Learning Parameterized Families of Games

Madelyn Gatchel[1] & Bryce Wiedenbeck[2]
[1]University of Michigan, [2]Davidson College

By learning a single model for a parameterized game family, we can:

1. Achieve higher payoff accuracy with less data
2. Conduct new types of robustness and sensitivity analysis
3. Better characterize parameter impact on strategic incentives

Link to Paper:

Motivation & New Terminology

Example: Deciding whether to launch a new product

- Environment Parameter: $V = \Pr(Recession)$

<u>Current Practice</u>
Separate game instances:
$\Gamma(0); \Gamma(.25); \Gamma(.5); \Gamma(.75); \Gamma(1)$

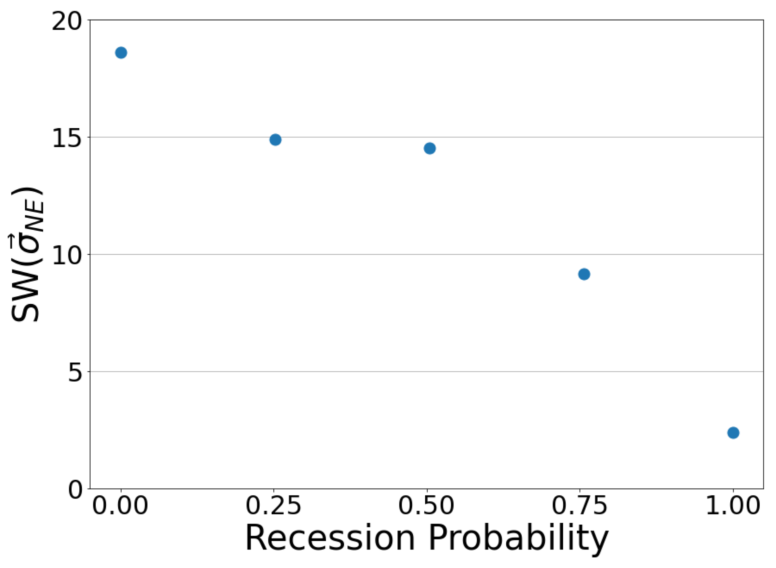

<u>This Work</u>
Parameterized Game Family:
$\{\gamma(V) = \{\Gamma(v): 0 \leq v \leq 1, v \in \mathbb{R}\}$

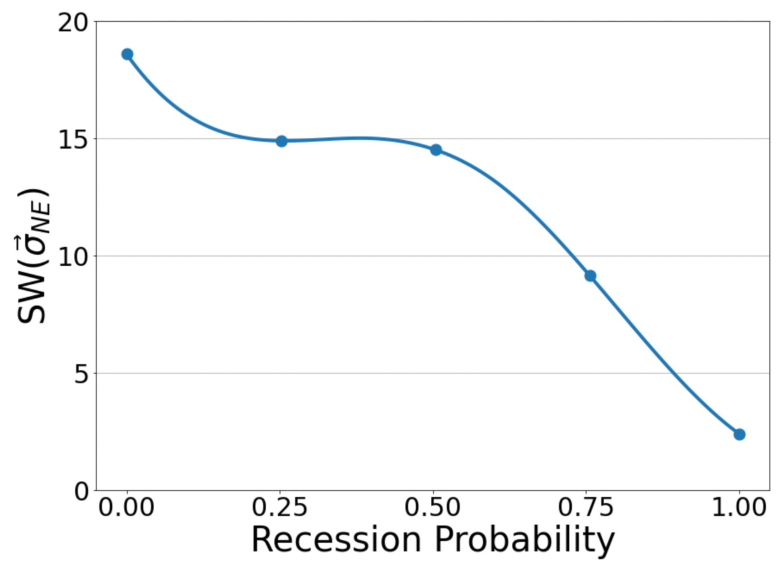

Hypothesis:

Games which are related by a common, ordinal environment parameter likely have related payoffs.

Simulation-Based Games

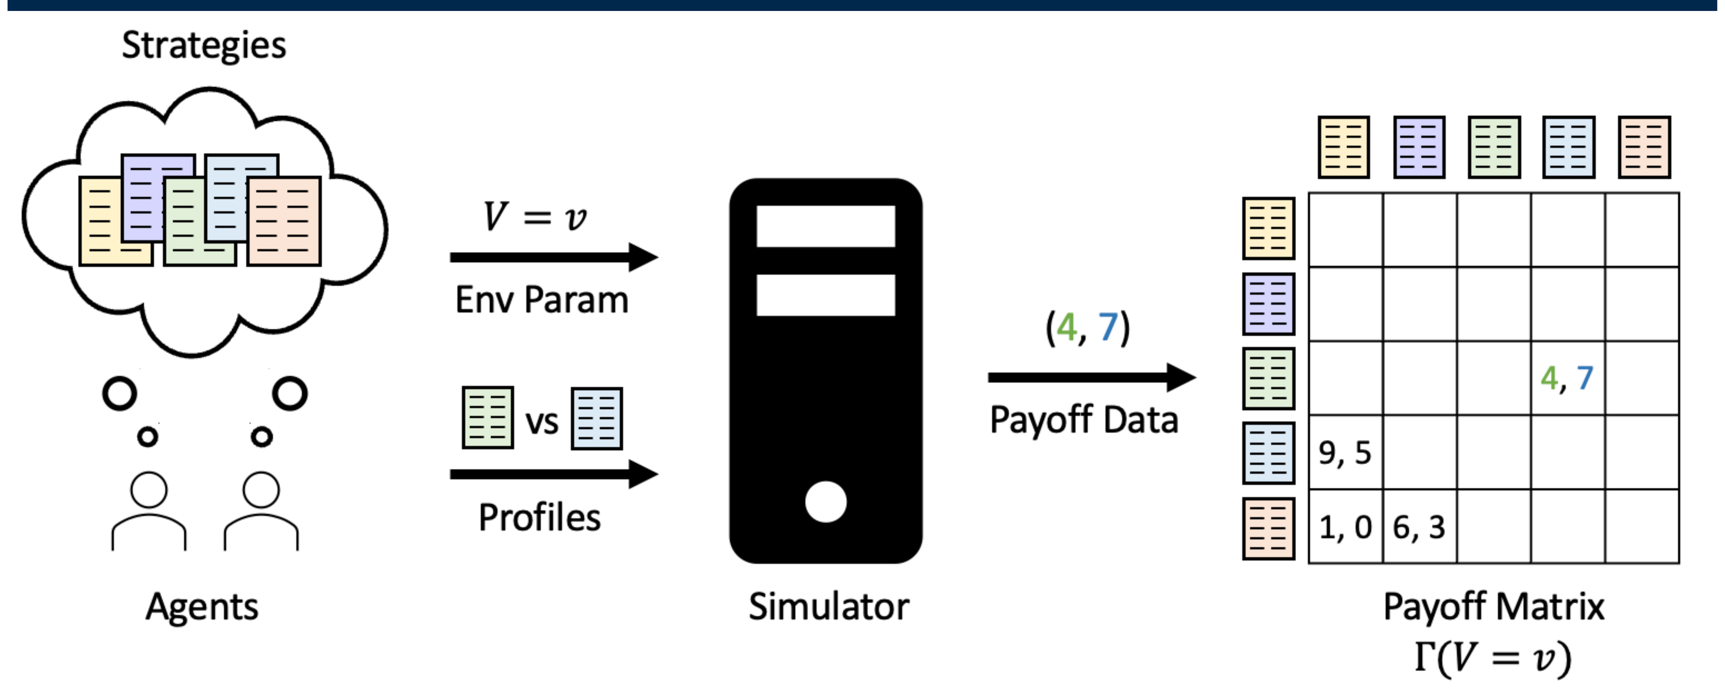

- Nearly all simulation-based games (SBGs) have a relevant environment parameter
  - Ex: number of background traders, asset recovery rate, probability of default
- Past SBG studies construct separate game models for each parameter value
- Player symmetries are common in SBGs, so this work focuses on finding **symmetric ε-Nash equilibria** in symmetric games
- **Deviation payoff:** the expected payoff a player would receive by deviating or changing strategies, given the mixed strategies everyone else is playing

Learning Parameterized Game Families

- We adapt prior work (Sokota et al. AAAI-19) to learn a mapping from symmetric mixed-strategy profiles ***and environment parameters*** to deviation payoffs
- This learned model may be used in an ε-Nash-finding algorithm to find equilibria for the parameterized game family

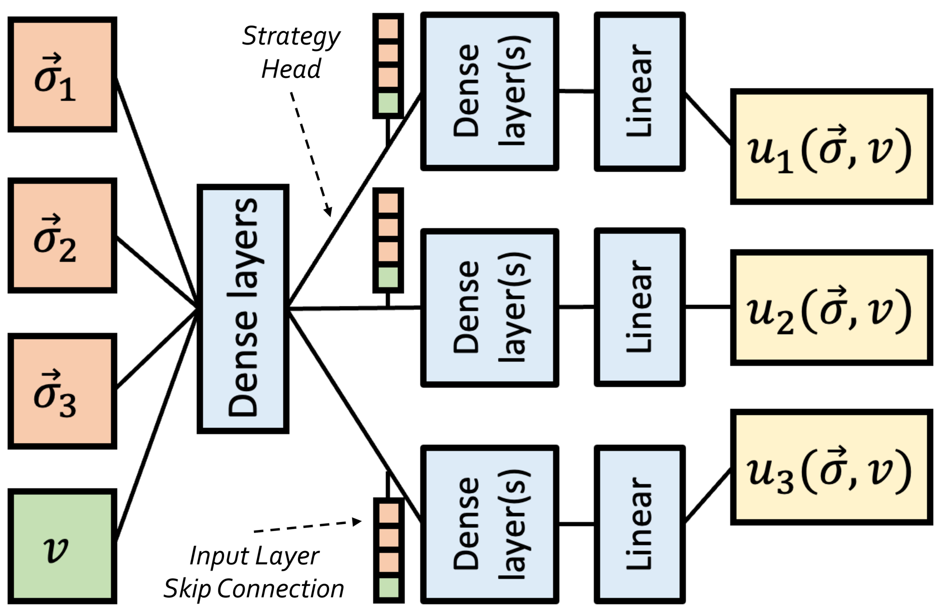

Comparison to Existing Work

***Fixed-Parameter Learning (FPL):** a separate learned model for each game instance*
***Variable-Parameter Learning (VPL):** a single learned model for entire game family*

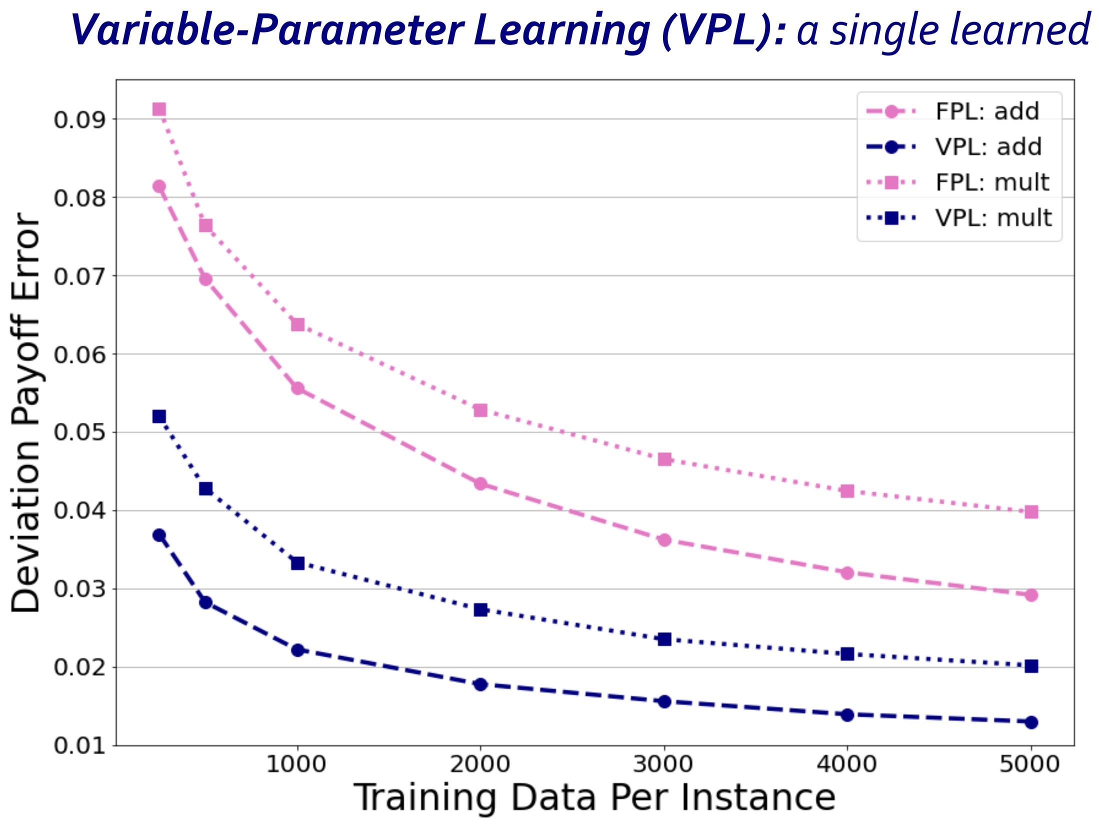

Varied Parameter
Number of players: [90, 100]

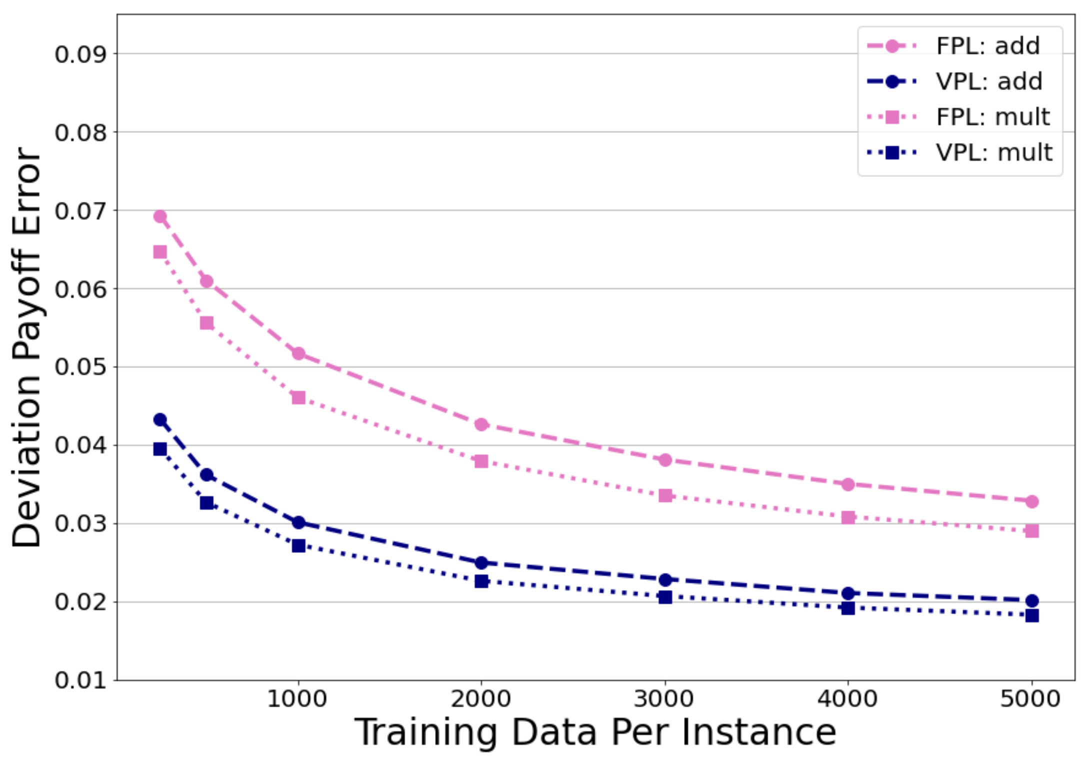

Varied Parameter
Erdős–Rényi Probability Threshold: [.15, .25]

For both discrete and continuous parameters, a single VPL model outperforms a collection of FPL models given identical training data.

Analyzing Parameterized Game Families

Identifying Robust Equilibria

Comparison of two regret robustness metrics on a random 3-strategy game family with 50 to 100 players:

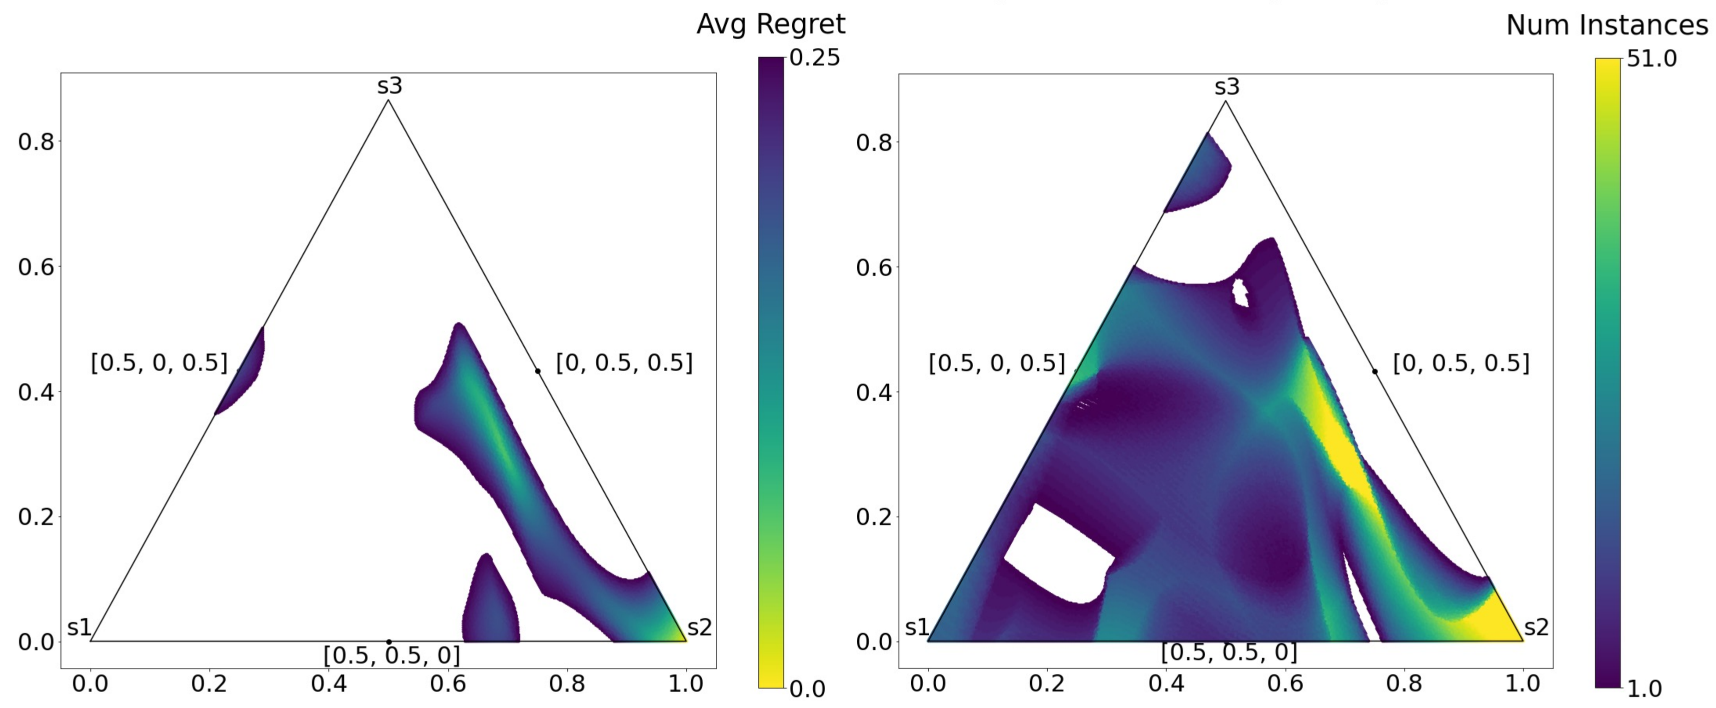

- Each point in the simplex corresponds to a symmetric mixed strategy
- White points correspond to mixtures that are not robust equilibria
- Brighter points correspond to mixtures that are considered more robust

Parameter Sensitivity Analysis

Approximate equilibrium probabilities for each strategy plotted as a function of the number of players for a random 5-strategy game family:

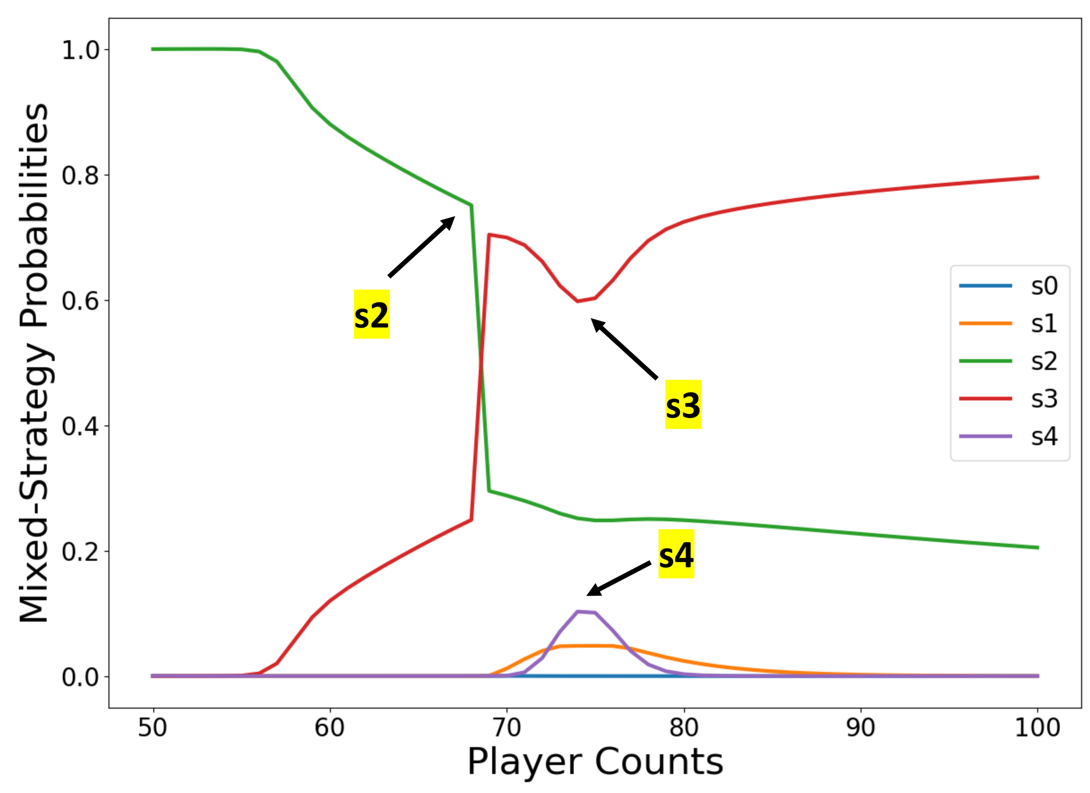

- This game family demonstrates the importance of analyzing the entire parameter space as opposed to "representative" parameter values (e.g., 50, 60,…,100)

Our single learned model allows for more tractable analysis of the full parameter space and enables new types of analysis.