

Learning Parameterized Families of Games

Madelyn Gatchel
University of Michigan
Ann Arbor, USA
gatchel@umich.edu

Bryce Wiedenbeck
Davidson College
Davidson, USA
brwiedenbeck@davidson.edu

ABSTRACT

Nearly all simulation-based games have environment parameters that affect incentives in the interaction but are not explicitly incorporated into the game model. To understand the impact of these parameters on strategic incentives, typical game-theoretic analysis involves selecting a small set of representative values, and constructing and analyzing separate game models for each value. We introduce a novel technique to learn a single model representing a family of closely related games that differ in the number of symmetric players or other ordinal environment parameters. Prior work trains a multi-headed neural network to output mixed-strategy deviation payoffs, which can be used to compute symmetric ϵ -Nash equilibria. We extend this work by making environment parameters into input dimensions of the regressor, enabling a single model to learn patterns which generalize across the parameter space. For continuous and discrete parameters, our results show that these generalized models outperform existing approaches, achieving better accuracy with far less data. This technique makes thorough analysis of the parameter space more tractable, and promotes analyses that capture relationships between parameters and incentives.

KEYWORDS

Simulation-Based Games; Equilibrium Computation; Deep Learning

ACM Reference Format:

Madelyn Gatchel and Bryce Wiedenbeck. 2023. Learning Parameterized Families of Games. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Game-theoretic analysis uses mathematical models to investigate incentives in multi-agent systems. Like all forms of mathematical modeling, this relies on numerous simplifying assumptions, but the constraints that arise from describing utilities and computing equilibria uniquely complicate game-theoretic modeling and analysis. In particular, the unobservability of utilities and out-of-equilibrium counterfactuals contributes to a preference in the field for seeking generalized insight from abstract games rather than placing too much trust in precisely tailored numerical models. One consequence of this imbalance is that game theory has, to its detriment, thus far derived less benefit than other fields from advances in data science, simulation, and machine learning. In this paper, we introduce a machine-learning-based approach to game-theoretic analysis that helps to bridge the gap between theoretical and numerical methods.

In pencil-and-paper approaches to game theory, researchers often analyze abstract classes of games and describe general properties of their equilibria. In contrast, existing computational tools typically focus on representing and solving a particular numerical instantiation, what we call a *game instance*, and may therefore struggle to demonstrate robust results or provide adequate comparative analyses. Our new approach uses a neural network to represent a parameterized *family* of game instances, enabling a rich variety of computational methods which generalize game-theoretic modeling and analysis across some space of environment parameters. By learning parameterized families, we provide a tractable computational model that can capture a greater portion of the expressiveness that would otherwise be the domain of purely theoretical analysis.

This work is motivated by the setting of simulation-based games [22] (SBGs), where a multi-agent interaction of interest is modeled via an agent-based simulation, and data from those simulations is used to construct a normal-form game. In typical SBGs the underlying agent-based model has many free parameters. Examples include the number of background traders in a study of latency arbitrage by high-frequency traders [20], the recovery rate in a model of debt consolidation among financial institutions [13], or the probability of defaults in a simulation of credit network liquidity [2]. In some cases, appropriate values for environment parameters can be chosen by empirical validation, but much more often, several possible settings for a parameter are plausible, and may require deliberate exploration. In each example above, distinct normal-form game models were constructed from separate sets of simulation data and analyzed independently for multiple values of the given parameter. Further, all of these simulation-based models have many other parameters of potential strategic relevance, and it seems likely that they were under-explored due to this burden of constructing entirely new SBGs.

Our method builds on the work of Sokota et al. [14], who demonstrated that a neural network mapping symmetric mixed strategies to deviation payoffs—the expected utilities for unilateral deviators—could be learned from sampled payoff data and could stand in for a normal-form payoff matrix in computing ϵ -Nash equilibria. Fundamentally, this is made possible by player symmetries in the environment, which are exhibited by nearly all large SBGs. Our core insight is that this approach can extend far beyond representing a single normal-form game by making other parameters of the studied interaction endogenous to the learned model. We show that by training a neural network to map symmetric mixed strategies *and environment parameters* to deviation payoffs, we can produce a single model that generalizes over a family of games, which can dramatically reduce the sample complexity of certain comparisons and enables other altogether new analyses. Our results show that this generalized model achieves better accuracy than existing approaches, using significantly less data.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2 RELATED WORK

This paper extends a long line of work on improving methods for analyzing simulation-based games, and connects with a range of studies applying machine learning techniques to game theory. It also addresses a clear need for better approaches to generalizing across environment parameters in simulation-based games.

2.1 Simulation-Based Games

In a typical simulation-based game, an agent-based model is used to capture the dynamics of the environment being studied, and different possible behaviors of the agents in that environment can be simulated [22].¹ An analyst implements a set of strategies that could govern agent behavior, and then a single run of the simulator assigns a strategy to each agent and produces a noisy sample of the payoffs for that particular profile of strategies. We take as given that the analyst has access to such an agent-based model, or some other function that maps an arbitrary pure-strategy profile to a (possibly noisy) sample of each player’s payoff. Of note, purely observational payoff data does not in general suffice, because our method requires sampling specific profiles on demand.

In principle, these simulations could be used to fill out a normal-form payoff matrix by sampling every possible profile enough times to produce accurate payoff estimates for that matrix cell; this approach was used by early simulation-based game studies [9]. However, this exhaustive sampling approach rapidly becomes infeasible as the model grows, because the number of profiles in a symmetric game grows combinatorially, and for any variation in environment parameters, an entirely new payoff matrix must be constructed.

2.2 Varying Environment Parameters

2.2.1 Simulation-Based Games. Essentially all simulation-based games employ agent-based models where certain parameters of the simulated environment can be varied, and many examples in the literature include variations over some parameter prominently in their analysis. For example, Wah and Wellman [20] construct an SBG to analyze the effects of latency arbitrage in financial markets. The authors analyze the same game with a variable number of background traders (24, 58, 238), but construct three independent model instances via player reduction [26] and analyze each one separately because the number of players is a parameter outside the scope of their game model. For each instance, they find role-symmetric equilibria in the game and then evaluate background-trader surplus and latency arbitrageur profit. In another study [19], the same authors compare the welfare of traders in frequent call markets versus continuous double auctions. They analyze four separate instances, varying the number of agents (8, 14, 42) and the mean-reversion parameter ($\kappa = 0.05, 0.01$).

Mayo and Wellman [13] use multi-agent simulation to model portfolio compression of debt cycles in a financial network as a strategic decision among firms. They define the recovery rate, α , as “the fraction of assets an insolvent node is able to recover and use to pay back creditors, with the remaining assets being lost to default.” They construct separate game-theoretic models for $\alpha \in \{0, 0.1, 0.3, 0.5, 0.7, 1.0\}$, and study the effect of the recovery rate on the incentives to compress debt cycles.

¹Also known as empirical games [21] or black-box games [28].

Other simulation-based game examples include Wellman et al. [23], who evaluate complex network routing protocols, Dandekar et al. [2], who study the formation of credit networks, and Mayer et al. [12] who explore prices in simultaneous sealed-bid auctions. Wellman et al. analyze separate instances that vary the number of non-attacking nodes (clients, ISPs, roots, and servers). Dandekar et al. analyze instances that vary the probability that a debtor will default, and the availability of information about those defaults. Mayer et al. vary buyer’s valuation distributions and degree of substitutability among the resources being sold. In all of these cases, the authors had to build separate models from distinct data sets for each game instance.

2.2.2 Other Variable-Parameter Games. Game-theoretic analyses that generalize over variable environment parameters are also common outside of simulation-based game settings. Tuffin and Maillé [17] model TCP sessions with variable cost and between 2 and 10 users. Thompson et al. [16] model voting in plurality elections with between 3 and 96 voters, and find different trends with even or odd numbers of players. Fatima [4] compares sequential and simultaneous auctions by varying the number of objects and bidders. Hanaki and Rouchier [6] study information in Cournot competition and vary the number of ignorant and informed agents. In all of these examples, the environment parameters are parameters of the model, meaning the authors have to construct separate game instances for each distinct parameter combination they want to consider. Ling et al. [11] study games with varied parameters, but focus on two-player zero-sum games and address parameter uncertainty from the agent’s perspective as opposed to the analyst’s.

2.3 Methods for Simulation-Based Games

In the vast majority of simulation-based game studies, sampling from the agent-based simulator is the primary bottleneck, so the literature on simulation-based game theory includes a wide range of approaches for improving data efficiency. Some tools employ statistical methods such as control variates [9] and bootstrapping [25] to reduce the number of samples required to confidently estimate a profile’s payoffs. Others focus on iteratively exploring the space of possible strategies [2, 8] to identify equilibria from partial payoff matrices. Games with a large number of players have been tackled via player-reduction approaches [24, 26] which analyze a small game that is hoped to roughly approximate a larger one. Various attempts [1, 7] have been made to derive theoretical bounds on sample complexity. We prioritize the practicality of our approach for SBGs over these sorts of theoretical guarantees that have limited applicability to large and complex game families.

Several machine learning methods have also been proposed for simulation-based games. Some of these aim to identify compact structure, such as role symmetry [5], a graphical game [3] or both [10], underlying a particular data set. But they can only help in the event that such structure is present but not known in advance. Others aim to learn models that can replace the payoff matrix data structure, such as Vorobeychik et al. [18] who propose regression approaches for learning payoff functions in games with real-valued strategies, and Wiedenbeck et al. [27] who use Gaussian process regression to learn the utility function in symmetric games with a large number of players.

Most directly relevant is the work of Sokota et al. [14], who use a neural network to learn a mapping from role-symmetric mixed-strategy profiles to deviation payoffs. They exploit the player symmetries common in simulation-based settings to learn the deviation payoff function in games with a large number of players. This learned deviation payoff function is used in Nash-finding algorithms to identify role-symmetric approximate equilibria in SBGs without constructing an explicit payoff table. Our approach extends that of Sokota et al. [14], adding environment parameters as an input to the network and allowing us to model parameterized game families.

3 BACKGROUND

A normal-form game consists of a set of players, indexed by i , a set of strategies S_i for each player, and a utility function for each player: $u_i : \prod_i S_i \rightarrow \mathbb{R}$. A game is *symmetric* if any permutation of the player set yields the same game, and *role-symmetric* if there exists a non-trivial partition of the player set into roles where permutations within a role yield the same game. Our learning technique relies on symmetries, and in this paper we focus on symmetric games, but all our methods also apply under role symmetry.

Exploiting symmetry, we can rewrite the model to be independent of players' identities. A symmetric normal-form game, Γ , has a number of players p and one strategy set S , indexed by j , common to all players. The utility function shared by all players depends only on an individual's choice of strategy and the aggregate choice of strategies by others, so it can be represented by a separate payoff function u_j for each strategy $j \in S$ that maps *opponent profiles* to utilities. An opponent profile \vec{s} is an integer vector of dimension $|S|$ representing the number of other players selecting each strategy. The set of all opponent profiles \vec{S} is thus an integer simplex:

$$\vec{S} = \left\{ \vec{s} \in \mathbb{Z}^{|S|} : \vec{s}_j \geq 0, \sum_j \vec{s}_j = p - 1 \right\}$$

and the payoff function for each strategy maps this integer simplex to the reals: $u_j : \vec{S} \rightarrow \mathbb{R}$.

The main solution concept used to predict behavior in normal-form games is the *mixed-strategy Nash equilibrium*. A mixed strategy σ is a probability distribution over a player's strategies, and a Nash equilibrium is a mixed strategy for each player where no player can increase their expected utility by unilaterally deviating to any other strategy. In symmetric and role-symmetric games, analysts are often interested in symmetric or role-symmetric Nash equilibria, where all identical players play the same mixed strategy. In a symmetric game, a symmetric mixed-strategy profile $\vec{\sigma}$ is a probability distribution over S , according to which all p players (or sometimes all $p - 1$ opponents) independently randomize their actions. Note that σ and $\vec{\sigma}$ have the same dimension, but the vector-accent emphasizes that in a symmetric profile, many players are following the same mixed strategy.

A symmetric mixture $\vec{\sigma}$ is a Nash equilibrium if a unilateral deviator's expected utility for playing any pure strategy is no higher than their expected utility for playing $\sigma = \vec{\sigma}$, when all opponents play according to $\vec{\sigma}$. To express this mathematically, we extend our notation for utility functions to map symmetric mixed-strategy profiles to *deviation payoffs*—the expected utility of playing strategy j against opponents playing $\vec{\sigma}$. This expectation $u_j(\vec{\sigma})$ is a

probability-weighted sum of $u_j(\vec{s})$ over all opponent profiles:

$$u_j(\vec{\sigma}) = \sum_{\vec{s} \in \vec{S}} \Pr(\vec{s}|\vec{\sigma}) u_j(\vec{s}) \quad (1)$$

We define $u : \Delta^{|S|} \rightarrow \mathbb{R}^{|S|}$ (with no subscript) as a function that produces a vector of the deviation payoff for each strategy, given a symmetric mixed-strategy profile $\vec{\sigma}$ from the $|S|$ -dimensional probability simplex $\Delta^{|S|}$. We can then express the Nash equilibrium condition in terms of deviation payoffs: a symmetric mixed-strategy profile $\vec{\sigma}$ is a Nash equilibrium if $\forall j \in S$, we have $u_j(\vec{\sigma}) \leq \sigma \cdot u(\vec{\sigma})$, where the dot product of σ and u gives the expected utility of playing the same mixed strategy as everyone else.

It is also useful to define *regret*, the maximum gain achievable by deviating from $\vec{\sigma}$ to any strategy $j \in S$, as:

$$\epsilon(\vec{\sigma}) = \max_{j \in S} u_j(\vec{\sigma}) - \sigma \cdot u(\vec{\sigma})$$

A Nash equilibrium has $\epsilon(\vec{\sigma}) = 0$, but in practice, analysts generally identify ϵ -equilibria—mixtures with $\epsilon(\vec{\sigma}) \leq \epsilon$, for suitably small ϵ .

3.1 New Terminology

The premise of this work is that for settings which vary an environment parameter, we can construct and learn a single game-theoretic model and then perform analyses to characterize trends for the entire game family. We capture this idea by first defining an *instance*, which is a normal-form game Γ with fixed environment parameters. Next, a *parameterized game family* is a set of game instances that are related by one or more ordinal parameters of the environment. Formally, we define a parameterized game family as $\gamma(V) = \{\Gamma(v) : v \in \mathbb{R}\}$ for some environment parameter variable V . Many environments will have multiple variable environment parameters of interest, and our method can simultaneously generalize across all of them, but for our initial validation we focus on parameterized game families with one varied parameter.

A symmetric mixed-strategy profile $\vec{\sigma}$ can be played by any number of symmetric players (regardless of environment parameter values), but because the utility functions vary with the environment parameter V , we need to specify $V = v$ to compute regrets and identify equilibria. The deviation payoff $u_j(\vec{\sigma}, v)$ for $\Gamma(v) \in \gamma(V)$ is the expected utility of a unilateral deviator playing strategy $j \in S$ while all other opponents randomize according to $\vec{\sigma}$, and is still calculated by equation 1. In a parameterized game family $\gamma(V)$, the regret $\epsilon(\vec{\sigma}, v)$ of a given mixture, as well as other functions of equilibrium, varies as a function of V , and as a result, a mixture may be an ϵ -Nash equilibrium for certain values of V and not others.

4 LEARNING GAME FAMILIES

We hypothesize that game instances which are related by a common environment parameter likely have related payoff and deviation payoff functions. With this hypothesis, we learn a mapping from symmetric mixed strategies and environment parameter V to deviation payoffs in a parameterized game family. Our approach generalizes that of Sokota et al. [14] by adding the environment parameter as an input dimension of the model. This enables a *single model* to learn patterns which generalize across the parameter space. In Section 4.1 we describe how to construct, train and refine this variable-parameter learning model.

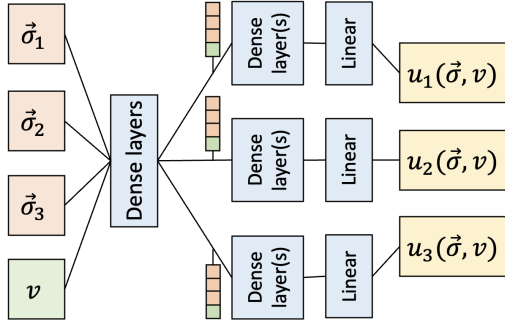


Figure 1: The neural net for variable-parameter learning (VPL) maps symmetric mixtures $\vec{\sigma}$ and environment parameter values $V = v$ to deviation payoff vectors $u(\vec{\sigma}, v)$.

4.1 Approximating Deviation Payoffs

We use data from an agent-based simulator or other black-box game model to train a neural network on the mapping $u(\vec{\sigma}, v)$ from a symmetric mixed-strategy profile $\vec{\sigma}$ and environment parameter v to the vector of deviation payoffs $[u_j(\vec{\sigma}, v) : j \in S]$. We adapt the technique developed by Sokota et al. [14] by adding the environment parameter V as an input dimension. Additionally, in the variable-parameter learning (VPL) model each strategy head has a skip connection from the input layer. We verified experimentally (not shown) that these skip connections slightly improved VPL performance. Figure 1 shows a general architecture for our variable-parameter neural network on a hypotheticalal game with 3 strategies. Note that $\vec{\sigma}_j$ denotes the probability that the j th strategy is played according to $\vec{\sigma}$ and that $u_j(\vec{\sigma}, v)$ denotes the deviation payoff of strategy j in instance $V = v$, the expected utility a player receives by playing strategy j when all opponents follow $\vec{\sigma}$. Observe the added input dimension for the environment parameter and the input layer skip connections for each strategy head.

The variable-parameter learning procedure is described in Algorithm 1. For input samples, we draw mixtures $\vec{\sigma} \sim \text{Dir}(\vec{\alpha})$ from a Dirichlet distribution with $\vec{\alpha} < 1$; this helps to ensure accurate estimates for small-support profiles. To generate the associated environment parameter values, we randomly select values v uniformly across the entire parameter space. For neighborhood sampling, near each returned candidate Nash equilibrium $\vec{\sigma}^*$ with associated environment parameter value v^* , we draw neighborhood mixtures $\vec{\sigma}' \sim \text{Dir}(\vec{\alpha} = \omega_{\vec{\sigma}} \cdot \vec{\sigma}^* + 1)$ where $\omega_{\vec{\sigma}} \gg 1$. We draw *normalized* neighborhood environment parameter values v' from a Beta (or other) distribution centered around v^* .

4.1.1 Targeted resampling in parameter space. Sokota et al. [14] demonstrated the importance of targeted resampling in the strategy-space neighborhood of candidate equilibria. Because of our hypothesis that neighboring game instances have similar payoff functions (and therefore equilibria), we also conduct targeted resampling in the neighborhood of the associated environment parameter value.

4.1.2 Avoiding Duplicate Queries. Unlike past approaches to simulation-based game analysis [14, 27], we avoid deliberately repeating queries to the simulator on the same profile. While this would

Algorithm 1 Finding ϵ -equilibria for a parameterized game family.

APPXNASHEQ($\gamma(V)$, initQueries, resampQueries, numIters):

```

([ $\vec{\sigma}$ ], [ $v$ ])  $\leftarrow$  initSamples( $\gamma(V)$ , initQueries)
[ $\vec{s}$ ]  $\leftarrow$  sampleOppProfiles([ $\vec{\sigma}$ ], [ $v$ ])
[ $\vec{u}$ ]  $\leftarrow$  simulator([ $\vec{s}$ ], [ $v$ ])
data  $\leftarrow$  ([ $\vec{\sigma}$ ], [ $v$ ], [ $\vec{u}$ ])
regressor.fit(data)
repeat
  ([ $\vec{\sigma}^*$ ], [ $v^*$ ])  $\leftarrow$  findNash(regressor) // parallelized
  ([ $\vec{\sigma}'$ ], [ $v'$ ])  $\leftarrow$  sampleNbhd([ $\vec{\sigma}^*$ ], [ $v^*$ ], resampQueries)
  [ $\vec{s}'$ ]  $\leftarrow$  sampleOppProfiles([ $\vec{\sigma}'$ ], [ $v'$ ])
  [ $\vec{u}'$ ]  $\leftarrow$  simulator([ $\vec{s}'$ ], [ $v'$ ])
  data  $\leftarrow$  data + ([ $\vec{\sigma}'$ ], [ $v'$ ], [ $\vec{u}'$ ])
  regressor.fit(data)
until numIters
return ([ $\vec{\sigma}^*$ ], [ $v^*$ ])  $\leftarrow$  findNash(regressor)

```

increase the accuracy of payoff estimates *for that profile*, we find that—given a limited sampling budget and the tiny fraction of mixtures being sampled in a typical game—additional queries are better allocated for a wider range of profiles.

4.1.3 Intermediate Regret Check. Because regret varies as a function of parameter V , a symmetric mixed strategy may be an approximate equilibrium in one instance but not others. As a result, we propose two modifications during the refinement process: filtering out returned candidate Nash that fall below some intermediate regret threshold for the associated instance, and having no intermediate regret filter. The intuition is that we do not want to use too many more queries for areas of the simplex which are less likely to contain ϵ -Nash equilibria. However, we also not want to have too strict of an intermediate regret filter such that the refinement process favors some game instances more than others. We evaluate these two algorithm variants in Section 5.2.

4.1.4 Parallelized Nash-Finding. After the model is trained, we run a Nash-finding algorithm using the learned deviation payoff function. This Nash-finding algorithm is run in parallel by using the neural network’s batch processing to compute deviation payoffs for many mixtures at once, even for different game instances. Given a matrix whose columns correspond to $(\vec{\sigma}, v)$ vectors, the neural network can output a matrix of deviation payoff estimates in time proportional to the size of the network as opposed to the size of the exponentially larger underlying normal-form game family. This matrix of deviation payoffs is then used to update the mixtures, also in parallel, as specified by the particular Nash-finding algorithm. The Nash-finding algorithm returns a matrix of candidate equilibria and associated parameter values. To our knowledge, this is the first algorithm in the literature that can compute approximate Nash equilibria in parallel for an entire game family.

4.1.5 Analyzing the Learned Model. After the last iteration, the analyst can use the newly refined model to construct a set of candidate equilibria for the parameterized game family $\gamma(V)$. First, the analyst could run the Nash-finding algorithm using the learned deviation

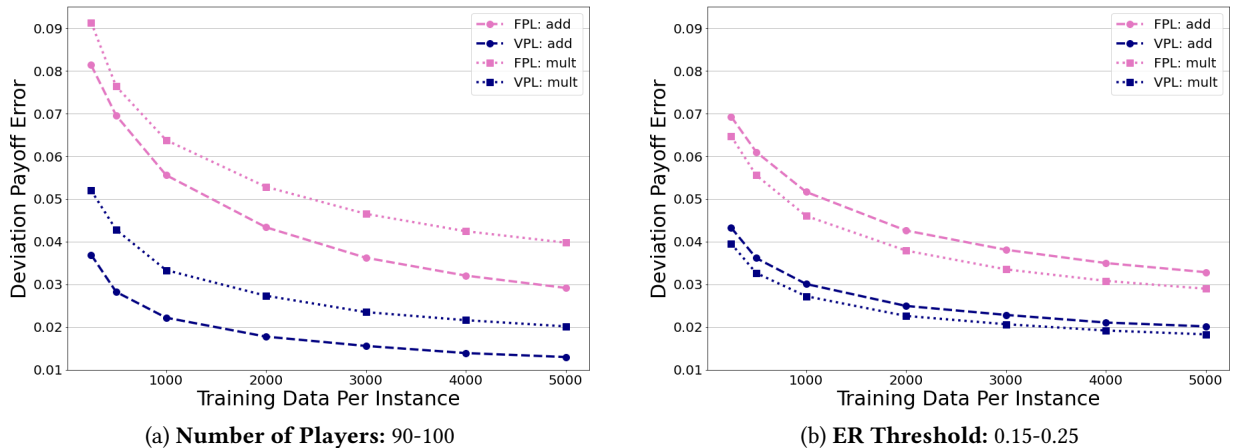


Figure 2: For both of the most difficult cases: (a) variable number of players and (b) a continuous environment parameter, a single Variable-Parameter Learning (VPL) model outperforms a collection of Fixed-Parameter Learning (FPL) [14] models on both classes of random games given identical training data. This suggests that the VPL model is more data efficient, utilizing payoff information from neighboring game instances to approximate deviation payoffs with greater accuracy.

payoff function to find approximate equilibria for each instance (as described above). Then they could apply any of the techniques proposed in Section 6 that use this set of candidate equilibria to perform analysis that generalizes across the parameter space.

5 EXPERIMENTS

The goal of our experiments is to validate that our variable-parameter learning (VPL) model can adequately learn the deviation payoff function for parameterized families of games. Since our method is the first to handle an entire game family, and prior work [14] is the current best approach for learning a single instance, our primary experiments compare our VPL model against a collection of fixed-parameter learning (FPL) models. The next experiment validates our approach to adapting the model refinement from [14] to the variable-parameter setting. Additional experiments with wider ranges of player counts (e.g., 50-100) and Erdős-Renyi threshold values (e.g., 0-1) confirm that VPL models continue to perform well, even as the parameter range width increases.

To serve as proxy for simulator data in all experiments, we generate random symmetric additive (or multiplicative) polynomial-sine bipartite action-graph games with additive function nodes (BAGGFN) as used in prior literature [14, 27]. These random games have complex but learnable payoff functions, particularly compared to common game distributions in related literature: substantially more challenging than congestion games, but much more structured than uniform random games. Refer to Appendices A and B for more details on random game generation and justification.²

Our experiments vary two key parameters of these types of games, chosen to represent some of the most relevant and challenging simulation-based game settings. First, we vary the number of players in the game. Because FPL and all previous methods are based around approximating a normal-form game, the number of players determines the dimension of the underlying data, meaning that

changing the number of players always requires starting over from scratch with zero carry-over of data. Second, we vary a continuous parameter: the Erdős-Renyi probability threshold used to determine edge inclusion in the underlying bipartite action-function graph.

5.1 Comparison to Existing Work

Because the number of simulator queries is most often the bottleneck in SBG studies, we first use *identical training sets* when comparing VPL and FPL techniques. In particular, we compare a single VPL model trained on the entire parameter space against the performance of several FPL models, one trained for each game instance. Thus the VPL model trains on the entire training set and the training dataset is partitioned so that each FPL model gets the appropriate data per instance. We evaluate the performance of the two techniques measured by the deviation payoff accuracy relative to the underlying game family. We assess model performance on 100 games from two classes: additive and multiplicative polynomial-sine games, each with 5 strategies. For a given game class, both models are evaluated on the same 100 games with identical training sets for each game. For example, if the total amount of data is 55,000 then the VPL model is trained with 55,000 training examples and each of the 11 FPL neural networks is trained with 5,000 training examples.

The FPL architecture is identical to that described in [14], and consists of a multi-headed neural network with 128-, 64-, and 32-node dense hidden layers and a head for each strategy with a 16-node dense layer followed by a linear layer. The VPL architecture consists of a multi-headed neural network with 256-, 128- and 64-dense hidden layers and a head for each strategy with a 32-node dense layer followed by a linear layer. Refer to Figure 1 to see the general VPL architecture visualization. The hyperparameters for the two models were optimized separately, and for expediency, hyperparameters were tuned using other random BAGGFN games. In a simulation-based game, such tuning would be performed on a hold-out set, which is validated in [14]. All 11 FPL neural networks use

²All appendices may be found in the arXiv version: <https://arxiv.org/pdf/2302.12969.pdf>

identical hyperparameters, but tuning them independently yields negligible improvement.

For each instance, we measure network performance on 495 mixtures corresponding to points on a lattice that evenly covers the simplex. We evaluate both models on all 11 instances of the parameterized game family. For each mixture we compute the error between the predicted deviation payoff and the ground truth deviation payoff for each strategy, and compute the mean absolute error across the 5 strategies and average across mixtures. These errors are computed on normalized deviation payoffs, so the average MAE tells us approximately what percentage we can expect our learned deviation payoffs to differ from the calculated deviation payoffs for each strategy. For all experiments we compute 95% confidence intervals for the deviation payoff error for each instance. For each model, we use the mean of MAEs from 495 mixtures for all 100 randomly generated games. Note that many of the confidence intervals are too small to show up in the plot.

Figure 2a shows FPL and VPL performance on additive and multiplicative polynomial-sine games with 5 function nodes, an Erdős-Renyi probability threshold of 0.2, and 90 to 100 players. Observe that increasing the training data size by the same amount for both models results in a greater increase in accuracy for FPL, but VPL still achieves better overall accuracy. This suggests that VPL is able to utilize neighboring instance payoff data to better approximate deviation payoffs, even with a smaller amount of training data.

Figure 2b shows FPL and VPL performance on additive and multiplicative polynomial-sine games with 50 function nodes, 100 players, and an Erdős-Renyi probability threshold ranging from 0.15 to 0.25. The increase from 5 to 50 function nodes in these random games results in 10 times as many possible edges compared to the games from Sokota et al. Even with this increased complexity and greater variation among neighboring game instances, our VPL method is able to generalize well across the parameter space.

Not surprisingly, for both varied parameters as the amount of data per instance increases, the deviation payoff error for both models decreases. VPL clearly outperforms FPL on both classes of games given the exact same training data, and FPL’s performance doesn’t approach VPL’s until the data-per-instance available to FPL approaches the *total* amount of data used by VPL. These results suggest that a VPL model is more sample efficient than a collection of FPL models, because it can also make use of data from neighboring instances to improve learning in areas of the simplex for which it has fewer samples. In Appendix C, Figure ?? validates that this trend is consistent across instances. Based on these results, we conclude our variable-parameter learning model is the better approach to learning the deviation payoff function for a parameterized family of games.

5.2 Model Refinement Validation

Sokota et al. [14] demonstrated the importance of targeted resampling in the neighborhood of candidate equilibria as a supplement to deviation payoff learning. Model refinement for parameterized game families also involves refinement over the space of the variable parameter, which is a non-trivial extension given that regret varies as a function of v , so an approximate equilibrium in one

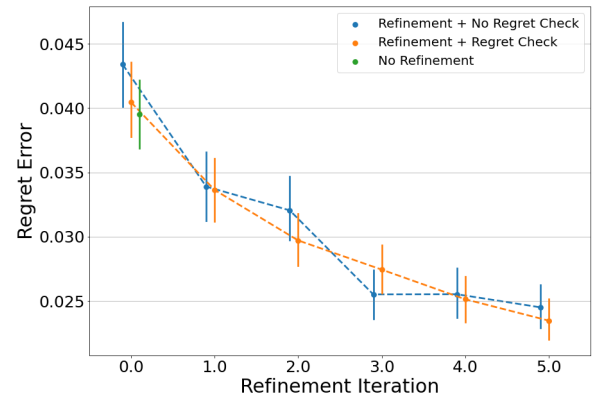


Figure 3: Partitioning the training data across several iterations of targeted model refinement reduces approximate equilibrium regret error compared to no model refinement with the same total amount of training data.

game instance may not be an approximate equilibrium in others. We consider two variations of refinement: one where the mixtures from the Nash-finding algorithm are filtered by a maximum regret threshold (as estimated by the model) prior to neighborhood sampling and another where all resulting mixtures are included. The number of neighborhood samples for each candidate equilibrium is adjusted accordingly (higher for filtered mixtures) to ensure that both refined models train on approximately the same amount of data as the model with no refinement.

We examine whether and how model refinement can improve the accuracy of equilibrium computations across the parameter space by measuring the absolute error of regret estimates for our reported equilibrium candidates. We compare the two refined VPL model variants to a model without refinement which trains on the same total amount of data, but gathered all at once rather than in refinement stages. For this experiment, we use replicator dynamics [15] as the Nash-finding algorithm. We run a large number of replicator dynamics updates in parallel from different initial mixtures for a fixed number of iterations.

For this experiment, the three models—no refinement, refinement with intermediate regret check, refinement with no regret check—are evaluated on the same 100 additive sine games with 5 strategies and 50 to 100 players. For each refinement iteration we run replicator dynamics on 100 random mixtures per instance (5100 mixtures total). All three models are trained on approximately 55,000 payoff samples each. The model without refinement is given all training data at once; the two refined models are given 44,000 initial payoff samples. One refined model uses a maximum regret of 0.1 to filter approximate equilibria, with 200 neighborhood samples per equilibrium. The other refined model does not apply a maximum regret filter, and generates 100 neighborhood samples per distinct mixture resulting from replicator dynamics. The regret error is measured as the mean absolute error of regret, normalized to the payoff scale, for all candidate equilibria across all instances for each iteration. As in previous experiments, we compute 95% confidence intervals around the average regret error.

Figure 3 shows that each iteration of targeted model refinement reduces approximate equilibrium regret error, and both refinement approaches significantly outperform the model with the same total training data but no refinement. In Appendix C, Figure ?? validates that this trend is consistent across instances. With no clear difference between the two refinement variations, we cannot draw any conclusions about the effectiveness of the intermediate regret threshold. We believe an analyst with more specific domain knowledge may be able to make an appropriate choice about whether to include a maximum regret filter for refinement.

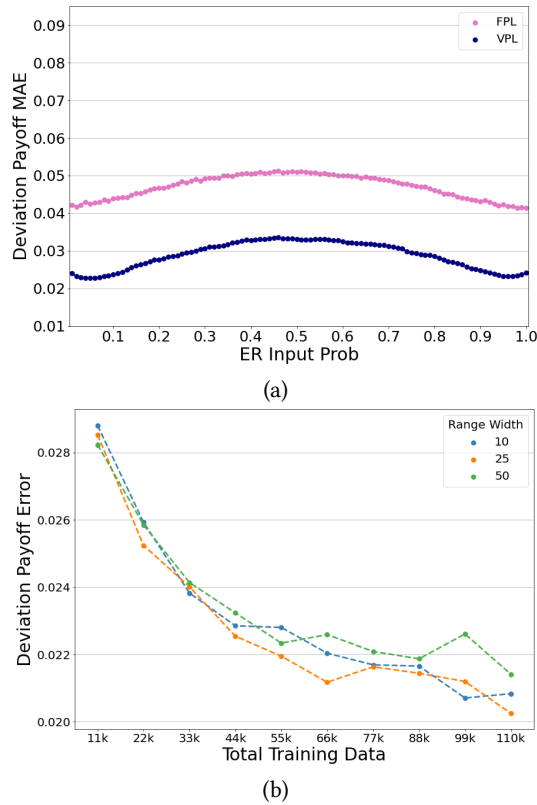


Figure 4: (a) VPL still outperforms FPL on random games with 100 function nodes and with Erdős-Rényi threshold varied from 0 to 1. (b) For the three models with different player-count range widths, the deviation payoff errors are roughly the same.

5.3 VPL Scalability Validation

5.3.1 Variable Erdős-Rényi Threshold. We assess VPL model scalability on random additive polynomial-sine game families with an Erdős-Rényi threshold ranging from 0.01 to 1. These random games have 5 strategies, 100 players, and 100 function nodes — the most complex random games out of all our experiments. Figure 4a compares a single VPL model against 100 FPL models (one per game instance) when trained on 100k total training examples. Not only does the VPL model still outperform the collection of FPL models,

but the overall VPL deviation payoff error across the entire range is comparable to the best overall VPL deviation payoff error in Figure 2b where the games had half as many function nodes!

5.3.2 Variable Number of Players. To further evaluate the scalability of our model for discrete parameters we investigate whether increasing the width of the parameter range affects the deviation payoff error, given the same amount of total training data. In this experiment we evaluate VPL models with player-count range widths equal to 10, 25 and 50 on the same set of 100 random 5-strategy, 5-function additive polynomial-sine games for different amounts of total training data. To account for varied deviation payoff function complexities associated with different magnitudes of player counts, we evaluate each model’s performance by averaging deviation payoff error on game instances with 50, 60, . . . , 100 players. This means that we evaluate one VPL model with range width equal to 50 ($50 \leq p \leq 100$ players), two VPL models with range width equal to 25 ($50 \leq p \leq 75$ and $75 \leq p \leq 100$ players), and three VPL models with range width equal to 10 ($50 \leq p \leq 60$, $70 \leq p \leq 80$, and $90 \leq p \leq 100$ players); each VPL model receives the same amount of training data. We compute the average normalized deviation payoff MAE across all strategies for 495 mixtures per instance, per game. Once again the magnitude of the confidence intervals is too small to see in the plot.

Figure 4b shows that increasing the total amount of training data improves performance regardless of parameter range size. Further, our results show no notable difference in performance between the models with differing player-count range widths, given the same amount of training data. This suggests that for wider parameter ranges, a single VPL model is still the better approach to learning the deviation payoff function for a parameterized family of games. Additionally, VPL is appealing for wider parameter ranges because it is far less tedious to train and refine a single neural network compared to 51 different FPL neural networks, and far more informative than just two snapshots at, say, 50 and 100 players.

6 PARAMETERIZED GAME FAMILY ANALYSIS

A limitation of work that studies a small number of separate game instances is that the analysis may not adequately depict trends for the entire parameter space. As a result, the relationship between the environment parameter and incentives in the game often remains understudied. Our single learned model allows for more tractable and complete analysis of the parameter space, and also enables several new types of analysis. We now present two possible avenues to help characterize the relationship between environment parameters and incentives. We emphasize that these are just a glimpse into the many new varieties of analysis that variable-parameter models will enable.

6.1 Robust Equilibria

Because a Nash equilibrium is only defined for a game instance with fixed environment parameters, in a parameterized game family $\gamma(V)$ any function of equilibrium, such as regret, price of anarchy or social welfare, is also parameterized by V . In some cases, an analyst with uncertainty about a parameter’s true value might want to identify robust equilibria. We can compute any statistic of the given function of equilibrium, and use the resulting value as a measure

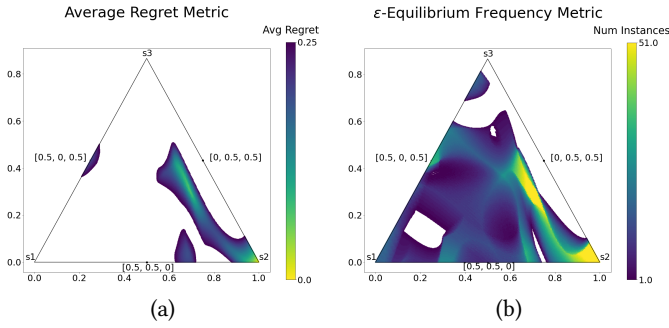


Figure 5: Comparison of two robustness metrics on a randomly generated 3-strategy game: (a) average regret metric, (b) ϵ -equilibrium frequency metric.

of the robustness of the equilibrium. For example, an analyst might evaluate average or max regret across all game instances for a given profile, and label the profile as a robust equilibrium if the regret statistic is below some threshold.

6.1.1 Example. Figure 5 compares robust equilibria found by two robustness metrics in a randomly generated 3-strategy multiplicative polynomial-sine game with 50 to 100 players using calculated regrets. Each point in the simplex corresponds to a symmetric mixed strategy. In Figure 5a, the color shows the mean regret of the corresponding profile across all game instances. Note that the white points correspond to profiles in which mean regret was greater than ϵ . In Figure 5b, the color shows how many times each mixture was an approximate equilibrium (for a fixed ϵ). In this plot, the white points correspond to profiles that were never approximate equilibria. In both plots, the brighter points correspond to profiles that are considered more robust. The similarities between these plots are typical for variable-parameter games we explored—in our experiments we have found that various robustness metrics tend to identify broadly similar sets of profiles as robust equilibria.

6.2 Parameter Sensitivity Analysis

For a given parameterized game family $\gamma(V)$, one important question an analyst might want to investigate is: As the parameter V is varied, how do the incentives in the game family change? One possible avenue to answer this question is to look at which strategies appear in the support of candidate equilibria across game instances. Strategies which appear in the support for many approximate equilibria across game instances may be less sensitive to changes in V and vice versa. Clustering instances which have (or lack) similar equilibria and then reasoning why the instances are clustered together based on domain-specific knowledge may provide additional insight. Finally, identifying equilibrium basins of attraction under a particular search algorithm may help describe how approximate Nash equilibria “move” in the simplex as the environment parameter is varied.

6.2.1 Example. Figure 6 shows how one ϵ -equilibrium varies across the player parameter space in a random multiplicative polynomial-sine game family with 5 strategies and 50 to 100 players. For each game instance, we ran replicator dynamics with a uniform initial

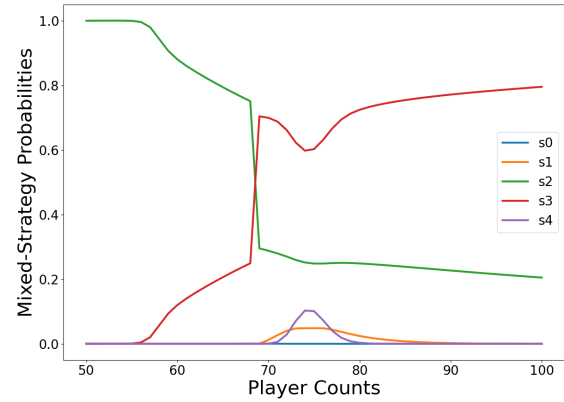


Figure 6: Approximate equilibrium probabilities for each strategy plotted as a function of the number of players for a random multiplicative polynomial-sine game family. This game family demonstrates the importance of parameter sensitivity analysis across the full parameter space to better understand the incentives in the game.

mixture and plotted the equilibrium probabilities for each strategy. Note that for the entire game family this ϵ -equilibrium has an average regret of 0.03 and a maximum regret of 0.15. If an analyst selects “representative” player count values 50, 75, and 100 and then constructs and analyzes separate game models for each game instance, they would have an incomplete view of how this ϵ -equilibrium varies across the parameter space, particularly between game instances with 50 and 75 players. Thus this game family demonstrates the importance of analyzing the entire parameter space to better understand an environment’s equilibria.

7 CONCLUSION

Our new approach to learning deviation payoff functions that generalize over parameterized game families demonstrates clear advantages over existing techniques for simulation-based games. Our experiments show that in the common case where analysts want to explore multiple settings for certain environment parameters, learning a single variable-parameter model can produce more accurate analysis from a smaller data set than previous state-of-the-art methods that rely on analyzing each game instance independently. This variable-parameter learning technique allows for more complete analysis of both continuous and discrete parameter spaces, and enables new types of robustness and sensitivity analysis that previously would have been intractable in simulation-based games.

ACKNOWLEDGMENTS

This work was supported in part by funding from the US Army Research Office (MURI grant W911NF-18-1-0208).

REFERENCES

- [1] Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. 2020. Improved Algorithms for Learning Equilibria in Simulation-Based Games. In *AAMAS*. 79–87.
- [2] Pranav Dandekar, Ashish Goel, Michael P. Wellman, and Bryce Wiedenbeck. 2015. Strategic Formation of Credit Networks. *ACM Transactions on Internet*

- Technology* 15, 1 (2015), 3:1–3:41.
- [3] Quang Duong, Yevgeniy Vorobeychik, Satinder Singh, and Michael P. Wellman. 2009. Learning Graphical Game Models. In *IJCAI*.
- [4] Shaheen Fatima. 2006. Sequential versus Simultaneous Auctions: A Case Study. In *EC*. 82–91.
- [5] Sevan G. Ficici, David C. Parkes, and Avi Pfeffer. 2008. Learning and Solving Many-Player Games through a Cluster-Based Representation. In *UAI*. 187–195.
- [6] Nobuyuki Hanaki and Juliette Rouchier. 2013. If You Are So Rich, Why Aren't You Smart?. In *Winter Simulation Conference (WSC)*. 1731–1741.
- [7] Steven Jecmen, Arunesh Sinha, Zun Li, and Long Tran-Thanh. 2020. Bounding Regret in Empirical Games. In *AAAI*, Vol. 34. 4280–4287.
- [8] Patrick R Jordan, L Julian Schwartzman, and Michael P Wellman. 2010. Strategy Exploration in Empirical Games. In *AAMAS*. 1131–1138.
- [9] Patrick R. Jordan, Michael P. Wellman, and Guha Balakrishnan. 2010. Strategy and Mechanism Lessons from the First Ad Auctions Trading Agent Competition. In *EC*. 287–296.
- [10] Zun Li and Michael P. Wellman. 2020. Structure Learning for Approximate Solution of Many-Player Games. In *AAAI*. 2119–2127.
- [11] Chun Kai Ling, Fei Fang, and J. Zico Kolter. 2018. What Game are We Playing? End-to-end Learning in Normal and Extensive Form Games. In *IJCAI*. 396–402.
- [12] Brandon A. Mayer, Eric Sodomka, Amy Greenwald, and Michael P. Wellman. 2013. Accounting for Price Dependencies in Simultaneous Sealed-Bid Auctions. In *EC*. 679–696.
- [13] Katherine Mayo and Michael P. Wellman. 2021. A Strategic Analysis of Portfolio Compression. In *ICAIF*. Article 20, 8 pages.
- [14] Sam Sokota, Caleb Ho, and Bryce Wiedenbeck. 2019. Learning Deviation Payoffs in Simulation-Based Games. In *AAAI*, Vol. 33. 2173–2180.
- [15] Peter D. Taylor and Leo B. Jonker. 1978. Evolutionary Stable Strategies and Game Dynamics. *Mathematical Biosciences* 40, 1-2, 145–156.
- [16] David R.M. Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey Rosenschein. 2013. Empirical Analysis of Plurality Election Equilibria. In *AAMAS*. 391–398.
- [17] Bruno Tuffin and Patrick Maillé. 2006. How Many Parallel TCP Sessions to Open: A Pricing Perspective. In *ICQT Workshop: Performability Has its Price*. Springer, 2–12.
- [18] Yevgeniy Vorobeychik, Michael P. Wellman, and Satinder Singh. 2007. Learning Payoff Functions in Infinite Games. *Machine Learning* 67 (05 2007), 145–168.
- [19] Elaine Wah, Dylan Hurd, and Michael P. Wellman. 2016. Strategic Market Choice: Frequent Call Markets vs. Continuous Double Auctions for Fast and Slow Traders. *EAI Endorsed Transactions on Serious Games* 3, 10 (2016).
- [20] Elaine Wah and Michael P. Wellman. 2016. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance* 5, 3-4, 69–93.
- [21] Michael P. Wellman. 2006. Methods for empirical game-theoretic analysis (extended abstract). In *AAAI*. 1152–1155.
- [22] Michael P. Wellman. 2020. Economic Reasoning from Simulation-Based Game Models. *Economia. History, Methodology, Philosophy* 2, 10 (2020), 257–278.
- [23] Michael P. Wellman, Tae Hyung Kim, and Quang Duong. 2013. Analyzing Incentives for Protocol Compliance in Complex Domains: A Case Study of Introduction-Based Routing. In *Workshop on the Economics of Information Security*.
- [24] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, Shih-Fen Cheng, and Rahul Suri. 2005. Approximate Strategic Reasoning through Hierarchical Reduction of Large Symmetric Games. In *AAAI*. 502–508.
- [25] Bryce Wiedenbeck, Ben-Alexander Cassell, and Michael P. Wellman. 2014. Bootstrap Statistics for Empirical Games. In *AAMAS*. 597–604.
- [26] Bryce Wiedenbeck and Michael P. Wellman. 2012. Scaling Simulation-Based Game Analysis through Deviation-Preserving Reduction. In *AAMAS*. 931–938.
- [27] Bryce Wiedenbeck, Fengjun Yang, and Michael P. Wellman. 2018. A Regression Approach for Modeling Games with Many Symmetric Players. In *AAAI*. 1266–1273.
- [28] Brian Zhang and Tuomas Sandholm. 2021. Finding and Certifying (Near-)Optimal Strategies in Black-Box Extensive-Form Games. In *AAAI*.