1986

Auctions and Bidding

R Preston McAfee

John McMillan

ISSN:0318-725X
ISBN:0-7714-0707-6

RESEARCH REPORT 8601

AUCTIONS AND BIDDING

R. Preston McAfee
John McMillan

University of Western Ontario

December, 1985

AUCTIONS AND BIDDING*

By

R. Preston McAfee and John McMillan

University of Western Ontario

December, 1985

ABSTRACT

This paper surveys recent developments in the theory of bidding mechanisms.

Correspondence to: John McMillan
Department of Economics
University of Western Ontario
London, Ontario N6A 5C2, Canada

1. WHY STUDY AUCTIONS?

One party to an exchange often knows something relevant to the transaction that the other party does not know. Such asymmetries of information are pervasive in economic activity: for example, in the relationship between employer and employee when the employee's effort cannot be monitored perfectly; between the stockholders and the manager of a firm; between insurer and insured; between a regulated firm and the regulatory agency; between the supplier and the consumers of a public good; between a socialist firm and the central planner; or (as is the subject of this paper) between buyer and seller when the value of the item is uncertain.

Forty years ago, F.A. Hayek criticized theories which purport to describe the price system but start from the assumption that individuals have symmetric information.

> The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus not merely a problem of how to allocate "given" resources--if "given" is taken to mean given to a single mind which deliberately solves the problem set by these "data". It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge which is not given to anyone in its totality (Hayek (1945, p. 519)).

Hayek went on to argue that to omit imperfections of information is to ignore the price system's chief advantage. The "marvel" of the price system is its efficiency in communicating information "in a system in which the knowledge of the relevant facts is dispersed among many people" (p. 525). All that a buyer or seller needs to know about a commodity's supply or demand is summarized by a single number, its price.

Hayek's critique of extant theories of the price system applies equally to the most thorough current theory, the Arrow-Debreu model, which assumes either that information is perfect or, what is essentially the same thing, that a full range of markets in contingent commodities exists. To paraphrase Hayek in modern terms, the constraints imposed by informational asymmetries can be as significant as any resource constraints.

Is the market an effective transmitter of information, as Hayek argued? The much-cited examples due to George Akerlof (1970) illustrate that it need not be. The inability of buyers of used cars to observe the quality of any one car might cause the used-car market to cease to function. Similarly, the working of medical-insurance markets is hindered by the inability of an insurance company to observe completely an individual's current state of health. Indeed, the general phenomenon of one individual's having information that is not available to others is called <u>adverse selection</u>. Either this is a misnomer or Hayek's claims for the informational efficiencies of the price system are unduly optimistic. Resolution of this question requires some systematic analysis of how economic agents behave when information is dispersed.

Some of the most exciting of the recent advances in microeconomic theory have been in the modelling of strategic behavior under asymmetric information. One part of this broad research program is the theory of bidding mechanisms. The modelling of auctions provides a narrowly defined set of questions with which to begin a rigorous examination of the implications for the price system of informational asymmetries.

The study of auctions provides one way of approaching the question of price formation. As was pointed out in a well-known article by Kenneth Arrow (1959), the standard economic model of small buyers and sellers, taking as given the market price, is lacking in that it fails to explain where prices come from. Once the deus ex machina of the Walrasian auctioneer is discarded, who sets prices? Auction theory provides one explicit model of price making.

A less fundamental but more practical reason for studying auctions is that auctions are of considerable empirical significance: the value of goods exchanged each year by auction is huge. This fact in itself indicates that some theoretical study of auctions is warranted. Moreover, as will be seen, the theory of auctions is closer to applications (in having less need of oversimplified assumptions) than is most frontier mathematical economics.

The theoretical results in auction theory can explain why sellers adopt particular practices, and perhaps can even suggest improvements which would on average increase the seller's revenue. Many of the results address the question: what is, from the point of view of the monopolist, the best form of selling mechanism to use in any particular set of circumstances? Other questions that can be answered include the following. Should the seller impose a reserve price? If so, at what level? Can the seller design the auction so as to achieve price discrimination among the bidders? Is it ever in the seller's interest to require payment from the unsuccessful bidders? If it is feasible to make payment depend not only on the bid but also on something correlated with the true value of the item (as is the effect of royalties, for example), should the seller do so? Should the seller release any information he has about the item's true value? What can the seller do to counter collusion among the bidders?

This paper surveys recent developments in the theory of bidding mechanisms and discusses the relevance of the theoretical results for auctions in practice.[1]

In what follows, theorems will be stated in italics. For ease of exposition, not only will the proofs of the stated results be omitted, but also some of the required technical assumptions. For precise statements of the assumptions, as well as proofs, see the cited papers.

2. THE TYPES AND USES OF AUCTIONS

What kinds of goods are sold by auction? The list is long: artwork, books, antiques, agricultural produce, timber, United States Treasury bills, corporations, and gold are some current examples; wives and slaves are historical examples (Ralph Cassady (1967); Martin Shubik (1983)). Why are auctions used rather than other selling devices such as posting a fixed price? According to Cassady (1967, p. 20): "One answer is, perhaps, that some products have no standard value. For example, the price of any catch of fish (at least of fish destined for the fresh fish market) depends on the demand and supply conditions at a specific moment of time, influenced possibly by prospective market developments. For manuscripts and antiques, too, prices must be remade for each transaction. For example, how can one discover the worth of an original copy of Lincoln's Gettysburg Address except by auction method?"

The selling by auction of mineral rights to government-owned land has motivated much of the recent theoretical analysis of auctions. However, this procedure is almost unique to the United States: governments in other countries usually allocate mineral rights to firms by discretionary procedures (James Ramsey (1980, pp. 56-62)). It is noteworthy, however, that in 1982 and 1985 the People's Republic of China called for bids for some offshore exploration contracts (New York Times, May 23, 1985).

Sometimes there is a single buyer who wishes to purchase some item from one of a set of potential suppliers. From a theoretical point of view, monopsony is essentially the same as monopoly apart from reversal of the signs of some variables. Thus, although the Oxford Dictionary defines an auction as a "public sale in which articles are sold to maker of highest bid", we shall use the term "auction" to describe both bidding for the right to sell and bidding for the right to buy. (Nevertheless, for the sake of brevity, we shall usually discuss auctions as mechanisms for selling.)

Governments are the most prominent users of procurement auctions. In a modern market economy, the government's purchases from private firms typically account for about 10 percent of gross domestic product. For many government contracts, firms submit sealed bids; the contract is required by law to be awarded to the lowest qualified bidder. Sealed-bid tenders are sometimes also used by firms procuring inputs from other firms. There is, however, an important difference between procurement in the private sector and procurement in the public sector. Governments favor the use of sealed-bid auctions for political reasons: in a sealed-bid auction, with the envelopes containing the bids opened in public at an appointed time, it is difficult for an illicit deal between a government official and a bidding firm to be arranged. The private sector, in contrast, is less constrained to use procurement mechanisms which can be seen to be conducted honestly. While formal auctions are sometimes used in private-sector procurement, more often informal, closed negotiations are used (McAfee and McMillan (1985c).

What are the types of auctions that are in use? There are four basic types: the English auction (also called the oral, open, or ascending-bid auction); the Dutch (or descending-bid) auction; the first-price sealed-bid auction; and the second-price sealed-bid (or Vickrey) auction.

The English auction is the auction form most commonly used for the selling of goods. In the English auction, the price is successively raised until only one bidder remains. This can be done by having an auctioneer announce prices, or by having bidders call the bids themselves, or by having bids submitted electronically with the current best bid posted. (The word "auction" is derived from the Latin augere, which means "to increase".) The essential feature of the English auction is that, at any point in time, each bidder knows the level of the current best bid. Antiques and artwork, for example, are often sold by English auction.

The Dutch auction is the converse of the English auction. The auctioneer calls an initial high price and then lowers the price until one bidder accepts the current price. The Dutch auction is used, for instance, for selling cut flowers in the Netherlands, fish in Israel, and tobacco in Canada.

With the first-price sealed-bid auction, potential buyers submit sealed bids and the highest bidder is awarded the item for the price he bid. The basic difference between the first-price sealed-bid auction and the English auction is that, with the English auction, bidders are able to observe their rival's bids and accordingly, if they choose, revise their own bids; with the sealed-bid auction, each bidder can submit only one bid. First-price sealed-bid auctions are used in the auctioning of mineral rights to U.S. Government-owned land; they are also sometimes used in the sales of artwork and real estate. Of greater quantitative significance is the use, already noted, of sealed-bid tendering for government procurement contracts.

Under the second-price sealed-bid auction, bidders submit sealed bids having been told that the highest bidder wins the item but pays a price equal not to his own bid but to the second-highest bid (William Vickrey (1961)). While this auction has useful theoretical properties, it is seldom used in practice.

Many variations upon these four basic auction forms are used. For example, the seller sometimes imposes a reserve price, discarding all bids if they are too low (Cassady (1967, Ch. 16)). Bidders may be allowed only a limited time for submitting bids (Cassady (1967, pp. 74-76); Shubik (1983, pp. 45-49)). The auctioneer may charge bidders an entry fee for the right to participate (Kenneth French and Robert McCormick (1984)). Payment may be made to depend not only on bids but also on something correlated with the true value of the item, as is achieved for example by using royalties (Walter Mead, Asbjorn Moseidjord, and Philip Sorenson (1984)). In an English auction, the auctioneer sometimes sets a minimum acceptable increment to the highest existing bid (B.S. Yamey (1972)). The seller might, instead of selling the item as a unit, offer for sale shares in the item (Robert Wilson (1979)).

Two broad questions are prompted by the foregoing description of the use of auctions. First, why is an auction used rather than some other selling (or buying) procedure? Second, given the diversity of types of auctions, what determines which particular auction form is chosen? In order to address these questions, some theoretical machinery is needed.

3. THE ABILITY TO MAKE COMMITMENTS

Auctions are often used by a monopolist (an individual selling a unique work of art, a government selling mineral rights, etc) or a monopsonist (a government contracting out the production of a public good); they are sometimes also used in a competitive setting (in the selling fish or agricultural produce, for example). This survey will follow the existing literature by considering the case of monopoly or monopsony; competition among the organizers of auctions will be little discussed.

It is presumed therefore that there is monopoly (or monopsony) on one side of the market. While it is possible in auctions that there are very many bidders so that perfect competition prevails, more usually there are only a few bidders: there is oligopsony (or oligopoly) on the other side of the market. In classical economics, monopoly-oligopsony problems were regarded as indeterminate: any outcome between all of the gains from trade going to the buyer and all of the gains going to the seller was seen as possible. The best that could be hoped for by way of a solution was to apply some *ad hoc* bargaining model.

Auction theory sidesteps bargaining problems by presuming that, in a sense, the monopolist (or monopsonist) has all of the bargaining power. More precisely, it is assumed that the organizer of the auction has the ability to commit himself in advance to a set of policies. He binds himself in such a way that all of the bidders know that he cannot change his procedures after observing the bids, even though it might be in his interest *ex post* to renege. In other words, the organizer of the auction acts as the Stackelberg leader or first mover.

Commitment matters because even as simple an institution as the first-price sealed-bid auction leaves the seller with a temptation to renege. As will be seen, the bidders submit bids which are functions of their valuations of the item for sale. Given the assumptions we shall make about the seller's knowledge, the seller is able to deduce from a bid the bidder's valuation of the item. Thus it would be in the seller's ex post interest to renege on his promise to charge a price equal to the highest bid: instead, he could offer the item at a price slightly less than the highest valuation, and it would be in the interest of the bidder who has that valuation to accept this offer. Of course, if the bidders knew in advance that the seller might renege on his announced policy, they would not bid as hypothesized.

The advantage of commitment is that procedures can be adopted which induce the bidders to bid in desirable ways. In <u>The Strategy of Conflict</u>, Thomas Schelling explained the advantages in general strategic situations of commitment power: "if the buyer can accept an irrevocable <u>commitment</u>, in a way that is unambiguously visible to the seller, he can squeeze the range of indeterminacy down to the point most favorable to him" (Schelling (1960, p. 24)).

There are many ways commitment can be achieved. For example, in the case of government contracting, the government official responsible for the decision is required to follow procedures which are explicitly and precisely set out in a publicly available book of rules. Also, a "potent means of commitment, and sometimes the only means, is the pledge of one's reputation" (Schelling (1960, p. 29)): the cost of reneging on a current commitment might be the inability credibly to commit oneself in future transactions, and therefore the loss of future bargaining power.

Nevertheless, it does not follow from that fact that one party has the ability to make commitments that he can extract all of the gains from trade. What limits his bargaining ability is the asymmetry of information: the seller does not know any bidder's valuation of the item for sale. If the seller were able to observe bidders' valuations, he could offer the item to the bidder who values it the most at a price slightly below this bidder's valuation, threatening to refuse to sell it if this offer is rejected. Given that the seller has so committed himself, it is in the bidder's interest to accept this take-it-or-leave-it offer. When information is asymmetric, the seller's ability to extract surplus is more limited. The seller can exploit competition among the bidders to drive up the price; but usually the seller will not be able to drive the price up so far as to equal the valuation of the bidder who values the item the most, because the seller does not know what this valuation is.

In the next section, we discuss in detail the asymmetry of information about bidder's valuations.

4. THE NATURE OF THE UNCERTAINTY

Uncertainty is the crucial element of the auction problem. In the case of perfect information, the auction problem is easily solved, as just noted: given the ability to make commitments, the organizer of the auction extracts all of the gains from trade. Indeed, the reason a monopolist chooses to sell by auction rather than, say, simply posting a price is that bidders' valuations cannot be observed.

How the bidders respond to uncertainty depends on whether they are risk neutral or risk averse. Thus one aspect of any particular bidding situation which the modeller must take into account is the bidders' attitudes towards risk.

Differences among the bidders' valuations of the item can arise for either of two distinct reasons. Which of these is relevant also affects how any particular bidding situation is to be modelled.

At one extreme, suppose that each bidder knows precisely how highly he values the item: he has no doubt about the true value of the item to him. He does not know anyone else's valuation of the item; instead, he perceives any other bidder's valuation as a draw from some probability distribution. Similarly, he knows that the other bidders (and the seller) regard his own valuation as being drawn from some probability distribution. Differences among the bidders' evaluations reflect actual differences in their tastes. More precisely, for bidder i, i=1,...,n, there is some probability distribution F_i from which he draws his valuation v_i. Only the bidder observes his own valuation v_i, but all the other bidders as well as the seller know the distribution F_i. Any one bidder's valuation is statistically independent from any other bidder's valuation. This is called the <u>independent-private-values model</u>. This model applies, for example, to an auction of an antique in which the bidders are consumers buying for their own use and not for resale. It also applies to government-contract bidding when each bidder knows what his own production cost will be if he wins the contract.

At the other extreme, consider the sale of an antique which is being bid for by dealers who intend to resell it, or the sale of mineral rights to a particular tract of land. Now the item being bid for has a single objective value, namely the amount the antique is worth on the market, or the amount of oil actually lying beneath the ground. However, no one knows this true value. The bidders, perhaps having access to different information, have different guesses about how much the item is objectively worth. If V is the unobserved true value, then the bidders' perceived values v_i, $i=1,\ldots,n$, are independent draws from some probability distribution $H(v_i|V)$. All agents know the distribution H. This is called the <u>common-value model</u>.

Suppose a bidder were somehow to learn another bidder's valuation. If the common-value model describes the situation, learning someone else's valuation provides useful information about the likely true value of the item: the bidder would probably change his own valuation in the light of this. In contrast, if the independent-private-values model describes the situation, the bidder knows his own mind; learning about another's valuation would not cause him to change his own valuation (although he might, for strategic reasons, change the amount of his bid).

The independent-private-values model and the common-value model should be interpreted as polar cases: real-world auction situations are likely to contain aspects of both simultaneously. For example, the bidders at an antiques auction may be dealers guessing about the ultimate market value of the item; but these dealers may differ in their selling abilities, so that the ultimate market value depends on which dealer wins the bidding. In the bidding for a government contract, there may be both inherent differences in the firms' production capabilities and a common element of technological uncertainty.

A general model which allows for correlations among the bidders' valuations and which includes as special cases both the common-value model and the independent-private-values model was developed by Paul Milgrom and Robert Weber (1982a). With n bidders, let x_i represent a private signal about the item's value observed by bidder i; let $x = (x_1,\ldots,x_n)$. Let $s = (s_1,\ldots,s_m)$ be a vector of variables which measure the quality of the item for sale.

The bidders cannot observe any of the components of s; however, some or all of the components of s may be observable by the seller. Now let the ith bidder's valuation of the item be $v_i(s,x)$. Thus any bidder's valuation may depend not only upon his own signal, but also upon what he cannot observe: namely, the other bidders' private signals and the true quality of the item. This formulation reduces to the independent-private-values model when $m = 0$ and $v_i = x_i$ for all i; and it reduces to the common-value model when $m = 1$ and $v_i = s_1$ for all i. The notion that bidders' valuations may to some extent be correlated is captured by the concept of affiliation: the vector of random variables (s,x) is affiliated if, roughly, some variables being large makes it likely that the other variables are large: if variables are affiliated, then they are positively correlated.[2]

When the independent-private-values model is applicable, there is a further choice to be made by the modeller. Are the bidders in some way recognizably different from each other? Is it appropriate to represent all bidders as drawing their valuations from the same probability distribution F, or should they be modelled as having different underlying distributions F_i, $i=1,\ldots,n$? The former case will be described for the sake of brevity as the case of symmetric bidders and the latter as the case of asymmetric bidders. An example of an asymmetric bidding situation arises in government procurement when both domestic and foreign firms submit bids and, for reasons of comparative advantage, there are systematic cost differences between domestic and foreign firms.

Yet another modelling consideration arising from uncertainty is that the amount of payment can only be made contingent upon variables that are observable to both buyer and seller. In some circumstances, the only such variables are the bids. In other circumstances, however, there are other mutually observable variables. If these other variables are correlated with the item's true value, it might be in the seller's interest to make payment depend on these other variables as well as the bids. (Clearly, however, the auction problem is trivial if this correlation is perfect.) For example, in mineral-rights auctions, royalties make the payment depend upon the amount of oil ultimately extracted as well as the winning bid.

The auction model which is the easiest to analyze is based on the following four assumptions.

(A1) The bidders are risk neutral.

(A2) The independent-private-values assumption applies.

(A3) The bidders are symmetric.

(A4) Payment is a function of bids alone.

This model will be referred to as the benchmark model; it will be discussed in Sections 5 and 6. However, many real-world auctions fail to satisfy these assumptions: the consequences of relaxing each of these assumptions, one at a time, will be discussed in Sections 7, 8, 9 and 10.[3]

The results to follow will describe bidding equilibria. Each bidder knows the rules of the auction that the seller has chosen and committed himself to. Bidder i knows his own valuation v_i (true valuation in the independent-private-values model; perceived valuation in the common-values model). Each bidder is assumed to know the number of bidders and the probability distributions of valuations, and to know everyone else knows that he knows this, and so on. (This is an instance of a game with incomplete information: John Harsanyi (1967, 1968).) Based on what he knows, each bidder decides how high to bid. At a Bayes-Nash equilibrium, each bidder bids an amount which is some function of his own valuation, such that, given that everyone else chooses his bid in this way, no individual bidder could do better by bidding differently.

One result can be obtained immediately: <u>the Dutch auction yields the same outcome as the first-price sealed-bid auction</u> (Vickrey (1961)). This is because the situation facing a bidder is exactly the same in each auction: the bidder must choose how high to bid without knowing the other bidders' decisions; if he wins, the price he pays equals his own bid. Because of this result, we do not need to discuss the Dutch auction in what follows.

5. <u>THE BENCHMARK MODEL: COMPARING AUCTIONS</u>

Which of the four simple auction types (English, Dutch, sealed-bid first-price, sealed-bid second-price) should a seller choose? In what we are referring to as the benchmark model (defined by assumptions (A1), (A2), (A3), and (A4)), this question has a surprising answer: it does not matter. Each of these auction forms yields on average the same revenue to the seller. At first glance, it may seem that this cannot be correct. For example, it might seem that receiving the value of the highest bid, as in the first-price sealed-bid auction, must be better for the seller than receiving the value of the second-highest bid, as in the second-price sealed-bid auction. The answer, of course, is that the bidders act differently in different auction situations; in particular, they bid higher in a second-price auction than in a first-price auction.

Consider first the English auction. When will the bidders stop bidding up the price in the English auction? The second-last bidder will drop out of the bidding as soon as the price exceeds his own valuation of the item.

Thus the highest-valuation individual wins the bidding and pays a price equal to the valuation of his last remaining rival. Usually this will be strictly below his own valuation of the item: the successful bidder earns some economic rent in spite of the monopoly power of the seller.

Only the bidder knows how much rent he receives because only he knows his own valuation. From the point of view of an outside observer or the seller, how large on average is the winner's rent? To answer this question, suppose that the n bidders' valuations are, in dollar terms, $v_1,\ldots,v_n$. Suppose v_1 is the highest valuation, v_2 is the second-highest, etc. (that is, v_1 is the first order statistic and v_2 is the second order statistic). From the previous argument, the winning bidder in the English auction earns a rent of $v_1 - v_2$. From the point of view of the winning bidder, the other bidders' valuations are independent draws from a probability distribution F (denote the density function by f). Thus the expected rent of the winning bidder is the expected difference between the first order statistic, v_1, and the second order statistic, v_2. The following result can be proven using the properties of order statistics:[4] <u>the expected difference between the first order statistic and the second order statistic is the expected value of</u> $[1-F(v_1)]/f(v_1)$.

The amount the seller is paid by the winning bidder is, by definition of economic rent, the buyer's valuation minus the buyer's rent. Thus, from the preceding argument, the expected payment to the seller in an English auction is the expected value of $J(v_1)$, defined by

$$(1) \qquad J(v_1) = v_1 - \frac{(1-F(v_1)}{f(v_1)} .$$

It will be assumed throughout that the distribution F is such that J is a strictly increasing function: this simply means that the winning bidder's expected payment is increasing in his own valuation.[5]

Consider now the second-price sealed-bid auction. In this, each bidder's equilibrium strategy is to submit a bid equal to his own valuation of the item. To see this, note that, since it is a second-price auction, the bidder's choice of bid determines only whether or not he wins; the amount he pays if he wins is beyond his control. Suppose the bidder considers lowering his bid below his valuation. The only case in which this changes the outcome is when this lowering of his bid results in his bid now being lower than someone else's and as a result this bidder now does not receive the item. Since he would have earned nonnegative rents if he won, lowering his bid below his valuation cannot make him better off. Conversely, suppose he considers raising his bid above his valuation. The only case in which this changes the outcome is when some other bidder has submitted a bid higher than the first bidder's valuation but lower than his new bid. Thus raising the bid causes this bidder to win, but he must pay more for the item than it is worth to him; raising his bid above his valuation cannot make him better off (Vickrey (1961)). This argument shows that, like the English auction, the second-price auction results in a payment equal to the actual valuation of the bidder with the second-highest valuation (that is, the realization of the second order statistic). Thus the expected payment is the expected value of $J(v_1)$.

The outcomes of the English and second-price auctions satisfy a strong criterion of equilibrium: they are dominant equilibria, that is, each bidder has a well-defined best bid regardless of how high he believes his rivals will bid. In a second-price auction, the dominant strategy is to bid true valuation; in an English auction, the dominant strategy is to remain in the bidding until the price reaches the bidder's own valuation. By contrast, as will be seen, a first-price sealed-bid auction does not have a dominant equilibrium. Instead, the equilibrium satisfies the weaker criterion of Nash equilibrium: each bidder chooses his best bid given his knowledge of the decision rules being followed by the other bidders.

A Nash equilibrium for a first-price sealed-bid auction is found as follows. Consider the decision of bidder i, whose valuation is v_i. He conjectures that, in equilibrium, the other bidders are following a decision rule given by a bidding function B: that is, he predicts that any other bidder j will bid an amount $B(v_j)$ if his valuation is v_j (although, of course, bidder i does not know this valuation). Assume that B is a monotonically increasing function. What is bidder i's best bid? If he bids an amount b_i and wins, he earns a surplus of $v_i - b_i$. The probability of winning with a bid b_i is the probability that all n-1 of the other bidders have valuations v_j such that $B(v_j) < b_i$; this probability is $[F(B^{-1}(b_i)]^{n-1}$, where, as before, F represents the distribution of valuations. Bidder i chooses his bid b_i to maximize his expected surplus $\pi = (v_i - b_i)[F(B^{-1}(b_i)]^{n-1}$. At a Nash equilibrium, the bidding function B must satisfy $b_i = B(v_i)$. From the Envelope Theorem,

$$(2) \quad \left.\frac{d\pi}{dv_i}\right|_{b_i = B(v_i)} = \frac{\partial\pi}{\partial v_i} = [F(v_i)]^{n-1}.$$

Integrating (2), using the fact that if a bidder has the lowest possible valuation v_ℓ then he earns zero surplus so that $B(v_\ell) = v_\ell$, yields

$$(3) \qquad B(v_i) = v_i - \frac{\pi(v_i)}{[F(v_i)]^{n-1}} = v_i - \frac{\int_{v_\ell}^{v_i} [F(x)]^{n-1} dx}{[F(v_i)]^{n-1}} .$$

The second term on the right-hand side of (3) shows how much the bidder shades his bid below his true valuation v_i.[6]

It can be shown that $B(v_1)$ as in (3) is equal to the expected value of the second order statistic conditional on the first order statistic's being v_1. That is, the winner of the first-price sealed-bid auction pays the expected value of the second order statistic conditional on his own information, namely that his own valuation is v_1. A risk-neutral bidder in a first-price auction estimates how far below his own valuation the next-highest valuation is on average, and then submits a sealed bid which is this amount below his own valuation. Thus, from the point of view of the seller, who does not know the winner's valuation v_1, the expected price is the expected value of $B(v_1)$, which in turn is equal to the expected value of $J(v_1)$ (defined by (1)). Hence, on average, the price reached in a first-price sealed-bid auction is the same as in an English or a second-price auction.

The foregoing argument establishes the Revenue-Equivalence Theorem: <u>For the benchmark model, each of the English auction, the Dutch auction, the first-price sealed-bid auction, and the second-price sealed-bid auction yield the same price on average</u> (Vickrey (1961); Charles Holt (1980); Milton Harris and Artur Raviv (1981a); Roger Myerson (1981); John Riley and William Samuelson (1981)).

The Revenue-Equivalence Theorem does not imply the outcomes of the four auction forms are always exactly the same. In an English or second-price auction, the price exactly equals the valuation of the bidder with the second highest valuation (that is, v_2). In a first-price sealed-bid or Dutch auction, the price is expectation of the second-highest valuation conditional on the winning bidder's own valuation (that is, $B(v_1)$). Only by accident, for particular highest and second-highest valuations v_1 and v_2, will these two prices be equal. They are, however, equal on average.

Although all four simple auctions yield the same price on average, there is an important practical difference between, on the one hand, the English and the second-price auctions and, on the other hand, the first-price and Dutch auctions. In the former case, any bidder can easily decide how high to bid; in an English auction, he remains in the bidding until the price reaches his valuation; while in the second-price auction, he submits a sealed bid equal to his valuation. In the case of a Dutch or a first-price auction, the bidder bids some amount less than his true valuation: exactly how much less depends upon the probability distribution of the other bidders' valuations and the number of competing bidders (as in equation (3)). Finding the Nash-equilibrium bid in the first-price or Dutch auction is a nontrivial computational problem.

The Revenue-Equivalence Theorem is devoid of empirical predictions about which type of auction will be chosen by the seller in any particular set of circumstances. However, as will be seen, when assumptions that underlie the benchmark model are relaxed, particular auction forms emerge as being superior.

Despite the monopoly-oligopsony nature of the problem, the outcome of the auctions is Pareto efficient: the bidder with the highest valuation receives the item (provided J is strictly monotonic). However, as will be seen, if the monopolist is given more instruments, he will in general distort the outcome away from efficiency.

What is the effect of increasing the amount of competition among the bidders? The more bidders there are, the higher on average is the valuation of the second-highest-valuation bidder. Hence: <u>Increasing the number of bidders increases the revenue on average of the seller</u> (Holt (1979); Harris and Raviv (1981)). This is a testable proposition: the effect of the number of bidders on price has been found to be statistically significant by Louis Ederington (1978) in a study of bidding for new bond issues; by Kenneth Gaver and Jerold Zimmermann (1977), in a study of bidding for construction contracts for San Francisco's BART subway system; by Otis Gilley and Gordon Karels (1981), in a study of bidding for mineral rights; and by Lance Brannman, Douglass Klein, and Leonard Weiss (1984), using data from auctions of tax-exempt bonds, government-owned timber, and mineral rights. Large gains can be had from introducing bidding competition where none formerly existed: Larry Yuspeh (1976) found price differences averaging 50 percent between identical military contracts let successively on a sole-source basis and under competitive bidding.[7]

Provided the number of bidders is finite, each bidder bids an amount strictly less than his own valuation of the item, so that the winning bidder earns a strictly positive economic rent. But if there is perfect competition among the bidders, then all of the gains from trade go to the seller. <u>As the</u>

number of bidders approaches infinity, the price tends to the highest possible valuation (Holt (1979)). To understand this, note that, with a continuum of bidders, there is with probability one a bidder having the highest possible valuation. Moreover, this bidder is constrained to bid no less than his valuation by his knowledge that there is another bidder with valuation very close to his who will beat him if he shades his bid.

In addition to the number of bidders, another determinant of the strength of the bidding competition is the variance of the distribution of valuations. The larger is this variance, the larger on average is the difference between the highest valuation and the second-highest valuation, and so the larger is the economic rent accruing to the winning bidder. However, an increase in the variance of a distribution, holding the mean constant, usually increases the second-highest valuation as well. Hence: an increase in the variance of valuations increases both the average revenue of the seller and the rents of the successful bidder (McAfee and McMillan (1984)).

The essence of the auction problem is the unobservability of bidders' valuations. Suppose the seller wishes to learn the bidders' valuations. Can he design the auction so as to induce the bidders directly to reveal their preferences? The following result was already obtained above: In the second-price sealed-bid auction, each bidder bids his true valuation (Vickrey (1961)). This is closely related to the fundamental insight underlying the Groves mechanism for inducing revelation of preferences for public goods (Theodore Groves and John Ledyard (1977)). Note that, in the benchmark model, the seller can obtain this information for free, because on average the revenue he receives from the second-price auction is just as high as the revenue from any of the other auctions.

6. OPTIMAL AUCTIONS

The Revenue-Equivalence Theorem compares the expected revenues accruing from each of the commonly used auction forms. Given that the monopolist has the power to design his selling mechanism, a more fundamental question to ask is: What is the best of all possible auction forms from the point of view of the seller?

The tool used to address this question is the Revelation Principle. Use the word mechanism to describe any process which takes as inputs the bidders' bids and produces as its output the decision as to which bidder receives the item and how much any of the bidders will be required to pay. Each of the auction forms so far described is an example of a mechanism. In a direct mechanism, each bidder is asked simply to report his valuation of the item. A mechanism is incentive compatible if the mechanism is structured such that each bidder finds it in his interest to report his valuation honestly. The Revelation Principle asserts the following: For any mechanism, there is a direct, incentive-compatible mechanism with the same outcome. Thus, in particular, the optimal mechanism can be mimicked by some direct, incentive-compatible mechanism. (For useful expositions of the Revelation Principle, see Harris and Robert Townsend (1985) and Myerson (1983)).

To exemplify the Revelation Principle, consider the direct, incentive-compatible mechanism that is equivalent to the first-price sealed-bid auction. In the first-price sealed-bid auction, the bidder with the highest valuation v wins and pays the amount of his bid, which was shown in the last section to be, in equilibrium, $B(v)$. Consider now a direct mechanism, in which the seller simply asks the bidders to report to him their valuations. Usually it will be in the bidders' interests to lie. Suppose, however, the seller announces that the mechanism is the following: the bidder who reports the highest valuation, $\hat{v}$, will win the item and be asked to pay $B(\hat{v})$. This particular direct mechanism is equivalent to the first-price sealed-bid auction. Since the bidders in the first-price sealed-bid auction are optimizing when they submit bids of $B(v)$, it must be optimal for them to report their valuations honestly (that is, $\hat{v} = v$) in this direct mechanism: the mechanism is incentive compatible. Also, as shown earlier, the second-price auction is incentive compatible.

The Revelation Principle achieves honest revelation in the direct mechanism by designing the payoff structure in such a way that it is in the bidders' interests to be honest. In effect, the computations that go on within the mind of any bidder in the nondirect mechanism are shifted to become part of the mechanism in the direct mechanism. Instead of having the bidder compute his own bid in the first-price sealed-bid mechanism, all of the computations are done inside the mechanism in the direct mechanism.

The significance of the Revelation Principle is that it shows that the modeller can limit his search for the optimal mechanism to the class of direct, incentive-compatible mechanisms. The number of possible selling procedures is huge; hence it is useful to be able to restrict attention to one relatively simple class of mechanisms. The Revelation Principle is purely a theoretical technique: few, if any, resource-allocation procedures in practical use are direct, incentive-compatible mechanisms. But using the Revelation Principle does facilitate solving for that resource-allocation mechanism which is optimal subject to the constraints imposed by the asymmetry of information. The optimal direct mechanism is found as the solution to a mathematical programming problem involving two kinds of constraints:

first, incentive-compatibility constraints, which state that the bidders cannot gain by misrepresenting their valuations; and second, individual-rationality or free-exit constraints, which state that the bidders would not be better off if they refused to participate.

Returning to the auction model stated in the last section, suppose that the seller himself attaches a value of v_0 to the item being offered for sale. (Since v_0 may be zero, we are not requiring that the seller necessarily has some use for the item.) Applying the Revelation Principle can be shown to yield the following result: <u>For the benchmark model, the auction which maximizes the expected price has the following characteristics: (a) If</u> $J(v_i) < v_0$ <u>for all bidders' valuations</u> v_i, <u>then the seller refuses to sell the item; (b) otherwise he offers it to the bidder whose valuation</u> v <u>is highest at a price equal to</u> $J(v)$ (Jean-Jacques Laffont and Maskin (1980); Harris and Raviv (1981); Milgrom (1985b); Myerson (1981); Riley and Samuelson (1981)).

The first part of this result says that the seller optimally sets a reserve price, not selling the item if all bidders' valuations are too low. Notice that this policy introduces the possibility of an inefficient outcome. Since $J(v) < v$, it is possible that the seller keeps the item despite the presence of some bidder with a valuation which is greater than the seller's own valuation. Like the elementary-textbook monopolist, the seller finds it in his interest to distort the outcome away from Pareto optimality.

Because the seller will expect to earn $J(v)$ if he sells to an individual with value f, or v_0 if he does not sell, the set of values (including his own) facing the seller is $\{v_0, J(v_1),\ldots,J(v_n)\}$. The first order statistic of earnings is, therefore, the maximum of $J(v)$ over v and v_0. Consequently, the seller is indifferent to keeping the good when $v_0 = J(v_1)$ (where v_1 is the highest valuation), defining the reserve price. A useful way of preserving the second-order-statistic intuition provided earlier is that seller earns the second order statistic of $\{J^{-1}(v_0), v_1, \ldots, v_n\}$, where $J^{-1}(v_0)$ is the value (if held by a bidder) that produces a payment to the seller of v_0.

This description of the optimal auction is in terms of a direct, incentive-compatible auction, with the seller asking the bidders how much they value the item and the bidders responding by honestly reporting their valuations to the seller. To convert back to a more familiar-looking mechanism, consider an English auction. As was shown in the last section, the seller expects to earn $J(v)$ from a winner with value v. Thus, if the reserve price is not binding, part (b) of the optimal-auction result shows that the optimal auction is equivalent to the English auction. The English-auction equivalent of the reserve price (part (a) of the result) is that the seller sets a reserve price r which strictly exceeds his own valuation v_0 (namely, $r = J^{-1}(v_0) > v_0$).

To understand why the reserve price increases the average selling price, suppose there is at least one bidder whose valuation v exceeds the seller's valuation. The reserve price is binding in the English auction only if the second-last bidder drops out before the reserve price is reached. The optimal level of the reserve price is determined by a tradeoff. The disadvantage of setting a reserve price, already noted, is that it is possible that the remaining bidder has a valuation which lies between the seller's valuation and the reserve price (that is, $v_0 < v < r = J^{-1}(v_0)$). In this case, the monopolist loses the sale despite the fact that the bidder would have been willing to pay more than the good is worth to the monopolist. On the other hand, the advantage of the reserve price is that it is possible that the bidder's valuation exceeds the reserve price, so that he pays at least the reserve price. If the reserve price is above the second-highest bidder's valuation, this means that the bidder pays more than he would have in the absence of the reserve price.

The case of a single bidder provides a simple example of this result. If the seller sets a reserve price r, the buyer will never pay more than r, as he faces no competition. The buyer will pay r if his value of the good exceeds r, which occurs with probability 1-F(r). Thus, the seller expects to earn:

(4) $\pi(r) = r(1-F(r)) + v_0F(r)$.

Maximizing π with respect to r yields $v_0 = J(r)$, and the second order condition that J is nondecreasing.

Note the simplicity of the formula for the optimal reserve price: in particular, it is independent of the number of bidders. For the case of a uniform distribution of valuations, the optimal reserve price is especially easy to compute: it is the average of the seller's own valuation and the highest possible valuation that a bidder could have.

Since, as was shown in the last section, the four common auction forms are essentially equivalent in the benchmark model, we can conclude as follows: <u>For the benchmark model, any of the English, Dutch, first-price sealed-bid, and second-price sealed-bid auctions is the optimal selling mechanism provided it is supplemented by the optimally-set reserve price. The optimal level of the reserve price for any of these auctions is</u> $J^{-1}(v_0)$ (Harris and Raviv (1981a), Myerson (1981), Riley and Samuelson (1981)).

This is a powerful result. No restriction has been placed on the types of policies the seller could use. The seller could, for example, have several rounds of bidding, or charge bidders entry fees, or subsidize bidders, or require losing bidders to pay an amount related to their bids, or allow only a limited time for the submission of bids. But none of these more complicated strategies would increase the expected price: the simple auction forms are the best out of the huge set of possible selling mechanisms.

Actual practice does not seem to be in accord with the theoretical result that it is in the seller's interest to announce a reserve price. In practice, reserve prices are often not used; when they are used, their existence is often not announced; and even when their existence is announced, the seller usually keeps the level of the reserve price secret (Cassady (1967, pp. 226-227)). Thus there appears to be a discrepancy between theory and practice. It is likely, however, that factors so far omitted from the analysis, such as risk aversion, correlations among the bidders' valuations, or collusion among the bidders, are significant in these observations. Consistently with the reserve price rationally being set at a level which exceeds the seller's own valuation, in the Netherlands cut flowers not attaining the reserve price are destroyed, indicating that their value to the seller is zero despite being given a positive reserve price (Cassady (1967, p. 230)).

This completes the analysis of the benchmark model. The rest of this survey examines the effects of changing the assumptions upon which the benchmark is based.

7. ASYMMETRIC BIDDERS: PRICE DISCRIMINATION

Instead of assuming that all bidders appear the same to the seller and to each other (assumption (A3)), suppose that the bidders fall into one of two recognizably different classes. Thus, instead of there being a single distribution F from which the bidders draw their valuations, there are two distributions, F_1 and F_2; bidders of type i draw their valuations independently from the distribution F_i (with density function f_i). For example, bidders at an antiques auction might be classifiable as either dealers or collectors, with the average demand price among dealers differing from that among collectors, or bidders for a government contract might be divided into and foreign firms, with systematic production-cost differences.[8]

The English auction in this asymmetric case operates much as in the benchmark model: the bids rise until the price reaches the second-highest valuation. In particular, the highest-valuation bidder wins, so that the outcome is efficient.

With the first-price sealed-bid auction, in contrast, the outcome is usually not efficient when bidders are asymmetric. This is because bidders from different classes perceive themselves to be facing different degrees of bidding competition. While it remains the case that, within a class, higher-valuation individuals bid higher, this is in general not the case across classes. Thus it is possible that, say, a type 1 bidder submits a higher bid than a type 2 bidder despite the type 2 bidder's having the higher valuation.

Hence the first-price sealed-bid auction in general yields a different price than the English auction when bidders are asymmetric. Examples have been constructed (by Vickrey (1961), Griesmer, Levitan and Shubik (1967), and Maskin and Riley (1983a, 1985)) which show that the English auction's price can be either higher or lower than the first-price sealed-bid auction's price.

Neither of these auctions is optimal. Analogous to the function J of Section 5, define functions $J_i(v_i)$, where v_i represents the valuation of a bidder of type i, by

$$(5) \quad J_i(v_i) = v_i - \frac{(1-F_i(v_i))}{f_i(v_i)}, \quad i = 1,2.$$

The following theorem is due to Myerson (1981): <u>In the auction which maximizes the expected selling price, the seller awards the item to the individual with the highest value of</u> $J_i(v_i)$. (In addition, the seller sets a different reserve price for each type of bidder, computed in exactly the same way as in Section 6.)

Since it is assumed in this section that the bidders differ from each other so that $F_1 \neq F_2$, the J_i's are different functions; thus Myerson's theorem shows that the optimal auction is discriminatory, in the sense that there will be a possibility that one bidder wins despite another bidder's having a higher valuation. To understand the nature of this discrimination, define a function $z(v_1)$ to compare a type 1 bidder with a type 2 bidder: a bidder of type 1 with valuation v_1 wins against a bidder of type 2 with valuation v_2 if and only if $z(v_1) > v_2$. Thus, for example, $z(v) > v$ for some v means that bidders of type 2 are discriminated against in favor of bidders of type 1, in the sense that it is possible for a type 1 bidder to beat a type 2 bidder despite having a lower valuation. It follows from Myerson's theorem that the optimal z function is implicitly defined by

$$(6) \quad J_1(v_1) = J_2(z(v_1)).$$

The optimal reserve-price policy described in the previous section can now be seen to be a special instance of this optimal discriminatory policy, with the seller discriminating between himself, as an implicit bidder, and the actual bidders.

Since the seller's optimal policy leaves a positive probability of the item being awarded to someone other than the bidder who values it the most, the policy is not Pareto-efficient. For this policy to be workable, it must be the case, as for the price-discriminating monopolist of elementary economic theory, either that the seller can prevent the successful bidder from reselling the item to some other bidder, or that the item being sold is inherently nontransferable. Arbitrage among the bidders, if it were possible, would sabotage any discriminatory selling scheme.

The probability that a bidder of type i has a valuation of at least v_i is $(1-F_i(v_i))$. Define η_i to be the elasticity of this probability with respect to v_i: that is, $\eta_i(v_i) = v_i f_i(v_i)/(1-F_i(v_i))$. Then it can be shown that: <u>The optimal discriminatory policy satisfies</u>

$$(7) \qquad \frac{z(v_1)}{v_1} = \frac{1 - 1/\eta_1}{1 - 1/\eta_2}$$

(McAfee and McMillan (1985b)). This superficially resembles the standard formula for optimal price discrimination in the elementary-textbook monopoly model, showing the analogy between the optimally discriminatory auction and more familiar notions of price discrimination.

Which type of bidders receive preferential treatment? The answer depends upon the relative shapes of the valuation distribution functions F_1 and F_2. However, one special case is useful in aiding understanding. <u>If</u>

<u>the distributions of valuations are identical except for their means, then the class of bidders with the lower average valuation are favored in the optimal auction</u> (McAfee and McMillan (1985b)). There is a tradeoff. By favoring the low-valuation type of bidders, the seller raises the probability of awarding the item to someone other than the bidder who values it the most. The benefit from this policy, however, is that the favoritism forces the bidders from the high-valuation class to bid higher than they otherwise would.

An important application of these results is to government procurement. Governments often favor local suppliers over foreign suppliers. For example, under buy-American legislation, the United States federal government offers a 6 percent price preference for domestic content: if a local firm's bid is no more than 6 percent higher than the lowest foreign bid, the local bid will be accepted. The results just cited show that there are circumstances in which a policy of this type can be optimal: if the foreign firms have on average lower production costs because they have a comparative advantage, then the government minimizes its expected payment by favoring the local firms. However, these considerations do not explain the existing policies: in an industry in which the local firms have a comparative advantage, minimizing the government's expected payment requires that the foreign firms be favored, which seems unlikely to occur. Undoubtedly the existing government-procurement preferences were introduced for political reasons and not to increase the amount of bidding competition. However, this analysis shows that it is not appropriate to evaluate such policies using as a benchmark the absence of preferences: an ostensibly nondiscriminatory sealed-bid auction results in ad hoc discrimination when the bidders are asymmetric.

Another instance of a price-discriminating auction occurs when a buyer has a sequence of projects: for example, a government offers a research-and-development contract followed by a production contract. The winner of the first auction reveals, by his winning, that he has a cost advantage. Thus the buyer should discriminate against the incumbent in the second auction (Richard Luton and McAfee (1985)).

8. ROYALTIES AND INCENTIVE PAYMENTS

In the last section we examined an asymmetric-information equivalent of the elementary textbook's concept of price discrimination by market segmentation. Another form of price discrimination discussed in elementary textbooks is multi-part pricing: in this section we examine an asymmetric-information analogue of multi-part pricing.

It has been assumed so far that the seller is able to make payment depend upon only the bids. The bids give the seller some information about how highly the bidders value the item for sale. In many circumstances, however, the seller has, or can obtain, additional information about valuations. In this section, we relax the assumption (A4) that payment can be a function only of bids, and show that it is in the seller's interest to condition the bidders' payments on any additional available information about the winner's valuation. (Of course, if the seller has perfect information on the bidders' valuations, the auction problem is trivial).

For example, in an auction of oil rights to government-owned land, the government can observe, *ex post*, how much oil is actually extracted: this provides additional information on the true value of the tract. The payment by the successful bidder equals the amount he bids plus a royalty based on the amount of oil extracted (Mead, et al. (1984)). Publishing rights for books are sometimes auctioned, with payment to the author depending both on the bid and, via a royalty, on the book's ultimate sales (John Dessauer (1981)). For weapons procurement, the U.S. Department of Defense increasingly often uses incentive contracts, which make payment to the contractor depend not only on his bid but also on the production costs he actually incurs (Peter deMayo (1983)). Incentive contracts are also used in the private sector when a firm procures inputs from another firm (Seiichi Kawasaki and McMillan (1985)).

All of these examples have the following properties. The seller observes *ex post* some variable $\tilde{v}$ which is an estimate of the winning bidder's true valuation v. The payment p to the seller by the winning bidder is a linear function:

$$(8) \quad p = b + r\tilde{v},$$

where b is the bid and r is the royalty rate. (In the case of contract bidding, the payment to the successful bidder is $p = b + \alpha(c-b)$, where c is realized production cost and α, the sharing parameter, is the fraction of any cost overrun or underrun $(c - b)$ that the winning bidder is responsible for. In the extreme case of $\alpha = 1$, the contract is cost-plus; with $\alpha = 0$, the contract is fixed-price.)

Three bidding mechanisms can be used with payment functions of the form (8). First, the seller can set the royalty rate and call for bids b. Second, the seller can set the fixed payment b and call for bids on the royalty rate r. Third, the seller can call for bids on both the fixed payment b and the royalty rate r simultaneously. Both the first mechanism and the second mechanism are used by the U.S. Government in auctioning offshore oil tracts, with the first being the more commonly used (Mead _et al_. (1984)). In what follows, we shall discuss bidding mechanisms of the first type. (See Robert Hansen (1985a), Douglas Reece (1979), and Riley (1985) for analyses of bidding mechanisms of the second type, and Samuelson (1983, 1984) for an analysis of bidding mechanisms of the third type.)

What is the reason for using royalties? _The seller's expected revenue is an increasing function of the royalty rate_ (McAfee and McMillan (1984), Riley (1985)). The intuition behind this is that an increase in the royalty rate lessens the significance for the bidding of inherent differences in the bidders' valuation. As was noted earlier (in Section 5), a decrease in the variance of the bidders' valuations generates more aggressive bidding and therefore a higher expected revenue for the seller. An increase in the royalty rate has a similar effect on the bidding to a decrease in the variance of valuations. The royalty serves to transfer rents from the successful bidder to the seller.[9]

If expected revenue monotonically increases with the royalty rate, why are royalties not always set at 100 percent? The answer is that the winning bidder, by his actions after the auction, often is able to affect the signal about his true valuation that the seller receives: there is a moral-hazard problem because the organizer of the auction cannot control what the winning bidder does afterwards. This is the case in each of the three examples above. The amount of oil extracted from a tract is decided by the extractor. Eventually, diminishing returns set in, and the higher the royalty rate, the less oil will be extracted; that is, the lower $\tilde{v}$ will be. The sales of a book vary with the amount of publicity the publisher chooses to give it.

The production costs incurred by a contractor in part depend on how much effort he makes to hold costs down. Such moral-hazard considerations must be weighed against the effects on bidding competition in the choice of what royalty rate to set.

When there is moral hazard, the optimal royalty is determined by tradeoff. Increasing the royalty rate serves to increase the bidding competition and raise the bids, as already argued. But an increase in the royalty rate reduces the return to the winning bidder on his own actions after the auction: the royalty has the effect of transferring part of the benefit of these actions to the seller. Thus the higher the royalty rate, the less the *ex post* effort made by the winning bidder; this tends to lower the seller's expected revenue. *With moral hazard, the optimal royalty r is less than 100 percent*. Thus moral hazard results in the seller not making payment fully dependent on his ex post information. *The royalty r is zero if and only if there are infinitely many bidders* (McAfee and McMillan (1984, 1985d)). This is because, when there are enough bidders that perfect competition prevails, there is no need to use royalties to stimulate bidding competition. Thus the contract is used only to address moral hazard; and moral hazard is most effectively addressed when the successful bidder keeps all of any marginal increases in the item's value; that is, when $r = 0$ in (8).

While the simplicity of the linear payment function (8) means that it is commonly used in practice, in general there is a nonlinear payment function which would yield a higher expected revenue for the seller. *With moral hazard, the optimal contract is linear in observed valuation* $\tilde{v}$ *but nonlinear in the winning firm's bid* b (McAfee and McMillan (1985d)).

Making payments conditional on ex post observations of valuations serves not only to stimulate bidding competition: it also shifts risk from the bidders to the seller. If the bidders are risk averse while the seller is risk neutral, then some amount of risk shifting is mutually beneficial. <u>The more risk averse are the bidders relative to the seller, the higher is the optimal royalty rate</u> (Hayne Leland (1978), Samuelson (1983, 1984)).

9. <u>RISK-AVERSE BIDDERS</u>

Auctions generally confront bidders with risk. Typically, a bidder obtains nothing and pays nothing if he loses, and earns positive rents if he wins. Thus if the bidders are risk averse, the extent of their aversion to risk will influence their bidding behavior. In this section we relax assumption (A1) and suppose the bidders have von Neumann-Morgenstern utility functions, while maintaining the other assumptions of the benchmark model. We continue to assume that the seller is risk neutral and therefore wishes to maximize his expected earnings.

The seller can do at least as well as in the risk-neutral-bidders case, for if he sells the good using an English auction it remains the case that buyers will remain in the bidding so long as the price is less than their value. Thus, the seller can expect to earn at least as much when the buyers are risk averse as when they are not. Indeed, the seller can do strictly better, for: <u>with risk-averse bidders, the first-price sealed-bid auction produces a larger expected revenue than the English or second-price auction</u> (Harris and Raviv (1981a), Holt (1980), Maskin and Riley (1980), Riley and Samuelson (1981)). The intuition behind this result is seen by examining the problem facing an agent in the first-price sealed-bid auction. If he loses, he gets nothing, while if he wins, he obtains a positive profit. Thus he is facing risk. By marginally increasing his bid, he lowers his profit if he wins, but increases the probability of this event. By smoothing his utility, he increases his expected utility (up to a point); but this also increases his payment to the seller. Thus the seller prefers the buyers to be risk averse.[10]

The first-price sealed-bid auction is not the optimal auction, however; it does not maximize the expected revenue of the seller when the bidders are risk averse. Because the seller is risk neutral, there may be gains from trade in risk. The seller is not fully exploiting his comparative advantage in risk-bearing when he uses a first-price sealed-bid auction. For example, if the seller makes low bids risky, he might encourage higher bids.

There are two ways in which the seller can impose risk on the bidders: first, the risk of losing; and second, random payments. It can be shown, however, that it is not in the seller's interest to use the second of these instruments. <u>The seller will not make losing bidders' payments random</u> (Maskin and Riley (1984b), Steven Matthews (1983)). The reason is straightforward. Instead of requiring a risky payment from a loser, the winner could require that payment's certainty equivalent, leaving the probability of winning and the payment upon winning unchanged. Then the payoff to the bidder is unchanged in utility terms. However, the certainty-equivalent payment exceeds the expected risky payment by the risk premium, a positive number. Thus the seller gains the risk premium; he is better off not making the losers' payments random. Should the seller make the winner's payment random? The foregoing argument does not in general apply to the payment made by the winner; the argument does, however, carry over when a bidder's risk aversion does not vary with his income. <u>If the bidders have constant absolute</u>

risk aversion, the payment required of the winning bidder is not random (Matthews (1983)).

The optimal auction is very complicated, in marked contrast to the simplicity of the optimal auction in the case of risk-netural bidders. However, some broad features can be described. The optimal auction with risk-averse bidders involves subsidizing high bidders who lose and penalizing low bidders (Maskin and Riley (1984b), Matthews (1983, 1984), John Moore (1984)). This is done by making the bidders' certainty-equivalent payment positive for bidders with low valuations and negative for bidders with high valuations. Thus the seller absorbs some of the risk faced by high bidders. Except in one case however, he does not absorb all of the risk: the seller provides full insurance only for the highest possible bid (Maskin and Riley (1984b)). Thus bidders prefer to win than to lose, despite the subsidies to some of the losers. Finally, in the case of constant absolute risk aversion, payment by the winner and the probability of winning are nondecreasing functions of the winner's valuation of the item (Matthews (1983)). The fact that the bidders' certainty-equivalent payment decreases as valuation rises provides high bidders with more insurance than low bidders. Thus the seller rewards high bids by means of insurance, which is costless for him to provide because of his risk neutrality. Because the seller does not offer full insurance, the bidders prefer to winning to losing; moreover, an increase in bid does not decrease the bidder's probability of winning. The seller compensates for his payments to high, losing bidders by requiring a large payment from the winning bidder.

Because the optimal auction with risk averse bidders is so complicated, it is unlikely to arise in practice. However, if the risk aversion is not very strong, the optimal auction is closely approximated by a sealed bid auction with a bidding fee that is a decreasing function of the bid. Bidding fees are not uncommon in contract bidding, although they do not depend on the bid. However, insofar as bidders with high but losing bids can be rewarded on other contracts, perhaps with favorable treatment, it is possible that the optimal auction could be approximated in practice.

Consider now another instrument which becomes useful to the seller when the bidders are risk averse. A standard assumption in auction theory is that each bidder knows exactly how many other bidders he is competing with. When bidders are risk averse, this is a nontrivial assumption, for the bidders behave differently when they have this knowledge than when they do not. This is a consequence of the fact that the seller rationally has a different expectation about the number of bidders than does any of the bidders, because the bidder conditions his probabilities on his knowledge that he himself is one of the bidders: it follows that, with indentical priors, any bidder always expects more bids than the seller. Under constant or decreasing absolute risk aversion, in a first-price sealed-bid auction the expected selling price is strictly higher when the bidders do not know how many other bidders there are than when they do know this (Matthews (1985), McAfee and McMillan (1985a)). It follows from this result that, if the seller can somehow organize the auction in such a way as to leave each bidder ignorant about the number of bidders, then he should do so. In some cases of government-contract bidding, the government agency has a policy of concealing information about how many firms it has invited to submit bids: the foregoing result provides some justification for such a policy.

10. CORRELATED VALUES

In many auctions, the uncertainty about each bidder's valuation of the item being sold does not result from inherent differences in the bidders' tastes, as has so far been assumed. Instead, it arises because each bidder, having access to different information, has a different estimate of the value of the item. In this section, we relax assumption (A2), the independent-private-values assumption, and allow interactions among the different bidders' valuations.

Consider first the extreme case of the common-value auction, in which the bidders guess about the unique true value of the item. When the item being bid for has a common value, the phenomenon dramatically named the "winner's curse" can arise. Each bidder in sealed-bid auction makes his own estimate of the true value of the item. The bidder who wins is the bidder who makes the highest estimate. Thus there is a sense in which winning conveys bad news to the winner, because it means that everyone else estimated the item's value to be less. The winner's curse has been noted in the book-publishing industry. One observer, commenting on the high prices fetched in the auctioning of manuscripts among publishers, said: "The problem is, simply, that most of the auctioned books are not earning their advances. In fact, very often such books have turned out to be dismal failures whose value was more perceived than real" (Dessauer (1981)). The winner's curse has also been claimed to exist in auctions of offshore oil rights (E.C. Capen, R.V. Clapp, and W.M. Campbell (1971)), in the market for baseball players (James Cassing and Richard Douglas (1980)), and in the bidding for contracts which have a common element of technological uncertainty (James Quirk and Katsuaki Terasawa (1984)).

Statements about the winner's curse such as that quoted come close to asserting that bidders are repeatedly surprised by the outcomes of auctions, which would violate basic notions of rationality. (James Cox and Mark Isaacs (1984) pursued this straw man.)[11] A more reasonable interpretation of the winner's curse is that sophisticated bidders, when deciding their bidding strategies, take into account the fact that winning reveals to the winner that his estimate of the item's value was the highest estimate; as a result, they bid more cautiously than if they adopted naive strategies. The basis for such sophisticated bidding strategies is the following result in probability theory. <u>Suppose the ith bidder's information about the item's true value</u> v <u>can be represented by a number</u> x_i, <u>such that a bigger value of</u> x_i <u>implies a bigger true value v</u>. <u>Then</u>

(9) $$E[v|x_i] \geq E[v|x_i, x_i > x_j \text{ for all } j \neq i]$$

(Milgrom (1979b, pp. 60-63; 1981a)). The left side of this inequality shows the bidder's expectation about the item's value before the bidding; the right side shows his expectation after he knows that he has won. Thus the mere knowledge that he has won will cause a naive bidder to revise downwards his estimate of the item's true worth.

The rational bidder in a common-value sealed-bid auction avoids becoming a victim of the winner's curse by presuming that his own estimate of the item's value is higher than any other bidder's; that is, by presuming that he is going to be the winner. He then sets his bid equal to what he estimates to be the second-highest perceived valuation given that all the other bidders are making the same presumption. There is no cost to making this presumption when it is wrong, because losing bidders pay nothing (James Smith (1981)).

It is often pointed out (for example, by Hayek in the paper cited earlier) that one of the remarkable and important features of the price system is its ability to convey information efficiently. All that a buyer or a seller needs to know about a commodity's supply or demand is summarized by a single number, its price. Does the process of price formation by competitive bidding have such information efficiencies? In the common-value model, the bidders lack complete information about the item's true value; each bidder has different partial information. However, despite the fact that no single bidder has perfect information, it can be shown that, if there is perfect competition in the bidding, the selling price reflects all of the bidders' private information. <u>If information is sufficiently dispersed among the bidders then the selling price converges to the item's true value</u> v <u>as the number of</u> <u>bidders becomes arbitrarily large</u> (Milgrom (1979a, 1979b), Wilson (1977)). Thus the selling price conveys information about the item's true value. With perfect competition, the price is equal to the true value even though no individual in the economy knows what this true value is and no communication among the bidders takes place.[12]

There is a stylized fact from the sealed bidding for mineral rights: in common-value auctions, the distribution of bids tends to be approximately lognormal (Chester Pelto (1971), Reece (1978)).

Consider now the more general model, due to Milgrom and Weber (1982a), which allows correlations among the bidders' valuations, and of which the common-value model is a special case. Recall from Section 4 that bidders' valuations are said to be affiliated if the fact that one bidder perceives the item's value to be high makes it likely that other bidders also perceive the value to be high. The essential difference between, on the one hand, the English auction and, on the other hand, the first-price sealed-bid, second-price, and Dutch auctions is that the process of bidding in the

English auction conveys information to the bidders: the remaining bidders observe the prices at which the other bidders drop out of the bidding. It was shown for the independent-private-values auction that this extra information does not on average change the outcome, in the sense that the expected price reached is the same for each type of auction. When bidders' valuations are affiliated, in contrast, the bids in the English auction have the effect of partially making public each bidders' private information about the item's true value, thus lessening the effect of the winner's curse. As a result: <u>When bidders' valuations are affiliated, the English auction yields a higher expected revenue than the first-price sealed-bid auction, the second-price sealed-bid auction, or the Dutch auction</u> (Milgrom and Weber (1982a)). In addition, the other three auction forms can be ranked. <u>With affiliated valuations, the second-price sealed-bid auction yields a higher expected revenue than the first-price sealed-bid auction, which yields the same expected revenue as the Dutch auction</u> (Milgrom and Weber (1982a)).

Because of the Revenue-Equivalence Theorem, the benchmark model made no predictions about which auction form will be used. Now we have contradictory predictions. If the bidders' valuations are affiliated and the bidders are risk neutral, then the seller will choose the English auction ahead of the other simple auction forms. But if the bidders' valuations are independent and the bidders are risk averse then, as we saw in the previous section, the first-price sealed-bid auction is the best of the simple auctions from the seller's point of view. Risk aversion is likely to be important when the item being sold is very valuable so that the bids are large relative to any bidder's assets. Examples of valuable items which are auctioned include mineral rights, government contracts, and artwork. However, artwork, mineral rights and, in some cases, government contracts also have common-value aspects. Government-procurement auctions and mineral-rights auctions are usually sealed-bid. This may be a consequence of the result that the risk aversion makes the sealed-bid auction preferable to the English auction. Artwork, on the other hand, is usually, but not always, sold by English auction, despite the fact that risk aversion is likely to be significant in at least some cases. Presumably in the case of artwork the correlations among bidders' valuations outweigh the risk-aversion effects in the bidding. Low-value items like agricultural produce, for which risk-aversion effects are likely to be negligible, are usually sold by English auction. The U.S. Forest Service has used both first-price sealed-bid and English auctions to sell contracts for harvesting timber: it has been estimated that the sealed-bid auctions generated higher revenue than the English auctions (Mead (1967), Hansen (1985b)).

Sometimes the seller has independent information correlated with the item's value to any of the bidders. (For example, the government can do its own geological surveys before offering mineral rights for sale; the seller of a painting can obtain an expert's appraisal.) Should the seller conceal this information, or should he reveal it? <u>The seller can increase his expected revenue by publicizing any information he has about the item's true value</u> (Milgrom and Weber (1982a)). This is because the new information tends to increase the value estimates of those bidders who perceive the item's true value to be relatively low, causing them to bid more aggressively.

How important is the privacy of any bidder's information? <u>If one bidder's information is available to another bidder, his expected surplus is zero</u> (Engelbrecht-Wiggans, Milgrom, and Weber (1982), Milgrom (1981), Milgrom and Weber (1982a)). This is a striking result: it implies that it is more important to a bidder that his information be private than that it be accurate.

What is the optimal auction when bidders' valuations are correlated? Consider the special case in which bidders have private valuations which are not statistically independent of each other. Jacques Cremer and Richard McLean (1985a, 1985b) have provided a method for the seller to extract all of the gains from trade from the buyer. This requires a certain level of correlation among the bidders' valuations; it will not work with pure independent values. In addition, it is assumed that only a finite number of valuations are possible; the case of continuously distributed values is not considered. The mechanism design may be understood as follows. Represent by $\pi(v_i|v^{-i})$ the probability that the ith bidder's value is v_i, given the vector of other bidders' values v^{-i}. The seller offers each bidder a lottery plus participation in a Vickrey (second-price) auction. The trick is to design the lottery so that the lottery's expected value for the ith agent is precisely the expected value the agent receives from the auction, and the outcome of the lottery depends only on v^{-i} and not on v_i. This ensures that the agents continue to bid honestly in the Vickrey auction, since a change in bid by bidder i does not affect his lottery. The condition that permits such a lottery to be designed is that v_i is, in a probabilistic sense, recoverable from v^{-i}; more precisely, the matrix $\pi(v_i|v^{-i})$ is of full rank (equal to the number of possible values v_i can take on). Thus, the Cremer-McLean result asserts that, if the distribution of each agent's valuation is altered sufficiently by changes in the other agents' valuations, the full surplus can be extracted by the seller.[13]

11. FURTHER TOPICS

Each of the four main assumptions underlying the benchmark model has now been relaxed. This section more briefly considers some additional questions.

11.1 VARIABLE SUPPLY AND CAPACITY CONSTRAINTS

In the models considered so far, only a single unit of an item is being sold. The opposite polar case occurs when the monopolist can freely vary the amount offered for sale, producing under constant returns to scale. What then is the optimal selling scheme when, as before, the seller does not know the buyers' tastes?

Suppose for simplicity that each buyer wishes to buy at most one unit of the commodity. Each buyer's valuation of the single unit (or willingness to pay) is drawn independently from a distribution F, with density f. Denote by c the constant average production cost. <u>With unlimited capacity, constant returns to scale, and independent valuations, the monopolist's optimal selling mechanism is to post a fixed price</u> r <u>defined by</u>

$$(10) \qquad c = r - \frac{1 - F(r)}{f(r)}$$

(Harris and Raviv (1981b), Matthews (1983)). (To obtain (10), note that the probability that any one customer buys at price r is 1 - F(r), so that the expected profit per potential buyer is [1 - F(r)](r-c). This is maximized when (10) is satisfied.) The posted price when capacity is unlimited is exactly the same as the reserve price when only one unit is to be sold (as in Section 6).

Interpolation from the extreme cases of single unit and unlimited capacity suggests the prediction that a monopolist will post a fixed price when his capacity constraint is not binding, and will sell by auction when his capacity constraint is significant.

Between the cases of single unit and unlimited capacity is the case in which a fixed quantity is put up for sale and buyers may bid for some portion of the available units. One example of such an auction is the weekly United States Treasury bill auction (Smith (1967)). Another example is the New Zealand Government's auctioning of import quotas, which was introduced so that the Government could estimate, on the basis of the prices reached, the extent of protection being provided by the quotas (M. Pickford (1985)).

Two kinds of sealed-bid auctions are used to sell multiple units simultaneously. Bidders submit bids which consist of both a price and a desired number of units of the commodity. Suppose enough units are available that the h highest bidders can be awarded the item. In the discriminatory auction, each of these h bidders pays the amount he bid. In the uniform-price auction, each successful bidder pays a price equal to the highest unsuccessful bid (the (h + 1)st bid). Clearly the former corresponds, in the single-unit case, to the first-price auction, and the latter corresponds to the second-price auction.

Results similar to the single-unit case can be established for multiple-unit auctions. For example, for the benchmark model, the discriminatory auction yields on average the same revenue as the uniform-price auction. Risk aversion of bidders results in the discriminatory auction yielding higher average revenue than the uniform-price auction; while in the common-value case this ordering is reversed (Weber (1983)).[14]

11.2 COLLUSION AMONG THE BIDDERS

It has been assumed so far that the bidders act noncooperatively:

they do not coordinate their bids. This assumption may not be appropriate in some circumstances, especially when the same bidders compete with each other over many successive auctions. Collusion may consist of either explicit agreements about which bidder will be allowed to win any particular auction, or implicit understandings that restraint will be exercised in bidding.

The familiar analysis of repeated oligopoly games can be applied to repeated auctions. In an infinitely repeated game, a collusive outcome can be maintained as a noncooperative equilibrium if each of the oligopolists adopts a strategy of threatening to retaliate to any deviation from the collusive arrangement by reverting to noncooperative behavior in future periods. (This is the Folk Theorem of repeated games: see, for example, Robert Aumann (1981).) Anecdotal evidence suggests the empirical relevance of the repeated-game argument. It has been observed that collusion in antique and artwork auctions is enforced by retaliatory strategies: the response to a defection by a cartel (or "ring") member is that "vindictive competition leads to crazy prices" (Jeremy Cooper (1973, pp. 37-38)).[15]

For retaliatory strategies to be workable, it must be possible for the bidders to observe the other bidders' past actions. As George Stigler (1964, p. 48) remarked of government procurement auctions, "the system of sealed bids, publicly opened with full identification of each bidder's price and specifications, is the ideal instrument for the detection of price cutting...collusion will always be more effective against buyers who report correctly and fully the prices tendered to them." This may explain the tendency for private-sector firms to use closed negotiations rather than formal auctions for procurement.

A cartel may include some or all of the bidders. How does a cartel decide which member is to receive the item? In English auctions, a common method in practice is for one member arbitrarily to be assigned to bid for the item without competition from his fellow cartel members. Afterwards, the item is re-auctioned among the cartel members. Such behavior occurs in auctions of antiques, fish, timber, industrial machinery, and wool (Cassady (1967, Ch. 13), Cooper (1977, pp. 35-38), F.H. Gruen (1960)). The cartel member who values the item the most will win the bidding in the illicit auction at a price equal to the second-highest valuation among cartel members. The cartel shares among its members a sum of money equal to the difference between the price reached in the cartel's own auction and the price reached in the original auction.

What should a seller do if he believes he faces a buyers' cartel? According to Cassady (1967, pp. 228-230), reserve prices are commonly used to counter the activities of cartels. To understand this practice, assume for simplicity that all n bidders belong to the cartel, and that assumptions (A1), (A2), (A3), and (A4) of the benchmark model are satisfied. The cartel not only reduces to one of the effective number of bidders; it also changes the effective distribution of valuations. The relevant distribution for the seller faced with the cartel is the distribution of the maximum of n valuations, each drawn from the distribution $F(x)$; this distribution of maxima is $F^n(x)$, with density $nF^{n-1}(x)f(x)$. Thus an argument similar to that in Section 6 yields: <u>the optimal anti-cartel reserve price</u> r <u>satisfies</u>

$$(11) \qquad v_o = r - \frac{1 - F^n(r)}{nF^{n-1}(r)f(r)} ,$$

where v_o is the seller's own valuation (Daniel Graham and Robert Marshall (1984, 1985)). This implies that the anti-cartel reserve price increases with the number of cartel members. Also, the anti-cartel reserve price is higher than the optimal reserve price in the absence of collusion.

11.3 DOUBLE AUCTIONS

It was assumed in the foregoing analysis that there was only one seller and that his valuation of the item was known to all of the bidders. We now relax both of these assumptions, and suppose that there are several sellers who, like the buyers, draw their valuations independently from some probability distribution.

In a double auction, both sellers and buyers submit sealed bids, buyers for how much they are willing to pay and sellers for the price at which they are willing to sell. The London gold market works much like this idealized double auction. The double auction may also be taken as a stylized description of how prices are formed in competitive markets in the absence of the Walrasian auctioneer.

Ranking the buyers' bids from highest to lowest produces a demand function which is a step function. Similarly, ranking the sellers' bids produces a supply step function. The intersection of demand and supply generally gives a quantity and an interval of prices. This is the quantity exchanged; the price is chosen from the interval according to some arbitrary rule.

Few results on the double auction exist. The main result is an efficiency result. If there are sufficiently many buyers and sellers, then there is no other trading mechanism which would increase some traders' expected gains from trade without lowering other traders' expected gains from trade (Wilson (1985)). That is, the double auction would survive any attempt to get the traders to agree unanimously to change the trading institution. This result is an asymmetric-information analogue of the Arrow-Debreu theorem on the Pareto efficiency of competitive equilibria.

12. AN AGENDA FOR RESEARCH

To what extent have the foregoing results in auction theory provided answers to the broad questions posed by Hayek (1945) about the working of the price system in the presence of informational asymmetries?

The questions examined by auction theory form a subset of the class of adverse-selection problems; that is, problems in which one individual knows something that the others do not know. The foregoing theorems show, however, that bidders can be induced to reveal implicitly (or, in the case of the Vickrey auction, directly) their private information, namely, how highly they value the item. With some exceptions, the bidder who values the item the most is awarded it. The Revelation Principle shows that, in many cases, individuals can be given incentives to share their information. Markets can work despite the dispersion of information, as Hayek suggested. There need be nothing adverse about adverse selection.[16]

The results described above have generated some insights into how the price system works: into the nature of the process of bidding competition and price formation. We have arrived at some understanding of why particular trading institutions - the English auction, for example - arise in particular circumstances. (See the optimal-auctions analyses of Sections 6 and 7.) Do prices serve to aggregate dispersed information? The result of Wilson (1977) (summarized in Section 10) shows that they can: provided there are many bidders, and provided information is sufficiently dispersed among the bidders, the price equals the item's true value even though no individual knows what this true value is. Is it correct, as Hayek asserted, that the price summarizes all of the relevant information about supply and demand? The answer is no: it is in the seller's interest, if possible, to adjust the price after the sale in the light of any new information he obtains about the item's value to the buyer (as discussed in Section 8).

Much remains to be done. The auction models are partial-equilibrium models. The role of the price system in coordinating the actions of different people cannot be understood except within a general-equilibrium system. How to embed bidding models in a general-equilibrium context remains an open question. Questions of the existence and social optimality of competitive equilibrium with informational asymmetries await the resolution of this question.

One crucial step towards a general-equilibrium formulation is modelling competition among mechanism designers. Consider a seller who designs an auction-like mechanism. If the buyers have an alternative, that is, the seller faces a competing mechanism, competition may constrain the seller's choice of mechanism. If there is one phenomenon that economists understand to be important, it is competition. Solving the difficult technical problems of modelling competition among mechanisms is, in our view, the major problem facing the auction literature.

The advantages resulting from the seller's ability to commit himself to a mechanism have been made apparent. We know why the seller would want to commit himself. What is less well understood is how he is able to achieve the commitment. One way of rationalizing this is that the seller plays the game repeatedly, while the buyers play it only once. Then it might be in the seller's interest to adhere to his mechanism, for a deviation might destroy his commitment ability for the future. However plausible this argument appears, it remains to be formalized and its logical coherence remains to be established.[17]

13. MACHIAVELLIAN ADVICE TO A MONOPOLIST

You are the seller of some good or service in the fortunate position of having no competitors. How should you design your selling methods so as to squeeze the last possible cent from your customers?

The first rule is to make your customers believe that, whatever pricing strategy you have chosen, you will not under any circumstances depart from it. Once you are visibly committed, all that prevents you from completely exploiting your customers is your lack of knowledge of exactly how high you can drive the price to any particular buyer without losing the sale.

Should you post a take-it-or-leave-it price, or should you hold an auction? If your production capacity is large, fixing a price maximizes your expected profits. The price you should charge, if you believe your customers' valuations of your product are approximately uniformly distributed, is the average of your unit production cost and the highest possible valuation. If, on the other hand, you have only one or a few units to sell, you should sell by auction.

What kind of auction should you choose? To answer this question, you must know whether your customers would be prepared to pay higher prices in exchange for your sheltering them from risk. You must also know whether differences among the bidders' valuations of the item are due to inherent differences in their tastes or to their having made different guesses about the unique true value of the item.

If your customers would pay nothing to be sheltered from risk and their different valuations reflect their different tastes, then your best selling device is any of the simple auction forms: English, Dutch, first-price sealed-bid, or second-price sealed-bid. You should impose a reserve price. If the commodity is useless to you unsold, and if you estimate the distribution of your customers' valuations to be approximately uniform, the reserve price you should set is one-half of the maximum possible valuation.

If your customers prefer to avoid risk, then you are no longer indifferent among the simple auction forms: your revenue will on average be higher from a first-price sealed-bid auction than from an English auction. However, if you have very sophisticated computational capacities, you can do still better by announcing that you will require payment from bidders who bid too low and that you will subsidize bidders who bid high but not quite high enough to win. You should, if possible, keep secret from each bidder how many other bidders he is competing with.

If the bidders fall into several categories and you observe that there are systematic differences in valuations across categories, then you can exploit this to your advantage (provided you can somehow prevent resale by the winning bidder). You do this by discriminating in favor of bidders in the category with on average low valuations: you announce that you will accept a lower bid from a member of the favored category over a higher bid from a member of another category, provided the difference in bids is not too great.

If you can monitor the buyer's subsequent usage of the commodity, you should, by the use of a royalty scheme, require continuing payments from the buyer based on value in use: royalties induce the bidding to be more competitive.

If the item you are selling has a unique true value, but the bidders have different imperfect estimates of this value, the English auction will on average yield more revenue for you than any of the other simple auction forms. You can encourage the bidders to raise their bids by publicizing any information you yourself have about the item's true value.

Finally, if your monopoly power is being eroded by the formation of a countervailing buyers' cartel, you can regain some of your monopoly profits by increasing your reserve price, making it higher the larger the number of cartel members.[18]

FOOTNOTES

* We thank Knick Harley, Peter Kuhn, David Laidler, and John Whalley for useful comments.

[1]Since there already exists a survey (by Richard Engelbrecht-Wiggans (1980)) and a bibliography (by Robert Stark and Michael Rothkopf (1979)) of the earlier literature, this paper will focus on recent developments. The surveys by Paul Milgrom (1985a, 1985b) cover in more depth a narrower range of topics than the present survey. There is a large and growing literature on experiments with auctions which will not be surveyed here for space reasons. See Charles Plott (1982) and Vernon Smith (1982) for general surveys of experimental economics, including some discussion of auction experiments.

[2]More precisely, let z and z' represent a pair of (m+n) - vectors, and let g(z) denote the joint probability density of the random variables z. Denote by $z \vee z'$ the component-wise maximum of z and z', and by $z \wedge z'$ the component-wise minimum. Then the variables are defined to be affiliated if, for all z, z',

$$g(z \vee z')\, g(z \wedge z') \geq g(z)\, g(z').$$

Assuming differentiability of g, this is equivalent to

$$\frac{\partial^2}{\partial z_i \partial z_j} (\log g) \geq 0,$$

where z_i, z_j, $i \neq j$, are elements of z. See Milgrom and Weber (1982a) for more details.

[3]Eric Maskin and John Riley (1985) provided simple examples illustrating the effects of varying assumptions (A1), (A2), and (A3).

[4]A proof of this is given in McAfee and McMillan (1985e). For the case of bidders competing to sell an object rather than to buy (for example,

bidding for contracts), $[1-F(v)]/f(v)$ is replaced by $-G(c)/g(c)$, where c is a bidder's production cost and G and g are the probability distribution function and density function of bidders' costs--McAfee and McMillan (1985b).

[5]For J to be a decreasing function, the distribution function F must be sufficiently concave. Also, J is increasing if and only if $[1 - F(x)]^{-1}$ is convex. The consequences of relaxing the assumption that J is strictly increasing have been examined by Maskin and Riley (1980, 1983b), Myerson (1981), Ralph Haywood (1984), and McAfee and McMillan (1985c, Ch. 11, Appendix).

[6]Note that the foregoing discussion restricted attention to symmetric equilibria: it was assumed that bidder j would bid the same as bidder k if they valued the item equally highly. On existence, symmetry, and uniqueness of bidding equilibria, see Maskin and Riley (1983c) and Milgrom and Weber (1980).

[7]However, if the bidders must incur a cost in preparing their bids, or if the seller must incur a cost in checking the credentials of the bidders, it need not be the case that expected net price rises with the number of bidders: see respectively Samuelson (1985) and McAfee and McMillan (1985e).

[8]The results to follow extend to the case of m classes of bidders for any $m \leq n$, where n is the number of bidders. The discussion here has two classes of bidders solely for ease of exposition.

[9]This is related to the "linkage effect" discussed by Milgrom (1985a) and Milgrom and Weber (1982a).

[10]The effects of bidders' having different degrees of risk aversion were investigated by James Cox, Vernon Smith and James Walker (1982).

[11]In a similar vein to the winner's curse, there is a conventional wisdom that the excitement generated in an English auction causes bidders to bid too high. Such assertions should also be treated with caution: they may be explained by the observation that, by the nature of the auction process, all but two of the bidders believe the selling price exceeds the item's value.

[12]Note that this result, which takes as given the diversity of the bidders' information, ignores the possibility that bidders might be able, at some cost to themselves, to obtain extra information relating to the item's true value. On information acquisition in auctions, see Engelbrecht-Wiggans, Milgrom and Weber (1983), Tom Lee (1982, 1985), Matthews (1984b), Milgrom (1981b, 1985b), and Milgrom and Weber (1982b).

[13]Because the Vickrey auction is a dominant-strategy auction, the Cremer-McLean auction/lottery also has a dominant equilibrium. With slightly weaker assumptions on π, a Bayes-Nash equilibrium auction which extracts all of the surplus can be constructed.

[14]For further analysis of auctions of multiple units of a commodity, see Engelbrecht-Wiggans and Weber (1979), Harris and Raviv (1981a, 1981b), Maskin and Riley (1983b, 1984a), Matthews (1983), Milgrom (1985b), Thomas Palfrey (1980, 1983), Smith (1967), and Weber (1983).

[15]Milgrom (1985a) showed that there is a sense in which repeated English auctions are more susceptible to collusion than repeated first-price sealed-bid auctions. Marc Robinson (1985) showed that an implicitly collusive outcome could be reached as a Nash equilibrium in a single English auction. However, his analysis assumes that each bidder knows his rivals' valuations, which assumes away much of the auction problem.

[16]Arrow (1984) recently suggested that the term "adverse selection" be replaced by the more descriptive "hidden knowledge".

[17]Also, repeated games typically have many equilibria. It may be that, while commitment emerges as an equilibrium of a repeated mechanism-design game, noncommitment remains as another equilibrium.

[18]For the reader interested in pursuing auction theory more deeply than this nontechnical survey, we recommend beginning with the following papers. Vickrey's remarkable 1961 paper, two decades ahead of its time, is still worth reading as an introduction to the analysis of auctions. Milgrom and Weber (1982a) provide a very general framework for analyzing auctions and compare different auction forms and different seller policies. On the designing of optimal auctions, the central papers are Myerson (1981) and Riley and Samuelson (1981): the former paper is more general, but the latter is more readable.

BIBLIOGRAPHY

Akerlof, George, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." Quarterly Journal of Economics, Vol 84, No. 3 (August 1970): 488-500.

Arrow, Kenneth J., "Toward a Theory of Price Adjustment," in M. Abramovitz et al, The Allocation of Economic Resources (Stanford: Stanford University Press, 1959).

Arrow, Kenneth J., "The Economics of Agency," Technical Report No. 451, IMSSS, Stanford University, October 1984.

Aumann, Robert J., "Survey of Repeated Games" in V. Bohm and H.-H. Nachtkamp, eds., Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern (Mannheim: Bibliographisches Institut, 1981).

Brannman, Lance, Klein, J. Douglas, and Weiss, Leonard, "Concentration and Winning Bids in Auctions," Antitrust Bulletin 29 (Spring 1984): 27-31.

Capen, E.C., Clapp, R.V., and Campbell, W.M., "Competitive Bidding in High-Risk Situations," Journal of Petroleum Technology 23 (June 1971): 641-653.

Cassady, R., Jr., Auctions and Auctioneering (Berkeley: University of California Press, 1967).

Cassing, James and Douglas, Richard W., "Implications of the Auction Mechanism in Baseball's Free Agent Draft," Southern Economic Journal 47 (July 1980): 110-121.

Cooper, Jeremy, Under the Hammer: The Auctions and Auctioneers of London (London: Constable, 1977).

Cox, James C., and Isaac, R. Mark, "In Search of the Winner's Curse," Economic Inquiry, Vol. 22, No. 4 (October 1984): 579-592

Cox, James C., Smith, Vernon L., and Walker, James M., "Auction Market Theory of Heterogeneous Bidders," Economics Letters 9 (1982): 319-325.

Cremer, Jacques, and McLean, Richard, "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," CARESS Working Paper No. 85-17, University of Pennsylvania, June 1985.

Cremer, Jacques and McLean, Richard P., "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demand are Interdependent," Econometrica 53 (March 1985): 345-363.

DeMayo, Peter, "Bidding on New Ship Construction," in R. Engelbrecht-Wiggans, M.Shubik and R.M. Stark, eds., Auctions, Bidding and Contracting: Uses and Theory (New York: New York University Press, 1983).

Dessauer, John P., Book Publishing (New York: Bowker, 1981).

Ederington, Louis H., "Biddig for Securities: The Effect on the Issuer's Interest Costs," Journal of Business, Vol 51, No. 4 (October 1978): 673-686.

Engelbrecht-Wiggans, R., "Auctions and Bidding Models: A Survey," Management Science 25, No. 12 (February 1980): 119-142.

Engelbrecht-Wiggans, R., Milgrom, P.R., and Weber, R.J., "Competitive Bidding and Proprietary Information," Journal of Mathematical Economics 11 (April 1983): 161-169.

Engelbrecht-Wiggans, R. and Weber, Robert J., "An Example of a Multi-Object Auction Game," Management Science 25 (December 1979): 1272-1277.

French, Kenneth R. and McCormick, Robert E., "Sealed Bids, Sunk Costs, and the Process of Competition," Journal of Business 57 (October 1984): 417-442.

Gaver, Kenneth M. and Zimmermann, Jerold L., "An Analysis of Competitive Bidding on BART Contracts," Journal of Business 50 (1977): 279-295.

Gilley, Otis V. and Karels, Gordon V., "The Competitive Effect in Bonus Bidding: New Evidence," Bell Journal of Economics 12 (1981): 637-648.

Graham, Daniel A. and Marshall, Robert C., "Bidder Coalitions at Auctions," mimeo, Duke University, November 1984.

Graham, Daniel A. and Marshall, Robert C., "Collusive Behaviour at a Single Object English Auction," Working Paper 85-01, Duke University, February 1985.

Griesmer, J.H., Levitan, R.E., and Shubik, M., "Towards a Study of Bidding Processes, Part Four: Games with Unknown Costs," Naval Research Logistics Quarterly 14 (1967): 415-433.

Groves, Theodore, and Ledyard, John, "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." Econometrica 45 (1977):783-809.

Gruen, F.H., "Goulburn, Forward Prices, and Pies," Review of Marketing and Agricultural Economics 28 (June 1960): 85-96.

Hansen, Robert C., "Auctions with Contingent Payments," American Economic Review 75, 4 (September 1985a): 862-865.

Hansen, Robert C., "Empirical Testing of Auction Theory," American Economic Review: Papers and Proceedings, 75, 2 (May 1985b): 156-159.

Harris, Milton and Raviv, Artur., "Allocation Mechanisms and the Design of Auctions," Econometrica 49 (1981a): 1477-1500.

Harris, Milton and Raviv, Artur., "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," American Economic Review 71 (June 1981b): 347-365.

Harris, Milton and Townsend, Robert M., "Allocation Mechanisms, Asymmetric Information, and the 'Revelation Principle'," in G. Fiewel, ed., Issues in Contemporary Microeconomics and Welfare (Albany: State University of New York Press, 1985).

Harsanyi, John C., "Games with Incomplete Information Played by 'Bayesian' Players," Management Science 14 (November 1967): 159-189; 15 (January 1968): 320-334; 15 (March): 486-502.

Harstad, Ronald M. and Levin, Dan, "A Class of Dominance Solvable Common-Value Auctions," Review of Economic Studies 22, No. 3 (July 1985): 525-528.

Hayek, F.A., "The Use of Knowledge in Society," American Economic Review 35, No. 4 (September 1945): 519-530.

Haywood, Ralph, "Optimal High-Valuer Auctions," mimeo, University of California, San Diego, November, 1984.

Holt, Charles A., Jr., "Uncertainty and the Bidding for Incentive Contracts," American Economic Review 69 (1979): 697-705.

Holt, Charles A., Jr., "Competitive Bidding for Contracts under Alternative Auction Procedures," Journal of Political Economy 88 (1980): 433-445.

Holt, Charles A., Jr., "Bidding for Contracts," in Bayesian Analysis in Economic Theory and Time Series Analysis (Amsterdam: North-Holland, 1982).

Kawasaki, Seiichi, and McMillan, John,"The Design of Contracts: Evidence from Japanese Subcontracting," mimeo, University of Western Ontario, November 1985.

Laffont, Jean-Jacques, and Maskin, Eric., "Optimal Reservation Price in the Vickrey Auction," Economics Letters 6 (1980): 309-313.

Lee, Tom K., "Resource Information Policy and Federal Resource Leasing," Bell Journal of Economics 12 (1982): 561-568.

Lee, Tom K., "Competition and Information Acquisition in First Price Auctions, "Economics Letters, Vol. 18, Nos. 2-3 (1985): 129-132.

Lee, Tom K., "Incomplete Information, High-Low Bidding, and Public Information in First Price Auctions," Management Science to appear.

Leland, Hayne E., "Optimal Risk Sharing and the Leasing of Natural Resources, with Application to Oil and Gas Leasing on the OCS," Quarterly Journal of Economics 92, No. 3 (August 1978): 413-437.

Luton, Richard and McAfee, R. Preston, "Sequential Procurement Auctions," mimeo., University of Western Ontario, October 1985.

Maskin, Eric and Riley, John, "Auctioning an Indivisible Object," mimeo., UCLA, April 1980.

Maskin, Eric and Riley, John, "Auctions with Asymmetric Beliefs," Working Paper No. 254, UCLA, April 1983a.

Maskin, Eric and Riley, John, "Monopoly with Incomplete Information," Working Paper No. 268, UCLA, August 1983b.

Maskin, Eric and Riley, John, "Uniqueness of Equilibrium in Open and Sealed Bid Auctions," mimeo, M.I.T., April 1983c.

Maskin, Eric and Riley, John, "Monopoly with Incomplete Information," Rand Journal of Economics, Vol 15, No. 2 (Summer 1984a): 171-196.

Maskin, Eric and Riley, John, "Optimal Auctions with Risk Averse Buyers," Econometrica 52 (November 1984b): 1473-1518.

Maskin, Eric and Riley, John, "Auction Theory with Private Values," American Economic Review: Papers and Proceedings, 75, 2 (May 1985): 150-155.

Matthews, Steven A., "Selling to Risk Averse Buyers with Unobservable Tastes," Journal of Economic Theory 30 (1983): 370-400.

Matthews, Steven A., "On the Implementability of Reduced Form Auctions," Econometrica 52 (November 1984): 1519-1523.

Matthews, Steven A., "Information Acquisition in Discriminatory Auctions," in Bayesian Models in Economic Theory, edited by M. Boyer and R.E. Kihlstrom (Amsterdam: North-Holland, 1984), pp. 181-208.

Matthews, Steven A., "On Comparing Auctions for Risk Averse Buyers: A Buyer's Point of View," Discussion Paper No. 664, Northwestern University, October, 1985.

McAfee, R. Preston, and McMillan, John, "Bidding for Contracts: A Principal-Agent Analysis," presented at the Conference on Strategic Behaviour and the Theory of the Firm, University of Western Ontario, 1984.

McAfee, R. Preston, and McMillan, John, "Auctions with a Stochastic Number of Bidders," Technical Report No. 9, Centre for Decision Sciences and Econometrics, University of Western Ontario, December 1985a.

McAfee, R. Preston and McMillan, John, "Discrimination in Auctions," Discussion Paper No. 85-7, University of California, San Diego, Janury 1985b.

McAfee, R. Preston and McMillan, John, Incentives of Government Contracting (Toronto: University of Toronto Press, 1985c), to be published.

McAfee, R. Preston and McMillan, John, "Competition for Principal-Agent Contracts," mimeo., University of Western Ontario, October 1985d.

McAfee, R. Preston and McMillan, John, "Search Mechanisms," mimeo., University of Western Ontario, November 1985e.

Mead, Walter J., "Natural Resource Disposal Policy: Oral Auction versus Sealed Bids," Natural Resources Journal 7 (1967): 195-224.

Mead, Walter J., Asbjorn, Moseidjord, and Sorenson, Philip E., "Competitive Bidding under Asymmetric Information: Behaviour and Performance in Gulf of Mexico Drainage Lease Sales 1959-1969," Review of Economics and Statistics 66 (August 1984): 505-509.

Milgrom, Paul R., "Auction Theory," presented at the World Congress of the Econometric Society, September 1985a.

Milgrom, Paul R., "The Economics of Competitive Bidding: A Selective Survey," In Social Goals and Social Organization, edited by L. Hurwicz, D. Schmeidler, and H. Sonnenschein (Cambridge: Cambridge University Press, 1985b).

Milgrom, Paul R., "A Convergence Theorem for Competitive Bidding with Differential Information," Econometrica 47 (April 1979a): 679-688.

Milgrom, Paul R., The Structure of Information in Competitive Bidding (New York: Garland, 1979b).

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," Bell Journal of Economics 12 (Autumn 1981a): 380-391.

Milgrom, Paul R., "Rational Expectations, Information Acquisition, and Competitive Bidding," Econometrica 49 (1981b): 921-941.

Milgrom, Paul R., and Weber, Robert J., "Distributional Strategies for Games with Incomplete Information," Discussion Paper No. 428, Northwestern University, June 1980.

Milgrom, Paul R., and Weber, Robert J., "A Theory of Auctions and Competitive Bidding," <u>Econometrica</u> 50 (September 1982a): 1089-1122.

Milgrom, Paul R., and Weber, Robert J., "The Value of Information in a Sealed-Bid Auction," <u>Journal of Mathematical Economics</u> 10 (1982b): 105-114.

Moore, John, "Global Incentive Constraints in Auction Design," <u>Econometrica</u> 52 (November 1984): 1523-1536.

Myerson, Roger B., "Optimal Auction Design," <u>Mathematics of Operations Research</u> 6 (February 1981): 58-73.

Myerson, Roger B., "Bayesian Equilibrium and Incentive Compatibility: An Introduction," Discussion Paper No. 548, Northwestern University, February 1983.

Palfrey, Thomas R., "Multiple-Object, Discriminatory Auctions with Bidding Constraints: A Game-Theoretic Analysis," <u>Management Science</u> 26, No. 9 (September 1980): 935-946.

Palfrey, Thomas R., "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information," <u>Econometrica</u> 51 (March 1983): 463-484.

Pelto, Chester R., "That Statistical Structure of Bidding for Oil and Mineral Rights," <u>Journal of the American Statistical Association</u> 66 (1971): 456-460.

Pickford, M., "A New Test for Manufacturing Industry Efficiency: An Analysis of the Results of Import Licence Tendering in New Zealand, <u>International Journal of Industrial Organization</u> 3 (1985): 153-177.

Plott, Charles, "Industrial Organization Theory and Experimental Economics," <u>Journal of Economic Literature</u> 20 (December 1982): 1485-1527.

Quirk, James and Terasawa, Katsuaki, "The Winner's Curse and Cost Estimation Bias in Pioneer Projects," Working Paper No. 512, California Institute of Technology, 1984.

Ramsey, James B., Bidding and Oil Leases (Greenwich, Conn: J.A.I. Press, 1980).

Reece, Douglas K., "Competitive Bidding for Offshore Petroleum Leases," Bell Journal of Economics 9 (Autumn 1978): 369-384.

Reece, Douglas K., "Alternative Bidding Mechanisms for Offshore Petroleum Leases," Bell Journal of Economics 10 (1979): 659-669.

Riley, John G., "Ex Post Information in Auctions," Working Paper No. 367, U.C.L.A., May 1985.

Riley, John G. and Samuelson, William F., "Optimal Auctions," American Economic Review 71 (1981): 381-392.

Robinson, Marc S., "Collusion and the Choice of Auction," Rand Journal of Economics Vol. 116, No. 1 (Spring 1985): 141-145.

Samuelson, William F., "Competitive Bidding for Incentive Contracts," in R. Engelbrecht-Wiggans, M. Shubik, and R.M. Stark, eds., Auctions, Bidding, and Contracting: Uses and Theory (New York: New York University Press, 1983).

Samuelson, William F., "Bidding for Contracts," Discussion Paper No.32/84, School of Management, Boston University, April 1984.

Samuelson, William F., "Competitive Bidding with Entry Costs," Economics Letters Vol. 17, Nos. 1-2 (1985): 53-57.

Schelling, Thomas C., "The Strategy of Conflict (Cambridge: Harvard University Press, 1960).

Shubik, Martin, "Auctions, Bidding, and Markets: A Historical Sketch, in R. Engelbrecht-Wiggans, M. Shubik, and R. Stark, eds., Auctions, Bidding, and Contracting: Uses and Theory (New York: New York University Press, 1983).

Smith, James L., "Non-Aggressive Bidding Behavior and the 'Winner's Curse'", Economic Inquiry, Vol. 29, No. 3 (July 1981): 380-388.

Smith, Vernon L., "Experimental Studies of Discrimination versus Competition in Sealed-Bid Auction Markets," Journal of Business Vol. 40, No. 1 (January 1967): 56-86.

Smith, Vernon L., "Microeconomic Systems as Experimental Science," American Economic Review 72 (December 1982).

Stark, Robert and Rothkopf, Michael,H., "Competitive Bidding: A Comprehensive Bibliography," Operations Research 27 (March-April 1979): 364-490.

Stigler, George J., "A Theory of Oligopoly," Journal of Political Economy 72 (1964): 44-61.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," Journal of Finance 16 (March 1961): 8-37.

Weber, R.J., "Multiple-Object Auctions," in R. Engelbrecht-Wiggans, M. Shubik, and R.M. Stark, eds., Auctions, Bidding, and Contracting: Uses and Theory (New York: New York University Press, 1983).

Wilson, Robert B., "A Bidding Model of Perfect Competition," Review of Economic Studies 44 (October 1977): 511-518.

Wilson, Robert B, "Auctions of Shares," Quarterly Journal of Economics 93, No. 4 (November 1979): 675-690.

Yamey, B.S., "Why 2,310,000 for a Valazquez? An Auction Bidding Role," Journal of Political Economy 80 (November-December 1972): 1323-1327.

Yuspeh, Larry, "A Case for Increasing the Use of Competitive Procurement in the Department of Defense," in Y. Amihud, ed., Bidding and Auctioning for Procurement and Allocation (New York: New York University Press, 1976).