

# Conservative Equilibrium Discovery in Offline Game-Theoretic Multiagent Reinforcement Learning

AAAI Track

Austin A. Nguyen  
University of Michigan  
Ann Arbor, United States  
ngaustin@umich.edu

Michael P. Wellman  
University of Michigan  
Ann Arbor, United States  
wellman@umich.edu

## ABSTRACT

Offline learning of strategies takes data efficiency to its extreme by restricting algorithms to a fixed dataset of state-action trajectories. We consider the problem in a mixed-motive multiagent setting, where the goal is to solve a game under the offline learning constraint. We first frame this problem in terms of selecting among candidate equilibria. Since datasets may inform only a small fraction of game dynamics, it is generally infeasible in offline game-solving to even verify a proposed solution is a true equilibrium. Therefore, we consider the relative probability of low regret (i.e., closeness to equilibrium) across candidates based on the information available. Specifically, we extend Policy Space Response Oracles (PSRO), an online game-solving approach, by quantifying game dynamics uncertainty and modifying the RL objective to skew towards solutions more likely to have low regret in the true game. We further propose a novel meta-strategy solver, tailored for the offline setting, to guide strategy exploration in PSRO. Our incorporation of Conservatism principles from **Offline** reinforcement learning approaches for strategy Exploration gives our approach its name: COffeE-PSRO. Experiments demonstrate COffeE-PSRO’s ability to extract lower-regret solutions than state-of-the-art offline approaches and reveal relationships between algorithmic components, empirical game fidelity, and overall performance.

## KEYWORDS

Empirical game-theoretic analysis; game-solving; multiagent reinforcement learning; offline reinforcement learning

### ACM Reference Format:

Austin A. Nguyen and Michael P. Wellman. 2026. Conservative Equilibrium Discovery in Offline Game-Theoretic Multiagent Reinforcement Learning: AAAI Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 8 pages. <https://doi.org/10.65109/XEYY6214>

## 1 INTRODUCTION

Game-theoretic analysis is a powerful tool when evaluating multiagent systems because it focuses attention on configurations where agents behave rationally relative to each other. In complex games with large strategy spaces, identifying such equilibrium profiles through **multiagent reinforcement learning** (MARL) requires

substantial computation and *data*. Gathering ample data for game-theoretic analysis can thus be difficult and costly in real-world applications. A capability to perform such analysis *offline* from pre-generated, fixed-size datasets addresses the extreme case of limits on real-time data collection.

Strategy generation using offline reinforcement learning is well-studied in single-agent and cooperative multiagent domains. Many offline learning algorithms adopt the *conservatism principle*: strategies should stay in dataset-informed regions and assume unknown regions yield low returns. This can be interpreted loosely as *policy selection*: among policies that seem approximately optimal, favor ones for which the dataset provides strong evidence and thus better relative confidence of high actual (true) return. Analogously, we frame offline game-solving as *equilibrium selection*, where we choose among profiles apparently in approximate equilibrium, aiming to minimize actual (true-game) regret.

Our contribution is in the application and extension of conservatism to offline *game-solving*. This extension poses a special challenge: consideration of a candidate strategy profile’s potential strategic deviations. Our approach achieves this by combining offline RL algorithmic concepts and Policy Space Response Oracles (PSRO) [11], an online game-solving framework that iteratively extends a strategy set using deep RL best-responses to mixed strategies in an empirically estimated game model. Given a dataset of trajectories (sequences of true-game multiagent interactions), we first train an ensemble dynamics model to replace an online simulator and use prediction differences within the ensemble to quantify uncertainty under the dataset. We apply conservatism by introducing a response objective that balances reward, uncertainty given the current response target, and uncertainty of potential strategic deviations. Further, we propose a novel meta-strategy solver that minimizes a heuristic-based, pessimistic estimate of regret. This extension of Conservatism from **Offline** RL approaches to aid strategy Exploration in PSRO gives our approach its name: COffeE-PSRO. In experiments on a sequential bargaining game, we find that COffeE-PSRO tends to output lower-regret solutions than tested baselines. Our results highlight the role of conservatism in improving model fidelity, robustness to dataset changes, and solution regret.

## 2 RELATED WORK

Conservatism is a common theme in offline RL approaches [12]. Model-free methods have incorporated conservatism through policy regularization [8, 30], value penalties [9], and implicit behavioral constraints [17, 19], while model-based methods have used uncertainty quantification and reward penalties [7, 32]. Conservative



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/XEYY6214>

methods in offline MARL have addressed centralized coordination [25] and nonstationarity in decentralized optimization [5].

Yang and Wang [31] provide an overview of game-theoretic approaches in MARL. The framework of **Policy Space Response Oracles** (PSRO) [2, 11] employs an empirical game modeling approach [29], iteratively extending the strategy set using deep RL to derive policies maximizing return against an other-player **response target** defined by a **meta-strategy solver** (MSS). Substantial work has investigated strategy exploration [6] in PSRO, in particular through MSS selection of response targets [1, 15, 16, 22, 27] or through definition of response objectives [18, 20, 28].

In offline game-theoretic MARL, theoretical work has investigated the necessary conditions for equilibrium extraction [4, 34], defining **unilateral concentration** as the property where a dataset covers an equilibrium and all its unilateral strategic deviations. Other approaches focus on offline evaluation of given strategies to determine which are most likely equilibria [33, 34]. Shao et al. [21] employ conservative critics and low-rank matrix completions to select equilibria given a fixed dataset of payoff data for select policy profiles. Most relevant here is work on strategy generation for offline equilibrium extraction [14], which uses a trained dynamics model in place of a simulator and mixes the final solution with a behavior cloning policy. Chen et al. [3] adapts fictitious self-play using single-agent offline RL and importance sampling.

### 3 NOTATION

We model a multiagent environment as a **stochastic game** defined by  $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, Z, O, n, \gamma \rangle$ . The environment follows trajectories of states  $s \in \mathcal{S}$ , influenced by actions  $a \in \mathcal{A}$  taken alternately by each of  $n$  agents. Each action prompts a transition according to the state transition function  $\mathcal{T}(s' | s, a) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ . Agent rewards are given by  $\mathcal{R}(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$ , and  $\gamma \in [0, 1)$  is a discount factor. The environment is partially observable, meaning agents draw observations  $z \in Z$  from an observation function  $O(s) : \mathcal{S} \rightarrow Z$ . We denote an action-observation history by  $h_t$ , with  $\mathcal{H}_t \equiv (Z \times \mathcal{A})^t$  the space of conceivable  $t$ -length histories. A **policy** or **pure strategy**  $\pi_i$  for player  $i$  is defined as a mapping from  $h_t$  to a probability distribution over legal actions  $\pi_i : \mathcal{H} \rightarrow \Delta(\mathcal{A})$ . We omit subscript  $i$  to refer to a **joint policy** (or **strategy profile**)  $\pi = (\pi_1, \dots, \pi_n)$ . A **trajectory**  $\tau$  describes a rollout of game play and is represented by a sequence of tuples  $\tau = ((s_t, o_{i,t}, a_{i,t}, r_{i,t}, i))_{t=1}^{|\tau|}$ , where  $t$  indexes the timestep. A dataset  $\mathcal{D}$  is a collection of trajectories  $\{\tau^j\}_{j=1}^{|\mathcal{D}|}$ , indexed here by  $j$ . We may subscript a trajectory  $\tau_\pi^j$  to indicate the joint policy  $\pi$  generating the rollout. The joint probability of being in state  $s$  and executing action  $a$  under joint policy  $\pi$  is denoted by  $\mu^\pi(s, a)$ .

Multiagent interactions may also be expressed in **normal form**. A normal form **symmetric game**  $\Gamma$  is represented by a tuple  $(\Pi, U, n)$ , where  $n$  is the number of players,  $\Pi = \Pi_1 \times \dots \times \Pi_n$  is a set of joint strategies ( $\pi \in \Pi$ ), and **payoff function**  $U : \Pi \rightarrow \mathbb{R}^n$  maps strategy profiles  $\pi$  to vectors of payoffs (i.e., returns). A probability distribution  $\sigma_i \in \Delta(\Pi_i)$  over pure strategies is called a **mixed strategy**. We also use  $\sigma_{i,s}$  to indicate player  $i$ 's strategy applied at step  $s$ . A subscript  $-i$  indicates all players except  $i$ . We write player  $i$ 's payoff for playing  $\pi_i$  when the others play  $\pi_{-i}$  as  $u_i(\pi_i, \pi_{-i})$ , where  $u$  is a length- $n$  utility vector. For mixed strategies, the payoffs are

given by expectation:

$$\begin{aligned} u_i(\pi_i, \sigma_{-i}) &= \mathbb{E}_{\pi_{-i} \sim \sigma_{-i}} u_i(\pi_i, \pi_{-i}), \\ u_i(\sigma) &= \mathbb{E}_{\pi_i \sim \sigma_i} u_i(\pi_i, \sigma_{-i}). \end{aligned}$$

Player  $i$ 's **best response** to all others playing  $\sigma$  is given by  $BR_i(\sigma) = \arg \max_{\sigma_i^* \in \Delta(\Pi_i)} u_i(\sigma_i^*, \sigma_{-i})$ . Mixed strategy  $\sigma^*$  is a **Nash equilibrium** (NE) iff  $\forall i. \sigma_i^* \in BR_i(\sigma^*)$ . The potential gain to deviating from  $\sigma$  is called **regret**:  $Regret_i(\sigma) = u_i(\pi_i^*, \sigma_{-i}) - u_i(\sigma)$ , where  $\pi_i^* \in BR_i(\sigma)$ . By definition, if  $\sigma^*$  is a NE, its regret is zero. An **empirical game model**  $\hat{\Gamma} = (\hat{\Pi}, \hat{U}, n)$  is a normal-form game over a finite joint-strategy subspace  $\hat{\Pi} = \hat{\Pi}_1 \times \dots \times \hat{\Pi}_n$ ,  $\hat{\Pi}_i \subset \Pi_i$ , with payoff function  $\hat{U}$  estimated through simulation.

## 4 CONSERVATISM IN OFFLINE GAME-SOLVING

Empirical game updates and best-response training in PSRO require explicit access to a game simulator. The offline setting, in contrast, works with fixed datasets of unspecified quality. Assessing candidate solutions requires reliable information on the included strategies *and* relevant deviations, which is contingent on the dataset's behavior coverage, or lack thereof. Since we cannot draw assumptions on the extent of this coverage, offline game-solving must account for uncertainty in regret assessments. A joint strategy  $\pi$  is considered **covered** if the dataset contains all state-action pairs visitable under  $\pi$ :  $\forall s, a. \mu^\pi(s, a) > 0 \implies (s, a) \in \mathcal{D}$ . In large games, full coverage of any policy  $\pi$  and its unilateral deviations is infeasible. Therefore, we consider the *extent* to which a policy  $\pi$  is covered. Our approach frames offline game-solving as an equilibrium selection problem: among solutions that seem to have low regret offline, we prefer those more certain to have low regret in the true game. By skewing offline strategy exploration toward high-confidence policy spaces, we seek to produce lower-regret solutions than alternative exploration methods with some consistency.

### 4.1 Training the Ensemble Dynamics Model

Our model-based approach adapts analogous single-agent offline RL approaches. We train an ensemble of  $K$  multi-layer perceptrons as dynamics models, each estimating the transition function  $\mathcal{T}$ , reward function  $\mathcal{R}$ , observation mapping  $O$ , and state-dependent action space  $\mathcal{A}$ , from which dataset  $\mathcal{D}$  was generated. Henceforth, our references to these MDP components denote predictions of a learned dynamics model. Ensemble predictions are the average prediction across all  $K$  networks:  $\bar{\mathcal{R}}_i(s_t, a_t) = \frac{1}{K} \sum_{j=1}^K \mathcal{R}_i^j(s_t, a_t)$ , where  $\mathcal{R}_i^j(s_t, a_t)$  refers to model  $j$ 's prediction of player  $i$ 's reward given state-action pair  $(s_t, a_t)$ , and so forth for other MDP components. Additional details, including terminal handling, information state generation, and training parameters are provided in Supp. 4.1<sup>1</sup>.

Prediction differences between intra-ensemble models can serve as proxy for uncertainty quantification [7]. We follow Yu et al. [32] and define  $\rho$  based on the range of model predictions:<sup>2</sup>

$$\rho(s_t, a_t, \mathcal{R}) = \max_{j,k} \sum_{i=1}^n \left| \mathcal{R}_i^j(s_t, a_t) - \mathcal{R}_i^k(s_t, a_t) \right|. \quad (1)$$

<sup>1</sup>All appendices may be found in the arXiv version.

<sup>2</sup>The original approach of Yu et al. [32] used transition function differences. We instead employ reward differences, which achieved stronger results in our experiments.

Equation 1 reflects an assumption that models tend to provide different reward predictions on information-deprived transitions, and similar predictions for sufficiently informed transitions (whether informed explicitly or implicitly by dataset  $\mathcal{D}$ ). The maximum prediction difference is a conservative indicator [7, 32].

## 4.2 A Conservative Response Objective

At each iteration of online PSRO, strategies are trained by optimizing a single agent’s expected return against a fixed, stationary response target,  $\sigma_{-i}$ .

$$\pi_i = \arg \max_{\pi_\theta} J_i(\theta, \sigma_{-i}),$$

$$J_i(\theta, \sigma_{-i}) = \mathbb{E}_{a_t \sim (\pi_\theta, \sigma_{-i})} \sum_{t=1}^T \bar{\mathcal{R}}_i(s_t, a_t). \quad (2)$$

Under a fixed dataset, we cannot assume availability of covered best-responses to  $\sigma_{-i}$ . COffeE-PSRO modifies PSRO’s response objective by optimizing an augmented reward function,  $J'_i(\theta, \sigma_{-i})$

$$= \mathbb{E}_{a_t \sim (\pi_\theta, \sigma_{-i})} \left[ \sum_{t=1}^T \bar{\mathcal{R}}_i(s_t, a_t) - \lambda \rho(s_t, a_t, \mathcal{R}) \right], \quad (3)$$

where  $\lambda$  is a hyperparameter. Equation 3 incorporates model-based uncertainty given  $\sigma_{-i}$ . Although this objective is sufficient in single-agent RL problems, offline game-solving requires information on strategy deviations to additionally *verify* equilibrium properties [4]. We hypothesize that skewing strategy exploration toward equilibria with low  $\rho$  on *unilateral deviations* will also improve performance. Suppose player  $i$  is creating a strategy exclusively optimizing for unilateral concentration. This entails creating a strategy  $\pi_i$  that provides other players the least uncertainty  $\rho$  on all potential responses to  $\pi_i$ . Fig. 1 demonstrates we can achieve this by sampling a different other-player policy at each training episode and minimizing  $\rho$ . A roughly equivalent formulation<sup>3</sup> is to respond to a uniform-random action policy exclusively optimizing  $\rho$ :

$$J''_i(\theta, \sigma_{-i}) = (1 - \alpha) J'_i(\theta, \sigma_{-i}) - \alpha \mathbb{E}_{a_t \sim (\pi_\theta, \mathcal{U}(\mathcal{A}))} \sum_{t=1}^T \lambda \rho(s_t, a_t, \mathcal{R}), \quad (4)$$

where  $\alpha$  is a hyperparameter denoting the probability of optimizing strategic deviation coverage in each training episode. In practice, we linearly anneal  $\alpha$  to zero over PSRO iterations. In the context of strategy generation, this objective balances model-predicted reward, uncertainty given the response target  $\sigma_{-i}$ , and uncertainty of potential other-player deviations. Under the lens of equilibrium selection, Equation 4 can be viewed as skewing strategy exploration toward higher-certainty policy regions [28].

## 4.3 A Conservative Meta-Strategy Solver

We propose a novel MSS that uses ensemble model predictions to derive a conservative response target. This MSS takes concepts from the **Bellman-consistent equilibrium learning** (BCEL) method of Zhang et al. [33], incorporated in a game-solver based on **replicator**

<sup>3</sup>Equivalence requires that the uniform distribution over strategies be behaviorally equivalent to the uniform-random strategy, which does not generally hold. In  $N$ -player games, only one player can use uniform-random action at any given episode.

**dynamics** (RD) [23]. BCEL employs lower and upper utility bounds for given joint strategies and identifies the solution minimizing worst-case regret. We design an RD update that computes a stable point in heuristically estimated, worst-case regret. We define the utility predicted by model  $j$  for player  $i$  under joint policy  $\pi$  as  $u_i^j(\pi) = \frac{1}{N} \sum_{k=1}^N \sum_{(s_t, a_t) \in \tau_\pi^k} \mathcal{R}_i^j(s_t, a_t)$ , where returns are averaged over  $N$  trajectories  $\tau_\pi^k$ . Lower and upper utility bounds are replaced by the minimum and maximum utility estimates over the ensemble  $\mathcal{R}$ , respectively:  $\bar{u}_i(\pi) = \max_{j \in \{1 \dots K\}} u_i^j(\pi)$ . Pessimistic utilities  $\underline{u}(\pi)$  are minimized over  $K$  models. We modify the RD update by myopically optimizing for the current pessimistic regret estimate at each step:

$$UBDP \doteq \bar{u}_i(\pi_i^k, \sigma_{-i}) - \underline{u}_i(\sigma)$$

$$UBDR \doteq \max_{\pi_i^j \in (\hat{\Pi}_i \setminus \pi_i^k)} \bar{u}_i(\pi_i^j, \sigma_{-i}) - \underline{u}_i(\pi_i^k, \sigma_{-i})$$

$$\Delta \sigma_i^k = \sigma_i^k [UBDP - UBDR]$$

where  $\sigma_i^k$  is the probability player  $i$  plays strategy  $k$  and  $\Delta$  indicates the amount of change applied to a variable. The first term captures an upper-bound deviation payoff (UBDP) for playing  $\pi_i^k$  against  $\sigma_{-i}$ . The second represents an upper-bound deviating regret (UBDR) if the player deviates to  $\pi_i^k$ . Our approach accounts for robust offline evaluation, giving our MSS its name: **Robust Replicator Dynamics** (R2D).

Neither RD nor R2D offer theoretical convergence guarantees. Nevertheless, in practice, RD reliably reaches exact or approximate equilibria in empirical games [29] with demonstrated success within the PSRO framework. Analogously, we demonstrate R2D’s empirical benefits within COffeE-PSRO. A complementary RD variant, Regularized RD (RRD) [27], employs an early-stop condition based on a regret threshold to avoid overfitting to the current empirical game solution. R2D can analogously adopt a stopping threshold measured against worst-case regret, resulting in another MSS called **Robust Regularized Replicator Dynamics** (R3D).

## 4.4 Conservative Offline Exploration in PSRO (COffeE-PSRO)

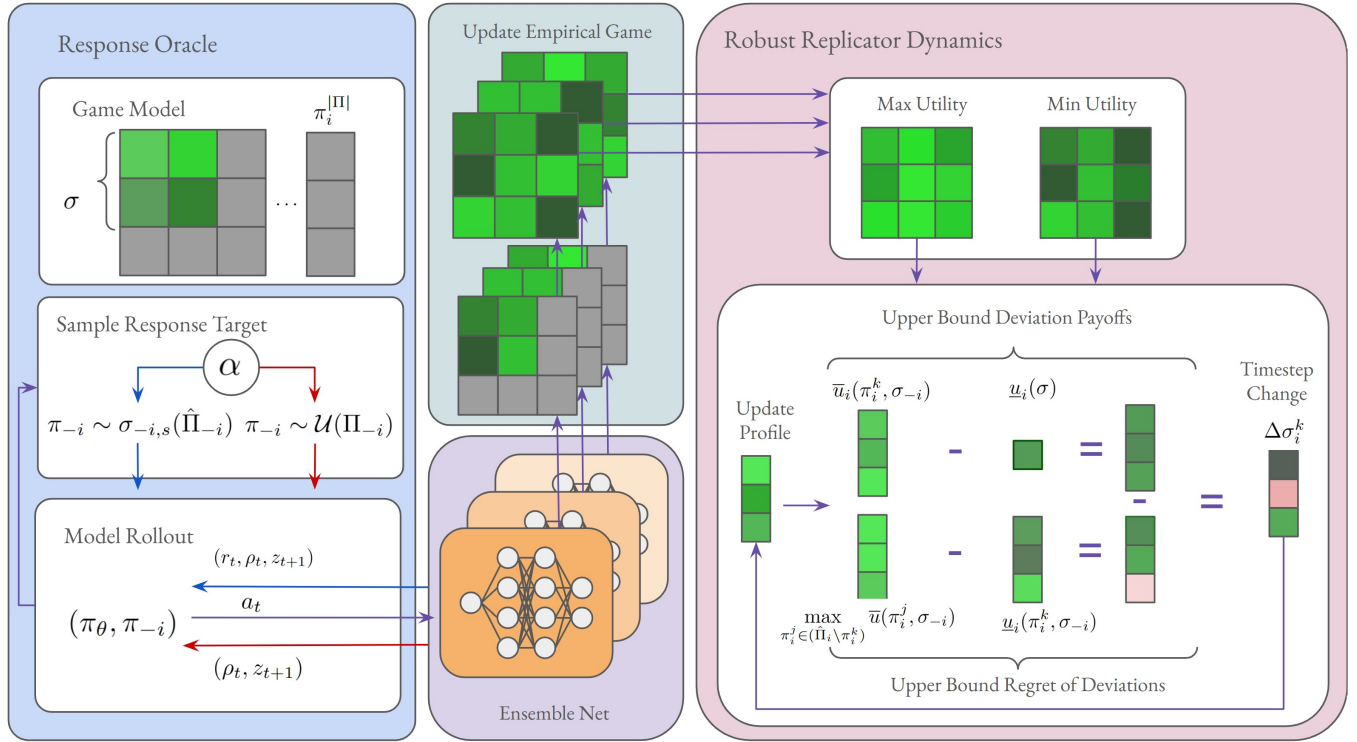
Algorithm 1 summarizes COffeE-PSRO, depicted schematically in Fig. 1. COffeE-PSRO follows the basic structure of PSRO, using a trained dynamics model in place of the online simulator, and using Equation 4 for strategy generation. Any meta-strategy solver  $\mathcal{M}$  can be used with COffeE-PSRO,<sup>4</sup> but our experiments suggest that R2D yields the best results.

# 5 EXPERIMENTS

## 5.1 Bargaining

We test our approach by adapting a 2-player, turn-based, imperfect information, symmetric bargaining game [13]. The game (which we refer to as *Bargaining*) features a pool of items that players must divide between them. Each player has a randomly drawn, private valuation that gives their per-unit utility value for each item. Players alternate making offers of division until one accepts

<sup>4</sup>For conventional MSSs, empirical game payoffs would be estimated using average reward predictions rather than bounds.



**Figure 1: An iteration of COFFe-PSRO.** During strategy generation, a rollout optimizes  $\rho$  against a uniformly sampled policy with probability  $\alpha$ , or the conservative response to the target (Equation 3) otherwise. The two cases are indicated by red and blue arrows. The newly trained policy is added to the strategy set, updating the empirical game using pessimistic and optimistic utility estimates. R2D determines the next iteration’s response target.

**Algorithm 1** Conservative Offline Exploration PSRO (COFFe-PSRO)

**Input:** Dataset  $\mathcal{D}$ , Penalty weight  $\lambda$ , Coverage weight  $\alpha_{init}$ , Steps anneal coverage weight  $S_\alpha$ , meta-strategy solver  $\mathcal{M}$ , PSRO iterations  $S$ , Number of simulations  $N$

**Output:** Player strategy sets  $\hat{\Pi}$ , Player profiles  $\sigma_S$

- 1: Initialize empty empirical game  $\hat{\Gamma} : (\hat{\Pi}, \hat{U}, n)$
- 2: Initialize strategy sets  $\hat{\Pi}_i \leftarrow \{\pi_i^0\}, i \in \{1 \dots n\}$
- 3: Train ensemble dynamics model of size  $K$  on  $\mathcal{D}$
- 4: Update  $\hat{\Gamma}$  using dynamics model with  $N$  simulations
- 5: **for** each PSRO iteration  $s$  from 1 to  $S$  **do**
- 6:   **for** each player  $i$  from 1 to  $n$  **do**
- 7:      $\pi_i \leftarrow \arg \max_{\pi_\theta} J''_i(\theta, \sigma_{-i})$  (4), using  $\lambda$  and  $\alpha$
- 8:     Update strategy sets  $\hat{\Pi}_i \leftarrow \hat{\Pi}_i \cup \{\pi_i\}$
- 9:   **end for**
- 10: Update  $\hat{\Gamma}$  using dynamics model with  $N$  simulations
- 11: Update joint profile  $\sigma_s \leftarrow \mathcal{M}(\hat{\Gamma})$
- 12: Anneal  $\alpha \leftarrow \max(0, \frac{\alpha_{init}(S_\alpha - i)}{S_\alpha})$
- 13: **end for**
- 14: **return**  $\hat{\Pi}, \sigma_S$

or the game reaches a maximum number of iterations. If a player

accepts, players receive utility according to the accepted allocations and respective value functions, discounted by  $\gamma^t$  where  $0 < \gamma < 1$  is a hyperparameter and  $t$  is the number of bargaining rounds so far. We set  $\gamma$  to .99. If no agreement is reached by the turn limit ( $T = 10$ ), both players receive zero utility. Supp. B provides further details.

## 5.2 Implementation and Reproducibility

Our experiments build on an existing codebase, Deepmind’s OpenSpiel [10]. We refer to datasets using tuples  $(N, \pi^\beta)$  where  $N$  is the number of trajectories and  $\pi^\beta$  is the joint behavior policy. Our experiments vary  $N \in \{500, 1000, 2000\}$  and  $\pi^\beta \in \{\pi^u, \sigma^{eq}\}$ , where  $\pi^u$  is a uniform-random action policy, and  $\sigma^{eq}$  is a uniform mixture over 5 equilibrium policies generated by online PSRO (Supp. C). Under these settings, coverage of Bargaining’s full strategy space is highly unlikely, providing benchmarks that measure ability to extract approximate equilibria under limited information. All trials use Double Deep Q-Networks (DDQN) [24] for best-response calculations, yielding deterministic, pure strategies  $\pi_i$ . Supp. A discusses hyperparameters for our best-response and dynamics model.

We compare our approach to two benchmarks, both variants of Offline Equilibrium Finding PSRO [14]. These approaches train a dynamics model, querying it in place of an online simulator (equivalent to COFFe-PSRO with  $\lambda = \alpha = 0$ ), and either use the output solution, or mix that solution with a behavior-cloned policy  $\pi_i^{bc}$  on

**Table 1: Regret of final solutions, averaged over 10 trials with corresponding standard deviation and P-values, each consisting of  $S = 40$  iterations. Bolded algorithmic components indicate our contributions. Datasets contain a varied number of trajectories  $N \in \{500, 1000, 2000\}$  and are generated by two different behavior policies: uniform random  $\pi^U$ , and a mixture over online-generated equilibria  $\sigma^{eq}$ . We distinguish between the MSS  $\mathcal{M}_{expl}$  used during training versus  $\mathcal{M}_{eval}$  during evaluation to isolate effects on strategy exploration and immediate regret [26]. All P-values are calculated with Welch’s t-test to determine whether COffeE-PSRO +R2D yields significant regret differences relative to the corresponding approach. True-game regret evaluations are discussed in Supp. E.**

Approach	(500, $\sigma^{eq}$ )	(500, $\pi^U$ )	(1000, $\sigma^{eq}$ )	(1000, $\pi^U$ )	(2000, $\sigma^{eq}$ )	(2000, $\pi^U$ )
Alg, $\mathcal{M}_{expl}$	$\mathcal{M}_{eval}$ : RD					
<b>COffeE-PSRO, R2D</b>	<b>2.02 ± 0.40</b>   N/A	4.23 ± 0.77   N/A	<b>1.10 ± 0.32</b>   N/A	<b>2.40 ± 0.55</b>   N/A	<b>0.88 ± 0.35</b>   N/A	<b>1.66 ± 0.57</b>   N/A
COffeE-PSRO, RD	2.58 ± 0.40   .01	4.20 ± 0.81   .53	1.47 ± 0.74   .09	2.64 ± 0.55   .17	1.64 ± 0.94   .02	3.15 ± 2.09   .03
OEF, <b>R2D</b>	2.71 ± 0.85   .02	5.01 ± 1.07   .04	1.42 ± 0.52   .05	3.20 ± 0.37   .01	1.03 ± 0.67   .27	2.24 ± 0.84   .01
OEF, RD	2.38 ± 0.55   .06	5.09 ± 1.19   .04	1.19 ± 0.29   .29	3.74 ± 2.10   .04	1.40 ± 0.79   .04	3.28 ± 2.16   .02
OEF-BC, <b>R2D</b>	2.91 ± 0.61   .01	4.00 ± 0.65   .76	2.00 ± 0.22   .01	2.62 ± 0.22   .13	1.85 ± 0.29   .01	2.15 ± 0.23   .01
OEF-BC, RD	2.82 ± 0.38   .01	<b>3.85 ± 0.54</b>   .89	1.92 ± 0.19   .01	2.96 ± 1.07   .08	1.92 ± 0.49   .01	2.83 ± 1.12   .01
Alg, $\mathcal{M}_{expl}$	$\mathcal{M}_{eval}$ : R2D					
<b>COffeE-PSRO, R2D</b>	<b>1.96 ± 0.45</b>   N/A	3.98 ± 0.60   N/A	1.22 ± 0.41   N/A	<b>2.12 ± 0.27</b>   N/A	<b>0.64 ± 0.12</b>   N/A	<b>1.40 ± 0.36</b>   N/A
COffeE-PSRO, RD	2.15 ± 0.43   .17	<b>3.76 ± 0.75</b>   .76	1.34 ± 0.42   .26	2.85 ± 0.91   .02	1.30 ± 0.59   .01	2.42 ± 1.43   .05
OEF, <b>R2D</b>	2.47 ± 0.82   .05	4.90 ± 1.11   .02	1.08 ± 0.37   .78	3.25 ± 1.00   .35	0.73 ± 0.52   .28	1.93 ± 0.49   .01
OEF, RD	2.11 ± 0.39   .22	4.88 ± 1.03   .02	<b>1.01 ± 0.19</b>   .92	3.68 ± 1.63   .01	0.96 ± 0.59   .06	2.79 ± 1.98   .03
OEF-BC, <b>R2D</b>	2.77 ± 0.33   .01	3.97 ± 0.65   .51	1.89 ± 0.16   .01	2.70 ± 0.39   .01	1.68 ± 0.22   .01	2.79 ± 1.98   .03
OEF-BC, RD	2.65 ± 0.20   .01	3.79 ± 0.46   .78	1.98 ± 0.19   .01	2.92 ± 0.75   .01	1.82 ± 0.18   .01	2.65 ± 0.94   .01

$\mathcal{D}$  by weight  $\alpha_{bc}$ . We refer to the former as OEF and the latter as OEF-BC. All approaches use symmetric implementations, where players share a strategy set  $\hat{\Pi}_i$  and profile  $\sigma_i$ ,  $\hat{\Gamma}$  is a symmetric game, and a single policy  $\pi_i$  is generated at each iteration. Analogous to how offline RL uses online simulations for return evaluation, we use online best responses (detailed in Supp. E) to evaluate true-game regret. All trials consist of  $S = 40$  iterations, each iteration corresponding to one strategy generation. We tune COffeE-PSRO and OEF-BC using  $\mathcal{D} = (1000, \pi^U)$ .<sup>5</sup> COffeE-PSRO is tuned across  $(\lambda, \alpha)$  for both  $\mathcal{M} \in \{\text{RD}, \text{R2D}\}$ , running 5 trials for each, and an additional 5 for the initially lowest regret configuration. OEF-BC is tuned exhaustively across  $\alpha_{bc}$ , where all algorithmic variants immediately run 10 trials and then select the lowest regret parameter setting. We fix tuned parameters across datasets. Note that this tuning scheme is more exhaustive than COffeE-PSRO’s. We chose to use R2D as opposed to R3D in order to isolate the effects of optimizing for conservative regret estimates over expected regret.<sup>6</sup> Hyperparameters are detailed in Supp. A.

### 5.3 Regret Analysis

Table 1 compares tuned versions of COffeE-PSRO, OEF, and OEF-BC. We isolate and measure the improvement from an MSS used in strategy exploration  $\mathcal{M}_{expl}$  by using a consistent *evaluation* MSS  $\mathcal{M}_{eval}$  across different approaches [26]. Results suggest that separately applying COffeE-PSRO or R2D yields variable improvement.

<sup>5</sup>OEF has no additional parameters.

<sup>6</sup>This additionally removes the need to tune additional parameters  $\lambda_{rd}$  and  $\lambda_{r3d}$  for fair comparison.

However, the combined use of COffeE-PSRO with R2D tends to outperform tested baselines, suggesting a strong, beneficial interplay between conservatism in strategy generation and MSS.

**Comparing with OEF: Benefits of Conservatism** - Except for  $(1000, \sigma^{eq})$ , COffeE-PSRO + R2D tends to outperform OEF, indicating that intentional conservatism and uncertainty quantification improve performance. Particularly, under a uniform policy  $\pi^U$ ,  $\mathcal{D}$  provides relatively little information on valuable policy spaces, widening the performance margin between COffeE-PSRO + R2D and OEF. Under the same argument of conservatism, OEF-BC also outperforms OEF when using  $\pi^U$  datasets.

**Comparing with OEF-BC: Strategic vs. Uninformed Conservatism** - COffeE-PSRO provides consistently stronger performance than OEF-BC under all datasets except  $(500, \pi^U)$ . These two approaches differ conceptually in how they apply conservatism. Through strategy exploration, COffeE-PSRO operates by prioritizing search in dataset-informed spaces while OEF-BC enforces similarity to dataset behavior. These results support the notion that low-data settings provide little information on how to be conservative in a calculated manner during exploration. Thus, COffeE-PSRO’s strategic conservatism is outperformed by OEF-BC’s simple method of conservatism in low data, weak behavior policy regimes.

**Behavior Cloning in General-Sum Games** - We highlight that OEF-BC is substantially outperformed by all approaches under  $\sigma^{eq}$  datasets, underscoring a core property of general-sum games. Although  $\sigma^{eq}$  provides information on several "strong" strategies, the mixture of equilibria in general-sum games is not necessarily an equilibrium itself, making mixing with a BC-trained policy arbitrarily effective. This emphasizes a vital consideration distinguishing conservative approaches in single and multiagent environments:




Figure 2: Relationships among model fidelity, performance, and conservatism. Each datapoint is a single trial ( $M_{eval} = R2D$ ).

imitating datasets that contain strong behavior does not imply the generation of strong strategies in the latter.

**Immediate Regret Benefits of R2D** - We now discuss  $M_{eval}$ 's effect on immediately evaluated regret independent of strategy exploration. For a given approach and dataset, R2D tends to produce lower regret solutions than RD, suggesting the optimization of conservative regret estimates benefits offline equilibrium extraction. These findings are consistent with prior work [33].

## 5.4 Analyzing Model Fidelity

While online strategy exploration focuses on creating a set of strategies that sufficiently capture game structure, a challenge of offline game-solving is the added uncertainty of utility estimations due to limited access to data. The extent to which our algorithm should trust its offline game model for equilibrium extraction should, intuitively, be a function of its utility approximation accuracy. We refer to a game model's utility estimation accuracy as **model fidelity** and hypothesize that it is a key component for robust offline equilibrium discovery. More precisely, since the optimization of Equation 4 can be interpreted as skewing strategy exploration to create a higher fidelity game model, we investigate the relationship between  $\rho$ , parameter strength, model fidelity, and solution quality.

For each experimental trial depicted in Fig. 2, we use offline-generated strategies  $\hat{\Pi}$  to reconstruct an empirical game  $\hat{\Gamma}$  with true game simulations, calculating the average uncertainty  $\rho$  across all simulations and average utility error across all entries in  $\hat{\Gamma}$ :  $\delta(\hat{\Pi}) = \frac{1}{|\hat{\Pi}|} \sum_{\pi \in \hat{\Pi}} \sum_{i=1}^n |u_i(\pi) - \hat{u}_i(\pi)|$ , where  $u_i(\pi)$  and  $\hat{u}_i(\pi)$  are player  $i$ 's true game and dynamics model estimated utilities, respectively. Figs. 2a and 2b present scatter plots relating  $\rho$  and  $\delta(\hat{\Pi})$ , color-coded proportionally to the magnitude of regret and parameters  $(\lambda, \alpha)$  applied, respectively.

Plots show a clear positive relationship between  $\rho$  and  $\delta(\hat{\Pi})$ , indicating  $\rho$  is an effective proxy for uncertainty quantification. Parameters  $\lambda$  and  $\alpha$  strongly correlate with improved model fidelity but not necessarily regret. Instead, moderate hyperparameter choices, and therefore modest model fidelity, yield the lowest regret, shown in Fig. 2b, suggesting that excessive conservatism prevents the discovery of valuable strategies even if the subgame's utility


estimates are relatively accurate. These experiments demonstrate how effective offline game-solving relies on policy generation balancing conservatism and potential strategic value (model fidelity vs. regret). This tradeoff between model fidelity and regret is reflected in Fig. 2c, where the use of R2D when  $\mathcal{D} = (1000, \pi^U)$  tends to decrease regret across all approaches despite compromising model fidelity.

## 5.5 Ablations

We provide experiments varying  $\lambda$  and  $\alpha$  in Fig. 3, showing the subset of plots that ablate from the strongest parameter settings:  $(\lambda = 4, \alpha = .2)$  and  $(\lambda = 4, \alpha = .3)$  for R2D and RD, respectively. All remaining ablations of COffeE-PSRO are provided in Supp. F. Moderate levels of  $\lambda$  tend to produce the strongest results, both in final regret and convergence speed. However, when using  $\mathcal{M} = R2D$ , a more conservative MSS, smaller values of both  $\lambda$  and  $\alpha$  are favored. This indicates a threshold at which excessive conservatism from all sources becomes detrimental. Generally, higher values of  $\alpha$  correlate to slower initial convergence, indicated by regret spikes in early iterations. These spikes are commonly followed by steep dips, sometimes finding lower regret solutions than configurations with initially stronger performance. This suggests that  $\alpha$ , by skewing strategy exploration toward regions with higher certainty on strategic deviations, benefits lower-regret, higher-certainty equilibrium selection in later iterations.

## 6 LIMITATIONS

An inherent limitation of PSRO, and by extension of COffeE-PSRO, is the computational cost of training a sequence of deep RL strategies in sequence, which can be exorbitantly expensive for complex games. Several approaches have focused on reducing this computational cost [2, 22]. Our experiments find that COffeE-PSRO improves over baselines if the dataset  $\mathcal{D}$  provides enough information to quantify or extrapolate uncertainty in uncovered strategy spaces (strategic conservatism). Otherwise, COffeE-PSRO provides no more benefit than uninformed conservatism. Dataset coverage is especially difficult for multiagent environments, which tend to have larger search spaces than single-agent. Lastly, experiments



**Figure 3: Regret over iterations plots under various parameter settings. Regret entries for  $(\mathcal{M} = \text{RD}, \lambda = 4, \alpha = .3)$  and  $(\mathcal{M} = \text{R2D}, \lambda = 4, \alpha = .2)$  are averaged over 10 trials; all others are over 5.**

$\mathcal{M}$	$\lambda \downarrow$	$\alpha = .1$	$\alpha = .2$	$\alpha = .3$
RD	2	4.15 ± 1.79	3.50 ± 0.78	<b>3.39 ± 0.98</b>
	4	<u>3.37 ± 0.80</u>	<u>3.33 ± 1.18</u>	<b>2.58 ± 0.56</b>
	6	4.91 ± 0.61	<b>4.03 ± 1.35</b>	5.37 ± 2.13
R2D	2	2.89 ± 1.51	<b>2.26 ± 0.47</b>	2.50 ± 0.93
	4	<u>2.49 ± 0.52</u>	<b>2.12 ± 0.27</b>	2.87 ± 0.61
	6	<b>3.43 ± 0.55</b>	3.75 ± 0.81	3.64 ± 1.01

**Figure 4: Final regret under various parameter settings. Regret entries for  $(\mathcal{M} = \text{RD}, \lambda = 4, \alpha = .3)$  and  $(\mathcal{M} = \text{R2D}, \lambda = 4, \alpha = .2)$  are averaged over 10 trials; all others are over 5. Bold and underlined numbers indicate the lowest regret for fixed  $\lambda$  and  $\alpha$ , respectively. During tuning, we use the same MSS  $\mathcal{M}$  for exploration and evaluation.**

demonstrate the efficacy of one hyperparameter setting  $(\lambda, \alpha)$  across different datasets. However, it is unclear whether these same parameters generalize to other games with different reward structures and scales.

### 7 CONCLUSION

We investigated a novel approach to offline game-solving, which addresses the discovery of equilibria from fixed datasets of multiagent trajectories. These datasets may provide incomplete information on relevant strategy spaces, making game-solving strictly more difficult in the offline setting. Under this added uncertainty, we framed this task as an equilibrium selection problem: of solutions apparently in approximate equilibrium, prefer those likely to minimize actual, true-game regret. We address equilibrium selection under the lens of strategy exploration, which refers to the process of deciding what strategies to include in an iteratively extended game model, by incorporating conservative algorithmic techniques from single-agent, offline RL algorithms.

Our approach, CoffeE-PSRO, extends PSRO by training a dynamics model to use in place of an online simulator. Then, we design a modified response objective that optimizes model-quantified uncertainty given the current response target and potential strategy deviations. Lastly, we propose a novel MSS, R2D, that optimizes conservative regret estimates. We test our algorithm in a general-sum, sequential Bargaining game against several SOTA baselines. Our results demonstrate that CoffeE-PSRO + R2D tends to produce lower regret solutions with higher consistency than SOTA approaches. We also investigated the unique relationship between algorithmic components, behavior policy, and performance in general-sum games, highlighting CoffeE-PSRO’s strategic method of applying conservatism. As a final solver ( $\mathcal{M}_{eval}$ ), R2D tends to identify lower regret solutions than the RD baseline, underscoring the benefits of optimizing for pessimistic regret estimates in offline settings. Lastly, experiments analyzing model fidelity and its relationship to regret demonstrate the importance of balancing conservatism and potential strategic value in strategy exploration.

### REFERENCES

[1] David Balduzzi, Marta Garnelo, Yoram Bachrach, Wojciech Czarnecki, Julien Perolat, Max Jaderberg, and Thore Graepel. 2019. Open-ended learning in symmetric

- zero-sum games. In *36th International Conference on Machine Learning*. 434–443.
- [2] Ariyan Bighashdel, Yongzhao Wang, Stephen McAleer, Rahul Savani, and Frans A. Oliehoek. 2024. Policy space response oracles: A survey. In *33rd International Joint Conference on Artificial Intelligence* (Jeju, Korea). 7951–7961.
  - [3] Jingxiao Chen, Weiji Xie, Weinan Zhang, Yong Yu, and Ying Wen. 2026. Offline fictitious self-play for competitive games. In *40th AAAI Conference on Artificial Intelligence* (Singapore).
  - [4] Qiwen Cui and Simon S. Du. 2022. When are offline two-player zero-sum Markov games solvable?. In *36th Annual Conference on Neural Information Processing Systems*. 25779–25791.
  - [5] Jiechuan Jiang and Zongqing Lu. 2023. Offline decentralized multi-agent reinforcement learning. In *26th European Conference on Artificial Intelligence*. 1148–1155.
  - [6] Patrick R. Jordan, L. Julian Schwartzman, and Michael P. Wellman. 2010. Strategy Exploration in Empirical Games. In *9th International Conference on Autonomous Agents and Multiagent Systems* (Toronto). 1131–1138.
  - [7] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. 2020. Morel: Model-based offline reinforcement learning. In *34th Annual Conference on Neural Information Processing Systems*. 21810–21823.
  - [8] Aviral Kumar, Justin Fu, George Tucker, and Sergey Levine. 2019. Stabilizing off-policy Q-learning via bootstrapping error reduction. In *33rd Annual Conference on Neural Information Processing Systems*.
  - [9] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-learning for offline reinforcement learning. In *34th Annual Conference on Neural Information Processing Systems*. 1179–1191.
  - [10] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. 2019. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453* (2019).
  - [11] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *31st Annual Conference on Neural Information Processing Systems*. 4190–4203.
  - [12] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* (2020).
  - [13] Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? End-to-end learning for negotiation dialogues. In *Conference on Empirical Methods in Natural Language Processing*. 2443–2453.
  - [14] Shuxin Li, Xinrun Wang, Youzhi Zhang, Jakub Cerny, Pengdeng Li, Hau Chan, and Bo An. 2022. Offline equilibrium finding. *arXiv preprint arXiv:2207.05285* (2022).
  - [15] Luke Marris, Paul Muller, Marc Lanctot, Karl Tuyls, and Thore Graepel. 2021. Multi-agent training beyond zero-sum with correlated equilibrium meta-solvers. In *38th International Conference on Machine Learning*. 7480–7491.
  - [16] Paul Muller, Shayegun Omidshafiei, Mark Rowland, Karl Tuyls, Julien Pérolat, Siqi Liu, Daniel Hennes, Luke Marris, Marc Lanctot, Edward Hughes, et al. 2020. A generalized training approach for multiagent learning. In *8th International Conference on Learning Representations*.
  - [17] Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. 2020. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359* (2020).
  - [18] Austin A. Nguyen, Anri Gu, and Michael P. Wellman. 2025. Explicit exploration for high-welfare equilibria in game-theoretic multiagent reinforcement learning. In *42nd International Conference on Machine Learning* (Vancouver). 45988–46007.
  - [19] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177* (2019).
  - [20] Nicolas Perez-Nieves, Yaodong Yang, Oliver Slumbers, David Henry Mguni, Ying Wen, and Jun Wang. 2021. Modelling behavioural diversity for learning in open-ended games. In *38th International Conference on Machine Learning*. 8514–8524.
  - [21] Zhengdao Shao, Liansheng Zhuang, Houqiang Li, and Shafei Wang. 2025. COP-SRO: An Offline Empirical Game Theoretic Method With Conservative Critic. *IEEE Transactions on Neural Networks and Learning Systems* (2025).
  - [22] Max Olan Smith, Thomas Anthony, and Michael P. Wellman. 2021. Iterative empirical game solving via single policy best response. In *9th International Conference on Learning Representations* (Virtual).
  - [23] Peter D. Taylor and Leo B. Jonker. 1978. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40, 1-2 (1978), 145–156.
  - [24] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double Q-learning. In *30th AAAI Conference on Artificial Intelligence*. 2094–2100.
  - [25] Xiangsen Wang, Haoran Xu, Yinan Zheng, and Xianyuan Zhan. 2024. Offline multi-agent reinforcement learning with implicit global-to-local value regularization. *Advances in Neural Information Processing Systems* 36 (2024).
  - [26] Yongzhao Wang, Qiurui Ma, and Michael P. Wellman. 2023. Evaluating strategy exploration in empirical game-theoretic analysis. In *21st International Conference on Autonomous Agents and Multiagent Systems*. 1346–1354.
  - [27] Yongzhao Wang and Michael P. Wellman. 2023. Regularization for Strategy Exploration in Empirical Game-Theoretic Analysis. In *22nd International Conference on Autonomous Agents and Multiagent Systems*. 2484–2486.
  - [28] Yongzhao Wang and Michael P. Wellman. 2024. Generalized response objectives for strategy exploration in empirical game-theoretic analysis. In *23rd International Conference on Autonomous Agents and Multiagent Systems*. 1892–1900.
  - [29] Michael P. Wellman, Karl Tuyls, and Amy Greenwald. 2025. Empirical Game-Theoretic Analysis: A Survey. *Journal of Artificial Intelligence Research* 82 (2025), 1017–1076.
  - [30] Yifan Wu, George Tucker, and Ofir Nachum. 2019. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361* (2019).
  - [31] Yaodong Yang and Jun Wang. 2021. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583v3* (2021).
  - [32] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. 2020. MOPO: Model-based offline policy optimization. In *34th Annual Conference on Neural Information Processing Systems*. 14129–14142.
  - [33] Yuheng Zhang, Yu Bai, and Nan Jiang. 2023. Offline learning in Markov games with general function approximation. In *40th International Conference on Machine Learning*. PMLR, 40804–40829.
  - [34] Han Zhong, Wei Xiong, Jiyuan Tan, Liwei Wang, Tong Zhang, Zhaoran Wang, and Zhuoran Yang. 2022. Pessimistic minimax value iteration: Provably efficient equilibrium learning from offline datasets. In *39th International Conference on Machine Learning*. PMLR, 27117–27142.