Auction Theory

Jonathan Levin

October 2004

Our next topic is auctions. Our objective will be to cover a few of the main ideas and highlights. Auction theory can be approached from different angles — from the perspective of game theory (auctions are bayesian games of incomplete information), contract or mechanism design theory (auctions are allocation mechanisms), market microstructure (auctions are models of price formation), as well as in the context of different applications (procurement, patent licensing, public finance, etc.). We're going to take a relatively game-theoretic approach, but some of this richness should be evident.

1 The Independent Private Value (IPV) Model

1.1 A Model

The basic auction environment consists of:

- Bidders $i = 1, ..., n$
- One object to be sold
- Bidder i observes a "signal" $S_i \sim F(\cdot)$, with typical realization $s_i \in [\underline{s}, \overline{s}]$, and assume F is continuous.
- Bidders' signals $S_1, ..., S_n$ are independent.
- Bidder i's value $v_i(s_i) = s_i$.

Given this basic set-up, specifying a set of auction rules will give rise to a game between the bidders. Before going on, observe two features of the model that turn out to be important. First, bidder i's information (her signal) is *independent* of bidder j's information. Second, bidder i's value is independent of bidder j's information — so bidder j's information is *private* in the sense that it doesn't affect anyone else's valuation.

1.2 Vickrey (Second-Price) Auction

In a Vickrey, or second price, auction, bidders are asked to submit sealed bids $b_1, ..., b_n$. The bidder who submits the highest bid is awarded the object, and pays the amount of the second highest bid.

Proposition 1 *In a second price auction, it is a weakly dominant strategy to bid one's value, $b_i(s_i) = s_i$.*

Proof. Suppose i's value is s_i, and she considers bidding $b_i > s_i$. Let $\hat{b}$ denote the highest bid of the other bidders $j \neq i$ (from i's perspective this is a random variable). There are three possible outcomes from i's perspective: (i) $\hat{b} > b_i, s_i$; (ii) $b_i > \hat{b} > s_i$; or (iii) $b_i, s_i > \hat{b}$. In the event of the first or third outcome, i would have done equally well to bid s_i rather than $b_i > s_i$. In (i) she won't win regardless, and in (ii) she will win, and will pay $\hat{b}$ regardless. However, in case (ii), i will win and pay more than her value if she bids $\hat{b}$, something that won't happen if she bids s_i. Thus, i does better to bid s_i than $b_i > s_i$. A similar argument shows that i also does better to bid s_i than to bid $b_i < s_i$. *Q.E.D.*

Since each bidder will bid their value, the seller's revenue (the amount paid in equilibrium) will be equal to the second highest value. Let $S^{i:n}$ denote the ith highest of n draws from distribution F (so $S^{i:n}$ is a random variable with typical realization $s^{i:n}$). Then the seller's expected revenue is $\mathbb{E}\left[S^{2:n}\right]$.

The truthful equilibrium described in Proposition 1 is the unique symmetric Bayesian Nash equilibrium of the second price auction. There are also asymmetric equilibria that involve players using weakly dominated strategies. One such equilibrium is for some player i to bid $b_i(s_i) = \overline{v}$ and all the other players to bid $b_j(s_j) = 0$.

While Vickrey auctions are not used very often in practice, open ascending (or English) auctions are used frequently. One way to model such auctions is to assume that the price rises continuously from zero and players each can push a button to "drop out" of the bidding. In an independent private values setting, the Nash equilibria of the English auction are the same as the Nash equilibria of the Vickrey auction. In particular, the unique symmetric equilibrium (or unique sequential equilibrium) of the English auction has each bidder drop out when the price reaches his value. In equilibrium, the auction ends when the bidder with the second-highest value drops out, so the winner pays an amount equal to the second highest value.

1.3 Sealed Bid (First-Price) Auction

In a sealed bid, or first price, auction, bidders submit sealed bids $b_1, ..., b_n$. The bidders who submits the highest bid is awarded the object, and pays his bid.

Under these rules, it should be clear that bidders will not want to bid their true values. By doing so, they would ensure a zero profit. By bidding somewhat below their values, they can potentially make a profit some of the time. We now consider two approaches to solving for symmetric equlibrium bidding strategies.

A. The "First Order Conditions" Approach

We will look for an equilibrium where each bidder uses a bid strategy that is a strictly increasing, continuous, and differentiable function of his value.[1] To do this, suppose that bidders $j \neq i$ use identical bidding strategies $b_j = b(s_j)$ with these properties and consider the problem facing bidder i.

Bidder i's expected payoff, as a function of his bid b_i and signal s_i is:

$$U(b_i, s_i) = (s_i - b_i) \cdot \Pr\left[b_j = b(S_j) \leq b_i, \ \forall j \neq i\right]$$

Thus, bidder i chooses b to solve:

$$\max_{b_i} (s_i - b_i) F^{n-1}\left(b^{-1}(b_i)\right).$$

The first order condition is:

$$(s_i - b_i)(n-1)F^{n-2}\left(b^{-1}(b_i)\right) f\left(b^{-1}(b_i)\right) \frac{1}{b'(b^{-1}(b_i))} - F^{n-1}\left(b^{-1}(b_i)\right) = 0$$

At a symmetric equilibrium, $b_i = b(s_i)$, so the first order condition reduces to a differential equation (here I'll drop the i subscript):

$$b'(s) = (s - b(s))(n-1)\frac{f(s)}{F(s)}.$$

This can be solved, using the boundary condition that $b(\underline{s}) = \underline{s}$, to obtain:

$$b(s) = s - \frac{\int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s}}{F^{n-1}(s)}.$$

[1] In fact, it is possible to prove that in any symmetric equilibrium each bidder *must* use a continuous and strictly increasing strategy. To prove this, one shows that in equilbrium there cannot be a "gap" in the range of bids offered in equilibrium (because then it would be sub-optimal to offer the bid just above the gap) and there cannot be an "atom" in the equilibrium distribution of bids (because then no bidder would make an offer just below the atom, leading to a gap). I'll skip the details though.

It is easy to check that $b(s)$ is increasing and differentiable. So any symmetric equilibrium with these properties must involve bidders using the strategy $b(s)$.

B. The "Envelope Theorem" Approach

A closely related, and often convenient, approach to identify necessary conditions for a symmetric equilibrium is to exploit the envelope theorem.

To this end, suppose $b(s)$ is a symmetric equilibrium in increasing differentiable strategies. Then i's equilibrium payoff given signal s_i is

$$U(s_i) = (s_i - b(s_i))\, F^{n-1}(s_i). \tag{1}$$

Alternatively, because i is playing a best-response in equilibrium:

$$U(s_i) = \max_{b_i} (s_i - b_i)\, F^{n-1}(b^{-1}(b_i)).$$

Applying the envelope theorem (Milgrom and Segal, 2002), we have:

$$\frac{d}{ds} U(s)\Big|_{s=s_i} = F^{n-1}(b^{-1}(b(s_i)) = F^{n-1}(s_i)$$

and also,

$$U(s_i) = U(\underline{s}) + \int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s}. \tag{2}$$

As $b(s)$ is increasing, a bidder with signal $\underline{s}$ will never win the auction — therefore, $U(\underline{s}) = 0$.

Combining (1) and (2), we solve for the equilibrium strategy (again dropping the i subscript):

$$b(s) = s - \frac{\int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s}}{F^{n-1}(s)}.$$

Again, we have showed necessary conditions for an equilibrium (i.e. any increasing differentiable symmetric equilibrium must involve the strategy $b(s)$). To check sufficiency (that $b(s)$ actually is an equilibrium), we can exploit the fact that $b(s)$ is increasing and satisfies the envelope formula to show that it must be a selection from i's best response given the other bidder's use the strategy $b(s)$. (For details, see Milgrom 2004, Theorems 4.2 amd 4.6).

Remark 1 *In most auction models, both the first order conditions and the envelope approach can be used to characterize an equilibrium. The trick is to figure out which is more convenient.*

What is the revenue from the first price auction? It is the expected winning bid, or the expected bid of the bidder with the highest signal, $\mathbb{E}\left[b(S^{1:n})\right]$. To sharpen this, define $G(s) = F^{n-1}(s)$. Then G is the probability that if you take $n-1$ draws from F, all will be below s (i.e. it is the cdf of $S^{1:n-1}$). Then,

$$b(s) = s - \frac{\int_{\underline{s}}^{s} F^{n-1}(\tilde{s})d\tilde{s}}{F^{n-1}(s)} = \frac{1}{F^{n-1}(s)} \int_{\underline{s}}^{s} \tilde{s} dF^{n-1}(\tilde{s}) = \mathbb{E}\left[S^{1:n-1} | S^{1:n-1} \leq s\right].$$

That is, if a bidder has signal s, he sets his bid equal to the expectation of the highest of the other $n-1$ values, conditional on all those values being less than his own.

Using this fact, the expected revenue is:

$$\mathbb{E}\left[b\left(S^{1:n}\right)\right] = \mathbb{E}\left[S^{1:n-1} | S^{1:n-1} \leq S^{1:n}\right] = \mathbb{E}\left[S^{2:n}\right],$$

equal to the expectation of the second highest value. We have shown:

Proposition 2 *The first and second price auction yield the same revenue in expectation.*

1.4 Revenue Equivalence

The result above is a special case of the celebrated "revenue equivalence theorem" due to Vickrey (1961), Myerson (1981), Riley and Samuelson (1981) and Harris and Raviv (1981).

Theorem 1 *(Revenue Equivalence) Suppose n bidders have values $s_1, ..., s_n$ identically and independently distributed with cdf $F(\cdot)$. Then all auction mechanisms that (i) always award the object to the bidder with highest value in equilibrium, and (ii) give a bidder with valuation $\underline{s}$ zero profits, generates the same revenue in expectation.*

Proof. We consider the general class of auctions where bidders submit bids $b_1, ..., b_n$. An auction rule specifies for all i,

$$\begin{aligned} x_i &: B_1 \times ... \times B_n \to [0,1] \\ t_i &: B_1 \times ... \times B_n \to \mathbb{R}, \end{aligned}$$

where $x_i(\cdot)$ gives the probability i will get the object and $t_i(\cdot)$ gives i's required payment as a function of the bids $(b_1, ..., b_n)$.[2]

Given the auction rule, bidder i's expected payoff as a function of his signal and bid is:

$$U_i(s_i, b_i) = s_i \mathbb{E}_{b_{-i}} [x_i(b_i, b_{-i})] - \mathbb{E}_{b_{-i}} [t_i(b_i, b_{-i})] .$$

Let $b_i(\cdot), b_{-i}(\cdot)$ denote an equilibrium of the auction game. Bidder i's *equilibrium* payoff is:

$$U_i(s_i) = U_i(s_i, b(s_i)) = s_i F^{n-1}(s_i) - \mathbb{E}_{s_{-i}} [t_i(b_i(s_i), b_{-i}(s_{-i})] ,$$

where we use *(i)* to write $\mathbb{E}_{s_{-i}} [x_i(b(s_i), b(s_{-i}))] = F^{n-1}(s_i)$.

Using the fact that $b(s_i)$ must maximize i's payoff given s_i and opponent strategies $b_{-i}(\cdot)$, the envelope theorem implies that:

$$\frac{d}{ds} U_i(s) \Big|_{s=s_i} = \mathbb{E}_{b_{-i}} [x_i(b_i(s_i), b_{-i}(s_{-i}))] = F^{n-1}(s_i),$$

and also

$$U_i(s_i) = U_i(\underline{s}) + \int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s} = \int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s},$$

where we use *(ii)* to write $U_i(\underline{s}) = 0$.

Combining our expressions for $U_i(s_i)$, we get bidder i's expected payment given his signal:

$$\mathbb{E}_{s_{-i}} [t_i(b_i, b_{-i})] = s_i F^{n-1}(s_i) - \int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s} = \int_{\underline{s}}^{s_i} \tilde{s} dF^{n-1}(\tilde{s}),$$

where the last equality is from integration by parts. Since $x_i(\cdot)$ does not enter into this expression, bidder i's expected equilibrium payment given his signal is the *same* under all auction rules that satisfy *(i)* and *(ii)*. Indeed, i's expected payment given s_i is equal to:

$$\mathbb{E}\left[S^{1:n-1} \mid S^{1:n-1} < s_i\right] = \mathbb{E}\left[S^{2:n} \mid S^{1:n} = s_i\right] .$$

So the seller's revenue is:

$$\mathbb{E}[\text{Revenue}] = n \mathbb{E}_{s_i} [i\text{'s expected payment} \mid s_i] = \mathbb{E}\left[S^{2:n}\right] ,$$

[2]So in a first price auction, $x_1(b_1, ..., b_n)$ equals 1 if b_1 is the highest bid, and otherwise zero. Meanwhile $t_1(b_1, ..., b_n)$ equals zero unless b_1 is highest, in which case $t_1 = b_1$. In a second price auction, $x_1(\cdot)$ is the same, and $t_1(\cdot)$ is zero unless b_1 is highest, in which case t_1 equals the highest of $(b_2, ..., b_n)$.

a constant. *Q.E.D.*

The revenue equivalence theorem has many applications. One useful trick is that it allows us to solve for the equilibrium of different auctions, so long as we know that the auction will satisfiy *(i)* and *(ii)*. Here's an example.

Application: The All-Pay Auction. Consider the same set-up — bidders $1, .., n$, with values $s_1, ..., s_n$, identically and independently distributed with cdf F — and consider the following rules. Bidders submit bids $b_1, ..., b_n$ and the bidder who submits the highest bid gets the object. However, bidders must pay their bid regardless of whether they win the auction. (These rules might seem a little strange — the all-pay auction is commonly used as a model of lobbying or political influence).

Suppose this auction has a symmetric equilibrium with an increasing strategy $b^A(s)$ used by all players. Then, bidder i's expected payoff given value s_i will be (if everyone plays the equilibrium strategies):

$$U(s_i) = s_i F^{n-1}(s_i) - b^A(s_i) = \int_{\underline{s}}^{s_i} F^{n-1}(\tilde{s}) d\tilde{s}$$

So

$$b^A(s) = sF^{n-1}(s) - \int_{\underline{s}}^{s} F^{n-1}(\tilde{s}) d\tilde{s}$$

In addition to the all-pay auction, many other auction rules also satisfy the revenue equivalence assumptions when bidder values are independently and identically distributed. Two examples are:

1. The English (oral ascending) auction. All bidders start in the auction with a price of zero. The price rises continuously, and bidders may drop out at any point in time. Once they drop out, they cannot re-enter. The auction ends when only one bidder is left, and this bidder pays the price at which the second-to-last bidder dropped out.

2. The Dutch (descending price) auction. The price starts at a very high level and drops continuously. At any point in time, a bidder can stop the auction, and pay the current price. Then the auction ends.

2 Common Value Auctions

We would now like to generalize the model to allow for the possibility that *(i)* learning bidder j's information could cause bidder i to re-assess his estimate of how much he values the object, and *(ii)* the information of i and j is not independent (when j's estimate is high, i's is also likely to be high).

These features are natural to incorporate in many situations. For instance, consider an auction for a natural resource like a tract of timber. In such a setting, bidders are likely to have different costs of harvesting or processing the timber. These costs may be independent across bidders and private, much like in the above model. But at the same time, bidders are likely to be unsure exactly how much merchanteable timber is on the tract, and use some sort of statistical sampling to estimate the quantity. Because these estimates will be based on limited sampling, they will be imperfect — so if i learned that j had sampled a different area and got a low estimate, she would likely revise her opinion of the tract's value. In addition, if the areas sampled overlap, the estimates are unlikely to be independent.

2.1 A General Model

- Bidders $i = 1, ..., n$
- Signals $S_1, ..., S_n$ with joint density $f(\cdot)$
- Signals are *(i)* exchangeable, and *(ii)* affiliated.
  - Signals are *exchangeable* if s' is a permutation of $s \Rightarrow f(s) = f(s')$.
  - Signals are *affiliated* if $f(s \wedge s')f(s \vee s') \geq f(s')f(s)$ (i.e. $s_j|s_i$ has monotone likelihood ratio property).
- Value to bidder i is $v(s_i, s_{-i})$.

Example 1 *The independent private value model above is a special case: just let $v(s_i, s_{-i}) = s_i$, and suppose that $S_1, ..., S_n$ are independent.*

Example 2 *Another common special case is the pure common value model with conditionally independent signals. In this model, all bidders have the same value, given by some random variable V. The signals $S_1, .., S_n$ are each correlated with V, but independent conditional on it (so for instance, $S_i = V + \varepsilon_i$, where $\varepsilon_1, .., \varepsilon_n$ are independent. Then $v(s_i, s_{-i}) = \mathbb{E}[V|s_1, ..., s_n]$.*

Example 3 *A commonly used, but somewhat hard to motivate, variant of the general model is obtained by letting* $v_i(s_i, s_{-i}) = s_i + \beta \sum_{j \neq i} s_j$, *with* $\beta \leq 1$, *and assuming that* $S_1, ..., S_n$ *are independent. This model has the feature that bidder's have independent information (so a version of the RET applies), but interdependent valuations (so there are winner's curse effects).*

A new feature of the more general auction environment is that each bidder i will want to account for the fact that her opponents' bids reveal something about their signals, information that is relevant for i's own valuation. In particular, if i wins, then the mere fact of winning reveals that her opponent's values were not that high — hence winning is "bad news" about i's valuation. This feature is called the winner's curse.

2.2 Second Price Auction

Let's look for a symmetric increasing equilibrium bid strategy $b(s)$ in this more general environment. We start by considering the bidding problem facing i:

- Let s^i denote the highest signal of bidders $j \neq i$.
- Bidder i will win if she bids $b_i \geq b(s^i)$, in which case she pays $b(s^i)$

Bidder i's problem is then:

$$\max_{b_i} \int_{\underline{s}}^{\overline{s}} \left[\mathbb{E}_{S_{-i}} \left[v(s_i, S_{-i}) \mid s_i, S^i = s^i \right] - b(s^i) \right] \mathbf{1}_{\{b(s^i) \leq b_i\}} f(s^i | s_i) ds^i$$

or

$$\max_{b_i} \int_{\underline{s}}^{b^{-1}(b_i)} \left[\mathbb{E}_{S_{-i}} \left[v(s_i, S_{-i}) \mid s_i, S^i = s^i \right] - b(s^i) \right] dF(s^i | s_i)$$

The first order condition for this problem is:

$$0 = -\frac{1}{b'(b^{-1}(b_i))} \left[\mathbb{E}_{S_{-i}} \left[v(s_i, S_{-i}) \mid s_i, S^i = b^{-1}(b_i) \right] - b(b^{-1}(b_i)) \right] f(b^{-1}(b_i) | s_i)$$

or simplifying:

$$b_i = b(s_i) = \mathbb{E}_{S_{-i}} \left[v(s_i, S_{-i}) \mid s_i, \max_{j \neq i} s_j = s_i \right]$$

That is, in equilibrium, bidder i will bid her expected value conditional on her own signal *and* conditional on all other bidders having a signal less than hers (with the best of the rest equal to hers).

Remark 2 *While we will not purse it here, in this more general environment the revenue equivalence theorem fails. Milgrom and Weber (1982) prove a very general result called the "linkage principle" which basically states that in this general symmetric setting, the more information on which the winner's payment is based, the higher will be the expected revenue. Thus, the first price auction will have lower expected revenue than the second price auction because the winner's payment in the first price auction is based only on her own signal, while in the second price auction it is based on her own signal and the second-highest signal.*

3 Large Auctions & Information Aggregation

An interesting question that has been studied in the auction literature arises if we think about auction models as a story about how prices are determined in Walrasian markets. In this context, we might then ask to what extent prices will aggregate the information of market participants.

We now consider a series of results along these lines. For each result, the basic set-up is the same. There are a lot of bidders, and the auction has pure common values with conditionally independent signals. We consider second price auctions (or more generally, with k objects, $k+1$ price auctions, where all winning bidders pay the $k+1$st highest bid). The basic question is: as the number of bidders gets large, will the auction price converge to the true value of the object(s) for sale?

3.1 Wilson (1977, RES)

Wilson considers information aggregation in a setting with a special information structure. In his setting, bidders learn a lower bound on the object's value. The model has:

- Bidder's $1, ..., n$
- A single object with common value $V \sim U[0, 1]$
- Bidder's signals $S_1, ..., S_n$ are iid with $S_i \sim U[0, v]$ (so signal $s_i \Rightarrow V \geq s_i$).

Wilson shows that as $n \to \infty$, the expected price converges to v. This is easy to see: with lots of bidders, someone will have a signal close to v, and you have to bid very close to your signal to have any chance of winning.

3.2 Milgrom (1979, EMA)

Milgrom goes on to identify a necessary and sufficient condition for information aggregation with many bidders and a fixed number of objects. Milgrom's basic requirement is, in the limit as $n \to \infty$, that for all $v' \in \Omega$ and all $v < v'$, and $M > 0$, there exists some $s' \in S$ such that

$$\frac{P(s'|v')}{P(s'|v)} > M.$$

That is, for any possible value v', there must be arbitrarily strong signals that effectively rule out any value $v < v'$.

This condition builds on Wilson, and the intuition is again quite easy. As $n \to \infty$, there is a very strong winner's curse. In a second price auction, the high bid is:

$$\mathbb{E}_V \left[V \mid s^{1:n}, \max_{j \neq i} S_j = s^{1:n} \right].$$

While it is clear that $s^{1:n}$ is likely to be very high as $n \to \infty$, the only way anyone would ever bid $\overline{v}$ would be to have a signal so strong that conditioning on millions of other signals being lower than your own would not push down your estimate too much.

3.3 Pesendorfer and Swinkels (1997, EMA)

Pesendorfer and Swinkels consider a model that is quite similar to Milgrom.

- n bidders, k objects
- Each bidder has value $V \sim F(\cdot)$ on $[0, 1]$
- Signals $S_i \sim G(\cdot|v)$ on $[0, 1]$, where $g(\cdot|\cdot)$ has the MLRP.
- Signals $S_1, ..., S_n$ are independent conditional on $V = v$.

Remark 3 *Clearly with a large number of independent signals, there is enough information to consistently estimate v. The question then is whether the bids will accurately aggregate this information — i.e. will the price be a consistent estimator?*

Pesendorfer and Swinkel's big idea is the following. While the winner's curse will make a bidder with $s = 1$ shade her bid as $n \to \infty$ with a fixed number of objects k, as $k \to \infty$, there is also a *loser's curse*: if lowering your bid ε matters there must be k people with *higher valuations*. This operates against the winner's curse, counterbalancing it.

Proposition 3 *The unique equilibrium of $k+1^{st}$ price auction is symmetric:*

$$b(s_i) = \mathbb{E}_V \left[V \mid s_i, kth \text{ highest of other signals is also } s_i \right].$$

Definition 1 *A sequence of auctions (n_m, k_m) has information aggregation if the price converges in probability to v as $m \to \infty$.*

PS first consider *necessary* conditions for information aggregation. Clearly, a requirement for information aggregation is that:

$$\lim_{m\to\infty} b_m(0) = 0 \qquad \text{and} \qquad \lim_{m\to\infty} b_m(1) = 1$$

Without this, there will be no convergence when $v = 0$ or when $v = 1$. Now,

$$b_m(1) = \mathbb{E}_V \left[V \mid s_i = 1, k\text{th highest of other signals is } 1 \right]$$

If $k_m \nrightarrow \infty$, then this expectation will shrink below 1. So we must have $k_m \to \infty$. And similarly, looking at $b_m(0)$, we must have $n_m - k_m \to \infty$.

PS go on to show that these conditions are not just necessary, but sufficient for information aggregation.

Proposition 4 *A sequence of auctions has information aggregation if and only if $n_m - k_m \to \infty$ and $k_m \to \infty$.*

The basic idea is that conditioning on $k_m - 1$ bidders having higher signals than s_i gives a very strong signal about true value. So the bidder who actually *has* the kth highest sighnal will tend to be right on when he bids $\mathbb{E}_V \left[V \mid s_i, k\text{th highest of other signals is } s_i \right]$. Thus, we will end up with information aggregation.

References

[1] Harris, Milton and Arthur Raviv (1981) "Allocation Mechanisms and the Design of Auctions," *Econometrica*, 49, 1477–1499.

[2] Kremer, Ilan (2002) "Information Aggregation in Common Value Auctions," *Econometrica*.

[3] Milgrom, Paul (1979) "A Convergence Theorem for Bidding with Differential Information," *Econometrica,* 47.

[4] Milgrom, Paul (2003) *Putting Auction Theory to Work*, Cambridge University Press.

[5] Milgrom, Paul and Robert Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, .

[6] Milgrom, Paul and Ilya Segal (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583-601.

[7] Myerson, Roger (1981) "Optimal Auction Design," *Math. Op. Res*, 6, 58–73.

[8] Pesendorfer, Wolfgang and Jeroen Swinkels (1997) "The Loser's Curse and Information Aggregation in Auctions," *Econometrica*, 65, .

[9] Riley, John and William Samuelson (1981) "Optimal Auction," *American Economic Review*, 71, 381–392.

[10] Vickrey, William (1961) "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–39.

[11] Wilson, Robert (1977) "A Bidding Model of Perfect Competition," *Review of Economic Studies.*

[12] Wilson, Robert (1992) "Strategic Analysis of Auctions," *Handbook of Game Theory.*