Economics Letters 26 (1988) 15–19
North-Holland

REVENUE EQUIVALENCE IN MULTI-OBJECT AUCTIONS

Richard ENGELBRECHT-WIGGANS

University of Illinois, Champaign, IL 61820, USA

Received 4 September 1987

The theory of auctions suggests that, under certain conditions, seemingly different methods for auctioning a single object will result in the same expected price. This result extends quite easily to the case of more than one object.

1. Introduction

The theory of auctions and competitive bidding suggests that, under certain conditions, seemingly different auction mechanisms result in the same expected cost or revenue to the bid-taker. In particular, the existing results assume bidders (i) to be risk-neutral, (ii) to obtain private information only by observing independently distributed signals, (iii) to face only the two possible outcomes of 'win' and 'lose', and (iv) to bid as if they were following Nash equilibrium bidding strategies. Under such conditions, the differences in how bidders bid in response to different pricing rules offsets the differences in the rules themselves.

The extent to which such a theory predicts what actually happens in the real world depends on how well the theory models the essence of actual situations, and on how sensitive the theory is to its assumptions. The literature has already established two of the assumptions – that of risk neutrality and that of independent signals – to be pivotal. This paper examines a third assumption.

In particular, we allow each bidder to be faced by more than two possible outcomes; for example, a bidder may be allowed to win any subset of the objects offered for sale. Then, for risk-neutral bidders with independent private values (for each possible outcome), the expected revenue or cost to the bid-taker at equilibrium depends only on (1) the number of bidders, (2) the distribution of each bidder's values, (3) the relationship between bidders' values and who wins what, and (4) the bidder's expected payments in some benchmark case. This establishes that revenue equivalence is not sensitive to the number of possible outcomes faced by each bidder.

Vickrey (1961) examines the possibility that different auction formats might give the same expected revenue to a seller of a single object. In particular, Vickrey considers symmetric models of risk-neutral bidders, models in which each of the known number of bidders knows his own value for the object being sold, but knows nothing about others' values except that they are independent samples from some known distribution. He then considers two auction formats. In the first, the high bidder wins and pays an amount equal to his winning bid. In the second – an approximation to the common oral auction – the high bidder wins, but pays an amount equal to the second highest bid. At equilibrium in the stated independent private values model, these two auction formats yield the same expected revenue for the seller.

Myerson (1981) examines the possibility that different single object auction formats might give the same expected revenue to the seller more generally. He suppresses the details of the auction rules themselves. Instead, he recognizes that each equilibrium in each auction game gives rise to some functional relationship between bidders' values and the outcome of the auction – who wins the object, and who pays how much. In order to arise from the equilibrium to some auction, this functional relationship must satisfy certain properties. Given these properties, Myerson shows that the expected revenue in an auction with independent private values depends only on (1) the number of bidders, (2) the distribution of a bidder's value, (3) the allocation rule – that is, the relationship between the bidders' values for the object and who wins the object – and (4) the expected price or profit in some benchmark situation (for example, when all bidders have the lowest possible value for the object). This generalizes Vickrey's result considerably.

Milgrom and Weber (1982) go a slightly different direction. Instead of assuming privately known values, they simply assume that each bidder knows how his value for the object depends on each bidder's private information. In addition, they only consider three specific auction mechanisms – the sealed bid first price mechanism and two models of the progressive oral auction. In the case of independent private values, all three mechanisms generate the same expected price for the object at equilibrium. This result gives up the sweeping generality of mechanisms allowed by Myerson, but in the process allows for dependent – or even common – values, and establishes that it is the independence of private information (rather than of values) on which the revenue equivalence really depends.

Vickrey (1962) also considers auctions with more than one, identical objects; each of the appropriate number of highest bidders wins one object. Three different auction formats receive attention: (i) each winner pays the amount of his bid, (ii) each winner pays an amount equal to the lowest winning bid, and (iii) each winner pays an amount equal to the highest losing bid. Again for the case of independent private values, at equilibrium, each of these formats yields the same expected revenue for the seller.

Weber (1983) goes on to a wider family of mechanisms – those in which each of the bidders with the highest values wins the object (in other words, mechanisms that efficiently award at most one object to a bidder). In addition, he allows a bidder's value to depend on others' private information, just as Milgrom and Weber allow for the case of a single object. Then, each such mechanism results in the same equilibrium price so long as it meets some appropriately specified boundary condition; for example, the condition might be that if a bidder has the lowest possible value for winning an object and still wins (which means that all others must also have the same lowest possible value for the object), then a bidder has an expected profit of zero from winning. Engelbrecht-Wiggans (1987) extends this result to the case of a random number of objects – the number being independent of bidder's private information – in analyzing alternatives to the first price mechanism used by the Department of Agriculture in the Dairy Termination Program.

This paper generalizes Vickrey's results for multi-object auctions in much the same manner and direction as Myerson did for Vickrey's results on single object auctions. We take a 'direct revelation' approach based on that of Myerson, and establish that if each risk-neutral bidder knows his own value for each possible allocation of objects to himself (and others' values are independent of his own values and are unknown to him), then the seller's expected revenue (given certain regularity) depends only on (1) the number of bidders, (2) the possible outcomes to each bidder, (3) the relationship between bidders' values and the final allocation, and (4) the expected price or profit in some specific benchmark situation. This generalizes the results of Vickrey in a direction different from that of Weber – we require the more restrictive independent private values assumption, but thereby can allow bidders to face more than one non-trivial possible outcome, and allow mechanisms that result in inefficient outcomes. In addition, were it not for the implicit regularity assumed (to drastically simplify the mathematical analysis), this paper would also generalize Myerson's results.

2. The model

This section defines our model of auctions a model both of the environment within which the auction takes place as well as of the bidding itself. In particular, imagine a fixed number of bidders, each risk-neutral; the number may be random, but its distribution must be fixed in the sense of being exogenously specified. We look at the problem from the viewpoint of bidder i. The problem looks the same to each other bidder, but lumping all of the other bidders together simplifies our notation and analysis.

As far as i is concerned, the auction may result in any one of m different non-trivial allocations of the objects. Perhaps i cares only about what objects he wins. On the other hand, perhaps i also cares about who else wins what objects. We allow either possibility.

Let $\boldsymbol{x} = (x_1, x_2, \dots, x_m)$ denote the outcome of a vector valued random variable $\boldsymbol{X}$. Assume that i knows the outcome $\boldsymbol{x}$, and that x_j denotes his value for outcome j. Of course, i need not truthfully reveal his actual $\boldsymbol{x}$ in the auction, so let $\boldsymbol{y}$ denote what he does reveal; $\boldsymbol{y}$ may, but need not, equal $\boldsymbol{x}$.

The other bidders also discover something about their own values. Let $\boldsymbol{w}$ denote the outcome of the vector valued random variable $\boldsymbol{W}$. Think of $\boldsymbol{w}$ as being the concatenation of all other bidders' value vectors. Assume that i does not know anything about $\boldsymbol{w}$ other than that (1) it is an outcome of $\boldsymbol{W}$, (2) the distribution $F(\cdot)$ of $\boldsymbol{W}$, and (3) that $\boldsymbol{W}$ is (statistically) independent of $\boldsymbol{X}$. Just as with i, the other bidders need not reveal $\boldsymbol{w}$ truthfully; let $\boldsymbol{z}$ denote what they do reveal.

Now imagine any bidding game – a game in which known rules translate the bidders' bids into who wins what and who pays how much. Bidders must decide how to bid, and presumably, their bids depend on what they know about their values for the possible outcomes; we call the relationship of a bidder's bids to his information his 'bidding strategy'. If we have a vector of bidding strategies, one for each bidder, such that no one bidder can do better than follow his specified strategy given that all other bidders follow their strategies, then we call this vector of strategies an 'equilibrium'. This equilibrium models the bidders' behavior; specifically, if a game has at least one equilibrium, then we presume that the bidders will bid as if they were following their respective strategies in some one of the equilibria. (Note that we are simply describing – not prescribing – how bidders bid; we do not presume that bidders explicitly calculate equilibrium strategies, but merely that they act as they would have had they calculated such strategies and followed them.)

Each equilibrium to any specific bidding game induces a specific relationship between the bidders' information and the outcome of the game. In particular, let $\boldsymbol{p}(\boldsymbol{y}, \boldsymbol{z})$ describe the allocation (as it affects i) if the bidders had bid as if they had seen $\boldsymbol{y}$ and $\boldsymbol{z}$ (rather than $\boldsymbol{x}$ and $\boldsymbol{w}$), but had still followed their respective equilibrium bidding strategies. Specifically, $p_j(\boldsymbol{y}, \boldsymbol{z})$ is the probability of outcome j to i. Likewise, let $c(\boldsymbol{y}, \boldsymbol{z})$ describe how the expected amount paid by i depends on the $\boldsymbol{y}$ and $\boldsymbol{z}$ that the bidders plug into their respective bidding strategies.

The fact that $\boldsymbol{p}(\boldsymbol{y}, \boldsymbol{z})$ and $c(\boldsymbol{y}, \boldsymbol{z})$ describe how the outcome of the auction to i depends on the $\boldsymbol{y}$ and $\boldsymbol{z}$ that bidders plug into their respective equilibrium bidding strategies allows us to suppress the inner workings of the auction itself. In fact, imagine a game in which specified functions $\boldsymbol{p}(\cdot)$ and $c(\cdot)$ directly determine the outcome as a function of the $\boldsymbol{y}$ and $\boldsymbol{z}$ that the bidders report; this is the 'direct revelation' game of Myerson. If the specified $\boldsymbol{p}(\cdot)$ and $c(\cdot)$ come from an actual equilibrium to some bidding game and if i presumes that the other bidders will report $\boldsymbol{z} = \boldsymbol{w}$, then i can do no better than to report $\boldsymbol{y} = \boldsymbol{x}$. Specifically, any direct revelation game arising from an equilibrium to a bidding game has an equilibrium in which each bidder truthfully reveals his information. Furthermore, any direct revelation game can itself be viewed as being a bidding game. Thus, the set of equilibria to bidding games corresponds to the set of equilibria to direct revelation games in which each bidder reveals his information truthfully. This allows us to study auctions quite generally by simply studying truthful equilibria in direct revelation games.

3. The analysis

This section examines the direct relation games arising from multi-object auctions with independent private values. In particular, we derive an expression for the expected profit to i, as a function of $\boldsymbol{p}(\cdot)$, $c(\cdot)$, and $F(\cdot)$, when i sees $\boldsymbol{x}$, but reports y. For this $\boldsymbol{p}(\cdot)$ and $c(\cdot)$ to have come from an equilibrium to some bidding game, the expected profit to i when $z = \boldsymbol{w}$ must be maximized at $y = \boldsymbol{x}$. This condition restricts how $c(\cdot)$ can depend on $\boldsymbol{p}(\cdot)$ and $F(\cdot)$. Indeed, given certain regularity, $c(\cdot)$ would now be uniquely specified as soon as we fix its value for some benchmark $\boldsymbol{x} = \boldsymbol{x}_0$ – say, the lowest value that bidder i can have for each outcome.

Specifically, start by deriving an expression for the expected profit $U(\boldsymbol{x}, y)$ to i of reporting y when he actually saw $\boldsymbol{x}$, assuming all along that others report $z = \boldsymbol{w}$.

$$U(\boldsymbol{x}, y) = \int_{\boldsymbol{w}} (\boldsymbol{x} \cdot \boldsymbol{p}(y, \boldsymbol{w}) - c(y, \boldsymbol{w}))\, \mathrm{d}F(\boldsymbol{w}).$$

Note that while the allocation and the payments depend on what i reports, his actual value for any specific allocation does not. If $\boldsymbol{P}(y)$ denotes the vector of integrals $\int_{\boldsymbol{w}} \boldsymbol{p}(y, \boldsymbol{w})\, \mathrm{d}F(\boldsymbol{w})$, and $C(y)$ denotes the integral $\int_{\boldsymbol{w}} c(y, \boldsymbol{w})\, \mathrm{d}F(\boldsymbol{w})$, then $U(\boldsymbol{x}, y) = \boldsymbol{x} \cdot \boldsymbol{P}(y) - C(y)$.

Now we assume that the set of possible $\boldsymbol{x}$ is connected. For a fixed benchmark $\boldsymbol{x}_0$ and any other $\boldsymbol{x}$, let $\boldsymbol{t}(s)$ denote a path from $\boldsymbol{x}_0 = t(0)$ to $\boldsymbol{x} = t(r)$ for some r. Assume that $\boldsymbol{P}(\boldsymbol{t}(s))$ is differentiable with respect to s. (For example, in Vickrey's k-object auction, the only non-trivial outcome to i is winning one object. For that outcome $P(x)$ is simply the probability that at most $k-1$ other bidders have a value of at least x for the object. Let $t(s) = s$. For continuously distributed values, the probability $P(t(s))$ will be differentiable with respect to s.)

For all bidders to truthfully reveal their actual $\boldsymbol{x}$ and $\boldsymbol{w}$ to be an equilibrium to the direct revelation game, $y = \boldsymbol{x}$ must maximize the expected profit $U(\boldsymbol{x}, y)$ to i. In particular,

$$\left.\frac{\mathrm{d}U(\boldsymbol{x}, \boldsymbol{t}(s))}{\mathrm{d}s}\right|_{s=r} = \boldsymbol{x} \cdot \left.\frac{\mathrm{d}\boldsymbol{P}(\boldsymbol{t}(s))}{\mathrm{d}s}\right|_{s=r} - \left.\frac{\mathrm{d}C(\boldsymbol{t}(s))}{\mathrm{d}s}\right|_{s=r} = 0.$$

Thus, if $C(\boldsymbol{x})$ is to be continuous in $\boldsymbol{x}$ (as it is in Vickrey's examples, and must in general be for it to correspond to some equilibrium), then

$$C(\boldsymbol{x}) = \int_{s=0}^{s=r} \left.\frac{\mathrm{d}C(t(q))}{\mathrm{d}q}\right|_{q=s} \mathrm{d}s + C(\boldsymbol{t}(0)),$$

where $\boldsymbol{t}(r) = \boldsymbol{x}$. Now, using the relationship between the derivatives of $\boldsymbol{P}(\cdot)$ and $C(\cdot)$ implied by the previous first-order condition,

$$C(\boldsymbol{x}) = \int_{s=0}^{s=t} t(s) \cdot \left.\frac{\mathrm{d}\boldsymbol{P}(\boldsymbol{t}(q))}{\mathrm{d}q}\right|_{q=s} + C(\boldsymbol{t}(0)).$$

However, we should expect the integral to be independent of the path taken. Thus, the expected amount $C(x)$ paid by i depends only on $\boldsymbol{p}(\cdot)$, $F(\cdot)$, the expected amount paid at some benchmark point $\boldsymbol{x}_0 = t(0)$, and the restriction that $C(\boldsymbol{x})$ be continuous. This gives the following theorem.

Theorem. For sufficiently regular allocation functions $\mathbf{p}(y, z)$ and $F(\mathbf{w})$, if the expected amount $C(\mathbf{x})$ paid by i at equilibrium is to be continuous in $\mathbf{x}$, then $C(\mathbf{x})$ depends only on the allocation function $\mathbf{p}(y, z)$, the distribution $F(\cdot)$ of others' values, and the value of $C(\mathbf{x}_0)$ for some fixed $\mathbf{x}_0$.

In particular, not only the three pricing mechanisms examined by Vickrey, but any other mechanism that always results in each of the k bidders who value the objects the most receiving one object each – at a price of zero if all bidders have the lowest possible value for the objects – will yield the same expected revenue to the seller at equilibrium. More generally, this theorem establishes that the revenue equivalence established (with somewhat more rigor) by Myerson for a single object depends not on the number of objects so much as on each bidder knowing his value for each possible allocation to himself, and that any one bidder's values are independently distributed from any other bidders' values.

References

Engelbrecht-Wiggans, R., 1987, On the expected cost of alternative auction mechanisms for the Dairy Termination Program, Working paper, Feb. (Department of Business Administration, University of Illinois at Urbana-Champaign, Urbana, IL).

Milgrom, P.R. and R.J. Weber, 1982, A theory of auctions and competitive bidding, Econometrica 50, 1089–1122.

Myerson, R.B., 1981, Optimal auctions, Mathematics of Operations Research 6, no. 1, 58–73.

Vickrey, W., 1961, Counterspeculation, auctions, and competitive bidding, Journal of Finance 16, no. 1, 8–37.

Vickrey, W., 1962, Auction and bidding games, Recent Advances in Game Theory, 15–27.

Weber, R.J., 1983, Multiple-object auctions, in: R. Engelbrecht-Wiggans, M. Shubik and R.M. Stark, eds., Auctions, bidding, and contracting: Uses and theory (New York University Press, New York).