Buy Many Mechanisms:

A new perspective on revenue-optimal mechanism design

SHUCHI CHAWLA

WISCONSIN UNIVERSITY OF WISCONSIN–MADISON

TEXAS The University of Texas at Austin

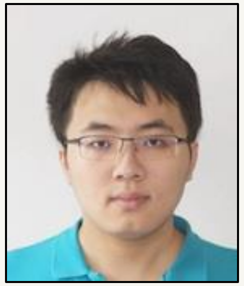

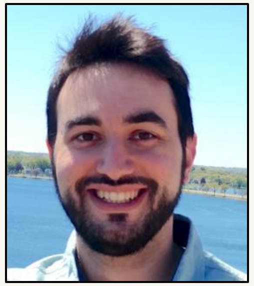

JOINT WORK WITH: YIFENG TENG & CHRISTOS TZAMOS

Revenue maximization with a single buyer

n items for sale

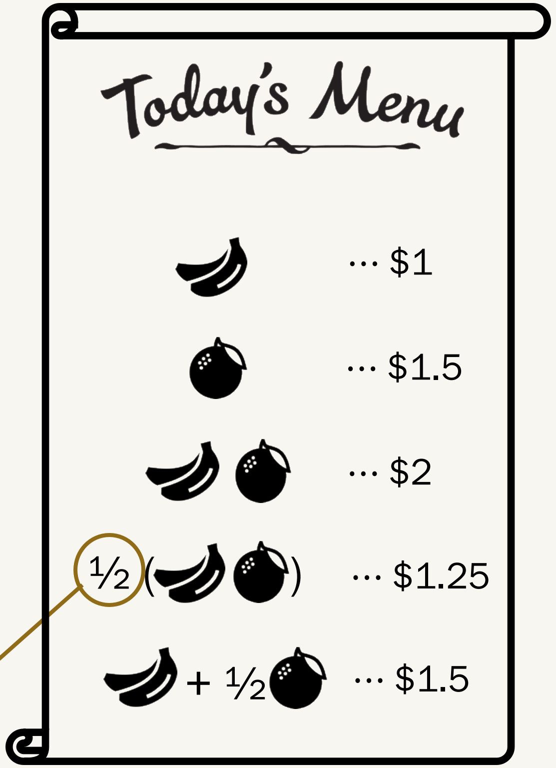

Probability of allocation

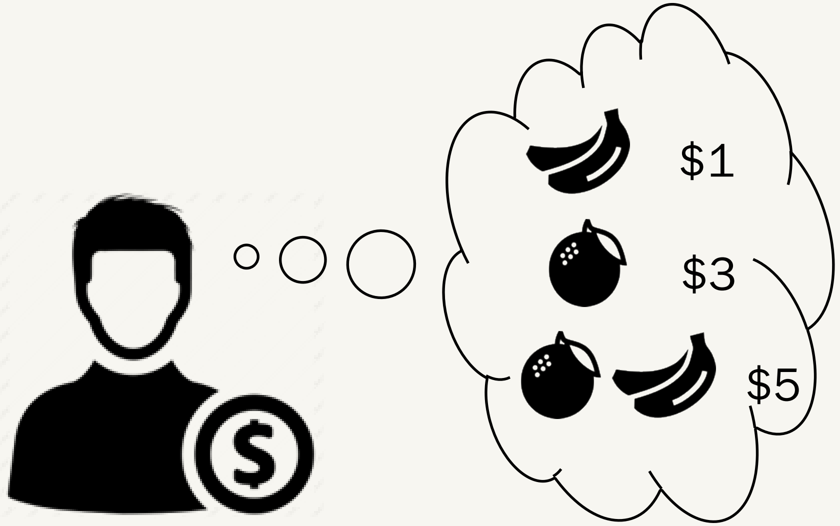

Drawn from some population
$v \sim D$

Value function
$v : 2^{[n]} \to \mathbb{R}^{+} \cup \{0\}$

Selling mechanism ≡ Menu of options / randomized options

Goal: maximize revenue $\mathrm{E}_{v\sim D}[\mathrm{Rev}_M(v)]$

What does the optimal menu look like? Can we find it?

Single item setting, $n = 1$ [Myerson'81]:

- Single menu option; no lotteries – just a fixed price

But for $n > 1$:

- Lotteries can improve revenue [Thanassoulis'04]
- Optimal mechanism has infinite number of lotteries [Hart-Nisan'13]
- Cannot hope to compute the optimal mechanism even in simple cases [Chen-Diakonikolas-Orfanou-Paparas-Sun-Yannakakis'15]

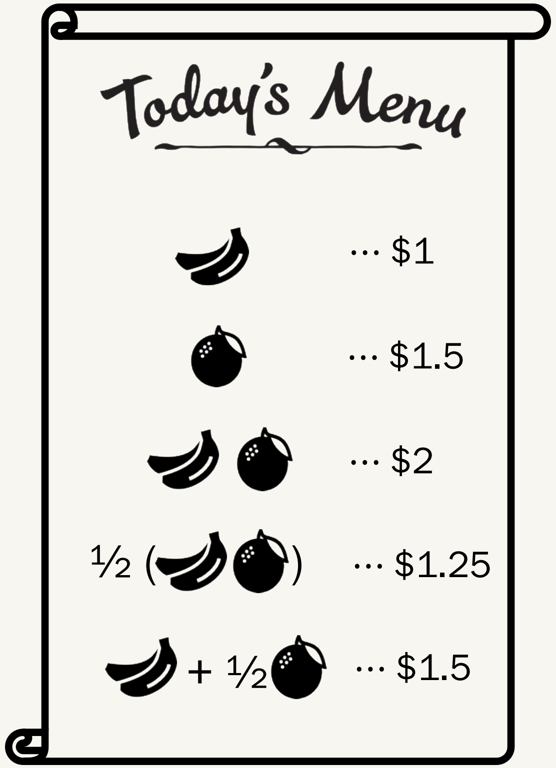

Goal: maximize revenue $\mathrm{E}_{v \sim D}[\mathrm{Rev}_M(v)]$

Can we get near-optimal revenue via a “simple” mechanism?

Item pricing : $p(S) = \sum_{i \in S} p_i$

The simplicity versus optimality tradeoff

Define:

$$\text{OPT}_D = \max_{\text{menu } M} \mathrm{E}_{v\sim D}[\text{Revenue of } M \text{ from } v]$$

$$\text{SRev}_D = \max_{\text{item pricings } p} \mathrm{E}_{v\sim D}[\text{Revenue of } p \text{ from } v]$$

Selling Separately

Approximation factor = $\max_{\text{distribution } D} \dfrac{\text{OPT}_D}{\text{SRev}_D}$

or $\dfrac{\text{OPT}_D}{\text{DRev}_D}$ or $\dfrac{\text{DRev}_D}{\text{SRev}_D}$

where DRev_D = $\max_{\text{determ. } M} \mathrm{E}_{v\sim D}[\text{Rev. } M \text{ from } v]$

We want the approximation factor to be as close to 1 as possible.

The simplicity versus optimality tradeoff

For a single item ($n = 1$), OPT = DRev = SRev [Myerson'81]

For $n > 1$, $\mathrm{OPT}/\mathrm{SRev}$ is small if:

- The value function is "nice" (e.g. additive or unit-demand)

  AND

- Values for different items are independent

[C. Hartline Kleinberg'07], [C. Malec Sivan'10], [Li Yao'13], [Babaioff Immorlica Lucier Weinberg '14], [Rubinstein Weinberg'15], ...

Without those assumptions:

- There exists an instance with a unit-demand buyer with $n = 2$ for which

  $\mathrm{OPT} = \infty$ and $\mathrm{SRev} < 1$

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

- There exists an instance with an additive buyer with $n = 2$ for which

  $\mathrm{OPT} = \infty$ and $\mathrm{Rev(any\ finite\ menu)} < \infty$ and $\mathrm{SRev} < 1$

  [Hart Nisan'13]

Is this the end of the story?

Optimal mechanisms can be "unreasonable"

Optimal mechanisms may charge super-additive prices.

Alternate approach: optimize over "reasonable" mechanisms

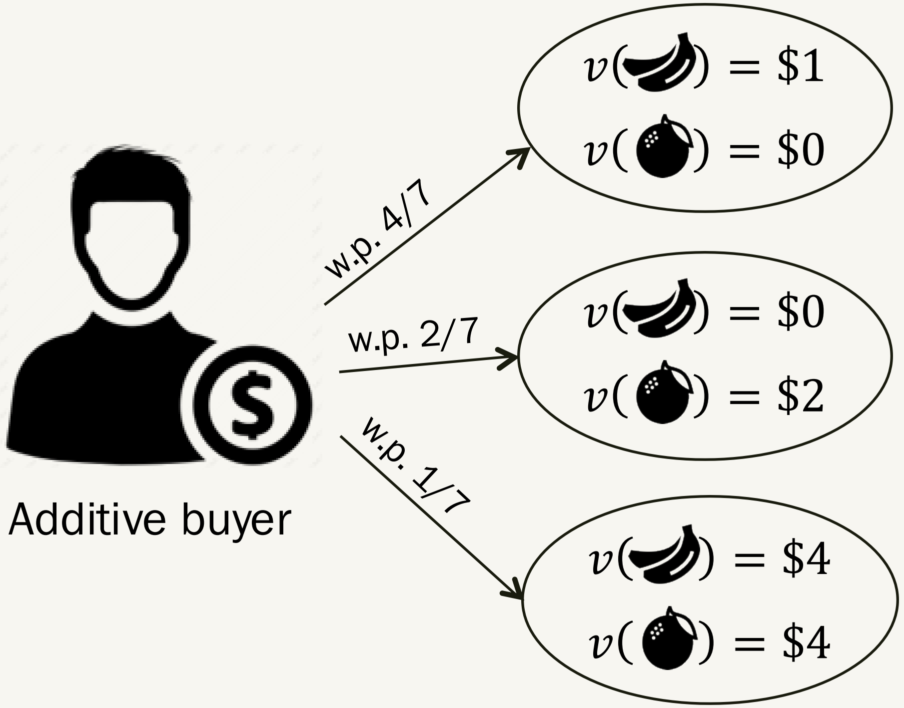

Additive buyer

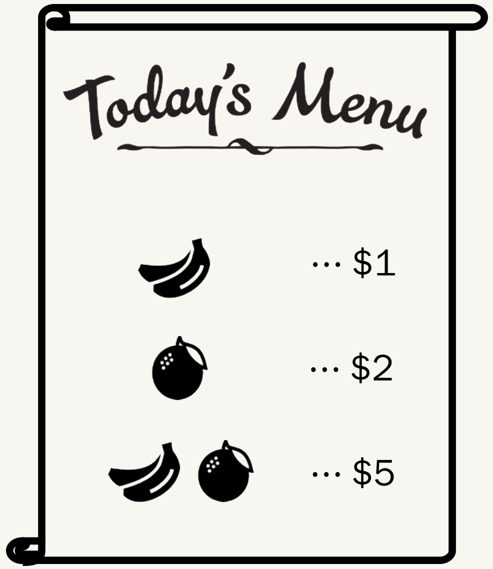

Optimal deterministic menu

Buy-many mechanisms

- In a buy-many strategy, a buyer can purchase any multi-set of menu options at the sum of their prices. The buyer obtains an independent draw from each option.
- A menu is "buy-many" if the random allocation resulting from any buy-many strategy is "dominated" by a single menu option.

Cheaper price; Bigger allocation

- For deterministic pricings, buy-many ≡ subadditivity

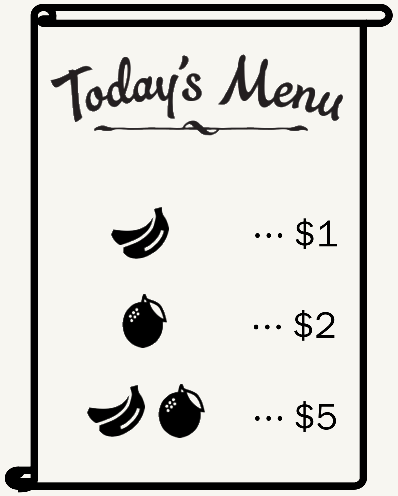

Not buy-many

Not buy-many

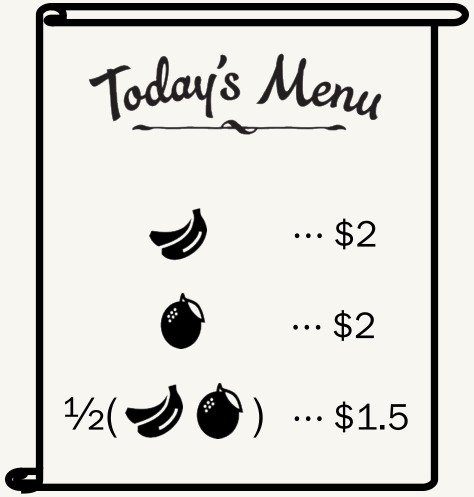

Buy-many

Buy-many mechanisms

- In a buy-many strategy, a buyer can purchase any multi-set of menu options at the sum of their prices. The buyer obtains an independent draw from each option.

New goal: Study the properties and approximability of optimal buy-many mechanisms

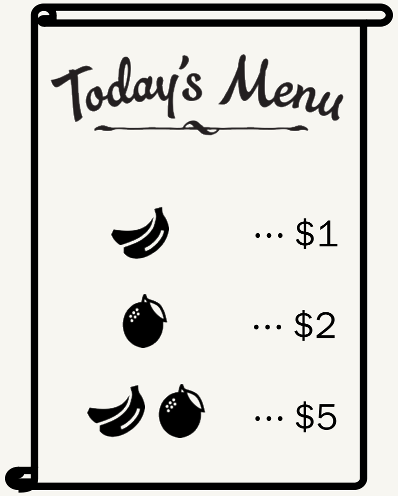

Not buy-many

Not buy-many

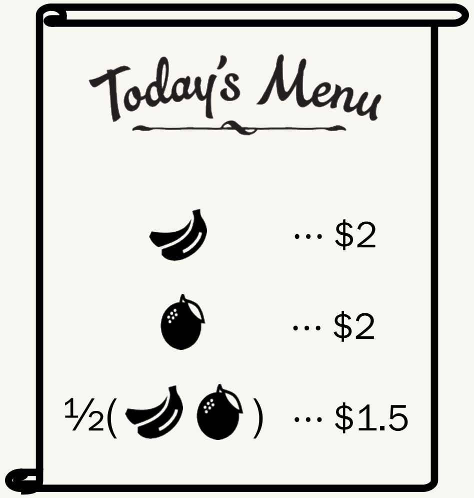

Buy-many

Approximability and other properties of Buy-Many mechanisms

Optimal buy-many mechanisms can be well approximated

[C. Teng Tzamos'19]

Theorem 1: For any value distribution D,

$$\text{Buy}-\text{many optimal revenue}_D \le 2\log(2n)\ \ \text{SRev}_D$$

For example, for $n = 2$ items, we can have $\text{OPT}_D = \infty$ and $\text{SRev}_D < \infty$

But we always have $\text{SRev}_D > 0.36\ \text{Buy}-\text{many Rev}_D$

Can get better bounds in some special cases e.g. "ordered" items [C. Rezvan Teng Tzamos'21]

Previous work showed…

[Briest C. Kleinberg Weinberg'10]: For any distribution D over unit-demand valuations, Buy-many Rev $\le O(\log n)$ SRev.

[Babaioff Nisan Rubinstein'18]: ∃ product distributions over additive values for which Buy-many Rev < OPT.

Optimal buy-many mechanisms can be well approximated

[C. Teng Tzamos'19]

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

$$\text{Buy-many optimal revenue}_D \leq 2\log(2n)\ \ \text{SRev}_D$$

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

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

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

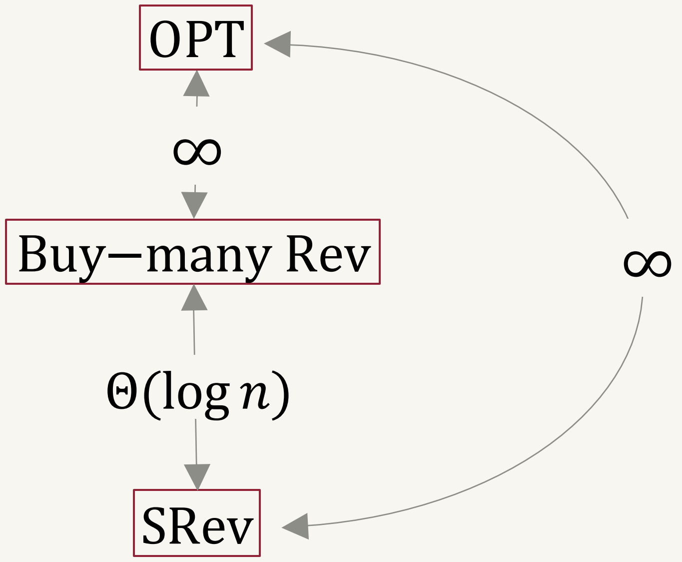

What about a 99% approximation to optimal revenue?

Menu size complexity: min number of menu options needed to describe the mechanism [Hart-Nisan'13]

How many menu options do we need to get 99% of the optimal revenue?

- Infinitely many in general [Hart-Nisan'13]
- Finite (but exponential in n) only known in settings where the buyer has "nice" values over independent items [Babaioff et al.'17, Kothari et al.'19, ...]

Theorem 3: For <u>any</u> value distribution D and $\epsilon \in [0,1]$, there exists a menu M of finite size $f(n, \epsilon)$, such that,

$$\mathrm{Rev}_D(M) \geq (1-\epsilon)\mathrm{Buy{-}many\ Rev}_D$$

[C. Teng Tzamos'20]

- Need $f(n, \epsilon) = (1/\epsilon)^{2^{O(n)}}$.
- Tight: any smaller menu will only get an $O(\log n)$ fraction of the revenue.

Revenue monotonicity

Suppose that values of all buyers in the market increase (but non-uniformly).

What happens to the optimal revenue?

- Single item: revenue increases
- General multi-item settings: revenue may decrease! [Hart-Reny'15]

What about buy-many mechanisms?

- Optimal revenue may decrease [C. Teng Tzamos'20]

  ... but not by much.

Revenue continuity

Suppose that values of all buyers in the market change by small multiplicative amounts:

Every $v \sim D$ is perturbed to v' such that $\forall S \subseteq [n], v'(S) \in (1 \pm \epsilon)v(S)$.

What happens to the optimal revenue?

- Single item: revenue changes slightly, by $1 \pm O(\epsilon)$
- General multi-item settings: revenue can change significantly!
  - $\mathrm{OPT}_D = \infty$ and $\mathrm{OPT}_{D'} < \infty$ [Psomas et al.'19]

Theorem 4: For <u>any</u> value distribution D and <u>any</u> multiplicative perturbation D':

$$\mathrm{Buy{-}many\ Rev}_{D'} \geq (1 - \mathrm{poly}(n, \epsilon))\mathrm{Buy{-}many\ Rev}_D$$

The dependence on n is necessary

What makes Buy-Many mechanisms so well behaved?

What makes buy-many menus well-behaved?

Observation 1:

- If x and x' are two "close enough" random allocations, they cannot be priced very differently.

  $\Longrightarrow$ mechanism can only price discriminate to a limited extent.

Observation 2:

- If v and v' are two "close enough" valuations resulting in very different payments, the buyer's payment at these values is much lower than his utility

  $\Longrightarrow$ such buyers cannot contribute too much to optimal revenue

Observation 3:

- Additive pricings point-wise n-approximate buy-many menus

A useful technical lemma

Point-wise approximation $\Longrightarrow$ approximation in revenue

Given any pricing functions f and g such that for all random allocations Λ,

$$\frac{1}{c}\, g(\Lambda) \le f(\Lambda) \le g(\Lambda).$$

there exists a distribution over scaling factors $\alpha > 0$, such that for any value function v,

$$\mathrm{E}_{\alpha}[\mathrm{Rev}_v(\alpha g)] \ge \frac{1}{2\log 2c}\mathrm{Rev}_v(f).$$

(Interpret any single buyer mechanism as a function that maps lotteries to prices.)

What makes buy-many menus well-behaved?

Observation 1:

- If x and x' are two "close enough" random allocations, they cannot be priced very differently.

  $\implies$ mechanism can only price discriminate to a limited extent.

Observation 2:

- If v and v' are two "close enough" valuations resulting in very different payments, the buyer's payment at these values is much lower than his utility

  $\implies$ such buyers cannot contribute too much to optimal revenue

Observation 3:

- Additive pricings point-wise n-approximate buy-many menus

  $\implies$ $O(\log n)$ approximation in revenue

Summary

Main idea: instead of restricting the market, simplify the optimization by introducing "reasonable" constraints

- Buy-many constraint is reasonable; frequently satisfied
- Buy-many mechanisms exhibit many nice properties
- Buy-many mechanisms can be well-approximated via item pricing

- Some interesting open directions:
  - Multiple buyers: what does the buy-many constraint mean in limited supply settings?
  - Exact computation? The buy-many constraint is not a linear constraint.

Thank you!