Online resource allocation, pricing, and prophet inequalities

SHUCHI CHAWLA

UNIV. WISCONSIN-MADISON

Based on:

- *Posted pricing for online resource allocation: intervals and paths*. C., Miller, and Teng. SODA'19.
- *Stability of service under time-of-use pricing*. C., Devanur, Holroyd, Karlin, Martin, and Sivan. STOC'17.
- *Combinatorial auctions via posted prices*, Feldman, Gravin, and Lucier. SODA'15.

Allocating limited resources

- Objective: maximize social welfare subject to supply constraints

Social Welfare = $\sum_{\text{buyers } i}$ value of buyer i from allocation

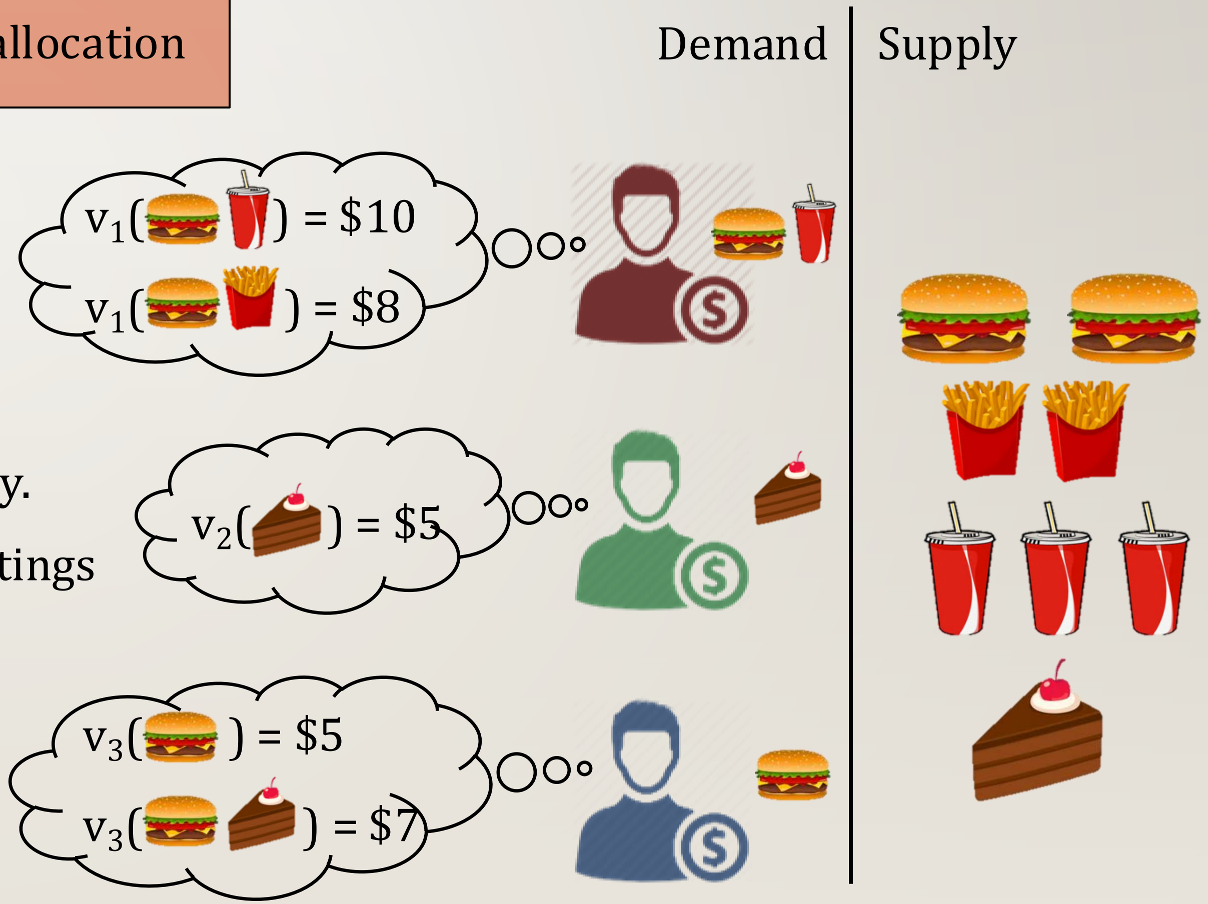

- Additional challenge:
  - Incentivize participants to report values truthfully.
  - Can achieve in computationally simple offline settings
    using payments
    via the "VCG mechanism"

Some computationally simple settings

- Matching a.k.a. unit-demand buyers

$$v(S) = \max_{i \in S} v_i$$

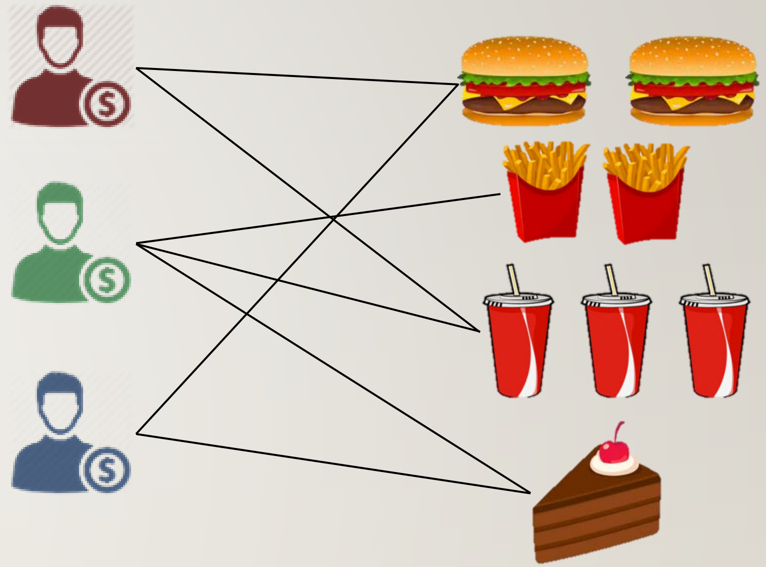

- Interval packing

  Items have a total ordering and buyers only assign values to intervals.

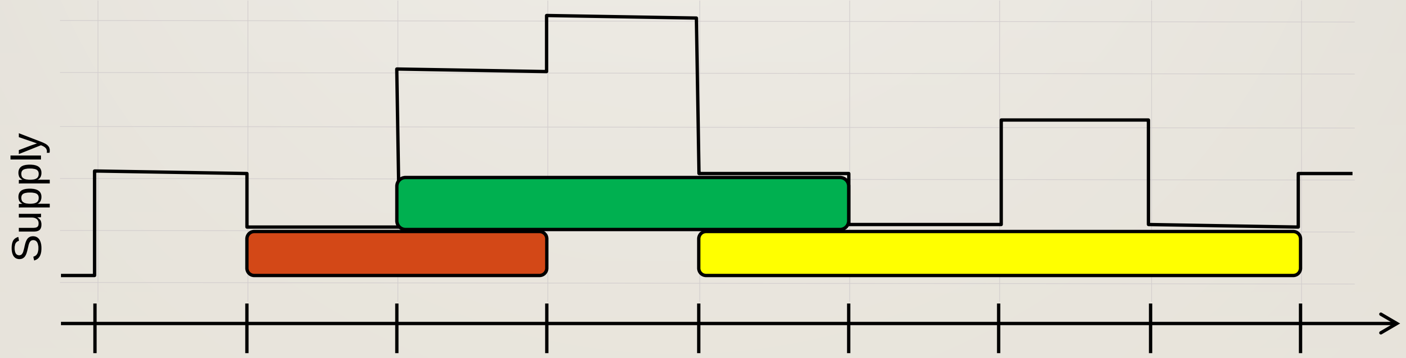

Focus of this work: online arrivals

- Upon each arrival, the algorithm makes an irrevocable allocation (and charges an irrevocable price)

Examples:

- Online shopping sites
- Airline/hotel reservations
- Spot markets for cloud resources

…

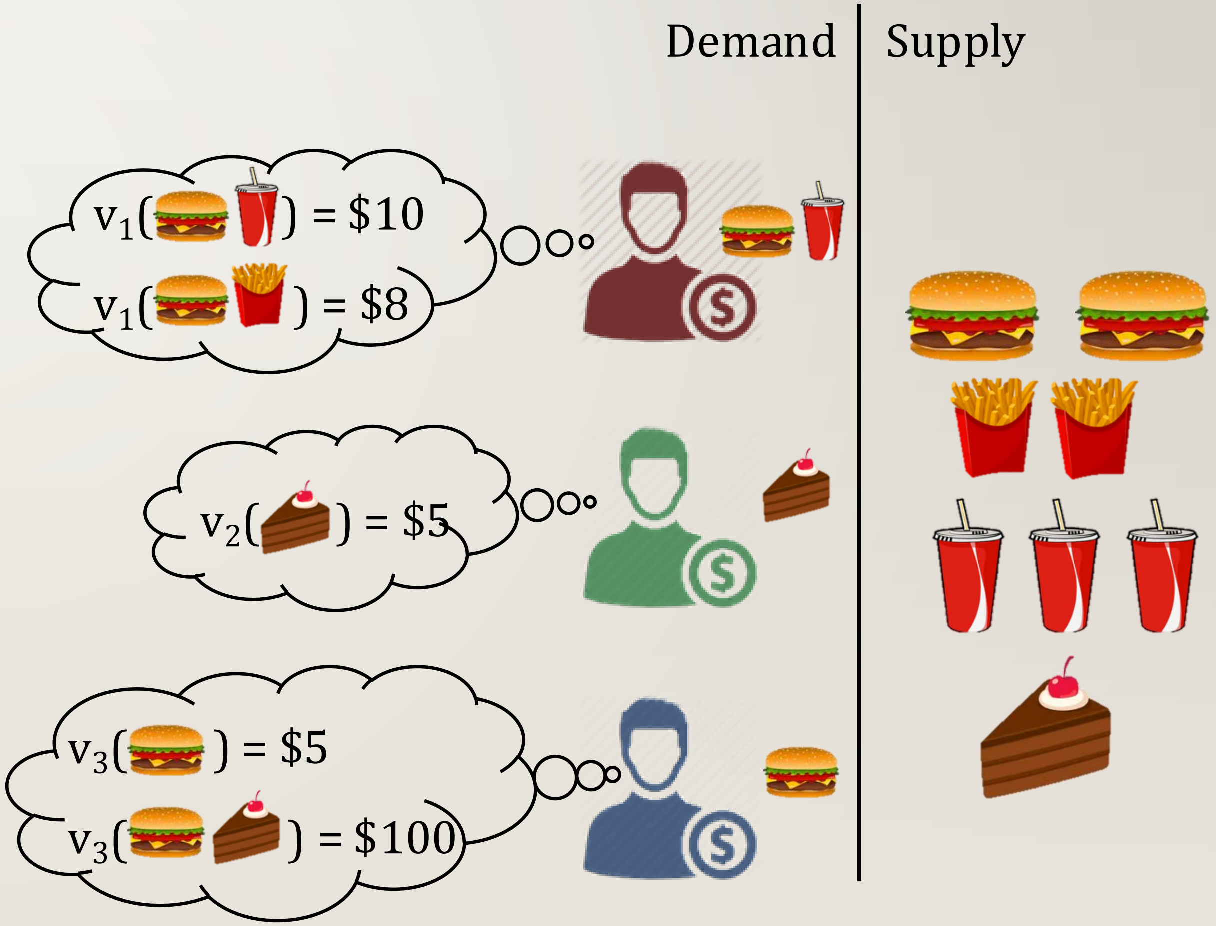

Focus of this work: online arrivals

- Upon each arrival, the algorithm makes an irrevocable allocation (and charges an irrevocable price)

Algorithmic challenge: online algorithm that is competitive against hindsight OPT

Economic challenge: buyers shouldn't misreport values <u>and</u> shouldn't delay arrival.

Without further assumptions, no reasonable solution.

A stochastic-online model

- n buyers; values drawn from known independent distributions; arrive in adversarial order.

Initial input:
n distributions

Stochastic step:
values instantiated

Adversarial step:
order of arrival

A simpler stochastic-online model

- n buyers with known values; "appear" independently with known probability; arrive in adv. order.

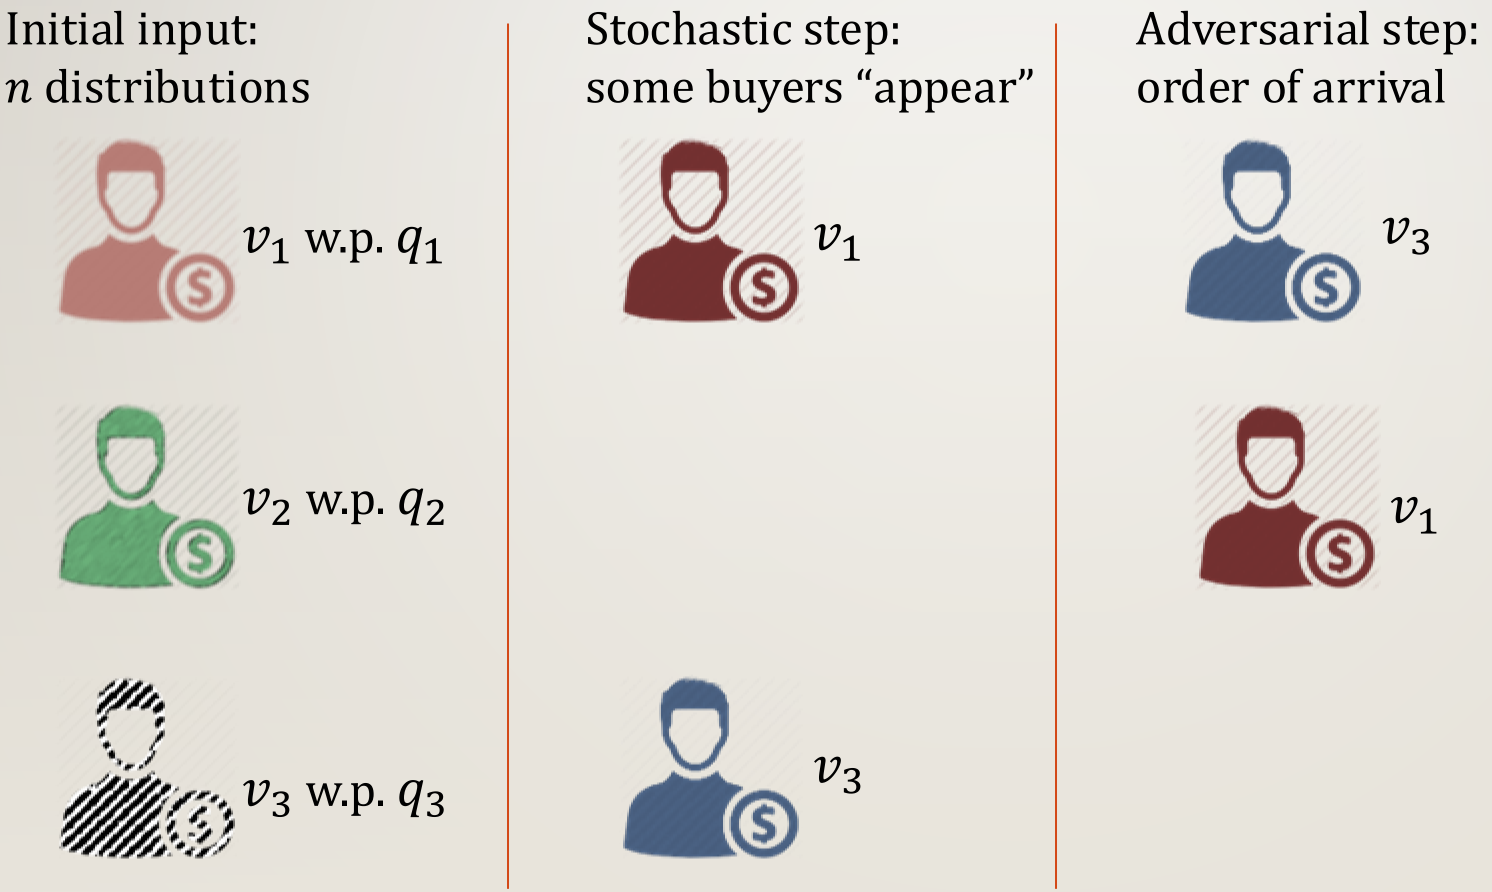

Questions

- Can we design an online allocation algorithm that is competitive against the hindsight OPT?

$$\boxed{\text{Hindsight OPT} = \mathrm{E}_{v\sim\mathcal{D}}[\mathrm{SW}(\text{optimal}-\text{alloc}(v))]}$$

Gap?

- Can we design an online mechanism that is incentive compatible and competitive?

An outline for the rest of the talk

- The simple single-item case a.k.a. prophet inequality
- Pricing as an online mechanism
- Two approaches:
  - Dual prices
  - Balanced prices
- Challenges
- Overcoming challenges for interval scheduling
- Results & open questions

The simplest setting: single item for sale a.k.a. prophet inequality

- An online algorithm is just a stopping rule.

U[0,2]

1 w.p. ½
3 w.p. ½

100 w.p. 0.01
0 w.p. 0.99

The simplest setting: single item for sale a.k.a. prophet inequality

- An online algorithm is just a stopping rule.

1

$^1/_\epsilon$ w.p. ϵ
0 w.p. $1 - \epsilon$

$$\mathrm{OPT} = \epsilon . \frac{1}{\epsilon} + (1 - \epsilon) . 1 = 2 - \epsilon$$

$$\mathrm{ALG} = 1$$

- No online algorithm can be better than 2-competitive.
- [Krengel Sucheston Garling'78; Samuel-Cahn'84]: There exists a threshold t such that accepting the first reward ≥ t is 2-competitive.

Economic interpretation: The threshold is like a price tag on the item.

- Obviously incentive compatible; No gap between online algorithms and online mechanisms!

(Aside) Further work on prophet inequalities

- Prophet inequality ≡ Stochastic online subset selection subject to a feasibility constraint.

  [Hajiaghayi Kleinberg Sandholm'07; C. Hartline Malec Sivan'10; Kleinberg Weinberg'12; Alaei Hajiaghayi Liaghat'12; Azar Kleinberg Weinberg'14; Duetting Kleinberg'15; Feldman Svensson Zenklusen'15; Rubinstein'16; Duetting Feldman Kesselheim Lucier'17; Rubinstein Singla'17; …]

- Online resource allocation is different:
  - Commit to allocation; not just accept/reject
  - Commit to payment

The grocery store mechanism

- Each buyer purchases her “favorite bundle” while supplies last.
- Obviously incentive compatible (assuming prices don’t change over time)

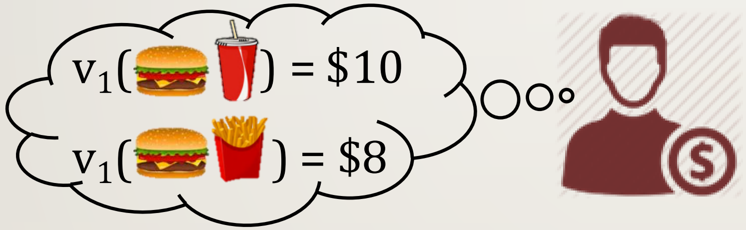

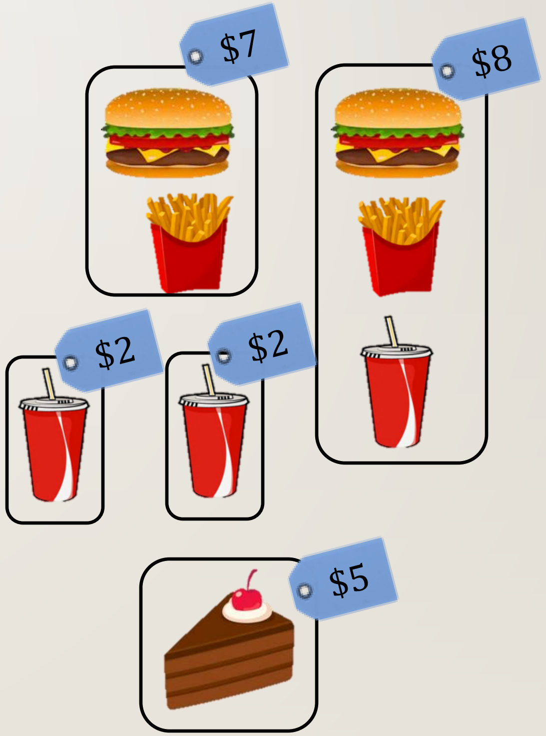

The buyer purchases $\underset{S}{\operatorname{argmax}}(v(S) - p(S))$

Prices as dual variables

$v_{i,S}$: buyer i's value for set S
q_i: buyer i's probability of arrival
$x_{i,S}$: buyer i's prob. of receiving set S

PRIMAL

$$\max \sum_{i,S} x_{i,S} v_{i,S}$$

subject to:

$$\sum_{S} x_{i,S} \le q_i \quad \text{for all buyers } i$$

$$\sum_{i,\, S \ni j} x_{i,S} \le 1 \quad \text{for all items } j$$

$$x_{i,S} \ge 0 \quad \text{for all } i \text{ and } S$$

DUAL

$$\min \sum_{j} p_j + \sum_{i} u_i q_i$$

Seller's revenue

Buyers' utility

subject to:

$$\sum_{j \in S} p_j + u_i \ge v_{i,S} \quad \text{for all } i, S$$

$$u_i, p_j \ge 0 \quad \text{for all } i, j$$

In an optimal solution, $u_i = \max_S \left(v_{i,S} - \sum_{j \in S} p_j \right)$

- Complementary slackness implies $x_{i,S} > 0$ iff S is one of i's favorite bundles under the pricing p.

How good are dual prices?

Problem 1: dual prices are usually too low.

1 w.p. $1-\epsilon$ $\quad$ $^{1}/_{\epsilon^2}$ w.p. ϵ

$\mathrm{LP} = \epsilon.\frac{1}{\epsilon^2} + (1-\epsilon).1 > {}^{1}/_{\epsilon}$

$\mathrm{OPT} = \epsilon.\frac{1}{\epsilon^2} + (1-\epsilon)^2.1 > {}^{1}/_{\epsilon}$

Dual price = 1

$\mathrm{ALG} = (1-\epsilon).1 + \epsilon^2.\frac{1}{\epsilon^2} < 2$

Problem 2: complementary slackness is not always useful due to the stochasticity of arrivals.

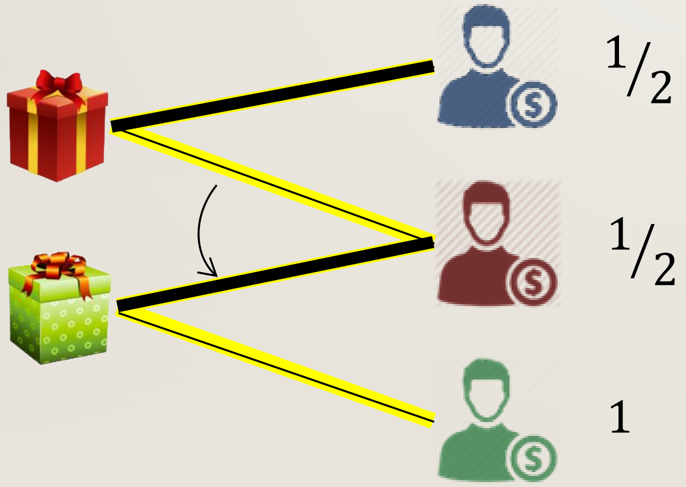

Buyer shifts preferences based on availability and has a new favorite set.

A second approach: balanced prices

[Feldman Gravin Lucier'15]
[Kleinberg Weinberg'12]

- Contribution of item j to optimal SW $= \sum_i v_{i,j} x_{i,j}$.
- Set the price for item j to $p_j = {}^1\!/_2 \sum_i v_{i,j} x_{i,j}$.

- The prices are not too low:

  If item j gets sold, then seller's revenue from j = p_j

- The prices are not too high:

  If item j does not get sold, then any buyer i's utility $\geq v_{i,j} - p_j$.

  $\Rightarrow$ Total utility "attributed to item j" $\geq \sum_i x_{i,j}(v_{i,j} - p_j) = p_j$.

- Social Welfare = Seller's revenue + buyers' utility

$$\max \sum_{i,j} x_{i,j} v_{i,j}$$

subject to:

$$\sum_j x_{i,j} \leq q_i \quad \text{for all buyers } i$$

$$\sum_i x_{i,j} \leq 1 \quad \text{for all items } j$$

$$x_{i,j} \geq 0 \quad \text{for all } i \text{ and } j$$

How good are balanced prices?

- Not very good when buyers desire bundles. [FGL’15]

n items

$v_1(\text{any single item}) = \$1$

$v_2(\text{all } n \text{ items}) = \$(n-1)$
$v_2(\text{any other set}) = \$0$

- $\text{OPT} = \$(n-1)$
- Any “static item pricing” must price every item at $> \$1$ to exclude buyer 1 but then also excludes buyer 2.
- $\text{ALG} = \$1$

The interval scheduling setting

- Items are totally ordered; buyers desire intervals.

Two main results:

- Can design competitive mechs using balanced prices if we allow bundling

[C. Miller Teng'19]

- Can design competitive mechs using dual prices if we have large supply (and some other assumptions)

[C. Devanur Holroyd Karlin Martin Sivan'17]

The static bundle pricing grocery store mechanism

- Partition supply into “bundles” and price each bundle.
- Each buyer purchases her “favorite bundle” while supplies last.

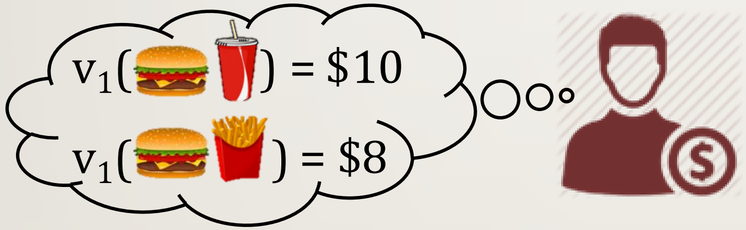

The buyer purchases $\operatorname*{argmax}_S (v(S) - p(S))$

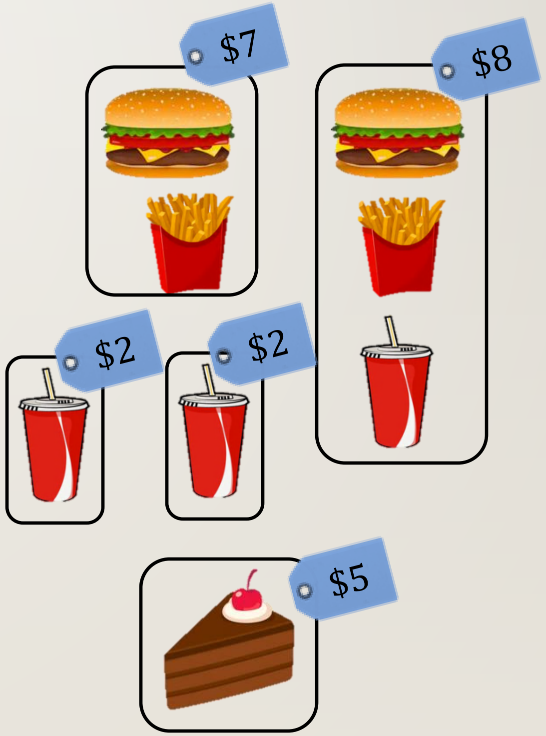

Leveraging balanced prices through bundling

- Key idea: partition items into bundles and pretend each buyer is unit-demand over the bundles. Then leverage FGL's balanced pricing approach.

- A fractional unit allocation is:

1. A partition of items into bundles
2. A fractional matching from buyers to bundles

- Question: how does the new value ($\sum_{i,S} y_{i,S} v_{i,S}$) compare to the original LP value ($\sum_{i,S} x_{i,S} v_{i,S}$)?

Original fractional solution

$$
\begin{aligned}
&\sum_{S} x_{i,S} \le q_i && \text{for all buyers } i \\
&\sum_{i,\, S \ni j} x_{i,S} \le 1 && \text{for all items } j \\
&x_{i,S} \ge 0 && \text{for all } i \text{ and } S
\end{aligned}
$$

Fractional unit allocation

$\mathcal{B}$ is a partition of items into bundles

$$
\begin{aligned}
&\sum_{S \in \mathcal{B}} y_{i,S} \le q_i && \text{for all buyers } i \\
&\sum_{i} y_{i,S} \le 1 && \text{for all sets } S \in \mathcal{B} \\
&x_{i,S} \ge 0 && \text{for all } i \text{ and } S
\end{aligned}
$$

Leveraging balanced prices through bundling

- Key idea: partition items into bundles and pretend each buyer is unit-demand over the bundles. Then leverage FGL's balanced pricing approach.

- A fractional unit allocation is:

1. A partition of items into bundles
2. A fractional matching from buyers to bundles

> Main claim: When all buyers desire bundles of size up to L, there exists a fractional unit allocation that loses at most an $O\left({}^{\log L}/_{\log\log L}\right)$ fraction of social welfare.

$\Longrightarrow O\left({}^{\log L}/_{\log\log L}\right)$-competitive mechanism using static bundle pricing.

- [Im Wang'13]: No online algorithm can be $o\left({}^{\log L}/_{\log\log L}\right)$-competitive for interval scheduling.

Large supply settings

- Can we obtain better competitive ratios when we have multiple copies of items?

<u>Prophet inequality for identical items</u>

Static pricing: $1 - O\left(\sqrt{\log k / k}\right)$ [Hajiaghayi Kleinberg Sandholm'07]

General online mechs: $1 - 1/\sqrt{k+3}$ [Alaei'11]

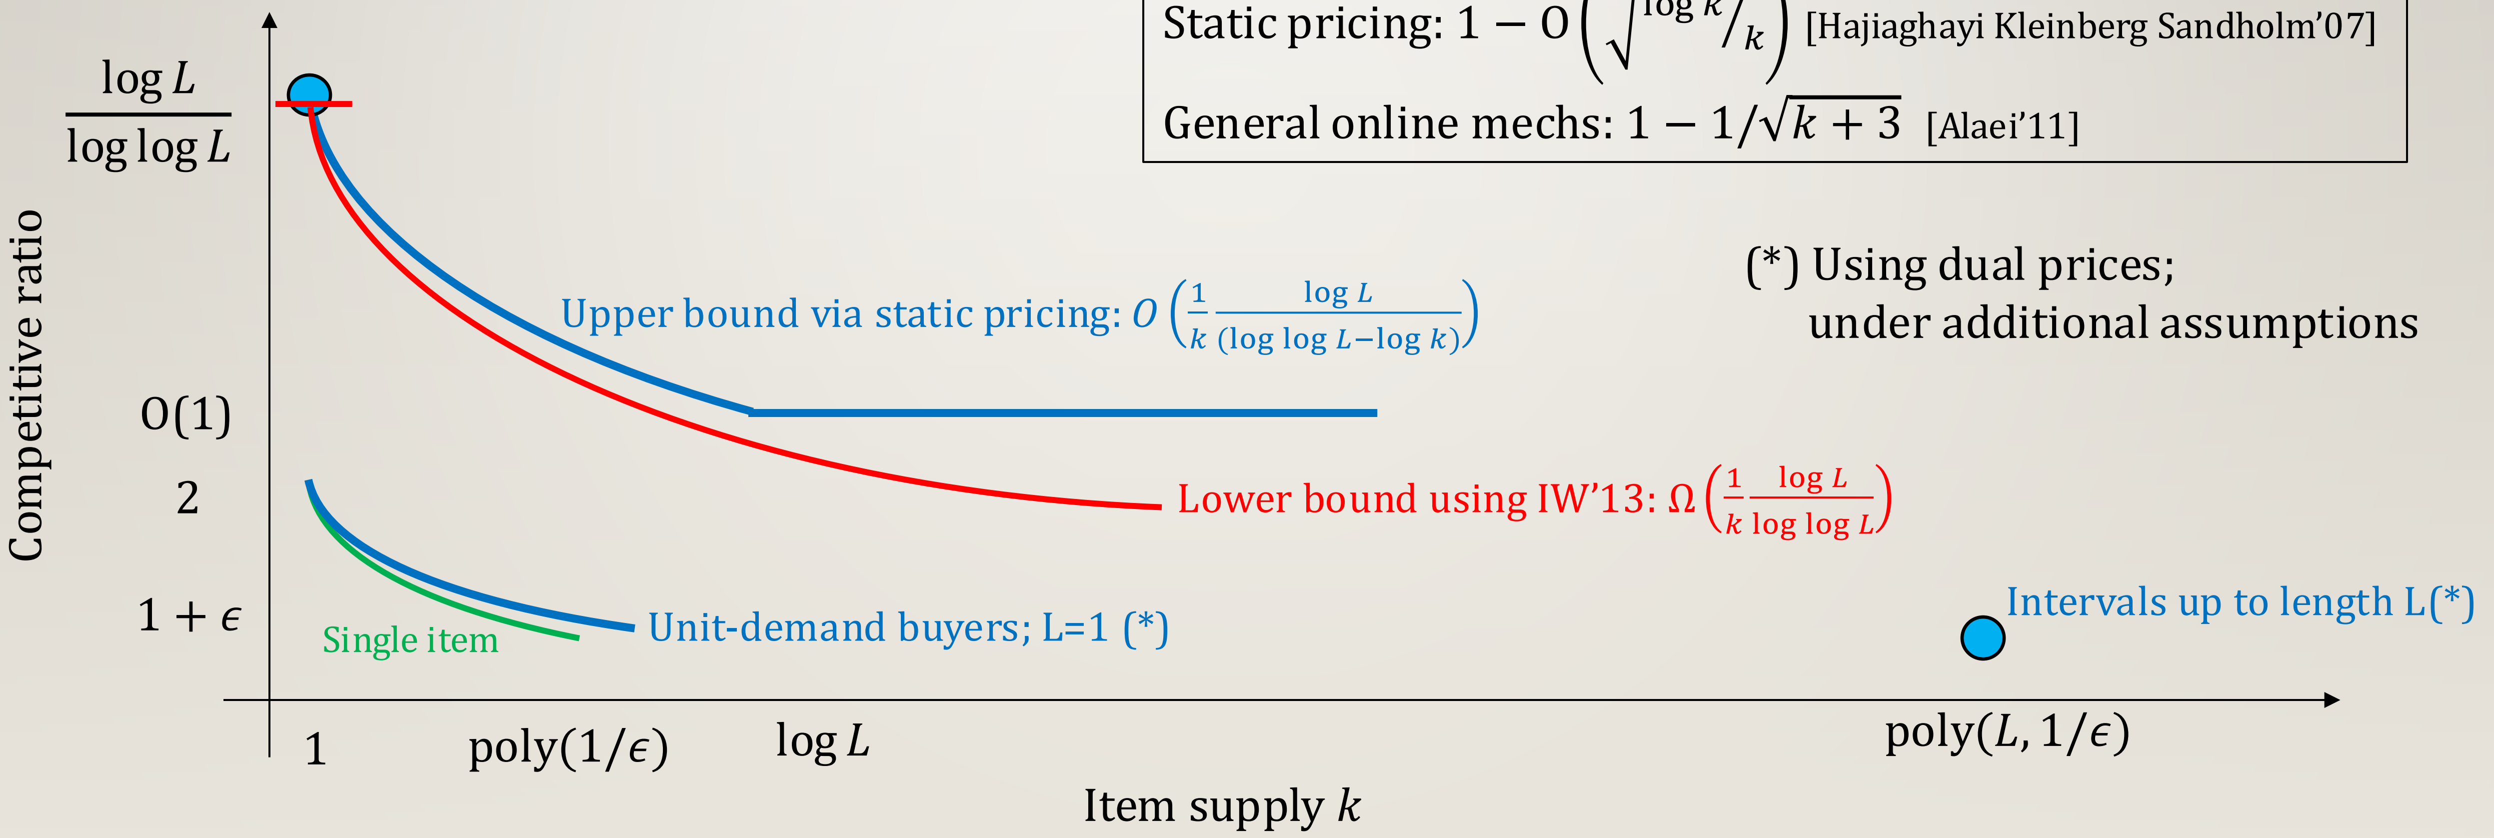

(*) Using dual prices; under additional assumptions

Leveraging dual prices by appealing to large supply

- Unit-demand buyers; n items; $k \approx \text{poly}(1/\epsilon)$ copies of every item.

Basic idea:

- Scale down the supply by a $(1-\epsilon)$ factor; compute dual prices; make primal-based assignments.
- Then for any individual item, $\Pr[\text{realized demand} > {}^{1}/_{1-\epsilon}\ \text{fractional demand}]$ is small.

failure event

- Problem: when the failure event happens for an item, future buyers' favorite item may change;
  We can no longer implement primal-based assignments.
  $\Longrightarrow$ Likelihood of other failure events may increase.

- Our approach: track buyers' preferences orderings over items.

The "forwarding" graph

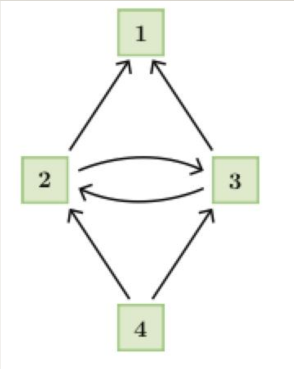

- The forwarding graph as a function of dual prices:
  - Nodes are items
  - $\exists$ a directed edge from j_1 to j_2 if j_1 and j_2 appear consecutively in some buyer's preference ordering.
- Issue: when j_1 has a failure event, it "forwards" buyers along its outgoing edges, potentially causing failure events at their endpoints.
- Assumption: Each buyer is single-valued over some contiguous set of items.

  That is, $v_i(j) = v_i \;\; \forall j \in S_i$ and $v_i(j) = 0 \; \forall j \notin S_i$, where S_i is an interval.
- Observation: Under these assumptions this forwarding graph behaves like a low in-degree tree.

Implication: For any item, Pr[realized demand $> {}^{1}/_{1-\epsilon}$ fractional demand] is $\leq \epsilon$.

$\Longrightarrow$ Every buyer gets their LP allocation with high probability.

Bounding the failure probabilities on a forwarding graph

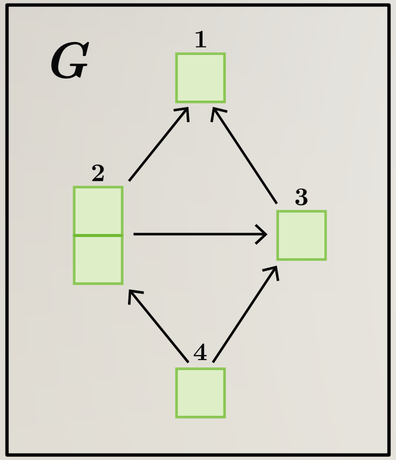

1. The forwarding graph

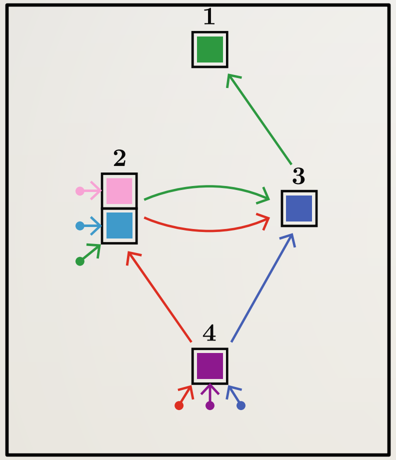

2. Instantiation of buyers; forwarding paths

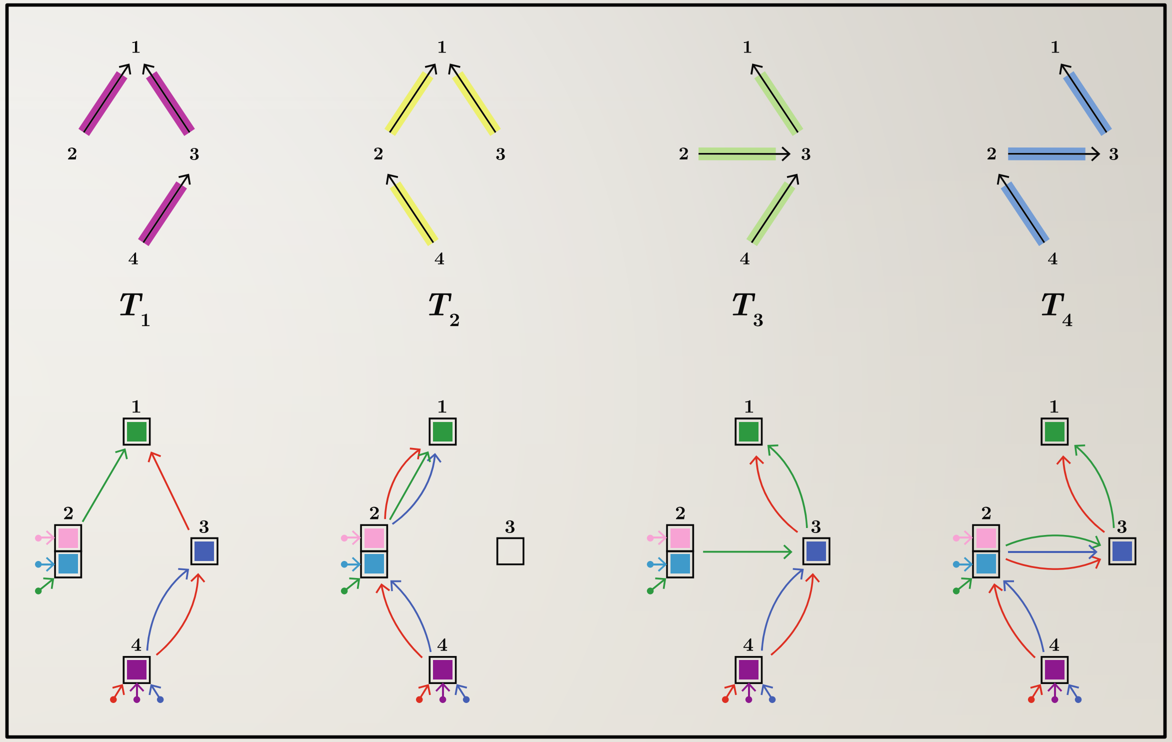

3. Consider all possible forwarding subtrees of G. The load in picture 2 can be bounded by the load in one of these subtrees.

4. Tree networks permit an inductive analysis. Failure probabilities depend on the in-degrees of nodes.

A summary of results for the interval scheduling setting

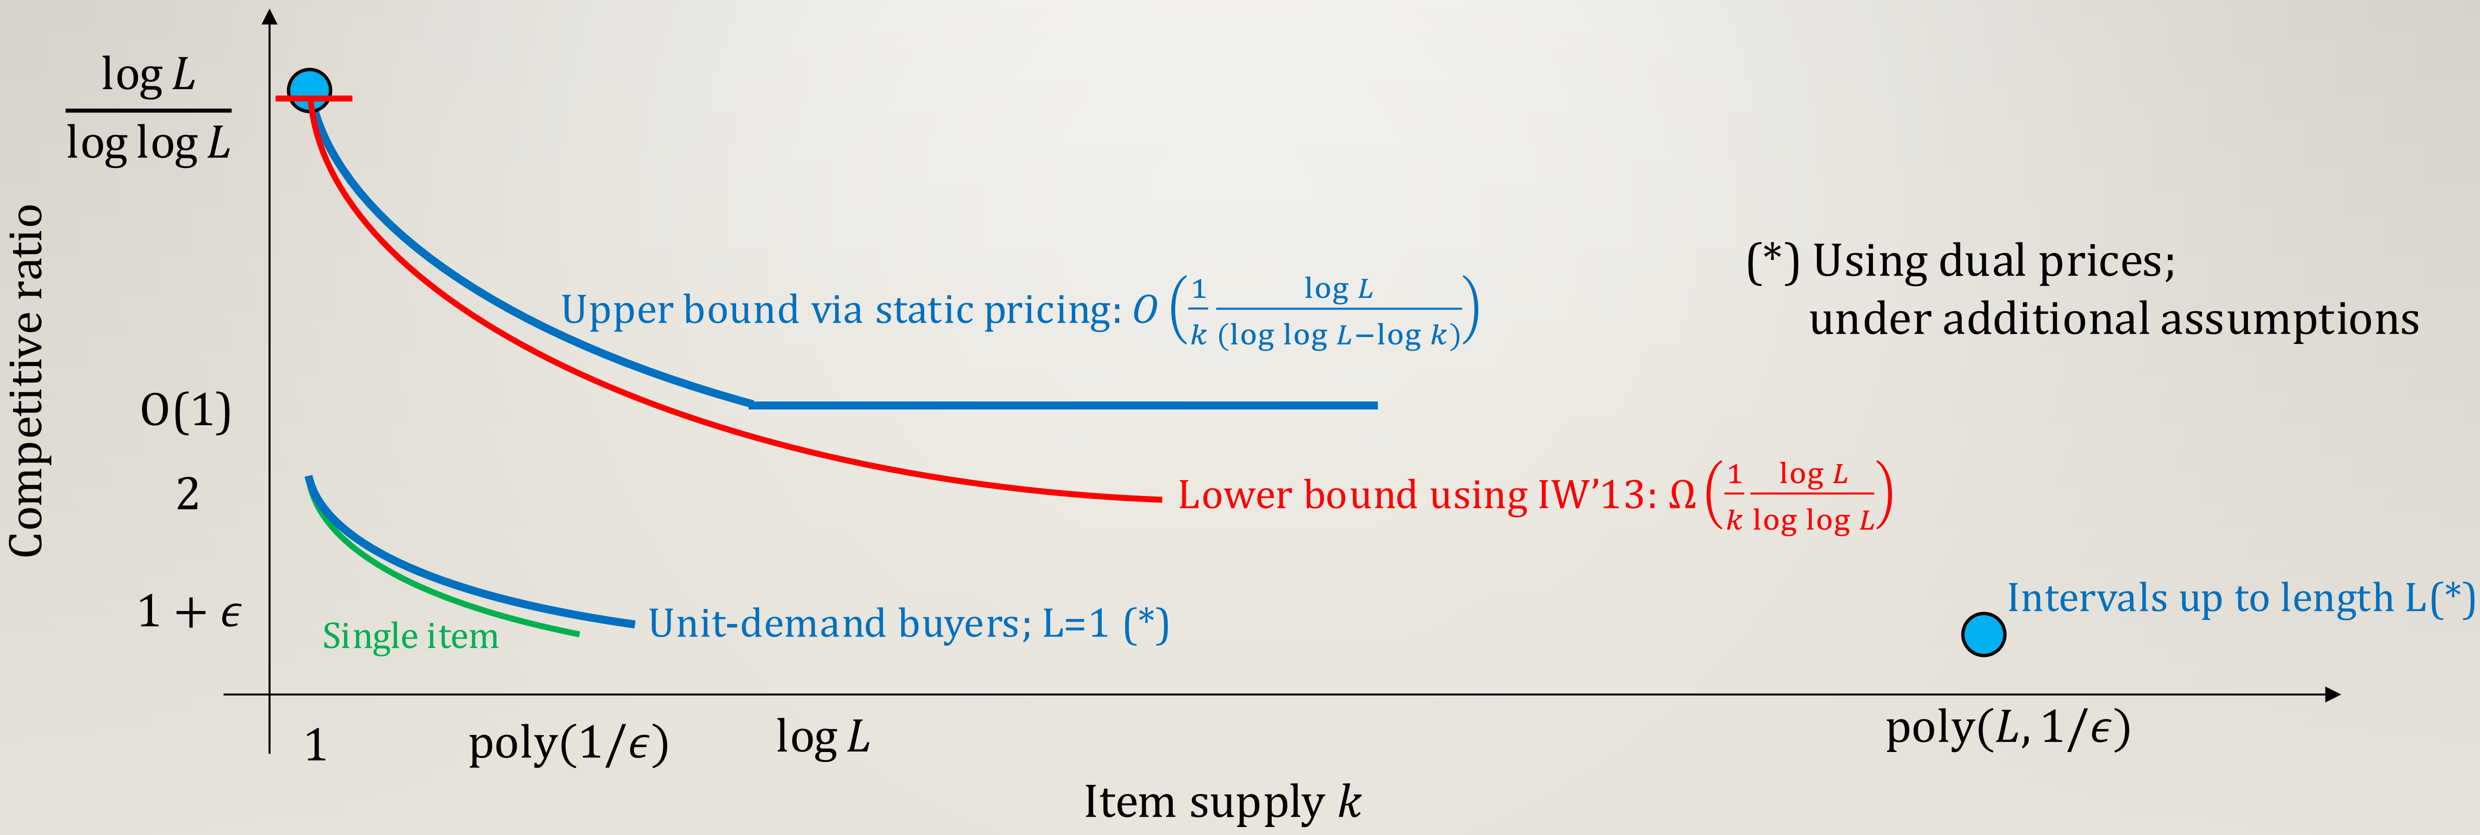

Other results…

- Balanced item prices:
  - Unit demand buyers – tight competitive ratio of 2 [Feldman Gravin Lucier'15]
  - XOS or fractionally subadditive valuations – tight competitive ratio of 2 [Feldman Gravin Lucier'15]
  - MPH-k values – factor of k; tight for static item pricing [Duetting Feldman Kesselheim Lucier'17]
- Balanced bundle prices:
  - Interval packing – tight competitive ratio [C. Miller Teng'19]
  - Packing paths in trees – nearly tight ratio [C. Miller Teng'19]
- Dual prices:
  - Interval packing with large supply and some other assumptions

In many settings, static pricing is near optimal!

(*) within constant factors

Some open directions

- Can bundling help obtain tight competitive ratios for other kinds of valuation functions?

  e.g. the MPH hierarchy?

  single-minded buyers? (The LP may be too weak.)

- A $(1-\epsilon)$ competitive mechanism for more general large supply settings?

  What about unit-demand buyers?

  Static pricings won't help. [C. Teng]

- Can we efficiently compute/learn prices?

- Other online mechanism design problems? E.g. revenue maximization.