
Learning with Bandit Feedback in Potential Games

Johanne Cohen

LRI-CNRS, Université Paris-Sud, Université Paris-Saclay, France
johanne.cohen@lri.fr

Amélie Héliou

LIX, Ecole Polytechnique, CNRS, AMIBio, Inria, Université Paris-Saclay
amelie.heliou@polytechnique.edu

Panayotis Mertikopoulos

Univ. Grenoble Alpes, CNRS, Inria, LIG, F-38000, Grenoble, France
panayotis.mertikopoulos@imag.fr

Abstract

This paper examines the equilibrium convergence properties of no-regret learning with exponential weights in potential games. To establish convergence with minimal information requirements on the players' side, we focus on two frameworks: the *semi-bandit case* (where players have access to a noisy estimate of their payoff vectors, including strategies they did not play), and the *bandit case* (where players are only able to observe their in-game, realized payoffs). In the semi-bandit case, we show that the induced sequence of play converges almost surely to a Nash equilibrium at a quasi-exponential rate. In the bandit case, the same result holds for ε -approximations of Nash equilibria if we introduce an exploration factor $\varepsilon > 0$ that guarantees that action choice probabilities never fall below ε . In particular, if the algorithm is run with a suitably decreasing exploration factor, the sequence of play converges to a bona fide Nash equilibrium with probability 1.

1 Introduction

Given the manifest complexity of computing Nash equilibria, a central question that arises is whether such outcomes could result from a dynamic process in which players act on empirical information on their strategies' performance over time. This question becomes particularly important when the players' view of the game is obstructed by situational uncertainty and the "fog of war": for instance, when deciding which route to take to work each morning, a commuter is typically unaware of how many other commuters there are at any given moment, what their possible strategies are, how to best respond to their choices, etc. In fact, in situations of this kind, players may not even know that they are involved in a game; as such, it does not seem reasonable to assume full rationality, common knowledge of rationality, flawless execution, etc. to justify the Nash equilibrium prediction.

A compelling alternative to this "rationalistic" viewpoint is provided by the framework of *online learning*, where players are treated as oblivious entities facing a repeated decision process with a priori unknown rules and outcomes. In this context, when the players have no Bayesian prior on their environment, the most widely used performance criterion is that of *regret minimization*, a worst-case guarantee that was first introduced by Hannan [1], and which has given rise to a vigorous literature at the interface of optimization, statistics and theoretical computer science – for a survey, see [2, 3]. By this token, our starting point in this paper is the following question:

*If all players of a repeated game follow a no-regret algorithm,
does the induced sequence of play converge to Nash equilibrium?*

For concreteness, we focus on the *exponential weights* (EW) scheme [4–7], one of the most popular and widely studied algorithms for no-regret learning. In a nutshell, the main idea of the method is that the optimizing agent tallies the cumulative payoffs of each action and then employs a pure strategy with probability proportional to the exponential of these cumulative “scores”. Under this scheme, players are guaranteed a universal, min-max $\mathcal{O}(T^{1/2})$ regret bound (with T denoting the horizon of play), and their empirical frequency of play is known to converge to the game’s set of *coarse correlated equilibria* (CCE) [8].

In this way, no-regret learning would seem to provide a positive partial answer to our original question: coarse correlated equilibria are indeed learnable if all players follow an exponential weights learning scheme. On the flip side however, the set of coarse correlated equilibria may contain highly non-rationalizable strategies, so the end prediction of empirical convergence to such equilibria is fairly lax. For instance, in a recent paper, Viossat and Zapechelnyuk constructed a 4×4 variant of Rock-Paper-Scissors with a coarse correlated equilibrium that assigns positive weight *only* on strictly dominated strategies [9]. Even more recently, [10] showed that the mean dynamics of the exponential weights method (and, more generally, any method “following the regularized leader”) may cycle in perpetuity in zero-sum games, precluding any possibility of convergence to equilibrium in this case. Thus, in view of these negative results, a more calibrated answer to the above question is “not always”: especially when the issue at hand is convergence to a Nash equilibrium (as opposed to coarser notions), “no regret” is a rather loose guarantee.

Paper outline and summary of results. To address the above limitations, we focus on two issues:

- a) Convergence to Nash equilibrium (as opposed to correlated equilibria, coarse or otherwise).
- b) The convergence of the actual sequence of play (as opposed to empirical frequencies).

The reason for focusing on the actual sequence of play is that time-averages provide a fairly weak convergence mode: a priori, a player could oscillate between non-equilibrium strategies with suboptimal payoffs, but time-averages might still converge to equilibrium. On the other hand, convergence of the actual sequence of play both implies empirical convergence and also guarantees that players will be playing a Nash equilibrium in the long run, so it is a much stronger notion.

To establish convergence, we focus throughout on the class of *potential games* [11] that has found widespread applications in theoretical computer science [12], transportation networks [13], wireless communications [14], biology [15], and many other fields. We then focus on two different feedback models: in the *semi-bandit framework* (Section 3), players are assumed to have some (possibly imperfect) estimate of their payoff vectors at each stage, including strategies that they did not play; in the full *bandit framework* (Section 4), this assumption is relaxed and players are only assumed to observe their realized, in-game payoff at each stage.

Starting with the semi-bandit case, our main result is that under fairly mild conditions for the errors affecting the players’ observations (zero-mean martingale noise with tame second-moment tails), learning with exponential weights converges to a Nash equilibrium of the game with probability 1 (or to an ε -equilibrium if the algorithm is implemented with a uniform exploration factor $\varepsilon > 0$).¹ We also show that this convergence occurs at a *quasi-exponential rate*, i.e. much faster than the algorithm’s $\mathcal{O}(\sqrt{T})$ regret minimization rate would suggest.

These conclusions also apply to the bandit framework when the algorithm is run with a positive exploration factor $\varepsilon > 0$. Thus, by choosing a sufficiently small exploration factor, the end state of the EW algorithm in potential games with bandit feedback is arbitrarily close to a Nash equilibrium.

On the other hand, extending the stochastic approximation and martingale limit arguments that underlie the bandit analysis to the $\varepsilon = 0$ case is not straightforward. However, by letting the exploration factor go to zero at a suitable rate (similar to the temperature parameter in simulated annealing schemes), we are able to recover convergence to the game’s exact Nash set (and not an approximation thereof). We find this property particularly appealing for practical applications because it shows that equilibrium can be achieved in a wide class of games with minimal information requirements.

¹Having a exploration factor $\varepsilon > 0$ simply means here that action selection probabilities never fall below ε .

Related work. No-regret learning has given rise to a vast corpus of literature in theoretical computer science and machine learning, and several well-known families of algorithms have been proposed for that purpose. The most popular of these methods is based on exponential/multiplicative weight update rules, and several variants of this general scheme have been studied under different names in the literature (Hedge, EXP3, etc.) [4–7].

When applied to games, the time-average of the resulting trajectory of play converges to equilibrium in two-player zero-sum games [6, 16, 17] and the players’ social welfare approaches an approximate optimum [18]. In a similar vein, focusing on the so-called “Hedge” variant of the multiplicative weights (MW) algorithm, Kleinberg et al. [19] proved that the dynamics’ long-term limit in load balancing games is exponentially better than the worst correlated equilibrium. The convergence rate to approximate efficiency and to coarse correlated equilibria was further improved by Syrgkanis et al. [20] for a wide class of N -player normal form games using a natural class of regularized learning algorithms. This result was then extended to a class of games known as *smooth games* [21] with good properties in terms of the game’s price of anarchy [22].

In the context of potential games, learning algorithms and dynamics have received significant attention and considerable efforts have been devoted to studying the long-term properties of the players’ actual sequence of play. To that end, Kleinberg et al. [23] showed that, after a polynomially small transient stage, players end up playing a pure equilibrium for a fraction of time that is arbitrarily close to 1 with probability also arbitrarily close to 1. Mehta et al. [24] obtained a stronger result for (generic) 2-player coordination games, showing that the multiplicative weights algorithm (a linearized variant of the EW algorithm) converges to a pure Nash equilibrium for all but a measure 0 of initial conditions. More recently, Palaiopanos et al. [25] showed that the MW update rule converges to equilibrium in potential games; however, if the EW algorithm is run with a constant step-size that is not small enough, the induced sequence of play may exhibit chaotic behavior, even in simple 2×2 games. On the other hand, if the same algorithm is run with a decreasing step-size, Krichene et al. [26] showed that play converges to Nash equilibrium in all nonatomic potential games with a convex potential (and hence, in all nonatomic congestion games).

In the above works, players are assumed to have full (though possibly imperfect) knowledge of their payoff vectors, including actions that were not chosen. Going beyond this semi-bandit framework, Coucheney et al. [27] showed that a “penalty-regulated” variant of the EW algorithm converges to ε -logit equilibria (and hence ε -approximate Nash equilibria) in congestion games with bandit feedback. As in [26], the results of Coucheney et al. [27] employ the powerful ordinary differential equation (ODE) method of Benaïm [28] which leverages the convergence of an underlying, continuous-time dynamical system to obtain convergence of the algorithm at hand. We also employ this method to compare the actual sequence of play to the replicator dynamics of evolutionary game theory [29]; however, finetuning the bias-variance trade-off that arises when estimating the payoff of actions that were not employed is a crucial difficulty in this case. Overcoming this hurdle is necessary when seeking convergence to actual Nash equilibria (as opposed to ε -approximations thereof), so a key contribution of our paper is an extension of Benaïm’s theory to account for estimators with (possibly) unbounded variance.

2 The setup

2.1 Game-theoretic preliminaries

An N -player *game in normal form* consists of a (finite) set of *players* $\mathcal{N} = \{1, \dots, N\}$, each with a finite set of *actions* (or *pure strategies*) \mathcal{A}_i . The preferences of the i -th player for one action over another are determined by an associated *payoff function* $u_i: \mathcal{A} \equiv \prod_i \mathcal{A}_i \rightarrow \mathbb{R}$ that maps the *profile* $(\alpha_i; \alpha_{-i})$ of all players’ actions to the player’s reward $u_i(\alpha_i; \alpha_{-i})$.² Putting all this together, a game will be denoted by the tuple $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$.

Players can also *mix* their strategies by playing probability distributions $x_i = (x_{i\alpha_i})_{\alpha_i \in \mathcal{A}_i} \in \Delta(\mathcal{A}_i)$ over their action sets \mathcal{A}_i . The resulting probability vector x_i is called a *mixed strategy* and we write $\mathcal{X}_i = \Delta(\mathcal{A}_i)$ for the mixed strategy space of player i . Aggregating over players, we also write $\mathcal{X} = \prod_i \mathcal{X}_i$ for the game’s *strategy space*, i.e. the space of all mixed strategy profiles $x = (x_i)_{i \in \mathcal{N}}$.

²In the above $(\alpha_i; \alpha_{-i})$ is shorthand for $(\alpha_1, \dots, \alpha_i, \dots, \alpha_N)$, used here to highlight the action of player i against that of all other players.

In this context (and in a slight abuse of notation), the expected payoff of the i -th player in the profile $x = (x_1, \dots, x_N)$ is

$$u_i(x) = \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N) x_{1\alpha_1} \cdots x_{N\alpha_N}. \quad (2.1)$$

To keep track of the payoff of each pure strategy, we also write $v_{i\alpha_i}(x) = u_i(\alpha_i; x_{-i})$ for the payoff of strategy $\alpha_i \in \mathcal{A}_i$ under the profile $x \in \mathcal{X}$ and

$$v_i(x) = (v_{i\alpha_i}(x))_{\alpha_i \in \mathcal{A}_i} \quad (2.2)$$

for the resulting *payoff vector* of player i . We thus have

$$u_i(x) = \langle v_i(x), x_i \rangle = \sum_{\alpha_i \in \mathcal{A}_i} x_{i\alpha_i} v_{i\alpha_i}(x), \quad (2.3)$$

where $\langle v, x \rangle \equiv v^\top x$ denotes the ordinary pairing between v and x .

The most widely used solution concept in game theory is that of a *Nash equilibrium* (NE), i.e. a state $x^* \in \mathcal{X}$ such that

$$u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_{-i}^*) \quad \text{for every deviation } x_i \in \mathcal{X}_i \text{ of player } i \text{ and all } i \in \mathcal{N}. \quad (\text{NE})$$

Equivalently, writing $\text{supp}(x_i) = \{\alpha_i \in \mathcal{A}_i : x_i > 0\}$ for the support of $x_i \in \mathcal{X}_i$, we have the characterization

$$v_{i\alpha_i}(x^*) \geq v_{i\beta_i}(x^*) \quad \text{for all } \alpha_i \in \text{supp}(x_i^*) \text{ and all } \beta_i \in \mathcal{A}_i, i \in \mathcal{N}. \quad (2.4)$$

A Nash equilibrium $x^* \in \mathcal{X}$ is further said to be *pure* if $\text{supp}(x_i^*) = \{\hat{\alpha}_i\}$ for some $\hat{\alpha}_i \in \mathcal{A}_i$ and all $i \in \mathcal{N}$. In *generic games* (that is, games where small changes to any payoff do not introduce new Nash equilibria or destroy existing ones), every pure Nash equilibrium is also *strict* in the sense that (2.4) holds as a strict inequality for all $\alpha_i \neq \hat{\alpha}_i$.

In our analysis, it will be important to consider the following relaxations of the notion of a Nash equilibrium: First, weakening the inequality (NE) leads to the notion of a δ -*equilibrium*, defined here as any mixed strategy profile $x^* \in \mathcal{X}$ such that

$$u_i(x_i^*; x_{-i}^*) + \delta \geq u_i(x_i; x_{-i}^*) \quad \text{for every deviation } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N}. \quad (\text{NE}_\delta)$$

Finally, we say that x^* is a *restricted equilibrium* (RE) of Γ if

$$v_{i\alpha_i}(x^*) \geq v_{i\beta_i}(x^*) \quad \text{for all } \alpha_i \in \text{supp}(x_i^*) \text{ and all } \beta_i \in \mathcal{A}'_i, i \in \mathcal{N}, \quad (\text{RE})$$

where \mathcal{A}'_i is some restricted subset of \mathcal{A}_i containing $\text{supp}(x_i^*)$. In words, restricted equilibria are Nash equilibria of Γ restricted to subgames where only a subset of the players' pure strategies are available at any given moment. Clearly, Nash equilibria are restricted equilibria but the converse does not hold: for instance, every pure strategy profile is a restricted equilibrium, but not necessarily a Nash equilibrium.

Throughout this paper, we will focus almost exclusively on the class of *potential games*, which have been studied extensively in the context of congestion, traffic networks, oligopolies, etc. Following Monderer and Shapley [11], Γ is a *potential game* if it admits a *potential function* $f: \prod_i \mathcal{A}_i \rightarrow \mathbb{R}$ such that

$$u_i(x_i; x_{-i}) - u_i(x'_i; x_{-i}) = f(x_i; x_{-i}) - f(x'_i; x_{-i}), \quad (2.5)$$

for all $x_i, x'_i \in \mathcal{X}_i$, $x_{-i} \in \mathcal{X}_{-i} \equiv \prod_{j \neq i} \mathcal{X}_j$, and all $i \in \mathcal{N}$. A simple differentiation of (2.1) then yields

$$v_i(x) = \nabla_{x_i} u_i(x) = \nabla_{x_i} f(x) \quad \text{for all } i \in \mathcal{N}. \quad (2.6)$$

Obviously, every local maximizer of f is a Nash equilibrium so potential games always admit Nash equilibria in pure strategies (which are also strict if the game is generic).

2.2 The exponential weights algorithm

Our basic learning framework is as follows: At each stage $n = 1, 2, \dots$, all players $i \in \mathcal{N}$ select an action $\alpha_i(n) \in \mathcal{A}_i$ based on their mixed strategies; subsequently, they receive some feedback on their chosen actions, they update their mixed strategies, and the process repeats.

A popular (and very widely studied) class of algorithms for no-regret learning in this setting is the *exponential weights* (EW) scheme introduced by Vovk [4] and studied further by Auer et al. [5], Freund and Schapire [6], Arora et al. [7], and many others. Somewhat informally, the main idea is that each player tallies the cumulative payoffs of each of their actions, and then employs a pure strategy $\alpha_i \in \mathcal{A}_i$ with probability roughly proportional to these cumulative payoff “scores”. Focusing on the so-called “ ε -HEDGE” variant of the EW algorithm [6], this process can be described in pseudocode form as follows:

Algorithm 1 ε -HEDGE with generic feedback

Require: step-size sequence $\gamma_n > 0$, exploration factor $\varepsilon \in [0, 1]$, initial scores $Y_i \in \mathbb{R}^{\mathcal{A}_i}$.

```

1: for  $n = 1, 2, \dots$  do
2:   for every player  $i \in \mathcal{N}$  do
3:     set mixed strategy:  $X_i \leftarrow \varepsilon \text{unif}_i + (1 - \varepsilon) \Lambda_i(Y_i)$ ;
4:     choose action  $\alpha_i \sim X_i$ ;
5:     acquire estimate  $\hat{v}_i$  of realized payoff vector  $v_i(\alpha_i; \alpha_{-i})$ ;
6:     update scores:  $Y_i \leftarrow Y_i + \gamma_n \hat{v}_i$ ;
7:   end for
8: end for

```

Mathematically, [Algorithm 1](#) represents the recursion

$$\begin{aligned} X_i(n) &= \varepsilon \text{unif}_i + (1 - \varepsilon) \Lambda_i(Y_i(n)), \\ Y_i(n+1) &= Y_i(n) + \gamma_{n+1} \hat{v}_i(n+1), \end{aligned} \tag{\varepsilon-Hedge}$$

where

$$\text{unif}_i = \frac{1}{|\mathcal{A}_i|} (1, \dots, 1) \tag{2.7}$$

stands for the uniform distribution over \mathcal{A}_i and $\Lambda_i: \mathbb{R}^{\mathcal{A}_i} \rightarrow \mathcal{X}_i$ denotes the *logit choice map*

$$\Lambda_i(y_i) = \frac{(\exp(y_i \alpha_i))_{\alpha_i \in \mathcal{A}_i}}{\sum_{\alpha_i \in \mathcal{A}_i} \exp(y_i \alpha_i)}, \tag{2.8}$$

which assigns exponentially higher probability to pure strategies with higher scores. Thus, action selection probabilities under (ε -Hedge) are a convex combination of uniform exploration (with total weight ε) and exponential weights (with total weight $1 - \varepsilon$).³ As a result, for $\varepsilon \approx 1$, action selection is essentially uniform; at the other extreme, when $\varepsilon = 0$, we obtain the original Hedge algorithm of Freund and Schapire [6] with feedback sequence $\hat{v}(n)$ and no explicit exploration.

The no-regret properties of (ε -Hedge) have been extensively studied in the literature as a function of the algorithm’s step-size sequence γ_n , exploration factor ε , and the statistical properties of the payoff estimates $\hat{v}(n)$ – for a survey, we refer the reader to [2, 3]. In our convergence analysis, we examine the role of each of these factors in detail, focusing in particular on the distinction between “*semi-bandit feedback*” (when it is possible to estimate the payoff of pure strategies that were not played) and “*bandit feedback*” (when players only observe the payoff of their chosen action).

3 Learning with semi-bandit feedback

3.1 The model

We begin with the *semi-bandit framework*, i.e. the case where each player has access to a possibly imperfect estimate of their entire payoff vector at stage n . More precisely, we assume here that the feedback sequence $\hat{v}_i(n)$ to [Algorithm 1](#) is of the general form

$$\hat{v}_i(n) = v_i(\alpha_i(n); \alpha_{-i}(n)) + \xi_i(n), \tag{3.1}$$

where $(\xi_i(n))_{i \in \mathcal{N}}$ is a martingale noise process representing the players’ estimation error and satisfying the following statistical hypotheses:

³Of course, the exploration factor ε could also be player-dependent. For simplicity, we state all our results here with the same ε for all players.

1. *Zero-mean:*

$$\mathbb{E}[\xi_i(n) | \mathcal{F}_{n-1}] = 0 \quad \text{for all } n = 1, 2, \dots \text{ (a.s.)} \quad (\text{H1})$$

2. *Tame tails:*

$$\mathbb{P}(\|\xi_i(n)\|_\infty^2 \geq z | \mathcal{F}_{n-1}) \leq A/z^q \quad \text{for some } q > 2, A > 0, \text{ and all } n = 1, 2, \dots \text{ (a.s.)} \quad (\text{H2})$$

In the above, the expectation $\mathbb{E}[\cdot]$ is taken with respect to some underlying filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n \in \mathbb{N}}, \mathbb{P})$ which serves as a stochastic basis for the process $(\alpha(n), \hat{v}(n), Y(n), X(n))_{n \geq 1}$.⁴ In words, Hypothesis (H1) simply means that the players' feedback sequence $\hat{v}(n)$ is *conditionally unbiased* with respect to the history of play, i.e.

$$\mathbb{E}[\hat{v}_i(n) | \mathcal{F}_{n-1}] = v_i(X(n-1)), \quad \text{for all } n = 1, 2, \dots \text{ (a.s.)} \quad (\text{3.2a})$$

Hypothesis (H2) further implies that the variance of the estimator \hat{v} is conditionally bounded, i.e.

$$\text{Var}[\hat{v}(n) | \mathcal{F}_{n-1}] \leq \sigma^2 \quad \text{for all } n = 1, 2, \dots \text{ (a.s.)} \quad (\text{3.2b})$$

By Chebyshev's inequality, an estimator with finite variance enjoys the tail bound $\mathbb{P}(\|\xi_i(n)\|_\infty \geq z | \mathcal{F}_{n-1}) = \mathcal{O}(1/z^2)$. At the expense of working with slightly more conservative step-size policies (see below), much of our analysis goes through with this weaker requirement for the tails of ξ . However, the extra control provided by the $\mathcal{O}(1/z^q)$ tail bound simplifies the presentation considerably, so we do not consider this relaxation here. In any event, Hypothesis (H2) is satisfied by a broad range of error noise distributions (including all compactly supported, sub-Gaussian and sub-exponential distributions), so the loss in generality is small compared to the gain in clarity and concision.

3.2 Convergence analysis

With all this at hand, our main result for the convergence of (ε -Hedge) with semi-bandit feedback of the form (3.1) is as follows:

Theorem 1. *Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative exploration factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. Then:*

1. $X(n)$ converges (a.s.) to a δ -equilibrium of Γ with $\delta \equiv \delta(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.
2. If $\lim_{n \rightarrow \infty} X(n)$ is an ε -pure state of the form $x_i^* = \varepsilon \text{unif}_i + (1 - \varepsilon)e_{\hat{\alpha}_i}$ for some $\hat{\alpha} \in \mathcal{A}$, then $\hat{\alpha}$ is a.s. a strict equilibrium of Γ and convergence occurs at a quasi-exponential rate:

$$X_{i\hat{\alpha}_i}(n) \geq 1 - \varepsilon - be^{-c \sum_{k=1}^n \gamma_k} \quad \text{for some positive } b, c > 0. \quad (\text{3.3})$$

Corollary 2. *If Algorithm 1 is run with assumptions as above and no exploration ($\varepsilon = 0$), $X(n)$ converges to a Nash equilibrium with probability 1. Moreover, if the limit of $X(n)$ is pure and $\beta < 1$, we have*

$$X_{i\hat{\alpha}_i}(n) \geq 1 - be^{-cn^{1-\beta}} \quad \text{for some positive } b, c > 0. \quad (\text{3.4})$$

Sketch of the proof. The proof of Theorem 1 is fairly convoluted, so we relegate the details to the paper's technical appendix and only present here a short sketch thereof.

Our main tool is the so-called *ordinary differential equation* (ODE) method, a powerful stochastic approximation scheme due to Benaïm and Hirsch [28, 30]. The key observation is that the mixed strategy sequence $X(n)$ generated by Algorithm 1 can be viewed as a ‘‘Robbins–Monro approximation’’ (an *asymptotic pseudotrajectory* to be precise) of the ε -perturbed exponential learning dynamics

$$\begin{aligned} \dot{y}_i &= v_i(x), \\ x_i &= \varepsilon \text{unif}_i + (1 - \varepsilon) \Lambda_i(y_i), \end{aligned} \quad (\text{XL}_\varepsilon)$$

By differentiating, it follows that $x_i(t)$ evolves according to the ε -perturbed replicator dynamics

$$\dot{x}_{i\alpha} = (x_{i\alpha} - |\mathcal{A}_i|^{-1}\varepsilon) \left[v_{i\alpha}(x) - (1 - \varepsilon)^{-1} \sum_{\beta \in \mathcal{A}_i} (x_{i\beta} - |\mathcal{A}_i|^{-1}\varepsilon) v_{i\beta}(x) \right], \quad (\text{RD}_\varepsilon)$$

⁴Notation-wise, this means that the players' actions at stage n are drawn based on their mixed strategies at stage $n - 1$. This slight discrepancy with the pseudocode representation of Algorithm 1 is only done to simplify notation later on.

which, for $\varepsilon = 0$, boil down to the ordinary *replicator dynamics* of Taylor and Jonker [29]:

$$\dot{x}_{i\alpha} = x_{i\alpha}[v_{i\alpha}(x) - \langle v_i(x), x_i \rangle], \quad (\text{RD})$$

A key property of the replicator dynamics that readily extends to the ε -perturbed variant (RD_ε) is that the game's potential f is a strict Lyapunov function – i.e. $f(x(t))$ is increasing under (RD_ε) unless $x(t)$ is stationary. By a standard result of Benaïm [28], this implies that the discrete-time process $X(n)$ converges (a.s.) to a connected set of rest points of (RD_ε), which are themselves approximate restricted equilibria of Γ .

Of course, since every ε -pure point of the form $(\varepsilon \text{unif}_i + (1 - \varepsilon)e_{\alpha_i})_{i \in \mathcal{N}}$ is also stationary under (RD_ε), the above does not imply that the limit of $X(n)$ is an approximate *equilibrium* of Γ . To rule out non-equilibrium outcomes, we first note that the set of rest points of (RD_ε) is finite (by genericity), so $X(n)$ must converge to a point. Then, the final step of our convergence proof is provided by a martingale recurrence argument which shows that when $X(n)$ converges to a point, this limit must be an approximate equilibrium of Γ . Finally, the rate of convergence (3.3) is obtained by comparing the payoff of a player's equilibrium strategy to that of the player's other strategies, and then “inverting” the logit choice map to translate this into an exponential decay rate for $\|X_{i\hat{\alpha}_i}(n) - x^*\|$. \square

We close this section with two remarks on [Theorem 1](#). First, we note that there is an inverse relationship between the tail exponent q in (H2) and the decay rate β of the algorithm's step-size sequence $\gamma_n \propto n^{-\beta}$. Specifically, higher values of q imply that the noise in the players' observations is smaller (on average and with high probability), so players can be more aggressive in their choice of step-size. This is reflected in the lower bound $1/q$ for β and the fact that the players' rate of convergence to Nash equilibrium increases for smaller β ; in particular, (3.3) shows that [Algorithm 1](#) enjoys a convergence bound which is just shy of $\mathcal{O}(\exp(-n^{1-1/q}))$. Thus, if the noise process ξ is sub-Gaussian/sub-exponential (so q can be taken arbitrarily large), a near-constant step-size sequence (small β) yields an almost linear convergence rate.

Second, if the noise process ξ is “isotropic” in the sense of Benaïm [28, Thm. 9.1], the instability of non-pure Nash equilibria under the replicator dynamics can be used to show that the limit of $X(n)$ is pure with probability 1.⁵ When this is the case, the quasi-exponential convergence rate (3.3) becomes universal in that it holds with probability 1 (as opposed to conditioning on $\lim_{n \rightarrow \infty} X(n)$ being pure). We find this property particularly appealing for practical applications because it shows that equilibrium is reached *exponentially faster* than the $\mathcal{O}(1/\sqrt{n})$ worst-case regret bound of (ε -Hedge) would suggest.

4 Payoff-based learning: the bandit case

We now turn to the *bandit framework*, a minimal-information setting where, at each stage of the process, players only observe their realized payoffs

$$\hat{u}_i(n) = u_i(\alpha_i(n); \alpha_{-i}(n)). \quad (4.1)$$

In this case, players have no clue about the payoffs of strategies that were not chosen, so they must *construct* an estimator for their payoff vector, including its missing components. A standard way to do this is via the *bandit estimator*

$$\hat{v}_{i\alpha_i}(n) = \frac{\mathbb{1}(\alpha_i(n) = \alpha_i)}{\mathbb{P}(\alpha_i(n) = \alpha_i | \mathcal{F}_{n-1})} \cdot \hat{u}_i(n) = \begin{cases} \hat{u}_i(n)/X_{i\alpha_i}(n-1) & \text{if } \alpha_i = \alpha_i(n), \\ 0 & \text{otherwise.} \end{cases} \quad (4.2)$$

Indeed, a straightforward calculation shows that

$$\begin{aligned} \mathbb{E}[\hat{v}_{i\alpha_i}(n) | \mathcal{F}_{n-1}] &= \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} X_{-i, \alpha_{-i}}(n-1) \sum_{\beta_i \in \mathcal{A}_i} X_{i\beta_i}(n-1) \frac{\mathbb{1}(\alpha_i = \beta_i)}{X_{i\alpha_i}(n-1)} u_i(\beta_i; \alpha_{-i}) \\ &= u_i(\alpha_i; X_{-i}(n-1)) \\ &= v_{i\alpha_i}(X(n-1)), \end{aligned} \quad (4.3)$$

⁵Specifically, we refer here to the so-called “folk theorem” of evolutionary game theory which states that x^* is asymptotically stable under (RD) if and only if it is a strict Nash equilibrium of Γ [15]. The extension of this result to the ε -replicator system (RD_ε) is immediate.

so the estimator (4.2) is unbiased in the sense of (H1)/(3.2a). On the other hand, a similar calculation shows that the variance of $\hat{v}_{i\alpha_i}(n)$ grows as $\mathcal{O}(1/X_{i\alpha_i}(n-1))$, implying that (H2)/(3.2b) may fail to hold if the players' action selection probabilities become arbitrarily small.

Importantly, this can never happen if (ε -Hedge) is run with a strictly positive exploration factor $\varepsilon > 0$. In that case, we can show that the bandit estimator (4.2) satisfies both (H1) and (H2), leading to the following result:

Theorem 3. *Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) the bandit estimator (4.2); ii) a strictly positive exploration factor $\varepsilon > 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (0, 1]$. Then:*

1. $X(n)$ converges (a.s.) to a δ -equilibrium of Γ with $\delta \equiv \delta(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.
2. If $\lim_{n \rightarrow \infty} X(n)$ is an ε -pure state of the form $x_i^* = \varepsilon \text{unif}_i + (1 - \varepsilon)e_{\hat{\alpha}_i}$ for some $\hat{\alpha} \in \mathcal{A}$, then $\hat{\alpha}$ is a.s. a strict equilibrium of Γ and convergence occurs at a quasi-exponential rate:

$$X_{i\hat{\alpha}_i}(n) \geq 1 - \varepsilon - be^{-c \sum_{k=1}^n \gamma^k} \quad \text{for some positive } b, c > 0. \quad (4.4)$$

Proof. Under Algorithm 1, the estimator (4.2) gives

$$\|\hat{v}_i(n)\| = \frac{|\hat{u}_i(n)|}{X_{i\alpha_i(n-1)}(n)} \leq \frac{|u_i(\alpha_i(n); \alpha_{-i}(n))|}{\varepsilon} \leq \frac{u_{\max}}{\varepsilon}, \quad (4.5)$$

where $u_{\max} = \max_{i \in \mathcal{N}} \max_{\alpha_1 \in \mathcal{A}_1} \cdots \max_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N)$ denotes the absolute maximum payoff in Γ . This implies that (H2) holds true for all $q > 2$, so our claim follows from Theorem 1. \square

Theorem 3 shows that the limit of Algorithm 1 is closer to the Nash set of the game if the exploration factor ε is taken as small as possible. On the other hand, the crucial limitation of this result is that it does not apply to the case $\varepsilon = 0$ which corresponds to the game's bona fide Nash equilibria. As we discussed above, the reason for this is that the variance of $\hat{v}(n)$ may grow without bound if action choice probabilities become arbitrarily small, in which case the main components of our proof break down.

With this ‘‘bias-variance’’ trade-off in mind, we introduce below a modified version of Algorithm 1 with an ‘‘annealing’’ schedule for the method's exploration factor:

Algorithm 2 Exponential weights with annealing

Require: step-size sequence $\gamma_n > 0$, vanishing exploration factor $\varepsilon_n > 0$, initial scores $Y_i \in \mathbb{R}^{\mathcal{A}_i}$

- 1: **for** $n = 1, 2, \dots$ **do**
 - 2: **for** every player $i \in \mathcal{N}$ **do**
 - 3: set mixed strategy: $X_i \leftarrow \varepsilon_n \text{unif}_i + (1 - \varepsilon_n) \Lambda_i(Y_i)$;
 - 4: choose action $\alpha_i \sim X_i$ and receive payoff $\hat{u}_i \leftarrow u_i(\alpha_i; \alpha_{-i})$;
 - 5: set $\hat{v}_{i\alpha_i} \leftarrow \hat{u}_i / X_{i\alpha_i}$ and $\hat{v}_{i\beta_i} \leftarrow 0$ for $\beta_i \neq \alpha_i$;
 - 6: update scores: $Y_i \leftarrow Y_i + \gamma_n \hat{v}_i$;
 - 7: **end for**
 - 8: **end for**
-

Of course, the convergence of Algorithm 2 depends heavily on the rate at which ε_n decays to 0 relative to the algorithm's step-size sequence γ_n . This can be seen clearly in our next result:

Theorem 4. *Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) the bandit estimator (4.2); ii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/2, 1]$; and iii) a decreasing exploration factor $\varepsilon_n \downarrow 0$ such that*

$$\lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_n^2} = 0, \quad \sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_n} < \infty, \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{\varepsilon_n - \varepsilon_{n+1}}{\gamma_n^2} = 0. \quad (4.6)$$

Then, $X(n)$ converges (a.s.) to a Nash equilibrium of Γ .

The main challenge in proving [Theorem 4](#) is that, unless the “innovation term” $U_i(n) = \hat{v}_i(n) - v_i(X(n-1))$ has bounded variance, Benaïm’s general theory does not imply that $X(n)$ forms an asymptotic pseudotrajectory of the underlying mean dynamics – here, the unperturbed replicator system (RD). Nevertheless, under the summability condition (4.6), it is possible to show that this is the case by using a martingale limit argument based on Burkholder’s inequality. Furthermore, under the stated conditions, it is also possible to show that, if $X(n)$ converges, its limit is necessarily a Nash equilibrium of Γ . Our proof then follows in roughly the same way as in the case of [Theorem 1](#); for the details, we refer the reader to the appendix.

We close this section by noting that the summability condition (4.6) imposes a lower bound on the step-size exponent β that is different from the lower bound in [Theorem 3](#). In particular, if $\beta = 1/2$, (4.6) cannot hold for any vanishing sequence of exploration factors $\varepsilon_n \downarrow 0$. Given that the innovation term U_i is bounded, we conjecture that this sufficient condition is not tight and can be relaxed further. We intend to address this issue in future work.

5 Conclusion and perspectives

The results of the previous sections show that no-regret learning via exponential weights enjoys appealing convergence properties in generic potential games. Specifically, in the semi-bandit case, the sequence of play converges to a Nash equilibrium with probability 1, and convergence to pure equilibria occurs at a quasi-exponential rate. In the bandit case, the same holds true for $\mathcal{O}(\varepsilon)$ -equilibria if the algorithm is run with a positive mixing factor $\varepsilon > 0$; and if the algorithm is run with a decreasing mixing schedule, the sequence of play converges to an actual Nash equilibrium (again, with probability 1). In future work, we intend to examine the algorithm’s convergence properties in other classes of games (such as smooth games), extend our analysis to the general “follow the regularized leader” (FTRL) class of policies (of which EW is a special case), and to examine the impact of asynchronicities and delays in the players’ feedback/update cycles.

Acknowledgments

Johanne Cohen was partially supported by the grant CNRS PEPS MASTODONS project ADOC 2017. Amélie Héliou and Panayotis Mertikopoulos gratefully acknowledge financial support from the Huawei Innovation Research Program ULTRON and the ANR JCJC project ORACLESS (grant no. ANR-16-CE33-0004-01).

References

- [1] James Hannan. Approximation to Bayes risk in repeated play. In Melvin Dresher, Albert William Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games, Volume III*, volume 39 of *Annals of Mathematics Studies*, pages 97–139. Princeton University Press, Princeton, NJ, 1957.
- [2] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [3] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [4] Volodimir G. Vovk. Aggregating strategies. In *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, pages 371–383, 1990.
- [5] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995.
- [6] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.
- [7] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [8] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, September 2000.
- [9] Yannick Viossat and Andriy Zapechelnyuk. No-regret dynamics and fictitious play. *Journal of Economic Theory*, 148(2):825–842, March 2013.

- [10] Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *SODA '18: Proceedings of the 29th annual ACM-SIAM symposium on discrete algorithms*, to appear.
- [11] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [12] Noam Nisan, Tim Roughgarden, Eva Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [13] William H. Sandholm. *Population Games and Evolutionary Dynamics*. Economic learning and social evolution. MIT Press, Cambridge, MA, 2010.
- [14] Samson Lasaulce and Hamidou Tembine. *Game Theory and Learning for Wireless Networks: Fundamentals and Applications*. Academic Press, Elsevier, 2010.
- [15] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.
- [16] Dean Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40–55, October 1997.
- [17] Avrim Blum and Yishay Mansour. Learning, regret minimization, and equilibria. In Noam Nisan, Tim Roughgarden, Eva Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007.
- [18] Avrim Blum, Mohammad Taghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *STOC '08: Proceedings of the 40th annual ACM symposium on the Theory of Computing*, pages 373–382. ACM, 2008.
- [19] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, 2011.
- [20] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, pages 2989–2997, 2015.
- [21] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)*, 62(5):32, 2015.
- [22] Dylan J Foster, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems*, pages 4727–4735, 2016.
- [23] Robert Kleinberg, Georgios Piliouras, and Eva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 533–542. ACM, 2009.
- [24] Ruta Mehta, Ioannis Panageas, and Georgios Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics. In *ITCS '15: Proceedings of the 6th Conference on Innovations in Theoretical Computer Science*, 2015.
- [25] Gerasimos Palaiopoulos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [26] Walid Krichene, Benjamin Drighès, and Alexandre M. Bayen. Online learning of Nash equilibria in congestion games. *SIAM Journal on Control and Optimization*, 53(2):1056–1081, 2015.
- [27] Pierre Coucheney, Bruno Gaujal, and Panayotis Mertikopoulos. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, August 2015.
- [28] Michel Benaïm. Dynamics of stochastic approximation algorithms. *Séminaire de probabilités de Strasbourg*, 33, 1999.
- [29] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.
- [30] Michel Benaïm and Morris W. Hirsch. Asymptotic pseudotrajectories and chain recurrent flows, with applications. *Journal of Dynamics and Differential Equations*, 8(1):141–176, 1996.
- [31] P. Hall and C. C. Heyde. *Martingale Limit Theory and Its Application*. Probability and Mathematical Statistics. Academic Press, New York, 1980.

A Learning with semi-bandit feedback

This section is devoted to the proof of Theorem 1. In Section 3, the recursion equation of Algorithm 1 depends on ε parameter. In this section, each player i can have its own mixing factor ε_i . This choice is done in order to simplify the notation in Sections 3 and 4. Thus the recursion equation of Algorithm 1 can be rewritten as

$$\begin{aligned} X_i(n) &= \varepsilon_i/|\mathcal{A}_i| \mathbf{1} + (1 - \varepsilon_i) \Lambda_i(Y_i(n)), \\ Y_i(n+1) &= Y_i(n) + \gamma_{n+1} \hat{v}_i(n+1), \end{aligned} \quad (\varepsilon\text{-Hedge})$$

To perform this, we introduce some other discrete-time process $(X_i^\Lambda(n))_{n \in \mathbb{N}}$ to help with this proof.

$$X_i^\Lambda(n) = \frac{1}{(1 - \varepsilon_i)} (X_i(n) - \varepsilon_i/|\mathcal{A}_i| \mathbf{1})$$

By definition, we have $X_i(n) = \varepsilon_i/|\mathcal{A}_i| \mathbf{1} + (1 - \varepsilon_i) X_i^\Lambda(n)$. Besides,

Observation 5. For any $n \in \mathbb{N}$, $\|X(n) - X^\Lambda(n)\|_2 \leq N\epsilon$, where $\epsilon = \max_{i \in \mathbb{N}} \frac{\epsilon_i}{\sqrt{|\mathcal{A}_i|}}$.

Moreover, we can establish an other relation between these two discrete-time processes:

Lemma 6. If $(X^\Lambda(n))_{n \in \mathbb{N}}$ converges to a Nash equilibrium x^* , then $(X(n))_{n \in \mathbb{N}}$, defined in (ε -Hedge), converges to a $\delta(\epsilon)$ -Nash equilibrium with $\epsilon = \max_{i \in \mathcal{N}} \frac{\epsilon_i}{|\mathcal{A}_i|}$ and $\delta(\epsilon)$ such that if ϵ tends to 0, then $\delta(\epsilon)$ also tends to 0.

Proof. If $(X^\Lambda(n))_{n \in \mathbb{N}}$ converges to a point x^* , then $(X(n))_{n \in \mathbb{N}}$ also converges to a limit \hat{x} . Observation 5 gives us that $\|X(n) - X^\Lambda(n)\|_2 \leq N\epsilon$, by continuity, $\|x^* - \hat{x}\|_2 \leq N\epsilon$.

It remains to show that if the two mixed strategy profiles p and p' are such that $\|p - p'\|_2 \leq \epsilon$ then $\|u_i(p) - u_i(p')\|_2 \leq \frac{\delta(\epsilon)}{2}$, for any player $i \in \mathcal{N}$.

$$\begin{aligned} |u_i(p) - u_i(p')| &= \left| \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N) (p_{1\alpha_1} \cdots p_{N\alpha_N} - p'_{1\alpha_1} \cdots p'_{N\alpha_N}) \right| \\ &\leq \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N) |p_{1\alpha_1} \cdots p_{N\alpha_N} - p'_{1\alpha_1} \cdots p'_{N\alpha_N}| \\ &\leq \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N) |(p_{1\alpha_1} \cdots p_{N\alpha_N} - (p_{1\alpha_1} - \epsilon) \cdots (p_{N\alpha_N} - \epsilon))| \\ &\leq \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} u_i(\alpha_1, \dots, \alpha_N) \sum_{k=1}^N (-\epsilon)^k \binom{N}{k} \\ &= \frac{\delta(\epsilon)}{2} \end{aligned} \quad (\text{A.1})$$

Moreover, observe that if $\epsilon \rightarrow 0$, then $\frac{\delta(\epsilon)}{2} \rightarrow 0$, and $|u_i(p) - u_i(p')| \rightarrow 0$.

Assume that $(X^\Lambda(n))_{n \in \mathbb{N}}$ converges to a Nash equilibrium x^* . By definition, for all $x_i \in X_i$, $i \in \mathbb{N}$ $u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_{-i}^*)$. In addition, we have $\|x^* - \hat{x}\|_2 \leq N\epsilon$ and $\|(x_i, x_{-i}^*) - (x_i, \hat{x}_{-i})\|_2 \leq N\epsilon \quad \forall x_i \in X_i$.

$$\begin{aligned} u_i(\hat{x}_i; \hat{x}_{-i}) + \frac{\delta(N\epsilon)}{2} &\geq u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_{-i}^*) \geq u_i(x_i; \hat{x}_{-i}) - \frac{\delta(N\epsilon)}{2} \\ u_i(\hat{x}_i; \hat{x}_{-i}) + \delta(N\epsilon) &\geq u_i(x_i; \hat{x}_{-i}) \text{ for all } x_i^* \in \mathcal{X}_i, i \in \mathcal{N} \end{aligned} \quad (\text{A.2})$$

So, the latter equation corresponds to the definition of $\delta(\epsilon)$ -equilibrium. We can conclude that $(X(n))_{n \in \mathbb{N}}$ converges to a $\delta(\epsilon)$ -Nash equilibrium with $\delta(\epsilon)$. \square

The remainder of this section is devoted to prove the convergence results of ε -HEDGE Algorithm, and it is split into two parts according to the feedback assumption. The proof is based on Benaim's study of stochastic approximations [28]. We follow 4 points:

1. Show that X is an asymptotic pseudo trajectory of a continuous dynamics,
2. Show that the potential function of the game is strict Lyapunov function of the dynamics,
3. Show that X converges toward a rest point of the dynamics,
4. Show that if X converges toward a point it is a Nash Equilibrium.

For the first step, we focus on the sequences $(X_i(n))_{n \in \mathbb{N}}$ and its linear interpolation $(x_i(t))$. We also define the continuous variables, $x_{i\alpha}^\Lambda(t) = \frac{x_{i\alpha}(t) - \varepsilon_i / |\mathcal{A}_i|}{1 - \varepsilon_i}$ and $y_{i\alpha}(t)$ such that $x_{i\alpha}^\Lambda(t) = \Lambda_i(y_{i\alpha}(t))$, with the same relations as the corresponding discret processes. We will prove that the sequences $(X_i(n))_{n \in \mathbb{N}}$ correspond to some classical family of stochastic approximation algorithms so-called *Approximate Robbins-Monro algorithm*.

Definition 7 (Approximate Robbins-Monro conditions). The general stochastic approximation algorithm $x(n+1) = x(n) + \gamma_n(F(x(n)) + U_n + \beta_n)$ is said to be an *approximate Robbins-Monro algorithm* if:

- $F : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a continuous map;
- $U_n \in \mathbb{R}^m$ are perturbations and U_n is a martingale difference noise;
- $\{\gamma_n\}_{n \geq 1}$ is a given sequence of nonnegative numbers such that $\sum_k \gamma_k = \infty$ and $\lim_{n \rightarrow \infty} \gamma_n = 0$;
- $\lim_{n \rightarrow \infty} b_n = 0$ almost surely.

Then, we will prove that the interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ is an asymptotic pseudo trajectory of the solutions of the following ordinary differential equation.

$$\begin{aligned} \dot{x}_{i\alpha} &= (x_{i\alpha} - |\mathcal{A}_i|^{-1} \varepsilon_i) \left[v_{i\alpha}(x) - (1 - \varepsilon_i)^{-1} \sum_{\beta \in \mathcal{A}_i} (x_{i\beta} - |\mathcal{A}_i|^{-1} \varepsilon_i) v_{i\beta}(x) \right] \\ &= (1 - \varepsilon_i) x_{i\alpha}^\Lambda \left[v_{i\alpha}(x) - \sum_{\beta \in \mathcal{A}_i} x_{i\beta}^\Lambda v_{i\beta}(x) \right] \end{aligned} \quad (\text{A.3})$$

To proceed, recall the definition of the asymptotic pseudo-trajectory.

Definition 8 (Asymptotic Pseudo-trajectories). Given a flow $\phi : \mathbb{R} \times M \rightarrow M$, $(n, x) \rightarrow \phi(n, x) = \phi_n(x)$ such that $\phi_0 = \text{Identity}$ and $\phi_{n+\alpha} = \phi_n \circ \phi_\alpha$, a continuous function $X : \mathbb{R} \rightarrow M$ is an asymptotic pseudo-trajectory if

$$\lim_{n \rightarrow \infty} \sup_{0 \leq k \leq T} d((X(n+k), \phi_n(X(n))), 0) = 0 \text{ for any } T > 0 \quad (\text{A.4})$$

Proposition 9. Suppose that *Algorithm 1* is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative mixing factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. The interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ is an asymptotic pseudo trajectory of the solutions of the following ordinary differential equation

$$\dot{x}_{i\alpha} = (x_{i\alpha} - |\mathcal{A}_i|^{-1} \varepsilon_i) \left[v_{i\alpha}(x) - (1 - \varepsilon_i)^{-1} \sum_{\beta \in \mathcal{A}_i} (x_{i\beta} - |\mathcal{A}_i|^{-1} \varepsilon_i) v_{i\beta}(x) \right], \quad (\text{RD}_\varepsilon)$$

Proof. The proof of Proposition 9 can be split into two parts. First, we will prove that the stochastic process $\{X_i(n)\}_{n \in \mathbb{N}}$ given by *Algorithm 1* is an approximate Robbins-Monro algorithm. Second, we will conclude by applying some classical results in [28].

Observe that for any $i \in \mathcal{N}$, for any $\alpha, \beta, \beta' \in \mathcal{A}_i$, $x_{i\alpha}^\Lambda = \Lambda_{i\alpha}(y_i) = \frac{\exp(y_{i\alpha})}{\sum_{s \in \mathcal{A}_i} \exp(y_{i\alpha})}$, we have

$$\frac{\partial \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta}} = x_{i\alpha}^\Lambda (\mathbb{1}_{s=\beta} - x_{i\beta}^\Lambda), \text{ and } \frac{\partial^2 \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta} \partial y_{i\beta'}} = x_{i\alpha}^\Lambda (\mathbb{1}_{\alpha=\beta=\beta'} - \mathbb{1}_{\alpha=\beta} x_{i\beta'}^\Lambda - x_{i\beta}^\Lambda (\mathbb{1}_{\alpha=\beta'} + \mathbb{1}_{\beta=\beta'} - 2x_{i\beta'}^\Lambda)).$$

Using Taylor's Remainder Theorem, let us rewrite the equation $X_{i\alpha}(n+1) = \frac{\epsilon_i}{|\mathcal{A}_i|} + (1 - \epsilon_i) \Lambda_{i\alpha}(Y_i(n+1))$ as

$$\begin{aligned}
X_{i\alpha}(n+1) &= \frac{\epsilon_i}{|\mathcal{A}_i|} + (1 - \epsilon_i) \Lambda_{i\alpha}(Y_i(n) + \gamma_{n+1} \hat{v}_i(n+1)) \\
&= \frac{\epsilon_i}{|\mathcal{A}_i|} + (1 - \epsilon_i) \Lambda_{i\alpha}(Y_i(n)) \\
&\quad + (1 - \epsilon_i) \gamma_{n+1} \left(\nabla \Lambda_{i\alpha}^T(Y_i(n)) \hat{v}_i(n+1) + \frac{1}{2} \gamma_{n+1} \hat{v}_i^T(n+1) \text{Hess} \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1) \right) \\
&= X_{i\alpha}(n) + (1 - \epsilon_i) \gamma_{n+1} \left(\nabla \Lambda_{i\alpha}^T(Y_i(n)) \hat{v}_i(n+1) + \frac{\gamma_{n+1}}{2} \hat{v}_i^T(n+1) \text{Hess} \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1) \right) \\
&= X_{i\alpha}(n) + (1 - \epsilon_i) \gamma_{n+1} \left(\nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) + \nabla \Lambda_{i\alpha}^T(Y_i(n)) (\hat{v}_i(n+1) - v_i(X(n))) + \gamma_{n+1} a_{n+1} \right)
\end{aligned} \tag{A.5}$$

where $\nabla \Lambda_{i\alpha}$ is the gradient vector of $\Lambda_{i\alpha}$, $\nabla \Lambda_{i\alpha}^T$ is its transposed, $\text{Hess} \Lambda_{i\alpha}$ is the Hessian matrix of $\Lambda_{i\alpha}$, $\psi_i(n)$ is in the line segment going out from $Y_i(n)$ to the point $Y_i(n+1)$, and $a_{n+1} = \frac{1}{2} \hat{v}_i^T(n+1) \text{Hess} \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1)$.

Next, we focus on $\gamma_n a_n$ (corresponding to parameter b_n in Definition 7). Since $\frac{\partial^2 \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta} \partial y_{i\beta'}} = x_{i\alpha}^\Lambda \left(\mathbb{1}_{\alpha=\beta=\beta'} - \mathbb{1}_{\alpha=\beta} x_{i\beta'}^\Lambda - x_{i\beta}^\Lambda (\mathbb{1}_{\alpha=\beta'} + \mathbb{1}_{\beta=\beta'} - 2x_{i\beta'}^\Lambda) \right)$, all components of $\text{Hess} \Lambda_{i\alpha}(\psi_i(n))$ are bounded. So, the limit of $\gamma_n a_n$ (when $n \rightarrow \infty$) depends on the limit of $\|\hat{v}_{i\alpha}(n)\|^2$.

Let $E_{i\alpha,n}$ be the event $\|\hat{v}_{i\alpha}(n)\|^2 \geq n^\alpha$ for $\frac{1}{q} < \alpha < \frac{1}{p}$, with q defined in Hypothesis (H2) and p such that $\gamma_n = \mathcal{O}(1/n^{\frac{1}{p}})$.

Hypothesis (H2) gives us that

$$\sum_{n=0}^{\infty} \mathbb{P}(E_{i\alpha,n}) = \sum_{n=0}^{\infty} \mathbb{P}(\|\hat{v}_{i\alpha}(n)\|^2 \geq n^\alpha \mid \mathcal{F}_{n-1}) = \sum_{n=0}^{\infty} \mathcal{O}\left(\frac{1}{n^{q\alpha}}\right) < \infty \tag{A.6}$$

The Borel-Cantelli lemma gives us that $E_{i\alpha,n}$ is true for only a finite number of $n \in \mathbb{N}$. Therefore for $n > \max_{m \in \mathbb{N}} \{m; \exists i \in \mathcal{N}, \alpha \in \mathcal{A}_i, E_{i\alpha,m} \text{ is true}\}$, $\|\hat{v}_{i\alpha}(n)\|^2 < n^\alpha$. By assumption, $\gamma_n = o(n^b)$ for any $b > -1/p$. In particular, $\gamma_n = o(n^{-\alpha})$ so $\lim_{n \rightarrow \infty} a_n \gamma_n = 0$.

$$\begin{aligned}
X_{i\alpha}(n+1) &= X_{i\alpha}(n) + (1 - \epsilon_i) \gamma_{n+1} \left(X_{i\alpha}^\Lambda(n) \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} v_{i\beta}(X(n)) X_{i\beta}^\Lambda(n) \right) \right) \\
&\quad + (1 - \epsilon_i) \gamma_{n+1} \left(X_{i\alpha}^\Lambda(n) (\hat{v}_{i\alpha}(n+1) - v_{i\alpha}(X(n))) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}^\Lambda(n) [\hat{v}_{i\beta}(n+1) - v_{i\beta}(X(n))] + \gamma_n a_n \right)
\end{aligned}$$

Let $U_{i\alpha,n} = (1 - \epsilon_i) X_{i\alpha}^\Lambda(n) (\hat{v}_{i\alpha}(n+1) - v_{i\alpha}(X(n))) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}^\Lambda(n) [\hat{v}_{i\beta}(n+1) - v_{i\beta}(X(n))]$.

Recall that (3.2a) and (3.2b) can be deduced from Hypotheses (H1) (H2). We get

1. $\mathbb{E}[U_{i\alpha,n} \mid \mathcal{F}_{n-1}] = 0$ for all n
2. $\mathbb{E}[\|U_{i\alpha,n}\|^2] < \infty$ for all n

So, $U_{i\alpha,n}$ is a martingale difference noise.

Since the function which to $x_{i\alpha}$ associates $(1 - \epsilon) x_{i\alpha}^\Lambda \left[v_{i\alpha}(x) - \sum_{\beta \in \mathcal{A}_i} x_{i\beta}^\Lambda v_{i\beta}(x) \right]$ is a continuous map, we can conclude that the stochastic process $\{X_i(n)\}_{n \in \mathbb{N}}$ is an approximate Robbins-Monro algorithm.

Second, Remark 4.5 and Propositions 4.2 and 4.1 of [28] allow us to conclude that the interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ is an asymptotic pseudo trajectory of the solutions of ODE (RD_ε) .

□

Dynamics (RD_ε) can be viewed as a ε -perturbed variant of the replicator dynamics. that readily extends to the ε -perturbed variant (RD_ε) . Moreover, the replicator dynamics have some good property in potential games. The next theorem expresses this property in our context:

Proposition 10. *Let Γ be a generic potential game. The potential function f of Γ is a strict increasing Lyapunov function of the flow induced by the dynamics (RD_ε) .*

Proof. We consider the variation of f . We have

$$\begin{aligned}
\dot{f}(x) &= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \frac{\partial f}{\partial x_{i\alpha}}(x) \dot{x}_{i\alpha} \\
&= \sum_{i \in \mathcal{N}} v_i^\Gamma(x(t)) \dot{x}_i(t) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} v_{i\alpha}(x(t)) \dot{x}_{i\alpha}(t) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} (1 - \varepsilon_i) v_{i\alpha}(x(t)) x_{i\alpha}^\Delta(t) \left(v_{i\alpha}(x(t)) - \sum_{\beta \in \mathcal{A}_i} v_{i\beta}(x(t)) x_{i\beta}^\Delta(t) \right) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \sum_{\beta \in \mathcal{A}_i, \beta > \alpha} (1 - \varepsilon_i) x_{i\alpha}^\Delta(t) x_{i\beta}^\Delta(t) [v_{i\alpha}(x(t)) - v_{i\beta}(x(t))]^2
\end{aligned}$$

The second equation is obtained by definitions of potential game and dynamic (RD_ε) (definition of $\dot{x}_{i\alpha}$). From the latter equation, we can conclude that $\dot{f}(x) \geq 0$.

Now we show that the rest points x of dynamics (RD_ε) are such that $\dot{f}(x) = 0$. Since $\dot{f}(x) = \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \frac{\partial f}{\partial x_{i\alpha}}(x) \dot{x}_{i\alpha}$, $\dot{f}(x) = 0$ when x is a rest point ($\dot{x} = 0$).

Conversely, we will prove that $\dot{f}(x) = 0$ implies that x is a rest point of dynamics (RD_ε) .

Observe that $\dot{f}(x) = 0$ implies that for all players i in \mathcal{N} , and pure strategies α and β in \mathcal{A}_i we have $x_{i\alpha}^\Delta = 0$ or $x_{i\beta}^\Delta = 0$ or $v_{i\alpha}(x) = v_{i\beta}(x)$.

Since the dynamics (RD_ε) is $\dot{x}_{i\alpha} = (1 - \varepsilon_i) x_{i\alpha}^\Delta \left(v_{i\alpha}(x) - \sum_{\beta \in \mathcal{A}_i} v_{i\beta}(x) x_{i\beta}^\Delta \right)$, if for all players i in \mathcal{N} , and pure strategies α and β in \mathcal{A}_i we have $x_{i\alpha}^\Delta = 0$ or $x_{i\beta}^\Delta = 0$, then we can deduce $\dot{x}_{i\alpha} = 0$. Otherwise, if $v_{i\alpha}(x) = v_{i\beta}(x)$ for all α and β in \mathcal{A}_i such that $x_{i\alpha}^\Delta \neq 0$ and $x_{i\beta}^\Delta \neq 0$, we have then $\dot{x}_{i\alpha} = 0$

To conclude, f is increasing, and its derivative is null if and only if it is evaluated on a rest point of the dynamics (RD_ε) . f is a strict increasing Lyapunov function of the dynamics (RD_ε) . □

Proposition 11. *Let Γ be a generic potential game. Suppose that Algorithm 1 is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative mixing factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. The interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ converges to a rest point of RD_ε .*

Proof. We showed that under these assumptions, $(X_i(n))_{n \in \mathbb{N}}$ is a pseudo asymptotic trajectory of the flow induced by the dynamics RD_ε .

We now show that RD_ε has a finite number of rest points. When $\forall i \in \mathcal{N}, \varepsilon_i = 0$, we have $x^\Delta = x$. Thus, we obtain from the previous proof that x is a rest point of RD_ε if and only if

$$\sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \sum_{\beta \in \mathcal{A}_i, \beta > \alpha} (1 - \varepsilon_i) x_{i\alpha}(t) x_{i\beta}(t) [v_{i\alpha}(x(t)) - v_{i\beta}(x(t))]^2 = 0 \quad (\text{A.7})$$

So x is a rest point of RD_ε if and only if $v_{i\alpha}(x(t)) = v_{i\beta}(x(t)) \quad \forall \alpha, \beta \in \text{supp}(x_i)$. Therefore x is a rest point of RD_ε if and only if it is a restricted equilibrium. The game is finite so they are a finite number of game restrictions. Each restricted game is a finite, generic, potential game so it has a finite number of Nash-equilibrium.

Now we address the other case : $\exists i \in \mathcal{N}, \varepsilon_i \neq 0$. Thus, x is a rest point of RD_ε if and only if $v_{i\alpha}(x(t)) = v_{i\beta}(x(t)) \quad \forall \alpha, \beta \in \text{supp}(x_i^\Delta)$. Let supp_i be the cardinal of $\text{supp}(x_i^\Delta)$. For each player i , we have $\text{supp}_i - 1$ unknown coordinates corresponding to the x_i^Δ non-zero coordinates. Thus, $\text{supp}_i - 1$ independent equations corresponding to $v_{i\alpha}(x(t)) = v_{i\beta}(x(t)) \quad \forall \alpha, \beta \in \text{supp}(x_i^\Delta)$ because the game is generic. Such a system has only an unique solution. Therefore for each possible support of x there is only one rest point, and dynamics RD_ε has a finite number of rest points.

Corollary 6.6 of [28] allows to conclude that the continuous-time process $x_i(t)$ converges to a rest point of RD_ε . \square

In order to prove Point 1 of Theorem 1, we need to have the following technical result. This lemma is about the properties on rationality properties (such as comparaisons among strategies corresponding to the elimination of dominated strategies).

Lemma 12. *Suppose that Algorithm 1 is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative mixing factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. If there exists some $a > 0$ such that $v_\beta(x) - v_\alpha(x) \geq a$ for all $x \in \mathcal{X}$ then for all $c \in (0, a)$, there exists some n_0 such that $Y_\beta(n) - Y_\alpha(n) \geq c \sum_{k=1}^n \gamma_k$ for all $n \geq n_0$ (a.s.).*

Proof. Let $\zeta_k = \hat{v}_\beta(k) - v_\beta(X(k-1)) - [\hat{v}_\alpha(k) - v_\alpha(X(k-1))]$. By assumption there exists $a > 0$ such that $v_\beta(X) - v_\alpha(X) \geq a$ for all $X \in \mathcal{X}$. Then,

$$\begin{aligned} Y_\beta(n) - Y_\alpha(n) &= Y_\beta(0) - Y_\alpha(0) + \sum_{k=1}^n \gamma_k (\hat{v}_\beta(k) - \hat{v}_\alpha(k)) \\ &= Y_\beta(0) - Y_\alpha(0) + \sum_{k=1}^n \gamma_k [v_\beta(X(k-1)) - v_\alpha(X(k-1))] + \sum_{k=1}^n \gamma_k \zeta_k \quad (\text{A.8}) \\ &\geq Y_\beta(0) - Y_\alpha(0) + \sum_{k=1}^n \gamma_k \left[a + \frac{\sum_{k=1}^n \gamma_k \zeta_k}{\sum_{k=1}^n \gamma_k} \right]. \end{aligned}$$

Now we will prove that $\frac{\sum_{k=1}^n \gamma_k \zeta_k}{\sum_{k=1}^n \gamma_k} \rightarrow 0$.

The reformulation of Hypothesis (H1) gives :

$$\begin{aligned} \mathbb{E}[\zeta_k | \mathcal{F}_{k-1}] &= \mathbb{E}[\xi_\beta(k) + v_\beta(s(k)) - v_\beta(X(k-1)) - \xi_\alpha(k) - v_\alpha(\alpha(k)) + v_\alpha(X(k-1)) | \mathcal{F}_{k-1}] \\ &= \mathbb{E}[\xi_\beta(k) - \xi_\alpha(k) | \mathcal{F}_{k-1}] + \mathbb{E}[v_\beta(\alpha(k)) | \mathcal{F}_{k-1}] - v_\beta(X(k-1)) - \mathbb{E}[v_\alpha(\alpha(k)) | \mathcal{F}_{k-1}] + v_\alpha(X(k-1)) \\ &= 0 \end{aligned} \quad (\text{A.9})$$

ζ_k is \mathcal{F}_k -measurable, meaning that it is fully determined by the information of \mathcal{F}_k .

With $S_n = \sum_{k=1}^n \gamma_k \zeta_k$, it follows that

$$\mathbb{E}[S_k | \mathcal{F}_{k-1}] = \gamma_k \mathbb{E}[\zeta_k | \mathcal{F}_{k-1}] + \mathbb{E}[S_{k-1} | \mathcal{F}_{k-1}] = S_{k-1} \quad (\text{A.10})$$

Therefore $\{S_n = \sum_{k=1}^n \gamma_k \zeta_k, \mathcal{F}_n, n \geq 1\}$ is a martingale, in addition as γ_t depends only on n and it is positive, $U_n = \sum_{k=1}^n \gamma_k$ is a nondecreasing sequence of positive random variable such that U_n is \mathcal{F}_{n-1} -measurable for each n . In addition $\beta \leq 1$ gives us that $\lim_{n \rightarrow \infty} U_n = \infty$.

We focus now on proving the last hypothesis of Theorem 2.18 of [31], i.e., $\sum_{k=1}^{\infty} \frac{\mathbb{E}[\|\gamma_k \zeta_k\|^2 | \mathcal{F}_{k-1}]}{U_k^2} < \infty$. First we show that $\mathbb{E}[\|\zeta(k)\|^2 | \mathcal{F}_{k-1}] \leq 4\sigma^2$:

$$\begin{aligned} \mathbb{E}[\|\zeta(k)\|^2 | \mathcal{F}_{k-1}] &= \mathbb{E}[\|\hat{v}_\beta(k) - v_\beta(X(k-1)) - [\hat{v}_\alpha(k) - v_\alpha(X(k-1))]\|^2 | \mathcal{F}_{k-1}] \\ &= 2\mathbb{E}[\|\hat{v}_\beta(k) - v_\beta(X(k-1))\|^2 | \mathcal{F}_{k-1}] + 2\mathbb{E}[\|\hat{v}_\alpha(k) - v_\alpha(X(k-1))\|^2 | \mathcal{F}_{k-1}] \\ &\leq 4\sigma^2 \end{aligned} \quad (\text{A.11})$$

according to Equation (3.2b). Second, γ_t is decreasing so $U_n = \sum_{k=1}^n \gamma_k \geq n\gamma_n$ and $U_n^{-2} \leq \frac{1}{n^2\gamma_n^2}$. Therefore

$$\sum_{k=1}^{\infty} \frac{\mathbb{E}(\|\gamma_k \zeta_k\|^2 | \mathcal{F}_{k-1})}{U_k^2} < \sum_{k=1}^{\infty} \frac{\gamma_k^2 4\sigma^2}{k^2 \gamma_k^2} = 4\sigma^2 \sum_{k=1}^{\infty} \frac{1}{k^2} < \infty \quad (\text{A.12})$$

Therefore all hypotheses are fulfilled and

$$\lim_{n \rightarrow \infty} \frac{\sum_{k=1}^n \gamma_k \zeta_k}{\sum_{k=1}^n \gamma_k} = 0 \quad (\text{a.s.}) \quad (\text{A.13})$$

So for all $c \in (0, a)$ there exist some t_0 such that $-\frac{\sum_{k=1}^n \gamma_k \zeta_k}{\sum_{k=1}^n \gamma_k} \leq \frac{Y_\beta(0) - Y_\alpha(0)}{\sum_{k=1}^n \gamma_k} + a - c$ for all $t > t_0$. Putting that in (A.8) we have :

$$Y_\beta(n) - Y_\alpha(n) \geq c \sum_{k=1}^n \gamma_k \quad \text{for all } t \geq t_0 \text{ (a.s.)} \quad (\text{A.14}) \quad \square$$

Now, we will prove the convergence of the discrete-time process $(X_i(n))_{n \in \mathbb{N}}$.

Theorem 1 (Part 1). *Let Γ be a generic potential game. Suppose that Algorithm 1 is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative mixing factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. $X(n)$ converges (a.s.) to a δ -equilibrium of Γ with $\delta \equiv \delta(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.*

Proof. When $X(n)$ converges, $X^\Lambda(n)$ also converges (by definition). Applying Lemma 6, it suffices to prove that if $X^\Lambda(n)$ converges to x^* then x^* is Nash Equilibrium. We prove by contradiction that $(X^\Lambda(n))_{n \in \mathbb{N}}$ converges to x^* a Nash Equilibrium. Assume that x^* is not a Nash Equilibrium. By definition, we have

$$\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), \text{ s.t., } v_{i\beta}(x^*) > v_{i\alpha}(x^*), \forall \alpha \in \text{supp}(x_i^*).$$

By continuity of utility u , there is a neighborhood U of x^* and $a > 0$ such that:

$\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), \text{ s.t., } v_{i\beta}(X^\Lambda) - v_{i\alpha}(X^\Lambda) > a, \forall \alpha \in \text{supp}(x_i^*), X^\Lambda \in U$ For ε small enough and for all n big enough, $X(n) \in U$ because $\|X(n) - x^*\| \leq \|X(n) - X^\Lambda(n)\| + \|X^\Lambda(n) - x^*\| \leq N\varepsilon + \|X^\Lambda(n) - x^*\|$. So, $\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), \text{ s.t., } v_{i\beta}(X(n)) - v_{i\alpha}(X(n)) > a, \forall \alpha \in \text{supp}(x_i^*)$ Using Lemma 12, for n_0 big enough and $n \geq n_0$:

$$Y_{i\beta}(n) - Y_{i\alpha}(n) \geq C + b \sum_{t=n_0}^n \gamma_t \rightarrow \infty \quad (\text{A.15})$$

Thus $\frac{X_{i\beta}(n)}{X_{i\alpha}(n)} = \exp(Y_{i\beta}(n) - Y_{i\alpha}(n)) \rightarrow \infty$.

$X_{i\alpha}(n) \rightarrow 0$ and α is not in the support of x_i^* which is a contradiction. Therefore x^* is a Nash Equilibrium and x^* is a $\delta(\varepsilon)$ -Nash Equilibrium. \square

Finally, we will prove the convergence rate of the discrete-time process $(X_i(n))_{n \in \mathbb{N}}$ corresponding to Part 2 of Theorem 1.

Theorem 1 (Part 2). *Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) semi-bandit feedback satisfying (H1) and (H2); ii) a nonnegative mixing factor $\varepsilon \geq 0$; and iii) a step-size sequence of the form $\gamma_n \propto 1/n^\beta$ for some $\beta \in (1/q, 1]$. Then If $\lim_{n \rightarrow \infty} X(n)$ is an ε -pure state of the form $x_i^* = \varepsilon / |\mathcal{A}_i| \mathbf{1} + (1 - \varepsilon) e_{\hat{\alpha}_i}$ for some $\hat{\alpha}_i \in \mathcal{A}$, then $\hat{\alpha}_i$ is a.s. a strict equilibrium of Γ and convergence occurs at a quasi-exponential rate:*

$$X_{i\hat{\alpha}_i}(n) \geq 1 - \varepsilon - b e^{-c \sum_{k=1}^n \gamma_k} \quad \text{for some positive } b, c > 0. \quad (\text{4.4})$$

Proof. We will focus on the sequence $(X_i^\Lambda(n))_{n \in \mathbb{N}}$ and its limit x^* . From the proof of Theorem 1 (Part 1), x^* is a Nash equilibrium. By continuity of u , there is a neighborhood U of x^* and $a' > 0$ such that:

$\forall i \in \mathcal{N}, \forall \beta \in \mathcal{A}_i, \beta \neq \alpha_i^*, v_{i\alpha_i^*}(x) - v_{i\beta}(x) > a', x \in U$.

Therefore, for ϵ small enough and for all n big enough, $X(n) \in U$ because $\|X(n) - x^*\| \leq \|X(n) - X^\Lambda(n)\| + \|X^\Lambda(n) - x^*\| \leq N\epsilon + \|X(n) - x^*\|$. So, $\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), s.t., v_{i\beta}(X^\Lambda(n)) - v_{i\alpha}(X^\Lambda(n)) > a, \forall \alpha \in \text{supp}(x_i^*)$ Using Lemma 12 for n_0 big enough:

$$Y_{i\alpha^*}(n) - Y_{i\alpha}(n) \geq C + b \sum_{k=n_0}^n \gamma_k$$

So, by computation, we can deduce that

$$\sum_{s \in \mathcal{A}_i, \alpha \neq \hat{\alpha}} \exp(Y_{i\alpha}(n) - Y_{i\hat{\alpha}}(n)) \leq \sum_{s \in \mathcal{A}_i, \alpha \neq \hat{\alpha}} \exp(-C_{n_0} - \sum_{k=n_0}^n \gamma_k a) \leq C \exp(-\sum_{k=n_0}^n \gamma_k a)$$

where we set $C = |\mathcal{A}_i| \exp(-C_{t_0})$. Now, we will focus on $x_{\hat{\alpha}}^\Lambda(n)$.

$$\begin{aligned} X_{i\hat{\alpha}}^\Lambda(n) &= \frac{\exp(Y_{i\hat{\alpha}}(n))}{\sum_{s \in \mathcal{A}} \exp(Y_{i\alpha}(n))} = \frac{1}{\sum_{s \in \mathcal{A}_i} \exp(Y_{i\alpha}(n) - Y_{i\hat{\alpha}}(n))} \\ &= \frac{1}{1 + \sum_{s \in \mathcal{A}_i, s \neq \hat{\alpha}} \exp(Y_{i\alpha}(n) - Y_{i\hat{\alpha}}(n))} \\ &\geq \frac{1}{1 + C \exp(-\sum_{k=n_0}^n \gamma_k a)} \end{aligned}$$

The first equation is due to the relation $\frac{\exp a}{\exp b} = \frac{1}{\exp a-b}$. The second equation is obtained by the fact that $\exp(\exp(Y_{i\hat{\alpha}}(n) - Y_{i\hat{\alpha}}(n))) = 1$. Since for any $z > 0, \frac{1}{1+z} \geq 1 - z$, we obtain $1 - X_{i\hat{\alpha}}^\Lambda(n) \leq C \exp(-\sum_{k=n_0}^n \gamma_k a)$. And, the theorem holds since $|X_{i\hat{\alpha}}^\Lambda(n) - X_{i\hat{\alpha}}(n)| \leq \epsilon$. \square

B Learning with bandit feedback

We begin our proof of Theorem 4 with some general properties of approximate Robbins–Monro algorithms. To state them, let $F : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a continuous map and consider a Robbins–Monro algorithm of the form

$$x_{n+1} - x_n = \gamma_{n+1} (F(x_n) + U_{n+1} + b_{n+1}) \quad (\text{B.1})$$

where

- $\{\gamma_n\}_{n \geq 1}$ is a given sequence of nonnegative numbers such that $\sum_{n=1}^{\infty} \gamma_n = \infty, \lim_{n \rightarrow \infty} \gamma_n = 0$,
- $U_n \in \mathbb{R}^m$ are martingale difference perturbations ,i.e., $\mathbb{E}[U_{n+1} | \mathcal{F}_n] = 0$ and $\mathbb{E}[\|U_n\|^2] < \infty \forall n$,
- $b_n \in \mathbb{R}^m$ is an error term such that $\lim_{n \rightarrow \infty} b_n = 0$ with probability 1.

We define:

- $\tau_0 = 0$ and $\tau_n = \sum_{i=1}^n \gamma_i$ for $n \geq 1$,
- $m : \mathbb{R}_+ \rightarrow \mathbb{N}$ by $m(t) = \sup\{k \geq 0 : t \geq \tau_k\}$,
- $\bar{U} : \mathbb{R}_+ \rightarrow \mathbb{N}$ by $\bar{U}(\tau_n + s) = U_{n+1}$, for all $0 \leq s < \gamma_{n+1}$,
- $\bar{\gamma} : \mathbb{R}_+ \rightarrow \mathbb{N}$ by $\bar{\gamma}(\tau_n + s) = \gamma_{n+1}$, for all $0 \leq s < \gamma_{n+1}$,

In order to prove that such an algorithm is an asymptotic pseudo trajectories of a given flow it is important to prove that for any $T > 0$, $\lim_{t \rightarrow \infty} \sup_{m(t) < h \leq m(t+T)} \|\sum_{i=m(t)}^{h-1} \gamma_{i+1} (U_{i+1} + b_{i+1})\| = 0$ with probability 1.

Lemma 13. *Let $\{x_n\}$ be a sequence given by (B.1). Suppose that:*

$$\sum_n \gamma_{n+1}^{\left(1+\frac{q}{2}\right)} \mathbb{E}[\|U_{n+1}\|^q] < \infty \text{ for a } q \geq 2$$

Then, for any $T > 0$, $\lim_{t \rightarrow \infty} \sup_{m(t) < h \leq m(t+T)} \|\sum_{i=m(t)}^{h-1} \gamma_{i+1} (U_{i+1} + b_{i+1})\| = 0$ with probability 1.

Proof of Lemma 13. To begin, observe that:

$$\begin{aligned} \sup_{m(t) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} (U_{i+1} + b_{i+1}) \right\| &\leq \sup_{m(t) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1} \right\| \\ &+ \sup_{m(t) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} b_{i+1} \right\| \\ &\leq \sup_{m(t) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1} \right\| \\ &+ (T+1) \max_{m(t) < h \leq m(t+T)} b_{h+1} \end{aligned}$$

By definition, $T \max_{m(t) < h \leq m(t+T)} b_{h+1}$ goes to 0. Thus, it remains to prove that $\lim_{t \rightarrow \infty} \sup_{m(t) < h \leq m(t+T)} \|\sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1}\| = 0$ with probability 1.

For any $t \geq 0$ Burkholder's inequality implies that, for some universal constant $C_q > 0$:

$$\mathbb{E} \left[\sup_{m(kT) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \right] \leq C_q \mathbb{E} \left[\left(\sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1}^2 \|U_{i+1}\|^2 \right)^{\frac{q}{2}} \right] \quad (\text{B.2})$$

If $q = 2$ we directly obtain:

$$\mathbb{E} \left[\sup_{m(kT) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^2 \right] \leq C_2 \sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1}^2 \mathbb{E}[\|U_{i+1}\|^2]$$

For $q > 2$ we need the Hölder's inequality :

$$\sum_i x_i y_i \leq \left(\sum_i x_i^u \right)^{\frac{1}{u}} \left(\sum_i y_i^{\frac{u}{u-1}} \right)^{\frac{u-1}{u}}$$

which applied to $x_i = \alpha_i^{1-\delta} |\beta_i|$ and $y_i = \alpha_i^\delta$ gives for any $\alpha_i \geq 0$, $\beta_i \in \mathbb{R}$, $u > 1$ and $0 < \delta < 1$:

$$\left(\sum_i |\alpha_i \beta_i| \right)^u \leq \left(\sum_i \alpha_i^{\frac{\delta u}{u-1}} \right)^{u-1} \sum_i \alpha_i^{(1-\delta)u} |\beta_i|^u \quad (\text{B.3})$$

Suppose $q > 2$ we apply B.3 to the second member of B.2 with $u = \frac{q}{2}$, $\delta = \frac{q-2}{2q}$, $\alpha_i = \gamma_{i+1}^2$ and $\beta_i = \|U_{i+1}\|^2$:

$$\begin{aligned} \mathbb{E} \left[\sup_{m(kT) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \right] &\leq C_q \mathbb{E} \left[\left(\sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1} \right)^{\frac{q}{2}-1} \sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1}^{\frac{q}{2}+1} \|U_{i+1}\|^q \right] \\ &\leq C_q T^{\frac{q}{2}-1} \sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1}^{\frac{q}{2}+1} \mathbb{E}[\|U_{i+1}\|^q] \end{aligned} \quad (\text{B.4})$$

Therefore, for $q \geq 2$, we have

$$\mathbb{E} \left[\sup_{m(kT) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \right] \leq C_q T^{\frac{q}{2}-1} \sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1}^{\frac{q}{2}+1} \mathbb{E} [\|U_{i+1}\|^q]. \quad (\text{B.5})$$

Let $E(k, \varepsilon)$ be the event $A(kT, T) = \sup_{m(kT) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} \right\| > \varepsilon$, for $\varepsilon > 0$.

Doob's martingal inequality yields:

$$\mathbb{P}(E(k, \varepsilon)) \leq \frac{\mathbb{E}[A(kT, T)^q]}{\varepsilon^q}$$

So

$$\begin{aligned} \sum_{k \in \mathbb{N}} \mathbb{P}(E(k, \varepsilon)) &\leq \frac{\sum_{k \geq 0} \mathbb{E}[A(kT, T)^q]}{\varepsilon^q} \\ &= \frac{C_q T^{\frac{q}{2}-1} \sum_{i=0}^{\infty} \gamma_{i+1}^{(q/2+1)} \mathbb{E} [\|\bar{U}_{i+1}\|^q]}{\varepsilon^q} \\ &< \infty \end{aligned} \quad (\text{B.6})$$

By the Borel-Cantelli Lemma this proves that for all $\varepsilon > 0$:

$$\mathbb{P}(\limsup_{k \rightarrow \infty} E(k, \varepsilon)) = 0$$

with probability one. So for all ε the set of k such that so $E(k, \varepsilon)$ is true is a null set.

Consider now $E(n) = \{k : E(k, \frac{1}{n}) \text{ is true}\}$. For all n , $E(n)$ is a null set. Therefore as a countable union of null set $\bigcup_{n \in \mathbb{N}} E(n)$ is a null set, so $\lim_{k \rightarrow \infty} A(kT, T) = 0$ with probability one.

On the other hand for $kT \leq t < (k+1)T$:

$$\begin{aligned} A(t, T) &= \sup_{m(t) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \\ &\leq \max \left(\sup_{m(t) < h \leq m((k+1)T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q, \sup_{m((k+1)T) < h \leq m(t+T)} \left\| \sum_{i=m(t)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \right) \\ &= \max \left(\sup_{m(t) < h \leq m((k+1)T)} \left\| \sum_{i=m(kT)}^{h-1} \gamma_{i+1} U_{i+1} - \sum_{i=m(kT)}^{m(t)-1} \gamma_{i+1} U_{i+1} \right\|^q, \right. \\ &\quad \left. \sup_{m((k+1)T) < h \leq m(t+T)} \left\| \sum_{i=m(kT)}^{m((k+1)T)-1} \gamma_{i+1} U_{i+1} - \sum_{i=m(kT)}^{m(t)-1} \gamma_{i+1} U_{i+1} + \sum_{i=m((k+1)T)}^{h-1} \gamma_{i+1} U_{i+1} \right\|^q \right) \\ &\leq 2A(kT, T) + A((k+1)T, T) \end{aligned} \quad (\text{B.7})$$

This concludes the proof. \square

Now, let $L = \max_{x \in \mathcal{X}, i \in \mathcal{N}} u_i(x)$, and $A_i = |\mathcal{A}_i|$.

$$\hat{v}_{i\alpha_i}(n) = \frac{\mathbf{1}_{\alpha_i = \alpha_i(n)}}{X_{i\alpha_i}(n-1)} u_i(\alpha(n)) \quad (\text{B.8})$$

First it is useful to remark that we have :

$$\mathbb{E}[\hat{v}_{i\alpha_i}(n) | \mathcal{F}_{n-1}] = \mathbb{E} \left[\frac{\mathbf{1}_{\alpha_i = \alpha_i(n)}}{X_{i\alpha_i}(n-1)} u_i(\alpha(n)) \middle| \mathcal{F}_{n-1} \right] = v_{i\alpha_i}(X(n-1))$$

$$\begin{aligned}\mathbb{E}\left[(\hat{v}_{i\alpha_i}(n) - v_{i\alpha_i}(X(n-1)))^2 \mid \mathcal{F}_{n-1}\right] &\leq 2\mathbb{E}\left[\frac{\mathbf{1}_{\alpha_i=\alpha_i(n)}}{X_{i\alpha_i}(n-1)^2} u_i^2(\alpha(n)) \mid \mathcal{F}_{n-1}\right] + 2\mathbb{E}[v_{i\alpha_i}^2(X(n-1)) \mid \mathcal{F}_{n-1}] \\ &\leq 2\frac{L^2 A_i}{\varepsilon_{n-1}} + 2L^2 = 2L^2\left(1 + \frac{A_i}{\varepsilon_{n-1}}\right)\end{aligned}\tag{B.9}$$

We begin by proving that when a strategies always leads to a better payoff, than an other its score becomes higher. So omitting the player index we can prove the following Lemma.

Lemma 14. *Suppose that Algorithm 1 is run with i) a bandit feedback estimator B.8; ii) a decreasing step-size γ_n ; iii) a decreasing mixing factor ε_n such that :*

$$\lim_{n \rightarrow \infty} \gamma_n = \infty \text{ and } \sum_{n=1}^{\infty} \frac{1}{\varepsilon_n n^2} < \infty \tag{B.10}$$

If there exists some $\mu > 0$ and a set $V \subset \mathcal{X}$ such that $v_\beta(x) - v_\alpha(x) \geq \mu$ for all $x \in V$ and if $X(n) \in V$, for all n , then for all $c \in (0, \mu)$, there exists some T_0 such that $Y_\beta(T) - Y_\alpha(T) \geq c \sum_{n=1}^T \gamma_n$ for all $T \geq T_0$ (a.s.).

Proof. Proof of Lemma 14 Let $\zeta_k = \hat{v}_\beta(k) - v_\beta(X(k-1)) - [\hat{v}_\alpha(k) - v_\alpha(X(k-1))]$. By assumption there exists $a > 0$ such that $v_\beta(x) - v_\alpha(x) \geq a$ for all $x \in V$. Then,

$$\begin{aligned}Y_\beta(T) - Y_\alpha(T) &= c_{\beta\alpha} + \sum_{n=1}^T \gamma_n [\hat{v}_\beta(n) - \hat{v}_\alpha(X(n-1))] \\ &= c_{\beta\alpha} + \sum_{n=1}^T \gamma_n [v_\beta(X(n-1)) - v_\alpha(X(n-1))] + \sum_{n=1}^T \gamma_n \zeta_n\end{aligned}\tag{B.11}$$

$$\geq c_{\beta\alpha} + \tau_T \left[\mu + \frac{\sum_{n=1}^T \gamma_n \zeta_n}{\tau_T} \right] \tag{B.12}$$

where we set $c_{\beta\alpha} = Y_\beta(0) - Y_\alpha(0)$ and $\tau_T = \sum_{n=1}^T \gamma_n$.

The law of large numbers for martingale difference sequences [31, Theorem 2.18] gives $\tau_T^{-1} \sum_{n=1}^T \gamma_n \zeta_n \rightarrow 0$ (a.s.), provided that $\sum_{n=1}^{\infty} \frac{\mathbb{E}[\|\gamma_n \zeta_n\|^2 \mid \mathcal{F}_{n-1}]}{\tau_n^2} < \infty$ and $\sum_{n=1}^{\infty} \gamma_n = \infty$.

We focus on an upper bound to $\mathbb{E}[\|\gamma_n \zeta_n\|^2 \mid \mathcal{F}_{n-1}]$, using (B.9) we obtain:

$$\begin{aligned}\mathbb{E}[\|\gamma_n \zeta_n\|^2 \mid \mathcal{F}_{n-1}] &\leq 4\gamma_n^2 \max_{\alpha \in \mathcal{A}} \mathbb{E}[(\hat{v}_\alpha(n) - v_\alpha)^2 \mid \mathcal{F}_{n-1}] \\ &\leq 8L^2 \gamma_n^2 \left(\frac{A}{\varepsilon_{n-1}} + 1 \right)\end{aligned}\tag{B.13}$$

Using that $\tau_n^2 \geq n^2 \gamma_n^2$ we obtain:

$$\frac{\mathbb{E}[\|\gamma_n \zeta_n\|^2 \mid \mathcal{F}_{n-1}]}{\tau_n^2} \leq \frac{8L^2}{n^2} \left(1 + \frac{A}{\varepsilon_{n-1}} \right)$$

and

$$\sum_{n=1}^T \frac{\mathbb{E}[\|\gamma_n \zeta_n\|^2 \mid \mathcal{F}_{n-1}]}{\tau_n^2} \leq 8L^2 \left(\sum_{n=1}^T \frac{1}{n^2} + \sum_{n=1}^T \frac{A}{\varepsilon_{n-1} n^2} \right) = \mathcal{O}\left(\sum_{n=1}^T \frac{1}{\varepsilon_{n-1} n^2}\right) < \infty$$

We readily get $\tau_T^{-1} \sum_{n=1}^T \gamma_n \zeta_n \rightarrow 0$ with probability 1. As a result, for all $c \in (0, \mu)$, there exists some random (but a.s. finite) T_0 such that if $T \geq T_0$, then $Y_\beta(T) - Y_\alpha(T) \geq c\tau_T$. \square

With this auxiliary result at hand, our proof of Theorem 4 relies on the following four steps:

1. Show that X is an asymptotic pseudo trajectory of a continuous dynamics (Proposition 17),
2. Show that the potential function of the game is a strict Lyapunov function of the dynamics (Proposition 18),
3. Show that X converges towards a rest point of the dynamics (Proposition 19),
4. Show that if X converged toward a point it is a Nash Equilibrium (Proposition 20).

Definition 15 (Asymptotic Pseudo-trajectories). Given a flow $\phi : \mathbb{R} \times M \rightarrow M$, $(n, x) \rightarrow \phi(n, x) = \phi_n(x)$ such that $\phi_0 = \text{Identity}$ and $\phi_{n+\alpha} = \phi_n \circ \phi_\alpha$, a continuous function $X : \mathbb{R} \rightarrow M$ is an asymptotic pseudo-trajectory if

$$\lim_{n \rightarrow \infty} \sup_{0 \leq k \leq T} d((X(n+k), \phi_h(X(n))) = 0 \text{ for any } T > 0 \quad (\text{B.14})$$

Before addressing the proof we show that specific conditions on the mixing factor and the step size, guarantee a few important relations.

Lemma 16. Suppose that Algorithm 1 is run with i) the bandit estimator (B.8), ii) a step size sequence γ_n decreasing to 0 and iii) a strictly positive mixing factor ε_n decreasing to 0 such that:

$$\lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_n^2} = 0, \quad \sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_n} < \infty \quad (\text{B.15})$$

Then for all players i and $\alpha \in \mathcal{A}_i$, $U_{i\alpha}(n+1) = \nabla \Lambda_{i\alpha}^T(Y_i(n)) (\hat{v}_i(n+1) - v_i(X(n)))$, is a martingale difference noise, $\sum_n \gamma_{n+1}^2 \mathbb{E}[\|U_{i\alpha}(n+1)\|^2] < \infty$ and $\lim_{n \rightarrow \infty} \gamma_{n+1} \hat{v}_{i\alpha}(n+1)^2 = 0$.

Proof. Proof of Lemma 16 For all players $i \in \mathcal{N}$, and for all $\alpha \in \mathcal{A}_i$, $U_{i\alpha}(n)$ is a martingale difference noise. In deed, $\mathbb{E}[U_{i\alpha}(n+1) | \mathcal{F}_n] = \nabla \Lambda_{i\alpha}^T(Y_i(n)) (v_i(X(n)) - v_i(X(n))) = 0$ (because $\mathbb{E}[\hat{v}_i(n+1) | \mathcal{F}_n] = v_i(X(n))$), and by definition of the bandit estimator, $\mathbb{E}[(\hat{v}_{i\alpha_i}(n+1) - v_{i\alpha_i}(X(n)))^2 | \mathcal{F}_n]$ is finite for all n , so $\mathbb{E}[\|U_{i\alpha}(n+1)\|^2]$ is also finite for all n , so $U_{i\alpha}(n)$ is a martingale difference noise.

To prove that $\sum_n \gamma_{n+1}^2 \mathbb{E}[\|U_{i\alpha}(n+1)\|^2] < \infty$, remark that $\mathbb{E}[\|U_{i\alpha}(n+1)\|^2 | \mathcal{F}_n] = \mathcal{O}(\frac{1}{\varepsilon_n})$. Therefore $\sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_{n-1}} < \infty$ allows us to conclude the second point.

For the last point, observe that $\hat{v}_{i\alpha}(n+1) = \mathcal{O}(\frac{1}{\varepsilon_n^2})$, so the hypothesis $\lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_{n-1}^2} = 0$ is sufficient to conclude that $\lim_{n \rightarrow \infty} \gamma_{n+1} \hat{v}_{i\alpha}(n+1)^2 = 0$. \square

Proposition 17. Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) the bandit estimator (B.8), ii) a step size sequence γ_n decreasing to 0 and iii) a strictly positive mixing factor ε_n decreasing to 0 such that:

$$\sum_{n=1}^{\infty} \gamma_n = \infty, \quad \lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_n^2} = 0, \quad \sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_n} < \infty \text{ and } \lim_{n \rightarrow \infty} \frac{\varepsilon_n - \varepsilon_{n+1}}{\gamma_{n+1}} = 0. \quad (\text{B.16})$$

Then, for all players $i \in \mathcal{N}$, the interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ are asymptotic pseudo trajectories of replicator dynamics:

$$\dot{x}_{i\alpha} = x_{i\alpha} \left[v_{i\alpha}(x) - \sum_{\beta \in \mathcal{A}_i} x_{i\beta} v_{i\beta}(x) \right], \quad (\text{RD})$$

Proof. Proof of Proposition 17 Observe that for any $i \in \mathcal{N}$, for any $\alpha, \beta, \beta' \in \mathcal{A}_i$, $\Lambda_{i\alpha}(y_i) = \frac{\exp(y_{i\alpha})}{\sum_{s \in \mathcal{A}_i} \exp(y_{is})}$, we have

$$\frac{\partial \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta}} = \Lambda_{i\alpha}(y_i) (\mathbb{1}_{\alpha=\beta} - \Lambda_{i\beta}(y_i)), \quad \text{and} \quad \frac{\partial^2 \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta} \partial y_{i\beta'}} = \Lambda_{i\alpha}(y_i) (\mathbb{1}_{\alpha=\beta=\beta'} - \mathbb{1}_{\alpha=\beta} \Lambda_{i\beta'}(y_i) - \Lambda_{i\beta}(y_i) (\mathbb{1}_{\alpha=\beta'} + \mathbb{1}_{\beta=\beta'} - 2 \Lambda_{i\beta'}(y_i))).$$

Using Taylor's Remainder Theorem, we rewrite the equation

$$\begin{aligned}
X_{i\alpha}(n+1) &= \frac{\varepsilon_{n+1}}{A_i} + (1 - \varepsilon_{n+1}) \Lambda_{i\alpha}(Y_i(n+1)) \text{ as} \\
X_{i\alpha}(n+1) &= \frac{\varepsilon_{n+1}}{A_i} + (1 - \varepsilon_{n+1}) \Lambda_{i\alpha}(Y_i(n) + \gamma_{n+1} \hat{v}_i(n+1)) \\
&= \frac{\varepsilon_{n+1}}{A_i} + (1 - \varepsilon_{n+1}) \Lambda_{i\alpha}(Y_i(n)) \\
&\quad + (1 - \varepsilon_{n+1}) \gamma_{n+1} \left(\nabla \Lambda_{i\alpha}^T(Y_i(n)) \hat{v}_i(n+1) + \frac{1}{2} \gamma_{n+1} \hat{v}_i^T(n+1) \text{Hess } \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1) \right) \\
&= X_{i\alpha}(n) + (\varepsilon_{n+1} - \varepsilon_n) \left(\frac{1}{A_i} - \Lambda_{i\alpha}(Y_i(n)) \right) \\
&\quad + (1 - \varepsilon_{n+1}) \gamma_{n+1} \left(\nabla \Lambda_{i\alpha}^T(Y_i(n)) \hat{v}_i(n+1) + \frac{\gamma_{n+1}}{2} \hat{v}_i^T(n+1) \text{Hess } \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1) \right) \\
&= X_{i\alpha}(n) + (1 - \varepsilon_{n+1}) \gamma_{n+1} (\nabla \Lambda_{i\alpha}^T(Y_i(n)) \hat{v}_i(n+1) + a_{in}) \\
&= X_{i\alpha}(n) + (1 - \varepsilon_{n+1}) \gamma_{n+1} (\nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) + \nabla \Lambda_{i\alpha}^T(Y_i(n)) (\hat{v}_i(n+1) - v_i(X(n))) + a_{in})
\end{aligned} \tag{B.17}$$

where $\nabla \Lambda_{i\alpha}$ is the gradient vector of $\Lambda_{i\alpha}$, $\nabla \Lambda_{i\alpha}^T$ is its transposed, $\text{Hess } \Lambda_{i\alpha}$ is the Hessian matrix of $\Lambda_{i\alpha}$, and $\psi_i(n)$ is in the line segment going out from $Y_i(n)$ to the point $Y_i(n+1)$ and $a_{in} = \frac{\gamma_{n+1}}{2} \hat{v}_i^T(n+1) \text{Hess } \Lambda_{i\alpha}(\psi_i(n)) \hat{v}_i(n+1) + \frac{\varepsilon_{n+1} - \varepsilon_n}{(1 - \varepsilon_{n+1}) \gamma_{n+1}} \left(\frac{1}{A_i} - \Lambda_{i\alpha}(Y_i(n)) \right)$.

Furthermore, $\nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) = \Lambda_{i\alpha}(Y_i(n)) \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} \Lambda_{i\beta}(Y_i(n)) v_{i\beta}(X(n)) \right)$ and we can rewrite

$$\Lambda_{i\alpha}(Y_i(n)) = \frac{X_{i\alpha}(n) - \frac{\varepsilon_n}{A_i}}{1 - \varepsilon_n} = X_{i\alpha}(n) + \frac{\varepsilon_n}{1 - \varepsilon_n} \left(X_{i\alpha}(n) - \frac{1}{A_i} \right)$$

Therefore,

$$\begin{aligned}
\nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) &= X_{i\alpha} \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}(n) v_{i\beta}(X(n)) \right) \\
&\quad + \frac{\varepsilon_n}{1 - \varepsilon_n} \left[\frac{\sum_{\beta \in \mathcal{A}_i} (X_{i\alpha}(n) + X_{i\beta}(n)) v_{i\beta}(X(n))}{A_i} \right] \\
&\quad - 2 \frac{\varepsilon_n}{1 - \varepsilon_n} \left[\sum_{\beta \in \mathcal{A}_i} X_{i\alpha}(n) X_{i\beta}(n) v_{i\beta}(X(n)) \right] \\
&\quad + \frac{\varepsilon_n}{1 - \varepsilon_n} \left[v_{i\alpha}(X(n)) X_{i\alpha}(n) - \frac{v_{i\alpha}(X(n))}{A_i} \right] \\
&\quad - \frac{\varepsilon_n^2}{(1 - \varepsilon_n)^2} \sum_{\beta \in \mathcal{A}_i} \left(X_{i\alpha}(n) - \frac{1}{A_i} \right) \left(X_{i\beta}(n) - \frac{1}{A_i} \right) v_{i\beta}(X(n))
\end{aligned} \tag{B.18}$$

Next, we focus on $b_{in} = a_{in} + \nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) - X_{i\alpha} \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}(n) v_{i\beta}(X(n)) \right)$ and show that it converges towards 0 as n grows. The term $\nabla \Lambda_{i\alpha}^T(Y_i(n)) v_i(X(n)) = X_{i\alpha} \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}(n) v_{i\beta}(X(n)) \right)$ goes to 0 because ε_n is decreasing. Since $\frac{\partial^2 \Lambda_{i\alpha}(y_i)}{\partial y_{i\beta} \partial y_{i\beta'}} = \Lambda_{i\alpha}(y_i) (\mathbb{1}_{\alpha=\beta=\beta'} - \mathbb{1}_{\alpha=\beta} \Lambda_{i\beta'}(y_i) - \Lambda_{i\beta}(y_i) (\mathbb{1}_{\alpha=\beta'} + \mathbb{1}_{\beta=\beta'} - 2 \Lambda_{i\beta'}(y_i)))$, all components of $\text{Hess } \Lambda_{i\alpha}(\psi_i(n))$ are bounded. In addition $\lim_{n \rightarrow \infty} \frac{\varepsilon_n - \varepsilon_{n+1}}{\gamma_{n+1}} = 0$. So, the limit of b_{in} (when $n \rightarrow \infty$) depends on the limit of $\lim_{n \rightarrow \infty} \gamma_n \|\hat{v}_{i\alpha}(n)\|^2 = 0$ as proved in Lemma 16.

We can write the dynamics of Algorithm 1 in the form:

$$X_{i\alpha}(n+1) = X_{i\alpha}(n) + \gamma_{n+1} (F_{i\alpha}(n) + U_{i\alpha}(n+1) + b_{in})$$

where $F_{i\alpha}(n) = X_{i\alpha} \left(v_{i\alpha}(X(n)) - \sum_{\beta \in \mathcal{A}_i} X_{i\beta}(n) v_{i\beta}(X(n)) \right)$ is the replicator dynamics, and $U_{i\alpha}(n+1) = \nabla \Lambda_{i\alpha}^T(Y_i(n)) (\hat{v}_i(n+1) - v_i(X(n)))$.

In Lemma 16 we also showed that $U_{i\alpha}(n)$ is a martingale difference noise and that $\sum_n \gamma_{n+1}^2 \mathbb{E}[\|U_{n+1}\|^2] < \infty$. We apply Lemma 13, and Proposition 4.1 of [28] allows us to conclude that the interpolated process of the sequences $(X_i(n))_{n \in \mathbb{N}}$ are asymptotic pseudo trajectories of the replicator dynamics. □

Proposition 18. *Let Γ be a generic potential game. The potential function f of Γ is a strict increasing Lyapunov function of the flow induced by the replicator dynamics.*

Proof. Proof of Proposition 18 We consider the variation of f to show that f is an increasing function. Using the definition of potential function, we have:

$$\begin{aligned}
\dot{f}(x) &= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \frac{\partial f}{\partial x_{i\alpha}}(x) \dot{x}_{i\alpha} \\
&= \sum_{i \in \mathcal{N}} v_i^T(x(t)) \dot{x}_i(t) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} v_{i\alpha}(x(t)) \dot{x}_{i\alpha}(t) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} v_{i\alpha}(x(t)) x_{i\alpha}(t) \left(v_{i\alpha}(x(t)) - \sum_{\beta \in \mathcal{A}_i} v_{i\beta}(x(t)) x_{i\beta}(t) \right) \\
&= \sum_{i \in \mathcal{N}} \sum_{\alpha \in \mathcal{A}_i} \sum_{\beta \in \mathcal{A}_i, \beta > \alpha} x_{i\alpha}(t) x_{i\beta}(t) [v_{i\alpha}(x(t)) - v_{i\beta}(x(t))]^2 \\
&\geq 0
\end{aligned} \tag{B.19}$$

Furthermore, when x is a rest point of the dynamics, $\dot{x} = 0$ and $\dot{f}(x) = 0$. Conversely, we prove that $\dot{f}(x) = 0$ implies that x is a rest point of the dynamics.

Observe that $\dot{f}(x) = 0$ then for all players i in \mathcal{N} , and for all pure strategies α and β in \mathcal{A}_i , $x_{i\alpha}(t) = 0$ or $x_{i\beta}(t) = 0$ or $v_{i\alpha}(x(t)) = v_{i\beta}(x(t))$. This also leads to $\dot{x}_{i\alpha}$ for all players i in \mathcal{N} , and for all pure strategies α in \mathcal{A}_i . Therefore x is a rest point of the dynamics, and f is a strict increasing Lyapunov function. □

Proposition 19. *Let Γ be a generic potential game and suppose that Algorithm 1 is run with i) the bandit estimator (B.8), ii) a step size sequence γ_n decreasing to 0 and iii) a strictly positive mixing factor ε_n decreasing to 0 such that:*

$$\sum_{n=1}^{\infty} \gamma_n = \infty, \quad \lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_n^2} = 0, \quad \sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_n} < \infty \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{\varepsilon_n - \varepsilon_{n+1}}{\gamma_{n+1}} = 0. \tag{B.20}$$

Then interpolated process of the sequences $X(n)$ converges (a.s.) to a rest point of the replicator dynamics.

Proof. Proof of Proposition 19 We showed in Proposition 17 that under the same assumptions interpolated process of the sequences $X(n)$ is a pseudo asymptotic trajectory of the flow induced by the replicator dynamics.

We now show that the replicator dynamics has a finite number of rest points. The dynamics is

$$\dot{x}_{i\alpha} = x_{i\alpha} \left[v_{i\alpha}(x) - \sum_{\beta \in \mathcal{A}_i} x_{i\beta} v_{i\beta}(x) \right]$$

So x is a rest point if and only if $v_{i\beta}(x) = v_{i\alpha}(x), \forall \alpha, \beta \in \text{supp}(x_i)$. Therefore x is a rest point of the replicator dynamics if and only if it is a restricted Nash-equilibrium. The game is finite so they are a finite number of game restrictions. Each restricted game is a finite, generic and potential game so it has a finite number of Nash-equilibrium. Therefore the number of rest points is finite.

Corollary 6.6 of [28] allows to conclude that the linear interpolation of the sequences $X(n)$ converges (a.s.) to a rest point of the replicator dynamics. □

Proposition 20. *Suppose that Algorithm 1 is run with i) the bandit estimator (B.8), ii) a step size sequence γ_n decreasing to 0 and iii) a strictly positive mixing factor ε_n decreasing to 0 such that:*

$$\sum_{n=1}^{\infty} \gamma_n = \infty, \lim_{n \rightarrow \infty} \frac{\gamma_n}{\varepsilon_n^2} = 0, \sum_{n=1}^{\infty} \frac{1}{\varepsilon_n n^2} < \infty, \sum_{n=1}^{\infty} \frac{\gamma_n^2}{\varepsilon_n} < \infty \text{ and } \lim_{n \rightarrow \infty} \frac{\varepsilon_n - \varepsilon_{n+1}}{\gamma_{n+1}} = 0. \quad (\text{B.21})$$

Then if $X(n)$ converges, its limit is a Nash Equilibrium of the game.

Proof. Proof of Proposition 20 Let $x^* \in X$ such that $\lim_{n \rightarrow \infty} X(n) = x^*$, we show by contradiction that x^* is a Nash Equilibrium. Suppose the contrary, by definition of a Nash Equilibrium we have :

$$\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), \text{ s.t.}, u_{i\beta}(x^*) > u_{i\alpha}(x^*), \forall \alpha \in \text{supp}(x_i^*)$$

By continuity of the utility function u , there exists a neighborhood U of x^* and $a > 0$ such that:

$$\exists i \in \mathcal{N}, \exists \beta \in \mathcal{A}_i, \beta \notin \text{supp}(x_i^*), \text{ s.t.}, u_{i\beta}(x) - u_{i\alpha}(x) > a, \forall \alpha \in \text{supp}(x_i^*), \forall x \in U$$

For n big enough, $X(n) \in U$, so using Lemma 14, for n_0 big enough and for $n \geq n_0$ we have :

$$Y_{i\beta}(T) - Y_{i\alpha}(T) \geq c \sum_{n=1}^T \gamma_n \quad \text{for all } T \geq T_0 \text{ (a.s.).}$$

Thus, with $X_{i\alpha}(T) = \frac{\varepsilon(T)}{A_i} + (1 - \varepsilon(T)) \Lambda_{i\alpha}(Y_i(T))$ by definition, we get

$$\begin{aligned} X_{\alpha}(T) &= \frac{\varepsilon(T)}{A_i} + (1 - \varepsilon(T)) \frac{e^{Y_{i\alpha}(T)}}{\sum_{\beta'} e^{Y_{i\beta'}(T)}} \\ &\leq \frac{\varepsilon(T)}{A_i} + (1 - \varepsilon(T)) \frac{e^{Y_{i\alpha}(T)}}{e^{Y_{i\beta}(T)}} \\ &= \frac{\varepsilon(T)}{A_i} + (1 - \varepsilon(T)) e^{Y_{i\alpha}(T) - Y_{i\beta}(T)} \\ &\leq \frac{\varepsilon(T)}{A_i} + (1 - \varepsilon(T)) e^{-c \sum_{n=1}^T \gamma_n} \text{ (a.s.)}, \end{aligned} \quad (\text{B.22})$$

So $\lim_{T \rightarrow \infty} X_{i\alpha}(T) = 0$ which is a contradiction with $\alpha \in \text{supp}(x_i^*)$. □