Revenue maximization with the Buy Many constraint

SHUCHI CHAWLA
UNIVERSITY OF WISCONSIN-MADISON

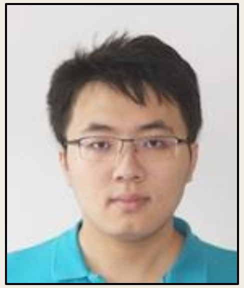

JOINT WORK WITH: YIFENG TENG & CHRISTOS TZAMOS

Revenue maximization with a single buyer

n items for sale

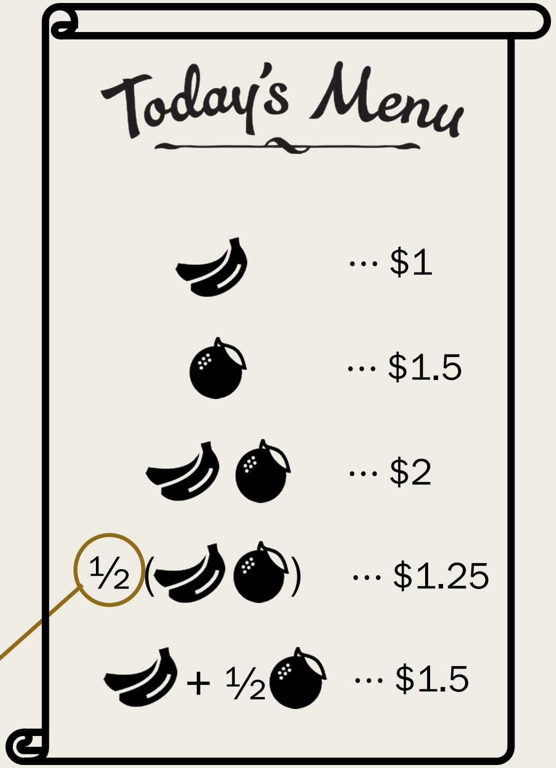

Probability of allocation

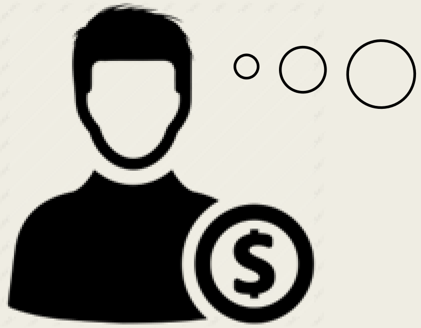

Drawn from some population

$v \sim D$

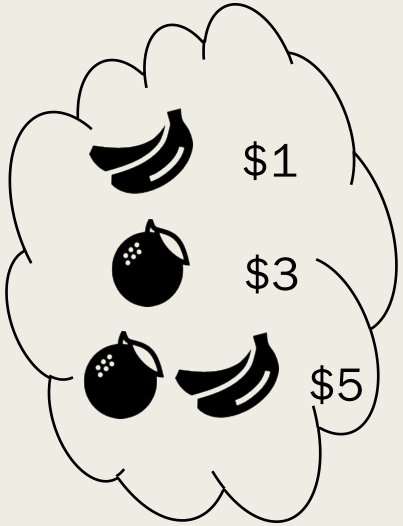

Value function

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

Selling mechanism ≡ Menu of randomized options

What does the optimal menu look like?

- Is randomness necessary?

  Yes for $n > 1$ [Thanassoulis'04]

  (No for $n = 1$) [Myerson'81]

- How many menu options?

  Unbounded in n for $n > 1$ [Hart-Nisan'13]

  (One for $n = 1$)

- Is the optimal mechanism easy to compute?

  No, not even in simple cases! [Chen-Diakonikolas-Orfanou-Paparas-Sun-Yannakakis'15]

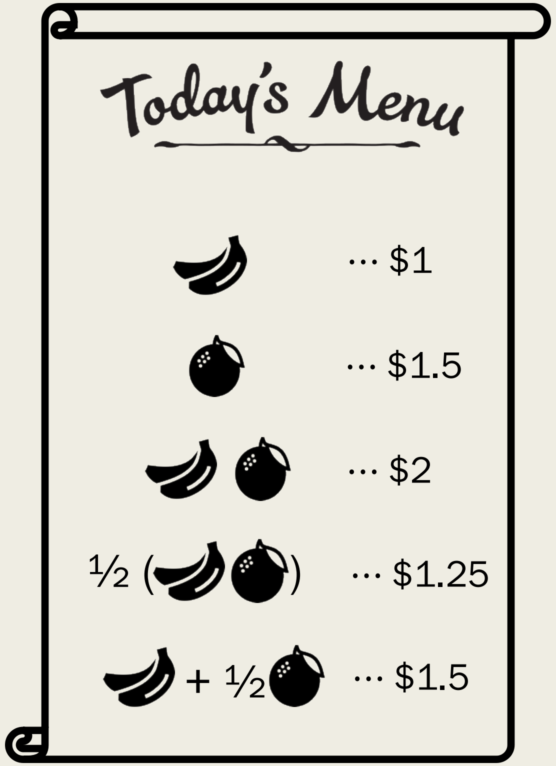

Can we approximate revenue?

Two principal approaches…

Approach # 1: computational approach

Exact optimization when the value distribution has small support

$\Longleftarrow$ the optimum is the solution to an LP

Exact!

Strong assumption

Complicated/impractical solution

Approach # 2: approximation for “nice” valuation functions

If values for different items are independent:

- Unit-demand valuations $\Longrightarrow$ item prices give a 4-approximation
- Additive valuations $\Longrightarrow$ item or grand bundle pricing gives a 6-approximation
- Subadditive valuations $\Longrightarrow$ item or grand bundle pricing gives an $O(1)$-approx

[C. Hartline Kleinberg’07, C. Hartline Malec Sivan’10, C. Malec Sivan’10, Li Yao’13, Babaioff Immorlica Lucier Weinberg’14, Rubinstein Weinberg’15, Kothari Mohan Schvartzman Singla Weinberg’19, …]

Posted pricing a.k.a. the grocery store mechanism

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

Grand bundle pricing : $p(S) = p([n])$

Approach # 2: approximation for “nice” valuation functions

If values for different items are independent:

- Unit-demand valuations ⟹ item prices give a 4-approximation
- Additive valuations ⟹ item or grand bundle pricing gives a 6-approximation
- Subadditive valuations ⟹ item or grand bundle pricing gives an $O(1)$-approx

Can relax a little bit
[C. Malec Sivan’10, Psomas Schvartzman Weinberg’19]

[C. Hartline Kleinberg’07, C. Hartline Malec Sivan’10, C. Malec Sivan’10, Li Yao’13, Babaioff Immorlica Lucier Weinberg’14, Rubinstein Weinberg’15, Kothari Mohan Schvartzman Singla Weinberg’19, …]

Simple practical solutions

Unrealistic assumptions

What about arbitrary value distributions?

With two items, there exists an instance with a unit-demand buyer for which:

- Optimal revenue = ∞
- Item pricing revenue < some constant
- Revenue of any deterministic mechanism < some constant

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

These large gaps do not go away we perturb values drawn from a worst case distribution by small amounts. [Psomas Schvartzman Weinberg'19]

A failure of theory!

What hope do we have for revenue maximization in a real-world setting?

Alternate approach: optimize over "reasonable" mechanisms

Optimal mechanisms can be "unreasonable": charge super-additive prices

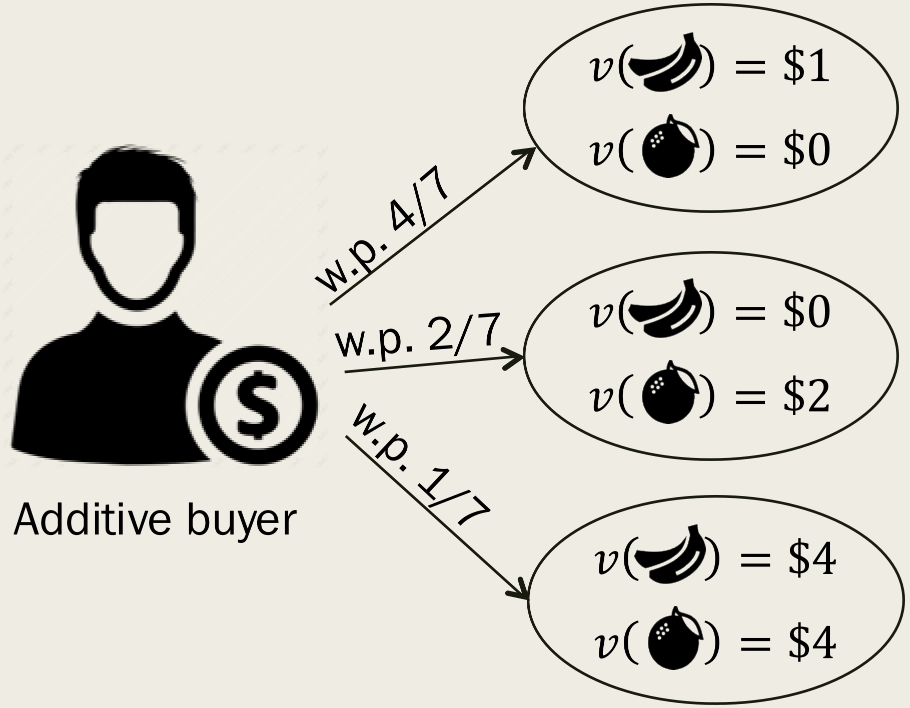

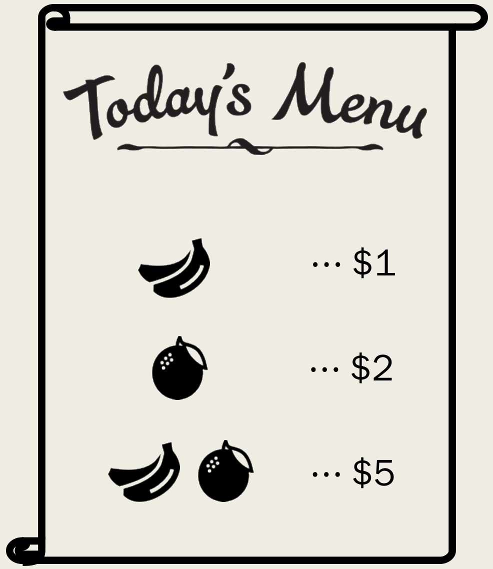

Optimal deterministic menu

Buy-many mechanisms, a.k.a. Sybil-proof mechanisms

"In a Sybil attack the attacker subverts … by creating a large number of pseudonymous identities, using them to gain a disproportionately large influence."

- In a Sybil 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.

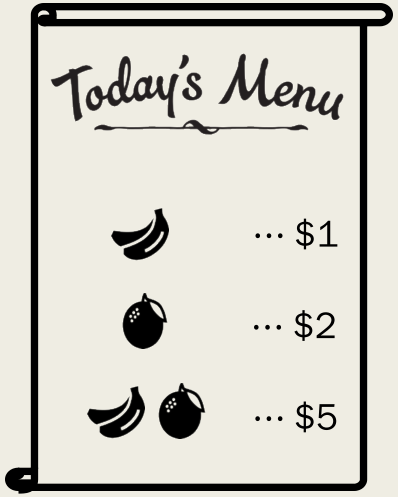

Not Sybil-proof

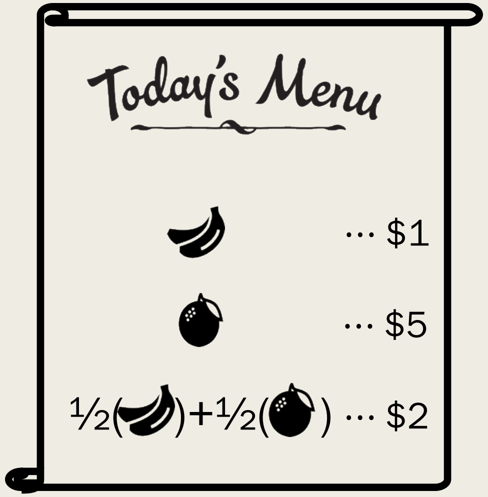

Not Sybil-proof

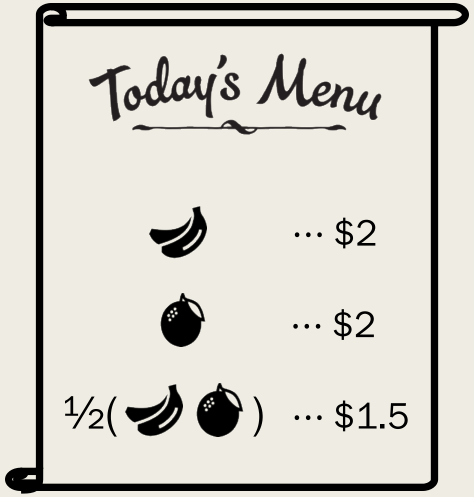

Sybil-proof

Buy-many mechanisms, a.k.a. Sybil-proof mechanisms

“In a Sybil attack the attacker subverts … by creating a large number of pseudonymous identities, using them to gain a disproportionately large influence.”

- In a Sybil 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 Sybil-proof if the random allocation resulting from any Sybil strategy is “dominated” by a single menu option.

Cheaper price; Bigger allocation

- For deterministic pricings, Sybil-proofness ≡ subadditivity

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,

Sybil-proof OPT $\leq O(\log n)$ Revenue of Item Pricing

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

Subadditive Deterministic OPT $\geq \Omega(\log n)$ Revenue of any "succinct" mechanism

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

Previous work showed…

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

[Briest Chawla Kleinberg Weinberg'10]: For any distribution D over unit-demand valuations,
Sybil-proof OPT $\leq O(\log n)$ Item Pricing Rev.

Optimal buy-many mechanisms can be well approximated

[C. Teng Tzamos'19]

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

Sybil-proof OPT $\leq O(\log n)$ Revenue of Item Pricing

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

Subadditive Deterministic OPT $\geq \Omega(\log n)$ Revenue of any "succinct" mechanism

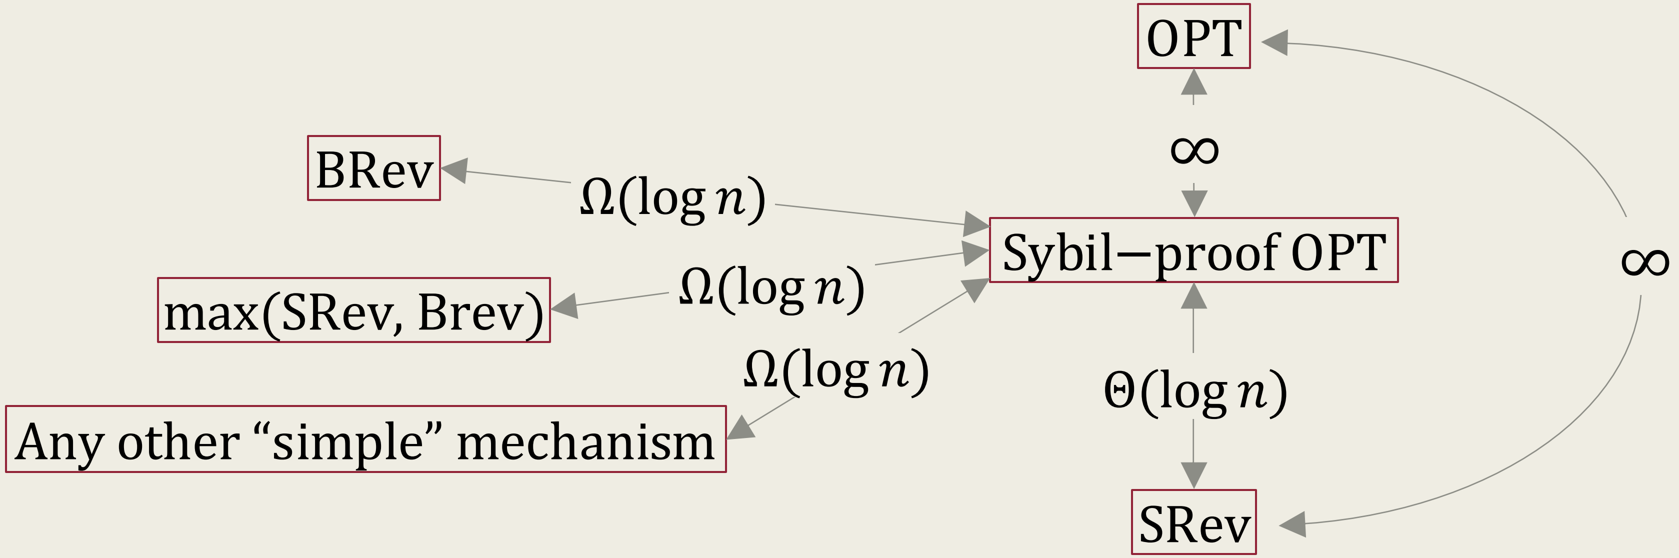

Other desirable properties…

"Small" menu sizes?

- Can get a finite bound over $(1-\epsilon)$-approximate menus

Revenue monotonicity for additive valuations?

- Likely doesn't hold

Revenue Lipschitzness?

[Psomas et al.'19] show that Lipschitzness doesn't hold for general mechanisms

- Holds for Buy-Many mechanisms!

What makes buy-many menus well-behaved?

- 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.

Key lemma: Additive pricings point-wise n-approximate buy-many menus

Additive pricings point-wise n-approximate subadditive pricings

- Sybil-proofness in deterministic pricings f requires:
  $f(A \cup B) \leq f(A) + f(B)$ for all subsets $A, B \subseteq [n]$

- Define additive pricing g as follows:

  $g_i = f_i$ for all $i \in [n]$ and $g(S) = \sum_{i \in S} g_i$ for all $S \subseteq [n]$

  Then:

$$\frac{1}{n} g(S) = \frac{1}{n} \sum_{i \in S} f_i \leq \max_{i \in S} f_i \leq f(S) \leq \sum_{i \in S} f_i = g(S)$$

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

- Given buy-many menu $\{(\lambda, f(\lambda))\}$

lottery a.k.a. randomized allocation

- Define additive pricing g as follows:

cheapest way to acquire item i under pricing f

$g_i = \min_{\lambda} \frac{f(\lambda)}{\Pr[i \in \lambda]}$ for all $i \in [n]$ and $g(S) = \sum_{i \in S} g_i$ for all $S \subseteq [n]$

- Extend g to lotteries: $g(\lambda) = \sum_S \Pr[S \sim \lambda] \{\sum_{i \in S} g_i\} = \sum_i g_i \Pr[i \in \lambda]$.

- Since f is buy-many, $g(\lambda) \geq f(\lambda)$.

- On the other hand, $f(\lambda) \geq g_i \Pr[i \in \lambda]\ \forall i$. Therefore, $f(\lambda) \geq \frac{1}{n} \sum_i g_i \Pr[i \in \lambda] = \frac{1}{n} g(\lambda)$.

$$\Rightarrow \frac{1}{n} g(\lambda) \leq f(\lambda) \leq g(\lambda)$$

A proof of the $O(\log n)$ approximation

Theorem 1: For any distribution D over valuations.

and any buy-many pricing function $f\colon \text{(random) allocations} \to \mathbb{R}^+ \cup \{0\}$,

there exists an additive pricing function g with

$$\text{Rev}_D(g) \geq \frac{1}{2\log(2n)} \text{Rev}_D(f)$$

Key technical claim: Point-wise approximation implies revenue approximation.

Theorem 3: Given any pricing functions f and g such that for all random allocations Λ,

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

Then for any value distribution D, there exists a scaling factor $\alpha > 0$, such that

$$\text{Rev}_D(\alpha g) \geq \frac{1}{2\log 2c} \text{Rev}_D(f).$$

Theorem 1: For any distribution D over valuations.

and any determ. subadditive pricing function f: (random) allocations $\rightarrow \mathbb{R}^+ \cup \{0\}$,

there exists an additive pricing function g with

$$\mathrm{Rev}_D(g) \geq \frac{1}{2\log(2n)} \mathrm{Rev}_D(f)$$

Key technical claim: Point-wise approximation implies revenue approximation.

Theorem 3: Given any det. pricing functions f and g such that for all subsets $S \subseteq [n]$,

$$\frac{1}{c} g(S) \leq f(S) \leq g(S).$$

Then for any value distribution D, there exists a scaling factor $\alpha > 0$, such that

$$\mathrm{Rev}_D(\alpha g) \geq \frac{1}{2\log 2c} \mathrm{Rev}_D(f).$$

Theorem 3: Given any det. pricing functions f and g such that for all subsets $S \subseteq [n]$,

$$\frac{1}{c} g(S) \leq f(S) \leq g(S).$$

Then for any value distribution D, there exists a scaling factor $\alpha > 0$, such that

$$\text{Rev}_D(\alpha g) \geq \frac{1}{2 \log 2c} \text{Rev}_D(f).$$

Restatement: Given any det. pricing functions f and g such that for all subsets $S \subseteq [n]$,

$$\frac{1}{c} g(S) \leq f(S) \leq g(S).$$

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

$$\text{Rev}_v(\alpha g) \geq \frac{1}{2 \log 2c} \text{Rev}_v(f).$$

Theorem 3: Given any det. pricing functions f and g such that for all subsets $S \subseteq [n]$,

$$\frac{1}{c} g(S) \leq f(S) \leq g(S).$$

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

$$\text{Rev}_v(\alpha g) \geq \frac{1}{2 \log 2c} \text{Rev}_v(f).$$

A scaling argument:

- Suppose a buyer purchases S^* under f. We want to recover $f(S^*)$.
- Consider varying α between ${}^1/_{2c}$ and 1. $\Longrightarrow O(\log c)$ scales of interest.
- At one of these scales, we have $\frac{1}{2} f(S^*) \leq \alpha g(S^*) \leq f(S^*)$.

  $\Longrightarrow$ Buyer can afford S^* and seller makes revenue at least $\frac{1}{2} f(S^*)$.

Problem: Buyer may buy something other than S^* at a much lower price than $f(S^*)$.

Observation: the buyer gets high utility under pricing $\alpha g \Longrightarrow$ can extract more revenue by raising prices.

As we raise prices, the buyer loses utility; with total utility loss comparable to $f(S^*)$.

- Our goal: recover this utility loss as revenue!

Recovering utility loss as revenue

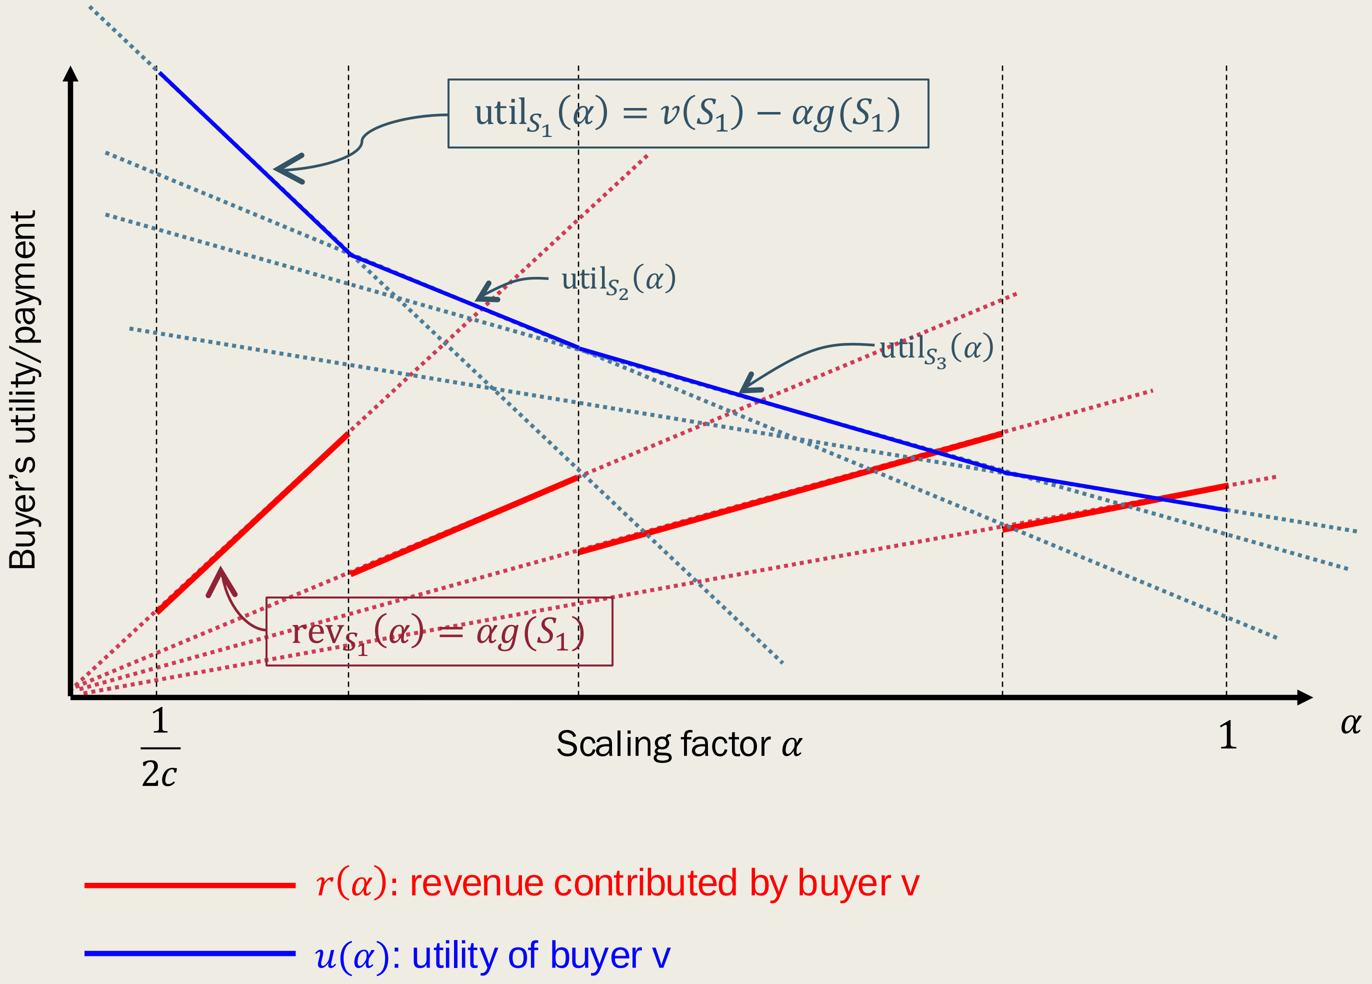

Define:

$$\text{util}_S(\alpha) = v(S) - \alpha g(S)$$
$$u(\alpha) = \max_S \{\text{util}_S(\alpha)\}$$

$$\text{rev}_S(\alpha) = \alpha g(S)$$
$$r(\alpha) = \alpha g(S_\alpha)$$

Observe:

$$\frac{d}{d\alpha} u(\alpha) = -r'(\alpha) = -\frac{1}{\alpha} r(\alpha)$$

Then, picking α with density $\propto {}^1/_\alpha$ gives:

$$E_\alpha[r(\alpha)] = \frac{1}{\log(2c)} \int \frac{r(\alpha)}{\alpha}\, d\alpha$$
$$= \frac{1}{\log(2c)} \int -\frac{d}{d\alpha} u(\alpha)\, d\alpha$$
$$= \frac{u({}^1/_{2c}) - u(1)}{\log(2c)}$$

Theorem 3: Given any det. pricing functions f and g such that for all subsets $S \subseteq [n]$,

$$\frac{1}{c} g(S) \leq f(S) \leq g(S).$$

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

$$\text{Rev}_v(\alpha g) \geq \frac{1}{2 \log 2c} \text{Rev}_v(f).$$

$$\text{util}_S(\alpha) = v(S) - \alpha g(S)$$
$$u(\alpha) = \max_S \{\text{util}_S(\alpha)\}$$

Outline:

- Pick α with density $\propto {}^1/_\alpha$.
- Then, $E_\alpha[r(\alpha)] = \frac{u({}^1/_{2c}) - u(1)}{\log(2c)}$
- $u(1) = \max_S \{v(S) - g(S)\} \leq \max_S \{v(S) - f(S)\} = v(S^*) - f(S^*)$
- $u({}^1/_{2c}) = \max_S \left\{v(S) - \frac{1}{2c} g(S)\right\} \geq \max_S \left\{v(S) - \frac{1}{2} f(S)\right\} \geq v(S^*) - \frac{1}{2} f(S^*)$
- Putting everything together, $E_\alpha[r(\alpha)] \geq \frac{1}{2 \log(2c)} f(S^*)$

Recap of approximation results

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

Sybil-proof OPT $\leq O(\log n)$ SRev

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

Subadditive Deterministic OPT $\geq \Omega(\log n)$ Revenue of any "succinct" mechanism

Theorem 3: For any two pricing functions, a pointwise c-approximation upon rescaling

implies an $O(\log c)$-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!