COMP/MATH 553 Algorithmic Game Theory
Lecture 10: Revenue Maximization in Multi-Dimensional Settings

Oct 06, 2014

Yang Cai

An overview of today's class

- *Unit-Demand Pricing (cont'd)*
- *Multi-bidder Multi-item Setting*
- *Basic LP formulation*

Two Scenarios

❖ **(a) UPP**

- ***One unit-demand bidder***
- ***n items***
- Bidder's value for the i-th item $\boldsymbol{v_i}$ is drawn independently from $\boldsymbol{F_i}$

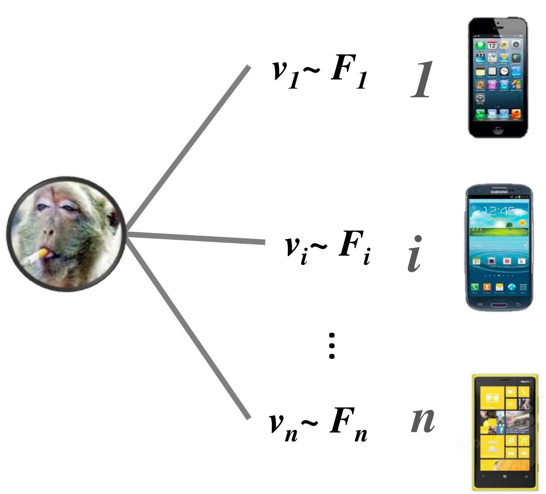

❖ **(b) Auction**

- ***n bidders***
- ***One item***
- Bidder I's value for the item $\boldsymbol{v_i}$ is drawn independently from $\boldsymbol{F_i}$

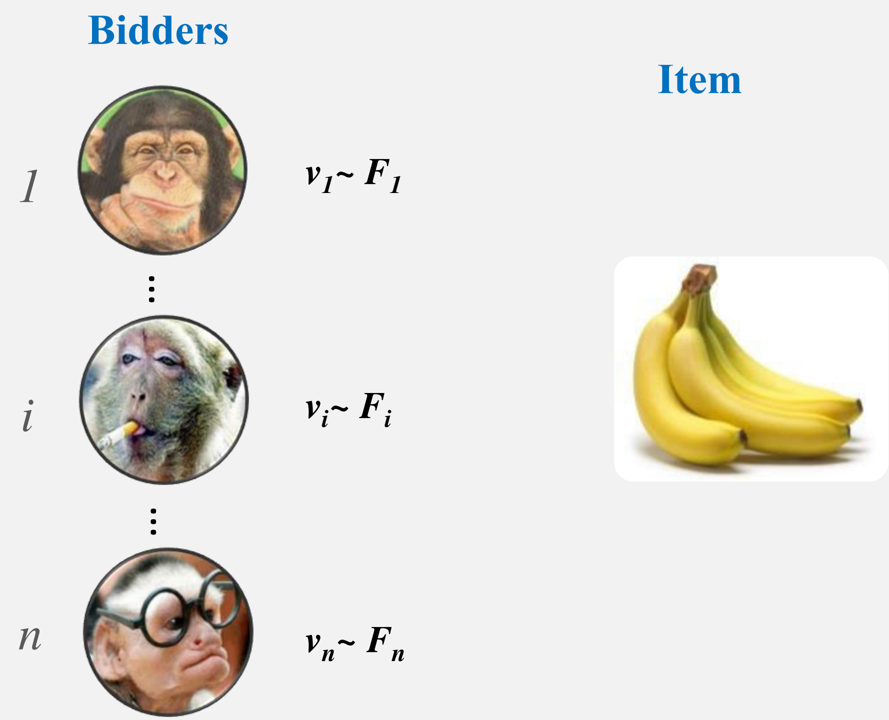

Benchmark

Lemma 1: The optimal revenue achievable in scenario (a) is always less than the optimal revenue achievable in scenario (b).

- Remark: This gives a natural benchmark for the revenue in (a).

A nearly-optimal auction (Lecture 6)

- In a single-item auction, the optimal expected revenue

$$\mathbf{E}_{v\sim F}[\max \Sigma_i x_i(v)\, \varphi_i(v_i)] = \mathbf{E}_{v\sim F}[\max_i \varphi_i(v_i)^+]$$

- Remember the following mechanism **RM** we learned in Lecture 6.
  1. Choose t such that $\Pr[\max_i \varphi_i(v_i)^+ \geq t] = ½$.
  2. Set a reserve price $r_i = \varphi_i^{-1}(t)$ for each bidder i with the t defined above.
  3. Give the item to the highest bidder that meets her reserve price (if any).
  4. Charge the payments according to Myerson's Lemma.

- By prophet inequality:

$$\textbf{ARev(RM)} = \mathbf{E}_{v\sim F}[\Sigma_i x_i(v)\, \varphi_i(v_i)] \geq ½\, \mathbf{E}_{v\sim F}[\max_i \varphi_i(v_i)^+] = ½\, \textbf{ARev(Myerson)}$$

- Let's use the revenue of RM as the benchmark.

Inherent loss of this approach

- Relaxing the benchmark to be Myerson’s revenue in (b)
- This step might lose a constant factor already.
- To get the real optimum, a different approach is needed.

Optimal Multidimensional Pricing

F_i is a **Monotone Hazard Rate** (**MHR**) distribution.

*** MHR Definition:**
f(x)/(1-F(x)) is non-decreasing.

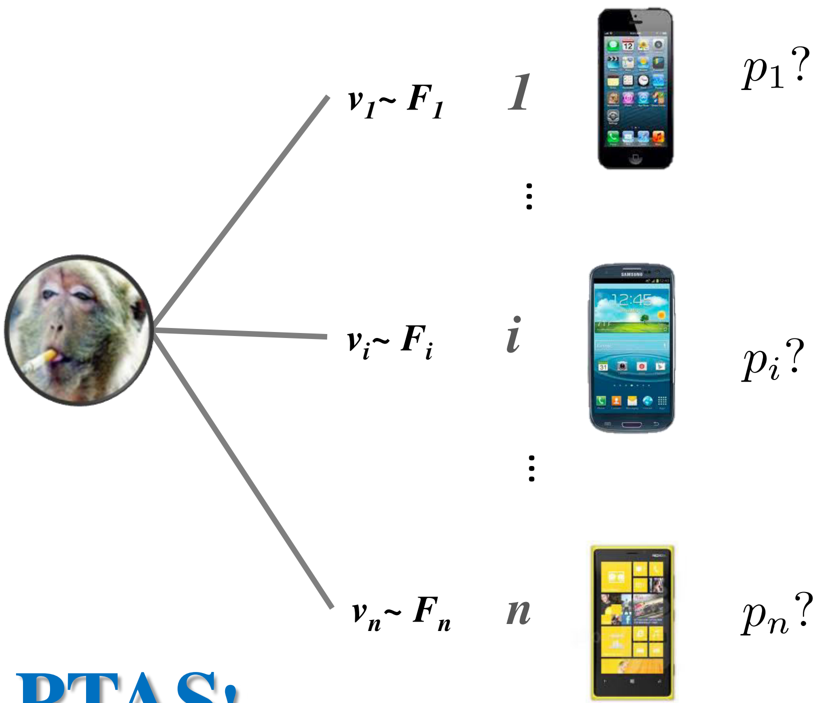

- Only constant factor appx are known [CHK '07, CHMS '10].
- [Cai-Daskalakis '11] **There is a PTAS!**
- PTAS: Polynomial-Time Approximation Scheme — for every constant ε in [0,1], there is a polynomial time algorithm that achieves (1- ε) fraction of the optimum (for maximization problems). The running time is required to be polynomial for every fixed ε, but could be different for different ε. For example, the running time could be $O(n^{1/\varepsilon})$

Extreme Value Theorem (MHR)

[Cai-Daskalakis '11]

Let $X_1, \ldots, X_n$ be independent (but not necessarily identically distributed) **MHR** random variables, Let $X = max_i X_i$. Then there exists ***anchoring point*** β such that:

COROLLARY:

$(1\text{-}\varepsilon)$ OPT is extracted from values in $(\varepsilon\,\beta,\ 1/\varepsilon \log 1/\varepsilon\,\beta)$.

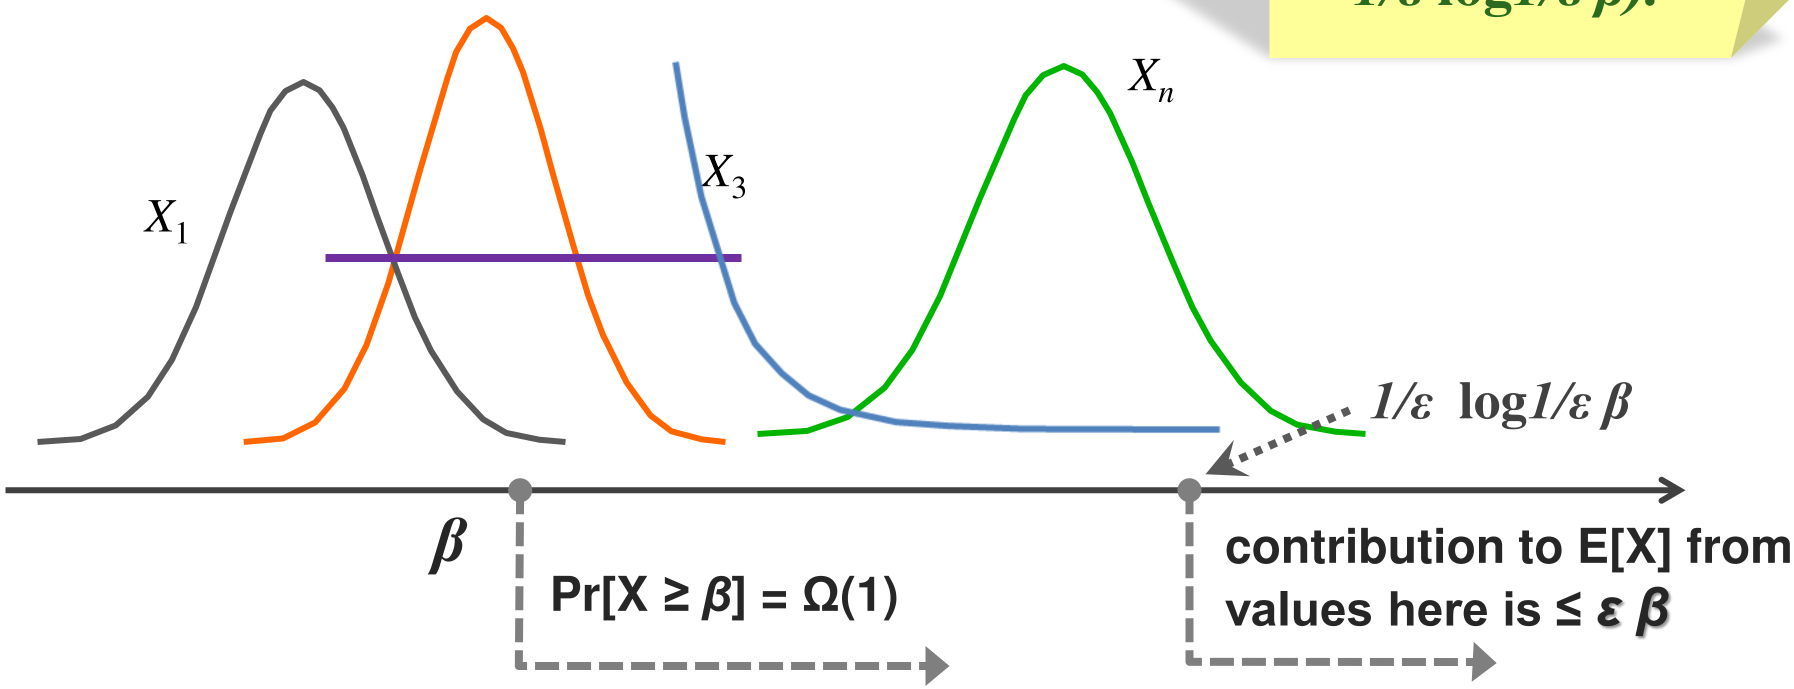

What if the items are i.i.d.?

- Say you know for each item there are only two prices 1 and 2, you can use.

- How many possible prices vectors are there?
  - 2^n
  - Do you really need to search over all of them?

- Only need to check O(n) different price vectors.

What if the items are i.i.d.?

- When you know you can use only c different prices on each item

- Only need to check $O(n^{c-1})$ different price vectors, when the distributions are i.i.d.

- Our theorem says you only need to consider poly($1/\varepsilon$) many different prices, so that gives you a PTAS for the i.i.d. case.

- When the distributions are not i.i.d., we need to use a more sophisticated Dynamic Programming algorithm to find the optimal price vector. But having only a constant number of prices is still crucial here.

Multi-item Multi-bidder Settings

Multi-item Multi-bidder Setting

- Remember the challenges. The optimal mechanism could have strange structure and uses randomization.
- Closed form solution (like Myerson's auction) seem impossible, even for a single bidder.
- More powerful machinery is required.
- Turn to ***Linear Programming*** for help.

Multi-item Multi-bidder Auctions: Set-up

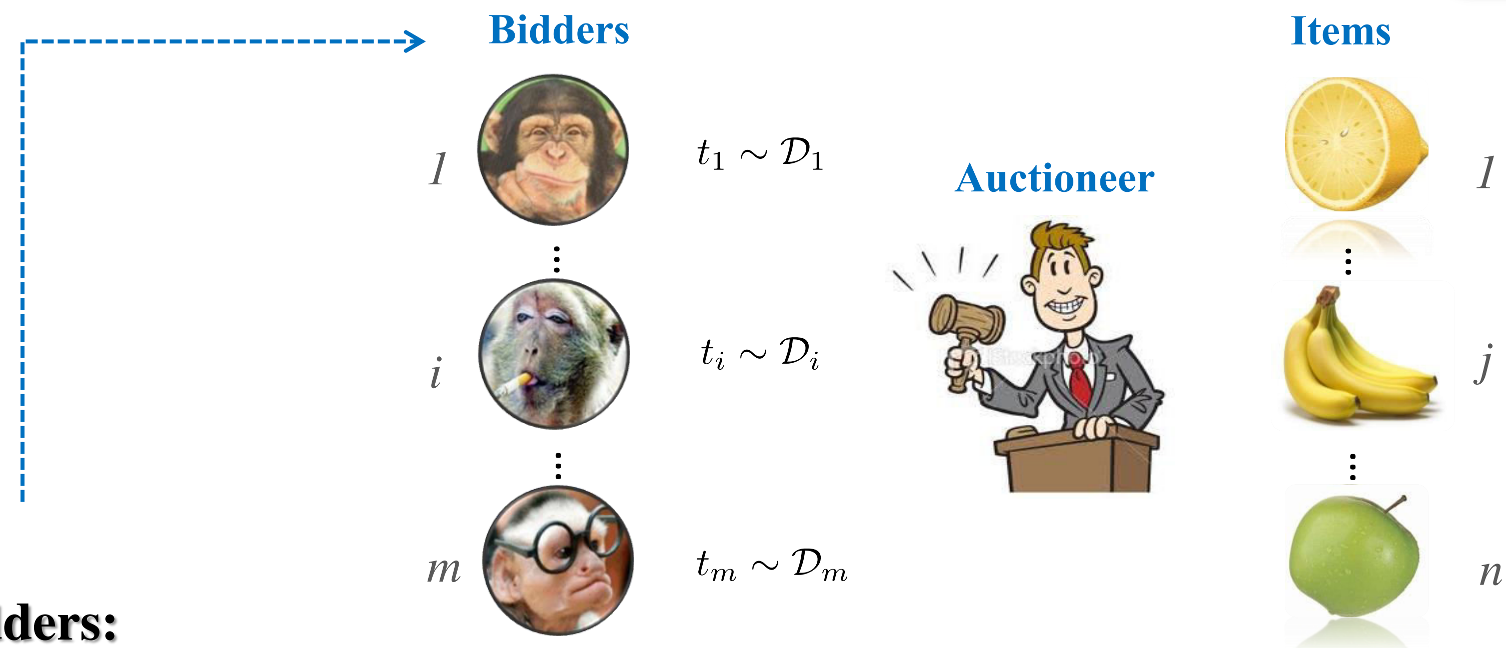

Bidders:

- have values on "items" and bundles of "items".
- ***Valuation*** aka ***type*** $t_i \in T_i$ encodes that information.
- **Common Prior:** Each t_i is sampled independently from $\mathcal{D}_i$.
  - Every bidder and the auctioneer knows $\mathcal{D}$
- **Additive:** Values for bundles of items = sum of values for each item.
  - **From now on,** $t_i = (v_{i1}, \ldots, v_{in})$.

A few remarks on the setting

- T_i is a subset of R^n
- Since we are designing algorithms, assume T_i is a discrete set.
- We know $\mathbf{Pr}[t_i = v]$ for all v in T_i and $\Sigma_v \mathbf{Pr}[t_i = v] = 1$.

Multi-item Multi-bidder Auctions: Execution

Each Bidder:

- **Uses as input:** the auction, own type, distributions about other bidders' types;
- Bids;

Goal: Optimize own utility (= expected value minus expected price).

Auctioneer:

- Designs auction, specifying allocation and payment rules;
- Asks bidders to bid;
- Implements the allocation and payment rule specified by the auction;

Goal: Find an auction that:

1) Encourages bidders to bid **truthfully** (w.l.o.g.)
2) Maximizes revenue, subject to 1)

LP Formulation

Single Bidder Case

- What are the decision variables?

- An auction is simply an allocation rule and a payment rule.

- Let's set the decision variables accordingly.

- Allocation rule: for each j in $[m]$, v in T, there is a variable $x_j(v)$: the probability that the buyer receives item j when his report is v.
  - if the mechanism is item pricing, and has price p_j for item j, then $x_j(v)=1$ if $v_j \geq p_j$ and 0 otherwise.
  - if the mechanism is grand bundling with price r. Then for all j, $x_j(v)=1$ if $\Sigma_j v_j \geq r$, otherwise all $x_j(v)=0$.
  - For deterministic mechanisms, $x_j(v)$ is either 0 or 1. But to include randomized mechanisms, we should allow $x_j(v)$ to be fractional.

Single Bidder Case

- Payment rule: for each $\boldsymbol{v}$ in $\boldsymbol{T}$, there is a variable $\boldsymbol{p(v)}$: the payment when the bid is $\boldsymbol{v}$.
- Objective function: $\mathbf{max}\ \Sigma_v\ \mathbf{Pr}[t = v]\ p(v)$
- Linear in the variables, since $\mathbf{Pr}[t = v]$ are constants (part of our input).
- Constraints:
  - incentive compatibility: $\Sigma_j v_j x_j(v) - p(v) \geq \Sigma_j v_j x_j(v') - p(v')$ for all v and v' in T
  - individual rationality (non-negative utility): $\Sigma_j v_j x_j(v) - p(v) \geq 0$ for all v in T
  - feasibility: $0 \leq x_j(v) \leq 1$ for all j in $[m]$ and v in T

Single Bidder Case

- We have a LP, we can solve it. But now what?
- ***What is the mechanism?***
- In this case, it's straightforward. Let x* and p* be the optimal solution of our LP.
- Then when the bid is $\boldsymbol{v}$, give the buyer item $\boldsymbol{j}$ with prob. $\boldsymbol{x_j(v)}$ and charge him $\boldsymbol{p(v)}$.
- This mechanism is feasible, incentive compatible and individual rational!
- So the buyer will bid truthfully, and thus the expected revenue of the mechanism is the same as the solution of our LP!