

NASH POLICY GRADIENT: A POLICY GRADIENT METHOD WITH ITERATIVELY REFINED REGULARIZATION FOR FINDING NASH EQUILIBRIA

Anonymous authors

Paper under double-blind review

ABSTRACT

Finding Nash equilibria in imperfect-information games remains a central challenge in multi-agent reinforcement learning. While regularization-based methods have recently achieved last-iteration convergence to a regularized equilibrium, they require the regularization strength to shrink toward zero to approximate a Nash equilibrium, often leading to unstable learning in practice. Instead, we fix the regularization strength at a large value for robustness and achieve convergence by iteratively refining the reference policy. Our main theoretical result shows that this procedure guarantees strictly monotonic improvement and convergence to an exact Nash equilibrium in two-player zero-sum games, without requiring a uniqueness assumption. Building on this framework, we develop a practical algorithm, *Nash Policy Gradient* (NASHPG), which preserves the generalizability of policy gradient methods while relying solely on the current and reference policies. Empirically, NASHPG achieves comparable or lower exploitability than prior model-free methods on classic benchmark games and scales to large domains such as *Battleship* and *No-Limit Texas Hold'em*, where NASHPG consistently attains higher Elo ratings.

1 INTRODUCTION

The success of Deep Reinforcement Learning has sparked significant interest in Multi-agent Reinforcement Learning (MARL) (Busoniu et al., 2008; Shoham et al., 2007), where algorithms must handle environments with multiple interacting agents. While impressive results have been achieved in settings with fully known environment models and perfect-information (Silver et al., 2016; 2018), learning in imperfect-information games remains challenging, as agents must reason under uncertainty about their opponents to approach Nash equilibria.

In multi-agent systems, Nash equilibria serve as the central solution concept: no agent can unilaterally improve its payoff. In two-player zero-sum games, these equilibria correspond to strategies that are unexploitable even against optimal adversaries. Traditionally, the most prominent approaches in these games have been *CFR-inspired methods* (Li et al., 2021) and *population-based methods* (Bighashdel et al., 2024). However, both come with drawbacks: CFR-inspired methods requires access to an exact environment model to perform well, while population-based methods are computationally expensive as they require training multiple agents from scratch. Further, both approaches typically offer only average-iteration convergence guarantees, which are inconvenient in deep learning contexts. These limitations hinder their adoption as general-purpose solvers.

Recently, a new class of regularization-based methods (Liu et al., 2023) has emerged. These approaches add a regularization term in objective function (Hofbauer & Hopkins, 2005), which guarantees last-iteration convergence to a regularized equilibrium. To approximate Nash equilibria, they typically rely on setting the regularization strength to be small or gradually decaying it. However, these methods have yet to demonstrate strong empirical success beyond tabular settings where exact, non-sampled feedback is accessible. A key challenge is that when the regularization strength becomes too small, noise in stochastic gradient estimations can dominate the diminishing regularization signal, leading to cycling or even divergence. As a result, training is highly sensitive to careful tuning of both the regularization parameter and the learning rate.

Motivated by these challenges, we propose an alternative: instead of decaying the regularization strength, we fix it at a large value for robustness and achieve convergence by iteratively refining the reference policy. Our main theoretical result shows that this iterative-refinement procedure ensures strictly monotonic improvement and convergence to a Nash equilibrium, without assuming uniqueness of the Nash equilibrium.

Theorem 1 (Informal). *For a two-player zero-sum game, let z^* be any Nash equilibrium. Starting from an initial policy z_0 , our proposed iterative-refinement procedure generates a sequence of policies $\{z_t\}$ such that*

$$B_\psi(z^*; z_t) > B_\psi(z^*; z_{t+1}), \quad (1)$$

whenever convergence has not yet occurred. Here, B_ψ denotes the Bregman divergence with generating function ψ . Moreover, the sequence is guaranteed to converge to a Nash equilibrium, i.e., $\lim_{t \rightarrow \infty} z_t$ is a Nash equilibrium.

Building on this theory, we adapt the procedure into a practical MARL algorithm, which we call *Nash Policy Gradient* (NASHPG). NASHPG retains the advantages of standard policy gradient methods, allowing decentralized updates where each agent updates its policy without explicitly reasoning about others. Additionally, it is lightweight, requiring only the maintenance of current and reference policies. Despite departing from the exact theoretical setting, NASHPG empirically converges to Nash equilibria. In experiments, we show that NASHPG achieves comparable or lower exploitability than prior model-free methods on classic benchmark games. To demonstrate scalability, we further evaluate NASHPG on large-scale real-world domains such as *Battleship* and *No-Limit Texas Hold'em*, where it consistently achieves higher Elo ratings.

2 RELATED WORK

Existing methods for solving two-player zero-sum imperfect-information games in MARL can be broadly partitioned into three categories.

Population-Based Methods. Population-based methods, such as Neural Fictitious Self-Play (NFSP) (Heinrich & Silver, 2016) and Policy Space Response Oracles (PSRO) (Lanctot et al., 2017), iteratively compute best-responses against mixtures of prior policies using deep reinforcement learning (DRL). These model-free approaches scale to complex domains like StarCraft (Vinyals et al., 2019). However, they require repeated best-response computations, resulting in high computational overhead. Moreover, they only guarantee average-iterate convergence where policy extraction is non-trivial in DRL settings, often requiring supervised learning or the storage of historical agents to obtain a final policy. Recent works have focused on accelerating empirical convergence (McAleer et al., 2024; 2020).

CFR-Inspired Methods. Counterfactual Regret Minimization (CFR) (Zinkevich et al., 2007), most notable for its success in poker games, enabled superhuman performance in heads-up limit Texas Hold'em (Brown & Sandholm, 2018). Deep CFR (Brown et al., 2019) integrates function approximation for larger state spaces. However, CFR relies on explicit game models for tree traversal, which limits its generalizability, and often relies on high-variance sampling to estimate counterfactual values. While some model-free variants do exist (Steinberger et al., 2020), they still suffer from large variance issues, and variance-reduction techniques (McAleer et al., 2023) can introduce exploration biases. Like population-based methods, CFR approaches typically guarantee only average-iterate convergence, requiring additional steps to extract a final policy.

Regularization-Based Methods. Recent regularization-based algorithms address last-iterate convergence by reformulating games with strongly convex-concave objectives, typically converging to entropy-regularized Nash equilibria such as Quantal Response Equilibria (QRE) (McKelvey & Palfrey, 1995). Several works (Cen et al., 2021; Sokota et al., 2023; Liu et al., 2023) establish linear convergence rates to QRE, demonstrating the efficiency of solving regularized games. This has spurred efforts to reach exact Nash equilibria, often through decaying regularization strengths or adaptive step sizes (Liu et al., 2023; Hsieh et al., 2021; Li et al., 2020). While effective in tabular settings with exact feedback, it remains unclear whether these techniques are robust to stochastic sampling errors, which risk instability or cycling (Daskalakis et al., 2018). We include an empirical study of this issue in Appendix C.1. Some theoretical convergence to Nash equilibria has been shown (Daskalakis & Panageas, 2019; Lee et al., 2021; Pérolat et al., 2021), but these results

require the assumption of a unique Nash equilibrium or that all Nash equilibria are interior. Empirically, regularization-based methods remain underexplored beyond normal-form games or simple extensive-form benchmarks (e.g., Kuhn or Leduc poker) where exact feedback is accessible.

3 PRELIMINARIES

In this paper, we focus on analyzing two-player zero-sum games. We first establish the necessary background and notation for extensive-form games, then present their connection to variational inequality problems and regularized variants that form the foundation of our approach. We use standard mathematical notation for concepts such as L-smoothness, strong convexity, monotonicity, and Bregman divergence (see Appendix A.1 for formal definitions).

3.1 GAME THEORY BACKGROUND

We consider two-player zero-sum **extensive-form games** (EFG) with perfect recall. The game is represented by a finite game tree with the following components: Let the players be denoted by $i \in \{1, 2\}$. The set of states is denoted by \mathcal{S} , with terminal states $\mathcal{Q} \subseteq \mathcal{S}$ determining payoffs, where player 1’s payoff is the negative of player 2’s payoff. Each non-terminal state $s \in \mathcal{S} \setminus \mathcal{Q}$ is controlled by either a player or chance. For each player i , decision nodes are partitioned into information sets $I \in \mathcal{I}_i$, where the player cannot distinguish between states within the same information set. The set of actions available at information set I is denoted by $A(I)$. A behavioral strategy for player i is a mapping $\pi^{(i)}: \mathcal{I}_i \rightarrow \Delta(A(I))$, assigning a probability distribution over actions at each information set. A strategy profile π is a collection of both players’ strategies, and $\pi^{(-i)}$ refers to player i ’s opponent strategy.

A two-player zero-sum extensive-form game induces an equivalent **normal-form game** (Kuhn, 1953) with payoff matrix \mathbf{A} , where each row and column correspond to a pure strategy (deterministic action plan across all information sets) for players 1 and 2, respectively. A mixed strategy is a probability distribution over pure strategies. Let \mathcal{X} and \mathcal{Y} denote the mixed strategy spaces for players 1 and 2, with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ representing their mixed strategies. The payoff function $f(x, y) = x^\top \mathbf{A}y$ gives player 1’s expected payoff, similarly, player 2’s expected payoff is $-f(x, y)$. When referring to mixed strategies in EFGs, we refer to those defined through this normal-form representation.

A **Nash equilibrium** (NE) is a strategy profile $\pi^* = (\pi^{(1)*}, \pi^{(2)*})$ in behavioral strategies or $z^* = (x^*, y^*)$ in mixed strategies, where each player’s strategy is a best-response to the other’s. Formally, for the normal-form representation, a Nash equilibrium satisfies:

$$x^* \in \arg \max_{x \in \mathcal{X}} f(x, y^*) \quad \text{and} \quad y^* \in \arg \max_{y \in \mathcal{Y}} -f(x^*, y)$$

That is, no player can improve their expected payoff by unilaterally deviating from their equilibrium strategy. Our goal is to develop efficient algorithms to find Nash equilibria in these games.

3.2 CONNECTION WITH VARIATIONAL INEQUALITY

Finding the Nash equilibrium is equivalent to solving the variational inequality (VI) problem, which forms the foundation for understanding both the challenges and opportunities in our approach.

Definition 1 (Variational Inequality Problem). *Given a closed convex set $\mathcal{Z} \subseteq \mathbb{R}^n$ and a mapping $G: \mathcal{Z} \rightarrow \mathbb{R}^n$, the variational inequality problem $\text{VI}(\mathcal{Z}, G)$ is to find $z^* \in \mathcal{Z}$ such that*

$$\langle G(z^*), z - z^* \rangle \geq 0 \quad \forall z \in \mathcal{Z}. \quad (2)$$

In a two-player zero-sum extensive-form game, let \mathcal{X} and \mathcal{Y} denote the mixed strategy spaces for players 1 and 2, respectively, and let \mathbf{A} be the payoff matrix of the equivalent normal-form game. We recall the fundamental connection between Nash equilibria and variational inequalities by defining $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ with strategy profile $z = (x, y) \in \mathcal{Z}$ and the operator $F: \mathcal{Z} \rightarrow \mathbb{R}^n$:

$$F(z) = \begin{bmatrix} -\nabla_x f(x, y) \\ \nabla_y f(x, y) \end{bmatrix} \quad (3)$$

where $f(x, y) = x^\top \mathbf{A}y$ is the payoff function.

Theorem 2 (Nash-VI Equivalence). *A strategy profile $z^* \in \mathcal{Z}$ is a Nash equilibrium if and only if it solves the variational inequality problem $\text{VI}(\mathcal{Z}, F)$ (Facchinei & Pang, 2003, Section 1.4.2):*

$$\langle F(z^*), z - z^* \rangle \geq 0 \quad \forall z \in \mathcal{Z}. \quad (4)$$

Thus, computing Nash equilibrium reduces to solving $\text{VI}(\mathcal{Z}, F)$. However, this creates a significant computational challenge: *The operator F is monotone but not strongly monotone, making $\text{VI}(\mathcal{Z}, F)$ difficult to solve directly.*

3.3 REGULARIZED VARIATIONAL INEQUALITY

To address the lack of strong monotonicity, Sokota et al. (2023) consider a *regularized* VI problem $\text{VI}(\mathcal{Z}, G_\rho)$, where a strongly convex regularizer is added to the operator F . The resulting operator G_ρ becomes *strongly monotone*, ensuring a unique solution (Larsson & Patriksson, 1994) and enabling efficient algorithms.

Formally, let $\rho = (\rho_1, \rho_2) \in \mathcal{Z}$ be a reference strategy profile. With a strongly convex function $\psi(z) = \psi_1(x) + \psi_2(y)$, the associated Bregman divergence is

$$B_\psi(z; \rho) = B_{\psi_1}(x; \rho_1) + B_{\psi_2}(y; \rho_2), \quad (5)$$

where ψ_1 and ψ_2 are strongly convex for each player. The regularized operator $G_\rho : \mathcal{Z} \rightarrow \mathbb{R}^n$ is

$$G_\rho(z) = F(z) + \alpha \nabla_z B_\psi(z; \rho), \quad (6)$$

with regularization parameter $\alpha > 0$. The corresponding variational inequality seeks $\hat{z} \in \mathcal{Z}$ such that

$$\langle G_\rho(\hat{z}), z - \hat{z} \rangle \geq 0 \quad \forall z \in \mathcal{Z}. \quad (7)$$

Under Assumption 1, The VI problem $\text{VI}(\mathcal{Z}, G_\rho)$ can be solved efficiently by Magnetic Mirror Descent (MMD) (Sokota et al., 2023):

Assumption 1 (MMD’s Convergence Conditions).

1. ψ is μ -strongly convex over \mathcal{Z} and differentiable on $\text{int dom } \psi$ for some $\mu > 0$.
2. $z_{t+1} \in \text{int dom } \psi$ for all iterations.
3. F is L -smooth and α, η satisfy $\alpha \geq \mu\eta L^2$.

Algorithm 1 (Magnetic Mirror Descent (MMD)). *Initialize $z_0 = (x_0, y_0) \in \text{int dom } \psi \cap \mathcal{Z}$, reference $\rho = (\rho_1, \rho_2) \in \text{int dom } \psi$, and step size $\eta > 0$. At each iteration t :*

$$\begin{aligned} x_{t+1} &= \arg \max_{x \in \mathcal{X}} \left\{ \eta \left(\langle \nabla_x f(x_t, y_t), x \rangle - \alpha B_{\psi_1}(x; \rho_1) \right) - B_{\psi_1}(x; x_t) \right\}, \\ y_{t+1} &= \arg \min_{y \in \mathcal{Y}} \left\{ \eta \left(\langle \nabla_y f(x_t, y_t), y \rangle + \alpha B_{\psi_2}(y; \rho_2) \right) + B_{\psi_2}(y; y_t) \right\}. \end{aligned}$$

With Assumption 1, MMD converges to the unique solution of $\text{VI}(\mathcal{Z}, G_\rho)$.

However, this creates a fundamental limitation of using MMD as a Nash equilibrium solver. To recover the exact equilibrium, one must shrink $\alpha \rightarrow 0$ so that G_ρ closely approximates the original operator F . Yet Assumption 1 requires $\alpha \geq \mu\eta L^2$, forcing either vanishingly small stepsizes (and thus negligible progress) or risking divergence when the condition is violated. Conversely, keeping α large ensures stability and efficient updates but biases the solution away from the true equilibrium, since G_ρ no longer approximates F . Thus, there is an inherent tension between approximation quality (small α) and algorithmic stability (large α). This motivates our approach, which breaks this tradeoff by allowing large α while still guaranteeing convergence to a Nash equilibrium.

4 METHODOLOGY

We propose a new approach to computing Nash equilibria based on *iteratively refined regularization*. Unlike prior methods that rely on shrinking the regularization strength α to approximate the original

VI problem $\text{VI}(\mathcal{Z}, F)$, we instead solve a sequence of regularized problems $\text{VI}(\mathcal{Z}, G_\rho)$, where the solution from each step becomes the reference strategy profile ρ for the next.

Our approach unfolds in two stages. First, we establish the theoretical basis of this iterative procedure (Sections 4.1–4.2), proving strictly monotone improvement and convergence to a Nash equilibrium without assuming uniqueness. Next, we derive a practical algorithm, *Nash Policy Gradient* (NASHPG), via a series of approximations (Section 4.3), yielding a method that integrates naturally with reinforcement learning (RL) frameworks.

4.1 REGULARIZED VI AS AN OPERATOR

We begin by viewing the regularized VI problem as an operator that maps a reference strategy profile ρ to the unique solution of $\text{VI}(\mathcal{Z}, G_\rho)$.

Definition 2 (Regularized VI Operator). *Define the operator $\mathcal{M} : \text{int dom } \psi \cap \mathcal{Z} \rightarrow \text{int dom } \psi \cap \mathcal{Z}$, where the input is the reference strategy profile ρ and the output $\mathcal{M}(\rho)$ is the unique solution of $\text{VI}(\mathcal{Z}, G_\rho)$. We assume $\mathcal{M}(\rho) \in \text{int dom } \psi \cap \mathcal{Z}$.*

Our theoretical results show that this operator satisfies three key properties.

First, applying \mathcal{M} never increases the Bregman divergence to a Nash equilibrium. This establishes the soundness of \mathcal{M} as a refinement step.

Lemma 1 (Distance Non-increase Property). *Given $\alpha > 0$, let $\rho \in \text{int dom } \psi \cap \mathcal{Z}$ and z^* be any Nash equilibrium (i.e., a solution of $\text{VI}(\mathcal{Z}, F)$). Then:*

$$B_\psi(z^*; \rho) \geq B_\psi(z^*; \mathcal{M}(\rho)) + B_\psi(\mathcal{M}(\rho); \rho). \quad (8)$$

(Proof in Appendix A.4.1.)

Second, the fixed points of \mathcal{M} are precisely the Nash equilibria. In other words, once the procedure reaches a Nash equilibrium, it remains there, this naturally characterizes the termination criterion for our refinement process.

Lemma 2 (Fixed Point Characterization). *A strategy profile $z^* \in \mathcal{Z}$ is a Nash equilibrium if and only if it is a fixed point of the regularized VI operator: $z^* = \mathcal{M}(z^*)$.*

(Proof in Appendix A.4.2.)

Finally, the operator \mathcal{M} is continuous, which will be useful in establishing convergence guarantees in the next section.

Lemma 3 (Continuity of \mathcal{M}). *Assume ψ is μ -strongly convex and continuously differentiable on $\text{int dom } \psi$ for some $\mu > 0$. Then the operator \mathcal{M} is continuous.*

(Proof in Appendix A.4.3.)

4.2 ITERATIVE REFINEMENT TOWARD EQUILIBRIUM

Motivated by the above properties, we propose the following procedure: starting from an arbitrary strategy profile, repeatedly apply the operator \mathcal{M} until convergence.

Algorithm 2 (Iterative \mathcal{M} Method). *Starting with $z_0 \in \text{int dom } \psi \cap \mathcal{Z}$, at each iteration t compute:*

$$z_{t+1} = \mathcal{M}(z_t) \quad (9)$$

Where we assume $z_t \in \text{int dom } \psi \cap \mathcal{Z}$ for all iterations.

The convergence guarantee follows from Lemmas 1 and 2 and the continuity of \mathcal{M} . We show that each refinement step *strictly* reduces the Bregman divergence to any Nash equilibrium until convergence is reached.

Theorem 3 (Convergence of Iterative \mathcal{M}). *Let z^* be any Nash equilibrium. Algorithm 2 generates a sequence $\{z_t\}$ such that*

$$B_\psi(z^*; z_t) > B_\psi(z^*; z_{t+1}), \quad (10)$$

whenever convergence has not yet occurred. Moreover, the sequence is guaranteed to converge to a Nash equilibrium, i.e., $\lim_{t \rightarrow \infty} z_t$ is a Nash equilibrium.

(Proof in Appendix A.4.4.)

To implement this iterative method, we employ MMD to solve each regularized VI subproblem. Under Assumption 1, this yields a theoretically well-justified algorithm for finding Nash equilibria in two-player zero-sum games (Algorithm 1). Crucially, our approach introduces no additional assumptions beyond those required for MMD convergence, and we allows α to be set to large values that easily satisfy the convergence constraint $\alpha \geq \mu\eta L^2$, while still guaranteeing convergence to a Nash equilibrium.

Algorithm 1 MMD-based Iterative \mathcal{M} Method

- 1: Initialize $x_0 \in \text{int dom } \psi_1 \cap \mathcal{X}$
 - 2: Initialize $y_0 \in \text{int dom } \psi_2 \cap \mathcal{Y}$
 - 3: Set $\rho_1 \leftarrow x_0, \rho_2 \leftarrow y_0$
 - 4: **for** $t = 0, 1, 2, \dots$ until convergence **do**
 - 5: Set $x_{t,0} \leftarrow \rho_1, y_{t,0} \leftarrow \rho_2$
 - 6: **for** $k = 0, 1, 2, \dots$ until convergence **do**
 - 7: $x_{t,k+1} \leftarrow \arg \max_{x \in \mathcal{X}} \{ \eta (\langle \nabla_{x_{t,k}} f(x_{t,k}, y_{t,k}), x \rangle - \alpha B_{\psi_1}(x; \rho_1)) - B_{\psi_1}(x; x_{t,k}) \}$
 - 8: $y_{t,k+1} \leftarrow \arg \min_{y \in \mathcal{Y}} \{ \eta (\langle \nabla_{y_{t,k}} f(x_{t,k}, y_{t,k}), y \rangle + \alpha B_{\psi_2}(y; \rho_2)) + B_{\psi_2}(y; y_{t,k}) \}$
 - 9: Update $\rho_1 \leftarrow x_{t,K}, \rho_2 \leftarrow y_{t,K}$ \triangleright where K is the final inner iteration
-

4.3 PRACTICAL IMPLEMENTATION: NASH POLICY GRADIENT ALGORITHM

Although Algorithm 1 provides theoretical guarantees, its reliance on mixed strategies makes it unclear how to mitigate the exponential complexity in the size of the game tree. To address this computational barrier, we develop a practical approximation that preserves the iterative refinement structure while enabling efficient implementation in RL frameworks.

From Mixed Strategies to Behavioral Strategies. We illustrate the transformation for player 1, the case for player 2 is symmetric. Consider the mirror-ascent subproblem in line 7 of Algorithm 1. Setting ψ_1 to the negative entropy (the canonical choice on the probability simplex) shows that, for sufficiently small step size η , the mirror step is first-order equivalent to a gradient ascent step on the regularized objective:

$$g(x_{t,k}) = f(x_{t,k}, y_{t,k}) - \alpha D_{\text{KL}}(x_{t,k} \parallel \rho_1). \quad (11)$$

This mixed-strategy formulation inspired a behavioral strategy analogue. For behavioral strategy profile $\pi = (\pi^{(1)}, \pi^{(2)})$ and reference behavioral strategy profile $\rho = (\rho^{(1)}, \rho^{(2)})$, we design the analogous objective:

$$g(\pi^{(1)}) = \mathbb{E}_{\tau \sim \pi} [R_1(\tau)] - \alpha \mathbb{E}_{o \sim \pi} \left[D_{\text{KL}}(\pi^{(1)}(\cdot \mid o) \parallel \rho^{(1)}(\cdot \mid o)) \right] \quad (12)$$

where $R_1(\tau)$ denotes player 1’s payoff along trajectory τ (a path of actions to a terminal node), o denotes an observation (information set) of player 1, and both trajectories τ and observations o are sampled under the strategy profile π . The first term preserves the expected payoff structure, while the regularization term aggregates Kullback–Leibler (KL) divergences across all observations. This design choice for regularization enables efficient estimation within RL frameworks, though it introduces a heuristic element by depending on the opponent’s strategy $\pi^{(2)}$.

Connection to Partially Observable MDPs. The shift to behavioral strategies allows a crucial reduction to single-agent RL. When opponent’s strategy $\pi^{(-i)}$ is fixed, the game reduces to a Partially Observable Markov Decision Process (POMDP) from player i ’s perspective (Greenwald et al., 2013). In this view, a behavioral strategy is equivalent to a stochastic policy in POMDP. Thus, we can leverage any single-agent policy gradient method to optimize the regularized objective in Equation 12. Specifically, the policy and KL regularization gradients can both be estimated via trajectory sampling.

Practical Algorithm. These insights lead to our main algorithmic contribution, *Nash Policy Gradient* (NASHPG) presented in Algorithm 2. At each iteration, trajectories are collected under the

current joint policy. Both players then use these trajectories to estimate their policy and KL regularization gradients, and perform regularized policy gradient updates. After K such updates, the reference policies are reset to the updated policies, and this process is repeated for T outer iterations. The parameters K and T can be chosen according to the available computational budget.

A key advantage of NASHPG is its modularity: it is agnostic to the choice of policy gradient estimator and can be seamlessly combined with modern single-agent RL methods. For instance, when adopting PPO, we simply replace the vanilla policy gradient update with PPO’s clipped surrogate objective, while leaving the regularization term unchanged. This design allows NASHPG to directly inherit algorithmic advances from the single-agent RL literature.

Algorithm 2 Nash Policy Gradient (NASHPG)

- 1: Initialize policy parameters $\{\theta^{(i)}\}_{i=1}^2$ for policies $\{\pi_{\theta^{(i)}}\}_{i=1}^2$
 - 2: Set reference policies $\rho^{(i)} \leftarrow \pi_{\theta^{(i)}}$ for $i \in \{1, 2\}$
 - 3: **for** $t = 0, 1, \dots, T - 1$ **do** ▷ Outer loop: refining reference policies
 - 4: **for** $k = 0, 1, \dots, K - 1$ **do** ▷ Inner loop: regularized policy gradient
 - 5: Sample trajectories $\{\tau\}$ by executing $\pi = (\pi_{\theta^{(1)}}, \pi_{\theta^{(2)}})$ in the environment
 - 6: **for** each player $i \in \{1, 2\}$ **do**
 - 7: Estimate policy gradient: $\hat{g}^{(i)} \approx \nabla_{\theta^{(i)}} \mathbb{E}_{\tau \sim \pi} [R_i(\tau)]$
 - 8: Estimate regularization gradient: $\hat{g}_{\text{reg}}^{(i)} \approx \nabla_{\theta^{(i)}} \mathbb{E}_{o \sim \pi} [D_{\text{KL}}(\pi_{\theta^{(i)}}(\cdot | o) \| \rho^{(i)}(\cdot | o))]$
 - 9: Update player i : $\theta^{(i)} \leftarrow \theta^{(i)} + \eta(\hat{g}^{(i)} - \alpha \hat{g}_{\text{reg}}^{(i)})$
 - 10: Update reference policies: $\rho^{(i)} \leftarrow \pi_{\theta^{(i)}}$ for $i \in \{1, 2\}$
-

5 EXPERIMENTS

This section outlines our experimental setup and highlights the main empirical results. Comprehensive settings and extra studies are deferred to Appendix B. All code needed to reproduce our experiments are publicly available at `nash_policy_gradient`.

5.1 EXPERIMENTAL SETUP

Environments: We evaluate all algorithms across seven two-player games spanning from classic benchmark environments to widely played real-world games. **Benchmarks** include Kuhn Poker, Leduc Poker, Dark Hex (3×3) and Phantom Tic-Tac-Toe (Lanctot et al., 2019). To probe scalability beyond synthetic benchmarks, we also evaluate on **real-world games**: Liar’s Dice, Battleship, and Heads-Up No-Limit Texas Hold’em. These environments reflect authentic gameplay dynamics and are dramatically more complex. Complete rules, parameter settings, and complexity statistics (e.g., number of non-terminal histories) are deferred to Appendix B.3.

Baseline Algorithms: We compare NASHPG against three established model-free baselines: NFSP (Heinrich & Silver, 2016), PSRO (Lanctot et al., 2017), and MMD (Sokota et al., 2023). NFSP and PSRO train their best-response policies against a mixture of past agents using PPO (Schulman et al., 2017). For MMD, we set its magnetic coefficient α to 0.05 and use a PPO-based policy gradient update.

Evaluation Metrics: We assess performance using *exploitability* and *Elo ratings*. Since exact best-response policies are intractable in large games, we approximate exploitability by training best-response agents using PPO. For Elo ratings, we save a checkpoint every 2,000 training steps for each algorithm and run a Swiss-style tournament across all checkpoints, yielding roughly 800 tournament entries (4 algorithms \times repeat 4 runs \times 50 checkpoints). Detailed definitions and implementation are provided in Appendix B.2.

Implementation: All environments and algorithms are implemented in JAX (Bradbury et al., 2018) for efficient training and evaluation. Within each game, all methods share identical model architectures and computational budgets to ensure fair comparison, and training follows standard PPO settings. Each experiment is repeated 4 times with different random seeds, and we report results as

mean \pm standard deviation. Detailed architectures, hyperparameters, and reward formulations are provided in Appendix B.1 and Appendix B.3.

5.2 HYPERPARAMETER TUNING

Regularization Strength	Kuhn Poker	Leduc Poker
$\alpha = 0.1$	0.0138 ± 0.0088	0.0075 ± 0.0068
$\alpha = 0.2$	0.0046 ± 0.0037	0.0028 ± 0.0037
$\alpha = 0.4$	0.0059 ± 0.0032	0.0029 ± 0.0026

Table 1: Final exploitability of NASHPG across different regularization strengths α on Kuhn Poker and Leduc Poker after 50,000 training steps (mean \pm standard deviation over 4 runs).

We first study the impact of the regularization coefficient α for NASHPG, varying α over the values 0.1, 0.2, and 0.4 in Kuhn Poker and Leduc Poker, which serve as representative small-scale domains. Table 1 reports the final exploitability for each setting. Both $\alpha = 0.2$ and $\alpha = 0.4$ yield stable training dynamics, with $\alpha = 0.2$ achieving the lowest exploitability. For all subsequent experiments, we use $\alpha = 0.2$ for NASHPG.

5.3 RESULTS IN BENCHMARKS

Figure 1 plots the exploitability curves for Kuhn Poker, Leduc Poker, Dark Hex 3×3 , and Phantom Tic-Tac-Toe. All algorithms converge in Kuhn and Leduc Poker but with different convergence speeds: both MMD and NASHPG converge faster than the population-based methods. In Dark Hex 3×3 and Phantom Tic-Tac-Toe, as the environments become more complex, NASHPG achieves the lowest exploitability, whereas MMD plateaus at a higher value, reflecting the influence of its regularization term, which biases the converged policy. Population-based algorithms exhibit slow convergence in these games, with NFSP attains the highest exploitability across all games and fails to converge entirely on Dark Hex 3×3 . Overall, NASHPG demonstrates scalability and consistently fast convergence as game complexity grows, motivating evaluation on more challenging, real-world environments.

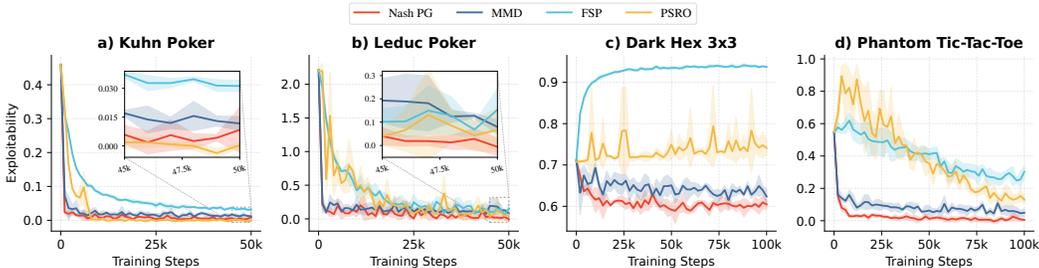


Figure 1: Exploitability curves for our benchmarks. NASHPG consistently achieves lower final exploitability across all tested environments except Kuhn Poker. Error bars show standard deviation over 4 runs.

5.4 EVALUATION ON REAL-WORLD GAMES

Limitations of Exploitability in Large Real-World Games. Figure 2 presents exploitability curves for Liar’s Dice, Battleship, and Heads-Up No-Limit Texas Hold’em. As game complexity grows, exploitability becomes increasingly unreliable under limited computational budgets. In Figure 2b-c, measured exploitability drops below zero; in particular, NASHPG plateaus near -1 in Figure 2b, indicating that our best-response approximation either receives diminishing learning signals against strong policies (where random play has a negligible winning probability) or that the training budget (2,000 steps) is insufficient to estimate a meaningful best-response in such domains. These observations motivate adopting Elo ratings as an alternative metric.

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

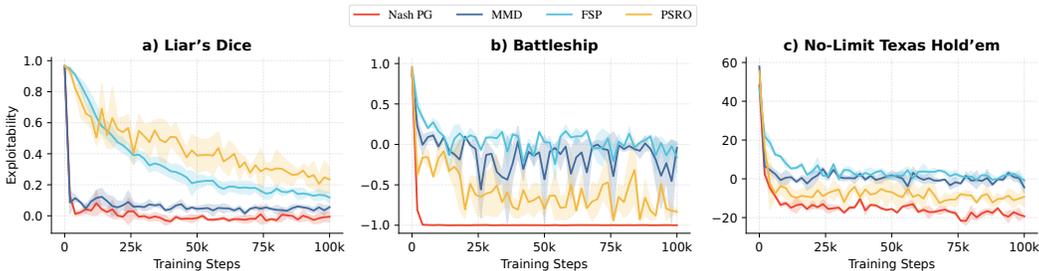


Figure 2: Exploitability in real-world games. Battleship and Texas Hold'em show negative values (theoretically impossible), highlighting the limitations of exploitability-based evaluation under limited computational budgets in complex domains.

Elo Ratings as a Measure of Relative Strength. Given the limitations of exploitability, we adopt Elo ratings as the alternative evaluation metric in large real-world games. Widely used in competitive games, this metric captures relative skill through head-to-head matchups without requiring best-response computation. Implementation details are provided in Appendix B.2. Figure 3 presents Elo rating curves for the these games. The results show clear and consistent trends favoring NASHPG, which achieves the highest final Elo ratings across all three games. Interestingly, PSRO performs well in Heads-Up Texas Hold'em, even surpassing NASHPG during the early iterations (< 25,000 steps), yet fails to maintain this trend in the other games, suggesting that population-based methods are more selective with respect to game type.

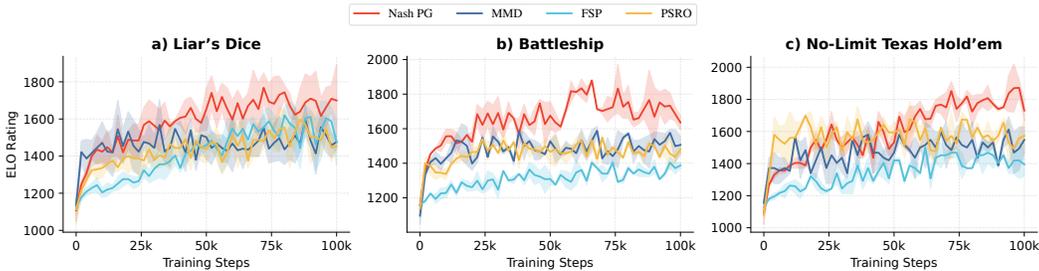


Figure 3: Elo ratings in large real-world games. NASHPG achieves consistently higher final Elo ratings across complex domains where exploitability is unreliable. Error bars denote standard deviation over 4 runs.

6 CONCLUSION AND FUTURE WORK

We proposed a regularization-based framework that achieves convergence to Nash equilibria by iteratively refining the reference policy. We proved strictly monotonic improvement of the solution over iterations, and established convergence guarantees for our algorithm. Inspired by these theoretical findings, we developed NASHPG, a practical algorithm that closely mirrors standard policy gradient methods while adding a KL regularization term and periodically updated reference policy. Empirically, NASHPG achieves low exploitability and scales effectively to large domains such as *Battleship* and *No-Limit Texas Hold'em*.

A central appeal of NASHPG is its accessibility: the method requires only minor modifications to conventional policy gradient algorithms, lowering the barrier to entry for practitioners in MARL. For future work, several directions remain open. On the theoretical side, our heuristic approximations would benefit from theoretically grounded justification. Empirically, benchmarking NASHPG against domain-specific state-of-the-art methods, such as Pluribus (Brown & Sandholm, 2019), in large poker settings would further clarify its strengths. Extending it to continuous-action domains and to high-dimensional environments, such as Atari, would provide a further test of scalability.

REFERENCES

- 486
487
488 Ariyan Bighashdel, Yongzhao Wang, Stephen McAleer, Rahul Savani, and Frans A. Oliehoek. Pol-
489 icy space response oracles: A survey. In *IJCAI*, pp. 7951–7961. ijcai.org, 2024.
- 490
491 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
492 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao
493 Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.
- 494
495 George W Brown. Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation*, 13(1):
496 374, 1951.
- 497
498 Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats
499 top professionals. *Science*, 359(6374):418–424, 2018. doi: 10.1126/science.aao1733. URL
500 <https://www.science.org/doi/abs/10.1126/science.aao1733>.
- 501
502 Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):
503 885–890, 2019. doi: 10.1126/science.aay2400. URL <https://www.science.org/doi/abs/10.1126/science.aay2400>.
- 504
505 Noam Brown, Adam Lerer, Sam Gross, and Tuomas Sandholm. Deep counterfactual regret min-
506 imization. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 793–802.
507 PMLR, 2019.
- 508
509 Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent rein-
510 forcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications
511 and Reviews)*, 38(2):156–172, 2008.
- 512
513 Shicong Cen, Yuting Wei, and Yuejie Chi. Fast policy extragradient methods for competitive games
514 with entropy regularization. In *NeurIPS*, pp. 27952–27964, 2021.
- 515
516 Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and
517 constrained min-max optimization. In *ITCS*, volume 124 of *LIPICs*, pp. 27:1–27:18. Schloss
518 Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 519
520 Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with
521 optimism. In *ICLR (Poster)*. OpenReview.net, 2018.
- 522
523 Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and comple-
524 mentarity problems*. Springer, 2003.
- 525
526 Amy Greenwald, Jiayu Li, Eric Sodomka, and Michael Littman. Solving for best responses in
527 extensive-form games using reinforcement learning methods. In *Proceedings of the conference
528 on reinforcement learning and decision making*, pp. 116–120. Citeseer, 2013.
- 529
530 Johannes Heinrich and David Silver. Deep reinforcement learning from self-play in imperfect-
531 information games. *CoRR*, abs/1603.01121, 2016.
- 532
533 Josef Hofbauer and Ed Hopkins. Learning in perturbed asymmetric games. *Games Econ. Behav.*,
534 52(1):133–152, 2005.
- 535
536 Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. Adaptive learning in con-
537 tinuous games: Optimal regret bounds and convergence to nash equilibrium. In *COLT*, volume
538 134 of *Proceedings of Machine Learning Research*, pp. 2388–2422. PMLR, 2021.
- 539
540 Harold W Kuhn. Extensive games and the problem of information. *Contributions to the Theory of
541 Games*, 2(28):193–216, 1953.
- 542
543 Marc Lanctot, Vinícius Flores Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien
544 Pérolat, David Silver, and Thore Graepel. A unified game-theoretic approach to multiagent rein-
545 forcement learning. In *NIPS*, pp. 4190–4203, 2017.

- 540 Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay,
541 Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, Daniel
542 Hennes, Dustin Morrill, Paul Muller, Timo Ewalds, Ryan Faulkner, János Kramár, Bart De
543 Vylder, Brennan Saeta, James Bradbury, David Ding, Sebastian Borgeaud, Matthew Lai, Julian
544 Schrittwieser, Thomas Anthony, Edward Hughes, Ivo Danihelka, and Jonah Ryan-Davis. Open-
545 Spiel: A framework for reinforcement learning in games. *CoRR*, abs/1908.09453, 2019. URL
546 <http://arxiv.org/abs/1908.09453>.
- 547 Torbjörn Larsson and Michael Patriksson. A class of gap functions for variational inequalities.
548 *Mathematical Programming*, 64(1):53–79, 1994.
- 549
550 Chung-Wei Lee, Christian Kroer, and Haipeng Luo. Last-iterate convergence in extensive-form
551 games. In *NeurIPS*, pp. 14293–14305, 2021.
- 552
553 Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. Sample complexity of asynchronous
554 q-learning: Sharper analysis and variance reduction. In *NeurIPS*, 2020.
- 555
556 Huale Li, Xuan Wang, Fengwei Jia, Yifan Li, and Qian Chen. A survey of nash equilibrium strategy
557 solving based on cfr. *Archives of Computational Methods in Engineering*, 28(4), 2021.
- 558
559 Mingyang Liu, Asuman E. Ozdaglar, Tiancheng Yu, and Kaiqing Zhang. The power of regulariza-
560 tion in solving extensive-form games. In *ICLR*. OpenReview.net, 2023.
- 561
562 Stephen McAleer, John B. Lanier, Roy Fox, and Pierre Baldi. Pipeline PSRO: A scalable approach
563 for finding approximate nash equilibria in large games. In *NeurIPS*, 2020.
- 564
565 Stephen Marcus McAleer, Gabriele Farina, Marc Lanctot, and Tuomas Sandholm. ESCHER: es-
566 chewing importance sampling in games by computing a history value function to estimate regret.
567 In *ICLR*. OpenReview.net, 2023.
- 568
569 Stephen Marcus McAleer, JB Lanier, Kevin A. Wang, Pierre Baldi, Tuomas Sandholm, and Roy
570 Fox. Toward optimal policy population growth in two-player zero-sum games. In *ICLR*. Open-
571 Review.net, 2024.
- 572
573 Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form
574 games. *Games and Economic Behavior*, 10(1):6–38, 1995. ISSN 0899-8256. doi: <https://doi.org/10.1006/game.1995.1023>. URL <https://www.sciencedirect.com/science/article/pii/S0899825685710238>.
- 575
576 Julien Pérolat, Rémi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, Mark Rowland, Pe-
577 dro A. Ortega, Neil Burch, Thomas W. Anthony, David Balduzzi, Bart De Vylder, Georgios
578 Piliouras, Marc Lanctot, and Karl Tuyls. From poincaré recurrence to convergence in imperfect
579 information games: Finding equilibrium via regularization. In *ICML*, volume 139 of *Proceedings
of Machine Learning Research*, pp. 8525–8535. PMLR, 2021.
- 580
581 Max Rudolph, Nathan Lichtle, Sobhan Mohammadpour, Alexandre Bayen, J. Zico Kolter, Amy
582 Zhang, Gabriele Farina, Eugene Vinitzky, and Samuel Sokota. Reevaluating policy gradient
583 methods for imperfect-information games, 2025. URL <https://arxiv.org/abs/2502.08938>.
- 584
585 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
586 optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- 587
588 Yoav Shoham, Rob Powers, and Trond Grenager. If multi-agent learning is the answer, what is the
589 question? *Artificial intelligence*, 171(7):365–377, 2007.
- 590
591 David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driess-
592 che, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneshelvam, Marc Lanctot, Sander
593 Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap,
Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game
of go with deep neural networks and tree search. *Nat.*, 529(7587):484–489, 2016.

594 David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur
595 Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap,
596 Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that mas-
597 ters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018. doi: 10.1126/
598 science.aar6404. URL [https://www.science.org/doi/abs/10.1126/science.](https://www.science.org/doi/abs/10.1126/science.aar6404)
599 [aar6404](https://www.science.org/doi/abs/10.1126/science.aar6404).

600 Samuel Sokota, Ryan D’Orazio, J. Zico Kolter, Nicolas Loizou, Marc Lanctot, Ioannis Mitliagkas,
601 Noam Brown, and Christian Kroer. A unified approach to reinforcement learning, quantal re-
602 sponse equilibria, and two-player zero-sum games. In *ICLR*. OpenReview.net, 2023.

603
604 Eric Steinberger, Adam Lerer, and Noam Brown. DREAM: deep regret minimization with advantage
605 baselines and model-free learning. *CoRR*, abs/2006.10410, 2020.

606 Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Juny-
607 oung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster
608 level in starcraft ii using multi-agent reinforcement learning. *nature*, 575(7782):350–354, 2019.

609
610 Martin Zinkevich, Michael Johanson, Michael H. Bowling, and Carmelo Piccione. Regret mini-
611 mization in games with incomplete information. In *NIPS*, pp. 1729–1736. Curran Associates,
612 Inc., 2007.

613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

A THEORETICAL ANALYSIS

In this appendix, we provide detailed proofs for the main theoretical results presented in Section 4.1 and Section 4.2. We begin by establishing the mathematical foundation and recall the problem setup.

A.1 MATHEMATICAL NOTATION

For a function f , we denote its domain by $\text{dom } f$ and for a set \mathcal{D} , we write $\text{int } \mathcal{D}$ for its interior.

Definition 3 (L -smooth function). *A differentiable function $f : \mathcal{Z} \rightarrow \mathbb{R}$ is L -smooth on $\mathcal{Z} \subseteq \mathbb{R}^n$ if its gradient is L -Lipschitz continuous:*

$$\|\nabla f(z) - \nabla f(z')\| \leq L\|z - z'\| \quad \forall z, z' \in \mathcal{Z}.$$

Definition 4 (μ -strong convex function). *A differentiable function $f : \mathcal{Z} \rightarrow \mathbb{R}$ is μ -strongly convex on \mathcal{Z} if there exists $\mu > 0$ such that*

$$f(z') \geq f(z) + \langle \nabla f(z), z' - z \rangle + \frac{\mu}{2}\|z' - z\|^2 \quad \forall z, z' \in \mathcal{Z}.$$

Definition 5 (Monotone operator). *An operator $G : \mathcal{Z} \rightarrow \mathbb{R}^n$ is monotone on \mathcal{Z} if*

$$\langle G(z) - G(z'), z - z' \rangle \geq 0 \quad \forall z, z' \in \mathcal{Z}.$$

It is μ -strongly monotone on \mathcal{Z} if there exists $\mu > 0$ such that

$$\langle G(z) - G(z'), z - z' \rangle \geq \mu\|z - z'\|^2 \quad \forall z, z' \in \mathcal{Z}.$$

Definition 6 (Bregman divergence). *Let $\psi : \mathcal{Z} \rightarrow \mathbb{R}$ be a differentiable, strongly convex function. For $z \in \text{dom } \psi$ and $z' \in \text{int } \text{dom } \psi$, the Bregman divergence is defined as*

$$B_\psi(z; z') = \psi(z) - \psi(z') - \langle \nabla \psi(z'), z - z' \rangle,$$

which satisfies $B_\psi(z; z') \geq 0$.

A.2 PROBLEM SETUP

Two-Player Zero-Sum Games. We consider two-player zero-sum extensive-form games with perfect recall. Their equivalent normal-form representation is given by a payoff matrix \mathbf{A} , where rows and columns correspond to the pure strategies of players 1 and 2. Let \mathcal{X} and \mathcal{Y} denote the mixed strategy spaces (probability simplices) of the two players, and define the strategy profile space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. A strategy profile is $z = (x, y) \in \mathcal{Z}$, and the payoff to player 1 is

$$f(x, y) = x^\top \mathbf{A} y.$$

Nash Equilibrium. A strategy profile $z^* = (x^*, y^*) \in \mathcal{Z}$ is a Nash equilibrium if neither player can improve unilaterally:

$$x^* \in \arg \max_{x \in \mathcal{X}} f(x, y^*) \quad \text{and} \quad y^* \in \arg \max_{y \in \mathcal{Y}} -f(x^*, y).$$

Variational Inequality Formulation. The Nash equilibrium problem can be equivalently expressed as a variational inequality (VI). Define the operator $F : \mathcal{Z} \rightarrow \mathbb{R}^n$ by

$$F(z) = \begin{bmatrix} -\nabla_x f(x, y) \\ \nabla_y f(x, y) \end{bmatrix} = \begin{bmatrix} -\mathbf{A}y \\ \mathbf{A}^\top x \end{bmatrix}. \quad (13)$$

Then z^* is a Nash equilibrium if and only if it solves $\text{VI}(\mathcal{Z}, F)$:

$$\langle F(z^*), z - z^* \rangle \geq 0 \quad \forall z \in \mathcal{Z}.$$

Regularized Variational Inequality. With a strongly convex regularizer $\psi(z) = \psi_1(x) + \psi_2(y)$ with associated Bregman divergence

$$B_\psi(z; \rho) = B_{\psi_1}(x; \rho_1) + B_{\psi_2}(y; \rho_2),$$

where $\rho = (\rho_1, \rho_2) \in \mathcal{Z}$ is a reference strategy profile. The regularized operator is defined as

$$G_\rho(z) = F(z) + \alpha \nabla_z B_\psi(z; \rho) \quad \alpha > 0, \quad (14)$$

or equivalently

$$G_\rho(z) = F(z) + \alpha (\nabla \psi(z) - \nabla \psi(\rho)) \quad \alpha > 0. \quad (15)$$

The corresponding regularized VI seeks $\hat{z} \in \mathcal{Z}$ such that

$$\langle G_\rho(\hat{z}), z - \hat{z} \rangle \geq 0 \quad \forall z \in \mathcal{Z}.$$

Regularized VI Operator. The central object of our analysis is the mapping

$$\mathcal{M} : \text{int dom } \psi \cap \mathcal{Z} \rightarrow \text{int dom } \psi \cap \mathcal{Z},$$

which takes a reference profile ρ to the unique solution $\mathcal{M}(\rho)$ of $\text{VI}(\mathcal{Z}, G_\rho)$. Formally, $\mathcal{M}(\rho)$ is the unique point satisfying

$$\langle G_\rho(\mathcal{M}(\rho)), z - \mathcal{M}(\rho) \rangle \geq 0 \quad \forall z \in \mathcal{Z}.$$

A.3 PRELIMINARY RESULTS

Lemma 4 (Three-Point Property). *For any strongly convex function ψ and points $a \in \text{dom } \psi$ and $\{b, c\} \subset \text{int dom } \psi$, the Bregman divergence satisfies:*

$$B_\psi(a; b) = B_\psi(a; c) + B_\psi(c; b) + \langle \nabla \psi(c) - \nabla \psi(b), a - c \rangle \quad (16)$$

Proof. By definition of Bregman divergence:

$$B_\psi(a; b) = \psi(a) - \psi(b) - \langle \nabla \psi(b), a - b \rangle \quad (17)$$

$$B_\psi(a; c) = \psi(a) - \psi(c) - \langle \nabla \psi(c), a - c \rangle \quad (18)$$

$$B_\psi(c; b) = \psi(c) - \psi(b) - \langle \nabla \psi(b), c - b \rangle \quad (19)$$

Direct algebraic manipulation yields the result. \square

Lemma 5 (Monotonicity of Operator F). *The operator $F : \mathcal{Z} \rightarrow \mathbb{R}^n$ defined in the problem setup is monotone on \mathcal{Z} .*

Proof. Let $z_1 = (x_1, y_1)$ and $z_2 = (x_2, y_2)$ be two strategy profiles in \mathcal{Z} . We have:

$$\langle F(z_1) - F(z_2), z_1 - z_2 \rangle \quad (20)$$

$$= \langle -\mathbf{A}y_1 + \mathbf{A}y_2, x_1 - x_2 \rangle + \langle \mathbf{A}^\top x_1 - \mathbf{A}^\top x_2, y_1 - y_2 \rangle \quad (21)$$

$$= -(x_1 - x_2)^\top \mathbf{A}(y_1 - y_2) + (x_1 - x_2)^\top \mathbf{A}(y_1 - y_2) \quad (22)$$

$$= 0 \geq 0 \quad (23)$$

where we used the definition $F(z) = [-\mathbf{A}y; \mathbf{A}^\top x]$ from the problem setup. \square

Lemma 6 (Strong Monotonicity of G_ρ). *Let ψ be μ -strongly convex on $\text{int dom } \psi$ with $\mu > 0$, then G_ρ is $\alpha\mu$ -strongly monotone on \mathcal{Z} , i.e.,*

$$\langle G_\rho(u) - G_\rho(v), u - v \rangle \geq \alpha\mu \|u - v\|^2 \quad \text{for all } u, v \in \mathcal{Z}.$$

Proof. For any $u, v \in \mathcal{Z}$,

$$\langle G_\rho(u) - G_\rho(v), u - v \rangle = \langle F(u) - F(v), u - v \rangle + \alpha \langle \nabla \psi(u) - \nabla \psi(v), u - v \rangle.$$

By Lemma 5, F is monotone, so the first term is nonnegative. By μ -strong convexity of ψ , the second term is at least $\alpha\mu \|u - v\|^2$. This proves $\alpha\mu$ -strong monotonicity of G_ρ . \square

A.4 MAIN THEORETICAL RESULTS

A.4.1 PROOF OF LEMMA 1

Lemma 1 (Distance Non-increase Property). *Given $\alpha > 0$, let $\rho \in \text{int dom } \psi \cap \mathcal{Z}$ and z^* be any Nash equilibrium (i.e., a solution of $\text{VI}(\mathcal{Z}, F)$). Then:*

$$B_\psi(z^*; \rho) \geq B_\psi(z^*; \mathcal{M}(\rho)) + B_\psi(\mathcal{M}(\rho); \rho). \quad (8)$$

Proof. For notational convenience, define $z_\rho := \mathcal{M}(\rho)$. The Nash equilibrium z^* satisfies the VI optimality condition for F :

$$\langle F(z^*), z - z^* \rangle \geq 0, \quad \forall z \in \mathcal{Z}. \quad (24)$$

Similarly, z_ρ satisfies the VI optimality condition for the regularized operator G_ρ :

$$\langle G_\rho(z_\rho), z - z_\rho \rangle \geq 0, \quad \forall z \in \mathcal{Z}. \quad (25)$$

756 Substitute $z = z_\rho$ in equation 24 and $z = z^*$ in equation 25, then add the two inequalities:

$$757 \quad \langle F(z^*), z_\rho - z^* \rangle + \langle F(z_\rho) + \alpha \nabla_{z_\rho} B_\psi(z_\rho; \rho), z^* - z_\rho \rangle \geq 0. \quad (26)$$

759 Rearranging equation 26 gives

$$760 \quad \alpha \langle \nabla_{z_\rho} B_\psi(z_\rho; \rho), z^* - z_\rho \rangle \geq \langle F(z^*) - F(z_\rho), z^* - z_\rho \rangle. \quad (27)$$

762 By Lemma 5, F is monotone on \mathcal{Z} , so the right-hand side of equation 27 is nonnegative. Hence, since $\alpha > 0$, we obtain

$$763 \quad \langle \nabla_{z_\rho} B_\psi(z_\rho; \rho), z^* - z_\rho \rangle \geq 0. \quad (28)$$

766 And since $\nabla_{z_\rho} B_\psi(z_\rho; \rho) = \nabla \psi(z_\rho) - \nabla \psi(\rho)$, equation 28 becomes

$$767 \quad \langle \nabla \psi(z_\rho) - \nabla \psi(\rho), z^* - z_\rho \rangle \geq 0. \quad (29)$$

769 Finally, applying the three-point property of Bregman divergences (Lemma 4) with $a = z^*$, $c = z_\rho$, and $b = \rho$, we have

$$772 \quad B_\psi(z^*; \rho) = B_\psi(z^*; z_\rho) + B_\psi(z_\rho; \rho) + \langle \nabla \psi(z_\rho) - \nabla \psi(\rho), z^* - z_\rho \rangle. \quad (30)$$

773 Combining equation 30 with equation 29 immediately gives

$$774 \quad B_\psi(z^*; \rho) \geq B_\psi(z^*; z_\rho) + B_\psi(z_\rho; \rho).$$

776 □

778 A.4.2 PROOF OF LEMMA 2

779 **Lemma 2** (Fixed Point Characterization). *A strategy profile $z^* \in \mathcal{Z}$ is a Nash equilibrium if and only if it is a fixed point of the regularized VI operator: $z^* = \mathcal{M}(z^*)$.*

782 *Proof.* (\Rightarrow) Suppose z^* is a Nash equilibrium. By definition, z^* solves $\text{VI}(\mathcal{Z}, F)$:

$$784 \quad \langle F(z^*), z - z^* \rangle \geq 0, \quad \forall z \in \mathcal{Z}. \quad (31)$$

786 Consider the regularized operator $G_{z^*}(z)$. By the VI optimality condition, z^* also solves $\text{VI}(\mathcal{Z}, G_{z^*})$, since

$$788 \quad \langle G_{z^*}(z^*), z - z^* \rangle = \langle F(z^*), z - z^* \rangle + \alpha \langle \nabla_z B_\psi(z^*; z^*), z - z^* \rangle = \langle F(z^*), z - z^* \rangle \geq 0,$$

789 using $\nabla_z B_\psi(z^*; z^*) = 0$. By definition of the operator \mathcal{M} , the unique solution of $\text{VI}(\mathcal{Z}, G_{z^*})$ is $\mathcal{M}(z^*)$. Therefore implies $z^* = \mathcal{M}(z^*)$.

792 (\Leftarrow) Conversely, suppose $z^* = \mathcal{M}(z^*)$. By definition of operator \mathcal{M} , z^* solves $\text{VI}(\mathcal{Z}, G_{z^*})$:

$$793 \quad \langle G_{z^*}(z^*), z - z^* \rangle \geq 0, \quad \forall z \in \mathcal{Z}. \quad (32)$$

795 Expanding $G_{z^*}(z^*) = F(z^*) + \alpha \nabla_z B_\psi(z^*; z^*) = F(z^*)$, we see that z^* satisfies

$$796 \quad \langle F(z^*), z - z^* \rangle \geq 0, \quad \forall z \in \mathcal{Z},$$

798 i.e., z^* is a Nash equilibrium. □

800 A.4.3 PROOF OF LEMMA 3

801 **Lemma 3** (Continuity of \mathcal{M}). *Assume ψ is μ -strongly convex and continuously differentiable on $\text{int dom } \psi$ for some $\mu > 0$. Then the operator \mathcal{M} is continuous.*

804 *Proof.* Fix $\rho \in \text{int dom } \psi \cap \mathcal{Z}$, let $\{\rho_n\} \subset \text{int dom } \psi \cap \mathcal{Z}$ be any sequence with $\rho_n \rightarrow \rho$, and for notational convenience we write

$$807 \quad z_n := \mathcal{M}(\rho_n), \quad z_\rho := \mathcal{M}(\rho).$$

808 By the VI optimality conditions for z_n and z_ρ we have

$$809 \quad \langle G_{\rho_n}(z_n), z_\rho - z_n \rangle \geq 0 \quad \text{and} \quad \langle G_\rho(z_\rho), z_n - z_\rho \rangle \geq 0.$$

810 Adding these two inequalities yields

$$811 \quad \langle G_{\rho_n}(z_n) - G_{\rho}(z_{\rho}), z_{\rho} - z_n \rangle \geq 0,$$

813 rearranges to get

$$814 \quad \langle G_{\rho_n}(z_n) - G_{\rho_n}(z_{\rho}), z_n - z_{\rho} \rangle \leq \langle G_{\rho_n}(z_{\rho}) - G_{\rho}(z_{\rho}), z_{\rho} - z_n \rangle. \quad (33)$$

816 By $\alpha\mu$ -strong monotonicity of G_{ρ_n} (Lemma 6), the left-hand side of equation 33 is bounded below
817 by $\alpha\mu\|z_n - z_{\rho}\|^2$. Applying the Cauchy–Schwarz inequality to the right-hand side of equation 33
818 yields

$$819 \quad \alpha\mu\|z_n - z_{\rho}\|^2 \leq \|G_{\rho_n}(z_{\rho}) - G_{\rho}(z_{\rho})\| \cdot \|z_n - z_{\rho}\|.$$

820 If $z_n = z_{\rho}$ the desired convergence holds trivially. Otherwise we divide both sides by $\|z_n - z_{\rho}\| > 0$
821 to obtain

$$822 \quad \|z_n - z_{\rho}\| \leq \frac{1}{\alpha\mu} \|G_{\rho_n}(z_{\rho}) - G_{\rho}(z_{\rho})\|. \quad (34)$$

824 Since

$$825 \quad G_{\rho_n}(z_{\rho}) - G_{\rho}(z_{\rho}) = \alpha(\nabla\psi(\rho) - \nabla\psi(\rho_n)),$$

826 we have that equation 34 simplifies to

$$827 \quad \|z_n - z_{\rho}\| \leq \frac{1}{\mu} \|\nabla\psi(\rho_n) - \nabla\psi(\rho)\|. \quad (35)$$

830 Since $\nabla\psi$ is continuous on $\text{int dom } \psi$ and $\rho_n \rightarrow \rho$, the right-hand side of equation 35 vanishes as
831 $n \rightarrow \infty$. Therefore $\mathcal{M}(\rho_n) \rightarrow \mathcal{M}(\rho)$. As this holds for every sequence $\rho_n \rightarrow \rho$, the mapping \mathcal{M} is
832 continuous on $\text{int dom } \psi \cap \mathcal{Z}$. \square

834 A.4.4 PROOF CONVERGENCE OF ITERATIVE REFINEMENT PROCEDURE

835 **Theorem 3** (Convergence of Iterative \mathcal{M}). *Let z^* be any Nash equilibrium. Algorithm 2 generates*
836 *a sequence $\{z_t\}$ such that*

$$837 \quad B_{\psi}(z^*; z_t) > B_{\psi}(z^*; z_{t+1}), \quad (10)$$

838 *whenever convergence has not yet occurred. Moreover, the sequence is guaranteed to converge to a*
839 *Nash equilibrium, i.e., $\lim_{t \rightarrow \infty} z_t$ is a Nash equilibrium.*

841 *Proof.* Fix any Nash equilibrium z^* . For the first part of the theorem, we show the strictly inequality
842 holds for all iterations before convergence occur. Apply Lemma 1 with $\rho = z_t$ and recall $z_{t+1} =$
843 $\mathcal{M}(z_t)$ to obtain

$$844 \quad B_{\psi}(z^*; z_t) \geq B_{\psi}(z^*; z_{t+1}) + B_{\psi}(z_{t+1}; z_t).$$

845 If the iteration has not converged then $z_{t+1} \neq z_t$, therefore $B_{\psi}(z_{t+1}; z_t) > 0$. Hence the preceding
846 inequality is strict, yielding

$$847 \quad B_{\psi}(z^*; z_t) > B_{\psi}(z^*; z_{t+1}).$$

849 Next, we show the sequence converge to a Nash equilibrium. Iterating the inequality from Lemma 1,
850 we get, for any $k \geq 1$,

$$851 \quad \begin{aligned} 852 \quad B_{\psi}(z^*; z_0) &\geq B_{\psi}(z^*; z_1) + B_{\psi}(z_1; z_0) \\ 853 \quad &\geq B_{\psi}(z^*; z_2) + B_{\psi}(z_2; z_1) + B_{\psi}(z_1; z_0) \\ 854 \quad &\geq B_{\psi}(z^*; z_k) + \sum_{\ell=0}^{k-1} B_{\psi}(z_{\ell+1}; z_{\ell}). \end{aligned} \quad (36)$$

857 As all terms in the RHS are non-negative, this implies that the partial sums of the series of non-
858 negative terms $\sum_{\ell=0}^{k-1} B_{\psi}(z_{\ell+1}; z_{\ell})$ are uniformly bounded, and so the series converges. In partic-
859 ular $B_{\psi}(z_{k+1}; z_k) \rightarrow 0$, which we will use below. Moreover, Equation equation 36 also implies
860 $B_{\psi}(z^*; z_k) \leq B_{\psi}(z^*; z_0)$, and the sequence $\{z_k\}_k$ belongs to the closed, bounded Bregman ball
861 of radius $r := B_{\psi}(z^*; z_0)$. As a result, there exists a convergent subsequence $\{z_{\phi(k)}\}_k$ converges
862 (where $\phi: \mathbb{N} \rightarrow \mathbb{N}$ is increasing): denote its limit by

$$863 \quad z_{\infty} := \lim_{k \rightarrow \infty} z_{\phi(k)}.$$

Then, we claim that z_∞ is a Nash equilibrium. Indeed, we have, from the above,

$$B_\psi(\mathcal{M}(z_{\phi(k)}); z_{\phi(k)}) = B_\psi(z_{\phi(k)+1}; z_{\phi(k)}) \xrightarrow[k \rightarrow \infty]{} 0$$

while, by continuity of the Bregman divergence and of the \mathcal{M} operator (Lemma 3),

$$B_\psi(\mathcal{M}(z_{\phi(k)}); z_{\phi(k)}) \xrightarrow[k \rightarrow \infty]{} B_\psi(\mathcal{M}(z_\infty); z_\infty)$$

and so, by uniqueness of the limit, $B_\psi(\mathcal{M}(z_\infty); z_\infty) = 0$. Lemma 2 then allows us to conclude that z_∞ is a Nash equilibrium.

We can then conclude that z_t converges to z_∞ : indeed, applying Lemma 1 to the Nash equilibrium z_∞ , we get as the beginning of the proof that, for all $t, s \geq 0$,

$$B_\psi(z_\infty; z_t) \geq B_\psi(z_\infty; z_{t+s})$$

(i.e., $B_\psi(z_\infty; z_t)$ is a non-increasing, non-negative sequence). Since $\lim_{k \rightarrow \infty} z_{\phi(k)} = z_\infty$, this easily allows us to conclude that the sequence $\{z_t\}_t$ converges to z_∞ as well. (In more detail: fixing $\varepsilon > 0$, there exists $k_\varepsilon \geq 0$ such that, for all $k \geq k_\varepsilon$, $B_\psi(z_\infty; z_{\phi(k)}) \leq \varepsilon$. Then, for all $t \geq \phi(k_\varepsilon)$, we have $B_\psi(z_\infty; z_t) \leq B_\psi(z_\infty; z_{\phi(k_\varepsilon)}) \leq \varepsilon$.) This concludes the proof. \square

B EXPERIMENT DETAILS

B.1 ALGORITHM SETTINGS

We evaluate NASHPG with three established model-free baselines: Neural Fictitious Self-Play (NFSP) (Heinrich & Silver, 2016), Policy Space Response Oracles (PSRO) (Lanctot et al., 2017), and Magnetic Mirror Descent (MMD) (Sokota et al., 2023). All methods employ PPO as the primary training framework for best-response learning and policy-gradient updates, with hyperparameters listed in Table 2. Detailed implementation notes for each algorithm are provided below.

NFSP. Maintains a uniform mixture of policies comprising all previously trained agents. Each iteration, a new agent is initialized and use PPO to train a approximate best-response agent against the mixture of prior policies. The newly trained agent is added to the population after each iteration.

PSRO. The setup follows NFSP but replaces uniform mixing with a meta-game constructed from pairwise interactions among historical agents. Each entry of the payoff matrix is estimated from rollouts with 128 parallel environments and 1,000 steps per environment. A Nash equilibrium of this meta-game is then approximated via fictitious play (Brown, 1951) for 10,000 iterations, and the resulting mixed strategy defines the new population distribution. Training then continues with a newly initialized agent.

MMD. A single policy is trained in self-play using PPO. The original paper states they use PPO but change forward KL to reversed KL; however, we hypothesize their intention is to add an additional KL regularization term. Based on their setup, the reference policy is set to a uniform policy. Following Sokota et al. (2023), we evaluate magnet coefficients in $\{0.05, 0.1, 0.2\}$ and find most of the time 0.05 yields better performance, hence use it for all experiments.

B.2 EVALUATION METRICS

In our work, we employ both *exploitability* and *Elo ratings* to evaluate the policies of all baselines and NASHPG.

Exploitability: Consider a two-player zero-sum game with player’s behavioural strategies π_1 and π_2 . Let $\mu_i(\pi_i, \pi_{-i})$ denote the expected payoff for player i , where $-i$ denotes the opponent of player i . The best-response strategy for player i is defined as

$$BR_i(\pi_{-i}) = \arg \max_{\pi'_i} \mu_i(\pi'_i, \pi_{-i})$$

Given a suboptimal strategy π_i , the incentive for player i to deviate is

$$\delta_i = \mu_i(BR_i(\pi_{-i}), \pi_{-i}) - \mu_i(\pi_i, \pi_{-i})$$

Hyperparameter	Value
Optimizer	AdamW
Step size(η)	3.0×10^{-4}
Num. epochs	4
Minibatch size	4
Discount (γ)	1
GAE parameter (λ)	0.95
Clipping parameter (ϵ)	0.2
Entropy coeff.	0.05
Number of envs.	32
Rollout length	128

Table 2: PPO hyperparameters

The exploitability is then defined as

$$\text{Exploitability} = \mathbb{E}_{i \sim \text{Uniform}} [\delta_i]$$

In our experiments, we approximate $BR_i(\pi_{-i})$ by training a newly initialized agent with PPO against the fixed strategy π_{-i} at each evaluation checkpoint. For convenience of implementation, we augment each environment so at the beginning of each game, we randomly assign player positions. The exploiter is trained for 2,000 steps. All other optimization and algorithmic hyperparameters follow the settings reported in Table 2.

Elo ratings: This system produces a scalar ranking based on empirical match outcomes. Given two players A and B with current Elo ratings Elo_A and Elo_B , the expected score of player A is

$$E_A = \frac{1}{1 + 10^{(\text{Elo}_B - \text{Elo}_A)/400}}$$

After a match, the realized outcome $R_A \in \{0, 0.5, 1\}$ is recorded, where 1 denotes a win for A, 0.5 a draw, and 0 a loss. Ratings are then updated as

$$\text{Elo}_A \leftarrow \text{Elo}_A + K \cdot (R_A - E_A) \quad \text{Elo}_B \leftarrow \text{Elo}_B + K \cdot ((1 - R_A) - (1 - E_A))$$

where K is the update factor. In our experiments, we set $K = 32$ and initialize all agents at an Elo rating of 1500. We evaluate policies from different algorithms by running a Swiss-style tournament consisting of 100 rounds. For each of the four algorithms, we save a checkpoint every 2,000 steps (50 checkpoints per run) and perform four independent runs, yielding approximately 800 tournament entries (4 algorithms \times 4 runs \times 50 checkpoints). At the beginning of each round, all players are sorted by their current Elo rating and paired with their nearest neighbor. Each pairing is evaluated over 100 repeated games. The cumulative reward difference across these 100 games determines the match result: if one player achieves higher total reward, they are assigned $R = 1$ and the opponent $R = 0$; if totals are equal, both receive $R = 0.5$. Elo ratings are then updated according to the above rule.

B.3 ENVIRONMENT DOMAIN

We evaluate NASHPG across seven two-player zero-sum games, Our benchmark covers **Kuhn Poker**, **Leduc Poker**, **Dark Hex (3 \times 3)**, and **Phantom Tic-Tac-Toe**, with approximately 54, 9,300, 2.9×10^{10} , and 2.7×10^{10} non-terminal histories, respectively (Sokota et al., 2023; Rudolph et al., 2025). We also evaluate on three real-world domains: **Liar’s Dice**, **Battleship**, and **Heads-Up No-Limit Texas Hold’em**. Liar’s Dice alone comprises roughly 10^{20} non-terminal histories, while Battleship and Heads-Up No-Limit Texas Hold’em are vastly larger. All environments use lightweight neural architectures comprising a feature extractor, a policy head, and a critic head, keeping capacity low to emphasise learning dynamics. For the two imperfect-information board games (Dark Hex and Phantom Tic-Tac-Toe), we apply the *abrupt* turn rule, i.e. any blocked move immediately ends the acting player’s turn, and we follow the configurations provided by OpenSpiel (Lanctot et al., 2019). Architectural details are summarised in Table 3. We denote by Δ the true gain or loss within a game or hand in poker-related environments. For instance, in Heads-Up No-Limit Texas Hold’em,

Environment	Kuhn Poker	Leduc Poker	Dark Hex (3×3)	Phantom Tic-Tac-Toe	Liar’s Dice	Battleship	Heads-Up Hold’em
Players	2	2	2	2	2	2	2
Obs. Dim.	7	49	9	9	54	111	285
Action spaces	2	3	9	9	61	100	16
Feature Extractor	Linear	Linear	Linear	Linear	Transformer	CNN	Transformer
Layers	2	2	2	2	1	2	2
Hidden Dim.	16	64	64	64	64	128	256
Attention Heads	–	–	–	–	1	–	8
History Length	–	–	–	–	16	–	64
Policy Head	Linear	Linear	Linear	Linear	Linear	CNN	Linear
Policy Layers	1	1	1	1	2	2	2
Critic Head	Linear	Linear	Linear	Linear	Linear	Linear	Linear
Critic Layers	1	1	1	1	2	1	2
Reward Design	Δ	$\Delta/20$	± 1	$\pm 1, 0$	± 1	± 1	Δ/Stack
Outer Iteration	50	50	50	50	50	50	50
Inner Iteration	1000	1000	2000	2000	2000	2000	2000

Table 3: Summary of environments and model architectures

if two players build a \$50 pot through betting and calling, and Player 1 bets an additional \$30 and Player 2 folds, then $\Delta = +25$ for Player 1 and -25 for Player 2. Comprehensive environment rules, architectural specifications, and reward formulations are listed below for reproducibility.

Kuhn Poker: A three-card, two-player imperfect-information poker game in which both players ante one chip and act sequentially by betting or passing. We encode the private card and betting sequence into a vector, process it with two linear layers, and apply linear policy and critic heads. Rewards are zero during play; at termination, the winner receives the pot and the loser forfeits their contribution. Observations include three dimensions for one-hot encoding of the private card ($J=0, Q=1, K=2$) and four dimensions for the complete betting history from the current player’s perspective.

Leduc Poker: An variant of Poker with six cards (two suits, three ranks), two betting rounds, and a public card revealed in the second round; raises are of fixed size. Rewards are zero during play; a fold immediately ends the hand in favour of the non-folding player, otherwise showdown is decided by pair strength or high card. In our implementation, we normalize the reward by factor of 20. Observations include 14 dimensions for one-hot encoding of the public and private cards ($J\spadesuit=0, Q\spadesuit=1, K\spadesuit=2, J\heartsuit=3, Q\heartsuit=4, K\heartsuit=5$), one for the revealed public-card indicator, two for the round indicator, two for player position, and 32 for the full betting history.

Dark Hex (3×3): A two-player imperfect-information variant of Hex on a 3×3 board: Player 0 aims to connect top to bottom, and Player 1 aims to connect left to right. Each player observes only their own successful placements and opponent stones discovered by attempting an already-occupied cell (a blocked move). We apply the *abrupt* rule. Rewards are zero during play; a successful placement that completes the acting player’s connection immediately ends the game with ± 1 (Hex admits no draws on a filled board). Observations are a length-9 vector with entries $\{-1, 0, 1\}$ denoting unknown, own stone, and discovered opponent stone.

Phantom Tic-Tac-Toe: A two-player imperfect-information version of Tic-Tac-Toe on a 3×3 grid: players alternate attempting to place a mark but observe only their own successful placements and cells where a previous attempt was blocked (revealing an opponent mark). We also use the *abrupt* rule. Rewards are zero during play; a successful placement forming three-in-a-row yields ± 1 , and a full board with no line results in a draw with 0 to both players. Observations are a length-9 vector with $\{-1, 0, 1\}$ encoding unknown, own mark, and discovered opponent mark.

Liar’s Dice: We implement a single-hand version with 5 dice, each having 6 faces. This is a sequential bidding game in which each player privately rolls dice and alternately bids on the total count of a face value across both players until someone issues a challenge; the dice are then revealed to resolve the claim. Rewards are zero until a challenge occurs; the previous bidder wins if the claim is correct, and the challenger wins otherwise, yielding ± 1 . Observations contain 16 historical

bidding states, each using three dimensions to record player ID, bid quantity, and bid face, plus six dimensions for the counts of each face in the player’s private dice.

Battleship: A two-stage, two-player game on a 10×10 grid: players first place ships, then alternate attacks until one fleet is destroyed. Rewards are $+1$ for victory and -1 for defeat; termination occurs when a player’s fleet is eliminated. Observations include 100 dimensions for the grid state, 10 for remaining ship statuses, and one for the current stage.

Heads-Up No-Limit Texas Hold’em: A two-player no-limit poker variant with blinds and four betting streets (preflop, flop, turn, river) and up to five community cards. Actions include folding, checking/calling, a 16 different fractional pot-sized raises, and all-in moves. At the beginning, the small blind/big blind is set to $1/2$ and stacks are initialized at 200 chips (100 big blinds) per player. Rewards are zero during play and correspond to normalized stack changes at the end of each hand; termination within a hand occurs upon a fold or at showdown (the overall match ends if a player goes bankrupt). Observations include two dimensions for the private and five for community cards, two for each player’s stacks, two for each player’s bets, one for pot size, one for the current street ID, and a sequence of historical actions. We use a history window of 64, each entry containing four dimensions (player ID, street ID, action taken, and betting amount).

C ADDITIONAL EXPERIMENTS

C.1 ANNEAL REGULARIZATION STRENGTH AS NASH EQUILIBRIA SOLVER

Sokota et al. (2023) introduced *Magnetic Mirror Descent* (MMD) and suggested two potential strategies for applying MMD to compute Nash equilibria: (i) *moving magnet*, which interpolates between the magnet policy and the current policy, and (ii) *annealing the regularization strength*. While the moving magnet approach is conceptually appealing, it requires geometric blending of policies at every information set, which is computationally infeasible in large games. This leaves annealing as the more practical alternative.

Here, we examine the feasibility of annealing regularization strength in the presence of stochastic sampling noise, i.e., when gradients are estimated from sampled trajectories rather than exact feedback. We conduct experiments on *Kuhn Poker*, a small imperfect-information game that exact computation of exploitability is feasible, making it a suitable testbed.

We use the same PPO hyperparameters as all other experiments and investigate the effect of linearly decaying the magnet coefficient α . Specifically, we test initial values $\alpha \in \{0.2, 0.4\}$, annealed linearly to 0.001. We additionally consider jointly decaying both α and the learning rate η , with η decayed linearly to zero. Each setting is repeated over four runs, and we report the mean and standard deviation of exploitability. Results are shown in Figure 4.

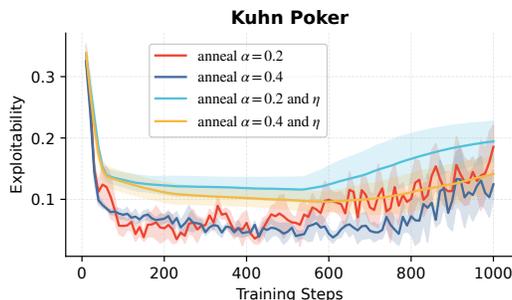


Figure 4: Exploitability under annealing different initial values for α (optionally annealing η). In all cases, exploitability diverges as α decreases.

The results reveal a consistent pattern: exploitability decreases steadily during the early stages of training, but as α becomes small, performance deteriorates and eventually diverges. Even when annealing η jointly with α , training proceeds more slowly and still exhibits divergence at later stages. These findings indicate that annealing regularization strength is inherently unstable in the stochastic

Component	Specification
CPU	AMD Ryzen Threadripper 3970X, 32 cores / 64 threads, 3.7 GHz (boost enabled), 128 MB L3 cache
RAM	256 GB DDR4
GPU	4 × NVIDIA GeForce RTX 4090 (24 GB GDDR6X each, PCIe)
Storage	2 × 2 TB NVMe SSDs
OS	Ubuntu 20.04 LTS

Table 4: Hardware specifications for all experiments.

setting. While careful hyperparameter tuning may mitigate these issues, the dependence of the dynamics on both α and η (recall the MMD convergence constraint $\alpha \geq \mu\eta L^2$) makes annealing regularization an unattractive choice as a general-purpose solver for Nash equilibria.

C.2 RUNTIME AND MEMORY ANALYSIS

To complement our performance analysis, we measured both real (wall-clock) training time and peak GPU memory allocation throughout the experiments. The hardware configuration used in our experiments is listed in Table 4, and all experiments were run with a single algorithm assigned to a single GPU at a time to ensure consistency. For brevity, we report results for four environments: *Kuhn Poker*, *Leduc Poker*, *Battleship*, and *Heads-Up No-Limit Texas Hold'em*. Detailed settings are summarized in Table 3. However, our results should be interpreted with caution, as population-based methods are sensitive to implementation details and may vary across studies. We aimed to provide a fair comparison but note that our implementation may not be the most efficient. For example, prior work such as Pipeline PSRO (McAleer et al., 2020) makes PSRO parallelizable, whereas we implemented the vanilla version.

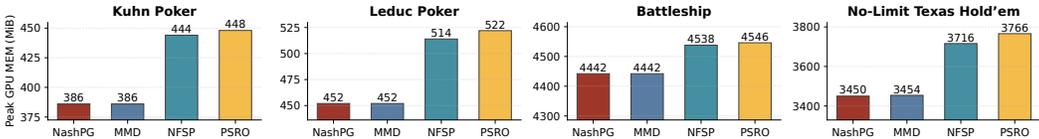


Figure 5: Maximum GPU memory usage during training across environments and algorithms.

Memory Footprint: Figure 5 shows the peak GPU memory usage during training under the settings described in Appendix B.1. Population-based methods (NFSP and PSRO) consume substantially more memory than single-agent approaches, as our implementation requires storing all historical agents. This limits scalability with larger models in complex environments. For example, in No-Limit Texas Hold'em, our current model size is approximately 4.6 MiB, yet PSRO and NASHPG differ by about 316 MiB—roughly $67\times$ the model size—after 50 meta-updates.

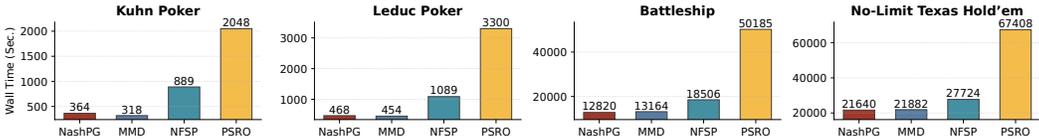


Figure 6: Total wall-clock time during training across all evaluated environments and algorithms.

Runtime: As shown in Figure 6, *PSRO* incurs significantly higher wall-clock costs, requiring approximately $3.1\text{--}7.3\times$ more time than *NASHPG* and *MMD* across all tested environments. *NFSP* is also slower, with $1.3\text{--}2.8\times$ higher runtime. One contributing factor to this overhead is that the meta-update stage in *PSRO* requires evaluating pairwise interactions with all previously trained agents, resulting in an $O(n^2)$ runtime complexity as the population grows. Additionally, as the population increases in *NFSP* and *PSRO*, the computation graph changes, causing *JAX* to recompile every outer iteration, which introduces further overhead.