Mechanisms for resource allocation

SHUCHI CHAWLA

UNIVERSITY OF WISCONSIN-MADISON

Question: how to allocate scarce resources among multiple parties?

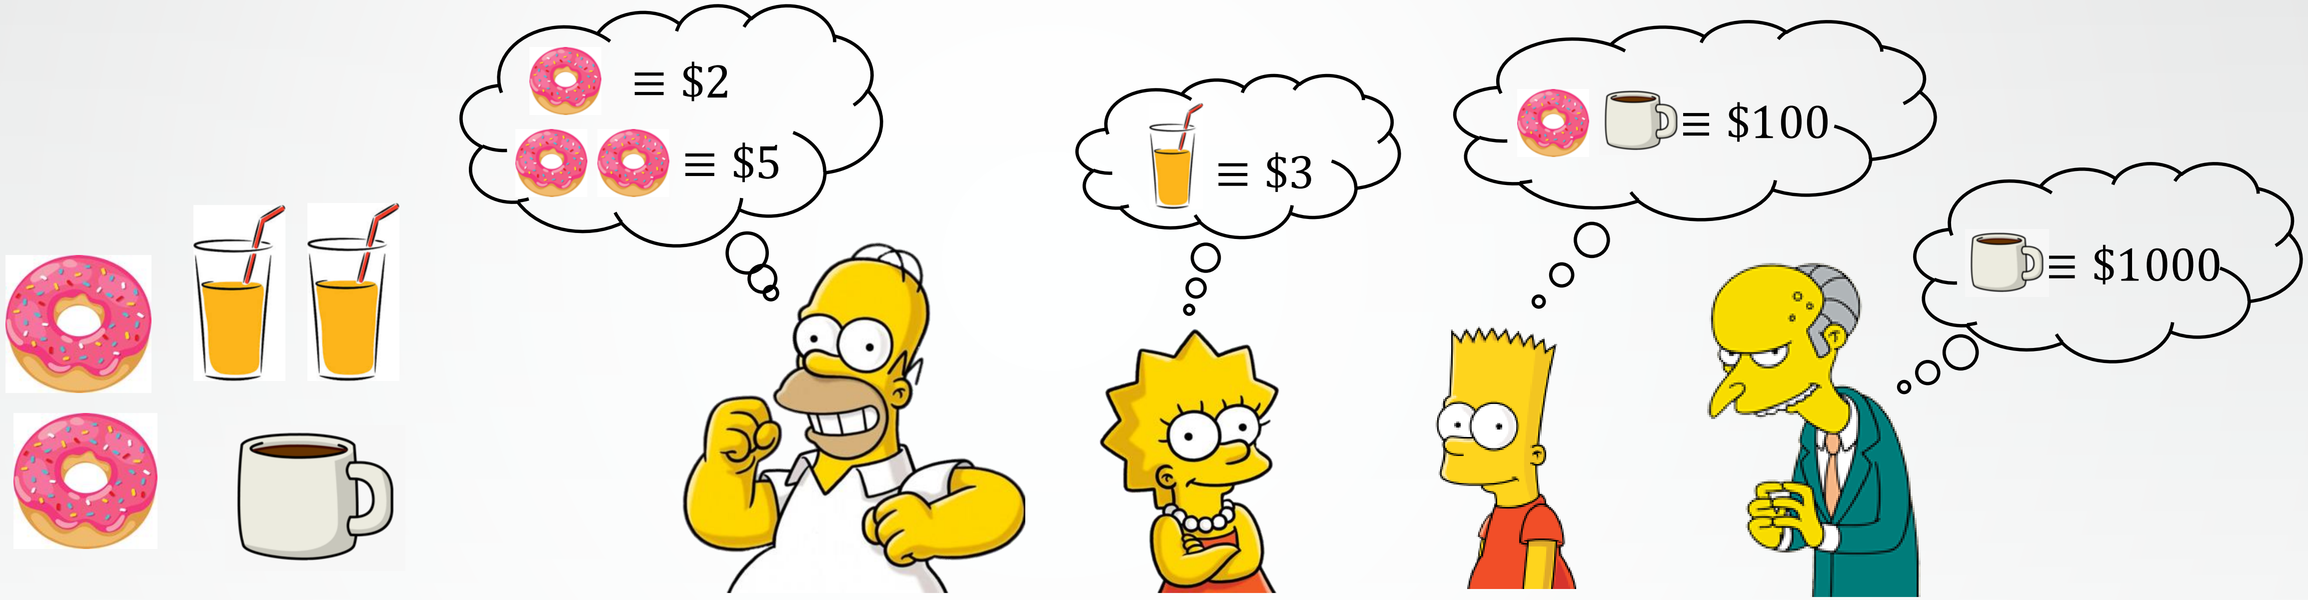

What if participants can lie and subvert rules?

What if participants arrive over time and future demand is unknown?

Objectives

SOCIAL WELFARE

$= \sum_{\text{participants } i} (\text{value } i \text{ gets from allocation})$

Competitive analysis:
compare against hindsight optimal allocation

REVENUE

$= \sum_{\text{participants } i} (\text{payment made by } i)$

Approximation:
compare against revenue-optimal mechanism

Some applications

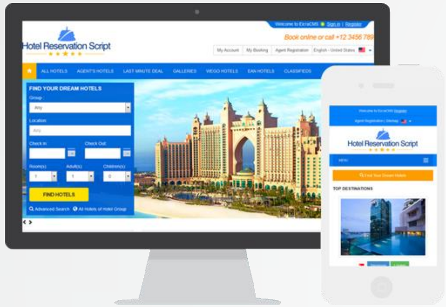

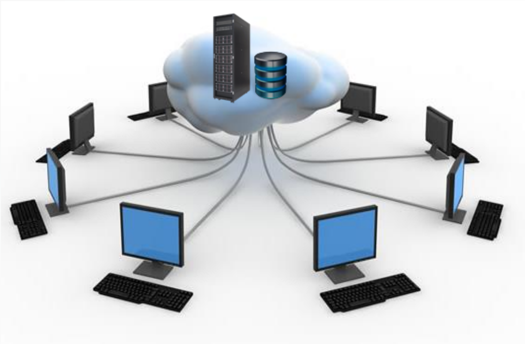

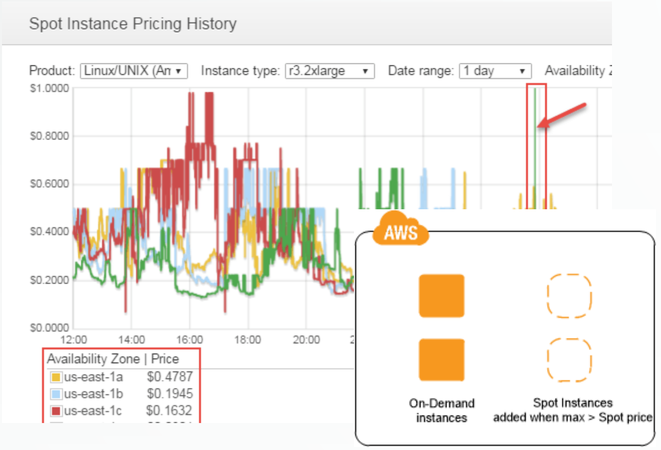

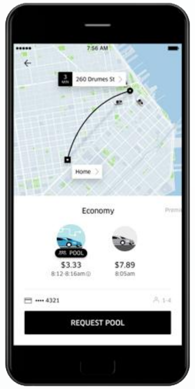

Two important settings:

- Scheduling jobs on a machine
  - Items ≡ "time slots"
  - Buyers ≡ jobs

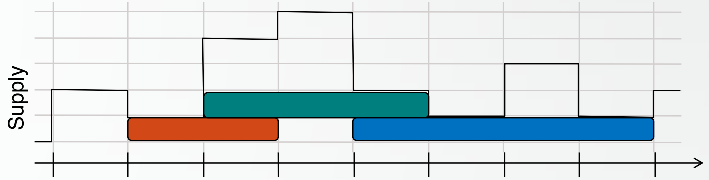

- Routing on a network
  - Items ≡ edges
  - Buyers ≡ paths

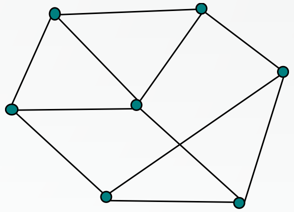

Assumptions

- Buyers' true values are unknown but their value distributions are known

$$\text{Hindsight OPT} = \mathrm{E}_{v_i \sim F_i}\left[\max_{(S_1,\dots,S_n)} \sum_i v_i(S_i)\right]$$

- Buyers arrive in an online fashion
- Buyers can lie about their values and delay their arrival

Value function of buyer i: $v_i \sim F_i$.

Adversarial order of arrival.
When buyer i arrives, his identity and distribution are revealed.

Algorithm solicits values from buyers when they arrive.
Buyers are rational:
maximize (value from alloc - payment)

We will think of truthful mechanisms as algorithms with structural constraints.

A simple class of algorithms: posted pricing

- When each buyer arrives, algorithm offers each subset of items at a certain price.
- The buyer purchases $\underset{S}{\operatorname{argmax}}(v(S) - p(S))$.

Always truthful!

Special types of pricings:

Anonymous: prices don't depend on buyers' identity

Non-adaptive: prices don't evolve over time

Order-oblivious: prices don't depend on ordering of buyers

(Anonymous, Non-adaptive, Order-oblivious) → Static pricing

Item pricing: additive pricing function

Some questions

- How well does simple posted pricing approximate welfare/revenue?
- Are there better (truthful) mechanisms?
- Are there better (non-truthful) algorithms?
- Can we optimize over the class of all pricings?

Maximizing social welfare

Key takeaway:

In many settings, static pricings are optimal-within-constant-factors across all online algorithms.

Outline

- Why do prices perform well?
  - A primal-dual view
  - Issues with dual prices
- Fix # 1: balanced prices
  - Warm up: single item prophet inequality.
  - Feldman-Gravin-Lucier generalization.
  - Extension to scheduling & routing
- Fix # 2: dual prices for large supply settings
  - Warm up: single item with copies.
  - Extension to scheduling
- Summary of results; open questions

Approach # 1: 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 $\sum_{i,S}$ **Social Welfare**

subject to:

$\sum_{S}$ **No overprovisioning** ...rs i

$\sum_{i,S\ni j}$ **Demand ≤ Supply** ...ems j

$x_{i,S} \geq 0$ for all i and S

DUAL

$$\min \sum_j p_j + \sum_i u_i q_i$$

Seller's revenue

Buyers' utility

subject to:

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

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

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.

value $= 1$
Prob. Arrival$= 1-\epsilon$

value $= 1/\epsilon^2$
Prob. Arrival$= \epsilon$

$$\text{LP} = \epsilon.\frac{1}{\epsilon^2} + (1-\epsilon).1 \approx 1/\epsilon$$

$$\text{OPT} = \epsilon.\frac{1}{\epsilon^2} + (1-\epsilon)^2.1 \approx 1/\epsilon$$

Dual price = 1

$$\text{ALG} = (1-\epsilon).1 + \epsilon^2.1/\epsilon^2 < 2$$

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

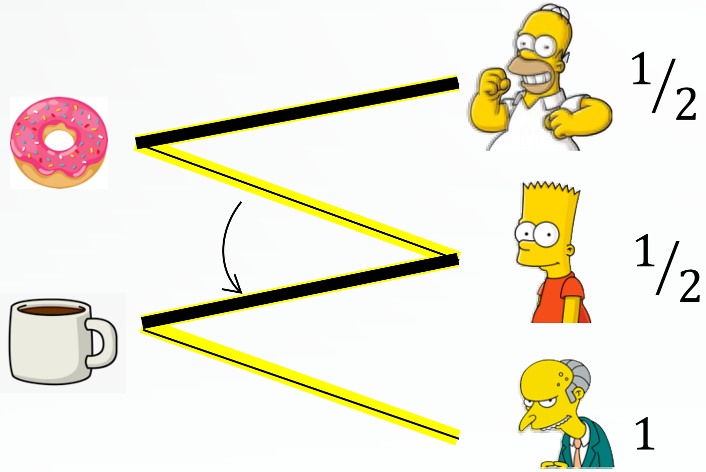

Buyer shifts preferences based on availability and has a new favorite set.

The single item prophet inequality




- Samuel-Cahn'84: There exists a static price p such that allocating item to the first buyer with value above p gets a competitive ratio of 2.
- Set p so that $\Pr[\exists i : v_i \geq p] = 1/2$.

- $\mathrm{OPT} = \mathrm{E}\left[\max_i v_i\right] \leq \mathrm{E}\left[\max_i (p + (v_i - p)^+)\right] \leq p + \sum_i \mathrm{E}[(v_i - p)^+]$
- $\mathrm{ALG} \geq p \, \underbrace{\Pr[\text{item is sold}]}_{= 1/2} + \sum_i E[(v_i - p)^+] \, \Pr[\text{item is offered to } i]$ $\geq 1/2$

$\Pr[\text{item is offered to } i] \geq \Pr[\text{item is unsold at the end}] = 1 - \Pr[\text{item is sold}]$

$$\Longrightarrow \mathrm{ALG} \geq \frac{1}{2}\mathrm{OPT}$$

Observations:

- Can also pick $p = \frac{1}{2}\mathrm{OPT}$.
- Tight!


value $= 1$
Prob. Arrival$= 1$

value $= 1/\epsilon$
Prob. Arrival$= \epsilon$

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

$\mathrm{ALG} = 1$

General (combinatorial) prophet inequalities

- Each buyer has a value $v_i \sim F_i$.
- Buyers arrive online; algorithm observes v_i; makes accept/reject decisions.
- The algorithm faces a feasibility constraint $\mathcal{F}$. Must ensure: set of accepted agents $\in \mathcal{F}$.

- Constant factor competitive ratios in many settings: k-unit, matroids, knapsack, matching, … [Chawla Hartline Malec Sivan'10, Alaei'11, Kleinberg Weinberg'12, Feldman Svensson Zenklusen'15, Dutting Kleinberg'15], …

- Different from our setting:
  - We select the actual allocation, not just accept/reject decisions.
  - Want a simple pricing-based algorithm

Approach # 2: balanced prices (for unit-demand buyers)

[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:

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

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

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

Approach # 2: balanced prices

[Feldman Gravin Lucier'15]
[Dutting Feldman Kesselheim Lucier'17]

Type of value function	Competitive ratio	Lower bound
Unit-demand or additive	2	2
XOS (max over additive functions)	2	2
MPH-L	$4L$-2	L

Limitations of balanced item prices

- Poor approximation when values have complementarities

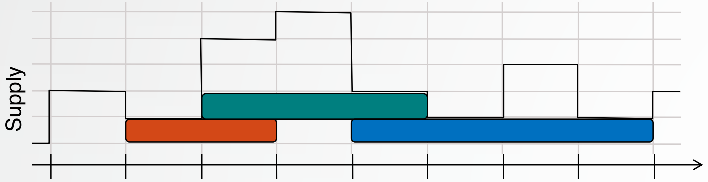

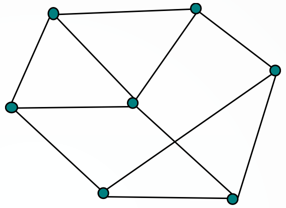

$v_1(\text{any single item}) = \$1$

$v_2(\text{all } n \text{ items}) = \$(n-1)$

$v_2(\text{any other set}) = \$0$

Any static item pricing must price every item at > 1 to exclude buyer 1 but then also excludes buyer 2.

$\text{OPT} = n - 1; \quad \text{ALG} = 1$

Approach # 3: Balanced bundle prices

[Chawla Miller Teng'19]

- 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

Original fractional solution

$$\sum_{S} x_{i,S} \leq q_i \quad \text{for all buyers } i$$
$$\sum_{i,\, S \ni j} x_{i,S} \leq 1 \quad \text{for all items } j$$
$$x_{i,S} \geq 0 \quad \text{for all } i \text{ and } S$$

Fractional unit allocation

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

$$\sum_{S \in \mathcal{B}} y_{i,S} \leq q_i \quad \text{for all buyers } i$$
$$\sum_{i} y_{i,S} \leq 1 \quad \text{for all sets } S \in \mathcal{B}$$
$$x_{i,S} \geq 0 \quad \text{for all } i \text{ and } S$$

Approach # 3: Balanced bundle prices

[Chawla Miller Teng'19]

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

Original fractional solution

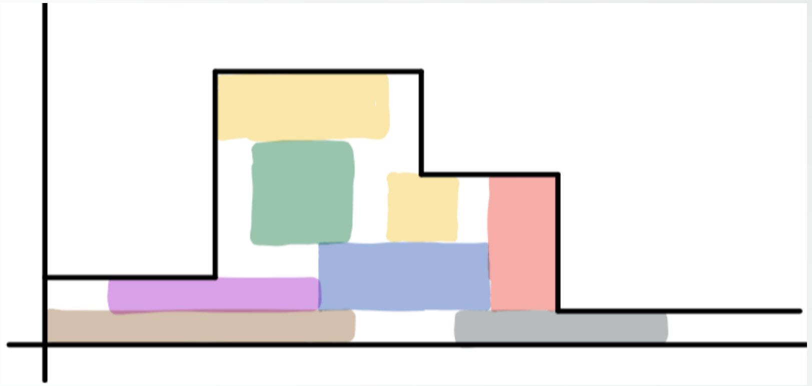

- A fractional unit allocation is:

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

Fractional unit allocation

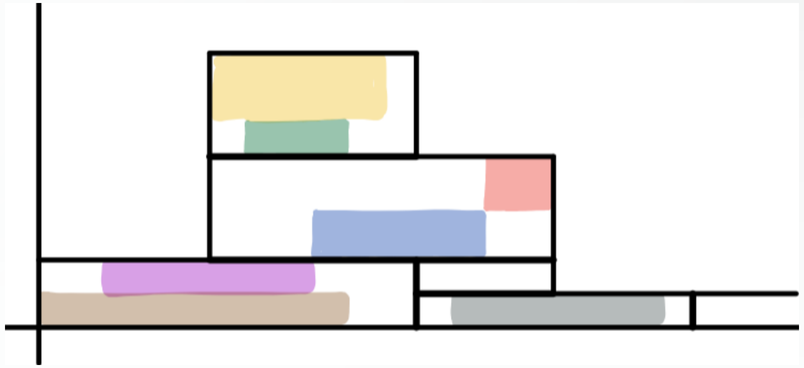

- Key lemma: show that the new value ($\sum_{i,S} y_{i,S} v_{i,S}$) is not much smaller than the original LP value ($\sum_{i,S} x_{i,S} v_{i,S}$).
- Can do for intervals and paths on trees while losing logarithmic factors.

Approaches #2 & #3: Balanced item and bundle prices

Value functions	Competitive ratio	Lower bound	Technique
Additive or unit-demand	2 [FGL'15]	2	Balanced item prices
XOS	2 [FGL'15]	2	Balanced item prices
MPH-L	4L-2 [DFKL'17]	L	Balanced item prices
Interval scheduling over intervals of size $\leq L$	$O\left(\frac{\log L}{\log\log L}\right)$ [CMT'19]	$\Omega\left(\frac{\log L}{\log\log L}\right)$	Balanced bundle prices
Routing on trees with values $\in [1, H]$	$O(\log H)$ [CMT'19]	$\Omega\left(\frac{\log H}{\log\log H}\right)$	Balanced bundle prices

Approaches #2 & #3: Balanced item and bundle prices

Value functions	Competitive ratio	Lower bound	Technique
Additive or unit-demand	2 [FGL'15]	2	Balanced item prices
XOS	2 [FGL'15]	2	Balanced item prices
MPH-L	$4L$-2 [DFKL'17]	L	Balanced item prices
Interval scheduling over intervals of size $\leq L$	$O\left(\frac{\log L}{\log\log L}\right)$ [CMT'19]	$\Omega\left(\frac{\log L}{\log\log L}\right)$	Balanced bundle prices
Routing on trees with values $\in [1, H]$	$O(\log H)$ [CMT'19]	$\Omega\left(\frac{\log H}{\log\log H}\right)$	Balanced bundle prices
Interval scheduling with capacities k	$O\left(\frac{\log L}{k\log\log L}\right)$ [CMT'19]	$\Omega\left(\frac{\log L}{k(\log\log L - \log k)}\right)$	Balanced bundle prices
Routing on trees with capacities k	$O\left(\frac{\log H}{k}\right)$ [CMT'19]	$\Omega\left(\frac{\log H}{k\log\log H}\right)$	Balanced bundle prices

Can we beat the 2 in large supply settings?

k-unit prophet inequality:

[Hajiaghayi Kleinberg Sandholm'07]

- Find price p such that $\mathrm{E}[|\{i: v_i \geq p\}|] \approx k - \sqrt{k \log k}$
- p is the dual price for the LP on the right
- w.h.p. item does not get sold out

$$\max \sum_i x_i v_i$$
subject to:
$$x_i \leq q_i \quad \text{for all } i$$
$$\sum_i x_i \leq k - \sqrt{k \log k}$$
$$x_i \geq 0 \quad \text{for all } i$$

$$\Rightarrow 1 - O\left(\sqrt{\frac{\log k}{k}}\right) \text{ competitive ratio.}$$

- Tight! [Ghosh Kleinberg'16]

(for pricings; for mechanisms, can get $1 - O(1/\sqrt{k})$ [Alaei'11])

Approach # 4: Dual prices for large supply interval scheduling

Assumptions:

- Each job has a fixed length; value.
- Wants to get scheduled within a certain time window.
- Supply at any time t is at least k

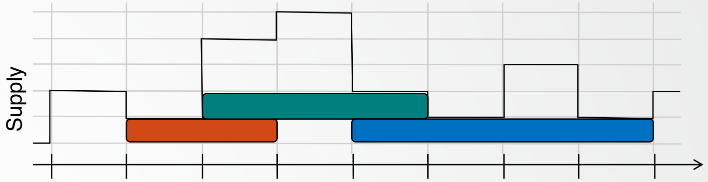

$v_{i,S}$: buyer i's value for set S

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:

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

$$\sum_{i,S \ni j} x_{i,S} \leq (1-\epsilon)k_j \quad \text{for all items } j$$

$$x_{i,S} \geq 0 \qquad \text{for all } i \text{ and } S$$

DUAL

$$\min \sum_{j} p_j (1-\epsilon)k_j + \sum_{i} u_i q_i$$

subject to:

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

$$u_i, p_j \geq 0 \qquad \text{for all } i, j$$

Approach # 4: Dual prices for large supply interval scheduling

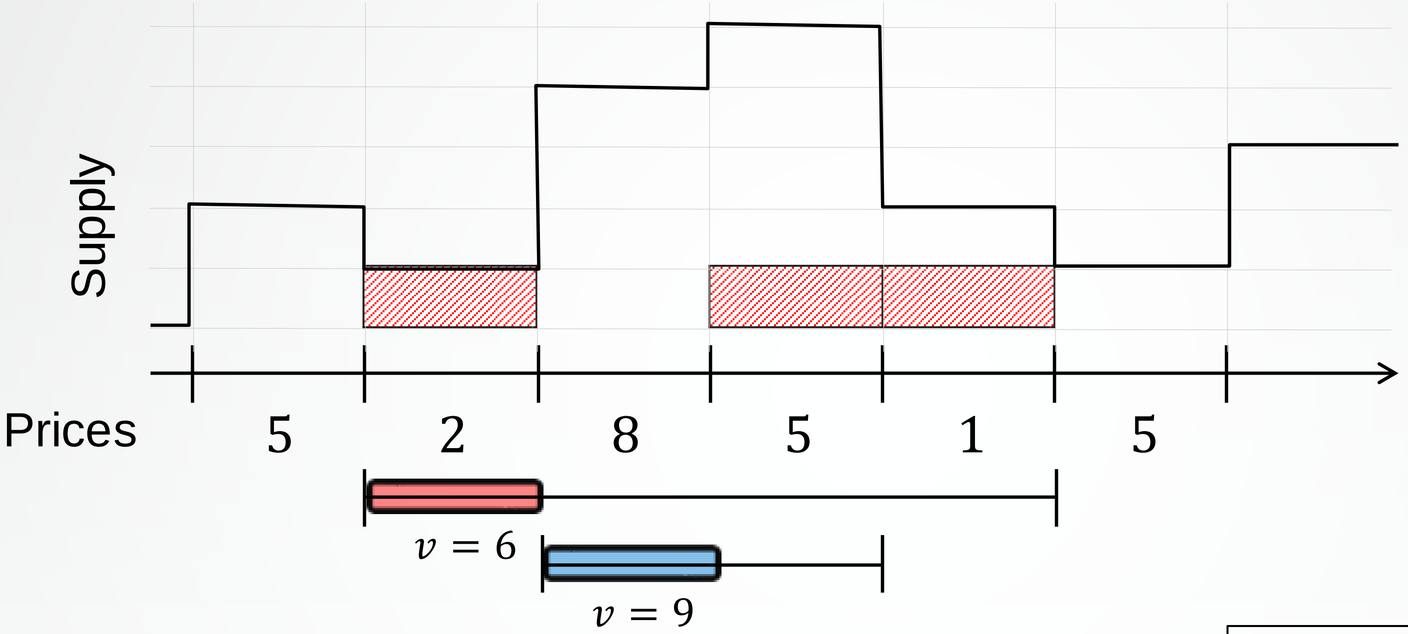

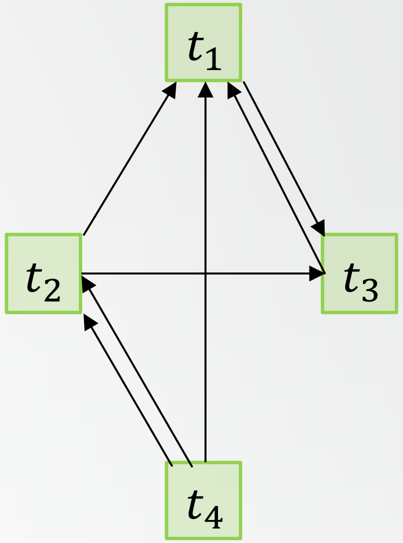

A_t = number of "arrivals" at time t

= # jobs instantiated that LP places at time t

= # jobs instantiated whose first try is t

Main claim: For all t, w.p. $1 - \epsilon'$, $X_t < B_t$.

Implication: Each job occupies its LP position w.p. $1 - \epsilon'$.

$$E[A_t] = \sum_j x_{jt} \leq (1 - \epsilon) B_t \ \forall t$$

$\Rightarrow$ w.p. $1 - \epsilon$, $A_t < B_t$

X_t = number of jobs that ever try t = arrivals + "forwards"

Want $X_t < B_t$ w.h.p.; Problem: bad events are correlated across t.

Dual prices: bounding the failure probabilities on a forwarding graph

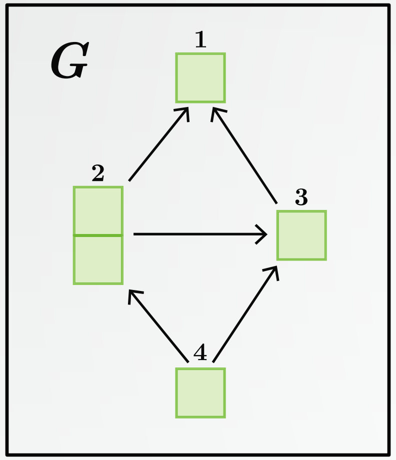

The forwarding graph

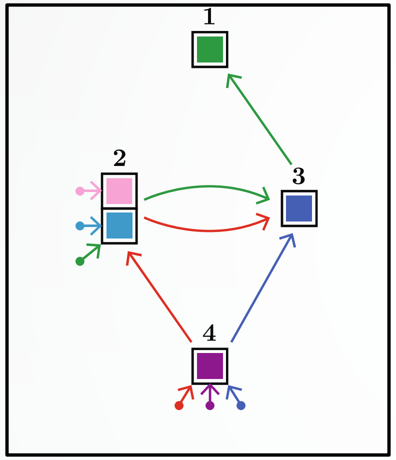

Instantiation of buyers;
forwarding paths

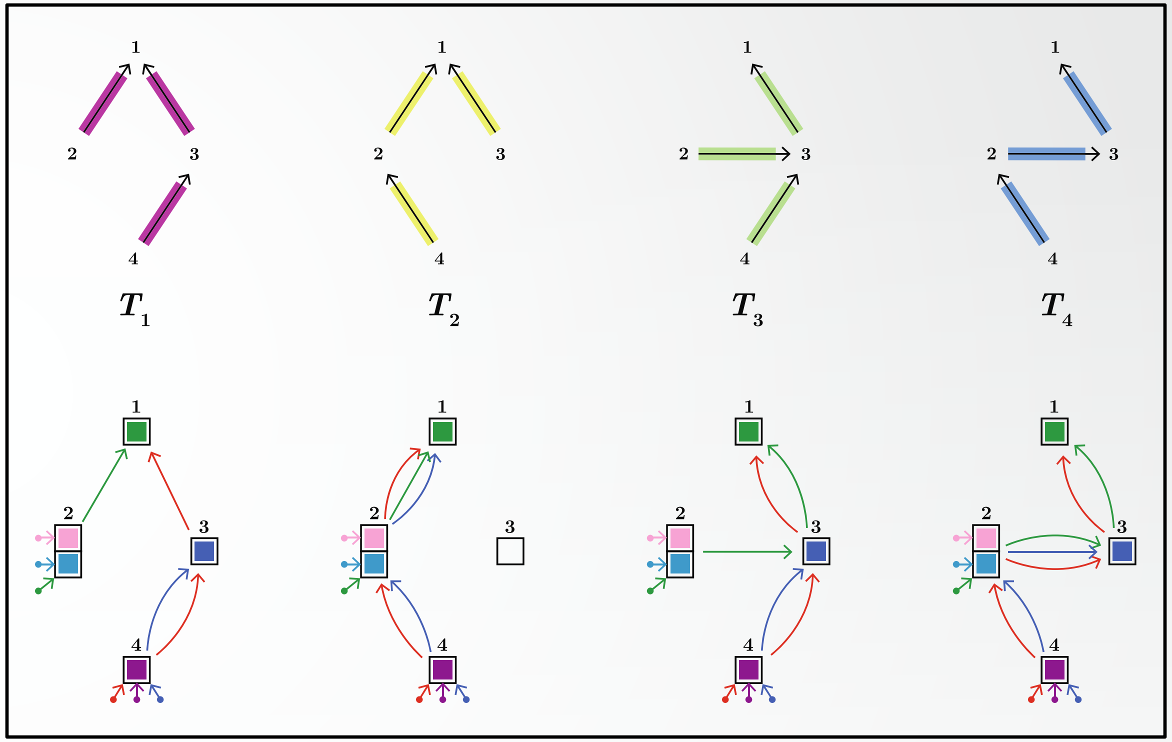

Consider all possible forwarding subtrees of G. The load in picture 2 can be bounded by the load in one of these subtrees.

Dual prices: combining subtrees into one tree

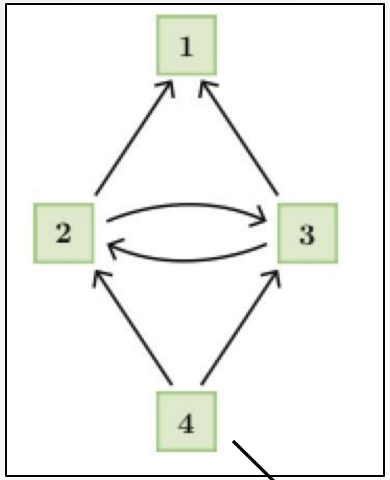

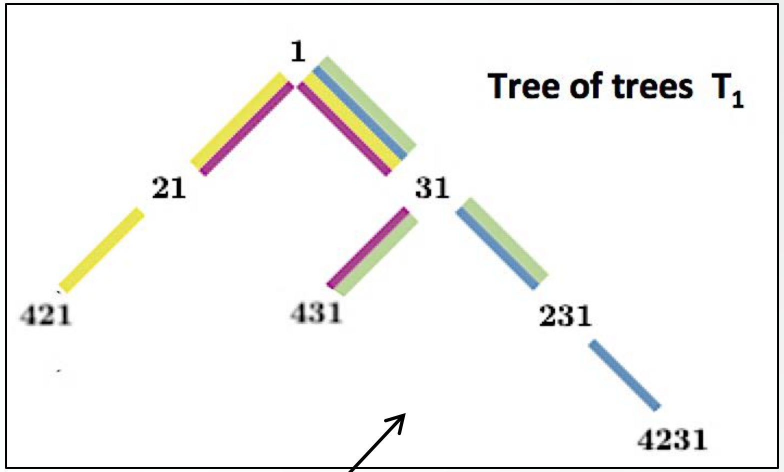

Worst case

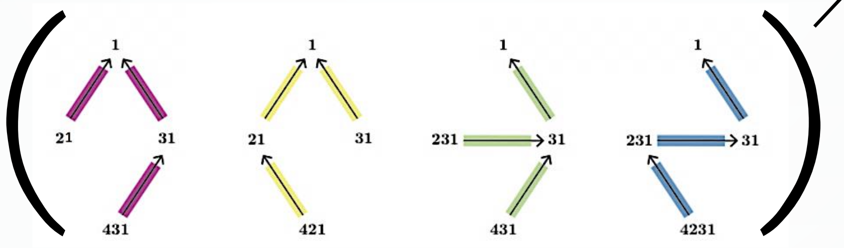

Tree networks permit an inductive analysis. Failure probabilities depend on the in-degrees of nodes.

Approach # 4: Dual prices for large supply interval scheduling

[Chawla Devanur Holroyd Karlin Martin Sivan '17]

There exist a price schedule such that if jobs are unit length(*), and,

$$k_j \geq \Omega\left(\frac{1}{\epsilon^2}\log\frac{1}{\epsilon}\right) \text{ for all } j$$

Then the expected social welfare achieved is at least $(1-\epsilon)$ times the Hindsight-OPT.

(*) Need $k_j \geq \Omega\left(\frac{L^6}{\epsilon^3}\log\frac{1}{\epsilon}\right)$ when jobs are of length up to L.

Approaches #2, #3, & #4: Balanced prices and dual prices

Value functions	Competitive ratio	Lower bound	Technique
XOS	2 [FGL'15]	2	Balanced item prices
MPH-L	$4L$-2 [DFKL'17]	L	Balanced item prices
Interval scheduling over intervals of size $\leq L$	$O(\log L / \log\log L)$ [CMT'19]	$\Omega(\log L / \log\log L)$	Balanced bundle prices
Routing on trees with values $\in [1, H]$	$O(\log H)$ [CMT'19]	$\Omega(\log H / \log\log H)$	Balanced bundle prices
Interval scheduling with capacities k	$O(\log L / k \log\log L)$ [CMT'19]	$\Omega(\log L / k\ (\log\log L - \log k))$	Balanced bundle prices
Routing on trees with capacities k	$O(\log H\ / k)$ [CMT'19]	$\Omega(\log H / k \log\log H)$	Balanced bundle prices
k-unit	$1 - O(\sqrt{\log k\ / k})$ [HKS'07]		Dual prices
Interval scheduling special case; capacity k	$1 - O(\text{poly}(L \log k / k))$ [CDH+'17]		Dual prices

Some open directions

- Beat the factor of 2 for unit-demand with large supply?
- Beat the factor of 2 for subadditive values with large supply?

- Routing on general graphs?
- General valuations on small sets?
  - LP is too weak

- Why do static prices perform so well?

Revenue maximization: a different story

Key takeaways:

Necessarily need non-anonymous mechanisms
Need to price random allocations
Even single-buyer setting is challenging

Revenue maximization: a different story

Simplest set-up: one buyer; two items

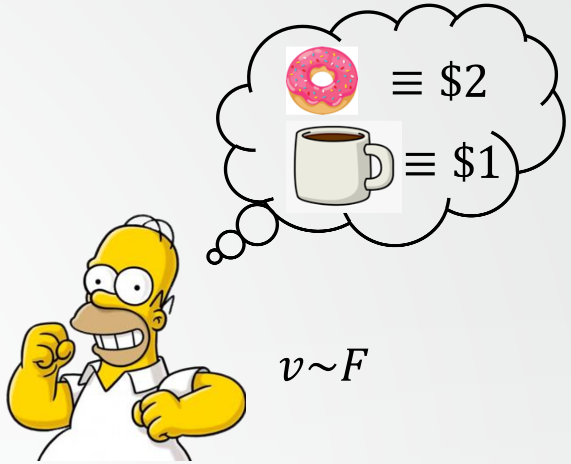

$v \sim F$

Optimal mechanism can be complicated:

- Offers random allocations, a.k.a. lotteries [Thannasoulis’05]
- Can have infinitely many options! [Hart Nisan’13]

Every near-optimal solution may be complicated

- No finite menu can provide a finite approximation!

[Briest C. Kleinberg Weinberg’10, Hart Nisan’13]

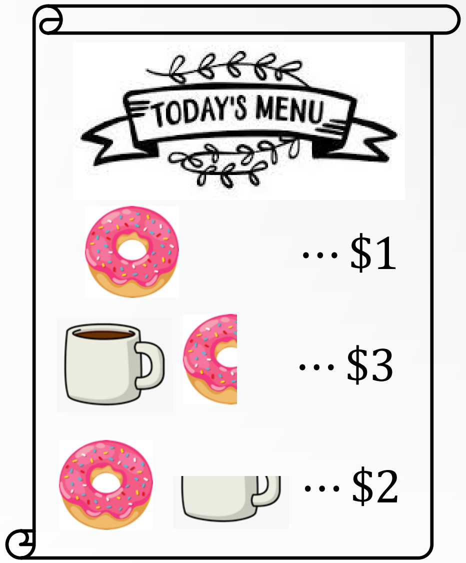

Revenue maximization: take two

[Chawla Teng Tzamos ’19]

Extra constraint on the mechanism:

cannot sell a bundle at a price higher than the sum of its constituents.

“Buy Many Constraint”

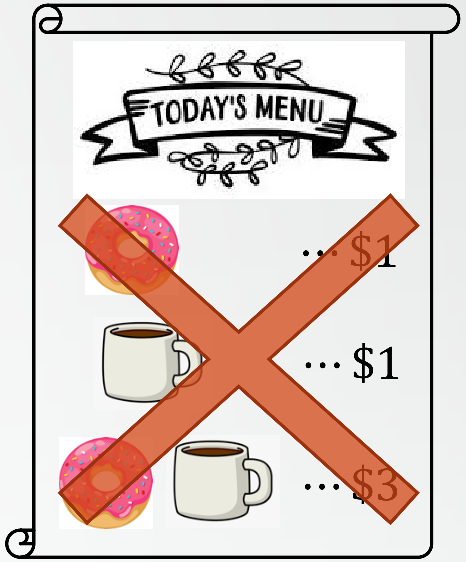

n = number of items

Theorem: Item-pricing is always an $O(\log n)$-approximation to the optimal buy-many mechanism.

(no matter the value distribution)

Some open directions

- When can pricing functions be approximated in revenue by simple pricing functions?
  - Any mechanism is a pricing function: f: random allocation → price.
  - Extend to $f(v) = f\left(\operatorname{argmax}_S\left(v(S) - f(S)\right)\right)$.
  - Want to find simple g such that $\mathrm{E}_v[g(v)] \geq$ (some fraction).$\mathrm{E}_v[f(v)]$

- Can we efficiently find an approximately revenue-optimal item pricing?

Pricing as a parameterized greedy algorithm

- Can prices be used to simplify algorithm design in non-strategic settings?
- Optimal prices depend on the instance – but can potentially be learned!

Acknowledgements










Buy-many mechanisms are not much better than item pricing.
C., Teng, and Tzamos. EC'19.

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.

Truth and regret in online scheduling.
C., Devanur, Kulkarni, and Niazadeh. EC'17.

Thanks for listening!