Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation

Christopher Avery
John F. Kennedy School of Government
Harvard University
Cambridge, MA 02138
chrisa@ksg1.harvard.edu

John H. Kagel
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
kagel+@pitt.edu

A series of two-player, second-price common-value auctions are reported. In symmetric auctions, bidders suffer from a winner's curse. In asymmetric auctions in which one bidder has a private value advantage, the effect on bids and prices is proportional rather than explosive (the prediction of Nash equilibrium bidding theory). Although advantaged bidders are close to making best responses to disadvantaged bidders, the latter bid much more aggressively than in equilibrium, thereby earning negative average profits. Experienced bidders consistently bid closer to the Nash equilibrium than inexperienced bidders, although these adjustments towards equilibrium are small and at times uneven.

1. Introduction

In the vast literature on auctions the effect of asymmetries between bidders has been little studied. Asymmetries are the norm and not the exception in many auctions. It is common for one bidder to be known to have a special interest beyond that of others in winning the auction. Some examples of this are an oil company bidding for a tract near its other properties, the current managers of a company bidding against

We are grateful for comments from Sushil Bikhchandani, Peter Cramton, Glenn Ellison, Drew Fudenberg, David Laibson, Dan Levin, Al Roth, and Bob Wilson. Robert Avery, Alexis Miller, and Wei Lo assisted in conducting the experiments. Christopher Avery was supported by a State Farm Insurance Dissertation Fellowship and by the Stanford Center on Conflict and Negotiation; John Kagel received support from the National Science Foundation.

Journal of Economics & Management Strategy, Volume 6, Number 3, Fall 1997, 573–603

outsiders for its takeover, and the recent FCC broadband MTA auctions.[1]

In the important situation of common values and correlated information, asymmetries can lead to a reversal of the standard revenue ranking different auction rules.[2] For example, Milgrom and Weber (1982) prove that symmetric equilibrium bid functions produce (weakly) greater expected revenue in second-price than in first-price auctions, given the assumptions of affiliated information and risk-neutral bidders. In contrast, Bikhchandani (1988) shows that when one player is known to have a payoff advantage K in a second-price common-value auction then (1) the advantaged bidder must win the auction with certainty in any Nash equilibrium, no matter how small the size of K, and (2) the disadvantaged bidder reduces the bid drastically in response to the addition of K, causing a large decline in expected revenue compared to the symmetric payoff case.[3] This large loss in revenue holds for any value of K, no matter how small, so we say that the addition of the private value K has an *explosive* effect in reducing the theoretical expected revenue for the auction. However, the addition of the same asymmetry has little effect on the bids or the expected revenues in first-price auctions (see the Appendix). In view of recent debates over the format of government auctions (such as the treasury-bill auctions and the FCC auctions), studying the effect of asymmetries on bidding outcomes is of practical as well as theoretical importance.

This experiment studies the effects of asymmetries on bidding in common-value second-price auctions. The question of interest is "Can payoff perturbations have explosive effects on bidding functions as predicted by an equilibrium analysis of second-price auctions?" Past experiments indicate that even though the point predictions of Nash equilibrium bidding models are rarely satisfied, the comparative static implications of the theory are likely to be upheld (see, for example, Kagel and Levin, 1993). Therefore, we study behavior in both symmetric and asymmetric second-price auctions, concentrating on the prediction that disadvantaged bidders will reduce their bids explosively in asymmetric auctions relative to symmetric auctions.

Second-price asymmetric auctions are derived from symmetric

1. For example, in the FCC auctions it was well known that PacTel had a particular interest in acquiring licenses in Los Angeles and San Francisco (Cramton, 1997).

2. Maskin and Riley (1984) demonstrate the effect of asymmetries on first- and second-price auctions for bidders with asymmetric private values.

3. There is an alternative class of equilibria in which the disadvantaged bidder wins each auction with certainty. Bikhchandani dismisses these equilibria because they rely on weakly dominated strategies.

second-price auctions by adding a separate private value component to the value of one bidder, giving that bidder a specialized, known advantage. We find evidence from the empirical bidding functions that disadvantaged bidders reduce their bids in response to the private-value advantage, but that the effect of the private-value advantage is proportional and not explosive. Thus, the expected revenue in second-price auctions is more robust to the addition of asymmetries than is predicted by equilibrium theory. The difference between the bids of advantaged and disadvantaged players is only slightly above the private-value advantage, rather than several times the private value as Nash equilibrium bidding theory predicts. Profits for advantaged bidders are held down from those of equilibrium by overly aggressive bidding by disadvantaged bidders, who lose money on average as a result of their overly aggressive bidding. These losses by disadvantaged bidders can largely be attributed to a winner's curse. A comparison of the inexperienced and experienced cases suggests that bidders are moving in the direction of Nash equilibrium, but very slowly.

The paper proceeds as follows. Section 2 describes the format of the two-bidder, second-price auctions utilized in the experiment and derives the theoretical equilibrium bid functions for the participants. Section 3 explains the experimental design and gives some particulars of the experimental sessions. Section 4 sets out the hypotheses to be tested, and Section 5 gives the main statistical results. Section 6 concludes.

2. The Base Model

We use the following general setting as the base case for experimental study. Two bidders participate in an auction for an object of value $\tilde{V}$, where $\tilde{V}$ is the sum of two independent symmetric random variables, X and Y, each uniform on the range (a, c). One bidder observes X, the other observes Y, and they compete in a sealed-bid second-price auction. Under these conditions a second-price auction is strategically equivalent to an English auction with fixed bidding increments.

One player may be known to have an additional private value of K for the object. When that is the case, we shall refer to that player as the advantaged bidder and to her opponent as the disadvantaged bidder. Then the advantaged bidder values the object at $\tilde{V} + K$, while her opponent continues to value it at $\tilde{V}$, which remains unknown. The value of K is common knowledge, as is the identity of the advantaged player. In the rest of the paper we refer to the case with no additional private value (i.e. $K = 0$) as a *standard second-price auction* and to the case with a positive value for K as a *private-value-advantage auction*.

2.1 Equilibrium Analysis

This subsection reviews the theoretical results that are relevant for our analysis. Although these results are not, for the most part, new, we record them as theorems to distinguish them from our experimental results, which we record as conclusions.

The standard second-price auction produces a class of equilibria as identified by Milgrom (1981) and Levin and Harstad (1986) among others. There is, however, just one equilibrium with symmetric bid functions for the bidders. Let $v(x, y) = E(\tilde{V}|X = x, Y = y) = x + y$. In the unique symmetric equilibrium of the standard second-price auction, both bidders follow the bidding function $B^*(x) = v(x,x) = 2x$, and the bidder with the higher private observation wins the auction.

Theorem 2.1: *There is no* ex post *regret in the symmetric equilibrium of the standard second-price auction. Even after learning the results of the auction, no bidder then wishes to change his bid.*

Proof. Suppose $x > y$. In equilibrium the winning bid, $2x$, falls above the true value $x + y$, which is in turn greater than the price set by the losing bid, $2y$. That is, the winning bidder is guaranteed a profit, while the loser could only lose money by raising his bid. The minimum price at which the loser can win the auction is $2x$, which is greater than the true value, $x + y$. □

The no-regret property occurs regularly in private-value auctions but is rare in common-value auctions. In effect, the loser turns down a price that must be greater than the true value, although that price is never stated formally. The no-regret property is important for experimental purposes, because it means that the symmetric equilibrium is unaffected by risk aversion. It also implies that there is no possibility that limited liability for losses can be responsible for bidding above the Nash equilibrium. In equilibrium, profits are always nonnegative. As in the Vickrey (private-value second-price) auction, each player achieves a profit in every instance in which it is available in equilibrium, and that profit does not depend on the winner's actual bid.

In addition to the symmetric equilibrium of the standard second-price auction, there are a continuum of asymmetric equilibria of the form $B_1(x) = v(x, f(x)) = x + f(x)$, $B_2(y) = y + f^{-1}(y)$. Any pair of functions $(B_1(x), B_2(y))$ constructed from an increasing function $f(\cdot)$ produces a bidding equilibrium. Each such pair has the same property as the symmetric equilibrium that the equilibrium bids are unaffected by risk aversion. It is possible to make a bid that falls above the upper bound for bids of one's opponent in a second-price equilibrium because that bid does not affect the price set in equilibrium.

We shall not be concerned with these asymmetric equilibria in the standard case, but they are vital in the private value advantage case. With $K > 0$, the discrepancy between the values for the two players invalidates most of the equilibria of the standard second-price auction. The following argument is based on related results from Bikhchandani (1988).[4] See Maskin and Riley (1996) for private value revenue comparisons between first-price and second-price asymmetric auctions.

THEOREM 2.2: *In any second-price bidding equilibrium with continuous (and increasing) strategies in the private advantage case, one player must win the auction with probability 1.*

Proof: Suppose not. Then there is an equilibrium with continuous bidding functions $(B_1(x), B_2(y))$ such that each player wins the auction with positive probability. Since the bidding functions are continuous, there must then be values x^*, y^* such that $B_1(x^*) = B_2(y^*)$. Denote $b^* = B_1(x^*)$.

Suppose that $b^* < x^* + y^* + K$. Then player 1 prefers to win the auction if both players bid b^*. The value of the item to him is $x^* + y^* + K$, which is greater than his prospective price of b^*. Further, since the strategies are continuous, player 1 also prefers to win the auction if player 2's observation is in the neighborhood of x^* so that player 2's bid is just above b^*. Therefore, player 1 prefers to increase his bid from $B_1(x^*)$ and the proposed equilibrium fails. A similar argument would imply that player 2 should reduce his bid from $B_2(y^*)$ if $b^* \geq x^* + y^* + K$. □

Ruling out the cases in which the advantaged player loses each auction, and restricting players to bid in the range of their own possible values conditional on their private signals, the natural set of remaining equilibria are those of the form $(B_A(x), B_a(y))$, where $B_A(x) \geq x + c$ is the bidding function for the advantaged player and $B_a(y) \leq y + a + K$ for the disadvantaged player and the bid functions satisfy $B_A(x) \geq B_a(y)$ for each (x,y). The advantaged player bids above the maximum value to her opponent, while the disadvantaged player bids below the minimum value to his opponent conditional on his private observation. These bid functions fulfill the conditions of equilibrium because of the implicit element of price discrimination in the auction. The relatively low bids of the disadvantaged player produce a low effective price for the advantaged player, which in turn give the advantaged player rea-

4. In addition, the same qualitative results hold with a single advantaged bidder and more than one disadvantaged bidder. Thus, our use of only two bidders in the experiment should be viewed as a procedural modification.

son to set a high bid and thus a high effective price for the disadvantaged player. This same result does not hold in a first-price auction, because the winner pays her own bid and thus cannot afford to make exorbitant bids.

While the conditions for a Nash equilibrium require the advantaged bidder to win every auction, there is some flexibility in the actual values of bids in equilibrium. In the most favorable outcome for the advantaged player, the equilibrium bid functions are $B_a(y) = y + a$ and $B_A(x) = x + c$ (or greater). In the least favorable outcome for the advantaged player, assuming that $K < (c - a)/2$, the equilibrium bid functions are $B_a(y) = y + a + K$ and $B_A(x) = x + c + K$. Although this flexibility produces a wider range of possible equilibria for larger values of K, it is important to remember that the equilibrium bid functions for advantaged and disadvantaged bidders are a matched pair. When the disadvantaged player bids more aggressively within the range of possible equilibria, the advantaged player increases her bids to compensate. Thus, the explosive effect of the private-value advantage holds in every equilibrium, and the privately advantaged player always wins the auction.

Since there are so many equilibria in the private-value-advantage auctions, it may seem desirable to allow for a failure of coordination in the players' actions. Thus, we consider the predictions of rationalizability, as well as those of Nash equilibrium. For two-player games, rationalizability is equivalent to solution by iterated strict dominance (Pearce, 1984). However, iterated strict dominance never eliminates any strategies in a second-price auction.

CONCLUSION 2.3 *All strategies are rationalizable in a second-price auction.*

Proof. Consider a strategy $\overline{B}_1$ in which player 1 bids above the maximum possible value for player 2. The best response for player 2 is to select any bid that will lose the auction. Thus, all strategies for player 2 are best responses to $\overline{B}_1$. Now consider a strategy for player 2, $\underline{B}_2$, in which player 2 bids 0. The best response for player 1 is to select any bid which will win the auction, meaning that $\overline{B}_1$ is a best response to $\underline{B}_2$ and vice versa. Both of these strategies will survive in each elimination of strictly dominated strategies, meaning that all other strategies survive as well. □

Since all strategies are rationalizable in any second-price auction, to produce a prediction related to rationalizability, we strengthen the definition to consider the elimination of weakly dominant strategies: we call strategies which satisfy this requirement *weak-dominance ration-*

alizable.[5] For our case, the logic behind Nash equilibrium bidding is sufficiently close to iterated weak dominance that the two requirements nearly coincide.

THEOREM 2.4 *All strategies that select bids in the range of conditional values (on observing one's signal) are weak-dominance rationalizable in the standard second-price auction. Any pair of strategies that are weak-dominance rationalizable in the private-advantage second-price auction cause the advantaged bidder to win the auction.*

Proof. See Appendix. □

In the remainder of the paper, we distinguish between iterated weak dominance and Nash equilibrium by examining predictions suggested by the logic underlying iterated weak dominance.

2.2 REVENUE COMPARISONS

Since the private observations are drawn independently, a standard second-price auction conforms to the requirements of the revenue equivalence theorem. Any auction rules and equilibrium strategies which give the same allocation of the good (and zero expected revenue to a player with the minimum observation) produce the same expected revenue to the auctioneer. In particular, since the symmetric equilibria of first- and second-price auctions allocate the good to the bidder with the highest observation, they must produce the same expected price.[6] The following theorem compares revenue between the standard and private value advantage second-price auctions.

THEOREM 2.5 *The expected revenue from the symmetric equilibrium of the standard second-price auction is* $2(2a + c)/3$. The expected revenue for the private-value-advantage auction is no more than $(3a + c)/2 + K$.

Proof. Calculating the revenue in the second-price symmetric equilibrium, we find that the expected price is $E(\min(B^*(x), B^*(y))) = E(\min(v(x, x), v(y, y))) = E(\min(2x, 2y)) = 2E(\min(x, y)) = 2(2a + c)/3$.

5. The procedure of eliminating weakly dominated strategies may give different results depending on the order of the elimination of weakly dominated strategies. We follow the usual convention of eliminating all weakly dominated strategies at each stage (Fudenberg and Tirole, 1991, Sec. 11.3).

6. Note that only the symmetric equilibrium of a standard second-price auction produces the same expected revenue as that of a first-price auction. The asymmetric equilibria of second-price auctions produce much less revenue than the symmetric outcomes.

In the private-value-advantage case, the second-price equilibrium predetermines the advantaged player as the winner. Then the expected revenue to the auctioneer is simply the expectation of the disadvantaged player's bid. Since $B_a(y) \le y + a + K$, expected revenue is bounded above by $E(y) + a + K = (3a + c)/2 + K$. □

For K near zero, there is a drop in expected revenue of $(c - a)/6$ from the standard case to the private-value-advantage case. Ironically, a small increase in the value of the object to one of the bidders creates a significant decline in the seller's expected revenue in a second-price auction. For larger K, revenue can rise or fall depending on the choice of new equilibrium.

In contrast to these results for the second-price auction, the equilibrium of a first-price auction is relatively unchanged by the addition of a private value for one of the players.

THEOREM 2.6 *In the private-value-advantage case, there is a first-price bidding equilibrium in which each player's bid is within K of the bid for the same observation in a standard first-price auction with no private value component for either player.*

The Appendix derives the first-price bidding equilibrium of the asymmetric case. See Klemperer (1997) for preliminary arguments towards generalizing this result beyond the specific case used in our experiments.

Theorems 2.5 and 2.6 imply that the theoretical revenue of the first-price auction dominates that of the second-price auction for the private-value-advantage case with K near zero. Our experiment is important for testing the relevance of this theoretical prediction. If equilibrium bidding is found in private-value-advantage second-price auctions, then sellers are likely to prefer a first-price auction when one bidder is known to have an unusual interest in the item up for auction.

3. Experimental Design

These experiments focus on comparing behavior for bidders in standard second-price auctions and private-value-advantage auctions. The base model was employed with the individual private observations X and Y independent and uniform on the range (1,4). Our choice of parameters was specifically designed to give bidders reasonable profit opportunities. In the standard case, the expected price is then \$4, giving the bidders an aggregate expected profit of \$1, or 50 cents each.

To study the private-value advantage, we chose $K = 1$, as we felt that a smaller value of K was unlikely to have much bite. Conditional on their observations, an advantaged bidder faced a uniform value on

$(x + 2, x + 5)$ as opposed to $(y + 1, y + 4)$ for his opponent; on average, the increment is 20 percent of the common value. Then there is a range of Nash equilibria in which the advantaged bidder wins the auction and pays a minimum expected price of \$3.50 against the bid function $B_a(x) = x + 1$ and a maximum expected price of \$4.50 against the bid function $B_a(x) = x + 2$. Therefore, the choice $K = 1$ does not yield a sharp price prediction in comparison with the \$4.00 expected price for the standard auction. The theory still makes a number of sharp predictions:

(N1) disadvantaged bidders should be reducing their bids on average compared to the symmetric auctions, with particularly sharp reductions in bids for higher values of x,

(N2) in equilibrium, for any given private information signal, the difference between bids of the advantaged and disadvantaged bidders should be at least \$3.00, and

(N3) the advantaged bidders win all the auctions (or at least a vast majority of them).

We ran two sessions with inexperienced subjects for the standard case and two for the private-value-advantage case, as summarized in Table I. After the sessions with inexperienced bidders, we brought back subsets of the individual groups to compete as experienced bidders under the same conditions. Participants began each session with a balance of \$10 and accrued profits and losses over a series of periods of

TABLE I.
EXPERIMENTAL SESSIONS

Experimental Session	Subject Experience	Number of Players	Number of Auction Periods	Auction Type
1	Inexperienced	12	16	Private advantage
2	Inexperienced	10	16	Private advantage
3	Experienced[a]	10	24	Private advantage
4	Experienced[b]	8	22	Private advantage
5	Inexperienced	12	18	Standard
6	Inexperienced	11[c]	18	Standard
7	Experienced	14	24	Standard

[a] Advantaged/disadvantaged bidder positions maintained for five consecutive periods.
[b] Results of all auctions were publicly posted. Two players who participated in session 3 participated in this session as well.
[c] With an odd number of players, one player each round was unpaired, but that player did not know that until told the results of the auction. Subjects received a fixed payment of \$1.00 for each such auction period.

bidding. Bidders alternated roles as advantaged and disadvantaged in the private-value-advantage bidding sessions. Switching roles was intended to speed up any learning that might be going on in the auctions and to minimize rivalrous bidding that might result from "fairness" considerations.

For each experiment, the subjects drew their own observations from containers filled with random values.[7] In each round, subjects were paired anonymously for a second-price auction based on the values that they drew.[8] In the standard auctions there was no additional private value. In the private-value-advantage auctions it was known by both players that one bidder's value was $\tilde{V} + K$ and that the other's was $\tilde{V}$, where K was known to be 1 and $\tilde{V} + K$ unknown. Although the players did not know the identity of their opponent, they did know the value of K and whether they were advantaged or disadvantaged in a particular auction.

After each round, players learned the bid of their competitor and the value of the object, and in all but one session (see Table I) did not learn the results of any other auction. Each session lasted two hours. Most of the participants were economics undergraduates recruited from advanced and introductory classes. To allow for some initial confusion, we began each session with inexperienced bidders with two practice rounds, and we discarded the first five periods of bidding. We also discarded the first few (four for asymmetric, two for symmetric) periods of bidding with experienced bidders.

The players were matched by prior assignment based on a round-robin format. Any pair of players were matched at most three times in a single session, and the players did not know the matching schedule. Since the players were matched so infrequently and there was no communication between rounds, it seems unlikely that there would be any incentive to alter one's strategy to attempt to create a bidding reputation.[9]

We changed formats slightly with experienced bidders in efforts to speed convergence to equilibrium. In the first experienced private-value-advantage session, bidders maintained an advantaged or disad-

7. These experiments were conducted by hand. Copies of the instructions employed are available on request.

8. Theory predicts the same experimental results with a single advantaged bidder and more than one disadvantaged bidder, since Nash equilibrium predicts that an advantaged bidder must win every auction, regardless of how many disadvantaged bidders participate in the auction. Further, use of more than one disadvantaged bidder increases the costs of the experiments substantially.

9. Subjects were told that they were matched according to a random matching plan designed to minimize the chances of repeated interactions.

vantaged position for sets of five periods in an effort to provide more time to adjust to the different circumstances of advantaged and disadvantaged bidders.[10] In the second experienced private-value-advantage bidders' session we employed a public information format in which players learned the results of the other auctions as well as their own. At the end of each round, the results of all the auctions were written on a blackboard rather than dispensed to each person privately. In addition, these outcomes were segregated so that cases where advantaged bidders won were posted separately from those where disadvantaged bidders won. Our purpose was to speed the transmission of information about the results of the game, since the disadvantaged players won relatively few auctions. Our conjecture was that a public information format would demonstrate that disadvantaged players tended to lose money when winning an auction, thus giving them additional incentive to adjust their bids towards the equilibrium where those players bid low enough to lose every auction.[11] As the analysis below indicates (see especially the support for Conclusion 5.3), these minor differences in treatment conditions had no material effect on behavior, so that we are fully justified in pooling the experienced subject data.

4. Experimental Hypotheses

The focus of the experiment was to test the Nash equilibrium prediction of Bikhchandani. For alternate hypotheses, we selected a number of other predictive models. While these alternate models predict disparate bidding results, Nash equilibrium is the only model that suggests that bidding should be systematically higher for bidders in the standard auction than for disadvantaged bidders in the private-value-advantage case. The other models (with the exception of the rivalrous model) conclude that since the distribution of individual signals and value are the same for symmetric and disadvantaged bidders, their bids should be the same in both cases. The expected-value model makes an explicit point prediction, while the other models merely make predictions for the general relationship between symmetric and asymmetric bids.

10. Player pairing continued to rotate between auction periods, so that the cost of this adjustment was that players could no longer be paired with all of the other participants evenly. To control for any confound this might have introduced, we used the same block design with the experienced symmetric bidders.

11. Even when losing an auction, the players received enough information to figure out their profits had they bid high enough to win. However, it seems clear that not all players will bother with such calculations (see below) so that we thought that observing all of the actual losses and gains might speed up the learning process.

The set of hypothesized models is as follows:

1. *Nash equilibrium:* There will be an explosive effect from the private-value advantage [prediction (N2)]. The advantaged bidder will win the vast majority of the auctions [prediction (N3)]. In comparison to the standard case, the disadvantaged ($K = 0$) bidder will reduce his bids while the advantaged player will increase her bids by more than the \$1 K-value [prediction (N1)].

1a. *Rationalizability and iterated weak dominance:* Rationalizability makes no prediction whatsoever about the bidders' choice of strategies because of the infinite set of strategies available to them. The concept of iterated weak dominance makes almost the same prediction for behavior as the Nash equilibrium requirement for the private-value-advantage case. We distinguish between the two by the logic behind iterated weak dominance, which eliminates some bidding strategies prior to others. As shown by the analysis in the Appendix, the first serious requirement of iterated dominance is that advantaged bidders should bid at least \$8.00 and win every auction with an observation above \$3.00 and that disadvantaged bidders with observations in the range (1.00, 2.00) should lose every auction and bid no more than \$3.00.

As noted above, iterated weak dominance makes no prediction about the comparison between standard and private-value-advantage auctions, because it does not eliminate any strategies in the standard second-price auction.

2. *Expected value:* Bidders will bid the expected value given their signal rather than the Nash equilibrium bid. In the standard case, that is the function $B_{EV}(x) = x + 2.5$. An advantaged bidder would add a dollar for the private-value-advantage case, while the disadvantaged bidder would not adjust at all. That is, the bid function for the disadvantaged bidders is the same as in the standard second-price auction.

Expected value is the classic example of bidders incurring a winner's curse. In the standard second-price auction, in a pure world of expected-value bidders, anyone with a private signal below the average value cannot win money in any auction: a bidder with a signal $x < 2.50$ wins the auction against a bidder with a signal $y \leq x$, getting an asset of value $x + y$ for the price $2.5 + y$, resulting in a certain loss. In contrast, with expected-value bidding, bidders with signals $x > 2.50$ will make positive profits. Thus, bidding according to expected value, particularly for $x \leq 2.50$, provides evidence of the winner's curse even when these bidders have the good fortune of not winning the auction. Judging from the results of earlier experiments (Kagel and Levin, 1986; Kagel et al., 1995), we expected that some form of winner's curse would prevail, at least in early bidding rounds.

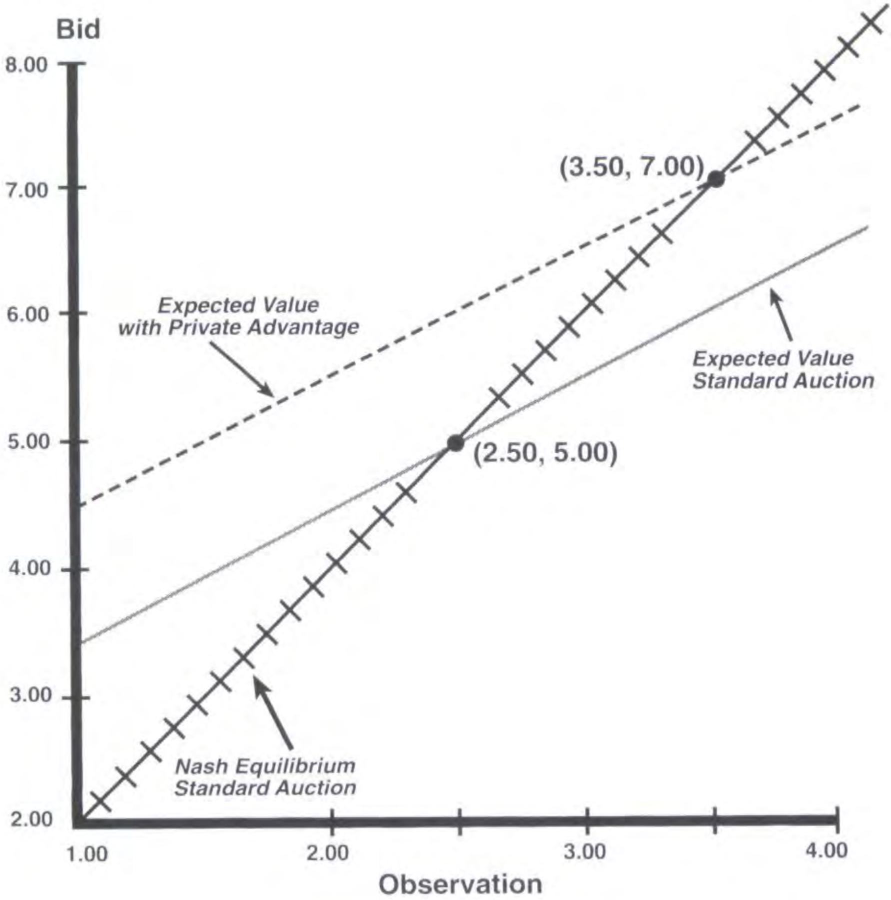

FIGURE 1. POSSIBLE BID FUNCTIONS

Figure 1 depicts the expected value and Nash equilibrium bid functions for the standard case and for the advantaged bidders in the private-value-advantage case. For the standard auction, note that expected-value bids fall below equilibrium bids for the lowest private observations and then cross the equilibrium bidding curve exactly at $x = 2.5$. These relationships offer good possibilities for comparing the models in relation to the data.

3. *Rivalrous bidding:* The strong form of the rivalrous bidding model predicts that disadvantaged bidders will confuse the desire for profits with a desire to win the auction, bidding more aggressively than

bidders in the standard-auction case to overcome their private-value disadvantages. An extreme version of the rivalrous bidding model predicts that the disadvantaged bidders will increase their bids sufficiently to completely offset the private-value advantage. In this case, the disadvantaged bidders will win 50% of the auctions, just as they would in the symmetric value case.

We will also consider a weaker version of the rivalrous bidding model in which disadvantaged bidders bid above equilibrium, not so much to win, but out of a rivalrous effort to deny large profits to the advantaged bidder.

4. *Epsilon-equilibrium:* Bidders will adjust to play approximate best responses to the aggregate set of strategies played by their opponents, but will not necessarily reach equilibrium. The ϵ-equilibrium calculations make no specific prediction about bidding strategies except that they will be close to (i.e. within ϵ of) best responses to rivals' play. Epsilon-equilibrium calculations provide measures of the (expected) cost of failing to respond optimally to the play of others.

In calculating best responses we constructed an empirical distribution of signals and bids from the experimental data and conducted an exhaustive Monte Carlo simulation for each combination of experimental conditions ({Experienced, Inexperienced} × {Standard, Private-Value Advantage}). In effect, we calculated the average payoff for each bid in the sample when matched up with every other bid and x-value in the distribution, including that player's other bids. This procedure is consistent with that of Fudenberg and Levine (1997) for estimating deviations from best responses in normal-form, complete-information games. Our experiment is complicated by uncertainty and the vast number of possible x-values. We weight each empirical observation equally in our simulation, with the result that the x-values which were drawn more frequently in the experiments are also given more weight in the simulation.[12] The empirical bidding function underlying the simulations is equivalent to an explicit mixed strategy that replicates the randomness of the environment faced by bidders.[13]

12. A Monte Carlo simulation which made all x-values equally likely would have to give more weight to the bids corresponding to less common x-values in the experiments.

13. In conducting the simulations it would be inappropriate to create a pure-strategy empirical bid function by averaging the set of bids for each observation, as this would disrupt the probabilities of winning and losing with a given bid. For example, the best response to a certain bid of $5.00 is likely to be much different than the best response to a bid of $8.00 with probability one-half and $2.00 with probability one-half.

5. Experimental Results

We now test the series of predictive models described in the previous section: Nash equilibrium [and the individual hypotheses (N1) to (N3)], weak-dominance rationalizability, expected-value bidding, rivalrous bidding (in two forms) and ϵ-equilibrium. We present our findings in the form of seven conclusions.

We begin by summarizing the bidding outcomes from the symmetric case in Conclusion 5.1 to set a baseline for comparison with the asymmetric auctions. Conclusion 5.2 rejects the strong version of the rivalrous bidding model. Conclusion 5.3 rejects the weak-dominance rationalizability model and hypothesis (N3) of the Nash model. Conclusion 5.4 rejects the remaining properties (N1) and (N2) of the Nash model, while providing support for the expected-value model. Conclusion 5.5 studies the ϵ-equilibrium predictions and shows that advantaged bidders are closer to using optimal strategies than are disadvantaged bidders. Conclusion 5.6 provides some counterevidence against the expected-value model and for prediction (N1) of the Nash model. Conclusion 5.7 rules out the weaker version of the rivalrous bidding model: that disadvantaged bidders increase their bids with the aim of reducing the profits of their opponents.

We now consider the results from the symmetric auctions. While we ran more sessions with the private-value advantage, we actually have a larger data set for standard auctions because bidders face the same game and the same situation in every period in that case. Further, the consistency of the standard auction format may also speed up learning (although we do not specifically test for this), since players do not have to learn how to play from the advantaged and disadvantaged positions. The equilibrium and expected-value predictions are the same in every period of the standard auction: $B^*(x) = 2x$, $B_{EV}(x) = x + 2.5$.

Conclusion 5.1 *There are strong traces of the winner's curse for the standard (control) case. Expected value is a better predictive model for the standard case than the Nash equilibrium prediction. Further, bidders almost invariably lose money, conditional on winning, for signal values of \$2.50 or below, consistent with the presence of a winner's curse. However, what adjustments there are between experienced and inexperienced cases move bidding closer to the predictions of the Nash-equilibrium bidding model.*

The first evidence of the winner's curse is the simple fact that the winning bidder frequently lost money: 39.8% of the auctions with inexperienced bidders and 29.2% of the auctions with experienced bidders resulted in losses, despite the fact that the winning bidder always makes a profit in any Nash equilibrium. As a result, for inexperienced bidders profits averaged 18 cents per player and prices averaged $4.62 per auction period, while for experienced bidders profits averaged 23 cents per player and prices averaged $4.63 per auction period. This contrasts with equilibrium predictions of an average profit of 50 cents per player and an average price of $4.00 in each auction.

While there is an element of randomness in the players' bids, this set of results is much more consistent with the expected-value model than with the Nash model. Both models predict that the player with the higher draw will win the auction. But in contrast to the Nash prediction, expected-value bidding produces an expected profit of 25 cents to each bidder, an average price of $4.50, and losses in 25 percent of the auctions, quite close to the results for the experienced bidders. Further, consistent with the expected-value model's predictions, when bidders won the auction with signal values of $2.50 or less, they usually earned negative profits (66% of the time for inexperienced bidders and 65% of the time for experienced bidders), while with signal values above $2.50 they usually earned positive profits (83% and 76% of the time for inexperienced and experienced bidders, respectively). Average profits show an even more dramatic effect of winning with low compared to high signal values: for experienced bidders these average −36 cents conditional on winning with a signal value of $2.50 or less, compared to +73 cents conditional on winning with a signal value greater then $2.50.

The empirical bid distribution gives further support to the expected-value model, though it demonstrates that expected value does not fully describe the actions of the bidders. Bids fall almost exclusively above the equilibrium prediction for draws below $2.25, and almost exclusively below the equilibrium prediction for draws above $3.25. That bias in the residuals relative to the equilibrium fit indicates that the slope of the empirical bid function is much less than that of the Nash equilibrium, just as implied by expected-value bidding. However, bids also tend to be higher than the expected-value prediction for the highest signal values, though the bias is much less than for the Nash bidding model.

Table II reports error-components estimates of bid functions for the standard auctions which confirm these results. The estimated bid function is

$$B_{it} = \alpha_0 + \alpha_1 x_{it} + \epsilon_{it}, \qquad i = 1, \ldots, N, \quad t = 1, \ldots, T,$$

where the error term $\epsilon_{it} = u_i + v_{it}$ is made up of a subject-specific

TABLE II.
ESTIMATES OF ERROR-COMPONENT BID FUNCTIONS FOR STANDARD AUCTIONS[a]

Bidders			R^2	*F*-Test: Nash	*F*-Test: Expected Value	No. of Observations
Inexperienced,	$B_{it} =$ 2.64	+ $1.13x_{it}$	0.47	59.8	1.55	299
	(0.68)[b]	(0.08)[b]		(<0.01)	(0.21)	
Experienced,	$B_{it} =$ 1.99	+ $1.34x_{it}$	0.75	80.9	22.6	308
	(0.35)[b]	(0.05)[b]		(<0.01)	(<0.01)	

[a] Standard errors in parentheses.
[b] Significantly different from 0 at the 0.01 level.

error term u_i and an auction-period error term v_{it}.[14] This model can accommodate the Nash bidding hypothesis ($\alpha_0 = 0$, $\alpha_1 = 2$) and the expected-value hypothesis ($\alpha_0 = 2.5$, $\alpha_1 = 1$). *F*-tests of these two hypotheses are reported in Table II along with the coefficient estimates.

For inexperienced subjects, the *F*-statistic shows that we cannot reject expected-value bidding. For experienced bidders, the results are sufficiently distinct from the two point predictions ($2x$ for Nash equilibrium, $x + 2.5$ for expected value) to reject them both immediately. Still, the results are closer to expected value than to equilibrium. As with the profit data, what changes there are in going from inexperienced to experienced bidders show movement in the direction of the Nash equilibrium model's prediction, as the slope is increasing and the intercept is decreasing.[15]

The ϵ-equilibrium calculations measure the expected cost of deviating from a best response to the sample population's behavior. Overall, experienced bidders were within 7.1 cents of the optimal payoff against the empirical distribution, compared to 12.3 cents for inexperi-

14. Standard assumptions were employed: $u_i \sim (0, \sigma^2_u)$ and $v_{it} \sim (0, \sigma^2_v)$ where the u_i and the v_{it} are independent of each other and among themselves. Baltagi's (1986) weighted least-squares computational procedure was used to invert the variance-covariance matrix. A fixed-effects error specification generated similar coefficient estimates and standard errors.

15. A regression model restricted to returning bidders shows this same pattern, with the changes in coefficient values being jointly significantly different from zero at just above the 5% level. Thus, the movement towards Nash equilibrium, although small, does not reflect a self-selection effect.

enced bidders.[16] Here too there are important differences in deviations from optimality for signal values of $2.50 or less compared to higher signal values, with the former showing losses of 12.4 cents for experienced subjects and 18.1 cents for inexperienced subjects relative to the expected payoffs from optimal bidding.

We now consider the results for the private-value-advantage auctions. Most of our comparative statics results will rely on comparisons with the standard (control) auctions just discussed.

CONCLUSION 5.2 *There is little support for the strong form of the rivalrous bidding model: Although disadvantaged players consistently bid more than predicted in equilibrium, they bid consistently less then the advantaged bidders. Further, disadvantaged bidders bid the same or less than in the standard auctions in the range of the highest signal values, those signal values where they are most likely to win the auctions.*

There is no evidence that disadvantaged bidders are completely rivalrous, increasing their bids to win 50% of the auctions. In fact, advantaged ($K = 1$) bidders win 62% of the auctions with inexperienced bidders and 71% of auctions with experienced bidders, both of which are significantly greater than 50%.

More generally, the rivalrous bidding model predicts that disadvantaged bidders will bid more than bidders in the standard second-price auctions. Tests of this hypothesis are offered in Table III, where we pool the data from the standard auctions with bids of disadvantaged bidders and estimate the error-components bid function

$$B_{it} = \alpha_0 + \alpha_1 x_{it} + \alpha_2 D x_{it} + \alpha_3 D_i + \epsilon_{it}.$$

In this equation, D is a dummy variable that takes on a value of 1 in the private advantage auctions and 0 in the standard auctions. An F-test of the joint hypothesis that α_2 and α_3 both equal 0 is reported in Table III along with the coefficient estimates.

For inexperienced bidders, the F-statistic indicates no significant differences between disadvantaged bidders and bidders in the standard auctions. This is, of course, inconsistent with the strong form of the rivalrous bidding hypothesis. For experienced bidders, the disadvantaged bidders bid more over lower signal values (the coefficient of the intercept dummy variable D_i is positive and statistically significant), but bid less over higher signal values (the coefficient for the slope

16. Deviations from optimality are, of course, considerably larger in the choice space (bids) then in payoff space (costs), as the former average $0.52 and $0.73 for inexperienced and experienced bidders respectively.

TABLE III.

BIDDING IN PRIVATE-ADVANTAGE AUCTIONS WITH $K = 0$ COMPARED WITH BIDDING IN STANDARD AUCTIONS[a]

Bidders		R^2	F-Test: No Difference[b] (prob F = 1.0)	Implied Bid Difference When $D = 1$: x = \$1.00	Implied Bid Difference When $D = 1$: \$4.00	No. of Observations
Inexperienced,	$B_{it} = 2.59 + 1.15x_{it} - 0.21Dx_{it} + 0.44D_i$	0.45	1.34	0.23	−0.40	419
	$(0.58)^c \quad (0.08)^c \quad (0.15) \quad (0.40)$		(>0.25)			
Experienced,	$B_{it} = 2.15 + 1.34 - .50Dx_{it} + 0.90D_i$	0.70	23.9	0.40	−1.10	483
	$(0.39)^c \quad (0.05)^c \quad (0.08)^d \quad (0.34)^c$		(<.001)			

[a] $D = 1$ in private advantage auctions; otherwise $D = 0$.
[b] Tests joint hypothesis that coefficients for Dx_{it} and D_i are both O.
[c] Significantly different from 0 at 1% level.
[d] Significantly different from 0 at 5% level.

dummy variable Dx_{it} is negative and statistically significant). The average disadvantaged bid is less than the average standard auction bid for signal values greater than \$1.80. Given that disadvantaged bidders rarely win auctions with these lower signal values (experienced disadvantaged bidders *lose* 83% of all auctions for which they have signal values of \$2.50 or less) and given that their bids are less over higher signal values, we do not count this as evidence for the strong form of the rivalrous bidding hypothesis. However, given that bidding is well above equilibrium in the standard auctions with lower signal values, the fact that bids of disadvantaged bidders are yet higher suggests the weaker form of rivalrous bidding—that disadvantaged bidders increase their bids to reduce the profits of their opponents. We discuss this possibility at the end of this section of the paper.

CONCLUSION 5.3 *We reject weak-dominance rationalizability and hypothesis (N3) of the Nash model. Advantaged bidders win more than 50% of the auctions, but far less than 100% as both Nash equilibrium bidding and weak-dominance rationalizability require. Further, there is little evidence for weak-dominance rationalizability, as both advantaged and disadvantaged bidders fail to satisfy the first serious requirements of iterated dominance.*

The advantaged bidders won 70.9% of the auctions for experienced bidders and 62.0% of the auctions for inexperienced bidders. Recall that weak-dominance-rationalizability requires that they win 100% of the auctions. Further, the first serious round of deletion of weakly dominated (rationalizable) strategies requires that disadvantaged bidders with signal values of 2.00 or less never bid above 3.00 (see the Appendix). Nevertheless, this fails to be satisfied 86% of the time with inexperienced bidders and 83% of the time for experienced bidders. In addition, the first serious round of deletion of weakly dominated (rationalizable) strategies also requires that advantaged bidders with signal values of 3.00 or more should never bid below 8.00. Nevertheless, this fails to be satisfied 73% of the time for inexperienced bidders and 71% of the time for experienced bidders.

The next three conclusions relate to the comparative static implications of the Nash equilibrium bidding model resulting from the introduction of asymmetries.

CONCLUSION 5.4 *Contrary to hypotheses (N1) and (N2) of the Nash model, the effect of the private-value advantage on bids and prices is proportional rather than explosive. The effect of the private-value advantage on bids and prices is closer to the predictions of the expected-value model than the Nash bidding model.*

Nash-equilibrium bidding theory requires advantaged bidders to bid \$3.00 more than disadvantaged bidders with the same signal, compared to the expected-value model's prediction of a \$1.00 difference in these bids. Table IV tests this prediction through error-components estimates of the bid function for private-value-advantage auctions. Two alternative specifications are employed. In the first specification we impose the restriction, implied by both the expected-value and Nash models, that the slope of the bid function does not vary as a function of being advantaged or disadvantaged. Instead, only the intercept changes in the equation

$$B_{it} = \alpha_0 + \alpha_1 x_{it} + \alpha_2 DK_{it} + \epsilon_{it},$$

where $DK_{it} = 1$ when $K = 0$ and $DK_{it} = 0$ when $K = 1$. Under this specification the expected-value model predicts $\alpha_0 = 3.5$, $\alpha_1 = 1.0$, and $\alpha_2 = -1.0$. For both inexperienced and experienced bidders, this is very close to the estimated coefficient values. We are unable to reject a null hypothesis of the expected-value model at conventional significance levels for both inexperienced and experienced bidders. For the Nash bidding model, $\alpha_0 = 4.0$, $\alpha_1 = 1.0$, and $\alpha_2 = -3.0$. An F-test decisively rejects these restrictions for both inexperienced and experienced bidders, primarily because the coefficient α_2 is too small.

In the second specification, we drop the restriction that the slope coefficient is the same for advantaged and disadvantaged bidders, giving the equation

$$B_{it} = \alpha_0 + \alpha_1 x_{it} + \alpha_2 DK_{it} + \alpha_3 DKx_{it} + \epsilon_{it},$$

where $DKx_{it} = x_{it}$ when $K = 0$ and $DKx_{it} = 0$ when $K = 1$. For inexperienced bidders, the value of α_3 is close to zero and not significant. For experienced bidders, α_3 is negative and statistically significant, while α_2 remains negative and statistically significant as well. So the difference between advantaged and disadvantaged bids grows with signal values. The minimum difference between these predicted bids is \$0.86 at the lowest signal value and the maximum difference is \$1.52 at the highest signal value, \$4.00. Thus, although there are significant differences in bids between advantaged and disadvantaged bidders, these differences are closer to the prediction of the expected-value model than to the prediction of the Nash model for almost the entire range of signal values. Finally, applying these regression specifications to individual subject data, only 1 of 16 experienced subjects consistently bid closer to the Nash than to the expected-value model's prediction.[17] Therefore,

17. Applying the second regression specification to the individual subject data, we fail to reject the null hypothesis (at the .10 level or better) that $\alpha_3 = 0$ for 12 of 16 subjects. For only one of these twelve subjects was α_2 less than $-\$2$, i.e., closer to the Nash equilibrium prediction. The remaining four bidders all had α's < 0, so that the difference

we conclude that the introduction of asymmetries does not produce anything approaching the explosive effect on bids and prices that the Nash model predicts. In what follows, we try to understand the mechanism behind this outcome.

CONCLUSION 5.5 *The ϵ-equilibrium calculations show that the advantaged bidders are close to making optimal responses. In contrast, the disadvantaged bidders show the strongest deviations from optimal responses: they consistently bid too much, earning negative average profits and sharply reducing the profit opportunities for the advantaged bidders.*

The ϵ-equilibrium calculations show that experienced advantaged bidders were quite close to optimal best responses: they were bidding slightly below the optimum, averaging 5.4 cents below the maximum average return (bids averaged 11 cents below the optimum). In contrast, experienced disadvantaged bidders earned negative average profits (-12.6 cents; -40.0 cents conditional on winning) and were losing an average of 16.1 cents relative to optimal bidding. Further, their bids were a full \$1.31 above the best response against the empirical distribution. For signal values below \$3.00, most bids by disadvantaged players incurred average losses of 20 cents or more in the simulation. Since they win only about 20% of the auctions for such observations, that implies an expected loss of \$1 per auction conditional on winning. In this range of signals, bids by disadvantaged players are commonly \$2 or more above the empirical best-response bid. But winning the auction for such observations may be a sufficiently rare event that there is little learning about the winner's curse. Finally, we see no noticeable differences in bidding with and without public information regarding auction outcomes. Either subjects did not notice or pay attention to the additional information, or they simply chose not to bid less when disadvantaged in spite of occasional losses.

A comparison between experienced and inexperienced bidders in the private-value-advantage auctions shows that experienced advantaged bidders increased their bids by 23 cents and experienced disadvantaged bidders reduced their bids by 16 cents on average. For the

between advantaged and disadvantaged bids grows with higher signal values. We then evaluated these four bid functions to find the value of x for which the slope coefficient declined to less than $-\$2$ (closer to Nash). In one case, there was no such x-value, while the average value in the other three cases was \$2.67, just above the midpoint of the interval for which signal values were drawn.

TABLE IV.
ESTIMATES OF ERROR-COMPONENT BID FUNCTIONS FOR PRIVATE-ADVANTAGE AUCTIONS

Bidders	B_{it}[a]							R^2	F-test: Nash	F-test: Expected Value	No. of Observations
Inexperienced	3.82	+	$0.92x_{it}$	−	$0.73DK_{it}$			0.49	161.8	1.85	242
	(0.80)[b]		(0.07)[b]		(0.13)[b]				(<0.001)	(0.14)	
	3.84	+	$0.91x_{it}$	−	$0.75DK_{it}$	+	$0.01DKx_{it}$	0.49	—	—	242
	(0.82)[c]		(0.10)[b]		(0.38)[c]		(0.14)				
Experienced	3.94	+	$0.98x_{it}$	−	$1.18DK_{it}$			0.069	255.7	1.98	344
	(0.51)[b]		(0.05)[b]		(0.08)[b]				(<0.001)	(0.12)	
	3.68	+	$1.09x_{it}$	−	$0.64DK_{it}$	−	$0.22DKx_{it}$	0.69	—	—	344
	(0.51)[b]		(0.06)[b]		(0.24)[b]		(0.09)[c]				

[a] $DK_{it} = 1$ if $K = 0$, $= 0$ if $K = 1$; $DKx_{it} = x_{it}$ if $K = 0$, $= 0$ if $K = 1$.
[b] Significantly different from 0 at the 1% level or better.
[c] Significantly different from 0 at the 5% level or better.

entire range of x-values, advantaged bidders won the auction more frequently in the experienced case than in the inexperienced case. For $x > 3.00$, the winning percentage increases from 79% to 98%. The increased bidding by experienced advantaged bidders moved them closer to optimal responses (inexperienced bidders averaged 11.2 cents below optimal earnings vs. 5.4 cents for the experienced case). In contrast, even though bidding less, the experienced disadvantaged bidders were worse off than in the inexperienced case (an average loss of 12.6 cents as opposed to 4.1 cents). The increased bidding by advantaged bidders simply provided many fewer profit opportunities than in the inexperienced case for disadvantaged bidders.

Conclusion 5.6 *Contrary to the expected-value model's prediction, there is a clear tendency for experienced disadvantaged bidders to bid less aggressively than in the standard auctions over higher signal values. This provides some weak support for hypothesis (N1) of the Nash model.*

With pure expected-value bidding, the disadvantaged bidders should bid no differently than in standard auctions. As the regression results in Table IV show, a null hypothesis of expected-value bidding for inexperienced bidders would not be rejected at standard significance levels. However, for experienced disadvantaged bidders, there is reduced bidding over higher signal values compared to the standard auctions. That effect is inconsistent with pure expected-value bidding.

Trends within the experienced sessions reinforce the conclusion that experience leads to less aggressive bidding by disadvantaged players with high signals. We divided the experienced data into two subcategories of *early* and *late* in the session. The estimated bid function for the second half of the session gives bids of up to 70 cents more for disadvantaged bidders at low signal values than in the first half of the session, and up to 35 cents less at high signal values than in the first half. In contrast, the standard auction bids seem to increase for all ranges of signals in the last half of the experienced session. One interpretation of the shift by disadvantaged bidders is that they learn to reduce their bids for high signals because that is the one situation where they most frequently win the auction and lose money. In contrast, the primary effect of higher bids by disadvantaged bidders with low signals is to reduce the profits of their advantaged opponents.

The question remains here as to what motivates these higher bids by disadvantaged bidders with lower signal values. One possibility that suggests itself is the weak rivalrous bidding hypothesis as described earlier in Section 4. It may be that disadvantaged bidders are reluctant to bid low enough that they will never win an auction, since such passive play would result in a very uneven distribution of earnings in favor of the advantaged bidder (minimum average profits of \$1.50 in each auction for the advantaged bidder versus \$0 for themselves). At least one player seems to have been motivated by this fact: paraphrasing the remarks of one subject on exiting the auction, "I know I should bid less with $K = 0$, but this just increases the profits of my opponent."

To examine this possibility, we looked at what disadvantaged bidders did following a failure to win the item when a winning bid would have earned negative profits [i.e., $x + y < B_A(x)$]. In general, disadvantaged bidders bid more aggressively following a failure to win the auction even when a winning bid would have lost money. This happened 63% of the time for inexperienced subjects and 56% of the time for experienced subjects.[18] This tendency to bid more aggressively was *not* strongly conditioned on the disadvantaged bidder's signal value, x. Rather, it was strongly conditioned on whether or not the advantaged bidder earned positive or negative profits: Experienced bidders bid more aggressively 69% of the time following the advantaged bidder winning and making a positive profit versus 20% of the time following the advantaged bidder winning and making a negative profit ($Z = 3.63$, $p < .01$), even though a higher winning bid would have resulted in losses in both cases.[19] This would appear to be symptomatic of confusion (failure to think through the implications of winning with a higher bid) rather than rivalrous bidding designed to deny advantaged bidders high earnings.[20] This interpretation is reinforced by the fact that the same pattern prevailed in the standard auctions: Losing bids were followed by relatively more aggressive bids 56% of the time for inexperienced bidders and 63% of the time for experienced bidders even though the losing bidders would have lost money had they won. Here too, bidding was more aggressive when the winner made positive profits compared to when the winner made negative profits (77% vs. 36%, $Z = 4.98$, $p < .01$ for inexperienced bidders; 62%

18. Bidding more or less aggressively between auction periods was measured by comparing the bid less the signal value in auction t with the bid less the signal value in auction $t + 1$. The slope of the aggregate bid function with respect to own signal value is close to 1 under all treatments, which is necessary to justify this calculation. To control for the slope coefficient for own signal value differing from 1, we also measured relative aggressiveness by subtracting the bid implied by the aggregate bid function from the actual bid in each auction period and compared these differences. Our results are robust to this alternative measure. In both cases, calculations are restricted to auctions in which a player was a disadvantaged bidder.

19. Data for inexperienced bidders is too thin to reach conclusions on this score, as there were relatively few cases in which the advantaged bidder won and lost money.

20. Garvin and Kagel (1994) report similar results in first-price common-value auctions.

vs. 52%, $Z = 1.28$, $p < .10$, one-tailed test for experienced bidders). Finally, note that for both disadvantaged bidders and bidders in the standard auctions, winning and losing money typically resulted in less aggressive bidding in the next auction period (72% of the time for disadvantaged bidders; 70% of the time in the standard auctions).

CONCLUSION 5.7 *In both the standard and private-value-advantage auctions, winning the item and losing money tends to result in less aggressive bidding for the next auction period. This tends to correct for the winner's curse. However, in both auctions, failure to win the item when the winner made positive profits promotes more aggressive bidding in the next auction period even in cases where the losing bidder would have succumbed to the winner's curse (lost money) had he won the item. The failure to fully appreciate the consequences of winning in the latter cases tends to perpetuate the winner's curse.*

The fact that the same phenomenon occurs in both the symmetric and asymmetric auctions leads us to reject the weaker form of the rivalrous bidding model in favor of a "confused bidder" model.

6. SUMMARY AND CONCLUSION

In our standard second-price common-value auctions, bidders suffer from a winner's curse, bidding closer to expected value than to the Nash equilibrium. Introduction of a private-value advantage generates changes in bidding. Among inexperienced bidders there are no significant differences between disadvantaged bidders and bidders in standard auctions. Among experienced bidders, disadvantaged bidders bid less than standard-auction bidders over higher signal values, those for which they were most likely to win the auction. However, disadvantaged players still bid substantially more than in equilibrium, earning average profits of −40 cents conditional on winning for experienced bidders. In contrast, experienced advantaged bidders, while not bidding as aggressively as the Nash model requires, are within 5 cents, on average, of maximum possible earnings, given the overly aggressive bidding of the disadvantaged bidders. The net result is that the existence of asymmetric valuations does not produce anything approaching the explosive change in bids, and reduction in revenues, that the Nash bidding model predicts, the primary impediment to this outcome being overly aggressive bidding by disadvantaged bidders.

Experienced subjects consistently bid closer to the Nash equilibrium than inexperienced bidders. But these changes are small and at times unsteady.[21] The introduction of a private-value advantage into the bidding might be expected to speed convergence to equilibrium (and elimination of the winner's curse), since the winner's curse will be exacerbated from the symmetric case unless disadvantaged players reduce their bids. However, a winner's curse remains, as disadvantaged bidders continue to lose money, primarily as a result of overly aggressive bidding with relatively low signal values.

APPENDIX

Proof of Iterated Weak-Dominance Results. Step one: Bidders must bid within the range of possible values conditional on their observations. This restricts the advantaged bidder to bids in the range $(x + 2, x + 5)$ given an observation of x, and the disadvantaged bidder to bids in the range $(y + 1, y + 4)$ given an observation of y.

Step two: Disadvantaged bidders with y less than \$2.00 cannot profit now from winning the auction. They face a price of at least $x + 2$ for an object whose value is $x + y < x + 2$ for $y < 2$. In this instance, they should bid no more than \$3.00 (the minimum advantaged bid retained from step one).

Similarly, advantaged bidders with observations greater than \$3.00 always profit from winning the auction. They face a price of at most $y + 4$ for an object whose value is $x + y + 1 > y + 4$ for $x > 3$. In this instance, they should bid at least \$8.00 (the maximum disadvantaged bid retained from step one).

Otherwise, bids may remain in the ranges $(y + 1, y + 4)$ for disadvantaged bidders with observations in (2,4) and $(x + 2, x + 5)$ for advantaged bidders with observations in (1,3).

Step three: The strategies remaining from step two yield competitive auctions for the cases where $x \in (1,3)$, $y \in (2,4)$. The minimum value for the advantaged bidder in these cases is $x + 3$, since $y \geq 2$, and the maximum value for the disadvantaged bidder in these cases is $y + 3$, since $x \leq 3$.

Therefore, the advantaged bidder must now bid in the range $(x + 3, x + 5)$ for observations in the range (1,3), and the disadvantaged bidder must bid in the range $(y + 1, y + 3)$ for observations in the range (2,4).

21. For example, a within-session analysis of bidding in the standard auctions shows that, on average, players *increased* their bids in the last half of the auction session, actually moving away from equilibrium rather than towards it.

Step four: Consider the same range of observations (1,3) for the advantaged bidder and (2,4) for the disadvantaged bidder. Now advantaged bidders always prefer to win the auction with observations of at least \$2.00, since the maximum price for them is $y + 3$ and the value is $x + y + 1$. Disadvantaged bidders always prefer to lose the auction with observations of \$3.00 or less. As a result, the advantaged bidder should bid at least \$7.00 for an observation in the range (2,3), and disadvantaged bidders should bid no more than \$4.00 for any observation in the range (2,3).

Step five: Now the auction is competitive only if the advantaged bidder has an observation x, in the range (1, 2), bidding in the range $(x + 3, x + 5)$, and the disadvantaged bidder has an observation y, in the range (3,4), bidding in the range $(y + 1, y + 3)$.

The value for an advantaged bidder under these conditions is at least $x + 4$, and the value for a disadvantaged bidder is at most $y + 2$. Thus they should adjust their bidding ranges to $(x + 4, x + 5)$ and $(y + 1, y + 2)$, respectively.

Step six: The maximum bid by a disadvantaged bidder is now $y + 2$, and the minimum bid by an advantaged bidder is now $x + 4$. Therefore, advantaged bidders always wish to win the auction and should bid at least \$6.00, while disadvantaged bidders always wish to lose the auction and should bid at most \$5.00.

At this point, the advantaged bidders always win the auction, the disadvantaged bidders always lose the auction and the maximum expected price is \$4.00 (with the disadvantaged bidders bidding no more than \$3.00 with observations less than \$2.00, no more than \$4.00 with observations less than \$3.00, and no more than \$5.00 for any observation). □

Proof of first-price auction result. We show that there is a first-price auction equilibrium for the asymmetric (perturbed) auction of the experiment with bids within ϵ of the bids for the symmetric game. For simplicity, we specialize to the case where each signal is $U(0,1)$, player 1's value is $x + y + \epsilon$, and player 2's value is $x + y$.

Fix the strategy $b_2(y)$ for player 2, and consider player 1's best response for an arbitrary draw of x. With a bid of $b_1(x)$, player 1 wins the auction at price b_1 whenever player 2's observation is less than $b_2^{-1}(b_1)$. Therefore, player 1 chooses the bid b_1 to solve

$$\max_{b_1} \int_0^{b_2^{-1}(b_1)} \left(x + y + \epsilon - b_1\right) f(y)\, dy$$

$$\max_{b_1} \frac{y^2}{2} + y\left[x + \epsilon - b_1\right]\Bigg|_0^{b_2^{-1}(b_1)}$$

We look for an equilibrium with linear bidding strategies: $b_1(x) = a_1 x + c_1$, $b_2(y) = a_2 y + c_2$. Then $b_2^{-1}(b_1) = (b_1 - c_2)/a_2$, and player 1's maximization problem simplifies to

$$\max_{b_1} \frac{1}{2}\left(\frac{b_1 - c_2}{a_2}\right)^2 + \left(x + \epsilon\right)\left(\frac{b_1 - c_2}{a_2}\right) - \frac{b_1(b_1 - c_2)}{a_2}$$

The first-order condition for this problem gives the outcome

$$b = \frac{a_2 x + a_2 \epsilon + (a_2 - 1)c_2}{2a_2 - 1}.$$

Similarly, $b_2 = [a_1 y + (a_1 - 1)c_1]/(2a_1 - 1)$. For these to hold simultaneously, it must be that

$$a_1 = \frac{a_2}{2a_2 - 1}, \tag{1}$$

$$a_2 = \frac{a_1}{2a_1 - 1}, \tag{2}$$

$$c_1 = \frac{[(a_2 - 1)\,c_2 + a_2\,\epsilon]}{2a_2 - 1} \tag{3}$$

$$c_2 = \frac{[(a_1 - 1)\,c_1]}{2a_1 - 1} \tag{4}$$

Note that (1) implies (2) and vice versa. Substituting (1) and (2) into (3) and (4) gives

$$c_1 = (1 - a_1)c_2 + a_1\epsilon, \tag{3'}$$

$$c_2 = (1 - a_2)c_1. \tag{4'}$$

Putting all equations in terms of a_2 gives a set of three equations:

$$c_1 = \epsilon/a_2, \tag{5}$$

$$c_2 = (1 - a_2)\epsilon/a_2, \tag{6}$$

$$a_2 = \frac{a_2}{2a_2 - 1}. \tag{7}$$

At this point, the choice of a_2 is arbitrary, but there is also a boundary condition. In a first-price auction, the top possible bid by each side must be the same. Otherwise, one player should reduce the top bid, since it is possible to win with probability 1 with a lesser bid. This gives the further condition

$$a_1 + c_1 = a_2 + c_2. \tag{8}$$

The final set of equations (5)–(8) gives a quadratic equation for a_2 with the positive root $a_2 = (1 + \epsilon) + \sqrt{1 + \epsilon^2}/2$. The remaining parameters, a_1, c_1, c_2 are given by the appropriate equations in (5)–(8) and the value of a_2.[22] These values satisfy $1 - \epsilon < a_1 < 1$, $0 < c_1 < \epsilon$, $-\epsilon < c_2 < 0$. All bids by either player are within ϵ of the symmetric equilibrium bid functions $b^*(x) = x$. □

References

Baltagi, B.H., 1986, "Pooling Cross-Sections with Unequal Time Lengths," *Economics Letters*, 18, 133–136.

Bikhchandani, S., 1988, "Reputations in Repeated Second-Price Auctions," *Journal of Economic Theory*, 46, 97–119.

Cramton, P., 1997, "The FCC Spectrum Auctions: An Early Assessment," *Journal of Economics and Management Strategy*, this issue.

Fudenberg, D. and D.K. Levine, 1997, "Measuring Players' Losses in Experimental Games," *Quarterly Journal of Economics*, 112, 507–536.

—— and J. Tirole, 1991, *Game Theory*, Cambridge, MA: The MIT Press.

Garvin, S. and J.H. Kagel, 1994, "Learning in Common Value Auctions: Some Initial Observations," *Journal of Economic Behavior and Organization*, 25, 351–372.

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 —— 1993, "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, 103, 868–879.

——, ——, and R.M. Harstad, 1995, "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions," *International Journal of Game Theory*, 24, 293–319.

Klemperer, P., 1997, "Almost Common Value Auctions: The 'Wallet' Game and Its Applications in Takeover Battles and PCS Auctions," Mimeo, Nuffield College, Oxford University.

Levin, D. and R. Harstad, 1986, "Symmetric Bidding in Second-Price, Common-Value Auctions," *Economic Letters*, 20, 315–319.

22. With these values $c_1 > c_2$. Some of player 2's bids fall below the minimum possible bid by player 1. For those bids below c_1 to be optimal, the equilibrium bid functions must satisfy an additional boundary condition $b_2(c_1) = c_1$. But this new boundary condition is guaranteed by the first-order equations and satisfied algebraically by the given values.

Maskin, E.S. and J.G. Riley, 1996, "Auction Theory with Private Values," *American Economic Review*, 75, 150–155.

—— and ——, 1991, "Asymmetric Auctions," Mimeo, Harvard University.

Milgrom, P., 1981, "Rational Expectations, Information Acquisition and Competitive Bidding," *Econometrica*, 49, 921–943.

—— and R. Weber, 1982, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

Pearce, D., 1984, "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.