Kenneth R. French
University of Chicago

Robert E. McCormick
Clemson University

Sealed Bids, Sunk Costs, and the Process of Competition*

Economic theory traditionally has relied on increasing marginal costs to explain the recovery of sunk costs. We develop a model of competition that does not rely on this assumption. The model predicts that each buyer offers to pay less for an asset than he thinks that it is worth —even when there is vigorous competition. However, competition does guarantee that the sale is like a fair game; the successful buyer's expected profit equals the sum of his competitor's sunk costs. We use the model to make predictions about pricing policies and marketing strategies in some retailing situations.

I. Introduction

How do firms recover precontract costs? For example, suppose 10 firms compete in a sealed-bid auction for the right to explore and develop an offshore oil field. Each spends $500,000 to estimate the value of the field and to prepare its bid. Suppose this investigation reveals that the tract is worth approximately $20 million. If sunk costs do not matter, competition dictates that the winner must bid $20 million, but if the winner is to recover its prebid costs the maximum offer can be only $19.5 million. Furthermore, long-run industry equilibrium considerations suggest that the highest bid will be only $15 million ($20 million minus the total bid preparation costs of $5 million).

If competition forces firms to charge marginal cost, why do firms ever invest in resources which become sunk? In some industries, inframarginal rents can be used to recover these costs. However, the example above, as well as many others, involves an all-or-none deci-

* The authors thank colleagues at UCLA, Clemson University, and the University of Rochester, especially Armen Alchian and Mike Jensen, for their help in developing this paper. They are also grateful to Yoram Barzel, an unidentified referee of this *Journal,* and seminar participants at MIT, Dartmouth, and Texas A&M for comments on earlier drafts. Financial support was provided by the Foundation for Research in Economics and Education and the Managerial Economics Research Center at the University of Rochester.

(*Journal of Business,* 1984, vol. 57, no. 4)

0021-9398/84/5704-0001$01.50

sion—there are no inframarginal rents. For firms in this industry to survive in the long run, the textbook prediction that price equals marginal cost cannot always be accurate. Clearly the conventional approach is useful in describing the ex ante equilibrium—before any resources are spent—but it cannot also be correct when the bids are tendered. In this paper we examine the process of competition to determine how firms recover precontract costs when there are no inframarginal rents.

The standard model of a first-price sealed-bid auction provides a useful framework for our initial analysis. In Section II we synthesize the results of Wilson (1977), Johnson (1979), Ramsey (1980), and others to show that, with a finite number of bidders competing for an asset, every buyer will offer less than he thinks it is worth. The expected value of the winning bid is lower than the expected value of the asset—there is some expected profit left over for the winner. Moreover, this expected profit is a decreasing function of the number of the competitors. In equilibrium, firms enter the auction until the winner's expected profit equals the industry's estimation costs. Precontract costs play an important role in determining the number of competitors. However, after a firm has decided to tender a bid, it can treat its estimation costs as sunk and still, on average, recover these costs.

In Section III we consider the effect of precontract costs on the other side of the market. If buyers expect to recover their estimation costs, the seller must expect to pay these costs; the asset owner's expected revenue is equal to the expected value of the asset minus all of the precontract costs. Section III examines three sales mechanisms the owner might use to reduce these costs.

The fourth section extends the model to several consumer goods markets. For example, the process of collecting estimates for home repair can be viewed as a sealed-bid auction, although there is an important difference between these auctions and the auctions examined in Section II. In the standard auctions in Section II, potential bidders determine how many bids will be submitted. In the consumer goods market this decision is often made by the customer. This distinction allows us to make several predictions about those markets. For example, repair shops in high-income neighborhoods are more likely to advertise while shops in low-income neighborhoods are more likely to charge for estimates. The fifth section summarizes the paper and discusses several implications of the model.

II. Precontract Costs in a Sealed-Bid Auction

In this section we consider the recovery of precontract costs within the framework of a first-place sealed-bid auction. This framework is particularly useful for our initial analysis because the model is well defined and because it has been studied by a number of economists.[1] Our analysis synthesizes the results of many of these economists, especially Wilson (1977), Johnson (1979), and Ramsey (1980).

In the first-place sealed-bid auction examined here, a number of identical firms compete for a valuable asset. Each firm's bid is based on both public and private information about the asset. Private information samples cost C dollars and they are specific to this auction. Firms maximize their expected profit. Ignoring the prebid cost, C, the winning bidder's profit is equal to the difference between the value of the asset, V, and the bid it submits, b_i, while the losers' profits are zero.

When the firms are determining their bids, we assume that (1) they all know the number of competing bidders, n; (2) none of them knows the actual value of the asset or the bid of any other firm; (3) all of the firms start with the same prior distribution on the value of the asset, $G(V)$; and (4) each firm purchases a private sample of information about that value, I_i. The private samples are drawn independently from the same distribution, $F(I_i|V)$, which is conditional on the actual value of the asset. We also assume that the owner of the asset does not know its exact value but that his expectation is based on the same prior distribution as the bidders, $G(V)$.

The bidding process can be broken into two steps. First, each firm estimates the value of the asset as if it has received the most optimistic sample of information. Since the firms are identical, this means that each firm assumes that it will win the bid. Obviously, all but one of the firms will be making this assumption incorrectly. However, the winning firm—the only one that correctly assumes that it will win—is the only bidder that affects its actual payoff by making this assumption.

Suppose firm i wins the auction. The act of winning provides valuable information; it implies that the firm observed the highest private estimate of the asset's value. If the winning bidder has not corrected for this "winner's curse" before submitting its bid, on average its estimate will exceed the asset's true value. Therefore, by correctly predicting that it will win the auction, firm i can form a more accurate estimate of the asset's value and improve both its bid and its expected payoff.

Usually the firm will be wrong when it assumes that it has received the most optimistic sample of information. This error will cause the firm to form a downward-biased estimate of the value of the asset. However, since the firm would not have won the bid anyway, this error has no effect on its payoff.

1. See Engelbrecht-Wiggans (1980) for a survey of the bidding literature. Vickery (1961), Wilson (1967, 1977), Baron (1972), Goldberg (1977), Reece (1978), Johnson (1979), Holt (1980), Ramsey (1980), Coppinger, Smith, and Titus (1980), Riley and Samuelson (1981), and Milgrom and Weber (1982) examine various aspects of the competitive bidding problem.

Before the bids are opened none of the firms knows whether it will win or lose. However, since the eventual winner improves his payoff by assuming he will win, while the eventual losers do not affect their payoffs, all of the firms improve their ex ante payoffs by assuming they will win. Therefore, the first step in each firm's bidding process is to form an estimate of the asset's value assuming that it has received the most optimistic sample of information.

The second step for firm i is to determine its bid based on its adjusted estimate of V. While it might appear that the bid should equal this estimate, in general this will not be the case. This would result in the winning bidder paying the expected value of the asset—which is equivalent to the competitive solution of price equal to marginal cost. Ignoring estimation costs, this strategy would also result in an expected profit of zero for all bidders. Firm i can improve its expected profit by submitting a bid that is slightly below its adjusted estimate of the asset's value. To see why, assume that the other firms do bid their adjusted estimates while firm i reduces its bid. As figure 1 indicates, this can lead to three different results. First, firm i may not have observed the highest sample of information. In this case, reducing its bid will not affect the firm's expected profit since it would not have won the auction anyway. Second, firm i may have observed the highest sample and submitted a bid below the next highest bid. In other words, reducing its bid below its adjusted estimate of the asset's value may cause firm i to lose the auction. Ignoring estimation costs, the firm's expected profits are zero in this case—exactly what they would be if the firm had bid its adjusted estimate. Finally, firm i may win the auction. This will happen if firm i has observed the highest sample and submits a reduced bid which is above the next highest bid. Since the firm bids below its adjusted estimate of the asset's value—the firm's expectation of the value, conditional on winning—its expected profit is positive. In both the first and second cases, firm i does not alter its expected payoff by bidding below its adjusted estimate. However, in the third case this strategy is more profitable than bidding the adjusted estimate. As long as there is some chance that the third case will occur, firm i will increase its profits by bidding below its adjusted estimate of the asset's value.

Figure 2 summarizes these effects. Firm i forms an initial estimate of the value of the asset, $E_0(V)$, using the prior distribution, $G(V)$. After drawing a sample of information which suggests that the asset is worth I_i, the firm updates its estimate to reflect this information, $E_1(V|I_i)$.[2] Firm i makes two more adjustments before submitting its bid. First, it

2. To simplify the description of the bidding process, we are assuming that I_i is an estimate of the asset's value. This assumption is stronger than necessary; the model only requires that I_i provide some information about the asset's value.

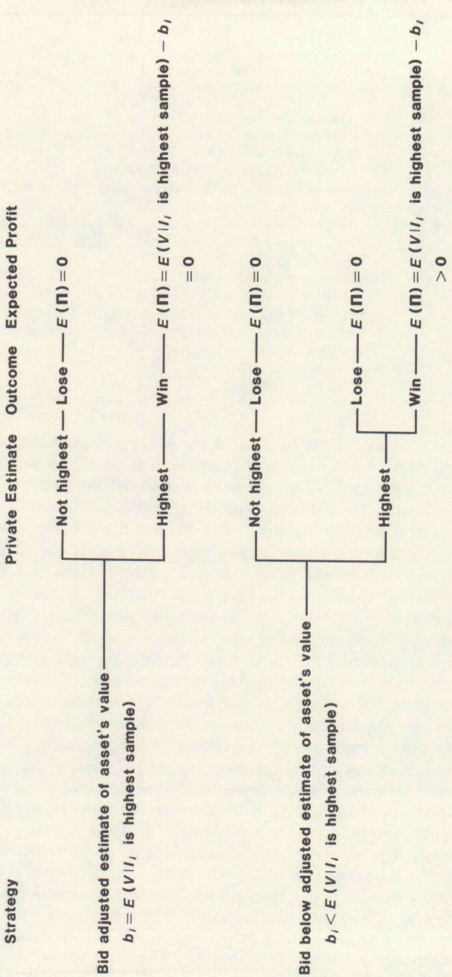

FIG. 1.—Expected profit under alternative bidding strategies when other firms bid their adjusted estimates of the asset's value. (The expected profit in this figure ignores the estimation costs.)

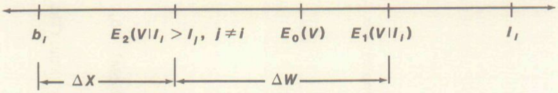

FIG. 2.—Sequential bid preparation:

$E_0(V)$	= initial estimate of the value of the asset, based on the prior distribution, $G(V)$.
I_i	= sample of information drawn by firm i.
$E_1(V\|I_i)$	= updated estimate of the asset's value. This estimate lies between $E_0(V)$ and I_i.
$E_2(V\|I_i > I_j, j{\neq}i)$	= adjusted estimate of the asset's value conditional on firm i winning the auction. This adjusted estimate is below the updated estimate.
ΔW	= winner's curse adjustment. $\Delta W = E_1(V\|I_i) - E_2(V\|I_i > I_j, j \neq i)$
b_i	= bid submitted by firm i. This bid is below the adjusted estimate so the firm's expected profit, ignoring estimation costs, is positive.
ΔX	= finite bidders' adjustment. $\Delta X = E_2(V\|I_i > I_j, j \neq i) - b_i$

reduces its estimate by ΔW because of the winner's curse. With this adjustment, firm i has an unbiased estimate of the asset's value if it wins the auction, $E_2(V|I_i > I_j, j \neq i)$. Then the firm reduces its bid below this adjusted estimate so that its expected payoff is positive.

This argument predicts that firm i will submit a bid that is below its adjusted estimate of the asset's value. However, it does not indicate how much firm i will reduce its bid, nor does it recognize that the other firms will also bid below their own adjusted estimates. In fact, since the firms are identical they must use the same bidding strategy. Wilson (1977) derives the optimal bidding strategy for this model. As we demonstrate in the Appendix, Wilson's profit-maximizing strategy has an appealing interpretation; the optimal bid for each firm occurs where the expected marginal revenue from lowering the bid 1 dollar equals the expected marginal cost. Firm i gains by lowering its bid because this reduces the price it must pay for the asset if it wins the auction. However, there is also a potential cost; lowering the bid reduces the firm's probability of winning. The firm lowers its bid until the expected marginal revenue from reducing the cost of the asset equals the expected marginal cost from reducing the probability of winning.

In summary, each firm obtains an initial estimate of the value of the asset. The firm then makes an adjustment to this estimate to correct for the possibility that its initial estimate was the highest—the winner's curse. The firm now has an unbiased estimate of the value of the asset conditional on being the winning bidder. Given this unbiased estimate the firm calculates its most profitable bid. The firm will not bid its adjusted estimate of the value of the asset—price is not equal to marginal cost. The firm uses a different marginal condition. Depending on the number of competitors, the firm reduces its bid below its estimate of the asset's value to the point where the expected cost of reducing the bid 1 more dollar equals the expected revenue.[3]

The fact that the winning bidder does not bid its adjusted estimate of the asset's value presents a partial solution to the question of how firms recover their prebid costs. Ignoring the estimation costs, each firm has a positive expected profit when its bid is tendered. However, as the model has been developed so far, there is nothing that requires this expected profit to equal the prebid costs. Firms still might not expect to recover all of their estimation costs. Alternatively, a firm's expected profit might be higher than these costs, even in the face of intense competition.

The second half of the estimation cost puzzle can be solved by observing that the number of bidders is endogenous. If each bidder's expected profit falls as the number of bidders increases, new firms will enter until the expected profit from bidding equals the expected cost. In equilibrium, there will be no incentive for new firms to enter and yet the bidding firms will, on average, recover their sunk costs.

This equilibrium argument relies on the proposition that each bidder's expected profit falls as the number of bidders increases. There are two reasons why this condition holds. First, as the number of bidders increases, the probability that any particular bidder will win the auction falls. Each bidder is less likely to receive a positive payoff. Second, the winner's payoff is likely to fall as the number of bidders increases.

The expected payoff to the winning bidder is positive because each firm bids below its adjusted estimate of the asset's value. As the number of bidders increases, each firm will narrow the gap between its bid and its adjusted estimate, reducing its payoff if it wins. To see why firms will behave this way, assume firm i has received the most optimistic sample of information, I_i. As the number of bidders increases, the expected difference between I_i and the next highest information sample falls. Therefore, there is a higher probability that firm i will lose the auction as it reduces its bid below its adjusted estimate. This increases the cost of lowering its bid without increasing the benefit. With more bidders, firm i will not reduce its bid below the perceived asset value as much as it will with fewer bidders. In fact, neither firm i nor any other firm knows whether it has observed the highest sample.

3. Of course, all of the competing firms are making their bids simultaneously. Each firm's optimal strategy must reflect the interdependencies among the competitors. Wilson's solution, which is discussed in the App., is a Nash equilibrium; it is optimal for each firm to follow the proposed strategy if its competitors do.

However, since each firm forms its bid under this assumption, this argument applies to all of them; increasing the number of bidders leads each firm to reduce the gap between its bid and its adjusted estimate of the asset's value.

To summarize, each bidder's expected profit falls as the number of bidders increases for two reasons. First, the expected profit for the winning bidder falls. Second, the probability that any particular firm will win and obtain this expected profit also falls.[4]

Figure 3 illustrates how a potential entrant decides whether to prepare a bid or not. Before beginning its estimation process, firm i has a prior estimate of the asset's value, $E(V)$, based on $G(V)$. It also forms an expectation of the winning bid, $E[B(n)]$, which increases with the number of bidders. If there are currently $n - 1$ firms preparing bids and firm i decided to enter, the expected profit for the winning bidder will be

$$E(\Pi_n|\text{Win}) = E(V) - E[B(n)]. \tag{1}$$

Since each bidder is equally likely to win, the potential entrant's expected profit from bidding, ignoring estimation costs, is

$$E(\Pi_n) = (1/n)\,\text{E}(\Pi_n|\text{Win}) \tag{2}$$

$$= (1/n)\,\{E(V) - E[B(n)]. \tag{3}$$

The equilibrium number of bidders, n^*, occurs where the expected profit from bidding equals the bid preparation costs,

$$E(\Pi_{n^*}) = C. \tag{4}$$

If fewer than n^* firms have committed resources to bidding, other firms will enter the auction because their expected profits, net of their prebid costs, are positive. However, once n^* firms are competing, no other firm has an incentive to enter because its expected profit will not cover its estimation costs. In fact, if a new firm did enter the auction it would also drive everyone else's expected profit below its (now sunk) estimation costs.

III. The Asset Owner's Problem

Before extending our model to other markets, it is helpful to examine the effect of estimation costs on the asset owner. If he uses the sealed-bid auction of Section II, equations (3) and (4) imply that his expected

4. See Ramsey (1980), esp. pp. 87–90, for further discussion of this result, as well as the necessary regularity conditions. It is conceivable that, under certain distributional assumptions, the expected profit will be an increasing function of the number of bidders over some range. However, this possibility is economically irrelevant. Firms will enter until the expected profit is in the decreasing range and the equilibrium condition holds.

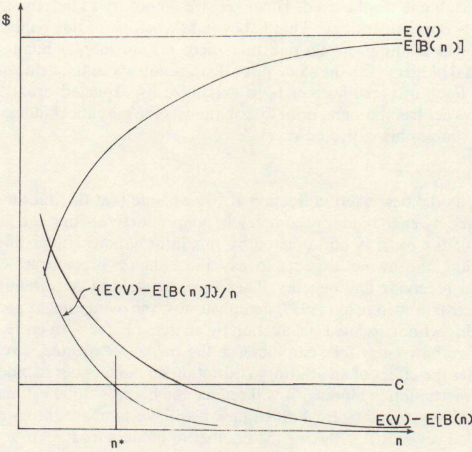

FIG. 3.—Determining the equilibrium number of bidders. $E(V)$ is the expected value of the asset, $E[B(n)]$ is the expected value of the winning bid when there are n bidders, and $E(V) - E[B(n)]$ is the expected profit for the winning bidder. Firms enter the auction until the expected profit from bidding, $\{E(V) - E[B(n)]\}/n$, equals the cost of bidding, C.

revenue is equal to the expected value of the asset minus all the bid preparation costs,

$$E(R) = E[B(n^*)] \\ = E(V) - Cn^* \tag{5}$$

In effect, the estimation costs are transferred from the bidders to the asset owner.

Although the asset owner can choose from a wide range of sales mechanisms, including a variety of auctions, retail arrangements, and negotiation strategies, the basic result in equation (5) is pervasive. The owner always expects to pay all of the potential buyers' precontract costs. This does not mean that the owner is indifferent about the sales mechanism. The institution he chooses determines the amount of information that is produced and the estimation costs he bears. In this section we examine the effect of three possible sales strategies on the owner's expected revenue. We also make some predictions about when each one will be used. However, we do not try to determine the owner's optimal strategy. That is beyond the scope of this paper.[5]

We will continue to assume that there are an infinite number of potential bidders with the same prior distribution about the value of the asset. Each of these firms acts to maximize its expected profit. The asset owner has the same prior distribution as the potential bidders and he maximizes his expected revenue.

A. *Entry Fees*

In the model described in Section II, we assume that the information collected by each bidder is completely nonproductive—that the actual value of the asset is not affected by this information. Under this assumption, the owner expects to pay the bidders' precontract costs without receiving any benefits. Therefore, he would like to minimize these costs. As Johnson (1979) demonstrates, the owner might accomplish this within a sealed-bid auction by charging bidders an entry fee.

To see how entry fees can increase the owner's expected revenue, consider the effect of an arbitrary restriction on the number of bidders. If the restriction is binding, it will reduce the bidders' total estimation costs. However, as figure 4 illustrates, it will not increase the owner's expected revenue.[6] Reducing the number of bidders from n^* to n' just reduces the expected value of the winning bid from $E(B^*)$ to $E(B')$. In fact, the firms that are allowed to bid are the benefactors of this strategy. They capture both the reduction in the expected value of the winning bid and the reduction in the estimation costs. When the number of bidders is restricted, each bidder's expected profit is positive,

$$\begin{aligned} E(\Pi_{n'}) &= \{E(V) - E[B(n')] - n'C\}/n' \\ &= \{E[B(n^*)] - E[B(n')] + (n^* - n')C\}/n' \qquad (6) \\ &> 0. \end{aligned}$$

This suggests a more profitable way for the owner to reduce the estimation costs. Potential bidders would be willing to pay for the right to enter the auction if the asset owner limits the number of bidders to less than n^*.

Suppose the owner charges an entry fee that is equal to the per firm profit in equation (6),

$$F = \{E(V) - E[B(n')] - n'C\}/n', \qquad (7)$$

5. A number of economists, including Riley and Samuelson (1981) and Milgrom and Weber (1982), have examined the issue of optimal auction design. However, they generally assume that there is an exogenously specified set of bidders and that these bidders are endowed with information; they ignore both potential entry and bid preparation costs.

6. Figure 4 is basically a reproduction from Johnson (1979, p. 316).

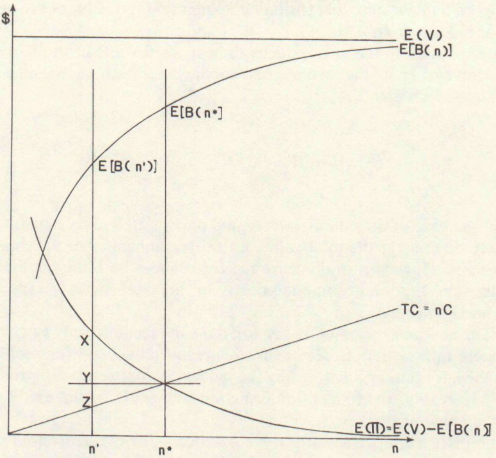

FIG. 4.—The effect of arbitrarily restricting the number of bidders. If the asset owner arbitrarily restricts the number of bidders to n', his expected revenue falls from $E[B(n^*)]$ to $E[B(n')]$. The bidding firms capture both the reduction in the expected value of the winning bid, $\overline{XY}$, and the reduction in the total estimation costs, $\overline{YZ}$.

instead of restricting entry. If this fee is paid when a firm purchases its sample of information, it will not affect the bidding strategy of any firm that does purchase information. Therefore, the function relating the expected value of the winning bid and the number of bidders, $E[B(n)]$, is unaffected.[7] However, the entry fee will change the number of bidders. Firms will enter until the expected profit from bidding equals the sum of the estimation costs and the entry fee,

$$\begin{aligned} \{E(V) - E[B(n)]\}/n &= C + F \\ &= C + \{E(V) - E[B(n')] - n'C\}/n' \qquad (8) \\ &= \{E(V) - E[B(n')]\}/n'. \end{aligned}$$

7. The problem becomes more complicated if firms can choose whether to pay the entry fee after they have purchased their information. In this case the relation between the expected value of the winning bid and the number of firms purchasing information is affected by the entry fee. For example, suppose the most optimistic firm's adjusted estimate is below the entry fee. If the fee is paid after firms purchase their information, no one will bid. Riley and Samuelson (1981) and Milgrom and Weber (1982) incorporate this effect in their examination of entry fees. However, they abstract from the effect of entry fees on the amount of information collected.

As figure 5 illustrates, this particular submission fee reduces the number of bidders from n^* to n'.

The submission fee also reduces the expected value of the winning bid. However, now the owner's revenue includes both the winning bid and the submission fees,

$$\begin{aligned} E(R) &= E[B(n')] + n'F \\ &= E[B(n')] + E(V) - E[B(n')] - n'C \qquad (9) \\ &= E(V) - n'C. \end{aligned}$$

As in the case of no submission fee and unrestricted entry, the owner expects to recover the total value of the asset minus all of the estimation costs. However, in this case the fee reduces both the number of bidders and the total estimation costs, so the asset owner's expected revenue is higher.

It appears that the asset owner will raise the fee until only two firms compete in the auction. However, this result reflects our assumption that the information generated by the bidders is completely nonproductive.[8] In reality, the bidders' information frequently increases the expected value of the asset.

The increase in the asset's value can occur for two reasons. First, the bidders' search may improve the allocation of the asset so that the highest valued user receives it. For example, if some oil companies specialize in deep wells while others specialize in shallow wells, prebid exploration can help identify which companies should bid aggressively for a particular oil lease.[9] Prebid information can also increase the expected value of the asset by improving the production plans of competing firms. Both of these factors imply that the expected value of the asset increases with the number of bidders. Each new bidder may discover that he is particularly well suited for this asset or he may find a particularly efficient plan for using the asset.[10]

8. If information is entirely nonproductive, the asset owner would actually like to eliminate all investment in information. In this case the owner would generally not choose an auction as the sales mechanism. He might post a price or negotiate with a single buyer. We discuss this topic in more detail below. See Barzel (1982) for a general discussion of this issue.

9. We would like to thank Robert Hansen for suggesting this example. The bidding strategy described in the App. is not the optimal strategy if the true value of the asset may be different for different bidders. Milgrom and Weber (1982) examine the optimal strategy for this case.

10. If these two types of information were produced separately, they would have different effects on the owner's expected revenue. Over some range of investment, the benefit of improved allocation may be larger than the cost of the information. On the other hand, prebid production of information about the efficient way to use the asset always reduces the owner's expected revenue. The winner can produce the "right" amount of information—without duplication—after the auction. However, the production of these two types of information usually cannot be divorced.

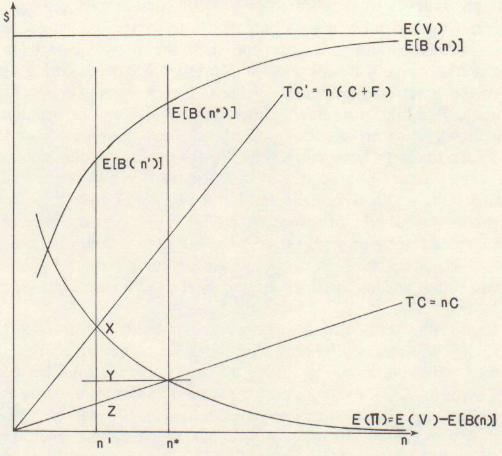

FIG. 5.—The effect of entry fees. If the owner charges an entry fee of F he reduces the number of bidders from n^* to n'. The expected value of the winning bid falls by $\overline{XY}$, from $E[B(n^*)]$ to $E[B(n')]$. However, the owner also receives $n'F$, or $\overline{XZ}$, in fees. Therefore, his total expected revenue increases by $\overline{YZ}$, the reduction in the estimation costs.

When the value of the asset varies with the number of bidders, the owner's selection of an entry fee—or, equivalently, the number of bidders—becomes more interesting. His expected revenue is still the expected value of the asset minus all of the estimation costs,

$$E[R(n)] = E[V(n)] - nC. \tag{10}$$

Maximizing expected revenue requires that

$$\partial E[R(n)]/\partial n = \partial E[V(n)]/\partial n - C = 0. \tag{11}$$

The asset owner will increase the entry fee and reduce the number of bidders until the marginal change in estimation costs equals the marginal change in the expected value of the asset.[11]

We have already detailed some of the effects that changing the number of bidders can have on the expected value of the asset. The value of

11. In reality the problem is more complicated than this. If buyers can choose the level of information to purchase, that is, if they have the option of purchasing any number of information samples, then C is not exogenous.

products which are complicated or unique, such as rare art objects, may vary considerably across bidders. In this case, increasing n increases the probability that a user for whom the true value of the asset is particularly high will bear the cost of preparing and submitting a bid. When the asset is simple or its value is easily estimated, the seller stands to gain little from having many buyers invest in information.[12] This allows us to make some predictions about the type of auctions which are likely to have entry fees. For example, ongoing auctions, where the product is standardized and the seller has a well-established brand name, are likely candidates for fees. In this case the cost of information acquired by prospective buyers is borne by the seller with little concomitant gain in revenues. On the other hand, auctions for unusual products are less likely to involve entry fees. By this reasoning, estate or land auctions are not likely to impose a fee for bidding.

The model in Section II also ignores the possibility of collusion. This possibility will almost certainly lead the asset owner to increase the number of bidding firms by setting a lower entry fee than he would otherwise. In fact, if it is very easy for the bidders to collude, the asset owner may seek other ways to hawk his wares. However, without building a model of the process of collusion we cannot make more specific predictions.

B. The Production of Information by the Asset Owner

On average, the asset owner pays for the information that is produced in a sealed-bid auction. Therefore, he would like that information to be produced efficiently. One way to accomplish this is by generating some information himself. This eliminates the duplication of effort that occurs when several bidders produce the same information. For example, the government might perform seismographic tests on an offshore oil tract and release this information to potential bidders.

This solution is not as simple as it may seem because of the conflict of interest between the owner and the bidders. The owner has an incentive to provide optimistic information about the value of the asset. In fact, unless this incentive can be controlled, the owner may be ignored.

There are several ways for the owner to establish his credibility. For example, he might post a bond that will be given to the winning bidder if his information is false. If the owner plans to hold other auctions, he may not even have to post a bond. As Klein and Leffler (1981) demonstrate, the adverse effects on his future revenues can act as a bonding

12. The cost of preparing a bid on a simple item will often be low. Therefore, even though the gains from having many bidders are small, the costs are also small. The conclusions which follow are based on the presumption that the former effect dominates.

mechanism. Alternatively, the owner might hire an independent firm, such as an auditor or a bond rating agency, to provide information.

The owner may also increase his credibility by retaining a partial claim to the asset. For example, on a mineral or timber sale, the owner may accept partial payment through a preestablished royalty fee. Selling the asset this way does not increase the cost of misrepresenting the truth, but it does reduce the gains. This reduces the owner's incentive to provide false information about the asset's value. In addition, it reduces the buyer's demand for information.

These factors allow us to make several predictions. First, the owner or his agent will produce objective information that can be communicated easily. For example, we expect an art dealer to estimate and reveal the age of an ancient Chinese vase. However, we do not expect him to spend many resources trying to inform potential bidders about the aesthetic qualities of the vase. We also predict that owners who return to the market frequently are more likely to provide information themselves while other owners will hire outside agencies for this purpose.

The Treasury uses a slightly different approach to reduce the production of duplicative information in T-bill auctions; small traders are allowed to free ride on information produced by large traders. Potential buyers can submit two different types of bids to the Treasury. The first, called a competitive bid, states both the price and the quantity the bidder is willing to buy at that price. The second type is called a noncompetitive bid. These bids, which are limited to less than $500,000, indicate the quantity of bills a buyer is willing to purchase at the average of the accepted competitive prices. This mitigates the incentive of noncompetitive bidders to collect information since there is no marginal action they can take. Moreover, since competitive bidders know that other buyers will free ride, their incentive to purchase information is also diminished. Essentially, the auction mechanism forces one group of bidders to release some of their information to the other group. This reduces the incentive of both groups to produce information. We conjecture that this institutional arrangement is employed by the Treasury because the asset is standardized and the auction is repetitive. Both of these features reduce the gain from the collection of information.

C. Negotiations

The asset owner might use a number of other strategies to reduce the bidders' investment in information. For example, he might offer to sell the asset for slightly less than $E(V)$, the mean of the prior distribution. Two conditions must be satisfied for this strategy to be successful. First, potential buyers must believe that the asset is randomly selected from the prior distribution. In other words, the owner cannot presort his assets, selling low-valued items at a fixed price of $E(V)$ and using some other technique to sell the rest. Second, the potential buyers must be prevented from collecting any information. Without this restriction, buyers will only purchase the asset when their adjusted estimate of its value is above the price.

De Beers diamond sales appear to satisfy both of these conditions.[13] Their long-standing reputation makes it possible for them to offer diamonds at a price slightly less than their expected value while limiting the opportunities of their customers to collect duplicative information. Invited customers submit a shopping list describing the number, weight, and character of diamonds they want to purchase. De Beers uses this information to prepare a collection of diamonds, called a "sight," for each customer. The sight is offered as a bundle on a take-it-or-leave-it basis. Buyers who reject their sight are not offered a substitute. They can only verify the classification of each diamond. If a buyer rejects his sight he is not invited to return. This institutional arrangement virtually eliminates the acquisition of information, save for the initial one by De Beers themselves. Presumably it is their long-standing reputation that warrants the quality of the information and prevents the duplication of information. This implies that an alternative arrangement was employed in the early days of the organization.

The De Beers approach is not common. Presumably this is because it is difficult to control the seller's incentive to offer only unexpectedly low-valued assets for sale and to control the buyer's incentive to identify underpriced goods. There is an alternative way for the owner to solve these adverse selection problems; he can negotiate with one of the potential bidders. In this case both the buyer and the seller collect some information.

There are costs associated with negotiation. First, the owner is giving some monopsony power to the firm with which he chooses to negotiate. Since the buyer has some information that is not available to the seller, the seller cannot capture the full updated estimate of the asset's value. A second cost arises if the actual value of the asset varies with the identity of the buyer. The buyer chosen by the owner might not place the highest value on the asset. Finally, the actual process of negotiation can be expensive. We expect firms to use negotiated sales instead of auctions if these costs are small when compared with the costs of bid preparation. For example, if there is not much dispersion in the true value of the asset across the potential buyers (or if the owner can determine the highest valued user ex ante), the asset owner is more

13. See Barzel (1982) for a general discussion of measurement cost and its specific application to diamond sales. Kenney and Klein (1983) also discuss the problem in the context of block booking.

likely to use a negotiated sale. Also, an ongoing relationship between the owner and a potential buyer will tend to reduce the negotiation costs and increase the probability that the owner will choose a negotiated sale.

IV. Sunk Costs in Other Markets

The bidding model of Section II can be extended to a number of consumer goods markets. For example, the usual process for purchasing home, television, and camera repairs is very similar to a sealed-bid auction; each repair shop forms an estimate of its costs and submits a bid without knowing its competitors' bids.[14] However, there is one important difference. In the auctions examined in Section II, the number of bidders is determined by the bidders; they enter until the next bidder's expected profit would be negative. In the consumer goods markets, the individual purchasing the service usually determines the number of bidders. This distinction leads to substantial differences in the operation of these markets.

Consider the problem from the point of view of the prospective buyer in a consumer goods market. He will collect bids until the marginal cost of search equals the expected marginal benefit. Assume that each buyer forms an expectation of the costs and benefits of search and calculates the optimal number of bids to collect before he begins searching. Further, assume that each bidder makes an accurate forecast of this number and uses it as an input in the bidding model of Section II.

The buyer's expected marginal benefit of search is equal to the expected reduction in the price he has to pay because there are more bidders. Two factors contribute to this reduction. First, there may be differences in the true cost of the product or service across bidders. For example, in the process of estimating their bids, different repairmen may discover different ways to fix a broken camera. Since each new bidder may discover a cheaper process, the expected cost falls with the number of bidders. The second benefit from search arises because of the finite-bidders adjustment discussed in Section II. If the buyer chooses to canvass n repair shops instead of $n - 1$ shops, the bids he receives will reflect the increased competition. Each bidder will reduce the amount he raises his bid above his adjusted estimate of the repair cost, so the winner's expected profit falls. Both of these factors

14. Preliminary research suggests that this model of the recovery of sunk costs may have more general applications. Specifically, the bidding model may make better predictions about price than conventional models even in markets that do not have this auction-like nature. We are currently working on this topic.

imply that the price the buyer expects to pay for the service falls as he increases the number of bidders.[15]

Figure 6 illustrates the buyer's problem. In this figure, nC is the total estimation costs for all of the bidders and $E(V)$ is the winner's expected cost of performing the service. This expected cost falls as the number of competitors increases because each new firm may discover a more efficient way to produce the service. The winning bidder's expected profit, ignoring estimation costs, is equal to the difference between the expected value of the winning bid and the winner's expected costs, $E(\Pi) = E[B(n)] - E(V)$.

In a normal sealed-bid auction, equilibrium occurs when there are n^* bidders and the winner's expected profit equals the industry's total estimation costs. However, the searcher will choose n^* only by accident. Instead, the buyer goes to the point where the cost of acquiring one more bid, S, equals the reduction in the expected value of the winning bid. Figure 6 depicts three of the many possibilities. If the buyer's marginal cost of search is $(-)S_1$, he will choose n_1 bidders; if the cost of search is S_2 he will choose n^*; and if the marginal cost is S_3, he will solicit n_3 bids.

In the first case, a profitable opportunity exists for the lucky n_1 bidders chosen by the buyer. Nonsolicited retailers would like to participate in the auction, but in order to do so they must entice the buyer into their stores by lowering his search costs. For example, they might increase their advertising, move to a more desirable location, or include a gift with each estimate. These strategies have two effects: they reduce the buyer's search costs so that more bids are collected, and they increase the seller's prebid costs. If this situation is common to many buyers, other firms will enter the industry and set up shop. This also reduces the buyer's search cost. When there is competition among the sellers, this multipronged process will continue until there are no marginal profits in the industry.

By contrast, in those situations where the search costs are relatively low the buyer will search too much from the sellers' point of view. This is the case at point n_3 in figure 6. The industry cannot survive in the long run because the expected profits from bidding are below the costs. There are several options available to the sellers. Some may leave the industry. Those remaining are inclined to reduce their advertising, hire less congenial salespeople (at lower cost), or charge for their estimates. Each of these responses increases the buyer's search costs and reduces the seller's estimation costs.

15. The benefit of search is slightly different if the retailers are not able to observe the number of bids being collected. In this case, each bid is drawn from a distribution that does not change when a particular individual solicits more bids. Although the benefit from search is now equal to the change in the expected value of the order statistic, the general results are the same.

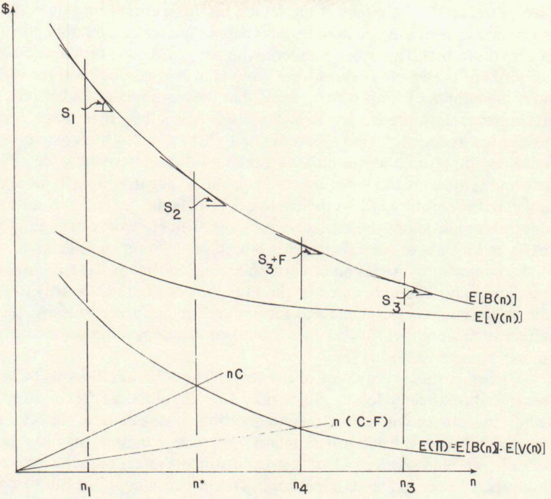

FIG. 6.—The relation between search and estimation costs. If search costs equal S_1, each bidding firm expects to earn a profit. Competitors will try to attract the customer by advertising, etc. If search costs equal S_2, the bidder's expected profit is zero. If search costs equal S_3, each bidder expects to lose money. Firms may respond by charging an estimation fee, F.

Consider the dynamics of the adjustment process. Suppose customers are currently canvassing too many repair shops such that the expected profit from tendering a bid is negative for each firm in the industry. Imagine that the shops and the customers are scattered over some area and that travel time is an important part of search costs. If one retailer tries to reduce his losses by charging an estimation fee, he will affect his and his competitors' profits. This retailer's current customers have three choices: they can pay the estimation fee and collect his bid; they can balk at the fee and visit another retailer; or they can simply choose to collect one less bid. The second and third strategies are likely to be popular with customers who are traveling from fairly distant locations. However, unless the fee is prohibitive, some (inframarginal) customers will still visit this particular repair shop.

The fee reduces the number of bids offered by the retailer, and it reduces the net cost of each bid. On both counts the retailer's expected loss is lowered. The effect of the fee on his competitors' profits is not so clear. Consider the customers who balk at the fee and do not collect a substitute bid. The other retailers who are canvassed by these customers are better off because they now face less competition for the winning bidder's profit, which means that they can tender higher bids. Because their finite-bidder adjustments are smaller, their expected losses are reduced. This result does not carry over to the customers who continue to collect n bids by seeking out an alternate store. By assumption, each bidder's expected profit is negative if n bids are collected. Therefore, when the customer rejects the fee and canvasses another repair shop, the replacement store is made worse off and the expected profits of the other stores who offer bids remain negative.

To summarize, if customers are collecting too many bids and one retailer begins charging an estimation fee, he does not necessarily lose all of his customers. While the fee reduces his expected losses, the effect on his competitors' profits is ambiguous, depending on the customer's response to his fee.

Of course, this is only part of the story. The other retailers will also react to the expected losses. Some may drop out of the market entirely, while the others will adjust their marketing strategies in a variety of ways. They may introduce estimation fees, reduce their advertising, or offer less prebid service. All of these adjustments affect both the initiating repair shop and its competitors. This process continues until no store expects to make a profit or lose money by tendering a bid.[16]

Three factors determine whether the buyer's optimal number of bids, n', is originally above or below the industry equilibrium: the cost of bid preparation, the cost of search, and the benefit of search. This model predicts that in industries where buyers face low search costs, where the payoffs to search are high, and where bid preparation costs are high, sellers are likely to charge for their estimates. Again, consider camera repairs. Because the product is complicated, each serviceman must spend a relatively large amount of time examining the camera before he provides a binding estimate. At the same time, the benefits from search are high because there is a large probability that an additional bidder will discover an unusually efficient repair technique. By contrast, some services, such as carpet installation, roof repair, and

16. There is yet another process which has evolved to solve the "too many bids" problem. Some repair shops provide a precursory, nonbinding bid for free. When the customer accepts a bid, the shop updates its estimate as it undertakes the repair. If this updated bid is above the original estimate, the customer can decide that the repair should not be completed. This process reduces the number of binding bids collected for two reasons. First, the direct search costs are increased. Second, the benefits of search are truncated because binding estimates must be collected sequentially; since a customer cannot return to a shop after balking at an updated estimate, he must sample without recall.

driveway paving, are relatively homogeneous and easy to estimate. Moreover, there is usually not a large chance that the buyer can uncover a particularly low-cost repair method. Both of these factors imply that camera repair shops are more likely to charge for binding estimates, while home repair shops are less likely to charge. Further, the home repair firms will probably do more advertising.[17]

This model also predicts that stores with different clienteles will use different marketing strategies. For example, repair shops in high-income neighborhoods are likely to face customers with higher search costs than shops in poor neighborhoods.[18] Therefore, high-income shops are more likely to advertise than low-income shops.[19] On the other hand, repair shops in poor neighborhoods are more likely to charge for estimates.

Firms can also affect a buyer's search in other ways. For example, when prospective customers are collecting too few bids, stores can increase their hours of operation, making it cheaper to collect information. Camera and auto body repair shops are more likely to be open at night and on weekends in rich, high-search-cost neighborhoods. A shop's location and general attractiveness also affects potential customers' search costs. We expect that stores catering to high-search-cost customers will be very visible and accessible. They will also be cleaner, less crowded, and more fully staffed than their low-search-cost counterparts.[20]

We note one other implication of the analysis for marketing strategy. In general, repair shops will not attract customers who have homogeneous search costs. Consider the case described in figure 6. There are three customers with search costs S_1, S_2, and S_3, soliciting bids from a

17. Again, the problem is confounded if the simple product is easy to estimate. In the examples chosen, prebid costs are relatively high even though the products are simple, because travel time is an important component of the repairman's bid preparation cost.

18. In a household production sense the effect of income on search costs is ambiguous. In general, rich people will have a high opportunity cost of time, but they may be efficient searchers. If the marginal productivity of time spent searching increases with income, the effect of income on search costs must be determined empirically. However, if the marginal productivity of search increases at a decreasing rate over the relevant income range, the opportunity cost of time will eventually dominate. This implies that rich people face higher search costs than poor people.

19. Obviously, there are other reasons why firms might advertise or charge for bids. For example, advertising might be used to develop brand name capital as a bonding device (see Klein and Leffler 1981). If a firm is charging for its bids, it may still advertise as a bonding device, but the advertising should take a particular form. It will not, for example, emphasize price information or other facts that reduce the customer's search cost. The ads may not even give the address of the establishment. Instead, the ads will tend to promote the reputation of the firm. Similarly, estimation fees might be used in a two-part pricing scheme (see Oi 1971; Murphy 1977). See Butters (1977) for an additional discussion of advertising in the context of bidding models.

20. For a discussion of product quality and the value of time, see Becker (1965) and De Vany (1976).

particular shop. If the store operator accurately forecasts the number of bids collected by each potential buyer, he expects that his profits will be positive from customer 1, zero from customer 2, and negative from customer 3. He has an incentive to discourage number 3. At the same time, the threat of losing the expected profits to his rivals induces him to find ways to lure and keep buyer 1 in his store. In other words, because of their differing search costs, the retailer will subsidize buyer 1 and tax buyer 3.

It is usually expensive to distinguish buyer 1 from buyer 3. However, the retailer may still be able to segment the market by allowing the customers to self-select. The firm can adjust its general marketing strategy so that each buyer is effectively charged the fee that is required to make the expected profits from buyer 3 positive. Then, by tailoring specific components of its marketing mix, such as advertising, the firm can rebate this fee to buyers 1 and 2. For example, suppose buyer 1 has high search costs because he must travel a great distance. The firm can reduce this person's cost of search by using locationally specific advertising. Alternatively, if buyer 1's cost of search is high because of his time cost and income, there may be certain magazines, television programs, or sections of the newspaper that are more likely to appeal to rich people.

By adjusting the marketing mix, the firm can effectively alter the search costs of most buyers so that bids are not tendered on repair jobs that are expected to lose money, and so that clients who represent a profit opportunity are enticed into the store to solicit a bid. Effectively, each buyer's search cost is altered so that it is S_2.

This strategy contrasts with the price discrimination incentive to segment markets. In general, it is anticipated that rich, high-search-cost persons have less elastic demand than their poorer, lower-search-cost counterparts. Therefore, the firm with monopoly power will raise its price for the rich person and lower it for the poor one. By contrast, our theory predicts that because of competition, market segmentation will lower the full price to the rich person. Prizes, promotions, and gifts will be aimed more at the rich, high-search-cost customer. The practice of offering private wine tastings prior to an auction which is open to the public would seem to be an example of this principle.

V. Conclusions

The textbook solution to the recovery of sunk costs depends on rising marginal costs. We have developed a model of the competitive process that does not rely on this assumption. The central feature of this model, which is based on the bidding models of Wilson (1977), Johnson (1979), and Ramsey (1980), is the distinction between competitive equilibria before and after resources are committed. Ex ante, there are an infinite number of potential competitors. This guarantees that a zero-profit equilibrium will prevail. However, after the initial investments are made only a finite number of competitors remain. It is this reduction that allows the remaining firms to recover their now-sunk estimation costs. With a finite number of competitors, each firm recognizes that there will be a gap between its adjusted estimate of the asset's value and the adjusted estimate of its closest rival. Each firm exploits this gap by offering less than its adjusted estimate. In other words, the winning firm is expected to earn a profit. This profit induces firms to invest in bid preparation. Firms enter the auction until the winner's expected profit equals the industry's estimation costs.

Since the winning firm captures the industry's precontract costs, the value of each of the $(n - 1)$ losing firms falls by its estimation costs. The fact that the stock price of a defense contractor, an offshore oil lessee, or a building contractor rises on announcement that it has won a sealed-bid auction does not, by itself, constitute evidence that the bidding-process was not competitive.

If the winning bidder captures the industry's estimation costs, the asset owner must expect to pay these costs. This simple result has some important implications. For example, consider the auctions that the U.S. government uses to award offshore oil leases. In a recent paper analyzing the distribution of rents in these auctions, Reece (1978) reports several simulations relating the number of bidders and the size of the winning bid. His results are consistent with our hypothesis in Section II; holding all other variables fixed, the expected value of the winning bid increases as the number of bidders increases. Further, he finds that "the expected value of the winning bid may be substantially less than the true value of the lease . . ." when the bidding firms make reasonable assumptions about the number of competitors they face (Reece 1978, p. 383).

Both our analysis and Reece's simulations are based on Wilson's (1977) model of competitive bidding. However, there is an important difference between the two applications of this model: Reece does not include bid preparation costs in his theory. This difference leads to considerably different interpretations of his results. Reece views the predicted divergence between the lease value and the winning bid as evidence of a lack of competition and a sign that the oil industry may be capturing a "remarkably large fraction of the economic rent" associated with the leases. Our interpretation of these results is significantly different. When prebid costs are included in the model, this predicted divergence provides no evidence that the oil industry is capturing rents. It simply reflects the fact that the winning bidder must, on average, recover the industry's estimation costs. In other words, even in an auction with identical firms and unrestricted entry, the government's expected revenue will not equal the expected lease value unless there are no bid preparation costs.

Appendix

In this Appendix we demonstrate that, under the assumption outlined in Section II, each firm's optimal bid occurs at the point where the expected marginal revenue from lowering the bid 1 dollar equals the expected marginal cost. Wilson (1977) shows that, for this bidding model, the optimal bid for each competitor must satisfy the condition[21]

$$\int_0^\infty [V - b(I_i, n)]Q_n'(I_i|V)\sigma_n'[b(I_i, n)]f(I_i|V)dG(V) = \int_0^\infty Q_n(I_i|V)f(I_i|V)dG(V), \tag{A1}$$

where V = the (unknown) value of the asset,
I_i = the information sample observed by firm i,
n = the total number of bidding firms,
$b(I_i,n)$ = the optimal bid based on the sample of information I_i when there are n bidders,
$G(V)$ = the prior distribution for the value of the asset,
$F(I_i|V)$ = the cumulative density function for I_i when the asset's value is V,
$f(I_i|V)$ = the probability density function of $F(I_i|V)$,
$Q_n(I_i|V)$ = $F(I_i|V)^{n-1}$ = the probability that no other firm receives a higher sample of information than firm i, conditional on the value of the asset being V,
$Q_n'(I_i|V)$ = $(n - 1)F(I_i|V)^{n-2}f(I_i|V)$
$\sigma_n(b)$ = the level of information a firm must receive to bid b, $\sigma_n[b(I_i, n)] = I_i$,
$\sigma_n'(b)$ = the derivative of $\sigma_n(b)$ with respect to b.

Consider the right-hand side of equation (A1). If the value of the asset is V, $Q_n(I_i|V)$ is the probability that no other firm receives a higher sample than I_i. Since the firms are identical except for their sample of information, this is also the probability that firm i wins the auction. This probability is weighted by the firm's marginal probability that the asset's value is V, $f(I_i|V)dG(V)$. Integrating this weighted probability over all possible values of V gives the firm's perceived probability of winning the bid. This is also the expected marginal revenue from lowering the bid 1 dollar since the gain from this reduction is only realized if firm i wins the bid.

The left-hand side of equation (A1) is the expected marginal cost of lowering the bid by 1 dollar. If firm i wins, its payoff is $V - b(I_i, n)$. The marginal

21. In fact, Wilson deals with a more general model in which the winner's payoff is not necessarily equal to V minus the winning bid. See Wilson for the regularity conditions required to prove eq. (A1).

probability that the most favorable sample observed by any other firm is I_i is given by $Q_n'(I_i|V)$. The function, $\sigma_n[b(I_i, n)]$, gives the minimum sample another firm would have to observe to bid higher than firm i, so $\sigma_n'[b(I_i, n)]$ is the amount firm i lowers this minimum when it reduces its bid by 1 dollar. Therefore, if the asset's value is V, $Q_n'(I_i|V)\sigma_n'[b(I_i, n)]$ is the amount firm i reduces its probability of winning by decreasing its bid 1 dollar and $[V - b(I_i, n)]Q_n'(I_i|V)\sigma_n'[b(I_i, n)]$ is the amount it lowers its expected payoff. Weighting this amount by the marginal probability that the asset's value is V, $f(I_i|V)dG(V)$, and integrating over all possible values of V gives the reduction in the firm's expected payoff when its bid is lowered by 1 dollar. In other words, equation (A1) says that firm i should decrease its bid until the marginal cost of lowering it 1 dollar more equals the marginal revenue.

References

Baron, D. P. 1972. Incentive contracts and competitive bidding. *American Economic Review* 62 (June): 384–94.

Barzel, Y. 1982. Measurement cost and the organization of markets. *Journal of Law and Economics* 2 (April): 27–48.

Becker, G. 1965. A theory of the allocation of time. *Economic Journal* 75 (September): 493–517.

Butters, G. R. 1977. Equilibrium distributions of sales and advertising prices. *Review of Economic Studies* 44 (October): 461–95.

Coppinger, V. M.; Smith, V. L.; and Titus, J. A. 1980. Incentives and behavior in English, Dutch, and sealed-bid auctions. *Economic Inquiry* 18 (January): 1–22.

De Vany, A. 1976. Uncertainty, waiting time, and capacity utilization: A stochastic theory of product quality. *Journal of Political Economy* 84 (June): 523–41.

Englebrecht-Wiggans, R. 1980. Auctions and bidding models: A survey. *Management Science* 26 (1980): 119–42.

Goldberg, V. P. 1977. Competitive bidding and the production of precontract information. *Bell Journal of Economics* 8 (Spring): 250–61.

Holt, C. A. 1980. Competitive bidding for contracts under alternative auction procedures. *Journal of Political Economy* 88 (June): 433–45.

Johnson, R. N. 1979. Auction markets, bid preparation costs and entrance fees. *Land Economics* 55 (August): 313–18.

Kenney, R. W., and Klein, B. 1983. The economics of block booking. *Journal of Law and Economics* 26 (October): 497–540.

Klein, B., and Leffler, K. B. 1981. The role of market forces in assuring contractual performance. *Journal of Political Economy* 89 (August): 615–41.

Milgrom, P. R., and Weber, R. J. 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (September): 1089–1122.

Murphy, M. 1977. Price discrimination, market separation, and the multi-part tariff. *Economic Inquiry* 15 (October): 587–99.

Oi, W. 1971. A Disneyland dilemma: Two-part tariffs for a Mickey Mouse monopoly. *Quarterly Journal of Economics* 85 (February): 77–96.

Ramsey, J. B. 1980. *Bidding and Oil Leases*. Greenwich, Conn.: JAI Press.

Reece, D. K. 1978. Competitive bidding for offshore petroleum leases. *Bell Journal of Economics* 9 (Autumn): 369–84.

Riley, J. G., and Samuelson, W. F. 1981. Optimal auctions. *American Economic Review* 72 (June): 381–92.

Vickery, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (May): 8–37.

Wilson, R. 1967. Competitive bidding with assymetric information. *Management Science* 13 (July): 816–20.

Wilson, R. 1977. A bidding model of perfect competition. *Review of Economic Studies* 4 (1977): 511–18.