Beyond Worst-Case Analysis in Algorithmic Game Theory

Inbal Talgam-Cohen, Technion CS

Games, Optimization & Optimism: Workshop in Honor of Uri Feige

Weizmann Institute, January 2020

Uri as an advisor

Q1: What did you appreciate most about Uri as an advisor?

Q2: What did you learn from him that has proved most meaningful over the years?

Uri as an advisor

- “Uri has scientific x-ray eyes. As a student, I observed with admiration his extraordinary capabilities of abstraction and presentation.
- Whenever I write a paper, or prepare a talk, I always use the Uri_Feige™ Latex/PowerPoint package.”

- Dan Vilenchik, BGU

Uri as an advisor

- "Working with Uri as an advisor was an inspiring experience, which helped me grow tremendously as a researcher.
- Privately, I used to call him "the oracle", for his tendency to spontaneously generate surprising insights and proof ideas almost mid-sentence, seemingly without any offline computational time."

- Eden Chlamtac, BGU

Uri as an advisor

- “My main insight from Uri is to keep it simple and look for simple and elegant solutions. His ability to simplify complicated problems never stopped amazing me.
- To see Uri solve mathematical questions was similar to listen to Glenn Gould play Bach: everything is so accurate and crystal clear.”

- Daniel Reichman, WPI

What I learned: Be accurate, be modest

From: Uriel.Feige@weizmann.ac.il

- “In Section 1.4 and elsewhere there are claims of the form `will be of independent interest’.
- I recommend to write instead `may be of independent interest’…
- …unless you know for sure that (a) it will be of interest, and (b) the interest will be independent of the application in the current paper.”

Beyond worst-case analysis in Uri’s Work

- Semi-random models:
  - A worst-case/average-case hybrid
  - Adversary and nature jointly produce problem instances
- [Feige-Krauthgamer’00, Feige-Kilian’01]:
  - Semi-random models for planted independent set
  - Insight into what properties of an IS make finding it easy
- Many additional works of Uri
  - Check out Uri’s forthcoming book chapter “Introduction to Semi-Random Models”

In this talk

- Some recent applications of the semi-random approach in algorithmic game theory (AGT)
  - [Carroll'17, Eden-Feldman-Friedler-T.C.-Weinberg'17, Duetting-Roughgarden-T.C.'19]
- A mystery in AGT:
  - Simple economic mechanisms are ubiquitous in practice…
  - … but suboptimal in the worst-case and average-case sense
- Semi-random models help explain, quantify and improve

Intersection of disciplinary approaches

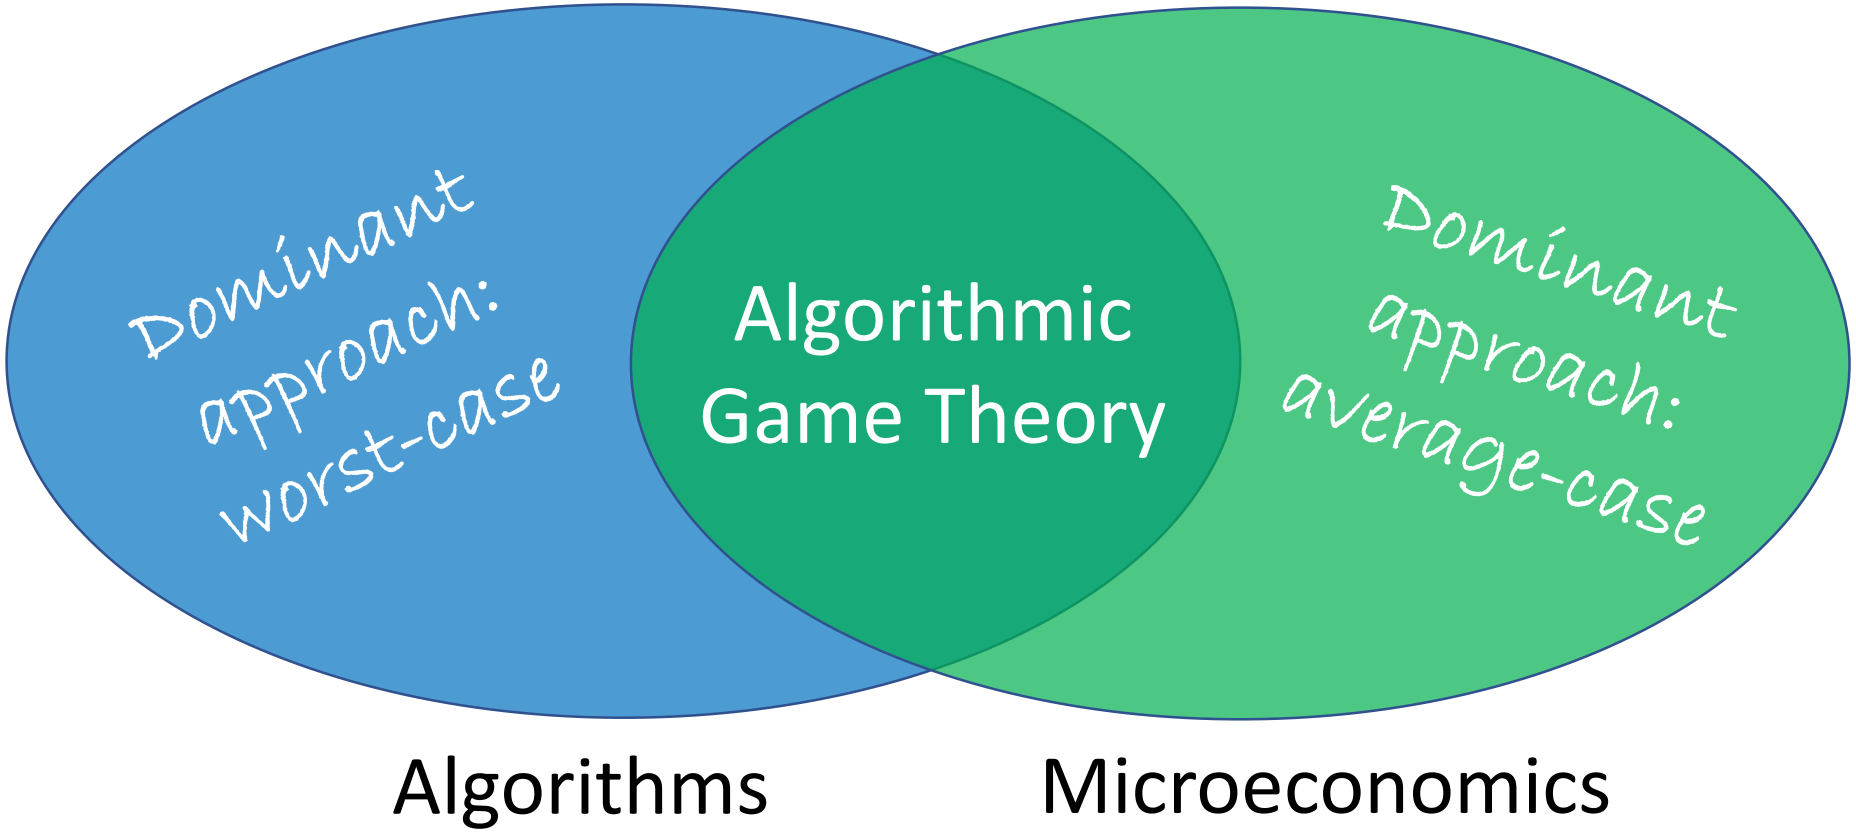

Mechanism design

Algorithm design with incentives, private information

- Agents use private information to maximize own utility
- Mechanisms use payments to maximize mechanism designer's utility a.k.a. revenue

Auction and contract design

1. Auctions:
    - Agents are buyers (e.g., online advertisers)
    - Private info: Buyers' values
    - Incentives: Auction induces buyers to bid their values
2. Contracts:
    - Agent hired to perform a task (e.g., online marketing)
    - Private info: Agent's effort level
    - Incentives: Contract induces efficient effort level

Simple ubiquitous mechanisms

1. Auctions:
   - 2nd-price auction – winner charged 2nd-highest bid
   - No incentive to underbid
   - As seen on: eBay
2. Contracts:
   - Linear contract – agent gets a cut of her effort's outcome
   - No incentive to slack off
   - As seen in: venture capital

Semi-random models for auctions

In what senses is the 2nd-price auction optimal for multi-item revenue?

Multi-item auction setting

n additive buyers

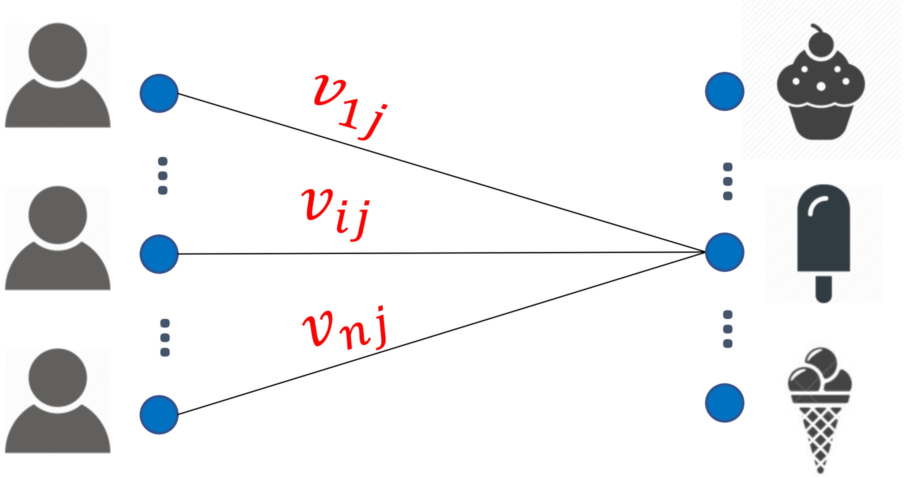

$m = \Theta(n)$ items

Bayesian (average-case) model

n additive buyers

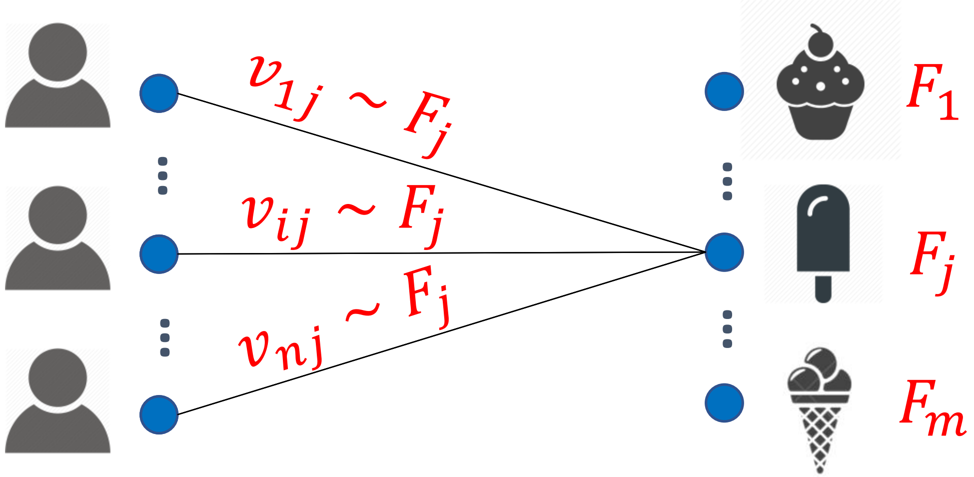

$m = \Theta(n)$ items

- Priors $F_1, \ldots, F_m$ known to auction
- Values sampled independently
- Auction gets bids, allocates items, charges payments

Average-case auction design?

- Design problem: Maximize expected revenue (total payment) subject to incentive compatibility (IC)
  - Expectation over priors $F_1, \dots, F_m$
  - IC = true bids maximize buyer utilities
- Notation: $\mathrm{OPT}_{F_1,\dots,F_m}$

- Auctions achieving $\mathrm{OPT}_{F_1,\dots,F_m}$ unrealistically complex for ≥ 2 items, and brittle even for 1 item

Worst-case auction design?

Nonstarter even for 1 item, 1 buyer with value v

- Design problem: Maximize revenue by setting reserve price p
- But: $\forall$ p $\exists$ worst-case value v s.t. revenue = 0

Semi-random to the rescue

- Semi-random models - recall:
  - A worst-case/average-case hybrid
  - Adversary and nature jointly produce problem instances

- In auctions:
  - Class of priors $\mathcal{F}$ known to auction
  - Adversary chooses worst-case prior $F \in \mathcal{F}$
  - Nature samples instance $v \sim F$

Semi-random instance generation

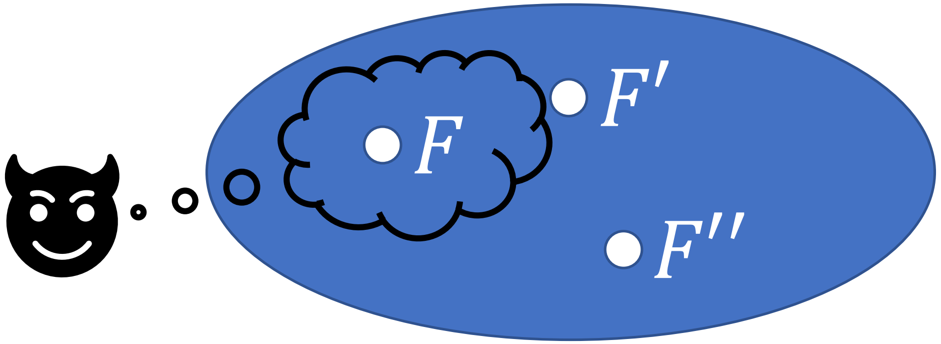

Class $\mathcal{F}$ of priors

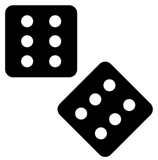

Instance v drawn from F

Two performance measures

Consider mechanism M

Recall OPT_F = $\mathbb{E}_F[\text{revenue of optimal mechanism for prior } F]$

Approximation ratio

1. Relative: $\min_{F \in \mathcal{F}} \left\{ \frac{\mathbb{E}_F[\text{revenue of } M]}{\mathrm{OPT}_F} \right\}$

2. Absolute: $\min_{F \in \mathcal{F}} \{\mathbb{E}_F[\text{revenue of } M]\}$

Two design goals

1. Maximize relative performance
   - Find M that approximates OPT_F simultaneously $\forall F \in \mathcal{F}$
   - Terminology: M is prior-independent [Dhangwatnotai'15]

2. Maximize absolute performance
   - Find M that achieves $\max_{M'} \min_{F \in \mathcal{F}} \{\mathbb{E}_F[\text{revenue of } M']\}$
   - Terminology: M is max-min optimal [Bertsimas'10, Carroll'19]

Choice of $\mathcal{F}$ is crucial

Recent results

- Prior-independent auctions
  1. Via extra buyers:
     - [Feldman-Friedler-Rubinstein EC’18] $(1-\epsilon)$-approximation
     - [Beyhaghi-Weinberg STOC’19] Improved and tight bounds
     - [Liu-Psomas SODA’18] Dynamic auctions
     - [Roughgarden-T.C.-Yan OR’19] Unit-demand buyers
  2. Via sampling + approximation:
     - [Allouah-Besbes EC’18] Lower bounds
     - [Babaioff-Gonczarowski-Mansour-Moran EC’18] Two samples
     - [Guo-Huang-Zhang STOC’19] Settling sample complexity
- Max-min optimal auctions
  - [Gravin-Lu SODA’18] With budgets
  - [Bei-Gravin-Lu-Tang SODA’19] Posted prices

Result 1: Max-min optimality [Carroll'17]

Setting: 1 buyer, m items with priors $F_1, \ldots, F_m$

$\mathcal{F}$ = all correlated distributions with marginals $F_1, \ldots, F_m$

Theorem [Carroll]: Selling each item j separately by 2nd-price auction with optimal reserve for F_j is max-min optimal wrt $\mathcal{F}$

Intuition: Selling separately is robust to correlation

Max-min optimality

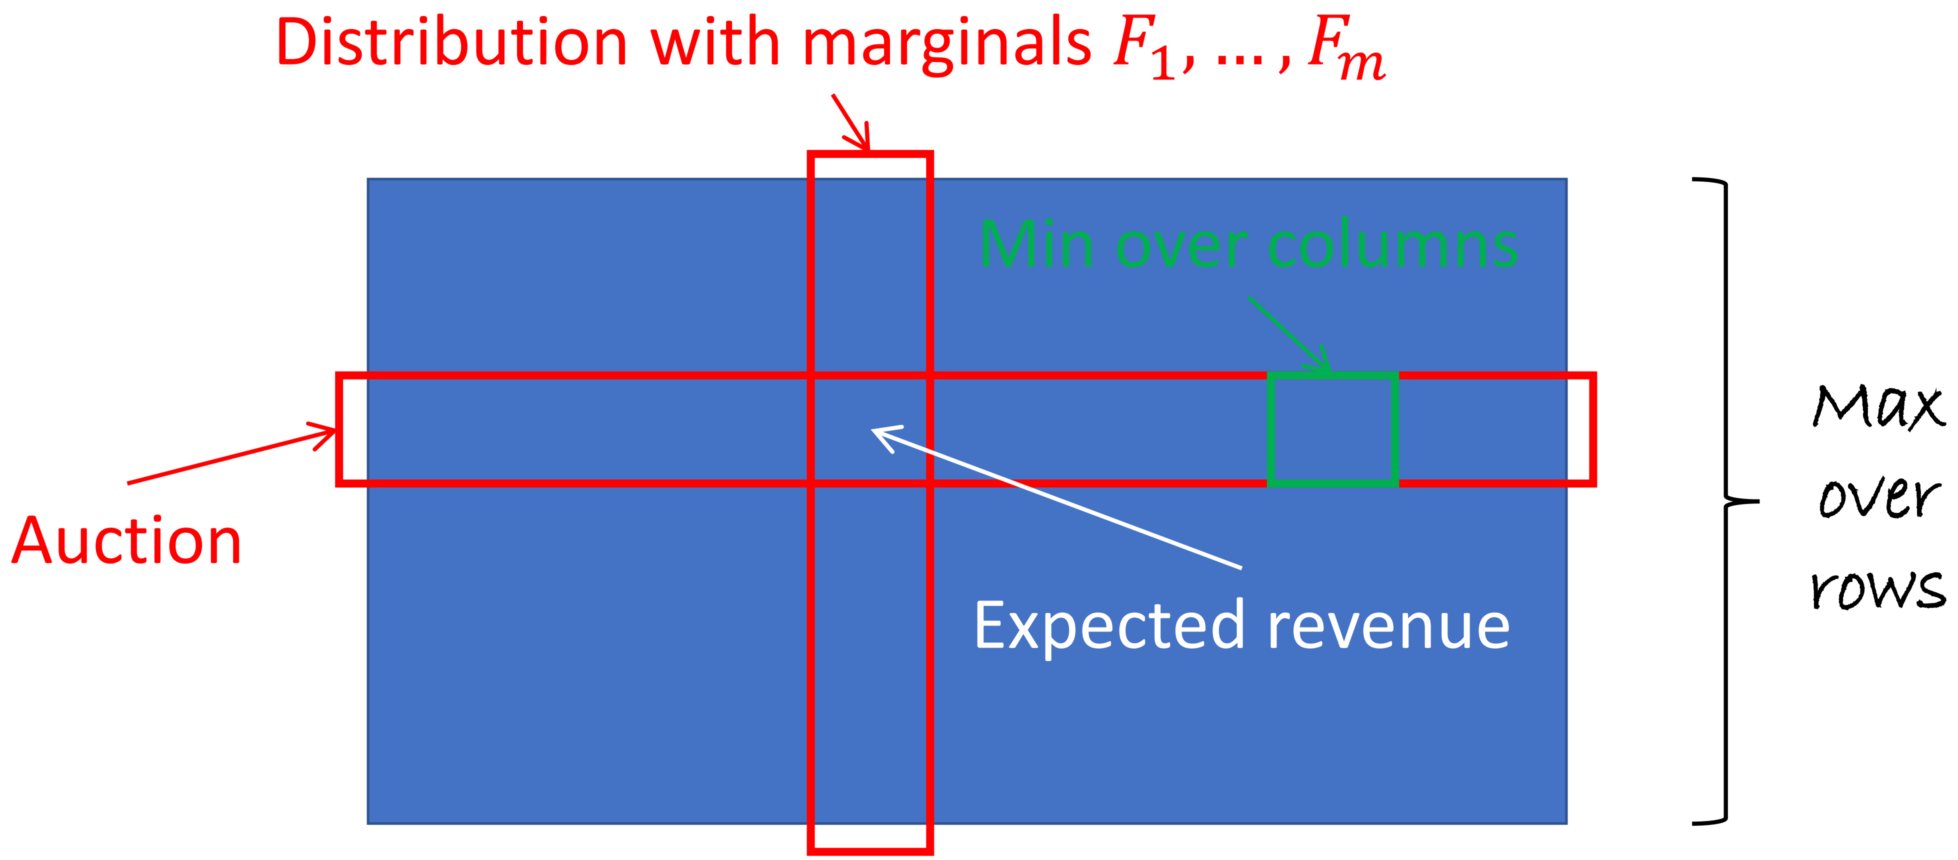

Robustness to correlation

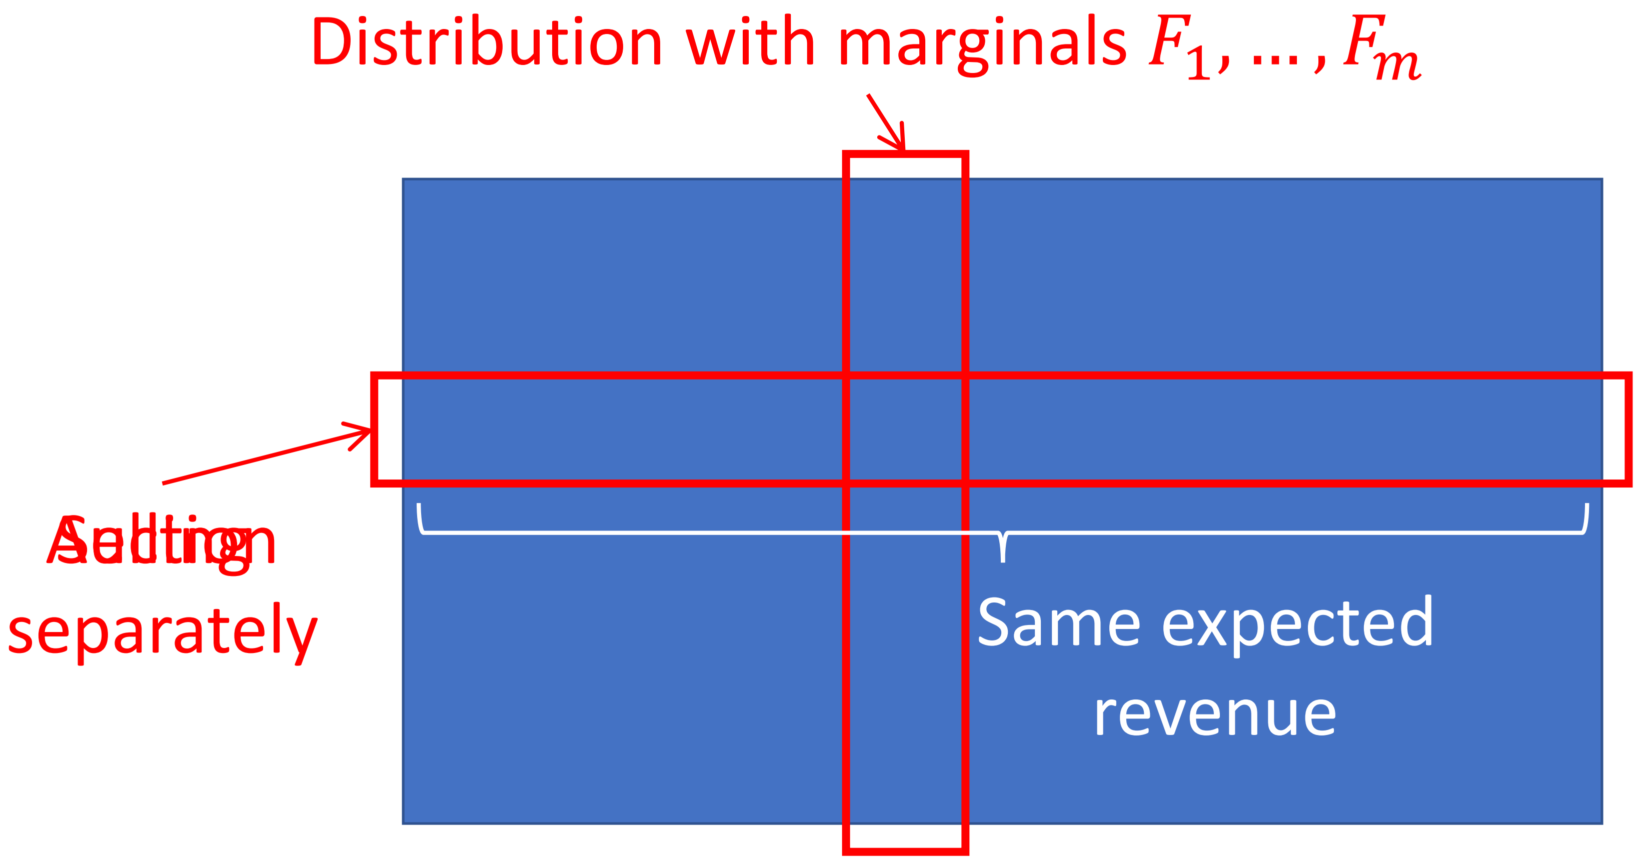

Towards result 2: What more do we want?

Recall theorem: Selling each item j separately by 2nd-price auction with optimal reserve for F_j is max-min optimal wrt $\mathcal{F}$

Want: Prior-independence

- No reserve price tailored to F_j
- Revenue guarantee relative to $\mathrm{OPT}_{F_1,\ldots,F_m}$

Willing to: assume values are independent

First attempt

Setting: n buyers, m items

$\mathcal{F}$ = all product distributions $F_1 \times \cdots \times F_m$ with regular marginals

"Theorem": Selling each item j separately by 2nd-price auction approximates $\mathrm{OPT}_{F_1,\ldots,F_m}$ simultaneously $\forall F_1 \times \cdots \times F_m \in \mathcal{F}$

Counterexample: 1 buyer

Resource augmentation

- Another beyond worst-case approach
- To compete with a powerful benchmark, the algorithm is allowed extra resources [Sleator-Tarjan'85]

- In our context [BulowKlemperer'96]:
  - Powerful benchmark is $\mathrm{OPT}_{F_1,\dots,F_m}$
  - Resources are buyers competing for the items

Result 2: Prior-independence [Eden+'17]

<u>Theorem</u>: With $O(m)$ extra buyers, selling each item j separately by 2nd-price auction matches $\mathrm{OPT}_{F_1,\dots,F_m}$ simultaneously $\forall F_1 \times \cdots \times F_m \in \mathcal{F}$

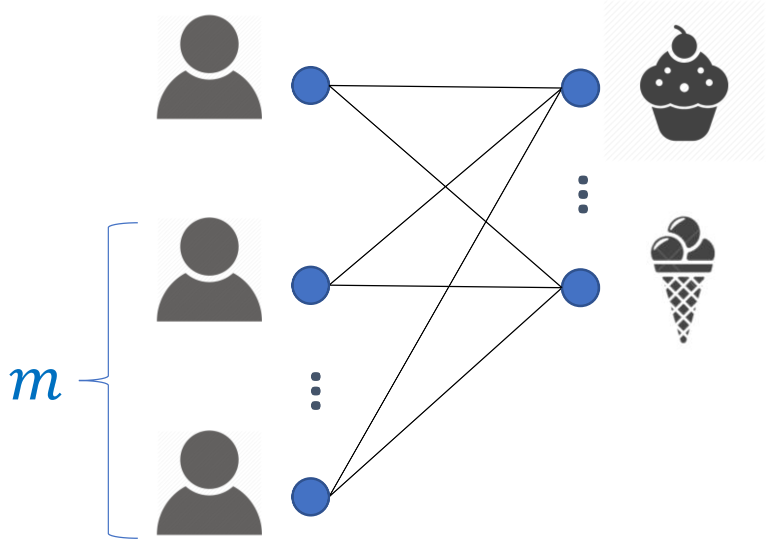

Result 2: Prior-independence [Eden+'17]

Theorem: With $O(m)$ extra buyers, selling each item j separately by 2nd-price auction matches $\mathrm{OPT}_{F_1,\dots,F_m}$ simultaneously $\forall F_1 \times \cdots \times F_m \in \mathcal{F}$

- [Feldman-Friedler-Rubinstein'18]: $\Omega(m)$ extra buyers necessary for $m = \Theta(n)$
- [Beyhaghi-Weinberg'19]: Additional tight results for other n, m regimes

Auctions Recap

- For the canonical problem of maximizing revenue from m items, semi-random models show that simple auctions are optimal
- Simple = selling each item by 2nd-price auction with reserve or more buyers
- Optimal =
  - Max-min optimal over adversarily chosen correlation or
  - Match OPT_F simultaneously for any regular product distribution F

Semi-random models for contracts

In what sense are linear contracts optimal?

Bayesian model for contracts

- Agent has n possible effort levels (hidden)
- Level i induces a distribution over m (observable) outcomes
  - μ_i = expected outcome
  - c_i = cost
- Example:

	Low outcome $4	Med. outcome $50	High outcome $100	
Low effort $0	0.6	0.3	0.1	$\mu_1 = 27.4$
Med. effort $2	0.4	0.4	0.2	$\mu_2 = 41.6$
High effort $9	0.1	0.5	0.4	$\mu_3 = 65.4$

Bayesian model for contracts

- Contract = non-negative payment for every outcome
- Revenue = outcome minus payment
  - Measured in expectation over outcome distribution given effort

Contract:	$2	$30	$45
	Low outcome $4	Med. outcome $50	High outcome $100
Low effort $0	0.6	0.3	0.1
Med. effort $2	0.4	0.4	0.2
High effort $9	0.1	0.5	0.4

Linear contracts

- A linear contract is defined by a parameter $\alpha \leq 1$
- Agent chooses level i^* that maximizes $\alpha\mu_i - c_i$
- Expected revenue is $(1-\alpha)\mu_{i^*}$

Contract:	4α =\$2	50α =\$25	100α =\$50
	Low outcome \$4	Med. outcome \$50	High outcome \$100
Low effort \$0	0.6	0.3	0.1
Med. effort \$2	0.4	0.4	0.2
High effort \$9	0.1	0.5	0.4

Result 3: Max-min optimality [Duetting+'19]

Setting: n effort levels with expected outcomes $\mu_1, \dots, \mu_n$

$\mathcal{F}$ = distributions $F_1, \dots, F_n$ with expectations $\mu_1, \dots, \mu_n$

Theorem: The optimal linear contract for $\mu_1, \dots, \mu_n$ is max-min optimal wrt $\mathcal{F}$

See also [Carroll'15, Dai-Toikka'18, Carroll-Walton'20]

Same intuition as Result 1 for auctions

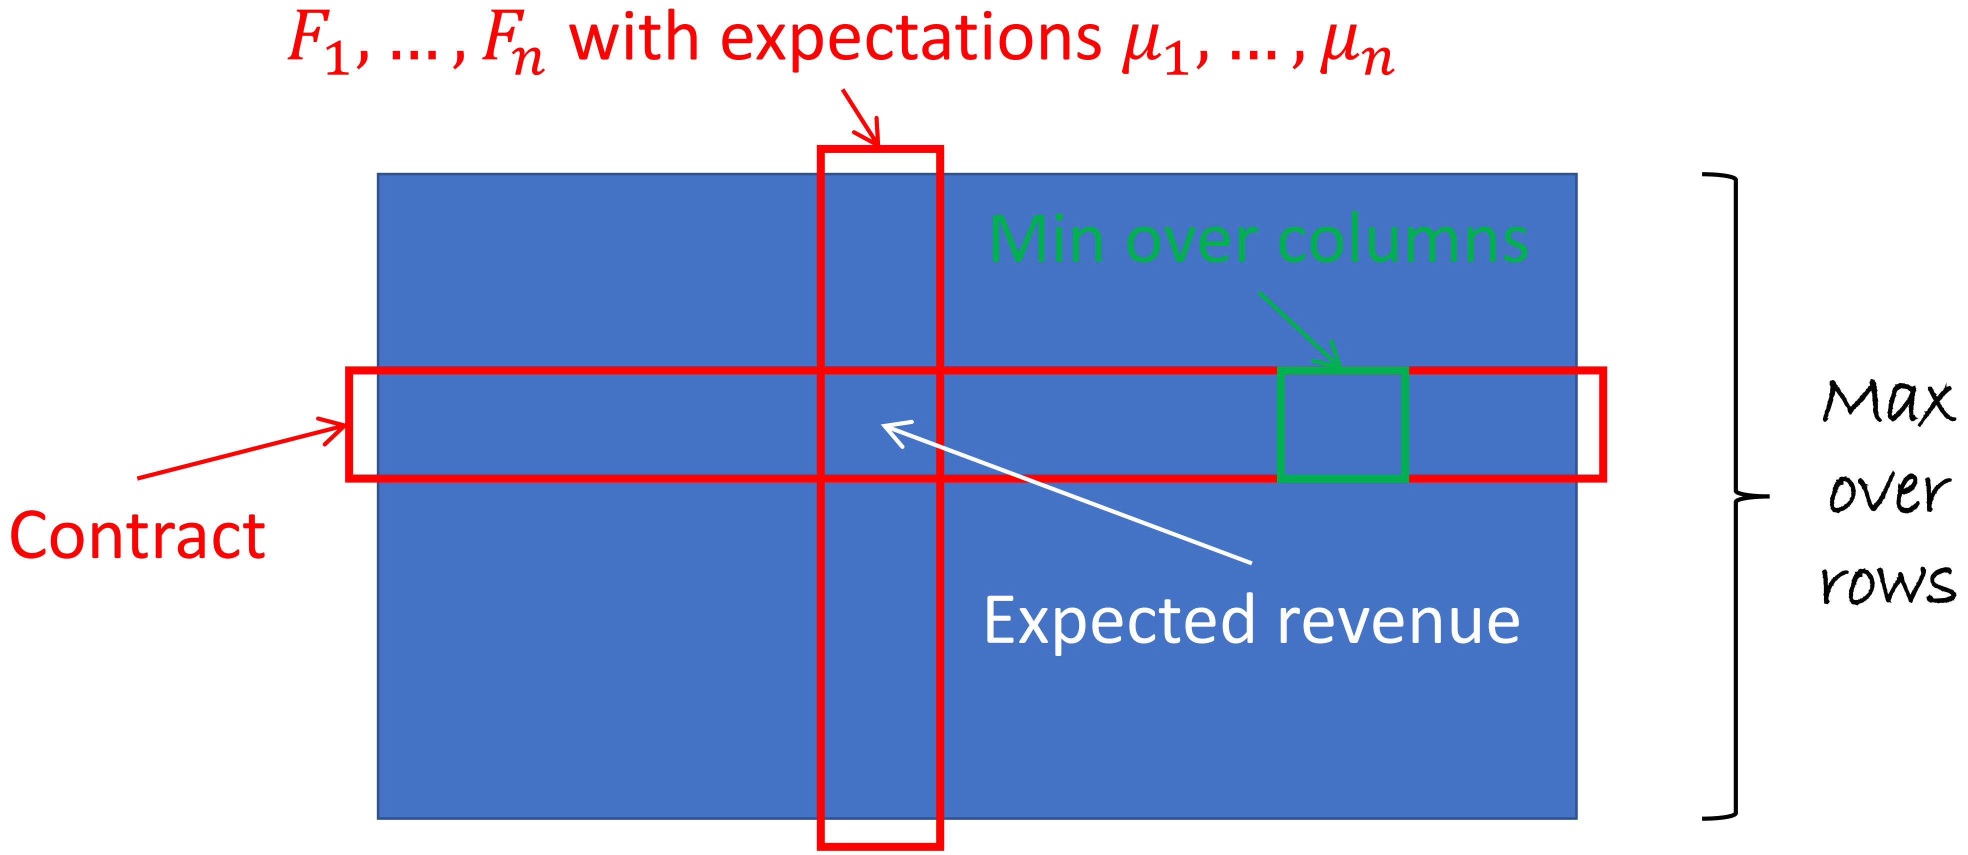

Same intuition as Result 1 for auctions

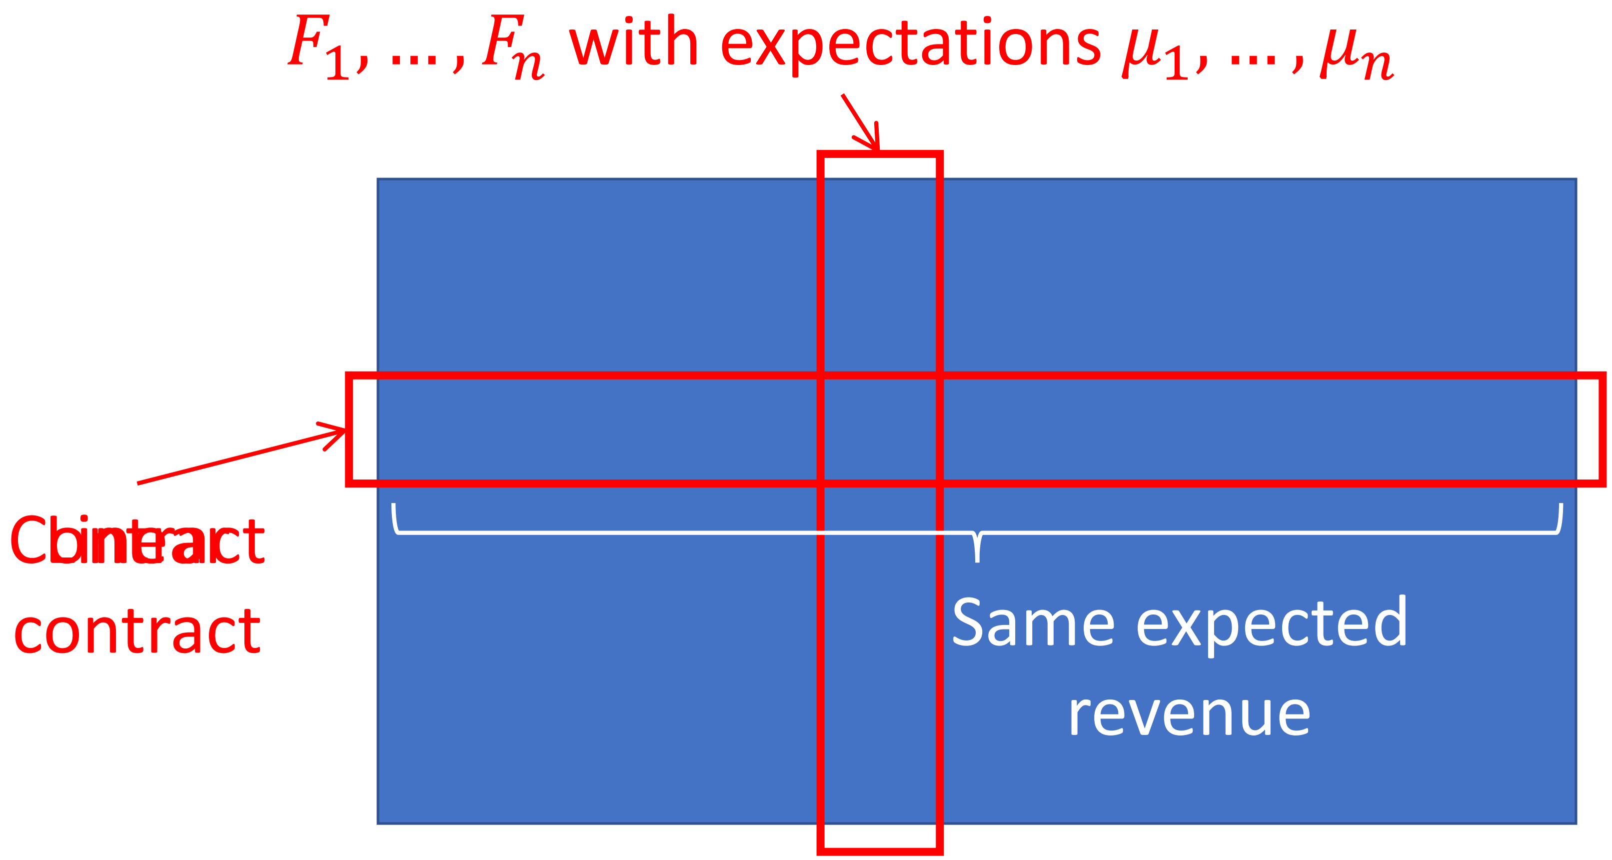

Adversary takes advantage of any non-robustness

Main lemma: ∀ contract, ∃ distributions with expectations $\mu_1, \ldots, \mu_n$ s.t. ∃ linear contract with better expected revenue

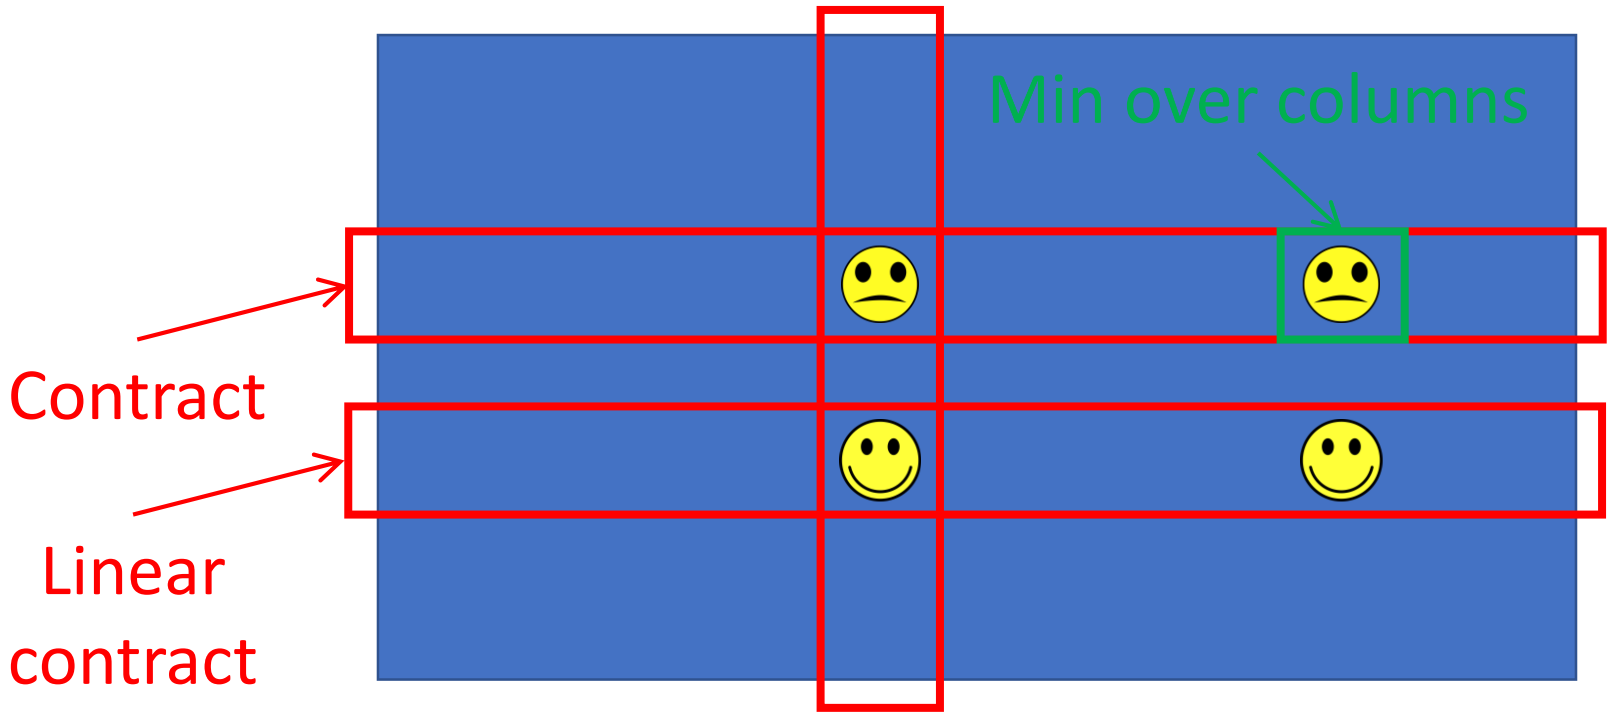

Take-away

Lots of recent beyond worst-case activity in AGT leading to new insights

"It is probably the great robustness of [simple mechanisms] that accounts for their popularity.

That point is not made as effectively as we would like by our model; we suspect that it cannot be made effectively in any traditional Bayesian model."

[Holmstrom-Milgrom'87]

Open problems

1. Auctions: beyond additive buyers?
2. Contracts: relative guarantees a la prior-independence?
3. General framework for max-min robustness?

Thanks for listening!