

Teaching an Old Dynamics New Tricks: Regularization-free Last-iterate Convergence in Zero-sum Games via BNN Dynamics

Tuo Zhang

Department of Computer Science
University of Birmingham
Birmingham, United Kingdom
t.zhang.8@bham.ac.uk

Leonardo Stella

Department of Computer Science
University of Birmingham
Birmingham, United Kingdom
l.stella@bham.ac.uk

ABSTRACT

Zero-sum games are a fundamental setting for adversarial training and decision-making in multi-agent learning (MAL). Existing methods often ensure convergence to (approximate) Nash equilibria by introducing a form of regularization. Yet, regularization requires additional hyperparameters, which must be carefully tuned—a challenging task when the payoff structure is known, and considerably harder when the structure is unknown or subject to change. Motivated by this problem, we repurpose a classical model in evolutionary game theory, i.e., the Brown-von Neumann-Nash (BNN) dynamics, by leveraging the intrinsic convergence of this dynamics in zero-sum games without regularization, and provide last-iterate convergence guarantees in noisy normal-form games (NFGs). Importantly, to make this approach more applicable, we develop a novel framework with theoretical guarantees that integrates the BNN dynamics in extensive-form games (EFGs) through counterfactual weighting. Furthermore, we implement an algorithm that instantiates our framework with neural function approximation, enabling scalable learning in both NFGs and EFGs. Empirical results show that our method quickly adapts to nonstationarities, outperforming the state-of-the-art regularization-based approach.

KEYWORDS

Multi-agent learning; evolutionary dynamics; game theory; last-iterate convergence

ACM Reference Format:

Tuo Zhang and Leonardo Stella. 2026. Teaching an Old Dynamics New Tricks: Regularization-free Last-iterate Convergence in Zero-sum Games via BNN Dynamics. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/ARUH6291>

1 INTRODUCTION

With the pervasive advent of generative models and agentic AI, the research community has turned towards multi-agent learning (MAL) to lay the foundations for these multi-agent systems [52], despite the growing criticism on their intrinsic nature [8]. Since Shoham, Powers and Grenager initiated a vibrant debate across AI, economics, and engineering with their famous inquiry “If multi-agent learning is the answer then what is the question?” [49], MAL

has largely trailed behind game theory since its inception, with early breakthroughs such as Brown and Robinson’s fictitious play [9, 47], designed to computationally validate von Neumann’s seminal results on mixed Nash equilibria in zero-sum games [54]. Yet, it is precisely in the high-stakes setting of learning zero-sum games with noisy feedback where obtaining theoretical guarantees of convergence to the Nash equilibria becomes crucial.

Existing approaches that provide such guarantees generally fall into two categories: time-average convergence, which ensures that the running average of strategies converges, and last-iterate convergence, a stronger form of convergence where the strategies themselves converge. Between these two notions of convergence, the former is often unsatisfactory in practice, and computationally costly in large-scale games. Moreover, in many realistic settings, an endogenous environmental nonstationarity is added to the intrinsic nonstationarity resulting from the agents’ interactions [21, 58]. In such cases, time averaging becomes ineffective as outdated historical information continues to influence the learning process. On the contrary, approaches that achieve last-iterate convergence provide a more stable and responsive notion of learning in dynamic environments. Existing approaches of this kind mainly rely on regularization or payoff perturbation, which stabilizes the learning dynamics through a policy-based regularization term to the objective. However, this additional term creates another problem, namely, the tuning of additional hyperparameters, which is challenging when the payoff structure is known, and becomes considerably harder when the environment is unknown or changes over time.

To further explain the issue caused by regularization, consider the following motivating example. The nonstationary Rock-Paper-Scissors (RPS) game is like the traditional RPS game but pairs of payoffs change continuously over time (within a time window of 2500 or 1250 iterations), ensuring that the zero-sum property of the game holds true at all time. Figure 1(a) shows the evolution of NASHCONV—a metric that captures the distance of the learning dynamics of an algorithm from the Nash equilibrium. In particular, we compare the Brown–von Neumann–Nash (BNN) dynamics, which our proposed approach is based on, with regularized replicator dynamics (RD), which forms the backbone of the state-of-the-art regularized Nash dynamics (R-NaD) algorithm [44]. The BNN dynamics is a well-studied dynamical system in evolutionary game theory (EGT), with its key properties established in the classical literature (see, e.g., [25, 26]), but this is the first time it is introduced into MAL. Although both sets of dynamics retain the last-iterate convergence guarantee, this experiment shows the superiority of the BNN dynamics, in terms of proximity and stability around the Nash equilibrium against large oscillations of the regularized RD.



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). <https://doi.org/10.65109/ARUH6291>

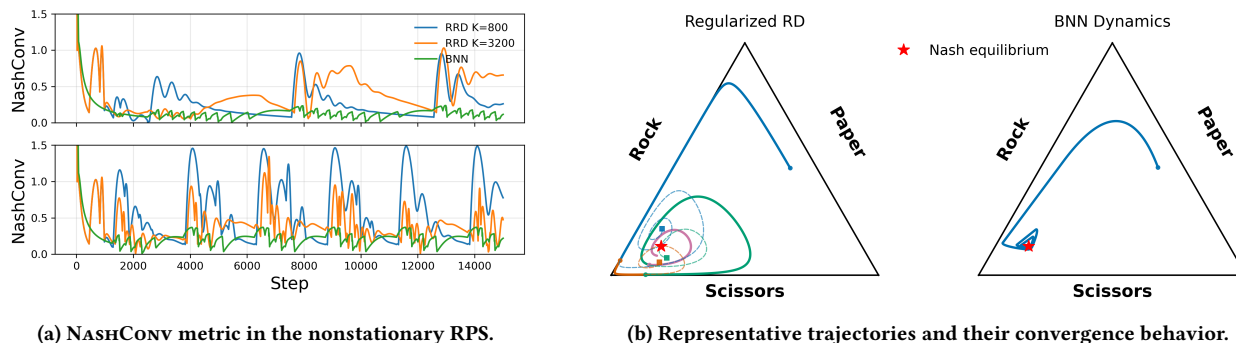


Figure 1: Comparison between the regularized RD and the BNN dynamics across different settings. NASHCONV metric in the nonstationary RPS environment with continuous changes in the payoffs (left). Representative trajectories under biased stationary payoffs in the simplex of the biased rock-paper-scissor (RPS) game (right).

This difference can be explained by the underlying mechanisms of the two sets of dynamics, as depicted on the simplex of the biased RPS game in Figure 1 (b). The structure of BNN dynamics inherently ensures convergence in negative semi-definite zero-sum games without requiring any additional regularization. In contrast, the regularized RD introduces a reference policy as a regularization term, effectively running the learning process on a perturbed zero-sum game whose payoffs are modified by the strength of the regularization. This modification shifts the system’s attractor to a new point, which is determined jointly by the original Nash equilibrium, the reference policy, and the regularization coefficient. By periodically updating the reference policy or decreasing the regularization strength, the new attractor gradually approaches the true Nash equilibrium. Consequently, the switching frequency and the regularization strength become two crucial hyperparameters. When the payoffs themselves change, this mechanism can become unsafe, as the coupling between the regularization dynamics and the evolving environment may drive the system towards unpredictable or undesired behaviors. To better visualize the described behavior, trajectories of the regularized RD are color-coded to indicate different reference stages of the process and different attractors. The BNN dynamics naturally converges to the Nash equilibrium without the need for regularization.

Motivated by this critical limitation in the state-of-the-art regularized RD methods, we propose the following contribution.

- (1) We establish a theoretical MAL framework grounded in BNN dynamics for two-player zero-sum games with noisy feedback, both in normal-form and in extensive-form settings.
- (2) Building on this framework, we develop an algorithm based on an actor-critic (AC) architecture that implements the population-level BNN dynamics.
- (3) We validate our algorithm on a comprehensive benchmark of zero-sum games, including both normal-form, extensive-form, and nonstationary settings. Our experiments show that its empirical behavior aligns with our theoretical predictions and that the algorithm outperforms state-of-the-art baselines based on regularized RD methods, showing superiority of our approach in terms of stability, convergence, and safety.

1.1 Related Work

Early connections between EGT and MAL were established by Börgers and Sarin [7], who provided the first formal link between the RD in EGT and the cross learning algorithm [15], later inspiring extensions to more sophisticated methods such as Q-learning and its variants [1, 31, 53], as well as regret-minimization approaches [29]. Subsequent work modified learning rules so that they inherit stability properties from evolutionary processes [30, 44], and extended the framework to a broader range of games and dynamics [22, 24, 41, 55]. Recently, EGT has also been used to provide theoretical grounding for algorithms with function approximation and deep learning [6, 23, 33, 43]. In our recent work, we have explored innovative dynamics, as opposed to imitation-based rules such as RD [59].

Parallel to the evolutionary approach, the line of work based on regularization originates from the idea of payoff perturbation [19]. By introducing a strongly convex penalty to the players’ payoff or utility functions, these methods ensure convergence towards approximate Nash equilibria [12, 13, 42]. For example, Sokota *et al.* [50] demonstrated that their perturbed mirror descent algorithm converges to the quantal response equilibrium [37, 38]. In continuous-time settings, numerous studies have also reported similar results [5, 14, 28, 35]. However, ensuring convergence requires careful adjustment of the perturbation strength [10, 36], with related studies exploring iterative Tikhonov regularization schemes [32, 51, 57]. Perolat *et al.* [44] introduced the idea of payoff perturbations that depend on a reference (or regularized) strategy, periodically updated to promote convergence. Later studies refined the switching mechanism to achieve faster convergence [2–4].

Another prominent class of methods to achieve last-iterate convergence relies on the idea of *optimistic learning* [45, 46]. Examples include the Optimistic Follow the Regularized Leader [48], Optimistic Mirror Descent algorithms [27, 60], and more recently Smooth Optimistic Gradient-Based Regret Matching+ [39]. Such approaches have been shown to attain last-iterate convergence across a wide range of games [16–18, 40], with established convergence rates [11, 20]. However, these results typically assume accurate and noiseless gradient feedback, and their performance deteriorates in the presence of stochastic or biased observations [34, 56]. In some cases, negative empirical results have also been reported [2].

2 PRELIMINARIES

This section provides the background and notation used throughout the paper. The remainder of this work is based on the concepts and formulations introduced here.

2.1 Normal-form Games

Normal-form games provide the static, population-level foundation for analysing learning dynamics in multi-agent systems. We consider two-player zero-sum games with finite action sets A_i and utility functions $u_i : A_1 \times A_2 \rightarrow [-u_{\max}, u_{\max}]$ for each player $i \in \{1, 2\}$, satisfying $u_1(a_1, a_2) = -u_2(a_1, a_2)$ for all $(a_1, a_2) \in A_1 \times A_2$. A mixed strategy of player i is a probability distribution $\pi_i \in \Delta(A_i)$ over its actions, where $\pi_i(a_j)$ denotes the probability of selecting action $a_j \in A_i$ under π_i .

The expected payoff under a strategy profile $\pi = (\pi_1, \pi_2)$ is

$$u_i(\pi_1, \pi_2) = \sum_{a_j \in A_1} \sum_{a_k \in A_2} \pi_1(a_j) \pi_2(a_k) u_i(a_j, a_k).$$

A strategy profile $\pi^* = (\pi_1^*, \pi_2^*)$ is a **Nash equilibrium** if, for each player i ,

$$u_i(\pi_i^*, \pi_{-i}^*) \geq u_i(\pi_i, \pi_{-i}^*), \quad \forall \pi_i \in \Delta(A_i).$$

For each pure action $a_j \in A_i$, we define its expected payoff against the opponent's strategy as

$$u_i(a_j | \pi_{-i}) = \sum_{a_k \in A_{-i}} \pi_{-i}(a_k) u_i(a_j, a_k).$$

For compactness, we write $u_i(a_j)$ to denote $u_i(a_j | \pi_{-i})$, and use $\bar{u}_i = u_i(\pi_1, \pi_2)$ to denote the expected payoff under the mixed strategy profile π .

2.2 Extensive-form Games

We now extend the normal-form notation introduced above to sequential games with imperfect information. An extensive-form game for two players is represented as a tuple

$$G = (H, Z, (A(h))_{h \in H \setminus Z}, (\tau(h))_{h \in H \setminus Z}, (u_i)_{i=1,2}, (I_i)_{i=1,2}),$$

where H is the finite set of all possible histories (action sequences), and $Z \subset H$ is the set of terminal histories. At each non-terminal history $h \in H \setminus Z$, player $\tau(h) \in \{1, 2\}$ chooses an action from the available set $A(h)$. Histories that player i cannot distinguish are grouped into information sets $I \in \mathcal{I}_i$, where \mathcal{I}_i denotes the collection of all such sets for player i .

A *behaviour strategy* of player i is a mapping $\sigma_i : \mathcal{I}_i \rightarrow \Delta(A(I))$, assigning a probability distribution over available actions at each information set I . The joint strategy profile is written $\sigma = (\sigma_1, \sigma_2)$.

Reach probabilities. Given σ , the probability of reaching a history $h = (a_1, \dots, a_t)$ is

$$\rho_\sigma(h) = \prod_{k=1}^t \sigma_{\tau(h_{k-1})}(a_k | I(h_{k-1})), \quad (1)$$

which can be factorised as $\rho_\sigma(h) = \rho_\sigma^i(h) \rho_\sigma^{-i}(h)$, representing the contributions of player i and its opponents, respectively. The *reach probability* of an information set $I \in \mathcal{I}_i$ is then

$$\rho_\sigma(I) = \sum_{h \in I} \rho_\sigma(h), \quad \rho_\sigma^{-i}(I) = \sum_{h \in I} \rho_\sigma^{-i}(h), \quad (2)$$

and under perfect recall we have $\rho_\sigma(I) = \rho_\sigma^i(I) \rho_\sigma^{-i}(I)$ for all I .

Counterfactual values. For each $I \in \mathcal{I}_i$ and action $a \in A(I)$, the *counterfactual value* of player i is

$$v_i^\sigma(I, a) = \sum_{h \in I} \rho_\sigma^{-i}(h) \sum_{\substack{z \in Z \\ ha \prec z}} \rho_\sigma(z | ha) u_i(z), \quad (3)$$

where $u_i(z)$ is the terminal payoff and $ha \prec z$ indicates that terminal history z extends the partial history ha . This expression isolates the expected return from taking a at I , weighted by the opponent's reach probability $\rho_\sigma^{-i}(h)$, and thus corresponds to the *counterfactual Q-function* in extensive-form games. The expected value at I is then

$$v_i^\sigma(I) = \sum_a \sigma_i(a|I) v_i^\sigma(I, a). \quad (4)$$

A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a *Nash equilibrium* if no player can unilaterally improve its expected payoff:

$$V_i(\sigma_i', \sigma_{-i}^*) \leq V_i(\sigma^*), \quad \forall \sigma_i'. \quad (5)$$

2.3 BNN Dynamics

The BNN dynamics belongs to the family of dynamics called innovative dynamics, as new strategies can enter the population through experimentation, exploration and mutation, and describes the continuous-time evolution of strategies. For player i and action $a \in A_i$, the BNN dynamics takes the form:

$$\dot{\pi}_i(a) = [u_i(a) - \bar{u}_i]_+ - \pi_i(a) \sum_{a' \in A_i} [u_i(a') - \bar{u}_i]_+.$$

These dynamics preserve the strategy simplex and guarantee no regret in the limit.

A discrete-time stochastic approximation with step size $\eta_t > 0$ takes the form

$$\pi_i^{t+1}(a) = \pi_i^t(a) + \eta_t \left([u_i(a) - \bar{u}_i]_+ - \pi_i^t(a) \sum_{a' \in A_i} [u_i(a') - \bar{u}_i]_+ \right).$$

2.4 Noisy Feedback

In practice, players often rely on estimated utilities rather than exact payoffs, due to sampling or inherent stochasticity of the environment. To capture this, we model feedback as a noisy observation. For each action $a \in A_i$, the observed payoff is

$$\tilde{u}_i(a) = u_i(a) + \xi_i(a),$$

and the corresponding average payoff under profile π is

$$\tilde{u}_i = \sum_{a_i \in A_i} \pi_i(a_i) \tilde{u}_i(a_i).$$

The noise $\xi_i(a)$ represents estimation errors or random perturbations. Following standard assumptions in stochastic approximation and learning in games, we require the noise to be unbiased with uniformly bounded variance.

ASSUMPTION 1. For each player i and action $a \in A_i$, the noise satisfies

$$\mathbb{E}[\xi_i(a)] = 0, \quad \text{Var}(\xi_i(a)) \leq \sigma^2 < \infty.$$

This ensures that observed payoffs remain correct on average, while the variance bound σ^2 provides a uniform control of uncertainty and will enter directly into our theoretical results.

3 BNN DYNAMICS UNDER NOISY FEEDBACK

We now analyze the BNN dynamics under noisy feedback in the normal-form setting. Our aim is to characterize how stochastic perturbations influence the learning trajectory and whether the convergence guarantees of the deterministic process persist in this stochastic regime. This section provides a foundation for examining stability and bias before extending the analysis to more general sequential games.

3.1 Discrete Update and SA Form

Having established the noiseless BNN dynamics in Section 2.3, we now turn to the setting with noisy feedback. To analyse the resulting stochastic recursion, it is convenient to introduce the vector-field notation. Let $H(\pi)$ denote the noiseless update direction defined in Section 2.3. Under noisy feedback (Assumption 1), the recursion becomes

$$\pi_{t+1} = \pi_t + \eta_t \widehat{H}(\pi_t),$$

with noisy field components

$$\widehat{H}_i(a; \pi) = [\tilde{u}_i(a) - \tilde{u}_i]_+ - \pi_i(a) \sum_{a' \in A_i} [\tilde{u}_i(a') - \tilde{u}_i]_+.$$

Because $[\cdot]_+$ is convex, the expectation of the noisy update differs from the noiseless field, creating a structural bias (a Jensen gap). Formally, the noisy field admits the decomposition

$$\mathbb{E}[\widehat{H}(\pi_t)] = H(\pi_t) + \beta(\pi_t),$$

where $\beta(\pi)$ denotes the bias.

To proceed with the stochastic approximation analysis, it is important to understand the size of this bias.

LEMMA 1. *For a player with action set A_i , the structural bias term $\beta(\pi)$ is uniformly bounded over the simplex as*

$$\|\beta(\pi)\|_\infty \leq (|A_i| - 1) \sigma,$$

where σ^2 is the variance bound in Assumption 1.

In addition to the bias, the decomposition also produces a residual noise term. Define

$$\zeta_{t+1} := \widehat{H}(\pi_t) - \mathbb{E}[\widehat{H}(\pi_t)].$$

This is a martingale-difference sequence with zero conditional mean, i.e. $\mathbb{E}[\zeta_{t+1} | \pi_t] = 0$. Since $[\cdot]_+$ is 1-Lipschitz and the payoff noise has bounded variance by Assumption 1, the variance of ζ_{t+1} is uniformly bounded. In particular,

$$\mathbb{E}[\|\zeta_{t+1}\|^2 | \pi_t] \leq (1 + |A_i|)^2 \sigma^2.$$

For brevity we henceforth write this bound as $C\sigma^2$.

Collecting the bias and noise terms, the recursion takes the stochastic approximation form

$$\pi_{t+1} = \pi_t + \eta_t (H(\pi_t) + \beta(\pi_t) + \zeta_{t+1}).$$

3.2 Convergence Results

To analyse convergence, we first recall that the continuous-time BNN flow $\dot{\pi} = H(\pi)$ admits a Lyapunov potential. Specifically,

$$\Gamma(\pi) := \frac{1}{2} \sum_{a \in A_i} [u_i(a; \pi) - u_i(\pi)]_+^2$$

serves as a strict Lyapunov function in normal-form zero-sum games (see, e.g., [ref]). For convenience, we also introduce

$$S(\pi) := \sum_{a \in A_i} [u_i(a; \pi) - u_i(\pi)]_+,$$

which represents the aggregate mass of profitable deviations. The two quantities are linked by the elementary inequality $S(\pi)^2 \geq 2\Gamma(\pi)$.

For the noiseless BNN flow, the Lyapunov potential satisfies the differential identity

$$\frac{d}{dt} \Gamma(\pi) = -2S(\pi)\Gamma(\pi).$$

Together with the inequality $S(\pi)^2 \geq 2\Gamma(\pi)$ established above, this implies

$$\frac{d}{dt} \Gamma(\pi) \leq -2\sqrt{2}\Gamma(\pi)^{3/2},$$

demonstrating strict dissipation whenever profitable deviations remain.

In the presence of noisy feedback, the exact dissipation identity no longer holds. The structural bias $\beta(\pi)$ perturbs the dynamics and the martingale noise adds further fluctuations, so the Lyapunov descent is weakened. In the discrete recursion this effect is captured by a one-step drift inequality in expectation, which combines the dissipative behaviour of the noiseless flow with the perturbations due to bias and noise.

Combining the bias bound of Lemma 1 with the bounded-variance noise, and under Robbins–Monro step size conditions, we obtain the following drift inequality.

LEMMA 2. *Denote $g_t := \mathbb{E}[\Gamma(\pi_t)]$. Suppose the step sizes $\{\eta_t\}$ satisfy $\sum_t \eta_t = \infty$ and $\sum_t \eta_t^2 < \infty$ (e.g. $\eta_t = c/(t + t_0)^{2/3}$). Then*

$$g_{t+1} \leq g_t - 2\sqrt{2}\eta_t g_t^{3/2} + (|A_i| - 1)\sqrt{2|A_i|}\sigma\eta_t g_t^{1/2} + C_3\eta_t^2,$$

where $C_3 > 0$ is a finite constant depending only on the game, but independent of σ and t .

The constant C_3 arises from the quadratic remainder of the Taylor expansion. It depends on game-specific Lipschitz constants and a uniform bound on the second moments of the noise. Since it always appears multiplied by η_t^2 and $\sum_t \eta_t^2 < \infty$, this contribution is summable and does not affect the asymptotic analysis.

The one-step inequality of Lemma 2 shows that the Lyapunov potential decreases in expectation, up to two perturbations: the quadratic remainder $C_3\eta_t^2$ and the linear bias term proportional to σ . Building on this descent, we now establish two global consequences: (i) almost sure stability within an $O(\sigma)$ neighbourhood, and (ii) a quantitative rate of approach to this noise floor. Under the standard Robbins–Monro step-size conditions for stochastic approximation, we obtain the following stability bound. These conditions ensure that the iterates asymptotically follow the mean-field dynamics while the accumulated variance remains finite.

THEOREM 1. *Suppose the step sizes $\{\eta_t\}$ satisfy $\sum_t \eta_t = \infty$ and $\sum_t \eta_t^2 < \infty$. Then*

$$\limsup_{t \rightarrow \infty} \Gamma(\pi_t) \leq \frac{1}{2}(|A_i| - 1)\sqrt{|A_i|}\sigma \quad \text{almost surely.}$$

This theorem shows that the iterates stabilise almost surely within an $O(\sigma)$ neighbourhood of the Nash set, so the asymptotic error floor scales linearly with the noise level. The proof relies on completing the square in the drift inequality and applying the Robbins–Siegmund supermartingale theorem; full details are provided in Appendix.

While Theorem 1 establishes asymptotic stability, it does not quantify the speed of convergence. To address this, we analyse the transient decay of Γ until it reaches the $O(\sigma)$ floor. This is achieved by rewriting the drift inequality as a discrete Riccati-type recursion and comparing it with an auxiliary differential equation.

THEOREM 2. *Let $\eta_t = c/(t + t_0)^{2/3}$ with $c > 0$ and $t_0 \geq 1$. Then*

$$\mathbb{E}[\Gamma(\pi_t)] = O(t^{-2/3})$$

until it stabilises at the $O(\sigma)$ floor described in Theorem 1. In the noiseless case ($\sigma = 0$), the rate $O(t^{-2/3})$ holds for all t .

Theorem 2 shows that the decay exponent $2/3$ coincides with the noiseless case; noise only determines the eventual $O(\sigma)$ floor. The proof is based on a discrete comparison argument using the substitution $y_t = \sqrt{\mathbb{E}[\Gamma(\pi_t)]}$; details are provided in Appendix. Two further remarks are worth noting: (i) with a constant step size η , the process attains a steady-state error of order $O(\sigma\eta) + O(\eta^2)$; (ii) the convergence is not in a standard norm but in a payoff-based measure, which coincides with exploitability and regret. This perspective makes the result more widely applicable, as elaborated in the next section.

3.3 Bias and Centroid Shift

Theorem 1 established that the noisy recursion stabilises within an $O(\sigma)$ neighbourhood in Γ , reflecting the stochastic fluctuations of the process. We now turn to a different question: how far the stationary points of the biased dynamics themselves are displaced from Nash equilibria. This effect is deterministic, arising solely from the structural bias β , and requires a separate line of analysis based on bounding the maximal advantage at a biased fixed point.

Formally, a biased fixed point π' is defined by

$$H(\pi') + \beta(\pi') = 0.$$

At such a point, the maximal advantage satisfies $\max_{a \in A_i} \text{Adv}(a) \leq (|A_i| - 1)\sigma$. Since the Lyapunov potential is quadratic in advantages, this yields the following bound.

THEOREM 3 (CENTROID SHIFT). *Let π' be a fixed point of the biased recursion, i.e. $H(\pi') + \beta(\pi') = 0$. Then*

$$\Gamma(\pi') \leq \frac{1}{2}|A_i|(|A_i| - 1)^2\sigma^2.$$

This shows that structural bias shifts the centroid of the dynamics only quadratically in σ : biased fixed points lie within an $O(\sigma^2)$ neighbourhood in Γ . In contrast, Theorem 1 describes a stochastic noise floor of order $O(\sigma)$. Taken together, these results imply that as the variance of the noise vanishes ($\sigma^2 \rightarrow 0$), both the fluctuation floor and the centroid shift disappear, and the iterates converge to exact Nash equilibria.

4 BNN DYNAMICS IN EXTENSIVE-FORM GAMES

Having established the properties of the BNN dynamics in normal-form games, we now extend the framework to extensive-form games with imperfect information. This setting captures sequential decision processes and therefore makes our framework more relevant to practical multi-agent learning and decision-making scenarios. The following section reformulates the instantaneous regret descent in terms of local information sets and introduces reach-weighted updates that generalize the normal-form dynamics to sequential play.

4.1 Local Regret Structure and Reach Weights

In normal-form games the dynamics are driven by the positive part of the payoff difference, $[\Delta_i(a | \pi)]_+$, which is exactly the instantaneous regret of action a . The aggregate $S_i(\pi) = \sum_a [\Delta_i(a | \pi)]_+$ is the total regret mass considered in regret-minimisation. Moreover, $S_i(\pi)$ is bounded (up to a normalisation) by the potential $\Gamma(\pi)$ introduced in Section 3, which is non-increasing along the flow and functions as a Lyapunov function for the dynamics. Because S_i is the sum of positive regrets, this identification provides the minimal bridge from the NFG analysis to the extensive-form case.

In the extensive-form setting the analogue of this signal is defined at each information set. For player i at information set x , let $u_i^\pi(x, a)$ be the expected payoff conditional on reaching x and choosing action a , and let $\bar{u}_i^\pi(x) = \langle u_i^\pi(x, \cdot), \pi_i(\cdot | x) \rangle$. The difference

$$A_i^\pi(x, a) = u_i^\pi(x, a) - \bar{u}_i^\pi(x)$$

is then used in place of $\Delta_i(a | \pi)$. Its positive part coincides with the local regret notion at information sets, making it a natural replacement for the normal-form signal.

Using only the payoff differences at information sets would still fail to capture their correct impact on the overall dynamics. A local deviation at x contributes to the game outcome only in proportion to the probability that the play actually reaches x , which is determined by the opponent’s reach probability $\rho_{-i}^\pi(x)$. To obtain the correct scaling, the update at each information set must therefore be multiplied by this factor. The resulting extensive-form dynamics take the form

$$\dot{\pi}_i(a | x) = \rho_{-i}^\pi(x) \left([A_i^\pi(x, a)]_+ - \pi_i(a | x) \sum_{b \in A(x)} [A_i^\pi(x, b)]_+ \right),$$

which is the direct analogue of the BNN dynamics in normal-form games, now lifted to the information-set decomposition with reach weighting.

4.2 Theoretical Results in EFG

We now turn to the theoretical analysis of the extensive-form dynamics introduced above. All results from Section 3 extend directly once local counterfactual advantages are weighted by the opponent’s reach probabilities. To avoid repetition, we state the reach-weighted counterparts of the normal-form results and highlight their parallels.

We first focus on the noiseless continuous-time dynamics. Define the reach-weighted potential

$$V(\pi) = \sum_i \sum_{x \in \mathcal{X}_i} \Gamma_i^x(\pi), \quad \Gamma_i^x(\pi) = \frac{1}{2} \|\rho_{-i}^\pi(x) A_{i,+}^\pi(x, \cdot)\|^2.$$

LEMMA 3. *Along the noiseless extensive-form dynamics, the potential $V(\pi)$ satisfies the exact identity*

$$\begin{aligned} \frac{d}{dt} V(\pi) &= - \sum_i \sum_{x \in \mathcal{X}_i} 2 S_i^x(\pi) \Gamma_i^x(\pi), \\ S_i^x(\pi) &= \sum_{a \in A(x)} \rho_{-i}^\pi(x) [A_i^\pi(x, a)]_+. \end{aligned}$$

In particular, $\dot{V}(\pi) < 0$ whenever $\pi \notin \Pi^*$, where Π^* denotes the equilibrium set.

Because BNN dynamics preserve the interior of the simplex, all actions maintain strictly positive probability, so every relevant reach probability $\rho_{-i}^\pi(x)$ remains bounded away from zero. This ensures that the weighted dissipation in the identity is strictly positive outside equilibrium, establishing V as a Lyapunov function.

We now turn to the stochastic recursion with noisy counterfactual feedback. To distinguish from the normal-form setting, we use a prime notation to denote the extensive-form counterparts of the quantities in Section 3, e.g., H' , β' , and ζ' .

With Robbins–Monro step sizes $\{\alpha_k\}$, the update at each information set $x \in \mathcal{X}_i$ takes the form

$$\begin{aligned} \pi_i^{k+1}(a|x) &= \pi_i^k(a|x) + \alpha_k \rho_{-i}^{\pi^k}(x) \\ &\quad \left([\widetilde{A}_i^{\pi^k}(x, a)]_+ - \pi_i^k(a|x) \sum_{b \in A(x)} [\widetilde{A}_i^{\pi^k}(x, b)]_+ \right), \end{aligned}$$

where $\widetilde{A}_i^{\pi^k}(x, a)$ is a possibly noisy estimator of the counterfactual advantage.

As in Section 3, the convexity of $[\cdot]_+$ creates a structural bias $\beta'(\pi)$, while the sampling error yields a martingale-difference noise ζ'_{k+1} with bounded variance. The form of these bounds is identical to the normal-form case, but the constants now depend on the local action set sizes $\{|A(x)| : x \in \mathcal{X}_i\}$.

LEMMA 4 (BIAS AND NOISE BOUNDS IN EFG). *There exist finite constants $C_1, C_2 > 0$, depending only on the information-set structure, such that for all π*

$$\|\beta'(\pi)\| \leq C_1 \sigma, \quad \mathbb{E}[\|\zeta'_{k+1}\|^2 | \pi^k] \leq C_2 \sigma^2,$$

where σ^2 is the uniform variance bound of the payoff noise.

Collecting these terms, the recursion admits the stochastic approximation representation

$$\pi^{k+1} = \pi^k + \alpha_k (H'(\pi^k) + \beta'(\pi^k) + \zeta'_{k+1}),$$

where H' denotes the noiseless extensive-form BNN field.

With these bounds in place, the Lyapunov potential admits a quantitative descent estimate in expectation, stated as follows.

LEMMA 5 (ONE-STEP EXPECTED DESCENT IN EFG). *Let $g_k := \mathbb{E}[V(\pi^k)]$, where V is the Lyapunov function defined in Lemma 4.2.1. Suppose the step sizes $\{\alpha_k\}$ satisfy $\sum_k \alpha_k = \infty$ and $\sum_k \alpha_k^2 < \infty$. Then there exist constants $C_1, C_2 > 0$, depending only on the information-set structure, such that*

$$g_{k+1} \leq g_k - 2\alpha_k \rho g_k^{3/2} + C_1 \sigma \alpha_k g_k^{1/2} + C_2 \alpha_k^2,$$

where $\rho > 0$ is a uniform lower bound on reach probabilities in the interior of the simplex, and σ^2 is the variance bound of the payoff noise.

Building on the one-step descent inequality, we now establish the asymptotic behaviour of the stochastic recursion in the extensive-form setting. The statements mirror those of Section 3, but with constants depending only on the information-set structure. For completeness, we collect the results here using the same Lyapunov potential $V(\pi)$.

THEOREM 4 (ASYMPTOTIC STABILITY IN EFG). *Let $\{\pi^k\}$ follow the stochastic recursion defined above with Robbins–Monro step sizes. Then the Lyapunov potential satisfies*

$$\limsup_{k \rightarrow \infty} \mathbb{E}[V(\pi^k)] = O(\sigma),$$

where σ^2 is the variance of the payoff noise.

THEOREM 5 (CONVERGENCE RATE IN EFG). *Suppose the step sizes are chosen as $\alpha_k = c/(k + t_0)^{2/3}$ with constants $c, t_0 > 0$. Then the expected potential decreases at rate*

$$\mathbb{E}[V(\pi^k)] = O(k^{-2/3}) + O(\sigma),$$

so the iterates converge to an $O(\sigma)$ neighbourhood of equilibrium at this rate.

THEOREM 6 (CENTROID SHIFT IN EFG). *In the presence of persistent noise, the stationary points of the recursion correspond to approximate equilibria where the Lyapunov potential satisfies*

$$V(\pi) = O(\sigma^2).$$

Thus the asymptotic bias induced by noise manifests only at order $O(\sigma^2)$.

These results confirm that the last-iterate guarantees obtained in the normal-form analysis extend directly to extensive-form games once counterfactual advantages are employed. In particular, the asymptotic stability and convergence rate are limited also by an $O(\sigma)$ noise floor, while the systematic bias introduced by noise appears only at order $O(\sigma^2)$. Thus the extensive-form recursion inherits the same robustness properties as in the normal-form case, with no additional complications beyond the dependence of constants on the local action set sizes.

4.3 Algorithmic Implementation

We implement the BNN Actor–Critic (BNNAC) within an actor–critic framework. A key feature of our design is that all three components—the actor, critic, and reach estimator—are represented by neural networks. While a tabular form would be sufficient for small games, extensive-form games can be extremely large, and neural networks provide the scalable function approximation required to handle such settings efficiently. The algorithm is distributed: both players run BNNAC independently and simultaneously, without direct communication, and their interaction occurs only through the trajectories generated by the underlying game.

The actor-logits network L_{θ^L} outputs logits, which are mapped to a policy via a softmax transformation. The critic network Q_{θ^Q} predicts payoffs and is used to compute local advantage estimates. The reach network ρ_{θ^R} estimates opponent reach probabilities at

Algorithm 1 BNN Actor–Critic (BNNAC)

```

1: Input: Initial parameters for three networks: actor-logits  $\theta_0^L$ 
   (logit network  $L_{\theta^L}$ , policy  $\pi = \text{softmax}(L_{\theta^L})$ ), critic  $\theta_0^Q$  ( $Q_{\theta^Q}$ 
   predicts payoffs), reach  $\theta_0^P$  ( $\rho_{\theta^P}$  estimates opponent reach). Step
   sizes  $\{\eta_t\}, \{\alpha_t\}, \{\beta_t\}$  and update interval  $K$ .
2: for  $t = 0, 1, 2, \dots$  do
3:   Obtain data  $\mathcal{D} = \{(x, a, r, x')\}$  from the game
4:   Construct loss  $\ell(\theta_t^Q; \mathcal{D})$  from the data
5:   Update critic:  $\theta_{t+1}^Q \leftarrow \theta_t^Q - \alpha_t \nabla_{\theta^Q} \ell(\theta_t^Q; \mathcal{D})$ 
6:   Update reach network:  $\theta_{t+1}^P \leftarrow \theta_t^P - \beta_t \nabla_{\theta^P} \ell_{\rho}(\theta_t^P; \mathcal{D})$ 
7:   if  $t \bmod K = 0$  then
8:     Estimate local advantages  $A_i^{\pi}(x, a)$  using  $Q_{\theta_{t+1}^Q}$ 
9:     Estimate reach probabilities  $\rho(x)$  using  $\rho_{\theta_{t+1}^P}$ 
10:    Compute logit update direction:
        
$$\Delta L(a|x) = \rho_{\theta_{t+1}^P}(x) \frac{A_i^{\pi}(x, a)}{\pi_{\theta_t^L}(a|x)}$$

11:    Update actor:  $\theta_{t+1}^L \leftarrow \theta_t^L + \eta_t \nabla_{\theta^L} \langle L_{\theta_t^L}(x), \Delta L(x) \rangle$ 
12:    Update policy:  $\pi_{t+1}(a|x) = \text{softmax}(L_{\theta_{t+1}^L}(x))$ 
13:   end if
14: end for

```

information sets. Together these three networks provide the quantities required by the extensive-form BNN update.

Each training iteration obtains data (x, a, u, x') from the game under the current policy. The critic parameters θ^Q are updated to reduce the prediction error of observed payoffs, while the reach parameters θ^P are adjusted according to the observed visitation frequencies. Every K iterations, the actor is updated: advantages are computed from the critic, reach probabilities from the reach network, and the actor logits are then moved in the BNN direction before being mapped through the softmax transformation.

This design follows the theoretical analysis in two ways. First, the critic and reach networks are updated more frequently to ensure that the feedback received by the actor can be treated as approximately unbiased. Second, the softmax ensures that the policy remains on the simplex. Since the actor is parameterized in logits, propagating the update through softmax introduces an additional $\frac{1}{\pi}$ factor in the update direction. Importantly, BNN dynamics never eliminate strategies completely, so the denominator is always well defined; in practice, a small positive floor is used to prevent numerical underflow. Finally, introducing a separate reach network cleanly separates structural counterfactual weighting from payoff learning, which both clarifies the theoretical correspondence and stabilizes training. If the reach probabilities are omitted, the same structure directly reduces to the normal-form BNN actor–critic framework.

5 EXPERIMENTS

For validation, we benchmark our proposed BNN-based framework against state-of-the-art regularized RD-based approaches in both normal-form and extensive-form games. In each game, we add unbiased noise directly into the rewards in the normal-form case or the counterfactual values in the extensive-form case. To this end, we consider Biased Rock–Paper–Scissors (BRPS) and its four-action variant (BRPS-W), as well as Kuhn Poker and Leduc Poker.

To further test adaptability, we construct non-stationary versions of BRPS and Kuhn Poker, where the underlying payoffs change either instantaneously (direct-shift mode) or gradually over time (continuous-shift mode). As a baseline, we compare against regularized RD-based approaches, including adaptively perturbed mirror descent (APMD) [2] for normal-form games and R-Nad [44] for extensive-form games. They are both state-of-the-art algorithms that achieve last-iterate convergence in those settings.

In stationary games, our method exhibits markedly faster convergence and smoother trajectories throughout training, achieving low exploitability earlier and maintaining stability over extended horizons. Whereas the baseline occasionally attains a lower asymptotic NASHCONV, this difference corresponds to the noise plateau predicted by our theoretical analysis. Under non-stationary conditions, our approach demonstrates superior adaptability: it rapidly converges after abrupt payoff changes and sustains low exploitability during gradual transitions.

Normal-form Games. The RPS family of games is defined by three parameters controlling the relative advantages between the pairwise matchups (a_{RP}, a_{PS}, a_{SR}). In the BRPS environment, a_{RP} is fixed at 12, creating a persistent bias that yields a unique interior Nash equilibrium. The BRPS-W variant extends this setting with a fourth action that interacts neutrally with all existing ones, producing payoffs of zero against every opponent action. For these experiments, we revert to a purely tabular version of the algorithm without any function approximation. This design isolates the effect of the underlying dynamics from the influence of neural approximation, allowing a direct comparison between the methods.

As depicted in Fig. 2, the BNN-based method maintains low exploitability and stable convergence throughout training until it reaches the noise plateau. Once the trajectory stabilizes near this region, the confidence intervals widen slightly in the later stages, indicating a modest increase in variance. This effect is consistent with our theoretical analysis, which predicts that the influence of stochastic noise becomes amplified only in the close vicinity of the Nash equilibrium. In contrast, the APMD displays large confidence intervals in early stages of training. These occur when its trajectories move too close to the simplex boundary, causing unstable probability mass distributions. This edge concentration leads to substantial fluctuations and irregular convergence behavior.

To assess adaptability under changing environments, we test a non-stationary extension of the Biased Rock–Paper–Scissors game in which the payoff parameters vary over time while preserving the zero-sum game structure shown in Fig. 3. The three matchup coefficients (a_{RP}, a_{PS}, a_{SR}) change sequentially such that the bias rotates between matchups: $(12, 1, 1) \rightarrow (6.5, 6.5, 1) \rightarrow (1, 12, 1)$.




Figure 2: BRPS and BRPS-W.

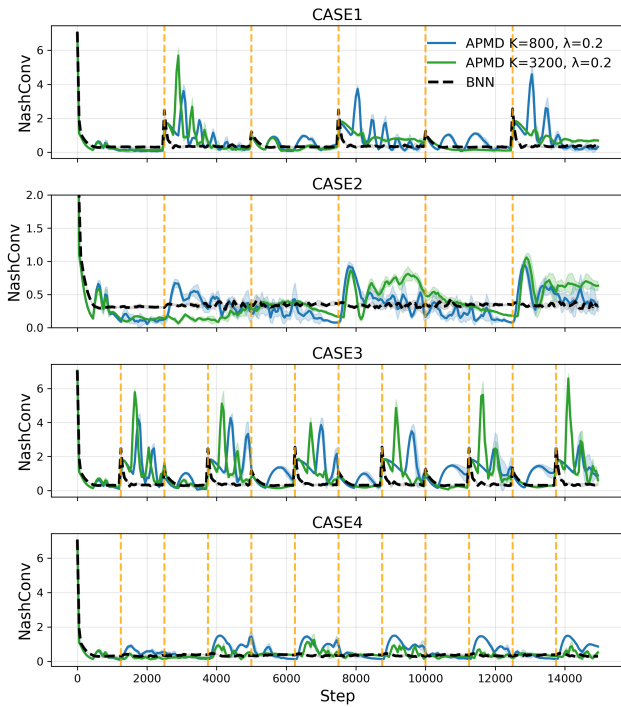


Figure 3: Nonstationary RPS.

We consider four variants of this process. In case 1, each stage lasts for 2500 steps. In case 2, the game first remains fixed for 2500 steps, and the subsequent transition between configurations is made gradually over another 2500 steps. Cases 3 and 4 follow the same patterns as cases 1 and 2, respectively, but each stage lasts for only 1250 steps. In this setting, we further evaluate the APMD under different hyperparameter configurations, varying the regularization update frequency (K) and the regularization strength (λ). Our BNN-based algorithm maintains low NASHCONV after each environmental transition, showing stable adaptation across all tested cases. In comparison, APMD adapts more slowly and exhibits transient instability following each change. Its performance also differs noticeably across transition speeds. Moreover, different combinations of hyperparameters yield inconsistent outcomes across transition modes, indicating that no single configuration performs robustly.

Extensive-form Games. We now consider stationary Kuhn and Leduc poker, and also construct two non-stationary variants of Kuhn Poker, in which the underlying payoffs evolve gradually over time as shown in Fig. 4. For all experiments, we employ the full neural implementation of our framework, called BNNAC, where both the actor and critic networks are trained with reach-weighted returns. The state-of-the-art baseline, R-NaD, is implemented under the same network architecture and optimization settings in the noisy environment. For all the environments, we sweep the regularization strength to examine the sensitivity of R-NaD to the rate and magnitude of environmental change. In the stationary settings, both algorithms exhibit rapid early-stage improvement in NASHCONV, with BNNAC showing a slightly faster initial descent. Overall, the

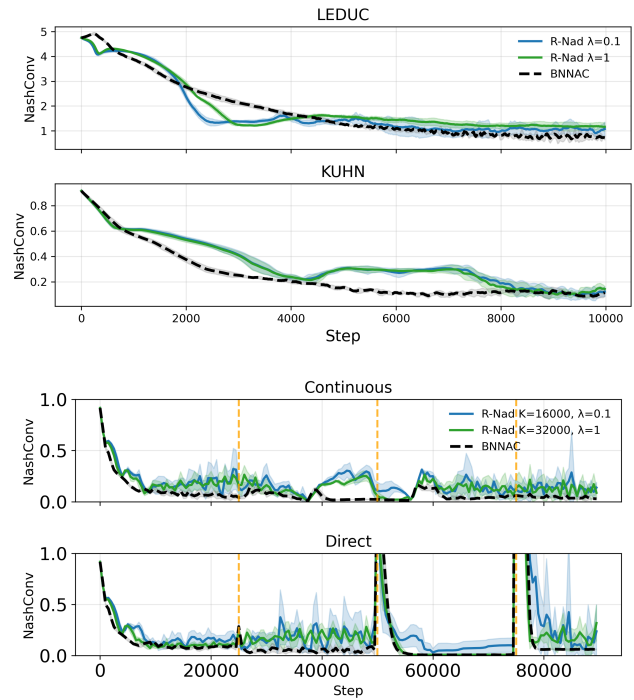


Figure 4: Stationary and nonstationary poker games.

results confirm that both approaches demonstrate stable convergence in extensive-form games under stationary conditions.

In the non-stationary settings, we modify the bet size over three distinct stages, changing from the initial value of 1 to 2, then to -2 , and finally to 6. In the continuous variant, the bet size smoothly returns to 1 in the last stage. Across all cases, we observe that the current strategies experience inevitable deviations in NASHCONV whenever the environment undergoes a major transition. Under these conditions, the two algorithms exhibit qualitatively different behaviors. While neither method can completely offset the environmental shifts, BNNAC consistently achieves lower NASHCONV, whereas R-NaD displays unstable oscillatory patterns.

6 CONCLUSION

This work revisits the classical BNN dynamics and establishes its theoretical and practical relevance in modern multi-agent learning. We show that BNN dynamics guarantee regularization-free last-iterate convergence in both normal-form and extensive-form zero-sum games, maintaining stability under noisy feedback. The resulting BNNAC algorithm scales to neural implementations and consistently outperforms baselines in both stationary and non-stationary environments. Future work will focus on accelerating convergence and further mitigating the effect of stochastic noise to enhance performance in large-scale and highly dynamic settings.

ACKNOWLEDGMENTS

LS and TZ have been supported by the Advanced Research and Invention Agency (ARIA), Safeguarded AI programme.

REFERENCES

- [1] Sherief Abdallah and Victor Lesser. 2008. A multiagent reinforcement learning algorithm with non-linear dynamics. *Journal of Artificial Intelligence Research* 33 (2008), 521–549.
- [2] Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, and Atsushi Iwasaki. 2023. Adaptively perturbed mirror descent for learning in games. *arXiv preprint arXiv:2305.16610* (2023).
- [3] Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, Kentaro Toyoshima, and Atsushi Iwasaki. 2022. Last-iterate convergence with full and noisy feedback in two-player zero-sum games. *arXiv preprint arXiv:2208.09855* (2022).
- [4] Kenshi Abe, Mitsuki Sakamoto, Kaito Ariu, and Atsushi Iwasaki. 2024. Boosting Perturbed Gradient Ascent for Last-Iterate Convergence in Games. *arXiv preprint arXiv:2410.02388* (2024).
- [5] Kenshi Abe, Mitsuki Sakamoto, and Atsushi Iwasaki. 2022. Mutation-driven follow the regularized leader for last-iterate convergence in zero-sum games. In *Uncertainty in Artificial Intelligence*. PMLR, 1–10.
- [6] Wolfram Barfuss, Jonathan F Donges, and Jürgen Kurths. 2019. Deterministic limit of temporal difference reinforcement learning for stochastic games. *Physical Review E* 99, 4 (2019), 043305.
- [7] Tilman Börgers and Rajiv Sarin. 1997. Learning through reinforcement and replicator dynamics. *Journal of economic theory* 77, 1 (1997), 1–14.
- [8] Vicent Botti. 2025. Agentic AI and Multiagent: Are We Reinventing the Wheel? *arXiv preprint arXiv:2506.01463* (2025).
- [9] George W Brown. 1951. Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation* 13, 1 (1951), 374.
- [10] Yang Cai, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. 2023. Uncoupled and convergent learning in two-player zero-sum markov games with bandit feedback. *Advances in Neural Information Processing Systems* 36 (2023), 36364–36406.
- [11] Yang Cai and Weiqiang Zheng. 2023. Doubly optimal no-regret learning in monotone games. In *International Conference on Machine Learning*. PMLR, 3507–3524.
- [12] Shicong Cen, Yuejie Chi, Simon S Du, and Lin Xiao. 2022. Faster last-iterate convergence of policy optimization in zero-sum Markov games. *arXiv preprint arXiv:2210.01050* (2022).
- [13] Shicong Cen, Yuting Wei, and Yuejie Chi. 2021. Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems* 34 (2021), 27952–27964.
- [14] Pierre Coucheny, Bruno Gaujal, and Panayotis Mertikopoulos. 2015. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research* 40, 3 (2015), 611–633.
- [15] John G Cross. 1973. A stochastic learning model of economic behavior. *The quarterly journal of economics* 87, 2 (1973), 239–266.
- [16] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. 2017. Training gans with optimism. *arXiv preprint arXiv:1711.00141* (2017).
- [17] Constantinos Daskalakis and Ioannis Panageas. 2018. Last-iterate convergence: Zero-sum games and constrained min-max optimization. *arXiv preprint arXiv:1807.04252* (2018).
- [18] Constantinos Daskalakis and Ioannis Panageas. 2018. The limit points of (optimistic) gradient descent in min-max optimization. *Advances in neural information processing systems* 31 (2018).
- [19] Francisco Facchinei and Jong-Shi Pang. 2003. *Finite-dimensional variational inequalities and complementarity problems*. Springer.
- [20] Eduard Gorbunov, Adrien Taylor, and Gauthier Gidel. 2022. Last-iterate convergence of optimistic gradient method for monotone variational inequalities. *Advances in neural information processing systems* 35 (2022), 21858–21870.
- [21] Yuzhen Han, Ruben Solozabal, Jing Dong, Xingyu Zhou, Martin Takac, and Bin Gu. 2022. Learning to control under time-varying environment. *arXiv preprint arXiv:2206.02507* (2022).
- [22] Daniel Hennes, Michael Kaisers, and Karl Tuyls. 2010. RESQ-learning in stochastic games. In *Adaptive and Learning Agents Workshop at AAMAS*. Citeseer, 8.
- [23] Daniel Hennes, Dustin Morrill, Shayegan Omidshafiei, Rémi Munos, Julien Perolat, Marc Lanctot, Audrunas Gruslys, Jean-Baptiste Lespiau, Paavo Parmas, Edgar Duñez-Guzmán, et al. 2020. Neural replicator dynamics: Multiagent learning via hedging policy gradients. In *Proceedings of the 19th international conference on autonomous agents and multiagent systems*. 492–501.
- [24] Daniel Hennes, Karl Tuyls, and Matthias Rauterberg. 2009. State-coupled replicator dynamics. In *AAMAS (2)*. 789–796.
- [25] Josef Hofbauer. 2011. Deterministic evolutionary game dynamics. (2011).
- [26] Josef Hofbauer, Jörg Oechssler, and Frank Riedel. 2009. Brown–von Neumann–Nash dynamics: the continuous strategy case. *Games and Economic Behavior* 65, 2 (2009), 406–429.
- [27] Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. 2021. Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium. In *Conference on Learning Theory*. PMLR, 2388–2422.
- [28] Aamal Abbas Hussain, Francesco Belardinelli, and Georgios Piliouras. 2023. Asymptotic convergence and performance of multi-agent Q-learning dynamics. *arXiv preprint arXiv:2301.09619* (2023).
- [29] Michael Kaisers, Daan Bloembergen, and Karl Tuyls. 2012. A common gradient in multi-agent reinforcement learning. In *AAMAS*. 1393–1394.
- [30] Michael Kaisers and Karl Tuyls. 2010. Frequency adjusted multi-agent Q-learning. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. 309–316.
- [31] Ardeshir Kianercy and Aram Galstyan. 2012. Dynamics of Boltzmann Q learning in two-player two-action games. *Physical Review E* 85, 4 (2012), 041145.
- [32] Jayash Koshal, Angelia Nedić, and Uday V Shanbhag. 2010. Single timescale regularized stochastic approximation schemes for monotone Nash games under uncertainty. In *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 231–236.
- [33] 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. *Advances in neural information processing systems* 30 (2017).
- [34] Sucheol Lee and Donghwan Kim. 2021. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. *Advances in Neural Information Processing Systems* 34 (2021), 22588–22600.
- [35] David S Leslie and Edmund J Collins. 2005. Individual Q-learning in normal form games. *SIAM Journal on Control and Optimization* 44, 2 (2005), 495–514.
- [36] Mingyang Liu, Asuman Ozdaglar, Tiancheng Yu, and Kaiqing Zhang. 2022. The power of regularization in solving extensive-form games. *arXiv preprint arXiv:2206.09495* (2022).
- [37] Richard D McKelvey and Thomas R Palfrey. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10, 1 (1995), 6–38.
- [38] Richard D McKelvey and Thomas R Palfrey. 1998. Quantal response equilibria for extensive form games. *Experimental economics* 1, 1 (1998), 9–41.
- [39] Linjiang Meng, Youzhi Zhang, Zhenxing Ge, Tianyu Ding, Shangdong Yang, Zheng Xu, Wenbin Li, and Yang Gao. 2025. Last-Iterate Convergence of Smooth Regret Matching+ Variants in Learning Nash Equilibria. *Advances in Neural Information Processing Systems* 38 (2025).
- [40] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. 2018. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. *arXiv preprint arXiv:1807.02629* (2018).
- [41] Panayotis Mertikopoulos and William H Sandholm. 2018. Riemannian game dynamics. *Journal of Economic Theory* 177 (2018), 315–364.
- [42] Sarath Pattathil, Kaiqing Zhang, and Asuman Ozdaglar. 2023. Symmetric (optimistic) natural policy gradient for multi-agent learning with parameter convergence. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 5641–5685.
- [43] Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T Connor, Neil Burch, Thomas Anthony, et al. 2022. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science* 378, 6623 (2022), 990–996.
- [44] Julien Perolat, Remi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, Mark Rowland, Pedro Ortega, Neil Burch, Thomas Anthony, David Balduzzi, Bart De Vylder, et al. 2021. From poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. In *International Conference on Machine Learning*. PMLR, 8525–8535.
- [45] Alexander Rakhlin and Karthik Sridharan. 2013. Online learning with predictable sequences. In *Conference on Learning Theory*. PMLR, 993–1019.
- [46] Sasha Rakhlin and Karthik Sridharan. 2013. Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems* 26 (2013).
- [47] Julia Robinson. 1951. An iterative method of solving a game. *Annals of mathematics* 54, 2 (1951), 296–301.
- [48] Shai Shalev-Shwartz and Yoram Singer. 2006. Convex repeated games and Fenchel duality. *Advances in neural information processing systems* 19 (2006).
- [49] Yoav Shoham, Rob Powers, and Trond Grenager. 2007. If multi-agent learning is the answer, what is the question? *Artificial intelligence* 171, 7 (2007), 365–377.
- [50] Samuel Sokota, Ryan D’Orazio, J Zico Kolter, Nicolas Loizou, Marc Lanctot, Ioannis Mitliagkas, Noam Brown, and Christian Kroer. 2022. A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. *arXiv preprint arXiv:2206.05825* (2022).
- [51] Tatiana Tatarenko and Maryam Kamgarpour. 2019. Learning Nash equilibria in monotone games. In *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 3104–3109.
- [52] Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D Nguyen. [n.d.]. Multi-agent collaboration mechanisms: A survey of LLMs, 2025. URL <https://arxiv.org/abs/2501.06322> ([n. d.]).
- [53] Karl Tuyls, Katja Verbeeck, and Tom Lenaerts. 2003. A selection-mutation model for q-learning in multi-agent systems. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. 693–700.
- [54] John von Neumann and Oskar Morgenstern. 2007. Theory of games and economic behavior: 60th anniversary commemorative edition. In *Theory of games and economic behavior*. Princeton university press.

- [55] Peter Vrancx, Karl Tuyls, Ronald L Westra, and Ann Nowé. 2008. Switching dynamics of multi-agent learning. *AAMAS (1) 2008 (2008)*, 307–313.
- [56] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. 2020. Linear last-iterate convergence in constrained saddle-point optimization. *arXiv preprint arXiv:2006.09517 (2020)*.
- [57] Farzad Yousefian, Angelia Nedić, and Uday V Shanbhag. 2017. On smoothing, regularization, and averaging in stochastic approximation methods for stochastic variational inequality problems. *Mathematical Programming* 165, 1 (2017), 391–431.
- [58] Tuo Zhang, Harsh Gupta, Kumar Suprabhat, and Leonardo Stella. 2023. A multi-agent reinforcement learning approach to promote cooperation in evolutionary games on networks with environmental feedback. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. 2196–2201.
- [59] Tuo Zhang, Leonardo Stella, and Julian Barreiro-Gomez. 2025. Experience-replay Innovative Dynamics. In *in Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*.
- [60] Zhengyuan Zhou, Panayotis Mertikopoulos, Aris L Moustakas, Nicholas Bambos, and Peter Glynn. 2017. Mirror descent learning in continuous games. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 5776–5783.