Multi-Item Mechanisms: Complexity, Simplicity, Menus & Communication

Yannai A. Gonczarowski

Harvard

SLMath (MSRI) Summer School
June 22, 2023

Based upon (but all typos are my own):

Bounding the Menu-Size of Approximately Optimal Auctions via Optimal-Transport Duality, Y.A.G., 2018

The Menu-Size Complexity of Revenue Approximation, Moshe Babaioff, Y.A.G., Noam Nisan, 2022

Strong Duality for a Multiple-Good Monopolist, Constantinos Daskalakis, Alan Deckelbaum, Christos Tzamos, 2017

A Classic Question in Economics

- A single **seller** has *n* **items** that she would like to sell to a single **buyer**. The seller has no other use for the items.
  - E.g., a market for **digital goods**.

A Classic Question in Economics

- A single **seller** has *n* **items** that she would like to sell to a single **buyer**. The seller has no other use for the items.
  - E.g., a market for **digital goods**.
- The buyer has a private **value** (the maximum price she is willing to pay) for each item (need not be the same for all items).
- The buyer's value for any subset of the items is the **sum** of the values of the items in the subset.

A Classic Question in Economics

- A single **seller** has *n* **items** that she would like to sell to a single **buyer**. The seller has no other use for the items.
  - E.g., a market for **digital goods**.
- The buyer has a private **value** (the maximum price she is willing to pay) for each item (need not be the same for all items).
- The buyer's value for any subset of the items is the **sum** of the values of the items in the subset.
- For each item, the seller has **prior knowledge** of a distribution from which the buyer's value for this item is drawn, **independently** of any other value.
- The seller can choose any selling **mechanism / auction** (as long as the buyer can both **opt out** and **strategize**. . .).

A Classic Question in Economics

- A single **seller** has n **items** that she would like to sell to a single **buyer**. The seller has no other use for the items.
  - E.g., a market for **digital goods**.
- The buyer has a private **value** (the maximum price she is willing to pay) for each item (need not be the same for all items).
- The buyer's value for any subset of the items is the **sum** of the values of the items in the subset.
- For each item, the seller has **prior knowledge** of a distribution from which the buyer's value for this item is drawn, **independently** of any other value.
- The seller can choose any selling **mechanism / auction** (as long as the buyer can both **opt out** and **strategize**. . .).
  - (The buyer would like to maximize her expected **utility** = value for bought items − payment.)

A Classic Question in Economics

- A single **seller** has n **items** that she would like to sell to a single **buyer**. The seller has no other use for the items.
  - E.g., a market for **digital goods**.
- The buyer has a private **value** (the maximum price she is willing to pay) for each item (need not be the same for all items).
- The buyer's value for any subset of the items is the **sum** of the values of the items in the subset.
- For each item, the seller has **prior knowledge** of a distribution from which the buyer's value for this item is drawn, **independently** of any other value.
- The seller can choose any selling **mechanism / auction** (as long as the buyer can both **opt out** and **strategize**. . .).
  - (The buyer would like to maximize her expected **utility** = value for bought items − payment.)

A fundamental question in mechanism design:
How can the seller maximize her (expected) revenue given the prior distribution over the buyer's values?

Earlier this Week, With Inbal: One Item

- A possible mechanism: choose a price, and offer the item for that price.
- The price that maximizes the revenue among all possible prices (the **monopoly price**) is

$$\arg\mathrm{Max}_p\, p \cdot \mathbb{P}_{v\sim F}\left[v \geq p\right].$$

- Other mechanisms are also possible (multiround, lottery tickets, etc.)

Earlier this Week, With Inbal: One Item

- A possible mechanism: choose a price, and offer the item for that price.
- The price that maximizes the revenue among all possible prices (the **monopoly price**) is

$$\arg \mathrm{Max}_p \, p \cdot \mathbb{P}_{v \sim F}\left[v \geq p\right].$$

- Other mechanisms are also possible (multiround, lottery tickets, etc.)

Theorem (Myerson, 1981; Riley and Zeckhauser, 1983)

In any single-item setting, no other mechanism can obtain higher revenue than posting the revenue-maximizing price.

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately!

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately?

Example

If both item values are uniformly distributed in $\{\$1, \$2\}$:

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately?

Example

If both item values are uniformly distributed in $\{\$1, \$2\}$:

- Pricing each item separately, seller obtains a revenue of \$1 for each item, for a total revenue of \$2.

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately?

Example

If both item values are uniformly distributed in $\{\$1, \$2\}$:

- Pricing each item separately, seller obtains a revenue of \$1 for each item, for a total revenue of \$2.
- Pricing only the bundle at \$3, seller obtains a revenue of $\$3 \cdot 0.75$

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately?

Example

If both item values are uniformly distributed in $\{\$1, \$2\}$:

- Pricing each item separately, seller obtains a revenue of \$1 for each item, for a total revenue of \$2.
- Pricing only the bundle at \$3, seller obtains a revenue of $\$3 \cdot 0.75 = 2.25 \;>\; 2!$

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately?

Example

If both item values are uniformly distributed in $\{\$1, \$2\}$:

- Pricing each item separately, seller obtains a revenue of \$1 for each item, for a total revenue of \$2.
- Pricing only the bundle at \$3, seller obtains a revenue of $\$3 \cdot 0.75 = 2.25 \ > \ 2!$
- So pricing each item separately does not always maximize revenue!

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items!

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle?

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at $3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately ($1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for $2 / both for $3

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗
- Post a price for each item and a price for the bundle?

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at $3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately ($1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for $2 / both for $3

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗
- Post a price for each item and a price for the bundle? ✗

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for \$2 / both for \$3
$\text{Unif}\{1,2\} \times \text{Unif}\{1,3\}$	Offers include lottery tickets (both for \$4 / for \$2.5: first w.p. 1, second w.p. $1/2$) T'04,DDT'14

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗
- Post a price for each item and a price for the bundle? ✗
- Choose between a few lotteries?

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for \$2 / both for \$3
$\text{Unif}\{1,2\} \times \text{Unif}\{1,3\}$	Offers include lottery tickets (both for \$4 / for \$2.5: first w.p. 1, second w.p. 1/2) T'04,DDT'14

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗
- Post a price for each item and a price for the bundle? ✗
- Choose between a few lotteries? ✗

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for \$2 / both for \$3
$\text{Unif}\{1,2\} \times \text{Unif}\{1,3\}$	Offers include lottery tickets (both for \$4 / for \$2.5: first w.p. 1, second w.p. $1/2$) T'04,DDT'14
$\text{Beta}(1,2) \times \text{Beta}(1,2)$	Offer infinitely many lotteries DDT'13,DDT'15

More than One Item: Complex!

How can the seller maximize the revenue from two items?

- Distributions independent, so optimally price each item separately? ✗
- Optimally price the bundle of both items? ✗
- Either price separately or bundle? ✗
- Post a price for each item and a price for the bundle? ✗
- Choose between a few lotteries? ✗

Generally: analytic solution not known, structure not understood.

Distribution	Unique Optimal Mechanism
$\text{Unif}\{1,2\} \times \text{Unif}\{1,2\}$	Price the bundle (at \$3)
$\text{Unif}\{0,1\} \times \text{Unif}\{0,1\}$	Price each separately (\$1 each)
$\text{Unif}\{0,1,2\} \times \text{Unif}\{0,1,2\}$	Offer: one for \$2 / both for \$3
$\text{Unif}\{1,2\} \times \text{Unif}\{1,3\}$	Offers include lottery tickets (both for \$4 / for \$2.5: first w.p. 1, second w.p. 1/2) T'04,DDT'14
$\text{Beta}(1,2) \times \text{Beta}(1,2)$	Offer infinitely many lotteries DDT'13,DDT'15

Not Merely Unaesthetic / Hard to Formally Analyze

- Cannot be computed in expected polynomial-time even for seemingly simple distributions (unless $\mathtt{ZPP} \supseteq \mathtt{P}^{\#\mathtt{P}}$). DDT'14
  - Even some simple questions about optimal mechanisms are $\#\mathtt{P}$-hard to answer, even for such simple distributions. DDT'14
- Harder to represent to the buyer.
- Harder for the buyer to find/verify optimal strategy.

Not Merely Unaesthetic / Hard to Formally Analyze

- Cannot be computed in expected polynomial-time even for seemingly simple distributions (unless $\mathtt{ZPP} \supseteq \mathtt{P}^{\#\mathtt{P}}$). DDT'14
    - Even some simple questions about optimal mechanisms are $\#\mathtt{P}$-hard to answer, even for such simple distributions. DDT'14
- Harder to represent to the buyer.
- Harder for the buyer to find/verify optimal strategy.

So what revenue can we get using simpler mechanisms?

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only pricing separately. HN'12, HR'19
- Allow only "packaging". BILW'14, R'16
- Disallow lotteries. BNR'18

An "all or nothing" approach...

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only pricing separately. HN'12, HR'19
- Allow only "packaging". BILW'14, R'16
- Disallow lotteries. BNR'18

Today with Konstantin

An "all or nothing" approach. . .

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only pricing separately. ← Today with Konstantin — HN'12, HR'19
- Allow only "packaging". — BILW'14, R'16
- Disallow lotteries. — BNR'18

An "all or nothing" approach...

Such studied features lose at least a constant fraction of the optimal revenue.

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only pricing separately. — HN'12, HR'19
- Allow only "packaging". — BILW'14, R'16
- Disallow lotteries. — BNR'18

Today with Konstantin

An "all or nothing" approach...

Such studied features lose at least a constant fraction of the optimal revenue.

Option 2: Quantitatively: limit a numeric complexity measure:

- Number of options presented to the buyer. — HN'13
- The communication requirements of the mechanism.
- Learning-theoretic dimensionality. — MR'15, MR'16, BSV'16, S'17, BSV'18

A " " approach...

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some “features”:

- Allow only pricing separately. HN’12, HR’19
- Allow only “packaging”. BILW’14, R’16
- Disallow lotteries. BNR’18

Today with Konstantin

An “all or nothing” approach. . .

Such studied features lose at least a constant fraction of the optimal revenue.

Option 2: Quantitatively: limit a numeric complexity measure:

- Number of options presented to the buyer. HN’13
- The communication requirements of the mechanism.
- Learning-theoretic dimensionality. MR’15, MR’16, BSV’16, S’17, BSV’18

A “ ” approach. . .

Suitable for a systematic study of the trade-offs between simplicity and quality.

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only pricing separately. HN'12, HR'19
- Allow only "packaging". BILW'14, R'16
- Disallow lotteries. BNR'18

Today with Konstantin

An "all or nothing" approach...

Such studied features lose at least a constant fraction of the optimal revenue.

Option 2: Quantitatively: limit a numeric complexity measure:

- Number of options presented to the buyer. HN'13
- The communication requirements of the mechanism.
- Learning-theoretic dimensionality. MR'15, MR'16, BSV'16, S'17, BSV'18

A "   " approach...

Suitable for a systematic study of the trade-offs between simplicity and quality.

This lecture.

Simple Mechanisms: Limiting Complexity

Option 1: Qualitatively: disallow some “features”:

- Allow only pricing separately. HN’12, HR’19
- Allow only “packaging”. BILW’14, R’16
- Disallow lotteries. BNR’18

Today with Konstantin

An “all or nothing” approach. . .

Such studied features lose at least a constant fraction of the optimal revenue.

Option 2: Quantitatively: limit a numeric complexity measure:

- Number of options presented to the buyer. HN’13
- The communication requirements of the mechanism.
- Learning-theoretic dimensionality. MR’15, MR’16, BSV’16, S’17, BSV’18

Later this morning

A “ ” approach. . .

Suitable for a systematic study of the trade-offs between simplicity and quality.

This lecture.

The Menu Size of a Selling Mechanism

The Menu Size of a Selling Mechanism

Well known: every truthful selling mechanism, however complex, is equivalent to specifying a **menu** of possible probabilistic outcomes for the buyer to choose from.

The Menu Size of a Selling Mechanism

Well known: every truthful selling mechanism, however complex, is equivalent to specifying a **menu** of possible probabilistic outcomes for the buyer to choose from.

Chez Seller

Items • Bundles • Lotteries

Today's Specials

$\mathbb{P}$[Item 1]	$\mathbb{P}$[Item 2]	Price
0%	100%	\$3
20%	30%	\$4
40%	60%	\$10
⋮	⋮	⋮
100%	100%	\$20

The Classic Choice

0%	0%	\$0

— One entry per buyer —

The Menu Size of a Selling Mechanism

Well known: every truthful selling mechanism, however complex, is equivalent to specifying a **menu** of possible probabilistic outcomes for the buyer to choose from.

Chez Seller

Items • Bundles • Lotteries

Today's Specials

$\mathbb{P}$[Item 1]	$\mathbb{P}$[Item 2]	Price
0%	100%	$3
20%	30%	$4
40%	60%	$10
⋮	⋮	⋮
100%	100%	$20

The Classic Choice

0%	0%	$0

— One entry per buyer —

Menu Size
HN'13
see also BCKW'10, D'11, DV'11, DV'12

The Menu Size of a Selling Mechanism

Well known: every truthful selling mechanism, however complex, is equivalent to specifying a **menu** of possible probabilistic outcomes for the buyer to choose from.

Menu Size
HN'13
see also BCKW'10, D'11, DV'11, DV'12

- The base-2 logarithm of the menu size is precisely the **deterministic communication complexity** of running the mechanism. BGN'21

Up-to-ε Optimality with a Finite Menu Size?

Open Question (Hart and Nisan, 2014)

Is there a **finite** menu **size** $C(n,\varepsilon)$ that suffices for attaining a $(1-\varepsilon)$ fraction of the optimal revenue when selling n items drawn from any given distributions?

(The menu entries can depend on the distributions; the menu size cannot.)

$$\left(\inf_{F_1,\ldots,F_n\in\Delta(\mathbb{R}_+)} \frac{\mathcal{R}ev_C(F_1\times\cdots\times F_n)}{\mathrm{OPT}(F_1\times\cdots\times F_n)}\right) \xrightarrow[C\to\infty]{???} 1$$

Up-to-ε Optimality with a Finite Menu Size?

Open Question (Hart and Nisan, 2014)

Is there a **finite** menu **size** $C(n,\varepsilon)$ that suffices for attaining a $(1-\varepsilon)$ fraction of the optimal revenue when selling n items drawn from any given distributions?

(The menu entries can depend on the distributions; the menu size cannot.)

$$\left(\inf_{F_1,\ldots,F_n\in\Delta(\mathbb{R}_+)} \frac{\mathcal{R}ev_C(F_1\times\cdots\times F_n)}{\mathrm{OPT}(F_1\times\cdots\times F_n)}\right) \xrightarrow[C\to\infty]{???} 1$$

- Proved some special cases.

Up-to-ε Optimality with a Finite Menu Size?

Open Question (Hart and Nisan, 2014)

Is there a **finite** menu **size** $C(n, \varepsilon)$ that suffices for attaining a $(1-\varepsilon)$ fraction of the optimal revenue when selling n items drawn from any given distributions?

(The menu entries can depend on the distributions; the menu size cannot.)

$$\left(\inf_{F_1,\ldots,F_n \in \Delta(\mathbb{R}_+)} \frac{\mathcal{R}ev_C(F_1 \times \cdots \times F_n)}{\text{OPT}(F_1 \times \cdots \times F_n)} \right) \xrightarrow[C\to\infty]{???} 1$$

- Proved some special cases.
- Challenge: Hart and Nisan, 2013: For correlated distributions, no!

Up-to-ε Optimality with a Finite Menu Size?

Open Question (Hart and Nisan, 2014)

Is there a **finite** menu **size** $C(n, \varepsilon)$ that suffices for attaining a $(1-\varepsilon)$ fraction of the optimal revenue when selling n items drawn from any given distributions? ✓

(The menu entries can depend on the distributions; the menu size cannot.)

$$\left(\inf_{F_1,\ldots,F_n \in \Delta(\mathbb{R}_+)} \frac{\mathcal{R}ev_C(F_1 \times \cdots \times F_n)}{\mathrm{OPT}(F_1 \times \cdots \times F_n)} \right) \xrightarrow[C \to \infty]{\checkmark} 1$$

- Proved some special cases.
- Challenge: Hart and Nisan, 2013: For correlated distributions, no!

Quantifying the Menu Size for Up-to-ε Optimality

Theorem (Babaioff, **G.**, Nisan, 2022)

For every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every n valuation distributions, some mechanism with menu size at most C obtains at least a $(1-\varepsilon)$ fraction of the optimal revenue.

Quantifying the Menu Size for Up-to-ε Optimality

Theorem (Babaioff, **G.**, Nisan, 2022)

For every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every n valuation distributions, some mechanism with menu size at most C obtains at least a $(1-\varepsilon)$ fraction of the optimal revenue.

But what is the rate of (uniform) convergence?

How fast must C grow as a function ε?

Quantifying the Menu Size for Up-to-ε Optimality

Theorem (Babaioff, **G.**, Nisan, 2022)

For every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every n valuation distributions, some mechanism with menu size at most C obtains at least a $(1-\varepsilon)$ fraction of the optimal revenue.

But what is the rate of (uniform) convergence?

How fast must C grow as a function ε?

I.e., how good can low-complexity mechanisms be?

How complex must high-revenue mechanisms be?

Quantifying the Menu Size for Up-to-ε Optimality

Theorem (Babaioff, G., Nisan, 2022)

For every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every n valuation distributions, some mechanism with menu size at most C obtains at least a $(1-\varepsilon)$ fraction of the optimal revenue.

But what is the rate of (uniform) convergence?

How fast must C grow as a function ε?

I.e., how good can low-complexity mechanisms be?

How complex must high-revenue mechanisms be?

Theorem (BGN, 2022)

For any fixed number of items n, a menu size polyomial in ${}^1/_\varepsilon$ is **sufficient**.

Quantifying the Menu Size for Up-to-ε Optimality

Theorem (Babaioff, G., Nisan, 2022)

For every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every n valuation distributions, some mechanism with menu size at most C obtains at least a $(1-\varepsilon)$ fraction of the optimal revenue.

But what is the rate of (uniform) convergence?

How fast must C grow as a function ε?

I.e., how good can low-complexity mechanisms be?

How complex must high-revenue mechanisms be?

Theorem (BGN, 2022)

For any fixed number of items n, a menu size polyomial in $1/\varepsilon$ is **sufficient**.

Theorem (G, 2018)

For any fixed number of items n, a menu size polyomial in $1/\varepsilon$ is **necessary**.

Communication Complexity of Up-to-ε Optimality

- Recall that the logarithm of the menu size is precisely the deterministic **communication complexity** of running the mechanism. BGN'22

Communication Complexity of Up-to-ε Optimality

- Recall that the logarithm of the menu size is precisely the deterministic **communication complexity** of running the mechanism. BGN'22

- While there still is a gap between our polynomial lower & upper bounds, they together **tightly** resolve the communication complexity question:

Corollary (**G.**, 2018)

For any fixed number of items n, the necessary and sufficient deterministic communication complexity of a mechanism for up-to-ε revenue maximization from any distribution is of the order of $\log {}^{1}/_{\varepsilon}$.

Communication Complexity of Up-to-ε Optimality

- Recall that the logarithm of the menu size is precisely the deterministic **communication complexity** of running the mechanism. BGN'22
- While there still is a gap between our polynomial lower & upper bounds, they together **tightly** resolve the communication complexity question:

Corollary (**G.**, 2018)

For any fixed number of items n, the necessary and sufficient deterministic communication complexity of a mechanism for up-to-ε revenue maximization from any distribution is of the order of $\log {}^{1}/_{\varepsilon}$.

- Main takeaway: **dichotomy** between one item (complexity 1) and any other fixed number of items (complexity $\Theta(\log {}^{1}/_{\varepsilon})$).
  - No further qualitative jump for larger n.

Communication Complexity of Up-to-ε Optimality

- Recall that the logarithm of the menu size is precisely the deterministic **communication complexity** of running the mechanism. BGN'22

- While there still is a gap between our polynomial lower & upper bounds, they together **tightly** resolve the communication complexity question:

Corollary (G., 2018)

For any fixed number of items n, the necessary and sufficient deterministic communication complexity of a mechanism for up-to-ε revenue maximization from any distribution is of the order of $\log {}^1/_\varepsilon$.

- Main takeaway: **dichotomy** between one item (complexity 1) and any other fixed number of items (complexity $\Theta(\log {}^1/_\varepsilon)$).
  - No further qualitative jump for larger n.
  - Communication complexity characterization despite mechanisms still not understood.

Lower Bound via Duality

- Lower bound proof already for two i.i.d. items, bounded, additive loss:

Theorem (**G.**, 2018)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and a distribution $F \in \Delta([0,1])$, such that for every $\varepsilon > 0$ it is the case that $\mathcal{R}ev_M(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$ for every mechanism M with menu-size at most $C(\varepsilon)$.

Lower Bound via Duality

- Lower bound proof already for two i.i.d. items, bounded, additive loss:

Theorem (**G.**, 2018)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and a distribution $F \in \Delta([0,1])$, such that for every $\varepsilon > 0$ it is the case that $\mathcal{R}ev_M(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$ for every mechanism M with menu-size at most $C(\varepsilon)$.

- Let's prove this!

Lower Bound via Duality

- Lower bound proof already for two i.i.d. items, bounded, additive loss:

Theorem (**G.**, 2018)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and a distribution $F \in \Delta([0,1])$, such that for every $\varepsilon > 0$ it is the case that $\mathcal{R}ev_M(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$ for every mechanism M with menu-size at most $C(\varepsilon)$.

- Let's prove this!
- Recall: Daskalakis, Deckelbaum, Tzamos (2013, 2015) prove that infinite menu-size required for precise revenue maximization with two items sampled i.i.d. from the Beta distribution $F = \text{Beta}(1, 2)$.

Lower Bound via Duality

- Lower bound proof already for two i.i.d. items, bounded, additive loss:

Theorem (**G.**, 2018)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and a distribution $F \in \Delta([0,1])$, such that for every $\varepsilon > 0$ it is the case that $\mathcal{R}ev_M(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$ for every mechanism M with menu-size at most $C(\varepsilon)$.

- Let's prove this!
- Recall: Daskalakis, Deckelbaum, Tzamos (2013, 2015) prove that infinite menu-size required for precise revenue maximization with two items sampled i.i.d. from the Beta distribution $F = \text{Beta}(1,2)$.
- They do so by identifying a (strong!) dual problem (an **optimal-transport** problem), identifying the optimal dual and primal solutions for this F, and showing that the optimal primal solution has infinite menu size.

Lower Bound via Duality

- Lower bound proof already for two i.i.d. items, bounded, additive loss:

Theorem (**G.**, 2018)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and a distribution $F \in \Delta([0,1])$, such that for every $\varepsilon > 0$ it is the case that $\mathcal{R}ev_M(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$ for every mechanism M with menu-size at most $C(\varepsilon)$.

- Let's prove this!
- Recall: Daskalakis, Deckelbaum, Tzamos (2013, 2015) prove that infinite menu-size required for precise revenue maximization with two items sampled i.i.d. from the Beta distribution $F = \text{Beta}(1,2)$.
- They do so by identifying a (strong!) dual problem (an **optimal-transport** problem), identifying the optimal dual and primal solutions for this F, and showing that the optimal primal solution has infinite menu size.
- We will start by reviewing their optimal-transport duality framework, and then see how to leverage it to reason about approximately optimal mechanisms.

A Mechanism as a Utility Function

- A single-item illustration:

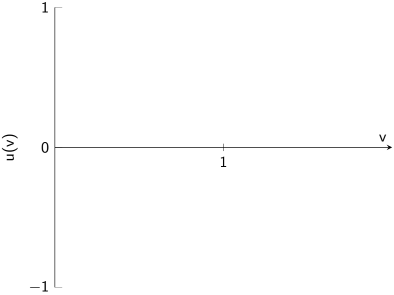

A Mechanism as a Utility Function

- A single-item illustration:

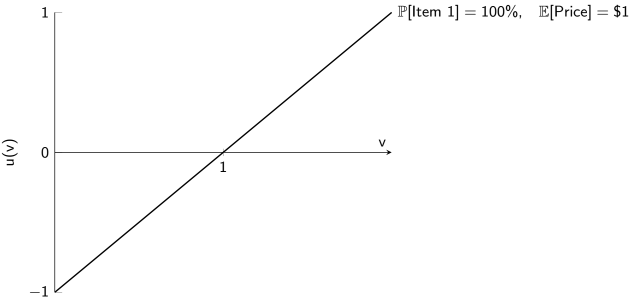

A Mechanism as a Utility Function

- A single-item illustration:

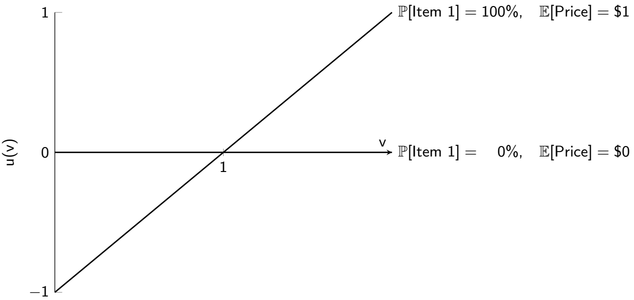

A Mechanism as a Utility Function

- A single-item illustration:

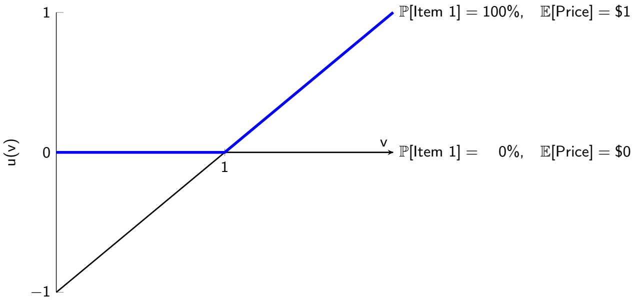

A Mechanism as a Utility Function

- A single-item illustration:

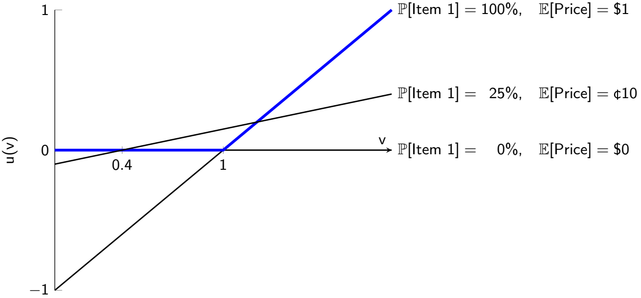

A Mechanism as a Utility Function

- A single-item illustration:

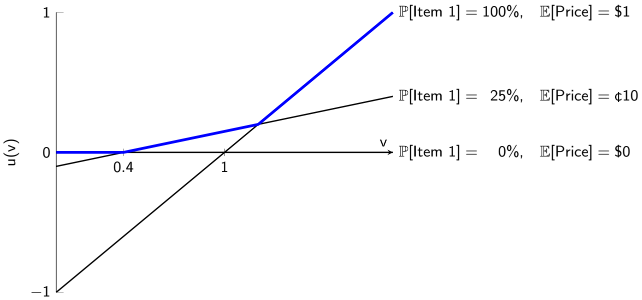

A Mechanism as a Utility Function

- A single-item illustration:

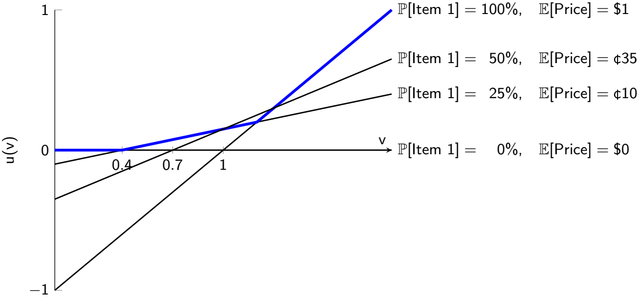

A Mechanism as a Utility Function

- A single-item illustration:

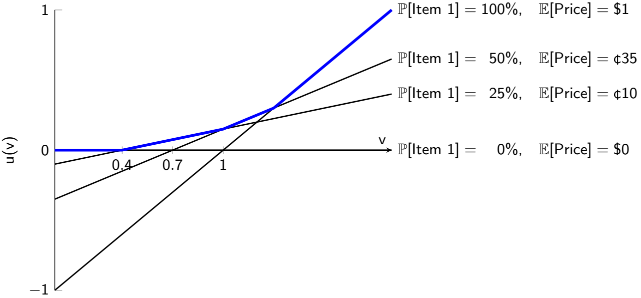

A Mechanism as a Utility Function

- A single-item illustration:

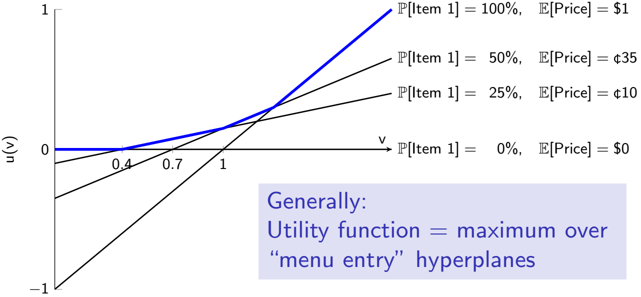

Generally:
Utility function = maximum over "menu entry" hyperplanes

A Mechanism as a Utility Function

- A single-item illustration:

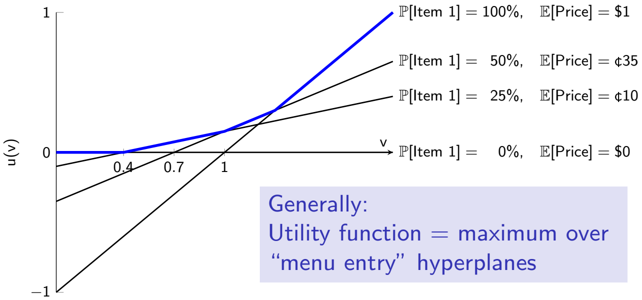

Generally:
Utility function = maximum over "menu entry" hyperplanes

Theorem (Rochet, 1987)

$u(\cdot)$ is the utility function of some mechanism iff it is nonnegative, nondecreasing, convex, 1-Lipschitz (ℓ_1 norm).

A Mechanism as a Utility Function

- A single-item illustration:

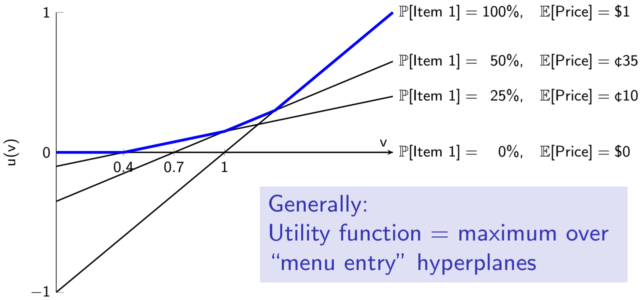

Generally:
Utility function = maximum over "menu entry" hyperplanes

Theorem (Rochet, 1987)

$u(\cdot)$ is the utility function of some mechanism iff it is nonnegative, nondecreasing, convex, 1-Lipschitz (ℓ_1 norm). For such $u(\cdot)$:

- *At every valuation v, the allocation probabilities form a subgradient.*

A Mechanism as a Utility Function

- A single-item illustration:

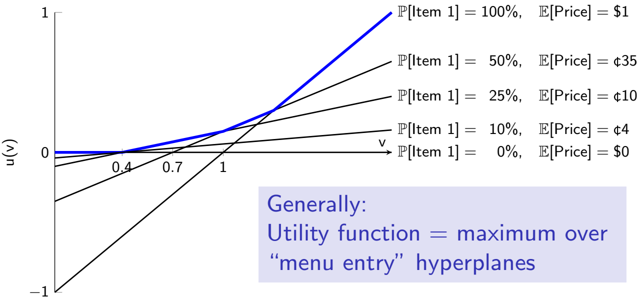

Generally:
Utility function = maximum over "menu entry" hyperplanes

Theorem (Rochet, 1987)

$u(\cdot)$ is the utility function of some mechanism iff it is nonnegative, nondecreasing, convex, 1-Lipschitz (ℓ_1 norm). For such $u(\cdot)$:

- *At every valuation v, the allocation probabilities form a subgradient.*

A Mechanism as a Utility Function

- A single-item illustration:

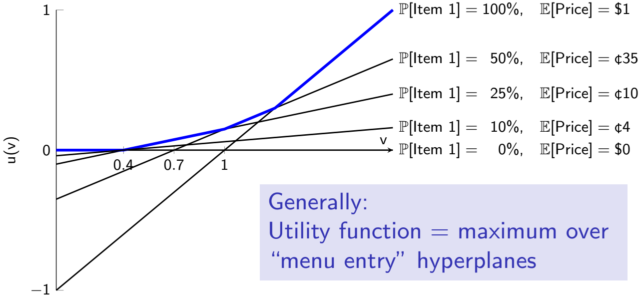

Generally:
Utility function = maximum over "menu entry" hyperplanes

Theorem (Rochet, 1987)

$u(\cdot)$ is the utility function of some mechanism iff it is nonnegative, nondecreasing, convex, 1-Lipschitz (ℓ_1 norm). For such $u(\cdot)$:

- *At every valuation v, the allocation probabilities form a subgradient.*
- *$\nabla u(v)$ exists almost everywhere, and for every v for which it exists, a buyer with valuation v pays $\nabla u(v) \cdot v - u(v)$.*

Massaging the Primal

$$\sup_{\substack{M: \\ \text{mechanism}}} \int \text{payment}_M(v) d\bar{F}(v)$$

Massaging the Primal

$$\sup_{\substack{M:\\ \text{mechanism}}} \int \text{payment}_M(v) d\bar{F}(v) =$$

Massaging the Primal

$$\sup_{\substack{M: \\ \text{mechanism}}} \int \text{payment}_M(v) d\bar{F}(v) =$$

...through the analysis of Rochet ('87) from the last slide...

Massaging the Primal

$$\sup_{\substack{M: \\ \text{mechanism}}} \int \text{payment}_M(v) d\bar{F}(v) =$$

. . . through the analysis of Rochet ('87) from the last slide. . .

$$= \sup_{\substack{u: \\ \text{nonnegative,} \\ \text{nondecreasing,} \\ \text{convex,} \\ \text{1-Lipschitz } (\ell_1)}} \int \Big(\nabla u(v) \cdot v - u(v)\Big) d\bar{F}(v)$$

Massaging the Primal

$$\sup_{\substack{M:\\ \text{mechanism}}} \int \text{payment}_M(v)\, d\bar{F}(v) =$$

...through the analysis of Rochet ('87) from the last slide...

$$= \sup_{\substack{u:\\ \text{nonnegative,}\\ \text{nondecreasing,}\\ \text{convex,}\\ \text{1-Lipschitz } (\ell_1)}} \int \Big(\nabla u(v) \cdot v - u(v)\Big) d\bar{F}(v) =$$

Massaging the Primal

$$\sup_{\substack{M:\\ \text{mechanism}}} \int \text{payment}_M(v)\, d\bar{F}(v) =$$

…through the analysis of Rochet ('87) from the last slide…

$$= \sup_{\substack{u:\\ \text{nonnegative,}\\ \text{nondecreasing,}\\ \text{convex,}\\ \text{1-Lipschitz } (\ell_1)}} \int \Big(\nabla u(v) \cdot v - u(v)\Big)\, d\bar{F}(v) =$$

…carefully applying (Daskalakis et al., '13,'15) the divergence theorem (think "high-dimensional integration by parts")…

Massaging the Primal

$$\sup_{\substack{M:\\ \text{mechanism}}} \int \text{payment}_M(v) d\bar{F}(v) =$$

...through the analysis of Rochet ('87) from the last slide...

$$= \sup_{\substack{u:\\ \text{nonnegative,}\\ \text{nondecreasing,}\\ \text{convex,}\\ \text{1-Lipschitz } (\ell_1)}} \int \Big(\nabla u(v) \cdot v - u(v)\Big) d\bar{F}(v) =$$

...carefully applying (Daskalakis et al., '13,'15) the divergence theorem (think "high-dimensional integration by parts")...

$$= \sup_{u:\cdots} \int u(v) d\mu(v)$$

where μ is a **signed** Radon measure of total mass 0 on the valuation space that depends only on $\bar{F}$ (and f, and ∇f)

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u: \\ u(0)\geq 0, \\ convex, \\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq$$

A Dual (Daskalakis et al., ’13,’15)

Theorem (Daskalakis et al., ’13)

$$\sup_{\substack{u: \\ u(0)\geq 0, \\ \textit{convex}, \\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \inf_{\substack{\gamma: \\ \textit{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-)$$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u: \\ u(0)\geq 0, \\ convex, \\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma: \\ coupling\ of\ \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ...

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u\,d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u\,d\mu = \int u\,d(\mu_+ - \mu_-) =$$

. . . by feasibility of γ. . . $$= \int \Big(u(v) - u(w)\Big) d\gamma(v,w)$$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

. . . by feasibility of γ. . . $$= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u\,d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u\,d\mu = \int u\,d(\mu_+ - \mu_-) =$$

. . . by feasibility of γ. . . $= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$

. . . by feasibility of u. . .

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u: \\ u(0)\geq 0, \\ \text{convex}, \\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma: \\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ... $\quad = \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$

... by feasibility of u... $\quad \leq \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ ... $= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$

... by feasibility of u ... $\leq \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$ □

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ convex,\\ u(v)-u(w)\leq|(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ coupling\ of\ \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

$$\ldots \text{by feasibility of } \gamma \ldots \quad = \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$$

$$\ldots \text{by feasibility of } u \ldots \quad \leq \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

□

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \mathrm{Beta}(1,2) \times \mathrm{Beta}(1,2)$.

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

Complementary slackness:
For equality, $\gamma(v, w)$-a.e.:

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ... $= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$

... by feasibility of u... $\leq \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$ □

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \text{Beta}(1,2) \times \text{Beta}(1,2)$.

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ convex,\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ coupling\ of\ \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

Complementary slackness:
For equality, $\gamma(v, w)$-a.e.:
$v_i < w_i \Rightarrow \nabla u_i = 0$ along segment

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ ... $$= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$$

... by feasibility of u ... $$\leq \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$ □

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \text{Beta}(1,2) \times \text{Beta}(1,2)$.

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u,γ:

Complementary slackness:
For equality, $\gamma(v,w)$-a.e.:
$v_i < w_i \Rightarrow \nabla u_i = 0$ along segment
$v_i > w_i \Rightarrow \nabla u_i = 1$ along segment

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

…by feasibility of γ… $$= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$$

…by feasibility of u… $$\leq \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$ □

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \text{Beta}(1,2) \times \text{Beta}(1,2)$.

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq |(v-w)_+|_1}} \int u d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \left|(v-w)_+\right|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

Complementary slackness:
For equality, $\gamma(v,w)$-a.e.:
$v_i < w_i \Rightarrow \nabla u_i = 0$ along segment
$v_i > w_i \Rightarrow \nabla u_i = 1$ along segment

$$\int u d\mu = \int u d(\mu_+ - \mu_-) =$$

... by feasibility of γ ... $= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$

... by feasibility of u ... $\leq \int \left|(v-w)_+\right|_1 d\gamma(v,w)$ □

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \text{Beta}(1,2) \times \text{Beta}(1,2)$.

A Dual (Daskalakis et al., '13,'15)

Theorem (Daskalakis et al., '13)

$$\sup_{\substack{u:\\ u(0)\geq 0,\\ \text{convex},\\ u(v)-u(w)\leq|(v-w)_+|_1}} \int u\,d\mu \quad \leq \quad \inf_{\substack{\gamma:\\ \text{coupling of } \mu_+,\mu_-}} \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$

Proof. For every feasible u, γ:

Complementary slackness:
For equality, $\gamma(v,w)$-a.e.:
$v_i < w_i \Rightarrow \nabla u_i = 0$ along segment
$v_i > w_i \Rightarrow \nabla u_i = 1$ along segment

$$\int u\,d\mu = \int u\,d(\mu_+ - \mu_-) =$$

. . . by feasibility of γ. . . $$= \int \Big(u(v) - u(w)\Big) d\gamma(v,w) \leq$$

. . . by feasibility of u. . . $$\leq \int \Big|(v-w)_+\Big|_1 d\gamma(v,w)$$ □

Daskalakis et al. then identified u,γ with equality for (μ of) $\bar{F} = \text{Beta}(1,2) \times \text{Beta}(1,2)$.

G.'18: lower-bound loss for u with small menu size and optimal γ

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. $\text{Beta}(1, 2)$.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.

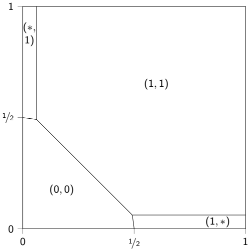

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.

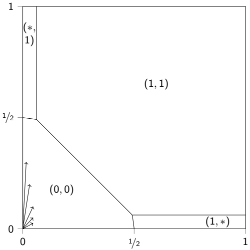

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.

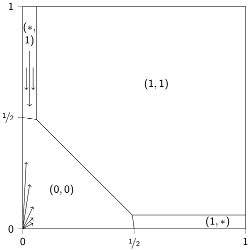

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.

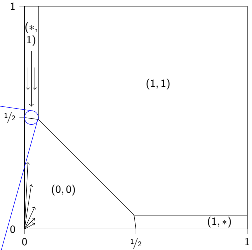

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.

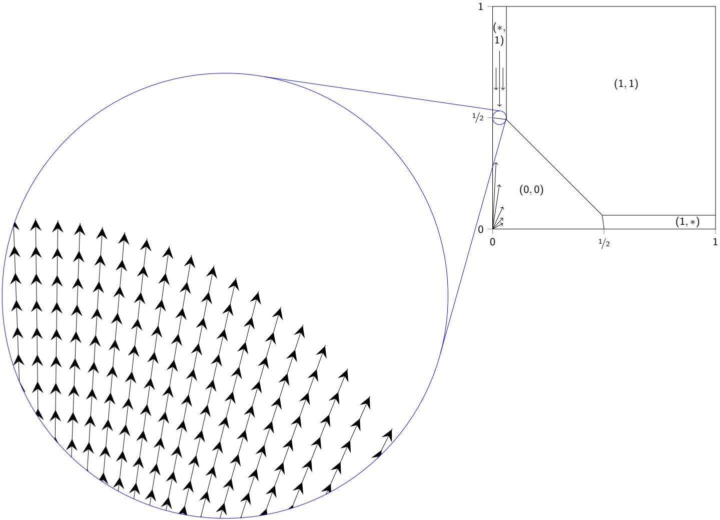

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta(1, 2).

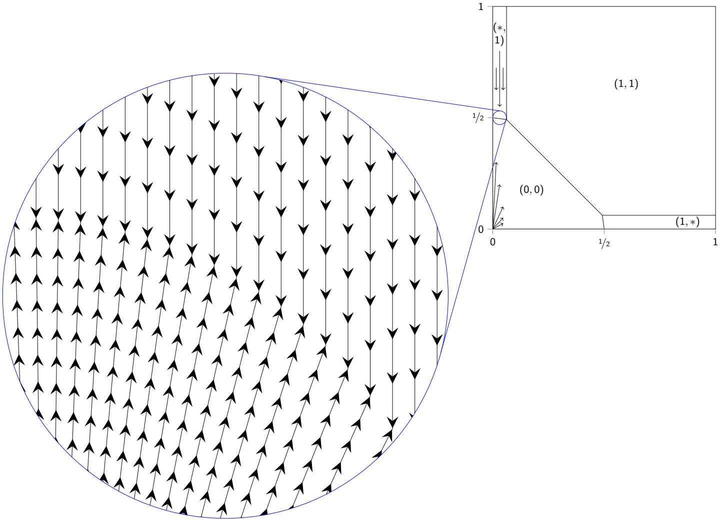

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

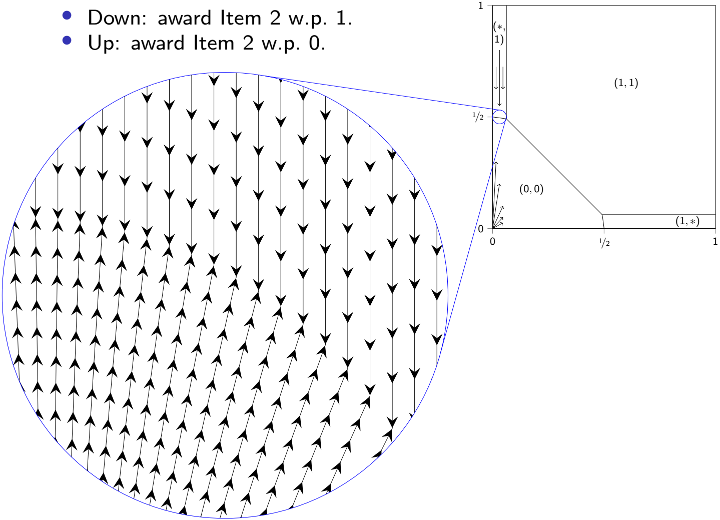

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

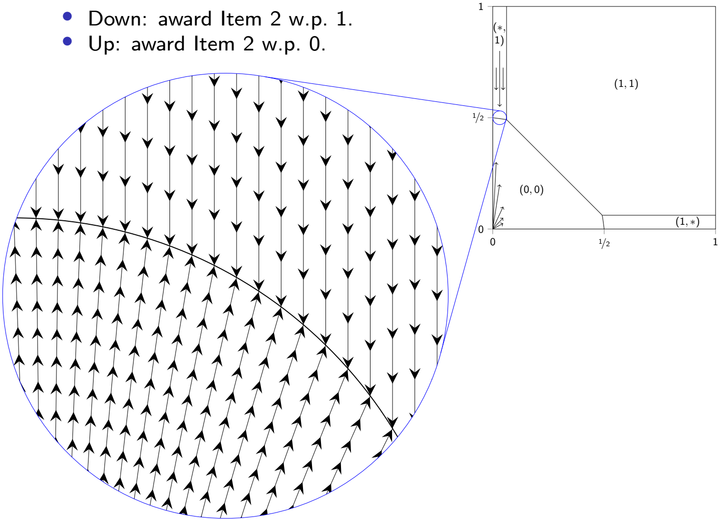

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

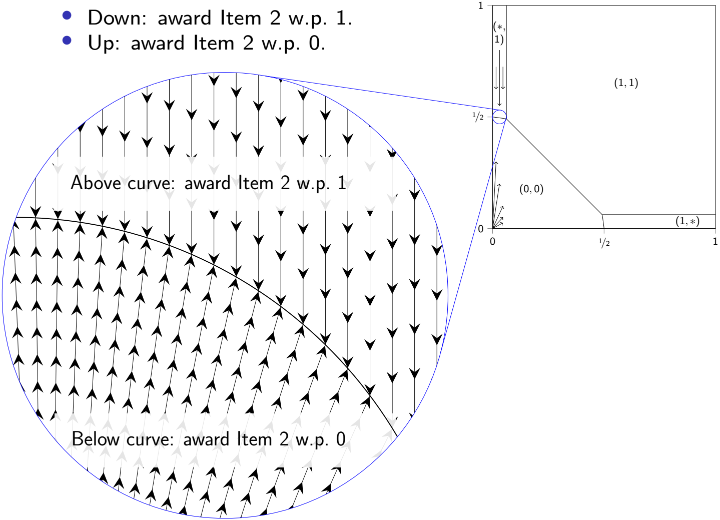

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

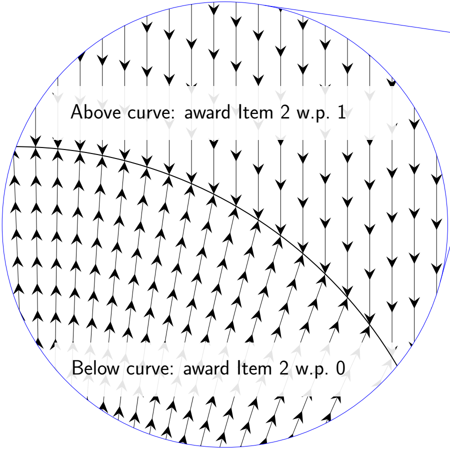

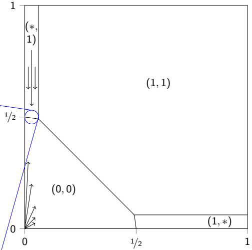

- Strictly concave ⇒ infinite menu size.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

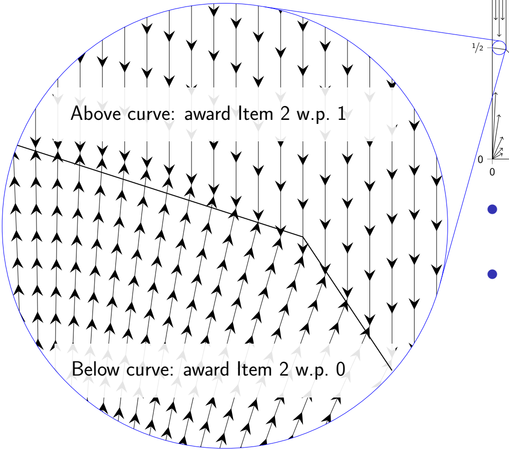

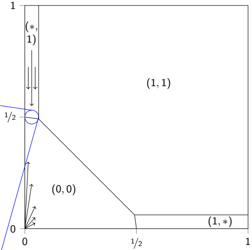

- Strictly concave ⇒ infinite menu size.
- Finite menu size ⇒ piecewise-linear.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

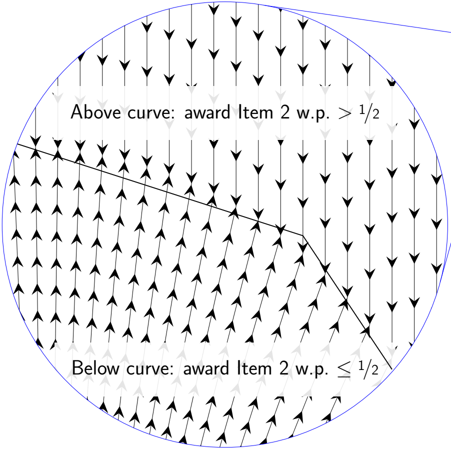

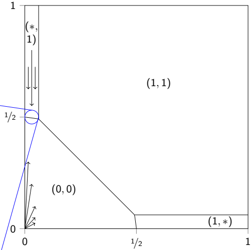

- Strictly concave ⇒ infinite menu size.
- Finite menu size ⇒ piecewise-linear.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

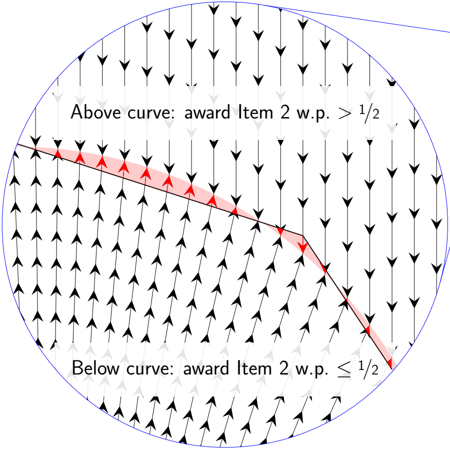

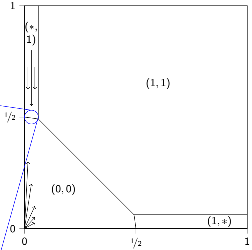

- Strictly concave ⇒ infinite menu size.
- Finite menu size ⇒ piecewise-linear.
  - Complementary slackness violated.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

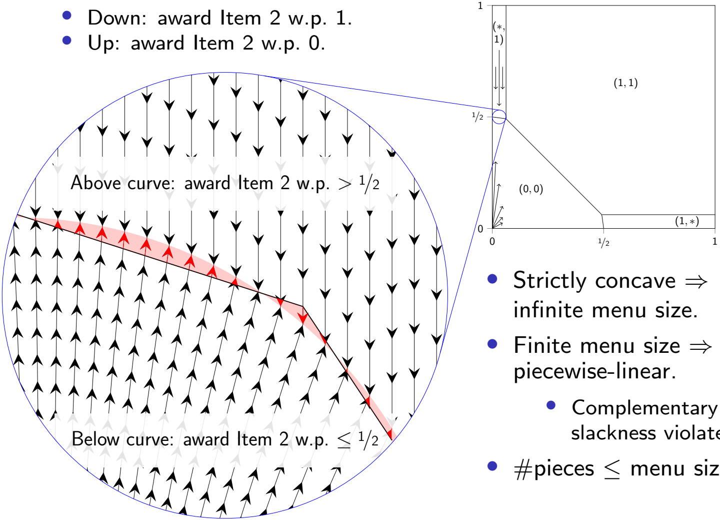

- Strictly concave ⇒ infinite menu size.
- Finite menu size ⇒ piecewise-linear.
  - Complementary slackness violated.
- #pieces ≤ menu size.

Wedging a Gap from the Optimal Dual

- DDT: optimal dual(&primal) for two items i.i.d. Beta$(1, 2)$.
- Complementary slackness:
  - Down: award Item 2 w.p. 1.
  - Up: award Item 2 w.p. 0.

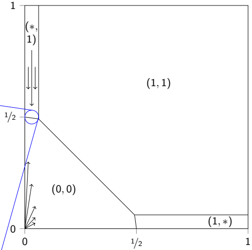

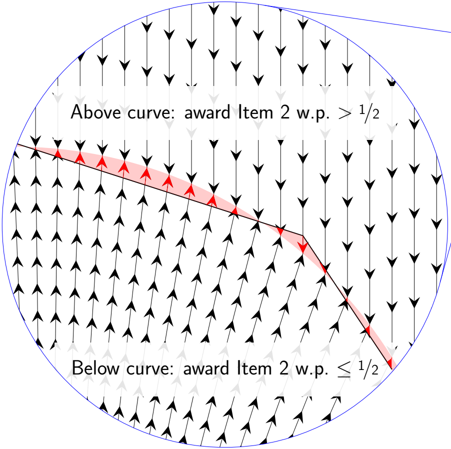

- Strictly concave $\Rightarrow$ infinite menu size.
- Finite menu size $\Rightarrow$ piecewise-linear.
  - Complementary slackness violated.
- #pieces $\leq$ menu size.
  - Quantify loss!

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

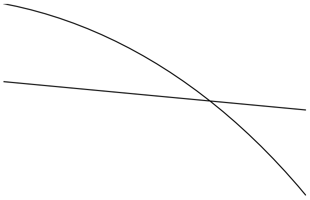

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

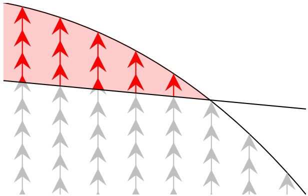

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

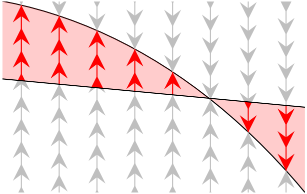

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

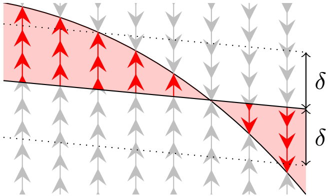

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

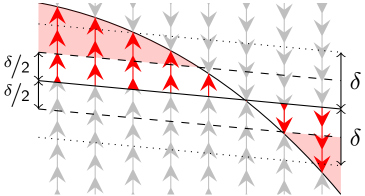

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

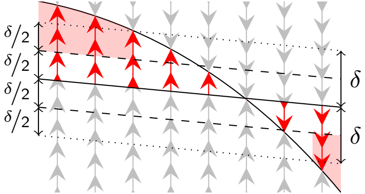

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

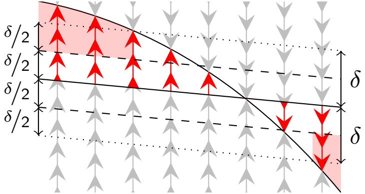

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

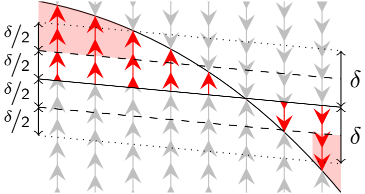

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

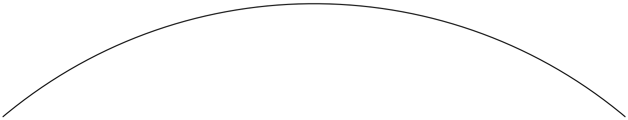

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

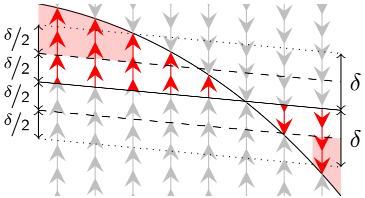

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

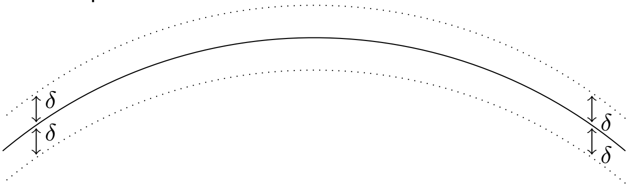

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

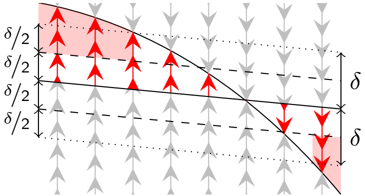

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

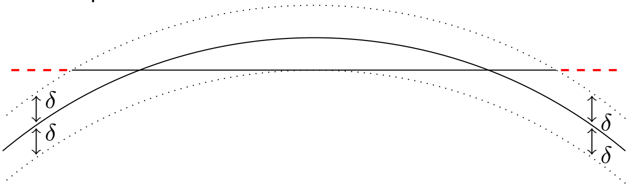

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

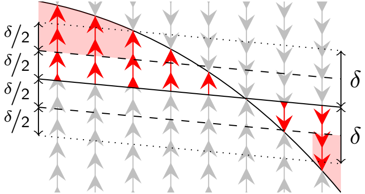

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

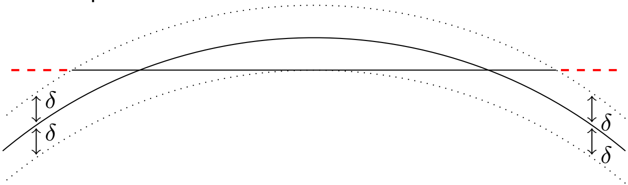

- Maximal "close" measure in one linear piece: circle chord of sagitta 2δ.

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

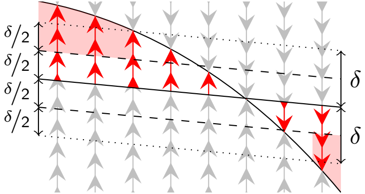

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

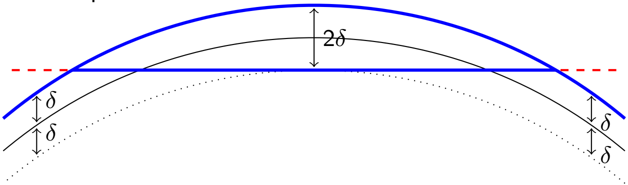

- Maximal "close" measure in one linear piece: circle chord of sagitta 2δ.

Quantifying the Gap from the Optimal Dual

- Quantifiable $\Omega(\delta^2)$ loss from each x-axis coordinate at which the piecewise-linear curve and the optimal curve are off by $\geq \delta$.

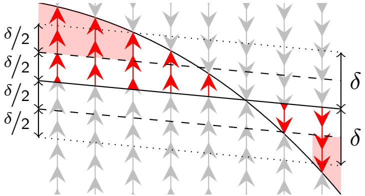

- Loss weighting "uniform enough" s.t. it suffices to show a constant measure of x-axis coordinates with distance $\geq \delta$.
- For circular opt.:

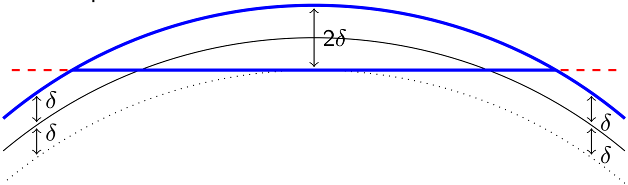

- Maximal "close" measure in one linear piece: circle chord of sagitta 2δ.
- Conclude: #pieces $\leq$ menu size; radius of curvature $\leq$ fixed r.

Menu Size Scalability as Market Grows

Menu Size Scalability as Market Grows

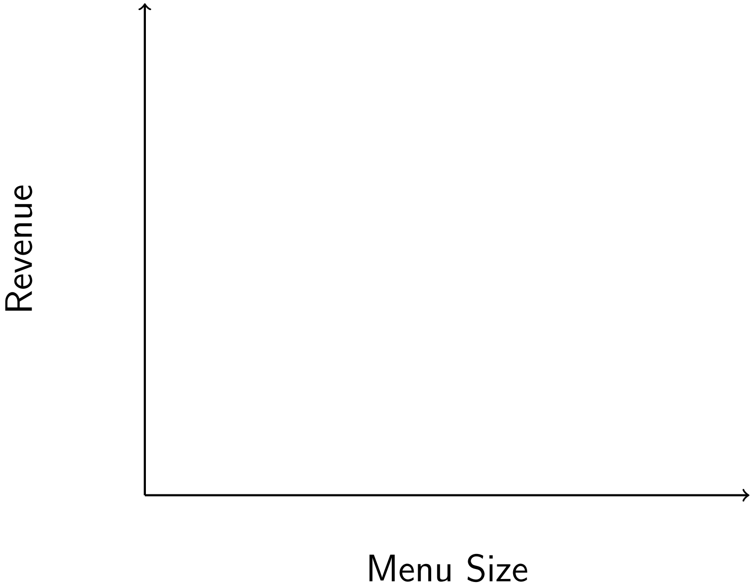

Menu Size Scalability as Market Grows

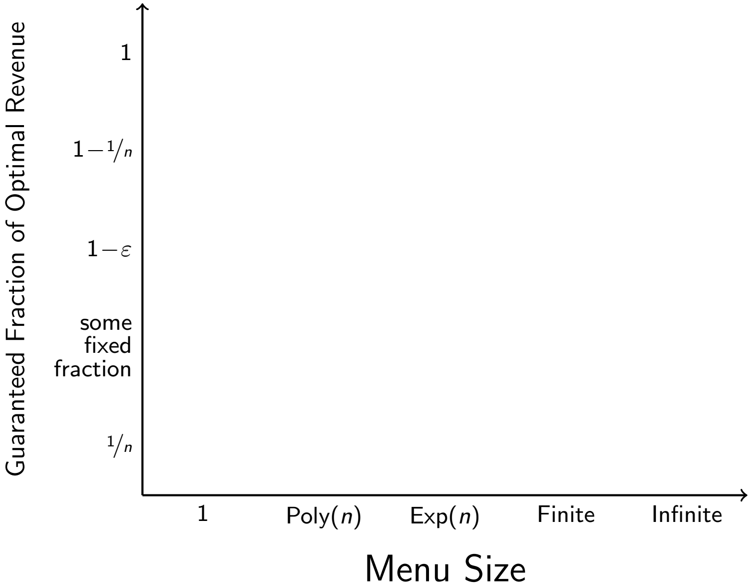

Menu Size Scalability as Market Grows

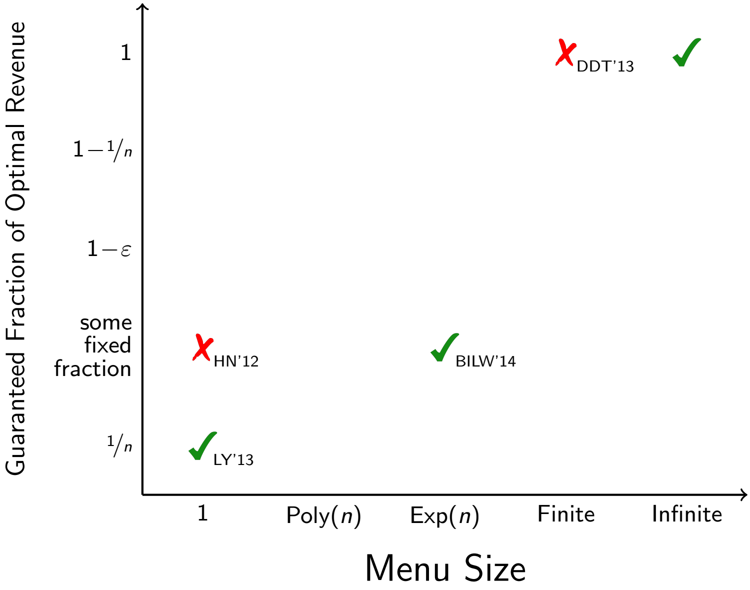

Menu Size Scalability as Market Grows

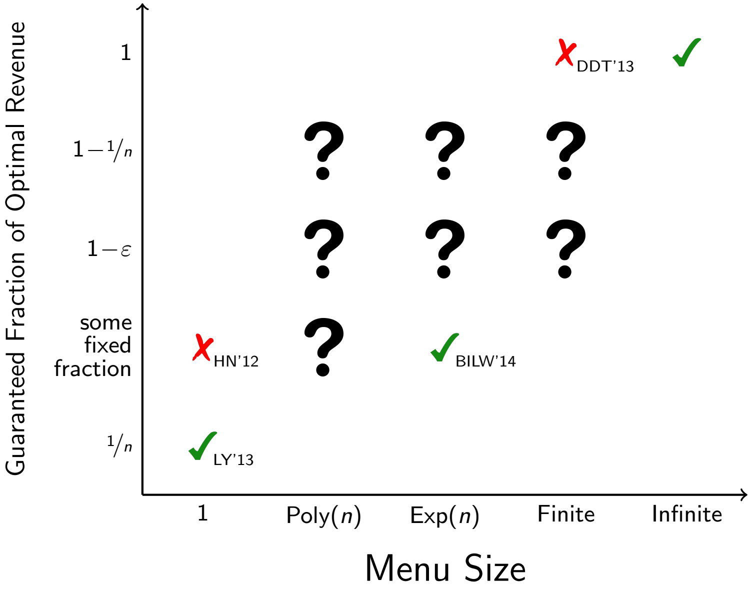

Menu Size Scalability as Market Grows

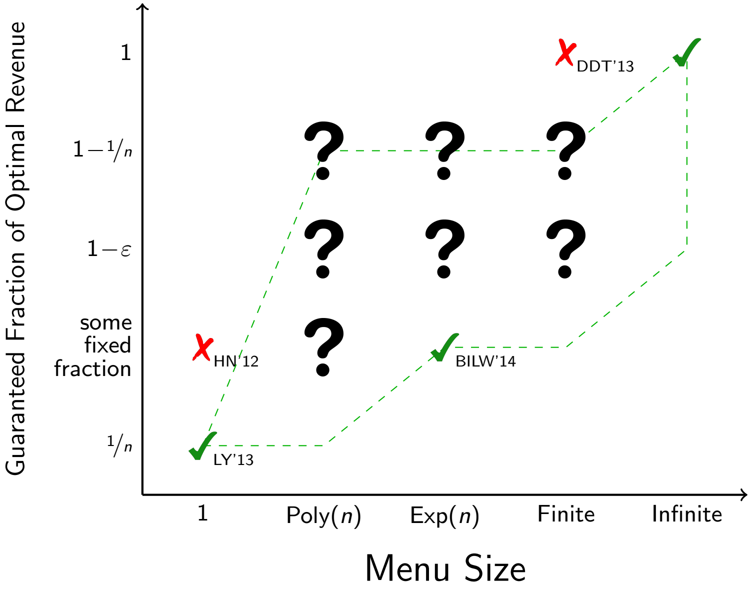

Menu Size Scalability as Market Grows

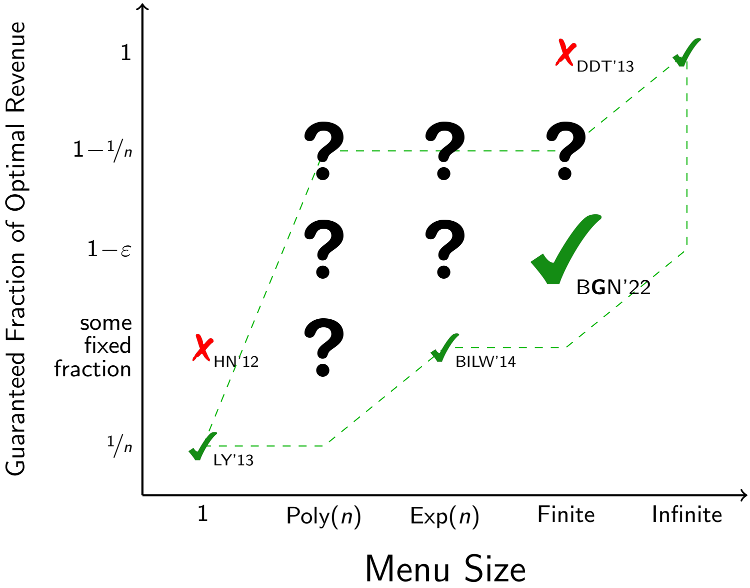

Menu Size Scalability as Market Grows

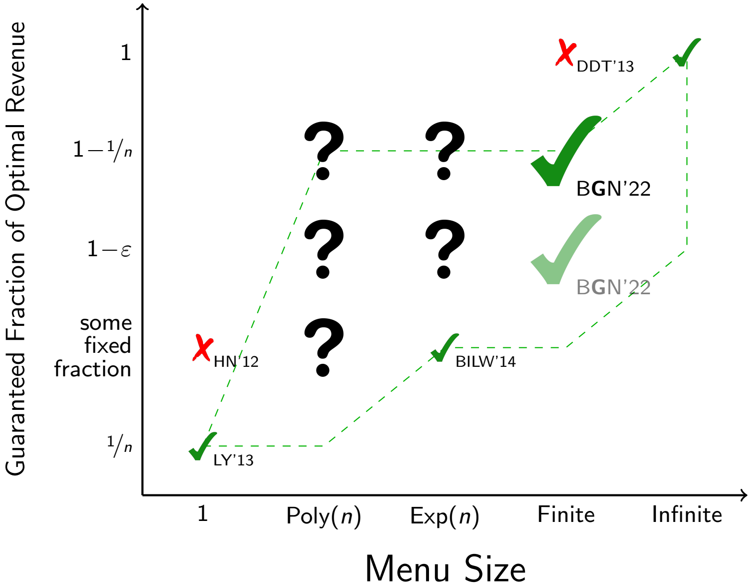

Menu Size Scalability as Market Grows

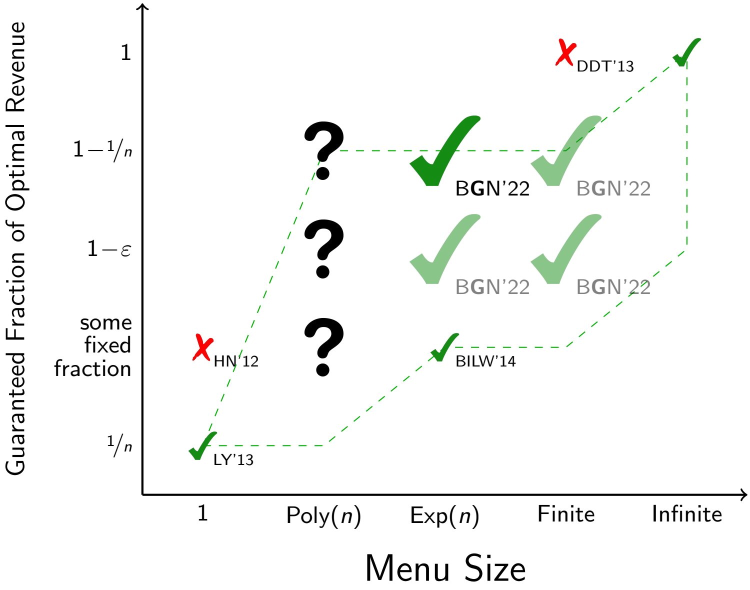

Menu Size Scalability as Market Grows

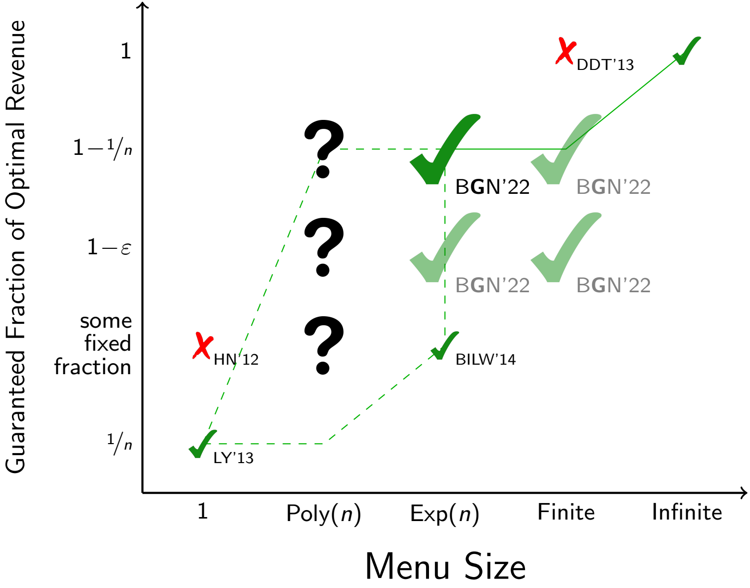

Menu Size Scalability as Market Grows

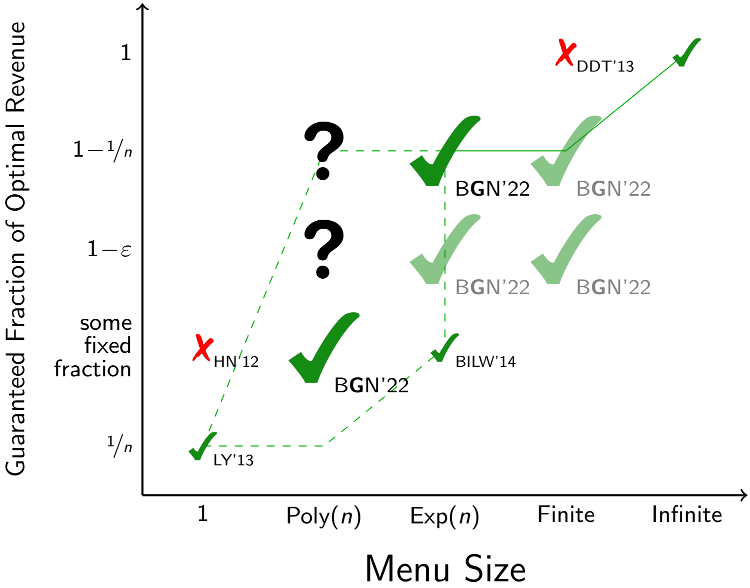

Menu Size Scalability as Market Grows

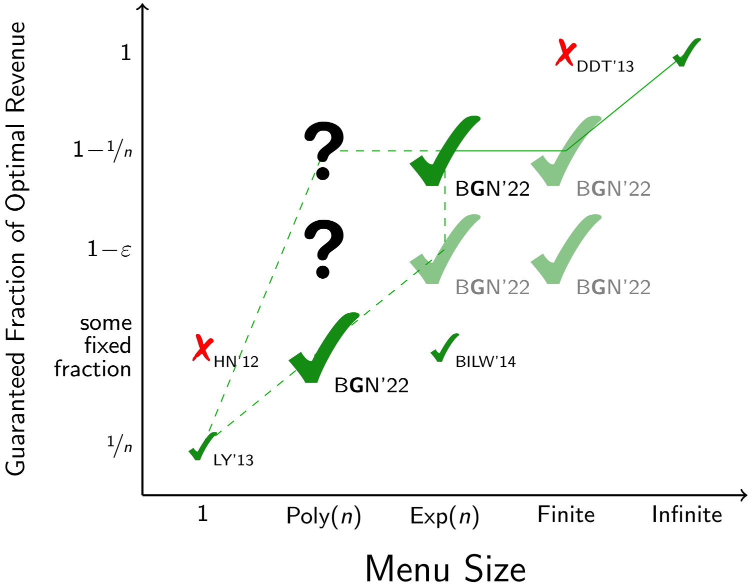

Menu Size Scalability as Market Grows

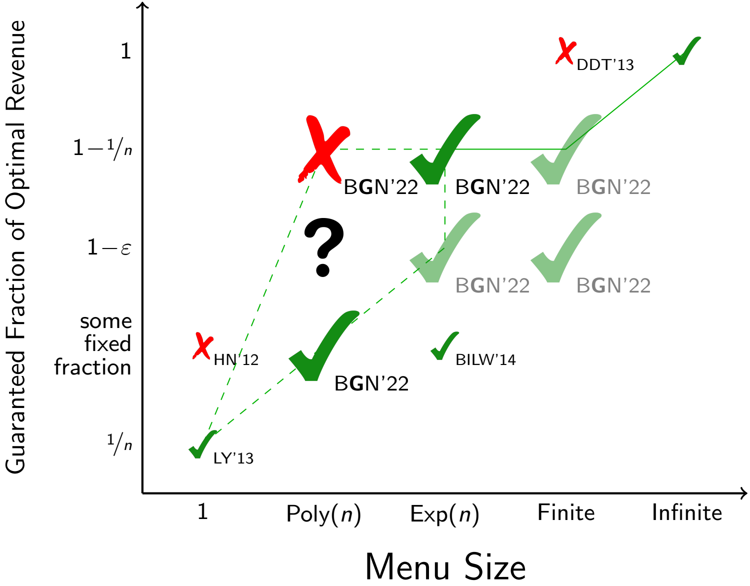

Menu Size Scalability as Market Grows

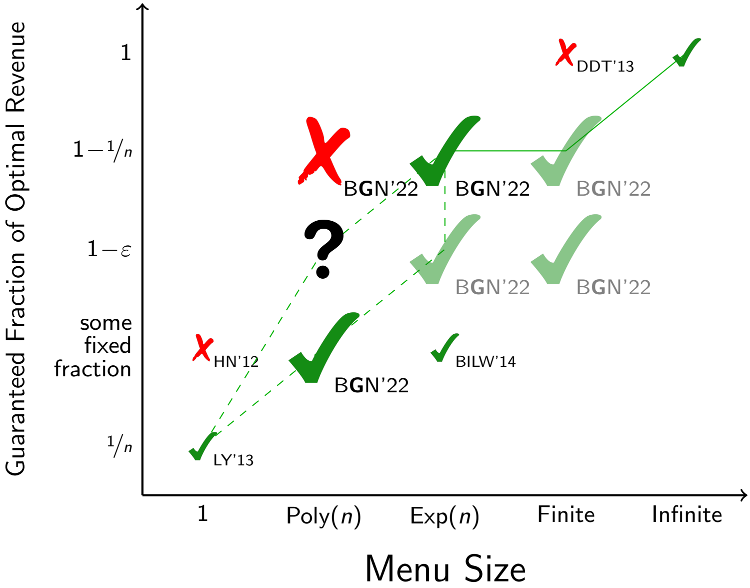

Correlated: Revenue Guarantee vs. Menu Size

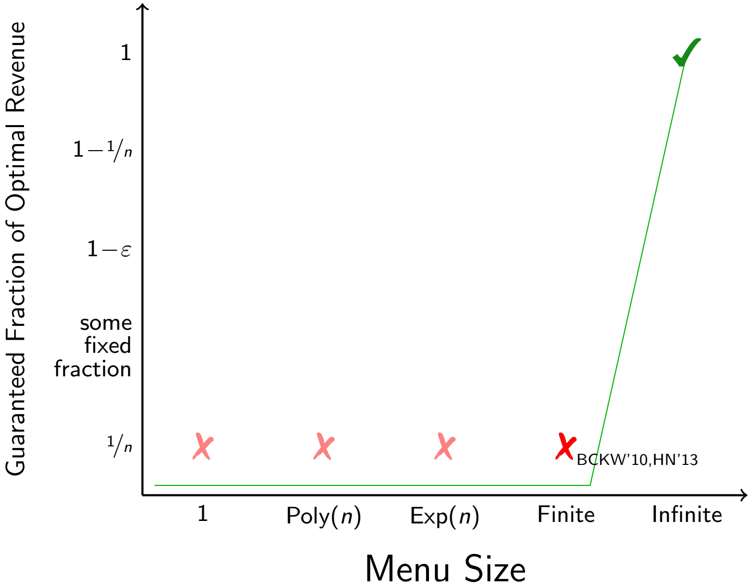

Menu Size Scalability as Market Grows

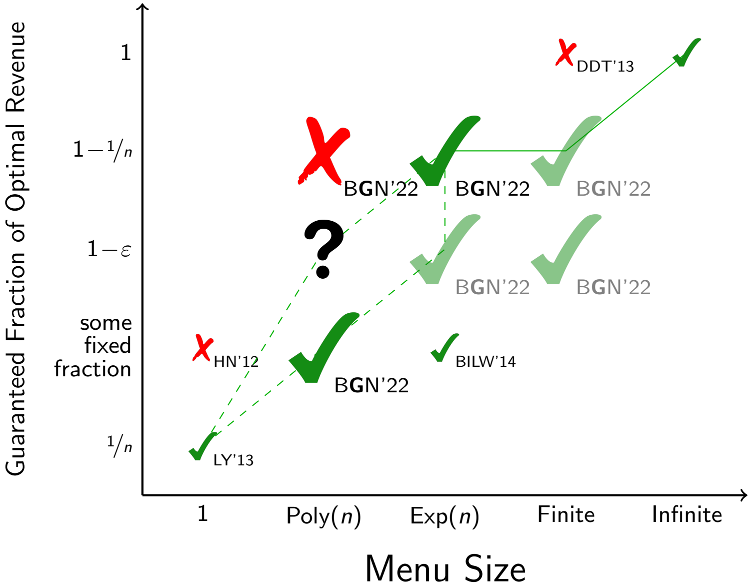

An Open Problem

- Main open problem: 99% of revenue via poly(n) menu-size, even for i.i.d. items, even for bounded distributions.

An Open Problem

- Main open problem: 99% of revenue via poly(n) menu-size, even for i.i.d. items, even for bounded distributions.
- The state-of-the-art literature seems to be a long way from identifying very-high-dimensional optimal mechanisms, and especially from identifying their duals (cf. GK'14).

An Open Problem

- Main open problem: 99% of revenue via poly(n) menu-size, even for i.i.d. items, even for bounded distributions.
- The state-of-the-art literature seems to be a long way from identifying very-high-dimensional optimal mechanisms, and especially from identifying their duals (cf. GK'14).
- One may hope that with time, it may be possible to do so.

An Open Problem

- Main open problem: 99% of revenue via poly(n) menu-size, even for i.i.d. items, even for bounded distributions.
- The state-of-the-art literature seems to be a long way from identifying very-high-dimensional optimal mechanisms, and especially from identifying their duals (cf. GK'14).
- One may hope that with time, it may be possible to do so.
- Plausibly, if one could generate high-dimensional optimal mechanisms (and corresponding duals) for which the high-dimensional analogue of the discussed strictly concave curve has large-enough measure (while maintaining a small-enough radius of curvature, etc.), then a proof similar to the above may be used to show that an exponential dependence on n is indeed required for sufficiently small, yet fixed, ε.

Questions?

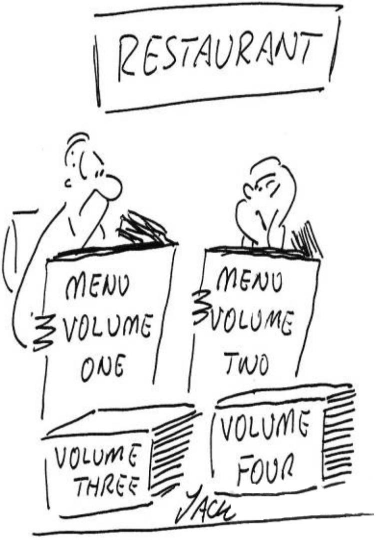

"Lots of choice, isn't there!"