# Learning Bayesian Game Families, with Application to Mechanism Design

Madelyn Gatchel & Michael P. Wellman
University of Michigan

1. Game-family learning enables generalization across the parameter space, supporting more effective mechanism optimization.
2. Learned interim models extrapolate more accurately and approximate equilibria more reliably than ex ante.
3. Learned interim models support the generation of beneficial piecewise best-response strategies.

**Link to Paper:**

## Motivation & Methodology

### Empirical Mechanism Design

**Goal:** Find the parameter setting that optimizes the design objective in equilibrium.

**Knowledge Source:** Agent-based simulation model

**Challenge:** Limited resources for simulation queries

### Example: Online Advertising Auctions

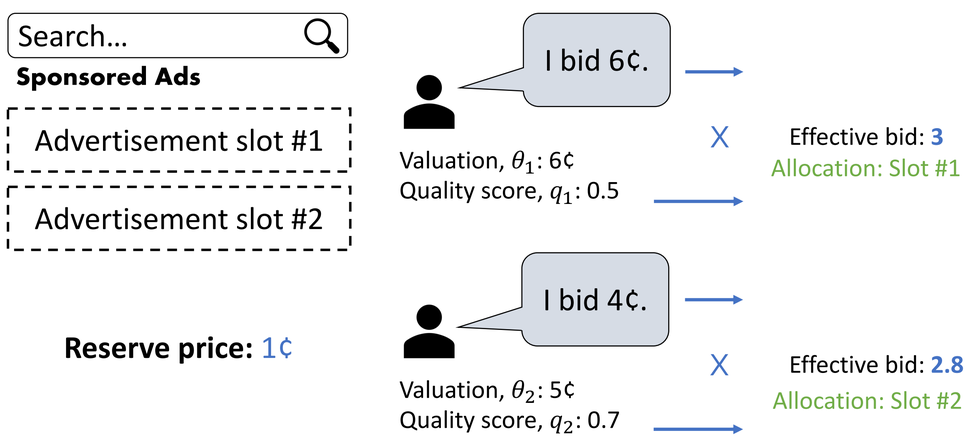


Players have ***private information*** that informs their choice of action and influences their payoff.

**Note:** *In our experiments, the auction is dynamic: advertisers may attempt to revise bids in a second stage, with attempts succeeding probabilistically.*

### Methodology Overview

Extend game-family learning [Gatchel and Wiedenbeck, 2023] to the Bayesian setting, exploiting the type-conditional form of strategies.

## Learning & Analyzing Bayesian Game Families

**Ex Ante Deviation Payoffs**

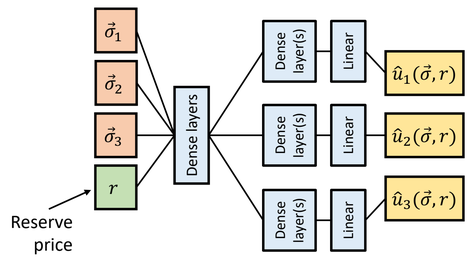


**Interim Deviation Payoffs**

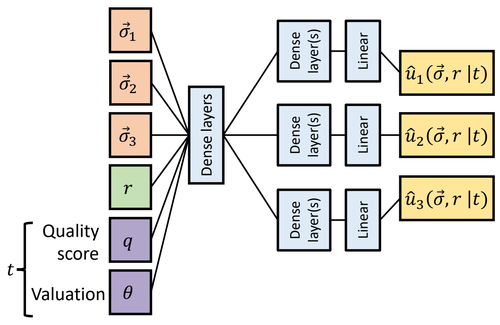


### Model Learning Results

With enough marginalization samples, interim model accuracy matches ex ante.

$r > 8$: Interim but not ex ante models extrapolate well beyond the trained range.

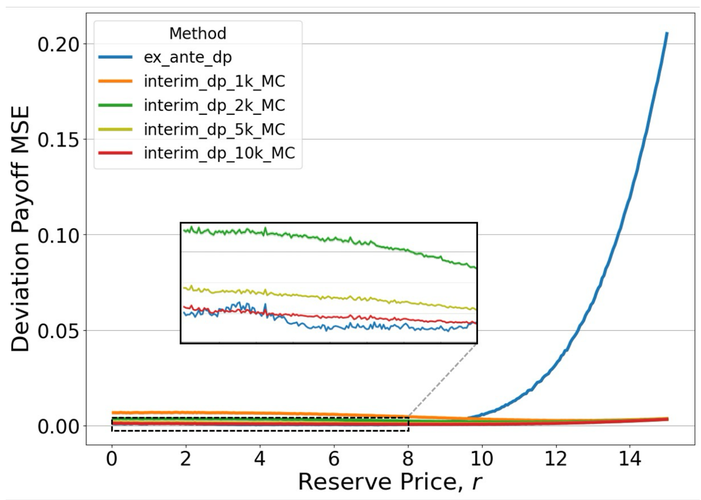


### Equilibrium Approximation Results

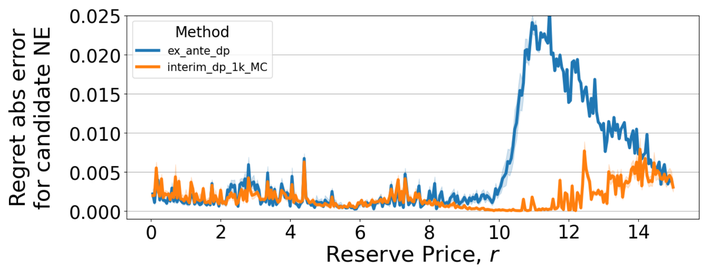


The interim method can better distinguish candidate equilibria from non-convergent mixtures.

## Empirical Mechanism Design (EMD) with Learned Game Families

### Grid Optimization

For each $\Gamma(r)$ and learned model:

- Use learned model to identify equilibria
- Compute expected revenue, averaged across all confirmed equilibria

Curves from (L) ex ante models exhibit more discontinuities than (R) interim models.

(Both) Expected revenue curves from model to model may vary significantly.

### Ex Ante

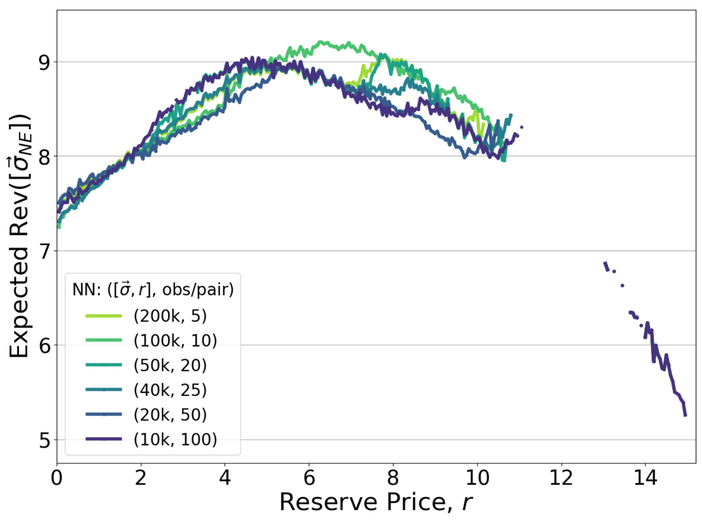


### Interim

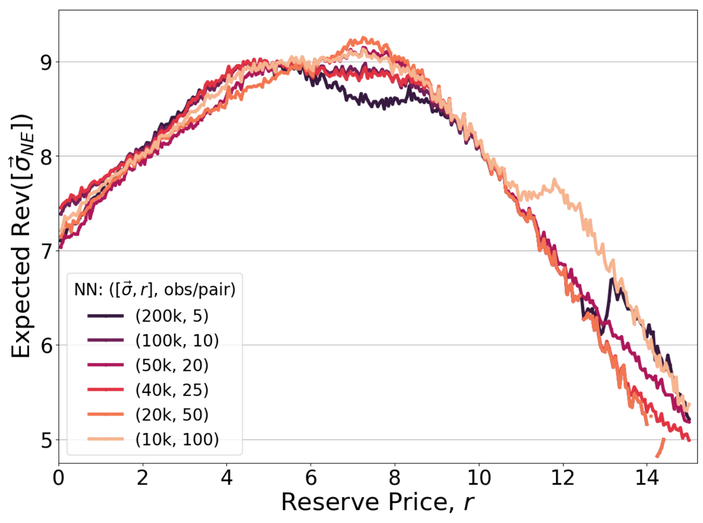


## Type-Conditional Piecewise Best-Response Strategies

### Piecewise Best-Response (BR) Strategies

Motivation:

$S$ is a highly restricted subset of all possible strategies. What if an omitted strategy is a beneficial deviation?

Example:

- $T = \{t_{LOW}, t_{HIGH}\}$

| $S$ | |
|---|---|
| $s_1$ | Bid truthfully |
| $s_2$ | Shade bid |

- Ex ante BNE: $s_1$

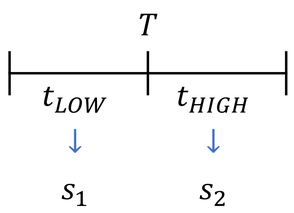


$$\phi = \begin{cases} t_{LOW} \mapsto s_1 \\ t_{HIGH} \mapsto s_2 \end{cases}$$

### Payoff Gain by Playing Piecewise BRs

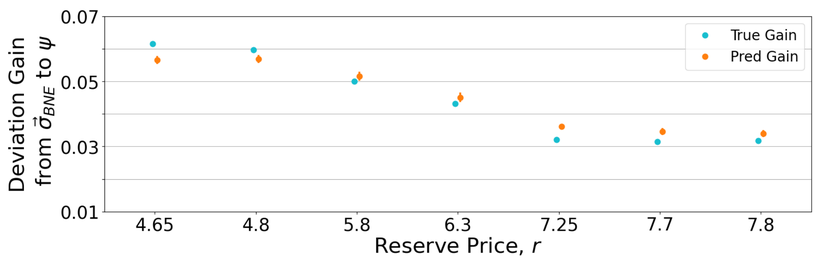


Profitable deviation refutes equilibrium over baseline strategies at candidate mechanism optimum.

### Future Work: Iterative EMD

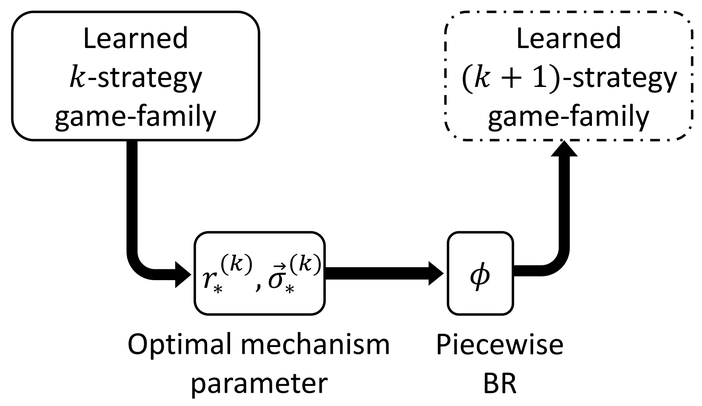


*…repeat until no beneficial piecewise deviation emerges.*