Common Value Auctions and the Winner's Curse:

Lessons from the Economics Laboratory*

John H. Kagel

and

Dan Levin

Department of Economics
Ohio State University

October, 2001

* Research support from the Economics and DRMS Divisions of NSF, the Sloan Foundation and the Russell Sage Foundation are gratefully acknowledged. Special thanks to my colleagues and my coauthors, especially Dan Levin, who have taught me so much. Much of the material here is taken from our more comprehensive survey titled "Bidding in Common Value Auctions: A Survey of Experimental Research," which appears as Chapter 1 in the collection of our published papers investigating common value auctions: John H. Kagel and Dan Levin, Common Value Auctions and the Winner's Curse, Princeton University Press, in press. We invite the interested reader to consult the full survey and the papers in the book for a more extensive and detailed review of common value auction experiments.

Auctions are of considerable practical and theoretical importance. In practical terms, the value of goods exchanged in auctions each year is huge. Governments routinely use auctions to purchase goods and services, to sell government assets, and to fund the national debt. Private sector auctions are common as well, and of growing importance in areas such as deregulated utility markets, allocation of pollution rights, and the large variety of items now being sold via Internet auctions. Auctions are commonly employed when one party to the exchange (for example the seller) is uncertain about the value that buyers place on the item. Auctions provide a mechanism, absent middlemen, to establish value in such situations. Auctions play a prominent role in the theory of exchange as they remain one of the simplest and most familiar means of price determination in the absence of intermediate market makers. In addition, auctions serve as valuable illustrations, and one of the most prominent applications, of games of incomplete information, as bidders' private information is the main factor affecting strategic behavior (Wilson, 1992).

Auctions have traditionally been classified as one of two types: Private-value auctions, where bidders know the value of the item to themselves with certainty but there is uncertainty regarding other bidders' values. Common-value auctions, where the value of the item is the same to everyone but different bidders have different estimates about the underlying value. Most (non-laboratory) auctions have both private value and common value elements. There are also many different methods for auctioning items, with first-price sealed-bid auctions and open outcry English auctions being the most common institutions. In analyzing auctions, economists have focused on questions of economic efficiency (getting items into the hands of the highest valued bidders), on maximizing sellers' revenue, and on how auctions aggregate information. The most developed branch of the literature deals with single unit auctions, where a single item is sold to a number of competing bidders or a number of sellers compete for the right to supply a single item. Recent Federal government spectrum (air wave rights) auctions have exposed many gaps in economists' knowledge about auctions in which multiple units of closely related items are sold, and in which individual bidders demand more than a single unit of the commodity.

The chapters in this book all deal with single unit common-value auctions. As noted, in a pure common-value auction the ex post value of the item is the same to all bidders. What makes the auction interesting is that bidders do not know the value at the time they bid. Instead they receive signal values that are correlated with the value of the item.[1] Mineral rights auctions, particularly the Federal government's outer continental shelf (OCS) oil lease auctions, are typically modeled as pure common-value auctions. There is a common-value element to most auctions. Bidders for an oil painting may purchase for their own pleasure, a private-value element, but they may also bid for investment and eventual resale, reflecting the common-value element.

There are no efficiency issues in pure common-value auctions as all bidders place equal value on the item[2]. What has been of overriding concern to both theorists and practitioners is the revenue raising effect of different auction institutions. A second key issue, one that provides much of the focus for the essays in this book, is the winner's curse, an *unpredicted* effect that was initially postulated on

[1]Or, more technically, signals that are affiliated with the value of the item. See Milgrom and Weber, 1982, for an excellent presentation, discussion and analysis of the statistical properties of affiliated variables in the context of auctions.

[2]However, once the seller uses a minimum bid requirement, and/or we consider entry to be determined endogenously, different auctions may induce different probabilities of an actual sale. Thus efficiency may become an issue (Levin and Smith, 1994).

the basis of field data, and whose existence has often been hotly debated among economists.

The winner's curse story begins with Capen, Clapp, and Campbell (1971), three petroleum engineers who claimed that oil companies had suffered unexpectedly low rates of return in the 1960's and 1970's on OCS lease sales "year after year."[3] They argued that these low rates of return resulted from the fact that winning bidders ignored the informational consequences of winning. That is, bidders naively based their bids on the unconditional expected value of the item (their own estimates of value) which, although correct on average, ignores the fact that you only win when your estimate happens to be the highest (or one of the highest) of those competing for the item. But winning against a number of rivals following similar bidding strategies implies that your estimate is an overestimate of the value of the lease *conditional on the event of winning*. Unless this adverse selection effect is accounted for in formulating a bidding strategy, it will result in winning bids that produce below normal or even negative profits. The systematic failure to account for this adverse selection effect is commonly referred to as the winner's curse: you win, you lose money, and you curse.

Terminological aside: Unfortunately, many economists, particularly theorists, when discussing the winner's curse use the term to refer to the difference between the expected value of the item conditional on the event of winning and the naive expectation (not conditioning on the event of winning). Further, their use of the term typically refers to players who fully account for this winner's curse, rather than those who fall prey to it.

The idea that oil companies suffered from a winner's curse in OCS lease sales was greeted with

[3] Unless of course, one argues that Groucho Marks statement "I do not wish to join any club that accepts me," is an earlier recognition of the winner's curse.

skepticism by many economists as it implies that bidders repeatedly err, violating basic economic notions of rationality and contrary to equilibrium predictions.[4] An alternative and simpler explanation as to why oil companies might claim that they fell prey to a winner's curse lies in cartel theory, as responsiveness to the winner's curse claim could serve as a coordination device to get rivals to reduce their bids in future sales. Nevertheless, claims that bidders fell prey to the winner's curse have arisen in a number of field settings. In addition to the oil industry (Capen, Clapp, and Campbell, 1971; Lorenz and Dougherty, 1983 and references cited therein), claims have been made in 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 in real estate auctions (Ashenfelter and Genesore, 1992).

It is exceedingly difficult to support claims of a winner's curse using field data because of reliability problems with the data and because alternative explanations for overbidding are often available. For example, Hendricks, Porter, and Boudreau (1987) found that in early OCS lease sales, average profits were negative in auctions with seven or more bidders. Hendricks et al. note that one possible explanation for this outcome is the increased severity of the adverse selection problem associated with more bidders. However, they note that the data could also be explained by bidder uncertainty regarding the number of firms competing on a given tract (their preferred explanation). That is, since most tracts received less than six bids, it seems likely that firms would expect this number or less. As a result, although firms might have fully accounted for the adverse selection effect based on the

[4] See, for example, the exchange between Cox and Isaac (1984, 1986) and Brown (1986).

expected number of firms bidding on a tract, they would nevertheless be incorrect for tracts that attracted above average numbers of bidders, and overbid on those tracts.

The ambiguity inherent in using field data, in conjunction with the controversial nature of claims regarding a winner's curse, provided the motivation for experimental studies of the winner's curse. Early laboratory experiments showed that inexperienced bidders are quite susceptible to the winner's curse (Bazerman and Samuelson, 1983; Kagel and Levin, 1986; Kagel, et al., 1989). In fact, the winner's curse has been such a pervasive phenomenon in the laboratory that most of these initial experiments have focused on its robustness and the features of the environment that might attenuate its effects. Additional interest has focused on public policy issues --- the effects of public information regarding the value of the auctioned item and the effects of different auction institutions on sellers' revenue.

This survey begins with a brief analysis of the first experimental demonstration of the winner's curse (Bazerman and Samuelson, 1983). This is followed by summaries of experiments investigating bidding in common-value auctions using an experimental design that I helped develop. These experiments also demonstrate the existence of a winner's curse even when allowing for extensive feedback and learning from past auction outcomes. They also address policy issues such as the effects of public information and different auction institutions (e.g., first price sealed-bid auctions versus open outcry English auctions) on sellers' revenue. I conclude with a brief summary the empirical findings from the experimental literature and the role experiments have played in the successful sale of government airwave rights (the spectrum auctions). In reviewing the experimental work on common value auctions, I hope to give the reader a flavor for how experiments proceed by successively narrowing down plausible explanations for the question at hand. This is done through a series of experiments rather than any single "critical" experiment, and is based on sorting out between competing explanations, and on following up on the logical implications of behavior observed in earlier experiments.

A. An initial experiment demonstrating the winner's curse:

Bazerman and Samuelson (1983) conducted the first experiment demonstrating a winner's curse. Using M.B.A. students at Boston University, the experiment was conducted in class, with students participating in four first-price sealed-bid auctions. Bidders formed their own estimates of the value of each of four commodities - jars containing 800 pennies, 160 nickels, 200 large paper clips each worth four cents, and 400 small paper clips each worth two cents. Unknown to subjects, each jar had a value of $8.00. (Subjects bid on the value of the commodity, not the commodity itself.) In addition to their bids, subjects provided their best estimate of the value of the commodities and a 90% confidence bound around these estimates. A prize of $2.00 was given for the closest estimate to the true value in each auction. The number of bidders varied between 4 and 26. Their analysis focused on bidder uncertainty about the value of the commodity and the size of the bidding population.

The average value estimate across all four commodities was $5.13 ($2.87 below the true value). As the authors note, this underestimation should reduce the likelihood and magnitude of the winner's curse. In contrast to the mean estimate, the average winning bid was $10.01 resulting in an average loss to the winner of $2.01.[5] The average winning bid generated losses in over half of all the auctions.

[5]Winning bidders paid these losses out of their own pockets or from earnings in other auctions.

Estimated bid functions, using individual bids as the unit of observation, showed that bids were positively, and significantly, related to individual estimates so that bidders indeed faced an adverse selection problem, only winning when they had higher estimates of the value of the item. Bids were inversely related to the uncertainty associated with individual estimates, but this effect was small (other things equal, a $1.00 increase in the 90% confidence interval reduced bids by 3¢). Numbers of bidders had no significant effect on individual bids.

In contrast, regressions employing the average winning bid showed that these bids were positively, and significantly, related to the winning bidder's estimate of uncertainty and to the number of bidders in the auction. This suggests that winning bidders are substantially more aggressive than other bidders. Indeed, Bazerman and Samuelson note that average winning bids were sensitive to a handful of grossly inflated bids.

The results of this experiment show that the winner's curse is easy to observe. However, many economists would object to the fact that subjects had no prior experience with the problem and no feedback regarding the outcomes of their decisions between auctions, so that the results could be attributed to the mistakes of totally inexperienced bidders. The robustness of these results is even more suspect given that their sensitivity to a handful of grossly inflated bids, which one might suppose would be eliminated as a result of bankruptcies or learning in response to losses incurred in earlier auctions. Common-value auction experiments conducted by Kagel and Levin and their associates explore these issues, along with a number of public policy implications of the theory.

B. Sealed-bid Auctions

Kagel and Levin and their associates conducted experiments in which bidders participated in a series of auctions with feedback regarding outcomes. Bidders were given starting cash balances from which losses were subtracted and profits were added. Bidders whose cash balances became negative were declared bankrupt and were no longer permitted to bid. Unlike the Bazerman and Samuelson experiment, Kagel and Levin (hereafter, KL) controlled the uncertainty associated with the value of the auctioned item rather than simply measuring it. They did this by conducting auctions in which the common value, x_o, was chosen randomly each period from a known uniform distribution with upper and lower bounds [$\underline{x}$, $\overline{x}$]. In auctions with a symmetric information structure each bidder is provided with a private information signal, x, drawn from a uniform distribution on [x_o-å, x_o+å], where å is known. In first-price sealed-bid auctions, bids are ranked from highest to lowest with the high bidder paying the amount bid and earning profits equal to x_o - b_1, where b_1 is the highest bid. Losing bidders neither gain nor lose money.

In this design the strategy of bidding, max [x - å, $\underline{x}$], is a risk-free strategy that fully protects a bidder from negative earnings since it is the lower bound estimate of x_o. This lower bound estimate for x_o was computed for subjects along with an upper bound estimate of x_o, (min [x + å, $\overline{x}$]). Bidders were provided with illustrative distributions of signal values relative to x_o and several dry runs were conducted before playing for cash. Following each auction period bidders were provided with the complete set of bids, listed from highest to lowest, along with the corresponding signal values, the value of x_o and the earnings of the high bidder.

Surviving bidders were paid their end-of-experiment balances in cash. To hold the number of bidders fixed while controlling for bankruptcies, m > n subjects were often recruited, with only n bidding at any given time (who bids in each period was determined randomly or by a fixed rotation rule). As bankruptcies occur m shrinks, but (hopefully) remains greater than or equal to the target value n.

B.1 Theoretical Considerations: First-Price Sealed-Bid Auctions

Wilson (1977) was the first to develop the Nash equilibrium solution for first-price common-value auctions, and Milgrom and Weber (1982) provide significant extensions and generalizations of the Wilson model.[6] In the analysis that follows, we restrict our attention to signals in region 2, the interval $\underline{x} + \varepsilon \le x \le \bar{x} - \varepsilon$, where the bulk of the observations lie.[7] Within region 2, bidders have no end point information to help in calculating the expected value of the item.[8]

For risk neutral bidders the symmetric risk neutral Nash equilibrium (RNNE) bid function $\gamma(x)$ is given by[9]

$$(1)\quad \gamma(x) = x - \varepsilon + h(x)$$

$$\textit{where}\ \ h(x) = \frac{2\varepsilon}{n+1} \exp\left[-\frac{n}{2\varepsilon}\left[x - (\underline{x} + \varepsilon)\right]\right]$$

This equilibrium bid function combines strategic considerations similar to those involved in first-price

[6]In a Nash equilibrium no bidder has any incentive to unilaterally deviate from the proposed outcome.

[7]For data outside this interval see Kagel and Richard (2001).

[8]For example, with a signal $x < \underline{x} + \varepsilon$ the bidder knows that $x - \varepsilon$ is smaller than $\underline{x}$, and can use this additional end point information to more precisely compute the expected value of the item (e.g., $x_o \in (\underline{x}, x + \varepsilon)$ which is smaller than the interval $(x - \varepsilon, x + \varepsilon)$.

[9]Derivation of the RNNE bid function for this design can be found in an appendix to Levin, Kagel, and Richard (1996) and Kagel and Richard (2001) A symmetric Nash equilibrium is one in which all bidders use the same bidding strategy but where actual bids are based on different private signals. An asymmetric Nash equilibrium is one in which different bidders employ different bidding strategies. Most equilibrium solutions assume symmetry as (i) it seems a natural assumption for most settings and (ii) it is often difficult to solve for asymmetric equilibria.

private-value auctions, and item valuation considerations resulting from the bias in the signal value conditional on the event of winning. We deal with the latter first.

In common-value auctions bidders usually win the item when they have the highest, or one of the highest estimates of value. Define $E[x_o| X = x_{1n}]$ to be the expected value of the item conditional on having x_{1n}, the highest among n signal values. For signals in region 2

(2) $E[x_o|X = x_{1n}] = x - [(n-1)/(n+1)]$ å.

This provides a convenient measure of the extent to which bidders suffer from the winner's curse since in auctions in which the high signal holder always wins the item, as bidding above $E[x_o| X = x_{1n}]$ results in negative expected profit. Further, even with zero correlation between bids and signal values, if everyone else bids above $E[x_o|X = x_{1n}]$, bidding above $E[x_o|X = x_{1n}]$ results in negative expected profit as well. As such, if the high signal holder frequently wins the auction, or a reasonably large number of rivals are bidding above $E[x_o|X = x_{1n}]$, bidding above $E[x_o|X = x_{1n}]$ is likely to earn negative expected profit.

Recall that within region 2, (x - å) is the smallest possible value for x_o, and that x is the unconditional expected value of x_o (the expected value, *independent* of winning the item), so that the expected value, conditional on winning, must be in-between (x - å) and x. Thus, from equation (2) it is clear that the amount bids ought to be reduced, relative to signal values (the "bid factor"), just to correct for the adverse selection effect from winning the auction, is quite large relative to the range of sensible corrections (å): with n = 4 the bid factor is 60% of å and with n = 7 it is 75% of å. Or put another way, for signals in the region 2 the RNNE bid function is well approximated by ã(x) = x - å (the negative exponential term $h(x)$ in equation 1, approaches zero rapidly as x moves beyond $\underline{x}$ + å).

Thus, the bid factor required just to avoid losing money, on average, represents 60% of the total bid factor with n = 4, and 75% with n = 7. Equation (2) also makes it is clear that the correction for the adverse selection effect is relatively large and increasing with increases in the number of bidders.

Strategic considerations account for the rest of the bid factor; 2å/(n+1). The strategic element results from the fact that if just correcting for the adverse selection effect, the winner would earn zero expected profits, which is not a very attractive outcome. As such, a bidder would find it profitable to lower her bid from this hypothetical benchmark (equation 2) since zero expected surplus is lost by doing so even if this causes her not to win the item, and strictly positive expected surplus is awarded should she win the item with the lower price. The interplay of these strategic considerations between different bidders results in the additional discounting of bids relative to signal values beyond equation (2).

B.2 Some Initial Experimental Results: Inexperienced Bidders

Auctions with inexperienced bidders show a pervasive winner's curse that results in numerous bankruptcies. Table1 provides illustrative data on this point. For the first nine auctions profits averaged -$2.57 compared to the RNNE prediction of $1.90, with only 17% of all auctions having positive profits. Note, this is after bidders had participated in 2-3 dry runs, with feedback of signal values, x_o, and bids following each auction, so that the results cannot be attributed to a total lack of experience. The negative profits are not a simple matter of bad luck either, or a handful of grossly inflated bids, as 59% of all bids and 82% of the high bids were above $E[x_o|X = x_{1n}]$. Further, 40% of all subjects starting these auctions went bankrupt. In short, the winner's curse is a genuinely pervasive problem for inexperienced bidders. It is remarkably robust being reported under a variety of treatment conditions

(Kagel et al, 1989; Lind and Plott, 1991, Goeree and Offerman, 2000) and for different subject populations, including professional bidders from the commercial construction industry (Dyer, Kagel and Levin, 1989).

[Insert Table1 here]

B.3 Auctions with Moderately Experienced Bidders and the Effects of Public Information on Sellers' Revenue

Kagel and Levin (1986) report auctions for moderately experienced bidders (those who had participated in at least one prior first price common-value auction experiment). Treatment variables of interest were the number of rival bidders and the effects of public information about x_o on revenue. Table 2 reports some of their results. For small groups (auctions with 3-4 bidders), the general pattern was one of positive profits averaging $4.32 per auction, which is significantly greater than zero, but still well below the RNNE prediction of $7.48 per auction. In contrast, for these same bidders, bidding in larger groups (auctions with 6-7 bidders) profits averaged -$0.54 per auction, compared to the RNNE prediction of $4.82. Thus, the profit picture had improved substantially compared to the inexperienced bidders discussed in the previous section.

[Insert Table 2 here]

However, comparing large and small group auctions, actual profit decreased substantially more than profit opportunities as measured by the RNNE criteria. This implies that subjects were bidding more aggressively, rather than less aggressively, as the number of rivals increased, contrary to the RNNE prediction. This is confirmed in regressions using individual subject bids as the dependent variable. Higher individual bids in response to increased numbers of rivals is often considered to be the hallmark characteristic of a winner's curse. Thus, although bidders had adjusted reasonably well to the adverse selection problem in auctions with 3-4 bidders, in auctions with 6-7 bidders, with its heightened adverse selection effect, the winner's curse reemerged as subjects confounded the heightened adverse selection effect by bidding more aggressively with more bidders. This result also suggests that the underlying learning processes is context specific rather than involving some sort of "theory absorption" that readily generalizes to new environments.[10]

Public information was provided to bidders in the form of announcing the lowest signal value, x_L. For the RNNE, public information about the value of the item raises expected revenue. The mechanism underlying this outcome works as follows: All bidders evaluate the additional public information assuming that their signal is the highest since, in equilibrium, they only win in this case. Evaluating additional information from this perspective, together with affiliation, induces all bidders other than the highest signal holder to, on average, revise their bids upward after an announcement of unbiased public information. This upward revision results from two factors:

1. Affiliation results in bidders without the highest signal systematically treating the public information as "good news." These bidders formulated their bids on the assumption they held the highest private information signal and would win the auction. As such, with affiliation, the public information tells them that, on average, the expected value of the item is higher than they had anticipated (i.e., the private information signal they are holding is somewhat lower than expected, conditional on winning, for this particular auction), which leads them to increase their

[10]There is a whole body of psychological literature indicating the difficulty of learning generalizing across different contexts (see, for example, Gick and Holyoak, 1980; Perkins and Salomon, 1988; Salomon and Perkins, 1989). Having read these papers provides me with some comfort when grading exams!

bids.

2. Bidders know that rivals with lower signal values are responding in this way. As such, other things equal, they will need to increase their bids in response to the anticipated increase in bids from lower signal holders.

The bidder with the highest signal is not, on average, subject to this first force. Thus, she does not, on average, revise her estimate of the true value. Nevertheless, she raises her bid in response to this second factor, the "domino" effect of bidders with lower signals raising their bids.

These strategic considerations hold for a wide variety of public information signals (Milgrom and Weber, 1982). There are, however, several methodological advantages to using x_L. First, the RNNE bid function can be readily solved for x_L, provided low signal holders are restricted to bidding x_L, so that the experimenter continues to have a benchmark model of fully rational behavior against which to compare actual bidding. Second, x_L provides a substantial dose of public information about x_o (it cuts expected profit in half), while still maintaining an interesting auction. As such it should have a substantial impact on prices, regardless of any inherent noise in behavior. Finally, the experimenter can always implement finer, more subtle probes of public information after seeing what happens with such a strong treatment effect.[11]

KL (1986) found that in auctions with small numbers of bidders (3 - 4), public information resulted in statistically significant increases in revenue that averaged 38% of the RNNE model's

[11]KL(1986) did not restrict low signal holders to bidding x_L, failing to recognize that without this restriction there is no pure strategy Nash equilibrium, but a much more complicated mixed strategy equilibrium so that their benchmark calculations are incorrect. However, the correct benchmark yields an even higher increase in revenue from announcing x_L so that the conclusions reached regarding public information receive even stronger support with the correct benchmark (Campbell, Kagel and Levin, 2000).

prediction. However, in auctions with larger numbers of bidders (6 - 7), public information reduced average sellers' revenue by $1.79 per auction, compared to the RNNE model's prediction of an increase of $1.78. KL attribute this reduction in revenue to the presence of a relatively strong winner's curse in auctions with large numbers of bidders. If bidders suffer from a winner's curse, the high bidder consistently overestimates the item's value, so that announcing x_L is likely to result in a downward revision of the most optimistic bidders' estimate. Thus, out of equilibrium, public information introduces a potentially powerful offset to the forces promoting increased bids discussed earlier, and will result in reduced revenue if the winner's curse is strong enough. This hypothesis is confirmed using detailed data from auctions with 6-7 bidders which shows that the RNNE model's prediction of an increase in sellers' revenue is critically dependent on whether or not there was a winner's curse in the corresponding private information market.

KL relate this public information result to anomalous findings from OCS auctions. Mead, Moseidjord and Sorensen (1983, 1984) (MMS) compared rates of return on wildcat and drainage leases in early OCS auctions. A wildcat lease is one for which no positive drilling data are available, so that bidders have symmetric information. On a drainage lease hydrocarbons have been located on an adjacent tract so that there is an asymmetric information structure, with companies who lease the adjacent tracts (neighbors) having superior information to other companies (non-neighbors). The anomaly reported by MMS is that both neighbors and non-neighbors earned a higher rate of return on drainage compared to wildcat leases. In other words, with the asymmetric information structure, even the less informed bidders (non-neighbors) received a higher rate of return on drainage leases than on leases with a symmetric information structure (wildcat tracts). In contrast, a fundamental prediction for models with “insider information” is that less informed bidders will earn smaller informational rents than they would in a corresponding symmetric information structure auction like the wildcat auctions. KL (1986) rationalize the MMS data by arguing that there is a considerable amount of public information associated with drainage tracts[12], and the public information may have corrected for a winner's curse that depressed rates of return on wildcat tracts. Although this is not the only possible explanation for the field data -- the leading alternative explanation is that the lower rate of return on wildcat leases reflects the option value of the proprietary information that will be realized on neighbor tracts if hydrocarbons are found – the KL explanation has the virtue of parsimony and consistency with the experimental data.

B.4 Is the Winner’s Curse a Laboratory Artifact: Limited-Liability for Losses

Results of experiments are often subject to alternative explanations. These alternative explanations typically provide the motivation for subsequent experiments which further refine our understanding of behavior. This section deals with one such alternative explanation and the responses to it.

In the KL (1986) design subjects enjoyed limited-liability as they could not lose more than their starting cash balances. Hansen and Lott (1991) (HL) argued that the overly aggressive bidding reported in KL *may* have been a rational response to this limited-liability rather than a result of the winner’s curse. In a one-shot auction, if a bidder's cash balance is zero, so that they are not liable for *any* losses, it indeed pays to overbid relative to the Nash equilibrium bidding strategy proposed in

[12]See Cooper (1998) for discussion of the extensive spying that goes on between rival companies once drilling starts on a tract and the difficulties involved in keeping drilling results out of the hands of competitors.

section B.1. With downside losses eliminated, the only constraint on more aggressive bidding is the opportunity cost of bidding more than is necessary to win the item. In exchange, higher bids increase the probability of winning the item and making positive profits. The net effect, in the case of zero or small cash balances, is an incentive to bid more than the Nash equilibrium prediction. HL's argument provides a possible alternative explanation to the overly aggressive bidding reported in KL (1986) and in Kagel et al (1989).

Responses to the limited-liability argument have been twofold. First, KL (1991) reevaluated their data in light of HL's arguments, demonstrating that for almost all bidders cash balances were *always* large enough so that it *never* paid to deviate from the Nash equilibrium bidding strategy in a one-shot auction. Second, subsequent empirical work has demonstrated a winner's curse in experimental designs where limited liability for losses could not logically account for overbidding. This provides experimental verification that limited-liability forces do not account for the overly aggressive bidding reported.

KL's design protects against limited-liability problems since bidding x - å insures against all losses and bidders have their own personal estimate of the maximum possible value of the item (min [x + å, $\overline{x}$]). The latter implies that it is never rational, limited-liability or not, to bid above this maximum possible value in a first price auction. Further, cash balances only have to be a fraction of the maximum possible loss for the limited-liability argument to lose its force in a first price auction. For example, KL (1991) report simulations for auctions with 4 or 7 bidders, with å = $30 and cash balances of $4.50 (which 48 out of the 50 bidders always had), for which unilateral deviations from the RNNE bid function were not profitable even when fully accounting for bidders limited liability. Further, limited-liability arguments imply more aggressive bidding in auctions with fewer rather than larger numbers of bidders, just the opposite of what the data shows.[13] As such overbidding in the KL experiment must be explained on some other grounds, such as the judgmental error underlying the winner's curse.

Empirical work on this issue has proceeded on several fronts. Lind and Plott (1991) (LP) replicated KL's results in auctions where bankruptcy problems were almost completely eliminated. One experimental treatment involved conducting private-value auctions where subjects were sure to make money simultaneously with the common-value auctions, thereby guaranteeing a steady cash inflow against which to charge any losses incurred in the common-value auctions. In addition, subjects agreed that if they ended the experiment with a negative cash balance, they would work losses off doing work-study type duties (photocopying, running departmental errands, etc.) at the prevailing market wage rate. A second treatment involved sellers' markets in which bidders tendered offers to sell an item of unknown value. Each bidder was given one item with the option to keep it and collect its value or to sell it. In this auction, all subjects earned positive profits, including the winner, but the winner could suffer an opportunity cost by selling the item for less than its true value.[14] LP's results largely confirm those reported by KL and their associates.

Cox, Dinkin, and Smith (1998) (CDS) conducted auctions using KL's design in which, under one treatment, they reinitialize bidders' cash balances in each auction period, with balances large

[13] The greater the number of rivals, the lower the probability of winning as a result of more aggressive bidding, hence the less likely it is to pay to deviate from the Nash strategy even with limited-liability. See also the calculations reported in Kagel and Richard (2001).

[14] To keep costs down, the seller's auctions were conducted in francs as opposed to dollars. The conversion rate from francs to dollars reduced the cost of the experiment, but reduced the marginal incentives for equilibrium behavior as well. There is no free lunch in designing experiments; gains on one dimension are usually offset by losses in other dimensions.

enough that subjects could not go bankrupt even if biding well above their signal values. In contrast to this unlimited liability treatment, their other treatments employed procedures where cash balances fluctuated, bidders could go bankrupt, and in some treatments, bidders with negative cash balances were permitted to continue to bid. Using data for all treatments and all levels of bidder experience, CDS find no significant differences in individual bid patterns in the unlimited liability treatment, contrary to HL's argument. Further, restricting their analysis to experiments with experienced subjects, and dropping data from an entire experiment if even one subject adopted a pattern of high bids when having a negative cash balance, CDS find that the unlimited liability treatment significantly *increased* individual bids, the exact opposite of HL's hypothesis. This seemly bizarre outcome is, however, consistent with KL's (1991) argument that in a multi-auction setting, where cash balances carry over from one auction to the next, there is a potentially powerful offset to any limited-liability forces present in a one-shot auction: Overly aggressive bidding due to low cash balances may be offset by the risk that such bids will result in bankruptcy, thereby preventing participation in later auctions with their positive expected profit opportunities. Unfortunately, it is consistent with the artifactual explanation that because subjects were paid off in only a few of the unlimited liability auctions (in order to keeps costs to a manageable level), subjects treated these auctions differently then those in which they were paid as a result of each outcome.[15]

B.5 Summing Up

Even after allowing for some learning as a result of feedback regarding past auction outcomes a

[15]For a completely different approach to the limited liability problem see Avery and Kagel (1997).

strong winner's curse is reported for inexperienced bidders in sealed-bid common-value auctions. High bidders earn negative average profits and consistently bid above the expected value of the item conditional on having the high signal value. Further, this is not the result of a handful of overly aggressive bidders but applies rather broadly across the sample population. Similar results are reported in low bid wins, supply auctions with both student subjects and professional bidders drawn from the commercial construction industry (Dyer, Kagel and Levin, 1989). Arguments that these results can be accounted for on the basis of limited-liability for losses have been shown to be incorrect. Further, a clever experiment by Holt and Sherman (1994) (also see Avery and Kagel, 1997) is able to rule out the idea the winner’s curse is a result of an added thrill, or extra utility, from winning.

Note that the overbidding associated with the winner’s curse is not simply a matter of miscalibrated bidders, but is associated with fundamental breakdowns of the comparative static predictions of the rational bidding model: With a winner's curse public information reduces revenue, contrary to the theory’s prediction, as the additional information helps high bidders to correct for overly optimistic estimates of the item's worth. In second-price sealed-bid auctions increased numbers of bidders produces no change in bidding, contrary to the robust Nash equilibrium prediction that bids will decrease (Kagel, Levin and Harstad, 1995).

We are still left with the puzzle, first expressed by Lind and Plott, that although many experiments report a clear winner’s curse (negative profits), comparing between the symmetric RNNE and totally naive bidding models offered in the literature (all players treat their signals as if they are private values and go on to bid as if in a private-value auction; KL, 1986), bidding is closer to the RNNE. One promising explanation for this phenomenon appears to be that bidders are cursed to different degrees. That is, agents may make partial, but incomplete, adjustments for the adverse selection effect associated with common-value auctions, with the perfectly rational and perfectly naive bidding models being polar cases. Depending on the extent to which players are “cursed” they may suffer losses, but bidding can, in fact, still be closer to the symmetric RNNE bidding model than the totally naive bidding model. (See Eyster and Rabin, 2000, for a formal model of this sort.)

C. English Auctions and First-Price Auctions with Insider Information

My colleagues and I have also studied English auctions and first-price auctions with insider information (one bidder knows the value of the item with certainty and this is common knowledge). These experiments were initially motivated by efforts to identify institutional structures that would eliminate, or mitigate, the winner’s curse for inexperienced bidders. The experiments also investigate the comparative static properties of Nash equilibrium bidding models for very experienced bidders. In both institutional settings the winner's curse is alive and well for inexperienced bidders, although it is clearly less severe in English than in first price auctions. In contrast, comparative static predictions of the Nash equilibrium bidding model are largely satisfied for more experienced bidders. However, in the case of English auctions, the information processing mechanism that the Nash bidding model specifies is not satisfied. Rather, bidders follow a relatively simple rule of thumb that results in almost identical prices and allocations as the Nash model’s predictions for the distribution of signal values employed in the experiment. In the insider information auctions less informed bidders (outsiders) have some proprietary information (i.e., the insider knows the value of the item with certainty, but does not know the outsiders’ signals). This results in marked differences in predicted outcomes compared to the standard insider information model in which the insider has a double informational advantage - she knows the value of the item and the signals the outsiders have (Wilson, 1967, Weverberg, 1979, Englebrecht-Wiggans, Milgrom, and Weber, 1983, Hendricks, Porter, and Wilson, 1994). Most notably, in our model the existence of an insider generates higher average revenue than in auctions with a symmetric information structure, a prediction that is satisfied in the data for experienced bidders. In contrast, in the double informational advantage model the existence of an insider reduces average revenue.

C. 1 English Auctions

Levin, Kagel, and Richard (1996) (LKR) implement an irrevocable exit, ascending-price (English) auction. Prices start at $\underline{x}$, 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 once they have dropped out. 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, also 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 their drop-out prices.

For signals in region 2, in a symmetric RNNE the bidder with the low signal value (x_L) drops out of the auction once the price reaches his signal value.[16] The price at which the low bidder drops out of the auction reveals his signal value to the remaining bidders. Thus, the public information, x_L, that was

[16]The intuition is roughly as follows: Given symmetry, the low signal holder knows that those remaining in the auction have higher signal values. But the low signal holder can't profit from this additional information since it is only revealed once the price is greater than these remaining signal values; i. e., price is already greater than the expected value of the item to the low signal holder.

provided by the experimenters in KL (1986) is provided endogenously here (at least in theory) by the first drop-out price. Given the uniform distribution of signal values around x_o, in a symmetric equilibrium, for any remaining bidder j, $(x_L + x_j)/2$ provides a sufficient statistic for x_o *conditional* on x_j being the highest signal, so that drop out prices other than x_L contain no additional information and should be ignored. This sufficient statistic is the equilibrium drop out price for j (d_j) in the symmetric RNNE

(3) $\quad d_j = (x_L + x_j)/2.$

This represents the maximum willingness to pay conditioned on all the information revealed by earlier drop-out prices and conditional on winning. As in first-price auctions with x_L publicly announced, expected profit in the English auction is sharply reduced (by about a half) compared to first-price auctions with strictly private information (as long as $n > 2$). As such, in equilibrium, the English auction is predicted to significantly raise average sellers' revenue compared to first-price sealed-bid auctions.

The key difference between the English auction and a first-price sealed-bid auction with x_L publicly announced is that in the English auction information dissemination is endogenous, rather than exogenous. 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. As such we would expect the information dissemination process to be nosier than with x_L publicly announced. Nevertheless, if bidders are able to correctly recognize and incorporate the public information inherent in other bidders' drop out prices we would predict that: (i) For inexperienced bidders, *contrary* to the Nash equilibrium bidding model's prediction, English auctions will reduce average sellers' revenue compared to first price sealed-bid auctions, as losses will be sharply reduced, or even be eliminated, on average, in the English auctions and (ii) For more experienced bidders, where negative average profits have been largely eliminated in the sealed-bid auctions, the English auctions will raise average revenue, as the theory predicts. The second prediction is the standard, equilibrium prediction. The first prediction follows directly from our experience with first-price auctions with x_L publicly announced.

Table 3 shows averages of predicted and actual changes in revenue between English and first-price auctions for inexperienced bidders, as well as averages of predicted and actual profit, with the results classified by numbers of bidders and å (t-statistics are reported in

[Insert Table 3 here]

parentheses).[17] 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 all values of å with n = 4 and for å = $12 with n = 7.[18] However, for these inexperienced bidders, with the exception of n = 4 and å = $24, actual revenue is lower in the English auctions in all cases, with significantly lower average revenue for n = 4 and 7 with å = $6, and with the reduction in revenue barely missing statistical significance (at the 10% level) with n = 7 and å = $12. Further, the revenue increase with n = 4 and å = $24 is statistically insignificant, and is well below the predicted increase.

---

[17] Common-value auctions involve pure surplus transfers so that revenue differences are calculated as: [ð$_E$ - ð$_F$] where ð$_E$ and ð$_F$ correspond to profits in English and first-price auctions, respectively. In this way we have effectively normalized for sampling variability in X_o by subtracting it from the price.

[18] t-tests are conducted for predicted revenue increases to measure the reliability of the prediction for the LKR sample data. One-tailed t-tests are used here 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 in practice there are forces promoting lower revenues in English auctions and we often observe this outcome.

These perverse revenue effects in terms of Nash equilibrium bidding theory are associated with negative average profit in both the first-price and English auctions. The negative average profits reported in Table 3 indicate that inexperienced bidders suffered from a winner's curse in both auction institutions, but that the curse was relatively stronger in the first-price auctions. These results serve to generalize those reported for first-price sealed-bid auctions with x_L publicly announced: Given a relatively strong winner's curse in sealed-bid auctions, public information reduces rather than raises sellers' average revenue. The major difference between the present results and the first-price auctions with x_L publicly announced are (i) here public information is generated endogenously in the form of drop-out prices and (ii) average profits in the English auctions were negative, but with the exogenous release of public information in te first-price auctions they were positive. This last result suggests that information dissemination in the English auction is noisier than with x_L publicly announced.[19]

For more experienced bidders, English auctions are capable of raising average sellers' revenue as the data in Table 4 demonstrate. With n = 4, actual revenue is higher in the English auctions for both values of å, with a statistically significant increase for å = \$18. However, for n = 7, there is essentially no difference in revenue between the first-price and English auctions. The significant increase in revenue in English auctions with n = 4 and å = \$18 is associated with 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. The importance of eliminating the winner's curse for the revenue raising prediction of

[19] To further investigate this question we have conducted some additional sessions with inexperienced bidders in which x_L was publicly announced prior to bidding in the English auction. In auctions with 6 bidders and å = \$12, average profits in the standard English auction (where x_L was not announced) 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; unpublished data).

the theory to hold is reinforced by the absence of any revenue increase with n = 7, in conjunction with the relatively low share of expected profit (21%) that was earned in these first price auctions.

[Insert Table 4]

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 LKR found, however, is that bidders placed weight on their own signal value 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 value, 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 value with the signal values underlying the drop out prices of *all* earlier bidders. LKR attribute the adoption of this signal averaging rule in favor of the Nash 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 averaging rule would still be used with distribution functions where it leads to markedly different outcomes from the Nash equilibrium. In this case, bidders would have more opportunity to recognize and respond to the profit opportunities inherent in abandoning the signal averaging rule.

C.2 Auctions with Insider Information

Kagel and Levin (1999) investigate bidding in first-price sealed-bid auctions with an asymmetric information structure (AIS). The asymmetry is introduced by choosing one bidder at random in each auction period - the insider (I) - to receive a private information signal x equal to x_o and being told that $x = x_o$. Each of the other bidders, the outsiders (Os), receive a private information signal from a uniform distribution on [x_o - å, x_o + å], as in the auctions with a symmetric information structure (SIS). The insider does *not* know the realizations of Os private information signals. Os know that they are Os, that there is a single I who knows x_o, and the way that all other Os got their private signals.

Note that this information structure differs substantially from the "standard" insider information model employed in the economics literature in which the insider has a double informational advantage - I knows x_o and Os only have access to public information about x_o (Englebrecht-Wiggans, Milgrom and Weber, 1983, Hendricks and Porter, 1988). In contrast, in our design Os have some proprietary information, which permits them to earn positive expected profit in equilibrium. In the double informational advantage model Os earn zero expected profit in equilibrium.

This experimental design has a number of interesting comparative static predictions that contrast sharply with the double informational advantage model. First, and foremost, the existence of an insider benefits the seller by increasing expected revenue relative to auctions with an SIS. In contrast, in the double informational advantage model the existence of an insider *unambiguously reduces* sellers' expected revenue.[20] Second, increases in the number of Os results in Is biding more aggressively in our model. In contrast, in the double informational advantage model, Is bidding strategy is unaffected by increases in the number of Os. Finally, both models imply that Is earn substantially larger expected

[20]Although one can readily demonstrate that increased revenue is *not* a general characteristic of AIS auctions in which Os have some proprietary information, it is a natural element in our design and can be found in other AIS structures as well (Campbell and Levin, 2000).

profit than Os (zero profit for Os in the double informational advantage model) and that Is earn higher expected profit, conditional on winning, than in SIS auctions, although the predicted increase in profit is relatively small in our design.

KL (1999) conjecture that for inexperienced bidders the existence of an insider might attenuate the winner's curse. Os in the AIS auctions who win against better informed Is face a stronger adverse selection effect than in SIS auctions. However, it is entirely plausible that the need to hedge against the existence of an insider is more intuitive and transparent than the adverse selection problem resulting from winning against symmetrically informed rivals. Thus, at least for inexperienced bidders, having an insider may actually reduce the severity of the winner's curse. This would be true, for example, if Os view the situation as similar to a lemon's market (Akerlof, 1970), where it seems reasonably clear there is no rampant winner's curse (our culture warns us to beware of used car salesmen). On the other hand, inexperienced subjects may bid higher in order to make up for their informational disadvantage, thus exacerbating the winner's curse.

KL employ two alternative definitions of the winner's curse for Os in the AIS auctions. In the first definition, KL ignore I's bid, and note that Os can expect to earn negative profits just competing against other Os when $\tilde{a}(x)$ is greater than

$$(4)\quad E[x_o \mid x = x_1^{n^o}] = x - \varepsilon \frac{n^o - 1}{n^o + 1}$$

where n^o is the number of Os bidding. Further, if all Os bid according to (4), and Is employ their best response to these bids then Os would earn average *losses* of more than \$1.50 per auction, conditional on winning. As such, bidding above (4) provides a first, very conservative, definition of the winner's curse. The second definition of the winner's curse accounts for Is best responding to Os' bids, and solves for the zero expected profit level for Os. Not surprisingly, this requires a somewhat larger bid factor (reduction of bids relative to private signals) than equation (2) requires for SIS auctions with equal numbers of total bidders.

[Insert Table 5 here]

Table 5 reports results for inexperienced bidders in these auctions. The data clearly indicate that the winner's curse is alive and well for inexperienced Os. Consider auctions with å = $6, which were used to start each session. With n = 4, almost 60% of the high Os' bids were above the conservative measure of the winner's curse (equation 4), so that these bids would have lost money, on average, just competing against other Os. Further, considering the behavior of both Is and Os (the second winner's curse measure), 94% of the high O bids were subject to the winner's curse. With n = 7, there is an even stronger adverse selection effect, with the result that the winner's curse was more pervasive: 100% of the high O bids and 85.2% of <u>all</u> O bids fell prey to the winner's curse, even with no accounting for Is' bids. The net result, in both cases, was large negative profits for Os when they won (-$1.68 per auction with n = 4; -$3.68 with n = 7). Although somewhat diminished in frequency, a strong winner's curse is also reported for higher values of å as Os continued to earn negative profits throughout, with at least 47% of all bids subject to the winner's curse for any value of å (when accounting for both Is' and Os' bids). Finally, regressions comparing bid functions for inexperienced Os in AIS auctions versus inexperienced bidders in SIS auctions show no significant difference between the two treatments. Thus, contrary to KL's original conjecture, the introduction of an insider did not induce significantly less aggressive bidding for inexperienced Os compared to SIS auctions.

Table 6 reports data for super-experienced bidders (subjects who had participated in at least two prior first-price sealed-bid auction sessions). For these bidders the winner's curse has been largely eliminated and the comparative static predictions of the theory are generally satisfied. Is earned significantly greater profits conditional on winning than did Os. For example, with å = $18 and n = 7, Os earned average profits of around $0.50 per auction conditional on winning. In contrast, Is earned around $3.25 per auction, conditional on winning. Further, Os earned substantially lower profits than in corresponding SIS auctions, for which profits averaged around $2.25 per auction. Also, as the theory predicts, Is increased their bids in the face of greater competition from more Os.

[Insert Table 6 here]

Last, but not least, as the theory predicts, for more experienced bidders, auctions with insider information consistently raised average sellers' revenue compared to SIS auctions (Table 7). The intuition underlying this prediction for our model is as follows: The seller would be unambiguously worse off in the AIS auction relative to the SIS auction if I's in the AIS auction won *all* the time while bidding according to the prescribed (AIS) equilibrium. However, I's do not win all the time, and when O's win (with their equilibrium bid) they win with relatively high signal values, yielding more revenue than when I's win. Further, the existence of the insider helps to "protect" the seller's revenue compared to an SIS auction when O's would have won with relatively low signal values in the SIS auction, since in this case I wins and pays more than O would have paid in the SIS auction. The net result is higher revenue for the seller and reduced variance in seller's revenue (holding x_o constant) compared to SIS auctions.[21]

---

[21]In our design the increase in revenue going from SIS to AIS varies with n, with revenue differences increasing starting from low n, reaching a maximum revenue differential for intermediate levels of n, and decreasing thereafter.

The increase in revenue resulting from an insider in our model is counterintuitive for those whose intuition has been honed on the double informational advantage model. This reversal of the double informational advantage model's prediction rests critically on the fact that less informed bidders have some proprietary information. Many "real world" cases are more realistically modeled with Os 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 Is and Os earn economic rents. This potential for insider information to raise average sellers' revenue had not been explicitly recognized in the auction literature prior to this.[22]

[Insert Table 7]

Concluding Remarks: Summary of Empirical Findings From the Laboratory and Policy Implications:

Experimental studies of common-value auctions have been going on for more than fifteen years now, paralleling the profession's interest in the theoretical and practical properties of these auctions. This research has established several facts about behavior relative to the theory.

For inexperienced bidders, Nash equilibrium bidding theory does not predict well. Inexperienced bidders suffer from a winner's curse, earning negative average profits and with relatively large numbers of bidders going bankrupt. Overbidding here represents a fundamental breakdown in the theory resulting in the reversal of a number of important comparative static predictions: Bidding does

[22]These results motivated Campbell and Levin (2000) to further investigate the role of insider information in first-price auctions compared to homogeneous information environments. This paper connects the revenue raising effects of an insider to more general propositions regarding the revenue raising effects of increased bidder information found in Milgrom and Weber (1982).

not decrease in response to increased numbers of bidders in second-price auctions as the theory predicts, and public information about the value of the item reduces, rather than raises, revenue in the presence of a winner's curse. This perverse effect of public information in the presence of a winner's curse extends to the endogenous release of public information in English clock auctions.

Experienced bidders in the lab eventually overcome the worst effects of the winner's curse, rarely bidding above the expected value of the item conditional on winning and earning positive average profits. Super-experienced bidders also satisfy key comparative static predictions of the theory: Release of public information in sealed-bid auctions raise revenue, and English clock auctions raise more revenue than do sealed-bid auctions. Further, average revenue increases in an experimental design where the existence of an informed insider is predicted to raise revenue compared to auctions with symmetrically informed bidders. Nevertheless these super-experienced bidders still earn well below equilibrium profits and, in the overwhelming majority of cases, are bidding not best responding to rivals' bids (they are bidding far more aggressively than they should; Kagel and Richard, 2001).

It is worth noting that these very experienced bidders in the lab have learned how to overcome the worst effects of the winner's curse in an environment with strong information feedback, substantially stronger than is likely to be present in field settings. As such, learning might not proceed as quickly in field settings. Further, there are dynamics of interactions within organizations that may retard adjustment to the winner's curse. These include, (i) payments of large salaries to petroleum geologists to estimate likely reserves, and then having to recognize that these estimates still have a very large variance and are not very precise, (ii) transfers of personnel within the firm and between firms prior to receiving feedback about the profitability of bids, and (iii) gaming that goes on within organizations.[23] Finally, even assuming that the winner's curse will be eliminated in the long run in field settings, it often takes some time before this happens, so this out-of-equilibrium behavior is important in its own right.

The winner's curse extends to a number of other settings as well: bilateral bargaining games (Samuelson and Bazerman, 1985; Ball, Bazerman and Carroll, 1991), blind bid auctions (Forsythe, Isaac, and Palfrey, 1989), markets where quality is endogenously determined (Lynch, Miller, Plott and Porter, 1986, 1991), and voting behavior (the swing voters curse; Feddersen and Pesendorfer, 1998, 1999).[24]

Experimental studies of auction markets have played a significant role in the design and execution of the recent wave of spectrum (air wave rights) auctions carried out in this country and abroad.[25] Auction experiments have served two principle functions in this work: (i) As a "wind tunnel" to test out the auction software, which implements a relatively complicated set of bidding rules (see, for example, Plott, 1997) and (ii) As a test bed against which to compare theory with behavior. In the latter role, a central design element has been to use ascending-price auctions (with price feedback for bidders) to both minimize the presence of the winner's curse and to generate increased revenue in the

[23]A friend of mine in Houston who was a geologist for a major oil company told me that there was such a broad range of legitimate value estimates for most tracts that when the bidding department started reducing bids relative to value estimates to the point that they were winning very few auctions, the geologists simply raised their estimates. (Geologists love to drill and failure to win tracts means they can't drill.)

[24]See Kagel and Levin, in press, for reviews of this work, or better yet, consult the original publications.

[25]Led by the Federal Communications Commission (FCC), the U. S. government has conducted a number of sales to date raising a total of $23.9 billion and selling over 10,000 licenses between July 1994 and July 2000. Even more spectacular, in an auction ending in April 2000, the British government raised 22.5 billion pounds ($35.53 billion) from the sale of "third generation" mobile phone licenses. See Klemperer (2000) and McAfee and McMillan (1996) for reviews and evaluations of these auctions.

absence of a winner's curse, central insights derived from the interaction between common-value auction theory and experiments:

> "An ascending auction ought to remove another common problem with auctions, the "winner's curse." This strikes when a successful bidder discovers too late that his prize is not worth what he paid for it. Some critics of the scale of the bids seem to see the curse at work [in Britain's third generation sales]. Yet the winner's curse is much likelier in sealed-bid auctions, where bidders lack an important piece of information about the value of the asset: the valuations of other, perhaps better-informed, bidders. In an ascending auction, however, that information is clearly revealed." *The Economist*, April 15, 2000 (p. 36)

> "...by allowing bidders to respond to each other bids, [an ascending-price auction] diminishes the winner's curse: that is, the tendency for naive bidders to bid up the price beyond the licenses's actual value, or for shrewd bidders to bid cautiously to avoid over paying." McAfee and McMillan (1996, p. 161)

References

Akerlof, G. 1970. "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," Quarterly Journal of Economics, 89:488-500.

Ashenfelter, O., and D. Genesore. 1992. "Testing for Price Anomalies in Real Estate Auctions," American Economic Review: Papers and Proceedings. 82:501-505.

Avery, C. and J. H. Kagel. 1997, "Second-Price Acutions with Asymmetric Payoffs: An Experimental Investigation," Journal of Economics and Management Strategy, 6: 573-604.

Ball, S. B., M.H. Bazerman, and J.S. Carroll.1991. "An Evaluation of Learning in the Bilateral Winner's Curse," Organizational Behavior and Human Decision Processes, 48:1-22.

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.

Blecherman, B. and C.F. Camerer. 1998. "Is There a Winner's Curse in the Market for Baseball Players?" mimeograph, Brooklyn Polytechnic University, Brooklyn, NY.

Brown, K.C. 1986. "In Search of the Winner's Curse: Comment," Economic Inquiry, 24:513-516.

Capen, E.C., R.V. Clapp, and W.M. Campbell. 1971. "Competitive Bidding in High-Risk Situations," Journal of Petroleum Technology, 23:641-53.

Campbell, C. and D. Levin. 2000. "Can the Seller Benefit from an Insider in Common Value Auctions?" Journal of Economic Theory, 91: 106-120.

_________, Kagel, J. H., and D. Levin. 1999. "The Winner's Curse and Public Information in Common Value Auctions: Reply." American Economic Review, 89, 325-334.

Cassing, J., and R.W. Douglas. 1980. "Implications of the Auction Mechanism in Baseballs' Free Agent Draft," Southern Economic Journal, 47:110-21.

Cooper, Christopher. 1998. "Oil Firms Still Rely on Corporate Spies to be Well-Informed." Wall Street Journal, Dec 7:1&23.

Cox, J.C. and R.M. Isaac. 1984. "In Search of the Winner's Curse," Economic Inquiry, 22:579-92.

_______, and _______. 1986. "In Search of the Winner's Curse: Reply," Economic Inquiry, 24:517-20.

________, S. H. Dinkin, and V.L. Smith. 1998. "Endogenous Entry and Exit in Common Value

Auctions." mimeograph, University of Arizona, Tucson, AZ.

Dessauer, J.P. 1981. Book Publishing, New York: Bowker.

Dyer, D., Kagel, J. H., and D. Levin. 1989. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," Economic Journal, 99:108-15.

Engelbrecht-Wiggans, R., P.R. Milgrom, and R.J. Weber. 1983. "Competitive Bidding and Proprietary Information," Journal of Mathematical Economics, 11:161-69.

Eyster, E. and M. Rabin. 2000. "Cursed Equilibrium," mimeograph, University of California at Berkeley, Berkeley, CA.

Feddersen, T. and W. Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts Under Strategic Voting," American Political Science Review, 92:23-36.

__________ and ____________. 1999. "Elections, Information Aggregation, and Strategic Voting," Proceedings of the National Academy of Science, 96, 10572-10574.

Forsythe, R. and R.M. Isaac., and T.R. Palfrey. 1989. "Theories and Tests of `Blind Bidding' in Sealed-Bid Auctions," Rand Journal of Economics, 20:214-238.

Gick, M.L. and K.J. Holyoak. 1980. "Analogical Problem Solving." Cognitive Psychology, 12: 306-355.

Goeree, J. K. and T. Offerman. 2000. "Efficiency in Auctions with Private and Common Values: An Experimental Study," mimeograph, University of Virginia, Charlottesville, VA.

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.

Hendricks, K., Porter R. H., and B. Boudreau. 1987. "Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969," The Journal of Industrial Economics, 35:517-542.

-----------, ------------, 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.

Kagel, J. H. and D. Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions," American Economic Review, 76:894-920.

_______ and _______. 1991. "The Winner's Curse and Public Information in Common Value Auctions: Reply," American Economic Review, 81:362-69.

________ and ________. 1999. "Common Value Auctions with Insider Information," Econometrica, 67 :1219-1238.

________ and ________. In press. "Bidding in Common Value Auctions: A Survey of Experimental Research." In J. H. Kagel and D. Levin, Common Value Auctions and the Winner's Curse. Princeton: Princeton University Press.

_______, _______, R. Battalio, and D.J. Meyer.1989."First-Price Common Value Auctions: Bidder Behavior and the Winner's Curse," Economic Inquiry, 27:241-58.

___________, 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.

Klemperer, P. 2000. "What Really Matters in Auction Design," mimeograph, Nuffield College, Oxford University, Oxford, UK.

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. L. Smith. 1994. "Equilibrium in Auctions with Entry," American Economic Review, 84: 585-99.

Lind, B., and C.R. Plott. 1991. "The Winner's Curse: Experiments with Buyers and with Sellers," American Economic Review, 81:335-46.

Lorenz, J., and E.L. Dougherty. 1983. "Bonus Bidding and Bottom Lines: Federal Off-shore Oil and Gas," SPE 12024, 58th Annual Fall Technical Conference.

Lynch, M., R.M. Miller, C.R. Plott, and R. Porter. 1986. "Product Quality, Consumer Information and `Lemons' in Experimental Markets," in P.M. Ippolito and D.T. Scheffman, (eds.), Empirical Approaches to Consumer Protection Economics, Washington DC: FTC Bureau of Economics. 251-306.

_______,_______,_______, and _______. 1991. "Product Quality, Informational Efficiency, and Regulations in Experimental Markets," in R. Mark Isaac (ed.), Research in Experimental Economics. Greenwich:CT, JAI Press. 4.

McAfee, R.P., and J. McMillan. 1996. "Analyzing the Airwaves Auction." Journal of Economic Perspectives, 10: 159-176.

Mead, W. J., A. Moseidjord, and P.E. Sorensen. 1983. "The Rate of Return Earned by Leases Under Cash Bonus Bidding in OCS Oil and Gas Leases," Energy Journal, 4:37-52.

_______, _______, and _______. 1984. "Competitive Bidding Under Asymmetrical Information: Behavior and Performance in Gulf of Mexico Drainage Lease Sales, 1954-1969," Review of Economics and Statistics, 66:505-08.

Milgrom, P., and R.J. Weber. 1982. "A Theory of Auctions and Competitive Bidding," Econometrica, 50:1485-527.

Perkins, D.N. and G. Salomon. 1988. "Teaching for Transfer." Educational Leadership, 46: 22-32.

Plott, C. 1997. "Laboratory Experimental Test Beds: Application to the PCS Auction," Journal of Economics and Management Strategy, 6: 605-638.

Roll, R. 1986. "The Hubris Hypothesis of Corporate Takeovers," Journal of Business, 59:197-216.

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, Vol. 3, Greenwich, CT: JAI Press.

Vickrey, W. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders," Journal of Finance, 16: 8-37.

Weverbergh, M. 1979. "Competitive Bidding with Asymmetric Information Reanalyzed," Management Science, 25, 291-94.

Wison, R. 1967. "Competitive Bidding with Asymmetric Information," Management Science, 13: 816-20.

_________ 1977. "A Bidding Model of Perfect Competition," Review of Economic Studies, 44: 511-18.

_______. 1992. "Strategic Analysis of Auctions," in R.J. Aumann and S. Hart, Handbook of Game Theory with Economic Applications, Vol. 1. Amsterdam: Elsevier Science Publishers.

Table 1
Profits and Bidding in First Nine Auctions for Inexperienced Bidders

Experiment	Percent of Auctions With Positive Profits	Average Actual Profits (t-statistic)	Average Predicted Profits Under RNNE (S_M)[a]	Percent of All Bids $b > E[x_0 \mid X = x_{1n}]$	Percent of Auctions Won by High Signal Holder	Percentage of High Bids $b > E[x_0 \mid X = x_{1n}]$	Percentage of Subjects Going Bankrupt[b]
1	0.0	-4.83 (-3.62)**	.72 (.21)	63.4	55.6	100	50.0
2	33.3	-2.19 (-1.66)	2.18 (1.02)	51.9	33.3	88.9	16.7
3	11.1	-6.57 (-2.80)*	1.12 (1.19)	74.6	44.4	88.9	62.5
4	11.1	-2.26 (-3.04)**	.85 (.43)	41.8	55.6	55.6	16.7
5	33.3	-.84 (-1.00)	3.60 (1.29)	48.1	44.4	88.9	50.0
6	22.2	-2.65 (-1.53)	2.55 (1.17)	67.3	66.7	100	33.3
7	11.1	-2.04 (-2.75)*	.57 (.25)	58.5	88.9	66.7	50.0
8	11.1	-1.40 (-2.43)*	1.59 (.34)	51.9	55.6	55.6	16.7
9	44.4	.32 (.30)	2.37 (.76)	35.2	88.6	66.7	16.7
10	0.0	-2.78 (-3.65)**	3.53 (.74)	77.2	66.7	100	20.0
11	11.1	-3.05 (-3.53)**	1.82 (.29)	81.5	55.6	88.9	37.5
Average	17.2	2.57	1.90	59.4	59.6	81.8	41.1

[a] S_M = standard error of mean.
[b] For all auctions.
* statistically significant at the 5 percent level, two-tailed test.
** statistically significant at the 1 percent level, two-tailed test.
From Kagel, Levin, Battalio and Meyer (1989).

Table 2 Profits and Bidding by Experiment and Number of Active Bidders: Private Information Conditions (Profits measured in dollars)					
Auction Series (No. of Periods)	No. of Active Bidders	Average Actual Profit (t-statistical)[a]	Average Profit Under RNNE (standard error of mean)	Percent of Auctions Won by High Signal Holder	Percent of High Bids $b_1 >$ $E[x_o \mid X=x_{1n}]$
6 (31)	3-4	3.73 (2.70)*	9.51 (1.70)	67.7	22.6
2 (18)	4	4.61 (4.35)**	4.99 (1.03)	88.9	0.0
3 small (14)	4	7.53 (2.07)	6.51 (2.65)	78.6	14.3
7 small (19)	4	5.83 (3.35)**	8.56 (2.07)	63.2	10.5
8 small (23)	4	1.70 (1.56)	6.38 (1.21)	82.6	39.1
1 (18)	5	2.89 (3.14)**	5.19 (.86)	72.2	27.8
3 large (11)	5-7	-2.92 (-1.49)	3.64 (.62)	81.8	63.6
7 large (18)	6	1.89 (1.67)	4.70 (1.03)	72.2	22.2
4 (25)	6-7	-.23 (-.15)	4.78 (.92)	69.2	48.0
5 (26)	7	-.41 (-.44)	5.25 (1.03)	42.3	65.4
8 large (14)	7	-2.74 (-2.04)	5.03 (1.40)	78.6	71.4
Small Market Average	3-4	4.32 (5.55)**	7.48 (0.77)	75.2	19.0
Large Market Average	6-7	-0.54 (0.87)	4.82 (0.50)	62.9	53.9

[a] Tests null hypothesis that mean is different from 0.0

* Significant at 5 percent level, 2-tailed t.test.

** Significant at 1 percent level, 2-tailed t-test.

From Kagel and Levin (1986)

	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	
,	(1) Actual	(2) Theoretica l	(3) Difference	(4) Actual	(5) Theoretical	(6) Actual	(7) Theoretical	(8) Actual	(9) Theoretica l	(10) Difference	(11) Actual	(12) Theoretical	(13) Actual	(14) Theoretical
$6	-1.54*	1.54**	-3.08**	-2.13	2.76	-0.58	1.23	-1.98*	0.10	-2.08*	-3.85	0.99	-I.87	0.89
	(0.72)	(0.49)	(0.71)	(0.52)	(0.38)	(0.50)	(0.30)	(0.87)	(0.34)	(0.78)	(0.71)	(0.19)	(0.51)	(0.29)
				[29]		[28]					[18]		[18]	
$12	-0.54	2.76**	-3.30**	-1.32	5.01	-0.78	2.25	-1.95	1.08	-3.03**	-3.75	2.76	-1.80	1.68
	(1.25)	(0.92)	(0.84)	(0.79)	(0.60)	(0.95)	(0.69)	(1.19)	(0.65)	(0.92)	(0.89)	(0.53)	(0.77)	(0.40)
				[41]		[45]					[30]		[43]	
$24	1.09	8.10**	-7.01*	1.20	9.83	0.11	1.73	ND	ND	ND	ND		ND	
	(3.29)	(2.32)	(3.05)	(1.93)	(1.25)	(2.64)	(2.14)							
				[25]		[13]								

TABLE 4

Super-Experienced Bidders: Actual vs. 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	
,	(1) Actual	(2) Theoretical	(3) Difference	(4) Actual	(5) Theoretical	(6) Actual	(7) Theoretical	(8) Actual	(9) Theoretical	(10) Difference	(11) Actual	(12) Theoretical	(13) Actual	(14) Theoretical
$18	2.21*	3.96**	-1.75*	3.37	6.77	1.16	2.82	-0.25	2.85**	-3.10**	0.76	3.86	1.01	1.01
	(0.95)	(0.73)	(0.68)	(0.50)	(0.48)	(0.88)	(0.53)	(0.86)	(0.61)	(0.59)	(0.65)	(0.50)	(0.56)	(0.37)
				[163]		[107]					[75]		[96]	
	1.20	2.98	-1.78	8.45	11.27	7.25	8.29							
$30	(3.10)	(2.30)	(2.19)	(1.28)	(1.34)	(2.76)	(l.93)	ND			ND		ND	
				[31]		[33]								

[a] All values reported in dollars
* 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).

Number of Bidders	,	Outsiders' Bids								Insiders' Bids	
		Average Earnings Conditional on Winning (S_m)	Frequency of Outsiders Winning (raw data)	Frequency of Winner's Curse (raw data)				Average Bid Factor[a] (S_m)	Frequency High Outsider Bid from High Outsider Signal Holder (raw data)	Average Earnings Conditional on Winning (S_m)	Average Bid Factor (S_m)
				Against Outsiders Only		Against Outsiders and Insiders					
				High Outsider Bid	All Bids	High Outsider Bid	All Bids				
4	6	-1.68 (0.93)	70.6% (12/17)	58.8% (10/17)	39.2% (20/51)	94.1% (16/17)	70.6% (36/51)	1.16 (0.62)	52.9% (9/17)	0.71 (0.35)	1.46[b] (0.26)
	12	-1.40 (0.50)*	65.2% (15/23)	39.1% (9/23)	23.2% (16/69)	65.2% (15/23)	47.8% (33/69)	6.00 (0.77)	73.9% (17/23)	2.74 (0.77)*	2.25 (0.35)
	24	-6.56 (3.07)	71.4% (5/7)	28.6% (2/7)	14.3% (3/21)	85.7% (6/7)	57.1% (12/21)	11.61 (2.78)	100% (7/7)	5.05 (3.50)	5.09 (1.27)
7	6	-3.68 (0.61)**	100% (9/9)	100% (9/9)	85.2% (46/54)	100% (9/9)	92.6% (50/54)	-0.61[c] (0.62)	66.7% (6/9)	_____	1.09[b] (0.29)
	12	-2.47 (1.03)*	78.9% (15/19)	89.5% (17/19)	69.7% (78/112)	89.5% (17/19)	79.8% (91/114)	4.85 (1.03)	73.7% (14/19)	1.93 (0.61)**	1.91[b] (0.33)

Table 6: Super-Experienced Bidders
Auctions with Asymmetric Information Structure (AIS)

Number of Bidders	,	Outsiders' Bids						Insider's Bids	
		Average Earnings Conditional on Winning (S_m)	Frequency of Outsiders Winning (raw data)	Frequency of Winner's Curse: Against Outsiders and Insiders (raw data)		Average Bid Factor[a] (S_m)	Frequency High Outsider Bid from High Outsider Signal Holder (raw data)	Average Earnings Conditional on Winning (S_m)	Average Bid Factor (S_m)
				High Outsider Bid	All Bids				
4	12	0.65 (0.43)	53.7% (29/54)	9.3% (5/54)	4.9% (8/162)	10.05 (0.23)	92.6% (50/54)	3.30 (0.23)**	3.60[c] (0.19)
	18	0.87 (0.68)	63.3% (19/30)	3.3% (1/30)	1.1% (1/90)	15.29 (0.26)	93.3% (28/30)	4.13 (0.37)**	5.80[c] (0.50)
	30	3.67 (2.32)	42.1% (8/19)	5.3% (1/19)	3.5% (2/57)	27.04 (0.65)	94.7% (18/19)	7.94 (0.69)**	8.24 (0.61)
7[b]	18	0.52 (0.34)	64.5% (49/76)	22.4% (17/76)	17.2% (77/453)	15.86 (0.26)	86.8% (66/76)	3.24 (0.36)**	4.35 (0.26)
	30	3.90 (3.07)	41.7% (5/12)	16.7% (2/12)	19.4% (14/72)	26.95 (0.85)	83.3% (10/12)	4.95 (0.80)**	5.98 (0.67)

S_m Standard error of the mean.
* Significantly different from 0 at the 5% level, two-tailed t-test.
** Significantly different from 0 at the 1% level, two-tailed t-test.
[a] High bids only.
[b] Includes several auctions with n = 6.
[c] A single outlier bid less than x_o-, was dropped.

From Kagel and Levin (1999).

Table 7

Change in Seller's Revenue: AIS versus SIS Auctions with Super-Experienced Bidders

	n=4			n=7		
	Change in Revenue: AIS *less* SIS Auctions[a] (t-stat)[b]	Mean Profits (σ^2)		Change in Revenue: AIS-SIS Auctions[a] (t-stat)[b]	Mean Profits (σ^2)	
		AIS	SIS		AIS	SIS
ε = \$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)
ε = \$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)