Tutorial:
Bayesian Mechanism Design

Shuchi Chawla & Jason Hartline
U.Wisc-Madison Northwestern U.

June 5, 2011

Mechanism Design

Basic question: How should an economic system be designed so that selfish agent behavior leads to good outcomes?

Internet applications: file sharing, reputation systems, web search, web advertising, email, Internet auctions, congestion control, etc.

General theme: resource allocation

Overview

Part 1: Intro to Bayesian Mechanism Design

- Classical mechanisms: First-price auction, Vickrey auction, Myerson's auction

- Focus on single-item auction
- Objective 1: Social welfare
- Objective 2: Revenue
- Generalize beyond single-item setting

Part 2 (after break): Recent results in Algorithmic BMD

Problem: single-item auction

Given:

- One item for sale
- n agents/bidders with unknown private values $v_1, \ldots, v_n$
- Agents' objective: max utility = value obtained – price paid

Design goal:

- Protocol to solicit bids; choose winner and payment

Possible objectives:

- Maximize **social surplus**, i.e. value of the winner
- Maximize **seller's revenue** i.e. payment of the winner

Objective 1: Maximize social surplus

Example auctions

First-price auction

1. Solicit sealed bids
2. Highest bidder wins
3. Winner pays his bid

Second-price auction

1. Solicit sealed bids
2. Highest bidder wins
3. Winner pays second-highest bid

Vickrey auction

Example input: $b = (2,6,4,1)$

Questions:

- What are the equilibrium strategies?
- What is the equilibrium outcome?
- Which one has higher surplus?
- Which one has higher revenue?

Second-price auction: equilibrium analysis

Second-price auction

1. Solicit sealed bids 2. Highest bidder wins
3. Winner pays second-highest bid

How should agent *i* bid?

- Let $t_i = \max_{j \neq i} b_j$
- If $b_i \geq t_i$, i wins and pays t_i; otherwise loses.

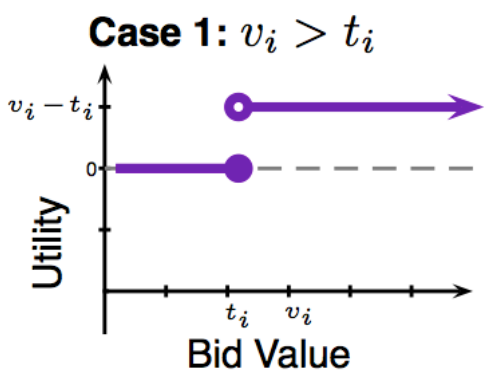

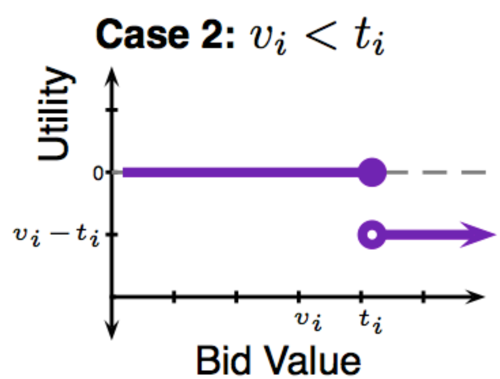

Result: Bidder i's dominant strategy is to bid $b_i = v_i$

Second-price auction: conclusion

Second-price auction

1. Solicit sealed bids 2. Highest bidder wins
3. Winner pays second-highest bid

Lemma: [Vickrey'61] Truthful bidding is a dominant strategy in the second-price auction

Corollary: Second-price auction maximizes social surplus, i.e. value of the winner.

First-price auction: equilibrium analysis

First-price auction

1. Solicit sealed bids
2. Highest bidder wins
3. Winner pays his bid

How would you bid?

Note: first-price auction has no dominant strategy equilibrium

Bayes-Nash equilibrium

Defn: the *common prior assumption*: bidders' values are drawn from a known distribution, i.e., $v_i \sim F_i$

Notation:

- $F_i(z) = \Pr[v_i \leq z]$ is the *cumulative distribution function*, (e.g. $F_i(z) = z$ for the uniform [0,1] distribution)
- $f_i(z) = \mathrm{d}F_i(z)/\mathrm{d}z$ is the *probability density function*, (e.g. $f_i(z) = 1$ for the uniform [0,1] distribution)

Defn: a *strategy* maps values to bids, i.e., $b_i = s_i(v_i)$

Defn: A strategy profile $(s_1, \ldots, s_n)$ is in *Bayes-Nash equilibrium* if for all i, $s_i(v_i)$ is a best response when others play $s_j(v_j)$ and $v_j \sim F_j$.

First-price auction: equilibrium analysis

Example: two bidders, values i.i.d. from U[0,1]

- Guess $s_i(z) = z/2$ is BNE and verify
- If agent 2 bids $b_2 \sim$ U[0,1/2], how should agent 1 bid?
- Compute agent 1's expected utility with bid b_1

$$\mathrm{E}[u_1] = (v_1 - b_1) \times \underbrace{\Pr[1 \text{ wins}]}_{\boxed{\begin{aligned}\Pr[b_1 > b_2] &= \Pr[b_1 > v_2/2] \\ &= \Pr[2b_1 > v_2] = F_2(2b_1) = 2b_1\end{aligned}}}$$

$$= (v_1 - b_1)2b_1$$

$$= 2(v_1 b_1 - b_1^2)$$

- To maximize, take derivative w.r.t. b_1 and set to zero; solve
- $b_1 = v_1/2$; guess is verified!

Conclusion: bidder with highest value wins, social surplus is maximized!

Surplus maximization conclusions

First-price auction

1. Solicit sealed bids
2. Highest bidder wins
3. Winner pays his bid

Second-price auction

1. Solicit sealed bids
2. Highest bidder wins
3. Winner pays second-highest bid

- Second-price auction maximizes surplus in DSE regardless of distribution
- First-price auction maximizes surplus in BNE for i.i.d. distributions

Surprising result: the auctions are optimal for any distribution

Objective 2: Maximize seller's revenue

• • •

Other objectives are similar

An example

Example: two bidders, values i.i.d. from U[0,1]

What is the revenue of the second-price auction?

- Draw values v_1, v_2 from the unit interval
- Sort values: $v_1 \geq v_2$
- Values divide the unit line equally
- E[revenue of 2^{nd} price auction] = E[v_2] = 1/3

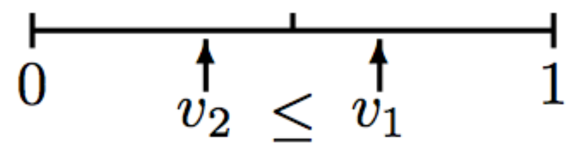

What is the revenue of the first-price auction?

- E[revenue of 1^{st} price auction] = E[b_1] = E[v_1]/2 = 1/3

Surprising result: both have the same expected revenue!

Can we get more?

Second-price auction with reserve price

Second-price auction with reserve price r
0. Place seller bid at r
1. Solicit sealed bids 2. Highest bidder wins
3. Winner pays second-highest bid

Lemma: Second-price auction with reserve *r* has a truthful DSE

What is the revenue of this auction?

Example: second-price with reserve

Example: two bidders, values i.i.d. from U[0,1]

What is the revenue of second-price with reserve ½?

- Draw values v_1, v_2 from unit interval
- Sort values: $v_1 \geq v_2$

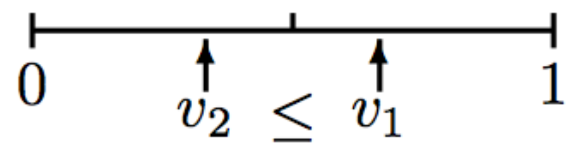

Case analysis	Probability	E[revenue]
$v_2 \leq v_1 < ½$	1/4	0

- E[Revenue of second-price with reserve ½]

$$= \frac{1}{4} \cdot 0 + \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{2}{3} \quad = \frac{5}{12} \quad > \frac{1}{3} = \text{E[Revenue of second-price]}$$

Can we do even better?

Characterizing Bayes-Nash equilibria

Notation

- $\boldsymbol{x}$ denotes allocation, x_i the allocation for agent i
- $\boldsymbol{x}(v)$ is the BNE allocation of mechanism on values v, i.e., the mechanism's outcome composed with agents' BNE strategies
- $v_{-i} = (v_1, \ldots, v_{i-1}, ?, v_{i+1}, \ldots, v_n)$
- $x_i(v_i) = \mathrm{E}_{v_{-i}}[x_i(v_i, v_{-i})]$

  is agent i's interim prob. of allocation when $v_{-i} \sim F_{-i}$

- Analogously define $\boldsymbol{p}$, $\boldsymbol{p}(v)$, $p_i(v_i)$ for payments

- Bidder i with value v_i mimicking strategy for value z has utility $u_i(v_i, z) = v_i x_i(z) - p_i(z)$

$$\boxed{\text{BNE} \Longrightarrow \text{for all } i,\ v_i, \text{ and } z,\ u_i(v_i, v_i) \geq u_i(v_i, z)}$$

Characterization of BNE

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$

   (Note: usually $p_i(0) = 0$.)

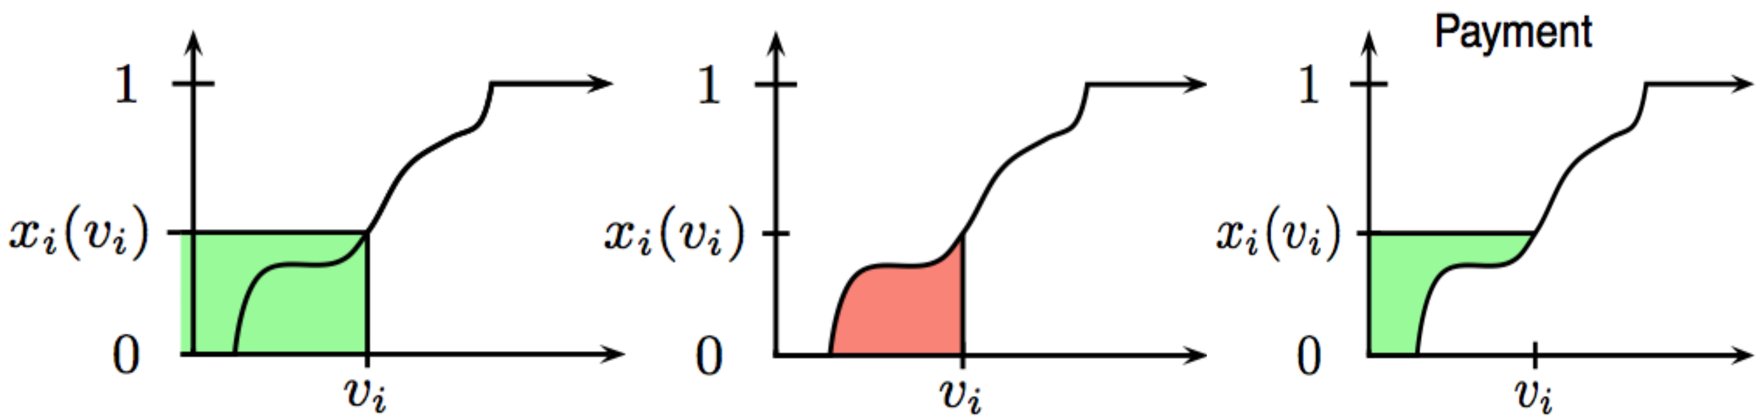

Characterization of BNE: proof outline

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$

   (Note: usually $p_i(0) = 0$.)

Proof approach:

1. BNE $\Longrightarrow$ M
2. BNE $\Longrightarrow$ PI
3. BNE $\Longleftarrow$ M & PI

BNE $\Longrightarrow$ M

Recall: BNE $\Longrightarrow u_i(v_i, v_i) \geq u_i(v_i, z)$ for all v_i and z

- Take $v_i = s$ and $z = t$ and vice versa:

$$sx_i(s) - p_i(s) \geq sx_i(t) - p_i(t)$$
$$tx_i(t) - p_i(t) \geq tx_i(s) - p_i(s)$$

- Adding and regrouping:

$$sx_i(s) + tx_i(t) \geq sx_i(t) + tx_i(s)$$
$$\Rightarrow (s - t)x_i(s) \geq (s - t)x_i(t)$$

- So x_i is monotone non-decreasing: $s > t \Rightarrow x_i(s) \geq x_i(t)$

Characterization of BNE: proof outline

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$$

   (Note: usually $p_i(0) = 0$.)

Proof approach:

1. BNE $\Longrightarrow$ M
2. BNE $\Longrightarrow$ PI
3. BNE $\Longleftarrow$ M & PI

BNE $\implies$ PI

Recall: BNE $\implies$ For all s and t:

$$sx_i(s) - p_i(s) \geq sx_i(t) - p_i(t)$$

$$tx_i(t) - p_i(t) \geq tx_i(s) - p_i(s)$$

- Rearranging:

$$t(x_i(s) - x_i(t)) \leq p_i(s) - p_i(t) \leq s(x_i(s) - x_i(t))$$

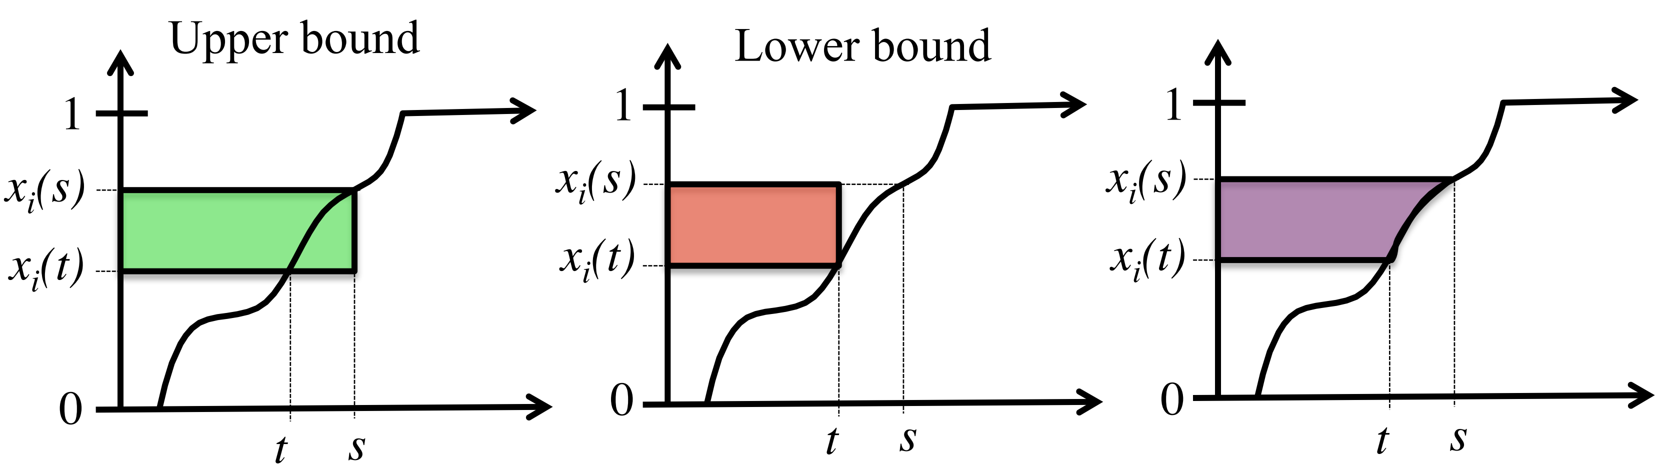

- Putting inequalities together for all pairs *s* and *t* implies PI

Characterization of BNE: proof outline

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$$
   (Note: usually $p_i(0) = 0$.)

Proof approach:

1. BNE $\Longrightarrow$ M
2. BNE $\Longrightarrow$ PI
3. BNE $\Longleftarrow$ M & PI

BNE $\Longleftarrow$ M & PI

Case 1: deviation from v_i to $z > v_i$

Claim: $u_i(v_i,v_i) \geq u_i(v_i,z)$

Recall: $u_i(v_i,v_i) = v_i x_i(v_i) - p_i(v_i)$

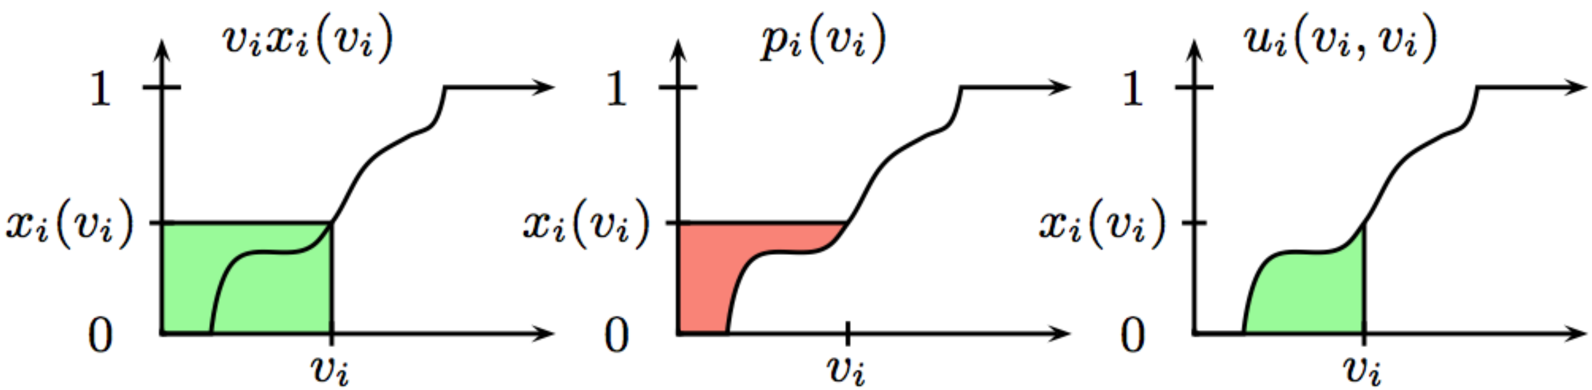

BNE $\Leftarrow$ M & PI

Case 1: deviation from v_i to $z > v_i$

Claim: $u_i(v_i,v_i) \geq u_i(v_i,z)$

Recall: $u_i(v_i,z) = v_i x_i(z) - p_i(z)$

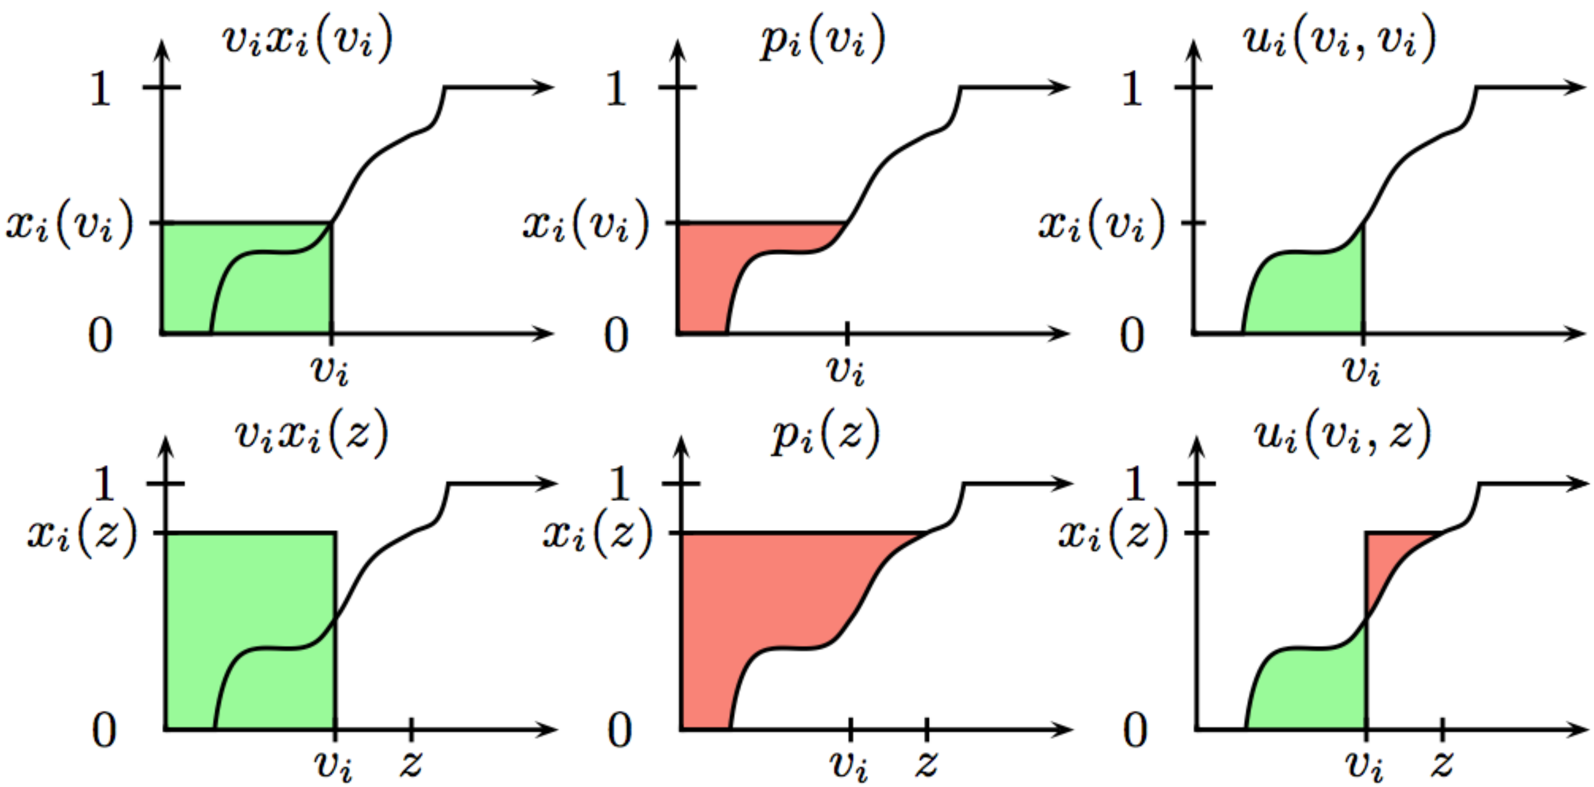

BNE $\Leftarrow$ M & PI

Case 2: deviation from v_i to $z < v_i$

Claim: $u_i(v_i,v_i) \geq u_i(v_i,z)$

Recall: $u_i(v_i,v_i) = v_i x_i(v_i) - p_i(v_i)$

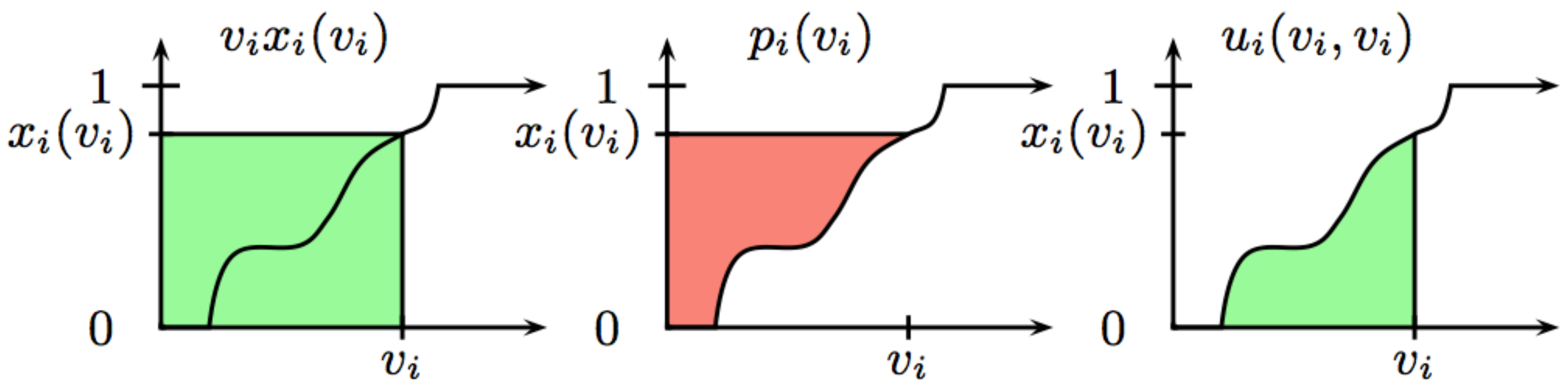

BNE $\Leftarrow$ M & PI

Case 2: deviation from v_i to $z < v_i$

Claim: $u_i(v_i,v_i) \geq u_i(v_i,z)$

Recall: $u_i(v_i,z) = v_i x_i(z) - p_i(z)$

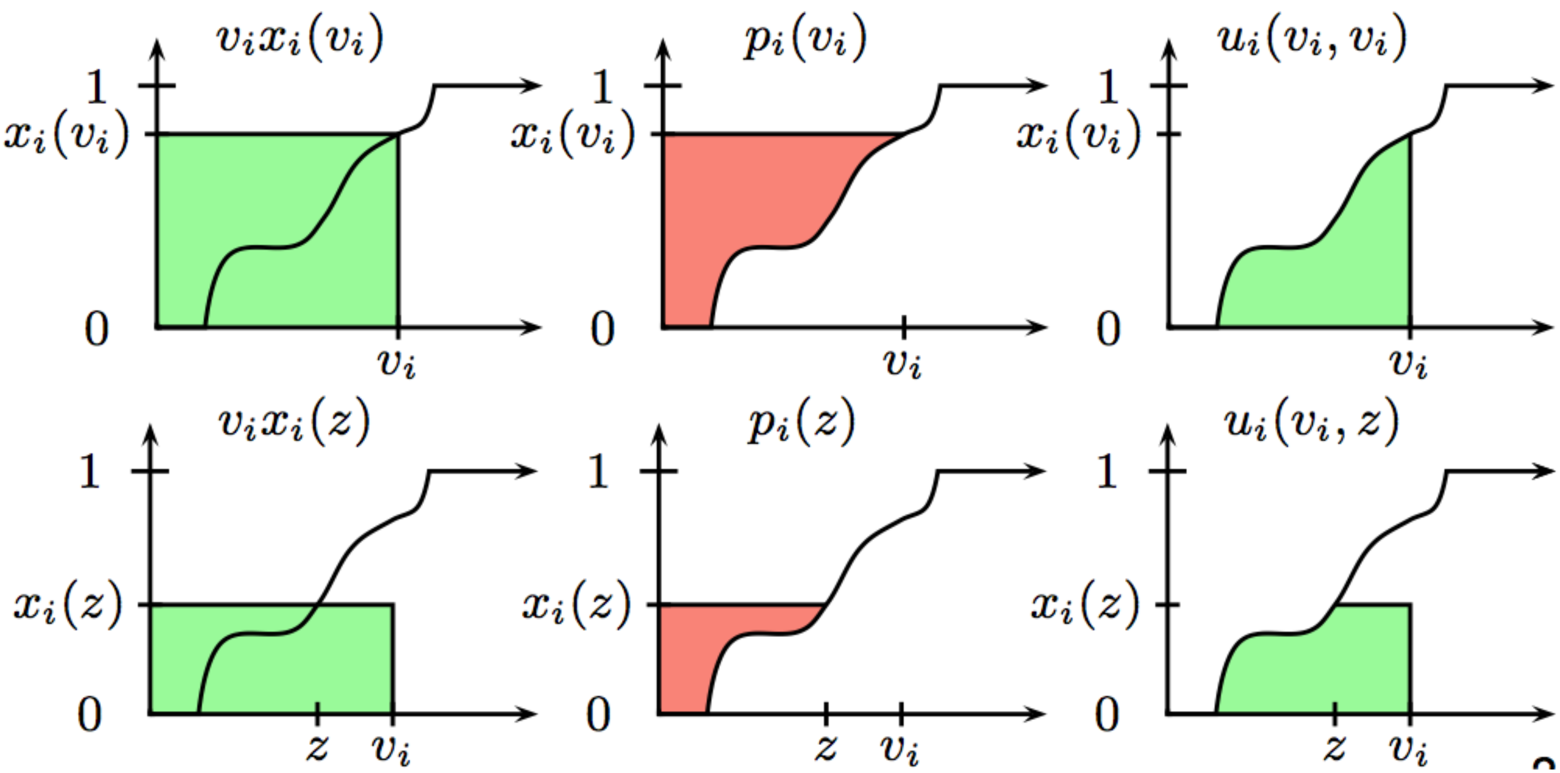

Characterization of BNE: implications

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$

   (Note: usually $p_i(0) = 0$.)

Implication: (Revenue Equivalence) Two auctions with the same outcome in BNE obtain the same expected revenue (e.g. first and second price auctions)

Implication: (strategy computation)

Characterization of BNE: implications

Thm: a mechanism and strategy profile are in BNE iff

1. Monotonicity (M): $x_i(v_i)$ is monotone non-decreasing in v_i
2. Payment identity (PI): $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz + p_i(0)$

   (Note: usually $p_i(0) = 0$.)

Implication: (strategy computation)

Example: two bidders, values i.i.d. from U[0,1]

- Expected payment of agent 1 at value v_1 in 2nd price auction
  $= \Pr[v_2 < v_1]\, E[v_2 \mid v_2 < v_1] = \Pr[v_2 < v_1]\, v_1/2$
- Expected payment of agent 1 at value v_1 in 1st price auction
  $= \Pr[v_2 < v_1]\, b_1(v_1)$

$\Longrightarrow$ In symmetric BNE, $b_1(v_1) = v_1/2$

Revisiting the revenue objective

Goal: find the auction that maximizes expected revenue

The Bayesian optimal auction

Objective: find monotone function $\boldsymbol{x}(v)$ to maximize $\mathrm{E}[\sum_i p_i(v_i)]$

Myerson's lemma: In BNE, $\mathrm{E}[\sum_i p_i(v_i)] = \mathrm{E}[\sum_i \phi_i(v_i)\, x_i(v_i)]$

where $\phi_i(v_i)$ is the virtual value function:

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

Proof sketch:

- Use characterization: $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,dz$
- Use definition of expectation: integrate payment x density
- Swap order of integration
- Simplify to get: $$\mathrm{E}\left[p_i(v_i)\right] = \mathrm{E}\left[\left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\right) x_i(v_i)\right]$$

The Bayesian optimal auction

Myerson's lemma: In BNE, $E[\sum_i p_i(v_i)] = E[\sum_i \phi_i(v_i)\, x_i(v_i)]$
where $\phi_i(v_i)$ is the virtual value function:

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

General approach for revenue maximization:

- Optimize revenue ignoring incentive constraints (i.e. monotonicity)
  - Winner is the agent with maximum virtual value
- Check to see if incentive constraints are satisfied
  - If $\phi_i(v_i)$ is monotone, then so is $x_i(v_i)$

Defn: A distribution F_i is *regular* if ϕ_i is monotone

Thm: [Myerson'81] If F is regular, the optimal auction is to allocate the item to the agent with the highest positive virtual value.

Myerson's mechanism: examples

Thm: [Myerson'81] If F is regular, the optimal auction is to allocate the item to the agent with the highest positive virtual value.

Example: n agents, i.i.d. regular values

- Virtual value functions are all identical: $\phi_i = \phi_j = \phi$ for all i, j
- Winner i satisfies $\phi(v_i) \geq \max_j(\phi(v_j),0)$
- That is, $v_i \geq \max_j(v_j, \phi^{-1}(0))$

- What is this auction?

  Second-price auction with reserve $\phi^{-1}(0)$!

Myerson's mechanism: examples

Thm: [Myerson'81] If F is regular, the optimal auction is to allocate the item to the agent with the highest positive virtual value.

Example: n agents, i.i.d. regular values

- Optimal auction: Second-price auction with reserve $\phi^{-1}(0)$!

Example: n agents, values i.i.d. from U[0,1]

- $F(v_i) = v_i; f(v_i) = 1$
- So, $\phi(v_i) = v_i - \dfrac{1 - F(v_i)}{f(v_i)} = v_i - \dfrac{1 - v_i}{1} = 2v_i - 1$
- Therefore, $\phi^{-1}(0) = 1/2$
- Optimal auction: Second-price auction with reserve ½!

Myerson's mechanism: non-regular case

Thm: [Myerson'81] If F is regular, the optimal auction is to allocate the item to the agent with the highest positive virtual value.

What if the distribution is non-regular?

- Convert virtual value to "ironed" virtual value
- "Ironed" virtual value is monotone non-decreasing
- Optimal mechanism: allocate item to the agent with the highest ironed virtual value breaking ties consistently

Note: Even with i.i.d. values, optimal mechanism is not necessarily second-price with reserve

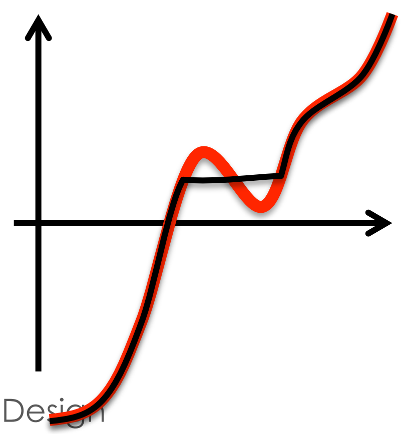

Beyond single-item auctions

…

Problem: general service providing

a.k.a. single-parameter MD

Given:

- A service to be provided
- n agents/bidders with unknown private values $v_1, \ldots, v_n$
- Agents' objective: max utility = value obtained – price paid
- General feasibility constraint on which subsets of agents can be simultaneously served

Design goal:

- Protocol to solicit bids; choose (feasible) winner(s) and payment(s)

Possible objectives:

- Maximize **social surplus**, i.e. sum of values of winners
- Maximize **seller's revenue** i.e. sum of payments of winners

General service providing: revenue

Thm: If F is regular, the optimal auction allocates to the feasible subset that maximizes "virtual surplus"

- Solicit bids, v
- Map bids to virtual bids $\phi_i(v_i)$
- Maximize over feasible sets S: $\sum_{i \in S} \phi_i(v_i)$
- Serve the set S
- Charge "critical prices"

Surprising result: the optimal auction is deterministic and dominant strategy truthful!

Observation: the theorem essentially gives a reduction from revenue maximization to surplus maximization

Part 1: conclusions

We saw:

- Characterization of BNE
- Revenue equivalence
- Optimal mechanism design via virtual values
- Reserve price based auctions are often but not always optimal

Issues:

- Optimal auctions are often too complicated; not seen in practice.
- Theory does not extend to "multi-dimensional" MD
- Theory requires knowledge of distribution
- Theory assumes we can solve optimization problems exactly

See part 2 for how to deal with these!