Learning Bayesian Game Families, with Application to Mechanism Design

Extended Abstract

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

Michael P. Wellman
University of Michigan
Ann Arbor, MI, USA
wellman@umich.edu

ABSTRACT

We demonstrate the advantages of learning an interim model for Bayesian game families through an in-depth study of empirical mechanism design for a dynamic sponsored search auction scenario. A full version of this paper, with additional background, methods, and results, is available at: https://arxiv.org/pdf/2502.14078.

KEYWORDS

Game-Model Learning; Empirical Mechanism Design

ACM Reference Format:
Madelyn Gatchel and Michael P. Wellman. 2025. Learning Bayesian Game Families, with Application to Mechanism Design: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Many real-world strategic interactions can be modeled as ***Bayesian games***, where payoffs depend on players' actions as well as their private information, or ***types***. Outcomes may also depend on parameters of the environment, such that each parameter setting induces a different Bayesian game. Given limited modeling resources, analysts must decide in advance the range of parameter settings to consider and the granularity of that range. While domain knowledge of the interaction may guide this selection, there is no guarantee that the most salient parameter settings are covered. What the analyst would ideally have is a model of the entire ***Bayesian game family***, from which they could reason about any relevant parameter setting.

A motivating application for Bayesian game families is ***mechanism design***, where a designer sets or influences an environment parameter that affects strategic incentives in the multi-agent interaction. Each value of this parameter results in a different ***game instance***. The mechanism designer's goal is to find the parameter setting that optimizes a relevant objective function, such as social welfare or revenue. In ***empirical mechanism design*** (EMD), game model instances are induced from simulation data. In past EMD studies [1, 3, 6], researchers selected a limited set of mechanism settings, separately modeling and analyzing each game instance.

Gatchel and Wiedenbeck [2] demonstrated that learning a single parameterized payoff model for families of related normal-form games is more data-efficient than training separate models for each game instance. We extend this approach to Bayesian game families, exploiting the type-conditional form of strategies for these games. Specifically, we investigate the learning of ***interim payoff functions***, which explicitly condition on a single player's type. By marginalizing out this type, we obtain the ***ex ante payoff functions***, which are essentially the payoffs learned in a normal-form model. We also explore learning ex ante payoff functions directly, and compare this approach with that of learning interim models.

We validate our method through an EMD case study in the domain of sponsored search, where the publisher sets an auction reserve requirement in order to maximize revenue in equilibrium. Our search auction model is designed to capture the dynamic nature of bidding, where advertisers can revise their bids based on provisional results of earlier bidding rounds. We do so in a two-stage scenario, in which the bidders can *attempt* to modify their bids given the state of bidding after the first round. These attempts succeed probabilistically, thus providing an incentive for the players to submit meaningful first-round bids. The scenario is simple to describe and design heuristic strategies, yet too complex for straightforward analytic solution. We implement an agent-based simulation model of the scenario, and from the simulation-generated data learn Bayesian game family models to support empirical game-theoretic analysis and mechanism design.

2 LEARNING BAYESIAN GAME FAMILIES

Simulator queries are costly, and the results are noisy due to randomness in strategies or in the game environment. We assume a fixed budget of simulator queries for learning and validation, so we must allocate queries across the parameter space, strategy space, and type space. This simulation data is used to train a model representing the deviation payoff function for a symmetric Bayesian game family. We experimentally compare ex ante and interim game-family learning methods. The ex ante method becomes equivalent to the normal-form approach developed by Gatchel and Wiedenbeck [2] once types are abstracted away, and thus serves as the baseline in the Bayesian setting.

We first train a neural network representing the ***ex ante deviation payoff function***, which takes a symmetric mixed strategy and parameter setting as input and outputs, for each strategy, the expected payoff a symmetric player would receive by deviating to that strategy, given that all other players follow the mixed strategy. In Bayesian games, sampling over types represents a distinct source of noise in payoff estimates. By conditioning estimates on the deviator's type, we can leverage type-specific information in each sample. We therefore also learn the ***interim deviation payoff***

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

function, which gives the vector of deviation payoffs conditional on the symmetric deviator's type.

3 EMPIRICAL MECHANISM DESIGN

Deriving Equilibria. We use the learned game-family model to derive approximate Bayes-Nash equilibria (BNE) in game instances corresponding to various settings of the environment parameter. Our method adapts existing techniques for deviation payoffs [2, 5] to the Bayesian case. For a given game instance, we run a Nash-finding algorithm using ex ante deviation payoffs that are either predicted directly by an ex ante model or computed via marginalizing an interim model's predictions. The output mixed-strategy profile is a candidate ϵ-BNE if the predicted regret is at most ε. We validate the candidate with a modest number of additional simulator queries to compute the true regret of the candidate approximate equilibrium. If the true regret is at most ε, the approximate equilibrium is confirmed.

Parameter Optimization. Once trained, the game-family model can evaluate any game instance within the trained range (and plausibly beyond), supporting a more granular parameter search than previous EMD approaches. When used in an optimization algorithm, it eliminates the need to train separate models at each iteration, reducing the algorithm's dependence on the sampling budget.

Iterative EMD with Piecewise Strategies. Analysis of Bayesian games typically focuses on identifying ex ante equilibria. In a BNE no player can gain by deviating to any other strategy in the strategy set, *in expectation over player types.* If the strategy set includes all mappings from type to action, then in an ex ante equilibrium the player would also not wish to deviate conditional on its own type; that is, the ex ante BNE would also be an interim equilibrium. Given a restricted set strategy set, however—the norm for empirical game models—a player can often benefit by deviating to an alternative strategy in the restricted strategy set once its own type is revealed. We consider a particular form of higher-order strategy that exploits such opportunities by selecting a base-level ***atomic strategy*** from the restricted strategy set conditional on revealed type. Given a set of contiguous type intervals which collectively partition the type space, a ***piecewise-conditional strategy*** maps each type interval to an atomic strategy from the base set. We can further integrate these strategies into an expanded model trained without requiring any additional simulation samples. These operations enable an iterative procedure that expands the game family model from an initial set of atomic strategies through a double oracle [4] approach: repeated generation of new (piecewise) strategies that best-respond to an equilibrium of the previous configuration. By choosing an equilibrium from the game instance that optimizes the design parameter, this becomes an iterative method for EMD.

4 OVERVIEW OF RESULTS

Fig. 1 shows that both ex ante and interim models achieve low payoff error across the trained parameter range. The interim model is able to maintain low error well beyond the trained range. Moreover, we find that equilibria approximated using the interim model consistently have equal or lower regret absolute error compared to ex ante, both within and beyond the trained parameter range (Fig 2). These improvements in extrapolation and equilibrium identification provide compelling evidence that exploiting type structure by learning an interim model is advantageous for Bayesian games.

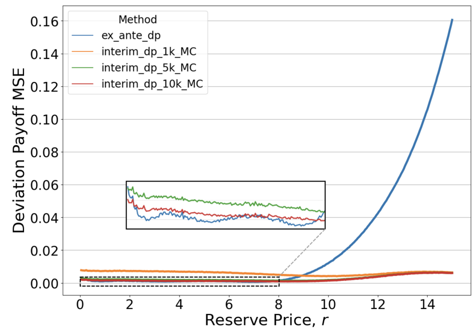

Figure 1: With enough marginalization samples, interim model accuracy matches ex ante on the trained range, $(0, 8]$. Interim, but not ex ante, models extrapolate well.

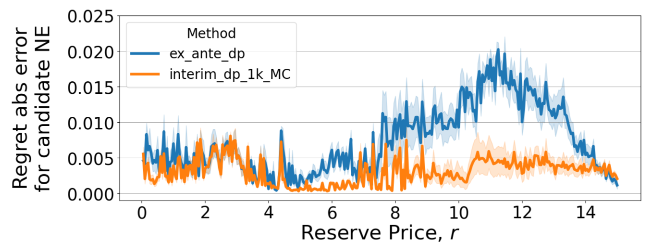

Figure 2: Candidate equilibria from the interim approach have equal or lower regret error compared to the ex ante approach.

In our application to mechanism design, we demonstrate that the learned models support effective EMD procedures. Analysis of a fine-grained grid over the game family reveals the benefit of learning multiple models from separate datasets, and further demonstrates the relative robustness of interim over ex ante. Local search methods reliably produce approximately optimal reserve settings, requiring only a modest number of restarts.

A final feature of the interim model is that it enables us to generate new strategies that outperform those in the original set. In applying this method to the dynamic search auction, we demonstrate that even a couple of iterations of piecewise strategy generation refines the model to produce decisions that improve revenue.

5 CONCLUSION

Our investigation produces new insights about alternative model forms for Bayesian game families, with compelling experimental evidence in favor of learning interim payoff functions. The interim model enables generation and integration of new strategies, based on optimal piecewise constructions. Overall, the methods developed here support an automated approach to empirical mechanism design, given an agent-based simulator and a seed set of strategies.

REFERENCES

[1] Erik Brinkman and Michael P. Wellman. 2017. Empirical Mechanism Design for Optimizing Clearing Interval in Frequent Call Markets. In *18th ACM Conference on Economics and Computation*. 205–221.

[2] Madelyn Gatchel and Bryce Wiedenbeck. 2023. Learning Parameterized Families of Games. In *22nd International Conference on Autonomous Agents and Multiagent Systems*. 1044–1052.

[3] Patrick R. Jordan, Michael P. Wellman, and Guha Balakrishnan. 2010. Strategy and Mechanism Lessons from the First Ad Auctions Trading Agent Competition. In *11th ACM Conference on Electronic Commerce*. 287–296.

[4] H. Brendan McMahan, Geoffrey J. Gordon, and Avrim Blum. 2003. Planning in the presence of cost functions controlled by an adversary. In *20th International Conference on Machine Learning* (Washington, DC). 536–543.

[5] Samuel Sokota, Caleb Ho, and Bryce Wiedenbeck. 2019. Learning Deviation Payoffs in Simulation-Based Games. In *33rd AAAI Conference on Artificial Intelligence*. 2173–2180.

[6] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P. Wellman. 2006. Empirical Mechanism Design: Methods, with Application to a Supply-Chain Scenario. In *7th ACM Conference on Electronic Commerce*. 306–315.