Mechanism Design: Private Value Auctions

Dilip Mookherjee

Boston University

Ec 703b Lectures 7/8 (text: FT Ch 7, pp 271-275)

Private Value Auctions: Introduction

- Consider the simplest setting with a seller S who owns an indivisible good, and:
  - S values the good personally at c dollars
  - there are n potential buyers/bidders $i = 1, \ldots, n$ where i values the object at θ_i, and is privately informed about realization of θ_i
  - *Private Values:* common knowledge that $\theta_1, \ldots, \theta_n$ are drawn from **independent** distributions $F_1, \ldots, F_n$, where support of F_i is $[\underline{\theta}_i, \bar{\theta}_i]$ with positive density f_i throughout

Private Value Auctions, contd.

- Bidders are all risk-neutral: payoff equals θ_i less price paid (in the event of winning), and 0 otherwise
- Participation in the auction is voluntary for all buyers; non-participation payoff is zero
- Seller is risk-neutral; payoff equals price received minus c, in the event of sale, and 0 otherwise

Common Forms of Auctions in the Real World

- Open bidding:
  - English auction (ascending prices/bids, last remaining bidder wins)
  - Dutch auctions (descending prices, first entering bidder wins)
- Sealed bids:
  - First price auctions (highest bidder wins, pays her bid)
  - Second price auctions (highest bidder wins, pays second highest bid)
- Most auctions also have a seller's reserve price: for sale to occur the winning bid must exceed the reserve price
- None of these auctions charge any entry fee, and only winners pay

Key Questions

- 1. How can we rank these commonly observed forms of auctions, from the standpoint of expected payoff of the seller?
- 2. What is the optimal auction, which maximizes seller's expected payoff?

1. Comparing Auctions

- First describe how to analyze the outcome of any given auction
- In any given auction with a given set of rules and reserve price, model bidder's behavior as a Bayesian Nash Equilibrium (BNE)
- Let $b_i(\theta_i)$ denote the bidding strategy of i (bid submitted in a sealed-bid auction, or last/first price in English/Dutch auction at which bidder is active)
- In any of these auctions, the highest bidder wins, and pays an amount $p(b_i, b_{-i})$ that depends on her own bid b_i, bids of others b_{-i} (if there is a reserve price r, treat this as a bid submitted by the seller)

Bayesian Equilibrium Bidding Strategies, contd.

- E.g, in first-price auction or Dutch auction, $p = b_i > \max_{j \neq i} b_j$; in second-price auction or English auction, $p = \max_{j \neq i} b_j < b_i$
- Let G_i denote bidder i's beliefs (cdf) over $z \equiv \max_{j \neq i}\{b_j(\theta_j)\}$, formed by composing the strategies of other bidders $b_j(\theta_j)$ with the distributions $F_j(\theta_j)$ over their types
- Then given a bid b_i, bidder i wins with probability $G_i(b_i)$
- A BNE $b_i(\theta_i)$ satisfies:

$$b_i = b_i(\theta_i) \quad \text{Max}_{b_i}[\theta_i G_i(b_i) - \int_{\underline{b}_{-i}}^{b_i} p(b_i, z) dG_i(z)]$$

Vickrey's Revenue Equivalence Theorem (VRET)

Vickrey (1961) focused on **symmetric private value contexts**, where valuations are i.i.d., all bidders use the same bidding strategy $b(\theta_i)$ which is strictly increasing (easy to check that these exist in any of these auction forms)

Theorem

In any symmetric private values context, the English, Dutch, first-price and second-price sealed bid auctions (with a zero reserve price) generate the same expected revenue.

Intuition underlying VRET

- Given private values, bidders do not learn anything about the object or their own values, from observing the bidding behavior of others
- Dutch auction is (strategically) equivalent to first price auction
- English auction is (strategically) equivalent to second price auction
- Compare first price and second price auctions:
  - In second price auction, dominant strategy to submit a bid equal to the true value, but winner pays the second-highest bid
  - In first price auction, winner pays own bid, but must bid below own true value
  - Extent to which bids are shaded below own true value in the first price auction, depends on what the bidder expects the second-highest bid to be (conditional on winning)

Steps in the proof of VRET

- We shall prove a more general version of this theorem (next slide)
- Notation: Consider any auction and any Bayesian Nash Equilibrium (BNE) among bidders in that auction
  - In *ex post* state (θ_i, θ_{-i}), let $x_i(\theta_i, \theta_{-i})$ denote the probability that i wins the object, and $t_i(\theta_i, \theta_{-i})$ be the payment made by i to the seller.
  - Let the corresponding (interim) expected probability of winning for i be denoted $X_i(\theta_i) \equiv E_{\theta_{-i}}[x_i(\theta_i, \theta_{-i})]$ and (interim) expected payment be $T_i(\theta_i) \equiv E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})]$.
  - Let the (interim) expected payoff of bidder i be denoted $W_i(\theta_i) \equiv \theta_i X_i(\theta_i) - T_i(\theta_i)$.

Generalized Revenue Equivalence Theorem

Theorem

Any two auctions with corresponding BNEs that result in the same (interim) winning probabilities for every bidder ($\{X_i(.)\}_i$) and the same payoffs for the lowest valuation types ($\{W_i(\underline{\theta}_i)\}_i$) generate the same expected payoff for the seller.

Proof of the GRET

- Use arguments analogous to those underlying the Revelation Principle, the seller's expected payoff from any BNE in any auction can be expressed as

$$\Pi \equiv \sum_i E_{\theta_i}[T_i(\theta_i)] - c \sum_i E_{\theta_i}[X_i(\theta_i)]$$

- and the BNE strategies must satisfy $(\forall i, \forall \theta_i :)$

$$\tilde{\theta}_i = \theta_i \quad \text{Max} \quad W_i(\tilde{\theta}_i|\theta_i) \equiv [\theta_i X_i(\tilde{\theta}_i) - T_i(\tilde{\theta}_i)] \qquad (BIC)$$

  i.e., truth-telling must be optimal for each bidder in a revelation mechanism where reports of $\{\tilde{\theta}_i\}_i$ are followed by i paying $t_i(\tilde{\theta}_i, \tilde{\theta}_{-i})$ and winning the object with probability $x_i(\tilde{\theta}_i, \tilde{\theta}_{-i})$

Proof of the GRET, contd.

- Now use the Mirrlees-Myerson theorem for single agent contexts, to argue that (BIC) is equivalent to:
  - (a) $W_i(\theta_i) = W_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i)d\tilde{\theta}_i$
  - and (b) $X_i(.)$ nondecreasing
- Equation (a) can be rewritten as:

$$\theta_i X_i(\theta_i) - T_i(\theta_i) = W_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i)d\tilde{\theta}_i$$

implying

$$T_i(\theta_i) = \theta_i X_i(\theta_i) - \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i)d\tilde{\theta}_i - W_i(\underline{\theta}_i)$$

i.e., the expected payment function for bidder i is entirely determined by the interim winning probability function $X_i(.)$ and the payoff $W_i(\underline{\theta}_i)$ of the lowest valuation type of i, QED

Proof of Vickrey's RET

- In any symmetric auction with zero reserve price where bidders use the same (strictly increasing) strategy $b(.)$, bidder i wins provided $b(\theta_i) > \max_{j \neq i}\{b(\theta_j)\}$
- i.e., provided $\theta_i > \max_{j \neq i}\{\theta_j\}$
- hence every symmetric auction has the same interim win probability function $X_i(\theta_i) = \text{Prob} \quad [\theta_i > \max_{j \neq i}\{\theta_j\}]$
- Finally, the lowest type bids the least and never wins, so obtains a zero payoff
- Now apply GRET

2. Optimal Private Value Auctions

- So far we considered specific auctions and compared the expected revenues that they would generate
- Now ask: what is the optimal auction for a profit maximizing seller?
- Allow asymmetric bidders, and reserve price to be set
- By the Revelation Principle, P can confine attention to revelation mechanisms specifying $\{x_i(\theta_i, \theta_{-i}), t_i(\theta_i, \theta_{-i})\}_i$ satisfying feasibiility constraints:

$$x_i(\theta_i, \theta_{-i}) \in [0, 1], \sum_i x_i(\theta_i, \theta_{-i}) \leq 1$$

besides Bayesian incentive compatibility and interim participation constraints

Optimal Private Value Auctions, contd.

With $X_i(\theta_i) \equiv E_{\theta_{-i}}[x_i(\theta_i, \theta_{-i})]$, $T_i(\theta_i) \equiv E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})]$, $W_i(\theta_i) \equiv \theta_i X_i(\theta_i) - T_i(\theta_i)$, problem is to maximize P's expected profit:

$$\Pi = \sum_i E_{\theta_i}[T_i(\theta_i) - cX_i(\theta_i)]$$

Constraints (for all i):

$$T_i(\theta_i) = \theta_i X_i(\theta_i) - \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i) d\tilde{\theta}_i - W_i(\underline{\theta}_i) \qquad (BIC1)$$

$$X_i(\theta_i) \quad \text{nondecreasing} \qquad (BIC2)$$

$$W_i(\underline{\theta}_i) \geq 0 \qquad (IPC)$$

$$x_i(\theta_i, \theta_{-i}) \in [0,1], \sum_i x_i(\theta_i, \theta_{-i}) \leq 1 \qquad (F)$$

Steps in Solving for Optimal Auction

- Substitute (BIC) into the objective function, to express Π depending only on $\{X_i(.), \underline{W}_i \equiv W_i(\underline{\theta}_i)\}_i$ (same argument to prove GRET):

$$\Pi = \sum_i E_{\theta_i}[\theta_i X_i(\theta_i) - \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i) d\tilde{\theta}_i - \underline{W}_i - cX_i(\theta_i)]$$

- Integrating by parts (just as in single agent problems):

$$\Pi = \sum_i E_{\theta_i}[\{v_i(\theta_i) - c\}X_i(\theta_i)] - \sum_i \underline{W}_i$$

where $v_i(\theta_i) \equiv \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$ is the *virtual value* of i

- Optimal to set $\underline{W}_i = 0$

Steps in Solving for Optimal Auction, contd.

- Consider problem of choosing functions $\{x_i(\theta_i, \theta_{-i})\}_i$ to maximize

$$\Pi = \sum_i E_{\theta_i, \theta_{-i}}[\{v_i(\theta_i) - c\} x_i(\theta_i, \theta_{-i})]$$

subject to $(\forall \theta_i, \theta_{-i}; \forall i{:})$

$$x_i(\theta_i, \theta_{-i}) \geq 0, \sum_i x_i(\theta_i, \theta_{-i}) \leq 1$$

- Call this the *relaxed* problem (dropping (BIC2): $X_i(.)$ is nondecreasing)
- Later check whether the solution to the relaxed problem satisfies (BIC2)

Solving the Relaxed Problem

- *Point-wise optimization:* fix any state $\theta = (\theta_i, \theta_{-i})$ and choose $x_i, i = 1, \ldots, n$ to maximize $\sum_i [v_i(\theta_i) - c]x_i$ subject to $x_i \geq 0, \sum_i x_i \leq 1$
- Define $s \equiv \sum_j x_j$ the probability of sale, and provided $s > 0$, define $a_i \equiv \frac{x_i}{s}$ the probability of selling it to i, conditional on selling it
- Set of controls equivalently written as $s, \{a_i\}_i$ (so $x_i \equiv sa_i$), objective is $s \sum_i [v_i(\theta_i) - c]a_i$, constraints: $s \in [0,1], a_i \in [0,1], \sum_j a_j = 1$
- **Solution:**
  $a_i = 1$ if $v_i(\theta_i) \geq \max_j\{v_j(\theta_j\}$, and 0 otherwise
  $s = 1$ if $\max_j\{v_j(\theta_j)\} > c$, and 0 otherwise
  (ignore ties as they would happen with zero probability)

Solution to the Relaxed Problem

$$x_i(\theta_i, \theta_{-i}) = 1 \quad \text{if} \quad v_i(\theta_i) \geq \max\{\max_j \{v_j(\theta_j)\}, c\} \quad \text{and 0 otherwise}$$

If bidders are ex ante symmetric ($F_i = F$ and $v_i = v$, all i) this reduces to:

$$x_i(\theta_i, \theta_{-i}) = 1 \quad \text{if} \quad \theta_i \geq \max\{\max_j \{\theta_j\}, v^{-1}(c)\} \quad \text{and 0 otherwise}$$

In words: treat seller as an additional (potential) bidder reporting a valuation of $v^{-1}(c)$ (*reserve price*), and then allocate the object to the bidder reporting the highest valuation

When is this the Solution to the Original Problem?

- If $v_i(\theta_i) \equiv \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$ is increasing, then $x_i(\theta_i, \theta_{-i})$ is increasing in θ_i, and $X_i(.) \equiv E_{\theta_{-i}}[x_i(\theta_i, \theta_{-i})]$ is increasing in θ_i
- Hence monotone hazard rates for all F_i's ensures this is the solution to the original problem

Optimal Prices in the Problem with Symmetric Bidders

- All buyers have same interim probability of winning: $X_i = X(\theta_i) \equiv \text{Prob}[z < \theta_i]$, where $r \equiv v^{-1}(c)$ and $z \equiv \max\{r, \max_{j \neq i}\{\theta_j\}\}$
- Letting G denote cdf of z, we have $X(\theta_i) = G(\theta_i)$
- Optimal transfers must satisfy

$$
\begin{aligned}
T_i(\theta_i) &= \theta_i X_i(\theta_i) - \int_{\underline{\theta}}^{\theta_i} X_i(\tilde{\theta}_i) d\tilde{\theta}_i \\
&= \theta_i G(\theta_i) - \int_{\underline{\theta}}^{\theta_i} G(\tilde{\theta}_i) d\tilde{\theta}_i \\
&= \int_{\underline{\theta}_i}^{\theta_i} z dG(z)
\end{aligned}
$$

Optimal Auction with Symmetric Bidders

- So it is optimal for i to pay z if $z < \theta_i$ and 0 otherwise
- *This is exactly the second-price auction!*
- By Vickrey's RE Theorem, all four auction forms (English, Dutch, first-price and second price auctions) are optimal (combined with a suitable reserve price, corresponding to $r = v^{-1}(c)$ in the revelation/second-price mechanism)
- The outcome is ex post Pareto efficient only if $v^{-1}(c) \leq \underline{\theta}_i$ or $c \leq v(\underline{\theta}_i)$, whence the object is sold with probability one; otherwise it is sold with too low a probability (reflecting the monopoly power of the seller)