
Auction theory

Paul R. Milgrom

1 Introduction

Auctions are one of the oldest surviving classes of economic institutions. The first historical record of an auction is usually attributed to Herodotus, who reported a custom in Babylonia in which men bid for women to wed.¹ Other observers have reported auctions throughout the ancient world – in Babylonia, Greece, the Roman Empire, China, and Japan.²

As impressive as the historical longevity of auctions is the remarkable range of situations in which they are currently used. There are auctions for livestock, a commodity for which many close substitutes are available. There are also auctions for rare and unusual items like large diamonds, works of art, and other collectibles. Durables (e.g., used machinery), perishables (e.g., fresh fish), financial assets (e.g., U.S. Treasury bills), and supply and construction contracts are all commonly bought or sold at auction. The auction sales of unique items have suggested to some that auctions are a good vehicle for monopolists. But it is not only those in a strong market position who use auctions. There are also auction sales of the land, equipment, and supplies of bankrupt firms and farms. These show that auctions are used by sellers who are desperate for cash and willing to sell even at prices far below replacement cost.

The first draft of this paper was written while I was a Fellow at the Institute for Advanced Studies of the Hebrew University of Jerusalem. Discussions with Charles Wilson, Motty Perry, and, especially, Ariel Rubinstein contributed enormously to my understanding of the relation between auctions and bargaining. Comments by Byung-Il Choi and Alvin Roth on a previous version of this manuscript led to improvements in the exposition.

¹ Herodotus may not have been the first to publish. Some scholars interpret the biblical account of the sale of Joseph (the great-grandson of Abraham) into slavery as being an auction sale.

² For a more detailed history of auctions and a description of some of the auctions used in the modern world, see Cassady (1967).

Indeed, the only clear common denominator for the kinds of objects that are sold at auction is the need to establish individual prices for each item sold. Used cars, whose condition varies over a wide range, are sold to dealers at auction; new cars are not. Livestock are sold at auction even though close substitutes are readily available, because individual animals differ in weight and health. The price of fresh fish needs to be determined daily, because the daily supply of fish varies so tremendously. Construction contracts are normally too complex to allow a simple pricing schedule to work; competitive bids sometimes provide a workable alternative.

In this essay, I review only a small part of auction theory – the part that claims to explain the long and widespread use of auctions and competitive bidding and to account for certain details of the way auctions are usually conducted. These details include the popular use of sealed-bid and ascending-bid auctions, the establishment of minimum prices, the preparation of expert appraisals of items being sold, and so forth.

Logically prior to explaining the use of auctions is defining just what an auction is. The characteristic feature of an auction is that there is an explicit *comparison* made among bids. In the ascending-bid (“English”) auction, a bidder’s offer remains open long enough for other bidders to make counteroffers, so that the seller can take the highest offer. In the sealed-bid auction, the bidders’ offers are all made simultaneously, so that the seller can compare them directly. In the descending-bid (“Dutch”) auction, the seller makes a series of price offers, declining over time. Each bidder has the opportunity to accept or reject the seller’s latest price offer; this affords the seller an opportunity to compare the timing of buyers’ offers, and to take the offer that is made earliest. Each of these auctions requires that all the bidding be completed within a relatively short period of time. They can be contrasted with, say, a sequential bargaining process in which the seller negotiates one-by-one with a series of buyers who make short-lived offers, so that the seller has no opportunity to compare the simultaneous offers of competing buyers. We shall develop the importance of this difference in more detail later.

The simplest explanation of the continuing popularity of auctions is that auctions often lead to outcomes that are efficient and stable. More formally, in a static deterministic model, the set of perfect equilibrium trading outcomes obtained in an auction game (as the minimum bid is varied) coincides with the set of core allocations. An outcome is in the *core* when there is no coalition of traders that can, by trading just among its members, make all coalition members better off.

To understand the significance of this conclusion, imagine a situation in which a single item is sold but the resulting allocation lies outside the core. There are two possibilities. First, the allocation may be inefficient;

in this case, the new owner will likely find it profitable to resell the item to a buyer who values it more. The second possibility is that, even though the allocation is efficient, there are other buyers around who were willing to pay a higher price (and after the auction are willing to tell the seller so). In either case, the seller may well resolve not to be so quick to sell the next time around and perhaps even to compare alternative offers – that is, to conduct some kind of auction.

A second explanation of the popularity of auctions highlights the advantages of an auction to a seller in a relatively poor bargaining position³ (such as the owner of a nearly bankrupt firm) when the goods sold at auction can later be resold. Consider the problem of such a seller. Suppose that there are two potential buyers: Mr. 1, who has a high valuation for the item being sold, and Mr. 2, whose valuation is lower. What happens if the seller conducts an auction with a low minimum price? At the equilibrium of the auction game, the item will be sold to Mr. 1 for approximately its value to Mr. 2. With the possibility of resale, that value cannot be less than the price that Mr. 2 could get by reselling to Mr. 1. By conducting an auction, the seller expects to get about the same price as Mr. 2 would get, even though Mr. 2 may be much better positioned for face-to-face bargaining with Mr. 1. Thus, a seller in a relatively weak bargaining position can do as well as a strong bargainer by conducting an auction.

These first two explanations of the prevalence of auctions are developed in detail in Section 2, which focuses on deterministic auction models. A third explanation, reviewed in Section 3, is that even a seller in a strong bargaining position will sometimes find it optimal to conduct an auction. That is, the seller will prefer to conduct some standard auction, such as the sealed-bid or ascending-bid auction with a suitably chosen minimum price, rather than to play any other exchange game⁴ with the bidders.

The three explanations just described are, of course, complementary. Together, they provide a cogent set of reasons for a seller to use an auction when selling an indivisible object over a wide range of circumstances.

In the auction models discussed so far, there is little that can be said about the details of how auctions are conducted. In those models, many kinds of auctions (including all the usual ones) lead to the same mean price. However, this “independence” result depends on the assumption that bidders have no private information about each other. Formally, the observations they make are assumed to be statistically independent. When

³ That is, a poor bargaining position relative to the potential buyers.

⁴ An exchange game is any game whose outcome determines an allocation and time of trade, and in which each player has a strategy of nonparticipation that leaves him with his initial allocation.

there is correlated uncertainty on the part of the bidders, different auction rules lead to different mean prices.

In Section 4, we introduce correlated uncertainty into the bidding model and focus on the strategies open even to a seller with no bargaining power, that is, one who cannot commit himself to withhold an item that attracts only low bids.⁵ What strategies can such a seller adopt? For one, he can normally choose which kind of auction to offer, provided the minimum bid is kept low, because buyers will always want to participate in the auction.⁶ Normally, the seller can also decide whether to reveal any information about the item being sold or about the potential buyers, because it always pays a buyer to listen if he can do so without being seen. Given these options, the seller's preferences are surprisingly systematic. In a wide range of circumstances,⁷ the seller will prefer (1) to conduct an ascending-bid auction rather than a sealed-bid auction, (2) to reveal all information that he has available, and (3) to link the price to any available exogenous indicators of value.

The analysis leading to these conclusions is founded on what has been called the *Linkage Principle*. Intuitively, a bidder's expected profits from an auction are greatest when he has private information that the item being sold is quite valuable. The intuition of the Linkage Principle is that the auctions yielding the highest average prices are those that are most effective at undermining the privacy of the winning bidder's information, thereby transferring some profits from the bidders to the seller. According to the principle, privacy is undermined by linking price to information other than (but correlated with) the winning bidder's private information.

The three conclusions described above all follow from the Linkage Principle. In an ascending-bid auction, the equilibrium price depends on the information of losing bidders through the bids they place. That dependence, or linkage, is absent in the sealed-bid auction. Its presence in the ascending-bid auction leads to a higher predicted price (provided that the bidders' information is correlated).

⁵ In Section 4, we review some game-theoretic arguments supporting the presumption that a "rational" seller cannot hold out for a high price when he is uncertain about the buyers' reservation prices.

⁶ No matter what strategies the other players adopt, each buyer does at least as well by entering the minimum bid as by abstaining from the auction. For some strategies – namely, when others refrain from bidding – he does better. (This argument is transparent for the case where resale is impossible, and can be extended also to the case with resale possibilities.)

⁷ The principal assumptions required include risk neutrality, symmetric uncertainty about the bidders' valuations, and a strong form of nonnegative correlation, known as *affiliation*, among the bidders' valuations.

In any kind of auction, the seller, by revealing information, influences the bids and therefore the price. So, by revealing his information, the seller links the price directly to his information. Thus, according to the Linkage Principle, a policy of revealing information raises the expected price that will result from the auction, provided that the information to be revealed is affiliated⁸ with the bidders' information. Similarly, basing the price in part on ex post indicators of value creates a linkage that on average increases the expected price (if these indicators are affiliated with the bidders' information). Examples of contracts let at auction where price is determined in part by ex post indicators include construction contracts with a cost-sharing provision and petroleum drilling contracts that provide for royalty payments based on actual production.

The main theme of explaining the prevalence and robustness of auctions is continued in Section 5, where the possibility of collusion is briefly studied. Collusion is widespread in real auctions, and there is little a one-time seller can do to prevent it when the bidders have a long-term relationship. However, it is shown that ascending-bid auctions are more vulnerable to collusive agreements among bidders in a long-term relationship than are sealed-bid auctions. This is an important reason for industrial firms to solicit sealed bids from suppliers, despite the general superiority of ascending-bid auctions in one-shot competitive situations.

2 Auctions, bargaining, and the core

We begin by formulating and proving the claim that the trading outcomes of the auction game coincide with the core of the corresponding exchange game. This result provides a simple, partial answer to the question of why auction institutions are so prevalent throughout the world and throughout history.

Consider a deterministic setting with a single seller and n (potential) buyers for some item. Let s be the monetary value of the item to the seller; this means that if the seller had the option of selling for some price p or not selling the item at all, he would choose to sell for p if and only if $p \geq s$. Similarly, the buyers have monetary valuations b_1, \dots, b_n . Our model is discrete: All the valuations and bids are multiples of some common unit. Here and throughout this chapter, we make the standard game-theoretic assumption that the deterministic parameters are common knowledge

⁸ Random variables are said to be *affiliated* when they are positively correlated conditional on lying in any small rectangle. For example, any pair of positively correlated joint normal random variables are affiliated. A precise formal definition of the concept is given in Section 4.

among the buyers and the seller.⁹ Without significant loss of generality, we may assume that $b_1 > \dots > b_n$ and limit attention to the case where there are some potential gains from trade: $b_1 > s$.

Now, if the seller offers the item for sale using a sealed-bid auction with minimum price $m < b_2$,¹⁰ what will happen? Using any sensible equilibrium concept (e.g., Nash equilibrium in undominated strategies,¹¹ perfect equilibrium, "rationalizable" strategies, or even correlated equilibrium), the item will be sold to bidder 1 for his bid of b_2 .¹² The same trade will occur if the seller sets any minimum price not exceeding b_2 . Again, the same will occur if the seller hires an auctioneer to conduct an ascending-bid auction, regardless of whether the bids are called by the bidders themselves or at a slow pace by the auctioneer.

If the seller sets a minimum price $m \in (b_2, b_1)$, the equilibrium outcome assigns the item to bidder 1 for a price of m . Of course, if $m > b_1$,

⁹ In the standard theory of games, the players need to know this structure in order to compute the equilibrium and determine how to play. An alternative view, relevant to auction theory, holds that players learn from experience about the reduced form of their decision problems and select their best bids for that problem. Equilibrium is then a state where all players have correctly learned and are using optimal strategies in their decision problems. Mathematically, this leads to the same definition of equilibrium as does the standard view, but it raises different stability questions and does not require as much knowledge among the players about the overall structure of the game.

¹⁰ We assume in this auction and all those considered hereafter that ties are broken by tossing a fair coin.

¹¹ Although the Nash equilibrium and its refinements are often justifiably criticized, they are particularly well suited to the analysis of auction games. A Nash equilibrium can be defined as a profile of strategies, one for each player, such that (1) each player is maximizing given his beliefs about how the others will play and (2) those beliefs are correct. The first condition is neither stronger nor weaker than the usual rationality assumption in economic models. The second ("rational expectations") condition is most plausible for institutions – such as auctions – that have existed for millennia and so for which expectations can be based on actual experience.

¹² Any perfect equilibrium (Selten, 1975) is a Nash equilibrium in undominated strategies, and in fact for this game the two concepts coincide. In the two-bidder game, the set of perfect equilibria are characterized as follows: Bidder 1 bids b_2 . Bidder 2 uses any mixed strategy F that satisfies two conditions. First, $F(b_2^-) = 1$. Second, let $G(x) = [F(x) + F(x-1)]/2$; then $G(x) \leq (b_1 - b_2)/(b_1 - x)$ for all $x \in (m, b_2)$. With more than two bidders, one can specify the strategies of the others arbitrarily, provided bidder j always bids less than b_j , and this remains a perfect equilibrium.

Rationalizable strategies are derived by eliminating weakly dominated strategies from the strategy set to form a reduced game. Then weakly dominated strategies are eliminated from the reduced game, and so on until the process ends. The strategies that survive are called *rationalizable*. The only such strategies for bidders 1 and 2 are to bid b_2 and $b_2 - 1$, respectively.

Correlated equilibria (Aumann, 1973) of bidding games employ only rationalizable strategies, so that concept is covered also.

no exchange takes place; in that case the seller's payoff is s and each buyer's payoff is zero. The case $m = b_1$ is somewhat degenerate; its equilibria include both the no-trade outcome and a trade at price b_1 . Our earlier choice of the phrase "equilibrium trading outcomes" was intended to denote all the equilibrium outcomes except the no-trade outcome. Our claim is then justified by the following proposition.

Proposition 1. The set of perfect equilibrium outcomes of the auction game, as the minimum price ranges from s to b_1 , consists of the core outcomes of the corresponding exchange game together with the no-trade outcome. The latter can only occur when the minimum price is b_1 .

Proof: Let $x = (x_0, x_1, \dots, x_n)$ be the vector of payoffs that are received by the seller and the n buyers, respectively. A vector of payoffs x is called an *imputation* if it is individually rational (i.e., nonnegative) and Pareto optimal and corresponds to some feasible allocation of the goods and money among the players. These imply:

$$x_0 + x_1 + \dots + x_n = b_1. \quad (2.1)$$

To be in the core, an imputation must also satisfy inequalities asserting that no coalition could, by agreeing to exchange among themselves, earn a higher total payoff:

$$x_0 + \sum_{i \in S} x_i \geq \max\{s, (b_i; i \in S)\}, \quad \text{for all } S \subset \{1, \dots, n\}. \quad (2.2)$$

In view of the preceding discussion, the proposition asserts that the core consists entirely of points of the form

$$(x_0, b_1 - x_0, 0, \dots, 0) \quad \text{for } \max(s, b_2) \leq x_0 \leq b_1.$$

It is easy to check that all such points satisfy (2.1) and the inequalities (2.2), and so in fact do lie in the core.

Conversely, suppose x lies in the core. From (2.1) and nonnegativity, $x_0 + x_1 \leq b_1$. From (2.2) for $S = \{1\}$, $x_0 + x_1 \geq b_1$. Hence, $x_0 + x_1 = b_1$ and, by (2.1) and nonnegativity, $x_2 = \dots = x_n = 0$. Therefore, all points in the core are of the form $(x_0, b_1 - x_0, 0, \dots, 0)$. Using (2.2) with $S = \{2\}$, one finds $x_0 + x_2 \geq \max(s, b_2)$; so $x_0 \geq \max(s, b_2)$. Nonnegativity of x_1 implies $x_0 \leq b_1$. ■

The strategic equivalence of the Dutch and sealed-bid auctions and the notion of perfect equilibrium do not transfer neatly to bidding games with continuous bid spaces. For discrete bid spaces with bid increment ϵ ,

the only perfect equilibrium in the Dutch auction is for the highest evaluator to stop the auction when the price reaches b_2 , and for each other player i to stop it at the price $b_i - \epsilon$. There are no corresponding strategies in the standard formulation of the continuous Dutch auction, because there is no possibility of bidding b_i "minus an infinitesimal." Indeed, in the standard formulation of the continuous Dutch auction, no subgame perfect equilibrium exists.

To avoid this problem, we formulate the extensive form Dutch auction game so that a bidder can claim the object whenever the price falls to p , which we call bidding p , or whenever the price falls strictly below p , which we call bidding p^- . If a player bids p , another bids p^- , and all others bid less, then the item is awarded to the one who bids p for price p . If a player bids p^- and nobody else bids more, then the item is awarded to that bidder for a price of p . This specifies a well-defined continuous Dutch auction game which suitably generalizes the game with discrete bid amounts. Moreover, like the discrete bids game, it does have a unique subgame perfect equilibrium: Player 1 bids b_2 and each $i \neq 1$ bids b_i^- .¹³

There still remains the problem that "trembling-hand" perfect equilibrium is undefined for sealed-bid auction games with a continuum of possible bids. To avoid unnecessary technical difficulties, we shall normally limit our analysis to equilibria of Dutch auctions.

From the perspective of cooperative game theory, the seller's ability to set any particular minimum price and stick to it measures his bargaining power.¹⁴ Indeed, the case $n = 1$ is just a bargaining problem, and auction theory predicts (as does core theory) only that the outcome will be efficient and that nobody will be worse off at equilibrium than if they did not trade. Evidently, a complete auction theory must be informed to some degree by bargaining theory. This leaves open the possibility that the predictions of auction theory could be quite sensitive to the bargaining model used.

Actually, when there are several viable bidders, auction theory is surprisingly insensitive to the bargaining theory used at its foundations. To show this, we embed the auction model in a general discounted, infinite-horizon, noncooperative model of bargaining in which an owner always

¹³ One could, of course, define a modified sealed-bid auction game that is strategically equivalent to our continuous Dutch auction game. However, comparing the subgame perfect equilibria of the Dutch auction game (identified in the text) with the trembling-hand perfect equilibria of the corresponding sealed-bid auction (identified in note 12) shows that the two games are not equivalent for the purposes of perfect equilibrium analysis.

¹⁴ The role of commitment in bargaining has been analyzed by Crawford (1982). The associated roles of patience and risk aversion have been given a particularly penetrating analysis by Binmore, Rubinstein, and Wolinsky (1986).

has the right to resell anything he has bought. Because a single player may be sometimes a buyer and sometimes a seller, we shall designate a player's valuation by v_i rather than by s or b_i . It is assumed that $v_1 > \dots > v_n > 0$.

Let Γ^i be a game form that is to be played when i is the owner of the durable good. Thus, $\Gamma^i = (\{\Sigma_j^i; j = 1, \dots, n\}, f^i)$, where Σ_j^i is the set of strategies available to j in the game form and f^i is a function mapping strategy profiles into outcomes. Time is modeled as discrete. An outcome involving trade specifies a date of trade $t \geq 1$, a (nonnegative) price p , and the next owner j . There is also an outcome called "no trade" that we identify as a trade at date $t = \infty$. To interpret the results that follow, it will be useful to think of t as the period of i 's ownership, rather than to associate t with any actual date.

Certain specified strategies are assumed to be available to the players in each game form Γ^i . First, the owner is permitted to keep the item for himself; that is, he may choose a strategy that always leads to no trade. Second, the owner is permitted to offer a Dutch auction with a zero minimum price. Such an offer, if made, is the first move in Γ^i and initiates an auction subgame (actually, a "subgame form"). If any non-owner bids in the auction, Γ^i ends at date 1 with the item being assigned according to the usual Dutch auction rules. Non-owners must decide simultaneously whether to bid. If no bids are made, play continues according to the continuation rules of Γ^i , whatever they may be. Each non-owner is assumed to have a strategy of refusing to be party to any trade, in which case no payment can be required of him. The assumption that the decision of whether to offer an auction immediately is the first move in Γ^i means that non-owners have no way, before an auction is offered, to commit themselves not to trade.

Using these very general game forms, which specify the rules governing trade given the owner's identity, we create a game in which the buyer can (if he chooses) resell the good. Let player i_0 be the initial owner. Then the game form Γ^{i_0} is played. If the outcome involves trade after a period of ownership of length t_0 , at price p_0 , and with next owner i_1 , we continue with game form Γ^{i_1} , which determines a period of ownership t_1 , price p_1 , and next owner i_2 . The outcome of this sequence of trades specifies that i_0 owns the item from date 0 to date $t_0 - 1$, i_1 owns the item from t_0 to $t_0 + t_1 - 1$, and generally i_j owns it from date $t_0 + \dots + t_{j-1}$ to date $t_0 + \dots + t_j - 1$. Payments are made on the dates of transfer of ownership. The number of actual times the good changes hands can be finite or infinite.

The payoff associated with any outcome for any fixed player j is the present value of the flow of benefits he receives plus the net present value of payments received minus payments made. To make this more precise,

fix an outcome path. Let $1_j(t)$ be one if player j owns the item on date t and zero otherwise. (In particular, $1_j(-1) = 0$.) Let $p(t)$ be the price paid in any trade at date t , or zero if there is no trade at t . Then, j 's payoff in the game is:

$$\sum_{t=0}^{\infty} \delta^t [(1-\delta)v_j 1_j(t) + p(t)[1_j(t-1) - 1_j(t)]]. \quad (2.3)$$

Thus, δ is the discount factor for the players' payoffs.

With this, the specification of the selling games is completed. Corresponding to each player i there is a game in which the identity i_0 of the initial owner is i . We shall call that game Γ_i^* .

The games Γ_i^* that can be constructed in this way for some choice of game forms Γ^i form a huge class. Included are games where the seller can conduct auctions with a positive minimum price, exclude some set of bidders, bargain effectively with some buyers, commit himself to take-it-or-leave-it offers, or do all of these. Indeed, the only important restrictions on the set of options available to a seller are that he can neither compel a non-owner to buy nor prevent a buyer from reselling the good, and that he can always offer an auction with a zero minimum price. An additional "stationarity" restriction will be imposed through the equilibria that we isolate for study.

In general, a strategy for a player specifies how to play at each date as a function of the date and the entire past history. For our analysis, we limit attention to equilibria in which the players adopt stationary strategies. A *stationary strategy* for player j is an n -tuple $\sigma_j = (\sigma_j^1, \dots, \sigma_j^n)$ such that $\sigma_j^i \in \Sigma_j^i$. Such a strategy specifies how player j should play in each game form Γ^i (he should play σ_j^i) without regard to the earlier history of play. By a *stationary perfect equilibrium*, we mean an n -tuple of stationary strategies $(\sigma_1, \dots, \sigma_n)$ that is a perfect equilibrium profile regardless of the identity of the initial owner (i.e., in each of the games Γ_i^*).

Given a strategy profile $(\sigma_1, \dots, \sigma_n)$, one can define for each player i a value v_i^* associated with owning the item, that is, with playing the game Γ_i^* . With stationary strategies, v_i^* is also the continuation payoff or value of acquiring ownership at any point in the game, regardless of the previous history of play. With nonstationary strategies, that value might depend on the history of play, because future play could also depend on the history.

Proposition 2. *Assume there are least three players, $n \geq 3$. Let v_i^* be the expected payoff to i at a stationary perfect equilibrium in the game Γ_i^* . Then $v_1^* = v_1 > v_i^*$ for all $i \neq 1$. Let a_i^* be the payoff to i in Γ_i^* if all players except i adhere to their equilibrium strategies while i deviates to adopt a strategy that entails conducting*

an auction and refusing to participate in any later sale. Then, for all $i, j \neq 1$, the following inequalities hold:

$$v_i^* > \delta v_j^*, \quad (2.4)$$

$$a_i^* > \delta^2 v_i^*. \quad (2.5)$$

According to Proposition 2, for δ close to 1 it is optimal or nearly optimal for any player other than the highest evaluator to offer the good at auction. Moreover, initial ownership is about equally valuable for all the players other than player 1, even though owners may differ in the strategies available to them. The proposition applies even when some owners are able to make credible take-it-or-leave-it offers while others can sell only at auction. The ability to conduct an auction allows a weak bargainer to benefit from the abilities of any stronger bargainers who may be present, forcing player 1 to bid just as if he were bargaining with a strong player.

Proof of Proposition 2: In the game Γ_1^* , player 1 can guarantee a payoff of v_1 by refusing to trade. Each other player can guarantee a payoff of zero. Hence, the equilibrium payoffs must be at least that high. But the maximum total utility from any outcome is v_1 , and that can be achieved only if player 1 is the owner in every period. Hence, in Γ_1^* , any equilibrium must specify that no trade occurs.

Suppose $i \neq 1$. In the game Γ_i^* , the total payoffs to all players at equilibrium cannot exceed the total payoff $(1-\delta)v_i + \delta v_1 < v_1$ that results from an efficient exchange. Since all players' payoffs are nonnegative, this implies that $v_i^* < v_1$.

Now, in Γ_i^* , i has the opportunity to conduct an auction with minimum price 0. If he does and the non-owners refuse to participate, let the expected continuation payoffs at equilibrium be $\bar{v} = (\bar{v}_1, \dots, \bar{v}_n)$. The auction offer will be accepted by someone with certainty (at equilibrium) unless each non-owner prefers (weakly) his continuation payoff to the payoff from placing the minimum bid of 0: $\bar{v}_j \geq \delta v_j^*$. Because the sale price cannot be negative, the seller's expected payoff can never be less than the first period flow: $\bar{v}_i \geq (1-\delta)v_i$. Thus, some bidder will participate in the auction at equilibrium unless the expected sum of payoffs when no auction takes place is at least $\delta \sum_{j \neq i} v_j^* + (1-\delta)v_i > \delta v_1 + (1-\delta)v_i$. The last expression, however, is the total payoff that results from efficient trade. Hence, there can be no strategies leading to continuation payoffs \bar{v} satisfying the requisite inequality.

So the auction, if offered, will be played with certainty. Let p be the lowest price in the support of the equilibrium price distribution when an auction is offered in Γ_i^* . Let $k \notin \{1, i\}$ be such that $v_k^* = \max_{j \in \{1, i\}} v_j^*$.

During a Dutch auction, when the price has fallen to any level $p' > p$, bidder 1 must expect a payoff of at least $\delta(v_1 - p')$ from allowing the auction to continue, and k must expect at least $\delta(v_k^* - p')$. Because the total expected payoff to all bidders in the continuation after p' is at most $\delta(v_1 - p)$ (since the price will be at least p), it follows that $\delta(v_1 - p') + \delta(v_k^* - p') \leq \delta(v_1 - p)$ for all $p' > p$. Hence, $p \geq v_k^*$. In the Dutch auction, no bidder j can benefit by bidding more than v_j^* , so bidder 1 never finds it optimal (at equilibrium) to bid strictly more than v_k^* . Hence, the highest equilibrium price cannot exceed v_k^* . Therefore the equilibrium price following an auction offer is v_k^* with probability one.

If i offers an auction in Γ_i^* and never repurchases the item, while the other players follow their stationary equilibrium strategies, i 's payoff will be

$$a_i^* = (1 - \delta)v_i + \delta v_k^* > \delta v_j^*, \quad (2.6)$$

for all $j \notin \{1, i\}$. Since $v_i^* \geq a_i^*$, (2.4) follows.

Next we make two applications of (2.6),

$$a_i^* > \delta v_j^* \geq \delta a_j^* > \delta^2 v_i^*, \quad (2.7)$$

which establishes (2.5). ■

So far, we have allowed the game forms $(\Gamma^i; i = 1, \dots, n)$ to be quite general. As an aid to intuition, let us now specify some simple game forms as follows. At odd-numbered dates, the owner chooses a non-owner to whom to make a price offer. The non-owner can accept the offer, in which case a trade is consummated. Or, the non-owner can reject the offer and, at the next (even) date, make a counteroffer. So far, this is the same game as used for the "telephone bargaining" model of Binmore (1983). Now comes a difference: We specify that the owner can, at time 0, offer a Dutch auction with a zero minimum price. If no buyer participates in the auction, then the seller can make a private offer at time 1, and the game continues in the Binmore fashion.

The games Γ_i^* constructed from the specified game forms differ from Binmore's telephone bargaining game in two ways: by allowing auctions and by including the possibility of resale. One can show that each Γ_i^* has a unique perfect equilibrium outcome. To describe it, define a function $B: \mathbb{R}^2 \rightarrow \mathbb{R}$ by:

$$B(x, y) \equiv x + (y - x)/(1 + \delta). \quad (2.8)$$

In the telephone bargaining model, if $i \neq 1$ then the perfect equilibrium outcome is that the item is sold to player 1 for the price $B \equiv B(v_i, v_1)$. Note that for δ close to one, the bargainers split the surplus almost equally. Note too that the price is not at all sensitive to the presence of additional

bargainers. For the games Γ_i^* , however, the equilibrium outcomes are quite different.

*Proposition 3. If the initial owner is $i = 1$, then no trade ever occurs at equilibrium. If the initial owner is any other player, then at equilibrium the item is sold at date 1 to player 1 and never resold. The sale is by private offer if the private offer price of $B(v_i, v_1)$ is larger than the auction price of $(1 - \delta)v_2 + \delta B(v_2, v_1)$. Otherwise, the sale is by auction.*¹⁵

The game has been structured so that the owner makes the first offer. Because delays are costly, this gives the owner some advantage in the bargaining. When the time between successive offers is long and delay costs are high, the advantage of making the first offer in negotiations is large, and it is not optimal then for the owner to give up that advantage by conducting a low minimum-price auction.

Probably more common is the situation where the time between offers is small enough that δ is nearly one.¹⁶ In this case, the auction price will exceed the private offer price if and only if there are at least two non-owners with higher valuations than the initial owner i . That is, at equilibrium the owner bargains if and only if there is only one real potential buyer. Otherwise, he conducts an auction!

Together, Propositions 2 and 3 provide a strong case for the desirability of conducting an auction.

3 Expected price-maximizing auctions

So far, we have shown that auctions lead to core outcomes and that when resale is possible and trading costs are low, it is almost optimal for almost every seller to conduct an auction with a low minimum price. This near optimality holds regardless of the other alternatives available to the seller, provided only that the buyers cannot be compelled to buy. However, the specific example with which we ended Section 2 establishes that conducting an auction with a low minimum price is not generally the best strategy for a seller in a strong bargaining position.

We therefore turn to the question: What is the best strategy for a seller with a hegemony of bargaining power? What we have in mind is a situation in which the seller, for some unspecified reason, has the power to

¹⁵ The proof is omitted. It follows the now familiar lines for alternating offer bargaining models (see Rubinstein, Chapter 5, this volume).

¹⁶ For example, if the annual real interest rate were as much as 5% and it took a week to arrange an auction, δ would be .999.

select any institution he likes for conducting trade. The seller assumes that the buyers will agree to participate if their expected payoffs are non-negative – the buyers are too weak to demand more. In the deterministic setting, the seller's optimal strategy is obvious: make a take-it-or-leave-it offer to the highest valuation buyer that extracts all the surplus from him. If buyers' valuations are private information, however, then the seller cannot implement such a strategy; he does not know what offer to make. What, then, should an expected price-maximizing seller do?

Let us begin with the simplest case. We suppose that there is only one buyer whose valuation V ($V \geq 0$) for the item is unknown and has distribution F . Suppose the seller's valuation is s , corresponding to a flow benefit from ownership of $(1-\delta)s$. These valuations mean that if the buyer acquires the item at date t for a price of p , his payoff (in von Neumann-Morgenstern utility) is $(V-p)\delta^t$ and the seller's is $p\delta^t + s(1-\delta^t)$.¹⁷ If no trade occurs, the buyer's payoff is zero and the seller's is s . Following Vickrey's style of formulation, let us suppose that the buyer observes private information X and has a valuation $V = u(X)$, where X is uniformly distributed on $(0, 1)$ and u is a nondecreasing function.¹⁸ For simplicity, we take u to be strictly increasing and continuously differentiable.

If the seller makes a take-it-or-leave-it offer at a price of $p = u(x)$, the buyer will accept if his valuation exceeds $u(x)$. The probability of that is $1-x$. The seller's expected payoff is then $(1-x)u(x) + xs$. Of course, the seller has other strategies available. He could require the buyer to play a game in which the buyer's choices determine a probability distribution over outcomes. An outcome specifies whether a trade occurs, when it occurs, and what payments are made at which dates. The content of the next result is that a simple take-it-or-leave-it offer is as good or better than any such game.¹⁹

Proposition 4. Let x^ solve $\max(1-x)u(x) + xs$. Then making an immediate take-it-or-leave-it offer to sell at the price $u(x^*)$, with a commitment never to make another offer, maximizes the seller's expected payoff (over the class of all exchange games).*

¹⁷ The assumption of identical discount rates can be weakened to an assumption that the seller is no more patient than the buyers, without upsetting any of our results.

¹⁸ This involves no loss of generality. One can reproduce any distribution F essentially by taking $u = F^{-1}$.

¹⁹ Propositions 4 and 5 synthesize results of Harris and Raviv (1982), Milgrom (1985), Myerson (1981), Riley and Samuelson (1981), and Rubinstein, Wilson, and Wolinsky (private communication, 1985). Rubinstein, Wilson, and Wolinsky were the first to extend the optimal auction results to models in which the seller could use the threat of delays to extract a higher price from the buyer. Their analysis makes clear that Proposition 5 depends on the assumption that the seller is no more patient than the buyers.

Moreover, if $(1-x)u(x)$ is strictly concave, then the seller's payoff is maximized only by games that sell at time 0 to all buyers for whom $X > x^$ and do not sell to other buyers at any date.*

The most surprising part of this conclusion is that it does not, in general, pay the seller to use time and uncertainty for purposes of price discrimination. As a corollary, the seller cannot benefit from private information about his own valuation; he would make the same take-it-or-leave-it offer as a function of s regardless of whether s is known ex ante.

The method of analysis used to prove Proposition 4 is important and worthy of detailed study. The heart of the method is the observation that it is possible to place substantive restrictions on the allocation that can result from any Nash equilibrium of any Bayesian game. The first restriction is the so-called *incentive compatibility* constraint: Each player must prefer his own equilibrium allocation to anything he could get by pretending to be a player of another type. A second type of restriction, the *participation constraint*, reflects the assumption that the buyer cannot be forced to participate: The buyer must actually prefer participation to non-participation. As applied to the problem at hand, the incentive compatibility constraint means that the seller cannot extract a higher price from a buyer with a higher valuation unless he gives that buyer something of corresponding value in return, such as a higher probability of receiving the item or the opportunity to receive it sooner. The participation constraint means that the buyer's expected payoff must be nonnegative, regardless of his valuation of the item. These constraints imply a bound on what the seller can expect to receive at any Bayesian-Nash equilibrium of any game. The proof of Proposition 4 amounts to computing the bound and showing that it is achieved by a take-it-or-leave-it offer.

This method is most fruitful when applied to a model in which the incentive constraints take a particularly simple form. In the problem at hand, the buyer cares only about the expected discounted date at which he acquires the item $E[\delta^T]$ (where $T = \infty$ if he makes no acquisition) and the expected discounted payments \bar{e} to be made. The payoff to a buyer of valuation v is $vE[\delta^T] - \bar{e}$, a linear function of the relevant variables. The seller cares about the same things; his payoff is $\bar{e} + s(1 - E[\delta^T])$. A proof of Proposition 4 using these ideas is given in the appendix to this chapter.

Next, we introduce multiple potential buyers into the environment. Assume that there are n bidders, and that bidder i 's private information is represented by the random variable X_i . Bidder i 's valuation $V_i = u(X_i)$ depends only on his private information. Assume that u is nondecreasing and that the X_i 's are independently uniformly distributed on $(0, 1)$. This combination of assumptions defines the *independent private values model*

adopted by Vickrey (1961, 1962), Griesmer, Levitan, and Shubik (1967), and Ortega-Reichert (1968) in their pioneering studies of auction theory. The independence assumption is particularly important for the results of this section; it means that an outside observer (or the seller) can never infer anything about X_1 by observing (X_2, \dots, X_n) . We relax this assumption in Section 4.

The analysis in the multiple buyer case follows the same lines as in the single buyer case. The conclusion, however, is even more striking.

Proposition 5. Assume that $(1-x)u(x)$ is concave and let x^ denote a maximizer of $(1-x)u(x)+xs$. Then among all possible games that the bidders might agree to play, the sealed-bid and ascending-bid auctions with minimum price $u(x^*)$ maximize the seller's expected payoff.²⁰*

Proposition 5 as stated applies only to a very limited set of auction environments. However, it can be (and has been) extended in many different directions. The case of risk-averse buyers has been treated by Matthews (1983) and Maskin and Riley (1984a). Cremer and McLean (1985b) have studied a variation involving some statistical dependence. Milgrom (1985) allowed the seller to have many objects for sale, subject to some convex cost of production. Other variants can also be found. Most often, the optimal selling strategies for these more complicated environments are not recognizable auctions nor, indeed, recognizable institutions of any kind. Thus, the optimal auction theory is inadequate, by itself, to explain why auctions are used.

What is perhaps most missed in the theory of optimal auctions is some indication of which institutions for selling an object are robust – that is, optimal or nearly so in a range of environments, or at least not weakly dominated across a range of environments. Also missing is some formalization of the idea that auctions are “simple” – for example, in most auctions, all the bids that can be made actually are made with positive probability at equilibrium in several simple environments. Properties like simplicity and robustness are interesting to think about but hard to formulate; almost nothing is known about this topic.

4 Strategies for a weak seller

The models of Sections 2 and 3 go a long way toward explaining the continuing widespread use of auctions for selling many goods. However,

²⁰ As shown in Section 4, other standard auctions with the same minimum price lead to the same expected price.

these theories tell us almost nothing about the details of auctions. The deterministic models predict that all the usual sorts of auction mechanisms lead to the same outcome. Vickrey, who first introduced and used the independent private values model of Section 3, found that all the common auctions lead to the same allocation of the item and the same average price for the seller. One of the main puzzles of auction theory since Vickrey's pioneering work has been to explain when different auctions can be expected to lead to substantially different outcomes. Our purpose in this section is to review a theory that offers such an explanation.

Our analysis is based on the symmetric auction model introduced by Milgrom and Weber (1982a), which extends and unifies the earlier models of Vickrey (1961, 1962) and Wilson (1977). In the Milgrom-Weber model, each bidder i observes some private information variable X_i in (x, \bar{x}) before bidding. These observations are assumed to be drawn from some symmetric joint distribution. The value of the item to bidder i is denoted by $V_i = u(X_i, X_{-i}, S)$, where X_{-i} is the list of valuations of the other bidders and S is some vector of unobserved random variables. It is assumed that u is nondecreasing in all its arguments and is a symmetric function of the components of X_{-i} . It is also assumed that the bidders are risk-neutral.

We have already seen in Section 2 that the value of the item at auction to any bidder can depend on the valuations of other bidders when there is a possibility for resale. The vector variable S , which generally represents unknown attributes of the item, could be interpreted as the valuations of bidders not present at the auction. Another interesting interpretation of the variable S is that it represents some unknown physical attributes of the item. For example, if the item being sold is the rights to timber on a tract of land in Oregon, then the potential yield of the tract in board feet of each species of timber is normally unknown. If the right to drill for oil on some underwater tract off the north coast of Alaska is being auctioned, the value will depend on the amount and grade of the oil, its depth, future oil prices, availability of transport facilities such as pipelines (which in turn depends on the productivity of nearby tracts), and so forth.

The presence of unknown attributes (physical or otherwise) gives rise to a curious phenomenon known as the *Winner's Curse*. The idea of the curse is that inexperienced bidders will often lose money, or earn less than expected, because a bidder is much more likely to place the highest bid when he has overestimated the value of the item than when he has underestimated it. Of course, experienced bidders are aware of this phenomenon and adjust their bids accordingly, which makes studying the bidding problem an interesting exercise. Before giving a more formal account of the Winner's Curse, we must finish specifying our modeling assumptions.

We shall assume that the variables (S, X) are pairwise positively correlated on all rectangles in \mathbb{R}^{n+m} . That is, they are positively correlated conditional on any information of the form $S_i \in (s_i, \bar{s}_i)$ and $X_j \in (x_j, \bar{x}_j)$, $i = 1, \dots, m$, $j = 1, \dots, n$. Such random variables are called *affiliated*. Some key facts about affiliated random variables are briefly summarized in the next paragraph. A complete treatment, including proofs, is given in Milgrom and Weber (1982a).

Suppose the random variables $Z \equiv (S, X)$ have a joint density $f(Z)$. Then affiliation can be expressed as a property of the density f as follows:

$$f(z)f(z') \leq f(z \wedge z')f(z \vee z'), \quad (4.1)$$

where $z \wedge z'$ is a vector whose i th component is $\min(z_i, z'_i)$ and $z \vee z'$ is a vector whose i th component is $\max(z_i, z'_i)$. Note, in particular, that independent random variables are affiliated. If f is smooth and everywhere positive, affiliation is equivalent to the requirement that the $\partial^2 \ln f / \partial z_i \partial z_j \geq 0$ for all $i \neq j$. A fact about an affiliated random vector Z which is used repeatedly in auction theory is that for any nondecreasing function g , the function G defined by

$$G[(z_i, \bar{z}_i; i = 1, \dots, n)] \equiv E[g(Z) | z_i < Z_i \leq \bar{z}_i; i = 1, \dots, n] \quad (4.2)$$

is nondecreasing. In particular, if (Z_1, Z_2) is affiliated then $P\{Z_1 \leq x | Z_2 = y\}$ is decreasing in y .²¹

Given our assumptions about the bidders' information, there is an especially nice way to formalize the Winner's Curse. Suppose all the bidders $j \neq 1$ choose bids in a sealed-bid auction as functions $\beta_j(X_j)$ of their information. Suppose each β_j is increasing. Finally, suppose that bidder 1 submits a bid of b and wins. When the bidder learns that he has won, how should he evaluate his winnings? The answer is that he should always revise his estimate of value downward from his initial estimate:

$$\begin{aligned} E[V_1 | X_1, \max_j \beta_j(X_j) < b] &\leq E[V_1 | X_1, \max_j \beta_j(X_j) < \infty] \\ &= E[V_1 | X_1] \end{aligned} \quad (4.3)$$

[where we have used the fact (4.2) that conditional expectations of monotone functions of affiliated variables are monotone functions of the conditioning variables]. In simple English, learning that others have bid less than b is bad news about the value of the item being acquired [see Milgrom (1981) for a more complete analysis]. Of course, at equilibrium, bidders will take this fact into account in advance when choosing their strategies.

²¹ That probability is equal to $1 - E[1_{\{Z_1 > x\}} | Z_2 = y]$ which, by (4.2), is decreasing in y .

Consider a sealed-bid auction with a zero minimum price. We wish to represent this formally as a game. The players are the n bidders. Each bidder i observes X_i and decides what to bid. A strategy is therefore a function $\beta_i(X_i)$ specifying how much to bid as a function of what the bidder knows. Given any realization of the vector X , the item will be sold for the price $\max_i \beta_i(X_i)$ to the player who submits that bid.

Next, consider a Dutch auction game. In a Dutch auction, the auctioneer starts the price at some very high level, and reduces it until some bidder shouts "Mine!" to claim the item. Although there are complicated ways to describe any bidder's strategic options, all amount to saying that, as a function of X_i , player i must decide how far to let the price fall before shouting "Mine!" Suppose that bidder i 's strategy is to let the price fall to $\beta_i(X_i)$ and then shout "Mine!" Then, the item will be sold for a price of $\max_i \beta_i(X_i)$ to the bidder who chose that maximum level. Remarkably, *in strategic form, the Dutch and sealed-bid auctions are the same game!*

It was Vickrey (1961, 1962) who first noted this equivalence; he also claimed that the standard ascending-bid auction is equivalent to a particular sealed-bid auction. He reasoned as follows. Suppose that in the ascending-bid auction bidder i decides to bid up to the level $\beta_i(X_i)$, which we shall call i 's "bid." Then the item will be sold for the second highest bid to the high bidder. This is the same as conducting a sealed-bid auction in which the item is awarded to the high bidder for the second highest bid. Actually, this analysis is not quite correct, because the bidders in an ascending-bid auction have additional strategies available: They can make their bids depend on the previous bids of the other bidders. Nevertheless, in the interest of brevity and simplicity, we shall adopt Vickrey's "second price auction" as a model of the ascending-bid auction. The results we obtain are not affected in any essential way by this modeling.²²

One can show for the model we have described that there are unique increasing strategies β_S and β_A such that $(\beta_S, \dots, \beta_S)$ is an equilibrium of the sealed-bid auction game, and $(\beta_A, \dots, \beta_A)$ is an equilibrium of the ascending-bid auction game, and that these strategies are characterized by solutions to first-order conditions.²³

²² Milgrom and Weber (1982a) distinguish the Vickrey second price sealed bid auction from the English ascending-bid auction. In the latter, they assume, bidders can base their bidding decisions on the levels where other bidders ceased to be active. Such strategies require making complicated inferences in real time, and in any case their equilibria share many of the same properties as those studied here.

²³ Here, I ignore the possibility of multiple equilibria, and focus attention exclusively on a monotone, symmetric equilibrium. The uniqueness problem and related issues are taken up by Milgrom (1981), Maskin and Riley (1983), and Harstad and Levin (1985).

How is one to compare the expected revenues from these two kinds of auctions? How can one evaluate the impact of revealing information on the expected selling price in either kind of auction? The main tool for this analysis is the Linkage Principle.

The sealed-bid and ascending-bid auctions are examples of what we shall call *standard* auctions – auctions in which (at equilibrium in the specified environments) the highest evaluator wins and pays a nonnegative price, and in which losers always pay zero. Given a standard auction A , let P^A be the resulting price; in general P^A might depend on all the bids. Let $\pi^A(z, x)$ be the expected payment by a bidder, say bidder 1, who bids in auction A as if his estimate were z when his actual estimate is x . Letting $Y = \max_{i \neq 1} X_i$, we may write $\pi^A(z, x) = E[P^A \cdot 1_{Y < z} | X_1 = x]$.

Proposition 6 (The Linkage Principle).²⁴ Let A and B be standard auctions. Suppose that for all x , if $\pi^A(x, x) < \pi^B(x, x)$ then $\pi_2^A(x, x) \geq \pi_2^B(x, x)$ (where the subscripts denote partial derivatives with respect to the second argument). Then, $\pi^A(x, x) \geq \pi^B(x, x)$ for all x , and the expected price is higher in auction A than in auction B .

Proof: Let $R(z, x) = E[V_1 \cdot 1_{Y < z} | X_1 = x]$, that is, the expected value received in a standard auction when player 1 observes $X_1 = x$ and bids as if he had observed $X_1 = z$, assuming that the other bidders adhere to their equilibrium strategies. At equilibrium, a bidder with estimate x cannot improve his payoff by bidding as if his estimate were some $z \neq x$. Hence, for any standard auction A , $R_1(x, x) - \pi_1^A(x, x) = 0$. Because this holds also for the standard auction B , $\pi_1^A(x, x) = \pi_1^B(x, x)$ for all x .

Now consider the function $\Delta(x) = \pi^A(x, x) - \pi^B(x, x)$. Since $\pi^A(\underline{x}, \underline{x}) = \pi^B(\underline{x}, \underline{x}) = 0$ (because $P\{Y < \underline{x}\} = 0$), $\Delta(\underline{x}) = 0$. For $x > \underline{x}$, by hypothesis either $\Delta(x) \geq 0$ or

$$\Delta'(x) = \frac{d}{dx} [\pi^A(x, x) - \pi^B(x, x)] = \pi_2^A(x, x) - \pi_2^B(x, x) \geq 0.$$

If there were some $\hat{x} > \underline{x}$ such that $\Delta(\hat{x}) < 0$, then by the mean value theorem there must be some $\bar{x} \in (\underline{x}, \hat{x})$ with $\Delta(\bar{x}) < 0$ and $\Delta'(\bar{x}) < 0$, which contradicts our last conclusion. ■

The Linkage Principle derives its name from the fact that it makes comparisons of revenues across auctions using the statistical linkages that exist

²⁴ The Linkage Principle was originally introduced by Milgrom and Weber (1982a), who described it as “the common thread running through” their results (pp. 110–11). The mathematics of the principle, which is buried in their arguments, was first made explicit in a second (unpublished) paper, “A Theory of Auctions, Part II.”

between a bidder's private information and any information on which the price is based. The partial derivative π_2 measures how much the expected price increases with increases in the bidder's private evaluation, holding his bidding behavior constant; that is, it measures changes in π that result from changes in the bidder's beliefs about how others will bid. When statistical linkages are used to align the price more closely to the buyer's willingness to pay, the rents associated with private information fall and the average price rises.

Our first application of the Linkage Principle is to explain the equivalence in expected prices that Vickrey first observed among the standard auctions in his model.

Proposition 7 (Revenue Equivalence). Assume that the bidders' observations X_1, \dots, X_n are independent. Then the expected price in any standard auction is the same as in the ascending-bid auction.²⁵

Proof: By independence, for any auction mechanism M :

$$\pi^M(z, x) = E[P^M \cdot 1_{Y < z} | X_1 = x] = E[P^M \cdot 1_{Y < z}].$$

Hence, for all auction games, $\pi_2^A(z, x) = \pi_2^B(z, x) = 0$. The result then follows from the Linkage Principle. ■

When we replace the hypothesis of statistical independence by one of affiliation, the Linkage Principle provides a powerful tool for making comparative statements about the expected prices under various alternative arrangements. For a first example, we can now compare the expected price in the sealed- and ascending-bid auctions.

Proposition 8. The expected price (at equilibrium) in the ascending-bid auction is never less, and is sometimes more, than for the sealed-bid auction.

Proof: Let A be the ascending-bid auction and B the sealed-bid auction, with equilibrium bidding strategies α and β , respectively. These auctions are standard, because α and β are increasing (Milgrom and Weber, 1982a). Now $\pi^B(z, x) = E[\beta(z) \cdot 1_{Y < z} | X_1 = x] = \beta(z) P\{Y < z | X_1 = x\}$, which we write as $\beta(z)F(z|x)$. Also,

$$\pi^A(z, x) = E[\alpha(Y) \cdot 1_{Y < z} | X_1 = x] = E[\alpha(Y) | X_1 = x, Y < z]F(z|x).$$

²⁵ For a more general revenue equivalence result and a helpful discussion of its implications, see Riley and Samuelson (1981).

Hence:

$$\frac{\partial}{\partial x} [\pi^A(z, x) - \pi^B(z, x)] = [\pi^A(z, x) - \pi^B(z, x)] \frac{F_2(z|x)}{F(z|x)} + F(z|x) \frac{d}{dx} E[\alpha(Y) | X_1 = x, Y < z].$$

By affiliation, $(d/dx)E[\alpha(Y) | X_1 = x, Y < z]$ is nonnegative and $F_2(z|x)$ is nonpositive, so the partial derivative is nonnegative unless $\pi^A(z, x) > \pi^B(z, x)$. Apply the Linkage Principle. ■

In the sealed-bid auction, the price the winning bidder pays depends only on his own bid; there is no direct or indirect linkage to other variables. In the ascending-bid auction, the price is an increasing function of the second highest bidder's observation, which is affiliated with the winning bidder's observation. This provides the price-increasing linkage.

When the seller has a private estimate that he can provide, the equilibrium price may be an increasing function of that estimate, providing yet another price-increasing linkage. Of course, this assumes that the seller can both provide information in a verifiable way and commit to a policy for revealing information.

Proposition 9. Verifiably revealing any information variable S_0 raises the expected price both in the sealed-bid auction and in the ascending-bid auction. Among all policies for full or partial revelation of information, the policy of full revelation maximizes the expected price.

Partial proof: For the sealed-bid auction A without information revealed and B with information revealed, the expected prices are

$$\pi^A(z, x) = E[\alpha(z) \cdot 1_{Y < z}] \quad \text{and} \quad \pi^B(z, x) = E[\beta(z, S_0) \cdot 1_{Y < z}],$$

where α and β are the equilibrium bidding strategies. As before,

$$\frac{\partial}{\partial x} [\pi^B(z, x) - \pi^A(z, x)] = [\pi^B(z, x) - \pi^A(z, x)] \frac{F_2(z|x)}{F(z|x)} + F(z|x) \frac{d}{dx} E[\beta(z, S_0) | X_1 = x, Y < z].$$

This is nonnegative whenever $\pi^B(z, x) < \pi^A(z, x)$. Thus, by the Linkage Principle, revealing information raises the average price.

Applying this result to the case where only partial information is revealed, the expected price can be raised by revealing the remaining information. Hence, no policy can be better than the policy of full revelation of the available information.

For the ascending bid auction, the price depends on both Y and S_0 . The proof given by Milgrom and Weber (1982a) uses the characterization of the equilibrium strategies to show that the linkage is greater when information is revealed than when it is withheld (even though the linkage to Y may be weakened). ■

The policy of full revelation can also be expected to emerge when the seller cannot commit to an information policy.

Proposition 10. If a seller must decide, after observing S_0 , whether to report it, and if his report is verifiable, then at a perfect equilibrium he always reports S_0 regardless of its value.

Proof sketch: Suppose that a seller observes S_0 and then decides whether to report it. At equilibrium, the buyers' bids (when the seller makes no report) depend on their beliefs about the value of S_0 . For any beliefs the buyers may have, one can show that the seller's best response is to make a report whenever S_0 is in fact sufficiently favorable; thus any equilibrium must have the property that the seller reports whenever $S_0 > s^*$. But then, when the seller reports nothing, at equilibrium the buyers must believe that $S_0 \leq s^*$ and bid accordingly. They would bid strictly more if they had the "more favorable" [in the sense of Milgrom (1981)] belief that $S_0 = s^*$, so it would be in the seller's interest to report S_0 even when it is slightly lower than s^* . Hence there can be no perfect equilibrium at which s^* exceeds the lower bound \underline{s} of the support of S_0 . Because $P\{S_0 > \underline{s}\} = 1$, the proposition is proved. ■

In a further application of the Linkage Principle, Riley (1985) argues that when value can be observed ex post (even imperfectly), the expected price is higher when part of the price is a royalty based on the observed value. McAfee and McMillan (1986) make the same observation in connection with incentive contracting, where bidding and moral hazard issues arise together. Consider a situation in which a buyer must select one of several contractors for, say, a construction project. Including a cost-sharing provision in a contract yields a lower price at the bidding stage (apply the Linkage Principle), but that must be balanced against the weakened incentive for cost control that such contracts may create. This examination of the relation between bidding and contract incentives is one of the most promising recent developments in bidding theory.

Even when a seller has little ability to enforce a high minimum price, he can still choose any auction with a zero minimum price because, as noted in Section 2, rational buyers cannot refuse to participate. If the bidders behave noncooperatively, the seller does better by using an ascending-bid

auction than by soliciting sealed bids. This accords well with Cassady's (1967) observation that ascending-bid auctions are by far the most popular kind worldwide. The seller also does well to provide any information he may have, since that, too, creates a price-increasing linkage.

In terms of incentive theory, the Linkage Principle is based on the observation that *a bidder's profits depend upon his ability to conceal information*. Linking the price to variables that are affiliated with the bidder's private information diminishes his ability to conceal information effectively, and so lowers his profits. When comparing auctions (with risk-neutral bidders and a risk-neutral seller) that allocate the good efficiently, any reduction in the bidders' payoffs is a gain to the seller.

Clearly, this reasoning depends on the assumption that bidders are risk-neutral. With risk aversion, there can be efficiency gains from making the bidders' payoffs less random. This idea has heretofore been studied only in connection with the independent private values model. For that model, Matthews (1980) noted that if the bidders' observations are statistically independent and the valuation function is $V_i = X_i$, then the sealed-bid auction is always preferred by the seller when either he or the buyers (or both) are risk-averse [see also Holt (1980)].

In the independent private values model, linkages inefficiently increase the randomness in the payoffs. In other models, however, linkages reduce the randomness in the payoffs. For example, suppose the observations are independent and the valuation function is $V_i = \min(X_i, \max X_j)$. Then the ascending-bid auction is always preferred by the seller when the bidders are risk-averse. In this example, the linkage of the price to other bidders' information reduces the fluctuations in the winning bidder's payoffs and makes him willing to pay more, to the seller's benefit. In general, in the presence of risk aversion, linkages may or may not enhance revenue and efficiency.

Finally, we turn to the question of whether uncertainty about the bidders' valuations makes it harder or easier for the seller to achieve commitment in his efforts to maintain a high minimum price. Consider a model in which the seller conducts a series of auctions, specifying any minimum price he chooses, and the buyers are limited to bidding in the auction; they cannot make extraneous offers. In the deterministic case, the seller can extract all the surplus from the highest valuation buyer by insistently setting the minimum price equal to that valuation. Indeed, in the discrete time version of this game model, the equilibrium just described is the only perfect equilibrium.

With uncertainty, however, the situation is quite different. The case of a single buyer and offers made by the seller has been analyzed by Stokey (1982), and by Gul, Sonnenschein, and Wilson (1986). In the discrete time

game where offers are made by the seller, there is a unique perfect equilibrium. If the buyer's reservation value is uncertain to the seller and distributed over an interval that includes the seller's reservation value, then at equilibrium the seller must sell for nearly his own reservation price: Uncertainty is the enemy of commitment. Gul and Sonnenschein (1985) have proved a variant of this result for the case where both the seller and the buyer can make offers.

Here we make a simple extension of this conclusion to the case of many buyers. Suppose that the seller can conduct an auction at each moment of time, and can vary the minimum price $m(t)$ over time. For technical convenience, we require that the seller choose a path $m(t)$ that is right continuous. A buyer's strategy specifies whether to bid and an amount to bid, for each moment in time, as a function of the path of minimum prices announced by the seller up to that point. We limit the bidders to strategies that determine a first moment to bid (possibly $+\infty$) for any feasible strategy of the seller. For example, the buyer cannot specify that he will make a bid whenever the minimum is below \$5 and has been below \$5 before, because that does not specify a first moment to bid if the seller sets a minimum of \$4 at all times. Finally, suppose that payoffs from trades conducted at any date $t > 0$ are discounted to time zero at the same rate for buyers and the seller.

Proposition II. An equilibrium of the continuous auction game is described as follows: The seller sets a minimum price equal to his reservation price s at every point in time. A buyer with observation $X_i = x$ bids $\beta_S(x)$ at the first moment that $m(t) \leq \beta_S(x)$, where β_S is the symmetric equilibrium strategy for the (static) sealed-bid auction.

During the play of the game, if the seller sets a minimum price other than s , a player using the prescribed strategy will bid if and only if his planned bid exceeds that minimum. Therefore, by watching the game progress, the seller and the other bidders could learn about the bidders' valuations – specifically, that they are too low to justify bidding at the prices named by the seller. This sort of learning is perfectly analogous to what occurs in a Dutch auction. As with Dutch auctions, no matter what upper bound on valuations a bidder may learn during the course of the auction, it will always be optimal for him to adhere to his equilibrium strategy, “because” he expects the seller to reduce the price to s very soon.

It is easy to see that the optimal auction, which requires that the seller make a once-and-for-all take-it-or-leave-it offer, is not a perfect equilibrium of this continuous time game. Suppose that the seller expects the

bidders to adhere to the optimal auction equilibrium strategies. The optimal auction offer always entails setting a minimum price in excess of the seller's reservation price. Hence, if the seller makes the optimal auction offer and no bidder bids, it will always pay the seller to make another, better offer to the buyers. Hence, the seller can conduct an optimal auction only if he can commit himself to refrain from making a profitable offer later. Such commitment is proscribed by the perfect equilibrium solution concept.

By imposing a plausible restriction on the buyers' strategy spaces, one can eliminate much more than just the optimal auction equilibrium in this continuous time game. Indeed, if the bidders are limited to strategies that satisfy the Gul and Sonnenschein (1985) stationarity property, then the equilibrium of Proposition 11 is the unique symmetric (among buyers) equilibrium. The stationarity property requires that a buyer's current decision depend only on his own type and on the "common knowledge" distribution of buyer types (which, at equilibrium, can be inferred from past bidding behavior); in particular, the behavior does not depend on the time nor on how those beliefs were reached. When the buyers are limited to stationary strategies, if all trading does not take place at time zero then it will always be in the seller's interest to "speed up the clock," replacing his strategy $m(t)$ by $\hat{m}(t) = m(2t)$. This stratagem makes all trades occur earlier, without reducing the price obtained by the seller. Therefore, at any equilibrium in stationary strategies, all trading occurs at time zero. Then, because the seller cannot adhere to a minimum price above his reservation price of s , the uniqueness result follows.

Thus, with uncertainty about the buyers' valuations, the seller's ability to achieve commitment is severely reduced. Whatever power the seller retains comes from his ability to generate competition among bidders by conducting an auction.

5 Collusion

Only recently has any theoretical attention been devoted to the problem of collusion in auctions. The fact that the outcome of auctions in deterministic settings lies in the core of the exchange game does not mean that auctions are immune to collusion; it means only that no subset of players could improve their lot by going off and trading among themselves. Auctions with low minimum prices are vulnerable to collusion among bidders. Graham and Marshall (1985), beginning with that observation and a claim that collusion is rampant in real auctions, have studied a variation of the independent private values model of Section 3 in which the bidders may have formed a cartel or *ring*. They find that the optimal minimum price

to be set by a seller is an increasing function of the likelihood that a ring has formed.

My purpose here is simply to examine the hypothesis that auction forms differ in their degrees of susceptibility to collusion. I will focus on Mead's (1967) hypothesis that ascending-bid auctions are more susceptible to collusion than are sealed-bid auctions. Such a conclusion would explain why a seller might choose a sealed-bid auction in preference to an ascending-bid auction, despite the latter's theoretical superiority when bidders behave competitively. The simplest model (not involving side payments) that I have found to study collusion is the following; it exploits the existence of multiple Nash equilibria in ascending-bid auctions to construct collusive perfect equilibria in repeated ascending-bid auctions.²⁶

The model is deterministic. We suppose that two bidders bid periodically against one another in an auction for items that both bidders value at x . Suppose they agree to take turns winning at a price of $b < x$. The discount factor that reflects the frequency of these periodic interactions is some number $\delta < 1$. How frequent must the interactions be to support this collusive arrangement? That is, how large must δ be to allow collusion of this sort to survive at an equilibrium?

In the case of a sealed-bid auction, suppose it is an equilibrium for the players to alternate bidding b , while the other bidder makes a show of it by bidding $b - \epsilon$. To be an equilibrium, it must be unprofitable for the scheduled loser to bid just more than b today and forgo future profits:²⁷

$$x - b < \delta \frac{x - b}{1 - \delta^2},$$

which reduces to $\delta > (\sqrt{5} - 1)/2$ or, approximately, $\delta > .62$. One corresponding collusive agreement in the ascending-bid auction has the scheduled winner bid x and the scheduled loser bid b . For the scheduled loser to find a deviation unprofitable requires only that $\delta > 0$.

Thus, collusion is easier to support in an ascending-bid auction than in a sealed-bid auction. The intuition for this result is a familiar one: Collusion is hardest to support when secret price concessions are possible, and easiest to support when all price offers must be made publicly.

²⁶ Bikhchandani (1984) has exploited the fact that the one-shot ascending bid auction has multiple equilibria to construct a repeated bidding game in which the equilibria have a collusive flavor. Robinson (1985) has also exploited the multiple equilibrium idea. He analyzes a simple one-shot game model, but the analysis assumes that colluders can share information verifiably - an assumption that naturally favors the formation of collusive rings.

²⁷ There are some minor issues here about whether the inequalities given below should be strict or weak. These depend on the solution concept used, and are not a matter of great importance here.

6 Conclusion

I have organized this paper around two central questions. Why do auction institutions continue to be so popular after thousands of years? And what accounts for such particular details as the popularity of sealed- and ascending-bid auctions? The answers to these questions were summarized in the introduction. The answers are plainly incomplete; indeed, they rely on fundamentally different models of the auction environment. It would be much better to have a single consistent model that explained both the use of auctions (in preference to other mechanisms when individual items are unique) as well as the widely observed preference for ascending-bid auctions over sealed-bid auctions.

There are large segments of auction theory that I have omitted from my review, partly because a surveyor must draw lines. One could survey very different territory by asking questions such as the following. (1) How do experimental subjects behave in auctions? Such a survey would feature the work of Vernon Smith and his colleagues, who have led the way in the study of bidding behavior with controlled laboratory experiments.²⁸ It would also cover studies of the implications of alternative models of choice under uncertainty for bidding behavior.²⁹ (2) What insights does auction theory offer into the problems of price discrimination by a monopolist? Many researchers who have contributed to the theory of expected price-maximizing auctions have extended their results to the problem of optimal price discrimination [Cremer and McLean (1985a), Harris and Raviv (1981), Maskin and Riley (1984b)]. (3) What are the relationships between bargaining theory, auction theory, and competitive equilibrium theory? Wilson's companion survey gives an introduction to this new and lively area of research (see Chapter 2, this volume), sometimes called the

²⁸ An excellent example of this line of research is the work reported in Cox, Smith, and Walker (1984).

The proper interpretation of their experimental results is controversial. The experimenters generally regard it to be evidence concerning how actual bidders behave in auctions. However, I tend to regard it as another kind of model: the subjects (instead of rational maximizers) are the model of actual bidders. This is analogous to comparing mathematical models of air foils with corresponding scale models tested in wind tunnels; often, the mathematical models are better predictors.

In auctions for mineral rights, the bidders (oil company executives) normally have access to professional consultants who can conduct formal analyses with much more proficiency than the typical subject in an experiment. It seems likely that these executives bid much more rationally than typical experimental subjects, although this supposition is of course subject to empirical refutation.

For a carefully articulated view of the role of experiments in economics, see Roth (Chapter 7, this volume).

²⁹ For example, Karni and Safra (1985) have a model of bidding behavior in which the bidders may behave differently in strategically equivalent games.

theory of market microstructure. Our Propositions 2, 3, and 11 introduce some of the main issues studied in that connection. (4) How are auctions used in contracting environments, where the risk of cost overruns may need to be shared and the bidder's performance in the contract properly motivated? This is a new subject of study that so far has limited itself mostly to private value auction models, but which is evidently generalizable to many important bidding situations.

The list of possible questions is endless, but this survey is not.

Appendix

Proof of Proposition 5: Suppose the seller designs some game Γ in which, at equilibrium, trade takes place at some date T that may depend (possibly probabilistically) on some choice made by the buyer. If no trade takes place, let us say that $T = \infty$. Suppose Σ is the set of strategies available to the buyer. For each $\sigma \in \Sigma$, let $p(\sigma) = E_\sigma[\delta^T]$ be the expected present value of 1 unit paid at the time of trade. This expectation depends, of course, on the strategy σ chosen by the buyer. Similarly, let $e(\sigma)$ be the expected present value of net payments made by the buyer over the course of the game. If the buyer's information is X , his expected payoff using σ is $u(X)p(\sigma) - e(\sigma)$. Let $\sigma^*(x)$ and $\Pi^*(x)$ denote, respectively, the buyer's optimal strategy and the maximum payoff in the game when $X = x$. Define $p^*(x) \equiv p(\sigma^*(x))$ and $e^*(x) \equiv e(\sigma^*(x))$. Then,

$$\Pi^*(x) = u(x)p^*(x) - e^*(x). \quad (\text{A.1})$$

As Vickrey originally argued, $p^*(x)$ must be nondecreasing; otherwise the buyer must be using a dominated strategy $\sigma^*(x)$. He could increase his ex ante expected payoff by rectifying his strategy to make $p^*(x)$ nondecreasing while holding the distribution of $\sigma^*(X)$ fixed, since that leaves his expected payment and probability of winning unchanged while increasing the expected value received. By the Envelope Theorem, $d\Pi^*/dx = u(x)p^*(x)$, using equation (A.1), $de^*(x) = u(x)dp^*(x)$. Hence, $e^*(x) = e^*(0) + \int_0^x u(t) dp^*(t)$. Now the seller's expected cash receipts, conditional on x , are $(1 - p^*(x))s + e^*(x)$. Because X is uniformly distributed on $(0, 1)$, the corresponding unconditional expectation is:

$$\begin{aligned} \int_0^1 e^*(x) dx &= e^*(0) + \int_0^1 \int_0^x u(t) dp^*(t) dx \\ &= e^*(0) + \int_0^1 \int_t^1 dx u(t) dp^*(t) \\ &= e^*(0) + \int_0^1 (1-t)u(t) dp^*(t). \end{aligned} \quad (\text{A.2})$$

In addition, the seller obtains value from keeping the item:

$$\begin{aligned} s \int_0^1 (1-p^*(x)) dx &= s \int_0^1 \left[1-p^*(1) + \int_x^1 dp^*(t) \right] dx \\ &= s \left[1-p^*(1) + \int_0^1 \int_0^t dx dp^*(t) \right] \\ &= s \left[1+p^*(1) + \int_0^1 t dp^*(t) \right]. \end{aligned} \quad (\text{A.3})$$

The seller's total expected payoff is therefore:

$$e^*(0) + s[1-p^*(1)] + \int_0^1 [(1-t)u(t) + st] dp^*(t). \quad (\text{A.4})$$

Because the buyer must have a strategy σ of nonparticipation, which leads to a payoff of zero, we may conclude:

$$e^*(0) \leq p^*(0)u(0). \quad (\text{A.5})$$

Inequality (A.5) constrains the seller in designing a game. Another constraint that (by Vickrey's argument cited earlier) must hold at equilibrium is:

$$p^*: [0, 1] \rightarrow [0, 1] \text{ is nondecreasing.} \quad (\text{A.6})$$

Because the objective (A.4) is linear in p^* and $e^*(0)$, and because the constraint set (A.5)–(A.6) is convex, the maximum must occur at an extreme point of the constraint set. Thus, at a maximum, (A.5) must hold with equality. Also, the optimal p^* function must be everywhere 0 or 1. So, by (A.6), there is a point x^* such that $p^*(x)$ is 0 for $x < x^*$ and 1 for $x \geq x^*$. The maximized value of the seller's payoff (A.4), subject to (A.5)–(A.6), is $sx^* + (1-x^*)u(x^*)$.

This maximum, which bounds what the seller can get at any equilibrium of any game, can be achieved by making the take-it-or-leave-it offer $u(x^*)$. Moreover, if $(1-x)u(x)$ is strictly concave, there is a unique x^* that maximizes $sx + (1-x)u(x)$. Hence, the unique $p^*(\cdot)$ function that attains the maximum is $p^*(x) = 0$ for $x < x^*$ and $p^*(x) = 1$ for $x \geq x^*$. Any institution as good as making a take-it-or-leave-it offer must lead to the same trading outcome. ■

References

- Aumann, Robert. 1973. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1: 67–96.
- Bikhchandani, Sushil. 1984. "Reputations in Repeated Second-Price Auctions." Working paper, Graduate School of Business, Stanford University.

- Binmore, Ken. 1983. "Bargaining and Coalitions I." London School of Economics, International Centre for Economics and Related Disciplines, Theoretical Economics Discussion Paper Series, 83/71.
- Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky. 1986. "The Nash Bargaining Solution in Economic Modelling." *Rand Journal of Economics* 17: 176–88.
- Cassady, R. 1967. *Auctions and Auctioneering*. Berkeley: University of California Press.
- Cox, James, Vernon Smith, and James Walker. 1984. "Theory and Behavior of Multiple Unit Discriminative Auctions." *The Journal of Finance* 39: 983–1010.
- Crawford, Vincent. 1982. "A Theory of Disagreement in Bargaining." *Econometrica* 50: 607–38.
- Cremer, Jacques, and Richard McLean. 1985a. "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent." *Econometrica* 53: 345–61.
- 1985b. "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions." CARESS Working Paper #85-17, University of Pennsylvania.
- Graham, Daniel, and Robert Marshall. 1985. "Collusive Bidder Behavior at a Single Object English Auction." Department of Economics Working Paper No. 85-01, Duke University.
- Griesmer, J., R. Levitan, and M. Shubik. 1967. "Toward a Study of Bidding Processes: Part IV: Games with Unknown Costs." *Naval Research Logistics Quarterly* 14: 415–33.
- Gul, Faruk, and Hugo Sonnenschein. 1985. "One-Sided Uncertainty Does Not Cause Delay." Mimeo, Stanford University.
- Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. 1986. "Foundations of Dynamic Monopoly and the Coase Conjecture." *Journal of Economic Theory* 39: 155–90.
- Harris, Milton, and Arthur Raviv. 1981. "A Theory of Monopoly Pricing With Uncertain Demand." *American Economic Review* 71: 347–65.
1982. "Allocation Mechanisms and the Design of Auctions." *Econometrica* 49: 1477–99.
- Harstad, Ronald, and Dan Levin. 1985. "A Dominant Strategy Argument for a Class of Common Value Auctions." *Review of Economic Studies* 52: 525–8.
- Holt, Charles A., Jr. 1980. "Competitive Bidding for Contracts Under Alternative Auction Procedures." *Journal of Political Economy* 88: 433–45.
- Karni, Edi, and Zvi Safra. 1985. "English and Vickrey Auctions in the Theory of Expected Utility With Rank-Dependent Probabilities." Working Paper, Tel-Aviv University.
- McAfee, Preston, and John McMillan. 1986. "Bidding for Contracts: A Principal-Agent Analysis." *Rand Journal of Economics* 17: 326–38.
- Maskin, Eric, and John Riley. 1983. "Uniqueness of Equilibrium in Open and Sealed Bid Auctions." Unpublished manuscript.
- 1984a. "Optimal Auctions with Risk Averse Buyers." *Econometrica* 52: 1473–1518.
- 1984b. "Monopoly With Incomplete Information." *Rand Journal of Economics* 15: 171–96.

- Matthews, Steven. 1980. "Risk Aversion and the Efficiency of First and Second Price Auctions." Unpublished manuscript, Northwestern University.
1983. "Selling to Risk Averse Buyers with Unobservable Tastes." *Journal of Economic Theory* 30: 370-400.
- Mead, Walter. 1967. "Natural Resource Disposal Policy - Oral Versus Sealed Bids." *Natural Resource Planning Journal* M: 195-224.
- Milgrom, Paul. 1981. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica* 49: 921-43.
1985. "The Economics of Competitive Bidding: A Selective Survey." In L. Hurwicz, D. Schmeidler, and H. Sonnenschein (eds.), *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, pp. 261-89. Cambridge: Cambridge University Press.
- Milgrom, Paul, and Robert Weber. 1982a. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50: 1089-1122.
- 1982b. "A Theory of Auctions: Part II." Unpublished manuscript.
- Myerson, Roger. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6: 58-73.
- Ortega-Reichert, Armando. 1968. "Models for Competitive Bidding Under Uncertainty." Technical Report No. 8, Department of Operations Research, Stanford University.
- Riley, John. 1985. "Ex Post Information in Auctions." Department of Economics Working Paper #367, University of California-Los Angeles.
- Riley, John, and William Samuelson. 1981. "Optimal Auctions." *American Economic Review* 71: 381-92.
- Robinson, Marc. 1985. "Collusion and the Choice of Auction." *RAND Journal of Economics* 16: 141-5.
- Selten, Reinhard. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory* 4: 25-55.
- Stokey, Nancy. 1982. "Rational Expectations and Durable Goods Pricing." *Bell Journal of Economics* 12: 112-28.
- Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16: 8-37.
1962. "Auctions and Bidding Games." In *Recent Advances in Game Theory*, pp. 15-27. Princeton: Princeton University Press.
- Wilson, Robert. 1977. "A Bidding Model of Perfect Competition." *Review of Economic Studies* 4: 511-18.

Game-theoretic analyses of trading processes

Robert Wilson

Abstract: Three topics are discussed. The first is a research program to establish whether the familiar trading rules, such as sealed-bid and oral double auctions, are incentive efficient over a wide class of economic environments. The second is a review of recent studies of dynamic trading processes, and particularly the effects of impatience and private information on the timing and terms of trade; the main emphasis is on models of bilateral bargaining. The third considers prospects for embedding bargaining and auction models in larger environments so as to endogenize traders' impatience as a consequence of competitive pressures; models of dispersed matching and bargaining and a model of oral bid-ask markets are mentioned.

Introduction

My aim in this chapter is to describe some developments in the theory of exchange. The topics I describe share a common focus, namely the determination of the terms of trade. They also share a common methodology: the application of game theory to finely detailed models of trading processes. The aim of this work is to establish substantially complete analyses of markets taking account of agents' strategic behavior. Typically the results enable two key comparisons. One is the effect of altering the trading rules, and the other is the effect of alterations in the environment, such as changes in the number, endowments, preferences or information of the participants. Beyond these comparisons, however, the results are building blocks in the construction of a genuine theory of price forma-

I am indebted to many colleagues for their interest, and to Peter Cramton, Drew Fudenberg, John Roberts, and Ariel Rubinstein for their specific comments on an earlier draft. I especially thank Faruk Gul and Hugo Sonnenschein for the pleasure of working together. Errors and omissions are mine. Research support from the National Science Foundation (SES-83-08-723) and the Office of Naval Research (N00014-79-C-0685) is gratefully acknowledged.