DATE DUE

BEBR

FACULTY WORKING
PAPER NO. 1243

Classes of Equivalent Auction Mechanisms:
The Case of Risk Neutral Bidders

Richard Engelbrecht-Wiggans

College of Commerce and Business Administration
Bureau of Economic and Business Research
University of Illinois, Urbana-Champaign

BEBR

Classes of Equivalent Auction Mechanisms:
The Case of Risk Neutral Bidders

Richard Engelbrecht-Wiggans, Associate Professor
Department of Business Administration

CLASSES OF EQUIVALENT AUCTION MECHANISMS:
THE CASE OF RISK NEUTRAL BIDDERS

Richard Engelbrecht-Wiggans

March 1986

Abstract: This paper addresses the question "what characteristics of
an auction mechanism affect the bid taker's expected payments?" To do
so, we define three variants of regret free mechanisms, mechanisms
with a corresponding direct revelation game that is individually
rational and is incentive compatible with respect to specified infor-
mation. By conditioning the incentive compatibility on factors not
explicitly included in previous studies, we establish special equiva-
lent revenue theorems for auctions with dependent information and for
multi-object auctions, as well as for a quite general family of
auctions with independent information.

Introduction

The theory of auction design attempts to understand, to describe, and to predict how the bid taker's expected revenue depends on the various factors that might affect the outcome of an auction. The outcome depends on both the bidders' bids, and on how these bids affect who wins what and who pays whom how much. The bids in turn depend on what the bidders know about the actual characteristics of the objects being auctioned, about the number of bidders, about their own and others' preferences, and how the bids will determine who wins what and who pays whom how much. In addition, the outcome of an auction depends on how the bids depend on what the bidders know. In short, the theory investigates how the outcome depends on the rules of the auction, on what the bidders know, and on how the bidders use what they know.

Note that the rules of the auction describe, but do not necessarily prescribe, who wins what and who pays whom how much as a function of the bids made. Two simple examples illustrate the distinction. On the one hand, a common form of sealed bid auction explicitly prescribes that the object (if sold at all) will be sold to the highest bidder at an amount equal to the winning bid. On the other hand, a model might describe a new car buyer as purchasing from the dealer offering the desired car at the lowest price even though there may well be no law (other than those of economics or of Nature) for the winning dealer to be so determined. Note also that this concept of "rules" allows us to define auctions quite generally to include any market mechanism in which all strategic actors know how the actions

they take will affect who wins what and who pays whom how much; this includes many mechanisms in addition to those commonly thought as auctions.

One might now define any particular auction method by describing what types of information the bidders will have, and also describing how the bids made will affect the auction's outcome. For example, the (symmetric) independent private values model assumes that bidders' values for the object being auctioned are independent identically distributed draws from a known distribution, that bidders are risk neutral, that each bidder knows his own value for the object precisely (but knows nothing about others' values except what can be inferred from the known distribution of the values), and that this entire description is common knowledge to all the bidders. The auction rules might specify that the highest bidder wins the object and pays the amount of the winning bid--the common first price sealed bid auction. Alternatively, the rules might specify that the highest bidder wins the object and pays an amount equal to the second highest bid--this second price sealed bid auction approximates the progressive oral auction in which the bidder willing to pay the most wins at a price just barely above the price at which the bidder willing to pay the second highest drops out of the bidding. The noisy progressive auction in which the auctioneer starts with a very low asking price and continuously raises it until all but one bidder have publicly indicated that (and at what price) they have dropped out of the bidding, and then awards the object to the remaining bidder at the current asking price provides yet another approximation to the common progressive oral auction.

To complete an auction model, one must describe--again describe rather than prescribe--how each bidder's bid depends on that bidder's information. For example, there might be a dominant strategy--a function translating a bidder's information into a bid with the property that regardless of how others bid, this bidder can do no better than follow the specified strategy. In the absence of dominant strategies, the theory typically resorts to some form of equilibrium, most commonly a Nash equilibrium--a set of strategies, one for each bidder, forms a Nash equilibrium if no one bidder can do better than by following the strategy specified for him so long as all other bidders follow the strategies specified for them. (Note that a set of dominant strategies, one for each bidder, always satisfies the conditions for a Nash equilibrium, and in addition, most other forms of equilibrium.)

Occasionally, auction models yield explicit characterizations of their equilibria, and in turn of the auctioneer's expected revenue, thereby allowing the direct comparison of one auction method to another. For example, Vickrey (1961) explicitly characterized the expected price paid by the winner in the first price and second price sealed bid independent private values auction models at their respective equilibria. (In the independent private values model, the second price auction has a dominant strategy Nash equilibrium.) The two models yield exactly the same expected revenue.

More recently, Milgrom and Weber (1982) relaxed both the independent values and known values assumptions of the independent private values model to obtain a more general affiliated information model, a

model that then yielded the following ranking (in order of decreasing expected revenue to the auctioneer at their corresponding equilibria): 1) the noisy progressive mechanism, 2) the second price sealed bid mechanism, and 3) the first price sealed bid mechanism. Roughly speaking, the more (affiliated) information that might affect the price paid by the winner--regardless of whether such information came from non-winning bidders during the auction itself, was revealed by the auctioneer before the auction started, or, as in the second price model, only affects the price after all bids have been submitted--the higher the auctioneer's expected revenue as predicted by the model.

Unfortunately, very few auction models have yielded explicit characterizations of their equilibria. Fortunately, however, an alternative approach to studying auctions suppresses the details of how bidders bid (and also of how the bids affect the outcome) by instead directly describing--not prescribing--how the bidders' information affects the expected outcome. In particular, each possible equilibrium to each possible auction model has a corresponding direct-revelation auction model--an auction model with an equilibrium at which each bidder truthfully reveals his actual information, and the outcome has the same relation to bidders' information as at the specified equilibrium in the original model. (To construct the direct revelation auction, simply require the auctioneer to calculate what bids the bidders would have made at the specified equilibrium using the information the bidders report to the auctioneer, and to then implement whatever outcome would have resulted from these bids in the original auction. Now, when asked to "bid" by telling the auctioneer

his actual information, no bidder can do better than responding truth-
fully no matter how other bidders respond.) Therefore, any relation-
ship between bidders' information and outcome that arises from some
equilibrium in some auction model also rises at a truth revealing
equilibrium of some direct-revelation auction. Since direct-
revelation auctions are themselves auctions, the class of relation-
ships possible between bidders' information and the outcome of an
auction at equilibrium coincides with the class of such relationships
for truth revealing equilibria of direct-revelation auctions. This
allows us to study the possible outcomes of a family of auctions by
focusing on the corresponding direct-revelation auctions—a class of
auctions for which the truth revealing strategy provides an explicit
characterization of an equilibrium.

Note that any failure to explictly characterize an equilibrium for
an auction translates into an inability to define precisely which
direct-revelation auction corresponds to the original auction. How-
ever, we may still study entire families of direct-revelation auctions.
Any auction that, for example, maximizes the auctioneer's expected
revenue over a family of direct-revelation auctions also maximizes the
auctioneer's expected revenue over the larger, corresponding family of
auctions in general. Thus, this approach allows one to construct a
practical auction with certain properties by first identifying a
direct-revelation auction with the desired properties, and then
finding a practical auction that may be expected to have the same
relationship between bidders' information and outcome as the chosen
direct-revelation auction.

Myerson (1981) first used this direct-revelation approach to study independent private values auction models. For any fixed distribution function and fixed number of risk neutral bidders, he considered all relations of outcome to bidders' actual values satisfying the following two conditions: 1) incentive compatibility--no bidder could do strictly better by lying about his value for the object than by telling the truth, and 2) individual rationality--each bidder, for each possible value of the object to him, could report a value with the ultimate effect of exactly zero profit (in other words, each bidder must always have an option in effect equivalent to not bidding). The result: for any fixed number of risk neutral bidders and fixed distribution of the bidder's values, the allocation rule--the part of the rules that describes who wins the object as a function of all the bidders' values--uniquely determines each bidder's expected payment as a function of all the bidder's values. In particular, at equilibrium, all independent private values auctions that always award the object to the bidder valuing it most highly generate precisely the same expected revenue for the auctioneer. (In addition, Myerson determined that the allocation rule that maximizes the auctioneer's expected revenue always awards the object to the highest valuer so long as that bidder's value exceeds an appropriately chosen critical value; otherwise the auctioneer destroys the object. This allocation rule gives a positive probability of the auctioneer destroying a valuable object. However, at least to some extent, this ex-post inefficiency of the auctioneer's expected revenue maximizing mechanism arises from the model's assuming a fixed number of bidders; Engelbrecht-Wiggans (1986)

provides an example in which increasing the probability of not award-
ing the object decreases the expected number of bidders with the
result that the auctioneer maximizes his expected revenue by always
awarding the object to the highest valuer).

While focusing on the independent values case, Myerson also pro-
vided an example to illustrate that with dependent values the auc-
tioneer may be able to expect a revenue equal to the full value of the
object to the bidder who values the object the most. Although the
example mechanism always awards the object to the highest valuer at a
price equal to his declared value, the mechanism differs from the
common first price sealed bid auction in that it effectively forces
bidders, as part of their bids, to bet on how others will bid. If all
bidders reveal their actual values, then each expects to break even on
these side bets, and the auctioneer extracts the full value of the
object; were a bidder to unilaterally deviate from revealing his
actual value, his expected loss from the bet would exceed any expected
gain possible from receiving the object at a more favorable price.
Unfortunately the success of such mechanisms depends entirely on being
very carefully tailored to the appropriate, specific joint distribu-
tion of bidders' values; this limits the practical applications of
full value extracting mechanisms.

Engelbrecht-Wiggans (1985a,b) defines a model of auctions with
possibly dependent information that differs from previous models in
focusing on appropriately chosen sub-family of the direct-revelation
auctions, thereby excluding the full value extracting mechanisms. In
particular, the model parameterizes how the price paid by one bidder

might depend on the information of other bidders. Specifically, let r_i be a function of the information held by bidders 1, 2, ..., i-1, i+1, ..., such that the amount eventually paid by bidder i depends on others' information only through this statistic r_i. Now, consider only regret free mechanisms, that is, auctions with corresponding direct-revelation auctions satisfying the following modified incentive compatibility condition: no bidder could do strictly better by lying about his value for the object even if he already knew (at the time of bidding) the specific statistic r_i through which his eventual payment would depend on others' information. This family of regret free mechanisms includes, for example, the common progressive oral auction, because in such an auction, at the instant the auctioneer says "sold," each bidder knows exactly how much he will pay and, yet, until that instant, each bidder still had the option to bid higher (that is, to indicate willingness to pay a higher price than he previously indicated). More generally, this family includes the many competitive contracting mechanisms in which, before the contract is officially let, each bidder knows whether or not he is the finalist, the finalist knows the precise terms of the contract to be let (here, the bidding determined the actual terms of the contract, rather than simply a one-shot price, although in some contracts the only terms affected by the bidding is the price), and yet no bidder advances a new proposal. On the other hand, the family excludes first price sealed bid auctions as well as the full value extracting auctions described above.

For regret free mechanisms with risk neutral bidders, the information structure (specifying the number of bidders, who has what type

of information, the joint distribution of the actual information, the
form of the r_i functions, and how each bidder's value of the object
relates to all the bidders' information) together with the allocation
rule uniquely determines how much more each bidder i pays (as a func-
tion of his information and the statistic r_i) than the minimum (taken
over all possible misrepresentations of his own information) he could
have paid in that specific situation. For example, if each bidder has
the option of withdrawing entirely from an auction of a valuable
object (by valuable, we mean that the auctioneer would never consider
paying a bidder a positive amount to·take the object), then the infor-
mation structure and allocation rule together uniquely determine the
amount paid by each bidder in any regret free mechanism. Engelbrecht-
Wiggans then proceeds to show that the amount paid will be independent
of the form of the r_i functions if the bidders have independent infor-
mation, a result closely related to that of Myerson. Furthermore, for
auctions with dependent information, the more refined the r_i functions—
that is, the more detail they capture about others' information--the
higher the price; as a corollary, this yields Milgrom and Weber's
result that at equilibrium, noisy progressive auctions generate at
least as much revenue for the auctioneer as second price sealed bid
auctions.

The current paper focuses on one part of the auction design
problem. In particular, it investigates the question "what character-
istics of an auction actually affect the auction's outcome?" To do
so, it defines three variations of the above mentioned concept of
regret free mechanisms, two of which include the first price sealed

bid auctions excluded by the previous definition. Applying these definitions to direct-revelation auctions yields three equivalence theorems, theorems stating that all regret free mechanisms that coincide on a specified set of characteristics will result in the same expected payments. In effect, this paper extends Myerson's equivalence theorem to multi-object auctions, and then obtains corresponding results for two appropriately restricted classes of auctions with dependent information.

Notation and Definitions

Let the subscript i (i=1, 2, ..., n) denote one of the n risk neutral bidders, while the subscript i=0 denotes the bid taker. Before the bidding starts, each individual i privately observes the outcome s_i of a random real valued vector S_i. The signals have a known joint or conditional--the context will indicate which-- cumulative probability distribution $F(\cdot)$. For convenience, s_{-i} denotes the vector $(s_0, s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$. Then, let x_i be the signal i claims to have observed.

Now, to characterize auction mechanisms, let $r_i(s_{-i})$ be some statistic of s_{-i}. If i claims to have observed x_i, and the others observed s_{-i}, individual i obtains the award (or allocation) $a_i(x_i, s_{-i})$; in the simplest case of a single object auction, a_i equals the object if i's claimed x_i would have him value the object at least as much as anyone else, and a_i is the empty set otherwise. The (expected) payment $p_i(x_i, \alpha_i, \rho_i)$ made by i when $a_i(x_i, s_{-i}) = \alpha_i$ and $r_i(s_{-i}) = \rho_i$ may depend on s_{-i} only through the $a_i(\cdot)$ and $r_i(\cdot)$

functions; for example, in a model with real value signals, $r_i(s_{-i})$ equal to a constant would correspond to a first price mechanism, $r_i(s_{-i}) = \max_{j \neq i} s_j$ would naturally suggest a second price mechanism, and $r_i(s_{-i}) = s_{-i}$ would naturally suggest a noisy progressive mechanism. The award α_i has an expected value (averaged over everything but s_i and s_{-i}) to i of $v_i(s_i, s_{-i}, \alpha_i)$.

For any specific n, and functions F, r_i, a_i, v_i, and p_i define $u_i(s_i, x_i, \alpha_i, \rho_i)$ as the expected utility

$$\int_{s_{-i}} [v_i(s_i, s_{-i}, a_i(x_i, s_{-i})) - p(x_i, a_i(x_i, s_{-i}), \rho_i)] dF(s_{-i} | S_i = s_i, r_i(S_{-i}) = \rho_i, a(s_i, S_{-i}) = \alpha_i);$$

define $U_i(s_i, x_i, \rho_i) = \sum_{\alpha_i} u_i(s_i, x_i, \alpha_i, \rho_i) \int_{s_{-i}:r_i(S_{-i})=\rho_i} dF(s_{-i} | S_i = s_i, a_i(s_i, S_{-i}) = \alpha_i)$

$$= V_i(s_i, x_i, \rho_i) - P_i(x_i, \rho_i)$$

where $V_i(s_i, x_i, \rho_i) = \int_{s_{-i}} v_i(s_i, s_{-i}, a_i(x_i, s_i)) dF(s_{-i} | r_i(S_{-i}) = \rho_i)$,

and $P_i(x_i, \rho_i) = \int_{s_{-i}} p_i(x_i, a_i(x_i, s_{-i}), \rho_i) dF(s_{-i} | r_i(S_{-i}) = \rho_i)$;

and define $U_i^*(s_i, x_i)$ as the expected utility

$$\int_{\rho_i} \sum_{\alpha_i} u_i(s_i, x_i, \alpha_i, \rho_i) \int_{s_{-i}:r_i(S_{-i})=\rho_i \& a_i(s_i, S_{-i})=\alpha_i} dF(s_{-i} | S_i = s_i)$$

$$= \int_{\rho_i} \sum_{\alpha_i} \int_{s_{-i}:r_i(S_{-i})=\rho_i \& a_i(s_i, S_{-i})=\alpha_i} [v_i(s_i, s_{-i}, a_i(x_i, s_{-i})) - p_i(x_i, a_i(x_i, s_{-i}), \rho_i)] dF(s_{-i} | S_i = s_i).$$

Note that if S_{-i} is statistically independent of S_i, then $U_i^*(s_i, x_i)$ may be rewritten as $V_i^*(s_i, x_i) - P_i^*(x_i)$, where $V_i^*(s_i, x_i) =$

$\int_{s_{-i}} v_i(s_i, s_{-i}, a_i(x_i, s_{-i})) dF(s_{-i})$ and $P_i^*(x_i) =$

$\int_{s_{-i}} p_i(x_i, a_i(x_i, s_i), r_i(s_{-i})) dF(s_{-i})$. Assume that for each s_i,

$u_i(s_i, x_i, \alpha_i, \rho_i)$ is differentiable with respect to x_i at x_i equal to

s_i for almost all α_i and ρ_i given $S_i = s_i$; assume that for each s_i,

$U_i(s_i, x_i, \rho_i)$ is differentiable with respect to x_i at x_i equal to

s_i for almost all ρ_i given $S_i = s_i$; and assume that for each s_i,

$U_i^*(s_i, x_i)$ is differentiable with respect to x_i at x_i equal to s_i.

Roughly speaking, the main import of of these assumptions is to

require zero probability of tied bids; since we allow s_i to have

multiple dimensions, any auction with a positive probability of ties

may be modeled by another auction with zero probability of ties (but

with higher dimensional information) using the techniques presented by

Engelbrecht-Wiggans, Milgrom and Weber (1983).

Finally, define a mechanism to be <u>weakly regret free</u> if

$\frac{d}{dx_i} U_i^*(s_i, x_i)\Big|_{x_i = s_i}$ is zero for all s_i, and $\min_{s_i} P_i^*(s_i) = 0$; <u>strongly</u>

<u>regret free</u> if for each s_i, $\frac{d}{dx_i} U_i(s_i, x_i, \rho_i)\Big|_{x_i = s_i}$ is zero for almost

all ρ_i given $S_i = s_i$, and for each ρ_i, $\min_{s_i} P_i(s_i, \rho_i) = 0$; and <u>totally</u>

<u>regret free</u> if for each s_i, $\frac{d}{dx_i} u_i(s_i, x_i, \alpha_i, \rho_i)$ is zero for almost all

α_i and ρ_i given $S_i = s_i$, and for each α_i and ρ_i, $\min_{x_i} p_i(s_i, \alpha_i, \rho_i) = 0$.

Note that to be weakly regret free corresponds to being incentive

compatible in the (now) traditional sense; any Nash equilibrium has a

corresponding weakly regret free mechanism. To be strongly regret

free, the mechanism must remain incentive compatible when bidder i

knows all the information ρ_i through which his payment might depend on

s_{-i} once the award has been determined. At equilibrium, most common auction mechanisms are strongly regret free; typical full revenue extracting mechanisms are not, thereby allowing the possibility of a revenue equivalence theorem that could not hold if full revenue extracting mechanisms were included. Finally, to be totally regret free requires incentive compatibility even after bidder i knows his award as well as the statistic ρ_i. This excludes first price auctions, but includes any mechanism in which the bid taker negotiates a tentative outcome (specifying awards and payments) known to all bidders, and this tentative outcome becomes the final outcome if and only if no bidder desires to negotiate further; when appropriately viewed, oral auctions serve as an example of such a negotiation mechanism.

Results

<u>Theorem 1</u>: For weakly regret free mechanisms with S_{-i} statistically independent of S_i, n and the functions F, a, and v uniquely determine $P_i^*(s_i)$.

<u>Proof</u>: Applying the definition of weakly regret free gives

$$\frac{d}{dx_i} V_i^*(s_i, x_i)\bigg|_{x_i=s_i} = \frac{d}{dx_i} P_i^*(x_i)\bigg|_{x_i=s_i} \quad \text{for all } s_i \text{ and } \min_{s_i} P_i^*(s_i) = 0.$$

The differential condition determines $P_i^*(x_i)$ except for an additive constant; $\min\limits_{s_i} P_i^*(s_i) = 0$ sets the constant.

<u>Theorem 2</u>: For strongly regret free mechanisms in which $p_i(x_i, \alpha_i, \rho_i)$ is exogenously specified for all values of α_i except one (say α_i^*), n and the functions F, a, r, and v uniquely determine $p_i(s_i, \alpha_i^*, \rho_i)$.

Proof: Since $p_i(x_i, \alpha_i, \rho_i)$ varies with x_i only for α_i equal to α_i*, the definition of strongly regret free gives $\dfrac{d}{dx_i} V_i(s_i, x_i, \rho_i)\Big|_{x_i = s_i} =$ $\dfrac{d}{dx_i} P_i(x_i, \rho_i)\Big|_{x_i = s_i}$ and $\min_{s_i} P_i(s_i, \rho_i) = 0$ for each ρ_i. The differential condition determines $P_i(s_i, \rho_i)$ except for an additive constant; $\min_{s_i} P_i(s_i, \rho_i) = 0$ sets the constant. Now, since we know $p_i(s_i, \alpha_i, \rho_i)$ for all α_i except for α_i equal to α_i^*, $P_i(s_i, \rho_i)$ determines $p_i(s_i, \alpha_i^*, \rho_i)$.

Theorem 3: For totally regret free mechanisms, n and the functions F, a, r, and v uniquely determine $p_i(x_i, \alpha_i, \rho_i)$ for each α_i and ρ_i.

Proof: The definition of totally regret free gives that

$$\frac{d}{dx_i} p_i(x_i, \alpha_i(x_i, s_{-i}), \rho_i)\Big|_{x_i = s_i}$$

$$= \frac{d}{dx_i} \int_{s_{-i}} v_i(s_i, s_{-i}, \alpha_i(x_i, s_{-i})) dF(s_{-i} | S_i = s_i, r_i(S_{-i}) =$$

$\rho_i, a_i(s_i, S_{-i}) = \alpha_i)$ and $\min_{s_i} p_i(s_i, \alpha_i, \rho_i) = 0$ for each α_i and ρ_i. The differential condition defines $p_i(x_i, \alpha_i, \rho_i)$ except for an additive constant; $\min_{s_i} p_i(s_i, \alpha_i, \rho_i) = 0$ sets the constant.

The first theorem shows that for independent signals, the expected payments do not depend on how the payments might depend on others' information. In other words, for the case of independent signals, at equilibrium, all auctions--first price, second price, noisy progressive, or what have you--generate the same expected revenue for the bid taker. This result extends the equivalence result of Myerson in two ways: first, our model allows bidders to have multi-dimensional private information, and second, our model covers multi-object auctions.

The second theorem shows that for a typical single object auction
in which loosing bidders pay nothing, the bid taker's expected revenue
hinges on how much of the other bidders' information gets factored into
the winner's price; given the existing results ordering the expected
revenue from mechanisms in the case of dependent information, we could
not have reasonably hoped for any more sweeping equivalence theorem.
The theorem also illuminates a critical difference between the first
price and the full value extracting mechanisms, two mechanisms that
typically result in different expected revenue levels for the bid
taker. While both mechanisms leave the winner totally ignorant of
others' information (other than what can be inferred through the known
joint distribution of the signals) up until the award and price are
announced, the first price auction is strongly regret free, while the
typical full revenue extracting mechanism is not. Therefore, for
dependent values, not only the amount of other bidders' information
reflected in the winner's expected payment affect the bid taker's
expected revenue, but even the remaining equivalence can hold only for
strongly regret free mechanisms.

Going to the third theorem involved trading off the diversity of
mechanisms covered versus the generality of mechanisms covered. In
particular, the theorem only applies to totally regret free mechanisms,
a smaller family than the strongly or weakly regret free mechanisms.
On the other hand, this theorem does apply to multi-object auctions in
which each bidder may have more than one piece of private information,
and in which one bidder's information may be statistically dependent
on the other bidders' information.

Summary

This paper defined three variations of the regret free mechanism concept previously defined by Engelbrecht-Wiggans. The different definitions include a varying diversity of auction mechanisms; all three definitions, however, include a greater diversity of mechanisms and models than previously considered. Finally, for each definition of regret free, a corresponding theorem establishes under what conditions what characteristics of the auction mechanism determine the expected payments.

References

Engelbrecht-Wiggans, R., "Optimal Auctions: The Role of Information," University of Illinois, Department of Business Administration working paper BEBR #1199, February 1985, revised October 1985(a).

_____, "Optimal Competitive Contracting," University of Illinois, Department of Business Administration working paper BEBR #1200, October 1985(b).

_____, "On Optimal Reservation Prices in Auction," University of Illinois, Department of Business Administration working paper, March 1986.

Engelbrecht-Wiggans, R., P. R. Milgrom, and R. J. Weber, "Competitive Bidding and Proprietary Information," Journal of Mathematical Economics, Vol. 11, 1983, pp. 161-169.

Milgrom, P. R. and R. J. Weber, "A Theory of Auctions and Competitive Bidding," Econometrica, Vol. 50, 1982, 1089-1122.

Myerson, R. B., "Optimal Auction Design," Mathematics of Operations Research, Vol. 6, No. 1, February 1981, 58-73.

Vickrey, W., "Counterspeculation, Auctions and Competitive Sealed Tenders," Journal of Finance, Vol. 41, No. 1, 1961, 8-37.