Auction Theory with Private Values

Eric S. Maskin; John G. Riley

The American Economic Review, Vol. 75, No. 2, Papers and Proceedings of the Ninety-Seventh Annual Meeting of the American Economic Association (May, 1985), 150-155.

Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198505%2975%3A2%3C150%3AATWPV%3E2.0.CO%3B2-A

The American Economic Review is currently published by American Economic Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/aea.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org/
Mon Jul 26 17:24:20 2004

Auction Theory with Private Values

By ERIC S. MASKIN AND JOHN G. RILEY*

For many centuries, auctions have been a common form of selling procedure. Although auction methods vary across country and product, the two most frequently observed are the open, ascending bid (or English) auction and the sealed-bid auction. Recent theoretical research has led to a theory of equilibrium bidding in these two auctions and a wide range of alternatives as well. As a result it has been possible to compare the revenue extracted by the seller under different auction methods and even to characterize the revenue-maximizing auction.

The Revenue Equivalence Theorem (see for example, William Vickrey, 1961, Roger Myerson, 1981, and Riley and William Samuelson, 1981) asserts that when each bidder's reservation price for a unit of an indivisible good is an independent draw from the same distribution, and bidders are risk neutral, the sealed-bid auction generates the same expected revenue as the open auction. Much recent research has involved weakening each of the main hypotheses—risk neutrality, identically distributed values, and independence of values—in turn. We shall illustrate some of the principal conclusions of this work by considering the properties of open and sealed-bid auctions in a model of two bidders whose reservation prices can assume only two values, and by comparing these auctions to the "optimal" or revenue-maximizing auction.

†*Discussant*: William F. Samuelson, School of Management, Boston University.

*Departments of Economics; Harvard University, Cambridge, MA 02138, and University of California, Los Angeles, CA 90024, respectively. We are indebted to William Samuelson for helpful suggestions. We thank the Sloan Foundation and the NSF for financial support.

I. Revenue Equivalence

Imagine that the reservation price of bidder i $(i=1,2)$ can assume the values v_H (with probability p) and v_L (with probability $1-p$), where $v_H > v_L \geq 0$. Bidders' values are private information and independently distributed. Bidders are risk neutral, that is, they maximize the expression

$$(1)\quad (\text{probability of winning})v - \text{expected payment}.$$

We suppose that an open auction proceeds by the auctioneer's continuously raising the asking price. The auction concludes when one of the bidders drops out. The remaining bidder is the winner and pays the dropout price (if both bidders drop out simultaneously, a coin is flipped to determine the winner). Given these rules, one can easily confirm that a bidder's unique (perfect) equilibrium strategy is to drop out when the asking price reaches his reservation price. (There are other "nonperfect" equilibria, see our 1983a paper). Thus the expected payoff of a v_L bidder (a bidder whose reservation price is v_L) is zero, and his probability of winning is $\frac{1}{2}(1-p)$. The expected payoff of a v_H bidder, by contrast, is his surplus if the other bidder is "low" (since then the asking prices only reaches v_L rather than v_H) times the probability of that event, that is, $(1-p)(v_H - v_L)$. Since a v_H bidder wins when the other bidder has a low value and wins half the time when the other bidder has a high value, his probability of winning is $\frac{1}{2}p + (1-p)$.

In the sealed-bid auction, bidders submit bids simultaneously. The higher bidder is the winner (ties again are resolved by coin flips) and he pays his bid. Consider a symmetric equilibrium. Because the distribution of values is discrete, the equilibrium will involve mixed strategies. Notice first that a v_L bidder (one whose reservation price is v_L) will never bid more than v_L because, if he did, the maximum of such bids (if bidders use mixed strategies that randomize over a variety of alternative bids) would win the auction with positive probability, inducing a negative expected payoff. Let $\underline{b}_L$ be the infimum of all bids submitted.

Suppose first that $\underline{b}_L < v_L$. Then bidders bid below v_L with positive probability and so a v_L bidder's expected payoff is positive. Suppose, furthermore, that bidder 1 bids $\underline{b}_L$ with positive probability. Then bidder 2's chances of winning increase discontinuously if he bids just more then $\underline{b}_L$ while his payment if he wins scarcely rises, thereby raising his expected payoff. But this is a violation of symmetry. On the other hand, if $\underline{b}_L$ is not bid with positive probability, then bids near $\underline{b}_L$ have almost no chance of winning, contradicting the positive expected payoff.[1]

Next let $\underline{b}_H$ be the infimum of bids made by a v_H bidder. If $\underline{b}_H > v_L$, then a bid strictly between $\underline{b}_H$ and v_L has the same chance of winning as $\underline{b}_H$, and so is preferable. Thus $\underline{b}_H = v_L$, and a v_H bidder's expected payoff must be $(v_H - v_L)(1-p)$. In equilibrium, any bid b made as part of a mixed strategy must generate the same expected payoff. Therefore if $F(b)$ is the cumulative distribution function of a v_H bidder's bid, it satisfies

$$(2) \quad [pF(b)+1-p](v_H - b) = (1-p)(v_H - v_L).$$

By symmetry, a v_H bidder's expected probability of winning is $\frac{1}{2}p+(1-p)$, whereas that of a v_L bidder is $\frac{1}{2}(1-p)$. Because a given type of bidder's probability of winning and expected payoff are the same in the open and sealed-bid auctions, formula (1) implies that his expected payment is the same in the two auctions. We have established the Revenue Equivalence Theorem for our model. Indeed, we obtain the same expected revenue from any other auction in which the high bidder wins, the expected payoff of a v_L bidder is zero, and the expected payoff of a v_H bidder is $(1-p)(v_H - v_L)$.

It is of some interest to compare the open and sealed-bid auctions with a revenue-maximizing auction (see Myerson and Riley-Samuelson). Suppose that bidders were offered the choice between bidding v_L or $b_H = (\frac{1}{2}v_H + \frac{1}{2}(1-p)v_L)/(\frac{1}{2}p+1-p)$, with, as always, the high bidder winning. Because b_H is greater than v_L, a v_L bidder will bid v_L. Since $(\frac{1}{2}p+1-p)(v_H - b_H) = \frac{1}{2}(1-p)(v_H - v_L)$, a v_H bidder is indifferent between bidding b_H and v_L, and so might as well choose the former. Since a v_H bidder bidding b_H has the same probability of winning as in an open or sealed-bid auction $(\frac{1}{2}p+(1-p))$, but has a lower expected payoff, $(\frac{1}{2}(1-p)(v_H - v_L)$ rather than $(1-p)(v_H - v_L)$, his expected payment must be higher. Thus, this alternative auction generates higher expected revenue. Indeed, it is *optimal* if $v_L > pv_H$. (If $v_L < pv_H$ it is optimal to set a reserve price at v_H, thereby rejecting all lower bids.) In either case, the optimal auction differs from the open and sealed-bid auctions by prohibiting bidders from making certain bids. This conclusion generalizes to more complicated models, including those with a continuum of possible reservation prices.

II. Risk Aversion

Let us modify the model of Section I only by supposing that bidders are risk averse. Let u be a strictly concave von Neumann-Morgenstern utility function, normalized so that $u(0)=0$. A v bidder's payoff if he wins and pays t is $u(v-t)$; his payoff if he loses and pays t is $u(-t)$.

Risk aversion does not alter the bidders' behavior in the open auction; it is still optimal for a bidder to drop out exactly when his reservation price is reached. Hence expected revenue is as before. In the sealed-bid auction, v_L bidders continue to bid v_L, and if F_R is the cumulative distribution function

[1] Our argument here presumes that the equilibrium in the sealed-bid auction is symmetric. One can show (see our 1983a paper) that there is no asymmetric equilibrium.

of a v_H bidder's bid, it satisfies the analogue of condition (2):

$$(3) \quad u(v_H - b)[(1-p) + pF_R(b)] = u(v_H - v_L)(1-p).$$

The strict concavity of u implies that $u(v_H - v_L)/u(v_H - b) < (v_H - v_L)/(v_H - b)$ for $v_L < b < v_H$. Hence, (2) and (3) imply that $F_R(b) \le F(b)$ with strict inequality for bids greater than v_L but less than the maximum. That is, F_R stochastically dominates F, and so the expected bid by a v_H bidder is higher with risk aversion than without. We conclude that, with risk aversion, a sealed-bid auction generates greater expected revenue than an open auction (see Gerard Butters, 1975, and Charles Holt, 1980). Intuitively, increasing a bidder's risk aversion heightens his fear of losing and so, in a sealed-bid auction, induces him to bid higher. Viewed alternatively, a sealed-bid auction, unlike an open auction, insures a winning bidder against fluctuations in the amount he has to pay, and a risk-averse bidder is willing to pay a premium—in the form of a higher bid—for this insurance.

By requiring payments even of losing bidders, an optimal auction (see our 1984 article, and Steven Matthews, 1983) can exploit the fact that a risk-averse bidder's marginal utility of income depends on whether he wins or loses. Let π_i be the probability of winning and b_i and a_i the payments by a winning and losing bidder, respectively, of type i $(i = L, H)$. An optimal auction chooses π_i, b_i, and a_i to maximize

$$(4) \quad p(\pi_H b_H + (1-\pi_H) a_H) + (1-p)(\pi_L b_L + (1-\pi_L) a_L),$$

subject to

$$(5) \quad \pi_H u(v_H - b_H) + (1-\pi_H) u(-a_H) \ge \pi_L u(v_H - b_L) + (1-\pi_L) u(-a_L)$$

$$(6) \quad \pi_L u(v_L - b_L) + (1-\pi_L) u(-a_L) \ge 0$$

$$(7) \quad \tfrac{1}{2}p + (1-p) \le \pi_H$$

$$(8) \quad \tfrac{1}{2} \ge p\pi_H + (1-p)\pi_L$$

$$(9) \quad \pi_H \ge 0 \quad \text{and} \quad \pi_L \ge 0.$$

Constraint (5), a self-selection constraint, ensures that a v_H bidder is at least as well off making a high as a low bid. We have omitted the analogous self-selection constraint for a v_L bidder since, as we shall see, it is satisfied automatically. Constraint (6) guarantees a v_L bidder a nonnegative expected payoff from participating. (Given (5), a v_H bidder's payoff will also be nonegative.) Condition (7) says that a v_H bidder can win with at most probability 1 if the other bidder has a low reservation price and, given the symmetry of the model, with at most probability $\frac{1}{2}$ if the other bidder's reservation price is high. Constraint (8) requires simply that each bidder's probability of winning, *unconditional* on his reservation price, not exceed $\frac{1}{2}$.

Letting α and β be the Lagrange multipliers for (5) and (6), respectively, we obtain the first-order conditions

$$(10) \quad p\pi_H - \alpha\pi_H u'(v_H - b_H) = 0$$

$$p(1-\pi_H) - \alpha(1-\pi_H) u'(-a_H) = 0$$

$$(11) \quad (1-p)\pi_L + \alpha\pi_L u'(v_H - b_L) - \beta\pi_L u'(v_L - b_L) = 0$$

$$(1-p)(1-\pi_L) + \alpha(1-\pi_L) u'(-a_L) - \beta(1-\pi_L) u'(-a_L) = 0.$$

From (10) we find that $v_H - b_H = -a_H$, that is, a high bidder is perfectly insured; he receives a monetary transfer $-a_H(>0)$, as compensation if he loses. From (11) and the fact that $u'(v_H - b_L) < u'(v_L - b_L)$,

$$(12) \quad (\beta - \alpha) u'(-a_L) = 1 - p > (\beta - \alpha) u'(v_L - b_L).$$

Thus a v_L bidder is better off winning than losing $(v_L - b_L > -a_L)$. Moreover, since (from (12)) (6) is binding, he must actually pay a penalty if he loses $(a_L > 0)$, which we can interpret a as an entry fee. Because (5) is binding and $v_H - b_L > -a_L$, we have $v_H - b_H < v_H - b_L$, that is, a v_H winner pays more than a v_L winner. If (8) is binding, as it will be if p is small enough, we can solve for π_L and rewrite (4) as $p\pi_H(b_H - a_H - b_L + a_L) + pa_H + (\frac{1}{2} - p)a_L$. From the above argument, $b_H - a_H - b_L + a_L > v_H - v_L > 0$. Hence, constraint (7) is binding: $\pi_H = \frac{1}{2}p + (1-p)$.

We conclude that an optimal auction with risk-averse bidders resembles that for risk-neutral bidders. Bidders are offered the choice between two prices b_H and b_L (if, as before, p is not too high), and the high bid wins. However, if a bidder loses with a bid of b_H, he is compensated for losing, whereas if he loses with a bid of b_L, he is penalized. Intuitively, introducing a penalty heightens a risk-averse bidder's fear of losing and therefore increases the revenue that can be extracted from a v_H bidder. Of course, this penalty, by increasing risk, reduces the payment that a v_L bidder makes. But the penalty has no effect to the first-order, since, with no penalty, a v_L bidder is perfectly insured.

It remains only to show that the solution to the program of maximizing (4) subject to (5)–(9) satisfies

$$(13)\quad \pi_L u(v_L - b_L) + (1-\pi_L)u(-a_L) \geq \pi_H u(v_L - b_H) + (1-\pi_H)u(-a_H),$$

the self-selection constraint for v_L bidders. But (13) follows immediately from the fact that (5) holds with equality and $\pi_H u'(v - b_H) > \pi_L u'(v - b_L)$ (since $\pi_H > \pi_L$ and $b_H > b_L$) for all v.

III. Asymmetry

Let us revert to risk neutrality but now drop the assumption that valuations are identically distributed. Specifically, assume that bidder 1's reservation price is distributed as in Section I, but that bidder 2's reservation price is either w_H or w_L with probabilities q and $1-q$, respectively. Continue to suppose that the two bidders' distributions are independent. For convenience, let us suppose that $v_L = w_L = 0$. Then the expected revenue generated by the open auction is

$$(14)\qquad pq\min\{v_H, w_H\}.$$

We wish to compare the difference in revenues, Δ, between the sealed-bid and open auctions.[2] To do this we shall consider two polar cases of asymmetry: (i) both bidders have the same probability of being high but have different high values, that is, $p = q$ and $v_H \neq w_H$, and (ii) both bidders have the same high values but different probabilities, that is, $v_H = w_H$ and $p \neq q$.

It is not difficult to see that in case (i), Δ is positive. We know from Section I that when $v_H = w_H$, Δ is zero. Now imagine raising w_H above v_H. This does not affect revenue from the open auction since there is no change in the distribution of the second highest reservation value. However, with a higher w_H, the optimal response in the sealed-bid auction by bidder 2 (when $v = w_H$) to bidder 1's equilibrium strategy is a higher bid. Bidder 2's higher bid, in turn, induces bidder 1 to bid higher than before (for details, see our 1983b paper). Hence, revenue from the sealed-bid auction rises, and Δ becomes positive.

In case (ii), expected revenue in the open auction is pqv_H. In the sealed-bid auction, the equilibrium cumulative distribution functions, F_1 and F_2, of the bids of bidders 1 and 2, when their reservation prices are v_H, satisfy the analogue of (2):

$$(15)\quad (1 - q + qF_2(b))(v_H - b) = (1 - q + qF_2(0))v_H;$$

$$(16)\quad (1 - p + pF_1(b))(v_H - b) = (1 - p + pF_1(0))v_H.$$

[2] As our model is formulated, an equilibrium in the sealed-bid auction may not exist. The nonexistence problem, however, is an artifact of our allowing literally a continuum of possible bids. In fact, we can restore existence even with a continuum by allowing the possibility of positive but infinitesimal bids, which we implicitly assume in our analysis.

Notice that right-hand sides of (15) and (16) allow for the possibility that a v_H bidder will bid zero (actually, infinitesimally more than zero) with positive probability. This will be the case if $p \neq q$ since both bidders must make the same maximum bid,[3] $\bar{b}_H$, when their reservation price equals v_H, and (15) and (16) can be satisfied for $b = \bar{b}_H$ only if one of $F_1(0)$ and $F_2(0)$ is nonzero. For example, if $p > q$, then (15) and (16) imply that

$$\bar{b}_H = qv_H = pv_H(1 - F_1(0)),$$

and so $F_1(0) = 1 - q/p$. Integrating (15), we obtain qzv_H as the expected payment by bidder 1 if his reservation price is v_H, where $z = \int F_2(b)\,dF_1(b)$. Similarly, from (16), the expected payment by bidder 2 is $(p(1-z)+q-p)v_H$. Hence total expected revenue is q^2v_H, which is less than the open auction revenue, pqv_H. Therefore, for case (*ii*), Δ is negative.

Roughly speaking, the sealed-bid auction generates more revenue than the open auction when bidders have distributions with the same shape (but different supports), whereas the open auction dominates when, across bidders, distributions have different shapes but approximately the same support.

IV. Correlation

Let us return to the model of Section I, except now assume that reservation prices are correlated across bidders. Specifically, let r_{ij} $(i, j \in \{L, H\})$ be the joint probability that bidder 1's value is v_i and that bidder 2's value is v_j. Correlation implies that

$$(17) \qquad r_{HH}r_{LL} - r_{HL}r_{LH} \neq 0.$$

As usual, behavior in the open auction remains the same, and so expected revenue is

$$(18) \qquad r_{HH}v_H + (1 - r_{HH})v_L.$$

Making the obvious modifications in the analysis of Section I, we conclude that expected revenue for the sealed-bid auction is also (18). This equivalence between the two auctions does not generalize to distributions with more than two point supports because, in general, with correlation, a higher reservation price does not imply a higher bid for the sealed-bid auction (although it does for the open auction).[4] Any condition sufficient to guarantee that bids are monotonic in reservation prices, however, ensures equivalence. One such condition is that the reservation prices be affiliated (see Paul Milgrom and Robert Weber, 1982).

When (17) holds, an optimal auction extracts all surplus from bidders (see Jacques Crémer and Richard McLean, 1985). To see this, let c_{ij} $(i, j \in \{L, H\})$ be the payment that bidder 1 makes when his $v = v_i$ and bidder 2's $v = v_j$. To extract all surplus, the c_{ij}s must satisfy

$$(19) \quad \tfrac{1}{2}r_{LL}v_L - r_{LH}c_{LH} - r_{HH}c_{HH} = 0$$

$$(20) \quad \left(\tfrac{1}{2}r_{LH} + r_{LL}\right)v_L - r_{LH}c_{HH} - r_{LL}c_{HL} < 0$$

$$(21) \quad \left(\tfrac{1}{2}r_{HH} + r_{HL}\right)v_H - r_{HH}c_{HH} - r_{HL}c_{HL} = 0$$

$$(22) \quad \tfrac{1}{2}r_{HL}v_L - r_{HH}c_{LH} - r_{HL}c_{LL} < 0.$$

Equations (19) and (21) require the surplus of v_L and v_H bidders, respectively, to be zero. Inequality (20) ensures that a v_L bidder is not better off bidding as a v_H bidder, and (22) imposes the corresponding constraint on a v_H bidder. But from (17), we can solve for c_{ij}s that satisfy (19)–(22).

[3] If, say, bidder 1's maximum bid were greater than that of bidder 2, bidder 1 could lower his bid without reducing his probability of winning.

[4] Suppose, for example, that v can take on three possible values: $v_H > v_M > v_L$. Assume that if $v = v_H$ for one bidder, then it is very likely that $v = v_L$ for the other bidder. Assume further that if $v = v_M$ for one bidder, then the other bidder in all likelihood has the same reservation price. In this case, a v_M bidder will bid higher on average than a v_L bidder in the sealed-bid auction. Furthermore, the sealed-bid auction, at least for some parameter values, generates strictly more revenue than does the open auction.

V. Concluding Remarks

We have discussed three major hypotheses of the Revenue Equivalence Theorem, but there remain two more implicit in our formulation. One is the assumption that only a single item is sold. If buyers have downward-sloping demand curves and there are multiple units for sale, the Revenue Equivalence Theorem again fails. Extrapolating from some simple examples, we conjecture that open bidding will tend to dominate sealed bidding in this environment.

The second assumption is that a bidder's reservation price does not affect the reservation price of any other bidder. This is the "private values" hypothesis: the assumption that reservation prices are a matter of taste rather than a reflection of information about the intrinsic value of the good. In the latter case, the "common values" model, the open auction tends to produce higher revenue than the sealed-bid auction when our other hypotheses are maintained (see Milgrom and Weber).

REFERENCES

Butters, G. R., "Equilibrium Price Distributions and the Economics of Information," unpublished doctoral dissertation, University of Chicago, 1975.

Crémer, J. and McLean, R. P., "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 1985, *53*, forthcoming.

Holt, C., "Competitive Bidding for Contracts under Alternative Auction Procedures," *Journal of Political Economy*, June 1980, *88*, 433–45.

Maskin, E. S. and Riley, J. G., (1983a) "Uniqueness of Equilibrium in Open and Sealed Bid Auctions," mimeo., April 1983.

______ **and** ______, (1983b) "Auctions With Asymmetric Beliefs," Discussion Paper No 254, University of California-Los Angeles, June 1983.

______ **and** ______, "Optimal Auctions With Risk Averse Buyers," *Econometrica*, November 1984, *52*, 1473–518.

Matthews, S. A., "Selling to Risk Averse Buyers With Unobservable Tastes," *Journal of Economic Theory*, August 1983, *30*, 370–400.

Milgrom, P. and Weber, R. J., "A Theory of Auctions and Competitive Bidding," *Econometrica*, September 1982, *50*, 1089–122.

Myerson, R., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6*, 58–73.

Riley, J. G. and Samuelson, W. F., "Optimal Auctions," *American Economic Review*, June 1981, *71*, 381–92.

Vickrey, W., "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, March 1961, *16*, 8–37.