Profit-Maximizing Envy-free Pricing

Venkat Guruswami
Jason Hartline
Anna Karlin
David Kempe
Claire Kenyon
Frank McSherry

Envy-free pricing

Seller: (m items)

- Sets price p_j for item j.

Consumers: (n consumers)

- Consumer i has *valuation* v_{ij} for item j.
- For item j at price p_j, consumer i has *utility*: $u_{ij} = v_{ij} - p_j$.
- Consumer desired item to maximize utility.

Agreement: (*envy-free pricing*)

- Must have allocation such that all consumers are happy.

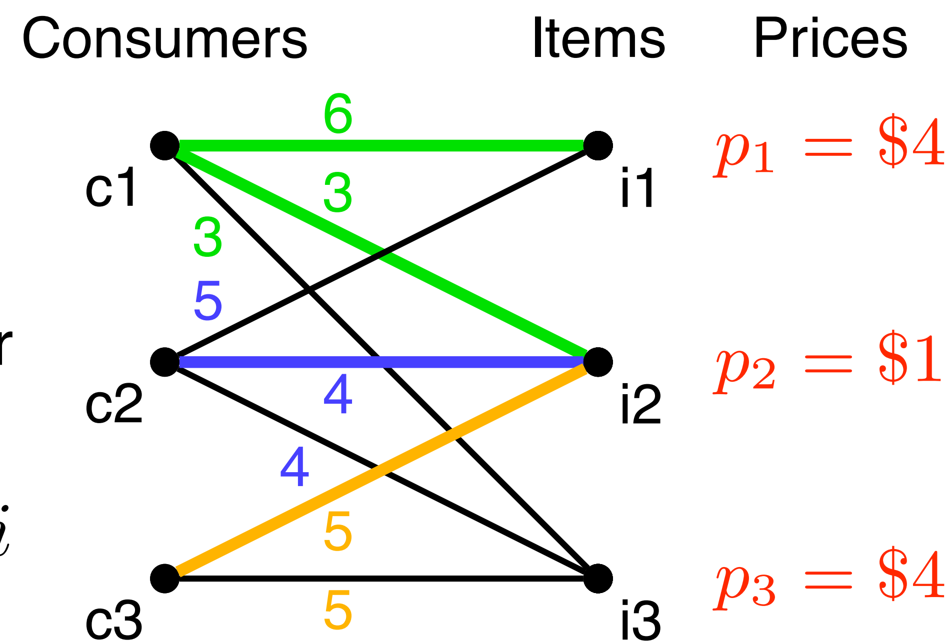

For Consumer 3:

$$u_{32} = \$4$$

$$u_{33} = \$1$$

Envy-free pricing

Given: the valuations v_{ij}

Find: prices p_j and allocation

so as: to maximize seller profit $\sum_{j\ allocated}\ p_j$

s.t.: envy-free constraint:

if customer i is allocated item j,

then $p_j \leq v_{ij}$

and for every j', we have $v_{ij'} - p_{j'} \leq v_{ij} - p_j$.

Motivation

- Envy-free pricing has been studied in Economics for 50 years [Walras 1954]
- Envy-free $\Rightarrow$ consumer has no incentive to change allocation; Maximum-profit $\Rightarrow$ seller has no incentive to change prices.
- "Price of truthfulness" in combinatorial auctions: to analyze performance, compare truthful mechanism to envy-free (full information) pricing

Main result

Envy-free pricing has an $O(\log n)$ approximation algorithm

Walrasian Equilibrium

Definition: *Walrasian Equilibrium*, an envy-free pricing with unallocated items at price zero.

Examples:

- for unlimited supply, all items at price zero.
- for limited supply, all items at price zero is not Walrasian Equilib.

Algorithm: *Vickrey-Clarke-Groves* (VCG):

1. Allocate items via *maximum weighted matching* (MM).
2. For (i, j) in matching, give item j to consumer i at price:

$$p_j = v_{i,j} - \mathrm{MM}(V) + \mathrm{MM}(V_{-i}).$$

Theorem: [Leonard 83] VCG outputs a Walrasian Equilibrium.

Reserve Prices

Definition: *reserve price*, a lower bound on the sale price of an item.

Reserve prices often used to obtain more revenue from VCG.
(e.g., Bayesian optimal auction [Myerson 81])

Walrasian + Reserve

Definition: *Walrasian Equilibrium with reserve price r*, a envy-free pricing with unallocated items at price r.

Algorithm: *Vickrey-Clarke-Groves with reserve price r* (VCG$_r$):

1. Construct V': add two dummy consumers for each item with valuation r.
2. Run VCG on V' and output prices.

(breaking ties in favor of real consumers)

Lemma: The VCG$_r$ prices are a Walrasian Equilibrium with reserve price r.

Proof: If price for unsold item j is less than r,

- in VCG(V'), one dummy consumer for item j would envy, thus
- prices of VCG(V') would not be Walrasian for V'.

Log n Approximation

Lemma: If MM sells k items at price $\geq p$, then VCG$_p \geq kp/2$.

Proof:

- Consider $(i, j) \in \mathrm{MM}$ at price above p.
- Either i or j is matched in VCG$_p$.
  (otherwise we could add (i, j) to VCG$_p$)
- Therefore, number of matched *vertices* in VCG$_p$ is $\geq k$.
- Number of matched items in VCG$_p$ is $\geq k/2$.

Algorithm: Limited Supply Logarithmic Approximation:

1. Run $\mathrm{MM}(V)$ to compute prices $p_1, \ldots, p_m$.
2. Output VCG_{p_i} with highest profit.

Analysis: Profit $\geq \max_i \frac{ip_i}{2} \geq \frac{\sum_i p_i}{2 \ln n} = \frac{\mathrm{MM}}{2 \ln n} \geq \frac{\mathrm{OPT}}{2 \ln n}$

Proof of lemma

if j is sold to a real consumer

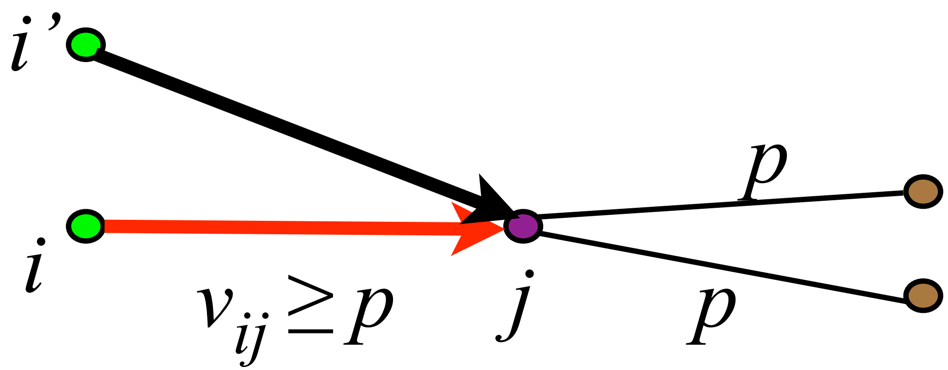

if j is not sold to a real consumer

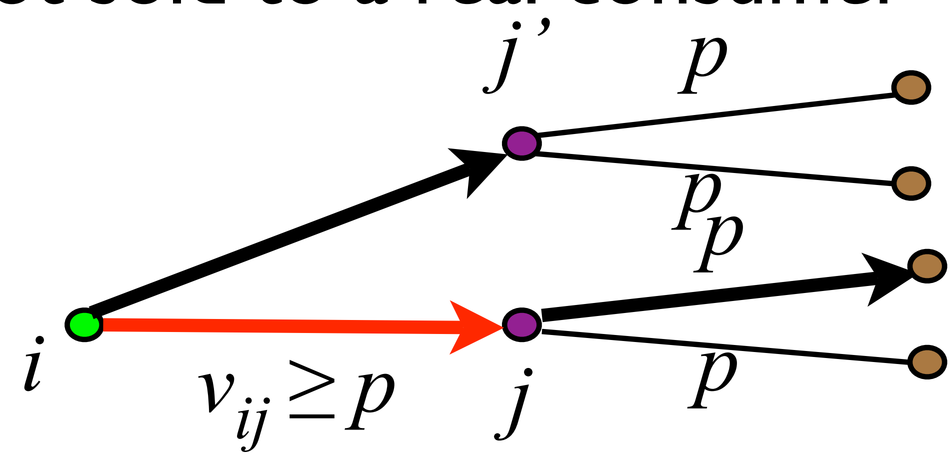

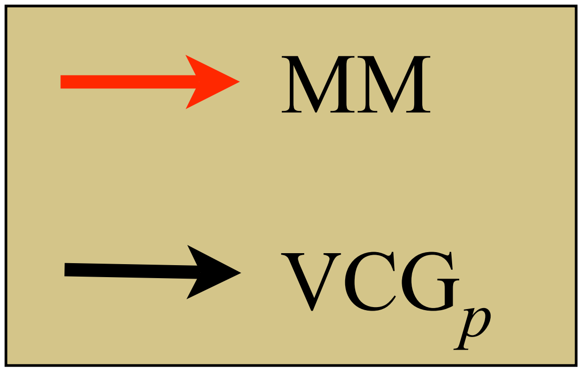

Open

- Open: Is there a matching lower bound?
- Known: APX-hard, by reduction from Vertex Cover in bounded degree graphs

Open Problem #1: Pricing over Time

A special case of Envy-free pricing.

- Customer: "I want to buy a Boston-Bologna ticket between June 16 and June 18 and pay at most $600."
- supply c_t of seats available at date t; customer = $[s_i, t_i]$, valuation v_i
- Open problem: design an algorithm that is better than the $\log n$ approximation (the case $c_t = 1$ ∀t would already be interesting)
- Not known: not known to be NP-hard
- Known: Solvable by dynamic programming if unlimited supply $c_t \geq n$ ∀t

Unlimited Supply

Definition: *unlimited supply* special case: the number of copies of each item is $> n$.
(E.g., pricing in-flight movies.)

Definition: Item j and j' are *identical* iff $v_{ij} = v_{ij'}$ for all i.

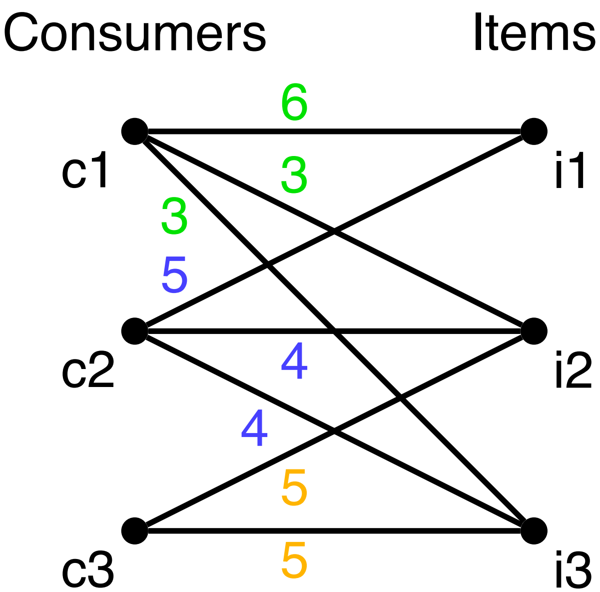

Unlimited supply implications:

- Identical items sold at same price $\Leftrightarrow$ envy-free.
- Given prices, consumers pick favorite item.

Open problem #2: Unlimited supply envy-free pricing

- Open problem: Is there an $O(1)$-approximation algorithm if all items are in unlimited supply?
- Known hardness: APX-hard
- Known algorithm: $O(\log n)$-approximation algorithm by our main result or by the best-uniform-price algorithm
- Failed attempt: Linear programming relaxation + randomized rounding. Variables $x(i, j, p) = 1$ if customer i buys item j at price p. And $y(j, p) = 1$ if item j is offered at price p. Difficulty: our models all seemed to have unbounded integrality gap!

Open problem #3: Unlimited Supply Tollbooth problem

- Variation: customers want to buy bundles instead of single items.
- The problem: Items are edges forming a tree, bundles are paths, customer i wants to buy a specific path (or nothing at all) and pay at most v_i to buy all the edges on the path. Unlimited supply: edges have infinite capacity.
- Known hardness: APX-hard
- Known algorithm: $O(\log n)$ approximation
- Known special case: solvable by dynamic programming if all paths end at same node
- Open problem: Is there an $O(1)$-approximation algorithm?

Open problem #4: Unlimited Supply Highway problem

- Customers want to buy bundles
- The problem: items are edges forming a path, bundles are subpaths, customer i wants to buy $[s_i, t_i]$ and pay at most an integer sum v_i to buy all the edges in the interval. Edges have infinite capacity.
- No known hardness.
- Known algorithm: $O(\log n)$ approximation
- Known special case: solvable by dynamic programming if all bundles have bounded length $t_i - s_i = O(1)$
- Open problem: Is there an $O(1)$-approximation algorithm, or maybe even an exact algorithm?