Mechanism Design, Machine Learning, and Pricing Problems

Maria-Florina Balcan

Carnegie Mellon University

Overview

Pricing and Revenue Maximization

Software Pricing

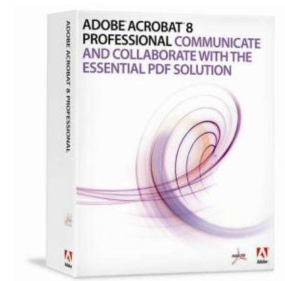

Digital Music

Pricing Problems

One Seller, Multiple Buyers with Complex Preferences.

Seller's Goal: maximize profit.

Algorithm Design Problem (AD)

Version 1: Seller knows the true values.

Incentive Compatible Auction (IC)

Version 2: values given by selfish agents.

BBHM'05: Generic Reduction based on ML techniques

Reduce IC to AD

Generic Framework for reducing problems of incentive-compatible mechanism design to standard algorithmic questions.

[Balcan-Blum-Hartline-Mansour, FOCS 2005, JCSS 2007]

- Focus on revenue-maximization, unlimited supply.
  - Digital Good Auction
  - Attribute Auctions
  - Combinatorial Auctions

Use ideas from Machine Learning.

–Sample Complexity techniques in ML both for design and analysis .

Outline

Part I: Generic Framework for reducing problems of incentive-compatible mechanism design to standard algorithmic questions.

[Balcan-Blum-Hartline-Mansour, FOCS 2005, JCSS 2007]

Part II: Approximation Algorithms for Item Pricing.

[Balcan-Blum, EC 2006, TCS 2007]

Revenue maximization in combinatorial auctions with single-minded consumers.

MP3 Selling Problem

- Seller of some digital good (or any item of fixed marginal cost), e.g. MP3 files.

Goal: **Profit Maximization**

MP3 Selling Problem

- Seller/producer of some digital good, e.g. MP3 files.

Goal: **Profit Maximization**

Digital Good Auction (e.g., [GHW01])

- Compete with fixed price.

or…

- Use bidders' attributes:
  - country, language, ZIP code, etc.
- Compete with best "simple" function.

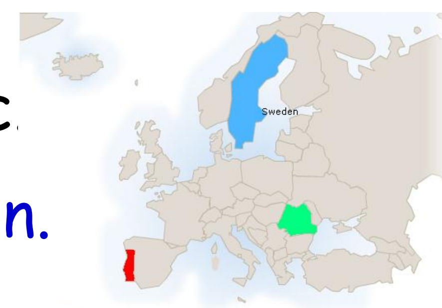

Attribute Auctions [BH05]

Example 2, Boutique Selling Problem

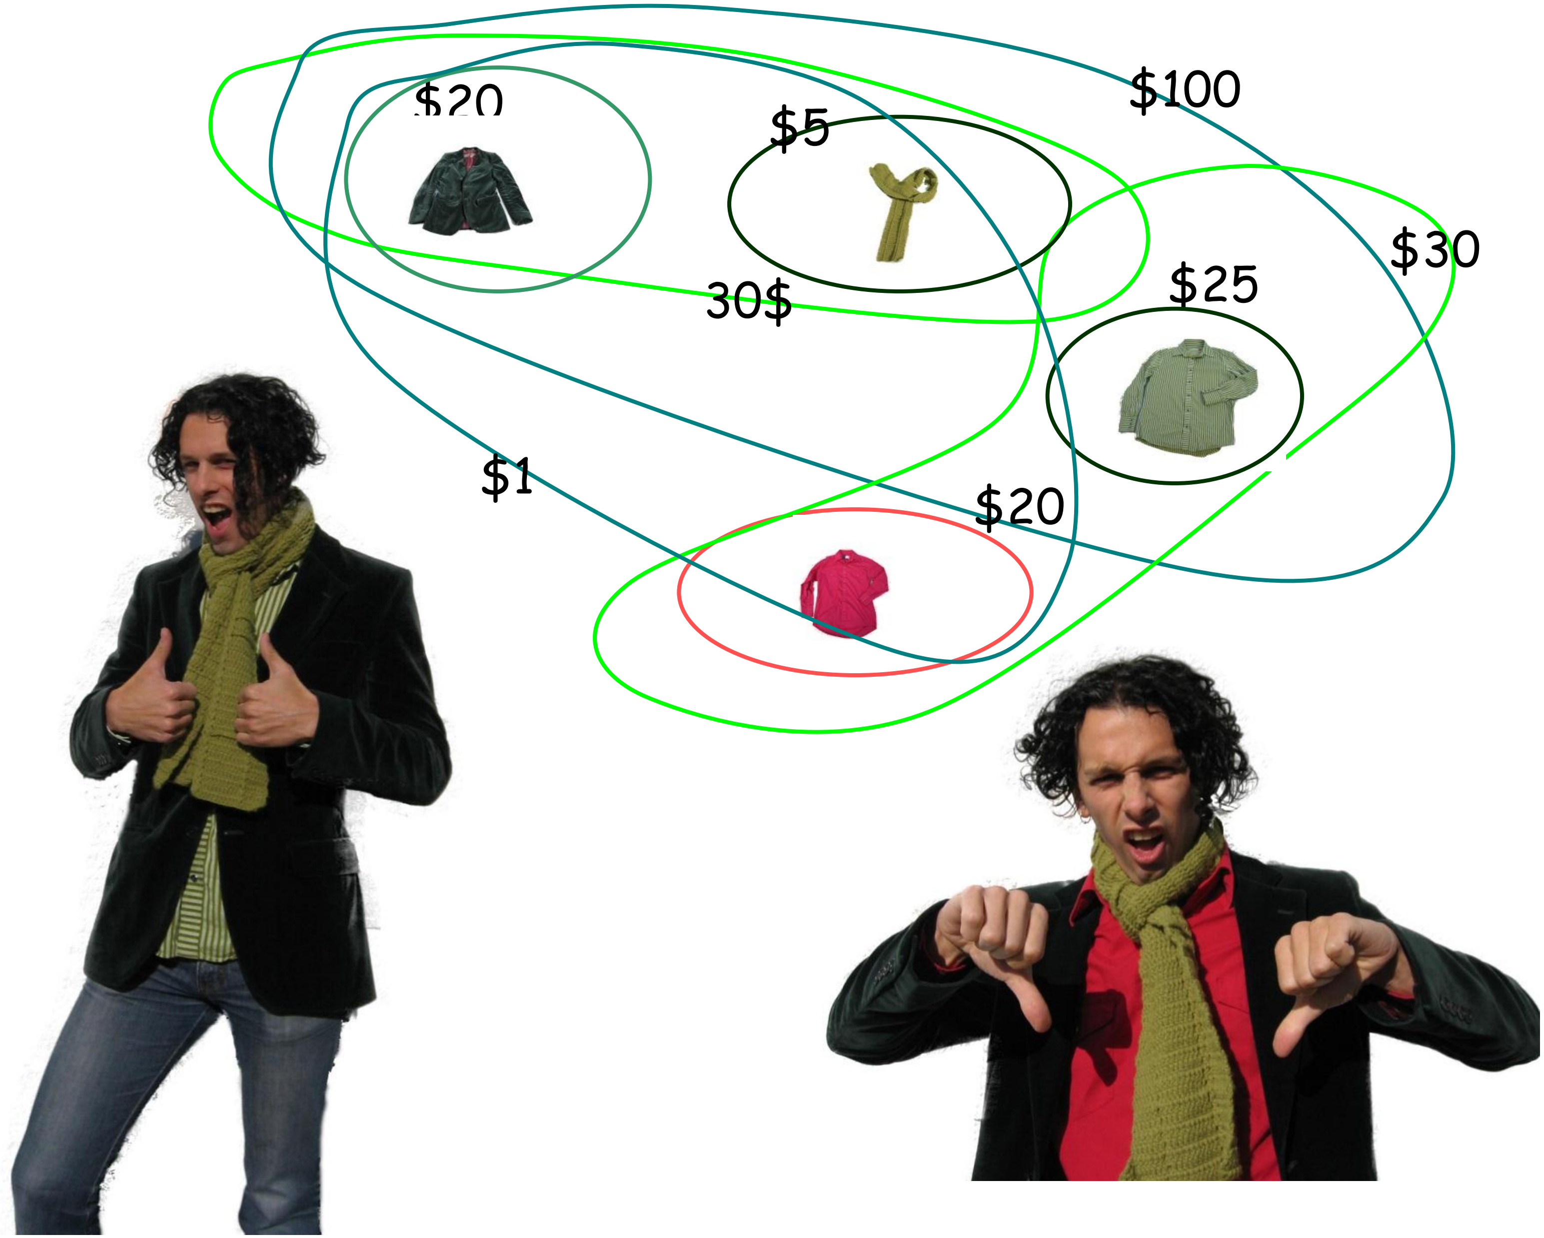

Example 2, Boutique Selling Problem

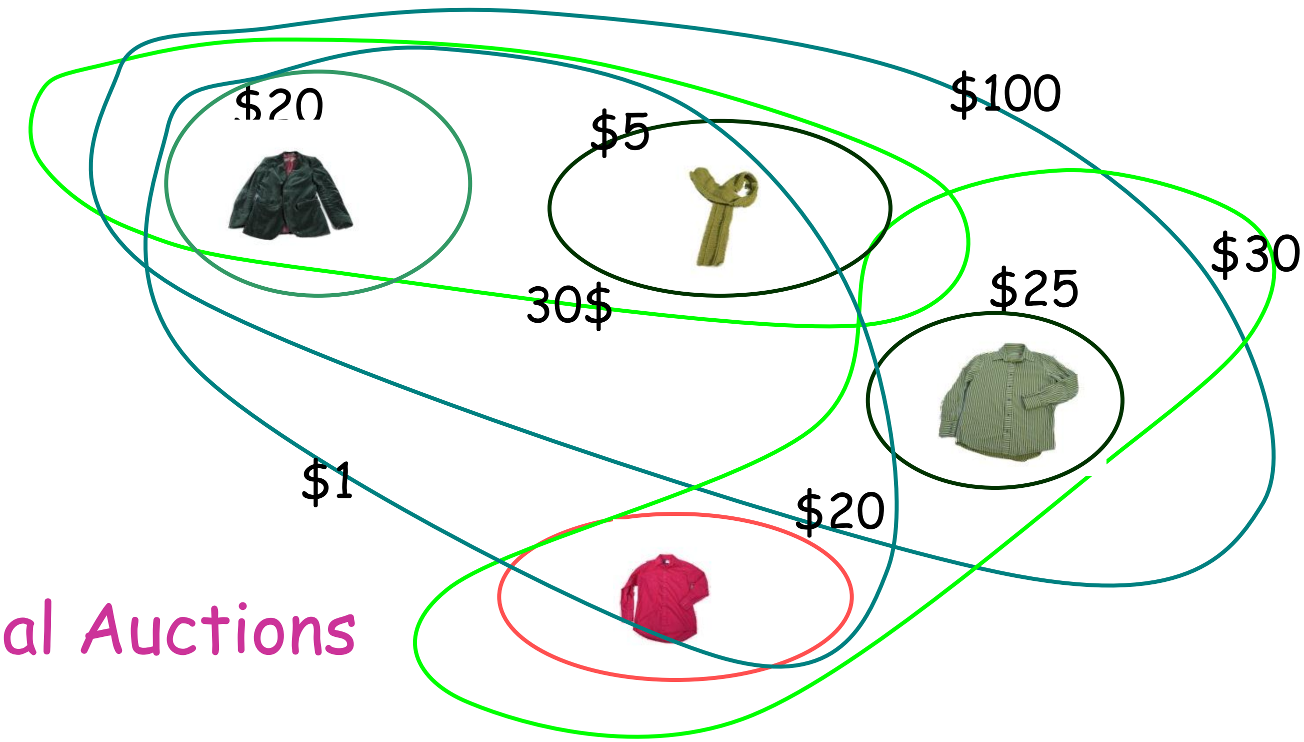

Combinatorial Auctions

Goal: Profit Maximization

- Compete with best item pricing [GH01].

  (unit demand consumers)

Generic Setting (I)

- S set of n bidders. O outcome space.
- Bidder i:
  - $priv_i$ (e.g., how much i is willing to pay for the MP3 file)
  - pub_i (e.g., ZIP code)
  - bid_i (reported $priv_i$)

Incentive Compatible: $bid_i = priv_i$

- Space of legal offers/pricing functions.
  - g maps the pub_i to pricing over the outcome space.
  - $g(i)$ – profit obtained from making offer g to bidder i

Digital Good g=" take the good for p, or leave it"

$g(i) = p$ if $p \le bid_i$

$g(i) = 0$ if $p > bid_i$

Generic Setting (I)

- S set of n bidders.
- Bidder i: $priv_i$, pub_i, bid_i
- Space of legal offers/pricing functions.
  - g maps the pub_i to pricing over the outcome space.
  - $g(i)$ – profit obtained from making offer g to bidder i

Goal: Profit Maximization

- G - pricing functions.
- Goal: Incentive Compatible mechanism to do nearly as well as the best $g \in G$.

Unlimited supply Profit of g: $\sum_i g(i)$

Attribute Auctions

- one item for sale in unlimited supply (e.g. MP3 files).
- bidder i has public attribute $a_i \in X$ (Attr. space)

- G - a class of "natural" pricing functions.

Example: $X = R^2$, G - linear functions over X

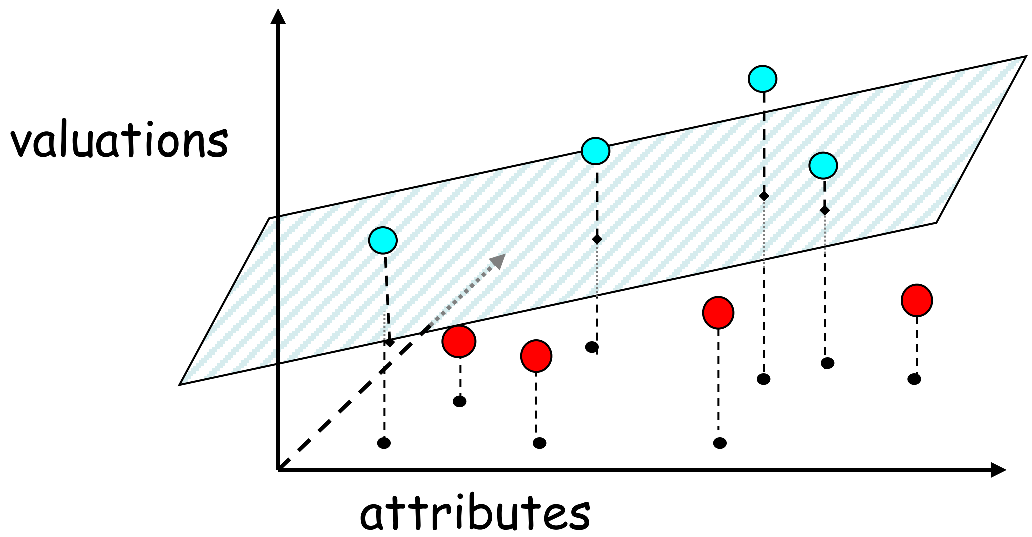

Generic Setting (II)

- Our results: reduce IC to AD.

- Algorithm Design: given ($priv_i$, pub_i), for all $i \in S$, find pricing function $g \in G$ of highest total profit.

- Incentive Compatible mechanism: $bid_i = priv_i$
  - offer for bidder i based on the public information of S and reported private info of $S \setminus \{i\}$.

- Focus on one-shot mechanisms, off-line setting.

Main Results [BBHM05]

- Generic Reductions, unified analysis.
- General Analysis of Attribute Auctions:
  - not just 1-dimensional
- Combinatorial Auctions:
  - First results for competing against opt item-pricing in general case (prev results only for "unit-demand"[GH01])
  - Unit demand case: improve prev bound by a factor of m.

Basic Reduction: Random Sampling Auction

$RSOPF_{(G,A)}$ Reduction

- Bidders submit bids.
- Randomly split the bidders into S_1 and S_2.
- Run A on S_i to get (nearly optimal) $g_i \in G$ w.r.t. S_i.
- Apply g_1 over S_2 and g_2 over S_1.

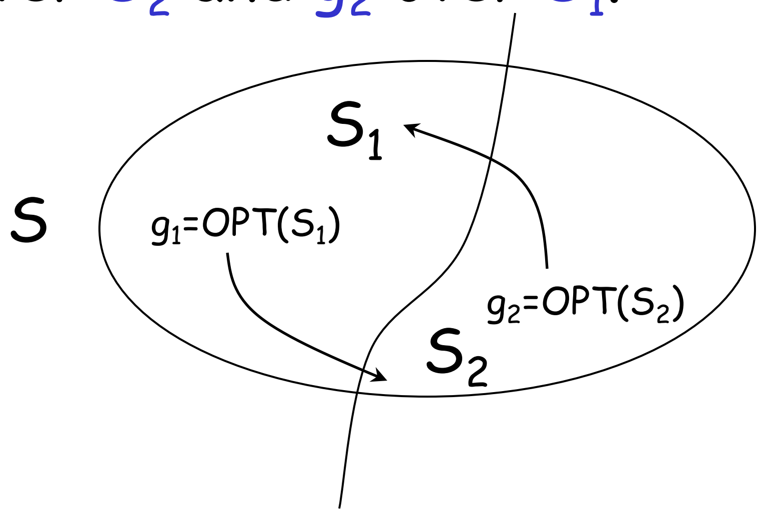

Basic Analysis, $RSOPF_{(G, A)}$

h - maximum valuation, G - finite

Theorem 1

Given a β-approximation algorithm A for optimizing over G, so long as $OPT_G \geq n$ and

$$n \geq \frac{18\beta h}{\epsilon^2} \ln(2|G|/\delta),$$

then whp $1 - \delta$, the profit is at least $(1 - \epsilon)OPT_G/\beta$.

Proof sketch

1) Fixed g and profit level p. Use a tail ineq. show:

Lemma 1

Randomly partition S into S_1 and S_2, then the probility that $|g(S_1) - g(S_2)| \geq \epsilon \max[g(S), p]$ is at most $2e^{-\epsilon^2 p/(2h)}$.

Basic Analysis, $RSOPF_{(G,A)}$, cont

2) Let g_i be the best over S_i. Know $g_i(S_i) \geq g_{OPT}(S_i)/\beta$.

Apply union bound, get whp $(1-\delta)$, every $g \in G$ satisfies $|g(S_1) - g(S_2)| \leq \frac{\epsilon}{2} \max[g(S), n]$.

In particular,

$$g_1(S_2) \geq g_1(S_1) - \frac{\epsilon}{2} \max[g_1(S), n]$$

$$g_2(S_1) \geq g_2(S_2) - \frac{\epsilon}{2} \max[g_2(S), n]$$

Using also $OPT_G \geq \beta$ n, get that our profit $g_1(S_2) + g_2(S_1)$ is at least $(1-\varepsilon)OPT_G/\beta$.

Attribute Auctions, $RSOPF_{(G_k, A)}$

G_k : k markets defined by Voronoi cells around k bidders & fixed price within each market.

Discretize prices to powers of $(1+\varepsilon)$.

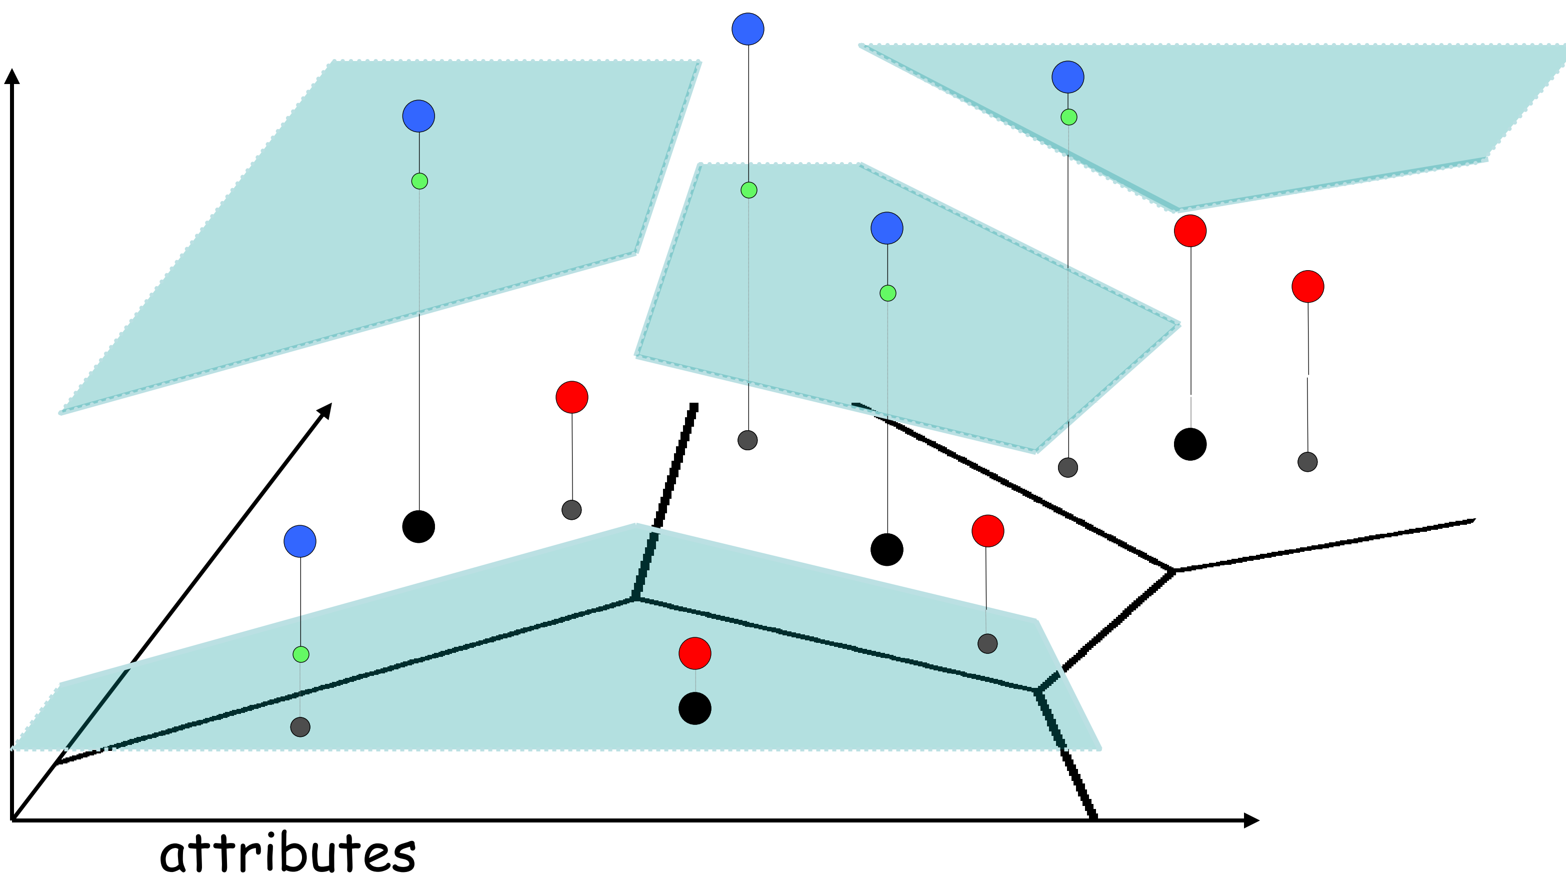

Attribute Auctions, $RSOPF_{(G_k, A)}$

G_k : k markets defined by Voronoi cells around k bidders & fixed price within each market.

Discretize prices to powers of $(1+\varepsilon)$.

Corollary (roughly)

So long as $OPT_{G_k} \geq \beta n$ and $n \geq \frac{kh}{\epsilon^2} \log\left(\frac{k}{\epsilon} h \log h\right)$, then whp the profit is at least $(1-\epsilon)OPT_{G_k}/\beta$.

Structural Risk Minimization Reduction

What if different functions at different levels of complexity?

Don't know best complexity level in advance.

SRM Reduction

Let $G_1 \subseteq G_2 \subseteq G_3 \subseteq \ldots$

- Randomly split the bidders into S_1 and S_2.
- Compute g_i to maximize $\max_k \max_{g \in G_k} [g(S_i) - \text{pen}(G_k)]$
- Apply g_1 over S_2 and g_2 over S_1.

Theorem

Let $\text{pen}(G_k) = \frac{8h}{\epsilon^2} \ln(8k^2|G_k|/\delta)$. Whp $1-\delta$, the profit is:

$$\max_k ((1-\epsilon)\text{OPT}_k - 2\text{pen}(G_k)).$$

Attribute Auctions, Linear Pricing Functions

Assume $X=R^d$. $N = (n+1)(1/\varepsilon) \ln h$.

$|G'| \leq N^{d+1}$

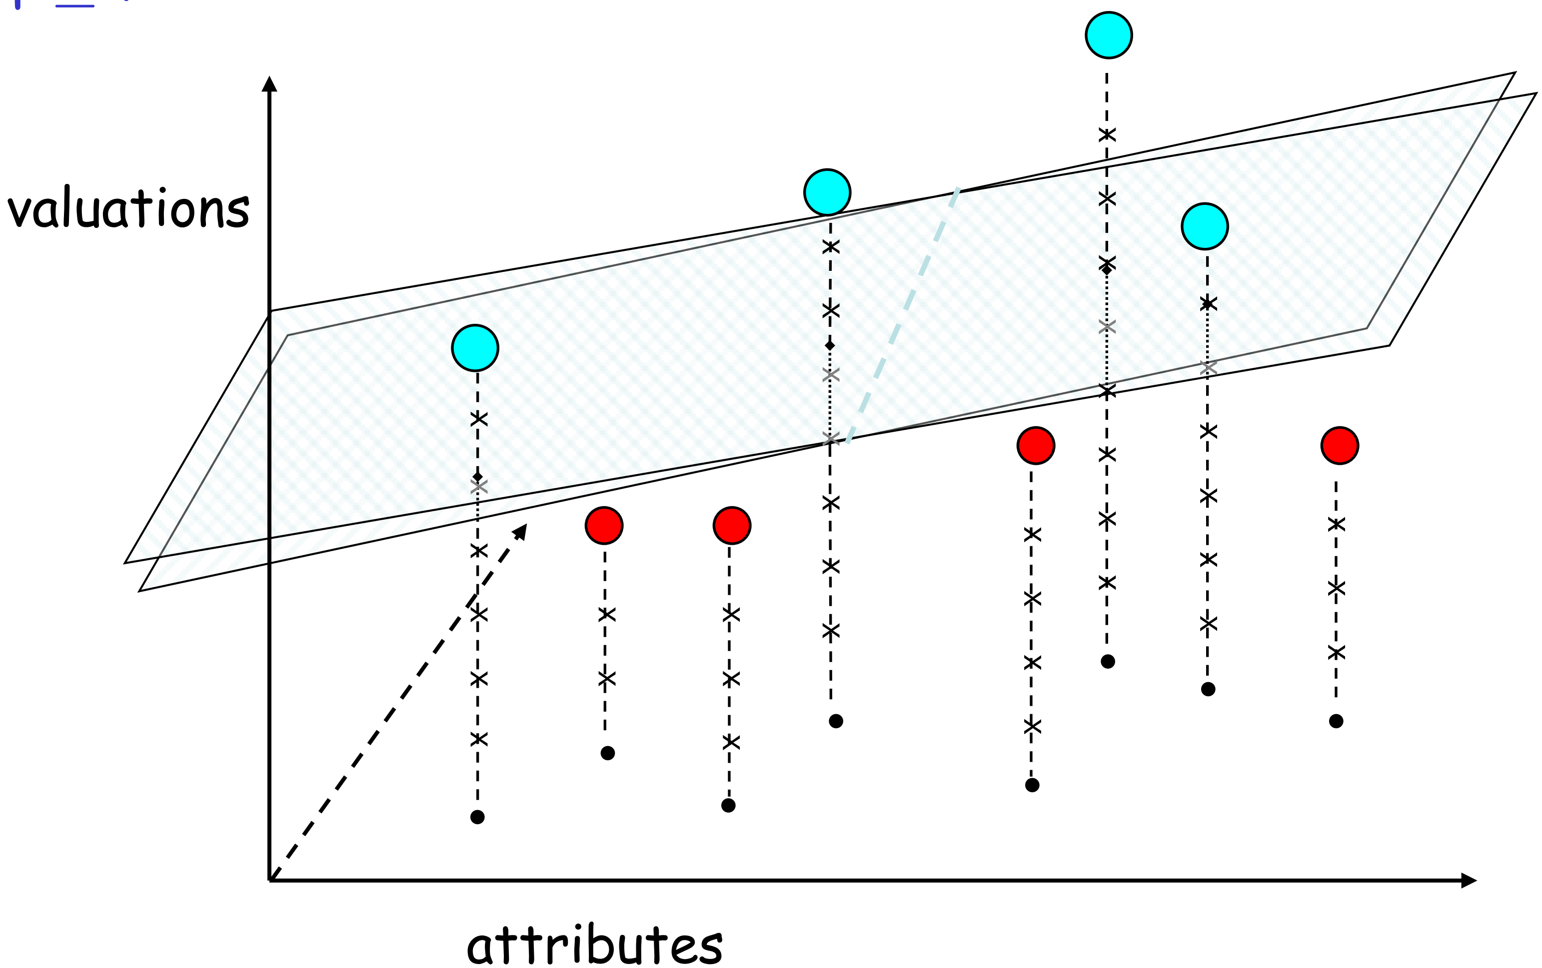

Covering Arguments

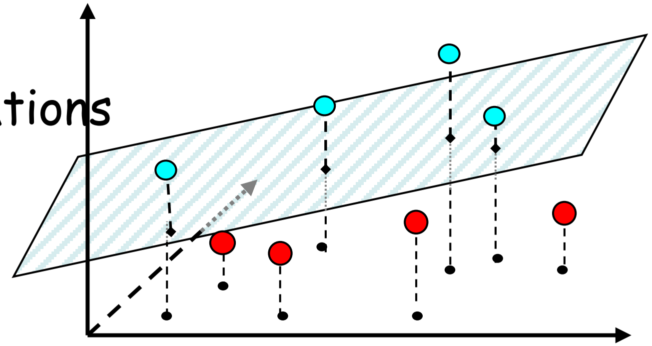

What if G is infinite w.r.t S?

Use covering arguments:

- find G' that covers G ,
- show that all functions in G' behave well

Definition:

G' γ-covers G wrt to S if for $\forall$ g $\exists$ $g' \in G'$ s.t.
$\forall$ i $|g(i)-g'(i)| \leq \gamma\, g(i)$.

Analysis Technique

Theorem (roughly)

If G' is γ-cover of G, then the previous theorems hold with $|G|$ replaced by $|G'|$.

Summary [BBHM05]

- Explicit connection between machine learning and mechanism design.

- Use MLT both for design and analysis in auction/pricing problems.

- Unique challenges & particularities:
  - Loss function discontinuous and asymmetric.
  - Range of valuations large.

- See also upcoming paper of [Morgenstern, Roughgarden, NIPS'15] for other settings (e.g., limited supply)!

Outline

Part I: Generic Framework for reducing problems of incentive-compatible mechanism design to standard algorithmic questions.

Part II: Approximation Algorithms for Item Pricing.

[Balcan-Blum, EC 2006, TCS 2007]

Revenue maximization in combinatorial auctions with single-minded consumers

Algorithmic Problem, Single-minded Bidders [BB'06]

- m item types with unlimited supply of each.
- n single-minded customers.
- Customer i: shopping list L_i, will only shop if the total cost of items in L_i is at most w_i
- All marginal costs are 0, and we know all the (L_i, w_i).

What prices on the items will make you the most money?

- Easy if all L_i are of size 1.
- What happens if all L_i are of size 2?

Algorithmic Problem, Single-minded Bidders [BB'06]

- A multigraph G with values w_e on edges e.

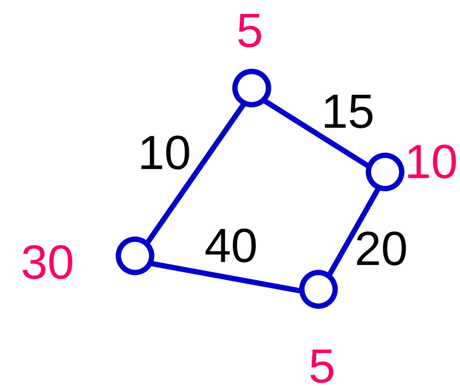

- **Goal:** assign prices on vertices $p_v \geq 0$ to maximize total profit, where:

$$\text{Profit}(p) = \sum_{\substack{e=(u,v) \\ p_u + p_v \leq w_e}} (p_u + p_v)$$

Unlimited supply

- APX hard [GHKKKM'05].

A Simple 2-Approx. in the Bipartite Case

- Given a multigraph G with values w_e on edges e.
- Goal: assign prices on vertices $p_v \geq 0$ to maximize total profit, where:

$$\text{Profit}(p) = \sum_{\substack{e=(u,v) \\ p_u + p_v \leq w_e}} (p_u + p_v)$$

Algorithm

- Set prices in R to 0 and separately fix prices for each node on L.
- Set prices in L to 0 and separately fix prices for each node on R.
- Take the best of both options.

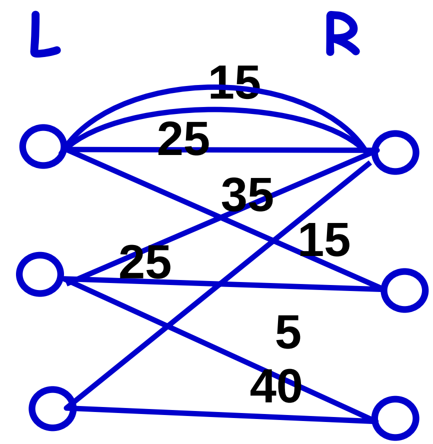

Proof

$OPT = OPT_L + OPT_R$

A 4-Approx. for Graph Vertex Pricing

- Given a multigraph G with values w_e on edges e.
- **Goal:** assign prices on vertices $p_v \geq 0$ to maximize total profit, where:

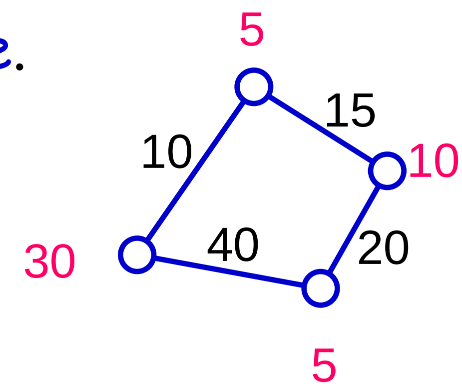

$$\text{Profit}(p) = \sum_{\substack{e=(u,v) \\ p_u + p_v \leq w_e}} (p_u + p_v)$$

Algorithm

- Randomly partition the vertices into two sets L and R.
- Ignore the edges whose endpoints are on the same side and run the alg. for the bipartite case.

Proof

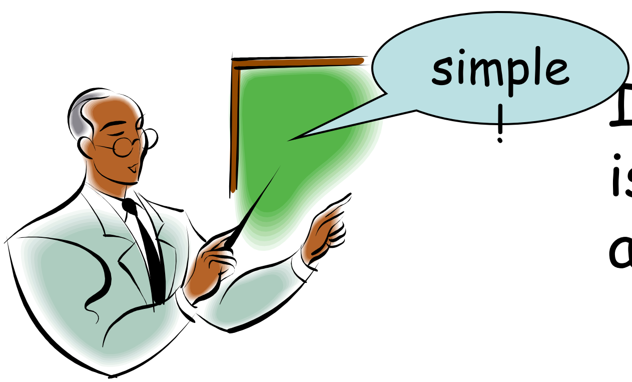

In expectation half of OPT's profit is from edges with one endpoint in L and one endpoint in R.

Algorithmic Pricing, Single-minded Bidders, k-hypergraph Problem

List of size $\leq$ k.

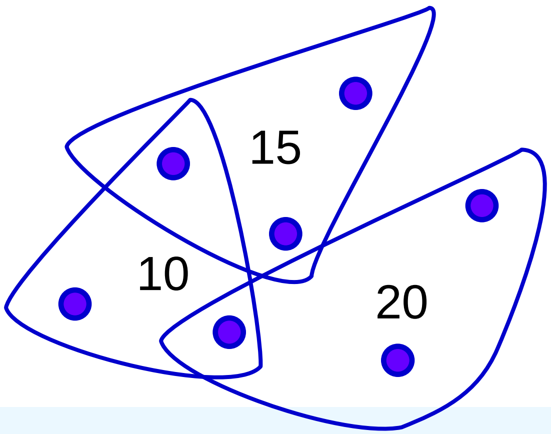

Algorithm

- Put each node in L with prob. 1/k, in R with prob. 1 – 1/k.
- Let GOOD = set of edges with exactly one endpoint in L. Set prices in R to 0 and optimize L wrt GOOD.

- Let $OPT_{j,e}$ be revenue OPT makes selling item j to customer e. Let $X_{j,e}$ be indicator RV for $j \in L$ & $e \in$ GOOD.
- Our expected profit at least:

$$\mathbf{E}\left[\sum X_{j,e} OPT_{j,e}\right] = \sum \mathbf{E}\left[X_{j,e}\right] OPT_{j,e} = \Omega(1/k)\, OPT$$

Summary [BB06]:

- 4 approx for graph case.
- O(k) approx for k-hypergraph case.

Improves the $O(k^2)$ approximation of Briest and Krysta, SODA'06.

- Also simpler and
  can be naturally adapted to the online setting.

Other known results:

- O(log mn) approx. by picking the best single price [GHKKKM05].
- $\Omega(\log^{\varepsilon} n)$ hardness for general case [DFHS06].

Overall Summary

Revenue Maximization in a wide range of settings.

- Both Algorithmic and Incentive Compatible Aspects.
- Natural Connections to Machine Learning.

Thank you !