

LEARNING IN NONATOMIC GAMES, PART I: FINITE ACTION SPACES AND POPULATION GAMES

SAEED HADIKHANLOO*, RIDA LARAKI†,
PANAYOTIS MERTIKOPOULOS‡, AND SYLVAIN SORIN‡

ABSTRACT. We examine the long-run behavior of a wide range of dynamics for learning in nonatomic games, in both discrete and continuous time. The class of dynamics under consideration includes fictitious play and its regularized variants, the best reply dynamics (again, possibly regularized), as well as the dynamics of dual averaging / “follow the regularized leader” (which themselves include as special cases the replicator dynamics and Friedman’s projection dynamics). Our analysis concerns both the actual trajectory of play and its time-average, and we cover potential and monotone games, as well as games with an evolutionarily stable state (global or otherwise). We focus exclusively on games with finite action spaces; nonatomic games with continuous action spaces are treated in detail in Part II.

1. INTRODUCTION

The study of evolutionary dynamics in population games has been the mainstay of evolutionary game theory ever since the inception of the field in the mid-1970’s. Such dynamics are typically derived from biological models or economic microfoundations that express the net growth rate of a type (or strategy) within a population as a function of the relative frequency of all other types. Thus, owing to the flexibility of this general framework, these considerations have generated an immense body of literature; for an introduction to the topic, we refer the reader to the masterful treatments of Hofbauer & Sigmund [20], Weibull [57], and the more recent, panoramic survey of Sandholm [45, 46].

Our paper takes a complementary approach and aims to examine the long-run behavior of *learning procedures* in nonatomic games. While the topic of learning in games with a *finite* number of players is fairly well-studied – see e.g., the classic works of Fudenberg & Levine [13] and Cesa-Bianchi & Lugosi [9] – the same cannot be said for learning in population games. On that account, our goal in this paper is to provide a synthetic treatment for a wide range of dynamics for learning in nonatomic games, unifying along the way a number of existing results in the literature.

We pursue this goal in two parts. In the current paper, we focus exclusively on games with shared, *finite* action spaces; games with *continuous* (and possibly player-dependent)

* TWEAG I/O, PARIS, FRANCE.

† CNRS (LAMSADE, UNIVERSITY OF DAUPHINE-PSL) AND UNIVERSITY OF LIVERPOOL (COMPUTER SCIENCE DEPARTMENT).

‡ UNIV. GRENoble ALPES, CNRS, INRIA, GRENoble INP, LIG, 38000, GRENoble, FRANCE.

‡ INSTITUT DE MATHÉMATIQUES JUSSIEU-PRG, SORBONNE UNIVERSITÉ, UPMC, CNRS UMR 7586.

E-mail addresses: saeed.hadikhanloo@tweag.io, rida.laraki@dauphine.fr,
panayotis.mertikopoulos@imag.fr, sylvain.sorin@imj-prg.fr.

2020 *Mathematics Subject Classification.* Primary 91A22, 91A26; secondary 37N40, 68Q32.

Key words and phrases. Learning; nonatomic games; fictitious play; dual averaging; evolutionary stability; Nash equilibrium; variational inequalities.

action spaces are significantly more complicated, so their study is deferred to the companion paper [16].

Overview of the paper. To specify the general framework under consideration one has to:

- a) Describe the family of games to be studied: These are anonymous nonatomic games and two concrete classes thereof, potential and monotone games; the relevant details are provided in Section 2.
- b) Identify the solution concepts involved: We will mainly focus on Nash equilibria and variants of evolutionarily stable states, expressed throughout as solutions of an associated variational inequality; this is also described in Section 2.
- c) Define the dynamics under study: This is the content of Section 3, where we begin by recalling the standard fictitious play procedure of Brown [8] and Robinson [41], its regularized variants [18], and its continuous-time analogue, the best reply dynamics of Gilboa & Matsui [15]. Subsequently, we introduce a general class of dynamics known as *dual averaging* – or “*follow the regularized leader*” – and which contains as special cases the replicator dynamics of Taylor & Jonker [55] and the projection dynamics of Friedman [12].

The dynamics we examine evolve in both discrete and continuous time: the continuous-time analysis is presented in Section 4 and comprises the template for the discrete-time analysis that follows in Section 5. In both cases, our analysis builds on previous work and results by Monderer & Shapley [37], Sandholm [44], Hofbauer & Sigmund [21], Hofbauer & Sandholm [18, 19], Hofbauer et al. [23] and Mertikopoulos & Sandholm [32, 33]. In this context, some of the new results that we present in our paper can be summarized as follows:

- a) We establish a precise link between (regularized) fictitious play and dual averaging: informally, fictitious play best-responds to the average population state whereas dual averaging best-responds to the average population payoff profile.
- b) We determine the equilibrium convergence properties of vanishingly regularized fictitious play in continuous time.
- c) We provide a range of conditions for the convergence of dual averaging trajectories to evolutionarily stable states, in both discrete and continuous time (and possibly with a variable learning rate).

Finally, in Section 6, we discuss several extensions of these results beyond single-population nonatomic games, and we provide a preview of the issues that arise in nonatomic games with continuous action spaces – the topic of interest of the follow-up paper [16].

2. PRELIMINARIES

2.1. Nonatomic games. A (single-population) *nonatomic game* is defined by the following primitives:

- A *population of players*. This is modeled by the unit interval $\mathcal{I} = [0, 1]$ endowed with the Lebesgue measure μ . The nonatomic aspect means that the behavior of a set of players with zero measure has no impact on the game (see below).
- A set of *pure strategies* $\mathcal{A} = \{1, \dots, A\}$. Following Schmeidler [47], the players’ *strategic behavior* is defined as a measurable function $\chi: \mathcal{I} \rightarrow \mathcal{A}$, with $\chi(i) \in \mathcal{A}$ denoting the pure strategy of player i . Given a pure strategy $\alpha \in \mathcal{A}$, the inverse image $\chi^{-1}(\alpha) = \{i \in \mathcal{I} : \chi(i) = \alpha\}$ will be called the set of “ α -strategists”, and the corresponding pushforward measure $x := \chi\#\mu \equiv \mu \circ \chi^{-1}$ on \mathcal{A} will be called the *state* of the population. This measure will be the main variable of interest; for convenience, we will view x as an

element of the simplex $\mathcal{X} := \Delta(\mathcal{A})$ and we will write $x_\alpha = \mu(\chi^{-1}(\alpha))$ for the *relative frequency* of α -strategists in the population.

- A family of continuous *payoff functions* $u_\alpha: \mathcal{X} \rightarrow \mathbb{R}$, $\alpha \in \mathcal{A}$. Specifically, $u_\alpha(x)$ denotes the payoff to α -strategists when the state of the population is $x \in \mathcal{X}$. Collectively, we will also write $F(x) = (u_\alpha(x))_{\alpha \in \mathcal{A}}$ for the population's *payoff vector* at state x .

The above definition specifies a *population game* $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, F)$, cf. Sandholm [45] and references therein. In this definition, players are *anonymous* in the sense that their payoffs factor through the state of the population and do not depend on specific players choosing specific strategies. For more intricate classes of nonatomic games (possibly involving idiosyncratic components), see Schmeidler [47] and Mas-Colell [27].

2.2. Equilibria, stability, and variational inequalities. In population games, an equilibrium is defined as a population state $x^* \in \mathcal{X}$ in which almost all players are satisfied with their choice of strategy. Formally, $x^* \in \mathcal{X}$ is an *equilibrium* of \mathcal{G} if

$$u_\alpha(x^*) \geq u_\beta(x^*) \quad \text{for all } \alpha \in \text{supp}(x^*), \beta \in \mathcal{A}, \quad (\text{Eq})$$

where $\text{supp}(x^*)$ denotes the support of x^* . Equivalently, this requirement for $x^* \in \mathcal{X}$ can be stated as a *Stampacchia variational inequality* (SVI) of the form

$$\langle F(x^*), x - x^* \rangle \leq 0 \quad \text{for all } x \in \mathcal{X}. \quad (\text{SVI})$$

Equilibria can also be seen as fixed points of the continuous endomorphism $x \mapsto \Pi_{\mathcal{X}}(x + F(x))$, where $\Pi_{\mathcal{X}}$ denotes the Euclidean projector on \mathcal{X} . By standard fixed point arguments, it follows that the game's set of equilibria is nonempty; we will denote this set as $\text{Eq}(\mathcal{G})$.

Remark. In operator theory and optimization, the direction of (SVI) is typically reversed because optimization problems are usually formulated as *minimization* problems. The utility maximization viewpoint is more common in the literature on population games, so we will maintain this sign convention throughout.

An alternative characterization of equilibria is by means of the *best-response correspondence* $\text{BR}: \mathcal{X} \rightrightarrows \mathcal{X}$, defined here as

$$\text{BR}(x) := \arg \max_{x' \in \mathcal{X}} \langle F(x), x' \rangle, \quad (1)$$

We then have the following string of equivalences:

$$x^* \in \text{BR}(x^*) \iff \text{supp}(x^*) \subseteq \text{BR}(x^*) \iff x^* \in \text{Eq}(\mathcal{G}) \quad (2)$$

i.e., the equilibria of the game correspond to the fixed points of BR . One can easily show that, under our assumptions, this correspondence has a closed graph and convex compact non-empty values. By Kakutani's fixed point theorem, $\text{Eq}(\mathcal{G})$ – that is, the set of fixed points of BR – is nonempty and compact.

We will also consider the associated *Minty variational inequality* (MVI) for $x^* \in \mathcal{X}$, namely

$$\langle F(x), x - x^* \rangle \leq 0 \quad \text{for all } x \in \mathcal{X}. \quad (\text{MVI})$$

For concreteness, we will write $\text{SVI}(\mathcal{G})$ and $\text{MVI}(\mathcal{G})$ for the set of solutions of (SVI) and (MVI) respectively. In terms of structural properties, $\text{MVI}(\mathcal{G})$ is convex and, under our continuity assumption for F , it is straightforward to verify that $\text{MVI}(\mathcal{G}) \subseteq \text{SVI}(\mathcal{G}) \equiv \text{Eq}(\mathcal{G})$. In this regard, (MVI) represents an “equilibrium refinement” criterion for \mathcal{G} (see also Section 2.3 below).

In more detail, (MVI) is intimately related to the seminal notion of *evolutionary stability* that was introduced by Maynard Smith & Price [28] and formed the starting point of

evolutionary game theory. To make this link precise, recall that a population state $x^* \in \mathcal{X}$ is an *evolutionarily stable state* of \mathcal{G} [28, 20, 45], if

$$\langle F(\delta x + (1 - \delta)x^*), x - x^* \rangle < 0 \quad (\text{ESS})$$

for all sufficiently small $\delta > 0$ and all states $x \neq x^*$. Analogously, $x^* \in \mathcal{X}$ is a *neutrally stable state* (NSS) of \mathcal{G} if (ESS) holds as a weak inequality, i.e.,

$$\langle F(\delta x + (1 - \delta)x^*), x - x^* \rangle \leq 0 \quad (\text{NSS})$$

for all sufficiently small $\delta > 0$ and all $x \in \mathcal{X}$. Finally, x^* is said to be a *globally evolutionarily stable state* (resp. *globally neutrally stable state*) of \mathcal{G} if (ESS) [resp. (NSS)] holds for all $\delta \in (0, 1]$. In obvious notation, we will write $\text{ESS}(\mathcal{G})$, $\text{NSS}(\mathcal{G})$, $\text{GESS}(\mathcal{G})$ and $\text{GNSS}(\mathcal{G})$ for the respective set of states of \mathcal{G} (by definition $\text{GESS}(\mathcal{G})$ consists of at most a single point).

Characterizations and implications. A concise characterization of these notions of stability was derived by Taylor [54] and Hofbauer et al. [22], who showed that $x^* \in \text{ESS}(\mathcal{G})$ if and only if there exists a neighborhood \mathcal{U} of x^* in \mathcal{X} such that

$$\langle F(x), x - x^* \rangle < 0 \quad \text{for all } x \in \mathcal{U} \setminus \{x^*\}, \quad (3)$$

and, furthermore, if we can take $\mathcal{U} = \mathcal{X}$ above, then $\text{GESS}(\mathcal{G}) = \{x^*\}$. For the corresponding neutral versions, we have $x^* \in \text{NSS}(\mathcal{G})$ if and only if (3) holds on \mathcal{U} as a weak inequality, and $x^* \in \text{GNSS}(\mathcal{G})$ if and only if we can take $\mathcal{U} = \mathcal{X}$ in this case. This characterization of $\text{GNSS}(\mathcal{G})$ corresponds precisely to the solution set of (MVI); in particular, we have the following string of inclusions:

$$\begin{array}{ccccccc} \text{GESS} & \subseteq & \text{GNSS} & \equiv & \text{MVI} & & \\ & \text{I}\cap & & \text{I}\cap & & \text{I}\cap & \\ \text{ESS} & \subseteq & \text{NSS} & \subseteq & \text{Eq} & \equiv & \text{SVI} \end{array} \quad (4)$$

In general, the above inclusions are all one-way; in the next section, we discuss two important cases where some of them become equivalences.

2.3. Classes of games. The following classes of games will play an important role in the sequel:

- (1) *Potential games*: these are games that admit a *potential function*, i.e., a function $P: \mathcal{X} \rightarrow \mathbb{R}$ such that $F = \nabla P$ on \mathcal{X} , or, more generally,

$$P'(x; x' - x) = \langle F(x), x' - x \rangle \quad \text{for all } x, x' \in \mathcal{X}, \quad (\text{Pot})$$

where $P'(x; z) := \lim_{t \rightarrow 0^+} [P(x + tz) - P(x)]/t$ denotes the one-sided directional derivative of P at x along z . If \mathcal{G} is a potential game, any local maximum of P is an equilibrium of \mathcal{G} , and any *strict* local maximum of P is an ESS of \mathcal{G} . For a detailed treatment of potential games in population games extending the initial study of Monderer & Shapley [37], see Sandholm [44, 45].

- (2) *Monotone games*: these are games that satisfy the monotonicity condition

$$\langle F(x') - F(x), x' - x \rangle \leq 0 \quad \text{for all } x, x' \in \mathcal{X}. \quad (\text{Mon})$$

This condition implies that every equilibrium of \mathcal{G} is neutrally stable, so we get the following equivalences:

$$\text{MVI}(\mathcal{G}) \equiv \text{GNSS}(\mathcal{G}) = \text{NSS}(\mathcal{G}) = \text{Eq}(\mathcal{G}) \equiv \text{SVI}(\mathcal{G}) \quad (5)$$

In this class, the existence of equilibria relies only on the minmax theorem [36].

Going further, a game is called *strictly monotone* if (Mon) holds as a strict inequality for all $x' \neq x$. In this case, we have the stronger equivalence

$$\text{GESS}(\mathcal{G}) = \text{ESS}(\mathcal{G}) = \text{Eq}(\mathcal{G}), \quad (6)$$

i.e., all inclusions in (4) become two-way. Moreover, $\text{Eq}(\mathcal{G})$ is a singleton which is also the unique GESS of the game.

The intersection of potential and monotone (resp. *strictly* monotone) games occurs when \mathcal{G} admits a *concave* (resp. *strictly concave*) potential. The most important example of such games is the class of nonatomic congestion games that arise in routing and transportation theory, cf. Dafermos [11].

2.4. Regularization. In the sequel, we will often consider “regularized” best responses that are single-valued. These are defined as follows:

Definition 2.1. A *regularizer on \mathcal{X}* is a function $h: \mathbb{R}^A \rightarrow \mathbb{R} \cup \{\infty\}$ such that:

- (1) h is supported on \mathcal{X} , i.e., $\text{dom } h = \{x \in \mathbb{R}^A : h(x) < \infty\} = \mathcal{X}$.
- (2) h is strictly convex and continuous on \mathcal{X} .

Given a regularizer on \mathcal{X} , we further define:

- a) The associated *choice map* $Q_h: \mathbb{R}^A \rightarrow \mathcal{X}$ given by

$$Q_h(v) := \arg \max_{x' \in \mathcal{X}} \{\langle v, x' \rangle - h(x')\}. \quad (7a)$$

- b) The *regularized best response* to a population state $x \in \mathcal{X}$:

$$\text{BR}_h(x) := \arg \max_{x' \in \mathcal{X}} \{F(x), x'\} - h(x') = Q_h(F(x)). \quad (7b)$$

Remark 2.1. By our assumptions for h (continuity and strict convexity), Q_h and BR_h are both well-defined and single-valued as correspondences. Moreover, by the strict convexity of h , it follows that the convex conjugate $h^*(y) = \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\}$ of h is differentiable and satisfies

$$Q_h(v) = \nabla h^*(v) \quad \text{for all } v \in \mathbb{R}^A. \quad (8)$$

Remark 2.2. By construction, BR_h (approximately) best-responds to strategies, while Q_h best-responds to payoff vectors. For this reason, we will sometimes refer to Q_h as a “payoff-based” regularized best response (as opposed to “strategy-based” regularized best response for BR_h).

Moving forward, for part of our analysis, we will also require one (or both) of the regularity conditions below:

- (1) h is *K-strongly convex* on \mathcal{X} , i.e.,

$$h(\lambda x + (1 - \lambda)x') \leq \lambda h(x) + (1 - \lambda)h(x') - \frac{K}{2} \lambda(1 - \lambda) \|x' - x\|^2 \quad (\text{StrCvx})$$

for all $x, x' \in \mathcal{X}$ and all $\lambda \in [0, 1]$.

- (2) The subdifferential ∂h of h admits a *continuous selection*; specifically, writing $\mathcal{X}_h := \text{dom } \partial h = \{x : \partial h(x) \neq \emptyset\}$ for the “prox-domain” of h , we posit that there exists a continuous map $\tilde{\nabla} h: \mathcal{X}_h \rightarrow \mathbb{R}^A$ such that

$$\tilde{\nabla} h(x) \in \partial h(x) \quad \text{for all } x \in \mathcal{X}_h. \quad (\text{Diff})$$

By replacing best responses with regularized best responses in the definition of equilibria, we also define the notion of a *regularized equilibrium* as follows:

Definition 2.2. Given a regularizer h on \mathcal{X} and a regularization weight $\varepsilon > 0$, we define an ε -regularized equilibrium of \mathcal{G} as a profile $x^* \in \mathcal{X}$ such that $x^* \in \text{BR}_{\varepsilon h}(x^*)$. The set of ε -regularized equilibria of \mathcal{G} will be denoted as $\text{Eq}_{\varepsilon h}(\mathcal{G})$.

Under (Diff), $x^* \in \mathcal{X}$ is an ε -regularized equilibrium if and only if

$$\langle F(x^*) - \varepsilon \tilde{\nabla} h(x^*), x^* - x \rangle \geq 0 \quad \text{for all } x \in \mathcal{X}. \quad (9)$$

In this regard, ε -regularized equilibria can be seen as Nash equilibria of the “ ε -regularized” game $\mathcal{G}_\varepsilon \equiv \mathcal{G}_\varepsilon(\mathcal{A}, F_\varepsilon)$ with payoff profile $F_\varepsilon(x) := F(x) - \varepsilon \tilde{\nabla} h(x)$; more concisely, $\text{Eq}_\varepsilon(\mathcal{G}) = \text{Eq}(\mathcal{G}_\varepsilon)$.

Equilibria and regularized equilibria are related by the next hemicontinuity property, itself a consequence of the maximum theorem of Berge [6].

Lemma 1. *Let $x_\varepsilon^* \in \text{Eq}_\varepsilon(\mathcal{G})$, $\varepsilon > 0$, be a family of regularized equilibria of \mathcal{G} . Then, in the limit $\varepsilon \rightarrow 0$, any accumulation point of x_ε^* is an equilibrium of \mathcal{G} .*

We close this section with a concrete illustration of the above concepts:

Example 2.1 (Logit choice). An important example of regularization corresponds to the entropic regularizer $h(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha \log x_\alpha$ which, by a standard calculation, leads to the choice map

$$\Lambda(v) = \frac{(e^{v_1}, \dots, e^{v_A})}{e^{v_1} + \dots + e^{v_A}}. \quad (10)$$

This map is commonly referred to as the *logit choice map* [29], and the corresponding fixed points $x^* = \Lambda(F(x^*)/\varepsilon)$ are known as *logit equilibria*; for a detailed discussion, see van Damme [56], McKelvey & Palfrey [30], Sandholm [45], and references therein. It is also well known that the entropic regularizer is 1-strongly convex relative to the L_1 -norm, so Λ is 1-Lipschitz continuous, cf. Shalev-Shwartz [48]. \diamond

3. DYNAMICS

In this section, we introduce a wide range of dynamics for learning in nonatomic games, in both continuous and discrete time. The unifying principle of the dynamics under study is that agents play a best response – regularized or otherwise – to some version of the history of play: the empirical population frequencies, or the players’ average gains. For historical reasons, we begin with the discrete-time framework in Section 3.1, and we present the corresponding continuous-time dynamics in Section 3.2.

3.1. Discrete-time processes. In the discrete-time case, we will consider the evolution of the population state $X_n \in \mathcal{X}$ over a series of epochs $n = 1, 2, \dots$ (corresponding for example to generations in an evolutionary context). In turn, this population state is determined by a recursive update rule – or *algorithm* – that defines the dynamics in question. Depending on the manner of “best-responding” and the signal that the agents are best-responding to, we have the following dynamics.

Fictitious play. Dating back to Brown [8] and Robinson [41], *fictitious play* (FP) posits that agents best-respond to the *empirical frequency* (or “*time average*”) $\bar{X}_n = (1/n) \sum_{k=1}^n X_k$ of their population state. Concretely, this means that the population state evolves according to the process

$$X_{n+1} \in \text{BR}(\bar{X}_n) = \arg \max_{x \in \mathcal{X}} \langle F(\bar{X}_n), x \rangle. \quad (\text{FP})$$

In terms of empirical frequencies, we have the recursive dynamics

$$\bar{X}_{n+1} - \bar{X}_n \in \frac{1}{n+1} [\text{BR}(\bar{X}_n) - \bar{X}_n]. \quad (11)$$

In contrast to the sequence of actual population states $X_n \in \mathcal{X}$, $n = 1, 2, \dots$, the stationary points of the time-averaged process (11) are straightforward to characterize. Indeed, any fixed point x^* of (11) satisfies

$$0 \in \text{BR}(x^*) - x^*, \quad (12)$$

i.e., x^* is stationary under (11) if and only if it is an equilibrium of \mathcal{G} .

Regularized fictitious play. A variant of the original fictitious play algorithm is the “regularized” version in which best responses are replaced by *regularized* best responses. In particular, given an underlying regularizer h on \mathcal{X} and a corresponding regularization weight $\varepsilon > 0$, the *regularized fictitious play* (RFP) process is defined as

$$X_{n+1} = \text{BR}_{\varepsilon h}(\bar{X}_n) = \arg \max_{x \in \mathcal{X}} V_{\varepsilon h}(\bar{X}_n, x) \quad (\text{RFP})$$

with

$$V_{\varepsilon h}(x', x) = \langle F(x'), x \rangle - \varepsilon h(x) \quad (13)$$

Then, as before, in terms of empirical frequencies, we have

$$\bar{X}_{n+1} - \bar{X}_n \in \frac{1}{n+1} [\text{BR}_{\varepsilon h}(\bar{X}_n) - \bar{X}_n]. \quad (14)$$

As in the case of (FP), the stationary points of the time-averaged dynamics (RFP) are given by the fixed point equation

$$0 \in \text{BR}_{\varepsilon h}(x^*) - x^*, \quad (15)$$

i.e., they are precisely the ε -regularized equilibria of \mathcal{G} (cf. Definition 2.2).

Remark. In the literature, this process is sometimes referred to as “smooth” or “perturbed” fictitious play, cf. Fudenberg & Levine [14], Hofbauer & Sandholm [18] and references therein. This difference in terminology is owed to a different set of assumptions for h , which gives rise to interior-valued choice maps $Q: \mathbb{R}^A \rightarrow \mathcal{X}$.

Vanishingly regularized fictitious play. When the regularization weight in (RFP) decreases over time, we obtain the *vanishingly regularized fictitious play* (VRFP) dynamics of Benaïm & Faure [3], viz.

$$X_{n+1} = \text{BR}_{\varepsilon_n h}(\bar{X}_n) = \arg \max_{x \in \mathcal{X}} V_{\varepsilon_n h}(\bar{X}_n, x) \quad (\text{VRFP})$$

with $\varepsilon_n > 0$ for all n and $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. In terms of empirical frequencies, we then have

$$\bar{X}_{n+1} - \bar{X}_n \in \frac{1}{n+1} [\text{BR}_{\varepsilon_n h}(\bar{X}_n) - \bar{X}_n]. \quad (16)$$

Because this process is non-autonomous (ε_n depends explicitly on n), it is no longer meaningful to discuss its rest points.

Dual averaging. Dually to the above, instead of best-responding to the aggregate history of *play* (perfectly or approximately), we can also consider the case where agents play a “regularized best response to their aggregate *payoff*”. In our context, this gives rise to the *dual averaging* dynamics

$$X_{n+1} = Q_h \left(\eta_n \sum_{k=1}^n F(X_k) \right), \quad (17)$$

with Q_h defined in (7a). Then, in iterative form, we get the update rule

$$\begin{aligned} S_n &= S_{n-1} + F(X_n) \\ X_{n+1} &= Q_h(\eta_n S_n) \end{aligned} \quad (\text{DA})$$

where $\eta_n > 0$ is a “learning rate” parameter in the spirit of Nesterov [40]. In the recursive formulation (DA) of the dynamics, the initialization $S_0 = 0$ is the most common one but, otherwise, it can be arbitrary.

Remark. In the literature on online learning and online optimization, the process (DA) is known as “follow the regularized leader” (FTRL) [48, 49]. The terminology “dual averaging” is due to Nesterov [40] and is more common in the theory of variational inequalities and offline convex optimization. We adopt the latter to highlight the link of our work to variational inequalities.

Links between the dynamics: the case of random matching. To illustrate the relations between the various learning processes discussed above, it will be instructive to consider nonatomic games generated by random matching [57, 45]. In such games, players are randomly matched to play a symmetric two-player game with payoff matrix $A \in \mathbb{R}^{A \times A}$, so the population’s mean payoff profile is $F(x) = Mx$ for all $x \in \mathcal{X}$. As a result, the dynamics (DA) may be rewritten as

$$X_{n+1} = Q_h \left(\eta_n \sum_{k=1}^n F(X_k) \right) = Q_h(n\eta_n F(\bar{X}_n)) = Q_{h/(n\eta_n)}(F(\bar{X}_n)), \quad (18)$$

which immediately allows us to recover the variants of fictitious play as follows:

- (1) For (VRFP): $\eta_n = 1/(n\varepsilon_n)$ for a sequence of weights $\varepsilon_n > 0$, $\lim_{n \rightarrow \infty} \varepsilon_n = 0$.
- (2) For (RFP): $\eta_n = 1/(n\varepsilon)$ for some fixed $\varepsilon > 0$.
- (3) For (FP): this corresponds to the limiting cases “ $n\eta_n = \infty$ ” or “ $h = 0$ ”. [Of course, these parameter choices are not formally allowed in (DA), so this limit is informal.]

Remark. The above is meaningful only when $F(x)$ is linear in x ; however, the relation between ε and η will be seen to underlie a large part of the sequel.

3.2. Continuous-time dynamics. We now proceed to define the corresponding continuous-time dynamics for each of the processes described above.

Best reply dynamics. The autonomous formulation (11) of fictitious play can be seen as an Euler discretization of the *best reply dynamics* (BRD) of Gilboa & Matsui [15], namely

$$\dot{\mu}_t \in \text{BR}(\mu_t) - \mu_t. \quad (\text{BRD})$$

In the above, $\mu_t \in \mathcal{X}$ is the continuous-time analogue of the empirical mean process \bar{X}_n in discrete time; we use the notation μ instead of x to highlight this link. Clearly, the rest points of (BRD) are again described by the fixed point equation (12), i.e., μ^* is stationary under (BRD) if and only if it is an equilibrium of \mathcal{G} .

Regularized best reply dynamics. Working as above, the *regularized best reply dynamics* (RBRD) are defined as

$$\dot{\mu}_t = \text{BR}_{\varepsilon h}(\mu_t) - \mu_t. \quad (\text{RBRD})$$

As in the case of (RFP), the rest points of (RBRD) are characterized by the fixed point equation (15), i.e., μ^* is stationary under (RBRD) if and only if it is an ε -regularized equilibrium of \mathcal{G} .

Example 3.1 (The logit dynamics). Building on the logit choice map discussed in Section 2 (cf. Example 2.1), the associated regularized best reply dynamics are known as *logit dynamics*, and are given by

$$\dot{\mu}_t = \Lambda(F(\mu_t)/\varepsilon) - \mu_t. \quad (19)$$

Vanishing regularized best reply dynamics. If we allow the regularization weight ε to vary in (RBRD), we obtain the non-autonomous dynamics

$$\dot{\mu}_t = \text{BR}_{\varepsilon_t h}(\mu_t) - \mu_t, \quad (\text{VBRD})$$

where the “V” indicates again that $\varepsilon_t \rightarrow 0$ as $t \rightarrow \infty$.

The dual averaging dynamics. Finally, given a regularizer h on \mathcal{X} as above, the dynamics of dual averaging in continuous time can be written as

$$\begin{aligned} \dot{y}_t &= F(x_t) \\ x_t &= Q_h(\eta_t y_t) \end{aligned} \quad (\text{DAD})$$

where $\eta_t \geq 0$ is a time-varying learning parameter, which we assume throughout to be non-increasing. More compactly, the above can also be written as

$$\dot{y}_t = F(Q_h(\eta_t y_t)) \quad (20a)$$

or

$$x_t = Q_h\left(\eta_t \int_0^t F(x_s) ds\right), \quad (20b)$$

depending on whether we take a differential- or integral-based viewpoint. Specifically, we note that (20a) is an autonomous ordinary differential equation evolving in the dual, payoff space of the game; by contrast, (20b) is an integral equation stated directly in the primal, strategy space of the game. We also note that, unlike (BRD), (RBRD) and (VBRD), the dynamics (DAD) are stated in terms of the actual population state x_t at time t , not its empirical mean μ_t (which is the driving variable of the previous dynamics).

Example 3.2. If $\eta_t \equiv 1$ and h is the entropic regularizer on \mathcal{X} (cf. Example 2.1), the dynamics (DAD) yield the replicator dynamics of Taylor & Jonker [55], viz.

$$\dot{x}_{\alpha,t} = x_{\alpha,t}[u_{\alpha}(x_t) - \langle F(x_t), x_t \rangle]. \quad (\text{RD})$$

In words, (RD) indicates that the per capita growth rate of the population of α -strategists is proportional to the difference between the payoff that they experience and the mean population payoff. For a detailed presentation of this derivation in different contexts, see Rustichini [43], Sorin [50, 51], Mertikopoulos & Moustakas [31] and references therein. \diamond

Example 3.3. If $\eta_t \equiv 1$ and $h(x) = (1/2) \sum_{\alpha \in \mathcal{A}} x_{\alpha}^2$, the integral form (20b) of (DAD) becomes

$$x_t = \Pi\left(\int_0^t F(x_s) ds\right) \quad (21)$$

where Π denotes the Euclidean projector on \mathcal{X} . The differential form of (21) is more delicate to describe because x_t may enter and exit different faces of \mathcal{X} in perpetuity. However, if we focus on an open interval $\mathcal{T} \subseteq \mathbb{R}$ over which the support of x_t remains constant, it can be shown that (21) follows the projection dynamics of Friedman [12], viz. for all $t \in \mathcal{T}$ we have

$$\dot{x}_{\alpha,t} = \begin{cases} u_{\alpha}(x) - \frac{1}{|\text{supp}(x_t)|} \sum_{\beta \in \text{supp}(x_t)} u_{\beta}(x) & \text{if } \alpha \in \text{supp}(x_t), \\ 0 & \text{otherwise.} \end{cases} \quad (\text{PD})$$

For the details of this derivation, see Mertikopoulos & Sandholm [32]. \diamond

Remark. For posterity, we note that both regularizers used in the above examples satisfy (StrCvx), (Diff), and the technical condition (Rec) that we describe later in the paper.

Random matching in continuous time. Following Hofbauer et al. [23], the various dynamics presented so far can be linked as follows when F is linear and $\mu_t = \frac{1}{t} \int_0^t x_s ds$ for $t > 0$:

$$\begin{aligned} x_t &= \arg \max_{x \in \mathcal{X}} \left\{ \eta_t \int_0^t \langle F(x_s), x \rangle ds - h(x) \right\} = \arg \max_{x \in \mathcal{X}} \{ \eta_t t \langle F(\mu_t), x \rangle - h(x) \} \\ &= \text{BR}_{[1/(t\eta_t)]h}(\mu_t). \end{aligned} \quad (22)$$

It thus follows that the empirical distribution of play μ_t under (DAD) follows the dynamics

$$\dot{\mu}_t = \frac{\text{BR}_{\varepsilon_t h}(\mu_t) - \mu_t}{t} \quad (23)$$

with $\varepsilon_t = 1/(t\eta_t)$. Hence, after the change of time $t \leftarrow \log t$, we get (RBRD) with a variable regularization parameter $\varepsilon_t = 1/(t\eta_t)$.

4. ANALYSIS AND RESULTS IN CONTINUOUS TIME

We now proceed to present our results for the class of dynamics under study. We begin with the continuous-time analysis which provides a template for the discrete-time analysis in the next section. We prove that under some assumptions depending on the dynamical system and the class of games under consideration (potential, monotone or strictly monotone), the dynamics converge to the set of equilibria or or – regularized equilibria – of the game.¹

4.1. Best reply dynamics. Our first result concerns the best reply dynamics (BRD). Albeit slightly more general, the results of this subsection and the next essentially follow Hofbauer & Sandholm [18, 19].

Theorem 4.1. *Suppose that one of the following holds:*

- (a) \mathcal{G} is potential.
- (b) \mathcal{G} is monotone and F is C^1 -smooth.

Then every solution orbit μ_t of (BRD) converges to $\text{Eq}(\mathcal{G})$.

Proof. Both cases rely on establishing a suitable Lyapunov function $E(t)$ for (BRD):

- (1) In the potential case, we simply take $E(t) = P(\mu_t)$.
- (2) In the monotone case, $E(t) = \text{Gap}(\mu_t)$ where Gap denotes the function

$$\text{Gap}(x) := \max_{x' \in \mathcal{X}} \langle F(x), x' - x \rangle = \langle F(x), \text{BR}(x) - x \rangle \quad (24)$$

Observe that $\text{Gap}(x) \geq 0$ with equality if and only if $x \in \text{Eq}(\mathcal{G})$.

We now proceed on a case-by-case basis:

Potential case. From (Pot) we obtain that

$$\frac{d}{dt} P(\mu_t) = \lim_{\delta \rightarrow 0^+} [P(\mu_t + \delta \dot{\mu}_t) - P(\mu_t)] / \delta = \langle F(\mu_t), \dot{\mu}_t \rangle \quad (25)$$

Since $\dot{\mu}_t \in \text{BR}(\mu_t) - \mu_t$, it follows that $\frac{d}{dt} P(\mu_t) = \text{Gap}(\mu_t) \geq 0$ with equality if and only if $\dot{\mu}_t = 0$. Consequently, from Lyapunov's first method [see e.g., 20, Theorem 2.6.1] we deduce that μ_t converges to the set of rest points of (BRD), which is $\text{Eq}(\mathcal{G})$.

¹Recall that a family $(x_t)_{t \geq 0}$ (resp. a sequence (X_n) , $n = 1, 2, \dots$) converges to a set B if the set of accumulation points $\mathbf{A}[x_t]$ (resp. $\mathbf{A}[x_n]$), as $t \rightarrow \infty$ (resp $n \rightarrow \infty$), satisfies: $\mathbf{A}[x_t] \subset B$ (resp. $\mathbf{A}[x_n] \subset B$).

Monotone case. Let JF denote the Jacobian of F . By the envelope theorem we get:

$$\frac{d}{dt} \text{Gap}(\mu_t) = \langle F(\mu_t), -\dot{\mu}_t \rangle + \dot{\mu}_t JF(\mu_t) (\text{BR}(\mu_t) - \mu_t) = -\text{Gap}(\mu_t) + \dot{\mu}_t JF(\mu_t) \dot{\mu}_t \quad (26)$$

As \mathcal{G} is monotone, $(x' - x)JF(x)(x' - x) \leq 0$ for all $x, x' \in \mathcal{X}$, so $\dot{\mu}_t JF(\mu_t) \dot{\mu}_t \leq 0$. Consequently:

$$\frac{d}{dt} \text{Gap}(\mu_t) \leq -\text{Gap}(\mu_t). \quad (27)$$

Hence $\text{Gap}(\mu_t)$ decreases exponentially to 0, so μ_t converges to $\text{Eq}(\mathcal{G})$. \blacksquare

4.2. Regularized best reply dynamics. In this section we will assume that the regularizer h satisfies (Diff). In a slight abuse of notation, we will skip the dependence on h and write F_ε for $F_{\varepsilon h} := F - \varepsilon \tilde{\nabla} h$, see (9), with $\tilde{\nabla} h$ defined as in (Diff). We also write $\text{BR}_\varepsilon(x)$ for $\text{BR}_{\varepsilon h}(x)$ and so on. The reason is that the regularizer h will be fixed throughout while the weight ε will be allowed to vary in the next subsection.

Lemma 2. *Under (Diff), we have $\langle F_\varepsilon(x), \text{BR}_\varepsilon(x) - x \rangle \geq 0$ for all $x \in \mathcal{X}$, and equality holds if and only if $x = \text{BR}_\varepsilon(x)$.*

Proof. Let $\phi(z) = \langle F(x), z \rangle - \varepsilon h(z)$. Then ϕ is strictly concave and so $\partial\phi$ defines a strictly monotone operator. Thus, $x' = \arg \max_{z \in \mathcal{X}} \phi(z)$ is a solution of the variational inequality (MVI) associated to $\partial\phi$: for all $d \in \partial\phi(x)$: $\langle d, x' - x \rangle \geq 0$ with equality if and only if $x = x'$. As $F_\varepsilon(x) \in \partial\phi(x)$ and $x' = \text{BR}_\varepsilon(x)$ the result follows. \blacksquare

Lemma 3. *Suppose that (Diff) holds. If \mathcal{G} with payoff field F is potential (resp. monotone) then the game \mathcal{G}_ε with payoff function F_ε is potential (resp. strictly monotone).*

Proof. If \mathcal{G} has a potential P , a potential function of \mathcal{G}_ε is $P - \varepsilon h$. If \mathcal{G} is a monotone game, the game \mathcal{G}_ε is strictly monotone because we have $\langle \tilde{\nabla} h(x) - \tilde{\nabla} h(x'), x - x' \rangle \geq 0$ for every $x, x' \in \mathcal{X}$, with equality if and only if $x = x'$. \blacksquare

We now turn to the asymptotic properties of regularized best reply dynamics.

Theorem 4.2. *Suppose that (Diff) is satisfied and assume one of the following holds:*

- (a) \mathcal{G} is potential.
- (b) \mathcal{G} is monotone and F is C^1 -smooth.

Then every solution orbit μ_t of (RBRD) converges to $\text{Eq}_\varepsilon(\mathcal{G})$.

Remark. Observe that when \mathcal{G} is monotone, \mathcal{G}_ε is strictly monotone, so $\text{Eq}_\varepsilon(\mathcal{G})$ is a singleton.

Proof. As above, both cases rely on establishing a suitable Lyapunov function $E(t)$ for (RBRD):

- (1) In the potential case: $E(t) = P_\varepsilon(\mu_t) := P(\mu_t) - \varepsilon h(\mu_t)$.
- (2) In the monotone case:

$$\begin{aligned} E(t) &= \text{Gap}_\varepsilon(\mu_t) = \max_{x' \in \mathcal{X}} \langle F(\mu_t), x' - \mu_t \rangle - \varepsilon(h(x') - h(\mu_t)) \\ &= \max_{x' \in \mathcal{X}} V_\varepsilon(\mu_t, x') - V_\varepsilon(\mu_t, \mu_t) \end{aligned} \quad (28)$$

Potential case. There is $z(t) \in \partial h(\mu_t)$ such that – similarly to (25) – we have:

$$\frac{d}{dt}P_\varepsilon(\mu_t) = \lim_{dt \rightarrow 0^+} [P_\varepsilon(\mu_t + \dot{\mu}_t dt) - P_\varepsilon(\mu_t)]/dt = \langle F(\mu_t), \dot{\mu}_t \rangle - \varepsilon \langle z(t), \dot{\mu}_t \rangle \quad (29)$$

As $\dot{\mu}_t = \text{BR}_\varepsilon(\mu_t) - \mu_t$, by Lemma A.1 in the appendix, we have $\langle z(t), \dot{\mu}_t \rangle \leq \langle \tilde{\nabla} h(\mu_t), \dot{\mu}_t \rangle$. Consequently

$$\frac{d}{dt}P_\varepsilon(\mu_t) \geq \langle F(\mu_t), \dot{\mu}_t \rangle - \varepsilon \langle \tilde{\nabla} h(\mu_t), \dot{\mu}_t \rangle = \langle F_\varepsilon(\mu_t), \dot{\mu}_t \rangle \quad (30)$$

By Lemma 2 and the above, we deduce that $\frac{d}{dt}P_\varepsilon(\mu_t) \geq 0$ with equality if and only if $\dot{\mu}_t = 0$. Consequently, by a Lyapunov argument, μ_t converges to $\text{Eq}_\varepsilon(\mathcal{G})$.

Monotone case. Using the envelope theorem as in (26) and, just as above, there is $z(t) \in \partial h(\mu_t)$ such that:

$$\frac{d}{dt}\text{Gap}_\varepsilon(\mu_t) = -\langle F(\mu_t) - \varepsilon z(t), \dot{\mu}_t \rangle + \dot{\mu}_t JF(\mu_t) \dot{\mu}_t \leq -\langle F_\varepsilon(\mu_t), \dot{\mu}_t \rangle + \dot{\mu}_t JF(\mu_t) \dot{\mu}_t \quad (31)$$

The inequality being a consequence of Lemma A.1 as it implies that $\langle z(t), \dot{\mu}_t \rangle \leq \langle \tilde{\nabla} h(\mu_t), \dot{\mu}_t \rangle$. Since \mathcal{G} is monotone, $\dot{\mu}_t JF(\mu_t) \dot{\mu}_t \leq 0$. Consequently, as $\dot{\mu}_t = \text{BR}_\varepsilon(\mu_t) - \mu_t$, by Lemma 2, $\frac{d}{dt}\text{Gap}_\varepsilon(\mu_t) \leq 0$ with equality if and only if $\dot{\mu}_t = 0$. Hence μ_t converges to $\text{Eq}_\varepsilon(\mathcal{G})$. ■

4.3. Vanishing regularized best reply dynamics. A similar technique will be used to obtain convergence to the set of Nash equilibria for vanishing regularized best reply dynamics. The additional difficulty here comes from the fact that the vanishing regularized best reply dynamics is not autonomous and so we should prove by hand the Lyapunov convergence argument. To the best of our knowledge, this is a new result.

Theorem 4.3. *Suppose that (Diff) holds and (VBRD) is run with a smoothly-varying and strictly decreasing regularization weight $\varepsilon_t \rightarrow 0$. Assume further that h takes values in $[0, 1]$, and one of the following holds:*

- (a) \mathcal{G} is potential.
- (b) \mathcal{G} is monotone, F is C^1 -smooth, and $\tilde{\nabla} h$ is bounded.

Then every solution orbit μ_t of (VBRD) converges to $\text{Eq}_\varepsilon(\mathcal{G})$.

Remark. The assumption that h takes value in $[0, 1]$ is only made for notational convenience and does not incur any loss of generality.

Proof. As before, we treat the potential and monotone cases separately.

Potential case. Let $P_{\varepsilon_t}(\mu_t) = P(\mu_t) - \varepsilon_t h(\mu_t)$. Then, a calculation similar to the above implies that there exists $z(t) \in \partial h(\mu_t)$ such that

$$\frac{d}{dt}P_{\varepsilon_t}(\mu_t) = \langle F(\mu_t) - \varepsilon_t z(t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle - \dot{\varepsilon}_t h(\mu_t) \geq \langle F_{\varepsilon_t}(\mu_t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle - \dot{\varepsilon}_t h(\mu_t) \quad (32)$$

The inequality being a consequence of Lemma A.1 as it implies that $\langle z(t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle \leq \langle \tilde{\nabla} h(\mu_t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle$. By Lemma 2, $\langle F_{\varepsilon_t}(x), \text{BR}_{\varepsilon_t}(x) - x \rangle \geq 0$. Since $-\dot{\varepsilon}_t h(\mu_t) \geq 0$ we deduce that $\frac{d}{dt}P_{\varepsilon_t}(\mu_t) \geq 0$ and consequently $P_{\varepsilon_t}(\mu_t)$ converges. As ε_t goes to zero and $h(\cdot)$ is bounded, $P(\mu_t)$ converges as well. We will deduce from this that μ_t converges to $\text{Eq}(\mathcal{G})$. By contradiction, if a limit point $\mu_0 = \lim_k \mu_{t_k}$ is not in $\text{Eq}(\mathcal{G})$ then:

$$\text{Gap}(\mu_0) = \max_{x' \in \mathcal{X}} \langle F(\mu_0), x' - \mu_0 \rangle = \langle F(\mu_0), \text{BR}(\mu_0) - \mu_0 \rangle = \alpha > 0. \quad (33)$$

Recall that, by definition, for all t and μ :

$$\begin{aligned} \text{Gap}_{\varepsilon_t}(\mu) &= \max_{x' \in \mathcal{X}} \langle F(\mu), x' - \mu \rangle - \varepsilon_t (h(x') - h(\mu)) \\ &= \langle F(\mu), \text{BR}_{\varepsilon_t}(\mu) - \mu \rangle - \varepsilon_t (h(\text{BR}_{\varepsilon_t}(\mu)) - h(\mu)) \\ &= \max_{x' \in \mathcal{X}} V_{\varepsilon_t}(\mu, x') - V_{\varepsilon_t}(\mu, \mu). \end{aligned} \quad (34)$$

Also, as h takes values in $[0, 1]$, $|V_{\varepsilon}(x, x') - V_{\varepsilon}(x, x) - (V_{\varepsilon'}(x, x') - V_{\varepsilon'}(x, x))| \leq |\varepsilon - \varepsilon'|$. Thus:

$$|\text{Gap}_{\varepsilon_t}(\mu) - \text{Gap}(\mu)| \leq \varepsilon_t. \quad (35)$$

As in (25) one has for all k and s :

$$\frac{d}{dt} P(\mu_{s+t_k}) = \langle F(\mu_{s+t_k}), \text{BR}_{\varepsilon_{s+t_k}}(\mu_{s+t_k}) - \mu_{s+t_k} \rangle \quad (36)$$

Thus, using (34) then (35) we deduce that:

$$\frac{d}{dt} P(\mu_{s+t_k}) \geq \text{Gap}_{\varepsilon_{s+t_k}}(\mu_{t_k+s}) - \varepsilon_{s+t_k} \geq \text{Gap}(\mu_{t_k+s}) - 2\varepsilon_{s+t_k}. \quad (37)$$

As the function $\mu \rightarrow \text{Gap}(\mu)$ is uniformly continuous and $t \rightarrow \mu_t$ is Lipschitz, recalling that $\text{Gap}(\mu_k) \rightarrow \alpha > 0$, we deduce that there is $\delta > 0$ and K_1 such that for $k \geq K_1$ and all $0 \leq s \leq \delta$ one has:

$$\text{Gap}(\mu_{t_k+s}) \geq \text{Gap}(\mu_{t_k}) - \alpha/4 \geq 3\alpha/4 \quad (38)$$

As ε_t decreases to 0 there is K_2 such that for all $k \geq K_2$, $\varepsilon_{s+t_k} \leq \alpha/4$. From this, (35) and (38), we obtain that $\frac{d}{dt} P(\mu_{s+t_k}) \geq \alpha/2$ for all $0 \leq s \leq \delta$, for all $k \geq \max(K_1, K_2)$. Thus $P(\mu_{\delta+t_k}) - P(\mu_{t_k}) = \int_0^{\delta} \frac{d}{dt} P(\mu_{s+t_k}) ds \geq \delta\alpha/2$, which contradicts the convergence of $P(\mu_t)$.

Monotone case. Using the envelope theorem, and a calculation similar to the above implies that there exists $z(t) \in \partial h(\mu_t)$ such that

$$\frac{d}{dt} \text{Gap}_{\varepsilon_t}(\mu_t) = -\langle F(\mu_t) - \varepsilon_t z(t), \dot{\mu}_t \rangle + \dot{\mu}_t JF(\mu_t) \dot{\mu}_t + \dot{\varepsilon}_t (h(\mu_t) - h(\text{BR}_{\varepsilon_t}(\mu_t))) \quad (39a)$$

$$\leq -\langle F_{\varepsilon_t}(\mu_t), \dot{\mu}_t \rangle + \dot{\mu}_t JF(\mu_t) \dot{\mu}_t + \dot{\varepsilon}_t (h(\mu_t) - h(\text{BR}_{\varepsilon_t}(\mu_t))) \quad (39b)$$

$$\leq -\langle F_{\varepsilon_t}(\mu_t), \dot{\mu}_t \rangle + \dot{\varepsilon}_t (h(\mu_t) - h(\text{BR}_{\varepsilon_t}(\mu_t))) \quad (39c)$$

$$= -\langle F_{\varepsilon_t}(\mu_t), \text{BR}_{\varepsilon_t} - \mu_t \rangle + \dot{\varepsilon}_t (h(\mu_t) - h(\text{BR}_{\varepsilon_t}(\mu_t))). \quad (39d)$$

The first inequality is a consequence of Lemma A.1, which implies that $\langle z(t), \dot{\mu}_t \rangle \leq \langle \tilde{\nabla} h(\mu_t), \dot{\mu}_t \rangle$. The second inequality is a consequence of $\dot{\mu}_t JF(\mu_t) \dot{\mu}_t \leq 0$ as \mathcal{G} is monotone.

We now proceed to prove that $\frac{d}{dt} (\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t) < 0$.

Case 1: If $\text{Gap}_{\varepsilon_t}(\mu_t) = 0$ then $\mu_t = \text{BR}_{\varepsilon_t}(\mu_t) = \text{Eq}_{\varepsilon_t}(\mathcal{G})$ thus $\dot{\mu}_t = 0$. Then (39a) implies that $\frac{d}{dt} \text{Gap}_{\varepsilon_t}(\mu_t) = 0$ and thus $\frac{d}{dt} (\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t) < 0$.

Case 2: If $\text{Gap}_{\varepsilon_t}(\mu_t) > 0$ then since $h(\cdot) \in [0, 1]$ and $\dot{\varepsilon}_t < 0$ we have:

$$\dot{\varepsilon}_t (h(\mu_t) - h(\text{BR}_{\varepsilon_t}(\mu_t))) \leq -\dot{\varepsilon}_t \quad (40)$$

and from Lemma 2 and (39d) we deduce that $\frac{d}{dt} (\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t) < 0$.

Consequently, $\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t$ converges to $\delta \geq 0$ and since ε_t decreases to zero we conclude that $\text{Gap}_{\varepsilon_t}(\mu_t) \rightarrow \delta$. If $\delta > 0$, define:

$$\begin{aligned} \Delta &:= -\langle F_{\varepsilon_t}(\mu_t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle \\ &= -\langle F(\mu_t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle + \varepsilon_t \langle \tilde{\nabla} h(\mu_t), \text{BR}_{\varepsilon_t}(\mu_t) - \mu_t \rangle. \end{aligned} \quad (41)$$

From (34) we further get:

$$\Delta = -\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t(h(\mu_t) - h(BR_{\varepsilon_t}(\mu_t))) + \varepsilon_t\langle \tilde{\nabla}h(\mu_t), BR_{\varepsilon_t}(\mu_t) - \mu_t \rangle \quad (42)$$

As $\text{Gap}_{\varepsilon_t}(\mu_t) \rightarrow \delta$, and by the assumptions, $h(\cdot)$ and $\tilde{\nabla}h$ are bounded and ε_t decreases to 0, we deduce that for t large enough:

$$\Delta = -\langle F_{\varepsilon_t}(\mu_t), BR_{\varepsilon_t}(\mu_t) - \mu_t \rangle \leq -\delta/2. \quad (43)$$

Consequently, for t large enough we deduce from (39d), (40) and the last inequality that

$$\frac{d}{dt}(\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t) \leq -\langle F_{\varepsilon_t}(\mu_t), BR_{\varepsilon_t}(\mu_t) - \mu_t \rangle < -\delta/2, \quad (44)$$

which is impossible because $\text{Gap}_{\varepsilon_t}(\mu_t) + \varepsilon_t \geq 0$. Hence $\delta = 0$ and so μ_t converges to $\text{Eq}(\mathcal{G})$. ■

4.4. Dual averaging dynamics with a constant learning rate. Before stating our main result we establish some useful lemmas. Define the regret associated to a trajectory x_t and a reference point $x \in \mathcal{X}$ as:

$$\text{Reg}_x(t) = \int_0^t \langle F(x_s), x - x_s \rangle ds \quad (45)$$

Lemma 4. *Let x_t be a solution trajectory of (DAD). Then:*

$$\text{Reg}_x(t) \leq \max h - \min h. \quad (46)$$

This is a particular case of a more general bound due to Kwon & Mertikopoulos [26] (see Lemma 7 below). The next more compact proof of this simpler bound follows Bravo & Mertikopoulos [7] and Sorin [52].

Proof. Define for $x \in \mathcal{X}$ and $y \in \mathbb{R}^A$ the Fenchel coupling

$$F_h(x, y) = h(x) + h^*(y) - \langle x, y \rangle. \quad (47)$$

By the Fenchel-Young inequality, $F_h(x, y) \geq 0$. Recall that from (8) and (20b) $x_t = \nabla h^*(y_t) = Q_h(y_t)$ with $y_t = \int_0^t F(x_s) ds$. Then, following Mertikopoulos & Zhou [35, Lemma A.2], a simple derivation gives that for all $x \in \mathcal{X}$:

$$\frac{d}{dt}F_h(x, y_t) = \langle F(x_t), \nabla h^*(y_t) - x \rangle = \langle F(x_t), x_t - x \rangle = -\frac{d}{dt} \text{Reg}_x(t), \quad (48)$$

Consequently, following the reasoning of Sorin [52], we get

$$\text{Reg}_x(t) \leq F_h(x, 0) = h(x) + h^*(0) \leq \max h + h^*(0) = \max h - \min h \quad \blacksquare$$

Lemma 5. *If a solution trajectory $x_t = x^*$ of (DAD) is stationary, then $x^* \in \text{Eq}(\mathcal{G})$.*

Proof. From the previous lemma, for all $t > 0$ and $x \in \mathcal{X}$ one has

$$\langle F(x^*), x - x^* \rangle = \frac{\text{Reg}_x(t)}{t} \leq \frac{\max h - \min h}{t}, \quad (49)$$

Letting $t \rightarrow \infty$ implies $x^* \in \text{Eq}(\mathcal{G})$, as was to be shown. ■

Lemma 6. *If h satisfies (StrCvx), then for all $x \in \mathcal{X}$ and all sequences y_n in \mathbb{R}^A .*

$$F_h(x, y_n) \rightarrow 0 \implies Q_h(y_n) \rightarrow x \quad (50)$$

Proof. See Lemma A.2 in the appendix. ■

We also need to assume the so called ‘‘reciprocity’’ condition which requires that for all $x \in \mathcal{X}$ and all sequences y_n in \mathbb{R}^A :

$$Q_h(y_n) \rightarrow x \implies F_h(x, y_n) \rightarrow 0 \quad (\text{Rec})$$

Taken together, (StrCvx) and (Rec) imply that the sequence $x_n = Q_h(y_n)$ converges to $x \in \mathcal{X}$ if and only if $F_h(x, y_n) \rightarrow 0$. The terminology ‘‘reciprocity’’ is justified by the fact that, under (Rec), the topology induced by the level sets of the Fenchel coupling becomes equivalent to the ordinary topology on \mathcal{X} ; for a precise statement, see Lemma A.2 in Appendix A. This is a subtle – but important – requirement, which is fortunately satisfied by all the standard regularization choices on \mathcal{X} , cf. [34, 35].

We can now state our convergence results for (DAD) when the learning rate is constant.

Theorem 4.4. *Let x_t be a solution trajectory of (DAD) with constant learning rate $\eta_t \equiv 1$. Assume further that one of the following holds:*

- (a) \mathcal{G} is potential and h^* is twice differentiable.
- (b) \mathcal{G} is strictly monotone, and h satisfies (StrCvx) and (Rec).

Then x_t converges to $\text{Eq}(\mathcal{G})$.

Proof. Again, we treat the potential and monotone cases separately.

Potential case. As $y_t = \int_0^t F(x_s) ds$ and $x_t = Q_h(y_t) = \nabla h^*(y_t)$ we obtain that:

$$\dot{x}_t = \langle F(x_t), \nabla^2 h^*(y_t) \rangle. \quad (51)$$

Since h^* is convex, $\nabla^2 h^*$ is semidefinite positive and so:

$$\frac{d}{dt} P(x_t) = \langle F(x_t), \dot{x}_t \rangle = F(x_t) \nabla^2 h^*(y_t) F(x_t) \geq 0, \quad (52)$$

the inequality being strict whenever $\dot{x}_t \neq 0$. As such, x_t converges to $\text{Eq}(\mathcal{G})$.

Monotone case. As \mathcal{G} is strictly monotone, $\text{Eq}(\mathcal{G})$ is a singleton which coincides with x^* the (necessarily unique) GESS of the game. The calculation in (48) gives:

$$\frac{d}{dt} F_h(x^*, y_t) = \langle F(x_t), x_t - x^* \rangle \leq 0 \quad (53)$$

As x^* is a GESS $\frac{d}{dt} F_h(x^*, y_t) \leq 0$, with equality if and only if $x_t = Q_h(y_t) = x^*$. Suppose there is $\delta > 0$ such that $\langle F(x_t), x_t - x^* \rangle < -\delta < 0$ for all t large enough. Then, $F_h(x^*, y_t)$ goes to $-\infty$ which is impossible because $F_h(\cdot) \geq 0$. Thus there is $t_k \rightarrow \infty$ such that $\langle F(x_{t_k}), x_{t_k} - x^* \rangle \rightarrow 0$ and so $x_{t_k} = Q_h(y_{t_k}) \rightarrow x^*$. By the reciprocity condition (Rec), we deduce that $F_h(x^*, y_{t_k}) \rightarrow 0$ since x^* is a GESS. As $\frac{d}{dt} F_h(x^*, y_t) \leq 0$ and $F_h(\cdot) \geq 0$, necessarily $F_h(x^*, y_t)$ converges and by the last argument, its limit is zero. Using the strong convexity of h , Lemma 6 implies that $x_t = Q_h(y_t)$ converges to x^* . ■

There are several things to note here. First, the proof of convergence to the unique equilibrium in strictly monotone games (part b of Theorem 4.4) extends to all games with a GESS. Second, the same proof implies that in all games, ESS (if exist) are asymptotically stable under (DAD) with constant learning rate (e.g., the dynamics converges to an ESS if it starts sufficiently close to it). Third, Theorem 4.4 concerns the ‘‘day-to-day’’ evolution of (DAD), not the corresponding empirical mean. Last, the theorem does not apply to monotone games which are not strict.

To complete the picture, the next proposition shows the convergence of the empirical mean to the set of Nash equilibria in all monotone games.

Proposition 4.1. *Let $x_t = Q_h(y_t)$ be a solution trajectory of (DAD) with constant learning rate $\eta_t \equiv 1$. If the underlying game is monotone, then $\mu_t = \frac{1}{t} \int_0^t x_s ds$ converges to $\text{Eq}(\mathcal{G})$.*

Proof. By Lemma 4 and monotonicity of the game, for all $x \in \mathcal{X}$:

$$\max h - \min h \geq \text{Reg}_x(t) = \int_0^t \langle F(x_s), x - x_s \rangle ds \geq \int_0^t \langle F(x), x - x_s \rangle ds \quad (54)$$

and hence

$$\frac{\max h - \min h}{t} \geq \langle F(x), x - \mu_t \rangle. \quad (55)$$

Consequently, μ_t converges to $\text{MVI}(\mathcal{G})$ which coincides with $\text{Eq}(\mathcal{G})$ as the game is monotone. \blacksquare

4.5. Dual averaging dynamics with a variable learning rate. This subsection extends the results of the previous section on monotone games in two directions: we consider GESSs instead of (strictly) monotone games, and we treat dynamics with a variable learning rate. Other extensions are also possible, but we do not treat them here.

We begin with a technical lemma on the regret incurred by (DAD):

Lemma 7. *If x_t is a solution trajectory of (DAD) then the regret is bounded as follows:*

$$\text{Reg}_x(t) \leq \frac{\max h - \min h}{\eta_t}. \quad (56)$$

Consequently, if $t\eta_t \rightarrow \infty$ and $x_t = x^$ for all $t \geq 0$, then $x^* \in \text{Eq}(\mathcal{G})$.*

Proof. The bound in the regret follows from Kwon & Mertikopoulos [26, Theorem 4.1], and the implication follows as in Lemma 5 above. \blacksquare

Lemma 7 allows to extend proposition 4.1 to variable learning rates.

Proposition 4.2. *Let $x_t = Q_h(y_t)$ be a solution trajectory of (DAD). If the underlying game is monotone and $t\eta_t \rightarrow \infty$ then $\mu_t = \frac{1}{t} \int_0^t x_s ds$ converges to $\text{Eq}(\mathcal{G})$.*

The next result extends part b) of theorem 4.4 to variable learning rates. The proof is a little more complicated.

Theorem 4.5. *Let x_t be a solution trajectory of (DAD) with a non-increasing learning rate. Suppose further that the underlying game \mathcal{G} admits a GESS x^* . If the regularizer satisfies (StrCvx) and (Rec), and the dynamics' learning rate satisfies $\lim_{t \rightarrow \infty} t\eta_t = \infty$, then x_t converges to x^* .*

Proof. Kwon & Mertikopoulos [26, Eq. 29] prove that

$$\frac{d}{dt} \frac{F_h(x^*, \eta_t y_t)}{\eta_t} \leq \langle F(x_t), x_t - x^* \rangle - \frac{\dot{\eta}}{\eta_t^2} (\max h - \min h). \quad (57)$$

Let $\mathcal{F}_\varepsilon = \{y \in \mathbb{R}^A : F_h(x^*, y) \leq \varepsilon\}$ denote the ‘‘Fenchel zone’’ of magnitude ε with respect to x^* . Then it suffices to show that

$$\text{For all } \varepsilon > 0, \text{ there exists } T(\varepsilon) > 0 \text{ such that } \eta_t y_t \in \mathcal{F}_\varepsilon \text{ for all } t \geq T(\varepsilon). \quad (\text{P})$$

Actually this means that $F_h(x^*, \eta_t y_t) \rightarrow 0$. By (StrCvx), (50) holds and thus x_t converges to x^* . The proof of (P) follows from the following two claims:

Claim 1. $\eta_t y_t \in \mathcal{F}_\varepsilon$ infinitely often.

Proof. If $\eta_t y_t \notin \mathcal{F}_\varepsilon$, then by (Rec), there exists a neighbourhood $\mathcal{U}(\varepsilon)$ of x^* such that $x_t \notin \mathcal{U}(\varepsilon)$. By definition of GESS, there is $c(\varepsilon) > 0$ such that $\langle F(x_t), x_t - x^* \rangle < -2c(\varepsilon)$. As $\frac{\eta_t}{\eta_t^*} \rightarrow 0$, there is $T_1(\varepsilon)$ such that for all $t \geq T_1(\varepsilon)$ we have $\frac{\dot{\eta}_t}{\eta_t}(\max h - \min h) < c(\varepsilon)$. Consequently by contradiction to the claim, $\frac{d}{dt} \frac{F_h(x^*, \eta_t y_t)}{\eta_t} < -c(\varepsilon)$ for all t large enough and so $\frac{F_h(x^*, \eta_t y_t)}{\eta_t}$ goes to $-\infty$. Since $F_h(x^*, \eta_t y_t) \geq 0$ this is impossible. \S

Claim 2. There exists $T_2(\varepsilon) \geq T_1(\varepsilon/2)$ such that if $t \geq T_2(\varepsilon)$ and $\eta_t y_t \in \mathcal{F}_\varepsilon$ then there exists $\tau > 0$ such that $\eta_s y_s \in \mathcal{F}_\varepsilon$ for all $s \in [t, t + \tau]$.

Proof. We need to consider two cases separately:

Case a). If $\eta_t y_t \in \mathcal{F}_{\varepsilon/2}$ then obviously, there is $\tau > 0$ such that $\eta_s y_s \in \mathcal{F}_\varepsilon$ for all $s \in [t, t + \tau]$.

Case b). If $\eta_t y_t \in \mathcal{F}_\varepsilon$ but $\eta_t y_t \notin \mathcal{F}_{\varepsilon/2}$ then, from claim 1, $\frac{d}{dt} \frac{F_h(x^*, \eta_t y_t)}{\eta_t} < -c(\varepsilon/2) < 0$. Thus there exists $\tau > 0$ such that for all $s \in [t, t + \tau]$: $\frac{F_h(x^*, \eta_s y_s)}{\eta_s} \leq \frac{F_h(x^*, \eta_t y_t)}{\eta_t}$. As $\eta_t y_t \in \mathcal{F}_\varepsilon$, $F_h(x^*, \eta_t y_t) \leq \varepsilon$ and since $\eta_s \leq \eta_t$, we deduce that $F_h(x^*, \eta_s y_s) \leq \varepsilon$. Thus $\eta_s y_s \in \mathcal{F}_\varepsilon$ for all $s \in [t, t + \tau]$. \S

From Claims 1 and 2, there exists $t \geq \max(T_2(\varepsilon), T(\varepsilon))$, such that $\eta_t y_t$ belongs to \mathcal{F}_ε and remains in it forever. \blacksquare

5. ANALYSIS AND RESULTS IN DISCRETE TIME

In this section, we continue with the analysis of the discrete-time dynamics presented in Section 3.1, starting with fictitious play and its variants in Section 5.1 before moving on to the dynamics of dual averaging in Section 5.2.

5.1. Fictitious play and regularized fictitious play. The convergence of (FP) and (RFP) can be obtained from the continuous-time analysis of Section 4 using the theory of stochastic approximation [2, 4, 5]. Informally, this theory links the asymptotic behavior of differential inclusions of the form

$$\dot{x}_t \in M(x_t) \tag{58}$$

to the limit sets of discrete-time processes satisfying the basic recursion

$$X_{n+1} - X_n \in \gamma_{n+1} M(X_n), \tag{59}$$

where $M: \mathcal{X} \rightrightarrows \mathbb{R}^A$ is an upper semicontinuous, compact-convex-valued correspondence on \mathcal{X} , and $\gamma_n > 0$ is a vanishing step-size sequence.

In our specific context, the role of M is to be played by the best-reply correspondence BR for (FP), and its regularized variant $\text{BR}_{\varepsilon h}$ for (RFP). Concretely, we have:

Proposition 5.1. *Let $\bar{X}_n = (1/n) \sum_{k=1}^n X_k$, $n = 1, 2, \dots$, be the empirical frequency of play under a sequence of population states $X_n \in \mathcal{X}$. Then:*

- (i) *If X_n is generated by (FP), \bar{X}_n is an asymptotic pseudotrajectory of (BRD).*
- (ii) *If X_n is generated by (RFP), \bar{X}_n is an asymptotic pseudotrajectory of (RBRD).*

Informally, the notion of an asymptotic pseudotrajectory (APT) means that \bar{X}_n asymptotically tracks the pseudo-flow of (BRD)/(RBRD) with arbitrary accuracy over windows of arbitrary length; for a precise definition, see Benaïm et al. [4, Definitions I and 4.1]. The proof of Proposition 5.1 follows immediately from the general theory of Benaïm et al. [4, Propositions 1.3, 1.4, and Theorem 4.2], so we do not present it here.

In view of this “asymptotic approximation” property, it is plausible to expect that \bar{X}_n exhibits the same asymptotic behavior as the corresponding continuous-time dynamics. Our next result makes this intuition precise in two classes of games:

- (1) Monotone games, i.e., games that satisfy **(Mon)**.
- (2) Potential games for which the ε -regularized potential $P_\varepsilon = P - \varepsilon h$, $\varepsilon \geq 0$, satisfies the Sard-type condition:

$$P_\varepsilon(\text{Eq}_\varepsilon(\mathcal{G})) \text{ has empty interior.} \quad (\text{S}_\varepsilon)$$

Our main results for **(FP)** and **(RFP)** may then be stated as follows:

Theorem 5.1. *Suppose that \mathcal{G} satisfies **(Mon)** or **(S $_\varepsilon$)** for $\varepsilon = 0$. Then the empirical frequency of play \bar{X}_n under **(FP)** converges to $\text{Eq}(\mathcal{G})$.*

Theorem 5.2. *Suppose **(Diff)** holds and \mathcal{G} satisfies **(Mon)** or **(S $_\varepsilon$)** for some $\varepsilon > 0$. Then the empirical frequency of play \bar{X}_n under **(RFP)** converges to $\text{Eq}_\varepsilon(\mathcal{G})$.*

Sketch of proof. The proof of **Theorems 5.1** and **5.2** follows a similar technical path starting with the fact that \bar{X}_n is an APT of **(BRD)** / **(RBRD)** respectively (cf. **Proposition 5.1** above). In both cases, Theorem 4.3 of Benaïm et al. [4] shows that the accumulation points of an APT lie in an internally chain transitive set of the underlying continuous-time dynamics.² The claims of **Theorem 5.1** may then be proved as follows:

- (1) For games satisfying **(Mon)**, invoke **Theorem 4.1** from **Section 4** and Theorem 3.23 and Proposition 3.25 of Benaïm et al. [4].
- (2) For games satisfying **(S $_\varepsilon$)** with $\varepsilon = 0$, invoke **Theorem 4.1** from **Section 4** and Proposition 3.27 of Benaïm et al. [4].

Finally, for $\varepsilon > 0$, the proof of **Theorem 5.2** is entirely analogous, and simply invokes **Theorem 4.2** in lieu of **Theorem 4.1** in the above chain of implications. ■

5.2. Dual averaging. We now proceed to state and prove our main convergence result for the recursion **(DA)**. To do so, we will require the “reciprocity” condition **(Rec)** that was introduced for the study of the continuous-time dynamics **(DAD)**, namely

$$Q_h(y_n) \rightarrow x \implies F_h(x, y_n) \rightarrow 0 \quad (\text{Rec})$$

for all $x \in \mathcal{X}$ and all sequences y_n in \mathbb{R}^A . We then have:

Theorem 5.3. *Suppose that **(DA)** is run with a regularizer satisfying **(StrCvx)** and **(Rec)**, and a vanishing non-increasing learning rate $\eta_n \searrow 0$ such that*

$$\lim_{n \rightarrow \infty} \left(\frac{1}{\eta_n} - \frac{1}{\eta_{n-1}} \right) = 0. \quad (60)$$

Then:

- If x^* is a GESS of \mathcal{G} , the sequence of population states X_n converges to x^* .
- If x^* is an ESS of \mathcal{G} , the same conclusion holds provided that $\max\{\eta_n, \eta_n^{-1} - \eta_{n-1}^{-1}\}$ is small enough and **(DA)** is initialized sufficiently close to x^* .

Corollary 5.1. *Suppose that \mathcal{G} is strictly monotone and **(DA)** is run with a regularizer satisfying **(StrCvx)** and **(Rec)**, and a learning rate of the form $\eta_n \propto 1/n^p$ for $p \in (0, 1)$. Then the sequence of population states induced by **(DA)** converges to the (necessarily unique) equilibrium of \mathcal{G} .*

²An internally chain transitive set is a compact invariant set that contains no proper attractors [10].

Our proof relies on the use of a suitable “energy inequality”, provided here by a deflated variant of the Fenchel coupling:

Lemma 8. *Fix a population state $x \in \mathcal{X}$, and let*

$$E_n := \frac{1}{\eta_n} F_h(x, \eta_n S_n) = \frac{1}{\eta_n} [h(x) - h(X_{n+1}) + \eta_n \langle S_n, X_{n+1} - x \rangle], \quad (61)$$

with S_n as in (DA). Then, for all $n = 1, 2, \dots$, we have:

$$E_n \leq E_{n-1} + \langle F(X_n), X_n - x \rangle + [h(x) - \min h] \delta_n + \frac{F_h(X_n, \eta_{n-1} S_n)}{\eta_{n-1}}, \quad (62)$$

where $\delta_n = 1/\eta_n - 1/\eta_{n-1}$. If, in addition, h satisfies (StrCvx), we have

$$E_n \leq E_{n-1} + \langle F(X_n), X_n - x \rangle + [h(x) - \min h] \delta_n + \frac{\eta_{n-1}}{2K} \|F(X_n)\|_*^2. \quad (63)$$

Energy inequalities of the form (62) are standard in the analysis of dual averaging schemes, and they go back at least as far as Nesterov [39, 40] (who introduced the method). The specific formulation above follows Héliou et al. [17, Lemma. C.1], so we omit the proof; for a range of similar computations, see Shalev-Shwartz [48, Lemma 2.20], Kwon [25, p. 42] and Sorin [52].

Now, to move forward with the proof of Theorem 5.3, we will require two complementary arguments. The first is a stability result: we show below that, if X_n is sufficiently close to an evolutionarily stable state x^* and the method’s learning rate is “small enough”, then X_{n+1} remains close to x^* . To state this precisely, let $\mathcal{F}_\varepsilon = \{y \in \mathbb{R}^A : F_h(x^*, y) \leq \varepsilon\}$ denote the “Fenchel zone” of magnitude ε with respect to x^* , and let $\mathcal{D}_\varepsilon = Q_h(\mathcal{F}_\varepsilon) = \{x = Q_h(y) : y \in \mathcal{F}_\varepsilon\}$ denote the image of \mathcal{F}_ε under Q_h . We then have:

Lemma 9. *Let $x^* \in \text{ESS}(\mathcal{G})$ and assume that (StrCvx) and (Rec) hold. Then, for all sufficiently small $\varepsilon > 0$, we have the following implication (valid for all $n = 1, 2, \dots$): if $X_n \in \mathcal{D}_\varepsilon$ and $\max\{\eta_n, \delta_n\}$ is small enough, we also have $X_{n+1} \in \mathcal{D}_\varepsilon$.*

The second component of the proof of Theorem 5.3 is a subsequence extraction argument: if the iterates of (DA) lie in a neighborhood of x^* where (ESS) holds, there exists a subsequence of X_n that converges to x^* .

Lemma 10. *Let $x^* \in \text{ESS}(\mathcal{G})$, let \mathcal{U} be the neighborhood of validity of (ESS), and assume further that (StrCvx) and (60) hold. If $X_n \in \mathcal{U}$ for all sufficiently large n , we have $\liminf_{n \rightarrow \infty} \|X_n - x^*\| = 0$.*

We prove these two ancillary results below.

Proof of Lemma 9. Let \mathcal{U} be the neighborhood whose existence is guaranteed by (ESS), i.e., $\langle F(x), x - x^* \rangle < 0$ for all $x \in \mathcal{U}$, $x \neq x^*$. By the relation between the Fenchel coupling and the ordinary norm topology on \mathcal{X} (cf. Lemma A.2 in Appendix A), it follows that $\mathcal{D}_\varepsilon \subseteq \mathcal{U}$ if ε is taken small enough; we will assume this to be the case throughout the rest of this proof. Moreover, by (Rec) and the continuity of u , there exists some $c \equiv c(\varepsilon)$ such that $\langle F(x), x - x^* \rangle \leq -c < 0$ for all $x \in \mathcal{D}_\varepsilon \setminus \mathcal{D}_{\varepsilon/2}$.

With all this in hand, assume that $X_n \in \mathcal{D}_\varepsilon$; we will show that $X_{n+1} \in \mathcal{D}_\varepsilon$ provided that $\max\{\eta_n, \delta_n\}$ is small enough. To do so, let $F_n = F_h(x^*, \eta_n S_n)$; then, by substituting $x \leftarrow x^*$ in the template inequality (63) of Lemma 8, we obtain:

Case 1. If $X_n \in \mathcal{D}_\varepsilon \setminus \mathcal{D}_{\varepsilon/2}$, then

$$F_n \leq F_{n-1} - \eta_n \left[c - \delta_n H - \eta_{n-1} \frac{G^2}{2K} \right], \quad (64)$$

where $H = \max h - \min h$ and $G = \max_{x \in \mathcal{X}} \|F(x)\|_*$. Hence, if $\delta_n \leq c/(2H)$ and $\eta_{n-1} \leq cK/G^2$, we conclude that $F_n \leq F_{n-1}$.

Case 2. If $X_n \in \mathcal{D}_{\varepsilon/2}$, then

$$F_n \leq \frac{\varepsilon}{2} + \delta_n H + \eta_{n-1} \frac{G^2}{2K}, \quad (65)$$

so $F_n \leq \varepsilon$ provided that $\delta_n < \varepsilon/(4H)$ and $\eta_{n-1} \leq \varepsilon K/(2G^2)$.

In both cases, we conclude that $X_{n+1} = Q_h(\eta_n S_n) \in \mathcal{D}_\varepsilon$, as claimed. \blacksquare

Proof of Lemma 10. Assume to the contrary that $\liminf_{n \rightarrow \infty} \|X_n - x^*\| > 0$. Then, by assumption, there exist $\varepsilon > 0$ and $n_0 \equiv n_0(\varepsilon)$ such that $X_n \in \mathcal{U}$ and $\|X_n - x^*\| \geq \varepsilon$ for all $n \geq n_0$. Since $x^* \in \text{ESS}(\mathcal{G})$, there exists a positive constant $c > 0$ such that $\langle F(X_n), X_n - x^* \rangle \leq -c < 0$ for all $n \geq n_0$. With this in mind, substituting $x \leftarrow x^*$ in the template inequality (63) and telescoping yields

$$\begin{aligned} E_n &\leq E_{n_0} - c(n - n_0) + \left(\frac{1}{\eta_n} - \frac{1}{\eta_{n_0}} \right) H + \frac{G^2}{2K} \sum_{k=n_0}^n \eta_k \\ &\leq E_{n_0} + cn_0 - H/n_0 - n \left[c - \frac{1}{n\eta_n} - \frac{1}{n} \sum_{k=n_0}^n \eta_k \right]. \end{aligned} \quad (66)$$

To proceed, note that the first part of (60) gives

$$\lim_{n \rightarrow \infty} \frac{(1/\eta_n - 1/\eta_{n-1})}{n - (n-1)} = \lim_{n \rightarrow \infty} \left(\frac{1}{\eta_n} - \frac{1}{\eta_{n-1}} \right) = \lim_{n \rightarrow \infty} \delta_n = 0, \quad (67)$$

so $\lim_{n \rightarrow \infty} 1/(n\eta_n) = 0$ by the Stolz–Cesàro theorem. By the second part of (60), we also have $(1/n) \sum_{k=n_0}^n \eta_k = 0$. Hence, by letting $n \rightarrow \infty$ in (66), we get $\lim_{n \rightarrow \infty} E_n = -\infty$, a contradiction which completes our proof. \blacksquare

The proof of our main convergence results follows by a tandem application of the above.

Proof of Theorem 5.3. We consider the two cases separately.

Case 1: $x^* \in \text{GESS}(\mathcal{G})$. Let $\varepsilon > 0$ be arbitrary. By (60), it follows that $\lim_{n \rightarrow \infty} \eta_n = \lim_{n \rightarrow \infty} \delta_n = 0$, and, by Lemma 10, there exists a subsequence X_{n_k} of X_n converging to x^* . This means that, in the long run, η_n and δ_n become arbitrarily small, and X_n comes arbitrarily close to x^* . Hence, by applying Lemma 9 inductively, we conclude that $X_n \in \mathcal{D}_\varepsilon$ for all sufficiently large n . Since $\varepsilon > 0$ has been chosen arbitrarily, this implies that X_n converges to x^* , as claimed.

Case 2: $x^* \in \text{ESS}(\mathcal{G})$. Let \mathcal{U} be the neighborhood whose existence is guaranteed by the evolutionary stability condition (3) for x^* . Then, by Lemma 9, if (DA) is initialized sufficiently close to x^* and $\max\{\eta_n, \delta_n\}$ is sufficiently small, we will have $X_n \in \mathcal{U}$ for all n . Hence, by Lemma 10, we conclude that $\liminf_{n \rightarrow \infty} \|X_n - x^*\| = 0$. Our proof then follows by the same inductive application of Lemma 9 as in the case of a GESS above. \blacksquare

Time averages under dual averaging. We close this section with a result on the empirical frequency of play under (DA). Indeed, telescoping (62) trivially yields the bound:

$$\sum_{k=1}^n \langle F(X_k), x - X_k \rangle \leq \frac{\max h - \min h}{\eta_n} + \frac{1}{2K} \sum_{k=1}^n \eta_{k-1} \|F(X_k)\|_*^2. \quad (68)$$

This regret-type bound echoes standard results in the literature – see e.g., Nesterov [40, Theorem 1] – and leads to the following time-average convergence result:

Proposition 5.2. *Suppose that (DA) is run with a regularizer satisfying (StrCvx) and a learning rate of the form $\eta_n \propto 1/\sqrt{n}$. If \mathcal{G} is monotone, the empirical frequency of play \bar{X}_n under (DA) converges to $\text{Eq}(\mathcal{G})$.*

The proof of Proposition 5.2 mimics that of Proposition 4.1 so we omit it. We only note that, in contrast to Theorem 5.3, Proposition 5.2 does not require strict monotonicity or reciprocity; however, it concerns the cruder, average frequency of play, so it provides a considerably weaker guarantee in this regard.

6. EXTENSIONS AND CONCLUDING REMARKS

In this concluding section, we provide a series of extensions and remarks that would have otherwise disrupted the flow of the rest of our paper.

6.1. Positive correlation. Sandholm [44] introduced the so-called “*positive correlation*” condition

$$\dot{x}_t \neq 0 \implies \langle F(x_t), \dot{x}_t \rangle > 0. \quad (\text{PC})$$

In the case of potential games, (PC) implies that

$$\frac{d}{dt} P(x_t) = \langle F(x_t), \dot{x}_t \rangle > 0 \quad (69)$$

whenever $\dot{x}_t \neq 0$. This in turn implies the convergence of the trajectories to the set of rest points of the dynamics [45, Theorem 7.1.2]. Some of our results in Section 4 for potential games can also be obtained by establishing (PC) for the dynamics under study.

6.2. No-regret dynamics. Our second remark concerns continuous- or discrete-time processes for which the average regret vanishes, i.e., for all $x \in \mathcal{X}$ we have:

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t \langle F(x_s), x - x_s \rangle ds = 0 \quad \text{in continuous time} \quad (70a)$$

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \langle F(X_k), x - X_k \rangle = 0 \quad \text{in discrete time} \quad (70b)$$

As we saw in Propositions 4.1 and 5.2, the dynamics of dual averaging have the no-regret property (70) in both continuous and discrete time. Looking more closely at the proofs of Propositions 4.1 and 5.2, we conclude that, in monotone games, the empirical frequency of play under any dynamical process that satisfies (70) converges to the game’s set of equilibria. This property forms the basis of many algorithms and dynamics designed to solve variational inequalities and equilibrium problems [38, 40, 24]; the above provides a game-theoretic interpretation of this property in the context of population games.

6.3. Extensions. We also note that our results can be extended to the following settings:

- (1) If \mathcal{X} is a convex compact subset of \mathbb{R}^A and $F: \mathcal{X} \rightarrow \mathbb{R}^A$ is a vector field on \mathcal{X} , we maintain the same results for convergence to solutions of (SVI) and (MVI); however, the population game interpretation disappears.
- (2) This interpretation is recovered in multi-population games: if there are several player populations indexed by $i = 1, \dots, N$, and if \mathcal{A}_i denotes the set of pure strategies of the i -th population and $u_{i\alpha}: \prod_{i=1}^N \Delta(\mathcal{A}_i) \rightarrow \mathbb{R}$ is the payoff function of α -strategists in the i -th population, it suffices to set $\mathcal{X} = \prod_i \Delta(\mathcal{A}_i)$ and $F(x) = (u_{i\alpha_i}(x))_{\alpha_i \in \mathcal{A}_i, i \in \mathcal{I}}$.
- (3) Our theory can also be extended to normal form games with a finite number of players and smooth concave payoff functions g_i , $i = 1, \dots, N$. In this case, the components of the corresponding vector field are $G_{i\alpha}(x) = \frac{\partial g_i}{\partial x_{i\alpha}}$, and the notion of evolutionary stability boils down to that of *variational stability*; for a detailed treatment, see Sorin & Wan [53] and Mertikopoulos & Zhou [35].

6.4. Links with variational inequalities and further extensions. Taking a more general point of view that focuses on the vector field F , we can also consider the following framework that brings to the forefront the associated variational inequalities. Let \mathcal{X}_i , $i = 1, \dots, N$ be a convex compact subset of a topological vector space \mathcal{V}_i with dual \mathcal{V}_i^* , write $\mathcal{V} = \prod_i \mathcal{V}_i$ and $\mathcal{X} = \prod_i \mathcal{X}_i$, and let $F_i: \mathcal{X} \rightarrow \mathcal{V}_i^*$ be a collection of continuous maps, $i = 1, \dots, N$. Finally, set $F(x) = (F_1(x), \dots, F_N(x))$ and, as usual, write

$$\langle F(x), z \rangle = \sum_{i=1}^N \langle F_i(x), z_i \rangle \quad (71)$$

for the dual pairing between $F(x) \in \mathcal{V}^* = \prod_i \mathcal{V}_i^*$ and $z = (z_1, \dots, z_N) \in \mathcal{V} := \prod_i \mathcal{V}_i$.

We may then define $\text{SVI}(\mathcal{X}, F)$ as the set of $x^* \in \mathcal{X}$ such that

$$\langle F(x^*), x - x^* \rangle \leq 0, \quad \text{for all } x \in \mathcal{X} \quad (72)$$

and $\text{MVI}(\mathcal{X}, F)$ as the set of $x^* \in \mathcal{X}$ such that

$$\langle F(x), x - x^* \rangle \leq 0 \quad \text{for all } x \in \mathcal{X}. \quad (73)$$

The counterparts for ESS and GESS may also be defined similarly (though the link with population games is no longer present).

Example 6.1. Consider a strategic game with N players, compact metric strategy spaces $\mathcal{A}_i \subseteq \mathbb{R}^d$ and continuous payoff functions $u_i: \mathcal{A} \rightarrow \mathbb{R}$, $i = 1, \dots, N$, with $\mathcal{A} := \prod_j \mathcal{A}_j$. Introduce the mixed extension of the game as usual: for a probability distribution (mixed strategy profile) $\chi = (\chi_1, \dots, \chi_N) \in \mathcal{X}$ with $\chi_i \in \mathcal{X}_i := \Delta(\mathcal{A}_i)$, let u_i be extended as

$$u_i(\chi) = \int_{\mathcal{A}} u_i(\alpha_1, \dots, \alpha_N) d\chi_1(\alpha_1) \cdots d\chi_N(\alpha_N). \quad (74)$$

Then, introduce $F_i(\chi) = u_i(\cdot, \chi_{-i})$ so that

$$\langle F_i(\chi), \chi'_i \rangle = F_i(\chi'_i; \chi_{-i}) \quad (75)$$

for all $\chi'_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i)$. In this way, the equilibrium condition for $\chi^* \in \mathcal{X}$ becomes

$$\langle F(\chi^*), \chi - \chi^* \rangle = \sum_{i=1}^N \langle F_i(\chi^*), \chi_i - \chi_i^* \rangle \leq 0 \quad \text{for all } \chi \in \mathcal{X}, \quad (76)$$

which is again an instance of $\text{SVI}(\mathcal{X}, F)$. \diamond

In this setting, F is monotone (dissipative) if

$$\langle F(x) - F(x'), x - x' \rangle \leq 0 \quad \text{for all } x, x' \in \mathcal{X} \quad (77)$$

As in the case of [Section 2](#), we have $\text{SVI}(\mathcal{X}, F) = \text{MVI}(\mathcal{X}, F)$ if F is monotone; furthermore, mutatis mutandis, all the properties presented in [Section 2.3](#) extend to this general setup.

The best-response map associated to F is defined again as

$$\text{BR}(x) = \arg \max_{x' \in \mathcal{X}} \langle F(x), x' \rangle. \quad (78)$$

In addition to monotonicity, we will also now make the “*unique best response*” assumption

$$\text{BR}(x) \text{ is a singleton for all } x \in \mathcal{X}. \quad (\text{U})$$

Proposition 6.1. *If F is monotone and (U) holds, the set $E := \text{MVI}(\mathcal{X}, F) = \text{SVI}(\mathcal{X}, F)$ is a singleton.*

Proof. If $x, x' \in E$, we have $\langle F(x), x' - x \rangle \leq 0$ and $\langle F(x'), x - x' \rangle \leq 0$, so

$$\langle F(x') - F(x), x' - x \rangle \geq 0 \quad (79)$$

implying first $\langle F(x') - F(x), x' - x \rangle = 0$ by monotonicity and then $\langle F(x'), x' - x \rangle = 0$. Thus, $x = x'$ by (U). ■

The regularized best response map is likewise defined as

$$\text{BR}_{\varepsilon h}(x) = \arg \max_{x' \in \mathcal{X}} \{\langle F(x), x' \rangle - \varepsilon h(x')\} \quad (80)$$

where $h: \mathcal{X} \rightarrow \mathbb{R}$ satisfies the same axioms as [Definition 2.1](#). Note that $\text{BR}_{\varepsilon h}$ satisfies (U) by construction.

Now, using these best-response maps, we may extend the definition of the dynamics (FP), (RFP) and (DA) to the more general setup considered here in the obvious way. We then have:

Proposition 6.2. *Suppose that (U) holds and let $X_n, n = 1, 2, \dots$, be a sequence of states generated by (FP). If the accumulation points of \bar{X}_n lie in $\text{SVI}(\mathcal{X}, F)$, the same holds for the accumulation points of X_n .*

Proof. Let x^* be an accumulation point of $X_{n+1} = \text{BR}(\bar{X}_n)$, and let X_{n_k+1} be a subsequence converging to x^* . By descending to a further subsequence if necessary, we may assume without loss of generality that \bar{X}_{n_k} converges to some limit point \bar{x} . Furthermore, by the definition of (FP), we have:

$$\langle F(\bar{X}_{n_k}), X_{n_k+1} - x \rangle = \langle F(\bar{X}_{n_k}), \text{BR}(\bar{X}_{n_k}) - x \rangle \geq 0 \quad \text{for all } x \in \mathcal{X}. \quad (81)$$

Hence, taking the limit $k \rightarrow \infty$, we get

$$\langle F(\bar{x}), x^* - x \rangle \geq 0 \quad \text{for all } x \in \mathcal{X}, \quad (82)$$

so $x^* \in \text{BR}(\bar{x})$. Since $\bar{x} \in \text{SVI}(\mathcal{X}, F)$ by assumption, we deduce further that $\bar{x} \in \text{BR}(\bar{x})$. Hence, by (U), we conclude that $\bar{x} = x^*$, and our assertion follows. ■

Obviously, under (U), a similar property holds for regularized fictitious play.

6.5. Preview for Part II. Part II of our paper [16] deals with games with continuous action spaces, in the spirit of [Example 6.1](#) (but with a continuous population of nonatomic players). The first step is to define the Stampacchia variational inequality when \mathcal{A} is a compact metric space, Θ is an appropriate compact convex subset of $\Delta(\mathcal{A})$, F is a continuous map $\theta \mapsto F_\theta(\cdot)$ from Θ to the space of continuous functions on \mathcal{A} , and the duality mapping is given by

$$\langle F_\theta, \nu \rangle = \int_{\mathcal{A}} F_\theta(\alpha) d\nu(\alpha) \quad \text{for all } \theta, \nu \in \Theta. \quad (83)$$

In this framework, we introduce potential and monotone maps, and we prove the convergence of the discrete-time dynamics of fictitious play.

We show that this study covers the case of nonatomic games with (player-dependent) compact action spaces. Finally, we introduce a class of first-order mean-field games and prove that solutions of the Stampacchia variational inequality correspond to the equilibria of the game in normal form where \mathcal{A} is the set of paths generated by the players.

A. PROPERTIES OF CHOICE MAPS AND THE FENCHEL COUPLING

Our aim in this appendix is to provide some useful background and properties for the regularized choice map h and the Fenchel coupling.

To recall the basic setup, we assume below that h is a regularizer on \mathcal{X} in the sense of [Definition 2.1](#). The convex conjugate $h^*: \mathbb{R}^A \rightarrow \mathbb{R}$ of h is then defined as

$$h^*(y) = \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\}. \quad (A.1)$$

Since \mathcal{X} is compact and h is strictly convex and continuous on \mathcal{X} , the maximum in [\(A.1\)](#) is attained at a unique point in \mathcal{X} , so $h^*(y)$ is finite for all $y \in \mathbb{R}^A$ [1]. Moreover, by standard results in convex analysis [42, Chap. 26], h^* is differentiable on \mathbb{R}^A and its gradient satisfies the identity

$$\nabla h^*(y) = \arg \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\} = Q(y) \quad (A.2)$$

where $Q: \mathbb{R}^A \rightarrow \mathcal{X}$ is the choice map induced by h (cf. [Definition 2.1](#)).

The lemma below establishes the inverse differentiability property for h , namely that $\partial h = Q^{-1}$:

Lemma A.1. *Let h be a regularizer on \mathcal{X} . Then, for all $x \in \mathcal{X}$ and all $y \in \mathbb{R}^A$, we have:*

$$x = Q(y) \iff y \in \partial h(x). \quad (A.3)$$

If, in addition, h satisfies [\(Diff\)](#), we also have

$$\langle \nabla h(x), x - p \rangle \leq \langle y, x - p \rangle \quad (A.4)$$

for all $p \in \mathcal{X}$ and all $y \in \partial h(x)$, $x \in \text{dom } \partial h$.

Proof of [Lemma A.1](#). To prove [\(A.3\)](#), note that x solves [\(A.2\)](#) if and only if $y - \partial h(x) \ni 0$, i.e., if and only if $y \in \partial h(x)$.

For the inequality [\(A.4\)](#), it suffices to show it holds for all $p \in \text{dom } \partial h$ (by continuity). To do so, let

$$\phi(t) = h(x + t(p - x)) - [h(x) + \langle y, x + t(p - x) \rangle]. \quad (A.5)$$

Since h is strictly convex and $y \in \partial h(x)$ by [\(A.3\)](#), it follows that $\phi(t) \geq 0$ with equality if and only if $t = 0$. Moreover, note that $\psi(t) = \langle \nabla h(x + t(p - x)) - y, p - x \rangle$ is a continuous subgradient selection of ϕ . Given that ϕ and ψ are both continuous on $[0, 1]$, it follows a fortiori that ϕ is continuously differentiable and $\phi' = \psi$ on $[0, 1]$. Thus, with ϕ convex and $\phi(t) \geq 0 = \phi(0)$ for all $t \in [0, 1]$, we conclude that $\phi'(0) = \langle \nabla h(x) - y, p - x \rangle \geq 0$, from which our claim follows. \blacksquare

Lemma A.2. *Let h be a regularizer on \mathcal{X} . Then, for all $p \in \mathcal{X}$ and all $y \in \mathbb{R}^A$, we have*

$$F_h(p, y) \geq 0 \quad \text{with equality if and only if } p = Q(y). \quad (\text{A.6})$$

If, in addition, if h satisfies (StrCvx), we also have:

$$F_h(p, y) \geq \frac{K}{2} \|Q(y) - p\|^2. \quad (\text{A.7})$$

Proof. Our first claim is a trivial consequence of the Fenchel-Young inequality and the strict convexity of h . For our second claim, let $x = Q(y)$. Then, by the definition (A.1) of the convex conjugate of h , we have:

$$\begin{aligned} F_h(p, y) &= h(p) + h^*(y) - \langle y, p \rangle \\ &= h(p) + \langle y, Q(y) \rangle - h(Q(y)) - \langle y, p \rangle \\ &= h(p) - h(x) - \langle y, p - x \rangle \end{aligned} \quad (\text{A.8})$$

Now, since h is K -strongly convex, we also have:

$$h((1 - \lambda)x + \lambda p) + \lambda(1 - \lambda)(K/2)\|x - p\|^2 \leq (1 - \lambda)h(x) + \lambda h(p) \quad (\text{A.9})$$

and, with $x = Q(y)$, we get:

$$\langle y, x \rangle - h(x) \geq \langle y, (1 - \lambda)x + \lambda p \rangle - h((1 - \lambda)x + \lambda p). \quad (\text{A.10})$$

Thus, after rearranging, dividing by λ , and letting $\lambda \rightarrow 0$, we obtain

$$h(p) - h(x) - \langle y, p - x \rangle \geq (K/2)\|x - p\|^2 \quad (\text{A.11})$$

and our assertion follows from (A.8). \blacksquare

By virtue of this lemma, we obtain the following convergence criterion in terms of the Fenchel coupling:

Corollary A.1. *Let h be a regularizer on \mathcal{X} satisfying (StrCvx), fix a base point $p \in \mathcal{X}$, and let y_n be a sequence in \mathbb{R}^A . Then, if $F_h(p, y_n) \rightarrow 0$, we also have $Q(y_n) \rightarrow p$.*

Proof. By Lemma A.2, we have $\|Q(y_n) - p\| \leq \sqrt{(2/K)F_h(p, y_n)}$, so our claim follows. \blacksquare

ACKNOWLEDGMENTS

The authors acknowledge support by the COST Action CA16228 ‘‘European Network for Game Theory’’ (GAMENET), the FMJH Program PGMO under grant Damper, and the French National Research Agency (ANR) in the framework of the ‘‘Investissements d’avenir’’ program (ANR-15-IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), MIAI@Grenoble Alpes (ANR-19-P3IA-0003), and the grants ORACLESS (ANR-16-CE33-0004) and ALIAS (ANR-19-CE48-0018-01).

REFERENCES

- [1] Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, New York, NY, USA, 2 edition, 2017.
- [2] Michel Benaïm. Dynamics of stochastic approximation algorithms. In Jacques Azéma, Michel Émery, Michel Ledoux, and Marc Yor (eds.), *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pp. 1–68. Springer Berlin Heidelberg, 1999.
- [3] Michel Benaïm and Mathieu Faure. Consistency of vanishingly smooth fictitious play. *Mathematics of Operations Research*, 38(3):437–450, 2013.
- [4] Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 2005.
- [5] Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Stochastic approximations and differential inclusions, part II: Applications. *Mathematics of Operations Research*, 31(4):673–695, 2006.
- [6] Claude Berge. *Topological Spaces*. Dover, New York, 1997.

- [7] Mario Bravo and Panayotis Mertikopoulos. On the robustness of learning in games with stochastically perturbed payoff observations. *Games and Economic Behavior*, 103, John Nash Memorial issue:41–66, May 2017.
- [8] George W. Brown. Iterative solutions of games by fictitious play. In T. C. Coopmans (ed.), *Activity Analysis of Productions and Allocation*, 374–376. Wiley, 1951.
- [9] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [10] Charles Cameron Conley. *Isolated Invariant Set and the Morse Index*. American Mathematical Society, Providence, RI, 1978.
- [11] Stella Dafermos. Traffic equilibrium and variational inequalities. *Transportation Science*, 14(1):42–54, February 1980.
- [12] Daniel Friedman. Evolutionary games in economics. *Econometrica*, 59(3):637–666, 1991.
- [13] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*, volume 2 of *Economic learning and social evolution*. MIT Press, Cambridge, MA, 1998.
- [14] Drew Fudenberg and David K. Levine. Conditional universal consistency. *Games and Economic Behavior*, 29(1):104–130, 1999.
- [15] Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica*, 59(3):859–867, May 1991.
- [16] Saeed Hadikhaneloo, Rida Laraki, Panayotis Mertikopoulos, and Sylvain Sorin. Learning in nonatomic games, Part II: Continuous action spaces and mean-field games. Mimeo, 2021.
- [17] Amélie Héliou, Matthieu Martin, Panayotis Mertikopoulos, and Thibaud Rahier. Online non-convex optimization with imperfect feedback. In *NeurIPS '20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [18] Josef Hofbauer and William H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, November 2002.
- [19] Josef Hofbauer and William H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693, July 2009.
- [20] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.
- [21] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.
- [22] Josef Hofbauer, Peter Schuster, and Karl Sigmund. A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81(3):609–612, 1979.
- [23] Josef Hofbauer, Sylvain Sorin, and Yannick Viossat. Time average replicator and best reply dynamics. *Mathematics of Operations Research*, 34(2):263–269, May 2009.
- [24] Anatoli Juditsky, Arkadi Semen Nemirovski, and Claire Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- [25] Joon Kwon. *Stratégies de descente miroir pour la minimisation du regret et l’approchabilité*. PhD thesis, Université Pierre-et-Marie-Curie, 2016.
- [26] Joon Kwon and Panayotis Mertikopoulos. A continuous-time approach to online optimization. *Journal of Dynamics and Games*, 4(2):125–148, April 2017.
- [27] Andreu Mas-Colell. On a theorem of Schmeidler. *Journal of Mathematical Economics*, 13:201–206, 1984.
- [28] John Maynard Smith and George R. Price. The logic of animal conflict. *Nature*, 246:15–18, November 1973.
- [29] Daniel L. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka (ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press, New York, NY, 1974.
- [30] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(6):6–38, 1995.
- [31] Panayotis Mertikopoulos and Aris L. Moustakas. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4):1359–1388, July 2010.
- [32] Panayotis Mertikopoulos and William H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, November 2016.
- [33] Panayotis Mertikopoulos and William H. Sandholm. Riemannian game dynamics. *Journal of Economic Theory*, 177:315–364, September 2018.

- [34] Panayotis Mertikopoulos and Mathias Staudigl. On the convergence of gradient-like flows with noisy gradient input. *SIAM Journal on Optimization*, 28(1):163–197, January 2018.
- [35] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.
- [36] George J. Minty. On the generalization of a direct method of the calculus of variations. *Bulletin of the American Mathematical Society*, 73(3):315–321, May 1967.
- [37] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [38] Arkadi Semen Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [39] Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2):319–344, 2007.
- [40] Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.
- [41] Julia Robinson. An iterative method for solving a game. *Annals of Mathematics*, 54:296–301, 1951.
- [42] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.
- [43] Aldo Rustichini. Optimal properties of stimulus-response learning models. *Games and Economic Behavior*, 29(1-2):244–273, 1999.
- [44] William H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97: 81–108, 2001.
- [45] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.
- [46] William H. Sandholm. Population games and deterministic evolutionary dynamics. In H. Peyton Young and Shmuel Zamir (eds.), *Handbook of Game Theory IV*, pp. 703–778. Elsevier, 2015.
- [47] David Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, April 1973.
- [48] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [49] Shai Shalev-Shwartz and Yoram Singer. Convex repeated games and Fenchel duality. In *NIPS’ 06: Proceedings of the 19th Annual Conference on Neural Information Processing Systems*, pp. 1265–1272. MIT Press, 2006.
- [50] Sylvain Sorin. Exponential weight algorithm in continuous time. *Mathematical Programming*, 116(1): 513–528, 2009.
- [51] Sylvain Sorin. Replicator dynamics: Old and new. *Journal of Dynamics and Games*, 7(4):365–386, 2020.
- [52] Sylvain Sorin. No-regret algorithms in online learning, games and convex optimization. Mimeo, 2021.
- [53] Sylvain Sorin and Cheng Wan. Finite composite games: Equilibria and dynamics. *Journal of Dynamics and Games*, 3(1):101–120, January 2016.
- [54] Peter D. Taylor. Evolutionarily stable strategies with two types of player. *Journal of Applied Probability*, 16(1):76–83, March 1979.
- [55] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.
- [56] Eric van Damme. *Stability and perfection of Nash equilibria*. Springer-Verlag, Berlin, 1987.
- [57] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.