Learning Bayesian Game Families, with Application to Mechanism Design

Madelyn Gatchel, Michael P. Wellman

COMPUTER SCIENCE AND ENGINEERING
UNIVERSITY OF MICHIGAN

SRG STRATEGIC REASONING GROUP

Outline

What is empirical mechanism design? 1

How can we learn a Bayesian game-family model? 2

How does the learned game-family model enable…

…comprehensive mechanism search? 3

…strategy generation without additional simulation? 4

[1]Wellman, AAMAS 2006
[2]Vorobeychik et al., EC 2006
[3]Wellman et al., *JAIR* 2025

Empirical Mechanism Design (EMD)

Goal

Find $v_* = \underset{v}{\operatorname{argmax}}\, \mathrm{f}(v)$

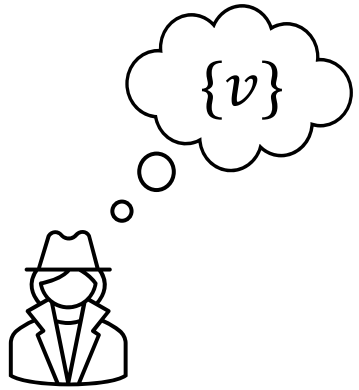

Mechanism Designer

Design Objective $\mathrm{f}(v)$

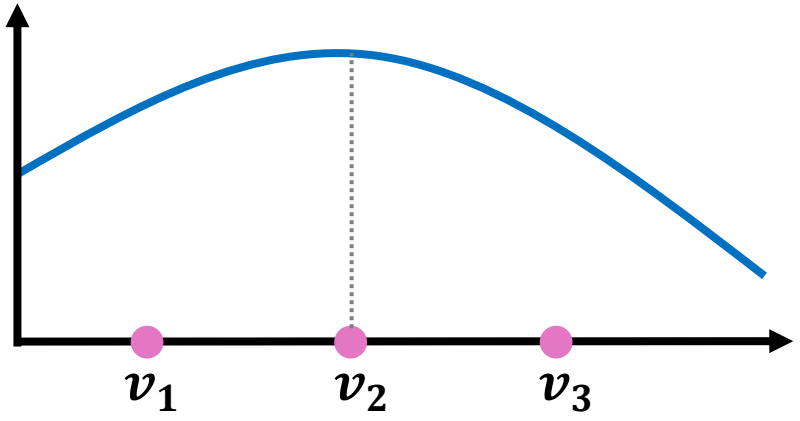

Mechanism Parameter

Q1. What is empirical mechanism design?

[1]Wellman, AAMAS 2006
[2]Vorobeychik et al., EC 2006
[3]Wellman et al., *JAIR* 2025

Empirical Mechanism Design (EMD)

Goal

Find $v_* = \underset{v}{\operatorname{argmax}}\, \mathrm{f}(v)$

Knowledge Source

Agent-based simulator

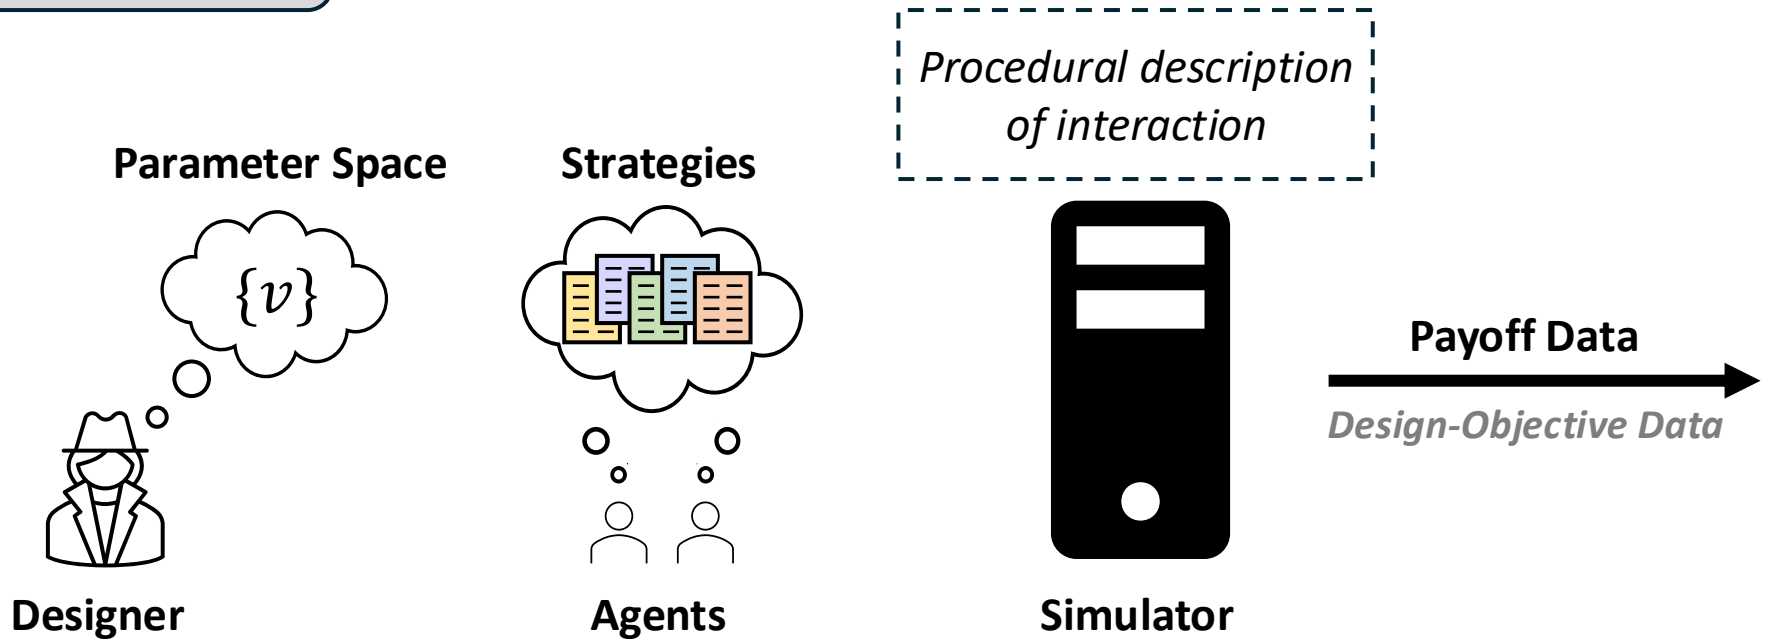

Q1. What is empirical mechanism design?

[1]Wellman, AAMAS 2006
[2]Vorobeychik et al., EC 2006
[3]Wellman et al., *JAIR* 2025

Empirical Mechanism Design (EMD)

Goal — Find $v_* = \operatorname{argmax}_v \mathrm{f}(v)$

Knowledge Source — Agent-based simulator

Reasoning Objects — Empirical game instances induced from simulation, $\{\Gamma(v)\}$

Game with mechanism parameter set to v

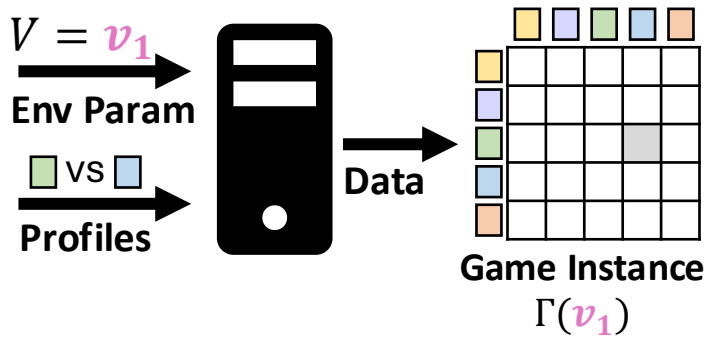

Q1. What is empirical mechanism design?

Motivating Example: Online Advertising Auctions

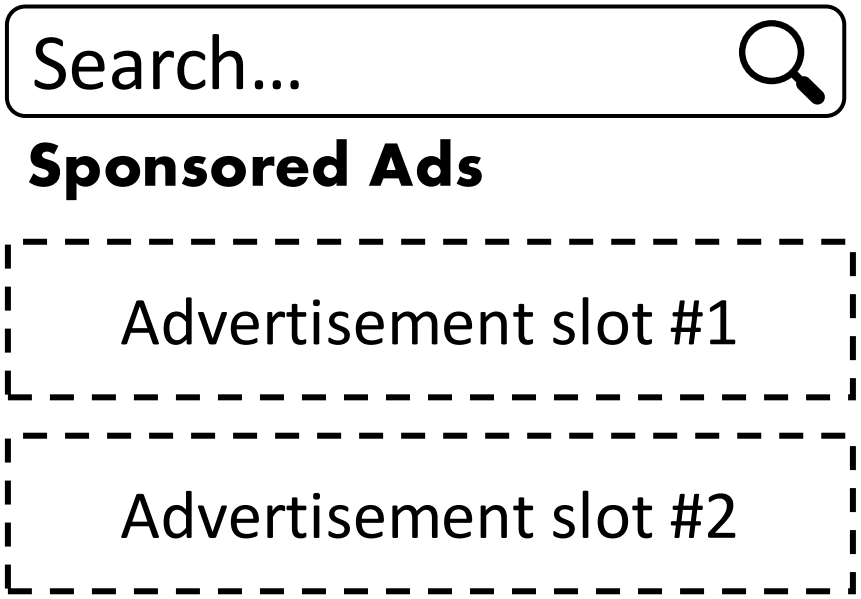

Reserve price: The minimum price at which a slot can be sold.

Mechanism Parameter	Reserve price
Design Objective	Expected revenue
Knowledge Source	Dynamic sponsored search auction simulator

Mechanism Design: Finding Optimal Reserve

Mechanism Design: Finding Optimal Reserve

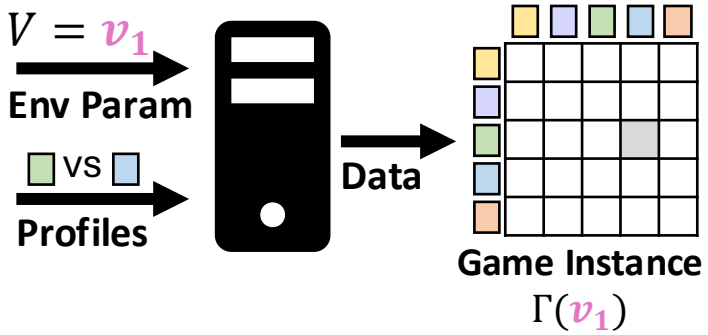

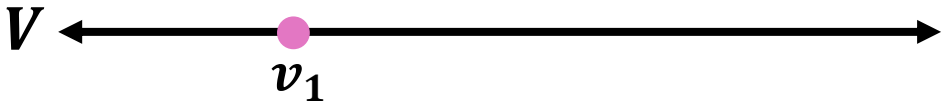

Q1. What is empirical mechanism design?

Mechanism Design: Finding Optimal Reserve

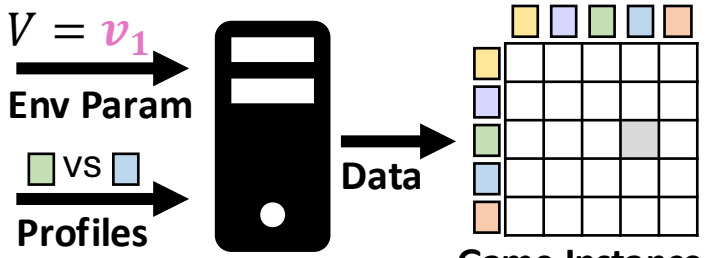

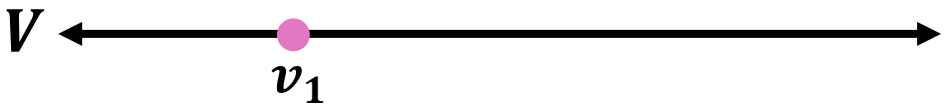

Q1. What is empirical mechanism design?

Mechanism Design: Finding Optimal Reserve

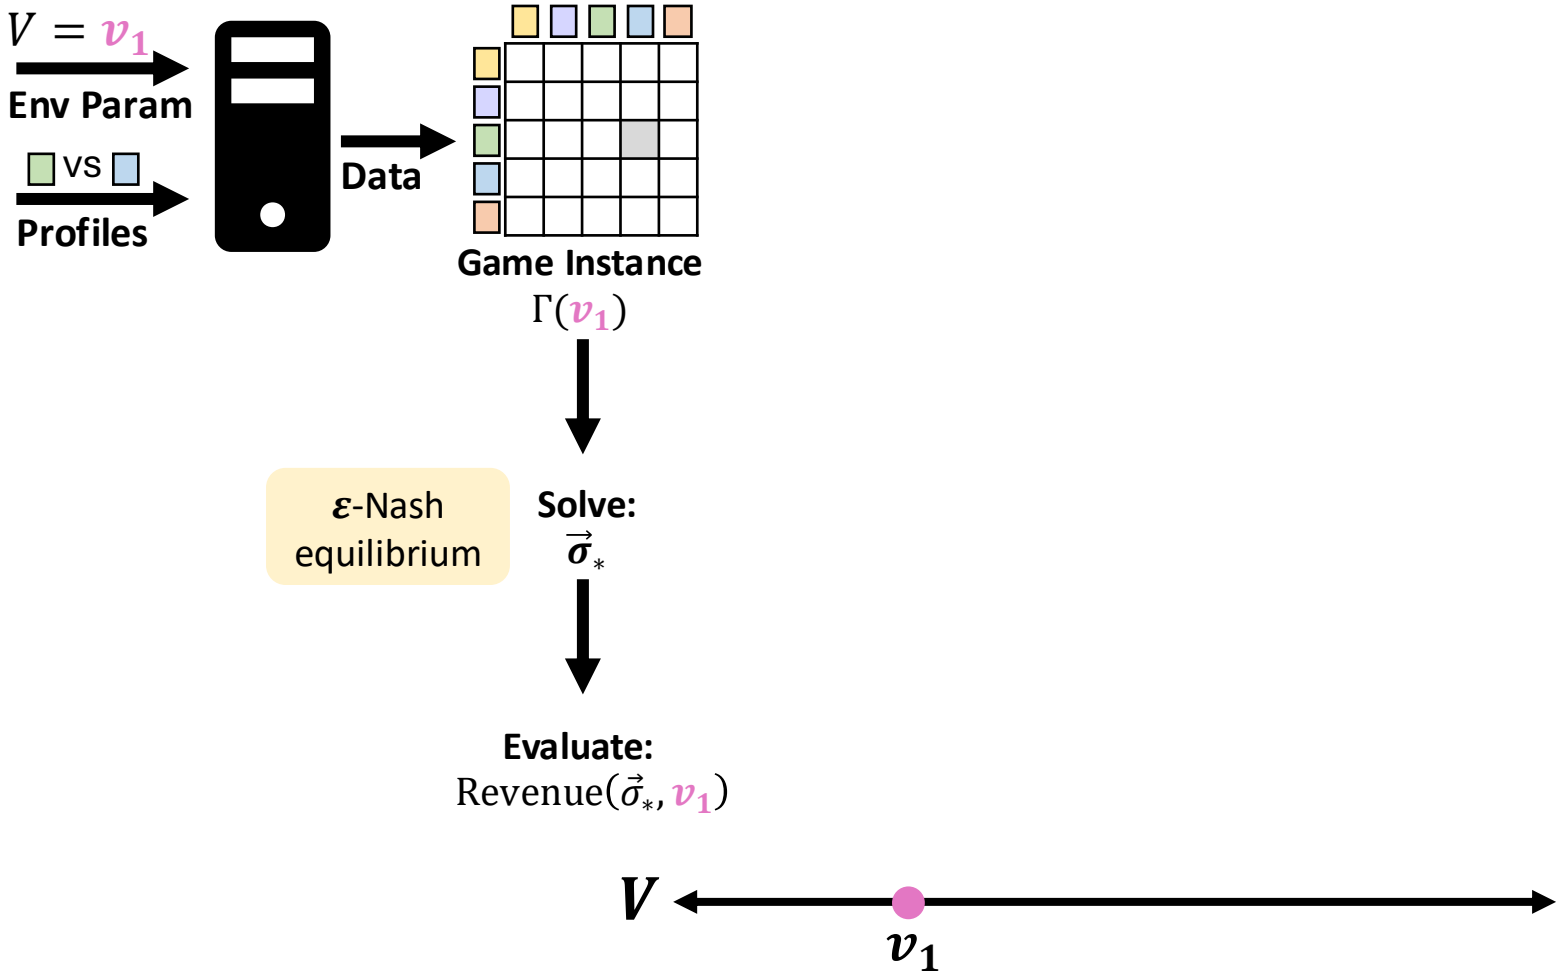

Q1. What is empirical mechanism design?

Mechanism Design: Finding Optimal Reserve

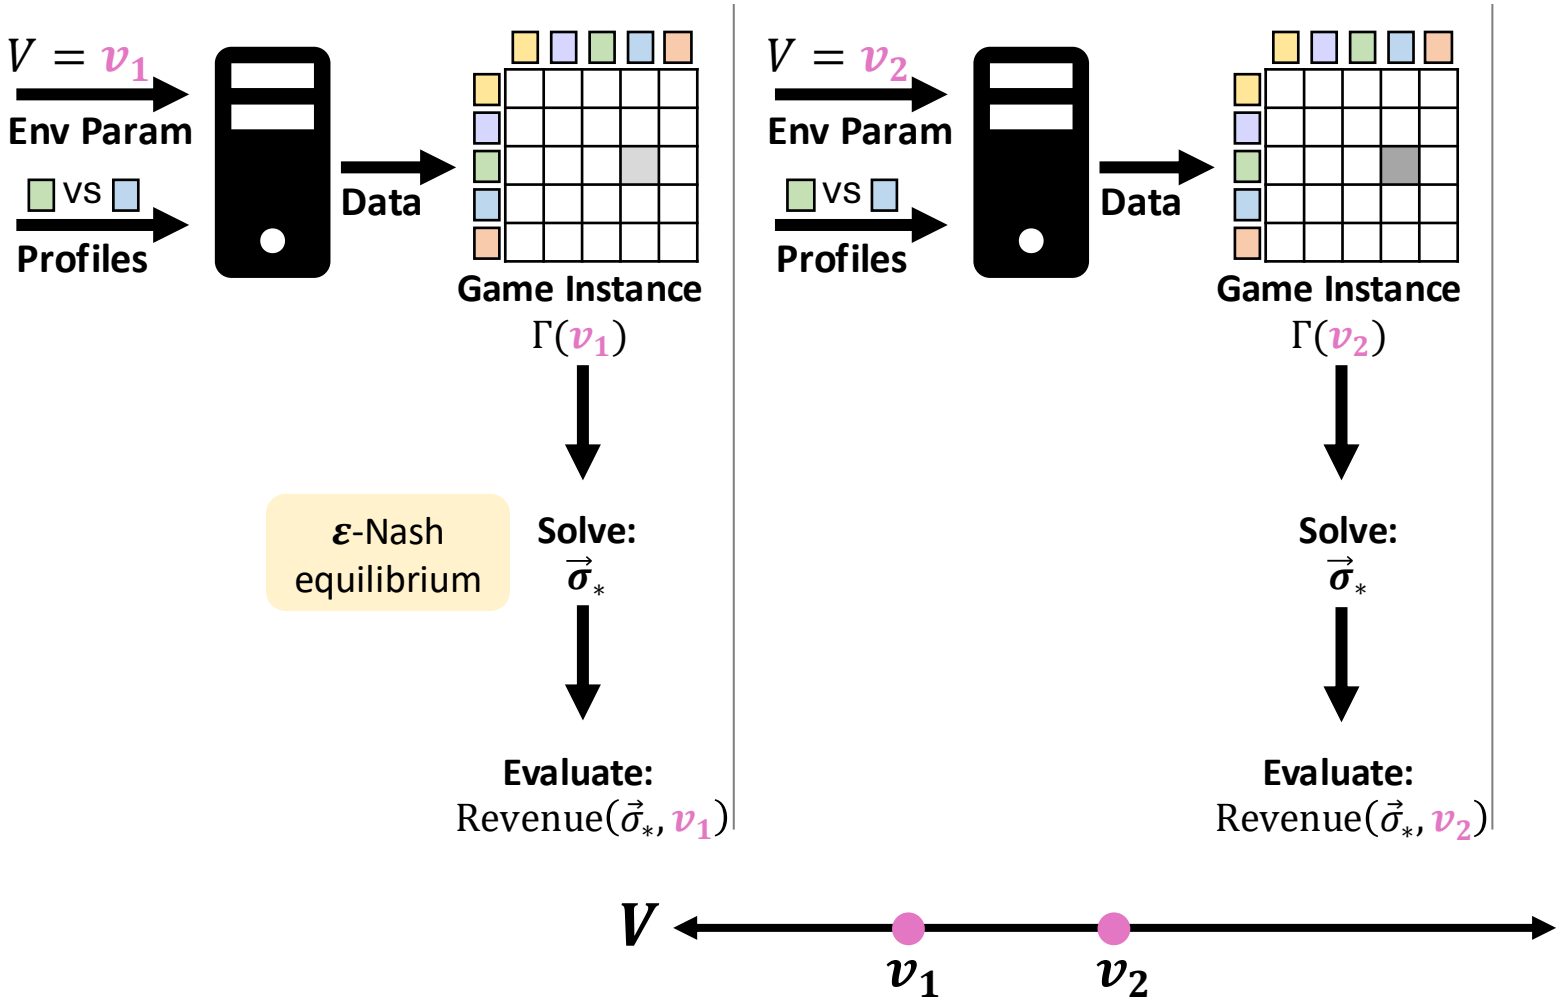

Q1. What is empirical mechanism design?

Mechanism Design: Finding Optimal Reserve

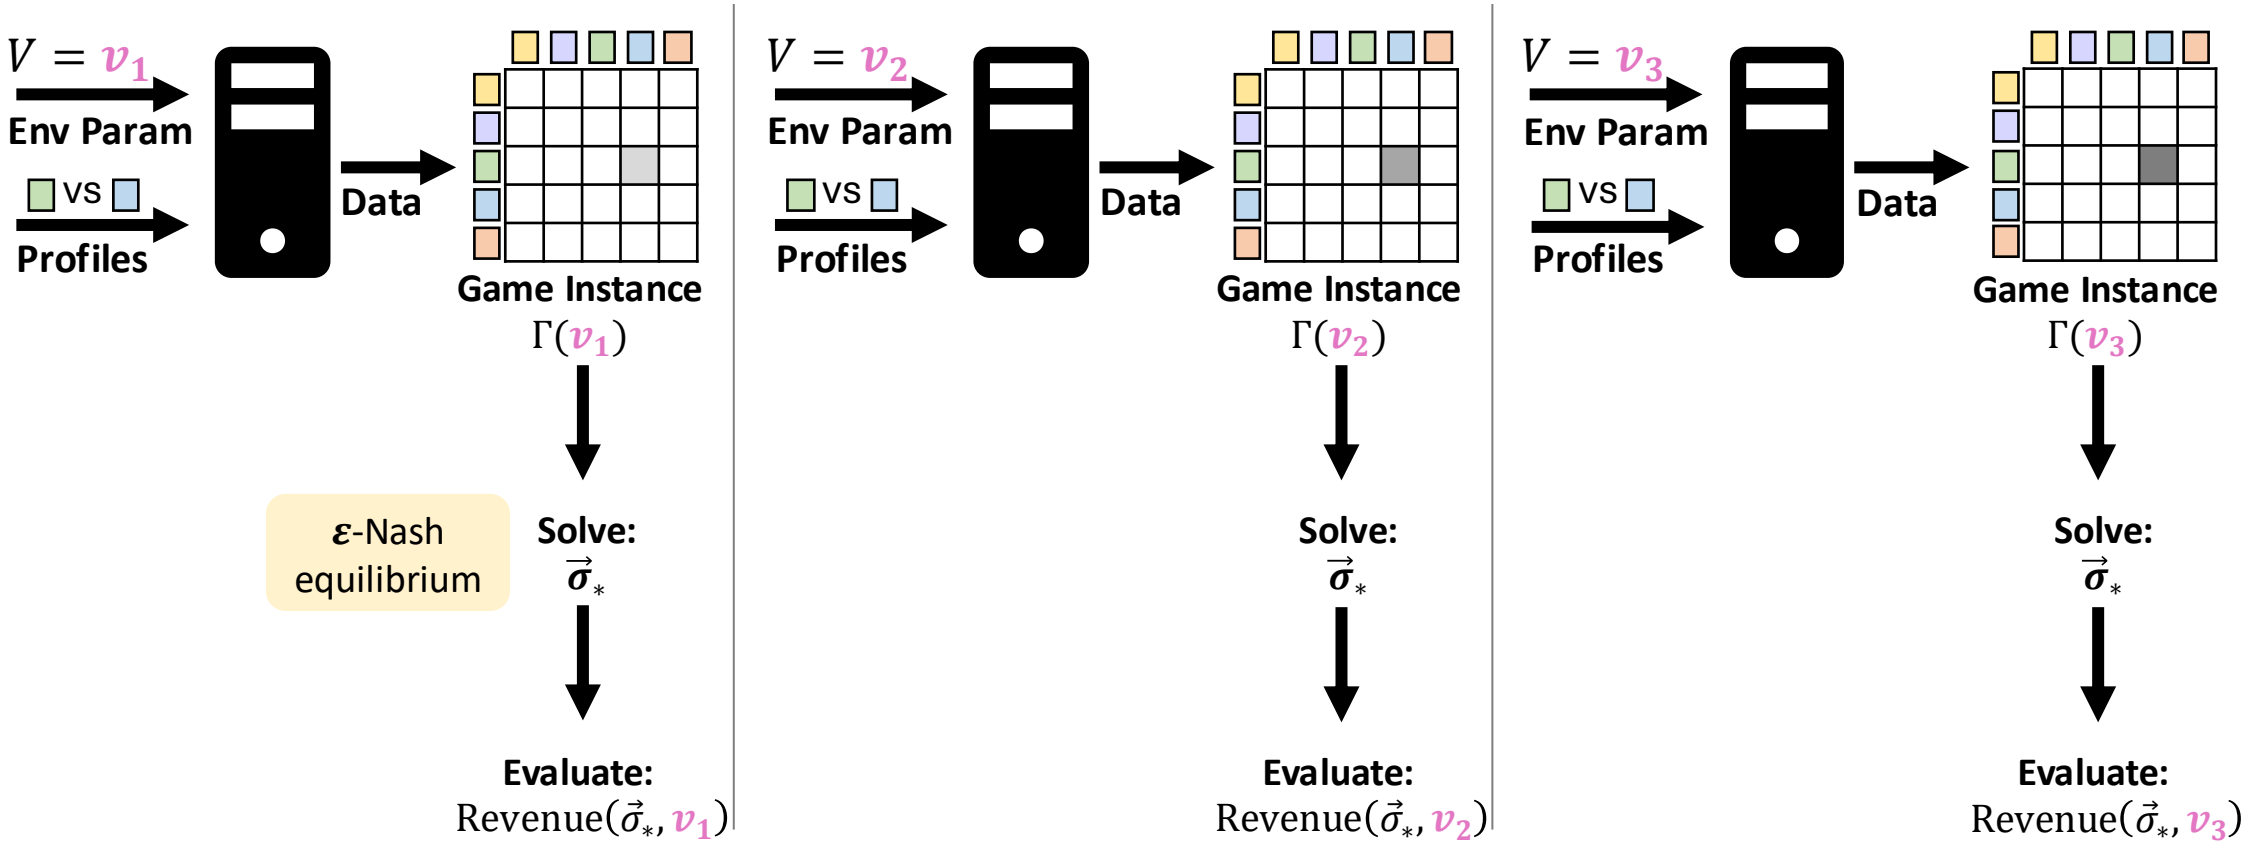

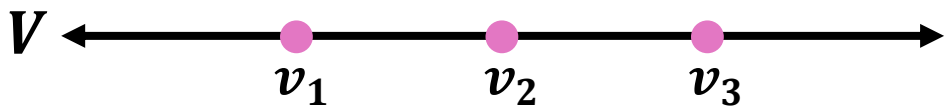

Q1. What is empirical mechanism design?

Mechanism Design: Finding Optimal Reserve

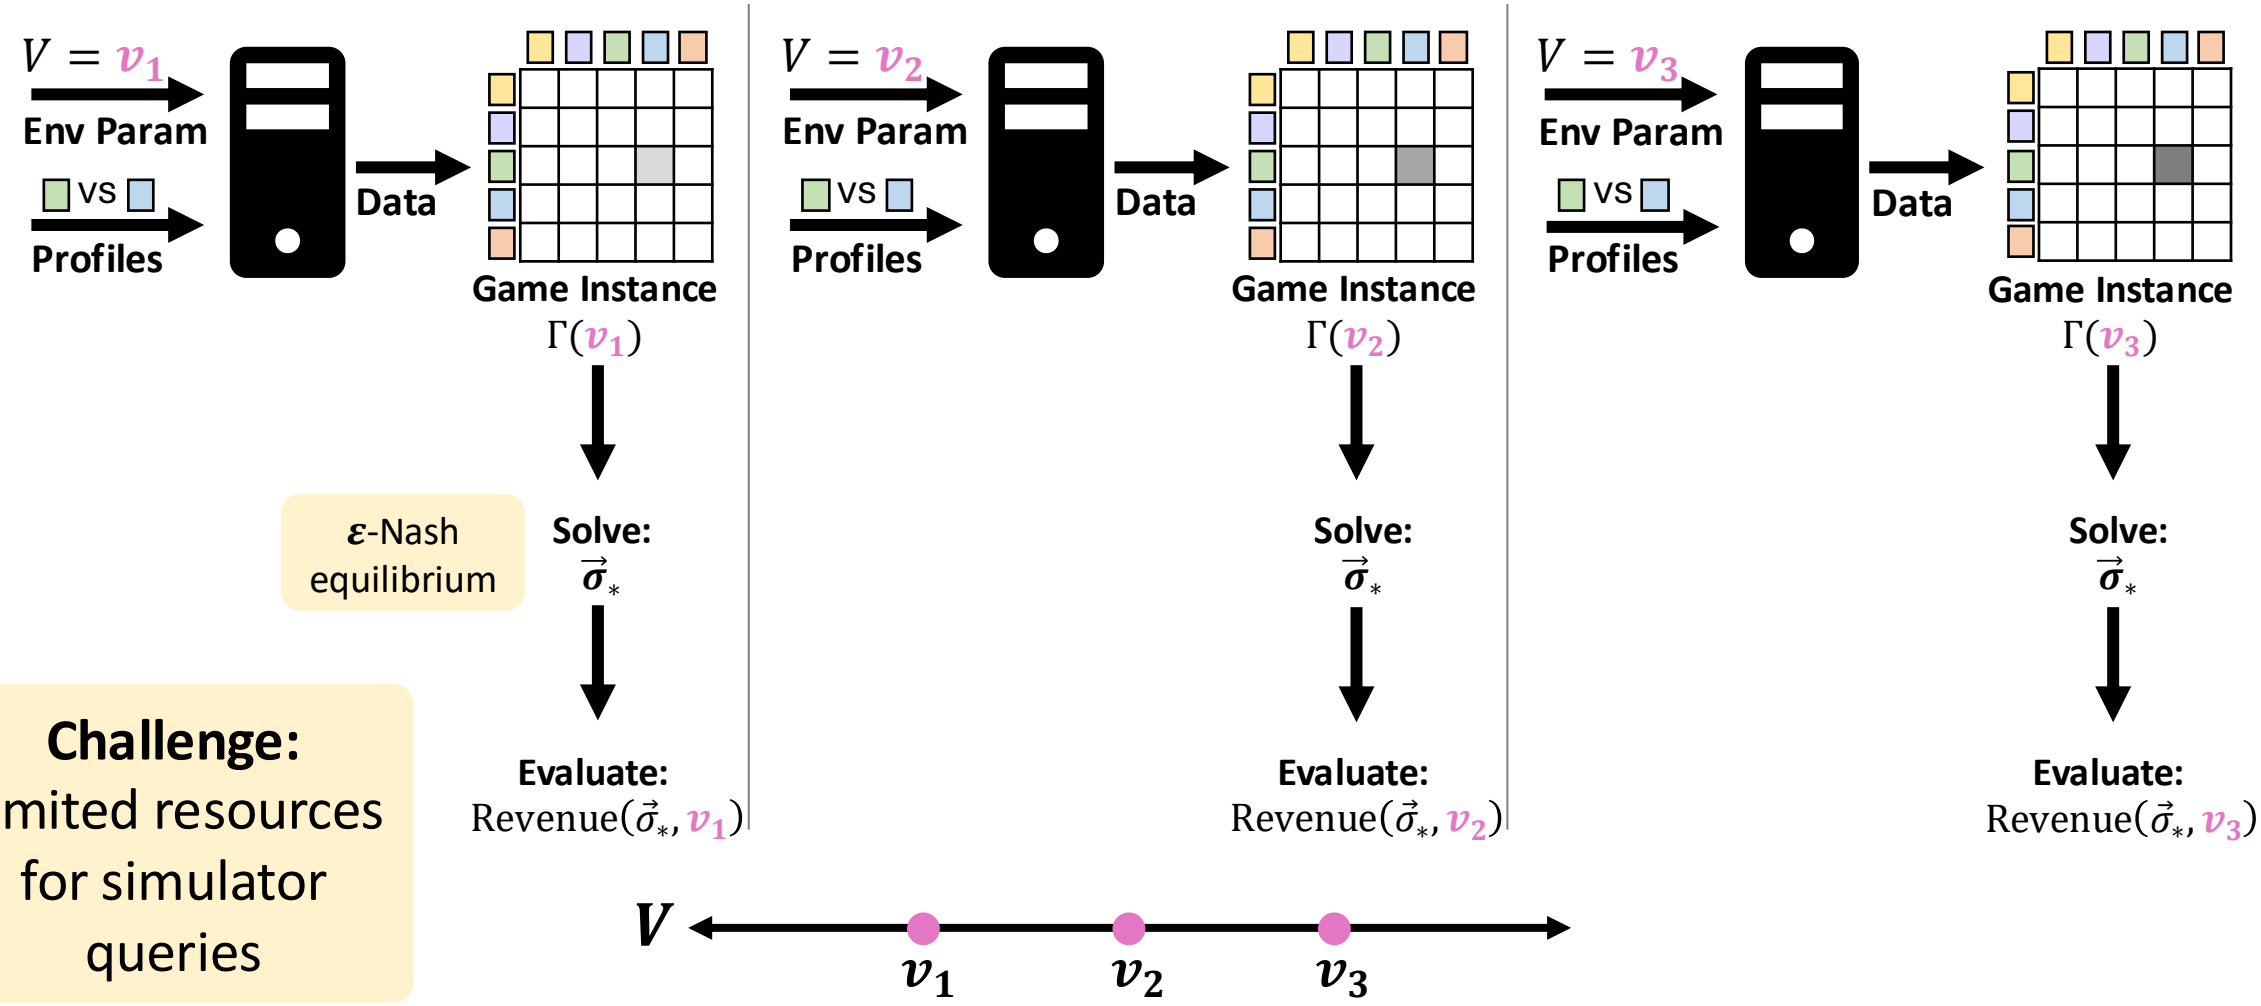

Challenge:
Limited resources for simulator queries

Q1. What is empirical mechanism design?

Outline

What is empirical mechanism design? 1

How can we learn a Bayesian game-family model? 2

How does the learned game-family model enable…

…comprehensive mechanism search? 3

…strategy generation without additional simulation? 4

[1]Sokota et al., *AAAI* 2019

How can we learn a Bayesian game-family model from data?

Game-Instance Learning[1]

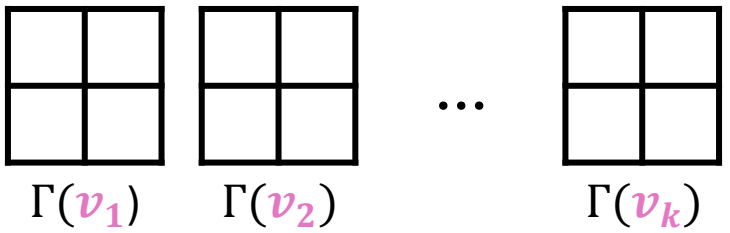

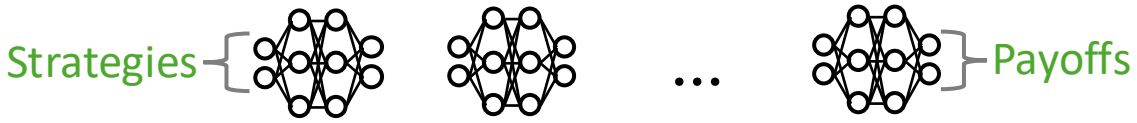

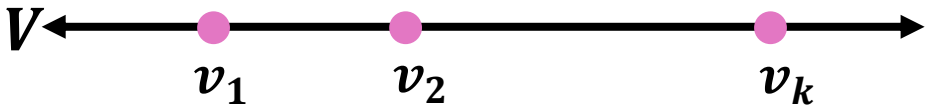

Payoff representation leverages player symmetry.

Q2. How can we learn a Bayesian game-family model?

[1]Sokota et al., *AAAI* 2019
[2]Gatchel and Wiedenbeck, AAMAS 2023.

How can we learn a Bayesian game-family model from data?

Game-Instance Learning[1]

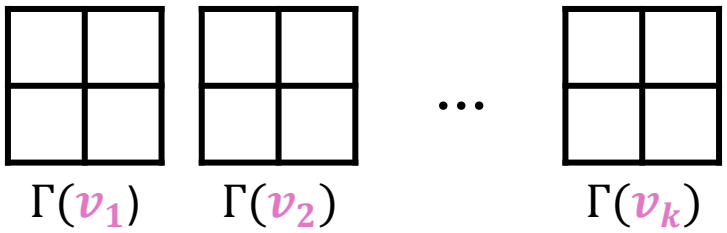

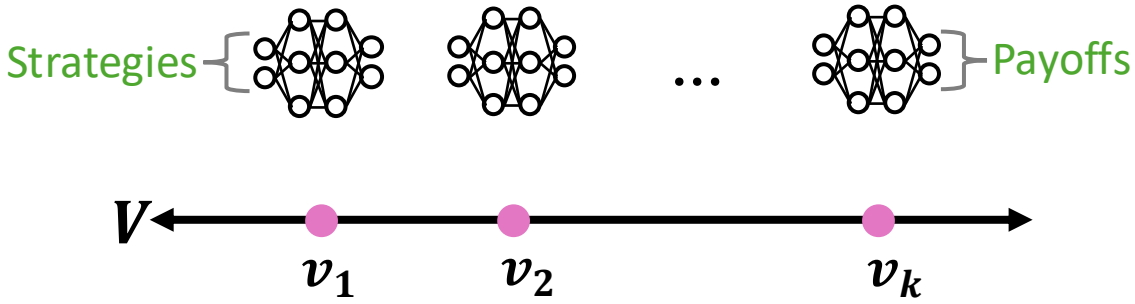

Payoff representation leverages player symmetry.

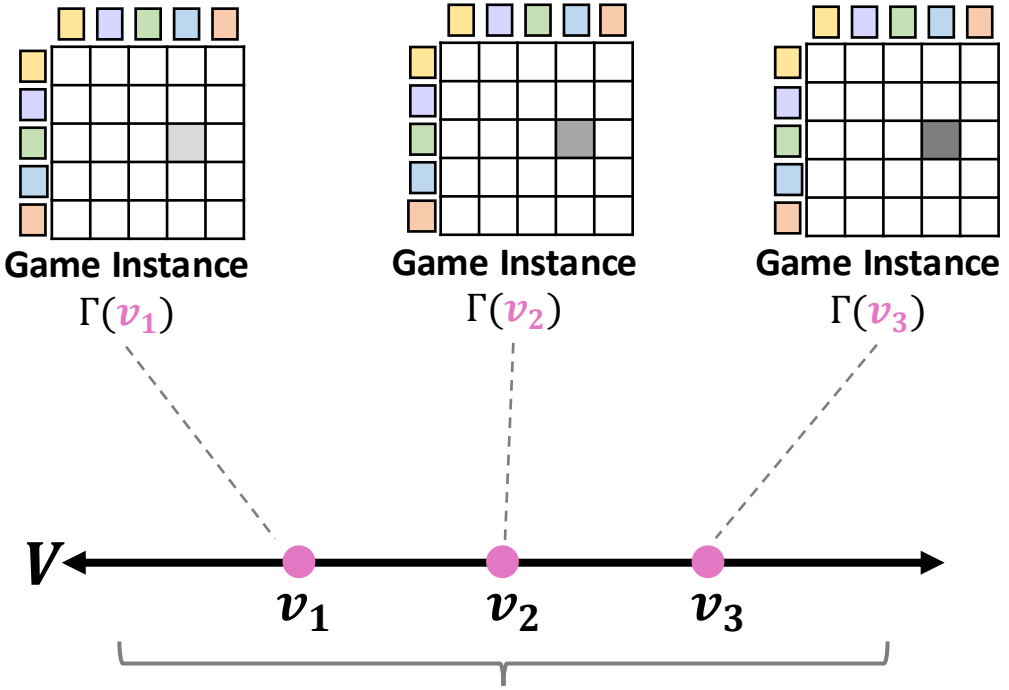

Family of game instances

Q2. How can we learn a Bayesian game-family model?

[1]Sokota et al., *AAAI* 2019
[2]Gatchel and Wiedenbeck, AAMAS 2023.

How can we learn a Bayesian game-family model from data?

Game-Instance Learning[1]

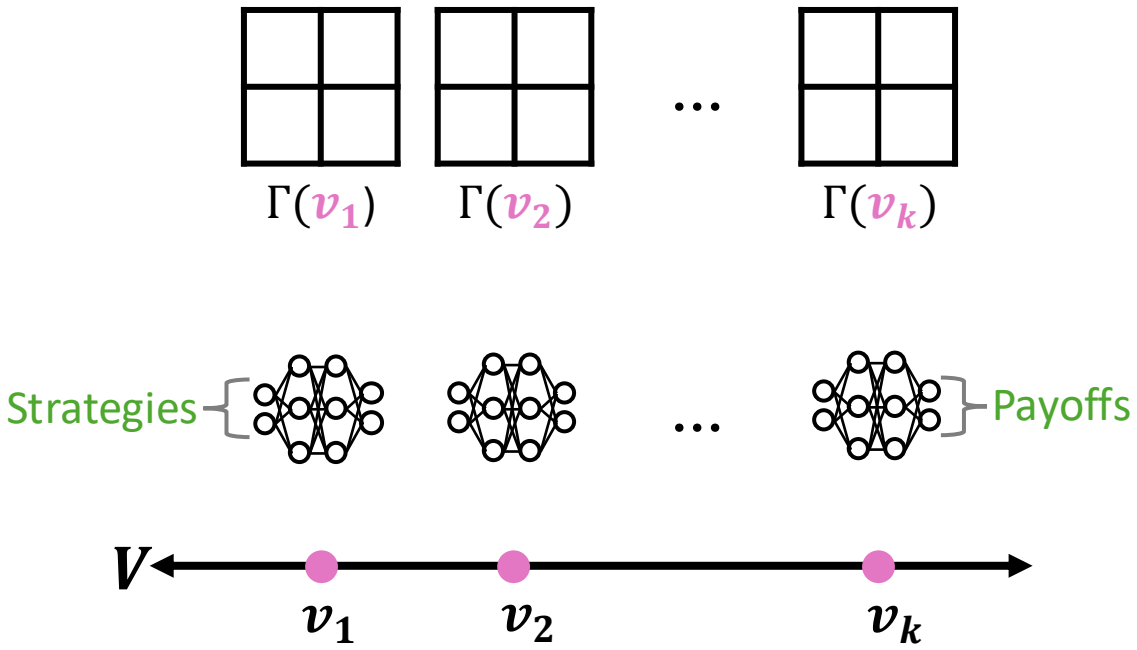

Payoff representation leverages player symmetry.

Game-Family Learning[2]

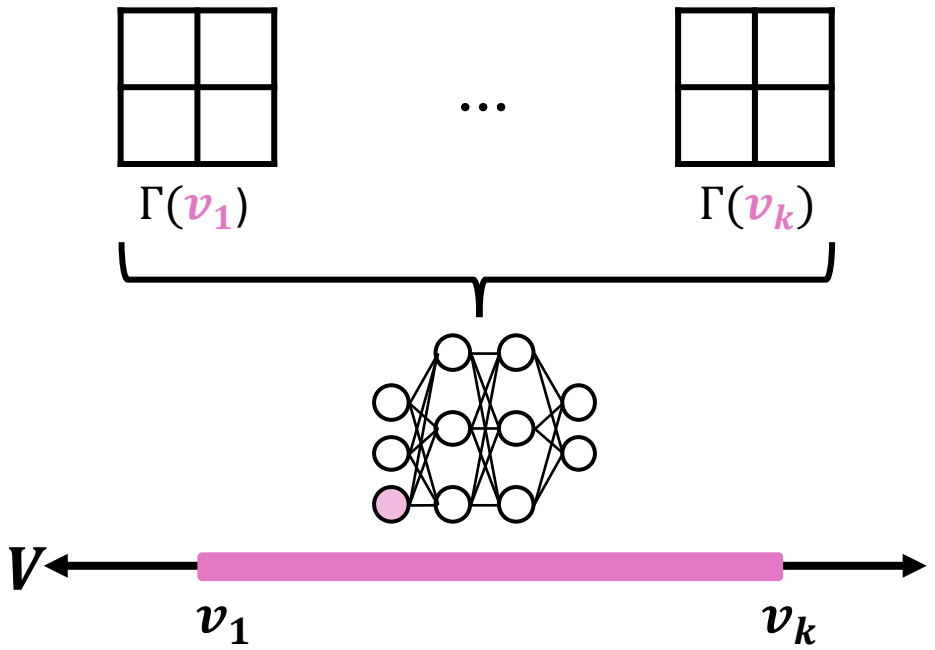

A game-family model achieves ***better accuracy with less data.***

Q2. How can we learn a Bayesian game-family model?

Modeling Opponent Uncertainty with Bayesian Games

Type = Private information:

- θ: Advertiser valuation
- q: Advertiser quality score

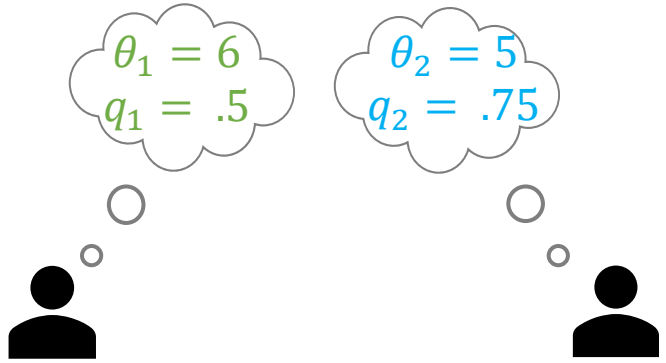

Q2. How can we learn a Bayesian game-family model?

CTR = Click-Through Rate

Modeling Opponent Uncertainty with Bayesian Games

Type = Private information:

- θ: Advertiser valuation
- q: Advertiser quality score

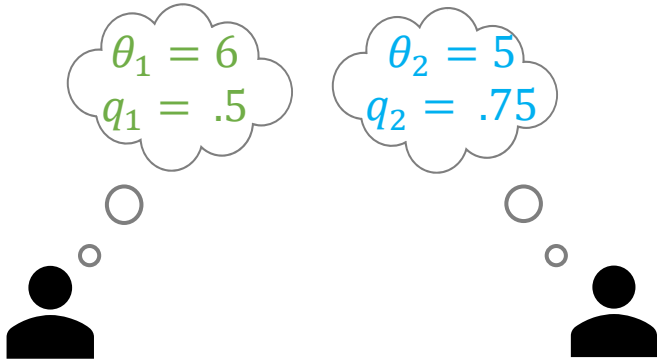

Payoff = CTR * (θ – Price)

A player's payoff depends on their ***own type*** and the ***types of others***.

Q2. How can we learn a Bayesian game-family model?

CTR = Click-Through Rate

Modeling Opponent Uncertainty with Bayesian Games

Type = Private information:
- θ: Advertiser valuation
- q: Advertiser quality score

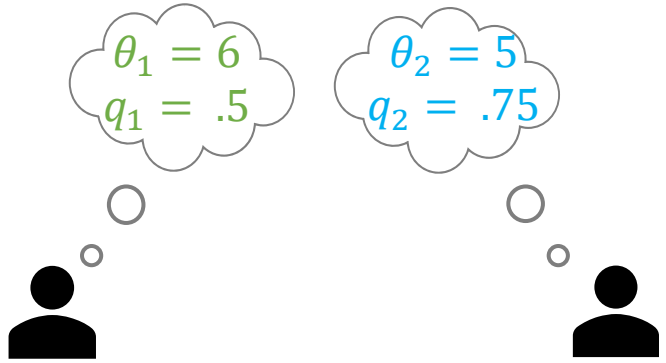

Payoff = CTR * (θ – Price)

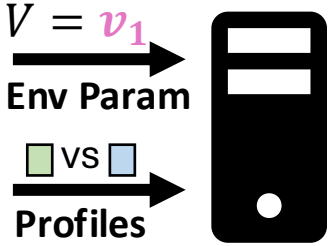

Ex Ante Payoff

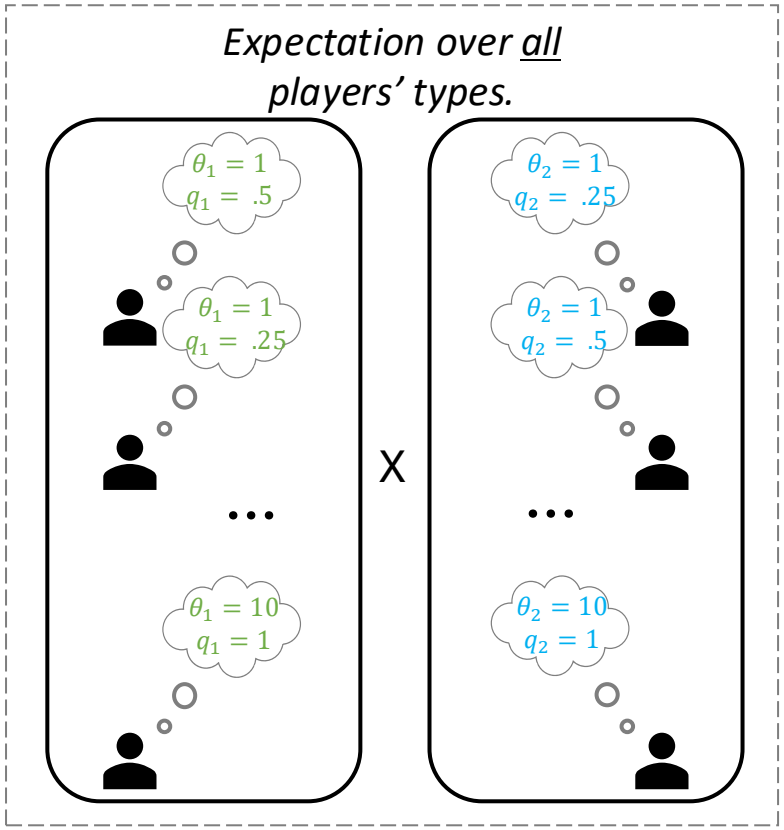

A player's payoff depends on their ***own type*** and the ***types of others***.

Q2. How can we learn a Bayesian game-family model?

CTR = Click-Through Rate

Modeling Opponent Uncertainty with Bayesian Games

Type = Private information:
- θ: Advertiser valuation
- q: Advertiser quality score

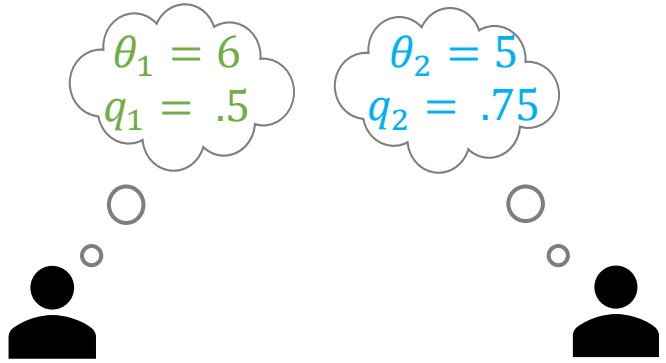

Payoff = CTR * (θ – Price)

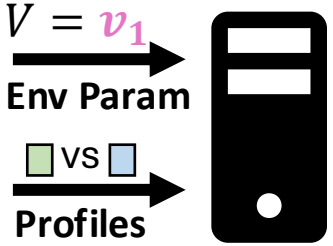

Interim Payoff

Expectation over opponents' types, and conditional on player's own type.

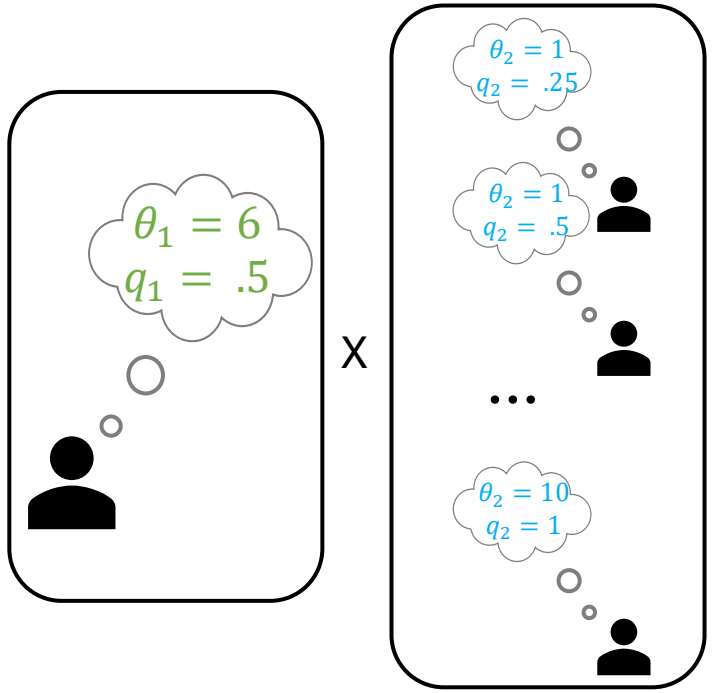

A player's payoff depends on their ***own type*** and the ***types of others***.

Q2. How can we learn a Bayesian game-family model?

Learning Bayesian Game Families

Ex Ante Payoff Function

Expectation over all players' types.

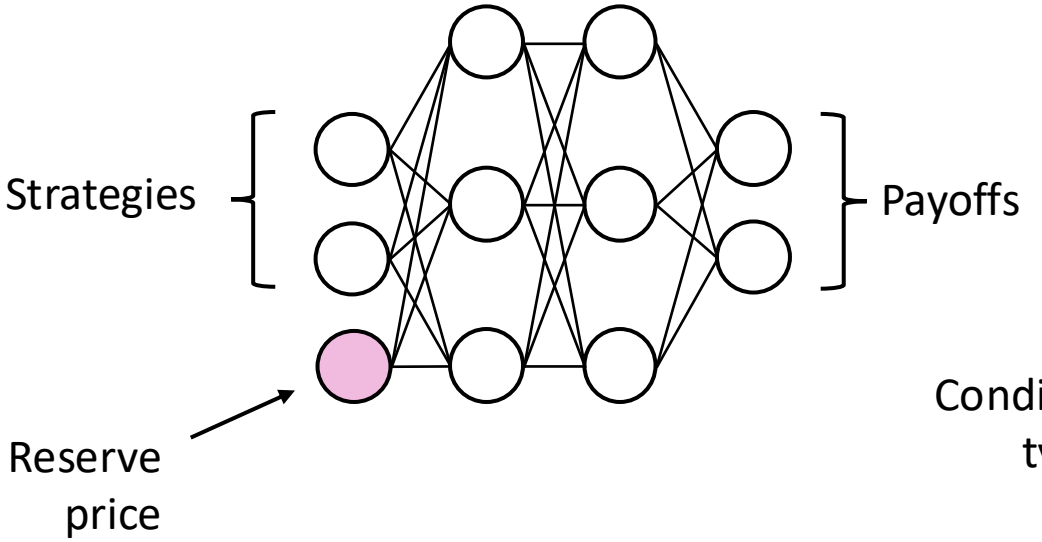

Interim Payoff Function

Expectation over opponents' types.

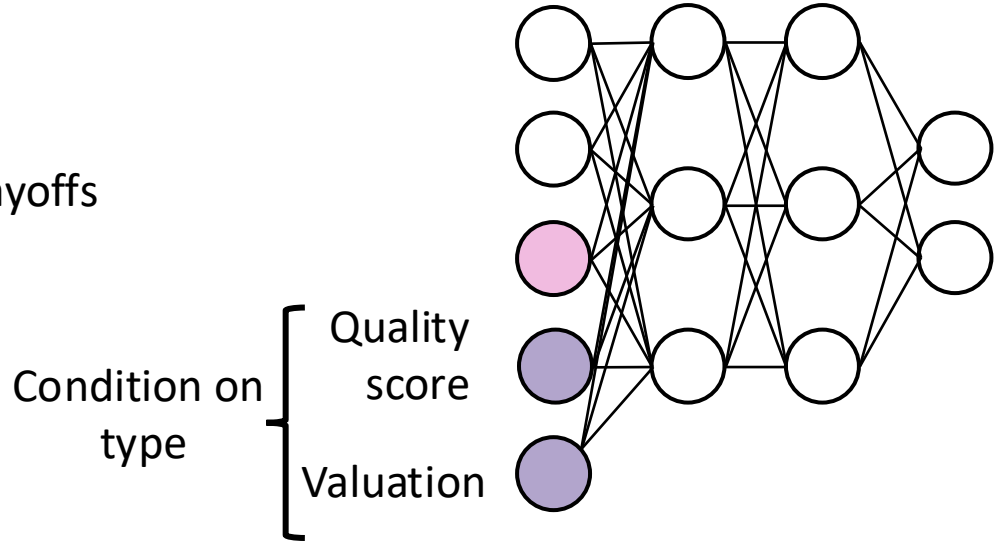

Q2. How can we learn a Bayesian game-family model?

Learning Bayesian Game Families

Ex Ante Payoff Function

Expectation over all players' types.

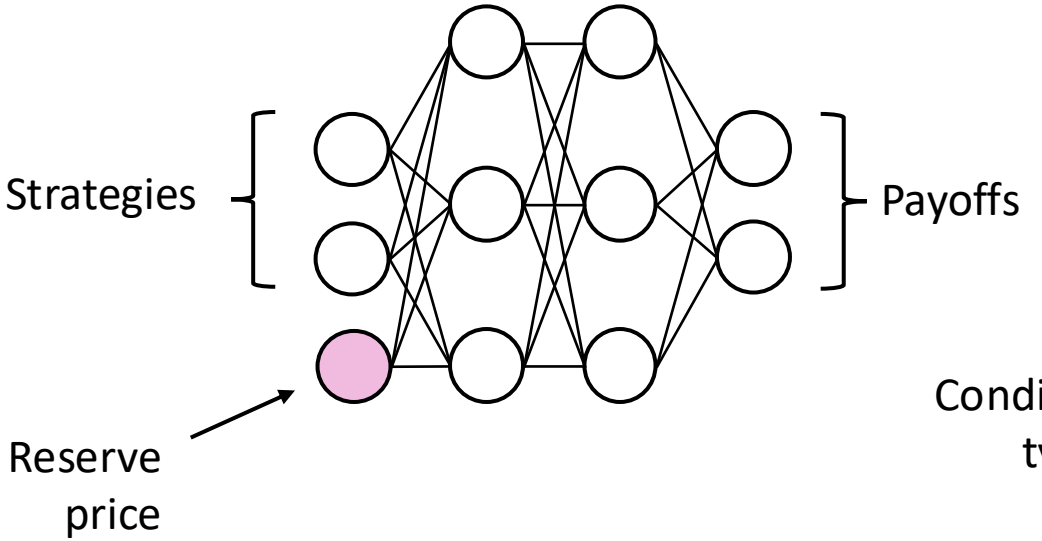

Interim Payoff Function

Expectation over opponents' types.

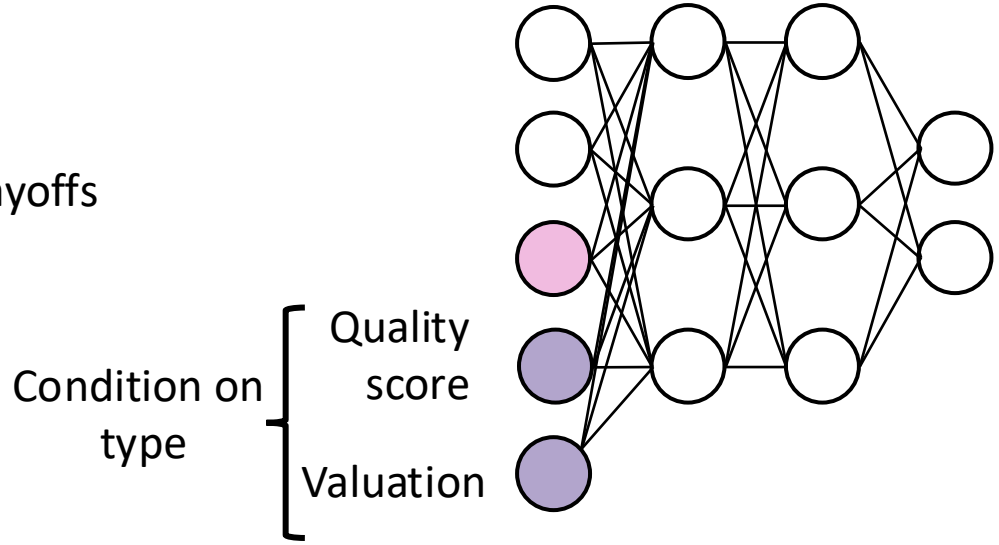

Given Q simulator queries…

Neural network is given **less training data**, but **data is more accurate**.

Neural network is given **more training data**, but **marginalization required** for inference.

Q2. How can we learn a Bayesian game-family model?

Learning Bayesian Game Families

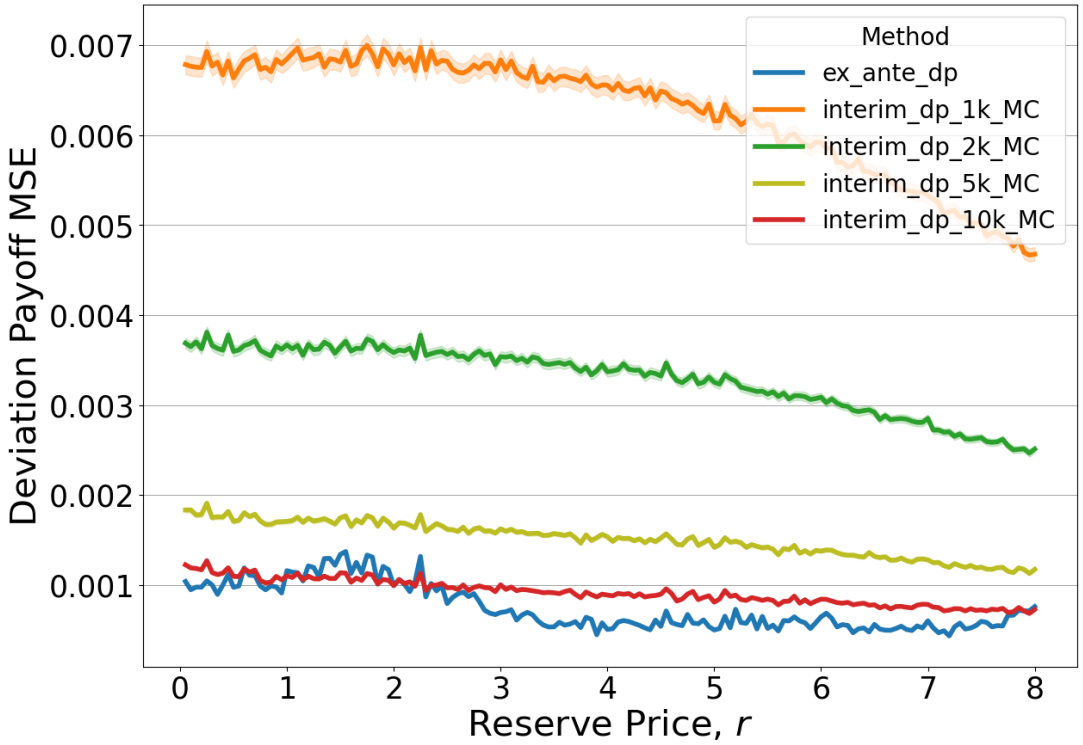

Interim, **1k** marginalization samples

Interim, **2k** marginalization samples

Interim, **5k** marginalization samples

Interim, **10k** marginalization samples
Ex ante

With enough marginalization samples, interim accuracy matches ex ante on trained range.

Q2. How can we learn a Bayesian game-family model?

Learning Bayesian Game Families

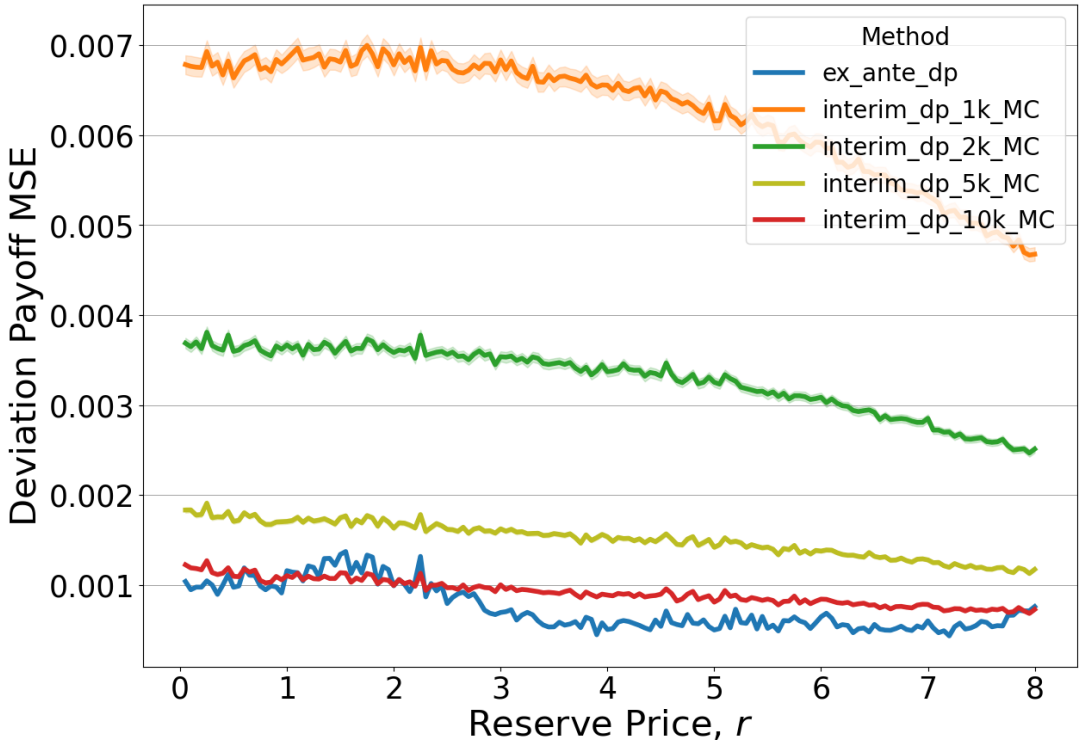

With enough marginalization samples, interim accuracy matches ex ante on trained range.

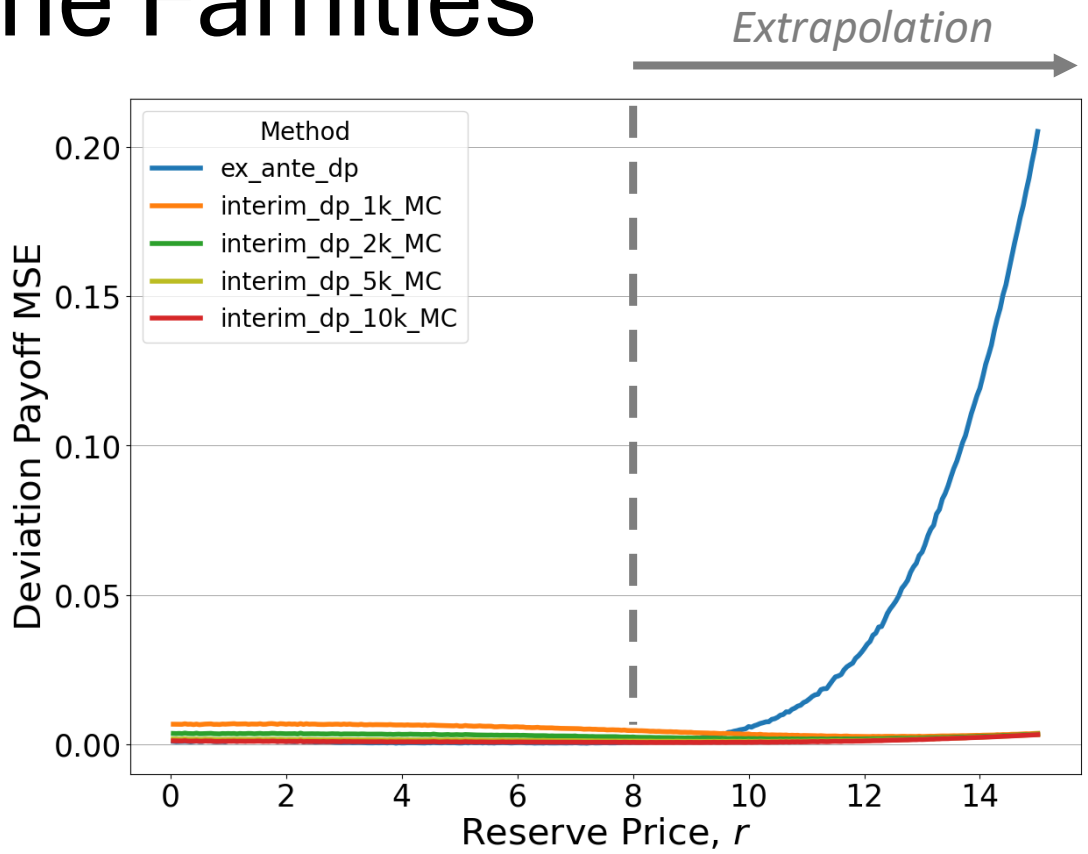

Only interim models extrapolate well.

Parameter range decided prior to any analysis.

Q2. How can we learn a Bayesian game-family model?

Outline

What is empirical mechanism design? 1

How can we learn a Bayesian game-family model? 2

How does the learned game-family model enable…

…comprehensive mechanism search? 3

…strategy generation without additional simulation? 4

EMD with Learned Bayesian Game Families

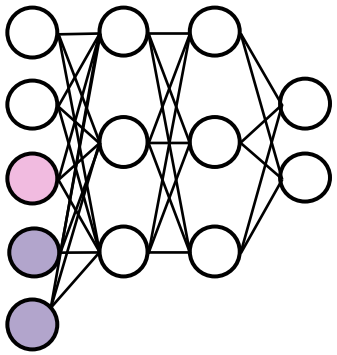

Learn Bayesian game-family model, $\hat{u}$

Q3. How does the learned game-family model enable comprehensive mechanism search?

EMD with Learned Bayesian Game Families

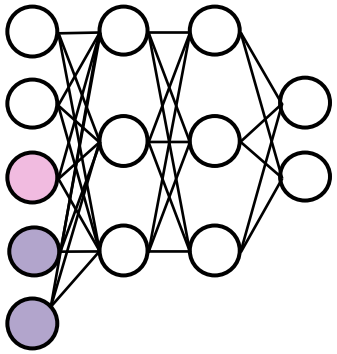

Evaluate game instance $\boldsymbol{\Gamma}(\boldsymbol{v_1})$:

$[\vec{\boldsymbol{\sigma}}_*] = \boldsymbol{solve}(\boldsymbol{v_1}, \hat{\boldsymbol{u}})$

$\mathrm{f}(\vec{\boldsymbol{\sigma}}_*, \boldsymbol{v_1})$

Learn Bayesian game-family model, $\hat{u}$ → Use $\hat{u}$ to approximate equilibria in relevant game instance(s) → Evaluate design objective in equilibrium

Q3. How does the learned game-family model enable comprehensive mechanism search?

EMD with Learned Bayesian Game Families

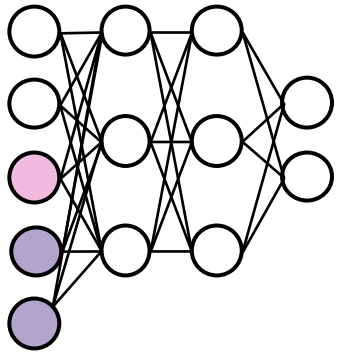

Selected by search procedure

Evaluate game instance $\boldsymbol{\Gamma}(v_m)$:

$[\vec{\sigma}_*] = \boldsymbol{solve}(v_m, \hat{u})$

$f(\vec{\sigma}_*, v_m)$

Learn Bayesian game-family model, $\hat{u}$

→

Use $\hat{u}$ to approximate equilibria in relevant game instance(s)

→

←

Evaluate design objective in equilibrium

Optimize via grid search, local search, Bayesian optimization, etc.

Q3. How does the learned game-family model enable comprehensive mechanism search?

EMD: Finding Optimal Reserve

Ex Ante

Interim

Q3. How does the learned game-family model enable comprehensive mechanism search?

EMD: Finding Optimal Reserve

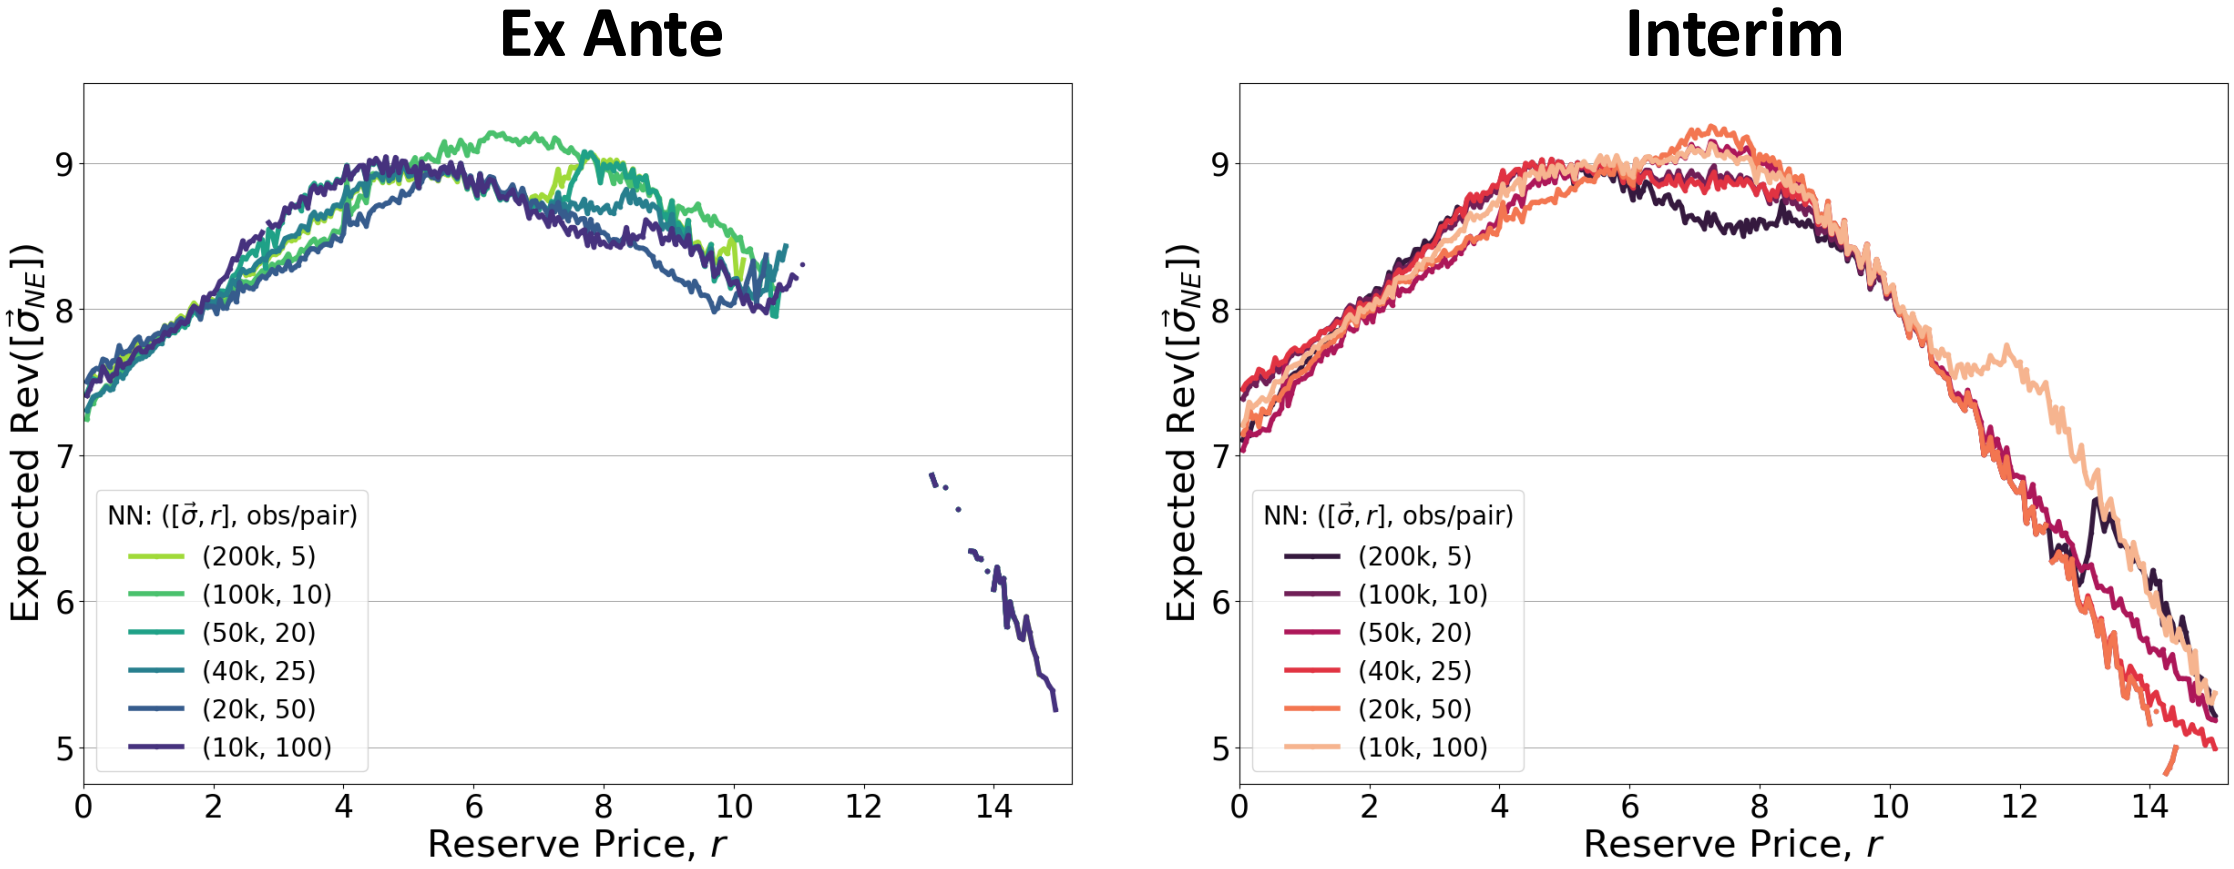

Curves from ex ante models exhibit more discontinuities than interim models.

Q3. How does the learned game-family model enable comprehensive mechanism search?

Outline

What is empirical mechanism design? 1

How can we learn a Bayesian game-family model? 2

How does the learned game-family model enable…

…comprehensive mechanism search? 3

…strategy generation without additional simulation? 4

Piecewise Best-Response Strategies

- Model only covers a limited strategy set
- Want to consider broader set of strategies

How can the interim model approximate piecewise BR strategies without additional simulation?

Piecewise Best-Response Strategies

- Model only covers a limited strategy set
- Want to consider broader set of strategies

How can the interim model approximate piecewise BR strategies without additional simulation?

S	
s_1	Bid θ
s_2	Bid $\theta - 4$

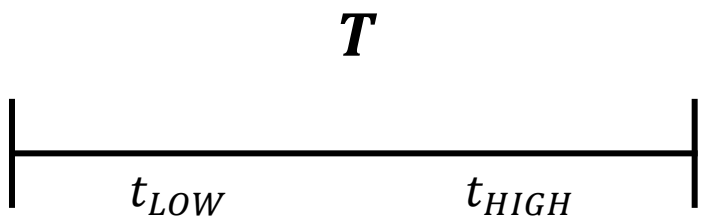

Equilibrium, $\vec{\sigma}_*$:

s_1

Piecewise Best-Response Strategies

- Model only covers a limited strategy set
- Want to consider broader set of strategies

How can the interim model approximate piecewise BR strategies without additional simulation?

S	
s_1	Bid θ
s_2	Bid $\theta - 4$

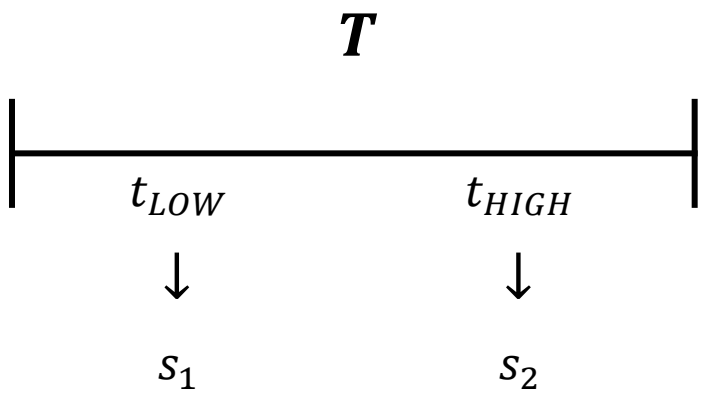

Compute BR conditioned on type

Equilibrium, $\vec{\sigma}_*$:

s_1

Q4. How does the learned game-family model enable strategy generation?

Piecewise Best-Response Strategies

- Model only covers a limited strategy set
- Want to consider broader set of strategies

How can the interim model approximate piecewise BR strategies without additional simulation?

S	
s_1	Bid θ
s_2	Bid $\theta - 4$

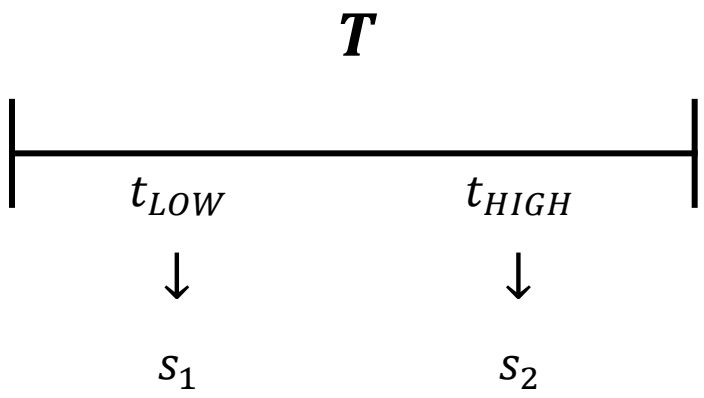

$$\psi = \begin{cases} t_{LOW} \mapsto s_1 \\ t_{HIGH} \mapsto s_2 \end{cases}$$

Q4. How does the learned game-family model enable strategy generation?

Conclusion

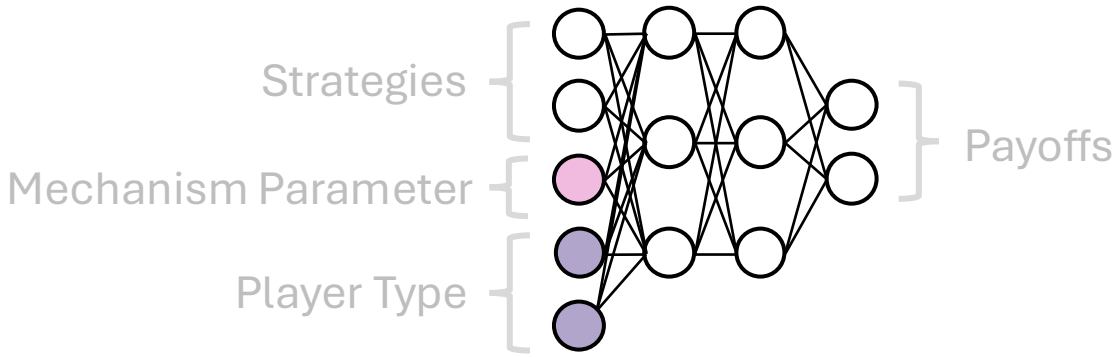

Learning an **interim** Bayesian game-family model enables…

Accurate payoff predictions, *even in extrapolation*.

Comprehensive mechanism parameter optimization.

Beneficial strategy generation, *without additional simulation*.

Thank you!

Paper Access

arxiv.org/pdf/2502.14078

Contact Info

Madelyn Gatchel

University of Michigan

Email: gatchel@umich.edu

Website: madelyngatchel.com

Advisor:

Prof. Michael P. Wellman