, and S. Onderstal. 2006. Fighting collusion in auctions. Discussion paper. Faculty of Economics and Business, Universiteit van Amsterdam.

Isaac, R. M. and D. James. 2000a. Just who are you calling risk averse? *Journal of Risk and Uncertainty*. 20:177-187.

Isaac, R. M. and D. James. 2000b. Robustness of the incentive compatible combinatorial auction. *Experimental Economics* 3:31-53.

Ivanova-Stenzal, R., and T. C. Salmon. 2007. Revenue equivalence revisited. Mimeograph. Florida State University.

Jofre-Bonet, M. and M. Pesendorfer, 2000. Bidding behavior in a repeated procurement auction. *Euorpean Economic Review* Revised 44:1006-1020.

Jofre-Bonet, M. and M. Pesendorfer, 2003. Estimation of a dynamic auction game. *Econometrica* 71:1443-1489.

Johnson-Laird, P. N. 1999. Deductive reasoning. *Annual Review of Psychology* 50:109-135.

Kagel, J. H. and D. Levin. 1986. The winner's curse and public information in common value auctions. *American Economic Review* 76:894-920.

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* 55:1275-1304.

Kagel, J. H. and D. Levin. 1991. The winner's curse and public information in common value auctions: Reply. *American Economic Review* 81:362-369.

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. *Economic Journal* 103:868-879.

Kagel, J. H. and D. Levin. 1999. Common value auctions with insider information. *Econometrica* 67:1219-1238.

Kagel, J. H., S. Kinross, and D. Levin. 2001. Comparing efficient multi-object auction institutions. Working paper, Ohio State University.

Kagel, J. H. and D. Levin. 2001. Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickery auctions. *Econometrica* 69:413-454.

Kagel, J. H. and J. F. Richard. 2001. Super-experienced bidders in first-price common value auctions: Rules of thumb, Nash equilibrium bidding and the winner's curse. *Review of Economics and Statistics* 83:408-419.

Kagel, J. H. and D. Levin. 2002. Bidding in common value auctions: A survey of experimental research. In *Common Value Auctions and the Winner's Curse.* Princeton University Press.

Kagel, J. H. and D. Levin. 2005. Multi-unit demand auctions with synergies: Behavior in sealed-bid versus ascending-bid uniform-price auctions. *Games and Economic Behavior* 53:170-207.

Kagel, J. H., S. Pevnitskaya, and L. Ye. 2008. Indicative bidding: An experimental analysis. *Games and Economic Behavior* 62:697-721.

Kagel, J. H., Y. Lien, and P. Milgrom. 2008 Ascending prices and package bidding: An experimental analysis. Working paper. Ohio State University.

Kagel, J. H., S. Pevnitskaya, and L. Ye. 2007. Survival auctions. *Economic Theory* 33:103-119.

Kagel, J. H. and D. Levin. in press. Implementing efficient multi-object auction institutions: An experimental study of the performance of boundedly rational agents. *Games and Economic Behavior*.

Kagel, J. H. 1995. Auction: Survey of experimental research. In Alvin E. Roth and John H. Kagel, Editors, Princeton University Press.

Katok, E. and A. E. Roth. 2004. Auctions of homogeneous goods with increasing returns: Experimental comparison of alternative "Dutch" auctions. *Management Science* 50:1044-1063.

Keser, C. and M. Olson. 1996. Experimental examination of the declining-price anomaly. In *Economics of the Arts: Selected Essays.* V. Ginsburg and P-M Menger (eds). Amsterdam: Elsevier.

Klemperer, P. 1998. Auctions with almost common values: The ‘Wallet Game’ and its applications. *European Economic Review* 42:757-769.

Klemperer, P. 2002. How (not) to run auctions: The European 3G Telecom auctions. *European Economic Review* 46:829-845.

Kwasnica, A. M. and K. Sherstyuk. In press. Collusion and equilibrium selection in auctions. *Economic Journal*.

Ledyard, John O., David P. Porter, and Antonio Rangel. 1997. Experiments testing multi-object allocation mechanisms. *Journal of Economics and Management Strategy* 6:639-75.

Leufkens, K., R. Peeters, and D. Vermeulen. 2006. Sequential auctions with synergies: The paradox of positive synergies. METEOR Research Memorandum 06/18, Universiteit Maastricht. 1-19.

Levin, D. 1990. Horizontal mergers: The 50 percent benchmark. *American Economic Review* 80:1238-1245.

Levin, D. 2005. Demand reduction in Multi-unit auctions: Evidence from a sportscard field experiment: A comment. *American Economic Review* 95:467-471.

Levin, D. and J. Smith. 1994. Equilibrium in auctions with entry. *American Economic Review* 84:585-599.

Levin, D., J. H. Kagel, and J. F. Richard. 1996. Revenue effects and information processing in English common value auctions. *American Economic Review* 86:442-460.

Levin, D. and J. H. Kagel. 2005. Almost common-value auctions revisited. *European Economic Review* 49:1125-1136.

Li, J. and C. R. Plott. 2005. Tacit collusion in auctions and conditions for its facilitation and prevention: Equilibrium selection in laboratory experimental markets. Unpublished Working Paper. Massachusetts Institute of Technology.

Lind, B. and C. R. Plott. 1991. The winner’s curse: Experiments with buyers and with sellers. *American Economic Review* 81:335-346.

List, J. A. and D. Lucking-Reiley. 2000. Demand reduction in multi-unit auctions: Evidence from a sportscard field experiment. *American Economic Review* 90:961-972.

Loomes, G. and R. Sugden. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* 92:805-824.

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *American Economic Review* 89:1063-1080.

Manelli, A. M., M. Sefton, and B. S. Wilner. 2006. Multi-unit auctions: A comparison of static and dynamic mechanisms. *Journal of Economic Behavior and Organization* 61: 304-23.

Maskin, E. and J. Riley. 2000. Asymmetric auctions. *Review of Economic Studies* 67:413-438.

Mathews, T. 2003. The impact of discounting on an auction with a buy-out option: A theoretical analysis motivated by eBay's Buy-It-Now feature. *Journal of Economics (Zietschrift für Nationalökonomie*) 81:25-52.

Mathews, T. and B. Katzman. 2006. The role of varying risk attitudes in an auction with a buyout option. *Economic Theory* 27:597-613.

McAfee, R. P. and J. McMillan. 1987. Auctions with entry. *Economic Letters* 23:343-347.

McAfee, R. P. and D. Vincent. 1993. The declining price anomaly. *Journal Economic Theory* 60:191-212.

Mellers, B. A. and A. D. J. Cooke. 1996. The role of task and context in preference measurement. *Psychological Science* 7:76-82.

Menezes, F. M. and P. K. Monteiro. 1998. Simultaneous pooled auctions. *Journal of Real Estate Finance and Economics* 17:219-232.

Menkhaus, D. J., O. R. Phillips and K.T. Coatney. 2003. Shared agents and competition in laboratory English auctions. *American Journal of Agricultural Economics* 85:829-839.

Milgrom P. R. and R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50:1089-1112.

Morgan, J. 2002. Combinatorial auctions in the information age: an experimental study. *Advances in Applied Microeconomics* 11:191-207.

Myerson, E. 1981. Optimal auction design. *Mathematics of Operations Research*. 6:58-73.

Neugebauer, T. and R. Selten. 2006. Individual behavior and first-price auctions: The importance of information feedback in computerized experimental markets. *Games and Economic Behavior* 54:183-204.

Neugebauer, T. and P. Pezanis-Christou. 2007. Bidding behavior at sequential first-price auctions with(out) supply uncertainty: A laboratory analysis. *Journal of Economic Behavior & Organization* 63: 55-72.

Niederle, M. and L. Vesterlund. 2007. Do women shy away from competition? *Quarterly Journal of Economics* 122:1067-1101.

Ochs, J. 1995. Coordination problems. In J. H. Kagel and A. E. Roth (ed) *Handbook of Experimental Economics.* Princeton University Press, Princeton. 195-251.

Offerman, T. and J. Potters. 2006. Does auctioning of entry licenses induce collusion? An experimental study. *Review of Economic Studies* 73:769-791.

Palfrey, T. and S. Pevnitskaya. 2008. Endogenous entry and self-selection in private value auctions: An experimental study. *Journal of Economic Behavior and Organization* 66:731-747.

Pezanis-Christou, P. 2002. On the impact of low-balling: Experimental results in asymmetric auctions. *International Journal of Game Theory* 31:69-89.

Phillips, O. R., D. J. Menkhaus, and K. T. Coatney. 2003. Collusive practices in repeated English auctions: Experimental evidence on bidding rings. *American Economic Review* 93:965-979.

Pitchik C. and A. Schotter. 1988. Perfect equilibria in budget-constrained sequential auctions: An experimental study. *Rand Journal of Economics* 19**:**363-388.

Porter, D. and R. Vragov. 2006. An experimental examination of demand reduction in multi-unit versions of the uniform-price, Vickery, and English auctions. *Managerial and Decision Economics* 27:445-458.

Rabin, M. 2000. Risk aversion and expected utility theory. *Econometrica* 68:1281-1292.

Rapoport, A., D. Seale, I. Erev, and J. Sundali. 1998. Equilibrium play in large group market entry games. *Management Science* 44:119-141.

Reynolds, S. S. and J. Wooders. Auctions with a buy price. Forthcoming in *Economic Theory*.

Riley, J. G. and W. F. Samuelson. 1981. Optimal auctions. *American Economic Review* 71:381-392.

Robert, J. and C. Montmarquette. 1999. Sequential auctions with multi-unit demand: Theory, experiments and simulations. CIRANO working paper, 99s-46.

Roll R. 1986. The hubris hypothesis of corporate takeovers. *Journal of Business* 59:197-216.

Rose, S. L. and J. H. Kagel. In press. Bidding in almost common value auctions: An experiment. *Journal of Economic Management Strategy.*

Rose, S. L. and D. Levin. In press. An experimental investigation of the explosive effect in almost common value auctions. *Journal of Economic Behavior and Organization.*

Roth, A. E. and A. Ockenfels. 2002. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the internet. *American Economic Review* 92:1093-1103.

Sade, O., C. Schnitzlein, and J. F. Zender. 2006. Competition and cooperation in divisible good auctions: An experimental examination. *Review of Financial Studies* 19:195-235.

Salmon, T. C. and B. J. Wilson. 2008. Second chance offers versus sequential auctions: Theory and behavior. *Economic Theory* 34:47-67.

Salmon, T. C. and M. Iachini. 2007. Continuous ascending vs. pooled multiple unit auctions. *Games and Economic Behavior* 61:67-85.

Salmon, T. C. and B. J. Wilson. 2008. Second chance offers vs. sequential auctions: Theory and behavior. *Economic Theory* 34:47-67.

Salomon, G. and D. N. Perkins. 1989. Rocky roads to transfer: Rethinking mechanisms of a neglected phenomenon. *Education Psychologist 24:113-142.*

Samuelson, W.F. and M.H. Bazerman. 1985. The winner's curse in bilateral negotiations, in V.L. Smith (ed.), *Research in Experimental Economics*, 3. Greenwich, CT.: JAI Press.

Shahriar, Q. and J. Wooders. 2007. An experimental study of auctions with a buy price under private and common values. Working Paper 07-19. University of Arizona.

Sherstyuk, K. 1999. Collusion without conspiracy: An experimental study of one-sided auctions. *Experimental Economics* 2:59-75.

Sherstyuk, K. 2002. Collusion in private value ascending price auctions. *Journal of Economic Behavior and Organization* 48:177-195.

Shogren, J. F., G. M. Parkhurst, and C. McIntosh. 2006. Second-price auction tournament. *Economics Letters* 92:99-107.

Smith, J. 1982. Equilibrium patterns of bidding in OCS lease sales. *Economic Inquiry* 20:180-190.

Smith, J. L. 1984. Further results on equilibrium patterns of bidding in OCS lease sales. *Economic Inquiry* 22:142-146.

Smith, J. and D. Levin. 2001. Entry coordination in auctions and social welfare: An experimental investigation. *International Journal of Game Theory* 30:321-350.

Tor, A and M. H. Bazerman. 2003. Focusing failures in competitive environments: Explaining decision errors in the Monty Hall game, the acquiring a company game, and multiparty ultimatums. *Journal of Behavioral Decision Making* 16:353-374.

Turocy, T. L. and E. Watson. 2007. Reservation values in laboratory auctions: Context and bidding behavior. Working Paper, Department of Economics, Texas A&M, 06-03. Revised 2007.

Turocy, T. L., E. Watson and R. C. Battalio. 2007. Framing the first-price auction. *Experimental Economics* 10:37-52.

Tversky, A. and D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Unvertainty* 5:297-323.

Walker, J. M., V. L. Smith, and J. C. Cox. 1987. Bidding behavior in first-price sealed-bid auctions: Use of computerized Nash competitors. *Economic Letters* 23:239-244.

Whinston, M. D. 2006. *Lectures on Antitrust Economics*. MIT Press.

Wilson, B. R. 1967. Insider competitive bidding with asymmetric information. *Management Science* 13:816-820.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8-37.

Ye, L. 2004. Optimal auctions with endogenous entry. *Countributions to Theoretical Economics* 4:1-27.

Table 1

Linear Estimations of Bidding Strategies under Regret
(standard errors in parentheses)

	Winner Regret	Loser Regret	No Regret
Slope	0.77	0.87	0.79
	(0.012)	(0.01)	(0.007)
Lower 95 percent	0.748	0.852	0.775
Upper 95 percent	0.796	0.893	0.805

From Felig and Ozbay (2007)

Table 2

Inexperienced Bidders: Actual versus Theoretical Revenue Changes and Profit Levels[a]
in English versus First-Price Auctions

	n=4						n=7					
	Average Change in Revenue: English Less First-Price (standard error)		Average Profit (standard error)				Average Change in Revenue: English Less First-Price (standard error)		Average Profit (standard error)			
			First-Price		English				First-Price		English	
ε	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical
$6	-1.54*	1.54**	-2.13	2.76	0.58	1.23	-1.98*	0.10	-3.85	0.99	-1.87	0.89
	(0.72)	(0.49)	(0.52)	(0.38)	(0.50)	(0.30)	(0.87)	(0.34)	(0.71)	(0.19)	(0.51)	(0.29)
			[29]		[28]				[18]		[18]	
$12	0.54	2.76**	-1.32	5.01	-0.78	2.25	-1.95	1.08	-3.75	2.76	-1.80	1.68
	(1.25)	(0.92)	(0.79)	(0.60)	(0.95)	(0.69)	(1.19)	(0.65)	(0.89)	(0.53)	(0.77)	(0.40)
			[41]		[45]				[30]		[43]	
$24	1.09	8.10**	1.20	9.83	0.11	1.73	ND	ND	ND		ND	
	(3.29)	(2.32)	(1.93)	(1.25)	(2.64)	(2.14)						
			[25]		[13]							

[a] All values reported in dollars.
[+] The null hypothesis that the value is greater than or equal to zero can be rejected at the 10% significance level.
* The null hypothesis that the value is greater than or equal to zero can be rejected at the 5% significance level.
** The null hypothesis that the value is greater than or equal to zero can be rejected at the 1% significance level.
ND No data

From Levin, Kagel and Richard (1996)

Table 3

Super-Experienced Bidders: Actual versus Theoretical Revenue Changes and Profit Levels[a] in English versus First-Price Auctions

	n=4						n=7					
	Average Change in Revenue: English Less First-Price (standard error)		Average Profit (standard error)				Average Change in Revenue: English Less First-Price (standard error)		Average Profit (standard error)			
			First-Price		English				First-Price		English	
ε	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical	Actual	Theoretical
$18	2.21*	3.96**	3.37	6.77	1.16	2.82	-0.25	2.85**	0.76	3.86	1.01	1.01
	(0.95)	(0.73)	(0.50)	(0.48)	(0.88)	(0.53)	(0.86)	(0.61)	(0.65)	(0.50)	(0.56)	(0.37)
			[163]		[107]				[75]		[96]	
	1.20	2.98	8.45	11.27	7.25	8.29						
$30	(3.10)	(2.30)	(1.28)	(1.34)	(2.76)	(1.93)	ND		ND		ND	
			[31]		[33]							

[a] All values reported in dollars. Super-experienced bidders had participated in at least two previous first-price common value auction sessions.
* The null hypothesis that the value is greater than or equal to zero can be rejected at the 5% significance level.
** The null hypothesis that the value is greater than or equal to zero can be rejected at the 1% significance level.
ND No data

From Levin, Kagel and Richard (1996)

Table 4

Change in Seller's Revenue: Auction with Insider versus No Insider
(Super-Experienced Bidders)

	n=4			n=7		
	Change in Revenue: Insider *less* No Insider [a] (t-stat)[b]	Mean Profits (σ^2)		Change in Revenue: Insider *less* No Insider[a] (t-stat)[b]	Mean Profits (σ^2)	
		Insiders	No Insiders		Insiders	No Insiders
$\varepsilon = \$18$	1.759 (2.057)*	2.063 (8.561)	3.822 (49.972)	0.739 (1.573)+	1.492 (6.770)	2.231 (19.221)
$\varepsilon = \$30$	2.734 (1.097)	6.148 (24.334)	8.876 (59.731)	0.919 (0.425)	4.517 (17.978)	5.436 (15.839)

* Significantly different from 0 at $p < .05$, one-tailed test.
\+ Significantly different from 0 at $p < .10$, one-tailed test.
From Kagel and Levin (1999)

Table 5
Comparing Frequency of Equilibrium Play Under Different Auction Institutions

v_h	No. Computers	Clock	Sealed Bid*
$3.00	3	46.3%	2.6%
$4.00	3	23.7%	1.6%
	5	22.3%	3.1%
$4.40	3	38.8%	27.7%
	5	35.8%	27.1%
$5.10	5	79.2%	40.6%

V_h is value for human bidder.

From Kagel and Levin (2005)

Table 6

Realized and Predicted Prices: First Price Sequential Auctions

	Unit 1	Unit 2	Unit 2	Unit 4
Predicted[a] Average (std)	444 (41)	446 (80)	449 (100)	426 (133)
Realized				
Mean (std) Median	500 (104) 492	474 (76) 470	463 (70) 461	454 (121) 456

a Based on bidders' realized valuations

std = standard deviations

From Keser and Olson (1996).

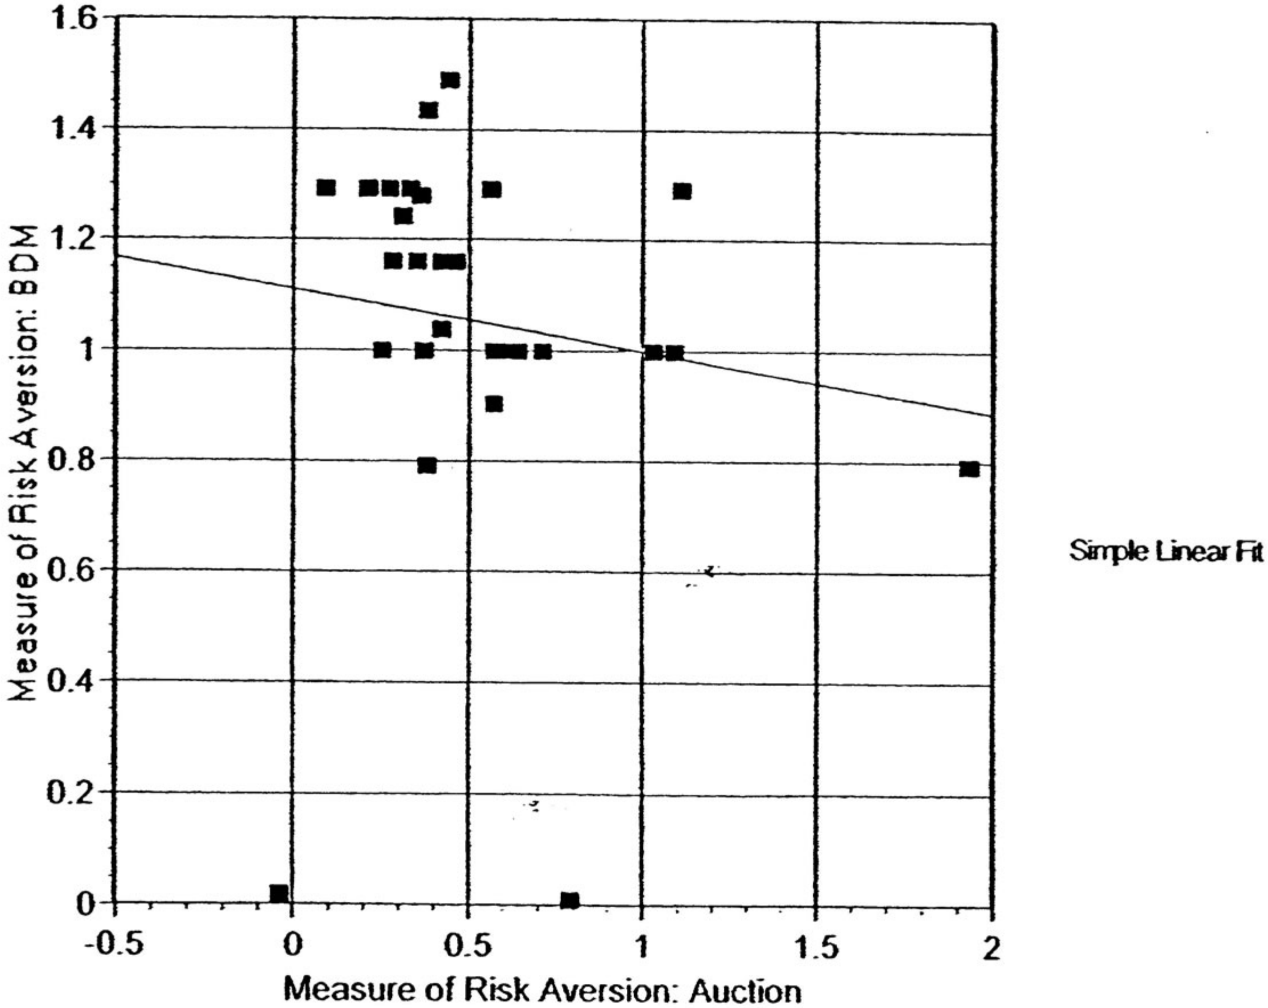

Figure 1: From Isaac and James (2000)

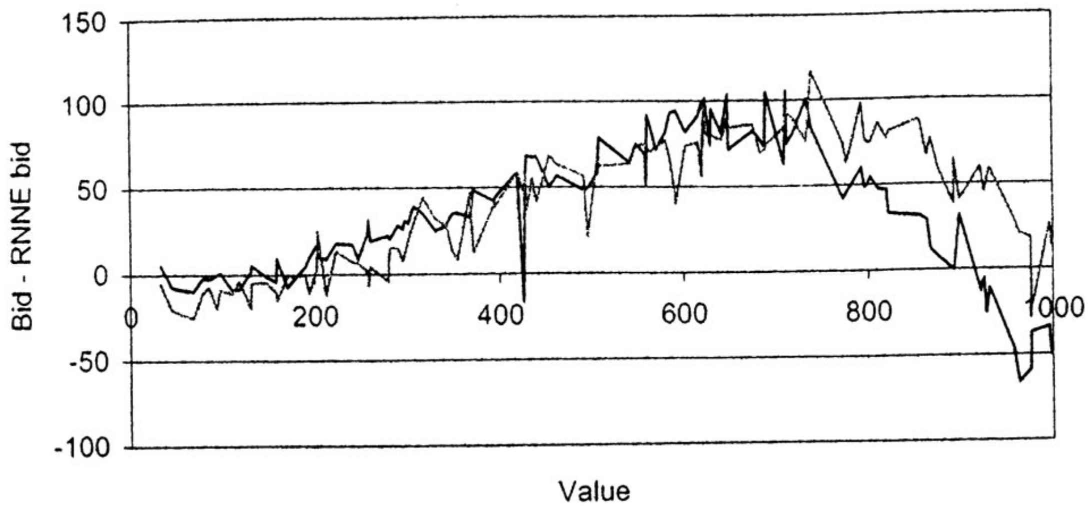

Deviations of mean auction bids (gray) and mean lottery-equivalent bids (black) from RNNE bid.

Figure 2: From Dorsey and Razzolini (2003)

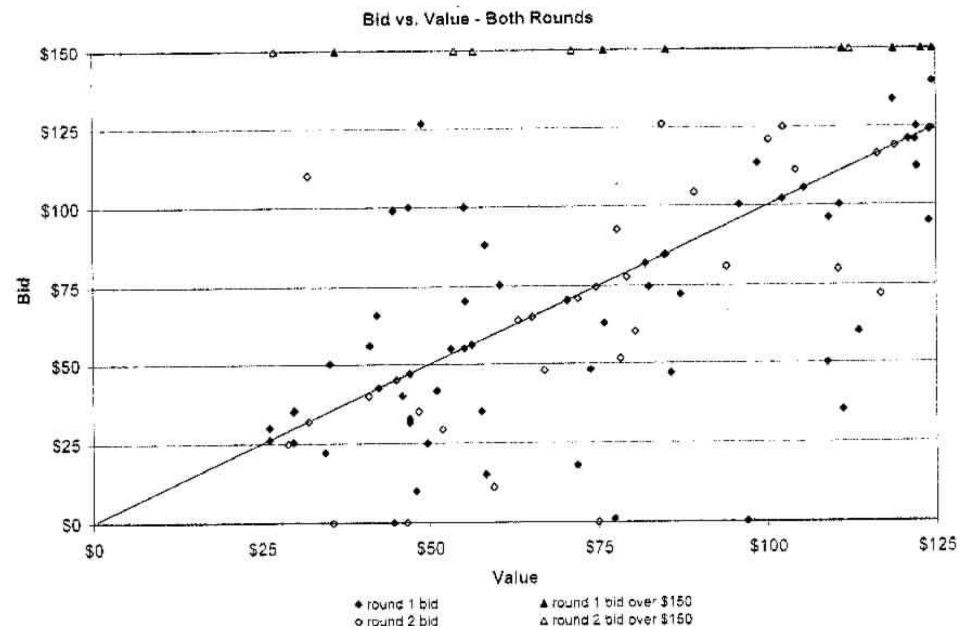

Figure 3: Bids and values in second-price internet auctions.

From Garrat, Walker and Wooders (2004)

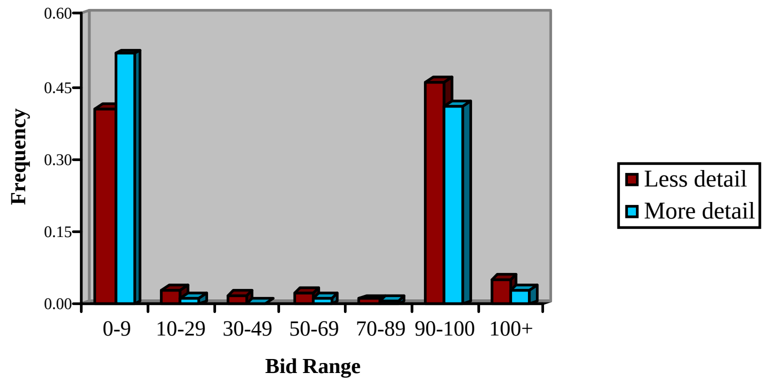

Figure 4: Bid frequencies in two value treatment in Charness and Levin (in press).

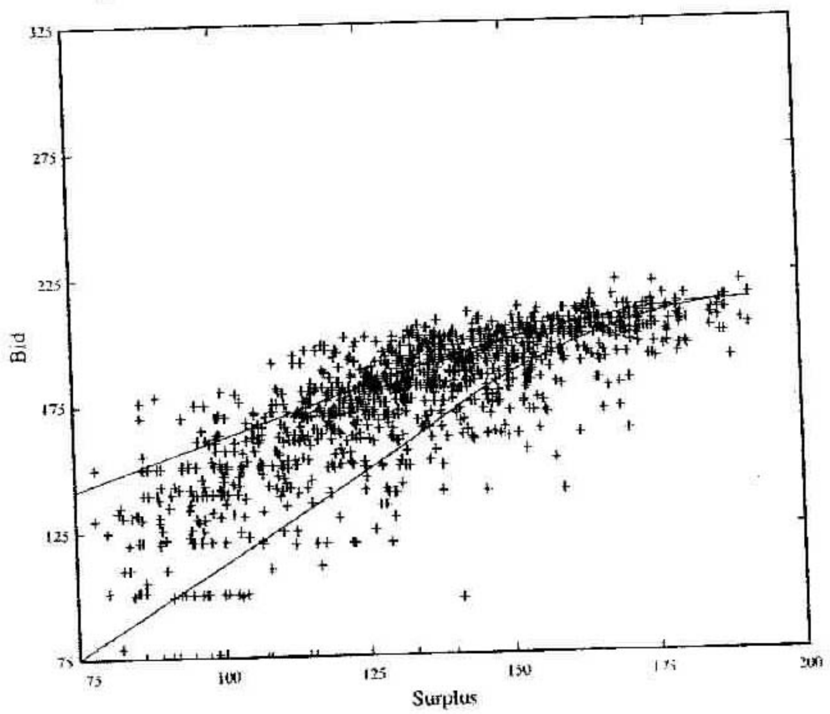

Figure 5: Bids (+) together with Nash Bids (lower line) and Naïve Bids (top line). Amount of overbidding (relative to Nash) tends to be higher when surplus is smaller as winning the auction is more informative about the common value in this case.

From Goeree and Offerman (2002).

Figure 6

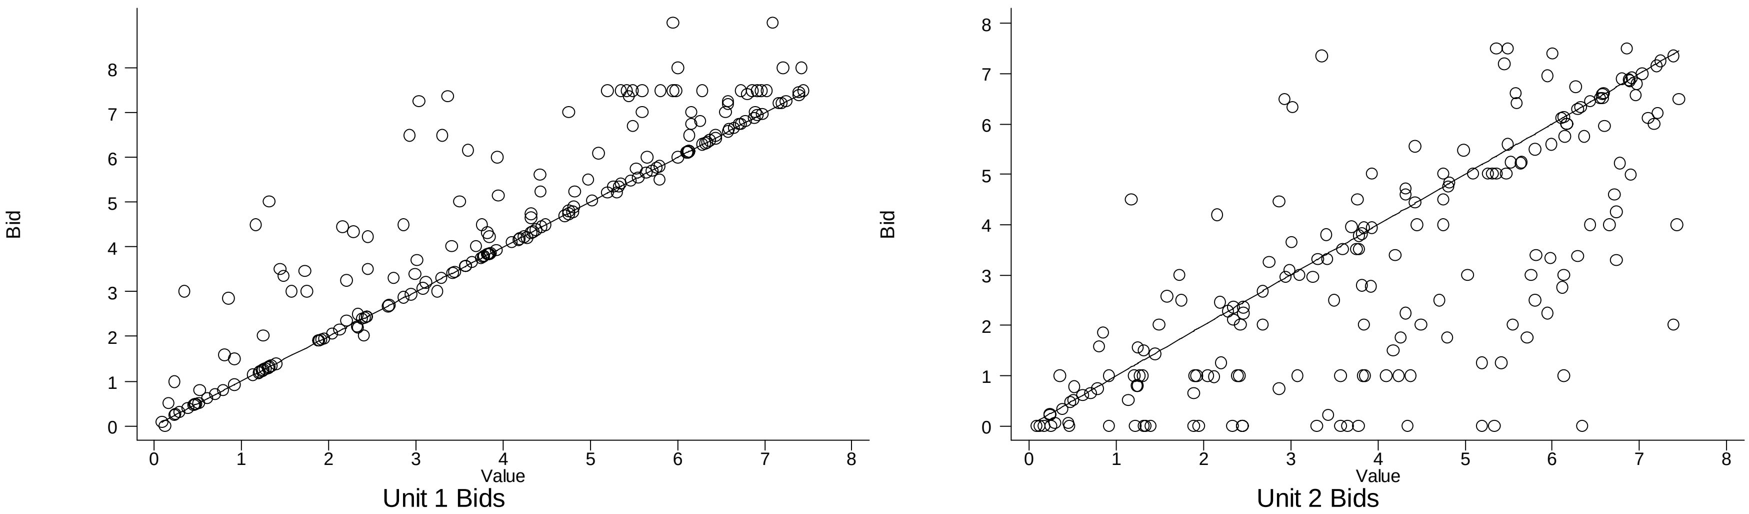

Bids relative to value for the last 12 auctions. Solid line bids equal to value.

From Kagel and Levin (2001).

Figure 7

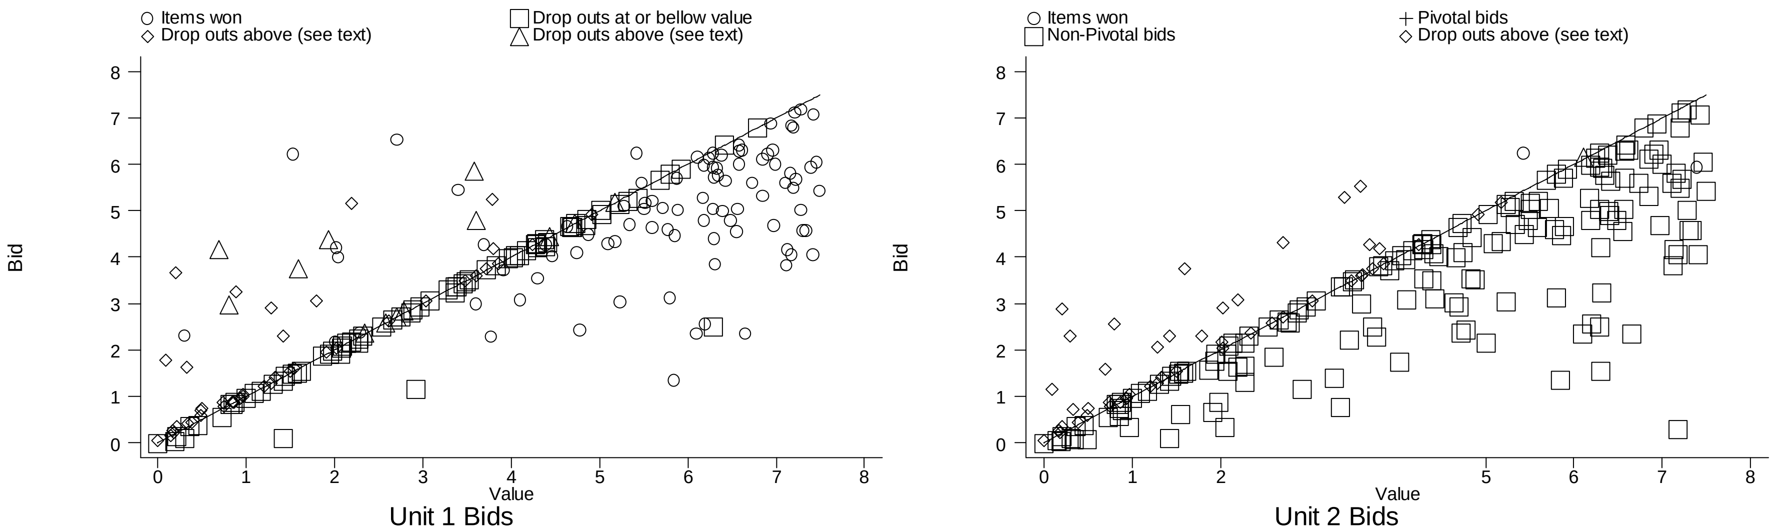

Left panel: Unit 1 bids. Circles are winning bids (these are censored). Squares are dropouts at or below resale values. Triangles are potentially harmful bids above value (had a single compute rival dropped out bidder would have earned negative profits). Diamonds are harmless bids above (bidder dropped out before potentially earning negative profits). Solid line: bids equal to value.

Right panel: Unit 2 bids. Circles are winning bids (these are censored). Squares are dropouts at or below second highest computer value (these represent optimal bids). Dropout that are pivotal (set the market price and are suboptimal) are +'s. Diamonds are harmless bids above (bidder dropped out before potentially earning negative profits). Solid line: bids equal to value.

From Kagel and Levin (2001).

Figure 8

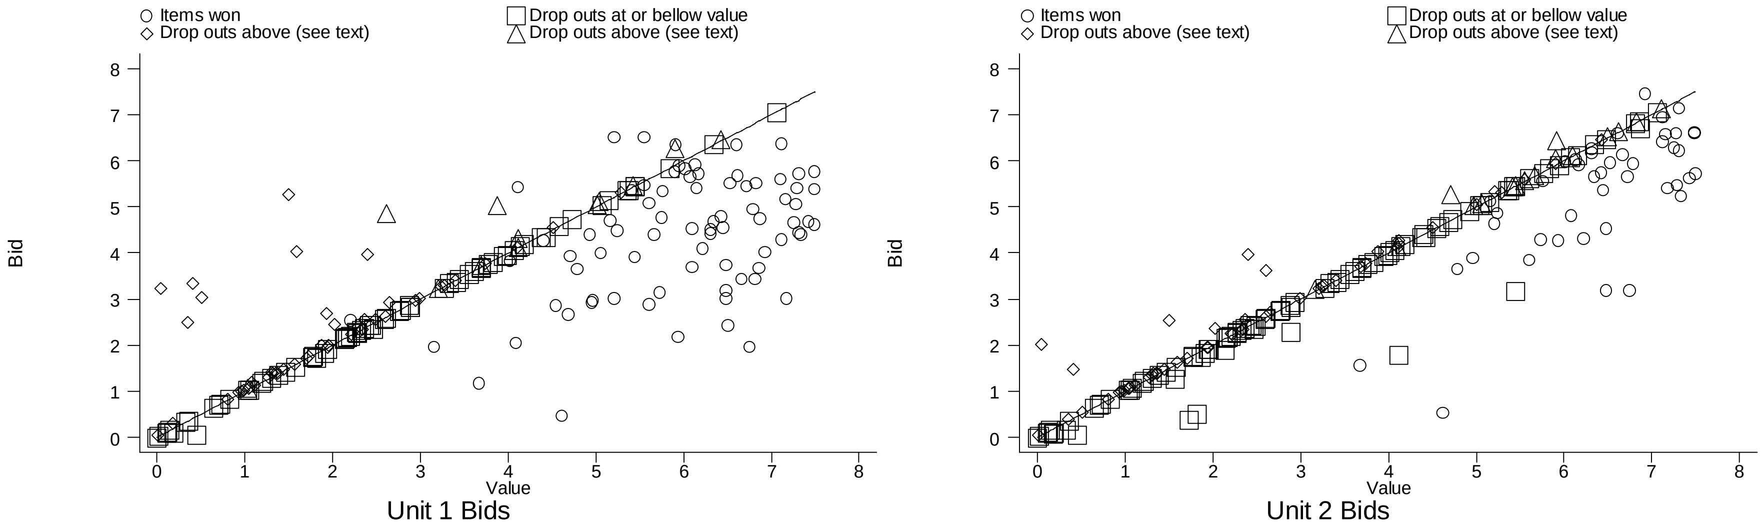

Bids relative to value in last 12 auctions. Circles are winning bids (these are censored). Squares are dropouts at or below resale values. Triangles are potentially harmful bids above value (had a single compute rival dropped out bidder would have earned negative profits). Diamonds are harmless bids above (bidder dropped out before potentially earning negative profits). Solid line: bids equal to value.

From Kagel and Levin (2001).

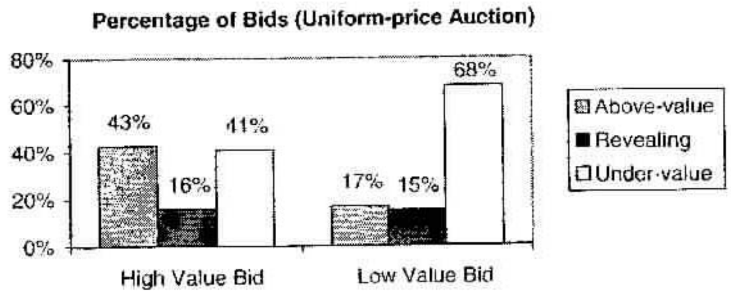

Classification of bids in the Uniform-price treatment. Bids within 5% of value are categorized as revealing.

Figure 9: From Porter and Vragov (2006)

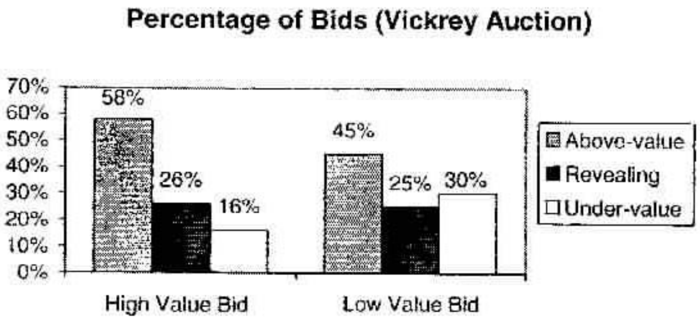

Classification of bids in the Vickery treatment. Bids within 5% of value are categorized as revealing.

Figure 10: From Porter and Vragov (2006)

Figure 11

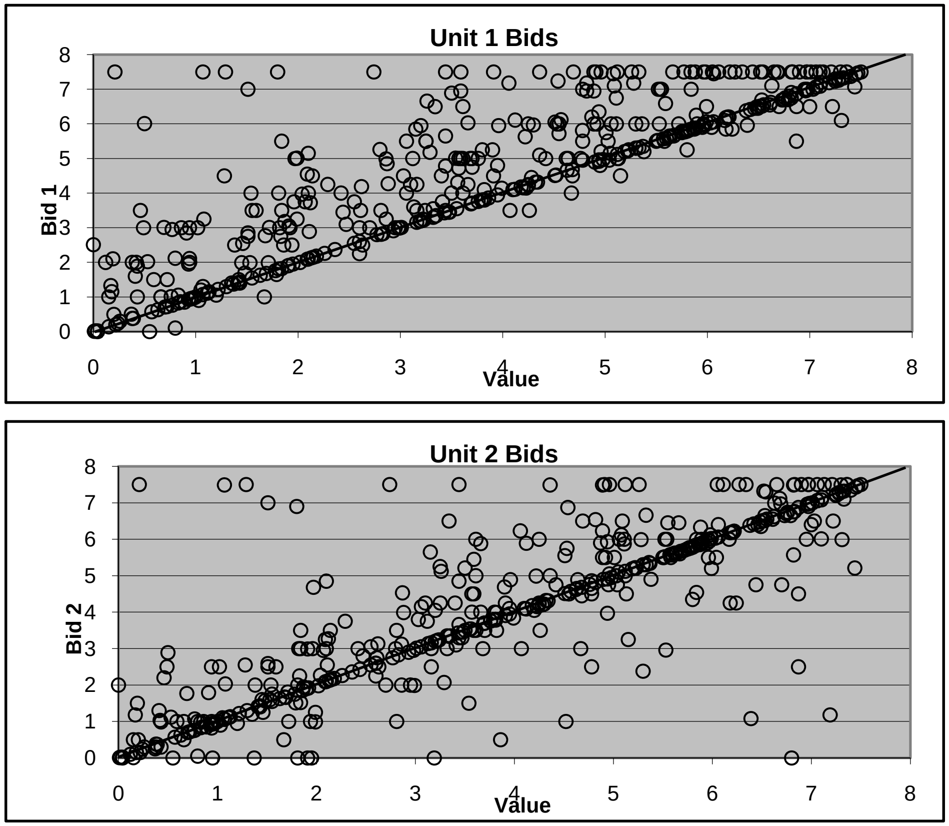

Sealed-bid Vickrey auctions with computer rivals.

From Kagel, Kinross and Levin (2001).

Figure 12

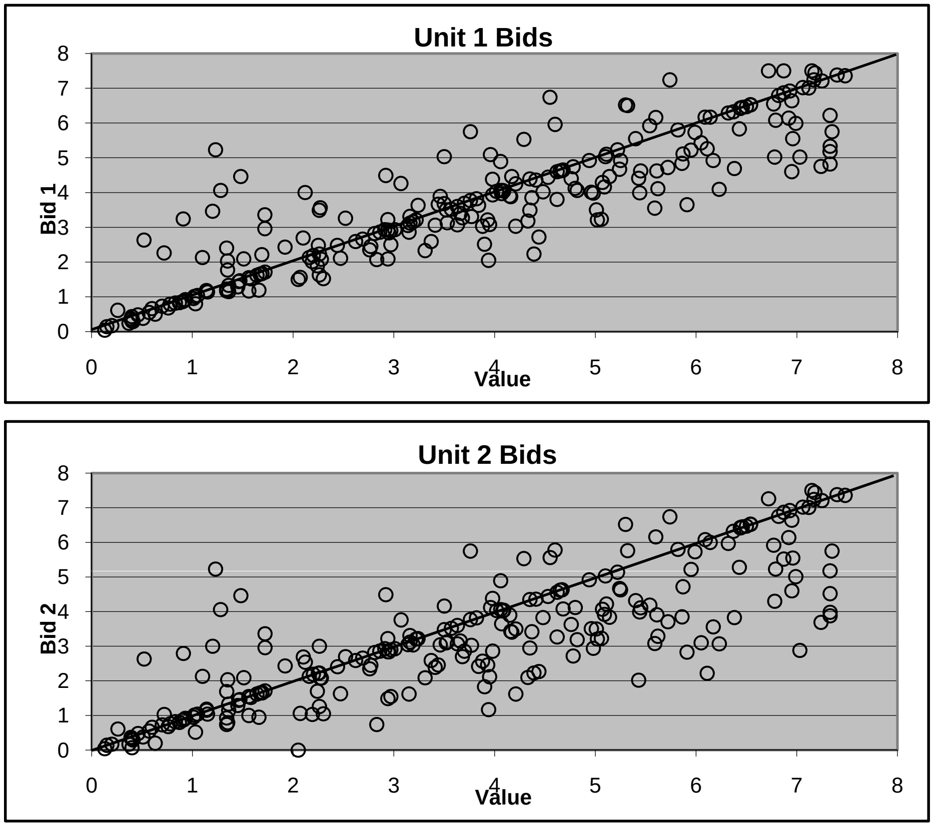

Ausubel auctions without drop-out information provided.

From Kagel, Kinross and Levin (2001).

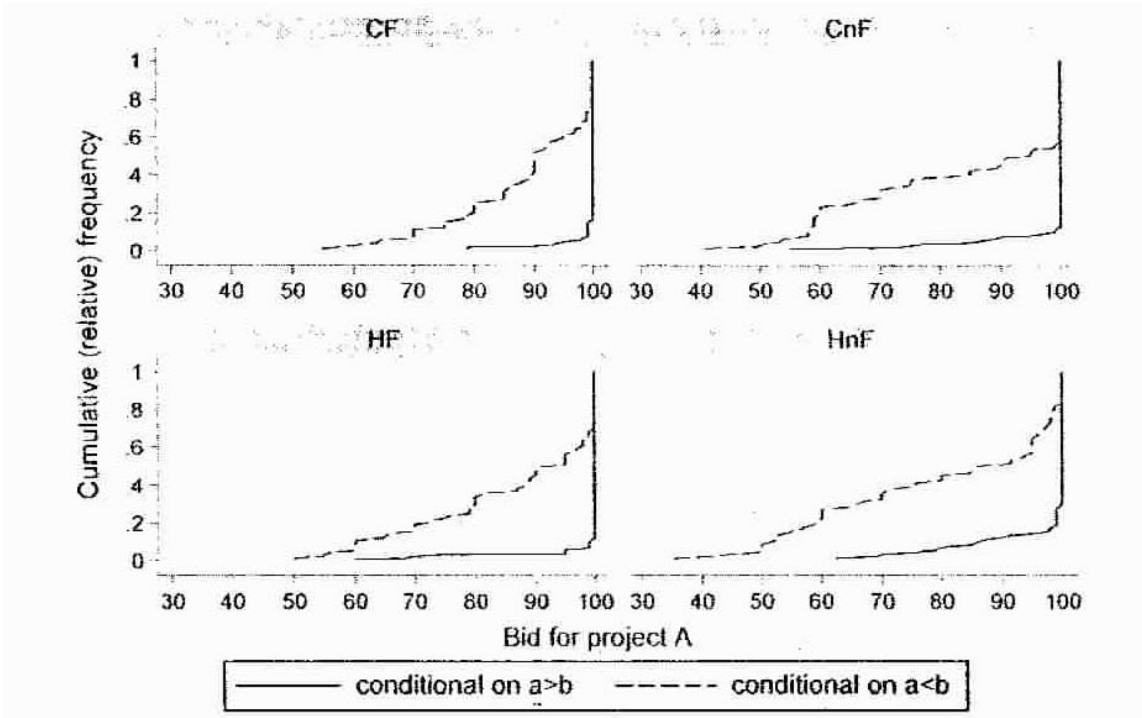

Figure 13: Empirical distributions of bids for project A conditional on its cost relative to B and the different treatments. CF and CnF = computerized rivals with and without feedback. HF and HnF = human rivals with and without feedback.

From Brosig and Reiβ (2007).

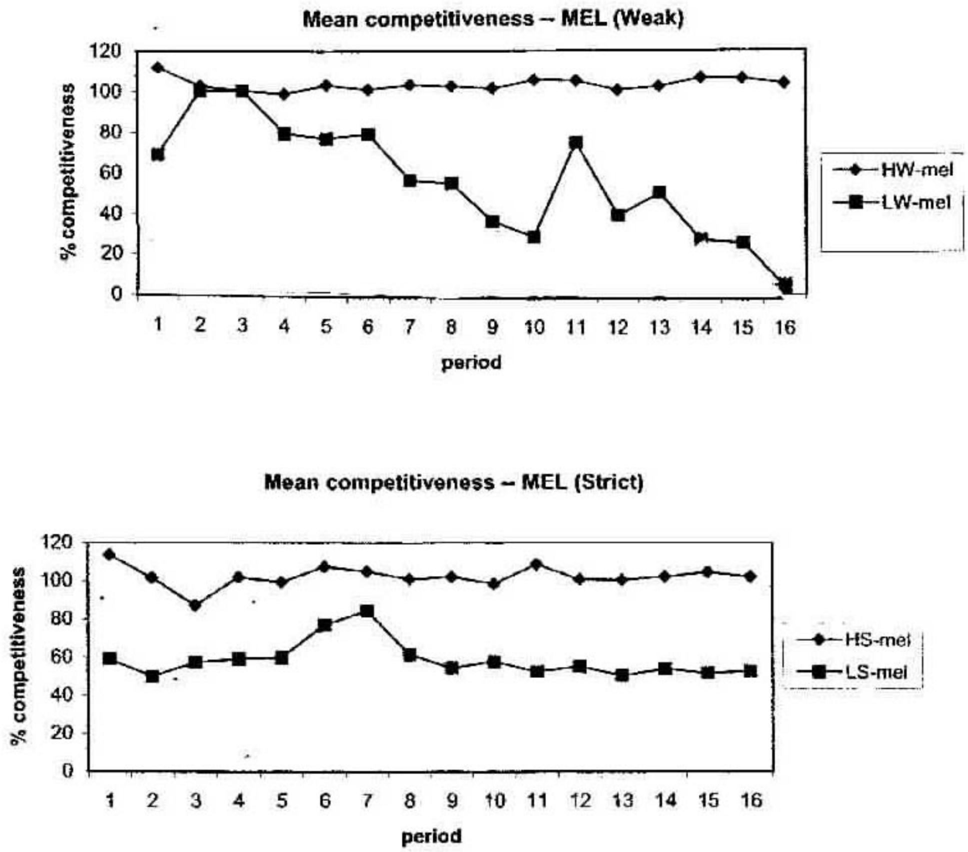

Figure 14: Bids under weal improvement rule (top panel) with support [50, 90] (HW–mel) and with support [90, 100] (LW-mel). Bids under strict improvement rule (bottom panel) with support [50, 90] (HS-mel) and with support [90, 100] (LS-mel). All subjects from University of Melbourne,

From Sherstyuk (2002).

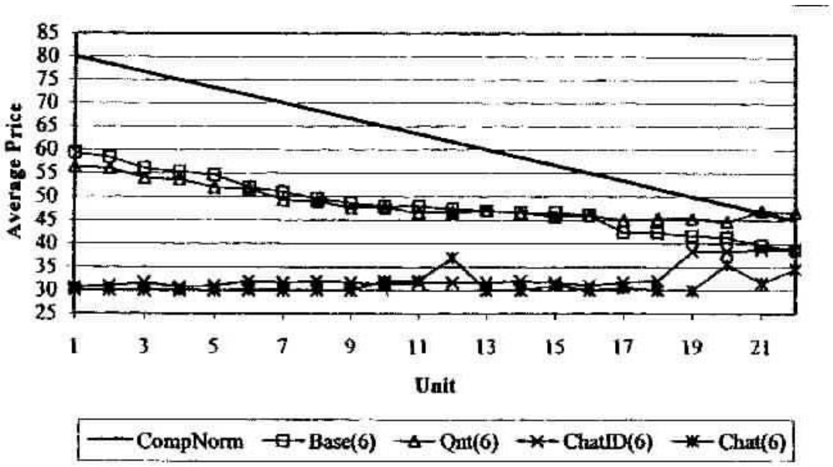

Figure 15: Average prices under different treatments in sequential auctions with six bidders. Qnt treatment bidders know the number of units for sale. Chat treatment bidders are allowed to chat and collude. ChatID is Chat but bidder IDs revealed in sales.

From Phillips.

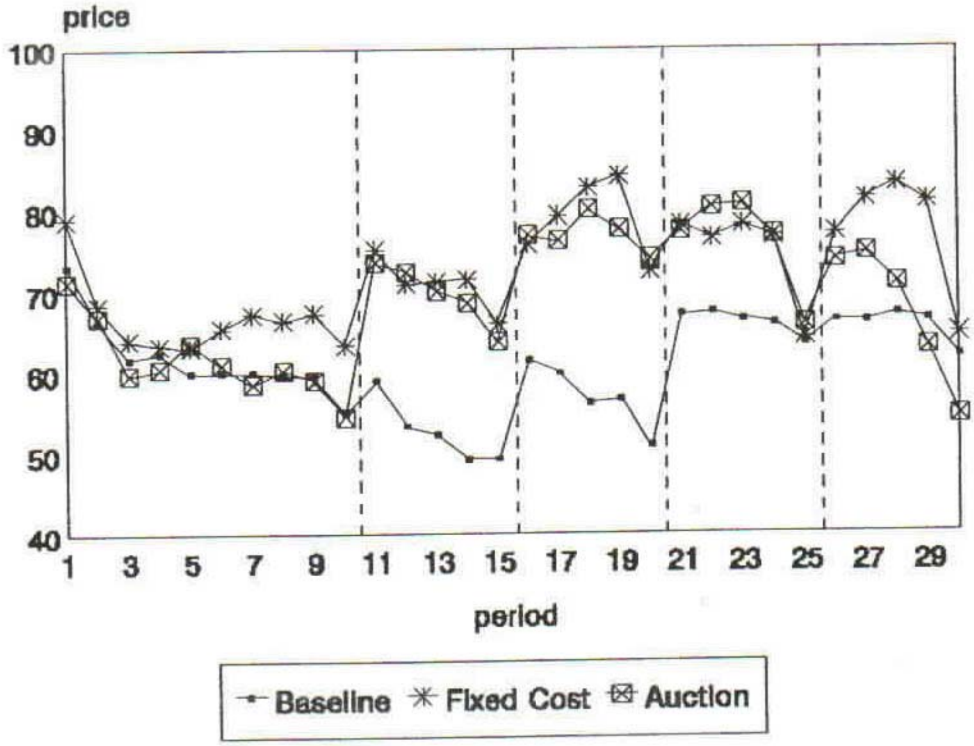

Figure 16: Effects on market prices of entry fees. Periods 1-10 no entry fees; Periods 11-30 entry fees. In Baseline right to produce for market randomly assigned. In Fixed Cost treatment entry right randomly assigned, with entrants paying exogenously determined fees. In Auction treatment potential entrants bid for the right to enter the market.

From Offerman and Potters (2006).