# Approximately Optimal Auction Design and Item Pricing

SHUCHI CHAWLA

TEXAS
The University of Texas at Austin

Optimization + Strategic participants ⟺ Mechanism Design

Allocation problem + Strategic agents ⟺ ~~Auction~~ Design

Pricing

TCS contributions:

+ computational considerations
+ robustness
+ learnability
+ simplicity
+ …


# This talk:
## some examples of pricing as a solution to an auction design problem

Part 1: Social Welfare Maximization

Part 2: Revenue Maximization

# A generic stochastic resource allocation setting

Many heterogenous items in limited supply

Buyers assign values to subsets of items

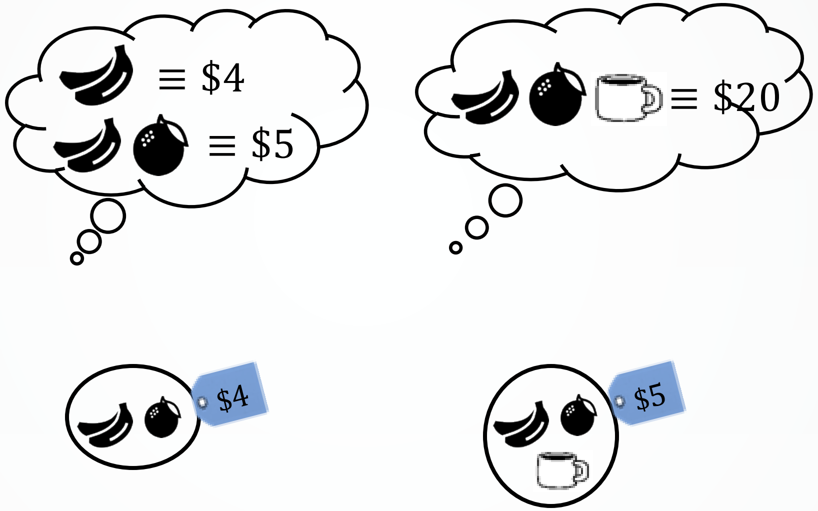


Known population of buyers

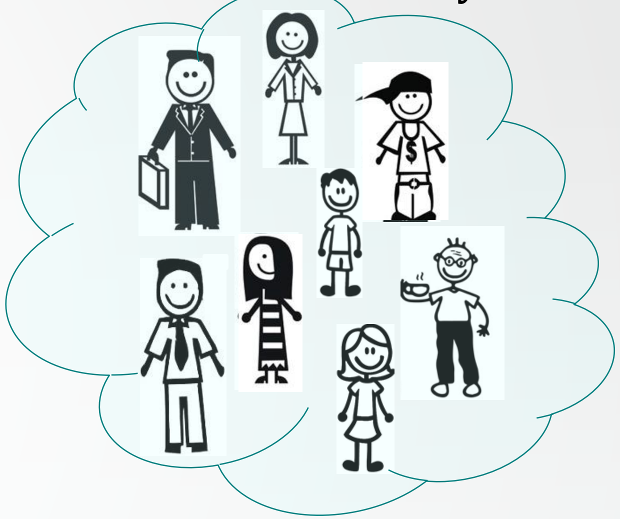

Buyers drawn randomly from population

Auction:
(all reported prefs, market info)
→ (allocation, payments)

Buyers' goal: obtain an allocation that maximizes their value – the price they pay.

# Part I: Social Welfare Maximization

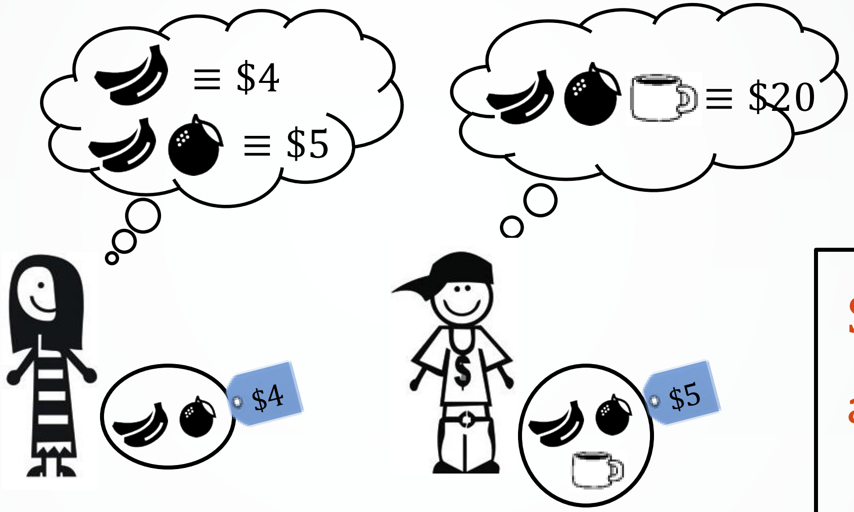


SOCIAL WELFARE,

a.k.a., Economic Efficiency

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

Vickrey Auction: assigns the optimal allocation and charges “supporting” prices. Always truthful.

# Part I: Online Stochastic Social Welfare Maximization

Challenge: determine the allocation and payment for each person before observing values of future arrivals.

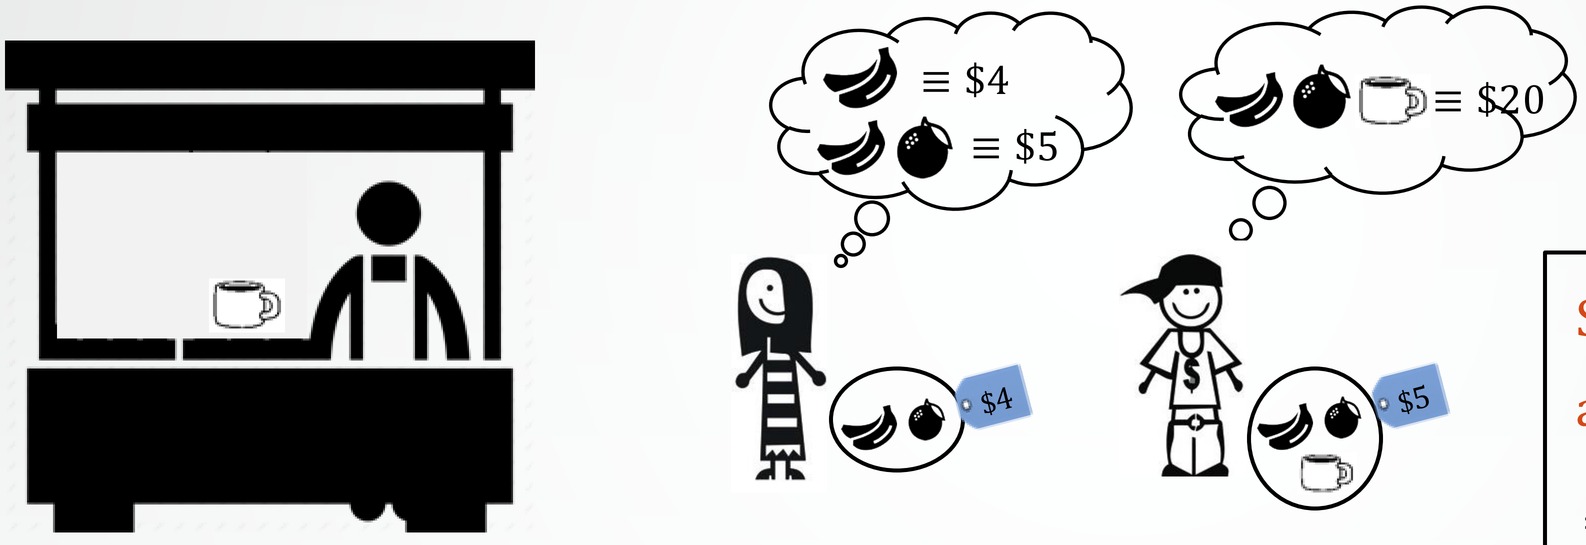


SOCIAL WELFARE,

a.k.a., Economic Efficiency

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

Vickrey Auction: assigns the optimal allocation and charges 'supporting' prices. Always truthful.

$$\text{Competitive Ratio} = \max_{\text{distributions}} \left\{ \frac{\mathrm{E}_{\text{instance}\sim\text{dist}}[\text{Hindsight-OPT}(\text{instance})]}{\mathrm{E}_{\text{instance}\sim\text{dist}}[\text{ALG}(\text{instance})]} \right\}$$

# The single item case: prophet inequality

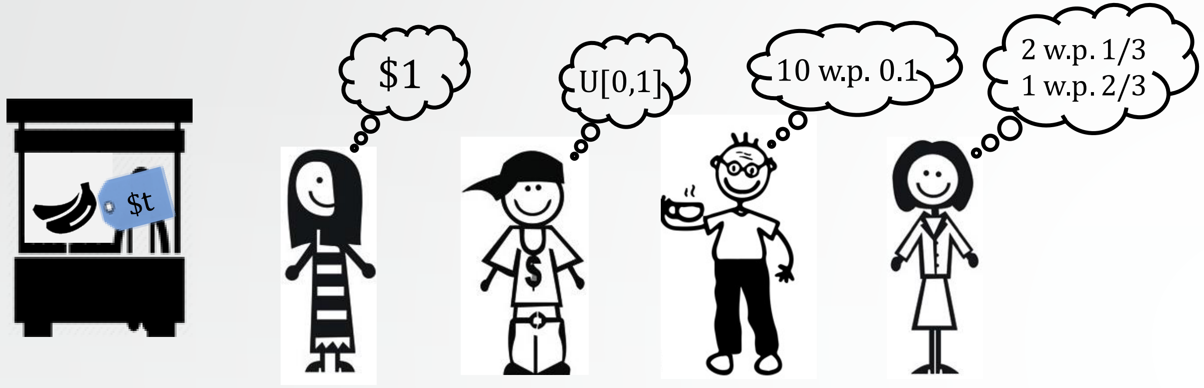


- Customers arrive in sequence and reveal their values
- At every step, the algorithm decides whether to allocate and stop; or to reject and move forward.
- Hindsight-OPT picks the maximum value

A threshold-based policy: allocate to the *first* value that crosses pre-determined threshold, a.k.a. price, t.

Samuel-Cahn'84: Threshold-based policies achieve a CR of 2.

- No other online algorithm can do better.
- Set price = $\frac{1}{2}$ Hindsight−OPT

Robust to different arrival orders!

# The single item case: prophet inequality

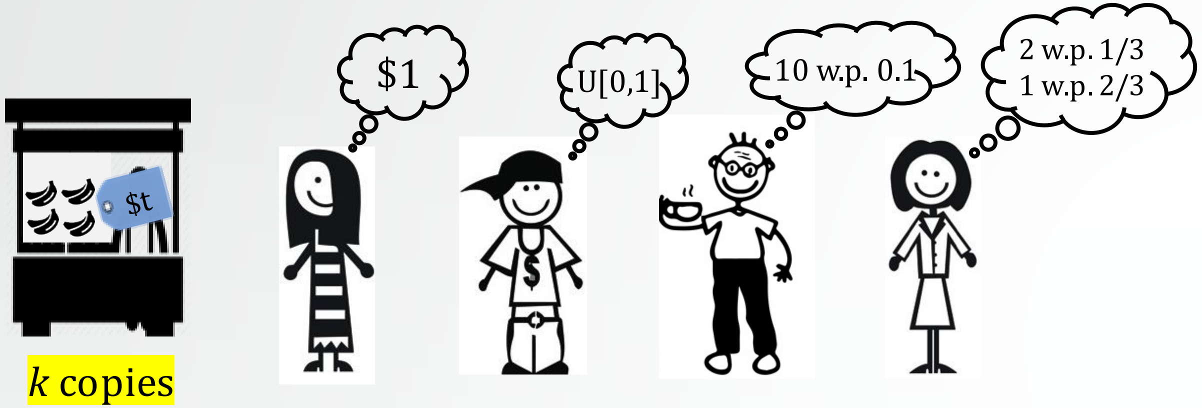


- Custor … lues
- At eve… allocat
- Hindsi

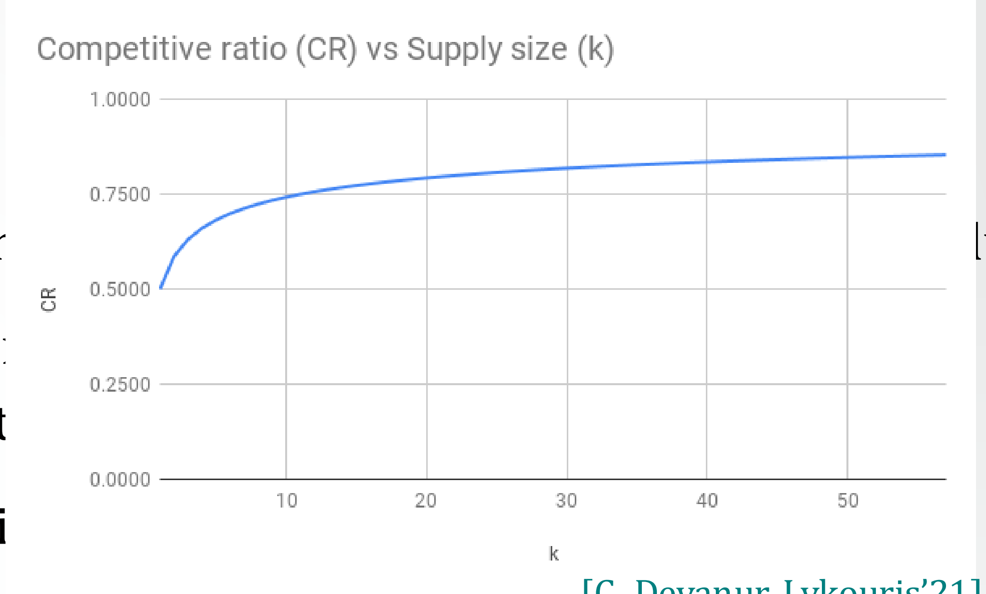


[C.-Devanur-Lykouris'21]

A threshold-based policy: allocate to the *first k* values that cross pre-determined threshold, a.k.a. price.

Hajiaghayi-Kleinberg-Sandholm'07
C.-Devanur-Lykouris'21

Samuel-Cahn'84: **Threshold-based policies achieve a CR of 2.**

**Threshold-based policies achieve a CR of** $1-\theta\left(\sqrt{\log k / k}\right)$.

- No other online algorithm can do better.

Ghosh-Kleinberg'16: No other online algorithm can do better asymptotically.

Robust to different arrival orders!

- Set price $= \frac{1}{2}$ Hindsight−OPT

# The “unit demand” case: balanced prices

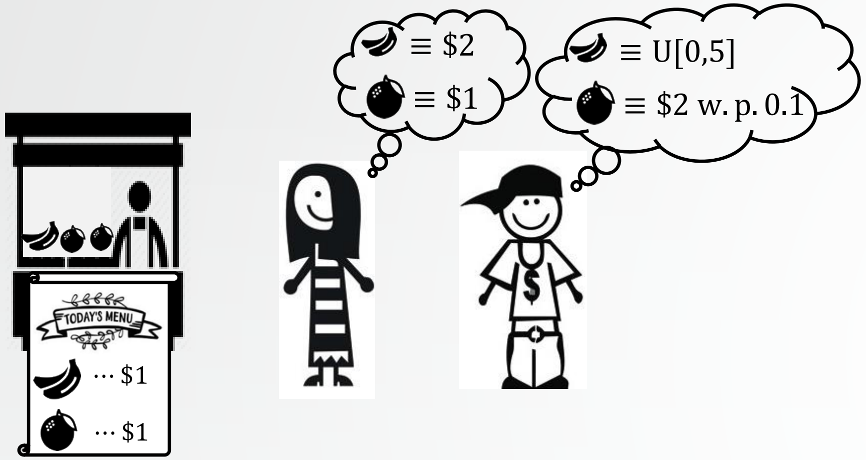


- Customers arrive in sequence and reveal their values
- At every step, the algorithm decides what to allocate and at what price; or to reject and move forward.
- Hindsight-OPT picks the SW maximizing allocation.

Item pricing: fix prices in advance; allow buyers to purchase their favorite item while supplies last.

Feldman-Gravin-Lucier’15: Item pricing achieves a CR of 2.

- No other online algorithm can do better.
- Set $\text{price}_i = \frac{1}{2}(\text{Contribution of i to Hindsight} - \text{OPT})$

Robust to different arrival orders!

# Item prices arise as dual variables

$v_{i,j}$: buyer $i$'s value for item $j$

$q_i$: buyer $i$'s probability of arrival

$x_{i,j}$: fraction of item $j$ allocated to buyer $i$

$k_j$: supply of item $j$

**PRIMAL a.k.a. expected case LP**

$$\max \sum_{i,j} x_{i,j} v_{i,j}$$ **Social Welfare**

subject to:

$$\sum_j x_{i,j} \le q_i \quad \text{for all buyers } i$$ **No overprovisioning**

$$\sum_i x_{i,j} \le k_j \quad \text{for all items } j$$ **Demand ≤ Supply**

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

One of these is "tight"

**DUAL**

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

subject to:

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

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

Optimal value of PRIMAL ≥ Hindsight-OPT

In an optimal solution, $u_i = \max_j (v_{i,j} - p_j)$

If $p_j$'s denote prices, then $u_i$'s are utilities!

Complementary slackness ⟹ LP allocates $j$ to $i$ iff $j$ is one of $i$'s favorite items under the pricing $p$.

# The "unit demand" case: balanced prices

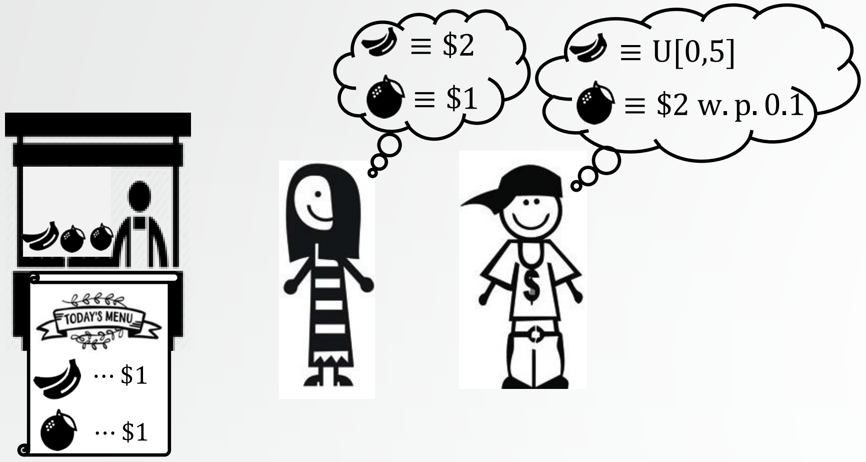


- Customers arrive in sequence and reveal their values
- At every step, the algorithm decides what to allocate and at what price; or to reject and move forward.
- Hindsight-OPT picks the SW maximizing allocation.

Item pricing: fix prices in advance; allow buyers to purchase their favorite item while supplies last.

Feldman-Gravin-Lucier'15: Item pricing achieves a CR of 2.

- No other online algorithm can do better.
- Set $\text{price}_i = \frac{1}{2}(\text{Contribution of } i \text{ to Hindsight}-\text{OPT})$

Use dual prices??

# But dual prices don’t work well in stochastic settings

- Problem 1: dual prices are too low.

Value = 1
Arrival prob. =0.9

Value = 100
Arrival prob.=0.1

Dual price = 1; Alg allocates to the second buyer w.p. 0.01.

- Problem 2: as supply diminishes, the correspondence between LP-allocation and buyer preferences breaks down.
  - Every buyer purchases her favorite of the remaining items

Coming up: two approaches to get around these problems…

# Approach 1: balanced prices

[Feldman-Gravin-Lucier'15]

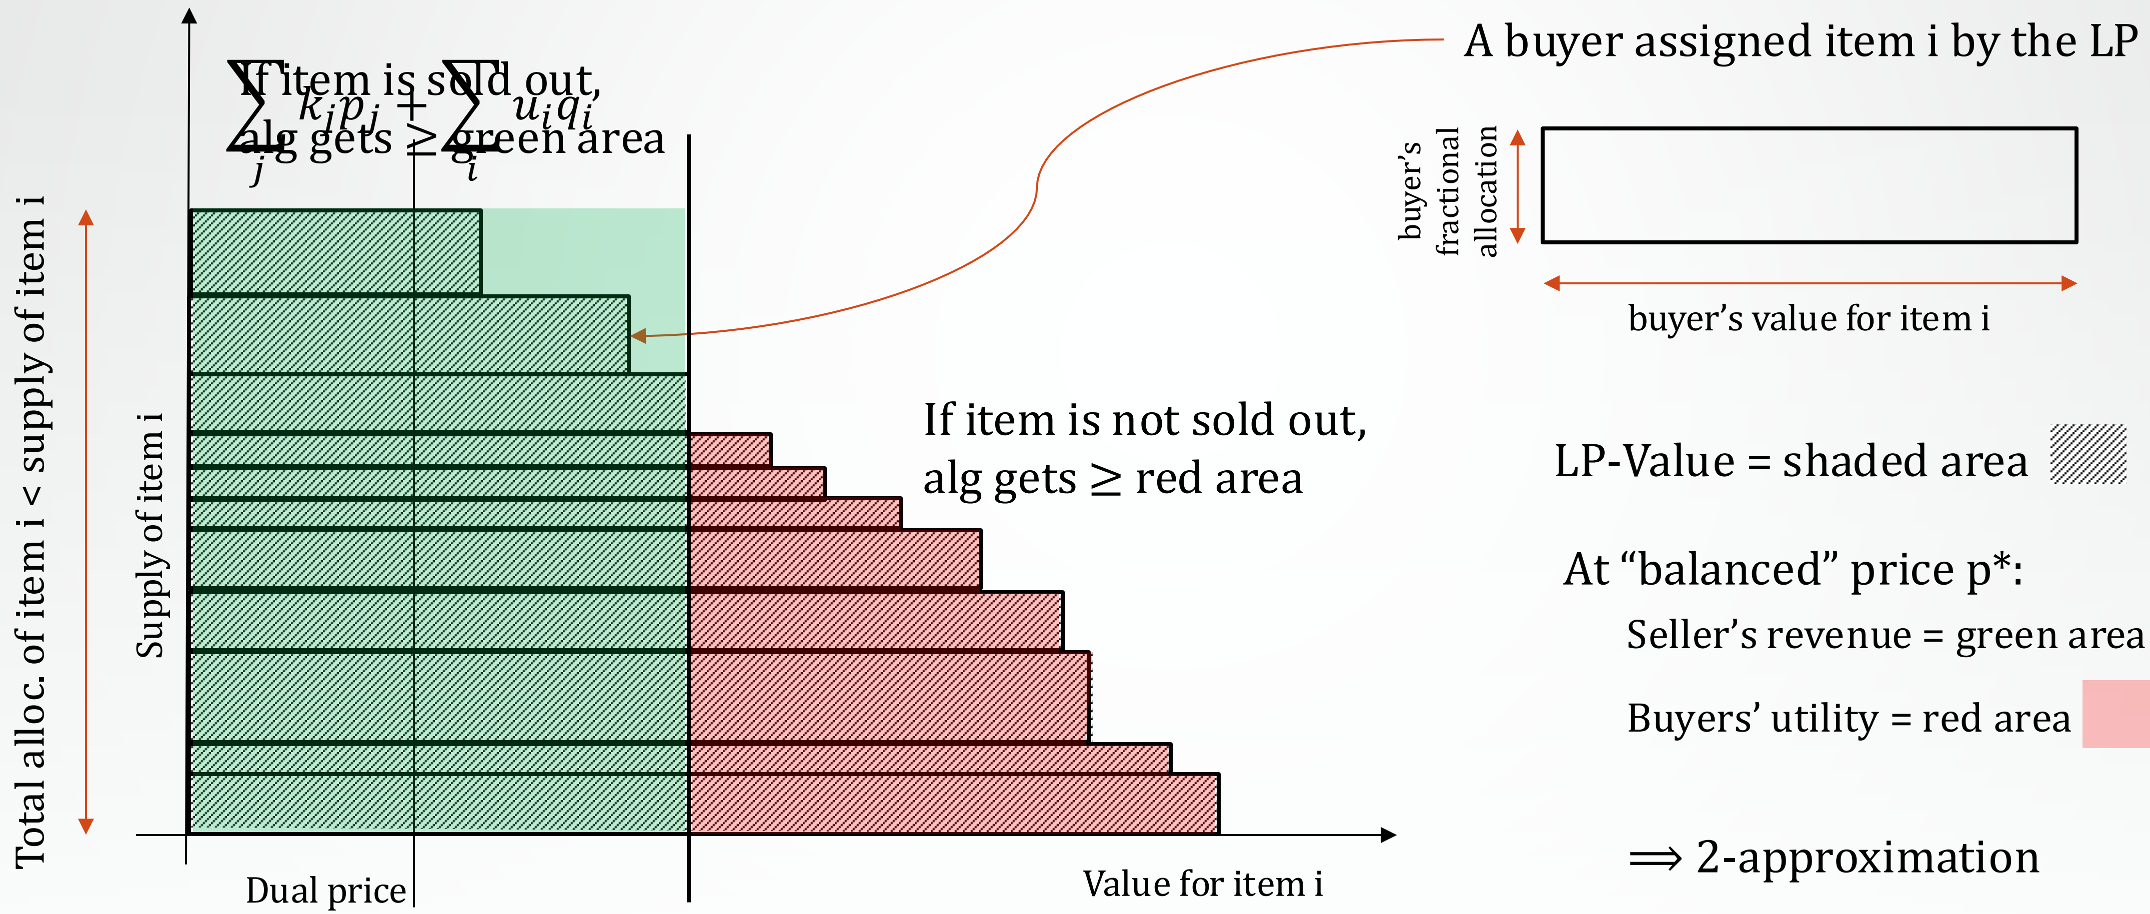


Green area = red area ≥ ½ LP-Value

LP-Value = shaded area

At "balanced" price p*:

Seller's revenue = green area

Buyers' utility = red area

⟹ 2-approximation

# Approach 2: tracking buyer preferences

[C. Devanur Holroyd Karlin Martin Sivan'17]

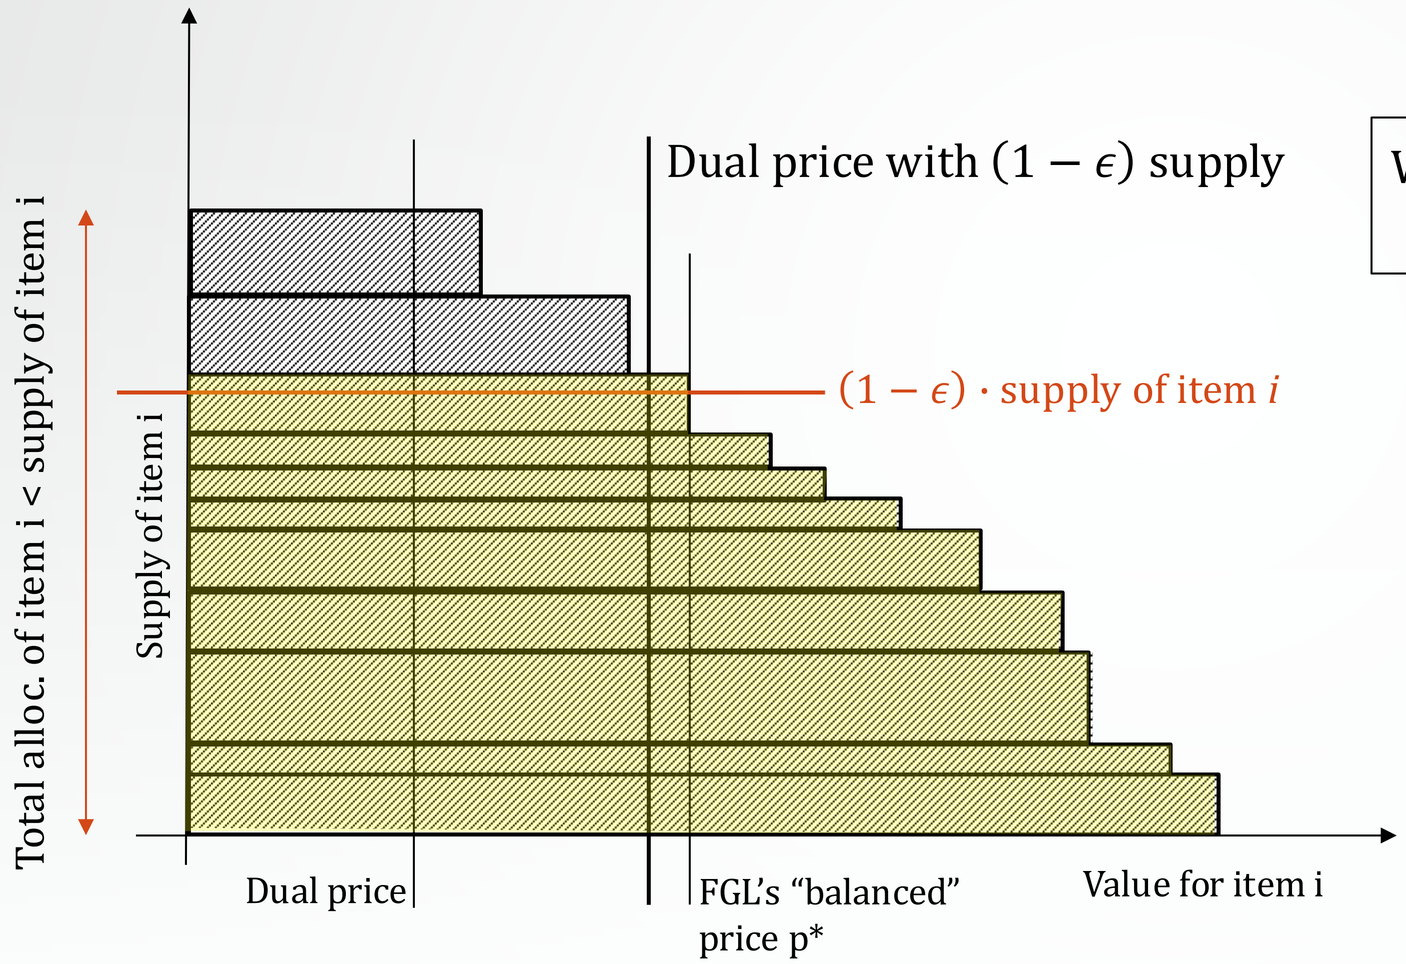


When supply is large enough,
$\Pr[\text{\#arrivals of yellow buyers} > \text{supply}] \leq \epsilon$

Failure event

Challenge: a failure event at one item can cause a failure event at another item.

**Question: how do failure events cascade?**

# Approach 2: tracking buyer preferences

[C. Devanur Holroyd Karlin Martin Sivan'17]

Forwarding graph

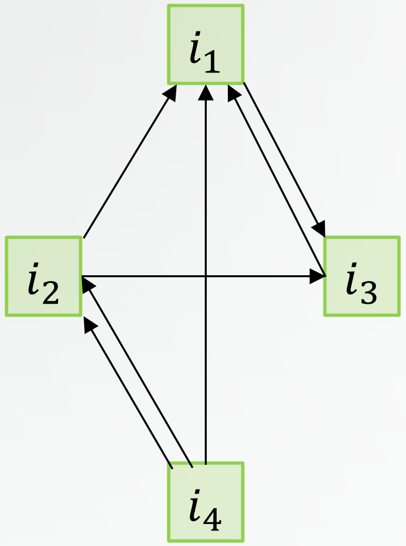


Vertices ≡ items
Edges ≡ movement of buyers from one item to the next

Failure events move along edges in the forwarding graph

When supply is large enough,
Pr[#arrivals of yellow buyers > supply] $\le \epsilon$

Failure event

Challenge: a failure event at one item can cause a failure event at another item.

**Question: how do failure events cascade?**

Theorem: If the graph has constant in-degree failure events cascade with low probability.
$\Rightarrow$ CR of $1 - O(\sqrt{\log k / k})$

# Approach 2: tracking buyer preferences; application to interval scheduling

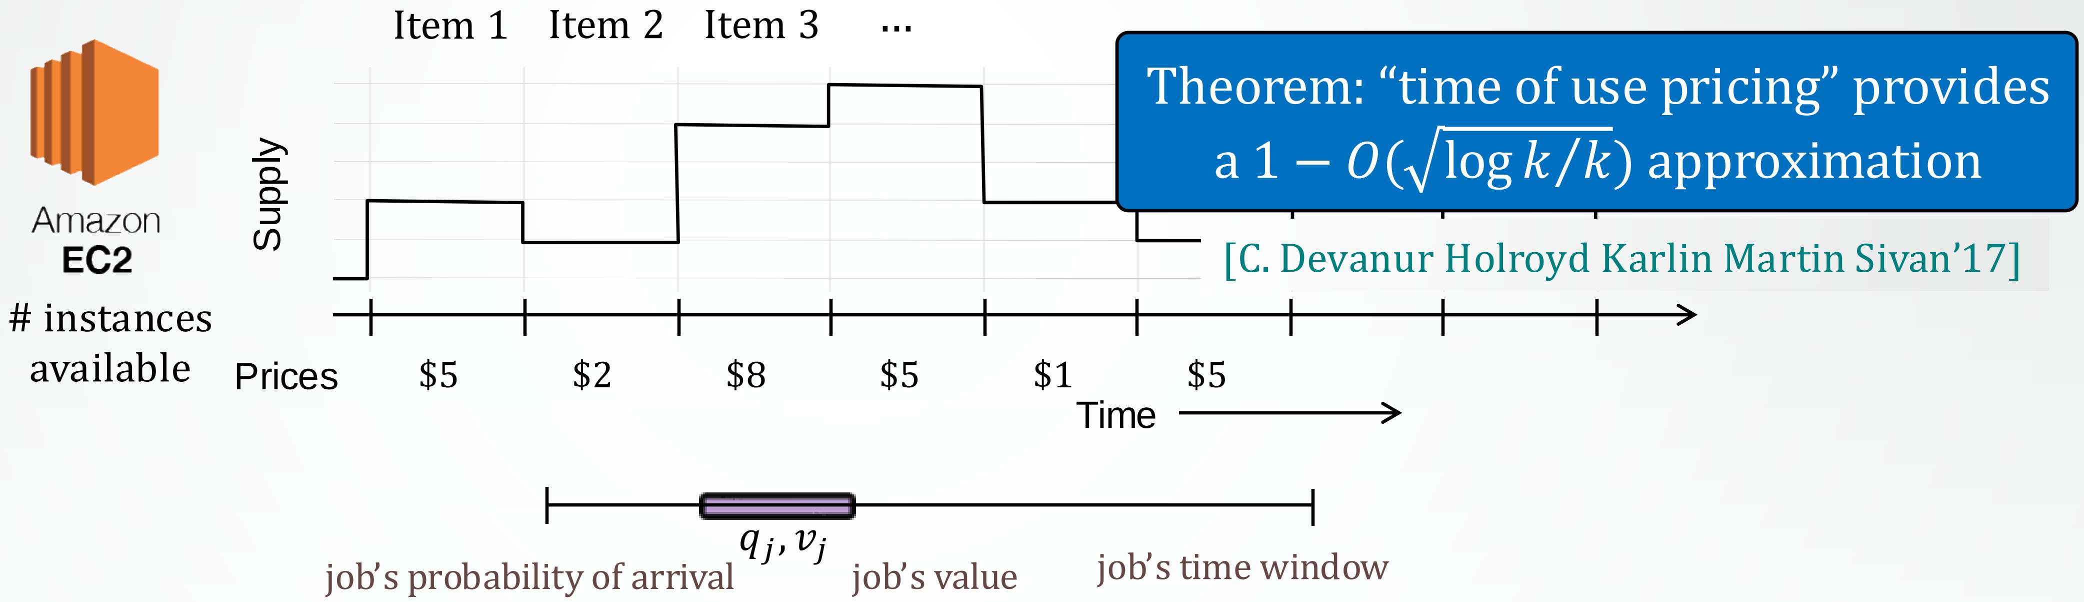


**Theorem:** "time of use pricing" provides a $1 - O(\sqrt{\log k / k})$ approximation

[C. Devanur Holroyd Karlin Martin Sivan'17]

Items ≡ compute instances at different points of time; Buyers ≡ jobs with requirements

# Part I: Online Stochastic Social Welfare Maximization – Summary

Posted (static) pricing is the best online truthful SW-maximizing algorithm known for many settings:

Item pricing; or
Bundle pricing

- Single item [Samuel-Cahn'84]
- Unit demand [Feldman Gravin Lucier'15]
- Job scheduling [C. Devanur Holroyd Karlin Martin Sivan'17, C. Miller Teng'19]
- Fractionally subadditive values [FGL'15]
- Subadditive values [Dutting Kesselheim Lucier'20]
- MPH hierarchy of values [FGL'15, Dutting Feldman Kesselheim Lucier'17]
- Bandwidth allocation [C. Miller Teng'19]

# A generic stochastic resource allocation setting

Many heterogenous items in limited supply

Buyer assigns values to subsets of items

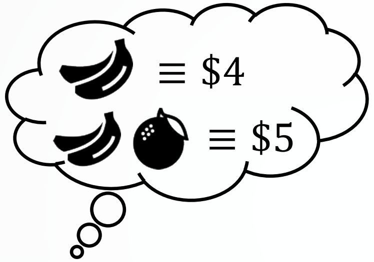


Known population of buyers

Simplifying assumption: single buyer

Buyer drawn randomly from population

Buyers' goal: obtain an allocation that maximizes their value – the price they pay.

Auction:
(all reported prefs, market info)
→ (allocation, payment)

# Part II: Revenue Maximization

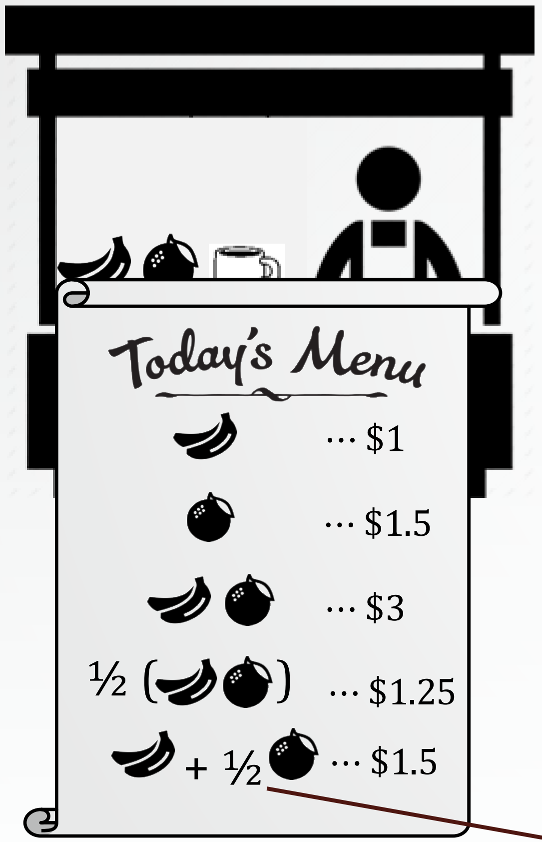


Prob. of allocation

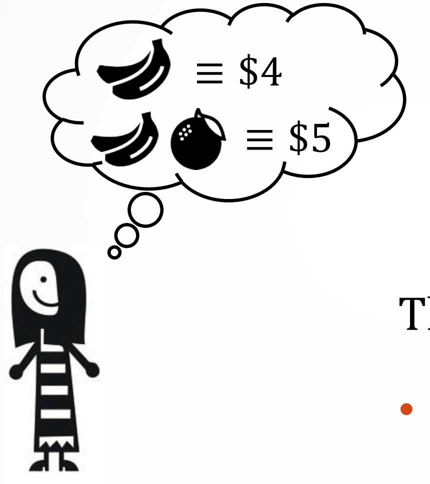


$$\text{REVENUE}^* = \sum_{\text{buyers } i} (\text{payment made by } i)$$

(*) Assumption: seller has a monopoly.

The optimal mechanism can be quite complicated:

- Offers items packaged into bundles
- Offers random allocations, a.k.a. lotteries [Thannasoulis'05]
- Can have infinitely many options! [Hart Nisan'13]
- Can be computationally hard to find. [Chen et al.'15]

# Part II: Revenue Maximization – Approximation

- If the buyer is unit-demand and his values for different items are independent, then

$$\text{Item Pricing} \geq {}^{1}\!/_{4}\,\text{OPT}$$

[C.-Hartline-Kleinberg'07]
[C.-Malec-Sivan'10]

- If the buyer has additive values and his values for different items are independent,

$$\max(\text{Item Pricing}, \text{Grand Bundle Pricing}) \geq {}^{1}\!/_{6}\,\text{OPT}$$

[Li-Yao'13]
[Babaioff-Immorlica-Lucier-Weinberg '14]

- If the buyer's value function is subadditive over independent item values,

$$\max(\text{Item Pricing}, \text{Grand Bundle Pricing}) \geq \Omega(1)\text{OPT}$$

[Rubinstein-Weinberg'15]

Best known approximations using any "simple" mechanisms

In the absence of independence, $\exists$ instances with $\text{OPT} = \infty$ and $\text{Revenue}(\text{any finite menu}) < \infty$

(even with just two items and unit-demand or additive values)

[Briest-C.-Kleinberg-Weinberg'10]
[Hart-Nisan'13]

Unit-demand: $v(S) = \max_{i \in S} v_i$
Additive: $v(S) = \sum_{i \in S} v_i$

Item pricing: $p(S) = \sum_{i \in S} p_i$
Grand Bundle pricing: $p(S) = p([n])$

# Part II: Revenue Maximization with a buy-many constraint

[C.-Tzamos-Teng'19]

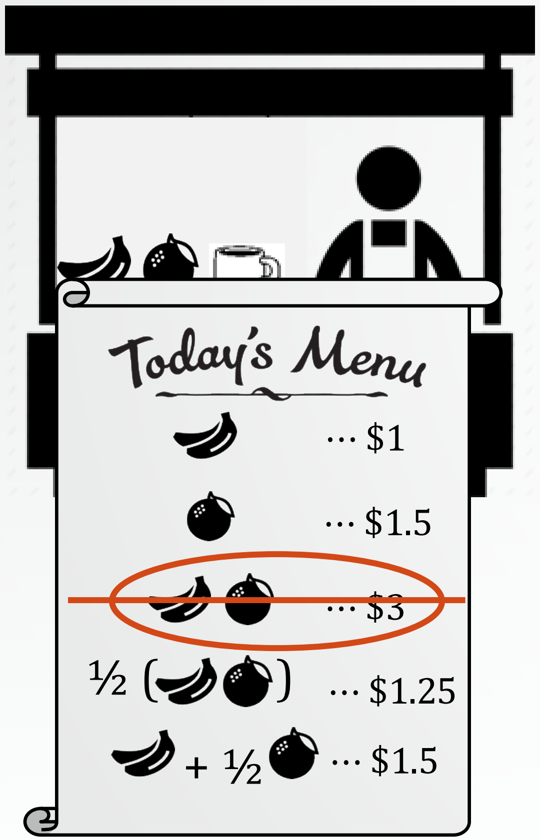


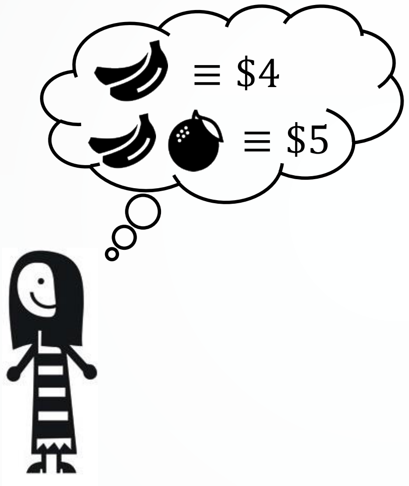


Buy-many constraint: cannot sell a bundle at a price higher than the sum of its constituents.

The optimal buy-many mechanism can be quite complicated:

- Offers random allocations, a.k.a. lotteries
- Can have infinitely many options!
- Can be computationally hard to find.

# Part II: Revenue Maximization with a buy-many constraint – Approximation

[C.-Tzamos-Teng'19]

Theorem 1: For any value distribution,

$n$: #items

$$\text{Buy-many OPT} \leq 2 \log 2n \cdot \text{Item Pricing}$$

Theorem 2: There exists a distribution over additive valuations such that

$$\text{Buy-many OPT} \geq \Omega(\log n) \text{ Revenue of any "succinct" mechanism}$$

One that can be described using $2^{o(n^{1/4})}$ bits

Theorem: Item Pricing is always a $2 \log 2n$-approximation to the optimal buy-many mechanism.

- Buy-many menus ≡ subadditive pricing function
- Item pricing ≡ additive pricing function
- Additive fns ($g$) pointwise $n$-approximate subadditive fns ($f$)

Pointwise approximation ⟹ Approximation in revenue

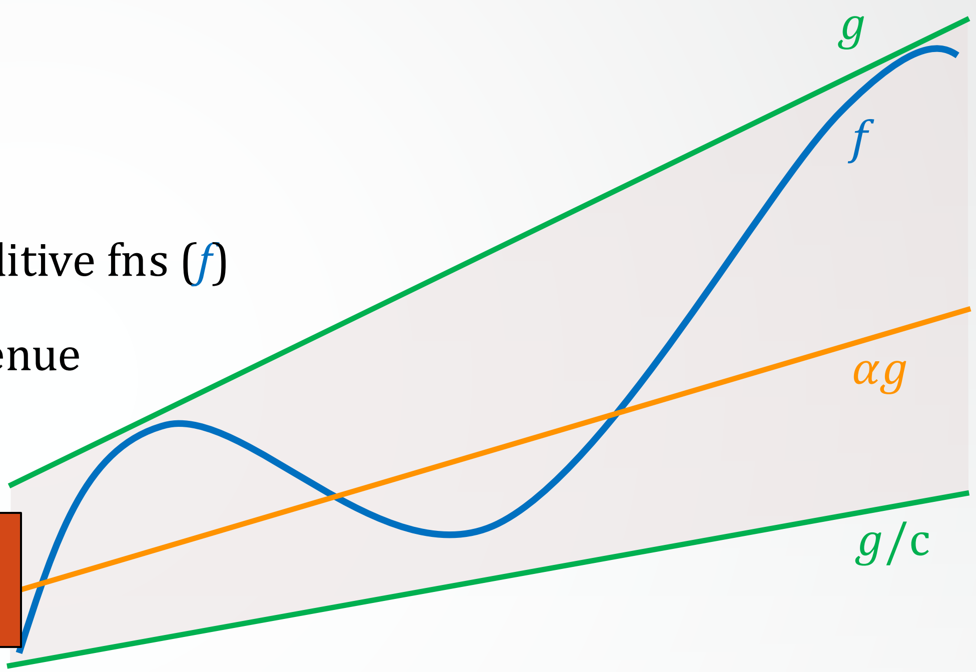


Additive functions are the succinct functions that best approximate an arbitrary subadditive function.

Lemma: Let $f$ and $g$ be any pricing functions such that $g$ **pointwise $c$-approximates** $f$.

Then there exists a distribution over scaling factors $\alpha > 0$, such that for any buyer,

The price paid by the buyer under $\alpha g \geq \frac{1}{2 \log 2c}$ (The price paid by the buyer under $f$)

$\frac{1}{2} f(S) \leq \text{scaled-}g(S) \leq f(S)$

$f(\text{argmax}_S\{v(S) - f(S)\})$

$\alpha g(\text{argmax}_S\{v(S) - \alpha g(S)\})$

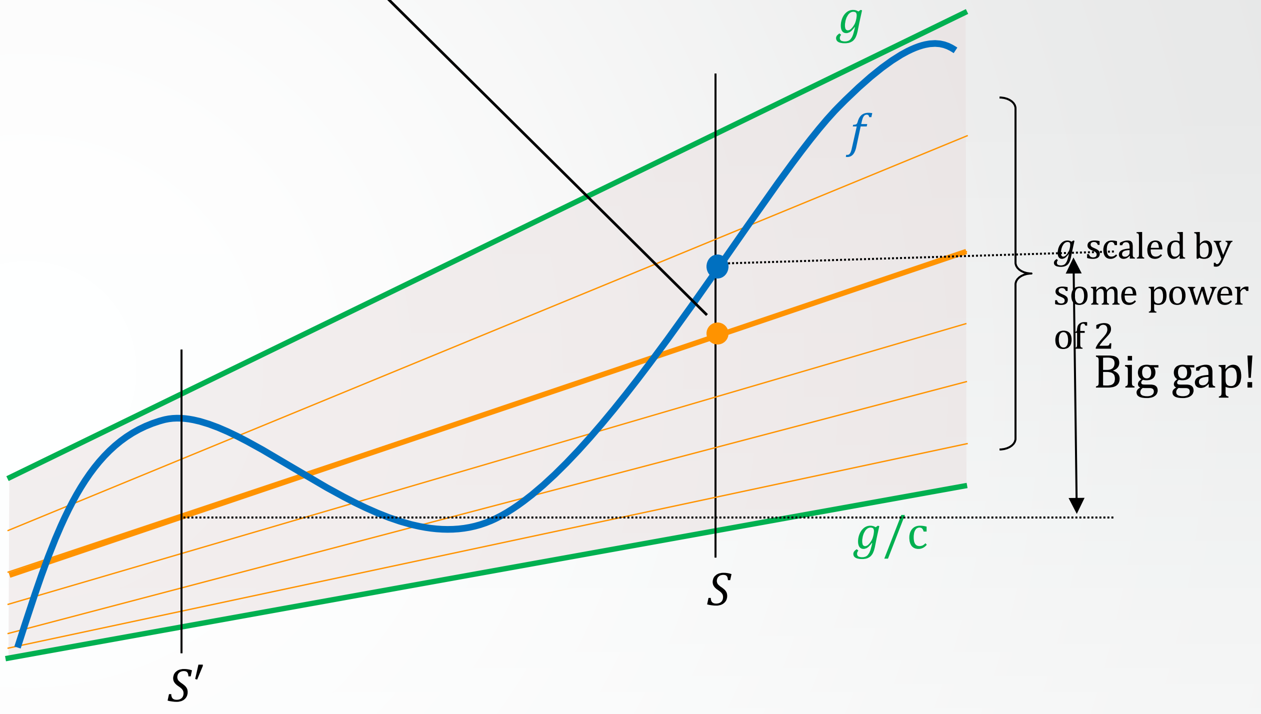


Lemma: Let $f$ and $g$ be any pricing functions such that $g$ **pointwise $c$-approximates** $f$.

Then there exists a distribution over scaling factors $\alpha > 0$, such that for any buyer,

The price paid by the buyer under $\alpha g \geq \frac{1}{2 \log 2c}$ (The price paid by the buyer under $f$)

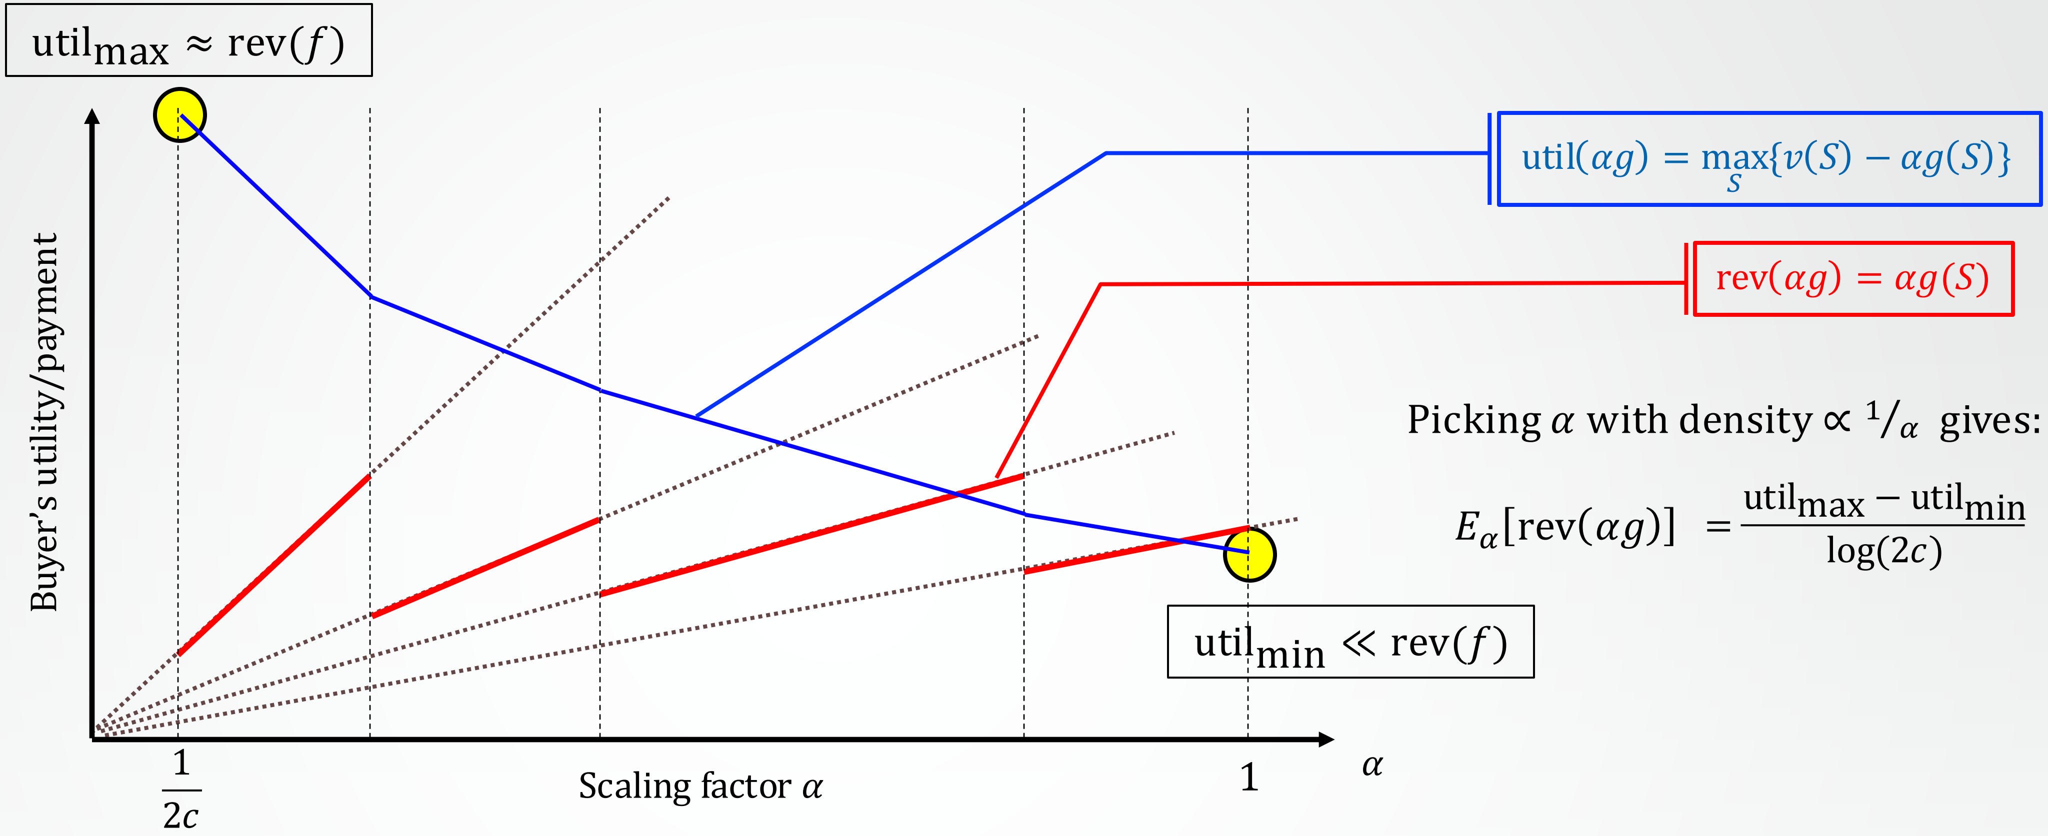


Picking $\alpha$ with density $\propto {}^1/_\alpha$ gives:

$$E_\alpha[\text{rev}(\alpha g)] = \frac{\text{util}_{max} - \text{util}_{min}}{\log(2c)}$$

Lemma: Let $f$ and $g$ be any pricing functions such that $g$ **pointwise $c$-approximates** $f$.

Then there exists a distribution over scaling factors $\alpha > 0$, such that for any buyer,

The price paid by the buyer under $\alpha g \geq \frac{\mathbf{1}}{\mathbf{2\,log\,2}c}$ (The price paid by the buyer under $f$)

# Part II: Revenue Maximization with a buy-many constraint – Approximation

[C.-Tzamos-Teng'19]

Theorem 1: For <u>any</u> value distribution,

$n$: #items

$$\text{Buy-many OPT} \le 2\log 2n \cdot \text{Item Pricing}$$

Theorem 2: There exists a distribution over additive valuations such that

$$\text{Buy-many OPT} \ge \Omega(\log n)\ \text{Revenue of any "succinct" mechanism}$$

Can get improved approximations for special valuation functions (e.g. "ordered" items)

Again, item pricing is the best "succinct" mechanism.

[C. Rezvan Tzamos Teng'21]

# Part II: Revenue Maximization – Summary

- For single buyer settings, item pricing or grand bundle pricing is the best "simple" mechanism.
- For multiple buyer settings:
  - Sequential posted price mechanisms
  - Price individual items as well as charge an "entry fee"
  - Generally not anonymous

[C. Hartline Malec Sivan'10, Yao'15, C. Miller'16, Cai-Zhao'17]

- For multiple buyer settings with buy-many constraint: nothing known yet!

# What else can posted prices do?

- Often the best simple/succinct mechanisms
- Suitable for online arrivals
- Robust – max-min optimal in some settings [Carrol'17]
- Learnable – polynomial pseudo-dimension [Morgenstern-Roughgarden'16]

Open direction: computing (approximately) optimal prices

# Thanks for your attention!

QUESTIONS?