



Counterspeculation, Auctions, and Competitive Sealed Tenders

William Vickrey

The Journal of Finance, Vol. 16, No. 1 (Mar., 1961), 8-37.

Stable URL:

<http://links.jstor.org/sici?sici=0022-1082%28196103%2916%3A1%3C8%3ACAACST%3E2.0.CO%3B2-L>

The Journal of Finance is currently published by American Finance Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/afina.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

COUNTERSPECULATION, AUCTIONS, AND COMPETITIVE SEALED TENDERS

WILLIAM VICKREY

Columbia University

IN HIS *Economics of Control*, A. P. Lerner threw out an interesting suggestion that where markets are imperfectly competitive, a state agency, through "counterspeculation," might be able to create the conditions whereby the marginal conditions for efficient resource allocation could be maintained. Unfortunately, it was not made clear just how this counterspeculation was to be carried out, and to many this term denotes just one more of the empty boxes that rattle around in the economist's cupboard of ideas. And there appears to have been, in the years since *Economics of Control* first appeared, no attempt to examine critically just what this intriguingly labeled box might in fact contain.

In Section I this counterspeculation box will be further examined; it turns out that most of the devices that most immediately suggest themselves under this heading prove to be inordinately expensive in terms of their demands on the fiscal resources of the state relative to the net benefits to be realized, at least where the commodity in question is finely divisible. The other extreme case, where there is only a single indivisible item to be allocated, is examined in Section II; in this case the possibilities for reaching an optimum solution in a market with a limited number of participants become considerably brighter: the common or progressive type of auction can be shown to provide better chances for optimal allocation than the regressive or "Dutch" auction. The implications of these findings for the more significant cases where contracts are let or sales made by competitive bids or tenders are examined in Section III; the analysis reveals a likelihood that certain modifications of current practices in these areas, more specifically by making the award price equal to the second highest (or lowest) bid price rather than the highest bid price, might prove generally beneficial in improving the allocation of resources without being as prejudicial to the interests of sellers (or buyers) as might at first seem to be the case. Section IV deals with the somewhat more complicated and general class of cases where there are several identical items to be auctioned, and Section V deals with the application of the concepts derived in Section IV to the sale

of a number of identical units under sealed-bid conditions; it turns out that here, too, significant gains can be expected from certain departures from currently prevalent practices.

I. THE EXCLUSIVE PUBLIC MARKETING AGENCY

To simplify the problem, let us consider the simple case of a standardized commodity in which the only imperfection in the market consists of the fact that either buyers or sellers or both are too few in number to ignore the repercussions of their actions on the market price but are either too numerous, too naïve, or too isolated from each other to engage in any overtly or tacitly concerted action. We will also assume that the individual marginal-cost and marginal-value curves of the sellers and buyers are well defined and have moderate positive and negative slopes, respectively. The normal result in such a case is that less than the optimal quantity will be produced and sold, and this will be true even though a "countervailing power" type of balance between buyers and sellers maintains the price at the same level as would result under strictly competitive conditions.

Let us now assume that there is established an exclusive public marketing agency to which all sales of this commodity must be made and from which all supplies of the commodity must be bought. A simple solution to the problem would be available if the public marketing agency could determine with confidence what the equilibrium competitive price would be and could then establish this price for its purchases and sales in such a way that neither buyers nor sellers could expect to have any influence over it. This price would then be a fixed datum to buyers and sellers, and competitive behavior could be expected. This is, indeed, the type of solution that comes most readily to mind on first meeting up with the concept of "counterspeculation."

The trouble with this as a workable solution is that much of the information that the marketing agency would need in determining the competitive equilibrium price would have to come from reports and actions of buyers and sellers, who would have an incentive to understate prospective demands and supplies or to curtail their actual sales and purchases in the hope of inducing the marketing agency to change the price in their favor. In a static situation, a marketing agency might conceivably manage not to be misled by such misinformation and to withstand the blackmail of curtailments in purchases or sales. But in a dynamic situation, where the equilibrium price is continually changing, it would be much more difficult to ascertain the

equilibrium price, keep the price at this level, and simultaneously persuade buyers and sellers that future changes in the published price will not be influenced by any tactical deviations on their part. Moreover, if the marketing agency should, under the guise of "stabilization" or otherwise, attempt to keep the price fixed over any extended period of time, even slight disequilibria at this price are likely to induce speculation against the pegged price, which, even if it does not succeed in inducing a change in the pegged price, will, unlike speculation over a market-determined price, necessarily of itself involve some misallocation of resources.

What the marketing agency needs, in order to determine the optimum pattern of transactions in its commodity, is an unbiased report of the marginal-cost (= competitive supply) curves of the sellers and of the marginal-value (= competitive demand) curves of the purchasers, or at least of the portions of these curves covering a range of prices that will be sure to contain the equilibrium price. The problem is then for the marketing agency to behave in such a way as to motivate the buyers and sellers to furnish such unbiased reports. One method, though an expensive one, is to arrange to purchase the commodity from suppliers and to sell it to purchasers on terms that are dependent on the reported supply and demand curves in such a way that the suppliers and purchasers will maximize their profits, individually at least, by reporting correctly, so that any misrepresentation will subject them to risk of loss (or at least offer no prospect of gain).

For example, the marketing agency might ask for the reporting of the individual demand and supply curves on the understanding that the subsequent transactions are to be determined as follows: The agency would first aggregate the reported supply and demand curves to determine the equilibrium marginal value, and apply this value to the individual demand and supply curves to determine the amounts to be supplied and purchased by the various individual buyers and sellers. The amount to be paid seller S_i would, however, somewhat exceed the amount calculated by applying this marginal value to his amount supplied; in effect for the r th unit supplied, S_i would be paid an amount equal to the equilibrium price that would have resulted if S_i had restricted his supply to r units, all other purchasers and sellers behaving competitively. In terms of Figure 1, D_n is the aggregate demand curve, S_n is the total supply curve, and the intersection at E indicates the equilibrium marginal value; S_{n-i} is the aggregate supply curve of the sellers other than S_i , and its intersection with the horizontal line PE at G indicates the amount GE to be supplied by i ;

the amount to be paid to i for this supply is indicated by the area $EFGMQ$, between the total demand curve D_n , the supply curve of the competitors S_{n-i} , and the quantity axis. Payments to other suppliers would be determined similarly, each supplier being considered in turn as the "last" supplier, giving rise to a total payment to suppliers that can be represented by the area $OQEFGF'G''G'' \dots PO$.

Given this method of determining payment, no individual supplier would have any direct incentive to misrepresent his supply schedule.

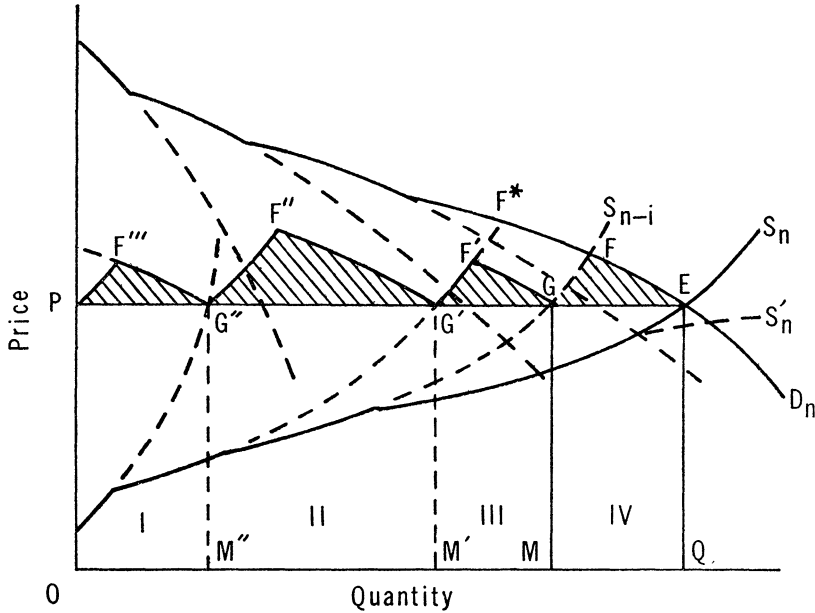


FIG. 1.—Counterspeculative payments to suppliers

If he misrepresents his true supply curve in a way that makes the total supply curve go through the correct point E , this will cause no change in the amount he is called upon to supply or to his own receipts, though it is likely to affect the receipts of (but not the demands on) other suppliers. On the other hand, if a misrepresentation is made that causes the aggregate supply curve to miss the point E —say, by falling below it as in the curve S'_n —this will entail his being called upon to supply additional units for which the marginal revenue received by this supplier will be less than the marginal cost to him of producing this increment of output, so that, on balance, he will be worse off. Conversely, for errors that would push the supply curve above the correct equilibrium point E , the amount that will be

ordered from him will be reduced, but the amount paid for his supply will be cut by more than the marginal-cost savings to him, and again he would be worse off, on balance.

An exactly symmetrical method could simultaneously be adopted for dealing with the demand side of the market: purchasers would receive the indicated equilibrium amounts of the commodity but would pay a price represented by the area $EKLNQ$ in Figure 2, determined by the aggregate supply curve S_n and the demand curve

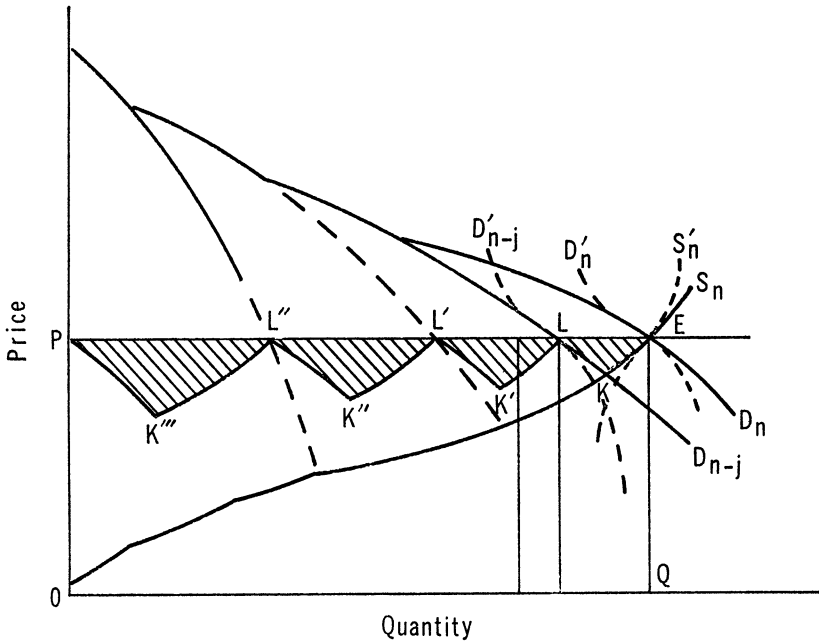


FIG. 2.—Counterspeculative charges to purchasers

D_{n-j} of the purchasers other than j . Again the individual purchasers would have no direct incentive for misrepresenting their demand schedule and would have a positive incentive for insuring that at the equilibrium value their demand would be correctly reported.

Since, in advance of filing their supply or demand schedules, suppliers and purchasers will in general be somewhat uncertain of the exact level of the eventual equilibrium price, incentives to report correctly at the equilibrium price actually cover a considerable range about this equilibrium price. On the other hand, if traders have a fairly confidently held expectation that the equilibrium price will fall within a certain narrow range, there may be an indirect community of

interest in shading the reported demand and supply curves outside this range in the direction of greater inelasticity, as indicated by the dotted curves S'_{n-i} , S'_n , D'_{n-j} , D'_n . Such shading for prices above the equilibrium value would result in higher payments to other suppliers, and below the equilibrium value would result in lower charges to other purchasers; no supplier or purchaser would benefit directly from his own misrepresentation, however, and optimum allocation would still be preserved.

The basic drawback to this scheme is, of course, that the marketing agency will be required to make payments to suppliers in an amount that exceeds, in the aggregate, the receipts from purchasers by the sum of the shaded areas in the two diagrams. The average price paid to suppliers will exceed the competitive equilibrium price OP , and the average price received from purchasers will be less than OP . This solution would indeed permit optimum allocation of resources to be achieved if there were a source of public funds that was without adverse influence on resource allocation in other directions. Even if such an ideal revenue source existed, such a scheme would still be open to criticism as discriminating in favor of larger units, since they would be obtaining a higher average price as producers and a lower average price as purchasers than would their smaller competitors. In Figure 1, for example, supplier *II* gets the amount $M''G''F''G'M'$ for his supply of $M''M'$, whereas supplier *III* gets $M'G'F'GM$ for a supply of $M'M$. Moreover, there remains under the scheme a positive incentive for firms to merge into larger units for the sake of obtaining more favorable treatment. Thus in Figure 1, if suppliers *III* and *IV* were to merge, they would automatically become entitled to an additional payment indicated by the quadrilateral area $F'F^*FG$. Considering the fact that public funds are obtainable only at a significant cost in terms of overhead expenses of collection as well as of misallocation of resources at other points, it is highly doubtful whether the carrying-out of such a scheme in full could ever be justified.

It is tempting to try to modify this scheme in various ways that would reduce or eliminate this cost of operation while still preserving the tendency to optimum allocation of resources. However, it seems that all modifications that do diminish the cost of the scheme either imply the use of some external information as to the true equilibrium price or reintroduce a direct incentive for misrepresentation of the marginal-cost or marginal-value curves. To be sure, in some cases the impairment of optimum allocation would be small relative to the reduction in cost, but, unfortunately, the analysis of such variations

is extremely difficult; considering the slight likelihood that any such scheme would be put into practice, further analysis will here be concentrated on situations more closely realized in practice.

II. SIMPLE AUCTIONS

Where the resource to be allocated comes in a small number of discrete indivisible units rather than consisting of a fungible commodity, the chances of insuring an optimum allocation of resources are considerably improved, since here there will be in general a certain range of prices, all leading to the same optimal allocation of the resources being traded, though, of course, to somewhat different distributions of income among the parties to the trade. The simplest case is one in which there is a *single unique indivisible object* to be sold to one of a number of potential purchasers. The purchaser and the price may be determined by any of a number of auctioning or bidding procedures, and the results will in general be significantly different according to the procedure adopted.

The simplest procedure to analyze is that of the ordinary or progressive auction, in which bids are freely made and announced until no purchaser wishes to make any further higher bid. The normal result (among rational bidders!) is that the bidding will stop at a level approximately equal to the second highest value among the values that the purchasers place on the item, since at that point there will be only one interested bidder left; the object will then be purchased at that price by the bidder to whom it has the highest value. (For simplicity, we shall assume that price can vary continuously and that there is no minimum increment between bids.) This result is obviously Pareto-optimal.

It is interesting to contrast this with the so-called "Dutch auction," in which the auctioneer announces prices in descending sequence, the first and only bid being the one that concludes the transaction. A mechanized form of this procedure is actually used in wholesale flower marketing in the Netherlands and has proved very economical of time and effort. In the analysis of this form of auction, however, we find that we are faced with what is essentially a "game" in the technical sense. Each bidder, in attempting to determine at what point he should be prepared to make a bid so as to obtain the greatest expectation of gain, will need to take into account whatever information he has concerning the probable bids that might be made by others, and the bids made by others will in turn depend on their expectations concerning the behavior of the first bidder. To put in a bid as soon as the price has come down to the full value of the object

to the bidder maximizes the probability of obtaining the object, but guarantees that the gain from securing it will be zero; as the announced price is progressively lowered, the possibility of a gain emerges, but as the gain thus sought increases with the lowering of the point at which a bid is to be made, the probability of securing this gain diminishes. Each bidder must thus attempt to balance these two factors in terms of whatever knowledge he has concerning the probable bids of the others.

The Dutch auction game.—To make this problem tractable, we shall suppose that the knowledge that each bidder has about the motives and probable behavior of the others can be derived from a set of probability distributions from which the value of the object to each of the bidders is conceived to be drawn. For simplicity, we shall assume that all bidders have the same conception of the probability distribution from which any given player is deemed to derive the value he places on the object; the given player, of course, knows the actual value he places on the object but is assumed also to know the distribution from which others consider his value to be drawn. These distributions need not be the same for all bidders, i.e., some may be considered to be more likely to place a high value on the object than others.

The situation can be considered analogous to a formal parlor game in which each player actually does draw a value at random from his own individual probability distribution of values (e.g., by drawing a card at random from a deck having a composition known to all the players), and then, knowing that value, but without directly revealing it to others, and knowing the probability distributions from which the other players are respectively drawing their values, but not knowing any of the values actually drawn, each player makes a bid without knowing the bids made by any of the other players; after all bids have been made, they are compared, and the player with the highest bid wins an amount equal to the excess of his value over his bid, the remaining players winning nothing. In the case of tie bids we can assume the tie to be broken by a random drawing giving each tied player the same probability of winning. Given a specific set of underlying probability distributions, the problem is then to determine how each bidder should determine his bid in order to maximize his gains. Unfortunately, the general theory of games gives an embarrassingly rich set of answers.

However, if we rule out such elements as collusion among bidders, side payments, communication, or signaling, the solution most appropriate to the situation being studied seems to be the non-co-op-

erative equilibrium point analyzed by Nash.¹ This equilibrium point is defined in this case as a set of functions $x_i(v_i)$, one such function for each player, relating the bid x_i to be made to the value v_i that he draws from his list, such that if any one player is able to determine (e.g., from observation over a sufficiently large number of plays of the game, or repetitions of analogous situations, or a study of an analysis such as this one!) just what functions are being used by the other players, or at least the resulting probability distribution of bids for each of the other players, which is all that concerns him directly, and considers these distributions to be fixed, at least for the time being, he can nevertheless find no way of changing his own function $x_i(v_i)$ in such a way as to increase his expected gain.

Unfortunately, while it can be shown that at least one such equilibrium point exists in each case and that we are not hunting a will-o'-the-wisp, the analytical determination of such equilibrium points involves extremely difficult mathematics in all but the simplest cases. However, some instructive hints can be derived from the examination of the more tractable cases.

The homogeneous rectangular case.—One simple case is that in which all of the individual values are drawn from the same rectangular distribution, which, by suitable choice of scale and origin, we can make the interval $(0, 1)$. Also we shall assume a linear utility function over the range of gains involved, so that we can speak in terms of maximizing expected money gains rather than having to allow for any "risk aversion" or "risk preference" that might be represented by a non-linear utility function. In this case the unique equilibrium strategy is for each player to determine his bid according to the relation

$$b_i = \frac{N-1}{N} v_i,$$

where N is the number of bidders. It is fairly easy to see that if all players behave in this way, no one player can gain by deviating from this pattern. It will be shown later, as part of the solution of a more general case, that this equilibrium point is indeed unique. If players conform to this norm, the highest bid will always be made by the player drawing the highest value for the object, so that the result will be Pareto-optimal (the seller being included, of course, among those to be preserved against loss in any proposed reallocation that would contradict this optimality).

1. See, e.g., Luce and Raiffa, *Games and Decisions* (New York: John Wiley & Sons, 1957), pp. 170-73.

It can also be shown (see Appendix I), that the two methods of auctioning produce, in this case, the same average expected price and hence the same average expected gains to the buyers and sellers, respectively. However, the variance of the price is greater under the common or progressive type of auction than with the Dutch auction by a factor of $2N/(N-1)$ while the variance of the gain to the buyer is greater by a factor of N^2 . If we introduce an element of risk aversion for purposes of evaluating these results (but not for deriving the results!), the Dutch auction thus proves slightly superior by reason of the smaller dispersion of the gains to each of the parties.

Non-homogeneous cases.—If the assumption of homogeneity among the bidders be abandoned, the mathematics of a complete treatment become intractable. It is fairly easy to show, however, that while the progressive auction still produces the Pareto-optimal allocation of the object, the Dutch auction will not, in general. Consider, for example, the case where there are two bidders, one drawing from a distribution rectangular between 0 and 1, the other drawing from a rectangular distribution ranging from a to b (see Appendix II). If $a \neq 0$, the essential asymmetry of the positions of the two players prevents their reacting similarly to similar value drawings, and the object may go to the bidder for whom it has the lower value.

An attempt at a complete solution for the above case runs into difficulties. In order to have at least one complete analysis of a non-homogeneous case, we can further simplify by supposing that the value to be drawn by the second bidder is fixed at a , rather than varying over a range from a to b . In this case we must allow the second bidder to use a mixed or randomized strategy, so as to distribute his bids over a range, even though he has only a fixed value, since if he were always to bid the same amount c , then bidder 1 would tend to just outbid him whenever v_1 is greater than c , giving bidder 2 in turn an incentive to change his fixed bid, and no equilibrium would be possible. On the other hand, for a bidder such as 1, whose values are drawn from a continuous distribution of positive range, there is no need for any randomization, and the bids can be determined as a single-valued function of the value drawn.

The analysis of this case is taken up in detail in Appendix III. The equilibrium strategies are illustrated in Figure 3, which gives the results for $a = 0.8$. The bidder with the fixed value a will have to distribute his bids over the range from $a/2$ to $a - \frac{1}{4} a^2$ according

to the cumulative frequency distribution $y_2(x)$, y_2 being the probability of a bid of less than x by player 2, where

$$y_2(x) = \frac{a(2-a)}{2(2x-a)} e^{[2/(2-a) - a/(2x-a)]}$$

The bidder with the variable value makes a bid x depending on the value v_1 that he draws given by

$$x = a \left(1 - \frac{a}{4v} \right)$$

the curve $y_1(x)$ in Figure 3 shows the cumulative probability of the various bids x , and the value corresponding to any given bid is obtained by looking for the point on the line $y_1 = v$ where the cumulative probability of the value corresponds to the cumulative probability of the given bid. If a value is drawn of less than $\frac{1}{2}a$, there is no possibility of gain for bidder 1, since bidder 2 always bids at least this amount; in this event bidder 1 may make bids falling anywhere within the shaded area without upsetting the equilibrium.

The difference between the mean results for the Dutch auction and the common progressive auction procedure in situations of this

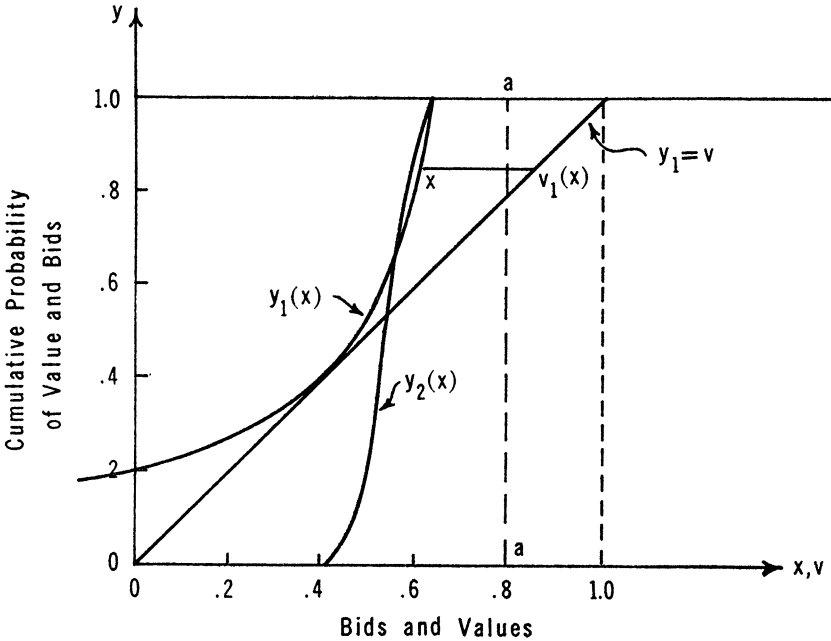


FIG. 3.—Equilibrium strategies for the asymmetrical case with $a = 0.8$

type varies according to where the second bidder's fixed value a lies in relation to the range of values over which the randomly assigned value of the first bidder may vary. In the common auction, for values of v_1 greater than a , bidder 1 gets the article at a price of a , while for a v_1 less than a , bidder 2 gets the article at a price of v_1 . The average price is thus $a - \frac{1}{2}a^2$. With the Dutch auction the average realized price turns out to be greater than this for a greater than about 0.43, and less for values below about 0.43. Figure 4 gives a general picture of how the average expected values for the price, the net gain to the two bidders, and the total gain differ for the two methods at various values of a .

To extrapolate rather boldly from these instances, one can perhaps hazard the guess that where the bidders are fairly homogeneous and sophisticated, the Dutch auction may produce results that are

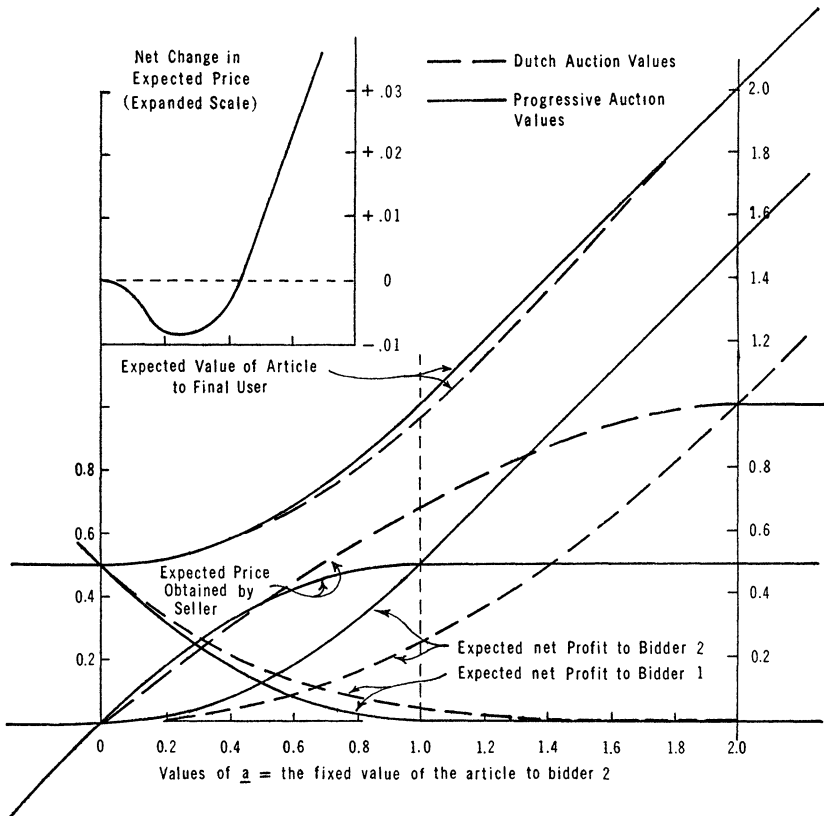


FIG. 4.—Comparison of average expected results of the Dutch auction with those of the progressive auction in an asymmetrical case.

reasonably close to the Pareto-optimal, but where there is much variation in the state of information or the generally expected intensity of desire of the various players for the object, or where the bidders are insufficiently sophisticated to discern the equilibrium-point strategy or for some other reason fail to use this strategy, then the Dutch auction is likely to prove relatively inefficient from the point of view of securing an optimum allocation. In the symmetrical case, the Dutch auction produces the same average price and the same average gain, but with a smaller dispersion of the gains to the bidders and of the realized price.

III. SALE OR PURCHASE OF A SINGLE LOT BY SEALED BIDS

A considerably greater degree of importance attaches, not to auctions as such, but to cases in which a contract is to be let on the basis of competitive sealed tenders. This may be a construction contract or the sale of a parcel of property or the underwriting of a security issue. Actually, the usual practice of calling for the tender of bids on the understanding that highest or lowest bid, as the case may be, will be accepted and executed in accordance with its own terms is isomorphic with the Dutch auction just discussed. The motivations, strategies, and results of such a procedure can be analyzed in exactly the same way as was done above with the Dutch auction.

Since it has been shown that the Dutch auction has certain characteristics in some circumstances that may be considered disadvantageous as compared with the more certainly Pareto-optimal results of the progressive auction, it is of interest to inquire whether there is not some sealed-bid procedure that would be logically isomorphic to the progressive auction. It is easily shown that the required procedure is to ask for bids on the understanding that the award will be made to the highest bidder, but on the basis of the price set by the second highest bidder. If this procedure is carried out, then the optimal strategy for each bidder (assuming, as is indeed necessary in the analysis of the progressive auction itself, the absence of collusion among bidders) will obviously be to make his bid equal to the full value of the article or contract to himself, i.e., to the highest amount he could afford to pay without incurring a net loss or to that price at which he would be on the margin of indifference as to whether he obtains the article or not. Bidding less than this full value could then only diminish his chances of winning at what would have been a profitable, or at least not unprofitable, price and could not, collusion aside, affect the price he would actually pay if he were the successful bidder. Bidding more than the full value, on the

other hand, would increase his chances of winning, but only under circumstances that would involve him in an unprofitable transaction, the price to be paid being greater than his value.

Judging from the preceding analysis of the Dutch auction paradigm, there would in many, if not most, cases be a considerable advantage to all concerned in shifting to this "second-price" method of handling sealed bids. In cases in which, by reason of asymmetry among the bidders, errors in evaluation, or mistakes in strategy, the result with the "top-price" method is non-optimal, a change to the "second-price" method will yield an increase in the aggregate profits to be shared among seller and buyers. The fact that in the symmetrical case when the correct equilibrium strategy is employed there is no change in the average realized price would tend to indicate that when there is a gain through a change from a non-optimal to an optimal one, the gain would, on the average, be shared between buyers and sellers. On the other hand, the study of the asymmetrical case indicates that there are some extremes where the Pareto-optimal progressive auction or second-price methods result in a lower average expected price to the seller than the non-optimal Dutch auction or top-price method, and other extremes where they result in a lower average expected gain to the buyers. On the whole, however, in a reasonably active market it seems likely that it would be quite rare for the asymmetry among the buyers to be so substantial and so apparent that one could say a priori that one party or the other would tend to lose from the shift. In the large majority of cases one could be fairly sure of at least some over-all gain from the shift (one can be certain of no over-all loss), with a fairly strong expectation that this over-all gain would be shared by both buyers and sellers in the long run.

In addition to the gain from the improved allocation of resources, there is another possible gain that is not covered in the above analysis, which abstracts from the costs involved in the negotiations. In the top-price method of negotiation, as in the Dutch auction, bidders, in order to maximize their expectation of profit, must concern themselves not only with their own appraisal of the article but also with their estimate of the value that others will place on it and their expectation of the bidding strategy that others will follow. This involves a considerable amount of appraisal of the market situation as a whole, in addition to an appraisal of what the article is worth to the particular bidder himself. Where the bidders are wholesalers who are purchasing the article or lot for retailing in a retail market shared by all the bidders, there may be no great difference between

appraising the article for one's self and estimating what others will value the article at. But especially where the various bidders want the article for different purposes or where the article will be retailed in different areas or by different methods by different bidders, the general appraisal of the market does involve substantial additional information-gathering activity. Moreover, failure to perform this general appraisal with reasonable uniformity is likely to increase the chances that the optimum allocation will not be achieved. It is one of the salient advantages of the second-price method that it makes any such general market appraisal entirely superfluous, whether considered from the standpoint of individual gain or from that of the over-all allocation of resources. Each bidder can confine his efforts and attention to an appraisal of the value the article would have in his own hands, at a considerable saving in mental strain and possibly in out-of-pocket expense. In the first instance this saving might redound largely to the benefit of the bidders; as a corollary, however, more bidders might be induced to put in bids, resulting in a better allocation of resources and a higher price for the seller.

The second-price method may not be automatically self-policing to quite the same extent as the top-price method, but there should be no real difficulty. It would be necessary to show the second-best bid to the successful top bidder so that he would be able to assure himself that the price he is being asked to pay is based upon a bona fide bid. To prevent the use of a "shill" to jack the price up by putting in a late bid just under the top bid, it would probably be desirable to have all bids delivered to and certified by a trustworthy holder, who would then deliver all bids simultaneously to the seller. Under these circumstances, the seller would have no incentive to do other than sell to the top bidder, showing him as his price the second-best bid. Where the seller is a governmental body or a large corporation, so that the agent handling the sale might not be adequately motivated to serve the interests of his principal, it would be desirable to publish the final terms of sale; if this is done, any bidder whose bid has been improperly overlooked would at least be on notice of this, though, unless he was actually the top bidder, he would have only an indirect interest in lodging a protest. If he were uncertain as to the amount of the bid put in by the successful bidder, his protest would be motivated by a hope of being top bidder, so that there would be some advantage at this point of not announcing the top bid, but only the effective price. Even this would not prevent the top bidder and the agent from showing the top bid to the second bidder, together with a quieting *douceur*, so as to be able to set the price at the third

highest bid. If corruption of this order cannot be prevented, then this would constitute a serious disadvantage of the second-price method.

If selling at the second bid price is better than selling at the top price, one is tempted to ask, as a matter of completeness, would not selling at the third bid price be even better? The answer is no, as might be expected from the fact that the second-price method is Pareto-optimal, and there is no further gain to be had from improvement in the allocation of resources. In the game paradigm, the equilibrium strategy for the third-bid-price game will be to make a bid somewhat higher than the value drawn, since the danger of a Pyrrhic victory in which the price to be paid exceeds the value is offset by the increased probability of gains in cases where the second bid exceeds this value but the third bid falls below it. The optimum strategy depends on the strategy of others, and the need for more information and the possibility of non-optimal allocation are reintroduced.

On the other hand the Dutch auction scheme is capable of being modified with advantage to a second-bid price basis, making it logically equivalent to the second-price sealed-bid procedure suggested above on page 20. As presently practiced, speed is achieved by having a motor-driven pointer or register started downward from a prohibitively high price by the auctioneer; each bidder may at any time press a button which will, if no other button has been pushed before, stop the register, thus indicating the price, flash a signal indicating the identity of the successful bidder, and disconnect all other buttons, preventing any further signals from being activated. There would be no particular difficulty in modifying the apparatus so that the first button pushed would merely preselect the signal to be flashed, but there would be no overt indication until the second button is pushed, whereupon the register would stop, indicating the price, and the signal would flash, indicating the purchaser. This would involve some increased difficulty in learning to control the timing of the button-pushing so as to indicate a desired bid, particularly if, in order to save time, the price register or indicator is made to move fairly rapidly. An even more rapid procedure could be developed, with relatively little increase in the apparatus required, if each bidder were provided with a set of dials or switches which could be set to any desired bid, with the electronic or relay apparatus arranged to search out the two top bids and indicate the person making the top bid and the amount of the second bid.

IV. MULTIPLE AUCTIONS

Another interesting case occurs where there is more than one identical object to be sold, but each bidder has use for at most one. Here there are two variations on the progressive type of auction: simultaneous and successive. In simultaneous auctioning the m items can be put up simultaneously, and each bidder permitted to raise his bid even when this does not make his bid the highest. When a point is reached such that no bidder wishes to raise his bid further, the items are awarded to the m highest bidders. If we assume the bid increment to be negligibly small, the results will be that the bidding will stop at a price equal to the $(m + 1)$ st highest value among those placed on the articles by the bidders, this being the bid of each of the bidders placing a higher value on the item—assuming that the $(m + 1)$ st bidder doesn't bother to make a profitless bid. Bidders with the top m values thus secure the article at a uniform price equal to the $(m + 1)$ st value; the result is again Pareto-optimal.

This method is applicable, however, only if the items are actually identical so that there is no problem of deciding who gets first choice and no variation in the value imputed to the various items by a given bidder. In part because of the possibility that there may be minor variations in quality among the items, the more frequent procedure in such cases is for the items to be auctioned off successively, one at a time. With this procedure, an element of speculation or strategy is present during the auctioning of all but the last item, as each bidder must consider whether he should push the bidding up higher on the current item or sign off in the hope that a subsequent item will become available at a lower price. This situation has characteristics similar to that of the Dutch auction.

Consider the game paradigm of the case where there are two items to be auctioned off among N similar bidders, their similarity being represented by each bidder drawing the respective value that any of the items is to have for him, if he secures one, from a distribution that is common to all the bidders. To exclude the complicating element of information that might be inferred from the way the bidding in the auction of the first item develops, it will be assumed that the first auction at least is by sealed bids, price being determined by the second highest bid. The problem is for each bidder i to determine the bid b_i to be made for this first item, as a function of the value v_i that he has drawn. By virtue of the symmetry among bidders, we can assume that this function $b_i(v_i) = x(v)$ is the same for all bidders.

Suppose this function $x(v)$ to have been established and that a bidder, having drawn a value v , is contemplating a deviation from this normal rule by raising his bid from $x(v)$ to $x(v) + dx$. Such a change will affect the outcome only in those cases where this causes the bidder in question to obtain the first item, whereas he would have failed to do so in the absence of this deviation; that is, the deviation will be consequential only where the highest of the other bids lies between $x(v)$ and $x(v) + dx$. The consequence, if any, will be that the increased bid secures the first item at a price between $x(v)$ and $x(v) + dx$, instead of the bidder being almost certain to be the top bidder for the second item auctioned, given that the values v_i will rank in the same order as the normal bids $x(v_i)$ made for the first item, and disregarding the vanishingly small relative probability that two other bids would fall in the range from $x(v)$ to $x(v) + dx$.

The price that would have been paid for the second item is the highest of the $N - 2$ values drawn at random from the common distribution by the $N - 2$ unsuccessful bidders in the second auction; for the case of a rectangular distribution over the interval from 0 to 1, this price has an expected value of $[(N - 2)/(N - 1)]v$ since, if the increment of bid dx being considered causes any change in the outcome at all, none of these unsuccessful bidders can have drawn a value greater than the value v drawn by the deliberating bidder. The bid for the first item must then be at least equal to $[(N - 2)/(N - 1)]v$, or there would be an expected gain from increasing the bid; a similar argument shows that it cannot be greater without creating an incentive to lower the bid.

The equilibrium situation then is that each bidder puts in a bid of

$$b_i = \frac{N - 2}{N - 1} v_i$$

for the first article, or alternatively, in the more usual form of bidding, competes in the bidding up to this level but no further. The second highest value among the N values drawn will average $(N - 1)/(N + 1)$, and the second highest bid, which determines the price of the first item, will then have an average expectation of $(N - 2)/(N + 1)$. This is also the average expected price for the second item, which is, indeed, the average expected price when the two items are auctioned simultaneously as described above. But, as with the Dutch auction of a single item, there is in the successive auction a slightly different dispersion of the prices. There is, also, a tendency in non-symmetrical cases for the results to be other than

Pareto-optimal. Unfortunately, even the simplest of the non-symmetrical successive auction cases would involve at least three bidders, and the complications of a complete analysis appear too formidable to go into here.

V. MULTIPLE SALES BY SEALED BIDS

A type of transaction that is of considerable practical interest is that of the sale of a number of identical items, say an issue of bonds, on the basis of sealed bids. The more usual practice is to accept a certain number of bids starting from those offering the highest price, the effective price for each transaction being the price in the individual bid. An alternative method is to set the effective price at the level of the last bid accepted and permit all successful bidders to benefit from this same uniform price. The usual rationale for this procedure is one of avoiding discrimination in the final price among the various buyers, even though the differential would be based on the bid submitted. The present analysis indicates that this method has the more material advantage of reducing the probability that a bidder's own bid will affect the price he receives, thus inducing bids closer to the full value to the bidder, improving the chances of obtaining or approaching the optimum allocation of resources, and reducing effort and expense devoted to socially superfluous investigation of the general market situation.

To obtain these advantages in full, however, it is necessary to go one step further than is usually done and make the uniform price to be charged the successful bidders equal to the first bid rejected rather than the last bid accepted; only in this way is it possible to insure that each bidder will be motivated to put in a bid at the full value of the article to himself, thus assuring an optimum allocation of resources, at least for the case where the number of items to be offered is absolutely fixed, and avoiding any incentive for wasteful individual expenditure on general market research. Again it appears that, in spite of what appears to be at first glance the establishment of a needlessly low price, this "first-rejected-bid" pricing can be expected in the long run to yield just as high an average price as the "greedier" method.

It is important to realize, however, that this result applies only to cases where each bidder is interested in at most a single unit, and there is no collusion among the bidders. As soon as we consider the more general case where an individual bidder may be interested in securing two or more of the units, while the number of bidders is still too few to produce a fully competitive market, the possibility

of so arranging things that the Pareto-optimal result is achieved without impairing the expectations of the seller disappears. It is not possible to consider a buyer wanting up to two units as merely an aggregation of two single-unit buyers: combining the two buyers into one introduces a built-in collusion and community of interest, and the bid offered for the second unit will be influenced by the possible effect of this bid on the price to be paid for the first, even under the first-rejected-bid method. Where individual bidders may buy more than one indivisible unit, we are, in effect, back in a variant of the exclusive marketing-agency case, where the interests of the marketing agency are merged with those of a single monopolistic seller. In such a case, while the marketing agency need have no concern for the amounts above the competitive equilibrium price which the Pareto-optimal marketing scheme of pages 10–12 would require to be paid to itself as seller, it would be concerned for the amounts by which the revenues from the purchasers would fall short of the competitive equilibrium price, or at least the amount by which these receipts fall short of the possibly somewhat smaller revenues which could in fact be secured on the basis of any other method of approximating the efficient allocation under imperfectly competitive conditions. Nor could optimal results be obtained merely by restricting all bids to an offer to take up to a given quantity at any price below a specified price, the final terms being a price equal to the price bid by the first unsuccessful bidder, each bidder bidding more than this being allotted the amount which he specified. Under such a scheme, for any quantity that a bidder might decide to specify, it would be advantageous for him to specify as his bid price the full *average* value of this quantity to him, since he would prefer this quantity to be allotted at any price lower than this bid rather than be excluded altogether, and a change in his bid price within the range in which he would be successful would not affect the contract price. If a particular bidder is sure that changing the *quantity* he specifies will not affect the contract price, as would be the case if the change is small enough so as not to change the identity of the first unsuccessful bidder and if his demand curve is linear over the relevant range, his quantity specifications would tend to equal the quantity he would demand at the mean of the prices that he expects to result. To the extent that he is mistaken as to the ultimate price, misallocation will result. Even more serious, the resulting bids do not provide in themselves the information necessary to enable the marketing agency to determine the Pareto-optimal result.

VI. SUMMARY

The problem of securing Pareto-optimal results in imperfect markets is thus a moderately difficult one. In the special class of cases where it is known that each purchaser will want a specified quantity or none at all (in particular, where the entire lot or contract must be taken up by a single bidder or where each bidder wants either one unit or none at all) and the total amount to be sold is fixed in advance, it is possible, by establishing in advance that the price is to be determined by the first rejected bid, to achieve the Pareto-optimal result. Moreover, in spite of this method's appearing to accord a lower price than necessary after the bids are in, the higher level of bids induced by this method results, on balance, in a price averaging-out at the same level as would be obtained under Dutch auction, individual bid pricing, or last-accepted-bid pricing methods, at least for cases where the bidders are symmetrical with respect to the *a priori* information which each one has about the probability distribution of the values or bids of the others. In such cases there is a rather strong presumption that a switch from other methods of negotiation to a first-rejected-bid pricing method would be to the long-run advantage of all concerned, the gain being derived from the greater certainty of obtaining a Pareto-optimal result and from the reduction in non-productive expenditure devoted to the sizing-up of the market by the bidders. To be sure, these conclusions are based on a model in which a high degree of rationality and sophistication is imputed to the bidders; nevertheless, in many markets the frequency of the dealings and the professional characteristics of the dealers are such as to make such an assumption not too far from reality; moreover, the change to the first-rejected-price method would substantially diminish the amount of sophistication required to achieve the optimum result.

Where there is a significant asymmetry in the *a priori* positions of the bidders, this conclusion must be somewhat modified: the one complete analysis of an asymmetrical case shows that in some cases the change to the first-rejected-bid method may be to the advantage of the seller but that in other cases it may be substantially to his disadvantage. To extrapolate rather rashly from a single example, one may hazard the generalization that, in the change to the first-rejected-bid or progressive method, bidders who have relatively greater knowledge of the probable behavior of other bidders, either through greater astuteness, more intensive research, or simply through their own position being inherently less patent than that of other bidders—as is exemplified in the example given by bidder 1 having

a value drawn from a distribution over the interval (0,1), and knowing the fixed value the article has for bidder 2, whereas bidder 2 is in relative ignorance concerning the value drawn by bidder 1—will tend to lose the advantage which their superior information gives them in the less determinate situation, whereas the less informed bidders tend to gain, sometimes rather substantially, as their lack of information becomes irrelevant to their behavior in the situation where all they are called upon to do is to make a bid equal to their full value. In situations where the relatively uninformed bidders are the ones more likely to be the successful bidders, the seller can expect to lose, whereas if the informed bidders are the ones more likely to become purchasers, the seller stands to gain from the change. The total gain from the change will always be positive, however.

When it comes to markets where the amounts which each trader might buy or sell are not predetermined but are to be determined by the negotiating procedure along with the amount to be paid, the prospects for achieving an optimum allocation of resources become much dimmer. A theoretical method exists, to be sure, which involves essentially paying each seller for his supply an amount equal to what he could extract as a perfectly discriminating monopolist faced with a demand curve constructed by subtracting the total supply of his competing suppliers from the total demand, and symmetrically for purchasers. But whether this method is thought of in terms of an exclusive state marketing agency operating as an intermediary between suppliers and sellers or in terms of sales or purchases by a government agency on its own account, the method is far too expensive in terms of the inflow of public funds that would be required, in a context where perfectly efficient sources of additional public revenues do not exist. Indeed, if speculation in the usual sense is trading motivated by the prospect of profit, it would hardly be expected that "counterspeculation," motivated by substantially contrary objectives, could be anything but a losing proposition. It may be that further analysis might reveal methods of dealing with imperfect markets that would produce a substantial improvement in the allocation of resources without incurring prohibitive costs, but the analysis of the relatively simple cases discussed here has already shown itself fairly intricate, so that the matter must be left here for the time being.

APPENDIX I

ANALYSIS OF THE HOMOGENEOUS RECTANGULAR SINGLE-PRIZE BIDDING GAME

Given N players, designated by $i = 1, 2, 3, \dots, N$, each drawing a value v_i from a rectangular distribution ranging from 0 to 1, and bidding for a single prize, the

probability that the first $N - 1$ players all draw values between 0 and v while the N th player draws a value between v and $v + dv$ is $v^{N-1} dv$; allowing for the possibility that any of the N players might have the top value, the probability that the highest value drawn lies between v and $v + dv$ is then given by the expression $dP_1(v) = N v^{N-1} dv$. In this event the price, assuming that each player puts in a bid $b_i = [(N - 1)/N] v_i$, will be $p_d = [(N - 1)/N] v$, and thus the expected price is

$$\bar{p}_d = \int p dP_1(v) = \int_0^1 \left(\frac{N-1}{N} \right) v N v^{N-1} dv = \frac{N-1}{N+1},$$

corresponding to an expected highest value drawn of $N/(N + 1)$. This result can be compared with the results under progressive or "common" auctioning, where the price is equal to the second highest value drawn. The probability that this second highest value lies between v and $v + dv$ is then given by the expression $dP_2(v) = N(N - 1) v^{N-2}(1 - v) dv$, so that the average price is

$$\begin{aligned} \bar{p}_c &= \int v dP_2(v) = \int_0^1 N(N-1)(v^{N-1} - v^N) dv \\ &= (N-1) - \frac{N(N-1)}{N+1} = \frac{N-1}{N+1}, \end{aligned}$$

which is the same as for the Dutch auction. The bidders' surplus again has the expected value of $1/(N + 1)$, and the expected value of the object in the hands of the successful bidder is again the same, $N/(N + 1)$. Thus in terms of average expected outcomes, the two methods of auctioning are equivalent.

The probability distributions of the gains realized by the buyers and sellers, however, are quite different under the two methods of auctioning. The variance of the price under the Dutch auctioning method is

$$\begin{aligned} \sigma_{pd}^2 &= \int (p_d - \bar{p}_d)^2 dP_1(v) \\ &= \int_0^1 \left(\frac{N-1}{N} v - \frac{N-1}{N+1} \right)^2 N v^{N-1} dv = \frac{(N-1)^2}{N(N+1)^2(N+2)}. \end{aligned}$$

Under the progressive auction method the variance is

$$\begin{aligned} \sigma_{pc}^2 &= \int (p_c - \bar{p}_c)^2 dP_2(v) \\ &= \int_0^1 \left(v - \frac{N-1}{N+1} \right)^2 N(N-1)(v^{N-2} - v^{N-1}) dv = \frac{2(N-1)}{(N+2)(N+1)^2}. \end{aligned}$$

The difference in variance for the gain to the buyers is even wider: in the Dutch auction the buyer's gain is $v - p = v - [(N - 1)/N] v = v/N$, v being the highest value drawn; the range of possible gains is from 0 to $1/N$, whereas with the progressive auction the gain can vary all the way from 0 to 1. Since with the Dutch auction the gain is always proportional to the price:

$$g_d = \frac{1}{N-1} p_d,$$

we have for the variance of the gain

$$\sigma_{gd}^2 = \frac{1}{(N-1)^2} \sigma_{pd}^2 = \frac{1}{N(N+1)^2(N+2)}.$$

With the progressive auction, for a given bidder to obtain a gain of g or more after drawing a value of v requires that all other bidders draw values less than $v - g$; the total probability of some bidder obtaining a gain of more than g thus is

$$P_c(g) = N \int_{v=g}^1 (v-g)^{N-1} dv = (1-g)^N.$$

The variance of the gain, then, is

$$\begin{aligned} \sigma_{gc}^2 &= \int (g-\bar{g})^2 dP_c(g) = \int_{g=1}^0 \left(g - \frac{1}{N+1}\right)^2 d(1-g)^N \\ &= \frac{N}{(N+1)^2(N+2)}. \end{aligned}$$

This is N^2 times as great as for the Dutch auction.

APPENDIX II

ANALYSIS OF ASYMMETRICAL RECTANGULAR 2-PERSON BIDDING GAMES

Bidder 1 draws his value v_1 from a rectangular distribution ranging from 0 to 1, while bidder 2 draws his value v_2 from a rectangular distribution ranging from a to $b \leq 1$. Let $y_1(x)$ and $y_2(x)$ be the respective probabilities of a bid of less than x by bidders 1 and 2, respectively, and $v_1(x)$ and $v_2(x)$ the value drawings by the two players that are to lead them to make a bid of x . Then the expected gain $E(g_1)$ to bidder 1, if he draws a value v_1 and makes a bid of x , will be the amount of gain if he is successful, $(v_1 - x)$, times the probability of success, which is, of course, the probability $y_2(x)$ that bidder 2 bids less than x .

A necessary condition for an equilibrium point such that neither bidder can gain by bidding other than according to the relation $v_i(x)$, then, is that $E(g_1) = y_2(x)(v_1 - x)$ shall be a maximum with respect to x , which requires, for all values of x for which $y_2(x)$ is continuous,

$$\frac{\partial E(g_1)}{\partial x} = -y_2(x) + (v_1 - x) \frac{dy_2}{dx} = 0, \quad (1)$$

and similarly

$$\frac{\partial E(g_2)}{\partial X} = -y_1(x) + (v_2 - x) \frac{dy_1}{dx} = 0. \quad (2)$$

From the way v_1 and v_2 are drawn, we can put

$$y_1(x) = v_1(x) \quad (3)$$

and

$$y_2(x) = \frac{1}{b-a} [v_2(x) - a] \quad (4)$$

(e.g., if the probability of a bid of less than x by bidder 2 is 0.3, then the value drawing that induces a bid of x must be $v_2(x) = a + 0.3[b - a]$).

If the solution is to be Pareto-optimal, we must have $v_1(x) = v_2(x) (= v(x))$ for all values of x for which both v_1 and v_2 lie within the range of possible drawings; otherwise there would be the possibility that the person drawing the lower value would obtain the object, leading to misallocation of resources. But if we make the indicated substitutions, this requires simultaneously that

$$-v + (v - x) \frac{dv}{dx} = 0, \quad (5)$$

and

$$-\frac{1}{b-a}(v-a) + (v-x) \frac{1}{b-a} \frac{dv}{dx} = 0, \quad (6)$$

which is possible only if $a = 0$. If we restrict the model to the case $a = 0$, equation (5) can be written $v dv = v dx + x dv = d(xv)$, which can be integrated directly to give $\frac{1}{2} v^2 = xv + c$, or $x = (v/2) - (c/v)$. If we assume some lower bound for bids as v approaches zero (having normalized the v_1 values to the range 0 to 1, we cannot rule out negative bids per se), this implies that $c = 0$, and we get, for the bidding rule, simply $x = v/2$ for $0 \leq v \leq b$.

However, not even this provides an equilibrium, since if player 2 follows this pattern, then for $v_1 > b$, bidder 1 will increase his gain without diminishing his chances of success by reducing his bids from $\frac{1}{2} v_1$ to $\frac{1}{2} b$; if he does this, however, there will then be a temptation for player 2 to bid more than $\frac{1}{2} b$ on some occasions, since by doing so his chances of winning are substantially increased. Hence Pareto-optimality is incompatible with Nash-equilibrium in this case also.

If we abandon the requirement of Pareto-optimality and look for a general Nash-equilibrium point without this stipulation, the solution runs into considerable mathematical difficulty. For purposes of simplification we can put

$$z_1 = \frac{y_1}{y'_1} \quad \text{and} \quad z_2 = \frac{y_2}{y'_2}, \quad \left(y'_1 = \frac{dy_1}{dx}, \text{ etc.} \right), \quad (7)$$

so that equations (1) and (2) become

$$v_1 - x = z_2 \quad (8)$$

$$v_2 - x = z_1; \quad (9)$$

while substituting equations (3) and (4) in equation (7) we get

$$z_1 = \frac{v_1}{v'_1} \quad (10)$$

and

$$z_2 = \frac{v_2 - a}{v'_2 - a}. \quad (11)$$

Using equation (8) to eliminate v_1 from equation (10) and using equation (9) to eliminate v_2 from equation (11), and clearing fractions, we have

$$z_1 z_2' + z_1 = z_2 + x \tag{12}$$

and

$$z_2 z_1' + z_2 = z_1 + x - a . \tag{13}$$

Adding equations (12) and (13), we get

$$z_1 z_2' + z_2 z_1' = 2x - a . \tag{14}$$

Since the left-hand side of equation (14) is an exact differential, we can now integrate:

$$z_1 z_2 = x^2 - ax + k . \tag{15}$$

Solving equation (15) for z_1 and putting the results in equation (12), we have

$$(x^2 - ax + k)(z_2' + 1) = z_2(z_2 + x) . \tag{16}$$

While equation (16) is now a relatively simple differential equation involving only $z_2(x)$, $z_2'(x)$, and x , it resists solution by analytical methods (even for $a = 0$) while if an approximate numerical quadrature is to be made, it is not immediately obvious what the required boundary conditions are to be that will determine k and the second constant of integration.

APPENDIX III

A SIMPLIFIED ASYMMETRICAL BIDDING GAME

In this game bidder 1 draws his value v_1 from a rectangular distribution of range 0 to 1, while bidder 2 has the fixed value $v_2 = a$.

If we allow bidder 2 to determine his bid x by a random drawing from a distribution selected by him, we can now let $y_2(x)$ be the probability that his bid will be less than x , this function now being determined directly by bidder 2. The equilibrium conditions then become

$$-y_2(x) + (v_1 - x)y_2'(x) = 0 \tag{17}$$

and

$$-y_1(x) + (a - x)y_1'(x) = 0 , \tag{18}$$

where for bidder 1 we again have

$$y_1(x) = v_1(x) . \tag{19}$$

Equation (18) can be integrated directly, after dividing through by $y_1(x)$ ($a - x$) to separate the variables, giving $\log(a - x) + \log y_1(x) = x_1 \log k$, or $(a - x)y_1(x) = k$, which by the use of equation (19) gives

$$x = a - \frac{k}{v} . \tag{20}$$

as the rule by which bidder 1 is to determine his bid, at least for any range of x where $y_2'(x)$ is continuous.

To determine an appropriate value for the constant of integration k , some general propositions concerning the necessary nature of the various probability distributions can be invoked. Obviously, the probability $p_i(x)$ that a bid of x by bidder i will be a winning bid for i , is a non-decreasing function of x for any given behavior pattern of the remaining bidders. Also $v_i(x)$ must be a non-decreasing function for each bidder i , for suppose v^* to be a value drawing greater than v , and x^* and x the corresponding bids, with $x^* < x$; then, if $p(x^*) = p(x)$, it would become profitable to make a bid of x^* rather than x in response to a drawing of v ; if $p(x^*) < p(x)$, it would become profitable to interchange the bids and bid x for a value of v^* and x^* for a value of v , the gain being

$$\begin{aligned} & [(v^* - x)p + (v - x^*)p^*] - [(v - x)p + (v^* - x^*)p^*] \\ & = (v^* - v)(p - p^*) > 0. \end{aligned}$$

If the value drawings of the different players are uncorrelated, this interchange will have no repercussions on the other players.

Let x_m be the greatest lower bound of the possible winning bids under equilibrium conditions. It can be shown that if $y_i(x)$ has a discontinuity at x_a , implying that the probability p_a that bidder i makes a bid of exactly x_a is positive, then $x_a \leq x_m$. This can be seen as follows: suppose $x_m < x_a$; let B be the least upper bound of all bids $b \leq x_a$ by bidders other than i , and let V be the least upper bound of all values corresponding to such bids by bidders whose possible bids individually have the least upper bound B . If $B < x_a$, bidder i could increase his expected gains by reducing all his bids of x_a to $B + \varepsilon$, since there are no competing bids in the interval. Hence $B = x_a$. If $V > x_a$, then, for any $\varepsilon > 0$, there exists a value v_ε capable of being drawn by some bidder j ($\neq i$) with a corresponding bid b_ε such that $x_a - \varepsilon < b_\varepsilon \leq x_a$ and $V - \varepsilon < v \leq V$; raising this bid from b_ε to $x_a + \varepsilon$ will increase the probability of winning by at least $p_a p_j(x_a)$, and this will be sufficient, for sufficiently small ε , to outweigh the reduction in the amount of profit in case of a win by something less than 2ε , so that this cannot be an equilibrium situation, and for equilibrium we must have $V \leq x_a$. If $V < x_a = B$, then all bids b^- such that $V < b^- < B = x_a$ have $v^- < b^-$ resulting in a loss if they are the winning bid; if no such bids win, we have $x_a \leq x_m$, Q.E.D. Any such bid that might win can be profitably reduced to a bid equal to v . If, finally, $V = x_a$, and there is some bid $b^* < x_a$ with $p_j(b^*) > 0$, then for any $\varepsilon > 0$, there will exist a value v^+ capable of being drawn by some bidder j ($\neq i$), with a corresponding bid b^+ , such that $x_a - \varepsilon < v^+ \leq x_a$ and $x_a - \varepsilon < b^+ \leq x_a$, implying $v^+ - b^+ < \varepsilon$; so that, for sufficiently small ε , $p_j(b^+)(v^+ - b^+) < \varepsilon < p_j(b^*)(v^+ - b^*)$, implying that in this case it would be profitable to reduce the bid from b^+ to b^* . Hence $p_j(b^*) = 0$ for all $b^* < x_a$, and $x_a \leq x_m$.

Any bids with lumped probabilities can thus occur, if at all, only at the bottom of the range of possibly successful bids, implying that for bids above the minimal winning bid and for value distributions that are dense, the function $v(x)$ must be continuous, while $x = v$ can occur only for $x = x_m$.

In the case at hand, equation (20) must hold at all points where $x \neq v_1$ and $y_2 \neq 0$; k must obviously be positive, otherwise the relationship would be perverse. In this two-bidder case, the maximum bid must be the same for the two bidders, since any bid by one higher than the maximum bid of the other could be

reduced with profit. Suppose $k < a^2/4$, so that the condition $x = v_1$ is realizable at the two roots r_1 and r_2 of the quadratic equation $x^2 - ax + k = 0$; if r_2 is the larger root, $x_m = r_2$ and $y_2(r_2) = 0$. Putting this in equation (17), we obtain

$$\frac{y_2'}{y_2} = \frac{a - x}{(x - r_1)(x - r_2)}, \quad \text{or} \quad (r_2 - r_1) \frac{y_2'}{y_2} = \frac{r_1}{x - r_2} - \frac{r_2}{x - r_1}. \quad (21)$$

This can be integrated to yield

$$(r_2 - r_1) \log y_2 = r_1 \log(x - r_2) - r_2 \log(x - r_1) + \log A, \quad (22)$$

or

$$y_2^{r_2 - r_1} = A (x - r_2)^{r_1} (x - r_1)^{r_2}. \quad (23)$$

This function runs from $y_2(r_2) = 0$ to $y_2(a - k) = 1$, by suitable choice of the constant of integration A , so that $a - k$ is the maximum bid for player 2 as it is for player 1, as indicated by putting $v_1 = 1$ in equation (20).

This equilibrium has only a precarious stability, however, for its stability depends, on the one hand, on player 1 actually bidding according to equation (20) for all drawings of v from 0 to 1, even though this means making potentially dangerous bids of more than v for drawings of v between r_1 and r_2 ; any reduction in such bids would create an inducement for player 2 to put in some bids of less than r_2 , and the equilibrium breaks down. Even if player one persists in following equation (20) to the letter, there is no positive short-run disadvantage to player 2 if he puts in bids of less than r_2 ; indeed, while neither player gains immediately by lowering his bid, he does not lose, and the resulting breakdown of the equilibrium tends to shift the equilibrium in the direction of a higher value of k , which is ultimately to the advantage of both bidders, so that one can hardly call the solutions reached for these values of k really stable.

On the other hand, if $k > a^2/4$, the function $v_1 = k/(a - x)$ lies entirely above the line $v_1 = x$, so that $v_1 - x > 1/h$, say, for some fixed positive h for all x ; equation (17) now implies

$$y_2' = \frac{1}{v_1 - x} y_2 < h y_2,$$

so that the curve $y_2(x)$ is always flatter than curves of the form $y = Ae^{hx}$ and thus can never reach the x -axis, no matter how far out on the negative x -axis we go. Nevertheless, in practice there will always be some lower bound on the bids that will be accepted. If the bidding distributions are terminated by a lumped probability at this lower bound, however, it immediately becomes profitable to shade the minimum bids upward, and no equilibrium is established.

There remains the case of $k = a^2/4$; now $v_1(x)$ is tangent to $v = x$ at $x = a/2$. Equation (20) now becomes

$$v_1 = \frac{a^2}{4(a - x)},$$

and putting this in equation (17) gives

$$y_2 = \frac{1}{4(a - x)} (a^2 - 4ax + 4x^2) y_2', \quad \text{or} \quad \frac{y_2'}{y_2} = \frac{4(a - x)}{(2x - a)^2}, \quad (24)$$

which can be integrated to give

$$\log y_2 = -\log(2x - a) - \frac{a}{2x - a} + C. \quad (25)$$

For all values of C we have $y_2 = 0$ for $x = a/2$. Bidder 2's minimum bid is thus $a/2$, and all bids of less than this by bidder 1 are non-winning. In order for the maximum bids to be the same, since we have, for $v_1 = 1$,

$$x = a - \frac{1}{4}a^2,$$

for $y_2 = 1$ we must have

$$\log y_2 = 0 = -\log(a - \frac{1}{2}a^2) - \frac{2}{2 - a} + C, \quad (26)$$

so that

$$C = \log \frac{a(2 - a)}{2} + \frac{2}{2 - a},$$

and equation (25) becomes

$$\log y_2 = \log \frac{a(2 - a)}{2(2x - a)} + \frac{2}{2 - a} - \frac{a}{2x - a}. \quad (27)$$

This solution for the case $a = 0.8$ is illustrated in Figure 3. The pro forma bids put in by player 1 when he draws a value of less than $a/2$ can take any form between the limits $x = v_1$ and $x = a - (1/4v_1)a^2$, without leading to any immediate breakdown of the equilibrium. Other bids would not be immediately disastrous, but, of course, a bid of greater than v would risk loss if the bid were successful, while consistent bids of less than $a - (1/4v_1)a^2$ would tempt bidder 2 to change his strategy.

The expected payment by bidder 1 would be

$$\int_{v_1=(1/2)a}^1 y_2(x) x \, d v_1(x),$$

and similarly for bidder 2 it would be

$$\int_{y_2=0}^1 v_1(x) x \, d y_2(x).$$

Integrating this latter expression by parts, we have

$$\begin{aligned} \int_{y_2=0}^1 v_1 x \, d y_2 &= v_1 x y_2 \Big|_{y_2=0}^{y_2=1} - \int_{y_2=0}^1 y_2 \, d(v_1 x) \\ &= v_1 x \Big|_{y_2=0}^{y_2=1} - \int_{y_2=0}^1 y_2 v_1 \, dx + y_2 x \, d v_1. \end{aligned} \quad (28)$$

Since the limits of integration are equivalent, the integral of the last term is equal to the expected payment of bidder 1, so that the total expected payment received by the seller is

$$\begin{aligned}
 a - \frac{a^4}{4} - \int_{x=a/2}^{a[1-(1/4)^a]} \frac{a(2-a)}{2(2x-a)} e^{[2/(2-a)-a/(2x-a)]} \frac{a^2}{4(a-x)} dx \\
 = 1 - b^2 - b(1-b)^2 e^{1/b} \int_0^b \frac{1}{(1-u)u} e^{-1/u} du,
 \end{aligned}
 \tag{29}$$

where

$$u = \frac{2x-a}{a} \quad \text{and} \quad b = \frac{2-a}{2}.
 \tag{30}$$

The integral in equations (29) cannot be evaluated in terms of standard functions, but it can readily be evaluated for specific values of a and b , for example, by Simpson's rule. The interesting comparison is with the progressive auction in which, for v_1 greater than a , bidder 1 gets the article at a price a , while for a v_1 less than a , bidder 2 gets the article at a price of v_1 . The average price thus is $a - \frac{1}{2}a^2$. With the Dutch auction the average realized price turns out to be greater than this for a greater than about 0.43, and less for values below 0.43.

The expected consumer's surplus, for the optimal allocation achieved with progressive auctioning is $\frac{1}{2}(1-a)^2$ for bidder 1 and $\frac{1}{2}a^2$ for bidder 2, yielding a total expected value of $\frac{1}{2}(1+a^2)$. Under the Dutch auction the expected consumer's surplus for the first bidder is $\int y_2(v_1-x)dv_1$ for the first bidder and $\int v_1(a-x)dy_2$ for the second. These integrals again must be evaluated by approximate numerical methods; it is possible, however, to express the various quantities in terms of a common integral function as is shown in the following summary table:

| AVERAGE EXPECTATION OF | PROGRESSIVE AUCTION | | | DUTCH AUCTION | |
|----------------------------------|---------------------|----------------------|-------------------|---|------------|
| | $a \leq 0$ | $0 \leq a \leq 1$ | $a \geq 1$ | $0 \leq a \leq 2$ | $a \geq 2$ |
| 1. Payment by 1... | a | $a - a^2$ | 0 | $(1 - [a/2])^2(a/2)(2 - [a/2]) - F(a)$ | 0 |
| 2. Payment by 2... | 0 | $\frac{1}{2}a^2$ | $\frac{1}{2}$ | $1/4a^2(3-a)$ | 1 |
| 3. Receipts of seller (1+2)..... | a | $a(1-a/2)$ | $\frac{1}{2}$ | $a/2(2-a + [a^2/2] - [a^3/8]) - F(a)$ | 1 |
| 4. Net gain of 1.... | $(1-2a)/2$ | $\frac{1}{2}(1-a)^2$ | 0 | $\frac{1}{2}(1 - [a/2])^4 + \frac{1}{2}F(a)$ | 0 |
| 5. Net gain of 2.... | 0 | $\frac{1}{2}a^2$ | $a - \frac{1}{2}$ | $\frac{1}{4}a^2$ | $a - 1$ |
| 6. Total value (3+4+5)..... | $\frac{1}{2}$ | $\frac{1}{2}(1+a^2)$ | a | $\frac{1}{2} + \frac{1}{2}a^2 - \frac{1}{8}a^4 - \frac{1}{2}F(a)$ | a |

where $F(a) = \frac{1}{4}a^2(1 - \frac{1}{2}a)e^{2/(2-a)} \int_0^{1-a/2} [u/(1-u)]e^{-1/u} du$

Figure 4 gives a general picture of how these various quantities behave as a changes.