Reprinted from JOURNAL OF ECONOMIC THEORY

Vol. 43, No. 1, October 1987
Printed in Belgium

Auctions with a Stochastic Number of Bidders*

R. PRESTON MCAFEE AND JOHN MCMILLAN

Department of Economics, University of Western Ontario, London, Ontario N6A 5C2, Canada

Received May 2, 1985; revised August 19, 1986

Auction theory is generalized by allowing the number of bidders to be stochastic. In a first-price sealed-bid auction with bidders having constant absolute risk aversion, the expected selling price is higher when the bidders do not know how many other bidders there are than when they do know this. Thus the seller should conceal the number of bidders if he can. Moreover, a bidder's ex ante expected utility is the same whether or not there is a policy of concealing the number of bidders: concealment therefore Pareto-dominates announcement. With risk-neutral bidders, the optimal auction is the same whether or not the bidders know who their competitors are. *Journal of Economic Literature* Classification Numbers: 022, 026. © 1987 Academic Press, Inc.

1. INTRODUCTION

Asymmetric information and imperfect competition are the two essential ingredients of the theory of auctions. But it is presumed in the existing auction models[1] that one piece of information is common knowledge: all bidders know how much competition they face. Is it appropriate to model the bidders as knowing who their competitors are?

In an English auction, a bidder often cannot identify his rivals. The other bidders may be acting on behalf of anonymous principals. Not all the people present are active bidders. Bidders use subtle signals to hide their bidding. "Such signals may be in the form of a wink, a nod, scratching an ear, lifting a pencil, tugging at the coat of the auctioneer, or even staring into the auctioneer's eyes—all of them perfectly legal. This method of communicating bids gives the process of bidding an aura of secrecy" (Cassady [1, pp. 149–150]).

In a sealed-bid auction, there is still less reason to suppose that bidders know the number or the identities of their competitors, since the bidders do

* We thank an Associate Editor and a referee for their perceptive comments, and the Ontario Economic Council and the Social Sciences and Humanities Research Council of Canada for research support.

[1] See McAfee and McMillan [9] and Milgrom [10, 11] for surveys of auction theory.

0022-0531/87 $3.00

not assemble together in one place. With some government-contract bidding, the government invites selected contractors to submit bids. In this case, the government knows in advance the number of bidders. The government therefore has available an extra policy instrument for fostering competition among the bidders, in addition to choosing the form of the auction, setting reserve prices, etc.: the government can choose either to conceal or to reveal the number of bidders.[2]

Are the results of auction theory sensitive to the assumption that each bidder knows exactly how many bidders there are? This paper will show that they are. First, if the bidders are risk averse (with constant absolute risk aversion) the seller's expected revenue in a first-price sealed-bid auction is higher if the bidders do not know how many bidders there are than if they do know this. Second, with risk-neutral bidders, the optimal direct, incentive-compatible auction is the same whether or not the bidders know the number and the identities of the other bidders. However, knowledge about the set of bidders still matters, because if different bidders have different (albeit Bayesian consistent) expectations over the set of bidders, then if the set of bidders is not known, the optimal auction cannot be implemented using a first-price sealed-bid auction, although it can be implemented using an English auction.

Milgrom and Weber [12] showed that, in many circumstances, it is in the seller's interest to reveal any information he has. This paper exhibits a different set of circumstances in which the seller should conceal information.

In Section 2, we examine probability distributions over subsets of potential bidders: this is necessary since the set of active bidders is a random variable. In Section 3 we consider, in an independent-private-values auction, the effects of the bidders' knowing how many bidders there are. If the bidders have constant absolute risk aversion, the seller's revenue is on average higher if he conceals the number of bidders than if he reveals it. Interestingly, the bidders' expected utility is the same in either regime, which means that the policy of concealment Pareto-dominates the policy of announcement.

In Section 4, we allow the bidders to be different ex ante, in the sense that their valuations of the good are drawn from different distributions. In addition, they may have different priors on how many bidders are present as long as these priors are Bayesian consistent. Thus revealing the number of bidders is not the only issue, for the seller could also reveal the identities of the bidders, as they are not ex ante the same. If the bidders are risk neutral, the seller is indifferent between revealing and concealing the bid-

[2] For example, Ontario Hydro, Ontario's electrical utility, has a policy of keeping secret the number of firms it has invited to submit bids; see McAfee and McMillan [8, Chap. 7].

In Section 3, but not in Section 4, symmetric priors among active bidders will be assumed:

$$(\forall n \in \mathfrak{N})(\exists\, p_n)(\forall k \in \mathfrak{N})\; p_n^k = p_n. \tag{5}$$

The remainder of this section develops some properties of (5).

EXAMPLE 1. Suppose at most m bidders can be present, and $\gamma_0, \ldots, \gamma_m$ are given, satisfying

$$\sum_{n=0}^{m} \gamma_n = 1. \tag{6}$$

Then one assignment of β_n's that satisfies (5) is

$$\beta_A = \begin{cases} \gamma_n / \binom{m}{n} & \text{if } A \subseteq \{1, \ldots, m\} \text{ and } |A| = n \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

In addition

$$\sum_{\substack{A \\ k \in A \\ |A| = n}} \beta_A = \binom{m-1}{n-1} \gamma_n \Big/ \binom{m}{n} = \frac{n\gamma_n}{m}. \tag{8}$$

Thus

$$p_n^k = \sum_{\substack{A \\ k \in A \\ |A| = n}} \beta_A \Big/ \sum_{k \in A} \beta_A = \frac{n\gamma_n}{m} \Big/ \sum_{i=0}^{m} \frac{i\gamma_i}{m} = \frac{n\gamma_n}{n^*}. \tag{9}$$

EXAMPLE 2.

$$\begin{aligned} \beta_{\{1\}} &= \frac{\gamma}{2(3-\gamma)} \\ \beta_{\{k\}} &= \frac{3\gamma 2^{-k}}{3-\gamma}, \qquad k \geqslant 2 \\ \beta_{\{k,k+1\}} &= \frac{3(1-\gamma)2^{-k}}{3-\gamma}, \qquad k \geqslant 1 \\ \beta_A &= 0 \qquad \text{otherwise.} \end{aligned} \tag{10}$$

ders' identities; his expected revenue is the same in either case. To prove this, we construct the optimal auction, extending the optimal-auctions literature to the case of unknown bidders.

2. Number of Bidders

We index the potential bidders with natural numbers $\mathfrak{N} = \{1, 2, 3, \ldots\}$. For any finite set $A \subseteq \mathfrak{N}$, let β_A represent the probability that A is the set of active bidders. We presume that the set of active bidders is finite with probability 1,

$$\sum_{n=0}^{\infty} \sum_{\substack{A \\ |A|=n}} \beta_A = 1, \tag{1}$$

where $|A|$ is the cardinality of A and $\sum_{A,|A|=n}$ denotes the sum over all sets A with cardinality n. Thus the probability that n bidders are present is

$$\gamma_n = \sum_{\substack{A \\ |A|=n}} \beta_A. \tag{2}$$

The expected number of bidders is

$$n^* = \sum_{n=1}^{\infty} n\gamma_n. \tag{3}$$

We assume $n^* < \infty$.

The process by which bidders are selected is taken to be exogenous, embodied in the probabilities β_A. For example, in government-contract bidding, bidders are selected from a list of qualified bidders on a rotating basis (McAfee and McMillan [8, Chap. 7]). This analysis generalizes the standard model, where $\beta_A = 1$ if $A = \{1, \ldots, n\}$ and $\beta_A = 0$ otherwise. It is presumed that all bidders are Bayesians and know the β_A's.

Once bidder k is selected, he updates his probability of the number of bidders being present

$$p_n^k = \sum_{\substack{A \\ |A|=n \\ k \in A}} \beta_A \Big/ \sum_{\substack{A \\ k \in A}} \beta_A. \tag{4}$$

It is easily verified that

$$p_2^k = \frac{\gamma}{3-2\gamma}$$
$$p_1^k = \frac{3(1-\gamma)}{3-2\gamma} \tag{11}$$
$$p_n^k = 0, \qquad n > 2.$$

Example 1 proves that, as long as there is an upper bound to the number of active bidders, there is a probability distribution on the finite subsets of $\mathfrak{N}$ which yields any desired distribution of the number of bidders and satisfies (5). In addition, it illustrates that $p_n = (n\gamma_n)/n^*$, which Lemma 1 will show to be true if (5) holds, and which is useful in the analysis of the symmetric-bidders case in Section 3. Example 2 provides an example with infinitely many potential bidders, and with (5) still satisfied.

LEMMA 1. *Given* (5),

$$n^* p_n = n\gamma_n. \tag{12}$$

Proof.

$$\begin{aligned} n\gamma_n &= \sum_{\substack{A \\ |A|=n}} n\beta_A \\ &= \sum_{\substack{A \\ |A|=n}} \sum_{\substack{k \\ k\in A}} \beta_A \\ &= \sum_{k=1}^{\infty} \sum_{\substack{A \\ k\in A \\ |A|=n}} \beta_A \\ &= \sum_{k=1}^{\infty} p_n \Big(\sum_{\substack{A \\ k\in A}} \beta_A \Big) \qquad \text{(by (4) and (5))} \\ &= p_n \sum_A \sum_{\substack{k \\ k\in A}} \beta_A \\ &= p_n \sum_A |A| \beta_A = p_n n^* \qquad \text{(by (1))}. \end{aligned}$$

Q.E.D.

Lemma 1 has an intuitive explanation. Any particular individual is n times as likely to be a participant when the number of bidders is n than when it is one. The posterior probability is thus proportional to $n\gamma_n$. For the probabilities to sum to one, the constant of proportionality is $1/n^*$.

Lemma 1 is illustrated by the following "classroom size" problem. Consider polling professors on the size of their classes, and suppose γ_n is the proportion of professors reporting a class size of n. If p_n is the proportion of students in the classes reporting class size of n, then γ_n and p_n are related by Lemma 1.[3]

3. Sealed-Bid Auction with Risk-Averse Bidders

In this section we consider a first-price sealed-bid auction. The bidders have independent private values, drawn from a distribution F with density f. Assume f is continuous and $f(x)$ is strictly positive if $x \in (0, \bar{x})$. Following Matthews [4] and Milgrom and Weber [12], we assume the bidders have the same constant-absolute-risk-aversion utility function

$$U(w) = \frac{1 - e^{-\lambda w}}{\lambda}, \tag{13}$$

for some constant $\lambda \geqslant 0$. We assume there exists a symmetric Bayes–Nash equilibrium, and denote by $B(x)$ the corresponding equilibrium bidding function, which we assume to be strictly increasing.

How does a policy of concealing the number of bidders affect the bidders' expected utility? Expected utility can be evaluated at two different times: first, before anyone knows whether he is a participant and what his value x is; and second, after the bidders have been selected. We shall denote an agent's expected utility at the former time-point ex ante utility, and at the latter time-point interim utility.

Theorem 1. *The interim utility of a bidder with given value is the same whether or not the bidders are to be informed about the number of bidders.*

Proof. The probability that any particular bidder wins with bid b is $F^{n-1}(B^{-1}(b))$. Thus the bidder's ex ante expected utility is

$$EAU(b) = \sum_{n=1}^{\infty} p_n \left[F^{n-1}(B^{-1}(b)) \frac{1}{\lambda} (1 - e^{-\lambda(x-b)}) \right]. \tag{14}$$

[3] An alternative to this model was suggested by Paul Milgrom. Suppose there are continuum many bidders, indexed on [0, 1]. Select the bidders by first selecting the number of bidders n using the probabilities γ_n, and then selecting n bidders using a uniform distribution. This satisfies (5) (by symmetry) and Lemma 1.

At a Nash bidding equilibrium,

$$\left.\frac{\partial EAU}{\partial b}\right|_{b=B(x)} = 0. \tag{15}$$

Thus

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x)\, e^{-\lambda(x-B(x))} B'(x)$$

$$= \sum_{n=1}^{\infty} p_n (n-1)\, F^{n-2}(x)\, f(x) \frac{1}{\lambda} (1 - e^{-\lambda(x-B(x))}). \tag{16}$$

Represent the bidder's equilibrium ex ante expected utility by $EEAU(x) \equiv EAU(B(x))$. Thus

$$EEAU(x) = \sum_{n=1}^{\infty} p_n F^{n-1}(x) \frac{1}{\lambda} (1 - e^{-\lambda(x-B(x))}). \tag{17}$$

Differentiate (17),

$$\frac{d}{dx}(EEAU(x)) = \sum_{n=1}^{\infty} p_n \Big[(n-1)\, F^{n-2}(x)\, f(x) \frac{1}{\lambda} (1 - e^{-\lambda(x-B(x))})$$

$$+ F^{n-1}(x)\, e^{-\lambda(x-B(x))}(1 - B'(x)) \Big]$$

$$= \sum_{n=1}^{\infty} p_n F^{n-1}(x) - \lambda EEAU, \tag{18}$$

using (15) and (16). Now solve this linear differential equation:

$$EEAU(x) = e^{-\lambda x} \left[K + \int_0^x \sum_{n=1}^{\infty} p_n F^{n-1}(s) e^{\lambda s}\, ds \right], \tag{19}$$

for some constant K. To evaluate K, note that if the minimum bid is zero, it follows that $EEAU|_{x=0} = 0$, or

$$EEAU(x) = e^{-\lambda x} \left[\int_0^x \sum_{n=1}^{\infty} p_n F^{n-1}(s) e^{\lambda s}\, ds \right]. \tag{20}$$

Thus, from (17) and (20),

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x) \frac{1}{\lambda}(1 - e^{-\lambda(x-B(x))})$$
$$= \sum_{n=1}^{\infty} p_n e^{-\lambda x} \int_0^x F^{n-1}(s) e^{\lambda s}\, ds$$
$$= \sum_{n=1}^{\infty} p_n e^{-\lambda x} \left[\frac{F^{n-1}(s) e^{\lambda s}}{\lambda} \Bigg|_0^x - \frac{1}{\lambda} \int_0^x (n-1) e^{\lambda s} F^{n-2}(s) f(s)\, ds \right]. \tag{21}$$

Thus

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x) e^{-\lambda(x-B(x))}$$
$$= e^{-\lambda x} \sum_{n=1}^{\infty} p_n \int_0^x (n-1) e^{\lambda s} F^{n-2}(s) f(s)\, ds \tag{22}$$

or

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x) e^{\lambda B(x)} = \sum_{n=1}^{\infty} p_n \int_0^x (n-1) e^{\lambda s} F^{n-2}(s) f(s)\, ds. \tag{23}$$

Consider now the equilibrium bids in the case in which the bidders know the number of bidders n. It follows from (23), by setting $p_n = 1$ for the announced n and $p_m = 0$ for $m \neq n$, that the equilibrium bid $B^n(x)$ when n is announced satisfies

$$F^{n-1}(x)\, e^{\lambda B^n(x)} = \int_0^x (n - 1) e^{\lambda s} F^{n-2}(s) f(s)\, ds. \tag{24}$$

Thus, from (23) and (24),

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x) e^{\lambda B(x)} = \sum_{n=1}^{\infty} p_n F^{n-1}(x) e^{\lambda B^n(x)}. \tag{25}$$

From (25),

$$\sum_{n=1}^{\infty} p_n F^{n-1}(x) \frac{1}{\lambda}(1 - e^{-\lambda(x-B(x))})$$
$$= \sum_{n=1}^{\infty} p_n F^{n-1}(x) \frac{1}{\lambda}(1 - e^{-\lambda(x-B^n(x))}). \tag{26}$$

The left side of (26) is (from (17)) equilibrium interim expected utility when n is not known. The right side of (26) is equilibrium interim expected utility when the policy of announcing n is in effect. Q.E.D.

THEOREM 2. *The ex ante utility of any bidder is the same whether or not the bidders are to be informed about the number of bidders.*

Proof. From the seller's point of view, the expected utility of all agents is

$$\begin{aligned}\sum_{n=1}^{\infty} \gamma_n \int_0^{\infty} \frac{1}{\lambda}(1-e^{-\lambda(x-B(x))})\, nF^{n-1}(x)\, f(x)\, dx &= \sum_{n=1}^{\infty} n^* p_n \int_0^{\infty} \frac{1}{\lambda}(1-e^{-\lambda(x-B(x))})\, F^{n-1}(x)\, f(x)\, dx \\ &= \sum_{n=1}^{\infty} n^* p_n \int_0^{\infty} \frac{1}{\lambda}(1-e^{-\lambda(x-B^n(x))})\, F^{n-1}(x)\, f(x)\, dx \\ &= \sum_{n=1}^{\infty} \gamma_n \int_0^{\infty} \frac{1}{\lambda}(1-e^{-\lambda(x-B^n(x))})\, nF^{n-1}(x)\, f(x)\, dx \qquad (27)\end{aligned}$$

(the first step using (12), the second (26), and the third (12)). Q.E.D.

The difference between Theorems 1 and 2 lies in the timing of the announcement of the number of bidders. In Theorem 1, the interim utility occurs after the bidder has learned his value and also learned he is a bidder. Thus, the posterior probabilities p_n^k are used in calculating expected utility. In Theorem 2, the prior probabilities γ_n are used, and the expectation is taken before the potential bidder knows his value. An alternative interpretation of Theorem 2 is that the expected utility of the winner of the auction is invariant to the policy of concealment.[4]

THEOREM 3. *Concealing the number of bidders in the first-price sealed-bid auction does not lower the ex ante expected selling price, and strictly raises it if the bidders are risk averse (i.e., $\lambda > 0$) and the distribution of the number of bidders is nontrivial (i.e., $0 < p_n < 1$ for some n).*

Proof. The seller obtains expected revenue $EB(x)$ (with concealment)

[4] The point at which the assumption of constant absolute risk aversion is used is Eq. (18), which is a differential equation defining the bidding function B: this is a linear differential equation if and only if the bidders have constant absolute risk aversion (McAfee and McMillan [7]).

or $EB^n(x)$ (with announcement), where the expectation is given by, for any β,

$$E\beta(n, x) = \sum_{n=1}^{\infty} \gamma_n \int_0^{\infty} \beta(n, x)\, nF^{n-1}(x)\, f(x)\, dx$$
$$= \sum_{n=1}^{\infty} n^* p_n \int_0^{\infty} \beta(n, x)\, F^{n-1}(x)\, f(x)\, dx. \tag{28}$$

From (25) and (12),

$$\sum_{n=1}^{\infty} n\gamma_n F^{n-1}(x) = \sum_{n=1}^{\infty} n\gamma_n F^{n-1}(x)\, e^{\lambda(B^n(x) - B(x))}. \tag{29}$$

Multiplying by $f(x)$ and integrating yields

$$1 = Ee^{\lambda(B^n(x) - B(x))}$$
$$\geqslant e^{\lambda E[B^n(x) - B(x)]}, \tag{30}$$

the inequality following because $e^{\lambda z}$ is convex in z. Thus

$$0 \geqslant \lambda E[B^n(x) - B(x)]. \tag{31}$$

Hence

$$EB^n(x) \leqslant EB(x). \tag{32}$$

Further, strict inequality holds if $\lambda > 0$ (so that $e^{\lambda z}$ is strictly convex) and the expectation is nontrivial ($\gamma_n \notin \{0, 1\}$ for some n). Q.E.D.

It should be stressed that what is essential in the proof of Theorem 3 is the difference in the objective and subjective probabilities over the number of bidders as shown in Lemma 1: the bidders do not act as though there are the objectively expected number of bidders. Theorem 3 is not merely a consequence of the fact that the expectation of a convex function exceeds the function of the expectation. In the next section, it will be shown that, when the bidders have symmetric expectations (as is assumed in this section) and are risk-neutral, the expected selling price is the same whether or not the number of bidders is announced.

To summarize: if the seller has the option of announcing in advance the number of bidders, should he do so? The answer is striking: concealment Pareto-dominates announcement. With constant absolute risk aversion, the bidders are neither worse off nor better off in ex ante terms with the policy of concealment; and, with risk-averse bidders, the seller is strictly better off.[5]

In an English auction, the price is the same whether or not the number of bidders is known, since in either case the bidding stops at the second-highest valuation; the same applies to a Vickrey auction. With independent private values and risk-averse bidders, the English auction and the Vickrey auction yield lower revenue for the seller than the first-price sealed-bid auction with the number of bidders known (Riley and Samuelson [15]); in turn, as we have seen, this yields lower revenue than the first-price sealed-bid auction with an unannounced number of bidders. Thus, with risk-averse bidders, the auction form considered in this section, the first-price sealed-bid auction with the number of bidders concealed, is the best of the simple auction forms from the point of view of the seller.[6]

Milgrom and Weber [12] examined the release of information by the seller in a model in which bidders' valuations are affiliated (which means, roughly, that bidders' valuations may be correlated). They obtained the following results. With risk-neutral bidders, the release of information by the seller raises the expected selling price in a first-price sealed-bid, Vickrey, or English auction. With risk-averse bidders, releasing information raises the price in an English or a Vickrey auction. The results obtained above imply that, in a first-price sealed-bid auction with risk-averse bidders (the case not examined by Milgrom and Weber), information release can result in either an increase or a decrease in the expected price.

The tendency for the selling price to rise after information release when bidders' valuations are affiliated is caused by what Milgrom and Weber [12] called the *linkage effect*. A rough intuition for this is as follows (for a more precise description, see Milgrom [11]). It is intuitively clear that, in general, any reduction in the variance of the bidders' estimates of the item's

[5] While the process by which potential bidders become actual bidders is exogenous, the analysis is consistent with an endogenous process of bidder selection. Suppose the probabilities β_A are determined by some exogenous stochastic costs incurred by the potential bidders in submitting bids. Then if some particular set of β_A's is an equilibrium distribution for the sealed-bid auction in which the auctioneer reveals the number of bidders, it is also an equilibrium distribution for the sealed-bid auction in which the auctioneer conceals the number of bidders. This is because the bidders' ex ante expected utility is the same for each auction; thus the payoff to learning one's value and then submitting a bid is the same for each auction. Thus concealing or revealing n would not affect any potential bidder's decision to submit a bid. Hence, even if the probability that any one bidder submitted a bid were endogenous, this probability would be the same for either auction.

[6] How large can the seller's gain from concealment be? In a simulation with $\gamma_2 = \gamma$, $\gamma_1 = 1 - \gamma$, $\lambda = 1$ and F uniform on $[0, 1]$, concealment increases revenue by 25% for small γ, and by 10% for $\gamma = \frac{1}{3}$. In addition, the percentage increase in revenue is a decreasing function of γ in this simulation.

value increases bidding competition and drives up the price. The release of affiliated information links bidders' estimates to the now-public information: it reduces the advantages from private information. In other words, the release of information has a similar effect to a reduction in the variance of perceived valuations.

With risk-averse bidders in a first-price sealed-bid auction, Theorem 3 above showed that there is a contrary tendency for the release of information to drive down the price. To understand this effect, note that a policy of revealing information would not on average change bidding behavior in an independent-private-values Vickrey or English auction: the expected second-highest valuation remains unchanged. In contrast, in a first-price sealed-bid auction with risk-averse bidders and independent private values, revealing information does change bidding behavior. If the revealed information is good news then each bidder knows that his rivals will bid more aggressively, so he must do likewise. Bad news similarly generates less aggressive bidding. Thus the policy of revealing information results in a higher variance of bids than the policy of concealing information. It therefore results in a lower price on average. We shall call this effect the *bid-dispersion effect.*

Milgrom [11] pointed out that linkages increase the randomness in bidders' payoffs. The foregoing results are consistent with this. From Theorem 2, bidders with independent private values in a first-price sealed-bid auction are indifferent between the policy of revealing information and the policy of concealing information. From Theorem 3, bids are on average lower with information revelation. It follows that the risk-averse bidders must be faced with more risk under revelation than under concealment.

Thus the bid-dispersion effect is absent from a Vickrey or English auction or when bidders are risk neutral; it operates by itself in a first-price sealed-bid auction with risk-averse bidders having independent valuations; and it operates in the opposite direction to the linkage effect in a first-price sealed-bid auction with risk-averse bidders having affiliated valuations.

The usual examples of information that can be revealed by the seller (an expert's appraisal of a painting, a geological survey of an oil well, etc.) do not apply to the independent-private-values model, since each bidder knows with certainty the value of the item to him. As discussed above, however, there is one type of information that is useful to a bidder when valuations are independent: information about the amount of competition.

In an extension of the present analysis, Matthews [6] has shown that the result that the seller prefers to conceal the number of bidders continues to hold when the bidders have decreasing absolute risk aversion. However, the bidders prefer a policy of revelation when they have decreasing absolute risk aversion, and a a policy of concealment when they have increasing absolute risk aversion.

4. The Optimal Auction with a Stochastic Set of Bidders

What is the optimal auction when the number of bidders is unknown? For the sake of tractability, we restrict attention now to the case of risk-neutral bidders.

Drop last section's assumption that potential bidders are identical. Suppose instead that potential bidder i independently draws his valuation x_i from a distribution F_i, which may vary from bidder to bidder; assume $F_i(0)=0$ and let $F_i'=f_i$. Thus informing bidders about the number of active bidders is now not the only issue; bidders also may or may not be informed about the identities of the other active bidders.

The problem just defined is the optimal-auction problem solved by Myerson [14], but generalized in one respect: in this problem the bidders know neither how many other active bidders there are, nor the identities (that is, the F_i's) of the other active bidders. Myerson's case is obtained in the analysis to follow by setting $\beta_A=1$ and $\beta_B=0$ for $B\neq A$, where A represents Myerson's known set of bidders.

We simplify the analysis relative to Myerson's by assuming the distributions F_i satisfy the regularity condition

$$\frac{d}{dx}\left(x-\frac{1-F_i(x)}{f_i(x)}\right)>0 \tag{33}$$

(cf. Myerson [14, p. 66]). This simplifies the analysis because, when (33) fails, the seller must randomize (Maskin and Riley [2]).

Suppose the seller values the object at $x_0 \geqslant 0$. The seller uses an incentive-compatible direct mechanism by announcing sets $\Gamma_i^A \subseteq \mathscr{R}^{|A|}$, $i\in A$, and functions $\alpha_i^A\colon \mathscr{R}^{|A|}\to\mathscr{R}$ so that, if the set of actual bidders is A and the bidders report valuations x^A, then each bidder $i\in A$ pays an amount $\alpha_i^A(x^A)$ and is awarded the good if $x^A\in\Gamma_i^A$. We use the notation

$$\{a_1,\ldots,a_{|A|}\}=A; \tag{34}$$

$$(x^A_{-a_i}, y)=(x_{a_1},\ldots,x_{a_{i-1}}, y, x_{a_{i+1}},\ldots,x_{a_{|A|}}); \tag{35}$$

$$\Gamma_0^A=\mathscr{R}^{|A|}\Big\backslash\left(\bigcup_{i\in A}\Gamma_i^A\right); \tag{36}$$

$$dx^A_{-a_i}=dx_{a_1}\cdots dx_{a_{i-1}}\,dx_{a_{i+1}}\cdots dx_{a_{|A|}}; \tag{37}$$

$$f^A(x^A)=\prod_{i\in A} f_i(x_i); \tag{38}$$

$$\Delta_i^A(z) = \{x_{-i}^A \mid (x_{-i}^A, z) \in \Gamma_i^A\} \qquad \text{if} \quad i \in A; \tag{39}$$

$$\mu_i^A(z) = \int_{\Delta_i^A(z)} \prod_{\substack{j \in A \\ j \neq i}} f_j(x_j)\, dx_{-i}^A \qquad \text{if} \quad i \in A; \tag{40}$$

$$\sigma_i^A(z) = \int_{\mathscr{R}^{|A|-1}} \alpha_i^A(x_{-i}^A, z) \prod_{\substack{j \neq i \\ j \in A}} f_j(x_j)\, dx_{-i}^A. \tag{41}$$

Here (36) defines the set of reports x for which the seller keeps the good. If the x's represent true evaluations, (35) gives the vector of responses when bidder i reports z and everyone else is honest. Equation (39) gives the set of others' valuations for which i wins the good with a report of z. Equation (40) gives the probability that i wins with the report z, while (41) gives i's expected payment with the report z. Let $\sum_{w,H(w)}$ denote the sum over w satisfying $H(w)$.

The incentive-compatibility constraints are

$$\sum_{\substack{A \\ i \in A}} \beta_A[z\mu_i^A(z) - \sigma_i^A(z)] \geqslant \sum_{\substack{A \\ i \in A}} \beta_A[z\mu_i^A(z_0) - \sigma_i^A(z_0)], \qquad \forall\, z, \forall\, z_0. \tag{42}$$

The free-exit constraints are

$$\sum_{\substack{A \\ i \in A}} \beta_A[z\mu_i^A(z) - \sigma_i^A(z)] \geqslant 0, \qquad \forall\, z. \tag{43}$$

THEOREM 4. *The seller maximizes his expected revenue by setting Γ_i^A to satisfy*

$$\Gamma_i^A = \left\{x \in \mathscr{R}^{|A|} \mid x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \geqslant \max\left[x_0, x_j - \frac{1 - F_j(x_j)}{f_j(x_j)}\right]\right\}. \tag{44}$$

Proof. Rewrite the incentive-compatibility constraints

$$\sum_{\substack{A \\ i \in A}} \beta_A[\sigma_i^A(z) - \sigma_i^A(z_0)] \leqslant \sum_{\substack{A \\ i \in A}} \beta_A z[\mu_i^A(z) - \mu_i^A(z_0)] \tag{45}$$

and

$$\sum_{\substack{A \\ i \in A}} \beta_A[\sigma_i^A(z) - \sigma_i^A(z_0)] \geqslant \sum_{\substack{A \\ i \in A}} \beta_A z_0[\mu_i^A(z) - \mu_i^A(z_0)]. \tag{46}$$

Divide by $z - z_0$ and take limits to yield

$$\sum_{\substack{A \\ i \in A}} \beta_A[\sigma_i^{A\prime}(z)] = \sum_{\substack{A \\ i \in A}} \beta_A z \mu_i^{A\prime}(z), \tag{47}$$

should these derivatives exist. The free-exit constraint (43) implies $\sigma_i^A(0) \leqslant 0$; but the seller wants $\sigma_i^A(0)$ to be as large as possible, so that $\sigma_i^A(0) = 0$. Hence

$$\sum_{\substack{A \\ i \in A}} \beta_A \sigma_i^A(z) = \sum_{\substack{A \\ i \in A}} \beta_A \left[z\mu_i^A(z) - \int_0^z \mu_i^A(t)\, dt \right]. \tag{48}$$

The seller expects to earn

$$\begin{aligned}
\Phi &= \sum_A \beta_A \left\{ x_0 \int_{\Gamma_0^A} f^A(x^A)\, dx^A + \sum_{\substack{i \\ i \in A}} \int_{\mathscr{R}^{|A|}} \alpha_i^A(x^A)\, f^A(x^A)\, dx^A \right\} \\
&= \sum_A \beta_A \left\{ x_0 \left[1 - \sum_{\substack{i \\ i \in A}} \int_{\Gamma_i^A} f^A(x^A)\, dx^A \right] \right. \\
&\qquad \left. + \sum_{\substack{i \\ i \in A}} \int_0^\infty \sigma_i^A(z)\, f_i^A(z)\, dz \right\} \\
&= x_0 - x_0 \sum_A \beta_A \sum_{i \in A} \int_{\Gamma_i^A} f^A(x^A)\, dx^A \\
&\qquad + \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{\substack{A \\ i \in A}} \beta_A \int_0^\infty \sigma_i^A(z)\, f_i^A(z)\, dz \\
&= x_0 - x_0 \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{\substack{A \\ i \in A}} \beta_A \int_{\Gamma_i^A} f^A(x^A)\, dx^A \\
&\qquad + \sum_{\substack{i \\ i \in \mathfrak{N}}} \left[\sum_{\substack{A \\ i \in A}} \beta_A \int_0^\infty \left[z\mu_i^A(z) - \int_0^z \mu_i^A(t)\, dt \right] f_i(z)\, dz \right]
\end{aligned}$$

$$= x_0 - x_0 \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{A} \beta_A \int_{\Gamma_i^A} f^A(x^A)\, dx^A$$

$$+ \sum_{\substack{i \\ i \in \mathfrak{N}}} \left[\sum_{\substack{A \\ i \in A}} \beta_A \left\{ \int_{\Gamma_i^A} x_i f^A(x^A)\, dx^A \right.\right.$$

$$+ \left[[1 - F_i(z)] \int_0^z \mu_i^A(t)\, dt \right] \Big|_0^\infty$$

$$\left.\left. - \int_0^\infty [1 - F_i(z)]\, \mu_i^A(z)\, dz \right\} \right]$$

$$= x_0 - x_0 \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{A} \beta_A \int_{\Gamma_i^A} f^A(x^A)\, dx^A$$

$$+ \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{\substack{A \\ i \in A}} \beta_A \left[\int_{\Gamma_i^A} \left[x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \right] f^A(x^A)\, dx^A \right]$$

$$= x_0 + \sum_{\substack{i \\ i \in \mathfrak{N}}} \sum_{\substack{A \\ A \subseteq \mathfrak{N}}} \beta_A \left\{ \int_{\Gamma_i^A} \left[x_i - x_0 - \frac{1 - F_i(x_i)}{f_i(x_i)} \right] f^A(x^A)\, dx^A \right\}$$

$$= x_0 + \sum_{A \subseteq \mathfrak{N}} \beta_A \sum_{i \in A} \int_{\Gamma_i^A} \left\{ x_i - x_0 - \frac{1 - F_i(x_i)}{f_i(x_i)} \right\} f^A(x^A)\, dx^A. \tag{49}$$

Maximizing Φ with respect to Γ_i^A, subject to $\Gamma_i^A \cap \Gamma_j^A = \phi$ for $i \neq j$, $i, j \in A$ yields, from Stokes' Theorem (Sagan [14, p. 542])

$$\Gamma_i^A = \left\{ x_A \mid x_i - x_0 - \frac{1 - F_i(x_i)}{f_i(x_i)} \geqslant \max \left[0, x_j - x_0 - \frac{1 - F_j(x_j)}{f_j(x_j)} \right] \right\}. \tag{50}$$

Condition (50) gives the unconstrained maximizer of Φ. However, Φ must be maximized subject to the constraint that the function $\sum_A \beta_A \mu_i^A(z)$ is nondecreasing in z (which is necessary for incentive compatibility). The regularity condition (33) ensures that the unconstrained maximizer in (50) does in fact satisfy the monotonicity constraint. Q.E.D.

Since the payment function σ_i^A is contingent on the set of bidders A, the seller implicitly reveals the set of bidders when he announces the payment function. However, if the seller wishes to conceal the set of bidders, he can make payment not contingent on A (provided he knows the probabilities β_A) by using the payment function

$$\sigma_i(z) = \sum_{\substack{A \\ i \in A}} \beta_A \sigma_i^A(z) \Big/ \sum_{\substack{A \\ i \in A}} \beta_A. \tag{51}$$

The noteworthy feature of Theorem 4 is that the optimal auction does not depend on the probabilities over the sets of active bidders, β_A. Thus the optimal auction derived by Myerson [14] when the set of bidders is common knowledge remains the optimal auction when the bidders do not know the set of bidders. This implies the following result.

COROLLARY. *The seller's maximum expected revenue with risk-neutral bidders having independent private values is the same whether or not the bidders know the set of bidders.*

Consider the case in which $F_i = F_j = F$ for all i, j: all bidders draw their valuations from the same distribution. Then Theorem 4 says that the seller optimally sets a reserve price r satisfying $x_0 = r - (1 - F(r))/f(r)$ and chooses the highest remaining bidder (as in Myerson [14] and Riley and Samuelson [15]). If $F_i \neq F_j$ for some i, j, then reserve prices are again used but the optimal auction discriminates against certain bidders, in that a lower-valuation bidder can win the item despite the presence of a higher-valuation bidder. (See Myerson [14] for details.)

How can the optimal auction be implemented in practice in the case of symmetric bidders (i.e., $F_i = F_j$ for all i, j)? When the set of bidders is common knowledge, it is well known that the optimal direct, incentive-compatible auction can be mimicked by either an English auction or a first-price sealed-bid auction, provided appropriate reserve prices are imposed (Milgrom [11], Myerson [14], Riley and Samuelson [15]). Does this remain the case when the bidders do not know the set of bidders? In the case of symmetric expectations, it does. If the bidders' expectations over the set of bidders, while Bayesian consistent, are not identical, then the optimal auction can be implemented using an English auction, because in the English auction, the second-last bidder drops out of the bidding when the bids reach the value of the second order statistic; this gives the seller the same expected revenue as the optimal direct auction of Theorem 4. However, it is important to note that the optimal auction cannot be implemented by a first-price sealed-bid auction when bidders have different expectations. This is because the bidding functions of different bidders will fail to coincide; as a result, it is possible that the highest bidder in the first-price sealed-bid auction is not the appropriate winner as defined by the optimal Γ_i^A. Note also that the converse of this applies: if the bidders have different expectations over the set of bidders, then only for a measure-zero set of F_i's is the first-price sealed-bid auction optimal. Thus we have a result contrary to the Revenue-Equivalence Theorem, even though bidders are symmetric, risk neutral, and have independent private values.

5. Conclusion

The results of auction theory are sensitive to the assumption that the set of bidders is common knowledge. In a first-price sealed-bid auction with bidders who have independent private values and are risk averse (with constant absolute risk aversion), the expected selling price is strictly higher when the bidders do not know how many other bidders there are than when they do know this. In an ex ante sense, any bidder is indifferent between the policy of being told and the policy of not being told the number of bidders. With risk-neutral bidders, the optimal auction is the same whether or not the bidders know who their competitors are. However, this optimal auction may not be implementable using a first-price sealed-bid auction, although it is implementable using an English auction.

More generally, in a first-price sealed-bid auction with bidders who are risk-averse and have affiliated private values, the release of any affiliated information by the seller generates two opposing tendencies: the linkage effect, identified by Milgrom and Weber [12], which tends to raise the selling price; and the bid-dispersion effect identified above, which tends to lower the selling price.

The limits of the foregoing results should be stressed. As Milgrom and Weber [12] argued, the independent-private-values assumption is restrictive: it requires that each bidder has no doubt about the value of the item to him, and that there be no possibility of reselling the item later at some as yet unknown price. The analysis of Section 3 assumed constant absolute risk aversion: Matthews [6] showed that some of the results change when this assumption is relaxed. The results of Section 4 depend upon the assumption of risk neutrality, as is shown by the results on optimal auctions with risk-averse bidders of Maskin and Riley [3], Matthews [4, 5], and Moore [13].

References

1. R. Cassady, Jr., "Auctions and Auctioneering," Univ. of California Press, Berkeley, 1967.
2. E. Maskin and J. Riley, "Monopoly with Incomplete Information," Working Paper No. 268, UCLA, August 1983.
3. E. Maskin and J. Riley, Optimal auctions with risk averse buyers, *Econometrica* **52** (1984), 1473–1518.
4. S. A. Matthews, Selling to risk averse buyers with unobservable tastes, *J. Econ. Theory* **30** (1983), 370–400.
5. S. A. Matthews, On the implementability of reduced form auctions, *Econometrica* **52** (1984), 1519–1522.
6. S. A. Matthews, Comparing auctions for risk averse buyers: a buyer's point of view, *Econometrica*, in press.

7. R. P. McAfee and J. McMillan, Bidding for contracts: A principal-agent analysis, *Rand J. Econ.* **17** (1986), 326–338.
8. R. P. McAfee and McMillan, "Incentives in Government Contracting," Univ. of Toronto Press, Toronto, 1987.
9. R. P. McAfee and J. McMillan, Auctions and bidding, *J. Econ. Lit.* **25** (1987), 708–747.
10. P. Milgrom, The economics of competitive bidding: A selective survey, *in* "Social Goals and Social Organization" (L. Hurwicz, D. Schmeidler, and H. Sonnenschein, Eds.), Cambridge Univ. Press, Cambridge, 1985.
11. P. Milgrom, Auction theory, *in* "Advances in Economic Theory 1985" (T. Bewley, Ed), Cambridge Univ. Press, Cambridge, 1987, in press.
12. P. Milgrom and R. J. Weber, A theory of auctions and competitive bidding, *Econometrica* **50** (1982), 1089–1122.
13. J. Moore, Global incentive constraints in auction design, *Econometrica* **52** (1984), 1523–1536.
14. R. B. Myerson, Optimal auction design, *Math. Oper. Res.* **6** (1981), 58–73.
15. J. Riley and W. Samuelson, Optimal auctions, *Amer. Econ. Rev.* **71** (1981), 381–392.
16. H. Sagan, "Advanced Calculus," Houghton–Mifflin, Boston, 1974.

Printed by Catherine Press, Ltd., Tempelhof 41, B-8000 Brugge, Belgium