COMP/MATH 553 Algorithmic Game Theory
Lecture 18: Spectrum Auctions and Revenue Maximization in Multi-item Settings

Nov 03, 2016

Yang Cai

Menu

- Case Study: Spectrum Auctions
- Revenue Maximization in Multi-item Settings
- Upper Bound of the Optimal Revenue

Sequential Single-Item Auctions

- Run some single-item auction (e.g. first-price/second-price auction) sequentially, one item at a time.
- Difficult to play/predict bidder behavior
- Example: Suppose that *k* identical copies are sold to unit-demand bidders.
  - VCG would give each of the top k bidders a copy of the item and charge them the (k+1)-th highest bid.
  - What if we run sequential second-price auctions?
    - Easy to see that truthful bidding is not a dominant strategy, as if everyone else is bidding truthfully, I should expect prices to drop
    - Bidders will try to shade their bids, but how?
    - Outcome is unpredictable.
- Moving to more general settings only exacerbates issue.

Simultaneous Single-Item Auctions

- Run some single-item auction (e.g. first-price/second-price auction) simultaneously for all items.
- Bidders submit one bid per item.

- Issues for bidders:
  - Bidding on all items aggressively, may win too many items and over-pay (if, e.g., the bidder only has value for a few items)
  - Bidding on items conservatively may not win enough items

- What to do?
  - Difficulty in bidding and coordinating gives low welfare and revenue.

Simultaneous Single-Item Auctions

- In 1990, the New Zealand government auctioned off essentially identical licenses for television broadcasting using simultaneous (sealed-bid) Vickrey auctions.
- The revenue was only $36 million, a small fraction of the projected $250 million.
- For one license, the highest bid was $100,000 while the second-highest bid (and selling price) was $6! For another, the highest bid was $7 million and the second-highest bid was $5,000.
- Even worse: the top bids were made public so everyone could see how much money was left on the table.
- They later switched to first-price auctions. Similar problems remain (but it is less embarrassing).

Simultaneous Single-Item Auctions

- How to analyze theoretically?
- Auction is not direct, has no dominant strategy equilibrium.
- Hence need to make some further modeling assumptions, resort to some equilibrium concept.
- E.g. assume a *complete information setting*: bidders know each other's valuations (but auctioneer does not)
- E.g.2 assume *Bayesian incomplete information setting*: bidders' valuations are drawn from distributions known to every other bidder and the auctioneer, but each bidder's realized valuation is private

Theorem [Feldman-Fu-Gravin-Lucier'13]: If bidders' valuations are subadditive, then the social welfare achieved at a mixed Nash equilibrium (under complete information), or a Bayesian Nash equilibrium (under incomplete information) of the simultaneous $1^{st}/2^{nd}$ price auction is within a factor of 2 or 4 of the optimal social welfare.

Theorem [Cai-Papadimitriou'14]: Finding a Bayesian Nash equilibrium in a Simultaneous Single-Item Auction is highly intractable.

Simultaneous Ascending Auctions (SAAs)

- Over the last 20 years, ***simultaneous ascending auctions*** (SAAs) form the basis of most spectrum auctions.

- Conceptually, comprise several single-item English auctions running in parallel.

- In every round, each bidder places a new bid on any subset of items that she wants, subject to an *activity rule* and some constraints on the bids.

- Essentially the activity rule says: the number of items you bid on should decrease over time as prices rise.

Simultaneous Ascending Auctions (SAAs)

- Big advantage: *price discovery.*
- This allows bidders to do mid-course correction.
- Another advantage: value discovery.
- Finding out valuations might be expensive. Only need to determine the value on a need-to-know basis.

Simultaneous Ascending Auctions (SAAs)

- Poorly designed auctions still have issues.
- E.g. in 1999 the German government auctioned 10 blocks of cell-phone spectrum
- 10 simultaneous ascending auctions, with the rule that each new bid on a license must be at least 10% larger than previous bid
- Bidders: T-Mobile, Mannesman
- Mannesman first bid: 20 million Deutsche marks on blocks 1-5 and 18.18 on blocks 6-10
- Interestingly 18.18 * 1.1 = 19.99
- T-Mobile interpreted those bids as an offer to split the blocks evenly for 20 million each.
- T-Mobile bid 20 million on licenses 6-10
- The auction ended; German government was unhappy.

Revenue Maximization in Multi-item Settings

Revenue Maximization

- Goal: design a **revenue-optimal truthful** mechanism for selling a few heterogeneous items to a few heterogeneous buyers.
  - 1 item and 1 buyer, buyer's value $v \sim D$.
    - Optimal auction: sell at $p = \mathrm{argmax}_x \; x \cdot \big(1 - F(x)\big)$ where F is the cdf of D.
  - [Myerson '81] provides an optimal single-item auction that is **simple**, **deterministic** and **dominant strategy incentive compatible (DSIC)**.

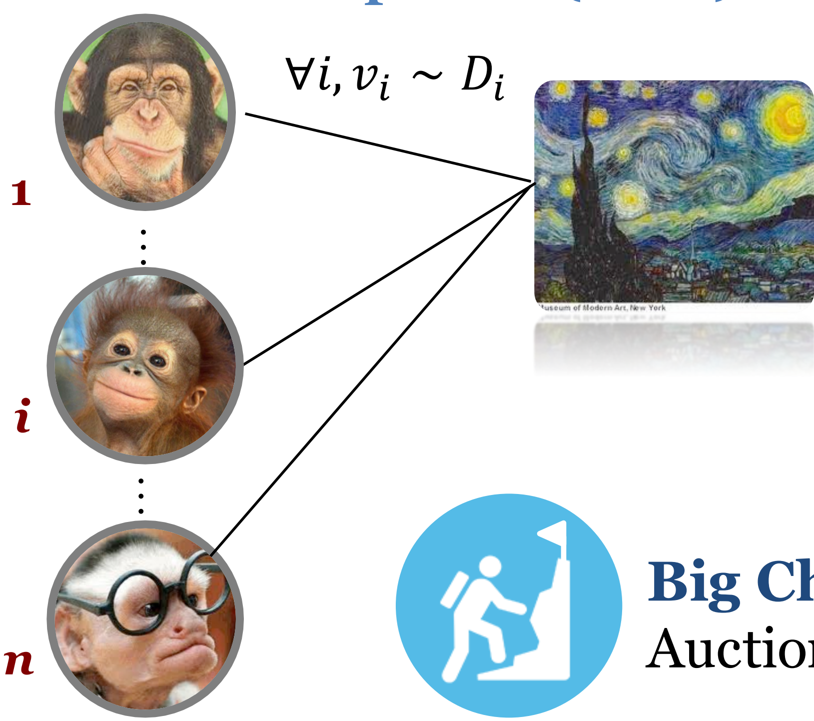

[Myerson' 81]: Optimal auction:
1. Collect bids $\boldsymbol{b_1, \ldots, b_n}$
2. For all $\boldsymbol{i}$: $\boldsymbol{b}_i \mapsto \boldsymbol{b}_i - \frac{1 - \boldsymbol{F}_i(\boldsymbol{b}_i)}{\boldsymbol{f}_i(\boldsymbol{b}_i)} \equiv \hat{\boldsymbol{b}}_i$
3. Choose $\boldsymbol{x}$ *maximizing* $\sum_i \boldsymbol{x}_i \hat{\boldsymbol{b}}_i$
4. Charge "Myerson payments"
   - ensures $\boldsymbol{b_i \equiv v_i}$

Big Challenge: Revenue-Optimal Multi-Item Auctions?

Optimal Multi-item Auctions

- Large body of work in the literature:
  - e.g. [Laffont-Maskin-Rochet'87], [McAfee-McMillan'88], [Wilson'93], [Armstrong'96], [Rochet-Chone'98], [Armstrong'99],[Zheng'00], [Basov'01], [Kazumori'01], [Thanassoulis'04],[Vincent-Manelli '06,'07], [Figalli-Kim-McCann'10], [Pavlov'11], [Hart-Nisan'12], …

- No general approach.

- Challenge already with selling 2 items to 1 bidder:

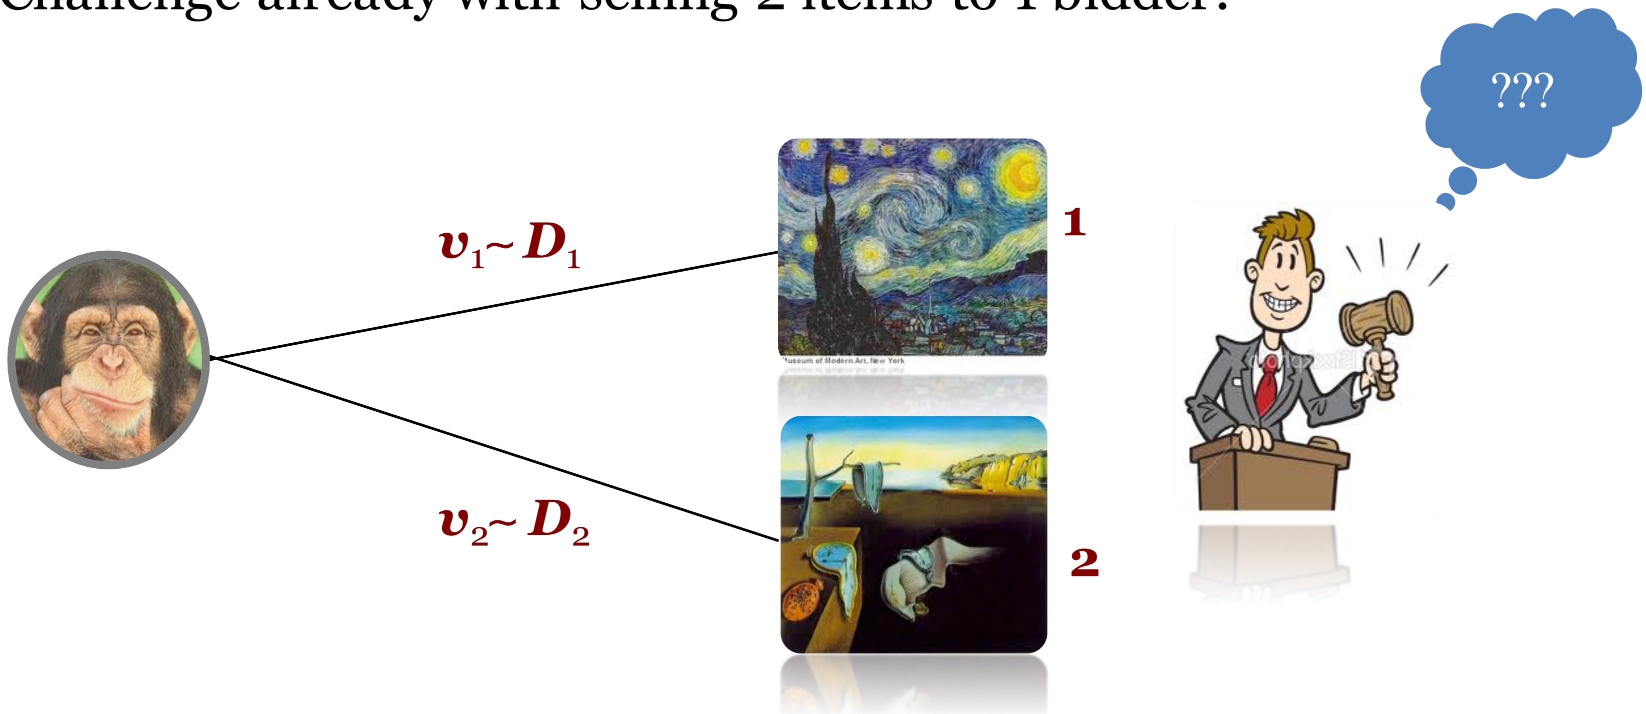

Example 1: Two IID Uniform Items

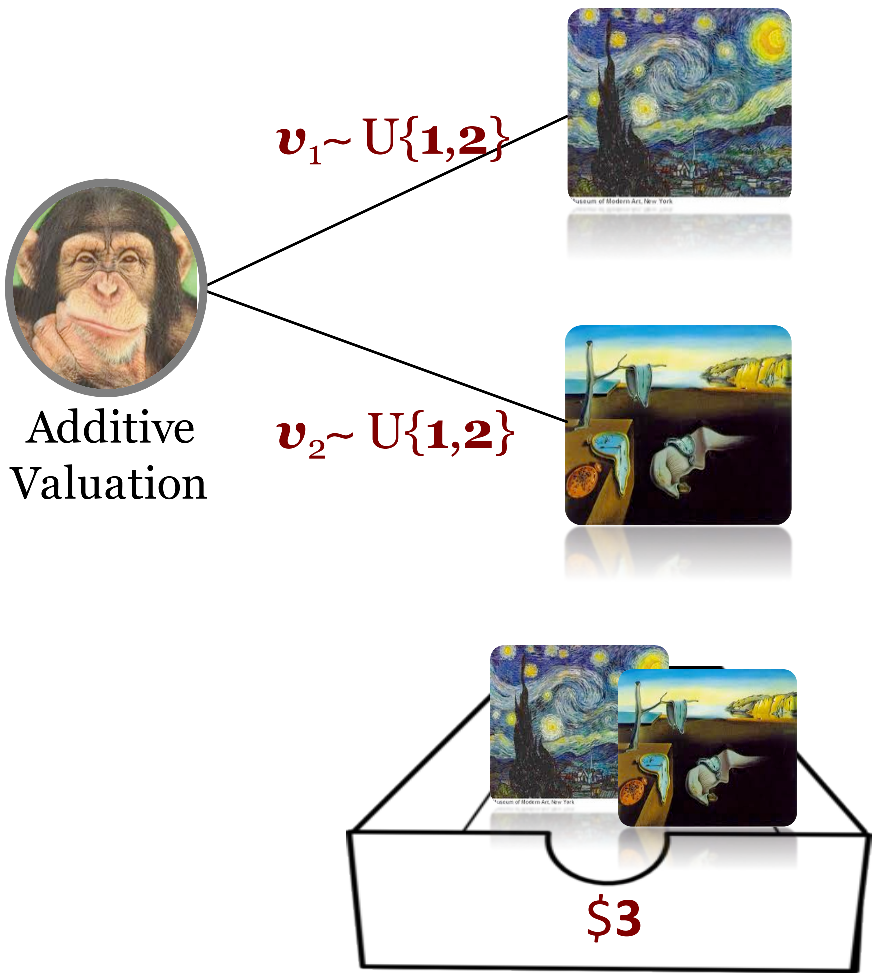

- Strawman approach:
  - Run Myerson for each item separately
  - Price each item at **1**
  - Each bought with probability **1**
  - Expected revenue: **2** × **1** = **2**
- Optimal auction:
  - Expected revenue: **3** × **¾** = **2.25**

∴ Selling items separately might not be optimal.
Bundling increases revenue.

Example 2: Two ID Uniform Items

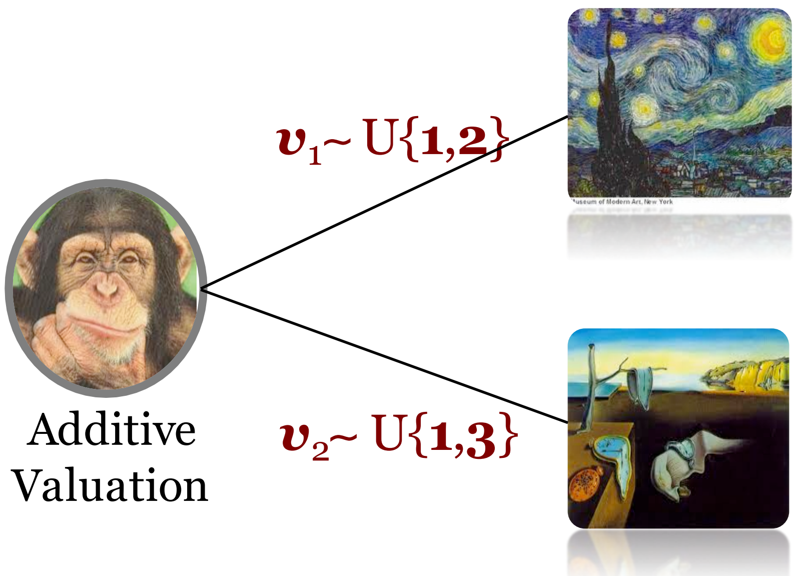

- ❑ Unique optimal auction:
  - expected revenue: **$2.625**

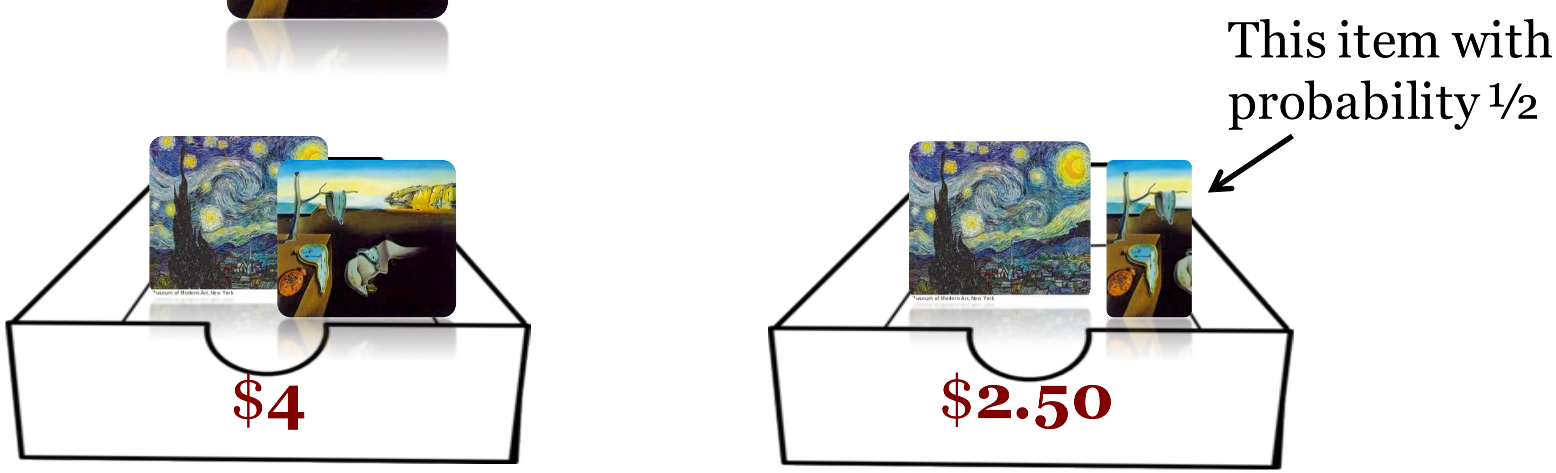

∴ The optimal mechanism may also use **randomization**.

Example 3: Two Beta Distributions

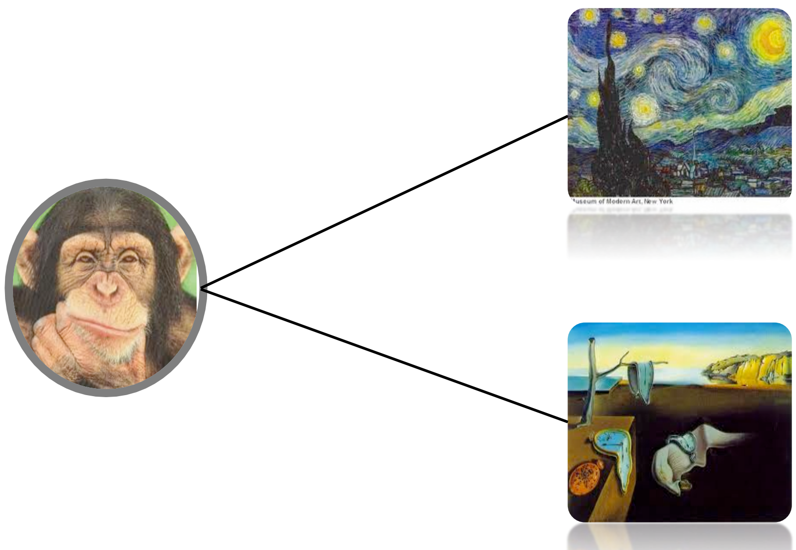

$$f_1(v_1) \propto v_1^2(1 - v_1)^2$$

$$f_2(v_2) \propto v_2^2(1 - v_2)^3$$

- **[Daskalakis-Deckelbaum-Tzamos '13]:** The optimal auction offers ***un-countably many*** randomized bundles.

∴ Can't even represent as a menu!

Example 4: Non-monotonicity

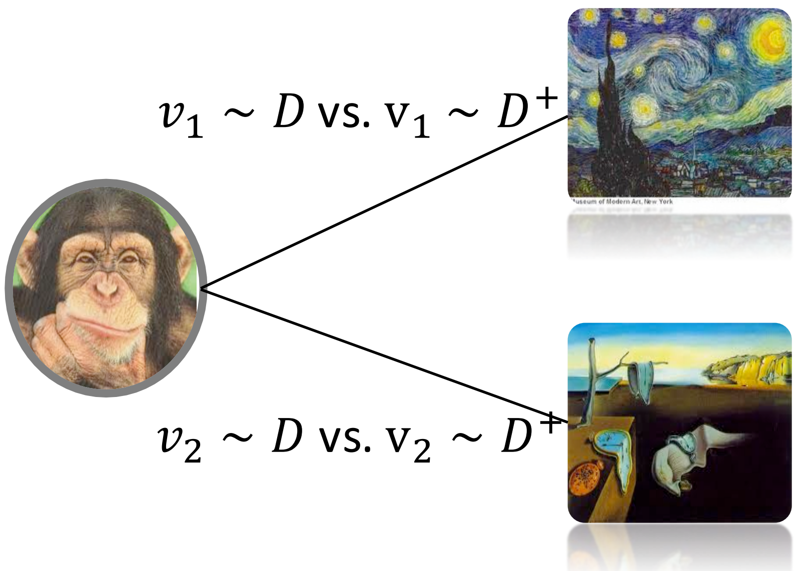

D^+ stochastically dominates D, meaning for any p, $1 - F^+(p) \geq 1 - F(p)$

Question: which is better, selling the paintings to $D \times D$ or $D^+ \times D^+$?

- **[Hart-Reny '13]:** Sometimes, selling to $D \times D$ is better!

$\therefore$ Selling to a worse distribution might generate higher revenue.

Optimal Multi-Item Auctions

- Large body of work in the literature :
  - e.g. [Laffont-Maskin-Rochet'87], [McAfee-McMillan'88], [Wilson'93], [Armstrong'96], [Rochet-Chone'98], [Armstrong'99],[Zheng'00], [Basov'01], [Kazumori'01], [Thanassoulis'04],[Vincent-Manelli '06,'07], [Figalli-Kim-McCann'10], [Pavlov'11], [Hart-Nisan'12], …

- No general approach.

- Challenge already with selling 2 items to 1 bidder:
  - Simple and closed-form solution seems unlikely to exist in general.

- Three possible ways to proceed:

1. **Special Cases:** Usually with assumptions on the distributions.
2. **Algorithmic Solution:** There are polynomial-time computable Revenue-optimal Multi-Item Auctions [Cai-Daskalakis-Weinberg '12 '13].
3. **Simple and Approximately Optimal Solution:** our focus.

Selling Separately and Grand Bundling

- **Theorem:** For a single additive bidder, either selling separately or grand bundling is a 6-approximation [Babaioff et. al. '14].
  - Selling separately: post a price for each item and let the bidder choose whatever he wants. Let SREV be the optimal revenue one can generate from this mechanism.
  - Grand bundling: bundle all the items together and sell the bundle. Let BREV be the optimal revenue one can generate from this mechanism.
  - We will show that Optimal Revenue $\leq$ 2BREV + 4SREV.

Upper Bound of the Optimal Revenue via Duality

Multi-item Auction: Set Up

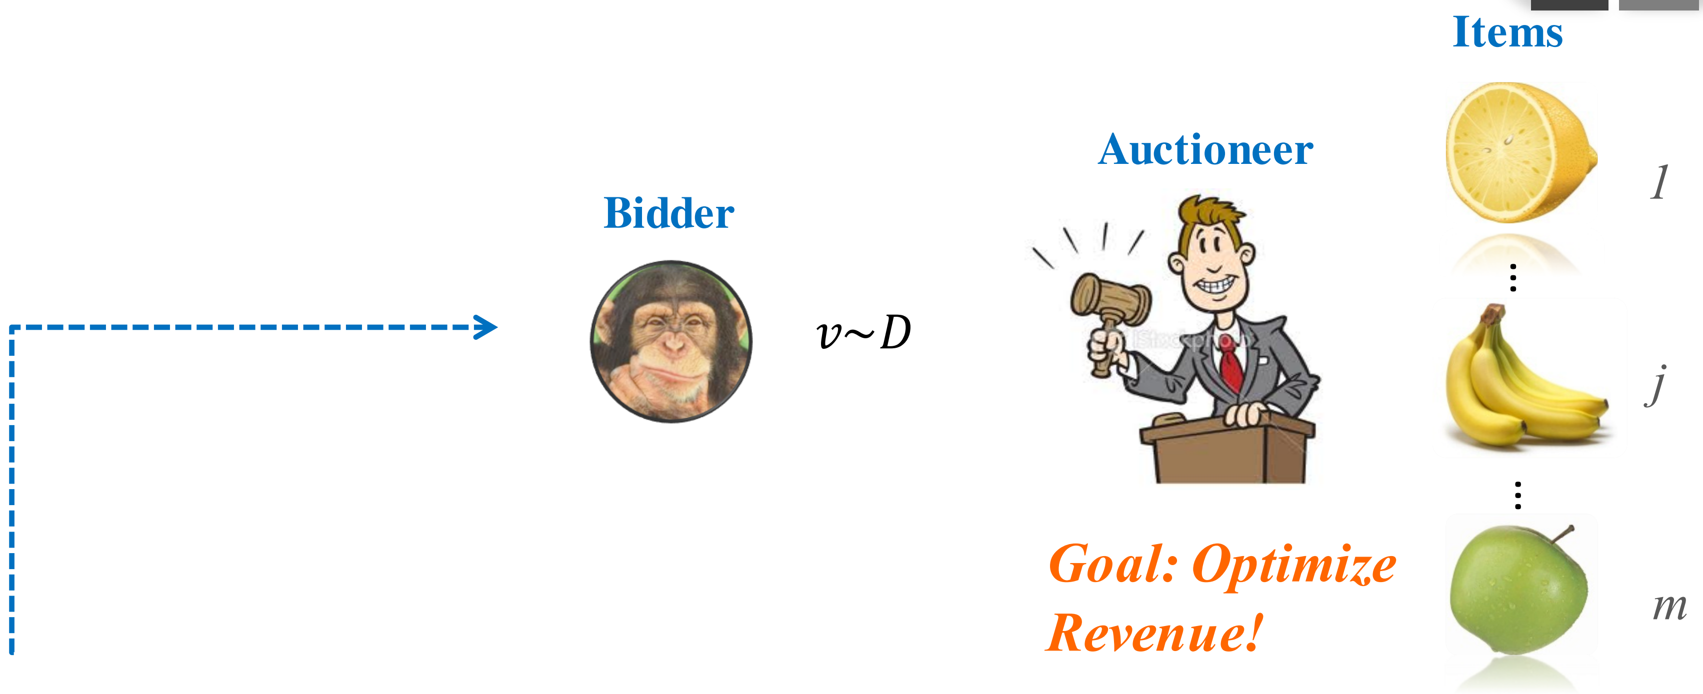

Bidder:

- ***Valuation*** aka ***type*** $v \sim D$. Let $\boldsymbol{V}$ be the support of D.
- **Additive and quasi-linear utility:**
  - $\boldsymbol{v} = (v_1, v_2, \ldots, v_m)$ and $v(S) = \sum_{j \in S} v_j$ for any set S.
- **Independent items:** $\boldsymbol{v} = (v_1, v_2, \ldots, v_m)$ is sampled from $D = \times_j D_j$.

Our Duality (Single Bidder)

Primal LP (Revenue Maximization for 1 bidder)

Variables:

$\boldsymbol{x_j}(\boldsymbol{v})$: the prob. for receiving item j when reporting v.

$\boldsymbol{p}(\boldsymbol{v})$: the price to pay when reporting v.

Constraints:

$\boldsymbol{v} \cdot \boldsymbol{x}(\boldsymbol{v}) - p(\boldsymbol{v}) \geq \boldsymbol{v} \cdot \boldsymbol{x}(\boldsymbol{v}') - p(\boldsymbol{v}'),\ \forall \boldsymbol{v}, \boldsymbol{v}' \in \boldsymbol{V}$ (Truthfulness Constraints)

$\boldsymbol{x}(\boldsymbol{v}) \in P = [0,1]^m, \forall \boldsymbol{v} \in \boldsymbol{V}$ (Feasibility Constraints)

Objective:

$$\max \sum_{\boldsymbol{v}} f(\boldsymbol{v}) p(\boldsymbol{v})$$

Partial Lagrangian

Primal LP:

$$\max \sum_{\boldsymbol{v}} f(v)p(v)$$

s.t. $v \cdot x(\boldsymbol{v}) - p(v) \geq v \cdot x(\boldsymbol{v'}) - p(v'), \ \forall v, v' \in \boldsymbol{V}$ (Truthfulness Constraints)

$x(v) \in P = [0,1]^m, \forall v \in \boldsymbol{V}$ (Feasibility Constraints)

Partial Lagrangian (Lagrangify only the truthfulness constraints):

$$\min_{\lambda > 0} \max_{x \in P, p} L(\lambda, x, p)$$

where

$$L(\lambda, x, p) = \sum_{v} f(v)p(v) + \sum_{v,v'} \lambda(v, v') \cdot (v \cdot (x(v) - x(v')) - (p(v) - p(v'))$$

$$= \sum_{v} p(v) \cdot \left(f(v) + \sum_{v'} \lambda(v', v) - \sum_{v} \lambda(v, v') \right)$$

$$+ \sum_{v} x(v) \cdot \left(v \cdot \sum_{v'} \lambda(v, v') - \left(\sum_{v'} v' \cdot \lambda(v', v) \right) \right)$$

Better be 0, o.w. dual = +∞

The Dual Variables as a Flow

- Observation: If the dual is finite, for every $\boldsymbol{v} \in \boldsymbol{V}$

$$\boldsymbol{f(v) + \sum_{v'} \lambda(v', v) - \sum_{v'} \lambda(v, v') = 0}$$

- This means λ is a flow on the following graph:
  - There is a super source s, a super sink $\emptyset$ and a node for each $\boldsymbol{v} \in \boldsymbol{V}$.
  - $f(\boldsymbol{v})$ flow from s to $\boldsymbol{v}$ for all $\boldsymbol{v} \in \boldsymbol{V}$.
  - $\lambda(\boldsymbol{v}, \boldsymbol{v}')$ flow from $\boldsymbol{v}$ to $\boldsymbol{v}'$, for all $\boldsymbol{v} \in \boldsymbol{V}$ and $\boldsymbol{v}' \in \boldsymbol{V} \cup \{\emptyset\}$.

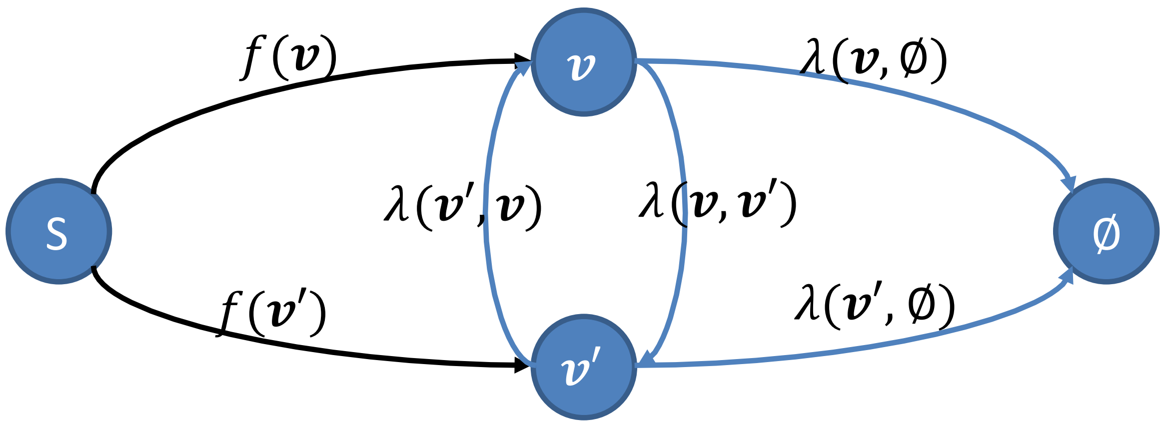

- Suffice to only consider λ that corresponds to a **flow**!

Duality: Interpretation

- Partial Lagrangian Dual (after simplification)

$$\min_{\boldsymbol{flow}\ \boldsymbol{\lambda}} \max_{\boldsymbol{x}\in P} L(\lambda, x, p)$$

where

$$L(\lambda, x, p) = \sum_v f(v) \cdot x(v) \left(v - \frac{1}{f(v)} \sum_{v'} \lambda(v', v)(v' - v) \right)$$

$$= \sum_v f(v) \cdot \sum_j x_j(v) \cdot \Phi_j^{(\lambda)}(v)$$

virtual welfare of allocation $\boldsymbol{x}$ w.r.t. $\Phi^{(\lambda)}(\cdot)$

virtual valuation of $\boldsymbol{v}$ (m-dimensional vector) w.r.t. λ

$$\boldsymbol{\Phi}^{(\lambda)}(\boldsymbol{v}) = \boldsymbol{v} - \frac{1}{f(\boldsymbol{v})} \sum_{\boldsymbol{v}'} \lambda(\boldsymbol{v}', \boldsymbol{v})(\boldsymbol{v}' - \boldsymbol{v})$$

where $\Phi_j^{(\lambda)}(v) = v_j - \frac{1}{f(\boldsymbol{v})} \sum_{\boldsymbol{v}'} \lambda(\boldsymbol{v}', \boldsymbol{v})(v_j' - v_j)$

Note: every flow λ corresponds to a virtual value function $\Phi^{(\lambda)}(\cdot)$

Optimal Revenue (Primal) $\leq$ Optimal Virtual Welfare w.r.t. any λ (Dual) (Weak Duality)

Optimal Revenue $=$ Optimal Virtual Welfare w.r.t. to optimal λ^* (Strong Duality)

Duality: Implication

- Strong duality implies Myerson's result in single-item setting.
  - $\Phi^{(\lambda^*)}(v_i) =$ Myerson's virtual value.

- Weak duality:

[Cai-Devanur-Weinberg '16]: A **canonical way** for deriving approximately tight upper bounds for the optimal revenue.

Single Bidder Flow

Single Bidder: Flow

- For simplicity, assume $\boldsymbol{V} = [H]^m \subseteq \mathbb{Z}^m$ for some integer H.
- Divide the bidder's type set into m regions
  - R_j contains all types that have j as the favorite item.
- **Our Flow**:
  - No cross-region flow ($\lambda(v', v) = 0$ if v, v' are not in the same region).
  - for any $v', v \in R_j$, $\lambda(v', v) > 0$ only if $v'_{-j} = v_{-j}$ and $v'_j = v_j + 1$.
- Our flow λ has the following two properties: for all j and $\boldsymbol{v} \in R_j$
  - $\Phi_{-j}^{(\lambda)}(\boldsymbol{v}) = v_{-j}$.
  - $\Phi_j^{(\lambda)}(v) = \varphi_j(v_j)$, where $\varphi_j(\cdot)$ is the Myerson's Virtual Value function for D_j.

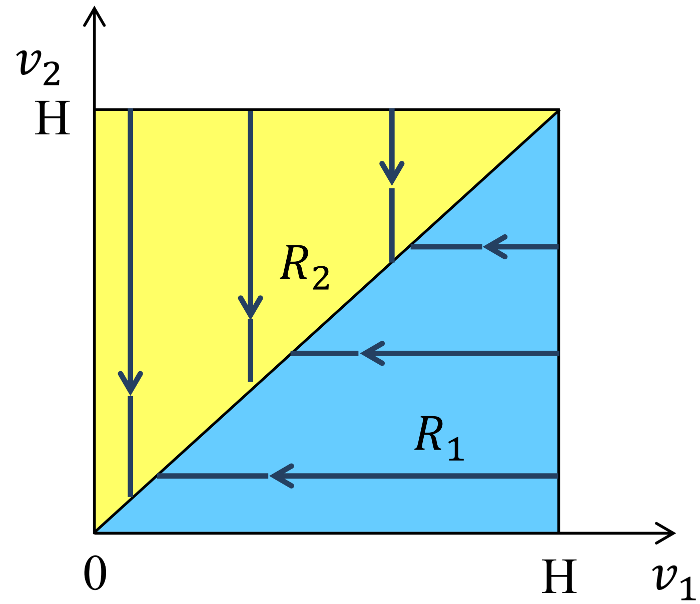

Virtual Valuation:

$$\Phi_j^{(\lambda)}(\boldsymbol{v})$$
$$= \boldsymbol{v_j} - \frac{1}{f(\boldsymbol{v})} \sum_{\boldsymbol{v}'} \lambda(\boldsymbol{v}', \boldsymbol{v})(\boldsymbol{v}'_j - \boldsymbol{v}_j)$$

Single Bidder: Flow (cont.)

For item j:

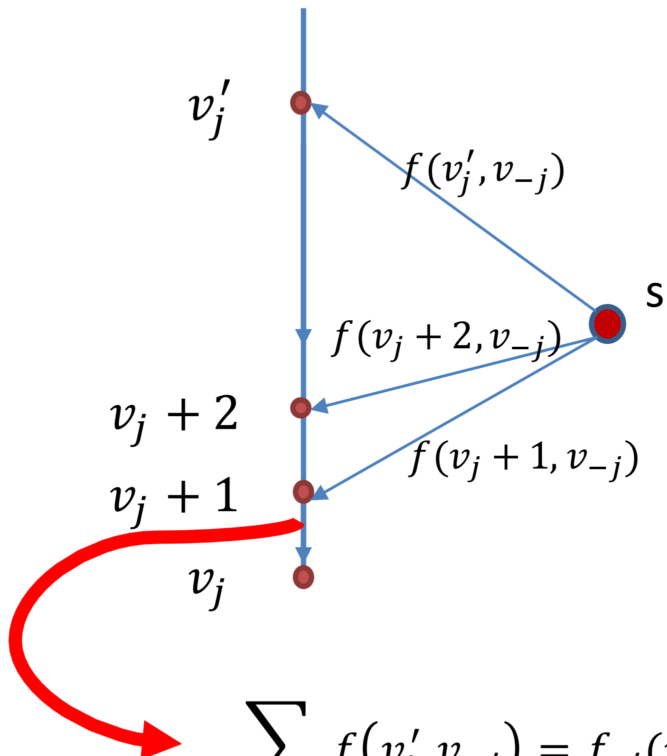

$$\sum_{v_j' > v_j} f\left(v_j', v_{-j}\right) = f_{-j}(v_{-j}) \cdot \left(1 - F_j(v_j)\right)$$

$$\Phi_j^{(\lambda)}(v) = v_j - \frac{1}{f(v)} \sum_{v_j' > v_j} f\left(v_j', v_{-j}\right) = v_j - \frac{1 - F_j(v_j)}{f_j(v_j)}$$

Myerson virtual value function for D_j.

Intuition behind Our Flow

- **Virtual Valuation:**

$$\Phi_j^{(\lambda)}(\boldsymbol{v})$$

$$= \boldsymbol{v}_j - \frac{1}{f(\boldsymbol{v})} \sum_{\boldsymbol{v}'} \lambda(\boldsymbol{v}', \boldsymbol{v})(\boldsymbol{v}'_j - \boldsymbol{v}_j)$$

- **Intuition:**

- Empty flow ➔ social welfare.
- Replace the terms that contribute the most to the social welfare with Myerson's virutal value.

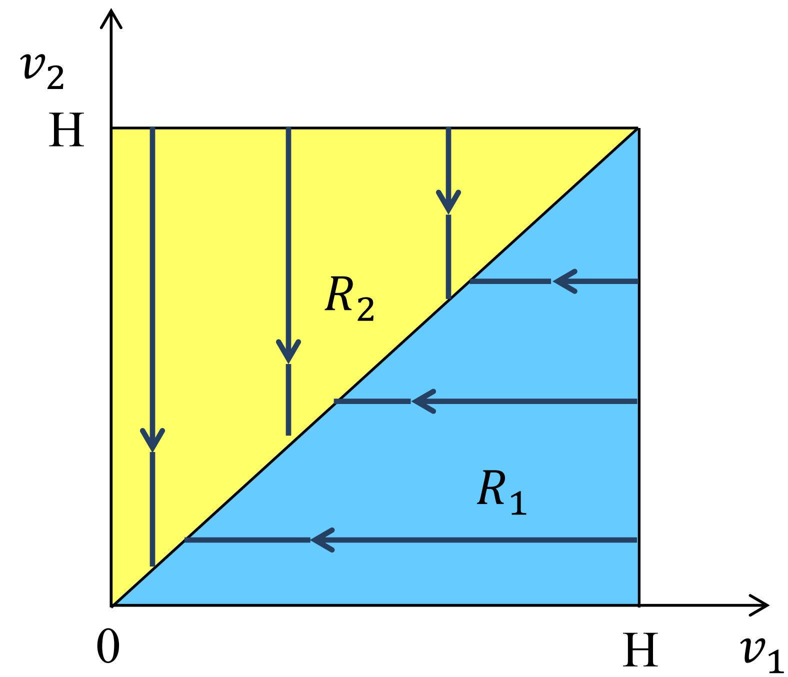

- Our flow λ has the following two properties: for all j and $\boldsymbol{v} \in R_j$
  - $\Phi_{-j}^{(\lambda)}(\boldsymbol{v}) = v_{-j}$.
  - $\Phi_j^{(\lambda)}(v) = \varphi_j(v_j)$, where $\varphi_j(\cdot)$ is the Myerson's Virtual Value function for D_j.

Upper Bound for a Single Bidder

Corollary: $\Phi_j^{(\lambda)}(\boldsymbol{v}) = v_j \cdot \mathbb{I}\left[\boldsymbol{v} \notin R_j\right] + \varphi_j(v_j) \cdot \mathbb{I}[\boldsymbol{v} \in R_j]$.

Upper Bound for Revenue (single-bidder):

$$\text{REV} \leq \max_{x \in P} L(\lambda, x, p) = \sum_{v} \sum_{j} f(\boldsymbol{v}) x_j(\boldsymbol{v}) \cdot \left(v_j \cdot \mathbb{I}\left[\boldsymbol{v} \notin R_j\right] + \varphi_j(v_j) \cdot \mathbb{I}[\boldsymbol{v} \in R_j]\right)$$

Interpretaion: the optimal attainable revenue is no more than the welfare of all non-favorite items plus some term related to the Myerson's single item virtual values.

Theorem: Selling separately or grand bundling achieves at least **1/6** of the upper bound above. This recovers the result by Babaioff et. al. [BILW '14].

Remark: the same upper bound can be easily extended to unit-demand valuations.

Theorem: Posted price mechanism achieves **1/4** of the upper bound above. This recovers the result by Chawla et. al. [CMS '10, '15].