TRUTH JUSTICE ALGOS

Mechanism Design II: Revenue

Teachers: Ariel Procaccia and Alex Psomas (this time)

LAST TIME

- Second price auctions:
  - Maximize social welfare $\sum_i v_i x_i(\vec{v})$
    - Can we give buyers more utility?
  - DSIC
  - Polytime computable
- Myerson's lemma:
  - In a single parameter environment, an allocation rule x is implementable iff it is monotone. Furthermore, there is a unique payment that makes (x, p) DSIC.

LAST TIME

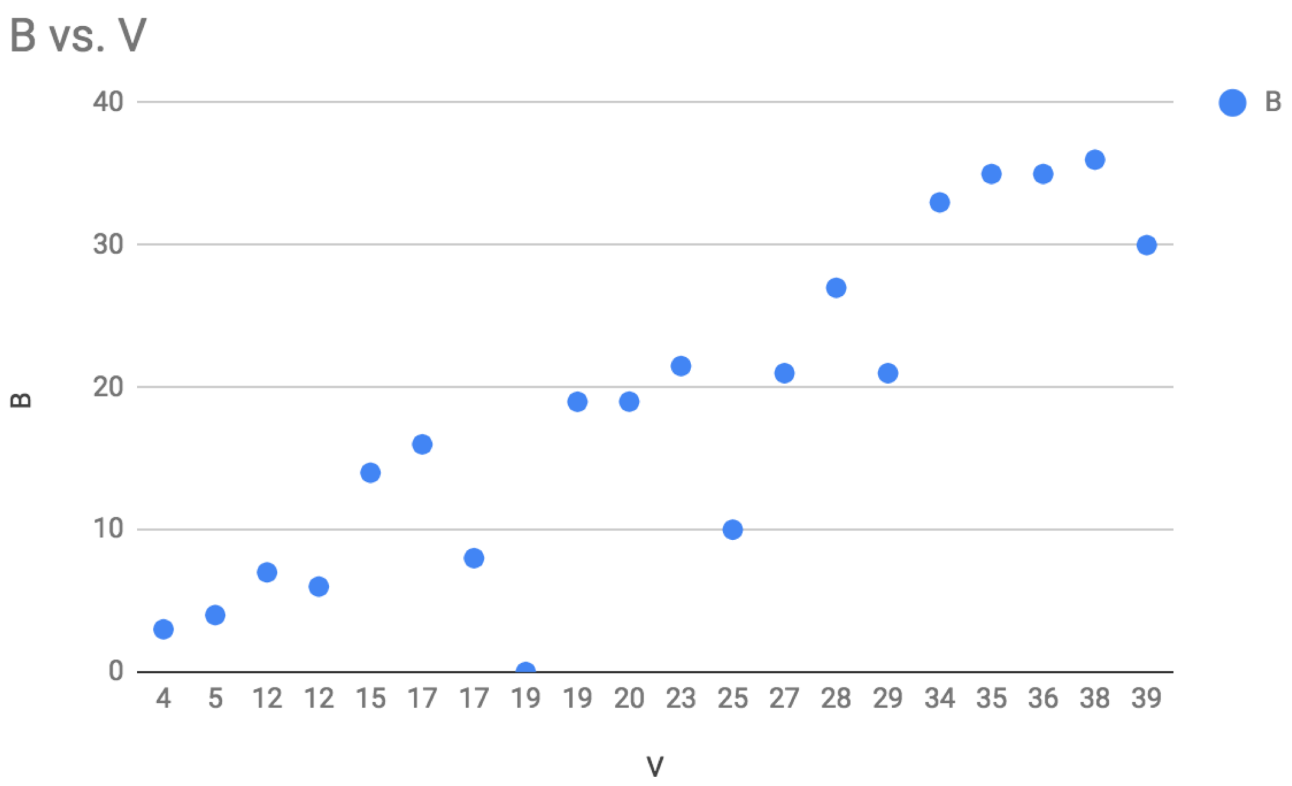
B vs. V
40
30
20
10
0
B
4 5 12 12 15 17 17 19 19 20 23 25 27 28 29 34 35 36 38 39
V
B

OBSERVATION: ALLOCATE TO THE BIDDER WITH THE HIGHEST VALUE

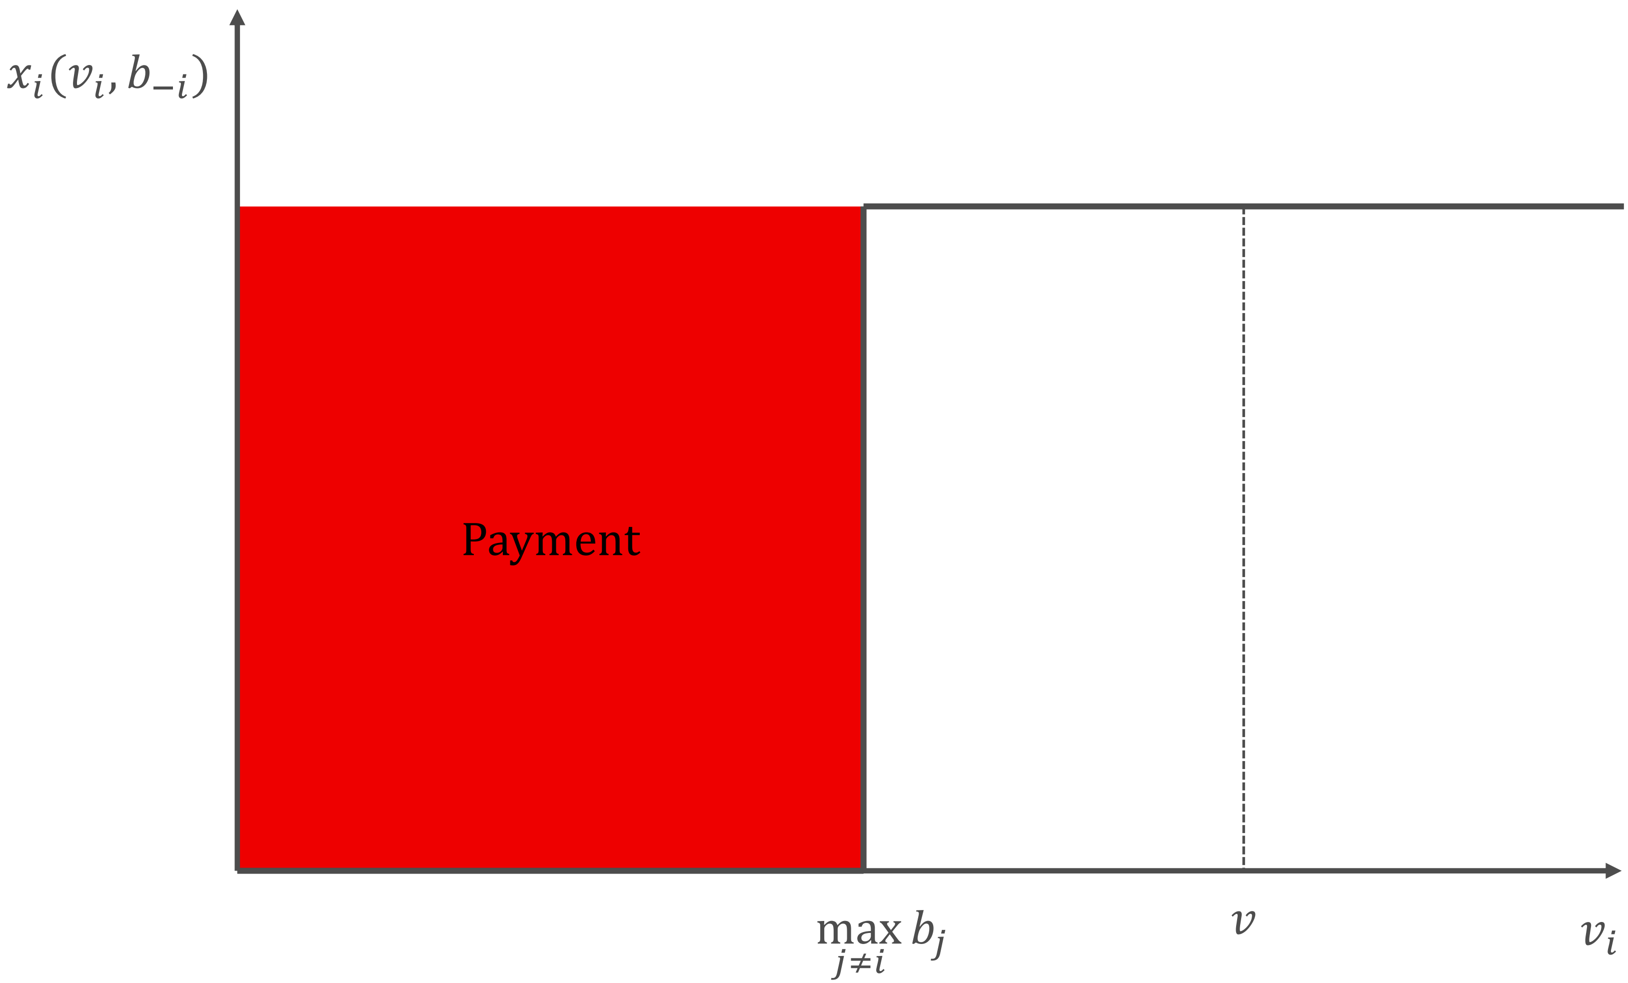

TODAY: REVENUE

- Why would we maximize social welfare?
- More reasonable to assume that sellers are trying to maximize revenue!
- For example, for $n = 1$ bidders, second price gives the item for free!
  - Pretty unreasonable…

ROGER MYERSON

MAXIMIZE REVENUE

- Focus on a single bidder, with private value v
- Make a take-it-or-leave-it offer
  - For a single bidder this is the only deterministic DSIC mechanism
- How much should we price the item at?
- If we magically knew v, we would set a price of v, but v is private…

EXAMPLE

Poll 1

How much would you price this boat?

EXAMPLE

- A price of r yields revenue r if $v \geq r$, and 0 otherwise
- A price of 10,000\$ is
  - Good if v is slightly higher than 10,000\$
  - Bad if v is a lot higher than 10,000\$
  - Horrible if v is 9,999\$

REVENUE

- Different auctions perform different on different inputs.
  - Contrast this with social welfare.
- We take a Bayesian approach!
- The private value v_i of bidder i is drawn from a **known** distribution D_i.
  - Today: distributions' support is $[0, v_{max}]$
- Goal: Maximize **expected** revenue over all DSIC and IR mechanisms.

WHY DSIC?

- Easy for participants to figure out what to bid
- The seller can predict what the bidders will do assuming only that they bid their dominant strategy
  - Pretty weak behavioral assumption
- Can you make more money with a non-DSIC mechanism??
  - Today: no!
  - Generally: yes!

REVELATION PRINCIPLE

- Optimize over the space of all DSIC mechanisms???
- That sounds super hard…
- It suffices to focus on **direct revelation mechanisms**!
  - You reveal your private information to the system.
  - As opposed to setting up a weird auction, where agents have dominant strategies

REVELATION PRINCIPLE

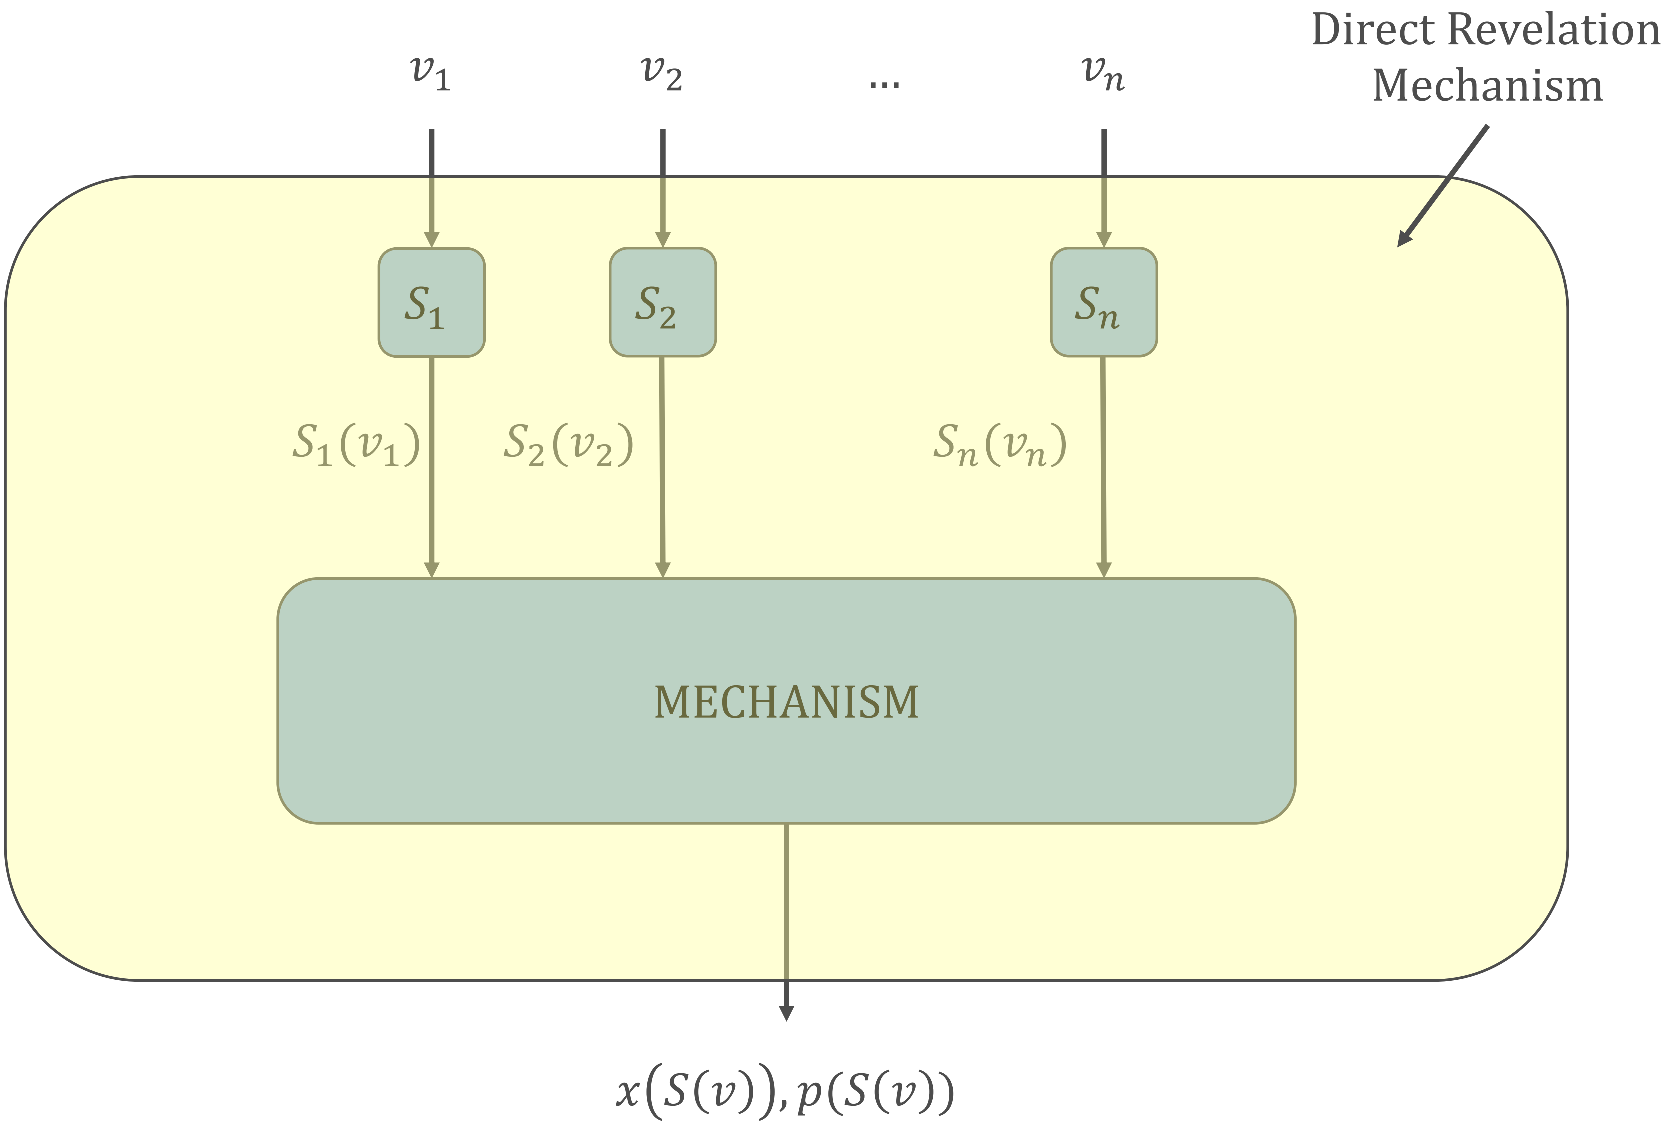

THE GAME

1. Seller is told distributions D_i for each buyer
2. Seller commits to a DSIC auction (x, p)
3. Nature draws v_i from D_i.
   - Today: independent D_is
4. Agent i learns v_i
5. Agent i submits bid b_i
6. Item is allocated according to $x(\vec{b})$, and payments are transferred according to $p(\vec{b})$

Goal:

- We take the seller's perspective.
- Design a DSIC and IR auction that maximizes **expected** revenue (expectation with respect to randomness in D and randomness in the auction)

SINGLE BUYER

- Expected revenue from setting a price of p

$$p \cdot \Pr[v \geq p] = p \cdot (1 - F(p))$$

- Say $D = U[0,1]$
- $Rev(p) = p \cdot \left(1 - F(p)\right) = p \cdot (1 - p)$
- $Rev'(p) = -2p + 1 = 0 \rightarrow p = \frac{1}{2}$
- Expected revenue $= \frac{1}{4}$
- This is optimal!
- What about two bidders??

TWO BIDDERS

- Say $D_1 = D_2 = D = U[0,1]$
- We could run a second price auction…
- What's the expected revenue?
- Observation: $E[Rev] = E[\min\{v_1, v_2\}]$
- $\Pr[\min\{v_1, v_2\} \geq x] = \Pr[v_1 \geq x \,\&\, v_2 \geq x]$

  $= \Pr[v_1 \geq x] \cdot \Pr[v_2 \geq \mathrm{x}]$

  $= (1 - x)^2$
- $E[min] = \int_{x=0}^{1} \Pr[min \geq x]\, dx = 1/3$

TWO BIDDERS

- $D_1 = D_2 = D = U[0,1]$
- Second price auction gives 1/3
- Can we do better?
- What if we never sell under ½?
  - Similar to what we did for one buyer.
- If highest bid $> \frac{1}{2}$: Highest bidder pays the maximum of ½ and the second highest bid
- If highest bid $< \frac{1}{2}$: No one gets the item
- Expected revenue of this auction is $\frac{5}{12} > \frac{1}{3}$
- Can we do better???

MYERSON

- The expected revenue of a DSIC auction (x, p) is equal to

$$\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} p_i(\vec{v})]$$

- For this results we assume **independent** buyer distributions.
- **Goal**: give a formula for the expected revenue that's easier to maximize!

MYERSON

- Step 0: Move things around:

$$\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} p_i(\vec{v})] = \sum_{i=1}^{n} \mathbb{E}_{v_{-i}}[\mathbb{E}_{v_i}[p_i(v_i, v_{-i})]]$$

MYERSON

- $\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} p_i(\vec{v})] = \sum_{i=1}^{n} \mathbb{E}_{v_{-i}}[\mathbb{E}_{v_i}[p_i(v_i, v_{-i})]]$
- Step 1: Apply Myerson's lemma

$$p_i(v, b_{-i}) = v x_i(v, b_{-i}) - \int_0^v x_i(z, b_{-i}) dz$$

- $\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] = \int_0^{v_{\max}} p_i(v_i, v_{-i}) f_i(v_i) dv_i$

$$= \int_0^{v_{\max}} \left(v_i x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(z, v_{-i}) dz \right) f_i(v_i) dv_i$$

$$= \int_0^{v_{max}} v_i \, x_i(v_i, v_{-i}) f_i(v_i) dv_i - \int_0^{v_{max}} \int_0^{v_i} x_i(z, v_{-i}) \, f_i(v_i) dz dv_i$$

MYERSON

$$\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] = \int_0^{v_{max}} v_i\, x_i(v_i, v_{-i}) f_i(v_i) dv_i - \int_0^{v_{max}} \int_0^{v_i} x_i(z, v_{-i})\, f_i(v_i) dz dv_i$$

- Step 2: Change order of integration
- $\int_0^{v_{max}} \int_0^{v_i} x_i(z, v_{-i})\, f_i(v_i) dz dv_i$

$$= \int_0^{v_{max}} x_i(z, v_{-i}) \int_z^{v_{max}} f_i(v_i) dv_i dz$$

$$= \int_0^{v_{max}} x_i(z, v_{-i})(1 - F_i(z)) dz$$

MYERSON

$$\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] = \int_0^{v_{max}} v_i\, x_i(v_i, v_{-i}) f_i(v_i) dv_i - \int_0^{v_{max}} x_i(v_i, v_{-i})(1 - F_i(v_i)) dv_i$$

- Step 3: Combine
- $\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] =$
$$\int_0^{v_{max}} f_i(v_i) x_i(v_i, v_{-i}) \left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right) dv_i$$

MYERSON

- $\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] =$
$$\int_0^{v_{max}} f_i(v_i) x_i(v_i, v_{-i}) \left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right) dv_i$$
- Step 4: A definition:

The **virtual value** of bidder i is

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

MYERSON

$$\mathbb{E}_{v_i}[p_i(v_i, v_{-i})] = \mathbb{E}_{v_i}[x_i(v_i, v_{-i}) \cdot \phi_i(v_i)]$$

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

- Step 5: Plug everything back:

$$\mathbb{E}_{\vec{v}\sim D}\left[\sum_{i=1}^{n} p_i(\vec{v})\right] = \sum_{i=1}^{n} \mathbb{E}_{v_{-i}}\left[\mathbb{E}_{v_i}[p_i(v_i, v_{-i})]\right]$$

$$= \mathbb{E}_{\vec{v}\sim D}\left[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})\right]$$

MYERSON

- $\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} p_i(\vec{v})] =$
$$\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})]$$
- Ok, let's parse this…
- Maximizing expected revenue is the same as maximizing the expected virtual welfare!
- We (kind of)already know how to solve that!
- **Second price auction** (but in virtual value space).

MYERSON

- Old problem:

$$\max \mathbb{E}_{\vec{v}\sim D}[\textstyle\sum_{i=1}^{n} p_i(\vec{v})]$$

Subject to

$$v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i x_i(v', v_{-i}) - p_i(v', v_{-i})$$

$$v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq 0$$

$$\sum_i x_i(\vec{v}) \leq 1$$

- New problem:

$$\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})]$$

Subject to

$$\sum_i x_i(\vec{v}) \leq 1$$

MYERSON

- Maximize $\mathbb{E}_{\vec{v}\sim D}\left[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})\right]$
- We can maximize this pointwise!

MYERSON

- Example: $n = 2$, D_1 and D_2 have support size $\{0,1\}$
- Maximize

$$\Pr[v_1 = 0, v_2 = 0](\phi_1(0)x_1(0,0) + \phi_2(0)x_2(0,0)) +$$
$$\Pr[v_1 = 0, v_2 = 1](\phi_1(0)x_1(0,1) + \phi_2(1)x_2(0,1)) +$$
$$\Pr[v_1 = 1, v_2 = 0](\phi_1(1)x_1(1,0) + \phi_2(0)x_2(1,0)) +$$
$$\Pr[v_1 = 1, v_2 = 1](\phi_1(1)x_1(1,1) + \phi_2(1)x_2(1,1))$$

- Subject to

$$x_1(i,j) + x_2(i,j) \leq 1, for\ all\ i,j \in \{0,1\}$$

MYERSON

- Maximize $\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})]$
- We can maximize this pointwise!
- Who gets the item?
  - **Highest virtual value!**
- How much do they pay?
  - Second highest virtual value??
  - **The value they would have to bid in order to lose!**

Kind of…

POLL

- Maximize $\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})]$
- $\phi_1(v_1) = v_1 - 1$
- $\phi_2(v_2) = v_2 - 1$
- $v_1 = 1/2$
- $v_2 = 1/4$

Poll 2

Who gets the item?

1. 1
2. 2
3. Half, half
4. Neither

POLL

- Maximize $\mathbb{E}_{\vec{v}\sim D}[\sum_{i=1}^{n} \phi_i(v_i) \cdot x_i(v_i, v_{-i})]$
- $\phi_1(v_1) = 2v_1 - 1$
- $\phi_2(v_2) = v_2 - 1$
- $v_1 = 1$
- $v_2 = 1/4$

Poll 3

Who gets the item? How much do they pay?

1. 1, -3/4
2. 1, 1/2
3. 1, 0
4. 1, 1/4

MYERSON

- Allocate to the agent with the highest virtual value (if it’s non-negative).
- **No!** The allocation rule might not be monotone!
  - $\phi_i(v)$ might decrease as v increases
- Myerson provided a solution to this: “iron” the virtual value function.
  - We won’t cover this.

MYERSON

- ***Definition***: A distribution with cdf F and pdf f is called **regular** if $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing
  - If $\frac{1-F(v)}{f(v)}$ is monotone non-increasing we say that the distribution has **monotone hazard rate** (MHR).
- Most distributions you know are regular (and MHR): uniform, exponential, Normal, Gamma, etc etc.
- Intuitively, regular = small tail

MYERSON: REGULAR DISTRIBUTIONS

- Give the item to the agent with the highest virtual value, or no one if all virtual values are negative.
  - Good news: monotone allocation rule
  - Weird news: Highest virtual value ≠ highest value!
- $\phi_1(v_1) = 2v_1 - 1, \phi_2(v_2) = 2v_2 - 100$
  - $v_1 = 0.6, v_2 = 50 \rightarrow$ Agent 1 wins!

MYERSON: REGULAR DISTRIBUTIONS

- Give the item to the agent with the highest virtual value, or no one if all virtual values are negative.
- If the item was given to agent i
  - Let j be the agent with the second highest virtual value
  - If $\phi_j(v_j) < 0$, i pays $\phi_i^{-1}(0)$
  - If $\phi_j(v_j) \geq 0$, i pays $\phi_i^{-1}\left(\phi_j(v_j)\right)$
- Different way to think about it:
  - Seller inserts her own bids (in v.v. space) $\phi_1^{-1}(0), \phi_2^{-1}(0), \ldots$

MYERSON: IDENTICAL REGULAR DISTRIBUTIONS

- Actually simple if all agents have the same distribution $D = D_i, \forall i$
- Give the item to the agent with the highest virtual value, or no one if all virtual values are negative.
  - Highest virtual value = Highest value
  - Rephrase: Give the item to the agent with the highest value, if her virtual value is non-negative.
- If the item was given to agent i, she pays the maximum of the second highest bid and $\phi^{-1}(0)$
- In other words, the optimal auction is a second price auction with a reserve of $\phi^{-1}(0)$
  - Does this look familiar?
  - Precisely the E-Bay format!

EXAMPLE

- 2 agents, $D_1 = D_2 = D = U[0,1]$
- $\phi(v) = v - \frac{1-F(v)}{f(v)} = v - \frac{1-v}{1} = 2v - 1$
- Allocation rule: give it to the person with the highest virtual value $\phi(v_i)$, if its non-negative
- Aka, give it to the person with the highest value v_i, if its at least ½
- Charge max{½, $other\ bid$}

SUMMARY

- Single parameter environments
  - Second price auctions
  - Myerson’s lemma
  - Myerson’s optimal auction