Lec 2: Myerson's Revenue-Optimal Auction

Guest Lectures at ZJU Computer Science
(Summer 2024)

Instructor: Haifeng Xu

Outline

- Problem Setup and Simplifications
- The Revenue-Optimal Auction
- Key Proof Ideas

Recap: Single-Item Allocation

- A single and indivisible item, n buyers $\{1, \cdots, n\} = [n]$
- Buyer i has a (private) value $v_i \in V_i$ about the item
- Outcome: choice of the winner of the item, and payment p_i from each buyer i
- Objectives: maximize revenue
  - Last lecture: VCG auction maximizes welfare (even for multiple items)

The Mechanism Design Problem

Mechanism Design for Single-Item Allocation

Described by $\langle n, V, X, u \rangle$ where:

- $[n] = \{1, \cdots, n\}$ is the set of n buyers
- $V = V_1 \times \cdots \times V_n$ is the set of all possible value profiles
- $X = \{x \in [0,1]^n \colon \sum_i x_i \leq 1\}$ = set of all possible allocation rules
- $u = (u_1, \cdots, u_n)$ where $u_i = v_i x_i - p_i$ is the utility function of i for any outcome $x \in X$ and payment p_i required from i

Objective: maximize revenue $\sum_{i \in [n]} p_i$

- Cannot have any guarantee without additional assumptions
- Will assume public prior knowledge on buyer values v_i's.
- For today, always assume $v_i \sim f_i$ independently

The Mechanism Design Problem

Mechanism Design for Single-Item Allocation

Described by $\langle n, V, X, u \rangle$ where:

- $[n] = \{1, \cdots, n\}$ is the set of n buyers
- $V = V_1 \times \cdots \times V_n$ is the set of all possible value profiles
- $X = \{x \in [0,1]^n : \sum_i x_i \leq 1\}$ = set of all possible allocation rules
- $u = (u_1, \cdots, u_n)$ where $u_i = v_i x_i - p_i$ is the utility function of i for any outcome $x \in X$ and payment p_i required from i

Remarks:

- (Naturally) assume all players maximize expected utilities
- Will guarantee $\mathbb{E}[u_i] \geq 0$ (a.k.a., individually rational or IR)
  - Otherwise, players would not even bother coming to your auction

The Design Space (i.e., Possible Mechanisms)

A mechanism (i.e., the game) is specified by $\langle A, g \rangle$ where:

- $A = A_1 \times \cdots \times A_n$ where A_i is allowable actions for buyer i
- $g: A \to [x, p]$ maps any $a = (a_1, \cdots, a_n) \in A$ to
  [an allocation outcome $x(a)$ & a vector of payments $p(a)$]

- That is, we will design a game $\langle A, g \rangle$
- Players' utility function will be fully determined by $\langle A, g \rangle$
- This is a game with incomplete information – v_i is privately known to player i; all other players only know its prior distribution

The Design Space (i.e., Possible Mechanisms)

A mechanism (i.e., the game) is specified by $\langle A, g \rangle$ where:

- $\mathrm{A} = A_1 \times \cdots \times A_n$ where A_i is allowable actions for buyer i
- $g: A \rightarrow [x, p]$ maps any $a = (a_1, \cdots, a_n) \in A$ to

  [an allocation outcome $x(a)$ & a vector of payments $p(a)$]

Example 1: first-price auction

- $A_i = \mathbb{R}_+$ for all i
- $g(a)$ allocates the item to the buyer $i^* = \arg\max_{i \in [n]} a_i$ and asks i^* to pay a_{i^*}, and all other buyers pay 0

The Design Space (i.e., Possible Mechanisms)

A mechanism (i.e., the game) is specified by $\langle A, g \rangle$ where:

- $\mathrm{A} = A_1 \times \cdots \times A_n$ where A_i is allowable actions for buyer i
- $g: A \to [x, p]$ maps any $a = (a_1, \cdots, a_n) \in A$ to
  [an allocation outcome $x(a)$ & a vector of payments $p(a)$]

Key Challenge 1

- Design space $\langle A, g \rangle$ generally can be arbitrary and too large
- E.g, the following is a valid – though weird – mechanism
  - $A_i = \{jump\ twice\ (J),\ kick\ in\ a\ penalty\ ball\ (L)\}$
  - $x(a)$ gives the item to anyone of L uniformly at random
  - $p(a)$ asks winner to pay #Jumps

The Solution: Revelation Principle

That is, an observation stating that focusing on certain natural class of mechanisms is **without loss of generality**

Direct Revelation Mechanisms

Definition. A mechanism $\langle A, g \rangle$ is a direct revelation mechanism if $A_i = V_i$ for all i. In this case, mechanism is described only by g.

- That is, the action for each player is to "report" her value (but they don't have to be honest…yet)
- Examples: second-price auction, first-price auction
- Note: this constrains our design space as it limits choice of A_i's
  - It will not reduce our best achievable revenue, as turned out

Incentive-Compatibility (IC)

Definition. A direct revelation mechanism g is Bayesian incentive-compatible (a.k.a., truthful or BIC) if truthful bidding forms a Bayes Nash equilibrium in the resulting game

- A similar but stronger IC requirement

Definition. A direct revelation mechanism g is Dominant-Strategy incentive-compatible (a.k.a., truthful or DSIC) if truthful bidding is a dominant-strategy equilibrium in the resulting game

- A DSIC mechanism is also BIC

Incentive-Compatibility: Examples

Second-price auction is DSIC, and thus also BIC.

First-price auction is not BIC, neither DSIC.

Example (posted price mechanism). Auctioneer simply posts a fixed price p to players in sequence until one buyer accepts.

- Not a direct revelation mechanism as buyer's action is only to accept or not accept, but not report their value (hence no BIC)
- But this can be converted to an equivalent direct BIC mechanism:
  1. Ask each buyer to report their value v_i
  2. Designer accepts iff $v_i \geq p$ and charges p (essentially, designer simulates buyer's actions in the original mechanism)

The Revelation Principle

Theorem. For any mechanism achieving revenue R at a Bayes Nash equilibrium [resp. dominant-strategy equilibrium], there is a direct revelation, Bayesian incentive-compatible [resp. DSIC] mechanism achieving revenue R.

Remarks

- Can be stated more generally, but this version is sufficient for our purpose of optimal auction design
  - The same proof idea: simulating buyer's actions in original mechanism
- Can thus focus on BIC mechanisms henceforth

The Revelation Principle

Theorem. For any mechanism achieving revenue R at a Bayes Nash equilibrium [resp. dominant-strategy equilibrium], there is a direct revelation, Bayesian incentive-compatible [resp. DSIC] mechanism achieving revenue R.

This simplifies our mechanism design task

Optimal Mechanism Design for Single-Item Allocation

Given instance $\langle n, V, X, u \rangle$, supplemented with prior $\{f_i\}_{i\in[n]}$, design the allocation function $x: V \to X$ and payment $p: V \to \mathbb{R}^n$ such that truthful bidding is a BNE in the following Bayesian game:

1. Solicit bid $b_1 \in V_1, \cdots, b_n \in V_n$
2. Select allocation $x(b_1, \cdots, b_n) \in X$ and payment $p(b_1, \cdots, b_n)$

Design goal: maximize expected revenue

Outline

- Problem Setup and Simplifications
- The Revenue-Optimal Auction
- Key Proof Ideas

Optimal (Bayesian) Mechanism Design

➢ Previous formulation and simplification leads to the following optimization problem

BIC constraints

$$\max_{x,p} \ \mathbb{E}_{v\sim f} \ \sum_{i=1}^{n} p_i(v_1, \cdots, v_n)$$

$$\text{s.t.} \quad \boxed{\begin{aligned} &\mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})\right] \\ &\quad \geq \mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(b_i, v_{-i}) - p_i(b_i, v_{-i})\right], \end{aligned}} \qquad \forall i \in [n], v_i, b_i \in V_i$$

$$\boxed{\mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})\right] \geq 0,} \qquad \forall i \in [n], v_i \in V_i$$

$$x(v) \in X, \qquad \forall v \in V$$

Individually rational (IR) constraints

Optimal (Bayesian) Mechanism Design

- Previous formulation and simplification leads to the following optimization problem

$$\max_{x,p} \ \mathbb{E}_{v\sim f} \ \sum_{i=1}^{n} p_i(v_1, \cdots, v_n)$$

$$\begin{aligned} \text{s.t.} \quad & \mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})\right] \\ & \quad \geq \mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(b_i, v_{-i}) - p_i(b_i, v_{-i})\right], && \forall i \in [n], v_i, b_i \in V_i \\ & \mathbb{E}_{v_{-i}\sim f_{-i}}\left[v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})\right] \geq 0, && \forall i \in [n], v_i \in V_i \\ & x(v) \in X, && \forall v \in V \end{aligned}$$

- This problem is challenging because we are optimizing over functions $x: V \to X$ and $p: V \to \mathbb{R}^n$

Optimal DSIC Mechanism Design

➢ Designing optimal **DSIC** mechanism is a strictly more constrained optimization problem

$$\max_{x,p} \; \mathbb{E}_{v\sim f} \sum_{i=1}^{n} p_i(v_1, \cdots, v_n)$$

$$\text{s.t.} \quad [v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})] \qquad \forall\, v_{-i}$$

$$\geq \quad [v_i x_i(b_i, v_{-i}) - p_i(b_i, v_{-i})], \qquad \forall i \in [n], v_i, b_i \in V_i$$

$$\mathbb{E}_{v_{-i}\sim f_{-i}}[v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i})] \geq 0, \qquad \forall i \in [n], v_i \in V_i$$

$$x(v) \in X, \qquad \forall v \in V$$

Corollary. Optimal DSIC mechanism achieves revenue at most that of optimal BIC mechanism.

Myerson's Optimal Auction

Theorem (informal). For single-item allocation with prior distribution $v_i \sim f_i$ independently, the following auction is BIC and optimal:

1. Solicit buyer values $v_1, \cdots, v_n$
2. Transform v_i to "virtual value" $\phi_i(v_i)$ where $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$
3. If $\phi_i(v_i) < 0$ for all i, keep the item and no payments
4. Otherwise, allocate item to $i^* = \arg\max_{i \in [n]} \phi_i(v_i)$ and charge him the minimum bid needed to win, i.e., $\phi_i^{-1}\left(\max\left(\max_{j \neq i^*} \phi_j(v_j), 0\right) \right)$; Other bidders pay 0.

It turns out to deeply relate to, though critically differ from, second price auction.

Remarks

Myerson's optimal auction is noteworthy for many reasons

- Matches practical experience: when buyer values are i.i.d, optimal auction is a second price auction with reserve $\phi^{-1}(0)$.
- Applies to "single parameter" problems more generally
- The optimal BIC mechanism just so happens to be DSIC and deterministic!!
  - Not true for multiple items – there exists revenue gap even when selling two items to two bidders

Outline

- Problem Setup and Simplifications
- The Revenue-Optimal Auction
- Key Proof Ideas

Key Quantities

- Expected winning probability of buyer i, as a function of bid b_i

$$x_i(b_i) = \mathbb{E}_{v_{-i} \sim f_{-i}} x_i(b_i, v_{-i})$$

- Similarly, expected payment

$$p_i(b_i) = \mathbb{E}_{v_{-i} \sim f_{-i}} p_i(b_i, v_{-i})$$

- Consequently,
  - Expected bidder utility of b_i

$$v_i x_i(b_i) - p_i(b_i)$$

  - If BIC, expected revenue

$$\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i} p_i(v_i)$$

Step 1: Myerson's Monotonicity Lemma

Lemma. Consider single-item allocation with prior distribution $v_i \sim f_i$ independently. A direct-revelation mechanism with interim allocation x and interim payment p is BIC if and only if for each buyer i:

1. $x_i(b_i)$ is a monotone non-decreasing function of b_i
2. $p_i(b_i)$ is uniquely determined as follows

$$p_i(b_i) = b_i \cdot x_i(b_i) - \int_{b=0}^{b_i} x_i(b)\, db\,.$$

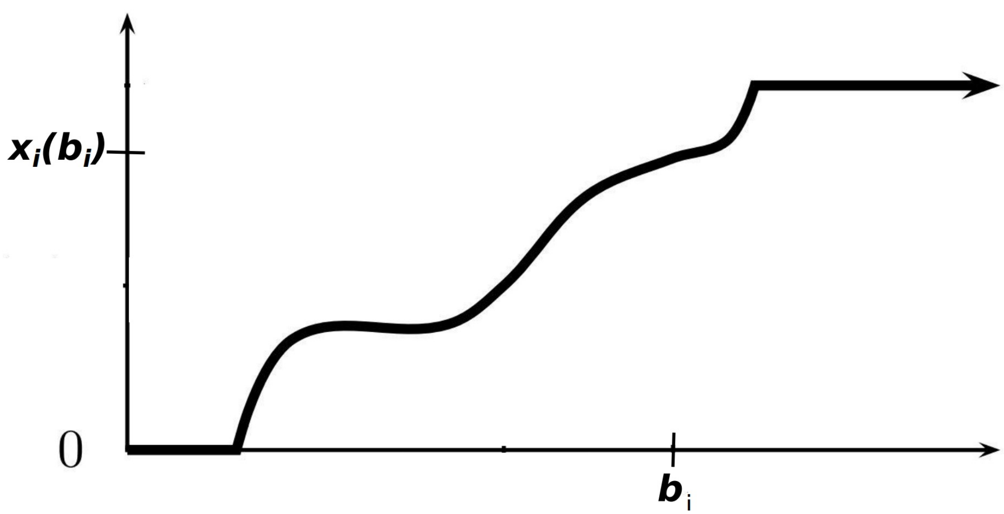

Step 1: Myerson's Monotonicity Lemma

Lemma. Consider single-item allocation with prior distribution $v_i \sim f_i$ independently. A direct-revelation mechanism with interim allocation x and interim payment p is BIC if and only if for each buyer i:

1. $x_i(b_i)$ is a monotone non-decreasing function of b_i
2. $p_i(b_i)$ is uniquely determined as follows

$$p_i(b_i) = b_i \cdot x_i(b_i) - \int_{b=0}^{b_i} x_i(b)\, db \, .$$

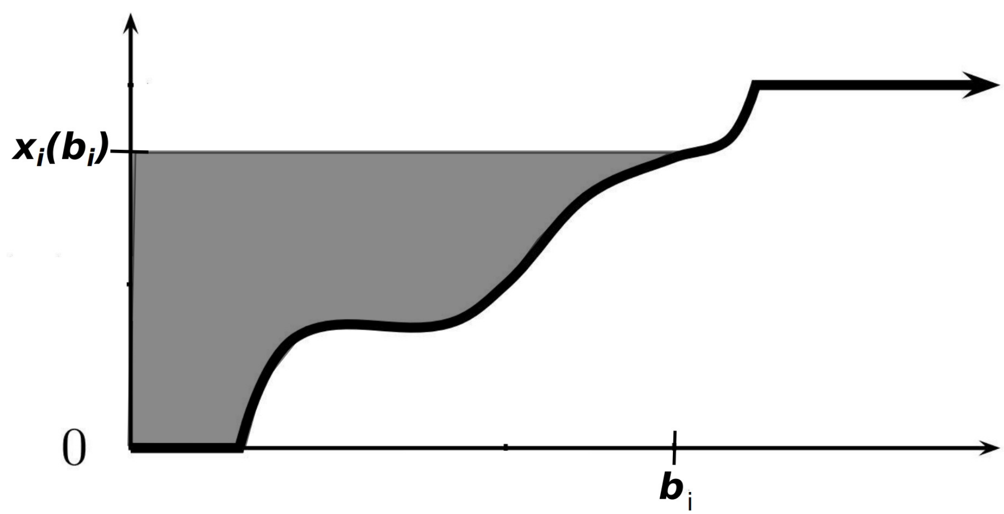

Interpretation of Myerson’s Lemma

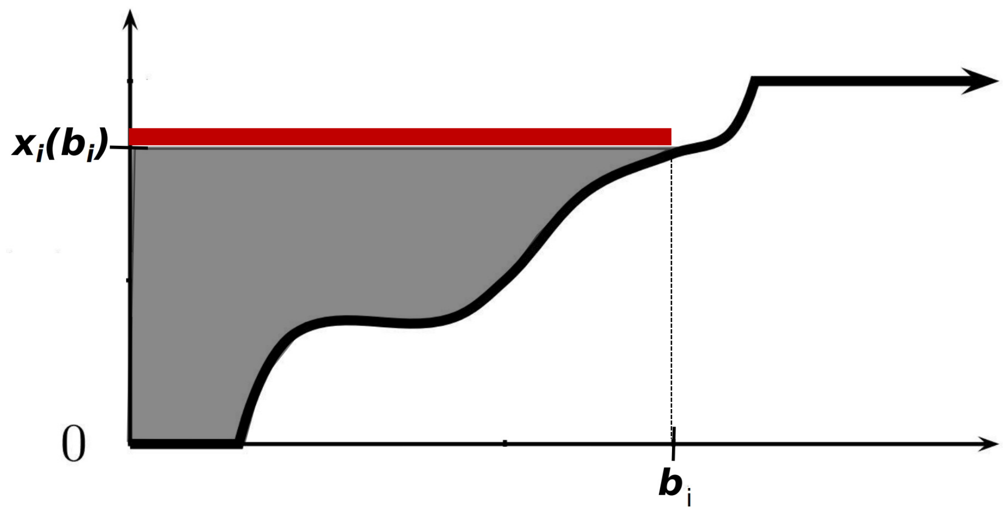

- The higher a player bids, the higher the probability of winning
- For each additional ϵ increase of winning probability, pay additionally at a rate equal to the current bid
- Proof omitted here

Corollaries of Myerson's Lemma

Corollaries.

1. Interim allocation uniquely determines interim payment
2. Expected revenue depends only on the allocation rule
3. Any two auctions with the same interim allocation rule at BNE have the same expected revenue at the same BNE

Step 2: Revenue as Virtual Welfare

- Define the virtual value of player i as a function of his value v_i:

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

Lemma. Consider any BIC mechanism M with interim allocation x and interim payment p, normalized to $p_i(0) = 0$. The expected revenue of M is equal to the expected virtual welfare served

$$\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}[\, \phi_i(v_i) x_i(v_i)]$$

- This is the expected virtual value of the winning bidder
- Proof is an application of Myerson's monotonicity lemma, plus algebraic calculations
- Recall the expected revenue is $\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}\, p_i(v_i)$

Step 2: Revenue as Virtual Welfare

- Define the virtual value of player i as a function of his value v_i:

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

Lemma. Consider any BIC mechanism M with interim allocation x and interim payment p, normalized to $p_i(0) = 0$. The expected revenue of M is equal to the expected virtual welfare served

$$\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}[\,\phi_i(v_i) x_i(v_i)]$$

- This is the expected virtual value of the winning bidder
- Proof is an application of Myerson's monotonicity lemma, plus algebraic calculations
- Recall the expected revenue is $\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}\, p_i(v_i)$

Proof

$$\mathbb{E}_{v_i \sim f_i}\,\overline{p_i}\,(v_i) = \int_{v_i} \left[v_i \cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\,db \right] f_i(v_i) dv_i$$

By Myerson's monotonicity lemma

Assumed bidder i bids truthfully

Proof

$$\mathbb{E}_{v_i \sim f_i}\, \overline{p_i}\,(v_i) = \int_{v_i} \left[v_i \cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\, db \right] f_i(v_i) dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{v_i} \int_{b=0}^{v_i} x_i(b)\, f_i(v_i) db\, dv_i$$

Rearrange terms

Proof

$$\mathbb{E}_{v_i \sim f_i}\, \overline{p_i}\,(v_i) = \int_{v_i} \left[v_i \cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\, db \right] f_i(v_i) dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{v_i} \int_{b=0}^{v_i} x_i(b)\, f_i(v_i) db\, dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{b} \int_{v_i \geq b} x_i(b)\, f_i(v_i) dv_i db$$

Exchange of integral variable order

Proof

$$\mathbb{E}_{v_i \sim f_i}\, \overline{p_i}\,(v_i) = \int_{v_i} \left[v_i \cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\, db \right] f_i(v_i) dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{v_i} \int_{b=0}^{v_i} x_i(b)\, f_i(v_i) db\, dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{b} \int_{v_i \geq b} x_i(b)\, f_i(v_i) dv_i db$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{b} x_i(b)(1 - F_i(b))\, db$$

Since $\int_{v_i \geq b} f_i(v_i) dv_i = 1 - F_i(b)$

Proof

$$\mathbb{E}_{v_i \sim f_i} \overline{p_i}(v_i) = \int_{v_i} \left[v_i \cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\, db \right] f_i(v_i) dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{v_i} \int_{b=0}^{v_i} x_i(b)\, f_i(v_i) db\, dv_i$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{b} \int_{v_i \geq b} x_i(b)\, f_i(v_i) dv_i db$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{b} x_i(b)(1 - F_i(b))\, db$$

$$= \int_{v_i} v_i \cdot x_i(v_i) f_i(v_i) dv_i - \int_{v_i} x_i(v_i)\big(1 - F_i(v_i)\big) dv_i$$

Proof

$$\mathbb{E}_{v_i\sim f_i}\,\overline{p_i}\,(v_i) = \int_{v_i}\left[v_i\cdot x_i(v_i) - \int_{b=0}^{v_i} x_i(b)\,db\right] f_i(v_i)dv_i$$

$$= \int_{v_i} v_i\cdot x_i(v_i)f_i(v_i)dv_i - \int_{v_i}\int_{b=0}^{v_i} x_i(b)\,f_i(v_i)db\,dv_i$$

$$= \int_{v_i} v_i\cdot x_i(v_i)f_i(v_i)dv_i - \int_{b}\int_{v_i\geq b} x_i(b)\,f_i(v_i)dv_i db$$

$$= \int_{v_i} v_i\cdot x_i(v_i)f_i(v_i)dv_i - \int_{b} x_i(b)(1-F_i(b))\,db$$

$$= \int_{v_i} v_i\cdot x_i(v_i)f_i(v_i)dv_i - \int_{v_i} x_i(v_i)\big(1-F_i(v_i)\big)dv_i$$

$$= \int_{v_i} x_i(v_i)\cdot\left[v_i f_i(v_i) - \big(1-F_i(v_i)\big)\right]dv_i$$

$$= \int_{v_i} x_i(v_i)\cdot f_i(v_i)\left[v_i - \frac{\big(1-F_i(v_i)\big)}{f_i(v_i)}\right]dv_i$$

$$= \mathbb{E}_{v_i\sim f_i}[\,\phi_i(v_i)x(v_i)]$$

Step 3: Deriving the Optimal Auction

➢Revenue of any BIC mechanism equals $\sum_{i=1}^{n} \mathbb{E}_{v_i \sim f_i}[\, \phi_i(v_i) x(v_i)]$

Q: how to extract the maximum revenue then?

1. Elicit buyer values $v_1, \cdots, v_n$ and calculate virtual values $\phi_i(v_i)$
2. If $\phi_i(v_i) < 0$ for all i, keep the item and no payments (why?)
3. Otherwise, allocate item to $i^* = \arg\max_{i \in [n]} \phi_i(v_i)$
4. How much to charge? Myerson's lemma says there is a unique interim payment
   - Charging minimum bid needed to win $\phi_i^{-1}\left(\max\left(\max_{j \neq i^*} \phi_j(v_j)\,, 0\right)\right)$ works.

The optimal auction

Step 3: Deriving the Optimal Auction

1. Elicit buyer values $v_1, \cdots, v_n$ and calculate virtual values $\phi_i(v_i)$
2. If $\phi_i(v_i) < 0$ for all i, keep the item and no payments (why?)
3. Otherwise, allocate item to $i^* = \arg\max_{i \in [n]} \phi_i(v_i)$, charge him the minimum bid needed to win $\phi_i^{-1}\left(\max\left(\max_{j \neq i^*} \phi_j(v_j), 0\right)\right)$; others pay 0

Observations.

- The allocation rule maximizes virtual welfare point-point, thus also maximizes expected virtual welfare
- By previous lemma, this is the maximum possible revenue
- Payment satisfies Myerson's lemma (check it)

Are we done?

A Wrinkle

- One more thing – Myerson lemma requires the interim allocation to be monotone
- When $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is monotone in v_i, allocation is monotone
- Fortunately, most natural distributions will lead to monotone VV function (e.g., Gaussian, uniform, exp, etc.)
  - Such a distribution is called regular

Conclusion. When values are drawn from regular distributions independently, previous auction (aka Myerson's optimal auction) is a revenue-optimal mechanism!

Can be extended to non-regular distributions via ironing, but won't cover

Thank You

Haifeng Xu
University of Chicago
haifengxu@uchicago.edu