Within-Instance Mechanism Design

Siddharth Prasad
Carnegie Mellon University

Based on joint work with:
Maria-Florina Balcan (CMU) and Tuomas Sandholm (CMU, Optimized Markets, Inc., Strategic Machine, Inc., Strategy Robot, Inc.)

Learning Within an Instance for Designing High-Revenue Combinatorial Auctions. *IJCAI 2021.*
Maximizing Revenue Under Market Shrinkage and Market Uncertainty. *NeurIPS 2022.*

INFORMS 2022

Background and motivation

- Classical mechanism design
  - 1981: Myerson showed how to sell a single item to maximize revenue (uses details of the distribution of buyers’ values for the item)
  - Today: don’t know how to sell two items optimally

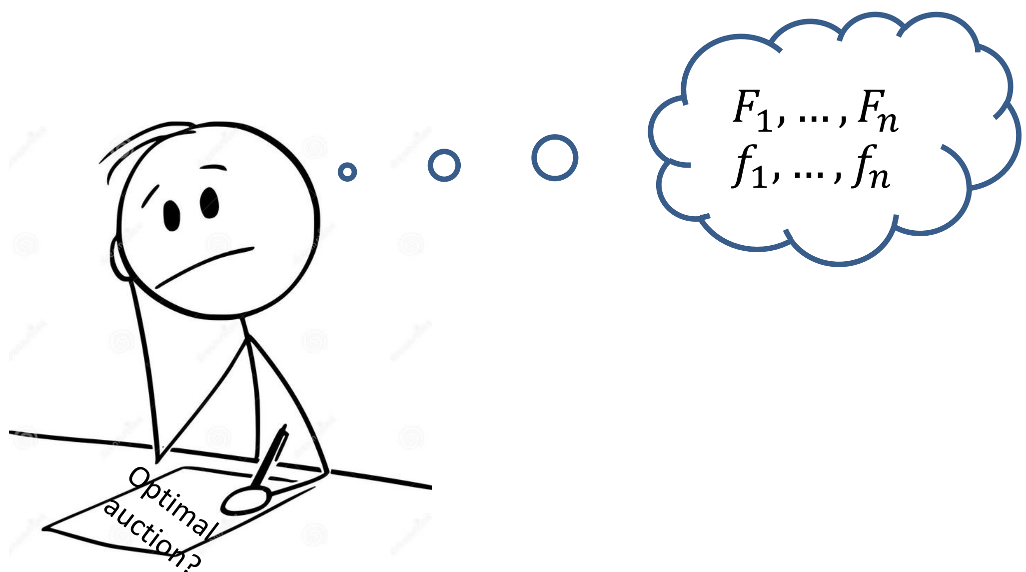

Background and motivation

- Automated mechanism design [Conitzer and Sandholm UAI '02]:
  - why struggle with the hard economics problem of designing explicit mechanisms when a computer program can do it for you?
  - requiring details of distributions -> computational hardness in many settings

Background and motivation

- Sample-based automated mechanism design
  - use machine learning, don't need details about distribution [Sandholm and Likhodedov AAAI '04, '05, Operations Research '15, Mohri and Medina ICML '14, Morgenstern and Roughgarden NIPS '15, Balcan, Sandholm, and Vitercik EC'18, Duetting, Feng, Narasimhan, Parkes, Ravindranath ICML'19, Balcan, **Prasad**, and Sandholm IJCAI'20]

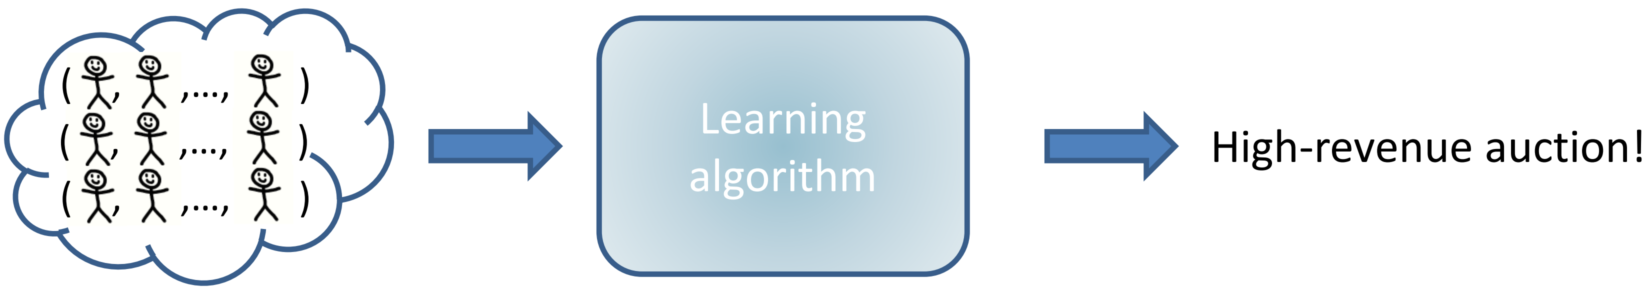

Background and motivation

- Within-instance mechanism design
  - adapting sample-based techniques to understand prior-free mechanism design [Balcan, **Prasad**, and Sandholm IJCAI’21, NeurIPS’22]

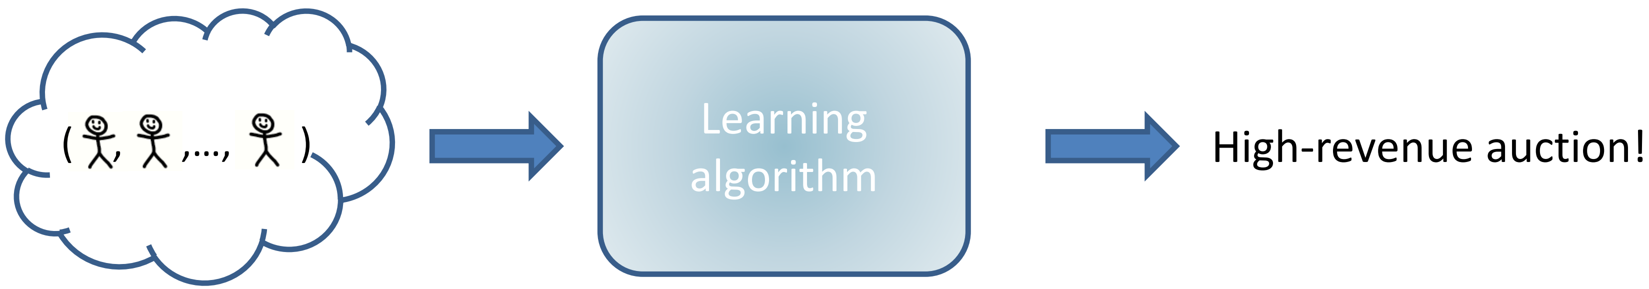

Outline

Part 1

Seller faces a population of buyers. Can he "learn within an instance" to find a high-revenue auction?

Part 2

Seller faces a (known) population of buyers. Can he learn an auction that extracts high revenue from a shrinking market?

Combinatorial auctions crash course

- Seller has m indivisible items to sell among set S of n bidders. *(limited supply)*
- Bidders have combinatorial valuations $v_i: 2^{\{1,\dots,m\}} \to \mathbf{R}_{\geq 0}$.

- For reported valuations $v_1, \dots, v_n$, an auction M specifies an allocation $\alpha(v_1, \dots, v_n)$ and payments $p_i(v_1, \dots, v_n)$.

- Seller wants to design M that extracts high revenue in an incentive compatible manner, that is,

$$v_i\big(\alpha(v_1, \dots, v_n)\big) - p_i(v_1, \dots, v_n) \geq v_i\big(\alpha(v_1, \dots, \hat{v}_i, \dots, v_n)\big) - p_i(v_1, \dots, \hat{v}_i, \dots, v_n)$$

Combinatorial auctions crash course

Vickrey-Clarke-Groves (VCG) auction:

- use allocation α^* that maximizes welfare

$$W(\alpha) = \sum_{i=1}^{n} v_i(\alpha)$$

- bidder i pays

$$\max_{\alpha} \sum_{j \neq i} v_j(\alpha) - \sum_{j \neq i} v_j(\alpha^*)$$

- VCG is incentive compatible

Combinatorial auctions crash course

<u>λ-auction:</u> parameterized by $\lambda \in \mathbf{R}^{(n+1)^m}$

- use allocation α^* that maximizes welfare plus boost

$$\sum_{i=1}^{n} v_i(\alpha) + \lambda(\alpha)$$

- Bidder i pays

$$\max_{\alpha}\left[\sum_{j \neq i} v_j(\alpha) + \lambda(\alpha)\right] - \left[\sum_{j \neq i} v_j(\alpha^*) + \lambda(\alpha^*)\right]$$

- Many other parameterized generalizations of VCG exist

Outline

Part 1

Seller faces a population of buyers. Can he "learn within an instance" to find a high-revenue auction?

Part 2

Seller faces a known population of buyers. Can he learn an auction that extracts high revenue from a shrinking market?

Random-sampling for unlimited supply

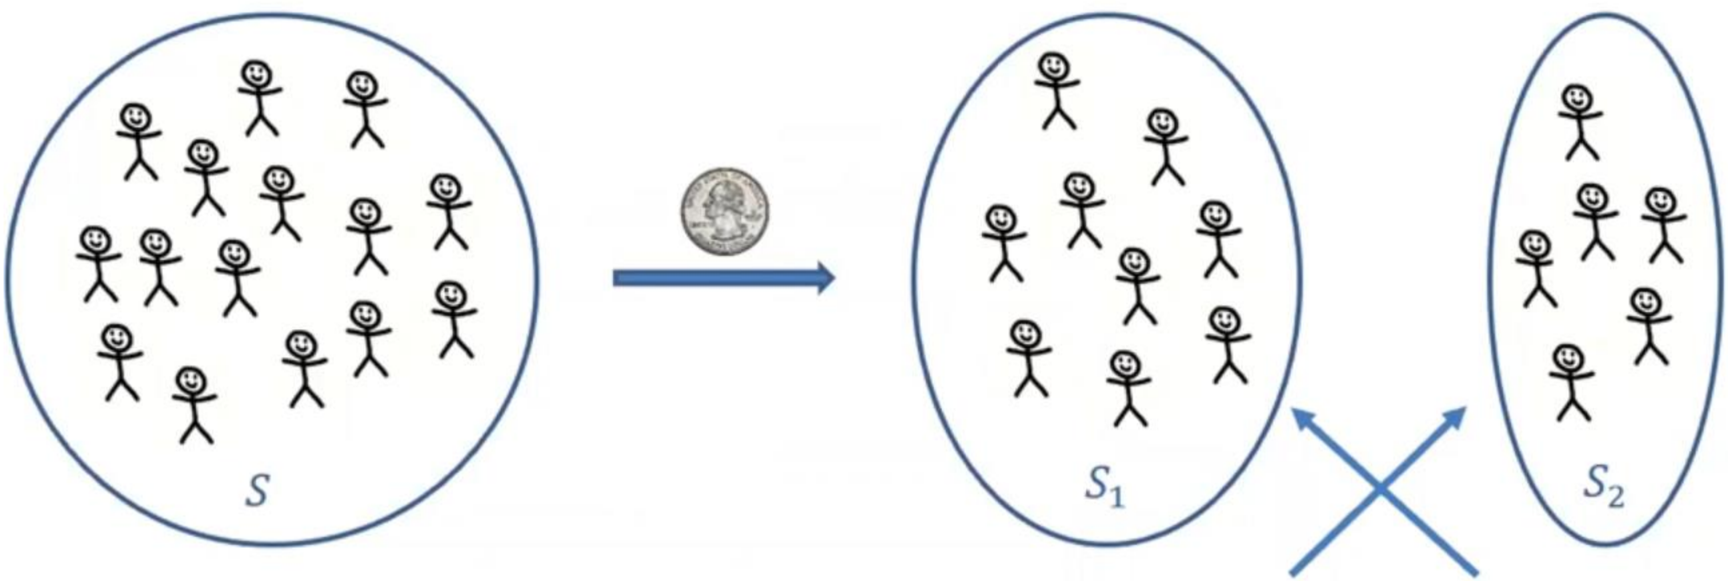

Revenue = $\mathrm{Rev}_{\mathrm{M}_1}(S_2) + \mathrm{Rev}_{\mathrm{M}_2}(S_1)$

M_1 optimal mechanism for S_1

M_2 optimal mechanism for S_2

- Converges to OPT as # bidders grows
- **Crucially relies on lack of supply constraints.**

Difficulties of limited supply

- Try to adapt random sampling to limited supply
- Attempt 1: Partition S into S_1, S_2 as before, compute optimal mechanism for S_1, run on S_2
  - Single item, sold via second-price auction with reserve price
  - Item goes to highest bidder v_1, payment = $\max\{v_2, r\}$

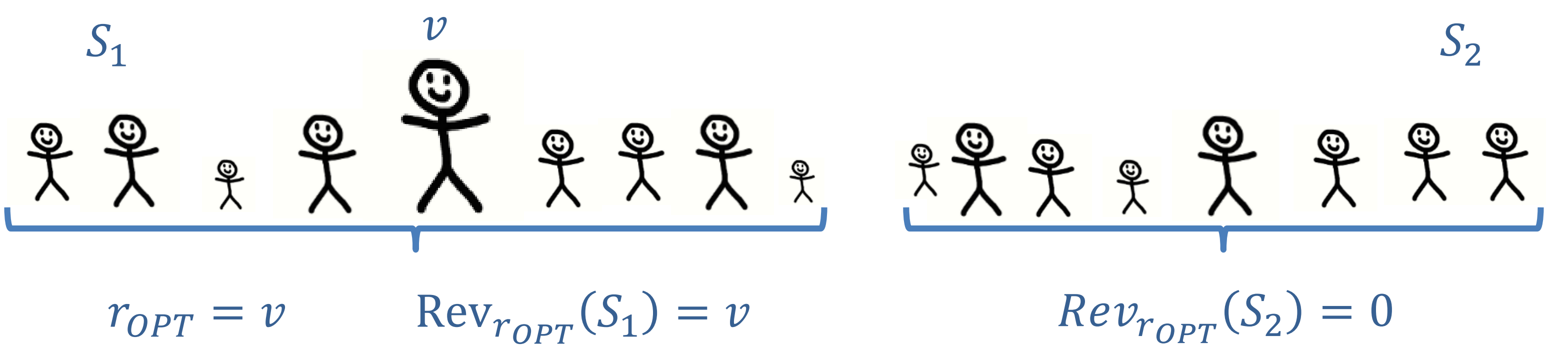

With probability ½, revenue = 0!

Difficulties of limited supply

- Try to adapt random sampling to limited supply
- Attempt 2: Partition S into $S_1,\ldots,S_N,S_{N+1}$. Compute empirically optimal mechanism for $S_1,\ldots,S_N$, run on S_{N+1}
  - If auction class is complex, N needs to be large for generalization guarantees to hold
  - S_{N+1} contains a tiny fraction of bidders, losing a lot of revenue

Learning Within an Instance (LWI) mechanism

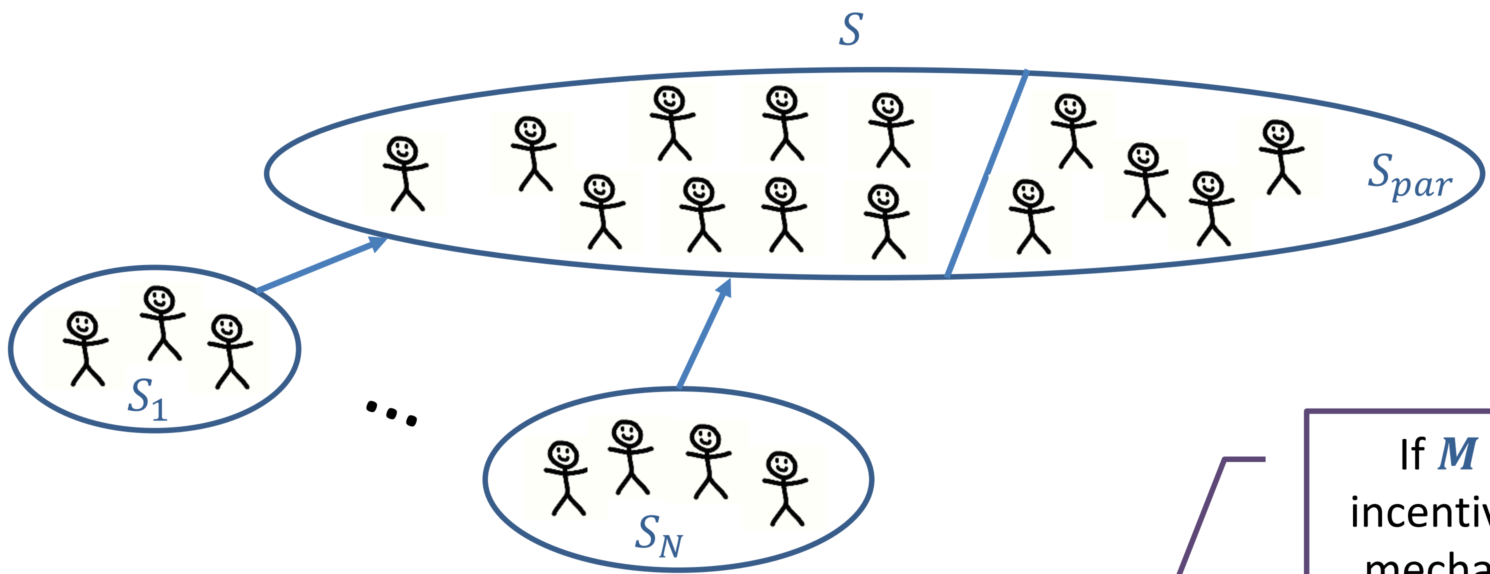

If $\boldsymbol{M}$ is a class of incentive compatible mechanisms, LWI is incentive compatible

1. Sample participatory group $S_{par} \sim_p S$
2. Sample learning groups $S_1, \dots, S_N \sim_q S \setminus S_{par}$
3. Compute ERM mechanism $\widehat{M} \in \boldsymbol{M}$ over learning groups
4. Run mechanism $\widehat{M}$ on S_{par}

Main guarantee

Theorem (Balcan, **Prasad**, Sandholm IJCAI'21)**.** For $N \geq N_{\boldsymbol{M}}(\varepsilon, \delta)$, and $W(S)$ sufficiently large,

$$Rev_{\widehat{M}}(S_{par}) \geq W(S)\big((1-\eta)p - \varepsilon\big) - 2\tau_{\boldsymbol{M}}(q, S_{par})$$

with probability $\geq 1 - 2\delta$ over the draw of the learning groups. $\boldsymbol{M}$ is the class of bundling-boosted auctions.

Parameterized class of VCG-like auctions that give boosts to specific bundlings

Term that captures uniformity in bidder set. Related to fundamental quantities in learning theory: pseudodimension, covering numbers, etc.

Comparison to existing results

- Subadditive valuations
  - $O(2^{\sqrt{\log m \log\log m}})$-approximation [Balcan, Blum, Mansour EC 2008]
  - $O(\log^2 m)$-approximation [Chakraborty, Huang, Khanna FOCS 2009; SICOMP 2013]
- Additive valuations
  - $O(\log(h/l))$-approximation, h, l highest and lowest values for any bundle [Sandholm and Likhodedov AAAI 2005; Operations Research 2015]
- Our guarantees
  - No assumptions on valuation functions
  - Fine tuned to structure in the set of bidders (other results are worst case)

Outline

Part 1

Seller faces a population of buyers. Can he "learn within an instance" to find a high-revenue auction?

Part 2

Seller faces a known population of buyers. Can he learn an auction that extracts high revenue from a shrinking market?

Examples of shrinking markets

Cord cutters

NETFLIX hulu
amazon Prime HBO NOW

Retail stores

Labor markets among a shrinking population

www.ere.net › labor-market-where-is-everybody-the-sh...

Labor Market: Where Is Everybody? (The Shrinking Labor ...

Sep 24, 2020 — Simply put, the **labor force** participation rate has been falling. The rate for men has been trending downward for nearly 60 years, from 86.7% in ...

www.epi.org › news › shrinking-labor-force-explains-d...

Shrinking labor force explains drop in unemployment

In her analysis of the report, labor economist Heidi Shierholz explained that most of that decline can be explained by the drop in the **labor force** participation rate ...

www.wsj.com › articles › covid-shrinks-the-labor-market-...

Covid Shrinks the Labor Market, Pushing Out Women and ...

Dec 3, 2020 — Nearly four million Americans have stopped working or looking for jobs, a 2.2% contraction of the U.S. work force. A smaller **labor market** leaves ...

Modeling a shrinking market

- Fixed set $S = \{v_1, \dots, v_n\}$ of bidder valuations
- Seller knows S
- Each bidder in S shows up independently with probability p

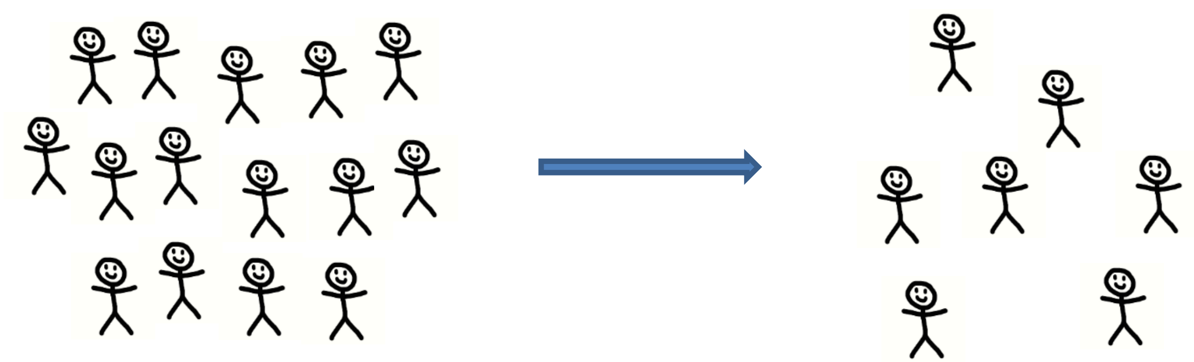

S $\longrightarrow$ S_0

What fraction of revenue can the seller guarantee?

$$\sup_M \mathbf{E}[\mathrm{Rev}_M(S_0)] \geq (???) \cdot W(S)$$

Revenue loss can be drastic

- At first glance answer might appear to be p (or even higher, if revenue thought to have diminishing returns in number of buyers)

- Example 1: $\mathbf{E}[\mathrm{Rev}_{VCG}(S_0)] = p^2 \mathrm{Rev}_{VCG}(S) = p^2(W(S) - \varepsilon)$
  - Due to reduced competition among buyers

c	0	0
$c - \varepsilon/m$	0	0
0	c	0
0	$c - \varepsilon/m$	0
0	0	c
0	0	$c - \varepsilon/m$

VCG gets payment of $c - \varepsilon/m$ for each item so $\mathrm{Rev}_{VCG}(S) = mc - \varepsilon = W(S) - \varepsilon$

But

$$\mathbf{E}[\mathrm{Rev}_{VCG}(S_0)] = \sum_{\text{item } i} \mathbf{E}[\text{Rev from item } i]$$

$$= p^2(mc - \varepsilon)$$

Revenue loss can be drastic

If valuations can depend on what other bidders receive, things are even worse

Theorem (Balcan, **Prasad**, Sandholm NeurIPS'22). For any $\varepsilon > 0$ there exists a set S of bidders with allocational valuations such that

$$\sup \mathbf{E}[\mathrm{Rev}_M(S_0)] \leq p^{m/2} \cdot (\mathrm{Rev}_{VCG}(S) + 2\varepsilon) + \varepsilon$$

where the supremum is over all possible auctions M.

Escaping large revenue loss

Enabled by two main assumptions:

- *Winner monotonicity*
  - if bidder i wins in VCG, and j leaves, i still wins in VCG
- *Welfare submodularity*
  - efficient welfare a submodular function

e.g. bidders with gross-substitutes valuations

How much revenue can be preserved?

General possibility result: rich enough set of mechanisms always contains one robust to shrinkage

Theorem (Balcan, **Prasad**, Sandholm NeurIPS'22). Exists auction *M* s.t.

$$\mathbf{E}[\mathrm{Rev}_M(S_0)] \geq \Omega\left(\frac{p^2}{k^{1+\log_{1/\gamma}(4/p)}}\right) \cdot W(S)$$

γ a constant depending on S, $k \approx$ max number of winners in VCG

A shrinkage-robust auction can be computed by sampling simulated shrunken markets and maximizing empirical revenue