ECONOMIC RESEARCH REPORTS

TWO STAGE AUCTIONS II: COMMON-VALUE STRATEGIES AND THE WINNER'S CURSE

BY

Steven J. Brams
and
Alan D. Taylor

RR # 92-02 January, 1992

C. V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003

Two-Stage Auctions II: Common-Value Strategies and the Winner's Curse*

Steven J. Brams
Department of Politics
New York University
New York, NY 10003

Alan D. Taylor
Department of Mathematics
Union College
Schenectady, NY 12308

*We thank Christopher Achen, Halina Frydman, Gary King, Dorothy E. Klotz, and Samuel Merrill, III for valuable discussions of ideas in this paper. Steven J. Brams gratefully acknowledges the support of the C. V. Starr Center for Applied Economics at New York University, and Alan D. Taylor gratefully acknowledges the support of the National Science Foundation under grant DMS-9101830.

Abstract

Two-stage auctions are defined by the following rules of play:

1. 1st stage: The players submit sealed bids, all of which are then opened and made public.

2. 2d stage: Each player chooses exactly one of the 1st-stage bids, either by affirming its own bid or by usurping another player's bid.

3. Payoffs: If there is only one player who makes the highest bid at the 2d stage, that player wins the auction, regardless of what it bid at the 1st stage. If there is more than one player who makes the highest bid at the 2d stage, the player who made the highest bid at the 1st stage wins the auction.

The "winner's curse," in which the highest bidder, because it wins, overpays, plagues common-value sealed-bid auctions, in which there is a single common value that an object is worth. A model for estimating this value in sealed-bid auctions, which requires that players take account of the curse and depreciate their valuations in making bids, is developed.

The curse is more easily averted in two-stage auctions. Because the highest bidder always has the opportunity to "bail out" in the 2d stage, and lower bidders to move up, players generally have an incentive to bid their valuations in the 1st stage. The 1st-stage bids can then be used to revise estimates of the common value, based on the median, mean, or a "weighted point estimate," which partially discounts outliers. Thereby the players can approximate the common value--less some reasonable profit--which benefits both them and the bid taker.

Two-Stage Auctions II: Common-Value Strategies and the Winner's Curse

1. Introduction

Two-stage auctions are defined by the following rules of play:

1. 1st stage: The players submit sealed bids, all of which are then opened and made public.

2. 2d stage: Each player chooses exactly one of the 1st-stage bids, either by affirming its own bid or by usurping another player's bid.

3. Payoffs: If there is only one player who makes the highest bid at the 2d stage, that player wins the auction, regardless of what it bid at the 1st stage. If there is more than one player who makes the highest bid at the 2d stage, the player who made the highest bid at the 1st stage wins the auction.

In Brams and Taylor (1991b), we analyzed private-value two-stage auctions, in which each player has a valuation of the object being auctioned off, based on private information, that does not depend on the values of other players: its valuation would not change if these other valuations became common knowledge. We showed that, under rather general conditions, the players would make "sincere" 1st-stage bids (i.e., bid their valuations, which are the amounts at which they would be indifferent beween winning and losing the auction), and the highest bidder would usurp the second-highest bid in the 2d stage and win.

This result duplicates the outcome of a Vickrey auction (Vickrey, 1961; Rothkopf, Teisberg, and Kahn, 1990), which is a sealed-bid auction in which the highest bidder wins but pays only the second-highest bid.

Two-stage auctions, we showed, are less vulnerable to cheating than Vickrey auctions and have the added advantage of allowing the players to decide--instead of having the procedure dictate--what bid will be chosen in the end.

In the present paper, we show that this latter advantage is far more significant in common-value auctions, in which there is a single common value that an object is worth (e.g., the price it commands in a market), but the players can only estimate its value. Such auctions are probably the norm; even art has its common-value aspects (McAfee and McMillan, 1987, p. 726).

We begin by modeling how players estimate common value in one-stage sealed-bid auctions. In section 2 we indicate how players might take account of the "winner's curse," in which a player, because it wins, overpays (Capen, Clapp, and Campbell, 1971; Thaler, 1988, 1992). To avert this curse--that is, the "bad news" about winning--we argue that sophisticated players will generally depreciate their sincere estimates in a process that stabilizes quickly (i.e., after each player factors into its calculation how other players revise their own initial estimates).

In section 3 we turn to two-stage auctions and show that sincere 1st-stage bids are not generally a Nash equilibrium. Which players have an incentive to deviate, however, depends on whether players use the median, mean, or some other measure of central tendency to estimate the common value.

For example, if the median is used, then the only players who will have an incentive to be insincere are those whose bids are at or above the median but not the top bid. On the other hand, if the mean is considered to be the best estimate of common value, other players (possibly all but the highest) may have an incentive to defect from their sincere bids.

We believe that both the median and mean are flawed as estimators of common value and propose a new estimator, the "weighted point estimate" (WPE) in section 4. WPE is neither as sensitive as the mean, nor as insensitive as the median, to outliers: the more extreme an outlier is, the less it counts in the estimate.

But WPE, as well as all other point estimates based on lst-stage bids, would appear to have the seeds of its own destruction built in, because one or more players will always be able to induce a preferrred outcome by bidding insincerely in the lst stage. Put another way, sincere lst-stage bids are not part of a Nash equilibrium, whatever procedure one uses to estimate the common value.

Yet this instability is not necessarily damning, because it is not biased in either direction. Although it would appear that players would always want to bid higher than their sincere estimates to be able to break ties in their favor in the 2d stage, we demonstrate that players may do even better by bidding lower.

The lack of a systematic bias, inducing the players either to overbid or to underbid in the lst stage, gives them good reason to be sincere in this stage. Although they can always determine, after their lst-stage bids are revealed, whether they could have benefitted by raising or lowering their bids, by then it is too late. Indeed, the very fact that players have incomplete information in the lst stage discourages them from trying, by being insincere, to be "too clever by half."

It is in the 2d stage, after the lst-stage bids are revealed, that the advantages of two-stage auctions are most striking. They not only give players an opportunity to "bail out" if they overbid, or usurp a higher bid if they underbid, but, more important, they give them a better opportunity to approximate the common value by seeing all the lst-stage bids at once.[1]

In section 5 we indicate how "split awards," without price discrimination and with the split made endogeneously, can be carried out in two-stage auctions. We conclude that experiments should be conducted with two-stage auctions in order to ascertain how players determine their lst-stage bids, how they make revisions in them in the 2d stage, and how the resulting outcomes compare with those under other auction procedures (Smith, 1987).

2. One-Stage Auctions: Taking Account of the Winner's Curse

Assume there are n players 1, 2, . . ., n in a sealed-bid auction. Their initial estimates of the common value of the object being auctioned off are $v_1(1), \ldots, v_1(n)$, where we use the subscript "1" to indicate that these estimates are only preliminary and will be revised later. If the auction is of an oil tract, for example, these estimates might

[1]The oral, ascending bids of English auctions also give players the opportunity to revise their bids, but not on the basis of seeing--at any one time before the end--all the bids together on the table. For recent reviews of the auction literature, see McAfee and McMillan (1987), Smith (1987), and Milgrom (1989). We first explored the idea of a 2d stage, though not in auctions, in Brams and Taylor (1991a).

be the best judgments of each company's geologists, based on their private information, of the tract's worth.

We assume that each company associates with its judgment a fraction α, where $0 \le \alpha \le 1$, which indicates its level of confidence in its geologists' estimate. The complementary fraction, $(1-\alpha)$, indicates its confidence in all the other companies' estimates, which are averaged and considered as a collectivity. For simplicity, we assume that all players have the same α, which may or may not be realistic.[2]

If $\alpha = 1$, the auction is one of independent private values. In this case, of course, the judgments of the other players do not affect a player's estimate.

It might be thought that these judgments also do not matter in a one-stage auction, because a player does not learn about the other players' bids--much less their geologists' estimates--until the auction is over, when it is too late for it to revise its own estimate. But, as we shall next show, simply an <u>appreciation</u> of the winner's curse can cause a player to revise its initial estimate and, as a consequence, alter its bidding strategy.

Suppose, initially, that each player bids <u>sincerely</u>--that is, it bids its initial estimate of the value of the oil tract--and player 1 wins with its bid of $v_1(1)$. Then player 1 knows that the other (n-1) bids--$v_1(2), \ldots, v_1(n)$--are all less than $v_1(1)$. This knowledge enables player 1 to make a revised calculation of the value of the tract, taking into account not only

---

[2]We shall consider the reasonableness of this assumption later in two-stage auctions, where we shall suggest different procedures for taking account of the estimates of others.

what its own geologists say but also what the geologists of the other (n-1) players must have said. To wit, player 1's revised estimate, denoted $v_2(1)$, is

$$v_2(1) = \alpha v_1(1) + (1-\alpha)\{[v_1(2) + \ldots + v_1(n)]/(n-1)\}. \tag{1}$$

To be sure, without making additional assumptions, player 1 does not know the values of the (n-1) terms in brackets on the right-hand side of (1). Nevertheless, assuming $\alpha < 1$, it can surmise that $v_2(1) < v_1(1)$. That is, if it bids its initial estimate of $v_1(1)$ and wins, it is paying more than the tract is worth, at least according to the collective wisdom of all the geologists involved. This is the winner's curse in our model: the winner is cursed by overpaying if it bids its initial estimate.

To avert the winner's curse, suppose that player 1 revises its initial estimate, based on the following thought experiment: it assumes, provisionally, that

(i) the other players bid sincerely [i.e., bid their estimates of $v_1(j)$ for j = 2, . . ., n]; and, for concreteness, that

(ii) these bids are uniformly distributed below its bid, should its bid be sincere and winning.

Then it can determine that the average of the other bids, given in braces in (1), is $v_1(1)/2$.[3]

[3]More generally, this will be the average if the bids are symmetrically distributed below $v_1(1)$. We make the more restrictive assumption of uniformity here because we use it later in the calculation of a revised estimate. Neither assumption is realistic, however, if bids tend to cluster (e.g., just below

Substituting this average into (1), player 1 will revise its initial estimate as follows:

$$v_2(1) = \alpha v_1(1) + (1-\alpha)[v_1(1)/2].$$

Because this calculation pertains to every player, we can delete the "(1)," obtaining

$$v_2 = \alpha v_1 + (1-\alpha)(v_1/2) = (v_1/2)(\alpha+1). \qquad (2)$$

This revised estimate, it would seem, will enable the players to avert the winner's curse--at least to the extent that any probabilistic considerations can provide help. Notice, in particular, the behavior of v_2 at the extremes: when $\alpha = 0$ (i.e., only the other players' bids matter), players will bid exactly one-half their sincere estimates; when $\alpha = 1$ (i.e., only one's own estimate matters), players will bid exactly this estimate. Because the estimates of players generally depend on both their own estimates and those of other players, this mixture will lead them to bid below their own initial estimates, according to (2).

But this finding depends on assumption (i)--that "the other players bid sincerely"--which is inconsistent with our assumption that player 1 will revise (i.e., lower) its v_1 estimate to v_2. One cannot have it both ways: either all players make a revised estimate v_2 (as assumed of player 1), or no players do.

---

player 1's winning bid), which is a phenomenon we discuss later in connection with two-stage auctions.

If no players do, we are back to our original estimates ($\alpha = 1$), and the winner's curse remains. If all players do, we must take this fact into account by amending assumption (i).

We do this by postulating that all players go through the same thought experiment as player 1. Thereby, we do not single out player 1 as special but instead assume that the (n-1) other players also revise their v_1 estimates to v_2's, and player 1 knows this.

This knowledge on the part of player 1 leads it to make a still higher-order estimate, whereby it revises $v_2(1)$. We call this revision $v_3(1)$; it assumes that every other player has made second-order estimates $v_2(j)$ for $j = 2, \ldots n$, which are then factored into player 1's third-order estimate of $v_3(1)$.

To calculate this estimate, we assume that player 1 gives a weight of α to v_1, not v_2, because v_1 is its geologist's estimate of true worth. But now player 1 must take into account that, if it wins, it does so because its bid exceeds the v_2's of the (n-1) other players, not their v_1's.

If all players bid their v_2's, and player 1 wins, the bids of the other players are no longer uniformly distributed over $[0, v_1(1)]$ but over the smaller interval $[0, v_2(1)]$. From (2) and the fact that this distribution is uniform, we can normalize the smaller interval by multiplying by the factor $[2/(\alpha+1)]$. Player 1's third-order estimate, based on its own first-order estimate and the second-order estimates of the other players, then becomes

$$v_3(1) = \alpha v_1(1) + (1-\alpha)[2/(\alpha+1)][\text{expected value of other players' bids, given player 1 wins with bid of } v_2(1)]$$

$$= \alpha v_1(1) + (1-\alpha)[2/(\alpha+1)][(1/2)v_2]$$
$$= \alpha v_1(1) + (1-\alpha)[1/(\alpha+1)]\{[v_1(1)/2](\alpha+1)\}$$
$$= \alpha v_1(1) + (1-\alpha)[v_1(1)/2] = v_2(1)$$

from (2). Lo and behold, the third-order estimate leaves the second-order estimate unchanged!

Thus, the whole process stabilizes once the players reach the third order of sophistication. Consequently, after the second order of sophistication, the players can rest assured that their downward revisions, as given by (2), probabilistically avert the winner's curse, and no further adjustments are necessary.

Is there a better way to avert this curse? We think, by providing players with information on other players' initial estimates, that there is.

3. Two-Stage Auctions: Median and Mean as Best Estimates

Two-stage auctions differ qualitatively from one-stage auctions in having a lst-stage that "does not count." Of course, this is not literally true, because the lst-stage bids provide the menu from which players select in the 2d stage as well as determine the winner if there is a tie in the 2d stage. And what gets on the menu, and who ties, certainly affect what gets chosen in the end.

Nevertheless, a player never has to pay what it bids at the lst stage, so it makes no binding commitment at this stage. Thus, if a player is the highest bidder at the lst stage, it can usurp any lower bid that it thinks might be a more reasonable estimate of the common value; if it wins, it will only have to pay this bid. Even if a player is the lowest bidder, it still is not committed, because affirming its own bid ensures that it loses, no matter what the other players do.

The evaporation of commitment at the lst stage would seem to afford every player P the opportunity to be sincere at this stage, rather than worry about the winner's curse and try to adjust for it in its initial bidding. After all, what is there to lose if P can always "bail out" at the 2d stage with a lower bid, even if it is the highest bidder at the lst stage, or move up and usurp a higher bid?

To investigate the robustness of sincerity at the lst stage, we need to make some assumptions about the common value of an object. Unlike one-stage auctions, we assume that P's initial estimate will not, in general, be determinative. On the contrary, we assume that the other bids that P observes at the 2d stage not only matter but also play a critical role in P's choice of a 2d-stage bid. After all, the estimate of P's geologists is only one of several; some of these might be just as credible, if not more credible, than P's.

To give the estimation problem greater structure, assume that players believe that the best estimate of the common value of an object is the median bid at the lst stage--that is, this is most likely to be its true value. Therefore, the most competitive yet profitable bid they can make at the 2d stage is the next-lower bid to the median.

The choice of the median as the best estimate in the 2d stage is arbitrary, even when all the other players are sincere in the lst stage. Moreover, the choice of the next-lower bid as an appropriate amount to bid in the 2d stage seems questionable in certain situations. For example, consider the two cases, Cl and C2, in both of which 95 is the median bid (underscored):

Cl: lst-stage bids are 99, 98, 97, 96, 95, 4, 3, 2, 1

C2: lst-stage bids are 99, 98, 97, 96, 95, 94, 3, 2, 1

In Cl, the 99-player would win at 4, whereas in C2 it would win at 94 if, as we provisionally assume, all players usurp (or affirm) the next-lower bid to the median.

Notice that these vastly different bids stem from only one different bid in the two cases. Intuition suggests that with this much on the line, players might invoke other estimation procedures, based on standards other than the median.

For example, what if the mean, which is 55 in Cl and 65 in C2, were the best estimate? As before, the 99-player would usurp 4 as its first profitable bid in C1; in C2, however, it would usurp 3, not 97. Even if one believed in the median as the appropriate standard, bidding 94 in C2 seems an extremely tough call because of its enormous significance for profits.

It is not surprising that the median and the mean may give different results. However, this problem is partially solved, as we shall argue in the next section, by choosing a more compelling standard on which to base one's bidding strategy.

Before pursuing these reasons, we first investigate implications of using the median of the 2d-stage bids as the best estimate of the common value of an object:

Theorem. Consider a two-stage auction in which all players use the following strategy:

> S: be sincere in stage I and usurp the highest bid that is below the median bid in stage II.

Assume the median sincere bid is the best estimate. Then the players who can benefit by a unilateral defection from S are precisely those whose sincere bids in stage I are strictly between the highest 1st-stage bid and the winning 2d-stage price (i.e., at least the median, but not the highest, 1st-stage bid). Moreover, all these players can equally benefit either by leapfrogging the lowest or leapfrogging the highest bid.

Proof. Suppose first that there are an odd number of bidders. Let m be the median bid, and let b be the bid immediately below m. We consider the effects of a unilateral defection by player P who bids x, which we let range over all the bids:

<u>Case 1</u>: P leapfrogs the lowest bid.

1.1. $x < b$. In this case, the highest bidder still wins at b.

1.2. $x = b$. If b is the lowest bid, this case is vacuous. If not, then the median is unchanged, but the bid b' immediately below b becomes the bid usurped by every player using S. Thus, P would have to usurp m in order to win (because there are no bids between b' and m once P has defected). But this win at m is the same as losing for P.

1.3. $m \leq x$ and x not the highest bid. In this case, b becomes the new median, so every player who uses S will usurp a bid below b. Thus, P can usurp b and win at a profit.

1.4. x is the highest bid. If P proceeds as in case 1.3, it wins at b exactly as it did with the highest bid.

Case 2: P leapfrogs the highest bid.

2.1. $x \leq b$. In this case, the bid above m becomes the new median, so every player will usurp m. P can win by usurping m, but this is the same as losing the auction.

2.2. $x = m$. In this case, the bid above m again becomes the new median, but now every player will usurp b. Thus, P (or the new highest bidder) can win by also usurping b.

2.3. $m < x$. If x is the highest bid, this case is vacuous. Otherwise, neither the median nor the price changes, but P becomes the new winner at b.

A similar argument applies if the number of bidders is even, with the median the average of the middle two bids. ▯

Corollary. In a two-stage auction with two players, S is a Nash equilibrium when the median (or the mean) is used as the best estimate of common value.

When there are more than two players, the fact that a player whose bid is the median or above (but not the highest) can, by departing either up or down, win suggests that there is no fundamental bias favoring upward or downward departures. To illustrate the Theorem, we consider a three-person example.

Let the set of 1st-stage sincere bids be {1, 3, 11}, so m = 3 is the median bid. Call the median bidder M. In stage I, M can insincerely leapfrog up to, say, 12; in stage II, it can then usurp 1, the bid immediately below the new median, m' = 11, and win. Alternatively, M can leapfrog down to 0, in which case m' = 1, and the 11-player, following S, will usurp 0 in stage II. But at this point M can usurp 1 and win.

Unlike M, the 11-player has no reason to depart from S when everybody else chooses this strategy, because it can win by usurping 1. Neither does the 1-player have reason to depart, because no insincere bid in stage I--either higher or lower than 1--can make it the winning bid in stage II at a lower price than 3, which is equivalent to losing the auction.

A comparison of the median with the mean as best estimator is instructive. In the two-person case, they give the same result: by the Corollary, strategy S is a Nash equilibrium.

Differences crop up between the median and the mean as best estimators when there are more than two players. Consider a strategy analogous to S when the mean is the best estimator:

> S': be sincere in stage I and usurp the highest bid that is below the median bid in stage II.

Returning to our previous example, the mean of {1, 3, 11} is 5. Hence, following S', the 11-player will usurp 3 (not 1, as in the case of the median) and win with a profit (5 - 3 = 2).

As in the case of the median, the 3-player, by insincerely bidding 12 in stage I, can benefit. The new mean will be 8, so following S' the insincere 3-player (now the 12-player) will usurp 1 and win with a profit (5 - 1 = 4). However, if 0 is a lower bound on bidding (i.e., no negative bids are allowed), the 3-player cannot win by underbidding. The reason is that by bidding 0 in stage I, the 3-player lowers the mean only to 4. Consequently, the 11-player will still usurp 3 and win, precluding a win for the insincere 3-player (now 0-player) at 1 (or 3).

But this apparent bias in favor of insincerely bidding up is simply an artifact of the particular numerical bids. If they were each increased by, say, 10, making them {11, 13, 21}, then if the 13-player insincerely bid 0 in stage I, its bid would lower the mean from 15 to 10 2/3. Then the 21-player, following S', would usurp 0 rather than 11, enabling the 13-player (now the 0-player) to usurp 11 and win with a profit (15 - 11 = 4).

The lesson in these cases is that insincere bidding in stage I may be profitable for the middle bidder, whether the best estimate is the median or the mean. By leapfrogging either the highest or the lowest bids in stage I, it can induce the other players, followintg S or S', to behave in a way that it can exploit. Indeed, our last example with the mean (and our earlier example with the median) demonstrated that the insincere middle bidder may even succeed in winning at the lowest sincere bid, giving it maximal profit.[4]

Whether the best estimator is the median or the mean, the lack of a Nash equilibrium for $n > 2$ would seem to jeopardize sincere bidding in stage I. As we have seen, however, it is by no means evident whether players, especially when information is incomplete (as we assume in

---

[4]If the mean is the best estimate, the <u>lowest</u> sincere bidder--not just the middle bidder--can also drag the mean below its own sincere bid by substantially underbidding in stage I. It can then win at the lowest sincere bid in stage II (i.e., its own) when the higher bidders usurp its lower insincere bid.

stage I), can do better by underbidding or overbidding. By Occam's razor, therefore, they may as well be sincere.[5]

This conclusion also holds for a new best-estimate procedure we shall propose in section 4. It may be viewed as a compromise between the insensitivity of the median, and the sensitivity of the mean, to outliers.

4. Two-Stage Auctions: Weighted-Point Estimates

So far we have shown that, using the median or the mean as a best estimator, sincerity is not in general optimal. Some players will have an incentive to dissemble, but it is not clear how, in a game of incomplete information, they will be able to ascertain whether they are in a position to benefit, and how to do so, when they make their lst-stage bids.

Even after players acquire information about the other players' lst-stage bids, it is also not clear that they will all rush to affirm or usurp the same bid just below the median or mean. If there are ten bidders, for example, and five bids are tightly clustered around the median or mean, which should a player choose? True, the highest lst-stage bidder has an advantage if there is a tie, but knowing this, the other players will have an incentive to choose one of the high bids in this cluster. But

---

[5]This may seem a perverse argument for sincerity, but it is one that has been persuasively made in the social-choice literature. That is, a voting system will induce sincerity if, though manipulable (as all voting systems are), voters have neither the incentive (e.g., because the calculations are too complicated to make) nor the ability (e.g., because of incomplete information) to effect outcomes that could benefit them.

then the highest bidder can anticipate this action and may also aim higher.

In light of this rather confusing picture at the 2d stage, consider the 1st stage. Perhaps the best recourse of a player is to be sincere--bid the initial estimate of its geologists in the case of the oil tracts. But as the bids for offshore oil tracts in Table 1 demonstrate, there may be little agreement on what the true values of these tracts are, assuming

---

Table 1 about here

---

these are sincere bids. If they are not --perhaps because they have been discounted according to a model like that presented in section 2--one would still not expect the estimates to be so wildly different unless the science of estimating oil potential is truly primitive.

This may have been so, at least as of the late 1960s. Moreover, these bids, which differ by a factor of as much as one hundred, were made before the winner's curse was first described by three Atlantic Richfield engineers, Capen, Clapp, and Campbell (1971).[6] In the auction of U.S. Treasury debt notes, the story is far different, where a spread of six basis points (hundredths of a percentage point) is considered very large, though it is only one-hundredth the value of the notes being auctioned off (Gilpin, 1991).

In a two-stage auction, once past the 1st stage the next problem is to come up with a reasonable point estimate of common value, like the

---

[6]It is perhaps not surprising that financial returns from these leases were quite dismal, strongly pointing to the existence of a winner's curse (Thaler, 1988, 1992).

TABLE 1. Bids (in Millions of Dollars) by "Serious Competitors" in Oil Auctions

Offshore Louisiana, 1967 (Tract SS 207)	Santa Barbara Channel, 1968 (Tract 375)	Offshore Texas, 1968 (Tract 506)	Alaska North Slope, 1969 (Tract 253)
32.5	43.5	43.5	10.5
17.7	32.1	15.5	5.2
11.1	18.1	11.6	2.1
7.1	10.2	8.5	1.4
5.6	6.3	8.1	0.5
4.1		5.6	0.4
3.3		4.7	
		2.8	
		2.6	
		0.7	
		0.7	
		0.4	

Source: Capen, Clapp, and Campbell (1971).

median or mean, based on the 1st-stage bids. Several methods, possessing different statistical properties, have been proposed for making point estimates, including trimmed means, linear combinations of order statistics (L-estimators), and, more generally, M-estimators (Lehmann, 1983; Goodall, 1983).

We propose here what seems to be a new method, which gives what we call a weighted point estimate (WPE). We will illustrate this estimate, which is defined inductively, with our earlier example. Then we will justify it as reasonable procedure to apply to 1st-stage bids.

WPE is a recursive function which associates to each finite set $\{x_1, \ldots, x_n\}$ a number, denoted WPE($\{x_1, \ldots, x_n\}$). This number represents a weighted point estimate, in which the fraction of weight given to, say, x_1 is inversely proportional to the distance d_1 from x_1 to WPE($\{x_2, \ldots, x_n\}$).

We could start with WPE($\{x_i\}$) = x_i and then proceed with the inductive step. However, it is easier to start with WPE($\{x_i, x_j\}$) = $(x_i + x_j)/2$, which is both the mean and the median; we illustrate the calculation of $n = 3$ with our example of bids $\{1, 3, 11\}$.

We begin by calculating the WPE of every pair of bids and the distance to the third bid:

WPE($\{3, 11\}$) = 7; $d_1 = 6$ (distance between $x_1 = 1$ and this WPE)
WPE($\{1, 11\}$) = 6; $d_2 = 3$ (distance between $x_2 = 3$ and this WPE)
WPE($\{1, 3\}$) = 2; $d_3 = 9$ (distance between $x_3 = 11$ and this WPE).

Define

$$\text{WPE}(\{1, 3, 11\}) = w_1(1) + w_2(3) + w_3(11), \tag{3}$$

where each weight w_i is proportional to $1/d_i$. To get weights that sum to 1, we must find a k such that

$$1 = w_1 + w_2 + w_3 = k[1/d_1 + 1/d_2 + 1/d_3],$$

or

$$k = 1/[1/d_1 + 1/d_2 + 1/d_3] = 1/[1/6 + 1/3 + 1/9] = 18/11.$$

We can now use k to find the w_i:

$$w_1 = k(1/d_1) = 3/11; \; w_2 = k(1/d_2) = 6/11; \; w_3 = k(1/d_3) = 2/11.$$

Substituting these w_i into (3), we obtain

$$\text{WPE}(\{1, 3, 11\}) = (3/11)(1) + (6/11)(3) + (2/11)(11) = 3\ 10/11.$$

This estimate of the common value lies between the median of 3 and the mean of 5. WPE partially discounts the outlier of 11, neither weighting it equally with other bids (as does the mean) nor discounting it totally (as does the median, in which only its position, but not its numerical value, has an effect).

The algorithm illustrated in our 3-player example generalizes to any finite number of players. Thus, if there were a 4th player, one would compute the four 3-player WPEs in the manner previously illustrated for bids of {1, 3, 11}. Then the distances of each player's bid from the WPE of the three other players can be determined, from which the appropriate k and the w_i can be found to compute the WPE for all four players.

We have, in fact, made this calculation if there is a fourth player who bids 5 (the mean bid of the first three and, of course, for all four players as well). The addition of this bid gives WPE({1, 3, 5, 11} ≅ 4.20, somewhat raising WPE({1, 3, 11}) ≅ 3.91. Both WPEs are above their respective medians of 3 and 4, because the outlier 11 has some upward "pull" in each case.

WPE effectively makes the α in the section 2 model of one-stage auctions endogeneous. That is, instead of players' attaching some predetermined weight to their own bids versus the bids of the other players, the weights are determined in the recursion, with more weight given to bids that are closer to other bids, based on the inverse distance function. This seems to us proper: if a player's own bid is an outlier, it probably <u>should</u> be weighted less than if it were in a cluster with the other bids.[7]

That a player does not know beforehand whether or not it will be an outlier is a good reason for <u>not</u> postulating it to have an a priori distribution over the common value, and then doing Bayesian updating once the bids of the other players are revealed. Judging from the very disparate bids for oil tracts, a priori assumptions may well be massively invalidated by the 1st-stage bids and hence may be imporper to assume.

---

[7]Of course, the w_i's could be made proportional to some other function of the d_i's, such as the squared inverse if one wished to discount outliers even more.

We therefore prefer to use the lst-stage bids to make a "best" estimate of the common value, letting one's own bid, as it were, speak endogeneously. Attaching special signficance to one's estimate, much less postulating an a priori distribution in order to make Bayesian calculations, seems to us unhelpful unless a player has private information that the lst-stage bids do not reflect.

It is important to bear in mind that no purely algorithmic or statistical approach is going to answer the question of exactly what to bid in either the lst or 2d stage of a two-stage auction. We previously argued that, in the face of incomplete information and the lack of a systematic bias to underbid or overbid, lst-stage bids are likely to be sincere.

At the 2d stage, the key question becomes how to use the information about bids revealed in the lst-stage. We have recommended a point estimation procedure that discounts, but only partially, outliers. It is only a starting point, to which strategic and other considerations can then be appended to give a more rounded estimate.

5. Conclusions

The great advantage of two-stage auctions, in our view, is that they provide players with <u>relevant</u> new information in the 2d stage. To be sure, this information may leave players in a quandary, as we suggested with our first example in section 3, but we do not believe that dilemmas of this kind undermine the procedure. Rather, they mirror the inherent difficulties of bidding in auctions, which is a risky business.[8]

A good auction procedure should enable players to exploit as much relevant information as possible before settling on a bidding strategy. Sealed-bid auctions, because they occur in only one stage, are flawed in this regard. By permitting only a priori calculations of the kind modeled in section 2, they deny players feedback on other players' valuations, on the basis of which they can revise.

Sophisticated players today probably have at least an intuitive understanding of the winner's curse (McAfee and McMillan, 1987, p. 730). Whether they depreciate their initial estimates in the manner we described or discount in other ways, there is good reason to believe that their bids are not sincere. While Vickrey auctions offer one solution to this problem, they lead to certain difficulties in the private-value case that we argued two-stage auctions ameliorate (Brams and Taylor, 1991b).

In common-value auctions, Vickrey auctions offer no help to players in acquiring additional information on which to base a common-value estimate. Although they may relieve players of the need to discount their initial estimates, they do not enable them to incorporate into these estimates information about the estimates of other players.

To be sure, if there are more than two players, sincere 1st-stage bids are not a Nash equilibrium in two-stage common-value auctions. But because of incomplete information, players cannot be sure whether

---

[8]The fact that oil drilling at offshore tracts has turned up mostly dry wells indicates that oil companies do indeed take recognizable risks (Thaler, 1988, 1992).

they are in a position to do better by bidding higher or lower than their sincere estimates. Not knowing whether underbidding or overbidding is preferable, we think that most players will be sincere--or, at worst, departures from sincerity will not be systematically biased.

Sincerity is reinforced by the fact that, unlike a Vickrey auction, the highest bidder need not pay the next-highest bid but can, instead, opt out or at least drop considerably lower. Or a low bidder can usurp a higher bid. Giving players such a choice, we believe, is not only sensible but also encourages sincerity.

The difficulties of making point estimates at the 2d stage were illustrated with examples. We argued that neither the median nor the mean was ideal and instead proposed a "weighted point estimate" (WPE), defined recursively. Players presumably want to discount outliers, but not totally, in making their 2d-stage bids, and WPE does this in a natural way, making the weightings endogeneous.

The new information that the revelation of lst-stage bids provide, especially if the bids are sincere, makes all players better informed about each other's initial estimates. In this sense, it is a public good.

More specifically, such information better enables players to avert the winner's curse. While a Vickrey auction partially alleviates this problem in a one-stage auction, if there are two high outliers among, say, ten bidders, the highest bidder may still end up overpaying at the second-highest bid, whereas in a two-stage auction such players might usurp the third-highest bid. In any event, players can make a choice in two-stage auctions.

Auctions that allow for "split awards," in which a divisible good is apportioned among two or more bidders, have recently been proposed (Anton and Yao, 1989; Klotz and Chatterjee, 1991). We believe that they might be better wedded to two-stage auctions than one-stage auctions. While splitting an award in a predetermined ratio between the two top bidders reduces the risk for each in a one-stage auction, it raises the problem of price discrimination: one player receives a portion of the award at a lower price than the other.

By contrast, with two-stage auctions, if there is a tie at the 2d stage, one can make an award to the tied players at the <u>same</u> price. Moreover, instead of fixing the proportions exogeneously, one can determine them endogeneously, based on the 1st-stage bids. Presumably, the highest 1st-stage bidder would win a greater portion of the award than lower 1st-stage bidders who tie at the 2d stage (we leave open the function of their 1st-stage bids that might be used). Two-stage auctions, then, would become a special case of two-stage split auctions, in which the split is 100 percent to the highest bidder.

We conclude that two-stage auctions, whether split or not, offer a promising alternative for players to approach, through their bidding, the common value of objects. We have not resolved all ambiguities at either the 1st stage (will players be sincere?) or the 2d stage (how will they estimate the common value?), which experimentation with such auctions should help to clarify. It would also be useful to compare two-stage outcomes with outcomes obtained experimentally under other procedures (Smith, 1987).

The new information that the 2d stage provides seems to us beneficial for both the players and the bid taker. The bid taker benefits if, indeed, the 1st-stage bids are sincere and not depreciated because of the winner's curse. The players benefit if the winning 2d-stage bid approximates the common value--less some reasonable profit--because then the procedure will be viewed as fair, legitimating the "rules of the game."

References

Anton, J., and D. Yao (1989). "Split Awards, Procurement, and Innovation." RAND Journal of Economics 20, no. 4 (Winter).

Brams, Steven J., and Alan D. Taylor (1991a). "Three Solutions to Divide the Dollar." Preprint, New York University.

Brams, Steven J., and Alan D. Taylor (1991b). "Two-Stage Auctions I: Private-Value Strategies." Preprint, New York University.

Capen, E., R. Clapp, and W. Campbell (1971). Competitive Bidding in High-Risk Situations." Journal of Petroleum Technology 23, no. 1 (June): 641-653.

Gilpin, Kenneth N. (1991). "3-Year Note Auction Goes Badly." New York Times (November 6), pp. D1, D15.

Goodall, Colin (1983). "M-Estimators of Location: An Outline of the Theory." In Understanding Robust and Exploratory Data Analysis, ed. David C. Hoaglin, Frederick Mosteller, and John W. Tukey. New York: Wiley, pp. 339-403.

Klotz, Dorothy E., and Kalyan Chatterjee (1991). "Variable Split Awards in a Single-Stage Model." Working Paper Number 91-6-1, Graduate School of Business Administration, Fordham University (January).

Lehmann, E. L. (1983). Theory of Point Estimation. New York: Wiley.

McAfee, R. Preston, and John McMillan (1987). "Auctions and Bidding." Journal of Economic Literature 25 (June): 699-738.

Milgrom, Paul (1989). "Auctions and Bidding: A Primer." Journal of EconomicPerspectives 3, no. 3 (Summer): 3-22.

Rothkopf, Michael H., Thomas J. Teisberg, and Edward P. Kahn (1990). Why Are Vickrey Auctions Rare?" Journal of Political Economy 98, no. 11 (February): 94-109.

Smith, Vernon L. (1987). "Auctions." The New Palgrave : A Dictionary of Economics, vol. 1. London: Macmillan, 1987, pp. 138-144.

Thaler, Richard H. (1988). "Anomalies: The Winner's Curse." Journal of Economic Perspectives 2, no. 1 (Winter): 191-202.

Thaler, Richard H. (1992). The Winner's Curse: Paradoxes and Anomalies of Economic Life. New York: Free Press, 1992.

Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." Journal of Finance 16, no. 1 (March): 87-37.