

RISK AVERSION AND THE WINNER'S CURSE

Charles A. Holt
Department of Economics
University of Virginia
Charlottesville, VA 22903
holt@virginia.edu

Roger Sherman
Department of Economics
University of Houston
Houston, TX 77204-5882
rsherman@uh.edu

revised, April 2000

Abstract

This paper analyzes an auction in which bidders see independent components of a common prize value. The Nash equilibrium for two rational bidders is shown to be independent of risk attitudes. The information structure allows explicit calculation of an alternative equilibrium in which naive bidders do not correctly discount the value of the prize, contingent on winning, and thus they suffer the winner's curse. Subjects in a laboratory experiment clearly fall prey to the winner's curse; the data conform most closely to the predictions of the naive model. Moreover, the level of risk aversion implied by fitting the naive model is quite similar to an independent risk aversion measure obtained in a separate (private-value) bidding exercise.

RISK AVERSION AND THE WINNER'S CURSE

Charles A. Holt and Roger Sherman*

I. Introduction

The value of an object at auction is often common to all bidders but unknown to each of them. Bidders, therefore, will seek and obtain their own information about the prize value. One bidder's idea of an object's value at auction often depends, at least in part, on inferences about information that is only seen by rival bidders. If a bidder can observe others' bids, as in an oral English auction, those other bids can provide some information about others' value estimates, and thus they can influence the bidder's own value assessment. Even if other bids cannot be observed, which is the case in a sealed-bid auction, it is possible for bids to be affected by the fact that winning the auction is itself informative. In particular, having the highest bid tends to mean that others' value estimates were relatively low, and a bidder who does not realize this may bid too high and end up paying more than the prize is worth. This possibility, known as the "winner's curse," was first considered by Robert Wilson (1969). Of course, it is also possible to model an equilibrium in which rational bidders recognize that winning is informative and make correct inferences from the equilibrium bids of others (Wilson, 1969, 1977).

There is considerable evidence that bidders fall prey to the winner's curse, both in field situations and in controlled laboratory experiments (John Kagel, 1995). In these experiments, each subject receives an independent signal that is uniformly distributed around the true common prize value. Subjects consistently lose money by bidding above the levels predicted by a Nash equilibrium that assumes rational information processing. One approach to explaining the winner's curse is to investigate whether losses are caused by artifacts of the experimental design. These could be subject inexperience, the inability of the experimenter to actually collect losses from laboratory subjects -- called a limited-liability effect -- or experimenter demand effects.

* Department of Economics, University of Virginia, and Department of Economics, University of Houston. This research is funded in part by the National Science Foundation (SBR-9818683). We wish to acknowledge helpful suggestions received from Simon Anderson, Jacob Goeree, Shiferaw Gurm, Ronald Harstad, Glenn Harrison, William Johnson, Aubbik Kahn, Jennifer Elise Novak, Steve Stern, and Nat Wilcox.

Overbidding may be reduced, but not eliminated, by giving subjects more experience (Kagel, 1995). The limited-liability argument is that subjects might bid aggressively in a risk-seeking manner if large losses would not have to be paid to the experimenter. The experimenter-demand explanation arises when subjects expect to be increasing their earnings during the experiment. These and other procedural explanations have not gone unchallenged, but it is prudent to use experimental procedures that deal with these explanations so results are not contaminated by them.¹

A second way to explain the winner's curse is to relax one or more of the behavioral assumptions in rational bidding theory. One approach is to keep the Nash equilibrium assumption about rational strategic responses to other's bids, but to relax the assumption that bidders correctly perceive that winning is an informative event. In particular, assume that bidders only base their estimates of the prize value on their own information and on some *average* of what the others' value estimates might be. This procedure is naive because it ignores the fact that a bid is only relevant when it wins, i.e. when the others' value estimates are relatively low. Kagel and Levin (1986) have termed this a "naive model", and Plott and Lind (1991) have termed it a "strategic discounting" model.² Kagel and Levin (1986) report that bidders' earnings were too low relative to the rational model, but still higher than predicted under a naive model. Lind and Plott (1991) conclude that the fully rational model fits their data better than their strategic discounting model, despite persistent overbidding and losses.

Lind and Plott (1991) point out a major impediment to evaluating the data from existing experiments: the equilibrium calculations involve an auxiliary assumption of risk neutrality, both with rational and with naive information processing:

¹ Robert Hansen and John Lott (1991) argued that the winner's curse is due to the limited liability of laboratory subjects, but this was sharply refuted by Kagel and Dan Levin (1991) and by Barry Lind and Charles Plott (1991). Cox and Smith (1992) reported that the "safe haven" treatment reduced the incidence of losses, especially with experienced subjects, but the conclusion that the winner's curse is largely due to experimenter demand effects has been disputed by Kagel (1995).

² Holt and Sherman (1994) formulate an analogous naive bidding model in a different context, with a single bidder and a single seller, where the seller is informed about the intrinsic prize value and the bidder is uninformed. For some parameter values, naive bidding results in excessively high bids (a winner's curse), and for other parameter values, naive bidding results in excessively low bids (a "loser's curse"). The reported bid data tend to conform more closely with the naive model than with a model that assumes rational information processing.

Part of the difficulty with further study stems from the lack of theory about the behavior of common-value auctions with risk aversion. Closed-form solutions which permit researchers to estimate models of "subrational" behavior have not been worked out. If the effect of risk aversion is to raise the bidding function as it does in private auctions, then risk aversion together with the strategic discounting model might resolve the puzzle; but, of course, this is only a conjecture.³

This paper presents a simple model of bidding in a first-price auction, with an information structure that permits direct evaluation of the Lind and Plott conjecture. In particular, the common prize value is determined by independent, distinct components that are privately observed by two competing bidders. Each bidder essentially sees half of the prize value. We show that with two bidders and rational information processing, the symmetric Nash equilibrium is independent of risk attitudes. But with naive information processing, the slope of the equilibrium bid function depends on a coefficient of constant relative risk aversion. Thus the framework allows a straightforward investigation of the separate effects of naive bidding and risk aversion.

The rest of the paper is organized as follows. Section II presents the model and derives testable propositions of rational bidding theory. Section III summarizes the laboratory procedures and the bid data, which exhibit the overbidding associated with naive behavior and the winner's curse. In Section IV a naive model of bidding behavior is developed that includes an effect of risk aversion. Section V contains a more detailed analysis of risk aversion, and the final section concludes.

II. Rational Bidding

The model is one in which two bidders observe separate components of a common prize value, i.e. bidder 1 observes v_1 and bidder 2 observes v_2 . The monetary value of the prize to the winning bidder is an average: $(v_1 + v_2)/2$. Each bidder essentially sees half of the common

³ Lind and Plott (1991, p.344). The comment on "private auctions" is a reference to the well known fact that risk aversion raises bids in auctions with known, independent private values.

value.⁴ The two bidder's value components (hereafter "values") are independent realizations of a uniform random variable on $[0, 1]$. Each bidder observes the actual realization of one value component, and knows the distribution used to determine the other component. Although this information structure is quite specialized, it has the advantage of permitting us to distinguish the separate effects of risk aversion and naive information processing in a first-price auction.⁵ We consider a first-price auction in which each bidder makes a (sealed) bid, and the winner must pay the winning (high) bid for the prize.

In this section, we begin by calculating a symmetric Nash equilibrium for the case of rational, risk-neutral bidders. The effects of risk aversion and naive information processing will be considered subsequently.

Rational Risk-Neutral Bidders

This is a game with incomplete information, and the equilibrium will specify the bid, b_i , as a function of the bidder's own value, v_i . The uniformity of the value distributions will produce a linear bidding function. To determine this function, let bidder 1 assume that bidder 2 uses:

$$(1) \quad b_2 = \lambda v_2,$$

where $\lambda > 0$ by assumption. It follows that the probability that a bid b_1 will win is calculated: $\Pr\{b_1 > b_2\} = \Pr\{b_1 > \lambda v_2\} = \Pr\{b_1/\lambda > v_2\} = b_1/\lambda$, since v_2 is uniformly distributed on $[0,1]$.⁶ Given (1), the rational bidder 1 realizes that winning the auction puts an upper limit, b_1/λ , on the range of possible values of v_2 . Winning the auction then means the expected value of v_2 would be $b_1/2\lambda$, since the conditional distribution of v_2 is uniform on $[0, b_1/\lambda]$. It follows

⁴ This value structure is a special case of the general model analyzed by Milgrom and Weber (1982). Also, see Harstad, Kagel, and Levin (1990) for an extension that permits the number of bidders to be uncertain. The particular formulation with a sum of independent common value components is used by Jeremy Bulow and Paul Klemperer (1994) in an analysis of market "frenzies" and crashes. A generalization with both common and private value components is developed by Goeree and Offerman (1999).

⁵ See Wolf Albers and Harstad (1991) for a theoretical and experimental analysis of bidding in an English auction with independent common value components and risk neutrality. Some results from English auctions carry over to first-price auctions via revenue equivalence, but revenue equivalence does not hold when bidders are risk averse (Holt, 1980) or naive. Therefore, this section will provide a direct analysis of bidding in the first-price auction.

⁶ It can be verified later that b_i/λ does not exceed 1 in equilibrium.

from the uniform distribution of v_2 that the expected prize value for bidder 1, conditional on winning, is

$$(2) \quad \frac{v_1 + E(v_2 \mid v_2 < b_1/\lambda)}{2} = \frac{v_1 + b_1/2\lambda}{2}.$$

Under risk neutrality, bidder 1's expected profit can be expressed as the probability of winning, b_1/λ , times the difference between the conditional expected prize value and the amount bid:

$$(3) \quad \text{expected profit} = \frac{b_1}{\lambda} \left[\frac{v_1 + b_1/2\lambda}{2} - b_1 \right].$$

The objective function in (3) is a concave, quadratic function of the bid b_1 , and setting the derivative equal to zero yields:

$$(4) \quad b_1 = \frac{\lambda}{4\lambda - 1} v_1,$$

which is linear in v_1 . In a symmetric equilibrium, the slope of bidder 1's best response in (4) is equal to the slope λ of bidder 2's bid function in (1), which implies that $\lambda = .5$. (Recall that the case of $\lambda = 0$ was ruled out by assumption.) Substituting this value into (4), we find that, in equilibrium, each bidder bids one half of the observed value estimate: $b_1 = v/2$.

Rational Risk Averse Bidders

Now suppose that each bidder has risk preferences represented by a concave utility function. As before, we assume that bidder 2 is using the bid function in (1), so that a bid b_1 will win if $b_1/\lambda > v_2$. Integrating bidder 1's utility of earnings over this range yields:

$$(5) \quad \int_0^{b_1/\lambda} U\left(\frac{v_1 + v_2}{2} - b_1\right) dv_2,$$

where the standard normalization that $U(0) = 0$ is used to eliminate the utility term for the event of a loss. The first-order condition for the choice of b_1 is:

$$(6) \quad \int_0^{b_1/\lambda} -U'\left(\frac{v_1 + v_2}{2} - b_1\right) dv_2 + \frac{1}{\lambda} U\left(\frac{v_1 + b_1/\lambda}{2} - b_1\right) = 0.$$

The left term in (6) can be integrated directly to obtain an expression in b_1 and v_1 only:

$$(7) \quad -2U\left(\frac{v_1 + b_1/\lambda}{2} - b_1\right) + 2U\left(\frac{v_1}{2} - b_1\right) + \frac{1}{\lambda}U\left(\frac{v_1 + b_1/\lambda}{2} - b_1\right) = 0.$$

Although (7) looks complicated, it is straightforward to show that it is satisfied if $b_1 = v_1/2$. In this case, the middle utility expression on the left side of (7) becomes $U(0)$, which is zero, and the other two utility expressions cancel out, leaving $\lambda = 1/2$. Thus bidding half of one's value estimate is a symmetric Nash equilibrium for the case of general concave utility functions. To summarize:

Proposition 1. When the common value is the average of (uniformly distributed) individual value signals, then it is a symmetric Nash equilibrium in a first price auction to bid half of the value signal, for any common concave utility function.

The intuition behind this invariance result is that risk aversion has two effects: a prize of uncertain value is worth less to a risk-averse bidder, which tends to cause lower bids, but the uncertainty associated with losing the auction tends to cause a risk averse bidder to bid higher. These two effects exactly balance in this model with two bidders.⁷

III. Experimental Evidence

Three groups of eight subjects were recruited from undergraduate classes at the University of Virginia (2 groups) and the University of Houston (1 group) "to earn money in a research experiment." The experiment began by reading aloud the instructions, included in Appendix A, and subjects were given a chance to ask questions. The initial common value auction lasted for 15 periods. Subjects were matched randomly into pairs at the start of each period, and did not know the identity of the other subject with whom they were matched. Thus four pairs of bidders were competing at auction in each period. The random matchings were intended to be a credible, flexible, and quick procedure for changing the pairing of subjects, to reduce the chance for tacit

⁷ In a model with independent private values, there is no prize value uncertainty, and the uncertainty associated with the possibility of a loss causes bids to be higher when bidders are risk averse.

cooperation across periods that could arise if the pairings were unchanged.⁸

The common-value auction model was implemented by throwing ten-sided dice privately for each subject, to determine independent value estimates. The participants then marked their bid decisions on record sheets, which were collected to record the other bidder's value and bid, and to record the subject's own earnings. The high bidder in each pair would pay the high bid and win a prize worth the average of the two bidders' value numbers. There was no public announcement of the results. Every subject began period 1 with an initial cumulative earnings of \$15 to ensure adequate earnings for the session and to avoid reckless, risk-seeking behavior that might occur if a earnings approached zero, since losses cannot be collected.⁹

At first, the ten-sided dice were thrown only twice for each bidder, with the first throw determining the "tens" digit and the second throw determining the "ones" digit. This yields value numbers between \$0.00 and \$0.99. Five periods were conducted with these payoffs in "part A," while subjects became familiar with the experimental task. In the ten periods of "part B" that followed, the dice were thrown three times for each bidder, with the first throw determining the dollars digit. These throws yield values in the range from \$0.00 to \$9.99, essentially making the potential earnings and losses ten times as great. These scale effects do not alter the previous section's Nash equilibrium calculations.

The part B bid data for all 24 subjects are shown in Figure 1, where the own value estimate in pennies is on the horizontal axis and the bid is on the vertical axis. The rational bid function is shown as the straight ray from the origin with a slope of 0.5. With few exceptions, bids are above this line, as predicted by the winner's curse. Many subjects suffered declines in cumulative earnings during the session. The regression with pooled data, allowing for a period effect, yields the following bid function:

⁸ As will be apparent from the data, there was no evidence of tacit collusion.

⁹ This was in addition to a \$6 participation payment that was made in cash at the beginning of the session. Participants were told that any subject would be paid and released immediately if his or her cumulative earnings fell below \$5 after any period (not counting the \$6 initial payment). In fact, no subject's cumulative earnings fell below \$11.92 during the bidding.

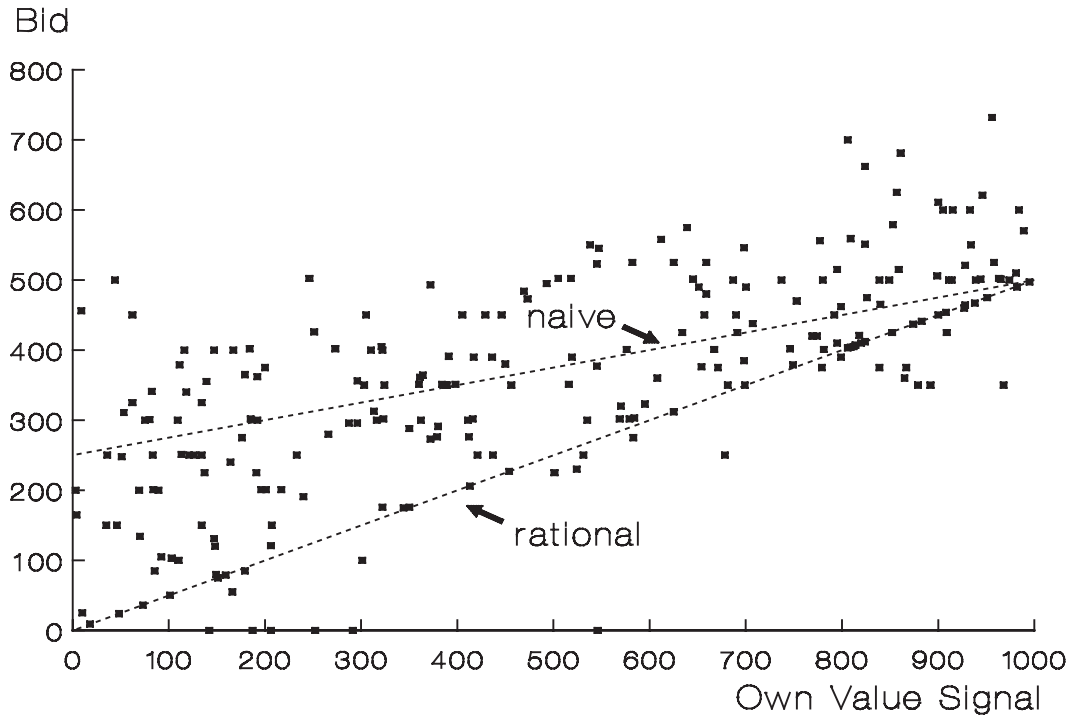


Figure 1. Bid Data for Common Value Auction

$$(8) \quad \text{Bid} = 200.8 + .360 \text{ Value} - 3.4 \text{ Period} \quad R^2 = .52,$$

$$\quad \quad (105.6) \quad (.022) \quad (2.4)$$

where the standard errors are shown in parentheses. The period variable was included to check for a time trend, which is small at a decline of 3.4 cents per period and has a sizeable error. Recall that the symmetric equilibrium for the rational bidding implies an intercept of 0 and a slope of .5 on value, regardless of risk aversion. The bid function slopes for the three sessions estimated independently are all within two percentage points of the slope in equation (8).

Conclusion 1. Bids are generally above the levels predicted by the rational behavior, and the empirical bid/value relationship is significantly flatter than the symmetric Nash bidding strategy (for any concave utility function).

Thus departures from rational bidding cannot be explained by risk aversion, or at least not by risk aversion alone.

IV. Naive Bidding

In any symmetric equilibrium, the bidder with the high value estimate will win the auction. In that case, on average, the lower value is only half as high as the value of the winning bidder. Thus if bidder 1 wins, the value of the prize conditional on winning in a symmetric equilibrium is the average of the bidder's own value, v_1 , and the conditional expected value of the lower value, $v_1/2$. The winner's curse can occur if a bidder *does not realize* that winning the auction implies the other bidder's value estimate is relatively low. The simplest naive model is obtained by using the unconditional expected value of the other bidder's value estimate, which is 0.5 when the values are uniform on $[0,1]$. Since 0.5 is typically greater than the conditional value, $v_1/2$, this bias will tend to produce higher bids, even with risk neutral bidders, and thus will give rise to a winner's curse.

With naive bidding, the equilibrium will no longer be to bid half of one's value, regardless of risk aversion, so we need to allow for nonlinear utility and non-proportional bid functions. Let bidder 2 use a differentiable strategy denoted by $b_2 = B(v_2)$, where $B' > 0$. The inverse of this bid function will be denoted by $v_2 = V(b_2)$, where $V' = 1/B'$. Bidder 1 wins the prize when $b_1 > B(v_2)$, or equivalently when $V(b_1) > v_2$, which occurs with probability $V(b_1)/L$ when v_2 is uniform on $[0, L]$, where $L = 999$ (pennies). The naive bidder is assumed to use the unconditional expected value of v_2 , i.e., $(0.5)L$, in evaluating the prize value: $(v_1 + 0.5L)/2$. Thus the naive bidder's objective is to maximize:

$$(9) \quad (0.5v_1 + 0.25L - b_1)^{1-r} V(b_1).$$

where r is a coefficient of constant relative risk aversion, with $0 < r < 1$.

The objective in (9) really incorporates two sources of bias, the use of the unconditional value distribution and the replacement of the random variable with its expected value. It is straightforward to express the first-order condition for the choice of b_1 :

$$(10) \quad -(1-r)V(b_1) + (0.5v_1 + 0.25L - b_1)V' = 0,$$

where V' denotes the derivative of the inverse bid function. In a symmetric equilibrium, bidder 1 uses the same strategy as bidder 2, so replace $V(b_1)$ and V' in equation (10) with v_1 and $1/B'$ respectively, to obtain:

$$(11) \quad -(1-r)v_1 + (0.5v_1 + 0.25L - b_1)/B' = 0.$$

This equation can then be solved for b_1 :

$$(12) \quad b_1 = 0.25L + (0.5 - B' + rB')v_1.$$

In equilibrium, the coefficient of v_1 in (12) must equal the slope of the bid function, B' , and equating these yields a solution for the slope: $B' = 1/(4-2r)$. Substituting back into (12) and using $L = 999$ (pennies), one obtains the bid function:

$$(13) \quad b_i = 250 + \left(\frac{1}{4-2r} \right) v_i.$$

Proposition 2. When the common value is the average of the (uniformly distributed) individual value signals and individuals ignore the fact that winning is informative, then the slope of the symmetric Nash bid strategy is 1/4 for risk neutral bidders, and it is 1/(4-2r) for bidders with a constant relative risk aversion, r .

When $r = 0$, equation (13) yields the naive bid function for the risk-neutral case, which shown as the upper straight line in figure 1. The naive bid function in (13) has an intercept that is close to the estimated bid function in (8), but the slope for the risk neutral case is too flat relative to the data in the figure. An increase in risk aversion increases the slope of the line, and the slopes of the theoretical bid function (13) will equal the slope of the regression line in (8) if $r = 0.583$, or nearly 0.6. This raises the issue of whether the implied level of risk aversion is reasonable.

V. Risk Aversion

To summarize up to this point, bidding in the common-value auction exhibits a clear winner's curse effect that cannot be explained by risk aversion alone, since the Nash bid function is independent of risk aversion in the setup used. Nor can the pattern of aggregate bids be explained by one particular type of naive information processing alone, since the naive bid line for risk neutral bidders in figure 1 is too flat to explain the data. Interpreted as for an individual, the estimated bid function does correspond to the naive model with a relative risk aversion parameter of about .6. Using risk aversion to explain bid patterns has a long and controversial

history for private-value auction experiments.¹⁰

To see how attitude toward risk might affect bidding we obtained an independent measure of risk aversion through a separate, private-value, auction. Each of the 24 subjects participated in the second auction with a known prize value.¹¹ In this auction, the subject knew that the "other bidder" was a simulated player whose bid was determined by the throw of dice (the instructions are attached in Appendix B). At the beginning of each period, we went to each subject's position and threw ten-sided dice to determine a prize value for that subject. Knowing this prize value, the subject would then choose and record a bid. Finally, we would throw the ten-sided dice to determine the bid of the "other bidder." The dice were thrown twice in the first 5 periods, with the first throw determining the "tens" digit, so that both the prize value and the other bid range from \$0.00 to \$0.99. The dice were thrown three times in the final 10 periods, so the prize value and other bid ranged from \$0.00 to \$9.99. We restrict our attention to these part B data.¹²

With constant relative risk aversion and a value of v_i , the expected utility is:

$$(14) \quad (v_i - b_i)^{1-r} (b_i/1000),$$

where $b_i/1000$ is the probability that the b_i will exceed the other bid (in pennies). The first-order condition for maximizing (14) with respect to b is $-(1-r)b + v - b = 0$, which yields an optimal bid that is proportional to the value: $b_i = v_i/(2-r)$. Thus the slope of the bid function will provide an estimate of the risk aversion coefficient. With risk neutrality, $r = 0$, and the optimal bid is one half of the private value. The data for all 16 subjects are shown in figure 2; most of the bids are above the risk neutral bid function (the lower straight line). A linear regression yields:¹³ The intercept is not significantly different from zero, but the slope coefficient of .667 is

¹⁰ For example, see Cox, Smith, and Walker (1983), Glenn Harrison (1989), and the exchange in the December 1992 issue of the *American Economic Review*.

¹¹ At Virginia the two different auctions, common value and private value, were carried out in separate sessions, and at Houston the two auctions were carried out in a single session.

¹² Since losses were not likely, we only gave subjects a fixed payment of \$5.00, in addition to the \$6 participation payment, to keep earnings for the session in the same range as in the common value auction.

¹³ The inclusion of a variable for the period number has little effect on the coefficient and is omitted.

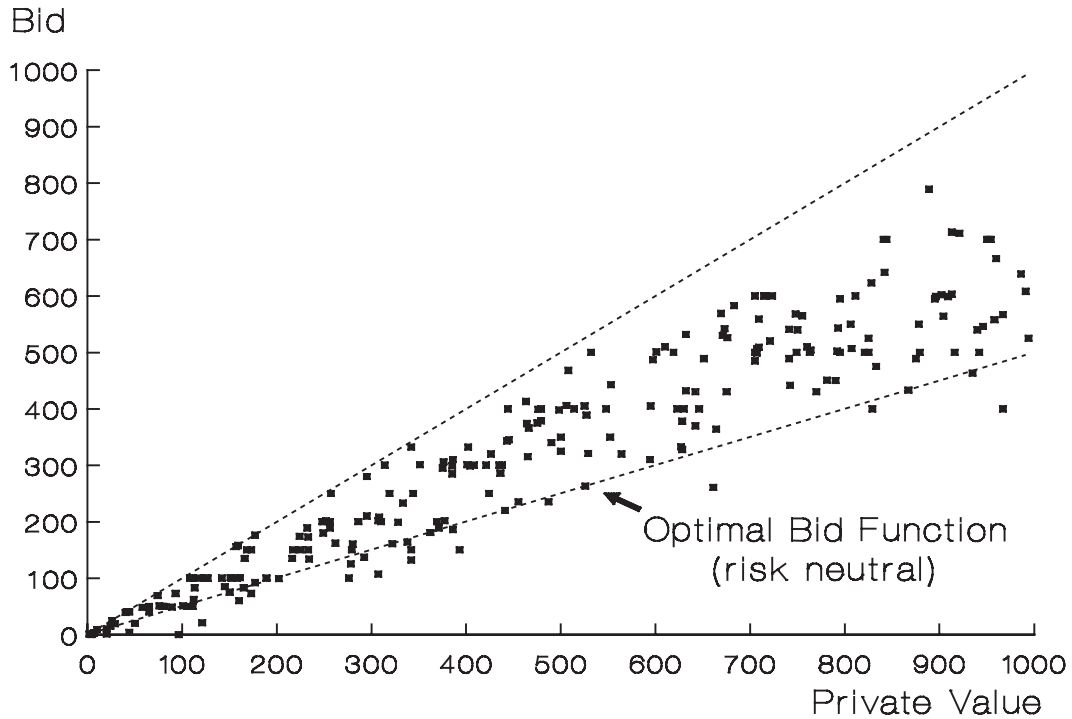


Figure 2. Bid Data for Private Value Bidding Against a Uniform Distribution of Other Bids

$$(15) \quad \text{Bid} = 0.076 + 0.654 \text{ Value} \quad R^2 = .89,$$

$$\quad \quad \quad (0.665) \quad (0.015)$$

significantly above the 0.5 prediction for risk neutrality. In fact, this estimated slope implies a level of constant relative risk aversion of about 0.47, which is essentially "square root" utility.

Conclusion 2. The level of constant relative risk aversion (0.61) implied by the symmetric naive bidding in the common value auction treatment is roughly consistent with the risk aversion estimate of (0.47) implied by aggregate bidding behavior against a uniform distribution of simulated bids in the private-value auction treatment.

The closeness of these two estimates of constant relative risk aversion is encouraging, given the different nature of the bidding tasks and the aggregate level of analysis.

The similarity of risk aversion measures implied by aggregate data for the two auction treatments raises the issue of whether behavior is related at the individual level. We ran the

regressions analogous to (15) to estimate bidding functions for each of the individuals in the private-value auctions, and equated the slope coefficients to $1/(2-r_i)$ to estimate individual risk aversion parameters. The median of these individual risk aversion estimates is 0.63, with two thirds between 0.19 and 0.88. These individual bidding tasks were separate decision problems against known bid distributions, so the estimates depend (only!) on an assumption of optimal behavior. Obtaining individual risk aversion estimates from the common-value data involves many more auxiliary assumptions, since individuals may differ not only in terms of risk aversion but also in terms of rationality and in their beliefs about others' bids based on their own particular experience. As a rough approximation, we assume that individuals are all perfectly naive (they ignore the fact that winning is informative) and that each individual's perceived probability of winning is a linear function of the bid, with a slope that can vary from one person to another.¹⁴ Consider the first-order condition (10) for individual optimization in the naive model, where r represents an individual-specific risk aversion parameter, hereafter r_i , and $V(b)$ summarizes the individual's subjective probability of winning with a bid b . Given the assumed linearity of this function, it is straightforward to show that the resulting relationship between bid and value has a slope of $1/(4-2r_i)$, just as in the case of the equilibrium model with homogeneous risk aversion (13). (This equivalence is due to the linearity of the probability of winning function.) In this way we obtained alternative risk aversion measures for the individuals. The median of these individual estimates is 0.40, with two thirds between 0.24 and 0.53, which is similar to the distribution of estimates implied by the private value bidding decisions. The relationship between these two sets of risk aversion estimates, however, is weak; a linear regression yields a reasonable slope coefficient of 0.92 (0.56), but with an R^2 of only 0.11, after eliminating an outlier subject with an estimated risk aversion of -8.4 in the common-value auction.¹⁵ This weak relationship is not surprising given the limited data for each person and

¹⁴ In fact, the cumulative distribution of actual bids in the common value auction is almost highly linear over the range from \$2.50 to \$5.40 where 70% of the bids occurred. Over this range, each \$1 increase in bid raises the probability of winning by about 25%. Of course, individuals do not see the population bid distribution, but linearity seems like a reasonable rule of thumb in this context.

¹⁵ This individual had also had a relatively low risk aversion estimate (0.30) from the private-value auction treatment. In the common-value treatment, this person had a string of high-payoff wins with bids of about \$3.00, so the person stayed with that bid and the resulting flat bid/value relationship implied risk loving behavior.

the auxiliary assumptions needed to obtain the common-value-auction estimates (identical naive behavior and linear probability of winning beliefs).

In general, the range of risk aversion estimates implied by our subjects' bidding behavior (0.2 to 0.8, centered around 0.4 to 0.6) is roughly consistent with some estimates obtained by others, in auctions, games, and direct lottery choice situations. In auction experiments, Chen and Plott find a mean risk aversion of 0.48 for first-price auction experiments with private values drawn from a non-uniform distribution. Harrison (1990, p.543) concludes that a constant relative risk aversion of 0.55 would "create virtual equality between observed and predicted behavior" in a four-person experiment reported in Cox, Roberson, and Smith (1982). Cox and Oaxaca (1994, table 1) estimate relative risk aversion parameters for 40 subjects; these parameters range from 0 to .98, with an average value of .67. This is fairly close to the estimate implied by the aggregate data here for the naive bidding model, especially when considering that it was obtained in a very different context with different subjects. Similarly, Goeree, Holt, and Palfrey report estimates of 0.55 and 0.51 in for two different groups of subjects in different treatments of a private value auction. They also report individual estimates, with three fourths being between 0.3 and 0.7. In an analysis of thirty-seven 2x2 matrix games, Goeree and Holt (2000) obtained constant relative risk aversions estimates of about 0.46; this risk aversion substantially improved theoretical predictions of behavior in games where one strategy yielded a much wider spread of payoffs than another. Data from direct binary choices involves fewer auxiliary assumptions. Binswinger (1980) uses high-stakes lottery choices made by rural farmers in India to conclude that "...at high payoff levels, virtually all individuals are moderately risk averse, with little variation according to personal characteristics." The most common levels of constant relative risk aversion were in a "moderate range" (.32 to .81) or higher, with risk aversion observed even in low-stakes gambles. More recently, Bosch-Domenech and Silvestre (1999) give subjects a chance to purchase actuarially fair insurance on losses, and they find that purchases are nearly universal for losses in the high (\$10-\$100) range. Holt and Laury (2000) use a menu of lottery choices to infer risk aversion from the "cross-over" point. The modal cross over implies a relative risk aversion of about 0.3, which is close to the 0.37 estimate for the same subjects who were put into

a private-value bidding task against a known distribution.¹⁶ When stakes were multiplied by a factor of 20, yielding prizes of up to \$77.00, the implied levels of risk aversion rose to about .5. Finally, the risk neutrality restrictions are rejected for all 53 individuals making the repeated lottery-choices over several days (Hey, 1999).

VI. Conclusion

Our aim in this paper is to evaluate the Lind and Plott (1991) conjecture that overbidding in common-value auctions may be due in part to risk aversion and in part to naive bidding. We first construct as a reference point a common-value auction in which risk aversion has no effect on rational bidding. Each of two bidders in this auction has information about roughly half of the prize value, and their rational bids are independent of risk aversion. The clear overbidding relative to this prediction is evidence of naive bidding that cannot be attributed to risk aversion alone. There is, however, overbidding even relative to the predictions of the naive model, especially for high value signals, which can be explained by allowing for risk aversion. The implied risk aversion across 24 subjects in common-value auctions is roughly consistent with independent estimates of risk aversion obtained for the same subjects in private value auctions against random other bidders. Thus risk aversion can affect naive bidding in common value auctions, even in low-stakes laboratory settings. As noted in the previous section, there is some evidence that risk aversion may increase for large stakes, which suggests that knowledge of risk aversion may help to predict overall bidding behavior in field situations.

¹⁶ The only estimates that indicate modal risk-preferring behavior for these subjects were obtained by giving them ownership of a 50/50 chance of winning \$8.50 or \$0.00 and asking for a willingness to accept in a Becker/DeGroot/Marshack mechanism. The modal price of \$5.00 was above the expected value and implied a risk aversion of -.72 (risk loving). This divergence from the other measures for the same subjects is probably due to the well known upward bias in willingness to accept prices, e.g., Shogren, Sgin, Hayes, and Kliebenstein (1994). A similar disparity between willingness to accept and other measures of risk aversion is reported in Kachelmeyer and Shehata (1992) and Isaac and James (1999). Curiously, Isaac and James find a negative relationship between auction-based and willingness-to-accept based measures of risk aversion.

References

- Albers, Wolf and Ronald M. Harstad (1991) "A Framing Effect Observed in a Market Game," in R. Selten, ed., *Game Equilibrium Models*, Berlin: Springer-Verlag, 308-336.
- Binswanger, Hans P. (1980) "Attitude Toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics*, 62, August, 395-407.
- Bosch-Domenech, Antoni and Joaquim Silvestre (1999) "Does Risk Aversion or Attraction Depend on Income?" *Economics Letters*, 65, 265-273.
- Bulow, Jeremy and Paul Klemperer (1994) "Rational Frenzies and Crashes," *Journal of Political Economy*, February, 102, 1-23.
- Chen, Kay-Yut and Charles R. Plott (1998) "Nonlinear Behavior in Sealed Bid First-Price Auctions," *Games and Economic Behavior*, 25, 34-78.
- Cox, James C., and Ronald L. Oaxaca (1996) "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions," in R. M. Isaac, ed., *Research in Experimental Economics*, Vol. 6, Greenwich, Conn.: JAI Press, 131-148.
- Cox, James C., B. Roberson, and Vernon L. Smith (1982) "Theory and Behavior of Single Price Auctions," in V. L. Smith, ed., *Research in Experimental Economics*, Vol. 2, Greenwich, Conn.: JAI Press, 1-43.
- Cox, James C., Smith, Vernon L., and James M. Walker (1983) "Tests of a Heterogeneous Bidder's Theory of First Price Auctions," *Economics Letters*, 12, 207-12.
- Goeree, Jacob K. and Charles A. Holt (2000) "A Model of Noisy Introspection," Discussion Paper, University of Virginia.
- Goeree, Jacob K., Charles A. Holt, and Thomas Palfrey (1999) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," Discussion Paper, California Institute of Technology.
- Goeree, Jacob K. and Theo Offerman (1999) "Competitive Bidding in Auctions with Private and Common Values," Discussion Paper, University of Virginia.
- Hansen, Robert G. and John R. Lott, Jr. (1991) "The Winner's Curse and Public Information in Common Value Auctions: Comment," *American Economic Review*, March, 81, 347-61.
- Harrison, Glenn W. (1989) "Theory and Misbehavior of First-Price Auctions," *American*

- Economic Review*, September, 79, 749-62.
- Harrison, Glenn W. (1990) "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis," *The Review of Economics and Statistics*, 541-546.
- Harstad, Ronald M., John H. Kagel, and Dan Levin (1990) "Equilibrium Bid Functions for Auctions with an Uncertain Number of Bidders," *Economics Letters*, 33, 35-40.
- Hey, John D. (1999) "Does Repetition Improve Consistency?" Discussion Paper, Universities of York and Bari.
- Holt, Charles A. and Susan K. Laury (2000) "Risk Aversion and Incentive Effects," Discussion Paper, Georgia State University.
- Holt, Charles A. (1980) "Competitive Bidding for Contracts Under Alternative Auction Procedures," *Journal of Political Economy*, June, 88, 433-45.
- Holt, Charles A. and Roger Sherman (1994) "The Loser's Curse," *American Economic Review*, June, 84, 642-52.
- Isaac, R. Mark and Duncan James (1999) "Just Who Are You Calling Risk Averse?" Discussion Paper, University of Arizona (forthcoming in the *Journal of Risk and Uncertainty*).
- Kachelmeier, Steven J. and Mohamed Shehata, "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, December, 82, 1120-1141.
- Kagel, John H. (1995) "Auctions: A Survey of Experimental Research," in J. H. Kagel and A. E. Roth, eds., *The Handbook of Experimental Economics*, Princeton: Princeton University Press, 501-85.
- Kagel, John H. and Dan Levine (1986) "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, December, 76, 894-920.
- Kagel, John H. and Dan Levine (1991) "The Winner's Curse and Public Information in Common Value Auctions: Reply," *American Economic Review*, March, 81, 362-69.
- Lind, Barry and Charles R. Plott (1991) "The Winner's Curse: Experiments with Buyers and with Sellers," *American Economic Review*, March, 81, 335-46.
- Milgrom, Paul R. and Robert J. Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, September, 50, 1089-1122.

- Shogren, Jason F., Shin, Seung Y., Hayes, Dermot J., and James B. Kliebenstein (1994) "Resolving Differences in Willingness to Pay and Willingness to Accept," *American Economic Review*, March, 84, 255-70.
- Wilson, Robert B. (1969) "Competitive Bidding with Disparate Options," *Management Science*, March, 15, 446-448.
- Wilson, Robert B. (1977) "A Bidding Model of Perfect Competition," *Review of Economic Studies*, October, 44, 511-18.

Appendix A. Instructions for the Winner's Curse Experiment

Introduction

You are going to take part in an experimental study of decision making. The funding for this study has been provided by several foundations. The instructions are simple, and by following them carefully, you may earn a considerable amount of money. The amount of money that you earn will depend partly on your decisions, partly on the decisions of other participants, and partly on chance. At this time, you will be paid \$6 for coming on time. All the money that you earn subsequently will be yours to keep, and your earnings will be paid to you in cash today at the end of this experiment. We will start by reading the instructions, and then you will have the opportunity to ask questions about the procedures described.

Earnings

The experiment consists of a series of separate periods in which each of you must choose an amount to bid for a prize of unknown value. You will be matched with another bidder, and the person with the high bid will obtain the prize and pay the amount that he or she bid. Thus the high bidder will earn a dollar amount that is the difference between the value of the prize and the amount of that person's bid. The low bidder will earn \$0 for the period.

The value of the prize in each period will be random. Here is a die with ten sides, marked 0, 1, ..., 8, 9. We will throw this die twice for each bidder to determine a random number between 0 and 99 cents (the first throw determines the "tens" digit, and the second throw determines the "pennies" digit). Any penny amount, \$0.00, ..., 0.99 is equally likely. The value of the prize is the *average* of the numbers seen by the two bidders, so the value of the prize will be between 0 and 99 cents. Thus each bidder sees two throws of a die that determine a number between 0 and 99 cents, but the value of the prize that is purchased by the high bidder is the average of the two numbers. Neither bidder can see the two throws of the die for the person with whom he or she is matched, and in this sense, each bidder only knows a part of the value of the prize when the bids are submitted. When you bid, you will know the number determined by your throws of the die, and you will only know that the other bidder's number is equally likely to be any number from 0 to 99. Thus the money value of the prize can be as low as half of your number (when the other person's number is 00), and it can be as high as your number plus 49.5 cents (when the other person's number is 99). For example, suppose I throw the die twice for you to get a ____ (first) and a ____ (second), which yields a number ____ . Since the other person's number can be between 00 and 99, all you would know is that the value of the prize is between $[00 + \text{____}]/2$ and $[99 + \text{____}]/2$. We will keep track of half pennies.

To explain the earnings more precisely, let V_1 denote the number (between 0 and 99 cents) observed by one bidder, and let V_2 denote the number (between 0 and 99 cents) observed by the other bidder. The two bidders will observe their own throws of the die (i.e. V_1 or V_2) and will be asked to make a bid for the prize without seeing the other person's bid or the other

person's throws of the die. The person with the highest bid will earn the average of V_1 and V_2 , minus the amount of the high bid, and the other person will earn zero. If the high bid is B , then the earnings will be computed:

$$\text{high bidder's earnings} = [V_1 + V_2]/2 - B$$

$$\text{low bidder's earnings} = 0.$$

Earnings will be divided equally between the two bidders in the case of a tie bid, so each will earn one half of the difference between $[V_1 + V_2]/2$ and the tie bid.

Obviously, the value of the prize will turn out to be between 0 and 99 cents, but neither bidder will have complete information about this value when the bids are submitted. Thus the value of the prize may be either above or below the highest (winning) bid. If the prize value turns out to be above the high bid, then the difference represents positive earnings that will be added to the cumulative earnings of the high bidder. If the prize value turns out to be below the high bid, then the difference represents a loss that will be subtracted from the cumulative earnings of the high bidder. The low bidder always earns \$0 and will have no change in cumulative earnings for the period.

Record of Results

Now, each of you should examine the decision sheet for part A. This sheet is the last one attached to these instructions. Please fill in the date on the upper left. Your identification number is written in the upper right part of this sheet. The column on the left side of your record sheet shows the period (the first row corresponds to period 1 and so on). Going from left to right, you will see columns for the number determined by your dice throws, for your bid, for the number determined by the other bidder's dice throws, and for their bid. These numbers will be used to calculate your earnings.

At the start of each period, one of us will come to your desk and throw the ten-sided die twice to determine a number between 0 and 99, which you will then enter in column (2). After seeing this number, you will choose a bid, which must be between \$0 and \$0.99. Your bid is to be written in column (3). Please use a decimal point to denote pennies.

After you make and record your bid for the period, we will collect all decision sheets and bring them to the front of the room. Then we will draw numbered ping pong balls to match each of you with another person. Here we have a container, and each ping pong ball in it has one of your identification numbers on it. We will draw the ping pong balls, two at a time, to determine who is matched with whom in each period.

After we have matched someone with you, we will write the other participant's number and their bid in columns (4) and (5) of your decision sheet. The relevant numbers for you and the other bidder will then be averaged to determine the value of the prize, which is entered in

column (6). Your earnings for the period are then calculated and entered in column (7). The low bidder earns \$0, and the high bidder in each pair earns the difference between the prize value and that person's bid. This difference, which can be either positive or negative, is then added to or subtracted from the previous cumulative earnings in column (8) to determine the new level of cumulative earnings for the high bidder. In the event of a tie bid, earnings or losses will be split between the two bidders.

Part A will consist of five periods. At the end of part A, we will distribute new decision sheets for part B, which will be the same as part A except that the die will be thrown three times for each bidder, with the first throw determining the number of dollars, and the second and third throws determining the number of pennies as before. Thus the range of prize values is about ten times as high in part B, and therefore, the possible gains and losses are larger.

Please keep track of your own cumulative earnings as you go. Notice that you begin period 1 with cumulative earnings of \$15. This is in addition to the \$6 already paid to you. Your cumulative earnings may go up or down during the course of the experiment. If your cumulative earnings fall below \$5 at the end of any period, you will be paid an amount that equals your cumulative earnings at that time (in addition to the \$6 participation payment that you already received earlier), and you will be released from the experiment. If your cumulative earnings level does not fall below \$5, then you will be paid your cumulative earnings at the end of the experiment today.

Final Remarks

At the end of today's session, we will pay to you, privately in cash, the amount that you have earned. You have already received the \$6 participation payment. Therefore, if you earn an amount X during the periods that follow, your total earnings for today's session will be $\$6.00 + X$. Your earnings are your own business, and you do not have to discuss them with anyone.

During the experiment, you are not permitted to speak or communicate with the other participants. If you have a question while the experiment is going on, please raise your hand and one of us will come to your desk to answer it. At this time, do you have any questions about the instructions or procedures?

The experiment is about to begin. You are not permitted to speak with other participants during the experiment. Please, don't ask public questions during the experiment. If you need to ask anything, just raise your hand and one of us will come to your desk to answer it.

(Instructions for other parts are available on request.)