

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 779

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he had access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

AUCTION THEORY

Paul R. Milgrom

November 1985

AUCTION THEORY

Paul R. Milgrom
Yale University
135 Prospect St.
New Haven, Connecticut 06520

Prepared for the 1985 World Congress
of the Econometric Society

October 30, 1985
Comments Welcomed!

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 and Ariel Rubinstein were invaluable to me in clarifying the relation between auctions and bargaining.

AUCTION THEORY

by 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

¹Herodotus 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.

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

who are desperate for cash and willing to sell even at prices far below replacement cost.

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 auction institutions 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 on.

Logically prior to explaining the use of auctions is defining just what an auction is. The characteristic feature of bidding institutions 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; that affords the seller an opportunity to compare the timing of buyers' offers, and to take the offer that is made earliest. Each of these bidding institutions can be contrasted with, say, a bargaining process in which the seller negotiates one-by-one with a sequence of buyers who make short-lived offers, so that the seller has no opportunity to compare the offers of different buyers.

The simplest explanation of the continuing popularity of auctions is that auctions 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 that 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 of this paper, which focuses on deterministic auction models. A third explanation, reviewed in section 3, is that 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

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

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 there is correlated uncertainty on the part of the bidders, different auctions 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 thing, he can normally choose what kind of auction to offer, provided the minimum bid is kept low, because buyers will always

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

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

want to participate in the auction.⁶ Normally, the seller can also decide whether to reveal any information he has 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 prefers (i) to conduct an ascending-bid auction rather than a sealed-bid auction, (ii) to reveal all the information that he has available, and (iii) to link the price to any available exogenous indicators of value.

The analysis leading to these conclusions is founded on 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 that yield 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, the way privacy is undermined is by linking the price to information other than the winning bidder's private information but which is correlated with his information.

⁶No 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 the 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.)

⁷The principal assumptions required include risk neutrality, symmetric uncertainty about the bidders' valuations, and non-negative correlation (actually, affiliation) among the bidders' valuations.

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).

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, which on average increases the expected price (if these indicators are affiliated with the bidders' information). Examples of contracts let at auction in which the price is determined in part by *ex post* indicators are 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

⁸Random 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.

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. By this I mean 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. All of this is assumed to be common knowledge among the buyers and the seller. 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 zero,⁹ what will happen? Using any sensible equilibrium concept, such as Nash equilibrium in undominated strategies,¹⁰ perfect equilibrium, "rationalizable" strategies, or even correlated equilibrium, the item will be sold to bidder 1 for his bid of b_2 .¹¹ The same trade will occur if the seller sets any minimum price

⁹We assume in this auction and all those considered below that ties are broken by tossing a fair coin.

¹⁰Although 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 (i) each player is maximizing given his beliefs about how the others will play and (ii) 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 like auctions, which have existed for millennia and so for which expectations can be based on actual experience.

¹¹Any 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 one bids b_2 . Bidder 2 uses any mixed strategy F satisfying two conditions. First, $F(b_2^-) = 1$. Second, let $G(x) = [F(x) + F(x-1)]/2$. Then,

$$G(x) \leq (b_1 - b_2)/(b_1 - x) \text{ for all } x \in (m, b_2).$$

With more than two bidders, one can specify the strategies of the other arbitrarily, provided bidder j always bids less than b_j , and this remains a perfect equilibrium.

Rationalizable strategies are derived by the 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, too.

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 $m > b_1$, 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 (that is, non-negative) and corresponds to some feasible allocation of the goods and money among the players. To be in the core, an imputation must be efficient and must satisfy inequalities asserting that no

coalition could, by agreeing to exchange among themselves, earn a higher total payoff:

$$(2.1) \quad x_0 + x_1 + \dots + x_n = b_1, \text{ and}$$

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

In view of the preceding discussion, the Proposition asserts that the core consists entirely of points of the form $(x_0, b_1 - x_0, 0, \dots, 0)$ for $\max(s, b_2) \leq x_0 \leq b_1$. It is easy to check that all such points satisfy efficiency (2.1) and the inequalities (2.2), and so do, in fact, lie in the core.

Conversely, suppose x lies in the core. Then $x_0 + x_1 \geq b_1$. This, together with efficiency and non-negativity, implies that $x_0 + x_1 = b_1$ and that $x_2 = \dots = x_n = 0$. So, all points in the core are of the form $(x_0, b_1 - x_0, 0, \dots, 0)$. Also, $x_0 + x_2 \geq \max(s, b_2)$, so $x_0 \geq \max(s, b_2)$. Non-negativity of x_1 implies $x_0 \leq b_1$. \square

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. With discrete bid spaces with bid increment ϵ , the only subgame 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^- .¹²

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

¹²One could, of course, define a modified sealed-bid auction game which 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 the previous footnote) shows that the two games are not "equivalent" for the purposes of perfect equilibrium analysis.

bargaining power.¹³ Indeed, the case $n = 1$ is just a bargaining problem, and auction theory, like the theory of the core, predicts only that the outcome will be efficient and 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, auction theory is surprisingly insensitive to the bargaining theory used at its foundations. To show that, we embed the auction model in a general noncooperative model of bargaining that allows the possibility of resale. This requires dropping the distinction between "seller" and "buyers"; we substitute the terms "owner" and "non-owners." The identity of the owner can vary over time. Similarly, we denote player i '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. An outcome involving trade specifies a date of trade $t \geq 1$, a (non-negative) price p , and the next owner j . There is also an outcome called "No Trade" which we identify as a trade at date $t = \infty$. To interpret the results that follow, it is

¹³The 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 [1985].

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 game form Γ^i . First, the owner is permitted to keep the item for himself, that is, he has 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. Nonowners must decide simultaneously whether to bid. If no bids are made, play continues according to the continuation rules of Γ^i . 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. However, a non-owner cannot commit himself not to trade before the owner makes the auction offer.

Now, we create quite a general bargaining model with auction offers and resale as follows. Let player i_0 be the initial owner. Then the game form Γ^{i_0} is played. If the outcome involves trade after period of ownership 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 $t_0+\dots+t_j-1$. Payments are made on the dates of transfer of ownership.

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:

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

Let the initial owner i_0 be player i . With this, we have a complete specification of an extensive form game, which we shall call Γ_i^* .

The games Γ_i^* that can be constructed in this way form a huge class. It may be that the seller can also conduct auctions with a positive minimum price, or can exclude some set of bidders, or can bargain effectively with some buyers, or can commit himself with take-it-or-leave-it offers. Indeed, the most important restriction on the class of admissible games is that an owner cannot prevent the next owner from engaging in resale. An additional important 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 $(\sigma_1, \dots, \sigma_n)$ which is a perfect equilibrium profile regardless of the identity of the initial owner (that is, 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, since future play could depend on the history.

Proposition 2. Assume there are at least 3 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 inequalities (2.4) and (2.5) hold.

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

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

Proof. 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 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 nonowner prefers (weakly) his continuation payoff to the payoff from placing the minimum bid of 0: $\bar{v}_j \geq \delta v_j^*$. Since 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 $k \in \{1, i\}$ 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 $\delta(v_k^* - p')$. Since 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 $\delta(v_1 - p') + \delta(v_k^* - p') \geq \delta(v_1 - p)$ 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^* . So, the equilibrium price is v_k^* , with probability one.

Suppose that, with positive probability, 1 is not the winning bidder; instead k is. When k wins, the total payoff to the nonowners in the game is at most $\delta[(1-\delta)v_k + \delta v_1 - v_k^*]$ which is strictly less than $\delta(v_1 - v_k^*)$. Hence, when k wins, 1's payoff is strictly less than that sum. But, for any $\epsilon > 0$, 1 can guarantee a payoff of $\delta(v_1 - v_k^* - \epsilon)$ by bidding $v_k^* - \epsilon$. Hence, if k wins with positive probability, then 1's bidding strategy is not optimal. Thus, following an auction offer by i in Γ_i^* , any equilibrium prescribes that 1 bid v_k^* and wins the auction.

Henceforth, we denote the price v_k^* by p . 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:

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

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

Next, we make two applications of (2.6):

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

which establishes (2.5). §§

Thus, 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 theorem 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.

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 one, 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

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

Then, in the telephone bargaining model, if $i \neq 1$, 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 $B(v_1, v_1)$ is larger than the auction price of $(1-\delta)v_2 + \delta B(v_2, v_1)$. Otherwise, the sale is by auction.¹⁴

The game has been structured so that the owner makes the first offer. Since 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 to 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.¹⁵ In that 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

¹⁴The proof is omitted. It follows the now familiar lines for alternating offer bargaining models.

¹⁵For example, if the annual real interest rate were as much as 5% and it took a week to arrange an auction, then δ would be .999.

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. The specific example with which we ended the section, however, 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 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 the 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 $p\delta^t + s(1 - \delta^t)$.¹⁶ 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.¹⁷ 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

¹⁶The 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.

¹⁷This involves no loss of generality. One can reproduce any distribution F essentially by taking $u = F^{-1}$.

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.¹⁸

Proposition 5. 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). 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 5 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

¹⁸The following two Propositions synthesize results of Harris and Raviv [1982], Milgrom [1985], Myerson [1981], Riley and Samuelson [1981], and Rubinstein, Wilson and Wolinsky [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.

first restrictions are the so-called incentive compatibility constraints: 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 constraints means that the seller cannot extract a higher price from a buyer with a higher valuation unless he gives that buyer something correspondingly valuable in return, such as a higher probability of receiving the item or the opportunity to receive it sooner. The participation constraints mean that the buyer's expected payoff must be non-negative, regardless of his valuation for 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 5 amounts to computing the bound and showing that it is achieved by a take-it-or-leave-it offer.

The method described above 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 5 using these ideas is given in an Appendix.

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 so-called independent private values model. The independence assumption is particularly important for the following results; it means that an outside observer (or the seller) could not 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 6. 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.¹⁹

Proposition 6 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 McClean [1985b] have studied a variation involving some statistical dependence. Milgrom [1985] allowed the seller to have many objects for

¹⁹Other auctions with the same minimum price that always allocate the object to the highest evaluator lead to the same expected price.

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 all the bids that can be made actually are made at equilibrium in several simple environments. Properties like simplicity and robustness are interesting to think about but hard to formulate; almost nothing has yet been done on this part of the 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, 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 which offers such an explanation.

Our analysis is based on the symmetric auction model introduced by Milgrom and Weber [1982], 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 (\underline{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 world oil

prices, availability of transport facilities like pipelines (which depends in turn on the productivity of nearby tracts), etc.

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 a study of the bidding problem quite an interesting exercise. Before giving a more formal account of the Winner's Curse, we must finish specifying our modelling 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 $S_i \in (\underline{s}_i, \bar{s}_i)$ and $X_j \in (\underline{x}_j, \bar{x}_j)$, $i=1, \dots, m$, $j=1, \dots, n$. Such random variables are called affiliated. The main facts about affiliated random variables are briefly summarized in the next paragraph.²⁰

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:

²⁰The theory of affiliated random variables is presented in the appendix of Milgrom and Weber [1982], where the cited results are proved. The property of the densities of affiliated random variables reported in the text has been studied by a number of authors, who have given it various names including the "FKG inequality" and the "MTP₂ property."

$$(4.1) \quad f(z) f(z') \leq f(z \cap z') f(z \cup z')$$

where $z \cap z'$ is a vector whose i^{th} component is $\min(z_i, z'_i)$ and $z \cup z'$ 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:

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

is nondecreasing.

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 downwards from his initial estimate:

$$(4.3) \quad E[V_1 \mid X_1, \max_j \beta_j(X_j) < b] \leq E[V_1 \mid X_1, \max_j \beta_j(X_j) < \infty] \\ = E[V_1 \mid X_1]$$

(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.²¹ Of course, at equilibrium, bidders will take this fact into account in advance in choosing their strategies.

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] who first noted this equivalence, and he also claimed that the standard ascending-bid auction is equivalent to a particular sealed-bid auction. He reasoned as follows. Suppose in the ascending-bid auction bidder i decides to bid up to the level $\beta_i(X_i)$,

²¹See Milgrom [1981] for a more complete analysis.

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 interests 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 an essential way by this modelling.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

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.

²²Milgrom and Weber [1982] 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 have the same properties as those of the model studied here.

²³Here, 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 [1984].

Consider any of the large family of auctions with the property that, at equilibrium, the bidder with the highest evaluation wins and acquires the object for some non-negative price, and where all losers pay zero. In general, the price paid may depend on how all the bidders behave. Given an auction "A", let $W^A(z,x)$ be the expected price paid by a bidder, say bidder 1, when $X_1 = x$, but bidder 1 bids as if X_1 were z , and he wins. For the sealed-bid auction, $W^A(z,x) = \beta_S(z)$. Notice that it doesn't depend on x at all, since the price a bidder pays depends only on the bid he makes. However, when the seller reveals information or when an ascending-bid auction is used, if the variables are not independent, $W^A(z,x)$ will normally depend upon x . For example, for the ascending-bid auction, $W^A(z,x) = E[\beta_A(Y,Y) \mid X_i=x, Y < z]$, where $Y = \max_{j \neq i} X_j$. When X_i and Y are correlated, this expectation depends on x as well as z . The message of the Linkage Principle is that such linkages raise the expected price.

Proposition 7 (The Linkage Principle).²⁴ Let "A" and "B" be a pair of auctions with the properties that prices are always non-negative and the bidder with the highest valuation wins (at equilibrium). Suppose that for all x , $W_2^A(x,x) \geq W_2^B(x,x)$, where the subscripts denote partial derivatives with respect to the second variable. Then, the expected price is higher in auction "A" than in auction "B".

²⁴The Linkage Principle was originally introduced by Milgrom and Weber [1982a], who described it as "the common thread running through" their results (pp. 110-111). 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."

Proof. Let $R(z,x) = E[V_1 1_{\{Y < z\}} | X_1 = x]$, that is, the expected value received 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. Let $F(z|x) = P\{Y < z | X_1 = x\}$, that is, the probability that such a bid will in fact win when $X_1 = x$. After observing $X_1 = x$, bidder 1, by choosing a bid, chooses a set $\{Y < z\}$ of situations in which to win. Thus, at equilibrium of auction game "A", x must be a solution to:

$$(4.4) \quad \max_z R(z,x) - F(z|x)W^A(z,x) .$$

Using subscripts to denote partial derivatives, the first-order necessary condition is:

$$(4.5) \quad 0 = R_1(x,x) - F_1(x|x)W^A(x,x) - F(x|x)W_1^A(x,x), \text{ or}$$

$$(4.6) \quad W_1^A(x,x) = R_1(x,x)/F(x,x) - [F_1(x,x)/F(x,x)]W^A(x,x),$$

and a similar expression holds for W^B . Define:

$$(4.7) \quad \Delta(x) = W^A(x,x) - W^B(x,x).$$

Now, in both auctions A and B, a bidder with information \underline{x} always loses at equilibrium, and so has an expected payoff of zero. Were he to bid as if his information were $\underline{x} + \epsilon$, he would win with positive probability. The expected price he pays in this deviation must, therefore, be at least $R(\underline{x} + \epsilon, \underline{x})/F(\underline{x} + \epsilon | \underline{x})$. It cannot be more (for ϵ small) since then the expected profits of a bidder with information $\underline{x} + \epsilon$ would be negative. Since this argument applies both to auctions A and B, $W^A(\underline{x}, \underline{x}) = W^B(\underline{x}, \underline{x})$. So, $\Delta(\underline{x}) = 0$.

Also, using (4.6) and (4.7),

$$(4.8) \quad \Delta'(x) = -[F_1(x|x)/F(x|x)] \Delta(x) + W_2^A(x,x) - W_2^B(x,x).$$

Now, whenever $\Delta(x) \leq 0$, the right-hand side of this expression is non-negative, by our hypothesis and the fact that $F_1/F \geq 0$.

Now suppose $\Delta(z) < 0$ for some $z > \underline{x}$. Let $\underline{z} = \sup\{x \mid x < z \text{ and } \Delta(x) \geq 0\}$. Since $\Delta(\underline{x}) = 0$, we must have $\Delta(\underline{z}) = 0 > \Delta(z)$. Hence, by the Mean Value Theorem, there exists $z' \in (\underline{z}, z)$ such that $\Delta'(z') < 0$. But by construction, $\Delta(z') < 0$, which using (4.8) implies that $\Delta'(z') \geq 0$, a contradiction. So, for all $x \geq \underline{x}$, $\Delta(x) \geq 0$.

The expected price in auction "A" when all bidders adhere to their equilibrium strategies is equal (by symmetry) to the expected price conditional on bidder 1 winning, which is:

$$(4.9) \quad E[W^A(X_1, X_1) \mid X_1 = \max_j X_j],$$

and similarly for B. Hence, the expected price difference is $E[\Delta(X_1, X_1) \mid X_1 = \max_j X_j] \geq 0$, as required. §§

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

Proposition 8 (Revenue Equivalence). Consider any auction game where X_1, \dots, X_n are independent and for which prices are restricted to be non-negative. Suppose that, at equilibrium, the high bidder always wins and losers neither make nor receive payments. Then, the expected price in every such auction is the same as for the ascending-bid auction.²⁵

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

Proof. When bidders $j \neq i$ adhere to their equilibrium strategies and bidder i , who has observed $X_i = x$, bids as if $X_i = z$, the price in auction "A" will be $p = p(z, X_{-i})$. Then, when i wins,

$$(4.10) \quad W^A(z, x) = E[p(z, X_{-i}) \mid X_i = x, \max_{j \neq i} X_j < z]$$

which, by the hypothesis of statistical independence,

$$= E[p(z, X_{-i}) \mid \max_{j \neq i} X_j < z] ,$$

and similarly for $W^B(z, x)$. Hence, $W_2^A(x, x) \equiv W_2^B(x, x) \equiv 0$. Apply the Linkage Principle. §§

When we replace the hypothesis of statistical independence by one of affiliation, the Linkage Principle provides a powerful tool for making positive 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 9. 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. It is clear that $W_2^B(x, x) \equiv 0$, since the price depends only on the winner's bid, as illustrated in the text above. The ascending bid auction "A" is also illustrated above. There, the equilibrium price paid by winning bidder i is a nondecreasing function of X_j , $j \neq i$, and these are affiliated to X_i . By the affiliation theorem cited earlier, this means that the expected price given i 's actual bid and X_i is a nondecreasing function of X_i . Hence, $W_2^A(x, x) \geq 0$. Apply the Linkage Principle. §§

In the sealed-bid auction, the winning bidder's payment depends only on his own observation; there is no linkage to other variables. In the ascending-bid auction, the price is determined by the second highest bidder's observation, which provides the price-increasing linkage.

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

Proposition 10. 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.

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

Proposition 11. If the 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 the 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, so any equilibrium must have the property that the seller

reports whenever $S_0 > s^*$. But then, at equilibrium, when the seller reports nothing, the buyers must believe that $S_0 \leq s^*$ and bid accordingly. They would bid strictly more if they had the more "favorable" 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 . §§

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. MacAfee and MacMillan [1984] 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 gets 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 the 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. With risk neutral bidders and a risk neutral seller, when comparing auctions that allocate the good efficiently, any reduction in the bidders' payoffs is a gain to the seller.

Plainly, that 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 up to now 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.²⁶

In the independent private values model, linkages inefficiently increase the randomness in the payoffs. However, in other models, linkages reduce the randomness in the payoffs. For example, suppose the

²⁶See also Holt [1980].

observations are independent and 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 effort 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 situation with a single buyer and offers made by the seller has been analyzed by Stokey [1982], and by Gul, Sonnenschein and Wilson [1985]. 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 for each moment in time, whether to bid and an amount to bid, 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, since that does not specify a first moment to bid when 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 12. 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 some 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 check is 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. For 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 12 is the unique symmetric (among buyers) equilibrium. The stationarity property requires that a buyer's current decision depends only on his own type and 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 or on how those beliefs were reached. When the buyers are limited to stationary strategies, if all trading does not take place at time zero, 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 strategem makes all trades occur earlier without reducing the price the seller gets. Therefore, at any equilibrium in stationary strategies, all trading occur at time zero. Then, since the seller cannot adhere to a minimum price above his reservation price of zero, 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 lie in the core of the exchange game does not mean that auctions are immune to collusion; it means only that no subset of the players could improve their lot by going off and trading among themselves. Auctions with low minimum prices are vulnerable to collusion among the 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 one, which exploits the existence of multiple Nash equilibria in ascending-bid auctions to construct collusive perfect equilibria in repeated ascending-bid auctions.²⁷

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

²⁷Bikhchandani [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 which naturally favors the formation of collusive rings.

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:²⁸

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

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.

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 particular details, like the popularity of sealed bid and ascending-bid auctions? The answers to these questions were summarized in the introduction. The answers are plainly incomplete;

²⁸There 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.

indeed, they rely on fundamentally different models of the auction environment. What is urgently needed is a single consistent model that explains both the use of auctions in preference to other mechanisms when individual items are unique and such details as the widely observed preference for ascending-bid auctions over sealed-bid auctions.

There are segments of auction theory that I have omitted from my review, partly because they do not fit neatly into the schema of my two questions, and partly because a surveyor must draw lines. One could survey very different territory by asking other questions, like: (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 studying bidding behavior with controlled laboratory experiments.²⁹ It would also cover studies of the implications of alternative

²⁹An 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, in which 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 predict better.

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, though this supposition is of course subject to empirical refutation.

models of choice under uncertainty for bidding behavior.³⁰ (2) What insights does auction theory offer into the problems of price discrimination? Many of the researchers who have contributed to the theory of expected-price maximizing auctions have extended their results to the problem of optimal price discrimination by a monopolist (Cremer and MacLean [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 [1985] companion survey gives an introduction to this new and lively area of research, sometimes called the theory of market microstructure. Our Propositions 2, 3, and 12 give some idea of the issues studied in that connection. (4) How are auctions used in contracting environments, where the bidder's performance in the contract needs to be properly motivated? This is a new subject of study that has so far limited itself unnecessarily 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.

³⁰For example, Karni and Safra [1985] have a model of bidding behavior in which the bidders may behave differently even in strategically equivalent games.

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," unpublished 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.
- , Ariel Rubinstein and Asher Wolinsky [1985], "The Nash Bargaining Solution in Economic Modelling." Working Paper No. 150, Department of Economics, Hebrew University of Jerusalem.
- 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, 50:607-638.
- Cremer, Jacques and Richard McLean [1985a], "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," Econometrica, 53:345-361.
- [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.
- Gul, Faruk and Hugo Sonnenschein [1985], "One-Sided Uncertainty Does Not Cause Delay," mimeo, Stanford University.
- and Robert Wilson [1985], "Foundations of Dynamic Monopoly and the Coase Conjecture," mimeo, Stanford University.
- Harris, Milton and Arthur Raviv [1981], "A Theory of Monopoly Pricing With Uncertain Demand," American Economic Review, 71:347-365.

- _____ [1982], "Allocation Mechanisms and the Design of Auctions," Econometrica, 49:1477-1499.
- Harstad, Ronald and Dan Levin [1984], "A Dominant Strategy Argument for a Class of Common Value Auctions," forthcoming in the Review of Economic Studies.
- Holt, Charles A., Jr. [1980], "Competitive Bidding for Contracts Under Alternative Auction Procedures," Journal of Political Economy, 88:433-445.
- Karni, Edi and Zvi Safra [1985], "English and Vickrey Auctions in the Theory of Expected Utility With Rank-Dependent Probabilities," Tel Aviv University Working Paper.
- MacAfee, Preston and John MacMillan [1984], "Bidding for Contracts: A Principal-Agent Analysis," mimeo, University of Western Ontario.
- Maskin, Eric, and John Riley [1984a], "Optimal Auctions with Risk Averse Buyers," Econometrica, 52:1473-1518.
- _____ [1984b], "Monopoly With Incomplete Information," Rand Journal of Economics, 15:171-196.
- _____ [1983], "Uniqueness of Equilibrium in Open and Sealed Bid Auctions," unpublished.
- Matthews, Steven [1980], "Risk Aversion and the Efficiency of First and Second Price Auctions," Northwestern University, unpublished.
- _____ [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, 49:921-943.
- _____ [1985], "The Economics of Competitive Bidding: A Selective Survey," in Social Goals and Social Organization, L. Hurwicz, D. Schmeidler, and H. Sonnenschein (eds). Cambridge: Cambridge University Press.
- _____ 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.
- Riley, John [1985], "Ex Post Information in Auctions," UCLA Department of Economics Working Paper #367.
- and William Samuelson [1981], "Optimal Auctions," American Economic Review, 71:381-392.
- Robinson, Marc [1985], "Collusion and the Choice of Auction," RAND Journal of Economics, 16:141-145.
- Rubinstein, Ariel, Charles Wilson and Asher Wolinsky [1985], private communication.
- 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, 12:112-128.
- 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, 15-27. Princeton: Princeton University Press.
- Wilson, Robert [1977], "A Bidding Model of Perfect Competition," Review of Economic Studies, 4:511-518.
- [1985], "Game-Theoretic Analyses of Trading Processes," Lecture for the Fifth World Congress of the Econometric Society.

APPENDIX

Proof of Proposition 5. Suppose the seller designs some game Γ in which, at equilibrium, trade takes place at some date T which 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 the buyer's optimal strategy and the maximum payoff, respectively, in the game when $X = x$. Define $p^*(x) \equiv p(\sigma^*(x))$ and $e^*(x) \equiv e(\sigma^*(x))$. Then,

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

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 but increases 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 $(1-p^*(x)) s + e^*(x)$.

Since X is uniformly distributed on $(0,1)$, the corresponding unconditional expectation is:

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

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

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

The seller's total expected payoff is therefore:

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

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

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

Inequality (A.5) constrains the seller in designing a game. Another constraint that must hold at equilibrium is:

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

Since the objective (A.4) is linear in p^* and $e^*(0)$ and the constraint set (A.5)-(A.6) is convex, the maximum must occur at an extreme point. Thus, at a maximum, (A.5) must hold with equality. Also, the p^* that

maximizes the seller's expected payoff (A.4) subject to (A.6) must either jump from zero to one at the point x^* where the integrand is maximized, or must be a constant, 0 or 1 (in case x^* is 1 or 0). 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, the unique $p^*(\cdot)$ function that attains the maximum is $p^*(x) = 0$ for $x < x^*$ and $= 1$ for $x \geq x^*$, so any institution as good as making a take-it-or-leave-it offer must lead to the same trading outcome. §§