Prophet Inequalities for Online Combinatorial Auctions

Jose Correa
U. de Chile

ALGORITHMIC GAME THEORY, MECHANISM DESIGN, AND LEARNING, Torino, November 2022

(Mostly) Joint work with Andres Cristi

Classic Prophet Inequality

Ticket for a concert

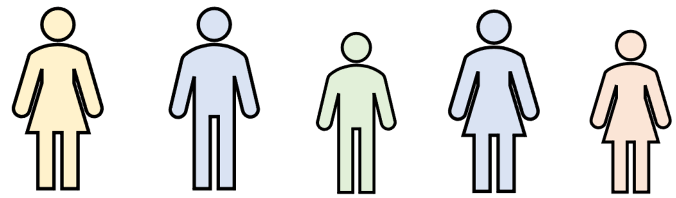

Sequence of n agents with independent valuations $v_i \sim F_i$

Theorem. [Krengel and Sucheston, Bull AMS'77]
We can get ½ of the expected optimal welfare.

E.g.

- Post price = Median of the r.v. $\max_i v_i$ [Samuel-Cahn, Ann Prob'84]
- Post price $p = \frac{1}{2} \cdot \mathbb{E}\left(\max_i v_i\right)$ [Kleinberg, Weinberg, STOC'12]
- Sample all distributions and use max as threshold [Rubinstein, Wang, Weinberg, ITCS'21]

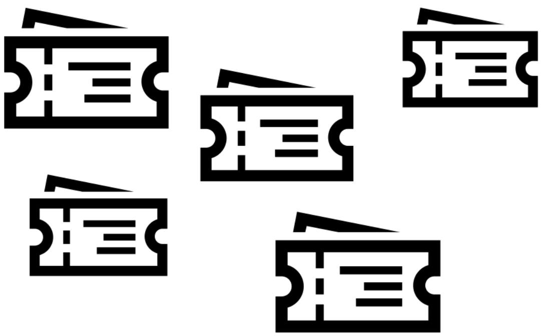

k tickets

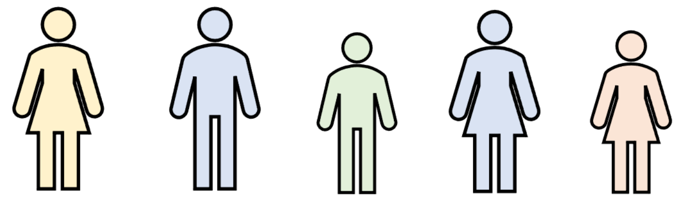

Sequence of n agents with independent valuations $v_i \sim F_i$

We can get $1 - O\left(\frac{1}{\sqrt{k}}\right)$ of the optimal welfare. [Alaei FOCS'11]

Tight fixed threshold algorithm recently found [Arnosti and Ma EC'22].

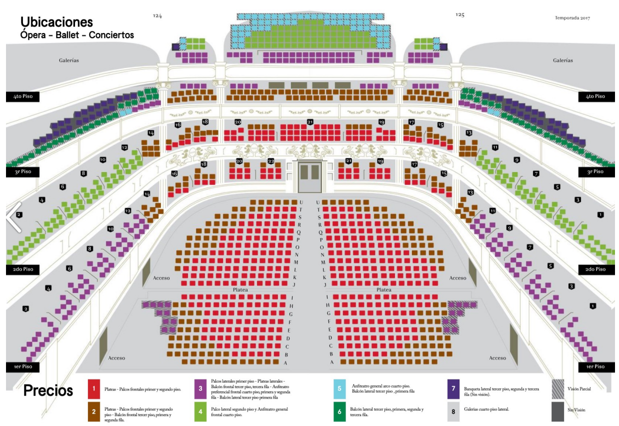

Set M with m heterogeneous items

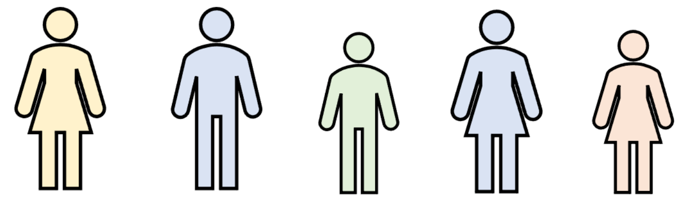

Sequence of n agents with independent valuations

$$v_i \sim F_i$$

$$v_i : 2^M \to \mathbb{R}_+$$

$v_i(S)$ is valuation of set $S \subseteq M$

Results (informal)

Theorem. [Correa and Cristi, 22+]

If there are no complementarities between items, then there is an online policy that gets $\frac{1}{6+\varepsilon}$ of the optimal welfare.

Theorem. [Correa, Cristi, Fielbaum, Pollner, Weinberg, IPCO'22]

If nobody wants more than d items, then there are *item prices* that guarantee $\frac{1}{d+1}$ of the optimal welfare (and we can compute them).

Online combinatorial auction

m heterogeneous items

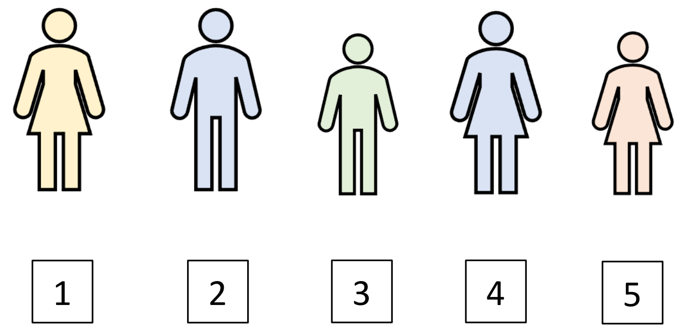

n agents with

monotone independent valuations

$v_i \sim F_i \qquad v_i: 2^M \to \mathbb{R}_+$

Online welfare

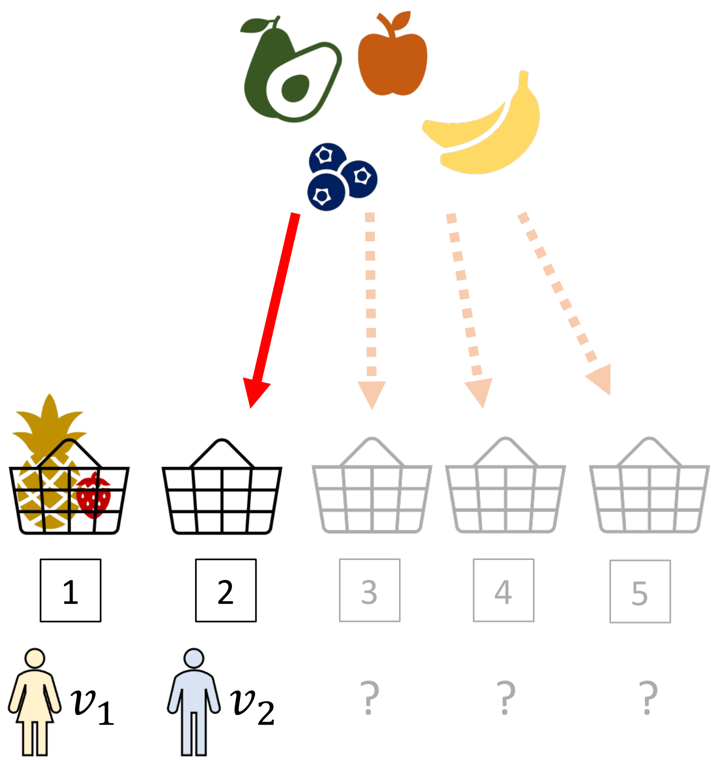

If agent i gets the set ALG_i
we want to **maximize**

$$\mathbb{E}(ALG) = \mathbb{E}\left(\sum_i v_i(ALG_i)\right)$$

Incentive Compatible Dynamic Program

Optimal online solution:

$$V_{n+1}(R) = 0$$
$$V_i(R) = \mathbb{E}\left(\max_{X \subseteq R} \{v_i(X) + V_{i+1}(R \setminus X)\} \right)$$

When set R is available, offer agent i **per-bundle prices**

$$p_i(X, R) = V_{i+1}(R) - V_{i+1}(R \setminus X)$$

If the agent maximizes utility, then she **selects the same as the DP**:

$$\max_{X \subseteq R}\{v_i(X) - p_i(X, R)\} = \max_{X \subseteq R}\{v_i(X) + V_{i+1}(R \setminus X)\} - V_{i+1}(R)$$

Benchmark: Optimal offline welfare

$$\mathbb{E}(OPT) = \mathbb{E}\left(\max_{\substack{X_1,\ldots,X_n \\ \text{partition}}} \sum v_i(X_i) \right)$$

Prophet Inequality

If agents arrive sequentially, is there a small number α such that

$$\alpha \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT) \quad ?$$

It can be proved that in general α is at least superconstant,

$$\alpha = \Omega\left(\frac{\log(m)}{\log\log(m)}\right)$$

Subadditive valuations
(a.k.a. complement-free)

$$v(A \cup B) \leq v(A) + v(B)$$

Gross-substitutes ⊆ Submodular ⊆ Fractionally-subadditive ⊆ Subadditive

Subadditive valuations

Offline:

Theorem. [Feige STOC'06]
If valuations are deterministic, we can find in polynomial time a 2-approximation.

Theorem. [Feldman, Fu, Gravin, Lucier STOC'13]
Simultaneous First-Price auctions result in a 2-approximation.

Online:

Theorem. [Dütting, Kesselheim, Lucier FOCS'20]
There is an $O(\log \log m)$ Prophet Inequality.

Theorem. [Correa and Cristi 2022+]
For every $\varepsilon > 0$, if valuations are subadditive, there is a $(6 + \varepsilon)$ Prophet Inequality, i.e., there is an online algorithm such that

$$(6 + \varepsilon) \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT)$$

Connection to single item

Ticket for a concert

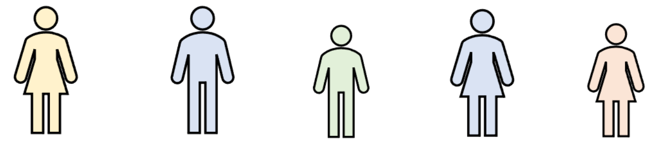

Sequence of n agents with independent valuations $v_i \sim F_i$

Theorem. We can get ½ of the expected optimal welfare.

- **Sample all distributions and use max as threshold** [Rubinstein, Wang, Weinberg, ITCS'21]
- **n pairs of numbers.**
  (say k largest come from different distributions)

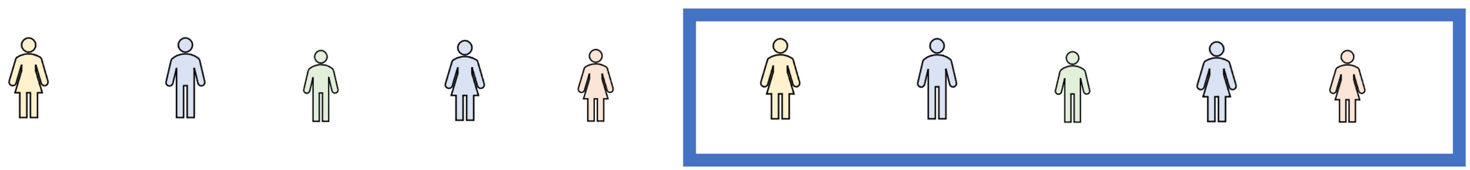

Connection to single item

Theorem. We can get ½ of the expected optimal welfare.

- **Sample all distributions and use max as threshold**

- **n pairs of numbers.**
  (say k largest come from different distributions)

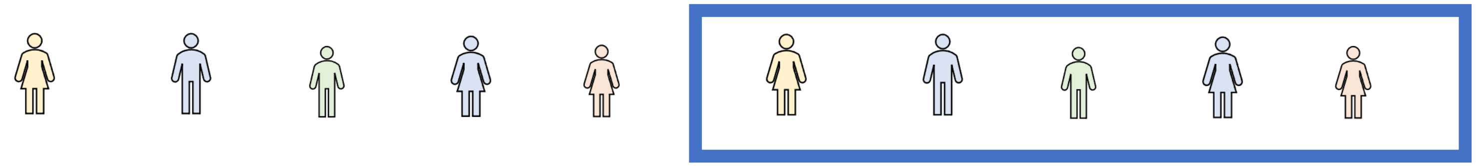

Blue box contains max w.p. 1/2
→ ALG gets max if max is in blue box and second max is not. Prob=1/4.

Blue box contains second max but not max w.p. 1/4.
→ ALG gets second max (or better) if max and second max are in blue box and third max is not. Prob=1/8.

Blue box contains third max but not max nor second max w.p. 1/8.
→ ALG gets third max if max, second max, and third max are in blue box, and fourth max is not. Prob=1/16.

Random Score Generators (RSG)

Imagine we could ask each agent how much they like each item

Formally, imagine there are functions $D_i: V_i \to \Delta(R_+^M)$

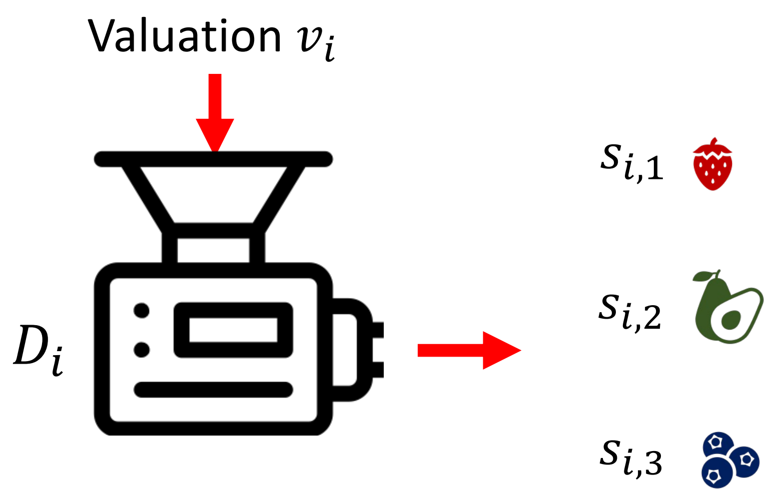

Random vector of scores

$s_i \sim D_i(v_i)$

Algorithm

Simulate valuations v_i' and scores $\left(s_{i,j}'\right) \sim D_i(v_i')$

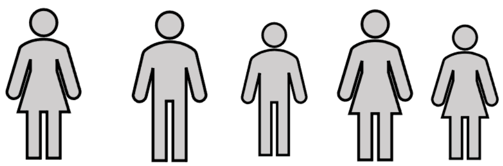

True valuations v_i and scores $\left(s_{i,j}\right) \sim D_i(v_i)$

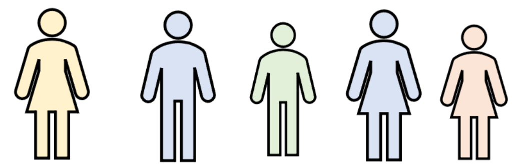

Set threshold
$T_j' = \max_i s_{i,j}'$

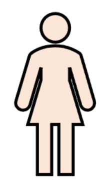

Give it to first agent such that
$s_{i,j} > T_j'$

Mirror Lemma. For every agent i,

$$\mathbb{E}(v_i(ALG_i)) \geq \frac{1}{2} \cdot \mathbb{E}\left(v_i\left(\left\{j: s_{i,j} > \max\left\{T_j', T_j''\right\}\right\}\right)\right)$$

Where T_j' and T_j'' are two independent samples of $\max_i s_{i,j}'$

$$\mathbb{E}(v_i(ALG_i)) \geq \frac{1}{2} \cdot \mathbb{E}\left(v_i\left(\left\{j: s_{i,j} > \max\left\{T_j', T_j''\right\}\right\}\right)\right)$$

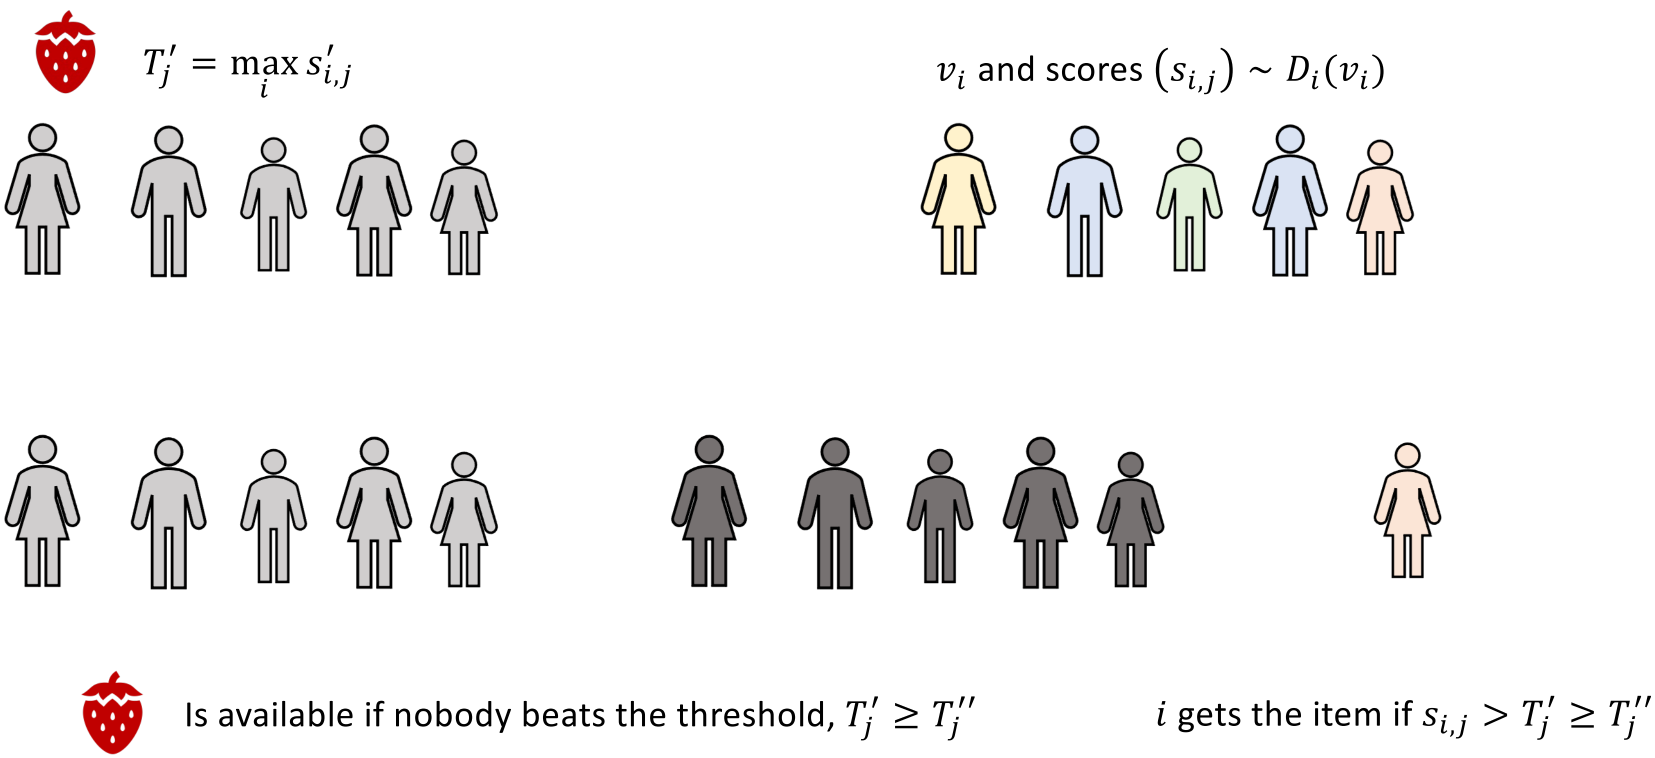

Key observation

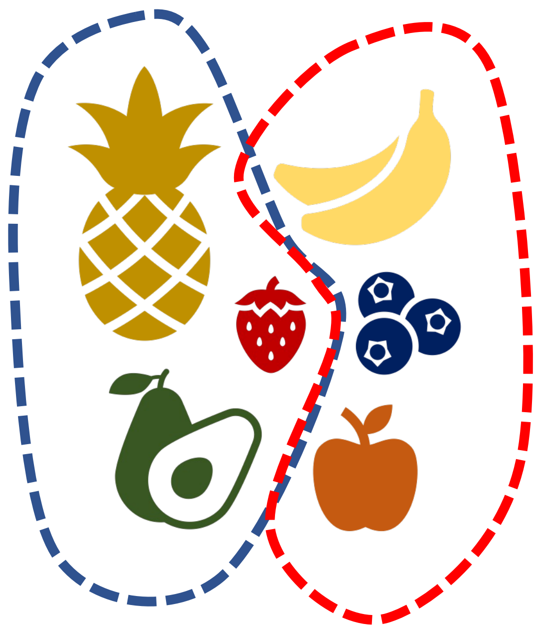

Set of available items and set of allocated items have the same distribution

$$T_j \leq T_j'$$

$$\left(\max s_{i,j} \leq \max s_{i,j}'\right)$$

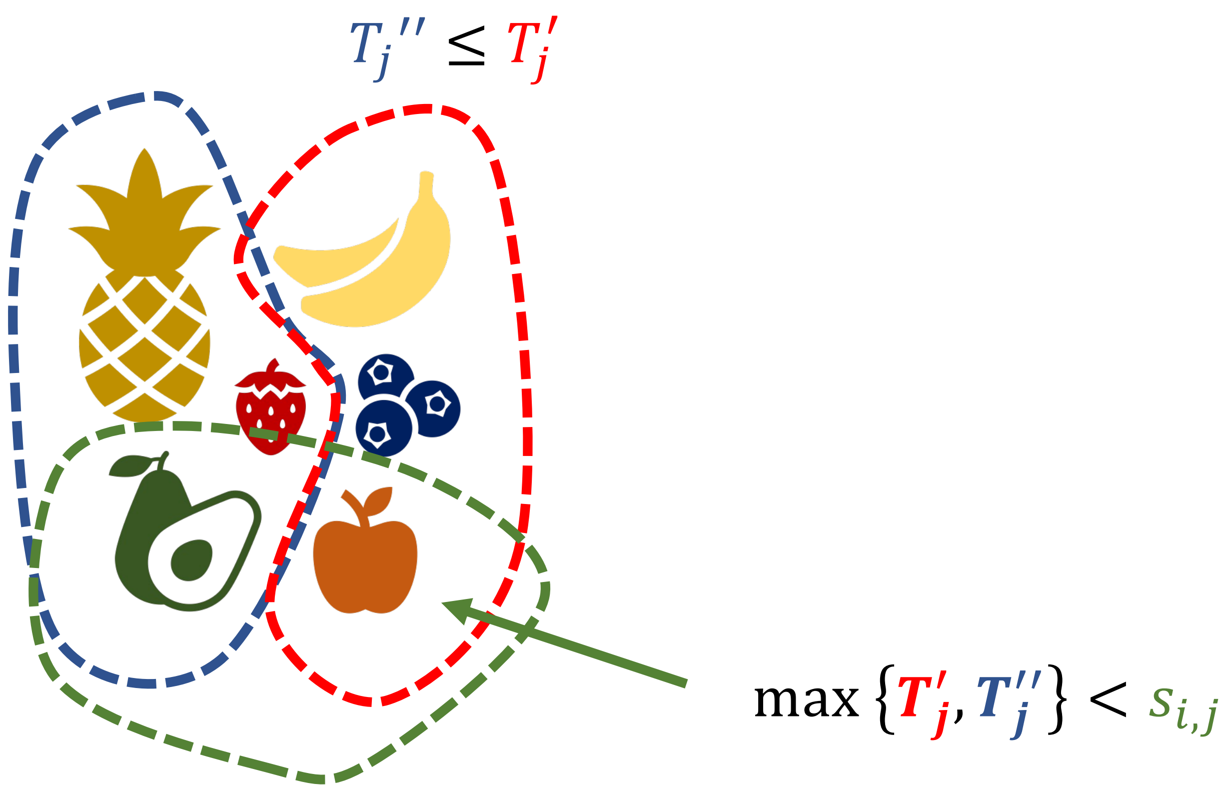
$T_j'' \leq T_j'$
$\max\{T_j', T_j''\} < s_{i,j}$

Lemma 2. There exist RSGs such that

$$\sum_i \mathbb{E}\left(v_i\left(\left\{j: s_{i,j} > \max\left\{T_j', T_j''\right\}\right\}\right)\right) \geq \frac{1}{3+\varepsilon} \cdot \mathbb{E}(OPT)$$

The proof uses a fixed-point argument.
Intuitively: we design a simultaneous auction with PoA $3+\varepsilon$, where each agent gets this set, and we take the equilibrium bids

Thus… for subadditive valuations

Theorem’.
For every $\varepsilon > 0$, there are RSGs such that

$$(6 + \varepsilon) \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT)$$

- Pricing implementation (Dynamic Program):
  - Uses item bundling
  - Uses dynamic pricing.
- Question: What if we cannot?

Demands of size $\boldsymbol{d}$

$$v(A) = \max_{X \subseteq A : |X| \leq d} v(X)$$

Theorem. [Correa, Cristi, Fielbaum, Pollner, Weinberg, IPCO'22]
If demands are of size at most d, there are *item prices* such that

$$(d+1) \cdot \mathbb{E}(\boldsymbol{ALG}) \geq \mathbb{E}(\boldsymbol{OPT})$$

and this is best possible. Moreover, we can compute them in polynomial time.

Matching: $\boldsymbol{d} = \mathbf{2}$

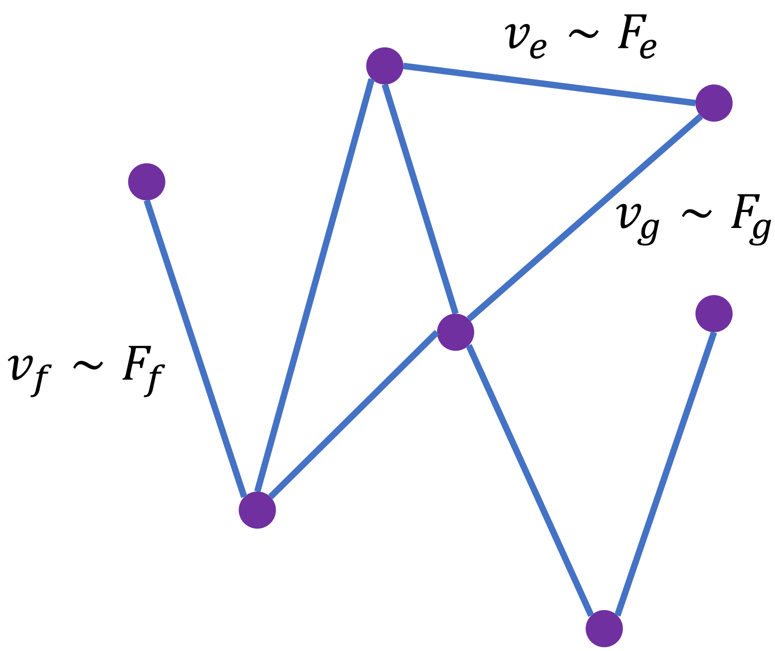

Edges come **one-by-one**

Select **matching** on the fly

Maximize expectation

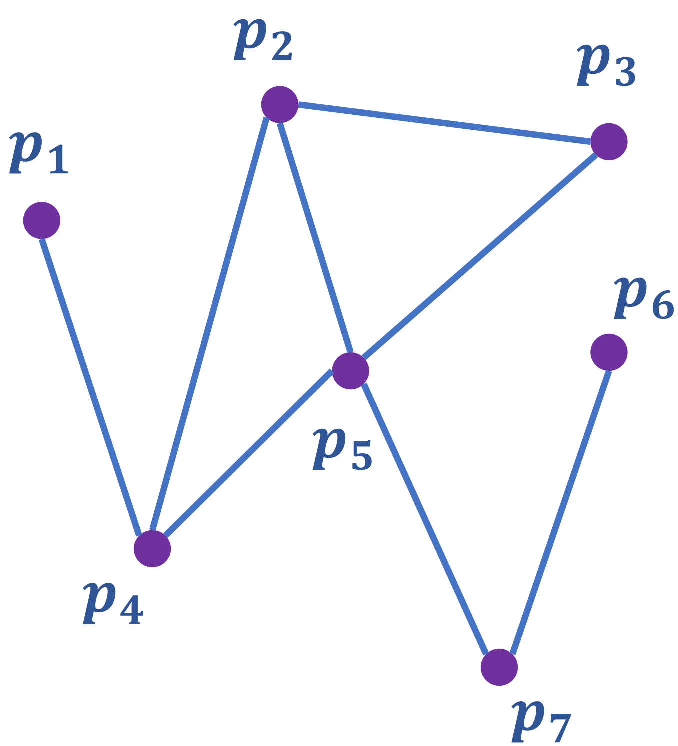

Algorithm:

$e = (u, w)$ arrives:

e buys u and w as long as they are not sold yet and $v_e \geq p_u + p_w$

$ALG(p)$ resulting matching

OPT optimal matching

Theorem. There is a vector of prices $\boldsymbol{p} \in \mathbb{R}_+^V$ s.t. for any arrival order,

$$3 \cdot \mathbb{E}\big(\boldsymbol{ALG}(\boldsymbol{p})\big) \geq \mathbb{E}\,(\boldsymbol{OPT})$$

To bound $\boldsymbol{OPT}$, imagine that edges in $\boldsymbol{OPT}$ had to pay the prices

$$\mathbb{E}(\boldsymbol{OPT}) = \mathbb{E}\left(\sum_{u \in V(OPT)} p_u + \sum_{e \in OPT} (v_e - p_u - p_w) \right)$$

$$\leq \sum_{u \in V} p_u + \sum_{e \in E} \mathbb{E}([v_e - p_u - p_w]_+)$$

$$:= \sum_{u \in V} p_u + \sum_{e \in E} \boldsymbol{z}_e(\boldsymbol{p})$$

$$\mathbb{E}(\boldsymbol{ALG(p)}) = \text{revenue} + \text{utility}$$

$$= \mathbb{E}\left(\sum_{u \in V(ALG(p))} p_u\right) + \mathbb{E}\left(\sum_{e \in ALG(p)} (v_e - p_u - p_w)\right)$$

We want balanced prices:
“high enough” so we get good revenue, yet “low enough” so buyers buy (and get good utility)

To lower bound $\mathbb{E}(ALG(p))$, utility is the tricky part:

$$\mathbb{E}\left(\sum_{e\in ALG(p)}(v_e - p_u - p_w)\right) = \sum_{e\in E}\mathbb{E}\left(I_{\{e\in ALG(p)\}}\cdot(v_e - p_u - p_w)\right)$$

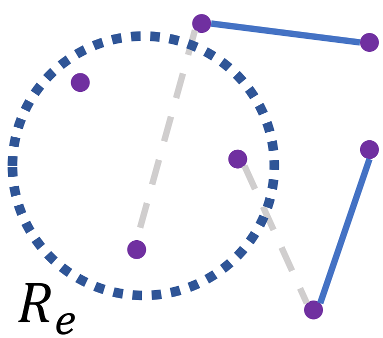

Recall that $\boldsymbol{ALG(p)}$ takes $e = (u, w)$ iff

- the two nodes are free, and
- $v_e \geq p_u + p_w$

R_e = set of remaining vertices when e arrives

R_e is independent of v_e

$$\text{Utility} = \sum_{e=(u,w)\in \mathrm{E}} \mathbb{E}\left(I_{\{u,w\in R_e\}} \cdot [v_e - p_u - p_w]_+\right)$$

$$= \sum_{e=(u,w)\in E} \mathbb{P}(u,w \in R_e) \cdot \mathbb{E}([v_e - p_u - p_w]_+)$$

$$\geq \sum_{e=(u,v)\in E} \mathbb{P}\left(u,w \notin V\left(ALG(p)\right)\right) \cdot \boldsymbol{z}_e(\boldsymbol{p})$$

$$= \mathbb{E}\left(\sum_{u,w\notin V\left(ALG(p)\right)} \boldsymbol{z}_e(\boldsymbol{p})\right)$$

$$
\begin{aligned}
\mathbb{E}\big(\boldsymbol{ALG}(\boldsymbol{p})\big) &= \text{revenue} + \text{utility} \\
&= \mathbb{E}\left(\sum_{u \in V(ALG(p))} p_u\right) + \mathbb{E}\left(\sum_{e \in ALG(p)} (v_e - p_u - p_w)\right) \\
&\geq \mathbb{E}\left(\sum_{u \in V(ALG(p))} p_u\right) + \mathbb{E}\left(\sum_{e=(u,w): u,w \notin V(ALG(p))} \boldsymbol{z}_e(p)\right) \\
&\geq \min_{X \subseteq V} \left\{ \sum_{u \notin X} p_u + \sum_{e \in E(X)} \boldsymbol{z}_e(p) \right\}
\end{aligned}
$$

$$\mathbb{E}(OPT) \leq \sum_{u \in V} p_u + \sum_{e \in E} z_e(\boldsymbol{p})$$

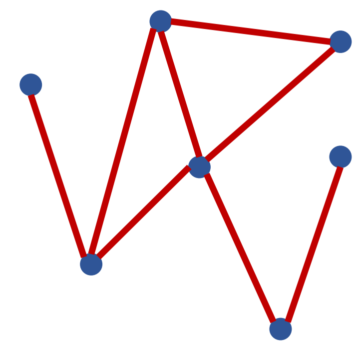

vs.

$$\mathbb{E}\big(ALG(\boldsymbol{p})\big) \geq \min_{X \subseteq V} \left\{ \sum_{u \notin X} p_u + \sum_{e \in E(X)} z_e(\boldsymbol{p}) \right\}$$

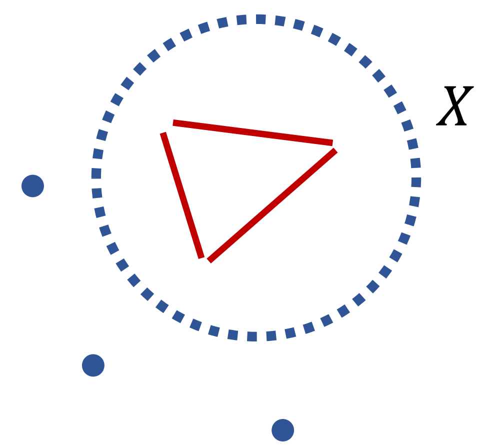

What if we set prices? $p_u = \sum_{e \in \delta(u)} \boldsymbol{z}_e(\boldsymbol{p})$

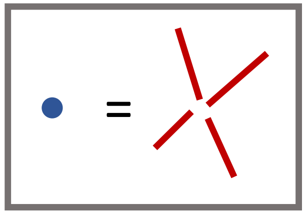

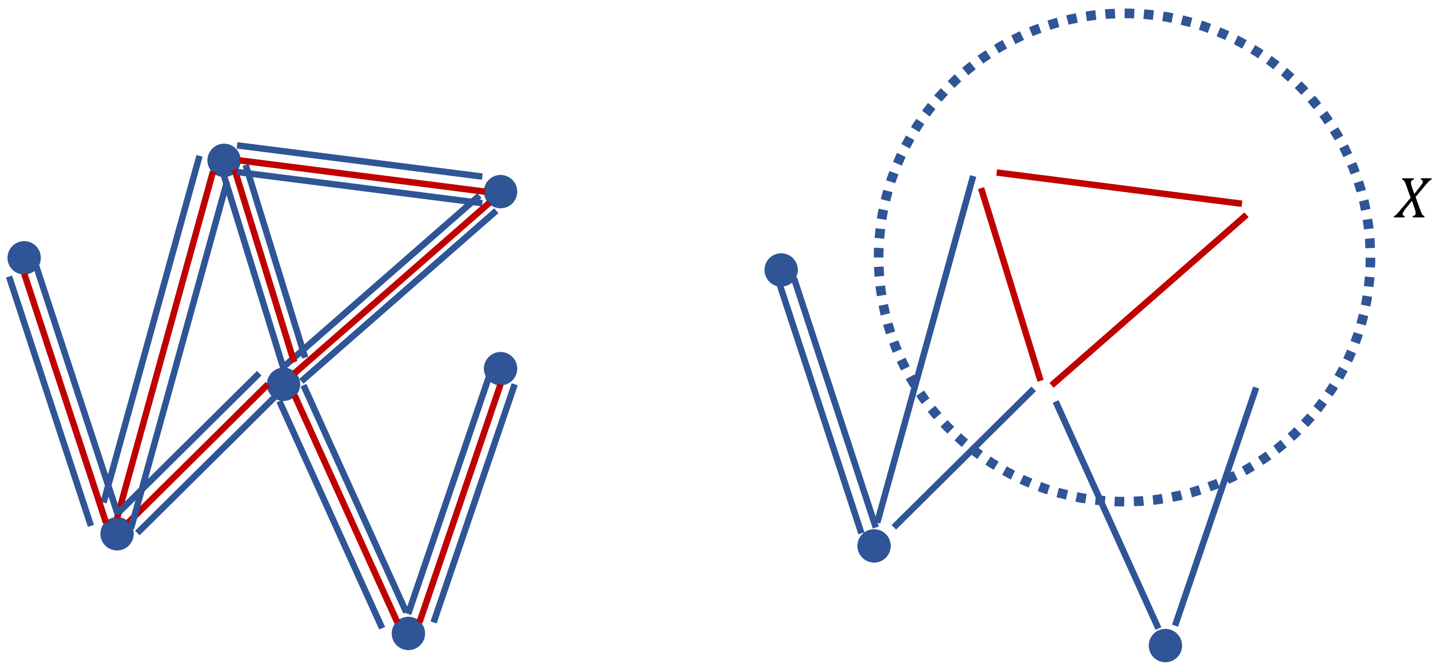

We want prices

$$p_u = \sum_{e \in \delta(u)} \boldsymbol{z}_e(\boldsymbol{p})$$

Define the operator: $\psi_u(\boldsymbol{p}) = \sum_{e \in \delta(u)} \boldsymbol{z}_e(\boldsymbol{p})$

Brouwer's thm $\Rightarrow$ there are prices $\boldsymbol{p} = \psi(\boldsymbol{p})$

Recall that $\boldsymbol{z}_e(\boldsymbol{p}) = \mathbb{E}([v_e - p_u - p_w]_+)$

■

Can we compute $\boldsymbol{p}$? Brouwer's only guarantees existence.

YES! For $\varepsilon > 0$, we can compute $\boldsymbol{p}$ in polynomial time s.t.

$$(3+\varepsilon) \cdot \mathbb{E}\left(\boldsymbol{ALG(p)}\right) \geq \mathbb{E}(\boldsymbol{OPT})$$

For $\varepsilon > 0$, m edges, n nodes and a bound $B \geq \frac{v_{\max}}{\mathbb{E}(OPT)}$, we can compute $\boldsymbol{p}$ in **time** $poly\left(m, n, \frac{1}{\varepsilon}, B\right)$, using $poly\left(m, n, \frac{1}{\varepsilon}, B\right)$ **samples.**

Sample:

$$\left(v_e^{(s)}\right)_{e \in E}$$

Calculate "empirical" ψ:

$$\hat{\psi}_u^{(s)}(\boldsymbol{p}) = \sum_{e \in \delta(u)} \left[v_e^{(s)} - p_u - p_w\right]_+$$

Repeat for $s = 1, \ldots, S$

Average

$$\hat{\psi} = \frac{1}{S} \sum_{s=1}^{S} \hat{\psi}^{(s)}$$

For large S, $\hat{\psi}$ is good approx (concentration bound)

Find

$$\boldsymbol{p} = \hat{\psi}(\boldsymbol{p})$$

convex QP

Example

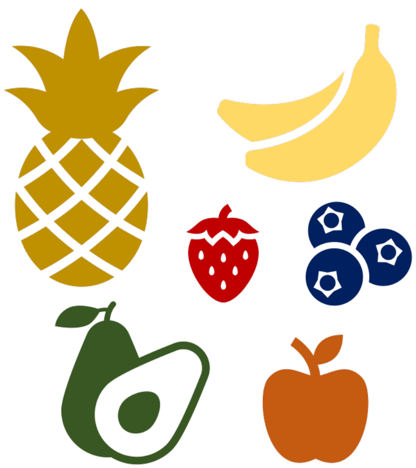

$|M| = d$

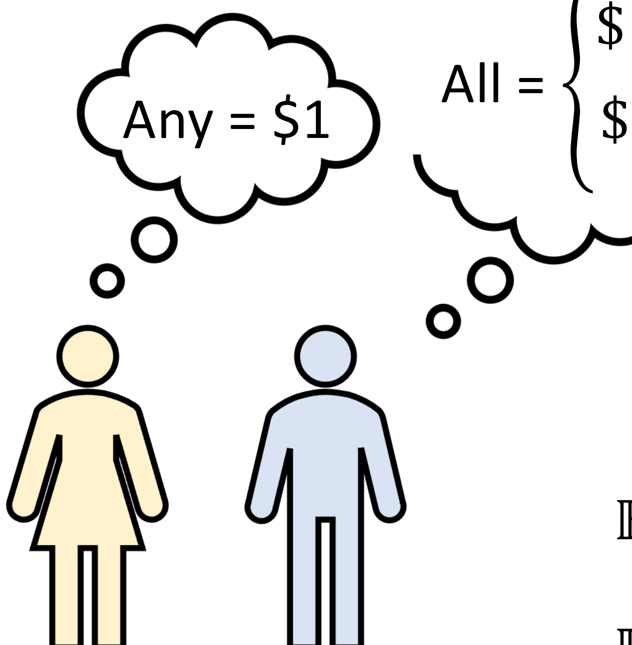

$$\text{All} = \begin{cases} \$\, d & w.p.\ 1-\varepsilon \\ \$\, \frac{1}{\varepsilon} & w.p.\ \varepsilon \end{cases}$$

$$\mathbb{E}(\boldsymbol{OPT}) = 1 + d(1-\varepsilon)$$

$$\mathbb{E}\big(\boldsymbol{ALG(p)}\big) \leq 1$$

SUMMARY

We show a $(6+\varepsilon)$-approx. for OCA with subadditive valuations

- Algorithm uses samples to "protect" items. We use simple scores to represent complex valuation functions and use a fixed-point argument to show existence of good scores.
- We improve upon the $O(\log\log m)$-approx solving an important open question.

We find the best possible prices for online combinatorial auctions with random valuation parametrized by d

- Existence follows by a fixed-point argument. Polynomial time computation follows by carefully analyzing the underlying function and classic optimization tools.
- The result improves upon some recent results in the literature:
  - Best-known factor of $(4d-2)$ [Dütting, Feldman, Kesselheim, Lucier, FOCS'20]
  - Single-minded and random valuations generalizes Prophet Inequality for intersection of d partition matroids. Best known approximation is $e(d-1)$. [Feldman, Svensson, Zenklusen, SODA'16]
  - For Prophet Inequality for matching ($d=2$) a 3-approx. is known, and a 2.96-approx. using adaptive thresholds (prices). [Gravin, Wang, EC'19], [Ezra, Feldman, Gravin, Tang, EC'20]