Online Mechanism Design with Predictions

To appear at **EC 2024**
(exemplary track award for theory track)

Vasilis Gkatzelis
Drexel University and Google Research

Eric Balkanski
Columbia University

Xizhi Tan
Drexel University

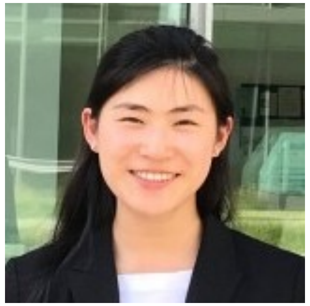

Cherlin Zhu
Columbia University

Learning-Augmented Algorithms

- **Tension** between classic analysis of algorithms and machine learning:

Worst-Case Analysis of Algorithms

Machine Learning Algorithms

- **Worst-case analysis** provides **robust guarantees**, but often **too pessimistic**
- **Machine learning** algorithms **work well**, but **lack robustness**

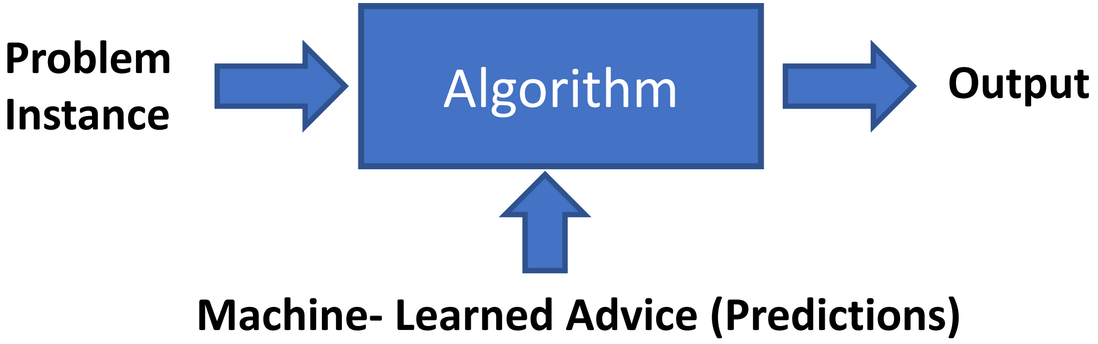

Learning-Augmented Algorithms

- Ideal algorithm with predictions:
  - Achieve **optimal** performance guarantees when predictions are accurate, **without sacrificing worst-case guarantees** when they are arbitrarily bad
- Framework originally proposed by Mahdian, Nazerzadeh, and Saberi [EC '07]
- Evaluation measures proposed by Lykouris and Vassilvitskii [ICML '18, JACM '21]:
  - **Robustness**: worst-case performance guarantee
  - **Consistency**: worst-case performance for instances with **accurate prediction**
- This provides a natural **refinement** of worst-case analysis

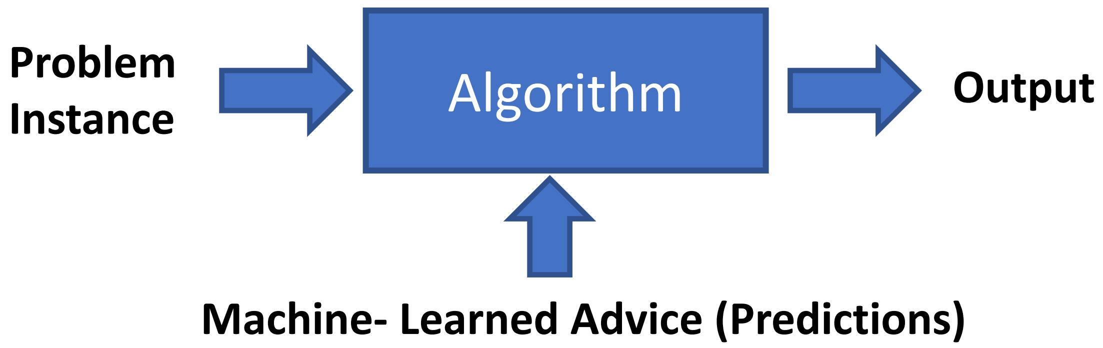

https://algorithms-with-predictions.github.io/

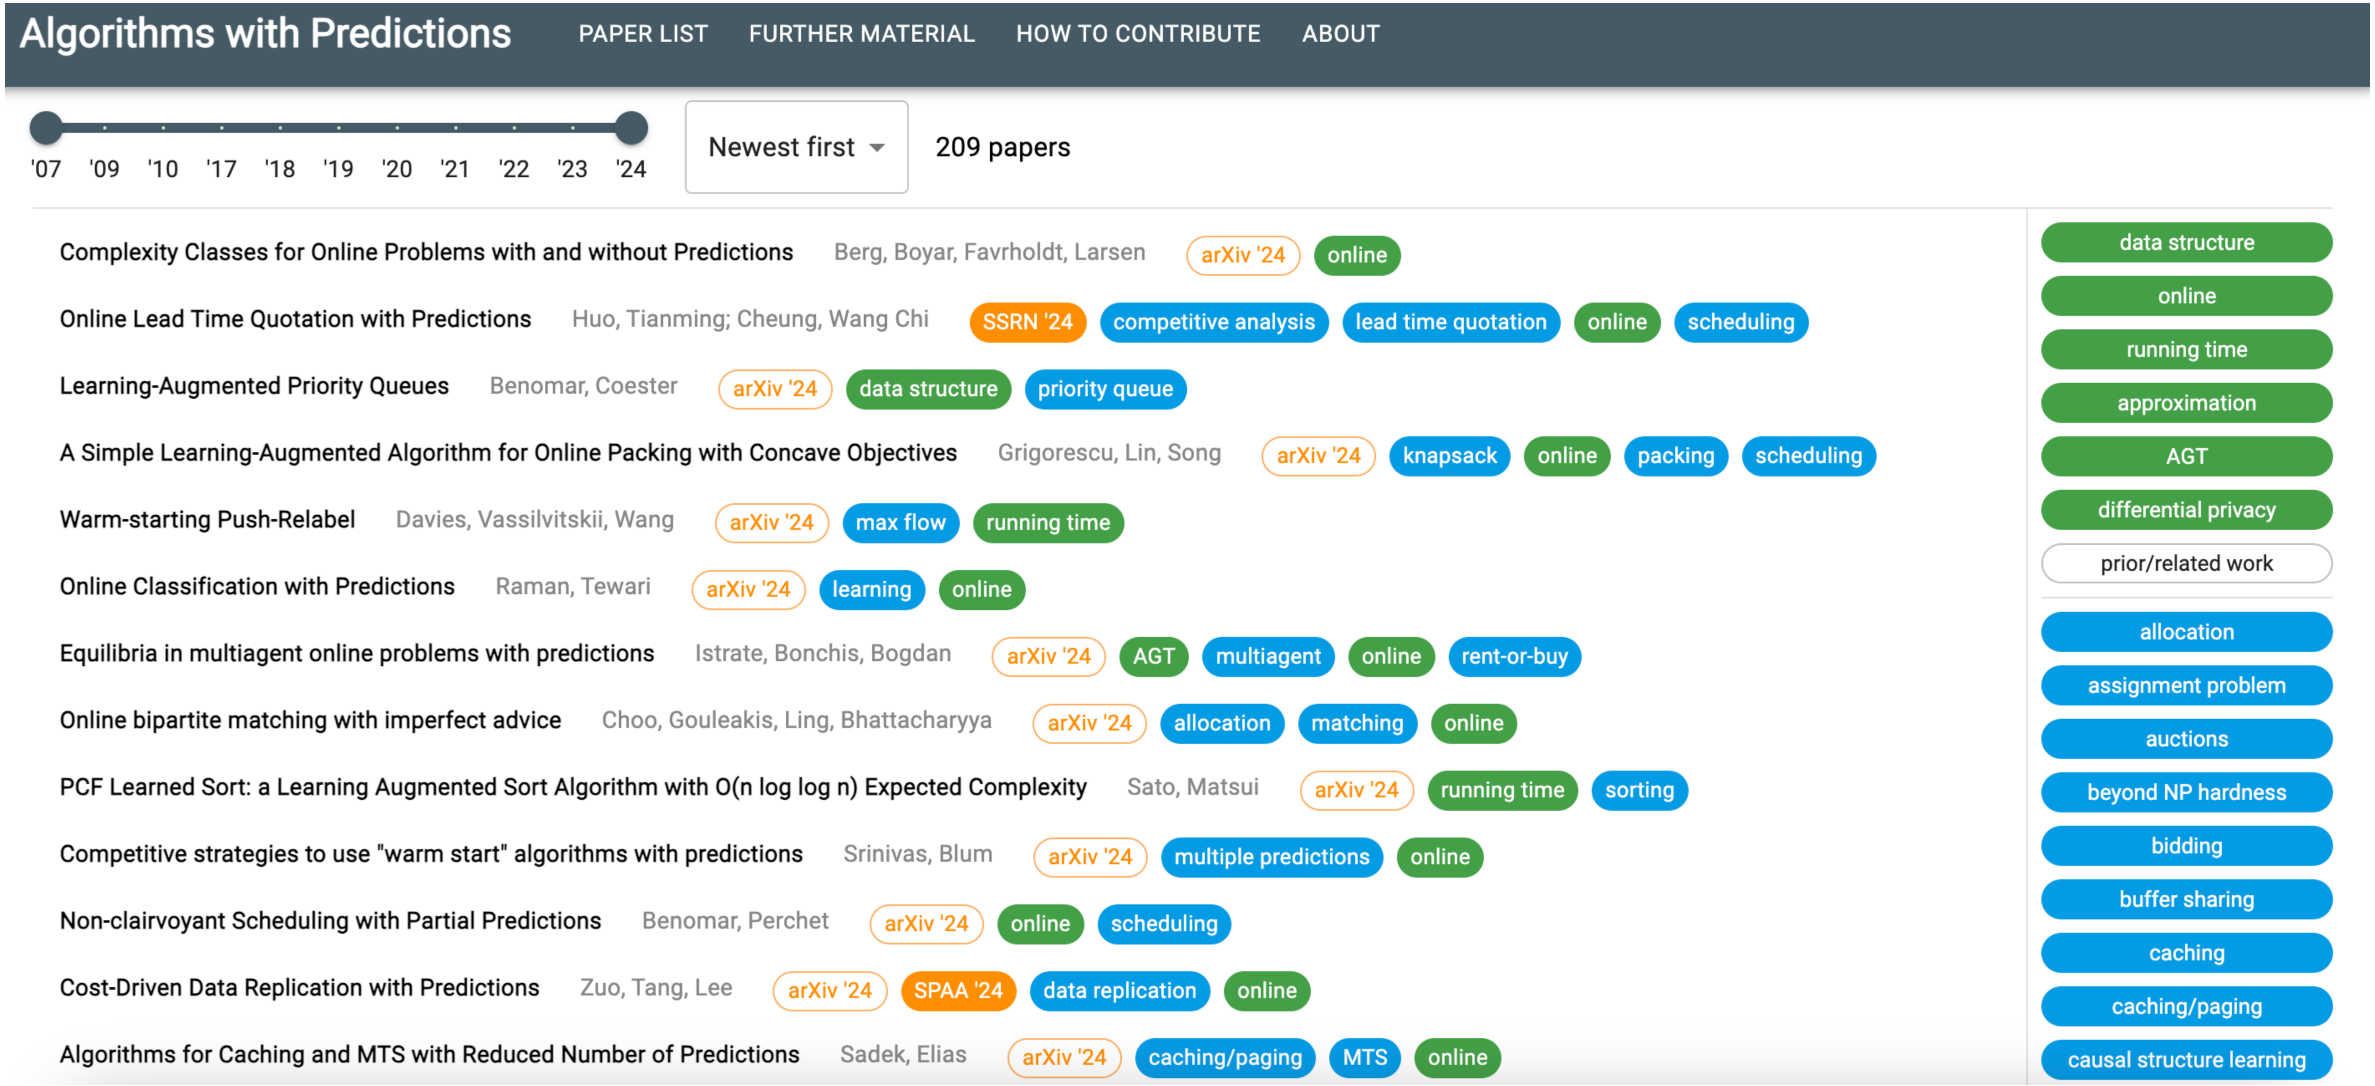

Learning-Augmented Mechanism Design

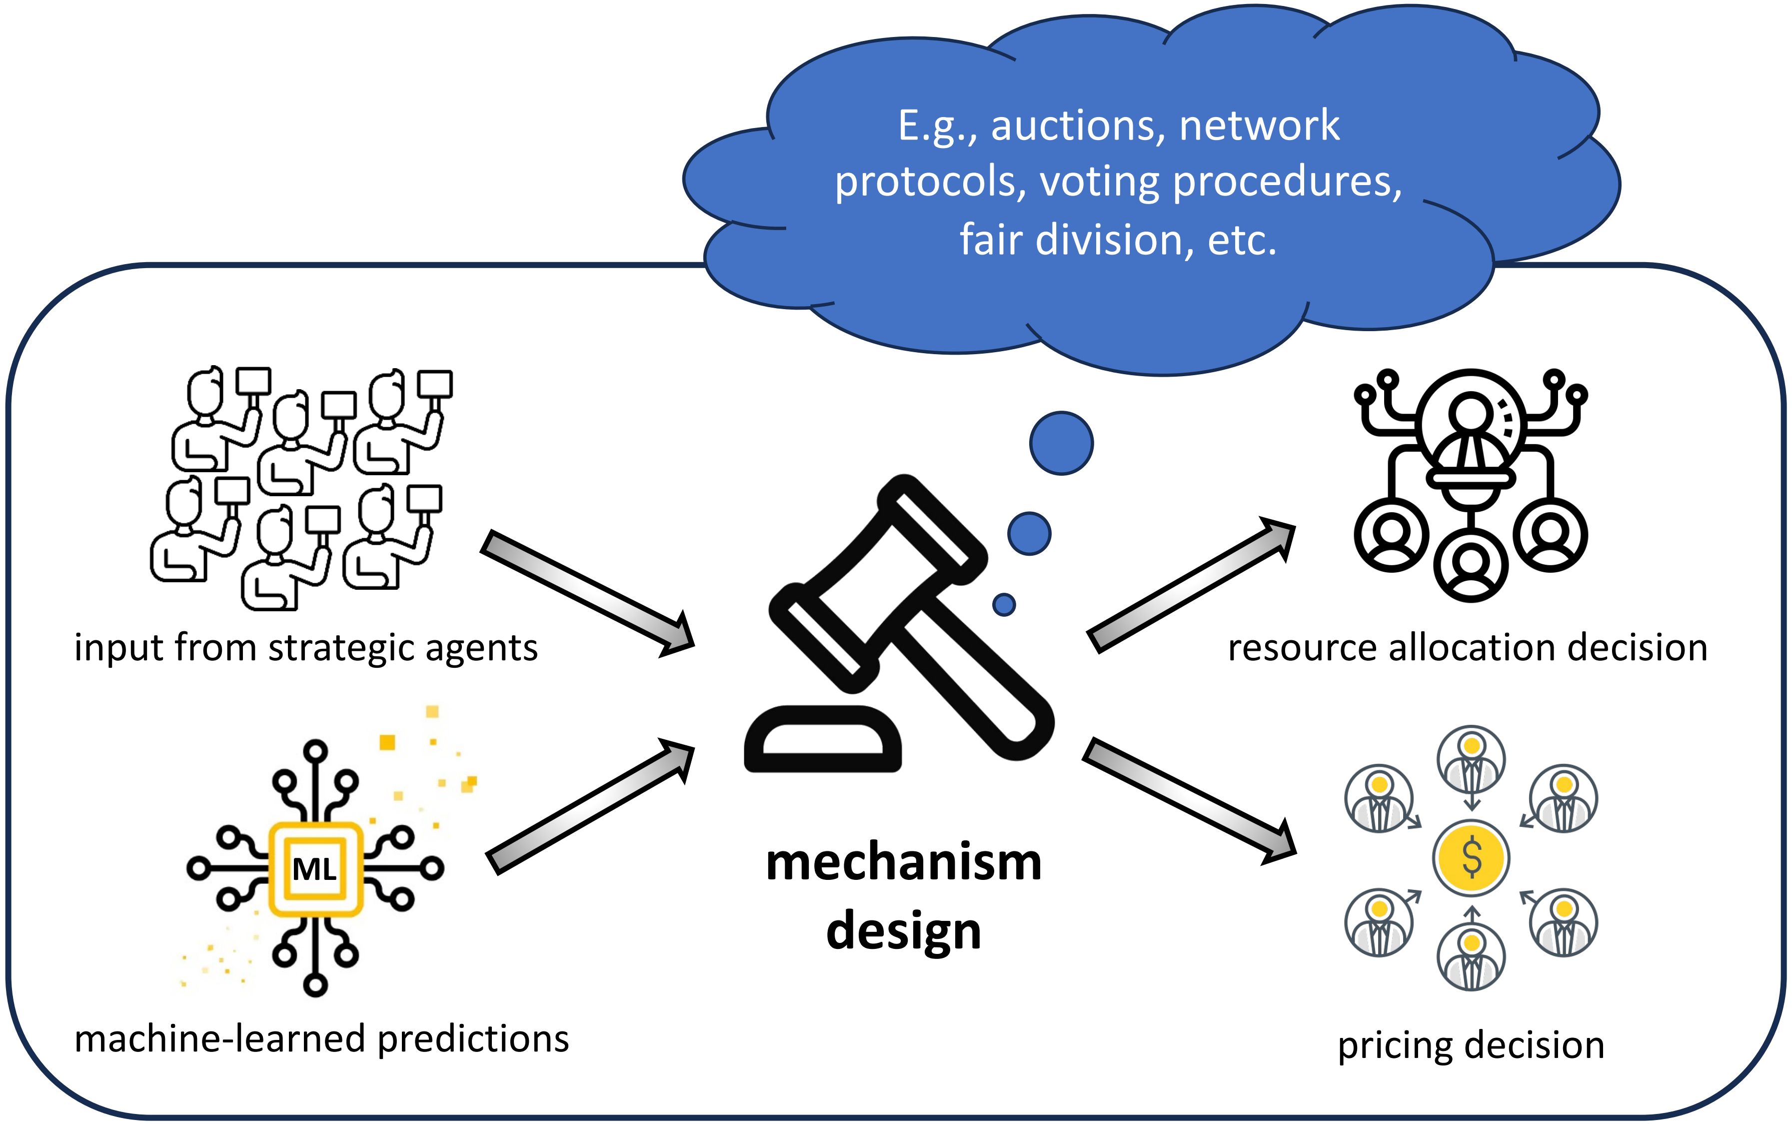

Algorithmic Game Theory papers

Algorithms with Predictions PAPER LIST FURTHER MATERIAL HOW TO CONTRIBUTE ABOUT

'07 '09 '10 '17 '18 '19 '20 '21 '22 '23 '24 | Newest first | 13 papers | RESET

Equilibria in multiagent online problems with predictions	Istrate, Bonchis, Bogdan	arXiv '24, AGT, multiagent, online, rent-or-buy
MAC Advice for Facility Location Mechanism Design	Barak, Gupta, Talgam-Cohen	arXiv '24, AGT, facility location, mechanism design
To Trust or Not to Trust: Assignment Mechanisms with Predictions in the Private Graph Model	Colini-Baldeschi, Klumper, Schäfer, Tsikiridis	arXiv '24, AGT, assignment problem, graph problems
Randomized learning-augmented auctions with revenue guarantees	Caragiannis, Kalantzis	arXiv '24, AGT, auctions, mechanism design
Online Mechanism Design with Predictions	Balkanski, Gkatzelis, Tan, Zhu	arXiv '23, AGT, auctions, mechanism design
Competitive Auctions with Imperfect Predictions	Lu, Wan, Zhang	arXiv '23, AGT, auctions
Optimal Metric Distortion with Predictions	Berger, Feldman, Gkatzelis, Tan	arXiv '23, AGT, metric distortion
Bicriteria Multidimensional Mechanism Design with Side Information	Balcan, Prasad, Sandholm	arXiv '23, NeurIPS '23, AGT, mechanism design
Mechanism Design With Predictions for Obnoxious Facility Location	Istrate, Bonchis	arXiv '22, AGT, mechanism design
Strategyproof Scheduling with Predictions	Balkanski, Gkatzelis, Tan	arXiv '22, ITCS '23, AGT, scheduling
Mechanism Design with Predictions	Xu, Lu	arXiv '22, IJCAI '22, AGT, auctions, scheduling
Improved Price of Anarchy via Predictions	Gkatzelis, Kollias, Sgouritsa, Tan	arXiv '22, EC '22, AGT
Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location	Agrawal, Balkanski, Gkatzelis, Ou, Tan	arXiv '22, EC '22, AGT, network design

data structure
online
running time
approximation
AGT ⊗
differential privacy
prior/related work

allocation
assignment problem
auctions
beyond NP hardness
bidding
buffer sharing
caching
caching/paging
causal structure learning
causality

Online Auctions for a Single Good

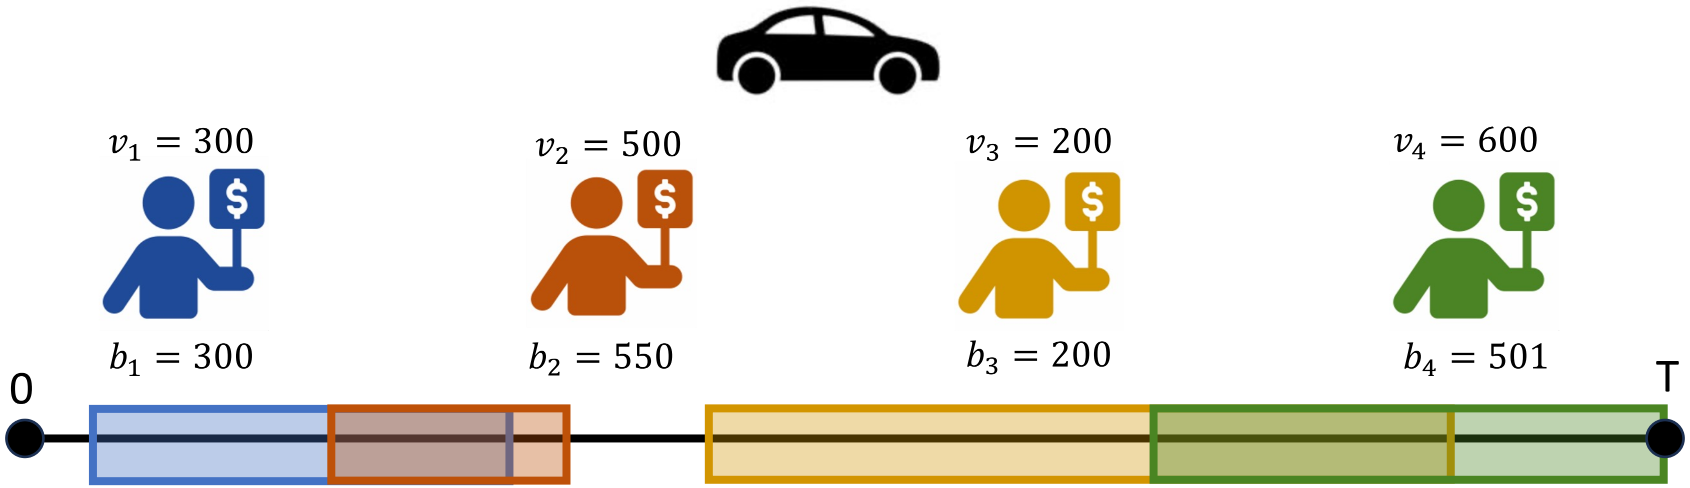

- Each bidder **announces their arrival and departure** and **reports their bid**
- A bidder can receive the good **only during their true active interval**
- Bidders can announce a **delayed arrival** time and an **earlier departure** time
- Bidders can also **arbitrarily misreport** their value when they bid
- The auctioneer must make irrevocable decisions based only on bids from agents that have already arrived, aiming to maximize **revenue**

[Hajiaghayi-Kleinberg-Parkes EC '04]

Connection to Secretary Problem

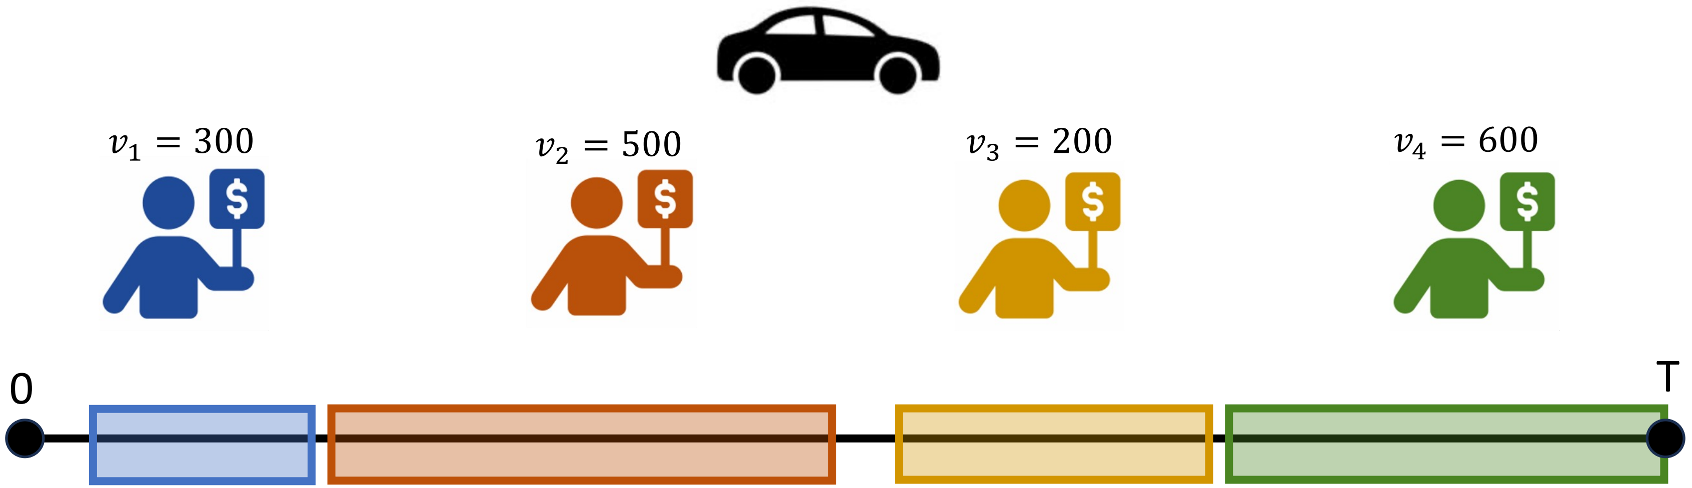

- If the **arrival-departure intervals are disjoint**, this closely resembles **secretary problem**
- The goal there is to **maximize the probability of choosing maximum value agent**
- Two crucial differences for secretary problem mechanisms:
  - The mechanism **only benefits if the highest value agent is selected**
  - The decisions of the mechanism **depend only on the ranking** of agent values
- The **design space for online auctions is richer** (so, harder to prove impossibility results)

Online Auctions for a Single Good

- The **”type”** $\boldsymbol{\theta_i}$ of each bidder i is determined by:
  - an **arrival time** $\boldsymbol{a_i}$ and **departure time** $\boldsymbol{d_i} \geq a_i$
  - a **value** $\boldsymbol{v_i}$ for the good being sold
- The **utility** of bidder i is equal to:
  - $\boldsymbol{v_i - p}$, if they receive the good at price p **within** $[\boldsymbol{a_i}, \boldsymbol{d_i}]$
  - $\leq \boldsymbol{0}$, otherwise
- The bidder can announce a **later arrival**, an **earlier departure**, and **bid** $\boldsymbol{b_i} \neq \boldsymbol{v_i}$
- **Value-strategyproofness:** it is a **dominant** strategy to report true value
- **Time-strategyproofness:** it is a **dominant** strategy to report true arrival/departure
- Adversary chooses **active intervals** $\boldsymbol{I} = \{(\boldsymbol{a_1}, \boldsymbol{d_1}), (\boldsymbol{a_2}, \boldsymbol{d_2}), \ldots (\boldsymbol{a_n}, \boldsymbol{d_n})\}$ and a **set** $\boldsymbol{V}$ of n **bidder values**. Each value is then assigned to a time interval **uniformly at random**
- Objective is to **maximize expected revenue** over the random arrival

Online Auctions for a Single Good

- In the **offline setting**, where all bidders are present at the same time:
  - it is impossible to extract a revenue approximating the **highest value, $v_{(1)}$**
  - But, second-price auction revenue is equal to the **second highest value, $v_{(2)}$**
- Can we approximate **second highest value, $v_{(2)}$**, in online setting [HKP '04]?
  - There exists a strategyproof auction that achieves a **0.25-approximation**
  - No strategyproof auction can achieve better than **0.66-approximation**
  - We prove a **tight lower bound of 0.25** for a large family of auctions
- [HKP '04] also considered **value (social welfare) maximization**, w.r.t., $v_{(1)}$
  - There exists a strategyproof auction that achieves a **$1/e$-approximation**
  - No strategyproof auction can achieve better than **0.5-approximation**
  - Correa, Duetting, Fischer, and Schewior [EC '19] recently showed **$1/e$ is tight**

Online Auctions with Predictions for a Single Good

- We are provided with a prediction , $\tilde{\boldsymbol{v}}_{(\mathbf{1})}$, regarding the highest value, $\boldsymbol{v}_{(\mathbf{1})}$
- **Goal:** design an online revenue-maximizing auction using this prediction
- An auction is ***β*-robust** if its expected revenue is always at least $\boldsymbol{\beta} \cdot \boldsymbol{v}_{(\mathbf{2})}$

$$\text{robustness}(M) = \min_{V,I,\tilde{v}_{(1)}} \frac{\mathbb{E}_{\Theta\sim\mu(V,I)}\left[\text{Rev}\left(M\left(\Theta, \tilde{v}_{(1)}\right)\right)\right]}{v_{(2)}}$$

- An auction is ***α*-consistent** if its expected revenue is at least $\boldsymbol{\alpha} \cdot \boldsymbol{v}_{(\mathbf{1})}$ whenever the prediction is accurate, i.e., $\boldsymbol{v}_{(\mathbf{1})} = \tilde{\boldsymbol{v}}_{(\mathbf{1})}$

$$\text{consistency}(M) = \min_{V,I} \frac{\mathbb{E}_{\Theta\sim\mu(V,I)}\left[\text{Rev}\left(M\left(\Theta, \boldsymbol{v}_{(1)}\right)\right)\right]}{\boldsymbol{v}_{(1)}}$$

- What are the best (α, β) pairs achievable by strategyproof online auctions augmented with a prediction $\tilde{\boldsymbol{v}}_{(\mathbf{1})}$ regarding the highest bidder value?

Online Auctions with Predictions for a Single Good

- We propose an auction that guarantees **α-consistency** and **$\frac{1-\alpha^2}{4}$ - robustness**
- The designer can choose the value of the **confidence parameter $\alpha \in [0, 1]$**
- We show that this tradeoff is **optimal within a large family of auctions**

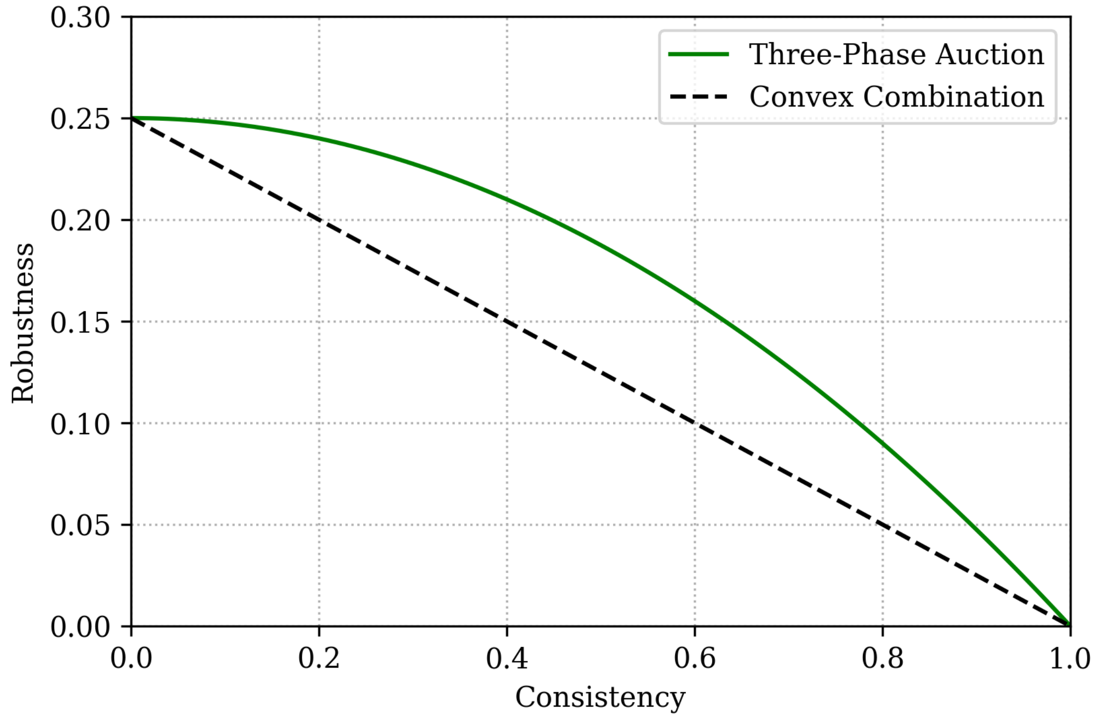

Three-Phase Auction for Disjoint Intervals

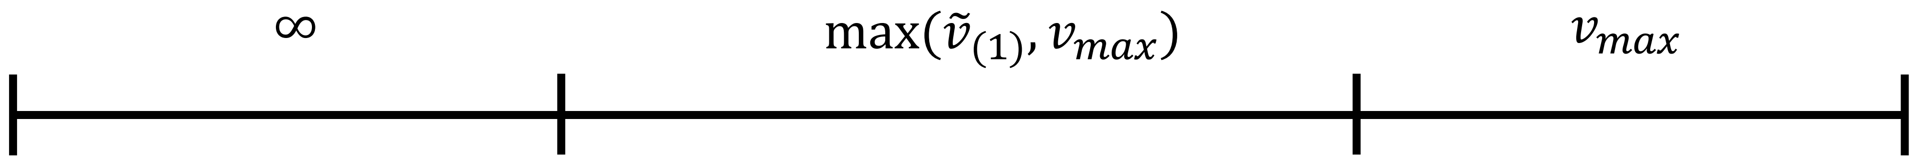

Simple case: if **all active intervals are disjoint**, we get a threshold-price auction

The phases:

1. **Learning phase**: only observe bids, never allocate item
2. **Prediction phase**: post maximum of prediction and highest bid so far
3. **Highest-so-far phase**: post highest bid so far

Phase 2 is skipped if prediction is shown to be inaccurate during phase 1

Three-Phase Auction for Disjoint Intervals

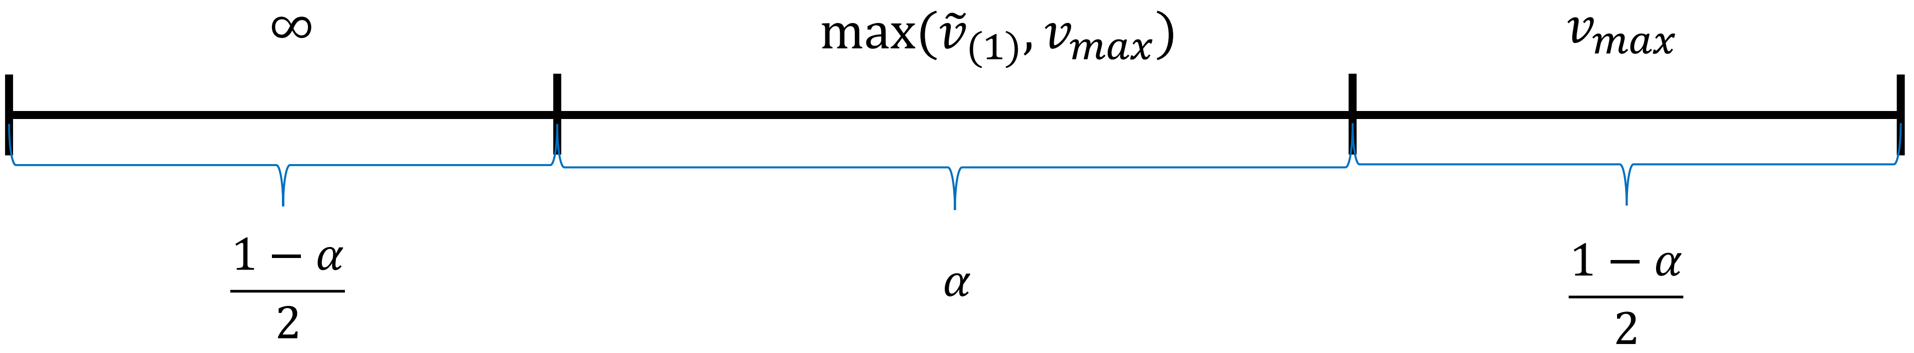

Phase lengths depend on the choice of **confidence parameter** $\alpha \in [0,1]$

Bidders are **ordered by their departure time**

The transition to the second phase takes place after $i_1 = \frac{1-\alpha}{2}n$ **departures**

The transition to the third phase takes place after $i_2 = \frac{1+\alpha}{2}n$ **departures**

Three-Phase Auction with Overlapping Intervals

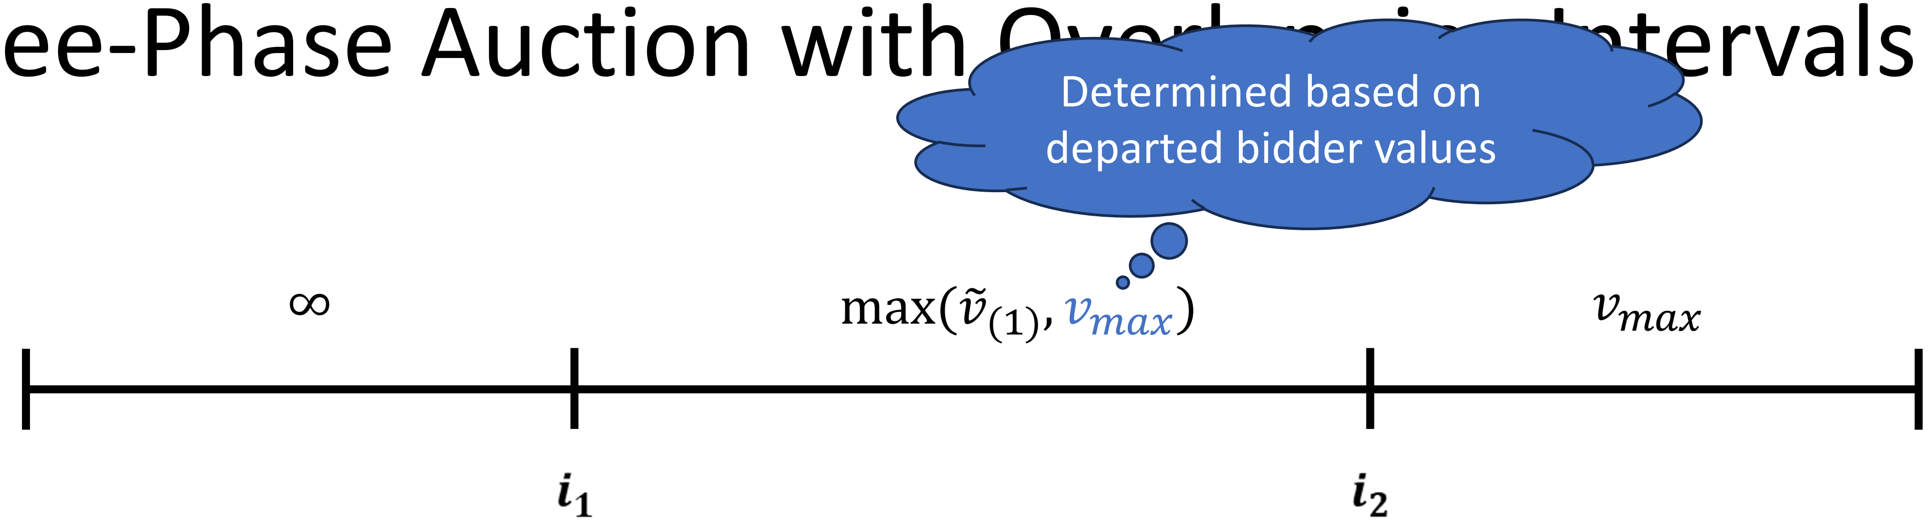

Allocation rule:

- Like before, there are three phases, each with a **threshold price τ**
- The winner is determined as soon as an active bidder has value at least τ
- If there are multiple such active bidders, **higher priority is given to bidders with an earlier arrival time** (ties broken arbitrarily)
- The good is always allocated to the winner **at the time of their departure**

Three-Phase Auction with Overlapping Intervals

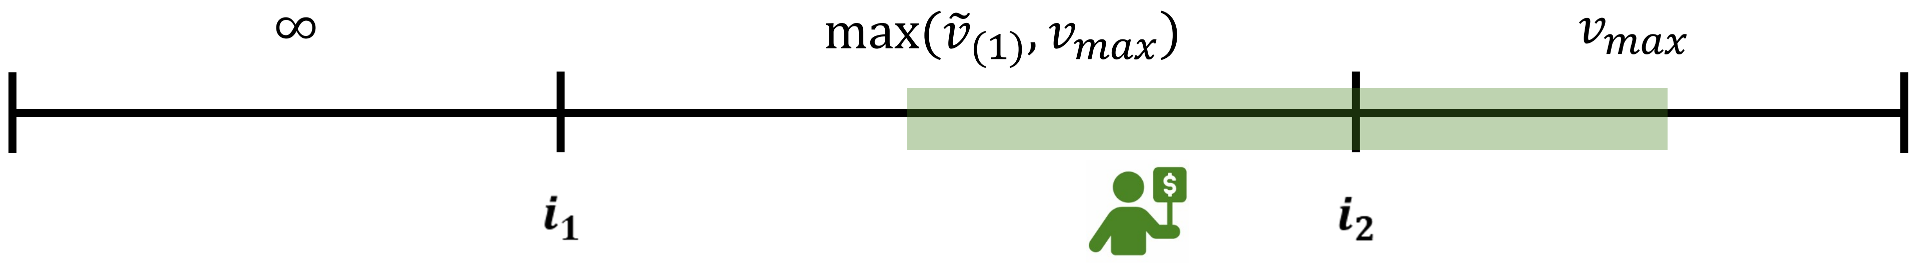

Payment rule:

- The winning bidder, i^*, pays **at most $\boldsymbol{\tau}$**, but may end up paying **less**
- If winner i^* **secures item during Phase 2** and **remains active in Phase 3**:
  - Simulate allocation rule with i^* removed to get winner i' and price τ'
  - If i' is inactive in Phase 3 or has lower priority than i^*, $\boldsymbol{i^*}$**pays price $\boldsymbol{\tau'}$**
  - Else $\boldsymbol{i^*}$ **pays price $\boldsymbol{\tau}$**

Impossibility Result (with Predictions)

- The robustness-consistency trade-off that we achieve is **optimal** over any auction in the **Prediction-or-Previously-Seen** family
- The price posted can be the **prediction**, a **previously seen bid**, or **infinite**
- The proof uses an **interchange argument** reducing any such auction to ours

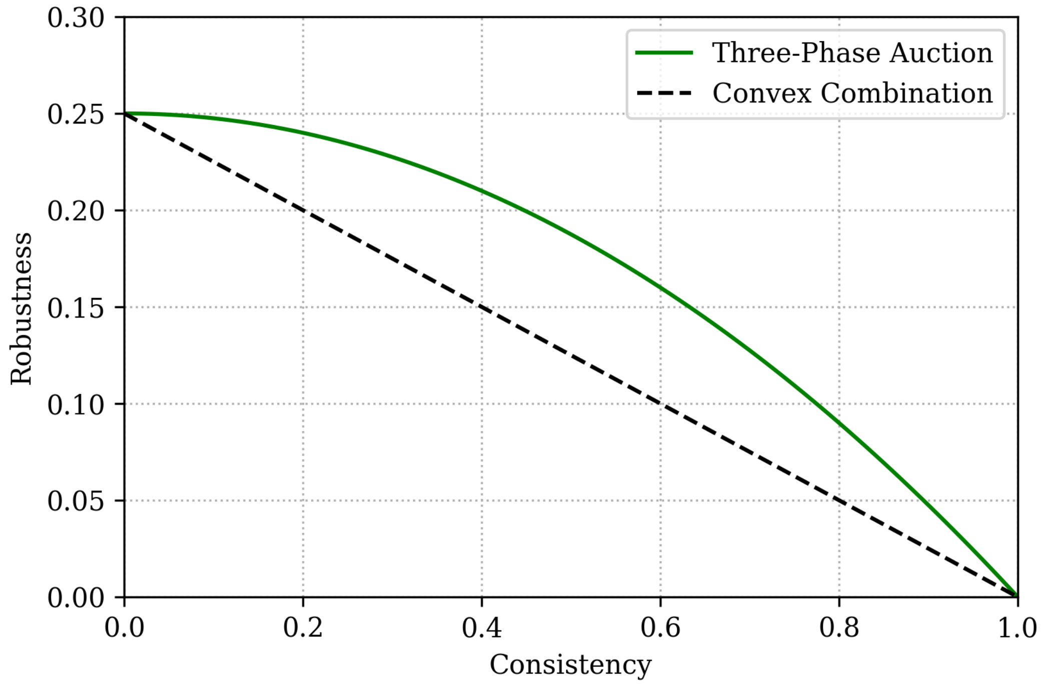

Impossibility Result (without Predictions)

- The 0.25 approximation is optimal for **Up-To-Max-Previously-Seen** auctions
- The price posted can be **at most the maximum bid seen so far** or **infinite**
- The proof uses tools from Correa, Duetting, Fischer, and Schewior [EC '19]
- Unlike their impossibility result, ours needs to use **strategyproofness**

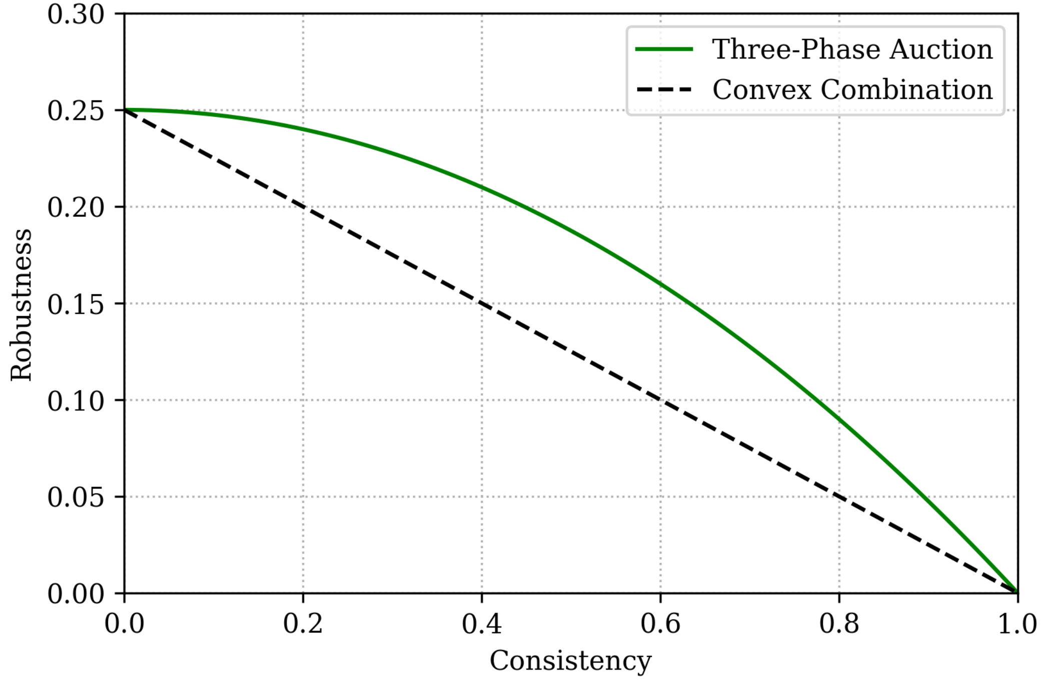

Open Problems and Future Directions

- **General lower bounds** for the single-good case
- What about online auctions for **multiple goods**?
- Many other open problems in learning-augmented mechanism design

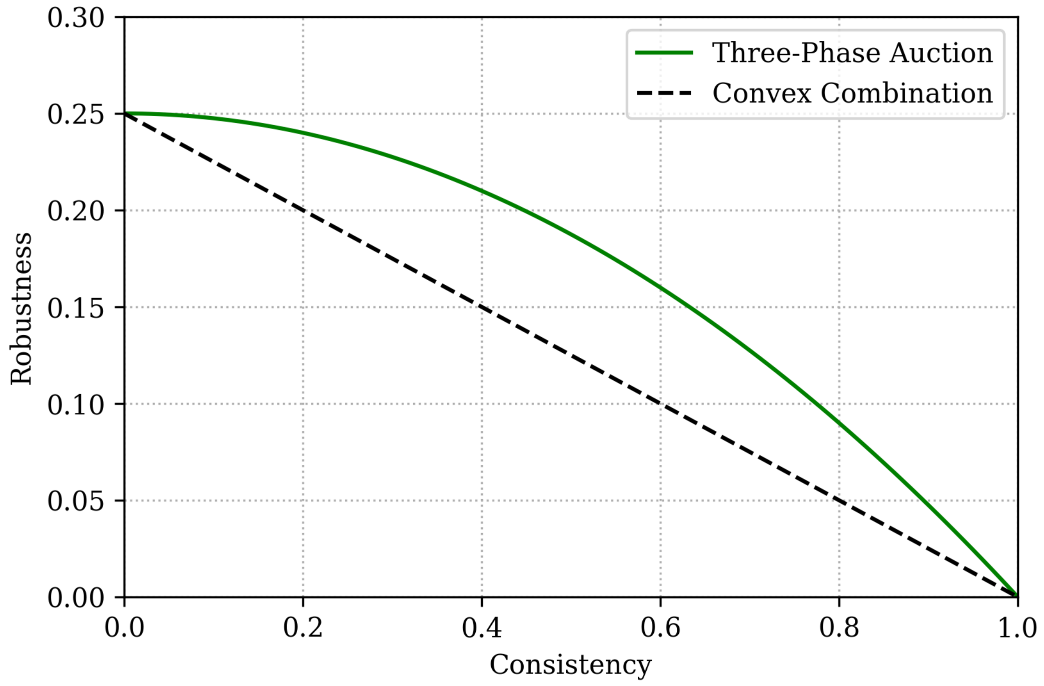

Other Recent Learning-Augmented Work

- **Online Algorithms:**
  - Allocating items that arrive over time, aiming to maximize fairness, with S. Banerjee, A. Gorokh, and B. Jin **(SODA 2022)**
  - Allocating a fixed budget on public goods in a dynamic fashion, with S. Banerjee, S. Hossain, B. Jin, E. Micha, and N. Shah **(IJCAI 2023)**
- **Mechanisms in Strategic Settings:**
  - Strategyproof mechanisms for facility location problems, with P. Agrawal, E. Balkanski, T. Ou, and X. Tan **(EC 2022)**
  - Improved price of anarchy bounds in decentralized systems, with K. Kollias, A. Sgouritsa, and X. Tan **(EC 2022)**
  - Strategyproof mechanisms for scheduling to minimize makespan, with E. Balkanksi and X. Tan **(ITCS 2023)**
  - Online mechanism design with predictions**,** with E. Balkanski, X. Tan, and C. Zhu **(EC 2024)**
  - Randomized strategic facility location with predictions, with E. Balkanski and G. Shahkarami **(Submitted 2024)**
  - Clock auctions augmented with unreliable advice, with D. Schoepflin and X. Tan **(Submitted 2024)**
- **Distortion in Voting:**
  - Optimal metric distortion with predictions**,** with B. Berger, M. Feldman, and X. Tan **(EC 2024)**
- **Robust Algorithmic Recourse in Machine Learning:**
  - Learning-augmented robust algorithmic recourse, with K. Kayastha and S. Jabbari **(Submitted 2024)**

Supported by NSF grant **"Mechanisms with Predictions"** with co-PI Eric Balkanski