Optimal Auctions through Deep Learning

Zhe Feng
Harvard SEAS

Joint work with
Paul Dütting (LSE), Harikrishna Narasimhan (Google), David C. Parkes (Harvard), Sai Srivatsa Ravindranath (Harvard)

Auction 101

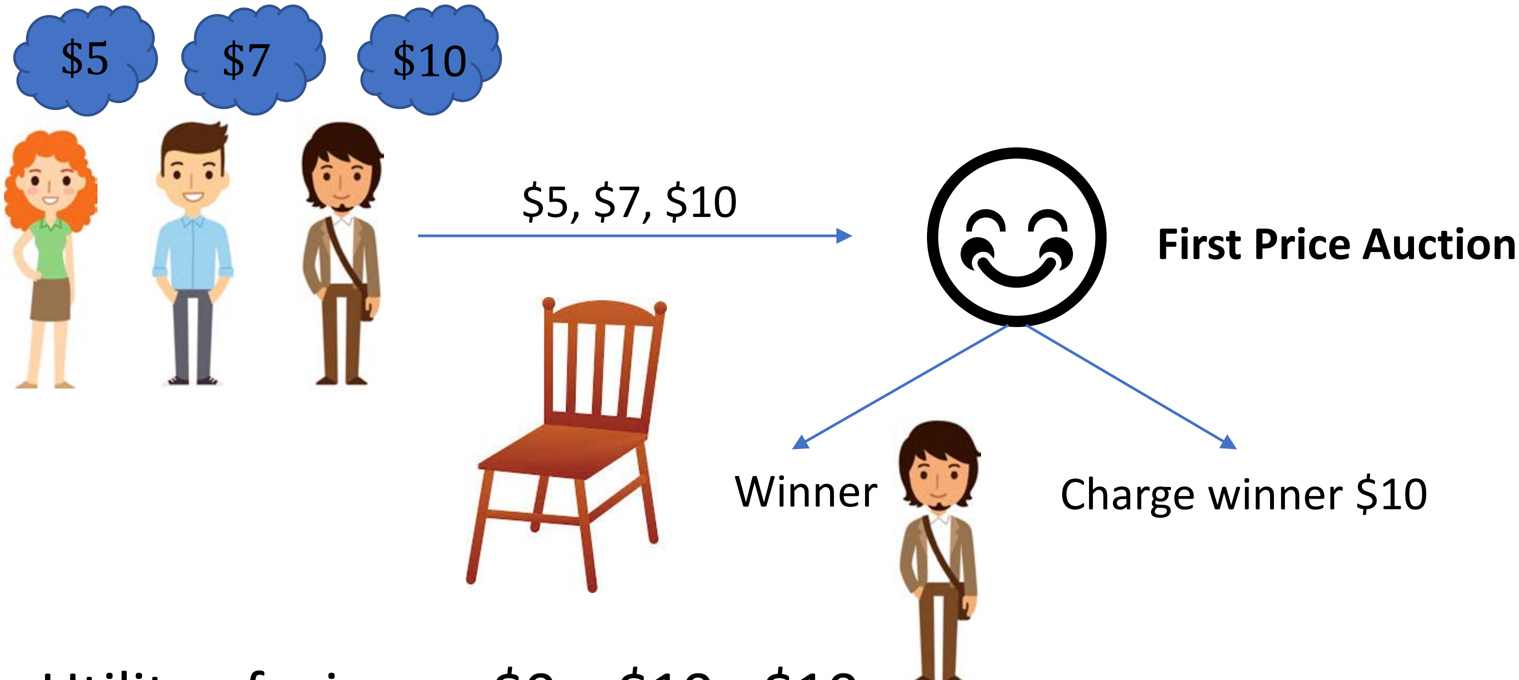

- Utility of winner: $0 = $10 - $10

Auction 101

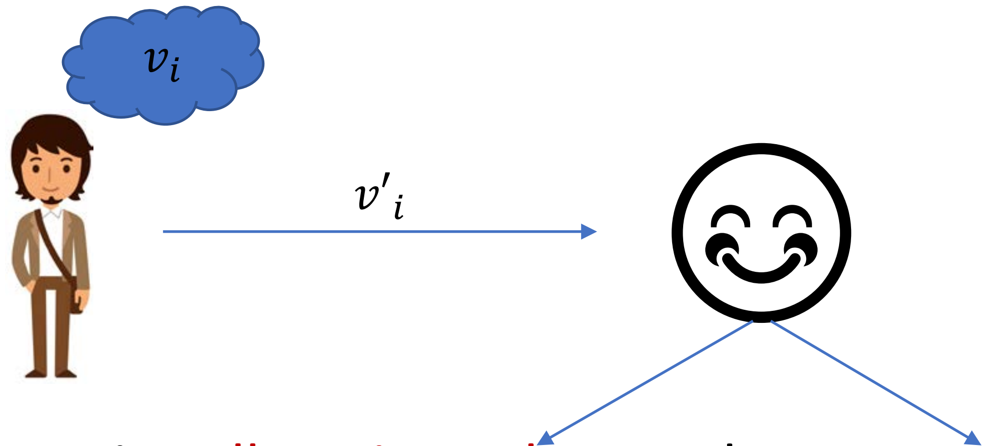

- Auction contains allocation rule g and payment rule p
- Utility function (quasilinear):

$$u_i(v'_i, v_{-i}) = v_i \cdot g_i(v'_i, v_{-i}) - p_i(v'_i, v_{-i})$$

True value Bid Others' bids

Auction 101

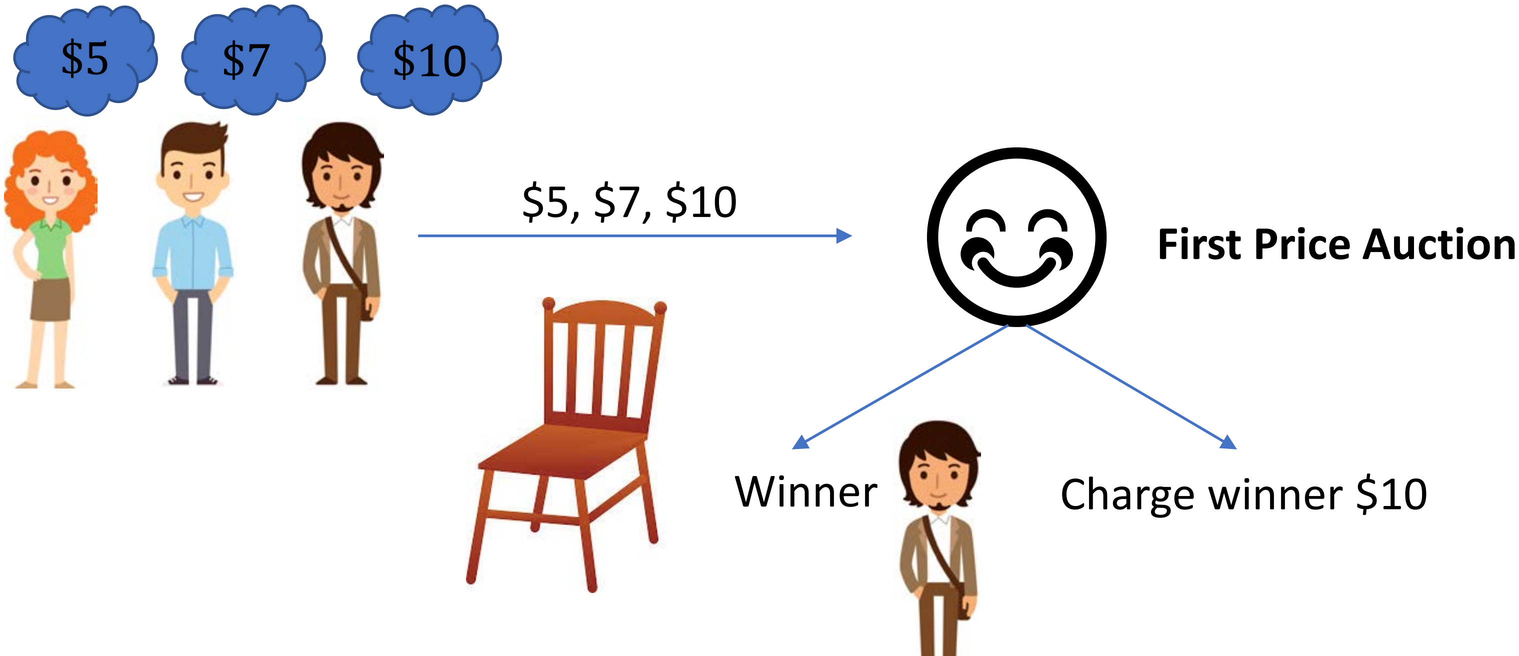

- Utility of winner: $0 = $10 - $10

Auction 101

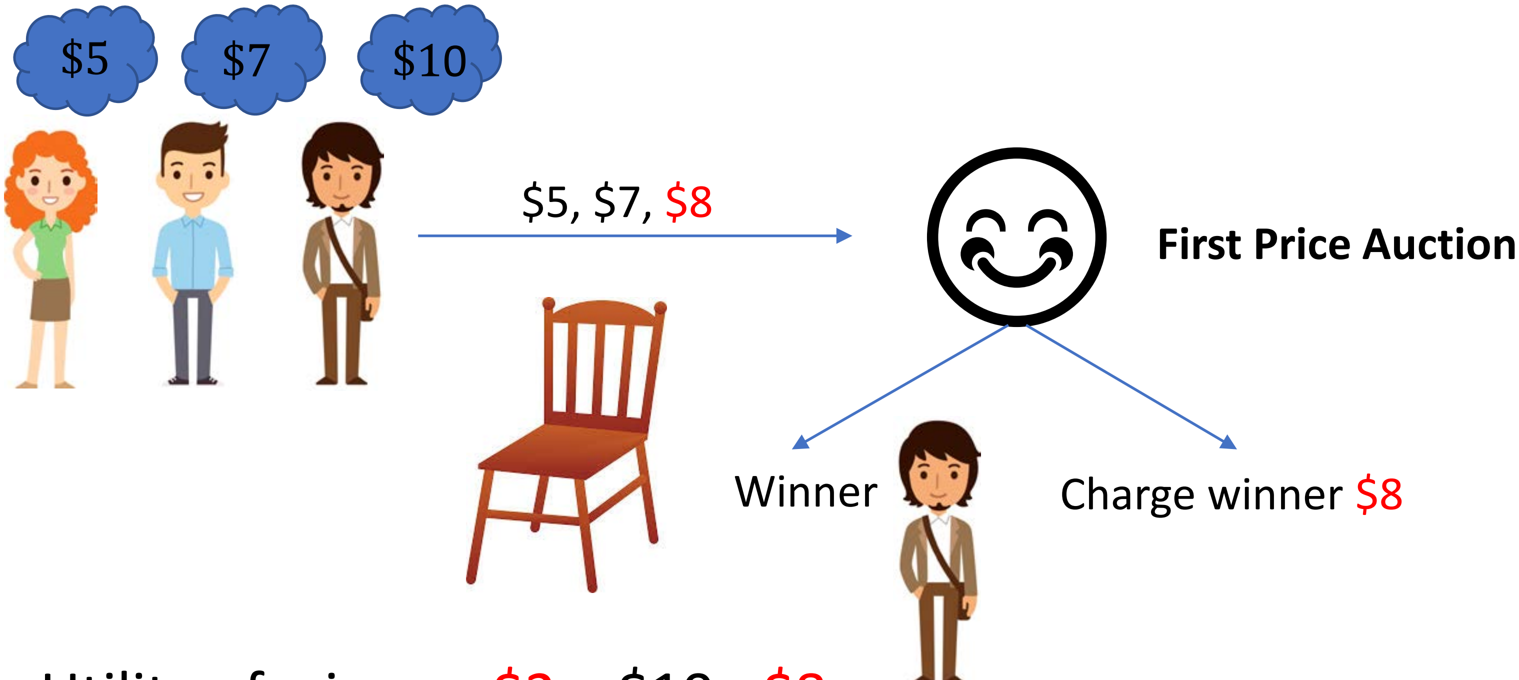

- Utility of winner: $2 = $10 - $8

Auction 101

- Incentive Constraints
  - Dominate Strategy IC (Strategy Proof): no matter what the other bidders report, truth-telling is always the weakly dominant strategy for this bidder.
    $\forall i, v_i, v_i', v_{-i}, u_i(v_i, v_{-i}) \geq u_i(v_i', v_{-i})$.
  - Individual Rationality: $u_i(v_i, v_{-i}) \geq 0$, for all v and i.
- Maximize Expected Revenue
  - $E_v[\sum_i p_i\,(v_i, v_{-i})]$

Optimal Auction Design

- Q: How to sell one item to maximize revenue?

Optimal Auction Design

- Q: How to sell one item to maximize revenue?
- A: Myerson Auction (Myerson’81)!

Optimal Auction Design

- Q: How to sell one item to maximize revenue?
- A: Myerson Auction (Myerson'81)!
- Q: How to sell two items to maximize revenue?

Optimal Auction Design

- Q: How to sell one item to maximize revenue?
- A: Myerson Auction (Myerson’81)!
- Q: How to sell two items to maximize revenue?
- A: No complete analytical understanding! 😢

Optimal Auction Design: Special cases

- (one additive buyer, two items) [Manelli & Vincent'06], [Haghpanah and Hartline'15], [Giannakopoulos and Koutsoupias'15,] [Daskalakis et al.'16]

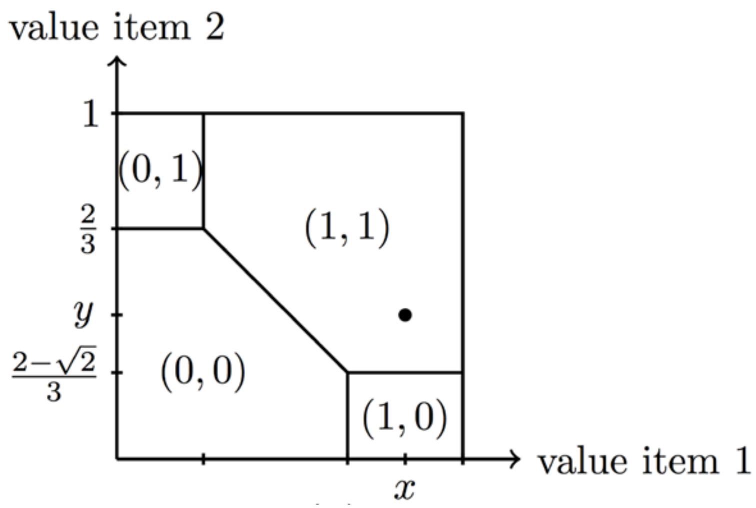

Optimal Auction Design: Special cases

- (one additive buyer, two items) [Manelli & Vincent'06], [Haghpanah and Hartline'15], [Giannakopoulos and Koutsoupias'15,] [Daskalakis et al.'16]
- (one unit demand buyer, two items) [Pavlov'11]

Optimal Auction Design: Special cases

- (one additive buyer, two items) [Manelli & Vincent'06], [Haghpanah and Hartline'15], [Giannakopoulos and Koutsoupias'15,] [Daskalakis et al.'16]
- (one unit demand buyer, two items) [Pavlov'11]
- (two items, $\geq$ 2 bidders, support of size two) [Yao'17]

Dominant-Strategy versus Bayesian Multi-item Auctions:
Maximum Revenue Determination and Comparison

ANDREW CHI-CHIH YAO, Tsinghua University

We address two related unanswered questions in maximum revenue multi-item auctions. Is dominant-strategy implementation equivalent to the semantically less stringent Bayesian one (as in the case of Myerson's 1-item auction)? Can one find explicit solutions for non-trivial families of multi-item auctions (as in the 1-item case)? In this paper, we present such natural families whose explicit solutions exhibit a revenue gap between the

Our Contribution

- Initiate the use of deep learning for optimal auction design. Motivate several follow-up works, [Feng et.al.'18], [Golowich et al.'18], [Manisha et al.'18], [Shen et al.'19]...
- Recreate state of the art analytical results of optimal auctions.
- Discover new auctions for settings where optimal solution is unknown.

Builds on Automated Mechanism Design [Conitzer & Sandholm'02] and Machine Learning for mechanism design [Dütting et al.'14, Narasimhan & Parkes'16].

The Problem

- A seller with a set of m items
- A set of n bidders with independent private valuations, $v_i: 2^M \to R_{\geq 0}, v_i \sim F_i$
- $F = (F_1, \cdots, F_n)$ is known.
- Design auction (g^w, p^w) that maximizes expected revenue, s.t. strategy-proofness.
  - g^w is parametrized allocation rule
  - p^w is parametrized payment rule

Our Approach: RegretNet

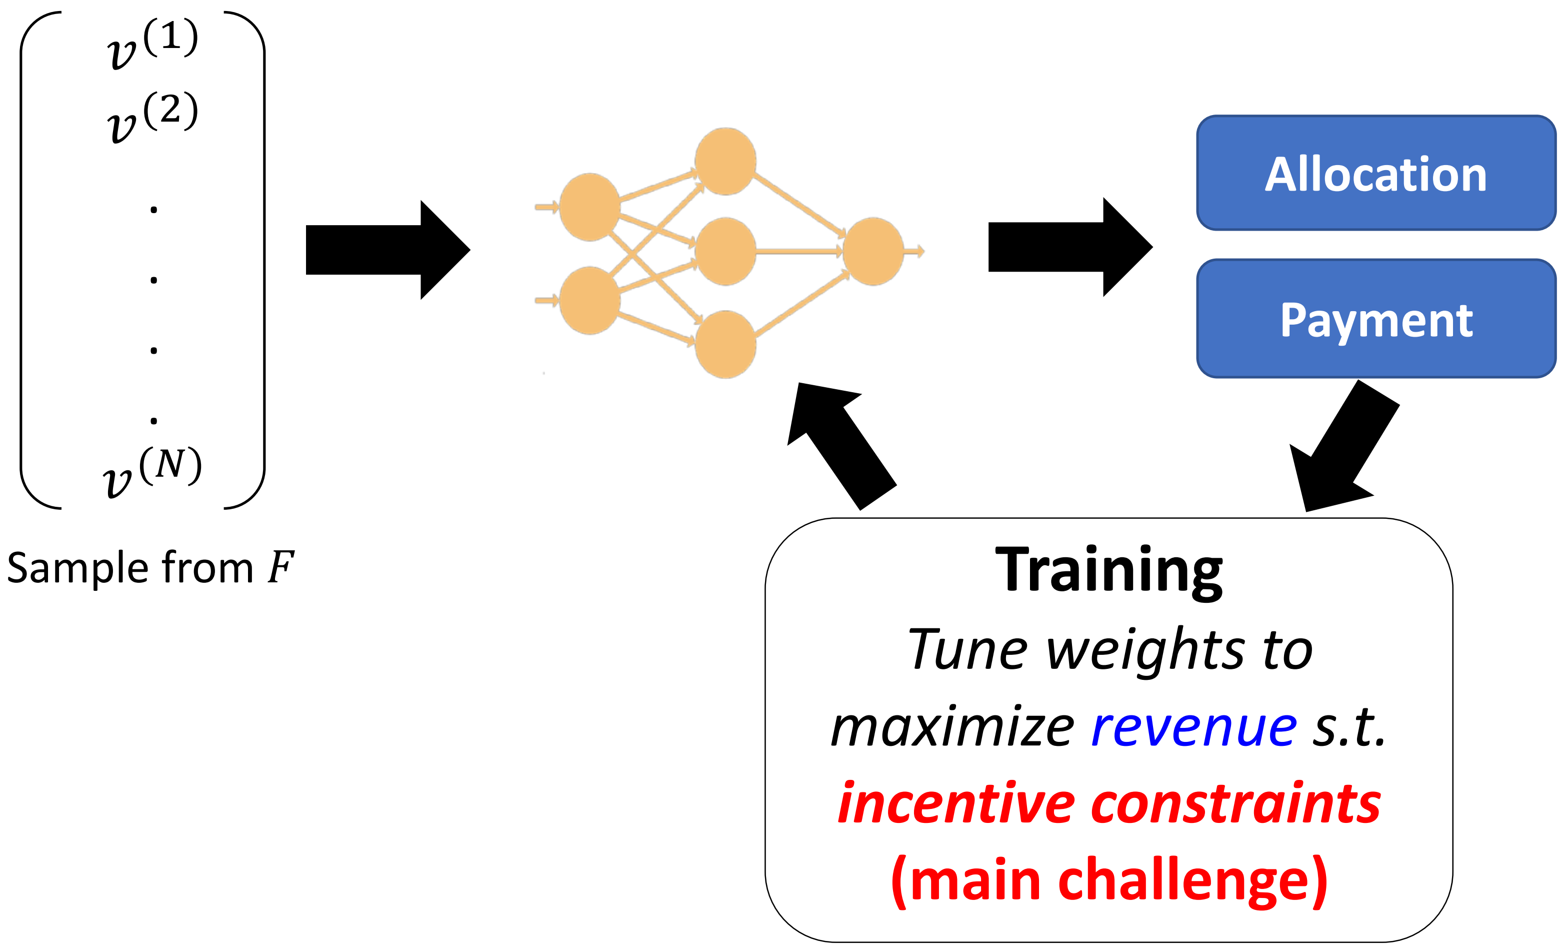

RegretNet: Regret

- What is Regret (expected ex post regret)?

$$rgt_i^w = E_v\left[\max_{v_i'} u_i^w(v_i', v_{-i}) - u_i^w(v_i, v_{-i})\right]$$

- Why Regret?
- Recap for Strategy-Proofness: Auction (g^w, p^w) is SP if $u_i^w(v_i', v_{-i}) \le u_i^w(v_i, v_{-i})$ for all i, all v, all v_i'.
- Ignoring measure zero events, Strategy-Proofness can be rewritten as: $rgt_i^w = 0$, for each bidder i.

RegretNet: Architecture

m items, n additive bidders, the bid of bidder i for item j is b_{ij}.
Parameters w.

Allocation Net

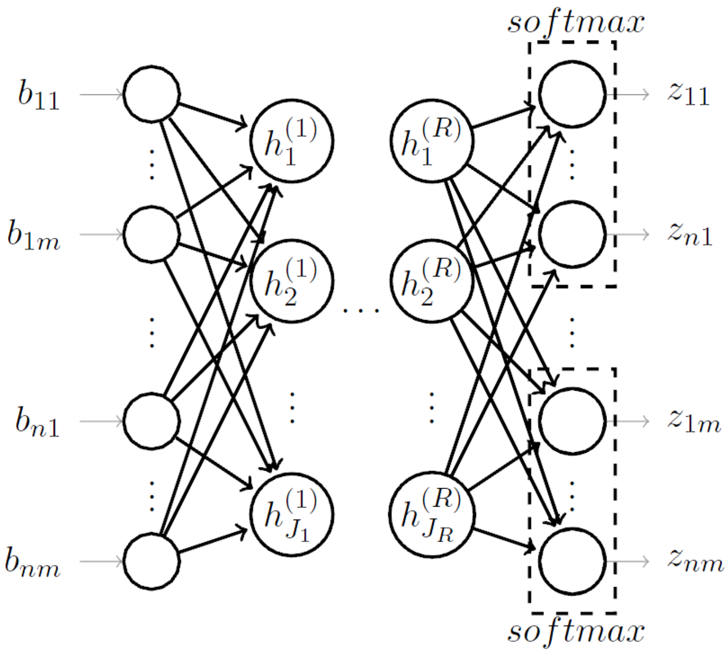

Payment Net

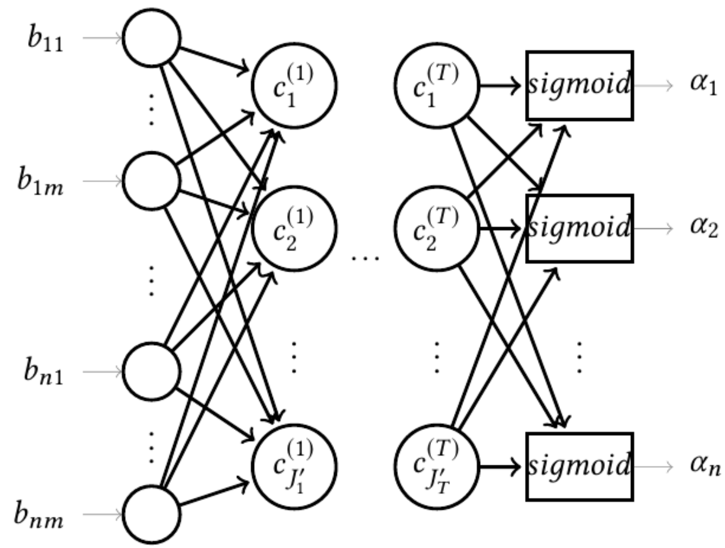

Allocation: $g^w\colon \mathbb{R}^{nm} \to \Delta_1 \times \cdots \times \Delta_m$

Payment: $p^w\colon \mathbb{R}^{nm} \to \mathbb{R}^n_{\geq 0}$

Fractional payment:

$p_i^w = \alpha_i \cdot (g_i^w \cdot v_i), \alpha_i \in [0,1]$

RegretNet: Training Problem

- Training problem:

$$\min_w \mathcal{L}(g^w, p^w) = -E_v[\sum_i p_i^w(v)]$$

$$\text{s.t. } \forall i \in [n],\ rgt_i^w = 0$$

- Train via augmented Lagrangian Method to handle regret constraints

$$w_t := \underset{w}{\operatorname{argmin}}[\mathcal{L}(g^w, p^w) + \sum_i \lambda_i^{t-1} \cdot rgt_i^w + \frac{\rho}{2}\sum_i (rgt_i^w)^2]$$

$$\forall i \in [n], \qquad \lambda_i^t := \lambda_i^{t-1} + \rho \cdot rgt_i^{w_t}$$

- Adaptively tune Lagrange multiplier. In our case, λ_i will always increase.

RegretNet: Inner Optimization

$$\underset{w}{\operatorname{argmin}}[\mathcal{L}(g^w, p^w) + \sum_i \lambda_i^{t-1} \cdot rgt_i^w + \frac{\rho}{2}\sum_i (rgt_i^w)^2]$$

- Use stochastic gradient descent (SGD)
- For each sample $v \sim F$
  - Loss: $-\sum_i p_i^w(v)$
  - Regret for bidder i:

$$rgt_i^w(v) = \max_{v_i'} u_i^w(v_i', v_{-i}) - u_i^w(v_i, v_{-i})$$

  - Adversarial approach: run gradient ascent to find optimal misreport v_i', given v and NN.

Theoretical Guarantee

- We show generalization bounds for both revenue and regret.
- Main challenge: The non-standard “max” structure in the regret .
- We measure the capacity of an auction class using a $L_{1,\infty}$ covering number.
- Bound the covering number for the NN architectures that we propose in the paper

Theoretical Guarantee

Thm 1 [Informal]. Fix $\delta \in (0,1)$, with high probability over sample of L profiles,

$$E_v\left[\sum_i p_i^w(v)\right] \geq \frac{1}{L}\sum_\ell \sum_i p_i^w\left(v^{(\ell)}\right) - 2\mathrm{n}\Delta_L - \tilde{O}(n\sqrt{1/L})$$

$$\frac{1}{n}\sum_{i=1}^{n} rgt_i^w \leq \frac{1}{n}\sum_{i=1}^{n} \widehat{rgt}_i^w + 2\Delta_L + \tilde{O}(\sqrt{1/L})$$

Theoretical Guarantee

Thm 1 [Informal]. Fix $\delta \in (0,1)$, with high probability over sample of L profiles,

$$E_v\left[\sum_i p_i^w(v)\right] \geq \frac{1}{L}\sum_\ell \sum_i p_i^w\left(v^{(\ell)}\right) - 2\mathrm{n}\Delta_L - \tilde{O}(n\sqrt{1/L})$$

$$\frac{1}{n}\sum_{i=1}^{n} rgt_i^w \leq \frac{1}{n}\sum_{i=1}^{n} \widehat{rgt}_i^w + 2\Delta_L + \tilde{O}(\sqrt{1/L})$$

Thm 2 [Informal]. Δ_L can be bounded as a function of $L, m, n,$ and # parameters in NN. $\Delta_L = O\left(\sqrt{\frac{log\ L}{L}}\right)$

Experiments

- TensorFlow library, Adam solver, NVDIA GPU core
- Learning rate 0.001 (fixed), min-batch size 128, parameter ρ initialized to 1.0, λ updates every 100 minibatches.
- Train on 640,000 valuation profiles, test on 10,000
- 2 hidden layers for smaller settings, 5 hidden layers for larger settings.

Experiments

Can RegretNet recover known auction designs?

2-item 1-buyer (additive)

- $v_1, v_2 \sim U[0,1]$, [Manelli & Vincent'06]

	Revenue	Regret
Optimal	0.550	-
RegretNet	**0.554**	<0.001

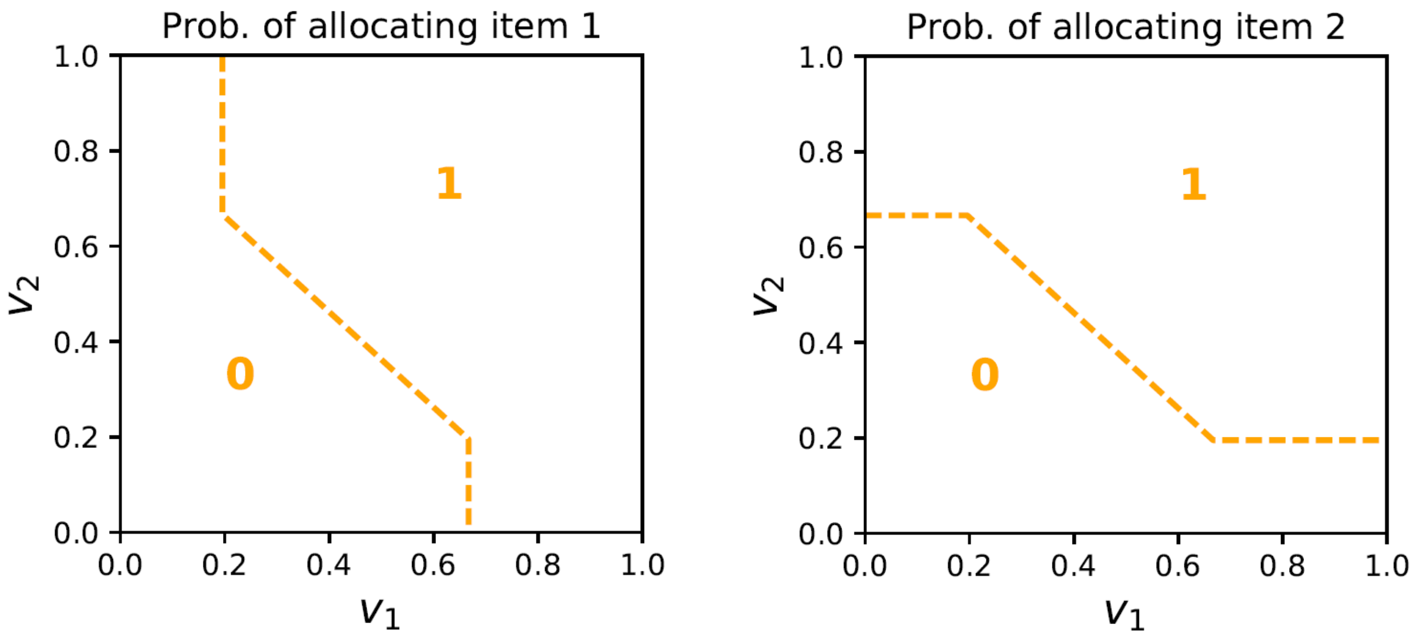

2-item 1-buyer (additive)

- $v_1, v_2 \sim U[0,1]$, [Manelli & Vincent'06]

	Revenue	Regret
Optimal	0.550	-
RegretNet	**0.554**	<0.001

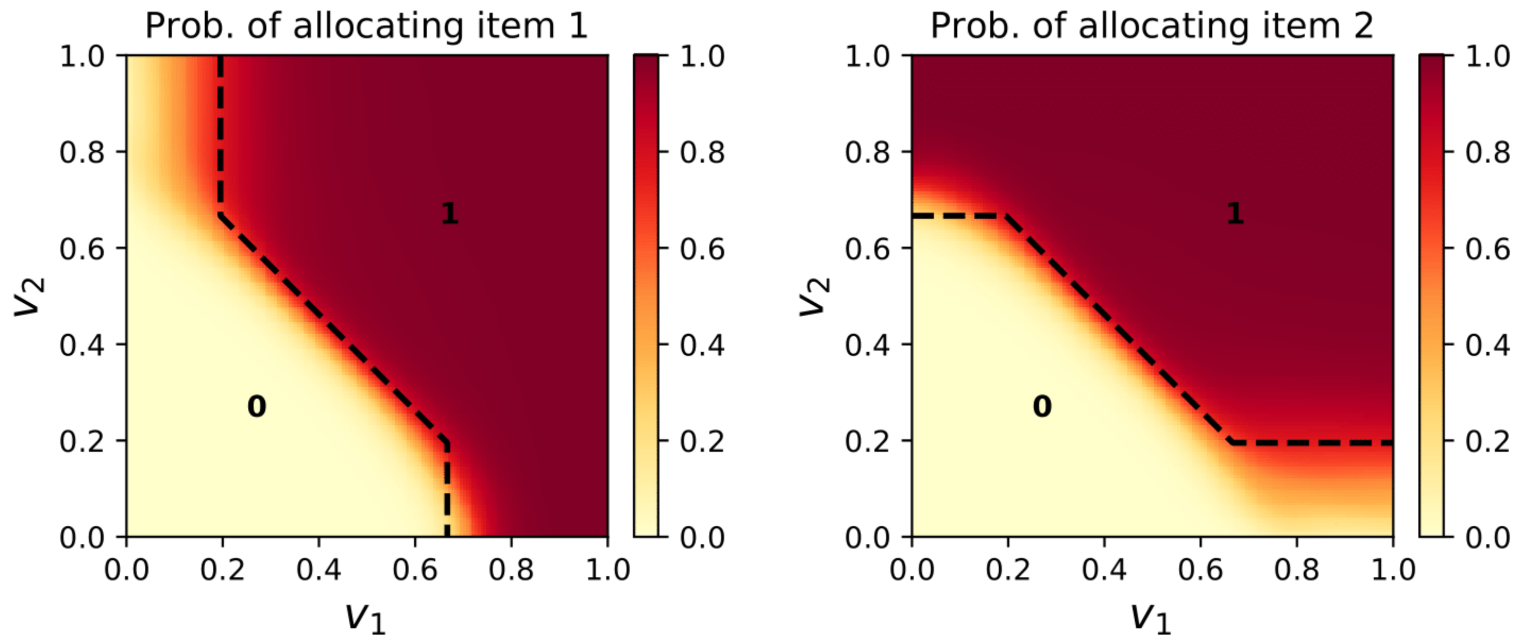

2-item 1-buyer (unit-demand)

- $v_1, v_2 \sim U[2,3]$, [Pavlov'11]

	Revenue	Regret
Optimal	2.137	-
RegretNet	2.137	<0.001

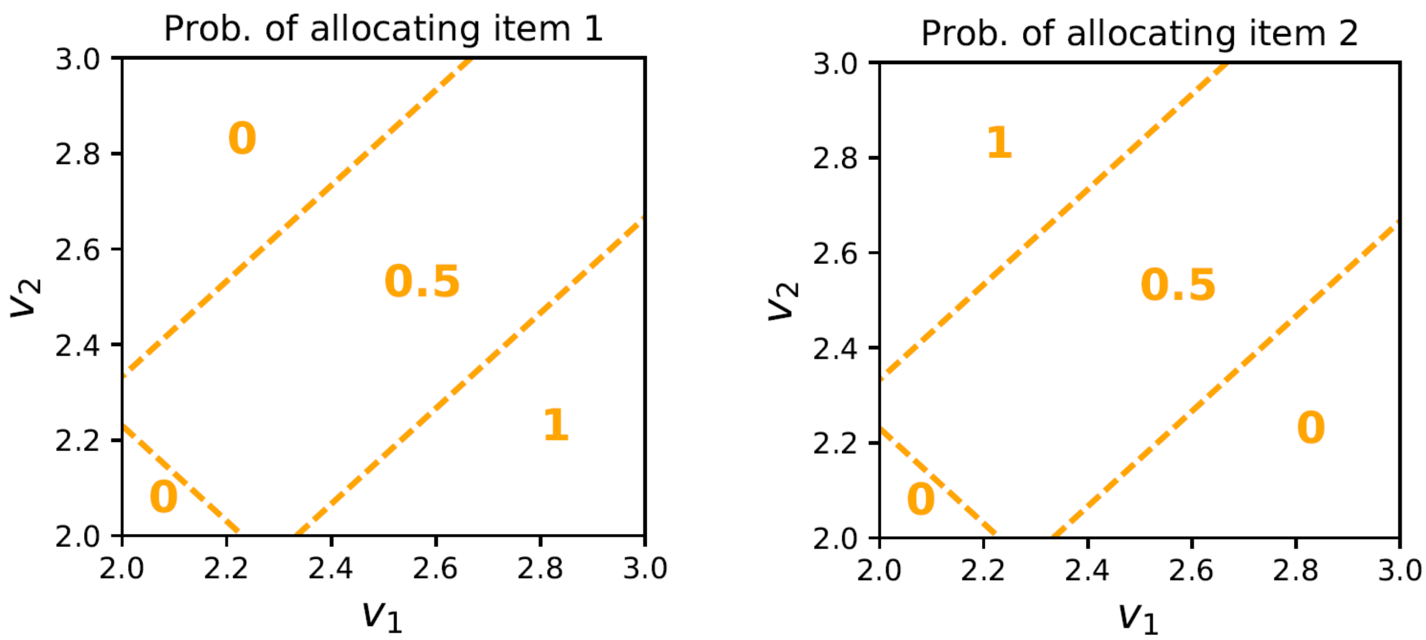

2-item 1-buyer (unit-demand)

- $v_1, v_2 \sim U[2,3]$, [Pavlov'11]

	Revenue	Regret
Optimal	2.137	-
RegretNet	2.137	<0.001

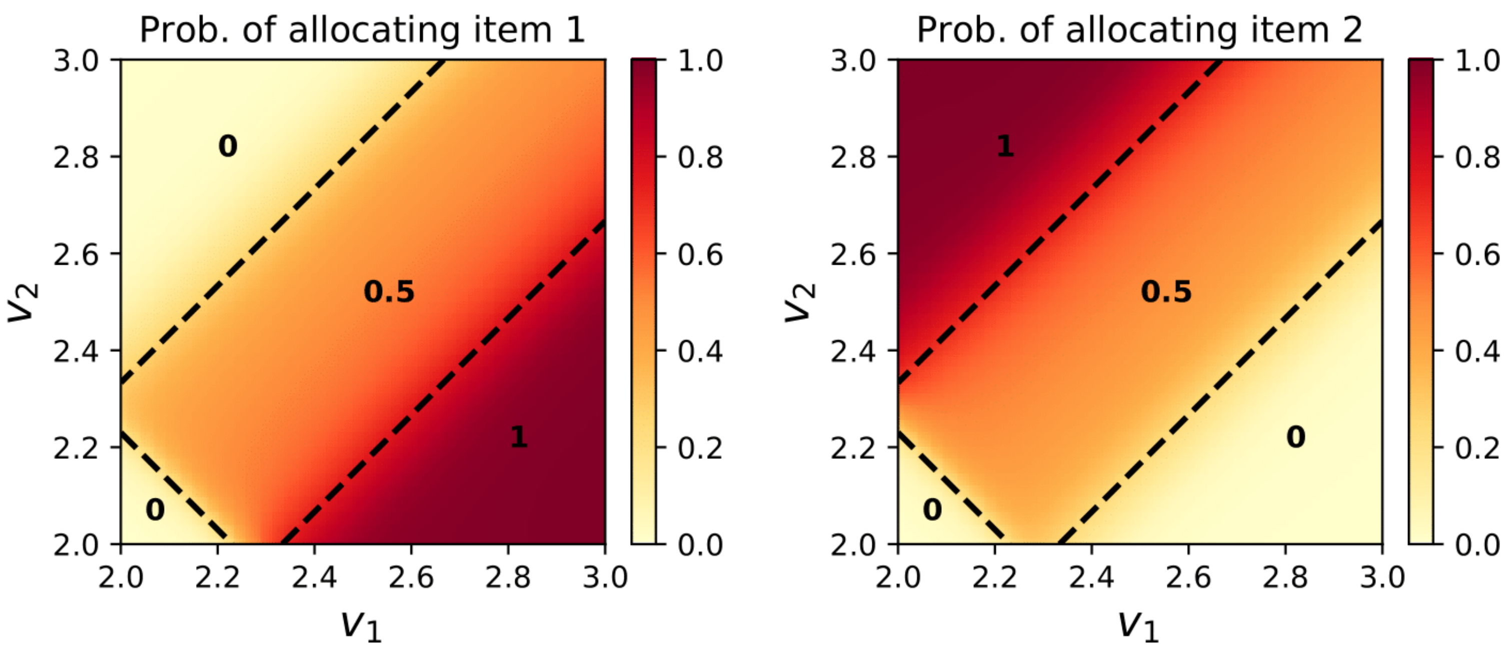

Experiments

Can RegretNet discover new auctions?

2-item 2-bidder, 3 support

- Extends [Yao'17]
- For each bidder i:
  - Additive valuation
  - $v_{i,1} \sim unif\{0.5,1.0,1.5\}, v_{i,2} \sim unif\{0.5,1.0,1.5\}$
- Experiments

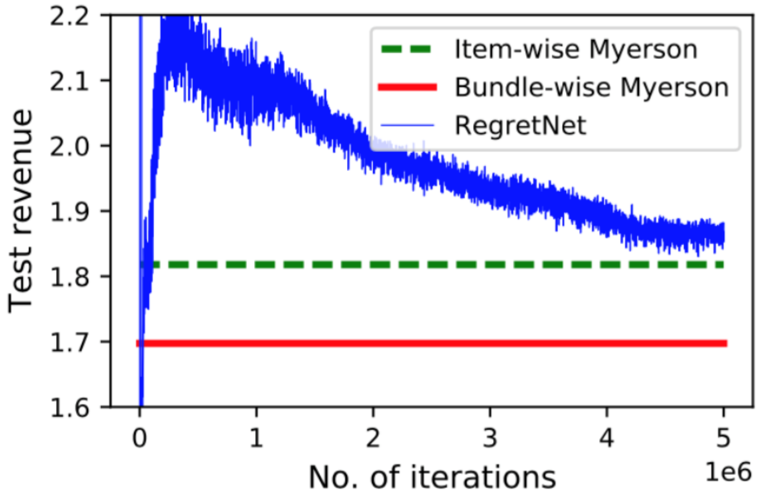

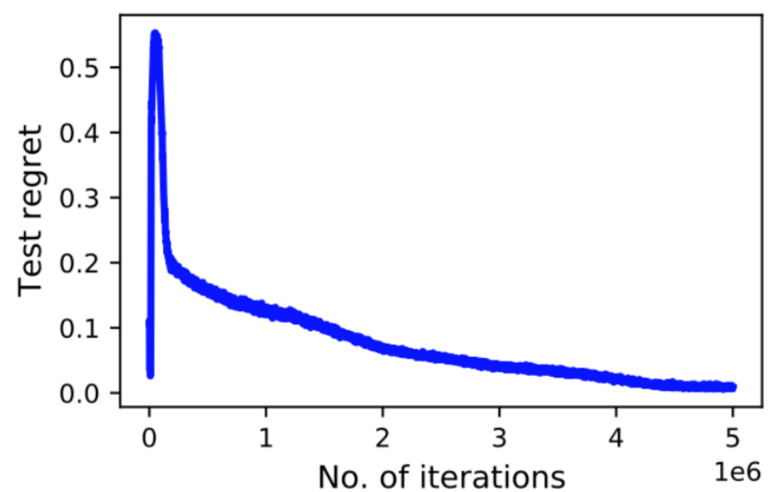

2-item 2-bidder, uniform values

- Compare to [Likhodedov and Sandholm'15].
- Additive valuation:
  - $v_{1,1}, v_{1,2} \sim U[0,1], v_{2,1}, v_{2,2} \sim U[0,1]$
- Experiments

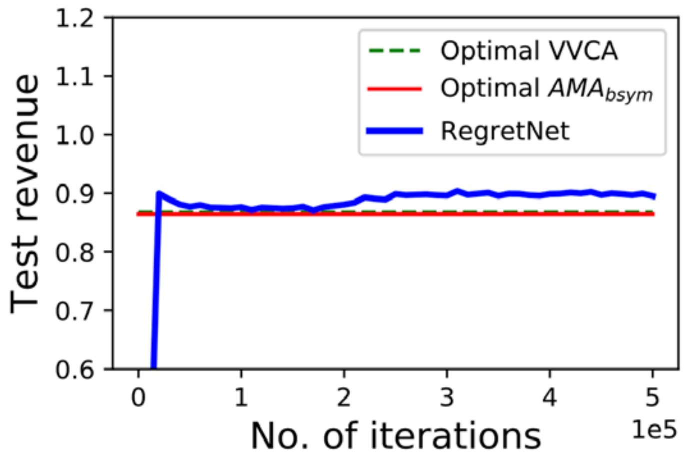

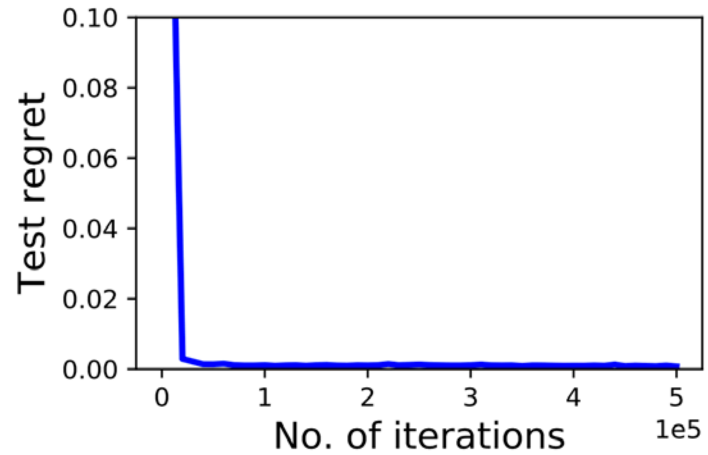

- In paper: combinatorial settings for 2-item, 2-bidder

Scaling up

- 3-bidder, 10-item, $v_{ij} \sim U[0,1]$
- 5-bidder, 10-item, $v_{ij} \sim U[0,1]$
- Experiments

Distribution	RegretNet		Item-wise Myerson	Bundle-wise Myerson
	Revenue	Regret	Revenue	Revenue
3 bidders, 10 items	**5.541**	<0.002	5.310	5.009
5 bidders, 10 items	**6.778**	**<0.005**	**6.716**	**5.453**

- Less than 13 hours training time, in contrast, LPs takes more than a week even for a 2-bidder, 3-item setting.

Future Work

- Scaling up
  - Universal network for different number of buyers and items.
  - Leverage economic structural results
- Guide economic theory: reveal gaps, test conjectures (e.g. “Pentagon conjecture” by Daskalakis et.al 13, “Revenue Monotonicity”)
- Other settings: stability, fairness, group-SP

Resource

- Poster: Pacific Ballroom #155, 6:30pm today!

- Full version: **arXiv:**1706.03459, for more experiments, theoretical analysis, and architectures (MyersonNet, RochetNet).

Thanks!