Auctions

Glenn Ellison

Massachusetts Insitute of Technology

Introduction

Many real-world applications

- Government procurement
- Treasury auctions
- Art and used-car auctions
- Online advertising
- Auctions for consumer goods

Lots of academic research in auctions in the past few decades

- How to most efficiently allocate resources or raise revenue
- Well-defined market mechanisms and questions
- Nice opportunities to combine theory and empirics

Recent surges in interest

- FCC spectrum auction
- Online advertising
- Structural estimation of bidder preferences

Independent Private Values (IPV) Model

- Bidders indexed by $i = 1, \ldots, n$
- One object to be sold
- Each bidder observes a signal $s_i \in [\underline{s}, \bar{s}]$. $s_i \sim F$, independently of s_{-i}
- Private Values: $v(s_i, s_{-i}) = v(s_i)$. WLOG to set $v(s_i) = s_i$
- Risk-Neutrality: Utility from winning and paying p is $s_i - p$. Utility from losing and paying pay p is $-p$.
- Private Information: Only observe your own signal

Second Price Sealed Bid Auction

Each bidder secretly submits a bid. Highest bidder wins and pays second highest bid

Proposition: In the second price sealed bid auction it is a weakly dominant to bid one's values, $b(s_i) = s_i$

Proof:
Consider a deviation to $b_i \neq s_i$ and let $\hat{b} = \max_{j \neq i} b_j$

- If $\hat{b} > b_i, s_i$: lose under b_i or s_i. No change in i's outcome
- If $b_i > \hat{b} > s_i$: b_i results in winning and paying more than the value
- If $s_i > \hat{b} > b_i$: b_i results in losing when would have gotten surplus
- If $b_i, s_i > \hat{b}$: win and pay $\hat{b}$ under b_i or s_i. No change in i's outcome

In all six possible orderings, the deviation is weakly worse.

Notes:

- Truthtelling is not the unique equilibrium strategy: e.g. i bids $\bar{s}$ and all other bidders bid 0
- Truthful equilibrium is the unique symmetric Bayesian Nash equilibrium if F is continuous and has full support on $[\underline{s}, \bar{s}]$

Second Price Sealed Bid Auction

Expected revenue is the expected second-highest value, denoted $s^{n-1:n}$

$$E[s^{n-1:n}] = \int_{\underline{s}}^{\bar{s}} xn(n-1)F(x)^{n-2}f(x)(1-F(x))dx$$

For example, for $s_i \in U[0,1]$, $F(x) = x$ for $x \in [0,1]$ and the expected revenue is

$$\int_0^1 xn(n-1)x^{n-2}(1-x)dx = \frac{n-1}{n+1}$$

The density calculation above is a an example of the density of the k-th order statistic:

$$f^{k:n}(x) = nf(x)\begin{pmatrix} n-1 \\ k-1 \end{pmatrix} F(x)^{k-1}(1-F(x))^{n-k}$$

English Auction

A game involving unrestricted sequential bids would be complicated to analyze so most models simplify to a "button auction".

- Bidders keep their finger on a button as prices continuously increase.
- Dropout is permanent.
- Auction ends when one bidder remains.

Propositon: A Perfect Bayesian Equilibrium of the button auction model is for all players to drop out at $b(s_i) = s_i$. The highest value bidder wins and pays the second highest price.

Proof:
Players could potentially deviate to a complex strategy where the dropout point is a function of the full history of dropouts. But, if we consider the realized drop-out point b_i given any realization of the opponent's bids, the same argument for the second-price auction shows that the bidder is weakly worse off.

Equivalence with second-price auction can be seen by thinking of the number submitted in a second price auction as dropout instructions to a proxy bidder.

Experiments suggest that button-auction is strategically simpler than the second-price sealed bid version.

First-Price Sealed Bid Auction

A standard mechanism in procurement and other auctions is:

- Bidders submit bids $b_1, \ldots, b_n$
- Highest bidder wins and pays his bid

Clearly, bidders do not want to bid their value. This guarantees zero.

Consider a symmetric equilibrium in strictly increasing strategies $b^*(s)$

The payoff from bidding b_i if others bid according to $b^*(s)$ is

$$Eu_i(b_i, b^*_{-i}; s_i) = (s_i - b_i) Pr\left[b^*(s_j) \leq b_i, \forall j \neq i\right]$$

$$\implies b^*(s_i) = \text{argmax}_{b_i}(s_i - b_i)F^{n-1}(b^{*-1}(b_i))$$

A trick for simplifying the analysis is to think about player i deviating to $b^*(s')$. This implies

$$s_i = \text{argmax}_s(s_i - b^*(s))F^{n-1}(s)$$

The FOC for this maximization is

$$0 = (s_i - b^*(s_i))(n-1)F^{n-2}(s_i)f(s_i) + F^{n-1}(s_i) \cdot (-b^{*\prime}(s_i))$$

First-Price Sealed Bid Auction

The FOC for this maximization is

$$0 = (s_i - b^*(s_i))(n-1)F^{n-2}(s_i)f(s_i) + F^{n-1}(s_i) \cdot (-b^{*\prime}(s_i))$$

One thing we can do with the FOC is to get intuition for bid shading and structural estimation.

Putting $b^*(s_i)$ alone on one side gives

$$\begin{aligned} b^*(s_i) &= s_i - \frac{F^{n-1}(s_i)b^{*\prime}(s_i)}{(n-1)F^{n-2}(s_i)f(s_i)} \\ &= s_i - \frac{G(b^*(s_i))}{g(b^*(s_i))}, \end{aligned}$$

where $G(x)$ is CDF of the highest of $n-1$ bids drawn from $b^*(s)$.
Some implications of this formula are:

- Players will be more aggressive in shading bids when the probability of winning is large and the density of the highest opposing bid low. This is the monopoly-pricing tradeoff.
- We can nonparametrically estimate $G(x)$ from data on opponents bids in repeated auctions. This formula then allows us to recover an estimate of each bidder's value.

First-Price Sealed Bid Auction

The FOC for this maximization is

$$0 = (s_i - b^*(s_i))(n-1)F^{n-2}(s_i)f(s_i) + F^{n-1}(s_i) \cdot (-b^{*\prime}(s_i))$$

We can also use this FOC to solve for the equilibrium.

It can be rearranged as

$$b^*(s_i)(n-1)F^{n-2}(s_i)f(s_i) + b^{*\prime}(s_i)F^{n-1}(s_i) = s_i(n-1)F^{n-2}(s_i)f(s_i)$$

The LHS is an exact derivative. Integrating gives

$$b^*(s_i)F^{n-1}(s_i) = \int_{\underline{s}}^{s_i} x(n-1)F^{n-2}(x)f(x)dx$$

Integrating this by parts (using $u = x$ and $dv = (n-1)F^{n-2}(x)f(x)$ we find

$$b^*(s_i) = s_i - \frac{\int_{\underline{s}}^{s_i} F^{n-1}(x)dx}{F^{n-1}(s_i)}$$

For example, when $s \sim U[0,1]$ this gives

$$b^*(s_i) = s_i - \frac{\int_0^{s_i} x^{n-1}dx}{s_i^{n-1}} = s_i - \frac{s_i^n/n}{s_i^{n-1}} = \frac{n-1}{n}s_i$$

First-Price Sealed Bid Auction

Another expression for the optimal bid helps to understand the auction revenue.

Proposition: $b^*(s)$ is the expected second highest value conditional on s being the highest value.

Proof:
The expression just after the integration step on the previous slide gives

$$b^*(s_i) = \frac{\int_{\underline{s}}^{s_i} x(n-1)F^{n-2}(x)f(x)dx}{F^{n-1}(s_i)}$$

Recognizing the density in the numerator as that of an order statistic gives

$$b^*(s_i) = \frac{\int_{\underline{s}}^{s_i} xf^{n-1:n-1}(x)dx}{F^{n-1}(s_i)} = E(s^{n-1:n-1}|s^{n-1:n-1} < s_i)$$

Corollary: First and second price auctions yield the same expected revenue.

Proof:
Expected revenue in the first price auction is $E_{s^{n:n}}[b^*(s^{n:n})]$. The proposition gives that this is equal to $E_{s^{n:n}}E[s^{n-1:n-1}|s^{n-1:n-1} \leq s^{n:n}] = E[s^{n-1:n}]$

Revenue Equivalence Theorem

Consider a more general auction mechanism. Bidders submit bids $b_1, \dots, b_n$. The good is allocated to bidder i with probability $x_i(b_1, \dots, b_n)$. Payments to the seller are $t_i(b_1, \dots, b_n)$.

Theorem: In the symmetric IPV model, suppose that a general auction mechanism has an equilibrium $b^*(s_i)$ in which (i) the object is awarded to the bidder with the highest value and (ii) a bidder with valuation $\underline{s}$ obtains zero profits. Then in that equilibrium
(a) The seller's expected revenue is $E(s^{n-1:n})$.
(b) The expected utility of a type s_i bidder is $\int_{\underline{s}}^{s_i} F^{n-1}(x)dx$.

The "revenue equivalence" name reflects that an implication of (a) is that all such mechanisms raise the same revenue for the seller.

The earlier revenue results on first- and second-price auctions are special cases. In each of those games the equilibrium bids $b^*(s_i)$ were monotonically increasing in s_i. When this happens the bidder with the highest value will win in equilibrium.

Revenue Equivalence Theorem

- Proof: Consider i's payoff in any equilibrium strategy profile $b_i(\cdot), b_{-i}(\cdot)$:

$$Eu_i(s_i) = \max_{b_i} Eu_i(b_i, b_{-i}(s_{-i}); s_i) = \max_{b_i} s_i E_{b_{-i}}[x_i(b_i, b_{-i})] - E_{b_{-i}}[t_i(b_i, b_{-i})]$$

  - The envelope theorem (b_i maximizes the payoff given s_i and b_{-i}) implies

$$\frac{dEu_i(s_i)}{ds_i} = E_{b_{-i}}[x_i(b_i(s_i), b_{-i}(s_{-i}))] = F^{n-1}(s_i)$$

    because the mechanism awards the object to the highest value bidder
  - Because $Eu_i(\underline{s}) = 0$, we have the formula in (b):

$$Eu_i(s_i) = \int_{\underline{s}}^{s_i} F^{n-1}(x)\mathrm{d}x$$

  - Rewrite the expected payment using as gross surplus minus utility:

$$\begin{aligned} E_{s_{-i}}[t_i(b_i(s_i), b_{-i}(s_{-i}))] &= s_i F^{n-1}(s_i) - Eu_i(s_i) \\ &= s_i F^{n-1}(s_i) - \int_{\underline{s}}^{s_i} F^{n-1}(x)\mathrm{d}x \end{aligned}$$

  - Note that neither term depends on the auction mechanism and that the expression is the same as one of the two we derived earlier for the expected payment in the first price auction. Hence, it is equal to $E[s^{n-1:n-1} | s^{n-1:n-1} < s_i]$. Therefore the expected revenue is $E[s^{n-1:n}]$

All Pay Auction

In an all-pay auction bidders submit bids $b_1, \ldots, b_n$, the high bidder is awarded the object, and all bidders pay what they bid.

Whle rarely used as an auction, the model can also be used to capture situations like lobbying to be awarded a government contract and patent races in which firms expend resources.

The revenue equivalence theorem gives us an easy way to derive equilibrium bidding. Utility is

$$Eu_i(s_i) = s_i F^{n-1}(s_i) - b^*(s_i)$$

For this to be equal to $\int_{\underline{s}}^{s_i} F^{n-1}(x)dx$ we must have

$$b^*(s_i) = s_i F^{n-1}(s_i) - \int_{\underline{s}}^{s_i} F^{n-1}(x)dx$$

Limitations of the Revenue Equivalence Theorem

The Revenue Equivalence Theorem is a beautiful and important result, but it is limited in both

1. The set of situations in which it applies.
   - Bidders are assumed to be risk neutral.
   - Private values only depend on a bidders own signal.
   - Values are independent.
2. The set of mechanisms that it considers.

Bidders could be risk-averse with utility $u(s_i - p_i)$ if they win and pay p_i and utility $u(-p_i)$ if they lose and pay p_i. The first-price auction raises more revenue than a second-price auction in this environment.

Private values are said to be "affiliated" if $f(s_{-i}|s_i)/f(s'_{-i}|s_i)$ is increasing in s_i whenever $s_{-i} > s'_{-i}$. The second price auction has a higher expected revenue in this environment.

A stark example of the limited set of mechanisms is the example of $n = 1$ with $s \sim U[0, 1]$. This is the classic monopoly pricing problem with $D(p) = 1 - p$. The RET only considers mechanism that always transfer the good regardless of the bid (or consumer's value). All such mechanisms produce zero revenue.

Optimal Reserve Prices

In the monopoly model we know the seller must set a reserve price to extract revenue. In an auction one can think that competition lessens the need for a reserve price. But it is also less likely that all bidders will have values below the reserve price and this lessens the downside.

To think about optimal reserve prices in an English auction we consider a relaxed problem in which the seller can choose the reserve price after it sees where the second-to-last bidder has dropped out at $s^{n-1:n}$.

Note that conditional on $s^{n-1:n}$ the probability that $s^{n:n}$ is greater than r (for $r \geq s^{n-1:n}$) is just $\frac{1-F(r)}{1-F(s^{n-1:n})}$. So the expected profit from using reserve price r is $\frac{1}{1-F(s^{n-1:n})} r(1-F(r))$.

Suppose the monopoly profit function $r(1-F(r))$ is single peaked with monoply price p^m. If $s^{n-1:n} < p^m$, then after it learns the second-highest value the optimal reserve is just the monopoly price. If $s^{n-1:n} > p^m$, then the optimal reserve price is $s^{n-1:n}$.

Setting the reserve price for an English auction at p^m will achieve the ex post optimal reserve price for any realization of $s^{n-1:n}$. Hence, this same reserve price is optimal for all n.

Optimal Reserve Price

Comments

- Optimal reserve prices are positive even if seller's value is zero
- Reserve prices reduce social welfare
- The above analysis ignores entry. In practice, bidders may not enter an auction with high reserve prices
- Bulow and Klemperer (1996) show that attracting one additional bidder is better than setting an optimal reserve price
- There may be a commitment problems in setting reserve prices.

Common Value Auctions

In a more general setup signals could be correlated and bidders values need not depend only on their own values. A general model would be:

- Signals $s_1, \ldots, s_n$ have joint density $f(\cdot)$
- Bidder i's expected value is $v(s_i, s_{-i})$

One special case is the pure common values model:

- The object has value of V to all bidders
- Signal $s_i = V + \varepsilon_i$ where $\varepsilon_i \perp \varepsilon_j$
- Bidder's expected value given all signals is $v(s_i, s_{-i}) = \mathbb{E}[V|s_i, s_{-i}]$
- Bidders directly observe only their own signal, s_i

In the common value there will be a "winner's curse."

- Winning the object reveals that other bidders had lower bids, indicating that all had lower signals.
- Hence, winning is "bad news" about i's valuation
- Rational players account for this in their bidding strategy so they are not disappointed to win

Multiunit Auctions

Many prominent auctions involve multiple units or mulitiple goods: Treasury bills, IPOs, Spectrum, Electricity

A simple base model would be a seller with K identical goods.

- Bidder incremental valuations are v_{ik} where k is the number of goods and $v_{ik+1} < v_{ik}$
- Each bidder submits a demand curve $d_i(p)$ or equivalently values for bundles of each size
- Multiple pricing rules are possible including discriminatory ("pay-your-bid") , uniform pricing, and a Vickrey-style rule: a winner of k units pays the k highest losing bids of other bidders

The designers of spectrum auctions have chosen somewhat complicated ascending-bid procedures. In part this reflects a desire to limit two common concerns: demand reduction and inefficient allocation.

Multiunit Auctions

- Consider uniform $K+1$st price sealed bid auction
- Observations
  - Truthful bidding is weakly dominant for bidder that wants only one unit i.e. a bidder with $v_{i2} = 0$ bids v_{i1} for 1 object
  - For bidder that values two units, the price of the first unit depends on the bid for the second, so second bid is shaded. This "demand reduction" can result in inefficient allocation and low revenues
- Example:

	1st unit	2nd unit
Consumer 1	60	40
Consumer 2	30	10

  - Proposition: With full information, in any equilibrium in strategies that are not weakly dominated, each bidder wins one unit.
  - Proof: Suppose bidder 1 gets both units. Price must be at least 30 and 1's surplus is at most 40. If 1 bids 10 for a second unit, then both consumers get one unit and 1's surplus is 50.
  - The model has an equilibrium with zero revenue: 1 bids (60, 0) and 2 bids (30, 0)

Complementarities

Standard auction designs can also lead to inefficiency when values are not additively separable.

- Examples
  - Electricity delivered to different places
  - Spectrum licenses
  - Emission reductions across time
- Example: Spectrum Licenses

	So.Cal.	No.Cal	Both
Bidder 1	v_1	0	v_1
Bidder 2	0	v_2	v_2
Bidder 3	w_1	w_2	$w_1 + w_2 + x$

  where $v_1, v_2 \sim U[10, 20]$, and say $w_1 = w_2 = x = 10$
  - Proposition: Assume that goods are allocated via sequential English auctions. Player 3 wins if and only if $v_1 \leq 15$. This is inefficient if $v_1 + v_2 > 30$.
  - Proof: If Bidder 3 has won the first auction, then she bids 20 and always wins second stage. 3's expected value of winning the first stage is $30 - E(v_2) = 15$. Therefore, she wins first stage if $v_1 < 15$.
- Designing package auctions is hard. Combinatorial bids are complex.

Hendricks and Porter, "An Empirical Study of an Auction with Asymmetric Information," *AER* 1988.

The US government auctions two types of offshore leases in US waters for oil drilling:

- Wildcat tracts are 5000 acre plots that do not abut any previous drilling. Only about 35% end up producing any oil.
- Drainage tracts are 2500 acre plots adjacent to previously drilled wildcat tracts. About 60% of auctioned drainage tracts have oil.

Hendricks and Porter, “An Empirical Study of an Auction”

Here’s a map of tracts in the Gulf of Mexico near New Orleans and Mobile

wildcat tract

drainage tracts

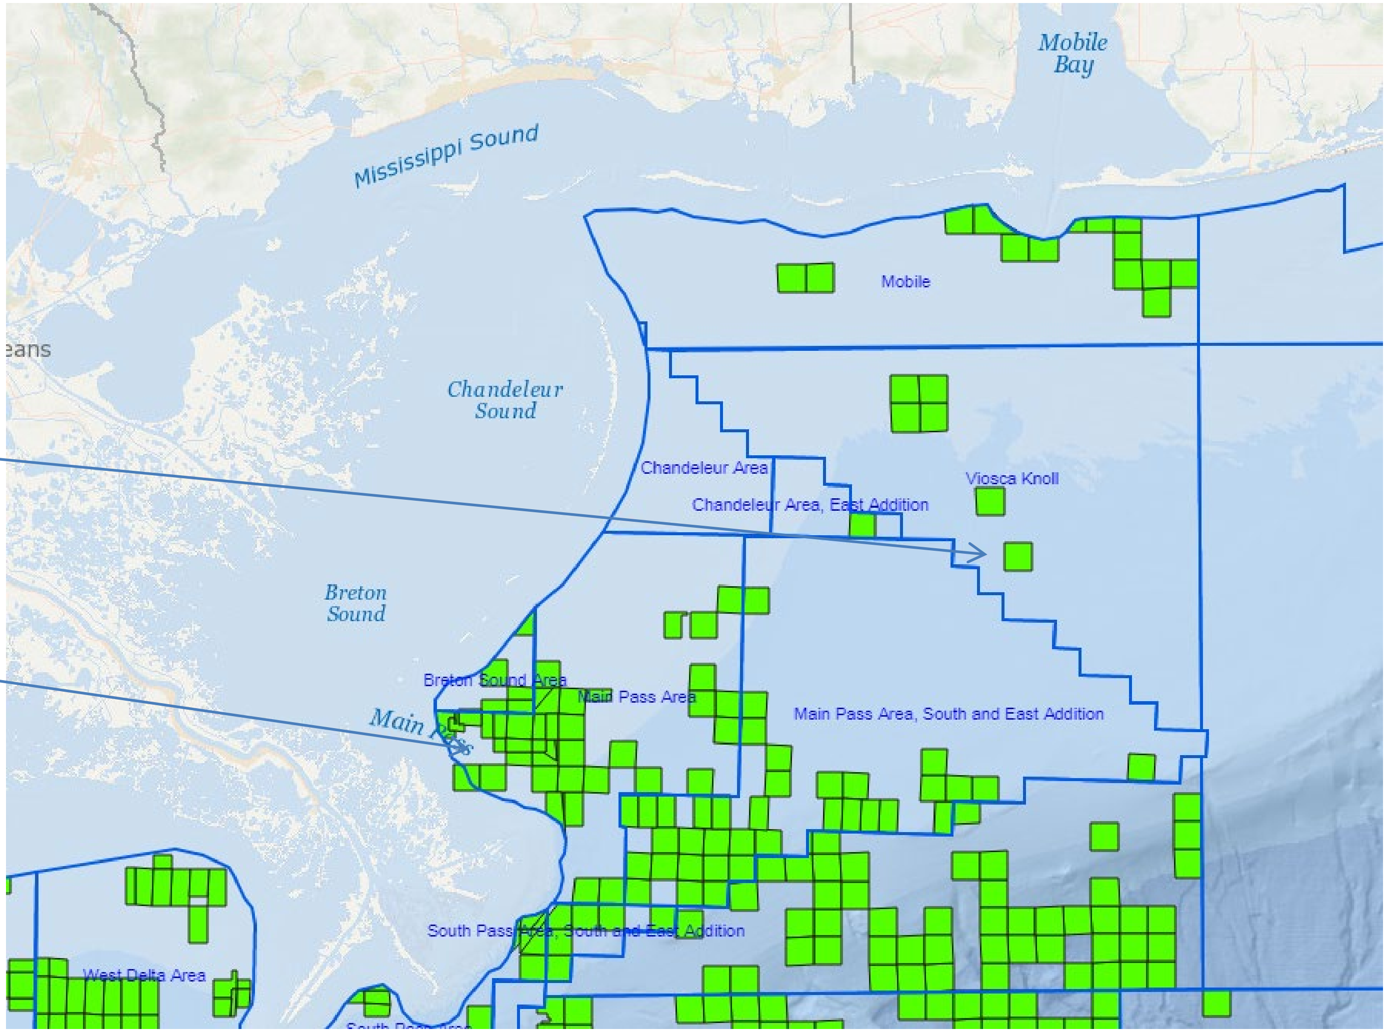

Some of the bidders on drainage tracts will be the current lessees of adjacent wildcat tracts. This creates the asymmetric info.

Hendricks and Porter, “An Empirical Study of an Auction”

Bidders submit sealed bids for a time-limited right to drill. Winners also pay a small percentage royalty on any oil extracted.

Hendricks and Porter note that participation is lower on drainage tracts despite the fact that the drainage tracts are more reliably productive and have higher values.

Drainage profits end up much higher.

H-P suggest that this could reflect that information asymmetries create a severe winner’s curse.

TABLE 1—SELECTED STATISTICS ON WILDCAT AND DRAINAGE TRACTS[a]

	Wildcat	Drainage
Number of Tracts	1056	144
Number of Tracts Drilled	748	124
Number of Productive Tracts	385	86
Average Winning Bid	2.67	5.76
	(0.18)	(1.07)
Average Net Profits	1.22	4.63
	(0.50)	(1.59)
Average Tract Value	5.27	13.51
	(0.64)	(2.84)
Average Number of Bidders	3.46	2.73

Hendricks Porter: "An Empirical Study of an Auction"

They present a simple model to explain the intuition and work out auxiliary predictions that one can examine to assess the model's relevance.

- A single drainage tract has value v to neighbors and $v - c$ to others.
- N uninformed bidders have signal z such that

$$E(v - c|z) \geq R, \frac{d}{dz} E(v|z) > 0, \Pr\{v - c < R|z\} > 0$$

- One informed neighbor sees z and v.
- First price sealed bid auction with reservation price R.

information asymmetry

reserve price

reflecting that it's easier to drill close to where you're already drilling

Hendricks Porter: “An Empirical Study of an Auction”

The game above clearly cannot have a pure strategy equilibrium:

Proposition: There is no pure strategy BNE where the uninformed bidders don’t bid.

Proof: If they bid 0, then we would then have $b_I^*(v) = \begin{cases} 0 \text{ if } v < R \\ R \text{ if } v \geq R \end{cases}$

If this happens, the uninformed bidder can earn positive profits at $R + \epsilon$.

Proposition: There is no pure strategy BNE with $b_u^*(z) \geq R$

Proof: With a pure $b_u^*(z)$ we would have $b_I^*(v, z) = \begin{cases} b_u^*(z) + \epsilon & \text{if } v > b_u^*(z) \\ 0 & \text{otherwise} \end{cases}$

Then the uninformed bidder wins only if $v \leq b_u^*(z)$, so he gets negative profits.

Hendricks Porter: “An Empirical Study of an Auction”

Hendricks and Porter show that the model has a mixed equilibrium where:

- The uninformed bidders mix over $\{0\}, [R, \overline{v}]$
- The informed bidder bids as a function of his private information.

$$b_I^*(v,z) = \begin{cases} 0 & \text{if } v < R \\ R & \text{if } v \text{ a little bigger than } R \text{ and } z \text{ not too high} \\ f(v,z) & \text{for some increasing } f \text{ if } v, z \text{ bigger} \end{cases}$$

Hendricks Porter: “An Empirical Study of an Auction”

A very nice feature of the paper is that this simple model yields several potentially testable predictions:

1. Having $b_I^* = 0$ is less common than $b_u^* = 0$
2. The informed bidder wins with probability $\geq \frac{1}{2}$
3. $E(\Pi_u^*) = 0$. Profits are negative if the informed bidder bids 0 and positive if the informed bidder bids at least R.
4. $E(\Pi_I^*) > 0$
5. For $c \approx 0$ the *ex ante* bid distributions of informed and uninformed bidders are identical
6. b_I^* is independent of N
7. $f(v, z)$ increasing in z and v.

Hendricks Porter: "An Empirical Study of an Auction"

The dataset has information on 114 drainage tracts auctioned between 1959 and 1969. All are adjacent to a wildcat tract for which they have production data. When multiple neighbors exist, the neighbors can submit a joint bid.

The data include bids of neighbors and non-neighbors and ex-post production.

TABLE 2—DEFINITION OF VARIABLES[a]

	Mean	Standard Deviation
B_I: maximum bid by neighbor	3.78	11.52
B_U: maximum bid by non-neighbor	3.60	9.57
N_I: number of neighbor bids	1.00	0.67
N_U: number of non-neighbor bids	1.69	2.09
N: number of neighbor tracts	3.01	1.98
NF: number of neighbor firms	2.06	1.08
π: *ex post* tract gross profitability	8.75	20.83
V: *ex post* gross profits of adjacent tract	14.51	20.16
A: tract acreage	2.679	1.533

[a] Dollar figures are in millions of $1972. Tract acreage is in thousands of acres.

Hendricks Porter: "An Empirical Study of an Auction"

Tests of Model Predictions:

1. Neighbors bid 83% of the time; non-neighbors 68% ✓
2. Neighbors won 52% of tracts √
3. $E(\Pi_u) \approx 0$ and $\begin{cases} \text{positive} & \text{if neighbor enters} \\ \text{negative} & \text{otherwise} \end{cases}$ ✓
4. $E(\Pi_I) > 0$ ✓
5. To test symmetry, run regressions:

   $$b_I/R = a_0 + a_1 V + a_2 V^2 + a_3 \text{Acreage} + a_4 \#\ \text{Neighbors} + \epsilon_I$$ √

   $$b_u/R = b_0 + b_1 V + b_2 V^2 + b_3 \text{Acreage} + b_4 \#\ \text{Neighbors} + \epsilon_u$$

   Can't reject coefficient equality. ϵ's have similar variances.
6. Find no effect in regression ✓
7. Bids increase in true value <u>and</u> wildcat value ✓

Hendricks Porter: "An Empirical Study of an Auction"

TABLE 3—SAMPLE STATISTICS ON TRACTS WON BY EACH TYPE OF FIRM[a]

	Wins by Neighbor Firms		Wins by Non-Neighbor Firms		
	A	Total	B	C	Total
No. of Tracts	35	59	19	36	55
No. of Tracts Drilled	23	47	18	33	51
No. of Productive Tracts	16	36	12	19	31
Average Winning Bid	3.28	6.04	2.15	6.30	4.87
	(0.56)	(2.00)	(0.67)	(1.31)	(0.92)
Average Gross Profits	10.05	12.75	−0.54	7.08	4.45
	(3.91)	(3.21)	(0.47)	(2.95)	(1.99)
Average Net Profits	6.76	6.71	−2.69	0.78	−0.42
	(3.02)	(2.69)	(0.86)	(2.64)	(1.76)

[a] Dollar figures are in millions of $1972. The numbers in parentheses are the standard deviations of the sample means. Column A refers to tracts which received no non-neighbor firm bid, column B refers to tracts which received no neighbor bid, and column C to those in which a neighbor firm bid, but a non-neighbor firm won the tract.

evidence for point 1) above.

evidence for point 3) above.

On Wednesday Tobias will discuss structural empirical auction work.

Papers that he'll cover will include

- *Guerre, Perrigne, and Vuong*
- *Athey, Levin, and Seira*

Next Monday I'll be back with empirical papers on advertising probably including

- *Milyo and Waldfogel*
- *Lewis and Reiley*
- *Aridor, Che, and Salz*
- *Shapiro*

See you then!