Contents

1

Bidding in Common-Value Auctions: A Survey of Experimental Research

John H. Kagel and Dan Levin[1]

Auctions are of considerable practical and theoretical importance.[1] 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, or 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 (nonlaboratory) 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 (airwave 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.[2] 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.[3] 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 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."[4] 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, in discussions of the winner's curse, many economists, particularly theorists, 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 conditional on the event of winning]. Further, their use of the term typically refers to the study of players who fully account for the 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 skepticism by many economists, as it implies that bidders repeatedly err, violating basic economic notions of rationality and contrary to equilibrium predictions.[5] 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), in professional baseball's free-agency market (Cassing and Douglas 1980; Blecherman and Camerer 1998), in 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 *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. (These results, along with other empirical studies of OCS oil-lease sales, are briefly reviewed in section 6.1 below.)

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, Levin, Battalio, and Meyer 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 we developed. 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. Experimental work on the winner's curse in other settings—in bilateral bargaining games with uncertainty, in "blind-bid" auctions, in two-sided auction markets with a lemon's problem, and in voting—are also reviewed. This is followed by reviews of experiments investigating whether and how bidders learn to overcome the winner's curse and a brief review of field data in relationship to findings from the experiments. The penultimate section of this survey summarizes the empirical findings from the experimental literature, and discusses several theoretical developments motivated by the experimental outcomes and the role this line of research has played in the successful sale of government airwave rights (the spectrum auctions). We conclude with an overview of the rest of the book.

1. 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 4¢, and 400 small paper clips each worth 2¢. Unknown to subjects, each jar had a value of $8.00. (Subjects bid on the value of the commodity, not on 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. Auction group size varied between four and twenty-six. The 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.[6] The average winning bid generated losses in over half of all the 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¢). Number of bidders had no significant effect on individual bids, although the sign was negative (but very small in absolute value).

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.[7] 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 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.

2. 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 $[\underline{x}, \bar{x}]$. In auctions with a symmetric information structure, each bidder is provided with a private information signal, x, drawn from a uniform distribution on $[x_o - \varepsilon, x_o + \varepsilon]$, where ε is known. (Given this informational structure, private signals are affiliated in the sense of Milgrom and Weber 1982.) 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 $[x - \varepsilon, \underline{x}]$, which we refer to as the risk-free strategy, 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 $[x + \varepsilon, \bar{x}]$). 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 (subject identification numbers were, however, suppressed).

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 bid 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. Alternative solutions to the bankruptcy problem are discussed below.

2.1 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. 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. (For data outside this interval, we direct the reader to chapters 3 and 6.) Within region 2, bidders have no endpoint 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]

$$\gamma(x) = x - \varepsilon + h(x) \tag{1}$$

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

This equilibrium bid function combines strategic considerations similar to those involved in first-price 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 $E[x_o \mid X = x_{1n}]$ to be the expected value of the item conditional on having x_{1n}, the highest among n signal values. For signals in region 2,

$$E[x_o \mid X = x_{1n}] = x - [(n - 1)/(n + 1)]\,\varepsilon. \tag{2}$$

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 $E[x_o \mid X = x_{1n}]$ results in negative expected profit. Further, even with zero correlation between bids and signal values, if everyone else bids above $E[x_o \mid X = x_{1n}]$, bidding above $E[x_o \mid X = x_{1n}]$ 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 $E[x_o \mid X = x_{1n}]$, bidding above $E[x_o \mid X = x_{1n}]$ is likely to earn negative expected profit.

Recall that within region 2, $(x - \varepsilon)$ 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 between $(x - \varepsilon)$ 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 (ε): with $n = 4$, the bid factor is 60% of ε, and with $n = 7$, it is 75% of ε. Or put another way, for signals in region 2, the RNNE bid function is well approximated by $\gamma(x) = x - \varepsilon$ (the negative exponential term $h(x)$ in equation [1] approaches zero rapidly as x moves beyond $\underline{x} + \varepsilon$). Thus, the bid factor required just to break even, on average, represents 60% of the total bid factor with $n = 4$, and 75% with $n = 7$.[10] Equation (2) also makes it 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, $2\varepsilon/(n + 1)$. 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 gains are lost by doing so even if this causes her not to win the item, and strictly positive expected gains are 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).

2.2 Some Initial Experimental Results: Inexperienced Bidders

Auctions with inexperienced bidders show a pervasive winner's curse that results in numerous bankruptcies. Table 1.1 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 two to three 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 $E[x_o \mid X = x_{1n}]$. Further, 41% 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, discussed in section 6.2 below).

2.3 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 auc-

TABLE 1.1

Profits and Bidding in First Nine Auctions for Inexperiences Bidders

Experiment	*Percentage of Auctions with Positive Profits*	*Average Actual Profits* (*t-statistic*)	*Average Predicted Profits under RNNE* (S_m)[a]	*Percentage of All Bids* $b > E[x_o \mid X = x_{1n}]$	*Percentage of Auctions Won by High Signal Holder*	*Percentage of High Bids* $b > E[x_o \mid X = x_{1n}]$	*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

Source: Kagel, Levin, Battalio, and Meyer 1989.
[a]S_M = standard error of mean.
[b]For all auctions.
*Statistically significant at the 5% level, 2-tailed test.
**Statistically significant at the 1% level, 2-tailed test.

TABLE 1.2
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)*	*Percentage of Auctions Won by High Signal Holder*	*Percentage 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

Source: Kagel and Levin 1986.
[a]Tests null hypothesis that mean is different from 0.0.
*Significant at the 5% level, 2-tailed *t*-test.
**Significant at the 1% level, 2-tailed *t*-test.

tion experiment). Treatment variables of interest were the number of rival bidders and the effects of public information about x_o on revenue. Table 1.2 reports some of their results. For small groups (auctions with three to four 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 six to seven bidders), profits averaged −\$0.54 per auction, compared to the RNNE prediction of \$4.82.[11] Thus, the profit picture had improved substantially compared to that of the inexperienced bidders discussed in the previous section.

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 of a winner's curse. Thus, although bidders had adjusted reasonably well to the adverse selection problem in auctions with three to four bidders, in auctions with six to seven 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 process is context-specific rather than involving some sort of "theory absorption" that readily generalizes to new environments.[12]

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 that 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 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.[13]

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.[14]

KL (1986) found that in auctions with a small number of bidders (three to four), public information resulted in statistically significant increases in revenue that averaged 38% of the RNNE model's prediction. However, in auctions with a larger number of bidders (six to seven), 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 a large number 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 data from auctions with six to seven 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.[15]

(*Methodological aside:* These experiments were conducted using a dual-market bidding procedure in which subjects first bid in a market with private information and then, before these bids were opened, bid again in a market with x_L announced. This maintains the *ceteris paribus* conditions under which the comparative static predictions of the theory are formulated. This procedure greatly facilitated understanding the basis for the breakdown in the model's predictions.)

KL relate this public information result to anomalous findings from OCS auctions. Mead, Moseidjord, and Sorensen (1983, 1984; hereafter 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 (see section 3.2 and chapter 7 below). KL (1986) rationalize the MMS data by arguing that there is a considerable amount of public information associated with drainage tracts,[16] and the public information may have corrected for a winner's curse that depressed rates of return on wildcat tracts.[17] 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.

KL also note that in markets with x_L announced, average profits were positive in all auction sessions and only slightly less than predicted on average. Further, there were no systematic differences in realized profits relative to predicted profits between auctions with small and large numbers of bidders. These two characteristics suggest that with the large dose of public information involved in announcing x_L, the winner's curse had been almost entirely eliminated.

2.4 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 that 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; hereafter 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 he is not liable for *any* losses, it indeed pays to overbid relative to the Nash equilibrium bidding strategy proposed in section 2.1. 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.[18]

Initial responses to this limited-liability argument were 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, Lind and Plott (1991) replicated KL's experiment in a design that eliminated limited-liability problems, and reproduced KL's primary results. 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 $x - \varepsilon$ insures against all losses, and bidders have their own personal estimate of the maximum possible value of the item (min $[x + \varepsilon, \bar{x}]$). 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 four or seven bidders, with $\varepsilon = \$30$ and cash balances of \$4.50 (which forty-eight out of the fifty 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 greater numbers of bidders, just the opposite of what the data shows.[19] As such, overbidding in the KL experiment must be explained on some other grounds, such as the judgmental error underlying the winner's curse.[20]

Lind and Plott (1991; hereafter 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.[21]

LP's results largely confirm those reported by KL and their associates. First, a winner's curse exists, and although the magnitude and frequency of losses decline with experience, it persists (see Table 1.3). Second, the winner's curse does not result from a few "irrational" bidders, but almost all agents experience the curse. Finally, LP test between alternative models of bidder behavior—comparing the RNNE bidding model with the naive bidding model offered in

TABLE 1.3

Frequency of Losses for Winning Bidders in the Lind and Plott Experiment, in dollars (Francs)

	Experiments				
	1. Buyers	*2. Buyers*	*3. Sellers*	*4. Sellers*	*5. Sellers*
Periods 1–10					
Number of periods of loss	8/10	8/10	5/10	6/10	5/10
Average profit per period	−7.90 (−7.90)	−8.31 (−8.31)	−0.075 (−29.80)	−0.048 (−48.20)	0.001 (1.10)
Average RNNE profit per period*	4.53 (4.53)	5.70 (5.70)	0.177 (70.96)	0.060 (60.44)	0.048 (68.71)
Periods 11–20					
Number of periods of loss	4/10	2/7	3/7	2/10	7/10
Average profit per period	4.57 (4.57)	3.12 (3.12)	0.053 (21.00)	0.032 (31.60)	−0.016 (−22.40
Average RNNE profit per period*	18.47 (18.47)	13.58 (13.58)	0.212 (84.85)	0.048 (48.15)	0.037 (52.68)
Periods 21–30					
Number of periods of loss	—	—	—	3/10	5/10
Average profit per period	—	—	—	0.058 (58.40)	−0.004 (−6.10)
Average RNNE profit per period*	—	—	—	0.104 (104.02)	0.090 (128.91)
Periods 31–40					
Number of periods of loss	—	—	—	2/5	8/10
Average profit per period	—	—	—	0.063 (62.80)	−0.033 (−46.80)
Average RNNE profit per period*	—	—	—	0.065 (65.34)	0.024 (33.72)

Source: Lind and Plott 1991.

*The given RNNE equation is valid only for $\underline{x} + \varepsilon \leq X_i \leq \bar{x} - \varepsilon$. Some of the winners' signals were not in this range, so no predicted RNNE is possible. Therefore, this average includes only periods for which the RNNE predicted profit can be calculated.

KL. Since these models imply different sets of parameter restrictions on a common functional form, LP compute F-statistics comparing the sum of squared errors of the unrestricted model with the restricted model, using the F-statistic as a measure of the relative goodness of fit of the competing models. They find that neither model organizes the data, but that the RNNE provides a better fit. This last result, in conjunction with the negative average profits reported, indicates that there was partial, but incomplete, adjustment to the adverse selection forces in LP's auctions.

Cox, Dinkin, and Smith (1998; hereafter CDS) conducted auctions using KL's design in which, under one treatment, they reinitialized bidders' cash balances in each auction period, with balances large enough that subjects could not go bankrupt even if bidding 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.[22] 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 seemingly bizarre outcome is, however, consistent with KL's (1991) argument that in a multiauction 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.[23]

2.5 Second-Price Sealed-Bid Auctions

LP are puzzled that even though there is a winner's curse in their first-price common-value auctions, the RNNE model provides the best fit to the data: "A major puzzle remains: of the models studied, the best is the risk-neutral Nash-equilibrium model, but that model predicts that the curse will not exist" (LP 1991, 344). They go on to comment that "part of the difficulty with further study stems from the lack of theory about [first price] common-value auctions with risk aversion. . . . If the effect of risk aversion is to raise the bidding function as it does in private [value] auctions, then risk-aversion . . . might resolve the puzzle; but, of course, this remains only a conjecture" (ibid.). Second-price auctions provide an ideal vehicle for exploring this conjecture.

In contrast to first-price auctions, the theory of risk-averse bidders is well understood in second-price auctions with both *symmetric* risk-averse bidders and with *asymmetric* risk-averse bidders.[24] Second-price sealed-bid auctions are similar to first-price auctions with one major difference: the price is determined by the second-highest (and not the winner's) bid. Matthews (1977) and Milgrom and Weber (1982) showed that the bid function $\gamma(x)$ implicitly defined by:

$$E[u(x_o - \gamma(x)) \mid X_i = x = Y_{1n}] = 0 \tag{3}$$

is the symmetric Nash equilibrium (symmetric in both risk preferences and strategy choices), where u(•) is a (common) concave utility function, x_i is the signal of bidder i, and Y_{1n} is the highest signal among $n - 1$ rival bidders.[25]

Under risk neutrality (RN), for signals in region 2, the bid function satisfying (3) is

$$\gamma^{RN}(x) = E\,[x_o \mid X_i = x = Y_{1n}] = x - \varepsilon(n - 2)/n, \tag{4}$$

where E $[x_o \mid X_i = x = Y_{1n}]$ is bidder i's expected value of the item conditional on having the highest signal value and conditional on the next highest rival having the *same* signal value. $\gamma(x)$ (or $\gamma^{RN}(x)$) measures a bidder's *maximum willingness to pay* conditional on winning in a symmetric equilibrium. It is the maximum in that it leaves that bidder just indifferent between winning and paying that price or not winning. In cases where $X_i = x$ is the highest signal, $X_i = x \geq Y_{1n}$ so that the winning bid is greater than the expected value (utility) of the item (since $\gamma(x)$ or $\gamma^{RN}(x)$ condition on $X_i = x = Y_{1n}$). Nevertheless, the winning bidder earns positive expected profit since she is paying the maximum willingness to pay of the second-highest signal holder, which is lower than hers. Losing bidders would earn negative expected profit by raising their bid enough to win the item, since in this event the highest signal holder sets the price (her maximum willingness to pay), which is higher than the expected utility of a deviating loser, who holds a lower signal.

With risk aversion, and symmetry, the Nash equilibrium bid factor is even larger than in (4), resulting in even larger profits than under risk neutrality, as the maximum willingness to pay given that risk is involved is lower than that of a risk-neutral bidder.[26] Even if bidders do not have identical risk preferences, they will bid below (4), provided they are all risk-averse. This result even extends to auctions where the strategy profile is not an equilibrium (Harstad 1991; Kagel, Levin, and Harstad 1995). Corresponding predictions in first-price auctions require symmetric bidders and are conditional on risk attitudes and the underlying distribution of information at bidders' disposal. Further, the symmetric Nash equilibrium bidding model has the important comparative static prediction that individual bids must decrease with more rivals.[27] This last prediction is also robust to assumptions regarding risk preferences and applies to best-response profiles in addition to equilibrium bid functions (Harstad 1991; Kagel, Levin, and Harstad 1995).

Kagel, Levin, and Harstad (1995; hereafter KLH) investigated these comparative static predictions, along with the effects of public information on sellers' revenue. They used moderately experienced bidders that had all participated in at least one prior series of second-price common-value auctions.[28] Table 1.4 reports results from this experiment. As with the first-price auctions, there are

TABLE 1.4

Market Outcomes under Private Information Conditions: Second-Price Auctions

Session (no. of auctions)	*No. of Active Bidders*	*Percentage of Auctions Won by High Signal Holder*	*Median Rank-Order Correlation Coefficient between Bids and Signals*	*Average Profit (standard, error mean)*			*Profit as a Percentage of RNNE*
				Naive Bidding	*Observed*	*RNNE*	
5A (20)	4	70.0	.90	−3.49 (1.83)	3.34** (1.48)	5.49 (2.09)	60.8
6A (22)	4	90.9	.80	−2.71 (1.82)	5.42** (1.73)	6.68 (1.51)	81.1
6B (7)	5	42.9	.70	−5.86 (2.33)	3.24 (4.84)	7.61 (2.51)	42.6
1 (25)	5	40.0	.80	−5.69 (1.04)	1.11 (1.18)	3.45 (.93)	32.2
5B (12)	5–6	58.3	.76	−9.03 (3.22)	−5.84* (2.89)	6.19 (3.91)	−94.5
2 (20)	6–7	70.0	.84	−5.10 (1.01)	−1.01 (1.28)	4.23 (1.14)	−23.9
4 (25)	6–7	44.0	.86	−10.80 (1.16)	−.50 (1.19)	2.04 (.95)	−24.5
3 (23)	7	43.5	.79	−9.44 (1.40)	−3.06** (1.45)	4.72 (1.25)	−64.8

Source: Kagel, Levin, and Harstad 1995.
*Significantly different from zero at 10% level in 2-tailed *t*-test.
**Significantly different from zero at 5% level in 2-tailed *t*-test.

substantial differences in profits conditional on the number of active bidders. In auctions with four or five active bidders, profits are positive, averaging 52.8% of the profits predicted under the symmetric RNNE. Such outcomes are closer to the RNNE benchmark than to a naive bidding model, in which bidders take their signal values to be equal to the value of the item (i.e., bidding as if in a second-price private-value auction). Note, however, that contrary to LP's conjecture, in this case profits below the RNNE benchmark cannot be attributed to risk aversion. Further, as in the first-price auctions experiments, in auctions with six or seven active bidders, average profits were consistently negative, averaging −$2.15 per auction period, compared with predicted profits of $3.97 under the symmetric RNNE benchmark.

Using a fixed-effect regression model, and comparing auctions with four and five bidders to those with six and seven bidders, KLH find no response to increasing numbers of rivals. This directly contradicts the comparative static prediction of the Nash bidding model, regardless of the degree of asymmetry in bidders' risk preferences or in their bidding strategies. However, it is consistent with a naive bidding model in which bidders fail to account for the adverse selection effect inherent in winning the auction.

As in the first-price auction experiments, the effects of public information on sellers' revenue are studied through publicly announcing x_L, the lowest private information signal.[29] In auctions with four or five bidders, announcing x_L raises average revenue by about 16% of the symmetric RNNE model's prediction (however, this increase is not statistically significant at conventional levels). In contrast, in auctions with six or seven bidders, announcing x_L *reduces* average revenue by \$4.00 per auction (which is significant at conventional levels), compared to a predicted *increase* of \$1.80 per auction under the symmetric RNNE. As in the first-price auctions, the ability of public information to increase revenue is conditional on eliminating the worst effects of the winner's curse. In the presence of a strong winner's curse, announcement of x_L serves to offset the high signal holders' overly optimistic estimate of the value of the item, thereby reducing, rather than raising, sellers' revenue.

Avery and Kagel (1997; hereafter AK) study a second-price common-value auction in which x_o, the value of the item, is the *sum* of two independent random variables, x_1 and x_2. In this "shoebox" auction, two bidders each bid on the item, with one signal for each bidder. Signals are drawn from a common uniform distribution with support $[\underline{x}, \overline{x}]$. In the unique symmetric equilibrium of this auction, each bidder bids $b_i = 2x_i$, which is implied by equation (4). There is no ex post regret in this symmetric equilibrium, so that even after learning the results of the auction, no bidder wishes to change his bid: When $x > y$, the winner is guaranteed a profit (in equilibrium), since she earns $x + y$ and pays only $2y$. Further, the loser is guaranteed to lose money if he deviates and bids above $2x$, since he would pay $2x$ and earn only $x + y$. Thus, there is no scope for limited liability for losses to affect bidding in equilibrium. In contrast, in the KLH experiment, the random variation in x_o relative to bidders' signal values means that with limited liability for losses, bidders can bid above (4) hoping to get lucky, while being shielded from all, or part, of the negative expected consequences of such overly aggressive bidding. Further, since the equilibrium bid function with only two bidders depends only on bidder i's signal value (x_i), there is no scope for risk preferences to affect bidding. Thus, this experiment rules out, by design, risk aversion and limited liability for losses as possible explanations for deviations from equilibrium bidding.

AK report strong traces of the winner's curse in this setting, as bids are closer to the unconditional expected value of the item ($x_i + [\overline{x} + \underline{x}]/2$) than to the symmetric Nash equilibrium. Expected-value bidding is a classic example of the winner's curse here. In a world with expected-value bidders, anyone with a signal below the mean value $[\overline{x} + \underline{x}]/2$ cannot make a positive profit in any auction: if a bidder with a signal $x_i < [\overline{x} + \underline{x}]/2$ wins, it implies that $x_j \leq x_i$, so that i earns $(x_i + x_j)$, and pays j's bid ($x_j + [\overline{x} + \underline{x}]/2$), so that i's profit is $(x_i - [\overline{x} + \underline{x}]/2) < 0$, a certain loss. In contrast, with expected-value bidding, bidders with signals greater than $[\overline{x} + \underline{x}]/2$ will make positive profits. Both inexperienced and once-experienced bidders earned negative profits in around two-thirds of all auctions won with signal values less than or equal to $[\bar{x} + \underline{x}]/2$.[30]

AK also investigate the effect of asymmetric payoffs on behavior in this environment. Consider the case where one of the two bidders is known to have a fixed extra payoff advantage, K, should she win the item. For example, in the FCC spectrum (airwave) auctions, it was well known that PacTel had a particular interest in acquiring licenses in Los Angeles and San Francisco (Cramton 1997). Theory suggests that in a second-price auction, (1) the advantaged bidder must win the auction with certainty in any Nash equilibrium, no matter how small K is, and (2) the disadvantaged bidder reduces her bid drastically in response to K, causing a large reduction in expected revenue compared to the symmetric payoff case (Bikhchandani 1988).[31] Essentially what the private-value advantage does is to destroy the symmetric equilibrium of the second-price auction. In the resulting 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. The second-price auction institution is crucial to this outcome—it does not emerge in a first-price auction—as the high bidder does not have to pay what he bids. In the experiment, the effect of the K value advantage on bids and prices was a proportional reduction in losing bids, not the explosive reduction anticipated by the theory. That is, the symmetric model continues to provide a reasonable approximation to behavior given the modest value of K employed ($1.00). This result has important potential public-policy implications, since there are typically some small asymmetries in auctions outside the laboratory, and these are often implemented using a second-price auction format (which includes open outcry, English auctions, close to the format employed in the FCC spectrum auctions). However, given that in virtually all experimental work, behavior is much closer to equilibrium predictions in open outcry, English auctions compared to sealed-bid auctions (Kagel, Harstad, and Levin 1987; Levin, Kagel, and Richard 1996; Kagel and Levin 1999), there is a clear need to explore asymmetries of this sort in English auctions and FCC multiple-round auctions before relying too heavily on the conclusion that small asymmetries do not matter very much in practice.

2.6 Group versus Individual Bids

Cox and Hayne (1998; hereafter CH) explore possible differences in bidding strategies between groups of bidders and individual bidders. The experiment is motivated by the fact that in many market settings, bids are made by groups of individuals in consultation with one another rather than by individuals acting on their own. Psychological research on group versus individual decision-making identifies classes of decision-making problems in which groups reduce judgmental errors to which individuals fall prey, but other types of problems where the opposite result holds (see the many references in CH). Further, it is not clear, a priori, into which of these two categories the winner's curse falls.

Using KL's experimental design, CH explore two different types of environment: one in which bidders in both the group and individual treatments receive one signal value, and a second in which they receive multiple signal values (each bidder in a group of five receives one independently drawn signal value, whereas in the individual treatment each bidder receives five independently drawn signals).[32] CH report that for inexperienced bidders, the winner's curse is alive and well regardless of treatment conditions. They also explore once- and twice-experienced bidders in some detail. In markets of size 7 (seven individual bidders or seven groups of five decision-makers) and 1 signal value, estimated bid functions in region 2 have slopes that are not significantly different from 1.0 with respect to their own signal value, and have bid factors (intercepts) that do not differ significantly from the minimum bid factor required to avoid the winner's curse. That is, neither individuals nor groups commit the winner's curse; both just about break even, failing to come anywhere close to earning the positive profits predicted in equilibrium. Although CH provide no direct comparisons of group and individual bid functions, differences in bid factors appear to be too small for there to be any statistically significant differences in bid patterns for the one signal case.[33]

In auctions with multiple signal values, estimated bid functions employ the midpoint of the winning bidders' signal values. However, this statistic ignores the information content inherent in the spread in signal values that bidders receive. The latter conveys considerable additional information not captured in using the midpoint of the signal values. To take an extreme example, suppose a bidder receives five signals. In one case, all five signals are the same, which is really no better than receiving one signal in identifying the expected value of the item. Alternatively, suppose that the midpoint of the five signals is the same as in the first case, but the spread between the maximum and minimum signal values is equal to 2ε. This set of signals identifies the value of the item precisely.

Thus, the underlying bid function CH estimate is incorrectly specified. Nevertheless, the results reported are suggestive, since slopes of the bid function with respect to this midpoint signal value are significantly less than 1 and the bid factors are consistently positive in sign. That is, in region 2, the estimated bid functions suggest that (1) with multiple signals, both groups and individuals bid less, other things equal, the higher the absolute signal value, which makes little sense (and has not been reported in previous studies of individual bidders receiving only one signal), and (2) there is no bid discount, on average, with lower signal values, in response to either residual item-valuation uncertainty or strategic considerations. Further, the positive bid factors are four to five times larger in the group bid functions, and significantly greater than zero for twice-experienced bidders, suggesting that groups handle the multiple signal value case much worse than individuals do. It will be interesting to see if these observations continue to hold up once there is a better specification for the benchmark bid function and, if so, to determine why groups mishandle the greater information content of the multiple signal case.

2.7 Summing Up

Even after allowing for some learning as a result of feedback regarding past auction outcomes, a 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; see section 6.2 below and chapter 11). Arguments that these results can be accounted for on the basis of limited liability for losses have been shown to be incorrect (KL 1991; LP 1991; CDS 1998; AK 1997).

In the absence of a winner's curse, public information tends to raise revenue, as the theory predicts. However, with a winner's curse, public information reduces revenue, as the additional information helps high bidders to correct for overly optimistic estimates of the item's worth. These results are found in both first- and second-price auctions. Increased numbers of bidders produce no change in bidding in second-price auctions, contrary to the robust Nash equilibrium prediction that bids will decrease. Second-price, "shoebox" auctions of the sort AK conducted also show classic traces of the winner's curse. Finally, there are no differences in bid patterns between individuals and groups when both receive one signal value, but the data suggest that groups bid more aggressively in the multiple-signal case.

We are still left with the puzzle expressed by Lind and Plott: 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, bidding is closer to the RNNE. Experiments in second-price auctions show that these differences between behavior and theory cannot be rationalized by risk aversion, as they can be in private-value auctions (see Kagel 1995a for a survey of private-value auction experiments). Rather, a more promising explanation 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. (In a perfectly naive bidding model, all players treat their signals as if they are private values and go on to bid as if in a private-value auction; see KL 1986 and KLH 1995 for development of naive bidding models for first- and second-price auctions, respectively.) 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 to the totally naive bidding model.

Eyster and Rabin (2000) have recently formalized a model of this sort, employing the concept of a cursed equilibrium: each player correctly predicts the distribution of other players' actions, but underestimates the degree to which these actions are correlated with these other players' signals. The formalization offered has at least two nice characteristics. First, it provides a ready-made, intuitively plausible measure of the extent of the winner's curse that can be applied in a variety of settings. Among other things, this enables one to identify the degree to which bidders must be cursed for negative profits to emerge.[34] Second, analyzing the comparative static properties of the model, one can readily identify predictions that are robust to the presence or absence of a winner's curse, evenly for mildly cursed agents. This in turn can enable experimenters (and those who consume results of experiments) to identify the crucial comparative static treatments that will provide rigorous tests of the theory in its many, related applications.

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

We have also studied bidding in 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; Weverbergh 1979; Engelbrecht-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.

3.1 English Auctions

Levin, Kagel, and Richard (1996; hereafter LKR) implemented an irrevocable-exit, ascending-price (English) auction. Prices start at 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.[35] 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.[36] 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 provided in KL (1986) exogenously is provided here endogenously (at least theoretically) 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

$$d_j = (x_L + x_j)/2. \tag{5}$$

The logic underlying this symmetric equilibrium is similar to the symmetric equilibrium for the second-price auction, as each bidder's dynamic (price-dependent) drop-out price is equal to her maximum willingness to pay conditioned on all the information revealed by earlier drop-out prices, and on winning. Conditioning on winning implies that a bidder's signal is the highest in the sample. Since the first drop-out price reveals x_L, with a uniform distribution of signals the average of the lowest and the highest sample signals is a sufficient statistic for x_o. This holds regardless of any other signal values, and serves as the relevant maximum willingness-to-pay benchmark. As in the first-price and second-price auctions with x_L publicly announced, expected profit in the English auction is sharply reduced (by about half) compared to first- and second-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- and second-price sealed-bid auctions.

The key difference between the English auction and the sealed-bid auctions with x_L publicly announced is that in the English auction, information dissemination is endogenous, rather than exogenous. 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. Thus, we would expect the information-dissemination process to be noisier 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 (1) 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 (2) 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- and second-price auctions with x_L publicly announced.

Table 1.5 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 number of bidders and ε (t-statistics are reported in parentheses).[37] 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 ε with $n = 4$ and for $\varepsilon = \$12$ with $n = 7$.[38] However, for these inexperienced bidders, with the exception of $n = 4$ and $\varepsilon = \$24$, actual revenue is lower in the English auctions in all cases, with significantly lower average revenue for $n = 4$ and 7 with $\varepsilon = \$6$, and with the reduction in revenue barely missing statistical significance (at the 10% level) with $n = 7$ and $\varepsilon = \$12$. Further, the revenue increase with $n = 4$ and $\varepsilon = \$24$ is statistically insignificant, and is well below the predicted increase.

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 5 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 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 differences between the present results and those with x_L publicly announced are (1) here public information is generated endogenously in the form of drop-out prices, and (2) average profits in the English auctions were negative, but with the exogenous release of public information they were positive. This last result suggests that information dissemination in the English auction is noisier than with x_L publicly announced.[39]

For more experienced bidders, English auctions are capable of raising aver-

TABLE 1.5

Inexperienced Bidders: Actual versus Theoretical Revenue Changes and Profit Levels in English versus First-Price Auctions (in dollars)

	n = 4								
	Average Change in Revenue: English Less First-Price (standard error)			*Average Profit (standard error)*					
				First-Price			*English*		
	Actual	*Theoretical*	*Difference*	*Actual*	*Theoretical*	*No. of Auctions*	*Actual*	*Theoretical*	*No. of Auctions*
ε = \$6	−1.54* (0.72)	1.54** (0.49)	−3.08** (0.71)	−2.13 (0.52)	2.76 (0.38)	29	−0.58 (0.50)	1.23 (0.30)	28
ε = \$12	−0.54 (1.25)	2.76** (0.92)	−3.30** (0.84)	−1.32 (0.79)	5.01 (0.60)	41	−0.78 (0.95)	2.25 (0.69)	45
ε = \$24	1.09 (3.29)	8.10** (2.32)	−7.01* (3.05)	1.20 (1.93)	9.83 (1.25)	25	0.11 (2.64)	1.73 (2.14)	13

Source: Levin, Kagel, and Richard 1996.
*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.

age sellers' revenue, as the data in Table 1.6 demonstrate. With $n = 4$, actual revenue is higher in the English auctions for both values of ε, with a statistically significant increase for ε = \$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 ε = \$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 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.[40]

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 (see chapter 6 for details). LKR explain the adoption of this signal-averaging rule in favor of the Nash rule by noting that (1) it is easy and quite natural to use, and

TABLE 1.5 (continued)

n = 7

Average Change in Revenue: English Less First-Price (standard error)			*Average Profit (standard error)*					
			First-Price			*English*		
Actual	*Theoretical*	*Difference*	*Actual*	*Theoretical*	*No. of Auctions*	*Actual*	*Theoretical*	*No. of Auctions*
−1.98* (0.87)	0.10 (0.34)	−2.08* (0.78)	−3.85 (0.71)	0.99 (0.19)	18	−1.87 (0.51)	0.89 (0.29)	18
−1.95 (1.19)	1.08 (0.65)	−3.03** (0.92)	−3.75 (0.89)	2.76 (0.53)	30	−1.80 (0.77)	1.68 (0.40)	43
— —	— —	— —	— —	— —	—	— —	— —	—

(2) 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 whether the signal-averaging rule would still be used with distribution functions where it leads to outcomes markedly different 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.

3.2 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 to be told that $x = x_o$. Each of the other bidders, the outsiders (*O*s), receive a private information signal from a uniform distribution on $[x_o - \varepsilon, x_o + \varepsilon]$, as in the auctions with a symmetric information structure (SIS). The insider does *not* know the realizations of *O*s' private information signals. *O*s know that they are *O*s, that there is a single *I* who knows x_o, and the way that all other *O*s got their private signals.

Note that this information structure differs substantially from the "standard" insider information structure in which the insider has a double informational advantage—*I* knows x_o and *O*s only have access to public information about x_o (Engelbrecht-Wiggans, Milgrom, and Weber 1983; Hendricks and Porter 1988). In contrast, in our design, *O*s have some proprietary information, which permits

TABLE 1.6
Super-Experienced Bidders: Actual versus Theoretical Revenue Changes and Profit Levels in English versus First-Price Auctions (in dollars)

	n = 4								
	Average Change in Revenue: English Less First-Price (standard error)			*Average Profit (standard error)*					
				First-Price			*English*		
	Actual	*Theoretical*	*Difference*	*Actual*	*Theoretical*	*No. of Auctions*	*Actual*	*Theoretical*	*No. of Auctions*
ε = \$18	2.21* (0.95)	3.96** (0.73)	−1.75* (0.68)	3.37 (0.50)	6.77 (0.48)	163	1.16 (0.88)	2.82 (0.53)	107
ε = \$30	1.20 (3.10)	2.98 (2.30)	−1.78 (2.19)	8.45 (1.28)	11.27 (1.34)	31	7.25 (2.76)	8.29 (1.93)	33

Source: Levin, Kagel, and Richard 1996.
*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.

them to earn positive expected profit in equilibrium. In the double informational advantage model, *O*s 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.[41] Second, increases in the number of *O*s result in *I*s bidding more aggressively in our model. In contrast, in the double informational advantage model, *I*'s bidding strategy is unaffected by increases in the number of *O*s. Finally, both models imply that *I*s earn substantially larger expected profit than *O*s (zero profit for *O*s in the double informational advantage model) and that *I*s 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. *O*s in the AIS auctions who win against better-informed *I*s 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 *O*s view the situation as similar to a lemon's market (Akerlof 1970), where it seems reasonably clear that there is no rampant winner's curse (our culture warns us to beware of used-car salesmen).[42] On the other hand, inexperienced subjects

TABLE 1.6 (continued)

n = 7								
Average Change in Revenue: English Less First-Price (standard error)			*Average Profit (standard error)*					
			First-Price			*English*		
Actual	*Theoretical*	*Difference*	*Actual*	*Theoretical*	*No. of Auctions*	*Actual*	*Theoretical*	*No. of Auctions*
−0.25* (0.86)	2.85** (0.61)	−3.10** (0.59)	0.76 (0.65)	3.86 (0.50)	75	1.01 (0.56)	1.01 (0.37)	96
— —	— —	— —	— —	— —	—	— —	— —	—

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 *O*s in the AIS auctions. In the first definition, KL ignore *I*'s bid, and note that *O*s can expect to earn negative profits just competing against other *O*s when $\gamma(x)$ is greater than

$$E[x_o \mid x_1^{n^o} = x] = x - \frac{n^o - 1}{n^o + 1}\varepsilon, \tag{6}$$

where n^o is the number of *O*s bidding. Further, if all *O*s bid according to equation (6), and *I*s employ their best response to these bids, then *O*s would earn average *losses* of more than \$1.50 per auction, conditional on winning. As such, bidding above (6) provides a first, very conservative definition of the winner's curse. The second definition of the winner's curse accounts for *I*s best responding to *O*s' bids, and solves for the zero expected profit level for *O*s. 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.

Table 1.7 reports results for inexperienced bidders in these auctions. The data clearly indicate that the winner's curse is alive and well for inexperienced *O*s. Consider auctions with ε = \$6, which were used to start each session. With $n = 4$, almost 60% of the high *O*s' bids were above the conservative measure of the winner's curse (equation [6]), so that these bids would have lost money, on average, just competing against other *O*s. Further, considering the behavior of both *I*s and *O*s (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 *all O* bids fell prey to the winner's curse, even with no accounting for *I*s' bids. The net result, in both cases, was

TABLE 1.7

Inexperienced Bidders: Auctions with Asymmetric Information Structure (AIS)

		Outsiders' Bids								*Insiders' Bids*	
				Frequency of Winner's Curse (raw data)							
				Against Outsiders Only		*Against Outsiders and Insiders*					
No. of Bidders	ε	*Average Earnings Conditional on Winning* (S_m)	*Frequency of Outsiders Winning (raw data)*	*High Outsider Bid*	*All Bids*	*High Outsider Bid*	*All Bids*	*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)
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)	68.4% (78/114)	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)

Source: Kagel and Levin 1999.
S_m = Standard error of the mean.
*Significantly different from zero at the 5% level, 2-tailed t-test.
**Significantly different from zero at the 1% level, 2-tailed t-test.
[a]High bids only.
[b]A single outlier bid less than $x_0 - \varepsilon$ was dropped.
[c]In this treatment high *O*s actually bid *above* their signal values, on average.

(continued...)

Index