Revisiting Revenue Maximization for Many Buyers

SHUCHI CHAWLA

TEXAS
The University of Texas at Austin

JOINT WORK WITH: DIMITRIS CHRISTOU, TRUNG DANG, ZHIYI HUANG,
GREG KEHNE, ROJIN REZVAN, YIFENG TENG & CHRISTOS TZAMOS

Revenue Maximization with Many Buyers

A seller with m items for sale

n buyers drawn from some population

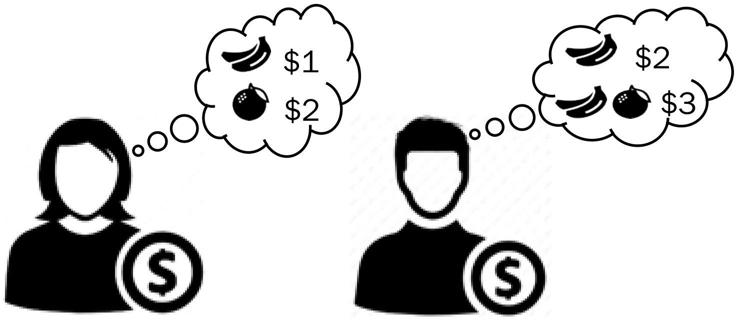

$v_1 \sim D_1$ $v_2 \sim D_2$

Optimal mechanisms can be complicated – even for just one buyer

- Sell random allocations a.k.a. lotteries
- Offer infinitely large menus

[Thanassoulis'04]
[Hart Nisan'13]

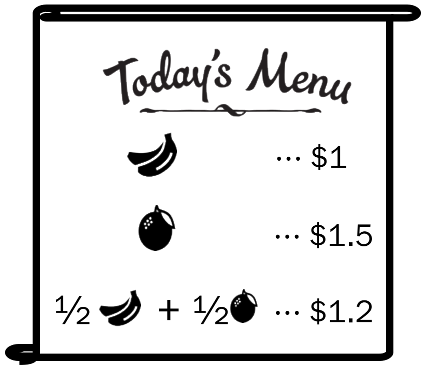

Mechanism:

$(v_1, \ldots, v_n) \rightarrow (S_1, \ldots, S_n; p_1, \ldots, p_n)$

Allocations Payments

Value functions

$v_i : 2^{[m]} \rightarrow \mathbb{R}^+ \cup \{0\}$

Buyer i's objective: maximize utility $:= v_i(S_i) - p_i$

Seller's objective: maximize revenue $:= \sum_i p_i$

Revenue Maximization with Many Buyers

A seller with m items for sale

n buyers drawn from some population

$v_1 \sim D_1$ $v_2 \sim D_2$

Mechanism:
$(v_1, \ldots, v_n) \rightarrow (S_1, \ldots, S_n; p_1, \ldots, p_n)$
Allocations Payments

Value functions
$v_i: 2^{[m]} \rightarrow \mathbb{R}^+ \cup \{0\}$

Buyer i's objective: maximize utility $:= v_i(S_i) - p_i$

Seller's objective: maximize revenue $:= \sum_i p_i$

Optimal mechanisms can be complicated – even for just one buyer

- Sell random allocations a.k.a. lotteries
- Offer infinitely large menus

Buyers impose externalities on each other

- Is it better to allocate a limited supply item to buyer 1 or buyer 2?
- Allocation and pricing can be inscrutable to buyers

Question: Is there a "simple" mechanism that approximates the optimal one?

The simplicity vs optimality tradeoff

- VCG with reserve pricing
- Additive pricing, a.k.a. item pricing
- Grand bundle pricing
- Two-part tariffs (i.e. subscription fees + item pricing)

…

Revenue Gap?

- random allocations a.k.a. lotteries
- infinitely large menus

Revenue Gap, a.k.a., approximation factor =

$$\max_{\text{distribution } D} \frac{\text{OPT}(D)}{\text{IP}-\text{Rev}(D)}$$

Item pricing : $p(S)$

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

IP-Rev = $\max_{\text{Item Pricing } p} \text{E}_{v\sim D}[\text{Revenue of } p \text{ from } v]$

…erent “simple” mechanisms

The simplicity vs optimality tradeoff

- VCG with reserve pricing
- Additive pricing, a.k.a. item pricing
- Grand bundle pricing
- Two-part tariffs (i.e. subscription fees + item pricing)

…

Revenue Gap?

- random allocations a.k.a. lotteries
- infinitely large menus

Good news: The gap is small (constant factor) in many settings

[Chawla Hartline Kleinberg'07] [Hartline Roughgarden'09] [Chawla Malec Sivan'10] [Li Yao'13] [Babaioff Immorlica Lucier Weinberg'14] [Yao'14] [Rubinstein Weinberg'15] [Chawla Miller'16] [Cai Devanur Weinberg'16] [Cai Zhao'17] [Feng Hartline Li'19] [Kothari Mohan Schvartzman Singla Weinberg'19] …

Question: What else can we say about the gap in the absence of item independence?

Bad news: **Requires** strong assumptions – independence of values across items – even for the single buyer setting

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

There exist value distributions for which $\frac{\text{OPT}}{\text{IP-Rev}} = \infty$

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

IP-Rev = $\max_{\text{Item Pricing } p} \mathrm{E}_{v \sim D}[\text{Revenue of } p \text{ from } v]$

Part 1: Single-buyer settings

A new benchmark
Approximation and other properties

Towards a new benchmark: limiting the power of the seller

[C. Teng Tzamos'19]

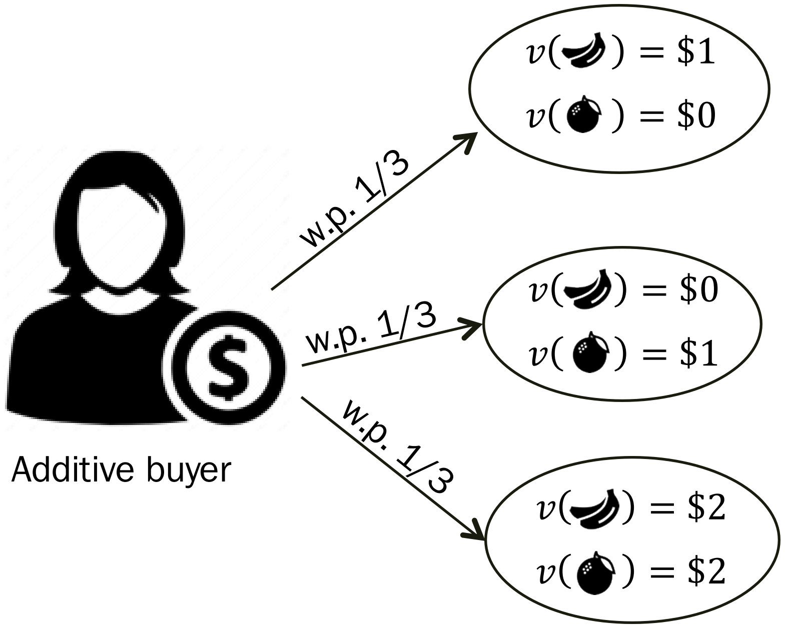

Buy-Many mechanisms: the buyer can purchase any multi-set of menu options at the sum of their prices. The buyer obtains an independent draw from each option.

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

Towards a new benchmark: limiting the power of the seller

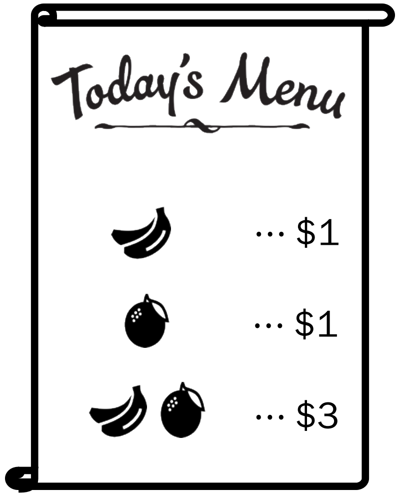

Not buy-many

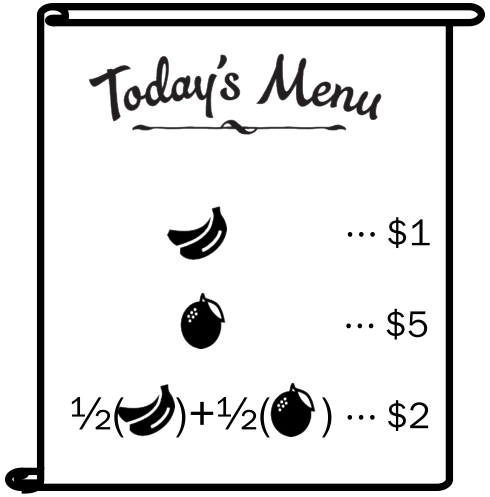

Not buy-many

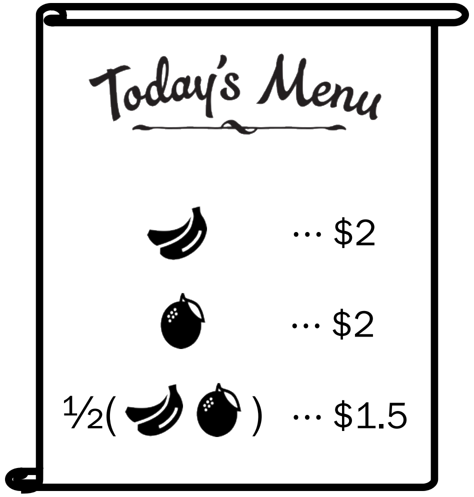

Buy-many

buy-many
$\approx$
subadditivity

Buy-Many mechanisms: the buyer can purchase any multi-set of menu options at the sum of their prices. The buyer obtains an independent draw from each option.

$$\text{BuyManyRev}(D) = \max_{\text{BuyMany menu } M} \mathbb{E}_{v \sim D}[\text{Revenue of } M \text{ from } v]$$

A menu is “buy-many” if the random allocation resulting from any buy-many strategy is “dominated” by a single menu option.

New question: Can simple mechanisms approximate BuyManyRev?

Optimal buy-many mechanisms can be well approximated

[Briest C. Kleinberg Weinberg'10]

[C. Teng Tzamos'19]

Theorem 1: For any value distribution D,

$$\text{BuyManyRev}_D \le 2\log(2n)\ \text{IP}-\text{Rev}_D$$

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

But we always have $\text{IP}-\text{Rev}_D > 0.36\ \text{BuyManyRev}_D$

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

Optimal buy-many mechanisms can be well approximated

[Briest C. Kleinberg Weinberg'10]
[C. Teng Tzamos'19]

Theorem 1: For any value distribution D,

$$\mathrm{BuyManyRev}_D \leq 2\log(2n)\ \ \mathrm{IP{-}Rev}_D$$

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

$\mathrm{BuyManyRev} \geq \Omega(\log n)$ Revenue of any "succinct" mechanism

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

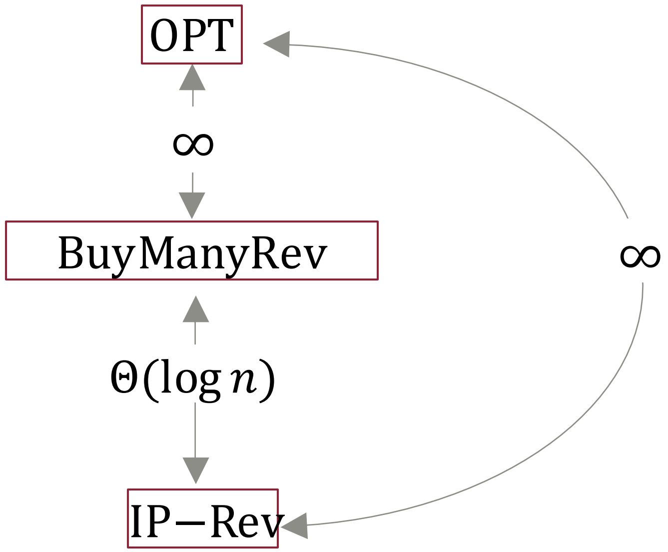

Item pricing : $p(S) = \sum_{i \in S} p_i$

Better approximations for structured unit-demand value functions

[C. Rezvan Teng Tzamos'22]

- FedEx setting: items are ordered and each buyer desires a minimum quality level

  i.e. for buyer of type (v, i), $v_1 = v_2 = \cdots = v_i = v$ and $v_{\{i+1\}} = \cdots = v_n = 0$.

  $\mathrm{IP{-}Rev} = \mathrm{BuyManyRev}$

- Ordered values: $v_1 \geq v_2 \geq \cdots \geq v_n$ for all buyers

  i.e. item 1 is better than item 2, which is better than item 3, and so on.

  $\mathrm{IP{-}Rev} \geq O(1)\ \mathrm{BuyManyRev}$

- k sets of ordered items (i.e. partial order of width k)

  $\mathrm{IP{-}Rev} \geq O(\log k)\ \mathrm{BuyManyRev}$

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,

[C. Teng Tzamos'20]

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

- 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':

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

The dependence on n is necessary

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 similar allocations but very different payments, either a slight perturbation of prices removes this price discrimination, or such buyers cannot contribute too much to optimal revenue

Observation 3:

- Additive pricings point-wise n-approximate buy-many menus (a.k.a. subadditive pricings)

Lemma: Let p and q be any pricing functions such that q pointwise c-approximates p.

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

(The price paid by the buyer under αq) $\geq \frac{1}{2\, log\, 2c}$ (The price paid by the buyer under p)

It’s just a benchmark…

- We don’t know how to compute an optimal buy-many menu

- We don’t know how a buyer would find the optimal buy-many strategy given a menu.

Part 2: Multi-buyer settings

Extending the benchmark
Approximation

An attempt to extend the benchmark

- Single-buyer definition: the buyer can purchase any multi-set of menu options at the sum of their prices. I.e., buyer can participate in the mechanism multiple times.
- Multi-buyer definition: allow each buyer to participate in the mechanism many times?

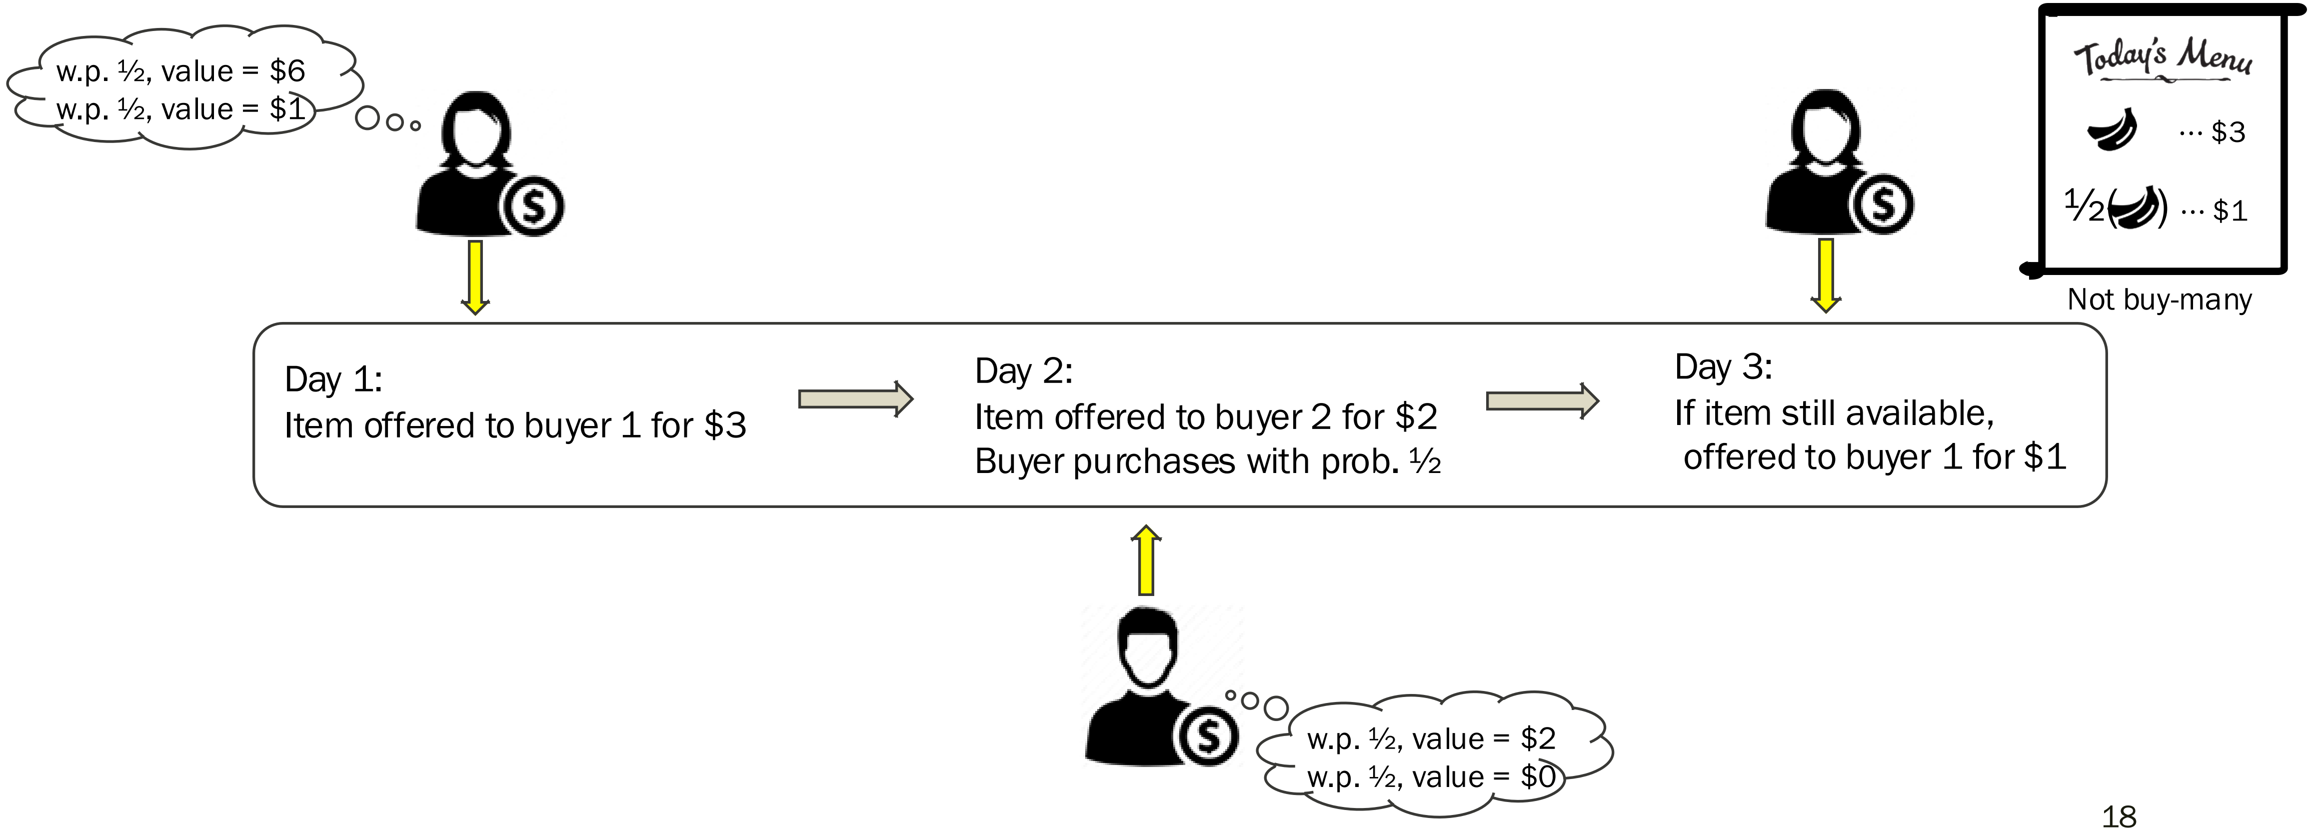

Two facets of the buy-many constraint

As a strengthening of buyers' power

- Buyer can interact with mechanism multiple times
- Unlimited supply $\Longrightarrow$ static menu across interactions
- Limited supply $\Longrightarrow$ seller can distinguish between and discriminate across different interactions

Offers no benefit over original OPT

As a weakening of the seller's power

- Single interaction between buyer and seller, but,
- Super-additive pricing is simply disallowed in buyer-seller interaction

We will use this as the basis of our new benchmark

Not the same in multiple buyer settings!

Buy-many mechanisms for multiple buyers: Take 1

1. Seller interacts with each buyer once; in some arbitrary sequence; offers buy-many menu

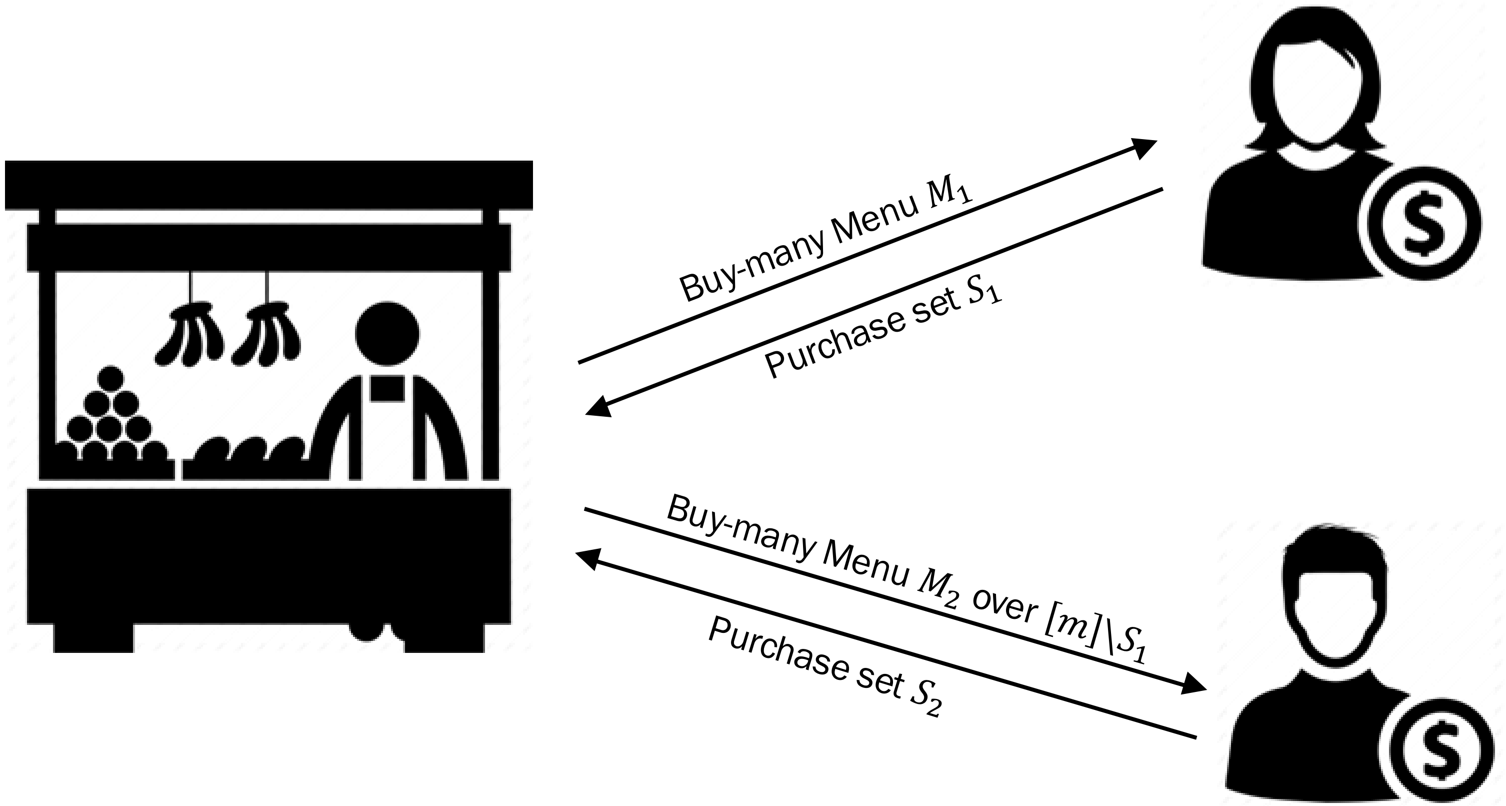

Buy-many mechanisms for multiple buyers: Take 2

1. Seller interacts with each buyer once; in some arbitrary sequence; offers buy-many menu
2. Seller runs a direct mechanism; "ex-post" pricing offered to each buyer is buy-many

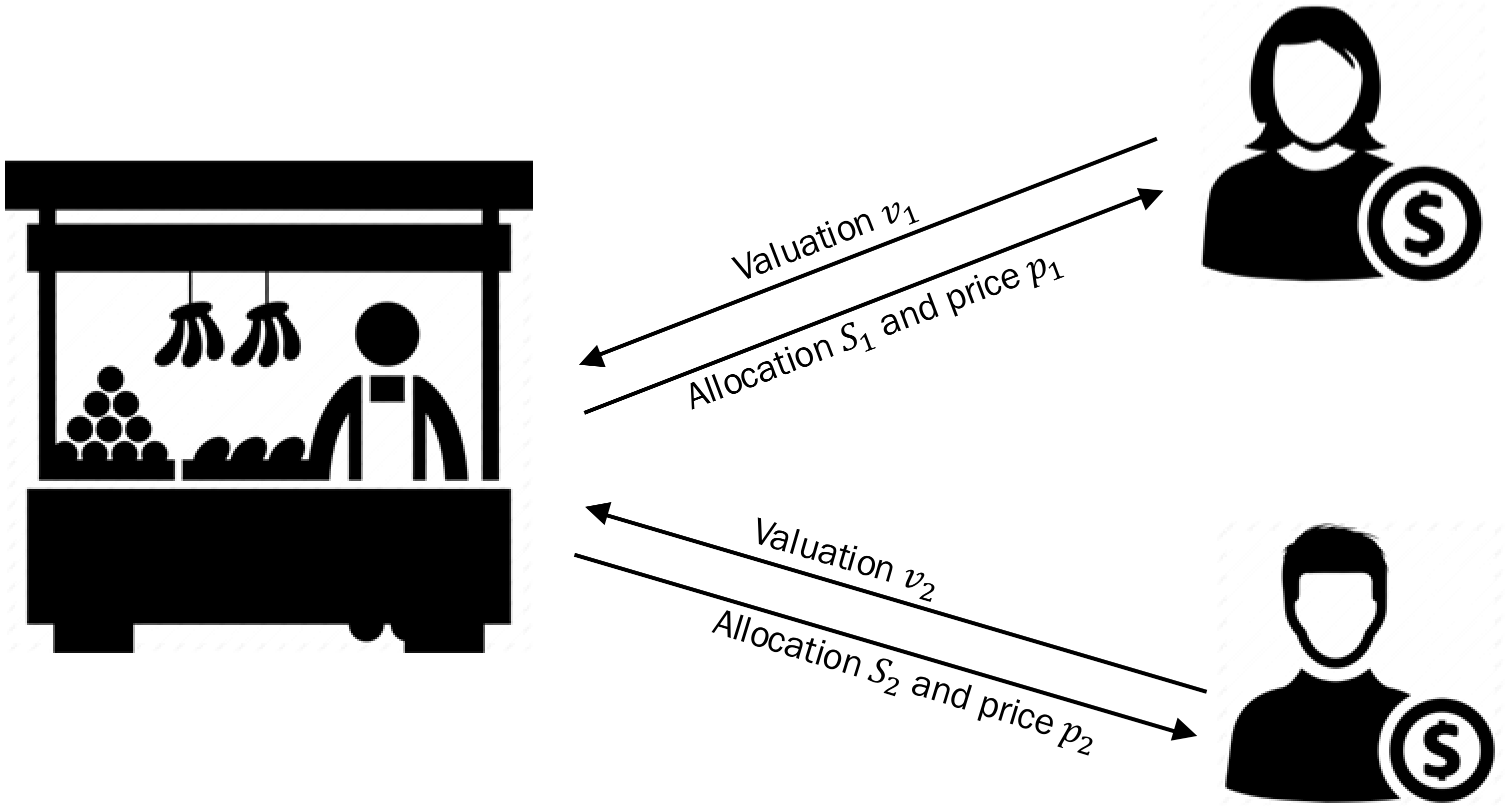

$\{(S_1, p_1)\}$ over possible choices of v_2 form a buy-many menu.

$\{(S_2, p_2)\}$ over possible choices of v_1 form a buy-many menu.

Buy-many mechanisms for multiple buyers: Take 3

1. Seller interacts with each buyer once; in some arbitrary sequence; offers buy-many menu
2. Seller runs a direct mechanism; “ex-post” pricing offered to each buyer is buy-many
3. Seller runs a direct mechanism; “ex-ante” pricing offered to each buyer is buy-many

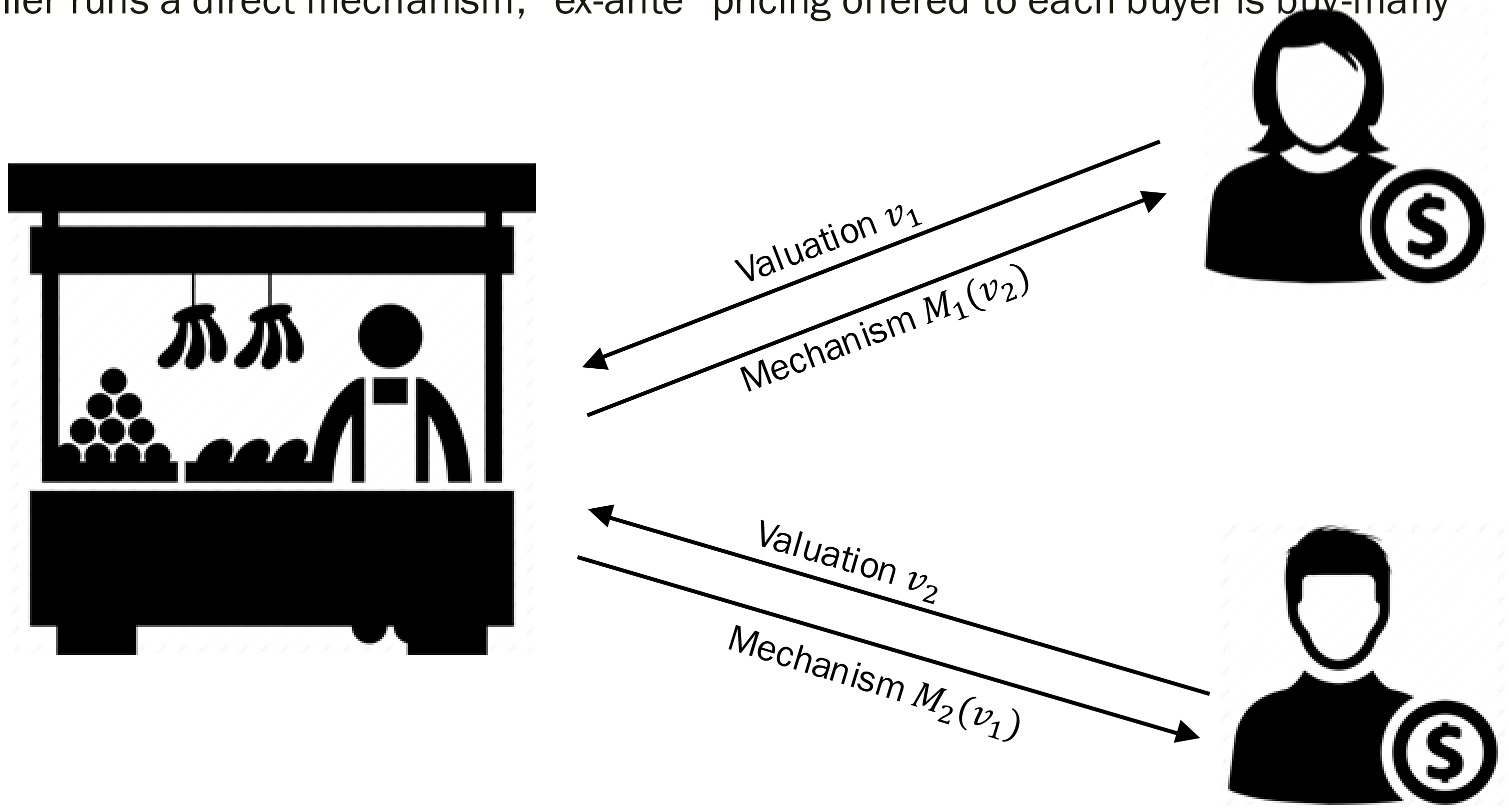

$E_{v_2}[M_1]$ forms a buy-many menu.

$E_{v_1}[M_2]$ forms a buy-many menu.

Buy-many mechanisms for multiple buyers: Take 3, …

1. Seller interacts with each buyer once; in some arbitrary sequence; offers buy-many menu
2. Seller runs a direct mechanism; “ex-post” pricing offered to each buyer is buy-many
3. Seller runs a direct mechanism; “ex-ante” pricing offered to each buyer is buy-many
4. …

Which definition should we use???

The Buy-Many benchmark for many buyers [C. Rezvan Teng Tzamos'23]

- Many ways to decompose a multi-buyer mechanism into its single buyer constituents
- Ex Ante relaxation – a convenient upper bound that captures many of these extensions. [Alaei'11]

Relax the ex post supply constraint to an ex ante supply constraint:

In expectation over

"~~For every instantiation of~~ $v_i \sim D_i$, each item is allocated to at most one buyer"

- Let $M_1, M_2, \ldots$ be buy-many mechanisms such that

$(x_1, x_2, \ldots, x_n)$ is ex ante feasible if $\sum_i x_{ij} \leq 1$ for all j

$\sum_i \Pr[M_i \text{ sells item } j \text{ to buyer } i] \leq 1$ for all items j.

- ExAnte-BuyMany-OPT is the most revenue that can be obtained by such a tuple.

The Buy-Many benchmark for many buyers [C. Rezvan Teng Tzamos'23]

- Many ways to decompose a multi-buyer mechanism into its single buyer constituents
- Ex Ante relaxation – a convenient upper bound that captures many of these extensions. [Alaei'11]

Relax the ex post supply constraint to an ex ante supply constraint:

In expectation over

"~~*For every instantiation of*~~ $v_i \sim D_i$, *each item is allocated to at most one buyer*"

Mechanism M satisfies ex ante supply constraint $y \in \mathbb{R}^m$ with respect to value distribution D if:
for all items j, $\Pr_{v \sim D}[\text{buyer with value } v \text{ buys } j \text{ in } M] \leq y_j$.

$\text{BMRev}(D, y)$ ≔ the most revenue achievable by a buy-many mechanism that satisfies y.

$(x_1, x_2, \dots, x_n)$ is ex ante feasible if $\sum_i x_{ij} \leq 1$ for all j

ExAnte-BM-OPT= $\max_{\text{Ex ante feasible } x} [\, \sum_i \text{BMRev}(D_i, x_i) \,]$

ExAnte-IP-OPT= $\max_{\text{Ex ante feasible } x} [\, \sum_i \text{ItemPricingRev}(D_i, x_i) \,]$

For any distributions $D_1, \dots, D_n$ over m items, ExAnte-BM-OPT $\leq O(\log m) \cdot$ ExAnte-IP-OPT

Tight! [C. Rezvan Teng Tzamos'23]

Ex Ante versus Ex Post feasible mechanisms

w.p. ½ : only for \$1
w.p. ½ : only for \$3

w.p. ½: for \$0; for \$0; Pair for \$5
w.p. ½: \$0 for all allocations

Ex Ante optimal item pricing:
- To Alice, offer $p_1 = (1,3)$. Allocation probabilities are (½, ½).
- To Bob, offer $p_2 = (3,2)$. Allocation probabilities are (½, ½).

Ex Ante supply constraints are met. ExAnte-IP-OPT = \$4.50.

Ex post feasible item pricing:
- If any item is sold to Alice, cannot extract any revenue from Bob.
- The above sequential pricing gets revenue \$2.

Better sequential pricing:
- To Alice, offer $p_1 = (3,3)$. is bought w.p. ½.
- To Bob, offer $p_2 = (3,2)$. Pair is bought w.p. ¼.

Revenue ${}^3/_2 + {}^5/_4 = \$2.75$.

Approximating the Ex Ante Relaxation: two components

ExAnte-BM-OPT = $\max_{\text{feasible } x} \sum_i \text{BMRev}(D_i, x_i)$

$\Theta(\log m)$

ExAnte-IP-OPT = $\max_{\text{feasible } x} \sum_i \text{ItemPricingRev}(D_i, x_i)$

??

Sequential item pricing

$\text{BMRev}(D_i, x_i) =$ Max revenue from BM mech M that satisfies x_i w.r.t. D_i

$\text{ItemPricingRev}(D_i, x_i) =$ Max revenue from item pricing p_i that satisfies x_i w.r.t. D_i

- Buyers arrive in sequence
- Buyer 1 is offered IP p_1 and purchases S_1.
- Buyer 2 is offered IP p_2 over $[m] \setminus S_1$ and buys S_2.
- Buyer 3 is offered IP p_3 over $[m] \setminus (S_1 \cup S_2)$ ...

...

Depends on buyers' valuation functions

- Unit demand or additive: 2
- Subadditive: $O(\log m)$
- ...

[C. Christou Dang Huang Kehne Rezvan'24]

Question: Is there a “simple” mechanism that approximates the “optimal” one?

Yes, under “mild” assumptions and against an “appropriate” benchmark.

Item Pricing approximates the Buy Many Optimum in arbitrary single buyer settings within a factor of $O(\log \#\text{items})$.

Sequential item pricing approximates the ex-ante Buy Many benchmark when:

- All buyers have gross substitute values – within a factor of $O(\log \#\text{items})$
- All buyers have subadditive values – within a factor of $O(\log^2 \#\text{items})$

No further assumptions necessary.

Cf.: under item independence, sequential two-part tariffs approximate the ex-ante opt revenue when:

- All buyers have gross substitute values – within a factor of $O(1)$ [C. Miller’16]
- All buyers have subadditive values – within a factor of $O(\log \log \#\text{items})$ [Cai Zhao’17, Duetting Kesselheim Lucier’20]

Thanks!