

Auctions and Bidding

By R. PRESTON MCAFEE

University of Western Ontario

and

JOHN McMILLAN

University of Western Ontario and University of California, San Diego

We thank Michael Carter, Knick Harley, Peter Howitt, Peter Kuhn, David Laidler, Glenn MacDonald, Robert Porter, William Samuelson, and John Whalley, as well as John Pencavel and four referees for useful comments. The Ontario Economic Council and the Social Sciences and Humanities Research Council of Canada provided research support.

I. *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 that 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 to make a rational decision is the vector of prices; he need not know the determinants of supply and demand that underlie the prices.

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 as effective at transmitting 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. Should we conclude that 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 modeling of strategic behavior under asymmetric information. One part of this broad research program is the theory of bidding mechanisms. The modeling 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 many 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? Even in an otherwise perfectly competitive market, Arrow argued, during the process of price formation there is considerable uncertainty, resulting in each seller facing a downward-sloping demand curve or each buyer facing an upward-sloping supply curve. During the adjustment to the competitive equilibrium, "the market consists of a number of monopolists facing a number of monopsonists" (p. 47). Auction theory provides one explicit model of price making (although it is a restricted model of price formation in that it ignores the bargaining aspects of the process sketched by Arrow).

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 than is most frontier mathematical economics.

The theoretical results in auction theory can explain the existence of certain trading institutions, and perhaps can even suggest improvements in the existing institutions: Thus auction theory has both positive and normative aspects. 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 assumptions, as well as proofs, see the cited papers.²

II. *The Types and Uses of Auctions*

What is an auction? An auction is a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants.

What kinds of goods are sold by auc-

tion? The list is long: Artwork, books, antiques, agricultural produce, mineral rights, 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?"

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 offering to sell and bidding 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.

¹ Because there already are 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 (1985, 1986) cover in more depth a narrower range of topics than the present survey. A brief survey of auction theory is given by Robert Wilson (1987).

² To the reader interested in pursuing auction theory more deeply than this nontechnical survey, we recommend beginning with the following papers. William 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 Robert 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 Roger Myerson (1981) and John Riley and William Samuelson (1981); the former paper is more general, but the latter is more readable.

Novel uses of auctions include the auctioning of the rights to a natural monopoly as a substitute for regulating it (Harold Demsetz 1968; Oliver Williamson 1976); the auctioning of import quotas as a way of generating data on the extent of protection afforded by the quotas as well as capturing the rents for the government (M. Pickford 1985); the use of an auction mechanism for selecting a location for noxious facilities like prisons and hazardous-waste disposal plants (Howard Kunreuther and Paul Kleindorfer 1986); and the auctioning of airport time slots to competing airlines, proposed as an improvement over the existing slot quotas (S. J. Rassenti, V. L. Smith, and R. L. Bulfin 1982).

In a double auction, several buyers and several sellers submit bids simultaneously. The double auction is a stylized representation of organized exchanges such as stock exchanges and commodity markets.

What are the types of auctions that are in use? Four basic types are used when a unique item is to be bought or sold: 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 (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 (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 by using royalties (James Ramsey 1980). 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 (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.

III. *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 of 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.³ (However, there may be any number, large or small, of bidders.)

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.

Auction theory sidesteps such 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 that 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 higher than the highest bid and yet 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 that induce the bidders to bid in desirable ways. In *The Strategy of Conflict*, Thomas Schelling explained the advantages in general strategic situations of commitment power: "If the buyer can accept an irrevoc-

³ An important exception, the double auction, will be discussed in Section XI.

cable *commitment*, 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). This follows from "the paradox that the power to constrain an adversary may depend on the power to bind oneself" (p. 22).

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 that 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 to commit oneself credibly 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; the commitment makes the threat credible. 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.⁴

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

IV. *The Nature of the Uncertainty*

Asymmetry of information 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 he does not know the bidders' valuations.⁵

How the bidders respond to uncertainty depends on their attitudes toward risk. Thus one aspect of any particular bidding situation that the modeler must take into account is the bidders' risk attitudes. (The risk attitudes of the seller may also matter; however, we shall assume throughout that the seller is risk neutral.)

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 modeled.

At one extreme, suppose that each bidder knows precisely how highly he values

⁴ Milgrom (1986) showed within a bargaining model that even if a seller is in a weak bargaining position (that is, if he does not have commitment ability), he might still choose to sell by auction. However, he then cannot use any of the strategic policies discussed in Sections VI through X to follow (such as imposing a reserve price). Hence we shall follow most of the existing literature by retaining the assumption of precommitment throughout this paper.

⁵ Note that informational asymmetries may or may not be removable in principle. On the one hand, for example, tastes are inherently unobservable. On the other hand, it might be feasible to monitor a firm's production costs, but simply too costly to do so. This distinction will be ignored in what follows.

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 , $i = 1, \dots, n$, 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 *independent-private-values model*. 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 that 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 , $i = 1, \dots, n$, are independent draws from some probability distribution $H(v_i|V)$. All agents know the distribution H . This is called the *common-value model*.

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 will not cause him to change his own valuation (although he might, for strategic reasons, change 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 that allows for correlations among the bidders' valuations and includes as special cases both the common-value model and the independent-private-values model was developed by Milgrom and Weber (1982a). With n bidders, let x_i represent a private signal about the item's value observed by bidder i ; let $x = (x_1, \dots, x_n)$. Let $s = (s_1, \dots, s_m)$ be a vector of variables that 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 i th 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.⁶

A further choice to be made by the modeler depends on the answer to the question: 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 modeled as having different distributions F_i , $i = 1, \dots, n$?⁷ 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.

⁶ 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.

Yet another modeling 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. 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 that 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 V and VI. 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 VII through X.⁷

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, their risk attitudes, and the probability

⁷ Eric Maskin and Riley (1985) provided simple examples illustrating the effects of varying assumptions A1, A2, and A3.

distributions of valuations, and to know everyone else knows that he knows this, and so on.⁸ Based on what he knows, each bidder decides how high to bid. At a Bayes-Nash equilibrium, each bidder bids an amount that 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, regardless of which of the assumptions about risk attitudes and value correlations apply: *The Dutch auction yields the same outcome as the first-price sealed-bid auction* (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 analyze the Dutch auction in what follows.

V. The Benchmark Model: Comparing Auctions

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 highest bid, as in the first-price sealed-bid auction, must be better for the seller than

receiving 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.⁹

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, $v_{(1)}, \dots, v_{(n)}$. 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

⁸ This is an instance of a game of incomplete information (John Harsanyi 1967, 1968). A strong assumption underlying the solution of any game of incomplete information is that the agents' subjective beliefs about the unknown parameters be mutually consistent; for expositions, see James Friedman (1986, pp. 48–56) and Myerson (1985).

⁹ According to conventional wisdom, the excitement generated by the open calling of bids in an English auction can cause bidders to bid too high. Such an assertion of irrational behavior should be treated with caution: It 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.

can be proven using the properties of order statistics:¹⁰ *The expected difference between the first order statistic and the second order statistic is the expected value of $[1 - F(v_{(1)})]/f(v_{(1)})$.* (This expectation is taken with respect to the distribution of the first order statistic $v_{(1)}$ because the winner has valuation $v_{(1)}$.)

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 payment expected by the seller in an English auction is the expected value of $J(v_{(1)})$, defined by

$$J(v_{(1)}) = v_{(1)} - \frac{(1 - F(v_{(1)}))}{f(v_{(1)})}. \quad (1)$$

(The seller does not know the value of $v_{(1)}$, but the expected value of $J(v_{(1)})$ is of course not a function of $v_{(1)}$.) 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.¹¹

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, because 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 occurs 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. Because he would have earned non-negative 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 occurs 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, 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 guess (correct in equilibrium) of the decision rules being followed by the other bidders.

A Nash equilibrium for a first-price

¹⁰ A proof of this is given in McAfee and McMillan (1987f). 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.

¹¹ For J to be an increasing function, the distribution function F must be sufficiently concave. Also, J is increasing if and only if $[1 - F(v)]^{-1}$ is convex. When J is decreasing, it is in the seller's interest to choose the winner according to a random process; see Maskin and Riley (1980, 1983b), Myerson (1981), and Ralph Haywood (1984). For some examples of nonmonotonic $J(v)$, see Maskin and Riley (1984a, p. 183).

sealed-bid auction is found as follows. Consider the decision of bidder i , whose valuation is v_i . He conjectures that 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 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:

$$\pi_i = (v_i - b_i)[F(B^{-1}(b_i))]^{n-1}. \quad (2)$$

Thus he chooses b_i such that $\partial\pi_i/\partial b_i = 0$. By differentiating π_i with respect to v_i , we obtain $d\pi_i/dv_i = \partial\pi_i/\partial v_i + (\partial\pi_i/\partial b_i)(db_i/dv_i) = \partial\pi_i/\partial v_i$. (This is, of course, the Envelope Theorem.) Thus, by differentiating (2), an optimally chosen bid b_i must satisfy

$$\frac{d\pi_i}{dv_i} = \frac{\partial\pi_i}{\partial v_i} = [F[B^{-1}(b_i)]]^{n-1} \quad (3)$$

So far, we have examined bidder i 's best response to an arbitrary decision rule B being used by his rivals. Now impose the Nash (or rational-expectations) requirement: The rivals' use of the decision rule B must be consistent with the rivals' themselves acting rationally. Together with an assumption of symmetry (any two bidders with the same valuation will submit the same bid), this implies that bidder i 's optimal bid b_i , satisfying (3), must be the bid implied by the decision rule B —in other words, at a Nash equilibrium, $b_i = B(v_i)$. When we substitute this Nash condition into (3), we obtain an

equation defining bidder i 's expected surplus at a Nash equilibrium:

$$\frac{d\pi_i}{dv_i} = [F(v_i)]^{n-1}. \quad (4)$$

At a Nash equilibrium, all n bidders must be maximizing simultaneously, so that the condition (4) must hold for all bidders $i = 1, \dots, n$. We solve the differential equation (4) for π_i simply by integrating (using the boundary condition that if a bidder has the lowest possible valuation v_ℓ then he earns zero surplus, implying $B(v_\ell) = v_\ell$). Then we use the definition of π_i (equation (2)) plus the Nash condition $b_i = B(v_i)$ to obtain each bidder's decision rule:

$$B(v_i) = v_i - \frac{\int_{v_\ell}^{v_i} [F(\xi)]^{n-1} d\xi}{[F(v_i)]^{n-1}}, \quad i = 1, \dots, n. \quad (5)$$

Note that this bidding function B is increasing, as was assumed. The second term on the right-hand side of (5) shows how much the bidder shades his bid below his true valuation v_i .

Note the two crucial steps in the foregoing argument. First, we found one bidder's best response to a particular decision rule that he arbitrarily conjectures his rivals to be using. The first order condition from this optimization gave rise to equation (3). Second, we imposed the Nash requirement that the conjectured decision rules are themselves consistent with optimizing behavior by the other bidders. This turned equation (3) into equation (4), which we then solved to yield (5).¹²

For the special case where the distribution of valuations F is uniform and the lowest possible valuation is zero, substi-

¹² On existence, symmetry, and uniqueness of bidding equilibria, see Maskin and Riley (1986) and Milgrom and Weber (1985).

tution into (5) shows that a bidder with value v in a first-price sealed-bid auction submits a bid of $B(v) = (n - 1)v/n$; that is, he bids a fraction $(n - 1)/n$ of his valuation.

The winning bidder is the individual with the highest valuation $v_{(1)}$. In choosing his bid, each bidder assumes his valuation is the highest, $v_{(1)}$, because this presumption is costless if incorrect as losers pay nothing. It can be shown that $B(v_{(1)})$ as in (5) is equal to the expected second-highest valuation conditional on the bidder's information, which is that his own valuation is $v_{(1)}$. (In other words, $B(v_{(1)})$ is the expected value of the second order statistic conditional on the first order statistic being $v_{(1)}$.) The bidder estimates how far below his own valuation the next highest valuation is on average, and then submits a bid that 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 can be shown to be equal to the expected value of $J(v_{(1)})$, defined by (1). But, as we have seen already, the expected value of $J(v_{(1)})$ equals the expected price in an English or second-price auction. 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: *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 yields the same price on average* (Vickrey 1961; Armando Ortega-Reichert 1968; Charles Holt 1980; Milton Harris and Artur Raviv 1981a; Myerson 1981a; Riley and Samuelson 1981).

The Revenue-Equivalence Theorem does not imply that 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, $v_{(2)}$. In a first-price sealed-bid or Dutch auction, the price is the expectation of the second-highest valuation conditional on the winning bidder's own valuation, $B(v_{(1)})$. 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 a 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 (5). 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.¹³

¹³ The variance of revenue is lower in an English or second-price auction than in a first-price or Dutch auction (Vickrey 1961). Hence, if the seller were risk averse (instead of, as assumed throughout this paper, risk neutral), he would choose one of the former auction forms.

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 and the seller values the item less than any bidder). 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: *Increasing the number of bidders increases the revenue on average of the seller* (Holt 1979; Harris and Raviv 1981a). In particular, from equation (2), if $n = 1$, $B(v_i) = v_\ell$: Without the pressure of competition, the sole bidder will bid the lowest possible valuation, v_ℓ . As the number of bidders increases, the bids, as given by (5), increase.¹⁴ Provided the number of bidders is finite, the winning bidder pays a price less than his own valuation of the item, so that he earns positive economic rent. But if there is perfect competition among the bidders,¹⁵ then all of the gains from trade go to the seller. *As the number of bidders approaches infinity, the price tends to the highest possible valuation* (Holt 1979). This is because, as the number of bidders increases, the second-highest valuation approaches the highest possible valuation.

¹⁴ 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 (1987f).

¹⁵ Here and elsewhere we use the term *perfect competition* to describe the case of unboundedly many bidders. Although in practice the number of bidders is of course always finite, the infinite-bidders case can be taken to represent a situation in which there are so many bidders that adding one extra bidder would not noticeably lower the probability of any particular bidder's having the highest valuation.

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 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, for particular distributions of valuations, such as normal and uniform, one can show that an increase in the variance of valuations increases both the average revenue of the seller and the rents of the successful bidders (McAfee and McMillan 1986).

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).¹⁶ 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. The seller cannot exploit this information, because he obtains it only after having committed himself to the second-price mechanism.

VI. Optimal Auctions

The Revenue-Equivalence Theorem compares the expected revenues accruing from each of the commonly used auc-

¹⁶ This is closely related to the idea underlying the Groves mechanism for inducing revelation of preferences for public goods (Theodore Groves and John Ledyard 1977).

tion forms. Given that the monopolist has the power to choose any selling mechanism, a more fundamental question to ask is, What is the best of all possible selling mechanisms 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 that takes as inputs the 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. Assume the absence of collusion among the bidders. 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 1985.)

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 $B(v)$. 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. Because 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 modeler 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 that 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 or self-selection 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. (Because v_0 may be zero, we are not requiring that the seller necessarily has some use for the item.) Continue to assume that the J function defined in equation (1) is increasing. Applying the Revelation Principle can be shown to yield the following result: *For the benchmark model, the auction that maximizes the expected price has the following characteristics: (a) If $J(v_i) < v_0$ for all bidders' valuations v_i , then the seller refuses to sell the item; (b) otherwise he offers it to the bidder whose valuation v is highest at a price equal to $B(v)$* (Laffont and Maskin 1980; Harris and Raviv 1981a; Milgrom 1985; Myerson 1981; Riley and Samuelson 1981; Engelbrecht-Wiggans 1986; McAfee and McMillan 1987a).

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. Because $J(v) < v$, it is possible that the seller keeps the item despite the presence of some bidder with a valuation that 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.

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 opti-

mal-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 trade-off. The disadvantage of setting a reserve price, already noted, is that it is possible that the remaining bidder has a valuation that lies between the seller's valuation and the reserve price, $v_0 < v < r = J^{-1}(v_0)$. In this case, the monopolist loses the sale even though the bidder would have been willing to pay more than the good is worth to the seller. 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, 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:

$$\pi(r) = r[1 - F(r)] + v_0 F(r). \quad (6)$$

Maximization of π with respect to r yields $v_0 = J(r)$ and the second order condition that J is nondecreasing.¹⁷

¹⁷ A useful way of preserving the second-order-statistic intuition provided earlier is to note that the

Note the simplicity of the formula for the optimal reserve price; in particular, it is independent of the number of bidders. The reason for this is that the seller imposes the reserve price in order to capture some of the informational rents that would otherwise go to the winning bidder; and these rents are equal to $[1 - F(v)]/f(v)$, which 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.

Because, as was shown in the last section, the four common auction forms are essentially equivalent in the benchmark model, we can conclude as follows: *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 $J^{-1}(v_0)$ (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.

This completes the analysis of the benchmark model. In the next four sections we examine the effects of changing

the assumptions upon which the benchmark model is based.

VII. 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. (In this section we retain the other assumptions A1, A2, and A4.) 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 domestic and foreign firms, with systematic production-cost differences.¹⁸

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.

When bidders are asymmetric, the first-price sealed-bid auction yields a different price from the English auction: Revenue equivalence breaks down. While it remains the case that, within a class, higher-valuation individuals bid higher, this is in general not the case across classes. Bidders from different classes perceive themselves to be facing different degrees of bidding competition. If n_i represents the number of bidders

seller, setting his reserve price optimally, earns the second order statistic of $[J^{-1}(v_0), v_1, \dots, v_n]$, where $J^{-1}(v_0)$ is the valuation (if held by a bidder) that produces a payment to the seller of v_0 .

¹⁸ 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.

of type i , then a type 1 bidder faces $n_1 - 1$ type 1 bidders and n_2 type 2 bidders, whereas a type 2 bidder faces n_1 type 1 bidders and $n_2 - 1$ type 2 bidders. Thus, in a first-price sealed-bid auction, a type 1 individual's bidding function, computed analogously to equation (5) above, differs from a type 2 individual's bidding function. A type 1 bidder's estimate of the gap between his own valuation and the second-highest valuation differs from that of a type 2 bidder with the same valuation.

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

Because the bidder with the highest valuation does not necessarily win when the bidders are asymmetric, even with no reserve price, the first-price sealed-bid auction can, unlike the English auction, yield an inefficient outcome.

Neither of these simple auctions is optimal when the bidders are asymmetric. Analogous to the function J of Section V, define functions $J_k(v_i^k)$, where v_i^k represents the valuation of a bidder of type k , by

$$J_k(v_i^k) = v_i^k - \frac{[1 - F_k(v_i^k)]}{f_k(v_i^k)}, \quad k = 1, 2. \quad (7)$$

The following theorem is due to Myerson (1981): *In the auction that maximizes the expected selling price, the seller awards the item to the individual with the highest value of $J_k(v_i^k)$.* (In addition, the seller sets a different reserve price for each type of bidder, computed in exactly the same way as in Section VI.)

This theorem shows that, when bidders are asymmetric, the optimal auction is discriminatory, in the sense that there is a possibility that one bidder wins despite another bidder's having a higher valuation. This is because asymmetry means that $F_1 \neq F_2$ so that $J_1(v) \neq J_2(v)$. Thus it is possible that, say, $J_1(v_{(1)}^1) > J_2(v_{(1)}^2)$ even though $v_{(1)}^2 > v_{(1)}^1$, so that the winner, with the highest $J_k(v_i^k)$, is not necessarily the bidder who values the item the most. This is analogous to second-degree price discrimination in the elementary-textbook monopoly model in that it involves discriminating across bidders with different demands.

Which type of bidder receives 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. *If 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* (McAfee and McMillan 1987d). There is a trade-off. 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 and receiving a relatively low payment. 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, driving up the price on average.¹⁹

Because 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

¹⁹ 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.

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.

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. Of course, 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-devel-

opment 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 1986).

VIII. *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 maintain the assumptions A1, A2, and A3 but relax the assumption A4 that payment can be a function only of bids. We 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 (Ramsey 1980). 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; McAfee and McMillan 1987e). Incentive contracts are also used in the private sector when a firm procures inputs from another firm.

All of these examples have the following properties. The seller observes ex post some variable $\bar{v}$ that 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:

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

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 under-run ($c - b$) 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 (Walter Mead, Asbjorn Moseidjord, and Philip Sorensen 1984). In what follows, we shall discuss bidding mechanisms of the first type. (See Robert Hansen 1985a, Douglas Reece 1979, and Riley 1986 for analyses of bidding mecha-

nisms of the second type, and Samuelson 1983, 1986b for an analysis of bidding mechanisms of the third type.)

What is the reason for using royalties? *If the distribution of the observed variable $\bar{v}$ is exogenous, the seller's expected revenue is an increasing function of the royalty rate* (McAfee and McMillan 1986, Riley 1986). 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 V), 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.

If expected revenue monotonically increases with the royalty rate, why are royalties not always set at 100 percent? One answer is that the distribution of $\bar{v}$ is not exogenous: 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 afterward. 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 $\bar{v}$ 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 trade-off. 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* (Engelbrecht-Wiggans 1985; Laffont and Jean Tirole 1985; McAfee and McMillan 1986, 1987c; Michael Riordan and David Sappington 1987). 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. This is because, when there are enough bidders that perfect competition prevails, there is no need to use royalties to stimulate bidding competition. Then 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).

When the simplicity of the linear payment function (8) means that it is commonly used in practice, in general when there is moral hazard a nonlinear payment function 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* (Laffont and Tirole 1985; McAfee and McMillan 1987c).

Another reason for not setting the royalty rate at 100 percent comes from the seller's inability to observe bidders' valuation ex ante. For simplicity, assume away moral hazard by assuming perfect ex post observability, so that $\tilde{v} = v$. The

successful bidder's rent is the difference between his valuation v and his payment (8). When the royalty r is 100 percent, this difference is the constant b ; thus the successful bidder's rent is independent of his valuation. This breaks the link between valuation and bids, so that now the highest bidder is not necessarily the individual who values the item the most. The bids fail to reveal relative valuations (McAfee and McMillan 1986; Samuelson 1986a).

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. *The more risk averse are the bidders relative to the seller, the higher is the optimal royalty rate* (Hayne Leland 1978; McAfee and McMillan 1986; Samuelson 1983, 1986b).

IX. Risk-Averse Bidders

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 *with risk-averse bidders, the first-price sealed-bid auction produces a larger expected revenue than the English or second-price auction* (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 bidders' risk aversion works to the seller's advantage.

The first-price sealed-bid auction is not the optimal auction, however; it fails to 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. *The seller will not make losing bidders' payments random* (Maskin and Riley 1984b; Steven Matthews 1983). The reason is straightforward. Instead of requiring a risky payment from a loser, the seller could require the 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, because winning is like receiving an increase in income, which will in general change the degree of risk aversion. The argument does, however, carry over when a bidder's risk aversion does not vary with his income: *If the bidders have constant absolute 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-neutral 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, 1984b; 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 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, requiring payments from some losing bidders and subsidies of others, it is unlikely to arise in practice. However, if the risk aversion is not very strong, the optimal auction is approximated by a sealed-bid auction with a bidding fee that is a decreasing function of the bid (Matthews 1983). 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 that 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 identical priors, any bidder always expects there to be more bidders than the seller. *Under constant or decreasing absolute risk aversion, in a first-price sealed-bid auction the expected sell-*

ing price is strictly higher when the bidders do not know how many other bidders there are than when they do know this (Matthews 1987; McAfee and McMillan 1987a). (Here expectations are taken over the number of bidders as well as their valuations.) 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: Concealing the number of bidders has the effect of making the bidding more competitive.²⁰

X. *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 maintain assumptions A1, A3, and A4. 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

²⁰ The effects of bidders' having different degrees of risk aversion were investigated by James Cox, Vernon Smith, and James Walker (1982), who derived equilibrium bidding functions for some particular auctions for the case of constant relative risk aversion.

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 that 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 Isaac (1984) pursued this straw man.) 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. *Suppose the i th bidder's information about the item's true value v can be represented by a number x_i , such that a bigger value of x_i implies a bigger true value v . Then*

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

(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 downward 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 (James Smith 1981). 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 (Milgrom and Weber 1982a).

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, even though 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. *If information is sufficiently dispersed among the bidders then the selling price converges to the item's true value v as the number of bidders becomes arbitrarily*

large (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.²¹

Consider now the more general model, due to Milgrom and Weber (1982a), that allows correlations among the bidders' valuations, and of which the common-value model is a special case. Recall from Section IV 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: *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 (Milgrom and Weber 1982a). In addition, the other three auction forms can be ranked. *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 revenue as the Dutch auction* (Milgrom and Weber 1982a).

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? *The seller can increase his expected revenue by having a policy of publicizing any information he has about the item's true value* (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? *If one bidder's information is available to another bidder, his expected surplus is zero* (Engelbrecht-Wiggans, Milgrom, and Weber 1983; Milgrom 1981b; 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 precise.

It is in the seller's interest in a common-value auction to impose a reserve price that is strictly above his own valuation of the item. In contrast to the independent-private-values case, *the reserve price in a common-value auction varies with the type of auction and with the number of bidders*; usually, but not always, it increases with the number of bidders (Milgrom and Weber 1982a; Marc Robinson 1984). This is because,

²¹ 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 (1984a), Milgrom (1981b, 1985), and Milgrom and Weber (1982b).

as implied by the winner's curse effects embodied in condition (9), any bidder's ex ante valuation of winning depends upon the auction form and the number of bidders. For example, the working of the English auction provides some information about others' valuations, mitigating winner's curse effects. The more bidders there are, the more heavily the bidder, following (9), discounts his private information.

What is the optimal auction when bidders' valuations are correlated? Consider the special case in which bidders have private valuations that are not statistically independent of each other. Jacques Crémer 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 type 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 i th 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 second-price auction. The trick is to design the lottery so that the lottery's expected value for the i th 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 second-price auction, because 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 Crémer-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.²²

XI. Further Topics

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

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 produces under constant returns to scale and can freely vary the amount offered for sale. 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 is drawn independently from a distribution F , with density f . Denote by c the constant average production cost. *With unlimited capacity, constant returns to scale, and independent valuations, the monopolist's optimal selling mechanism is to post a fixed price r defined by*

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

(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

²² Because the second-price auction is a dominant-strategy auction, the Crémer-McLean auction/lottery also has a dominant equilibrium (although participation is not necessarily a dominant strategy). With slightly weaker assumptions on π , a Bayes-Nash equilibrium auction that extracts all of the surplus can be constructed. Note also that the optimal-auction result of Myerson (1981) allows a limited kind of affiliation among valuations.

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 VI).

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 (Pickford 1985).

Two kinds of sealed-bid auctions are used to sell multiple units. Bidders submit bids that 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).²³

²³ For further analysis of auctions of multiple units of a commodity, see Engelbrecht-Wiggans and Weber (1979), Harris and Raviv (1981a, 1981b), Donald Hausch (1986), Maskin and Riley (1983b, 1984a), Matthews (1983, forthcoming), McAfee, McMillan and Whinston (1987), Milgrom (1985), Thomas Palfrey (1980, 1983), Vernon Smith (1967), Weber (1983), and Wilson (1979).

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 Robert Aumann 1981; in an incomplete-information setting, see Dilip Abreu, David Pearce, and Ennio Stacchetti 1986). 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 1977, pp. 37–38).²⁴

For retaliatory strategies to be workable, it must be possible for the bidders to infer the other bidders' past actions. As George Stigler (1964, p. 48) remarked

²⁴ Milgrom (1986) showed that there is a sense in which repeated English auctions are more susceptible to collusion than repeated first-price sealed-bid auctions. 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.

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."²⁵ This may explain the tendency for private-sector firms to use closed negotiations rather than formal auctions for procurement.

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. Afterward, the item is reaucted 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. 228–30), 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. Assume that (perhaps by the reaucting

process described above) the cartel can ensure that the member who values the item the most ultimately gets it. 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 the same as that in Section VI but for this density yields: *The optimal anticartel reserve price r satisfies*

$$v_0 = r - \frac{1 - F^n(r)}{nF^{n-1}(r)f(r)}, \quad (11)$$

where v_0 is the seller's own valuation (Daniel Graham and Robert Marshall 1984, 1985; Robinson 1984). This implies that the anticartel reserve price increases with the number of cartel members. Also, the anticartel reserve price is higher than the optimal reserve price in the absence of collusion.

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 sealed-bid double auction, both sellers and buyers submit bids, buyers for how much they are willing to pay and sellers for the price at which they are willing to sell. Ranking the buyers' bids from highest to lowest produces a bid-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. This resembles the

²⁵ On the various methods by which bidders for government contracts achieve collusion, see Organization for Economic Cooperation and Development (1976).

standard Walrasian model in which agents report their demand and supply curves to the auctioneer. However, the choice of bids reflects individuals' strategic attempts to manipulate the selling price, so that the quantity and price interval reached are not necessarily those of the competitive equilibrium.

Few results on the double auction exist, because of the difficulties of modeling strategic behavior on both sides of the market. The main result is an efficiency result concerning double auctions with sealed bids. *If there are sufficiently many buyers and sellers, then there is no other trading mechanism that would increase some traders' expected gains from trade without lowering other traders' expected gains from trade* (Wilson 1985). That is, the sealed-bid double auction would survive any attempt to get the traders to agree unanimously to change the trading institution.²⁶

The oral double auction, with the bids and offers openly called, is still more difficult to model because the process takes place over time and agents do not know what prices will be available if they wait instead of trading now. For some analysis of the oral double auction, see David Easley and Ledyard (1982), Daniel Friedman (1984), and Wilson (1986b).

XII. Empirical Studies of Auctions

What is the empirical content of auction theory? There are far fewer empirical than theoretical studies of auctions. Also, with the exception of some studies of common-value auctions (using mineral-rights bidding data), the existing empirical studies make little use of modern auction theory. As a result (again, except

in the mineral-rights studies), only the more commonsense propositions have been tested.

The independent-private-values model predicts revenue equivalence: Each of the four simple auctions yields the same price on average (Section V above). Unfortunately, this prediction is sensitive to changes in the underlying assumptions, making difficult the task of empirically testing it. If the bidders have independent private valuations but are risk averse, then the first-price sealed-bid auction yields a higher price on average than any of the other simple auctions (Section IX). This ranking is reversed if the bidders are risk neutral but have correlated valuations; then the English auction yields the highest price (Section X). If there are observable differences among the bidders' valuations, the ranking is indeterminate (Section VII). Thus any empirical test of revenue equivalence should carefully control for risk aversion of the bidders, correlations among their valuations, and asymmetries among the bidders.

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 that 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 risk aversion makes the sealed-bid auction preferable to the English auction. On the other hand, artwork is usually, but not always, sold by English auction, even though 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

²⁶ Also, Wilson (1986a) showed that, for the case of equal numbers of buyers and sellers with valuations distributed uniformly, the double auction satisfies the stronger criterion of ex ante efficiency: It maximizes the expected gains from trade.

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 auctions and English auctions to sell contracts for harvesting timber. This provides an opportunity for testing the Revenue-Equivalence Theorem: Do the two methods yield the same price on average? Walter Mead (1967) ran ordinary-least-squares regressions with, as the dependent variable, the logarithm of the ratio of bid price to appraised price and, as the independent variables, a dummy variable for the auction method and the logarithm of the quantity of timber for sale. The conclusion was that the sealed-bid auctions had yielded significantly higher revenue than the English auctions. Further unpublished regressions by Mead and coauthors (summarized by Hansen 1985b) supported this conclusion, finding that the sealed-bid auctions yielded about 10 percent more revenue. These results have been questioned by Hansen (1985b, 1986), who noted a selection bias caused by the way the Forest Service chose which auction to use.²⁷ (Evidently the Forest Service does not believe in revenue equivalence.) Some unobserved variables, such as the bidders' timber inventories, affected the choice of auction method. But to the extent that these variables also affected the size of the bids, ordinary-least-squares estimates are biased. After correcting for this selection bias by using a simultaneous-equations model, Hansen found that, although the sealed-bid auctions yielded slightly higher prices than the English auctions, the difference was statistically insignificant. (With high bids averaging \$130, the expected difference in revenue was be-

tween \$1 and \$6, with a standard error of \$5.) Revenue equivalence therefore cannot be rejected.

A question not addressed in these timber-bidding studies is, How appropriate is the independent-private-values assumption in this application? In other words, on a priori grounds, should we expect revenue equivalence to hold in this set of data? Because different bidders have different inventories of timber on hand and different stocks of equipment for harvesting timber, it is reasonable to suppose that there are inherent differences among the bidders' valuations at any particular sale. But because the timber is eventually to be resold at the as yet unknown market price, there might be some correlation among the bidders' valuations. To the extent that the second effect is empirically significant, the independent-private-values assumption breaks down and the appropriate model is the affiliation model of Milgrom and Weber (1982a) (see Section IV). We are now left with a puzzle: Contrary to the empirical results just cited, this model predicts that the English auction will produce the higher bids (Section X). The puzzle could be resolved by appealing to risk aversion of the bidders, but this remains an open empirical question. Given the sensitivity of the Revenue-Equivalence Theorem to its underlying assumptions, this theorem cannot be meaningfully tested until some way is found to test for independent private values against affiliated values.

A further reason, distinct from correlation of valuations, why revenue equivalence can break down is the existence of systematic differences among bidders' valuations that are observable by all of the bidders (as discussed in Section VII). For timber-rights bidding, Ronald Johnson (1979) and Hansen (1986) suggested that such variables as the required amount of road building or the required

²⁷ For the institutional details of the auction-selected process, see Hansen (1986, pp. 129–31).

harvest rate might cause harvesting costs and therefore valuations to differ in observable ways among the bidders. Johnson separated the data into auctions with symmetric bidders and auctions with asymmetric bidders by an ad hoc method: He hypothesized that asymmetry is more likely the larger the amount of road construction required, and took as his cutoff between symmetry and asymmetry sales involving \$20,000 worth of road building. He ran ordinary-least-squares regressions with, as dependent variable, the price, and as independent variables, a dummy representing the auction form as well as all of the components of the Forest Service's appraisal. For the symmetric auctions, consistent with the Revenue-Equivalence Theorem, there was no significant difference in price between the English auctions and the sealed-bid auctions. For the asymmetric case, there was a significant difference: The sealed-bid auctions yielded higher prices. Contrary to Johnson's interpretation, however, this does not test the theory, because in the asymmetric case there is only an ambiguous result: Either auction form can yield the higher price (see Section VII).²⁸

Another theoretical result on asymmetric auctions is that, in the absence of a reserve price, an English auction always yields an efficient outcome (the bidder with the highest valuation wins), whereas a sealed-bid auction may yield an inefficient outcome (Section VII). When the seller is the government (as in the timber-rights auctions), this provides a reason for using English auctions that is distinct from issues of revenue

raising. Johnson (1979) tested this proposition by noting that it implies that, if resale is allowed, it should occur only after a sealed-bid auction and not after an English auction. In Johnson's data, with a total of 379 sales, all seven of the resales that occurred followed sealed-bid auctions, corroborating the theory.

In value terms, one of the world's largest users of auctions is the U.S. Treasury in its selling of securities. This is an instance of a multiple-unit auction: Both the discriminatory auction (in which each successful bidder pays the amount he bid) and the uniform-price auction (in which all pay a price equal to the lowest accepted bid) have been used to sell Treasury bills. Because this is a common-value setting, theory predicts that the uniform-price auction, which is similar to the second-price auction, yields more revenue than the discriminatory auction, which corresponds to the first-price auction (Sections X and XI.1). Charles Baker (1976) summarized some Treasury studies that attempted to compare the performance of the two auction forms. Unfortunately the data are not strictly comparable: Uniform-price auctions were used to sell long-term bonds, whereas discriminatory auctions were used to sell short-term and medium-term bonds. Because of this, the empirical conclusions can only be tentative; however, they do support the theory, in that the uniform-price auctions seemed to generate the higher revenue.

An unsurprising prediction of auction theory is that, other things being equal, a bidder with a higher valuation will submit a higher bid (or, in the case of contract bidding, a firm with lower costs submits a lower bid). In their study of bidding for construction contracts for San Francisco's Bay Area Rapid Transit system, Kenneth Gaver and Jerold Zimmerman (1977) estimated a model with, as the dependent variable, the level of each

²⁸ Maskin and Riley (1985) conjectured that, in an asymmetric auction, a sealed-bid auction yields a higher price than an English auction if the bidders have distributions of valuations with roughly the same shape but different supports. While it remains to be shown that this is valid in general, it provides a potential explanation for Johnson's finding.

firm's bid for each contract, and, among the independent variables, various proxies for costs as well as the number of bidders. They found that bids were higher (a) the more BART work the bidder had in process; (b) the less time the contractor had to complete the project; and (c) the smaller the firm. Of these three variables, (a) reflects the bidder's opportunity cost, and (b) and (c) affect production costs if it is presumed respectively that faster production is more costly and that there are economies of scale.

Another prediction is that competition matters: The winning bid increases as the number of bidders increases (or, in the case of contract bidding, the winning bid falls). In their study of contract bidding, Gaver and Zimmermann (1977) found that the price fell as the number of bidders increased. One of their regressions, for example, implied that bids decline by about 2 percent when the number of bidders increases by one. Similarly, Lance Brannman, Douglas Klein, and Leonard Weiss (1984), using data from auctions of tax-exempt bonds, government-owned timber, and offshore oil rights, and attempting to control for the inherent value of the item being sold, found the effect of the number of bidders on the price to be statistically significant. Large gains can be obtained by 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. A cartel reduces the effective number of bidders: In a study of collusion among bidders for North Carolina highway-construction contracts, Jonathan Feinstein, Michael Block, and Frederick Nold (1985) found that the cartel not only significantly raised the bids but also reduced the variance of the bids.

While there appear to be no econometric analyses of the use of reserve prices (Section VI), there is some informal evidence. Practice does not seem to be in accord with the theoretical result that it is in a 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. 226–27). Thus there appears to be a discrepancy between theory and practice.²⁹ Consistent with the prediction that the reserve price will be set at a level that 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).

Studies of the bidding for offshore oil rights have focused on the existence and consequences of the winner's curse (Section X).³⁰ The first of these studies, by Capen, Clapp, and Campbell (1971), claimed that winning bidders earned on average less than the market rate of return on their investments. They found surprisingly large ranges of bids: The ratio of highest to lowest bid was as high as 100 and typically between 5 and 10, indicating much uncertainty about the amount of oil beneath any tract. They attributed the low rates of return to the fact that "the winner tends to be the

²⁹This need not, however, be inconsistent with theory. One explanation for the seller's not imposing a reserve price is that he lacks the necessary bargaining power, that is, the ability to precommit himself to his selling policy (see Milgrom 1986).

³⁰The auctioning of mineral rights to government-owned land is almost unique to the United States; governments in other countries usually allocate mineral rights to firms by discretionary procedures (Ramsey 1980, pp. 56–62). It is noteworthy, however, that in 1982 the People's Republic of China began auctioning offshore exploration contracts (*New York Times*, March 23, 1985).

player who most overestimates the true tract value" (p. 643). The conclusion of Capen, Clapp, and Campbell about low rates of return was based on casual empiricism. Subsequent, more careful studies have overturned this conclusion. Mead, Moseidjord, and Sorensen (1984) estimated after-tax internal rates of return of outer-continental shelf lessees to be roughly the same as the average rate of return on investments in the U.S. manufacturing industry. Kenneth Hendricks, Robert Porter, and Paul Gertler (1986) similarly found reasonable rates of return. Hendricks, Porter, and Gertler also discussed the practical importance of two aspects neglected by the theorists. Often, several tracts are auctioned simultaneously, so that a bidder's decision on how to bid on one tract may be influenced by his bids on other tracts. Moreover, as also emphasized by Otis Gilley and Gordon Karels (1981), there is a prior decision on whether to bid on a particular tract; only some of the potential bidders actually bid for any one tract. There tend to be more bidders for tracts that are perceived to be more valuable. Thus estimating the effect of competition by simply regressing price on the number of bidders (as discussed above) can be misleading.

Gilley and Karels (1981), after correcting for the sample-selection bias arising from the bidders' participation decisions, found that individual bids in an oil-rights auction decreased as the numbers of bidders increased. This apparently paradoxical result can be explained as resulting from the bidders' rationally taking account of the winner's curse effect of condition (9) above. Recall that a rational bidder in a common-value auction bases his bid not only on his own estimate of the item's worth, but also on the presumption that he has the highest estimate. The more bidders there are, the more marked is the second of these ef-

fects. This can mean that individual bids decline as the number of bidders rises. Hendricks, Porter, and Gertler (1986), generalizing the Gilley-Karels model by allowing nonlinear bidding functions, found that bids initially increased and then decreased as the number of bidders increased. (Note that these results are not inconsistent with the result that the winning bid increases with the number of bidders, for the more bidders there are, the more likely it is that some bidder perceives the item's value to be high.) More accurate information about the item's true value mitigates the effect of the winner's curse in causing the bidders to be cautious: Gilley and Karels found that the smaller the variance in the initial estimates of the tract's value, the higher the bids. With millions of dollars at stake, the oil firms evidently avoid falling victim to the winner's curse.

The results of Milgrom and Weber (1982b) (Section X above) create a presumption that having private information raises a bidder's profits. A test of this prediction is provided by the study of oil-rights bidding by Mead, Moseidjord, and Sorenson (1984) (although the study is not explicitly based on the Milgrom-Weber model). Two types of leases exist: drainage tracts (close to an existing well) and wildcat tracts (in an area for which no drilling data exist). For drainage tracts, denote by "neighbor" a bidder who has information because he has already drilled nearby wells. Mead, Moseidjord, and Sorenson estimated that there were gains from superior information. Drainage leases earned a higher after-tax internal rate of return than wildcat leases, and, among drainage leases, neighbors earned a higher rate of return than non-neighbors. Similar results were obtained by Hendricks, Porter, and Bryan Boudreau (1987). Theory suggests that the gains from superior information are lessened if a rival also has superior

information (Milgrom and Weber 1982a). Mead, Moseidjord, and Sorenson, however, found an unusually high rate of return when a neighbor competed with a neighbor (although a very small sample underlay this estimate).³¹

More empirical work is needed, both because several testable predictions remain untested and because usable sets of microlevel data have not yet been exploited. To what extent do royalties serve to extract surplus from the bidders, as discussed in Section VIII? Mineral-rights auctions, which sometimes use royalties and sometimes do not, might provide data for estimating this. Mineral-rights data might also be used to address the question: When reserve prices are used, are they set at a level that approximates the optimum (Sections VI, X)? Government contracting provides an underexploited source of data on bidding; questions such as the effect of incentive contracts as a rent-extracting device (Section VIII) and the effects on the bids of discriminating against foreign bidders (Section VII) could be examined. What types of auction are actually used when the bidders are markedly risk averse (Section IX)? In privatizing publicly owned firms, governments sometimes have auctioned the whole firm and sometimes have sold shares at preset prices. Which method has yielded more revenue for the government? Different auction forms are used in apparently similar circumstances; for example, at various times and in various places, tobacco has been sold by English, Dutch, and first-price

sealed-bid auctions. Can each use be explained, or is this merely due to revenue equivalence, which means that the seller is indifferent about the choice of auction form?

Another source of empirical evidence on bidding behavior is laboratory experiments. There is a large and growing literature on experiments with auctions, which is not 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; and, for representative papers, see Cox, Smith and Walker (1984) (on independent-private-value auctions) and Max Bazerman and Samuelson (1983) (on common-value auctions).

XIII. *An Agenda for Research*

Many questions about the working of the various auction institutions remain unanswered by the theory.

Much of the existing theory assumes independent private values. While this makes for elegant theorems, many real-world auctions fail to satisfy this assumption (as Milgrom and Weber 1982a persuasively argued). Independent private values requires, for example, that there be no resale possibilities, and that each bidder knows exactly how much the item would be worth to him. In the case of contract bidding, independent private values means that firms' production costs differ only because of differences in their production capabilities or in their alternative opportunities; however, if there is a common element of technological uncertainty, then the appropriate assumption is affiliation (Section IX). Despite the prevalence in practice of affiliation of valuations, little is known beyond the analysis of Crémer and McLean (1985a, 1985b) about optimal auctions with affiliated values.

³¹ There is a stylized fact from the sealed bidding for mineral rights: In common-value auctions, the distribution of bids is approximately lognormal (Chester Peltó 1971; Reece 1978). Albert Smiley (1979), however, found that the Weibull distribution fitted bidding data better than the lognormal distribution. The shape of the distribution of bids is of interest for simulating bidding behavior; also, under the assumption that the theory is correct, the distribution of valuations can be deduced from the distribution of bids.

In all of the existing models, the selling is in effect done by the owner. In practice, often the owner hires an agent, the auctioneer. How can the auctioneer be motivated to act in the owner's interest? Should he be paid a fixed fraction of the selling price? Might the buyer pay the auctioneer in some circumstances? Does the presence of a hired auctioneer affect the choice of auction form?

A question prompted by the empirical studies (Gilley and Karels 1981; Hendricks, Porter, and Gertler 1986) is, What determines the number of bidders? Although most existing models take the set of bidders to be exogenous, the bidders do in fact make a decision to learn their valuations. What are the determinants of this decision? The costs of collecting information and preparing a bid presumably are relevant here (compare with French and McCormick 1984, McAfee and McMillan 1987b, Samuelson 1985). How does the incorporation of this decision affect the performance of the different auction forms? How does the choice of auction form influence this decision?

In government contracting, bids are often multidimensional, involving quality as well as price. For example, in the competition for the contract for a new fighter aircraft, design features are at least as important as price. If the government has its own preferences over the various quality attributes and price, and the different firms have different technological trade-offs, what is the best procurement mechanism?³²

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. One might conjecture that

commitment is a natural concomitant of monopoly: A rational monopolist will find some way to achieve commitment. However, it is not well understood how this can be done. 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.³³

The opposite question should also be addressed. What happens if we drop the assumption that the seller is able to commit himself? The commitment assumption makes the theorist's life easier by providing a clear-cut solution in place of the indeterminacy of the bargaining problem: Apart from the rents due to the informational asymmetries, all of the gains from trade go to the individual who is able to precommit himself. Without the commitment assumption, the bargaining problem must be addressed. In an interesting analysis, Milgrom (1986) developed a bargaining model in which a seller finds it in his interest to sell by auction. This is an important area for further research.

Turn now to the broader questions posed by Hayek (1945) and discussed in Section I. To what extent has auction theory generated answers to questions 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 (Arrow 1985). The foregoing theorems

³² This question was posed by Paul Milgrom and by William Samuelson. The second-stage decision in Riordan and Sappington (1987) can be interpreted as a quality decision.

³³ 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.

show 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 (given in Sections VI and VII), 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. As Hayek suggested, markets can work despite the dispersion of information.

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 VI and VII.) Do prices serve to aggregate dispersed information? The result of Wilson (1977) and Milgrom (1979a) (summarized in Section X) 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: If competition is less than perfect 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 VIII).

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. Ques-

tions of the existence and social optimality of competitive equilibrium with informational asymmetries await the resolution of this question.

One crucial step toward a general-equilibrium formulation is modeling competition among mechanism designers. Consider a seller who designs an auction-like mechanism. If the buyers have an alternative, that is, if 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 modeling competition among mechanisms is, in our view, the major problem facing the asymmetric-information literature.

XIV. *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 val-

uation (provided this exceeds the lowest possible valuation). On the other hand, if 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 are no more reluctant to bear risk than you are 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 some-

how 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 having a policy of 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.

REFERENCES

- ABREU, DILIP; PEARCE, DAVID AND STACCHIETTI, ENRICO. "Optimal Cartel Equilibria with Imperfect Monitoring," *J. Econ. Theory*, June 1986, 39(1), pp. 251-69.
- AKERLOF, GEORGE. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quart. J. Econ.*, Aug. 1970, 84(3), pp. 488-500.
- ARROW, KENNETH J. "Toward a Theory of Price Adjustment," in *The allocation of economic resources*. Stanford: Stanford U. Press, 1959, pp. 41-51.
- . "The Economics of Agency," in *Principals and agents: The structure of business*. Eds.: JOHN W. PRATT AND RICHARD J. ZECKHAUSER. Boston: Harvard Business School Press, 1985, pp. 37-51.
- AUMANN, ROBERT J. "Survey of Repeated Games," in *Essays in game theory and mathematical economics in honor of Oskar Morgenstern*. Eds.:

- VOLKER BÖHM AND H.-H. NACHTKAMP. Mannheim: Bibliographisches Institut, 1981.
- BAKER, CHARLES C. "Auctioning Coupon-Bearing Securities: A Review of Treasury Experience," in *Bidding and auctioning for procurement and allocation*. Ed.: YAKOV AMIHUD. NY: New York U. Press, 1976, pp. 146-51.
- BAZERMAN, MAX AND SAMUELSON, WILLIAM. "The Winner's Curse: An Empirical Investigation," in *Aspiration levels in bargaining and economic decision making*. Ed.: REINHARD TEITZ. Berlin: Springer-Verlag, 1983, pp. 186-200.
- BRANNMAN, LANCE; KLEIN, J. DOUGLAS AND WEISS, LEONARD. "Concentration and Winning Bids in Auctions," *Antitrust Bull.*, Spring 1984, 29(1), pp. 27-31.
- CAPEN, E. C.; CLAPP, R. V. AND CAMPBELL, W. M. "Competitive Bidding in High-Risk Situations," *J. Petroleum Tech.*, June 1971, 23(6), pp. 641-53.
- CASSADY, RALPH, JR. *Auctions and auctioneering*. Berkeley: U. of Calif. Press, 1967.
- CASSING, JAMES AND DOUGLAS, RICHARD W. "Implications of the Auction Mechanism in Baseball's Free Agent Draft," *Southern Econ. J.*, July 1980, 47(1), pp. 110-21.
- 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," *Econ. Inquiry*, Oct. 1984, 22(4), pp. 579-92.
- COX, JAMES C.; SMITH, VERNON L. AND WALKER, JAMES M. "Auction Market Theory of Heterogeneous Bidders," *Econ. Letters*, 1982, 9(3), pp. 319-25.
- _____. "Theory and Behavior of Multiple Unit Discriminative Auctions," *J. Finance*, Sept. 1984, 39(4) pp. 983-1010.
- CRÉMER, JACQUES AND MCLEAN, RICHARD P. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." CARESS Working Paper No. 85-17. U. of Pa., June 1985a.
- _____. "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, Mar. 1985b, 53(2), pp. 345-61.
- DEMAYO, PETER. "Bidding on New Ship Construction," in R. ENGELBRECHT-WIGGANS ET AL., 1983, pp. 371-87.
- DEMSETZ, HAROLD. "Why Regulate Utilities?" *J. Law Econ.*, Apr. 1968, 11(1), pp. 55-66.
- DESSAUER, JOHN P. *Book publishing*. NY: Bowker, 1981.
- EASLEY, DAVID AND LEDYARD, JOHN. "A Theory of Price Formation and Exchange in Oral Auctions." Tech. Report No. 461, Northwestern U., Mar. 1982.
- ENGELBRECHT-WIGGANS, RICHARD. "Auctions and Bidding Models: A Survey," *Manage. Sci.*, Feb. 1980, 26(2), pp. 119-42.
- _____. "Optimal Competitive Contracting." Working Paper No. 1200, College of Commerce and Business Administration, U. of Ill. at Urbana-Champaign, Oct. 1985.
- _____. "Classes of Equivalent Auction Mechanisms: The Case of Risk Neutral Bidders." Working Paper No. 1243, College of Commerce and Business Administration, U. of Ill. at Urbana-Champaign, Apr. 1986.
- ENGELBRECHT-WIGGANS, RICHARD; MILGROM, PAUL R. AND WEBER, ROBERT J. "Competitive Bidding and Proprietary Information," *J. Math. Econ.*, Apr. 1983, 11(2), pp. 161-69.
- ENGELBRECHT-WIGGANS, RICHARD; SHUBIK, MARTIN AND STARK, ROBERT M., eds. *Auctions, bidding and contracting: Uses and theory*. NY: NYU Press, 1983.
- ENGELBRECHT-WIGGANS, RICHARD AND WEBER, ROBERT J. "An Example of a Multi-Object Auction Game," *Manage. Sci.*, Dec. 1979, 25(12), pp. 1272-77.
- FEINSTEIN, JONATHAN S.; BLOCK, MICHAEL K. AND NOLD, FREDERICK C. "Asymmetric Information and Collusive Behavior in Auction Markets," *Amer. Econ. Rev.*, June 1985, 75(3), pp. 441-60.
- FRENCH, KENNETH R. AND MCCORMICK, ROBERT E. "Sealed Bids, Sunk Costs, and the Process of Competition," *J. Bus.*, Oct. 1984, 57(4), pp. 417-41.
- FRIEDMAN, DANIEL. "On the Efficiency of Experimental Double Auction Markets," *Amer. Econ. Rev.*, Mar. 1984, 74(1), pp. 60-72.
- FRIEDMAN, JAMES W. *Game theory with applications to economics*. NY: Oxford U. Press, 1986.
- GAVER, KENNETH M. AND ZIMMERMAN, JEROLD L. "An Analysis of Competitive Bidding on BART Contracts," *J. Bus.*, July 1977, 50(3), pp. 279-95.
- GILLEY, OTIS W. AND KARELS, GORDON V. "The Competitive Effect in Bonus Bidding: New Evidence," *Bell J. Econ.*, Autumn 1981, 12(2), pp. 637-48.
- GRAHAM, DANIEL A. AND MARSHALL, ROBERT C. "Bidder Coalitions at Auctions." Mimeo. Duke U., Nov. 1984.
- _____. "Collusive Behaviour at a Single Object English Auction." Working Paper 85-01. Duke U., Feb. 1985.
- GRIESMER, J. H.; LEVITAN, R. E. AND SHUBIK, M. "Towards a Study of Bidding Processes, Part Four: Games with Unknown Costs," *Naval Res. Logistics Quart.*, Dec. 1967, 14(4), pp. 415-33.
- GROVES, THEODORE AND LEDYARD, JOHN. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem," *Econometrica*, May 1977, 45(4), pp. 783-809.
- GRUEN, FRED. "Coulburn, Forward Prices, and Pies," *Rev. Marketing Agr. Econ.*, June 1960, 28(1), pp. 85-96.
- HANSEN, ROBERT G. "Auctions with Contingent Payments," *Amer. Econ. Rev.*, Sept. 1985a, 75(4), pp. 862-65.
- _____. "Empirical Testing of Auction Theory," *Amer. Econ. Rev.*, May 1985b, 75(2), pp. 156-59.
- _____. "Sealed-Bid Versus Open Auctions: The Evidence," *Econ. Inquiry*, Jan. 1986, 24(1), pp. 125-42.
- HARRIS, MILTON AND RAVIV, ARTUR. "Allocation Mechanisms and the Design of Auctions," *Econometrica*, Nov. 1981a, 49(6), pp. 1477-99.
- _____. "A Theory of Monopoly Pricing Schemes

- with Demand Uncertainty," *Amer. Econ. Rev.*, June 1981b, 71(3), pp. 347-65.
- HARRIS, MILTON AND TOWNSEND, ROBERT M. "Allocation Mechanisms, Asymmetric Information, and the 'Revelation Principle'," in *Issues in contemporary microeconomics and welfare*. Ed.: GEORGE FIEWEL. Albany: State U. of New York Press, 1985, pp. 379-94.
- HARSANYI, JOHN C. "Games with Incomplete Information Played by 'Bayesian' Players," *Manage. Sci.*, Nov. 1967, 14(3), pp. 159-82; Jan. 1968, 14(5), pp. 320-34; Mar. 1968, 14(7), pp. 486-502.
- HARSTAD, RONALD M. AND LEVIN, DAN. "A Class of Dominance Solvable Common-Value Auctions," *Rev. Econ. Stud.*, July 1985, 52(3), pp. 525-28.
- HAUSCH, DONALD B. "Multi-Object Auctions: Sequential vs. Simultaneous Sales," *Manage. Sci.*, Dec. 1986, 32(12), pp. 1599-1610.
- HAYEK, F. A. "The Use of Knowledge in Society," *Amer. Econ. Rev.*, Sept. 1945, 35(4), pp. 519-30.
- HAYWOOD, RALPH. "Optimal High-Valuer Auctions." Mimeo, U. of Calif., San Diego, Nov. 1984.
- HENDRICKS, KENNETH; PORTER, ROBERT H. AND BOUTREAU, BRYAN. "Information and Returns in OCS Auctions: 1954-1969," *J. Ind. Econ.*, forthcoming.
- HENDRICKS, KENNETH; PORTER, ROBERT H. AND GERTLER, PAUL J. "Bidding Behavior and Joint Venture Formation in OCS Auctions." Mimeo. SUNY Stony Brook, Feb. 1986.
- HOLT, CHARLES A., JR. "Uncertainty and the Bidding for Incentive Contracts," *Amer. Econ. Rev.*, Sept. 1979, 69(4), pp. 697-705.
- _____. "Competitive Bidding for Contracts Under Alternative Auction Procedures," *J. Polit. Econ.*, June 1980, 88(3), pp. 433-45.
- _____. "Bidding for Contracts," in *Bayesian analysis in economic theory and time series analysis*. Amsterdam: North-Holland, 1982, pp. 1-62.
- JOHNSON, RONALD. "Oral Auction Versus Sealed Bids: An Empirical Investigation," *Natural Res. J.*, Apr. 1979, 19(2), pp. 315-35.
- KUNREUTHIER, HOWARD AND KLEINDORFER, PAUL R. "A Sealed-Bid Auction Mechanism for Siting Noxious Facilities," *Amer. Econ. Rev.*, May 1986, 76(2), pp. 295-99.
- LAFFONT, JEAN-JACQUES AND MASKIN, ERIC. "Optimal Reservation Price in the Vickrey Auction," *Econ. Letters*, 1980, 6(4), pp. 309-13.
- LAFFONT, JEAN-JACQUES AND TIROLE, JEAN. "Auctioning Incentive Contracts." Mimeo. MIT, Oct. 1985.
- LEE, TOM K. "Resource Information Policy and Federal Resource Leasing," *Bell J. Econ.*, Autumn 1982, 13(2), pp. 561-68.
- _____. "Competition and Information Acquisition in First Price Auctions," *Econ. Letters*, 1985, 18(2-3), pp. 129-32.
- LELAND, HAYNE E. "Optimal Risk Sharing and the Leasing of Natural Resources, with Application to Oil and Gas Leasing on the OCS," *Quart. J. Econ.*, Aug. 1978, 92(3), pp. 413-37.
- LUTON, RICHARD AND MCAFEE, R. PRESTON. "Sequential Procurement Auctions," *J. Public Econ.*, Nov. 1986, 31(2), pp. 181-95.
- MASKIN, ERIC AND RILEY, JOHN. "Auctioning an Indivisible Object." Mimeo. UCLA, April 1980.
- _____. "Auctions with Asymmetric Beliefs." Working Paper No. 254. UCLA, Apr. 1983a.
- _____. "Monopoly with Incomplete Information." Working Paper No. 268. UCLA, Aug. 1983b.
- _____. "Monopoly with Incomplete Information," *Rand J. Econ.*, Summer 1984a, 15(2), pp. 171-96.
- _____. "Optimal Auctions with Risk Averse Buyers," *Econometrica*, Nov. 1984b, 52(6), pp. 1473-1518.
- _____. "Auction Theory with Private Values," *Amer. Econ. Rev.*, May 1985, 75(2), pp. 150-55.
- _____. "Existence and Uniqueness of Equilibrium in Sealed High Bid Auctions." Mimeo, UCLA, Mar. 1986.
- MATTHEWS, STEVEN A. "Selling to Risk Averse Buyers with Unobservable Tastes," *J. Econ. Theory*, Aug. 1983, 30(2), pp. 370-400.
- _____. "Information Acquisition in Discriminatory Auctions," in *Bayesian models in economic theory*. Eds.: MARCEL BOYER AND RICHARD E. KIHILSTROM. Amsterdam: North-Holland, 1984a, pp. 181-208.
- _____. "On the Implementability of Reduced Form Auctions," *Econometrica*, Nov. 1984b, 52(6), pp. 1519-23.
- _____. "Comparing Auctions for Risk Averse Buyers: A Buyer's Point of View," *Econometrica*, forthcoming.
- MCAFEE, R. PRESTON AND MCMILLAN, JOHN. "Bidding for Contracts: A Principal-Agent Analysis," *Rand J. Econ.*, Autumn 1986, 17(3), pp. 326-38.
- _____. "Auctions with a Stochastic Number of Bidders," *J. Econ. Theory*, 1987a, forthcoming.
- _____. "Auctions with Entry," *Econ. Letters*, 1987b, forthcoming.
- _____. "Competition for Agency Contracts," *Rand J. Econ.*, forthcoming.
- _____. "Government Procurement and International Trade: Implications of Auction Theory." Mimeo, U. of Western Ontario, 1987d.
- _____. *Incentives in government contracting*. Toronto: U. of Toronto Press, 1987e.
- _____. "Search Mechanisms," *J. Econ. Theory*, 1987f, forthcoming.
- MCAFEE, R. PRESTON; MCMILLAN, JOHN AND WHINSTON, MICHAEL D. "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quart. J. Econ.*, 1987, forthcoming.
- MEAD, WALTER J. "Natural Resource Disposal Policy: Oral Auction Versus Sealed Bids," *Natural Res. J.*, Apr. 1967, 7(2), pp. 195-224.
- MEAD, WALTER J.; MOSEIDJORD, ASBJORN AND SORENSON, PHILIP E. "Competitive Bidding Under Asymmetrical Information: Behavior and Performance in Gulf of Mexico Drainage Lease Sales 1959-1969," *Rev. Econ. Statist.*, Aug. 1984, 66(3), pp. 505-08.
- MILGROM, PAUL R. "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica*, May 1979a, 47(3), pp. 679-88.
- _____. *The structure of information in competitive bidding*. NY: Garland, 1979b.

- _____. "Good News and Bad News: Representation Theorems and Applications," *Bell J. Econ.*, Autumn 1981a, 12(2), pp. 380-91.
- _____. "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, July 1981b, 49(4), pp. 921-43.
- _____. "The Economics of Competitive Bidding: A Selective Survey," in *Social goals and social organization*. Eds.: LEONID HURWICZ, DAVID SCHMEIDLER AND HUGO SONNENSCHNEIN. Cambridge: U. Press, 1985.
- _____. "Auction Theory," in *Advances in economic theory* 1985. Ed.: TRUMAN BEWLEY. Cambridge: Cambridge U. Press, 1986.
- MILGROM, PAUL R. AND WEBER, ROBERT J. "A Theory of Auctions and Competitive Bidding," *Econometrica*, Sept. 1982a, 50(5), pp. 1089-1122.
- _____. "The Value of Information in a Sealed-Bid Auction," *J. Math. Econ.*, June 1982b, 10(1), pp. 105-14.
- _____. "Distributional Strategies for Games with Incomplete Information," *Math. Operations Res.*, Nov. 1985, 10(4), pp. 619-32.
- MOORE, JOHN. "Global Incentive Constraints in Auction Design," *Econometrica*, Nov. 1984, 52(6), pp. 1523-36.
- MYERSON, ROGER B. "Optimal Auction Design," *Math. Operations Res.*, Feb. 1981, 6(1), pp. 58-73.
- _____. "Bayesian Equilibrium and Incentive Compatibility: An Introduction," in *Social goals and social organization*. Eds.: LEONID HURWICZ, DAVID SCHMEIDLER, AND HUGO SONNENSCHNEIN. Cambridge: Cambridge U. Press, 1985, pp. 229-59. *New York Times*, March 23, 1985.
- ORGANIZATION FOR ECONOMIC COOPERATION AND DEVELOPMENT. *Collusive tendering*. Paris: OECD, 1976.
- ORTEGA-REICHERT, ARMANDO. "Models for Competitive Bidding Under Uncertainty." Unpublished PhD thesis. Dept. of Operations Research, Stanford U., Jan. 1968.
- PALFREY, THOMAS R. "Multiple-Object, Discriminatory Auctions with Bidding Constraints: A Game-Theoretic Analysis," *Manage. Sci.*, Sept. 1980, 26(9), pp. 935-46.
- _____. "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information," *Econometrica*, Mar. 1983, 51(2), pp. 463-83.
- PELTO, CHESTER R. "The Statistical Structure of Bidding for Oil and Mineral Rights," *J. Amer. Statist. Assoc.*, Sept. 1971, 66(335), pp. 456-60.
- PICKFORD, M. "A New Test for Manufacturing Industry Efficiency: An Analysis of the Results of Import Licence Tendering in New Zealand," *Int. J. Ind. Organ.*, June 1985, 3(2), pp. 153-77.
- PLOTT, CHARLES. "Industrial Organization Theory and Experimental Economics," *J. Econ. Lit.*, Dec. 1982, 20(4), pp. 1485-1527.
- QUIRK, JAMES AND TERASAWA, KATSUAKI. "The Winner's Curse and Cost Estimation Bias in Pioneer Projects." Working Paper No. 512, Calif. Inst. of Tech., Apr. 1984.
- RAMSEY, JAMES B. *Bidding and oil leases*. Greenwich, CT: J.A.I. Press, 1980.
- RASSENTI, S. J.; SMITH, V. L. AND BULFIN, R. L. "A Combinational Auction Mechanism for Airport Time Slot Allocation," *Bell J. Econ.*, Autumn 1982, 13(2), pp. 402-17.
- REECE, DOUGLAS, K. "Competitive Bidding for Offshore Petroleum Leases," *Bell J. Econ.*, Autumn 1978, 9(2), pp. 369-84.
- _____. "An Analysis of Alternative Bidding Systems for Leasing Offshore Oil," *Bell J. Econ.*, Autumn 1979, 10(2), pp. 659-69.
- RILEY, JOHN G. "Ex Post Information in Auctions." Working Paper No. 367. UCLA, Mar. 1986.
- RILEY, JOHN G. AND SAMUELSON, WILLIAM F. "Optimal Auctions," *Amer. Econ. Rev.*, June 1981, 71(3), pp. 381-92.
- RIORDAN, MICHAEL H. AND SAPPINGTON, DAVID E. M. "Awarding Monopoly Franchises," *Amer. Econ. Rev.*, forthcoming.
- ROBINSON, MARC S. "Oil Lease Auctions: Reconciling Economic Theory with Practice." Discussion Paper No. 292. UCLA, Sept. 1984.
- _____. "Collusion and the Choice of Auction," *Rand J. Econ.*, Spring 1985, 16(1), pp. 141-43.
- SAMUELSON, WILLIAM F. "Competitive Bidding for Defense Contracts," in RICHARD ENGELBRECHT-WIGGANS ET AL., 1983, pp. 389-419.
- _____. "Competitive Bidding with Entry Costs," *Econ. Letters*, 1985, 17(1-2), pp. 53-57.
- _____. "Auctions with Contingent Payments: Comment." Mimeo. Boston U., Feb. 1986a.
- _____. "Bidding for Contracts," *Manage. Sci.*, Dec. 1986b, 32(12), pp. 1533-50.
- SCHELLING, THOMAS C. *The strategy of conflict*. Cambridge: Harvard U. Press, 1960.
- SHUBIK, MARTIN. "Auctions, Bidding, and Markets: A Historical Sketch," in R. ENGELBRECHT-WIGGANS ET AL., 1983, pp. 33-52.
- SMILEY, ALBERT K. *Competitive bidding under uncertainty: The case of offshore oil*. Cambridge, MA: Ballinger, 1979.
- SMITH, JAMES L. "Non-Aggressive Bidding Behavior and the 'Winner's Curse'," *Econ. Inquiry*, July 1981, 19(3), pp. 380-88.
- SMITH, VERNON L. "Experimental Studies of Discrimination Versus Competition in Sealed-Bid Auction Markets," *J. Bus.*, Jan. 1967, 40(1), pp. 56-84.
- _____. "Microeconomic Systems as an Experimental Science," *Amer. Econ. Rev.*, Dec. 1982, 72(5), pp. 923-55.
- STARK, ROBERT AND ROTHKOPF, MICHAEL H. "Competitive Bidding: A Comprehensive Bibliography," *Operations Res.*, Mar.-Apr. 1979, 27(2), pp. 364-91.
- STIGLER, GEORGE J. "A Theory of Oligopoly," *J. Polit. Econ.*, Feb. 1964, 72(1), pp. 44-61.
- VICKREY, WILLIAM. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *J. Finance*, Mar. 1961, 16(1), pp. 8-37.
- WEBER, ROBERT J. "Multiple-Object Auctions," in R. ENGELBRECHT-WIGGANS ET AL., 1983, pp. 165-91.
- WILLIAMSON, OLIVER E. "Franchise Bidding for Natural Monopolies: In General and with Respect to

- CATV," *Bell J. Econ.*, Spring 1976, 7(1), pp. 73-104.
- WILSON, ROBERT. "A Bidding Model of Perfect Competition," *Rev. Econ. Stud.*, Oct. 1977, 44(3), pp. 511-18.
- . "Auctions of Shares," *Quart. J. Econ.*, Nov. 1979, 93(4), pp. 675-90.
- . "Incentive Efficiency of Double Auctions," *Econometrica*, Sept. 1985, 53(5), pp. 1101-115.
- . "Double Auctions," in *Information, incentives, and economic mechanisms: Essays in honor of Leonid Hurwicz*. Eds.: T. GROVES, R. RADNER, AND S. REITER. Minneapolis: U. of Minnesota Press, 1986a.
- . "Equilibria of Bid-Ask Markets," in *Arrow and the ascent of economic theory: Essays in honor of Kenneth J. Arrow*. London: Macmillan, 1986b.
- . "Auction Theory," in *The new Palgrave: A dictionary of economic theory*. Eds.: J. EATWELL, M. MILGATE, AND P. NEWMAN. London: Macmillan, 1987.
- YAMEY, B. S. "Why £2,310,000 for a Velazquez? An Auction Bidding Rule," *J. Polit. Econ.*, Nov.-Dec. 1972, 80(6), pp. 1323-27.
- YUSPEL, LARRY. "A Case for Increasing the Use of Competitive Procurement in the Department of Defense," in *Bidding and auctioning for procurement and allocation*. Ed.: YAKOV AMIHUD. NY: NYU Press, 1976, pp. 104-26.