



Notes

No-regret dynamics and fictitious play [☆]

Yannick Viossat ^a, Andriy Zapechelnuk ^{b,*}

^a CEREMADE, Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, F-75775 Paris, France

^b School of Economics and Finance, Queen Mary, University of London, Mile End Road, London E1 4NS, UK

Received 1 October 2009; final version received 27 July 2012; accepted 31 July 2012

Available online 9 January 2013

Abstract

Potential based no-regret dynamics are shown to be related to fictitious play. Roughly, these are ε -best reply dynamics where ε is the maximal regret, which vanishes with time. This allows for alternative and sometimes much shorter proofs of known results on convergence of no-regret dynamics to the set of Nash equilibria.

© 2013 Elsevier Inc. All rights reserved.

JEL classification: C73; D81; D83

Keywords: Regret minimization; No-regret strategy; Fictitious play; Best-reply dynamics; Nash equilibrium; Hannan set; Curb set

1. Introduction

No-regret strategies are simple adaptive learning rules that recently received a lot of attention in the literature.¹ In a repeated game, a player has a *regret* for an action if, loosely speaking, she

[☆] The authors thank William Sandholm whose comments led to a substantial improvement of the paper, as well as Mathieu Faure, Sergiu Hart, Josef Hofbauer, Alexander Matros, Karl Schlag, Eilon Solan and Sylvain Sorin for helpful comments and suggestions. Y. Viossat gratefully acknowledges the support of ANR BLAN RISK.

* Corresponding author.

E-mail addresses: viossat@ceremade.dauphine.fr (Y. Viossat), a.zapechelnuk@qmul.ac.uk (A. Zapechelnuk).

¹ These rules have been used to investigate convergence to equilibria in the context of learning in games [19,20,24–26], for combining different forecasts [17,18] (for an overview of the forecast combination literature see [14,43]) and for combining opinions, which is also of interest to management science [33]. In finance this method has been used to derive bounds on the prices of financial instruments [13,15]. This method can be applied to various tasks in computer science, such as job scheduling [36] and routing [9] (for a survey of applicable problems in computer science see [31]).

could have obtained a greater average payoff had she played that action more often in the past. In the course of the game, the player reinforces actions that she regrets not having played enough, for instance, by choosing next action with probability proportional to the regret for that action, as in Hart and Mas-Colell's [24] *regret-matching* rule. Existence of *no-regret strategies* (i.e., strategies that guarantee no regrets almost surely in the long run) is known since Hannan [23]; wide classes of no-regret strategies are identified by Hart and Mas-Colell [25] and Cesa-Bianchi and Lugosi [11].²

A *no-regret dynamics* is a stochastic process that describes trajectories of the average correlated play of players and that emerges when every player follows a no-regret strategy (different players may play different strategies). By definition, it converges to the Hannan set (the set of all correlated actions that satisfy the no-regret condition first stated by Hannan [23]).³ This set is typically large. It contains the set of correlated equilibria of the game and we show that it may even contain correlated actions that put positive weight *only* on strictly dominated actions. Thus convergence of the average play to the Hannan set often provides very little information about what the players will actually play, as it does not even imply exclusion of strictly dominated actions.

In this paper we show that no-regret dynamics are intimately linked to the classical fictitious play process [10]. Drawing on Monderer et al. [38], we first show that contrary to the standard, discrete-time version, continuous fictitious play leads to no regret. We then show that, for a large class of no-regret dynamics, if a player's maximal regret is $\varepsilon > 0$, then she plays an ε -best reply to the average correlated play of the others. Since in this class the maximal regret vanishes (see Corollary 1 below), it follows that, for a good choice of behavior when all regrets are negative, the dynamics is a vanishingly perturbed version of fictitious play.

For two-player finite games, this observation and the theory of perturbed differential inclusions [4,5] allow us to relate formally the asymptotic behavior of no-regret dynamics and of continuous fictitious play (or its time-rescaled version, the best-reply dynamics [22]). In classes of games in which the behavior of continuous fictitious play is well known, this provides substantial information on the asymptotic behavior of no-regret dynamics. In particular, we recover most known convergence properties of no-regret dynamics. Our results do not just allow us to find new and sometimes much shorter proofs of convergence of no-regret dynamics towards the set of Nash equilibria in some classes of games, such as dominance solvable games or potential games. They also allow us to relate the asymptotic behavior of no-regret dynamics and continuous fictitious play in case of divergence, as in the famous Shapley game [41].

These results extend only partially to n -player games (though they fully extend to n -player games with linear incentives [40]). The issue is that in n -player games no-regret dynamics turn out to be related to the correlated version of continuous fictitious play, in which the players play a best reply to the *correlated* past play of the others. This version of fictitious play is defined through a correspondence which is not convex valued. This creates technical difficulties, because the theory of perturbed differential inclusions is not developed for nonconvex-valued correspondences.

A different way to analyze no-regret dynamics is to show that some sets attract nearby solution trajectories. We show that strict Nash equilibria and, more generally, the intersection of

² This paper deals with the simplest notion of regret known as *unconditional* (or *external*) regret [20,25,26]. For more sophisticated regret notions, see Hart and Mas-Colell [24], Lehrer [34], and Cesa-Bianchi and Lugosi [12].

³ The Hannan set of a game is also known as the set of *weak correlated equilibria* [39] or *coarse correlated equilibria* [48, Chapter 3].

the Hannan set and the sets that are *closed under rational behavior (curb)*⁴ are attracting for no-regret dynamics, in a sense to be defined in Section 4.

The remainder of the note is organized as follows. The next section introduces no-regret dynamics. Section 3 studies the links between no-regret dynamics and fictitious play. Section 4 shows that the intersection of the Hannan set and curb sets is attracting for no-regret dynamics. Section 5 studies the continuous-time version and the expected version of no-regret dynamics. Finally, Appendix A contains the proof of the main result.

2. Preliminaries

Consider a bimatrix game $\Gamma = (A_i, u_i)_{i=1,2}$, where A_i is the set of actions of player i and $u_i : A \rightarrow \mathbb{R}$ is her payoff function, with $A = A_1 \times A_2$. For any finite set B , denote by $\Delta(B)$ the set of probability distributions over B . A mixed action of player i is an element of $\Delta(A_i)$. A correlated action z is a probability distribution over the set of pure action profiles, i.e., $z \in \Delta(A)$. Given such a z , let $z_i \in \Delta(A_i)$ and $z_{-i} \in \Delta(A_{-i})$ denote its marginals for player i and her opponent, respectively. Thus, $z_i(a_i) = \sum_{a_{-i} \in A_{-i}} z(a_i, a_{-i})$. Throughout, $-i$ refers to i 's opponent. As usual, let $u_i(z) = \sum_{a \in A} z(a)u_i(a)$ and $u_i(k, z_{-i}) = \sum_{a_{-i} \in A_{-i}} z_{-i}(a_{-i})u_i(k, a_{-i})$ for $k \in A_i$. Depending on the context, a_i may refer to a pure action – an element of A_i – or to a vertex of $\Delta(A_i)$, i.e., a Dirac measure on a pure action.

The game is played repeatedly in discrete-time periods $t \in \mathbb{N}^* = \{1, 2, \dots\}$. In every period t each player i chooses an action $a_i(t) \in A_i$ and receives payoff $u_i(a(t))$ where $a(t) = (a_1(t), a_2(t))$. Denote by $h(t) = (a(1), a(2), \dots, a(t))$ the history of play up to t , and let \mathcal{H} be the set of all finite histories (including the empty history). A strategy of player i is a function $q_i : \mathcal{H} \rightarrow \Delta(A_i)$ that stipulates to play in every period $t = 1, 2, \dots$ a mixed action $q_i(t) \equiv q_i(h(t - 1))$ as a function of the history before t . The weight that this mixed action puts on action $k \in A_i$ is denoted by $q_{i,k}(t)$.

The *average correlated play* up to period t is $z(t) = \frac{1}{t} \sum_{\tau=1}^t a(\tau)$, where we identify $a(\tau)$ with the corresponding vertex of $\Delta(A)$. Since $z(t) = \frac{1}{t}[a(t) + (t - 1)z(t - 1)]$, it follows that for all $t > 1$,

$$z(t) - z(t - 1) = \frac{1}{t}(a(t) - z(t - 1)). \tag{1}$$

For a correlated action z , the *regret* of player i for action k is defined as $R_{i,k}(z) = u_i(k, z_{-i}) - u_i(z)$, and her maximal regret as $R_{i,\max}(z) = \max_{k \in A_i} R_{i,k}(z)$. Typically we deal with the regret based on the average correlated play, $z(t)$, up to some period t . In this case the regret of player i for action $k \in A_i$ is equal to the difference between the average payoff she would have obtained by always playing k (assuming that her opponent's play remains the same) and her average realized payoff:

$$R_{i,k}(z(t)) = u_i(k, z_{-i}(t)) - u_i(z(t)) = \frac{1}{t} \sum_{\tau=1}^t [u_i(k, a_{-i}(\tau)) - u_i(a(\tau))].$$

To simplify notations, we will often write $R_{i,k}(t)$ for $R_{i,k}(z(t))$ and $R_{i,\max}(t)$ for $R_{i,\max}(z(t))$.

⁴ A product set of action profiles is called *closed under rational behavior (curb)* [2] if it contains all best replies of each player whenever she believes that no actions outside this set are being played by the other players.

Player i has no asymptotic regret if her average realized payoff is asymptotically no less than her best-reply payoff against the empirical distribution of her opponent:

$$\limsup_{t \rightarrow \infty} R_{i,\max}(t) \leq 0. \tag{2}$$

A strategy of player i is a *no-regret strategy* if for any strategy of the other player, inequality (2) holds almost surely. This property is also called *Hannan consistency* [25] or *universal consistency* [20].

It is well known in the literature since Hannan [23] that there exist simple no-regret strategies. Hart and Mas-Colell [25] describe a wide class of *potential based* no-regret strategies. A twice differentiable, convex function $P_i : \mathbb{R}^{A_i} \rightarrow \mathbb{R}$ is called a *potential* if it satisfies the following conditions:

- (R1) $P_i(\cdot) \geq 0$, and $P_i(x) = 0$ for all $x \in \mathbb{R}_-^{A_i}$;
- (R2) $\nabla P_i(\cdot) \geq 0$, and $\nabla P_i(x) \cdot x > 0$ for all $x \notin \mathbb{R}_-^{A_i}$;
- (R3) if $x \notin \mathbb{R}_-^{A_i}$ and $x_k \leq 0$, then $\nabla_k P_i(x) = 0$,

where ∇_k denotes the partial derivative with respect to $x_i(k)$. The potential P_i can be viewed as a generalized distance function between a vector $x \in \mathbb{R}^{A_i}$ and the nonpositive orthant $\mathbb{R}_-^{A_i}$. Let $R_i(t) = (R_{i,k}(t))_{k \in A_i}$ denote player i 's regret vector.

Proposition 1. *Let P_i satisfy (R1)–(R3) and let strategy q_i satisfy*

$$q_{i,k}(t + 1) = \frac{\nabla_k P_i(R_i(t))}{\sum_{s \in A_i} \nabla_s P_i(R_i(t))}, \quad \forall k \in A_i, \tag{Q1}$$

whenever $R_{i,\max}(t) > 0$. Then q_i is a no-regret strategy.

Proof. This holds by Theorem 3.3 of Hart and Mas-Colell [25], whose conditions (R1) and (R2) are satisfied by our conditions (R1)–(Q1) and (R2), respectively, and whose proof is based on the Blackwell’s Approachability Theorem [8]. \square

A standard example of no-regret strategy satisfying the above conditions is obtained by letting P_i be the l_p -norm on $\mathbb{R}_+^{A_i}$, i.e. $P_i(x) = (\sum_{k \in A_i} [x_k]_+^p)^{1/p}$ with $1 < p < \infty$, where $[x_k]_+ = \max(0, x_k)$. The resulting strategy q_i is called the *l_p -norm strategy* [11,25]. It is defined by

$$q_{i,k}(t + 1) = \frac{[R_{i,k}(t)]_+^{p-1}}{\sum_{s \in A_i} [R_{i,s}(t)]_+^{p-1}}, \quad \forall k \in A_i,$$

whenever $R_{i,\max}(t) > 0$. The l_2 -norm strategy is the *regret-matching strategy* [24], that stipulates to play an action in the next period with probability proportional to the regret for that action. For large p , the l_p -norm strategies approximate fictitious play.

We say that the average correlated play $z(t)$ follows a *no-regret dynamics* if both players use (possibly different) no-regret strategies. A trajectory $(z(t))_{1 \leq t \leq +\infty}$ of a no-regret dynamics is thus a solution of (1) where $a(t)$ is a realization of $(q_1(t), q_2(t))$ and q_1, q_2 are no-regret strategies. We focus on the class \mathcal{R} of no-regret dynamics such that:

- (i) the no-regret strategies q_1, q_2 of the players are potential based: they satisfy (Q1) for some potentials P_1, P_2 satisfying (R1)–(R3);
- (ii) if a player has no regret then he takes some constant pure action: for each $i = 1, 2$, there exists $c \in A_i$ such that

$$a_i(t + 1) = c \quad \text{whenever } R_{i,\max}(t) \leq 0. \tag{Q2}$$

Our results are valid for a somewhat wider class of no-regret dynamics. What we really need, beside a no-regret dynamics, is that from some period t_0 on:

- (i') if a player has positive regret for some actions, then she plays one of these actions;
- (ii') if a player never has any positive regret, then she plays an $\varepsilon(t)$ -best reply to the empirical distribution of her opponent, where $\varepsilon(t) = \varepsilon(h(t)) \rightarrow 0$ almost surely.

Remark 1. Property (i') follows from (R3) and (Q1). This is a *better reply property* that stipulates to assign a positive probability only on better reply actions to the opponent's empirical distribution of play ("better" with respect to the realized payoff). Also it implies that if $R_{i,\max}(t) > 0$ in some period t , then $R_{i,\max}(t') > 0$ for all $t' > t$. Indeed, when an action k with positive regret is played, the sign of $R_{i,k}(t)$ does not change, hence the maximal regret remains positive [25, Proposition 4.3].

Remark 2. Assumption (Q2) is a simple way of ensuring (ii'), and in addition, that if $R_{i,\max}(t) \leq 0$ for all t , then $R_{i,\max}(t) \rightarrow 0$ as $t \rightarrow +\infty$.⁵ Indeed, if $R_{i,\max}(t) \leq 0$ for all $t > t_0$ then by (Q2), for all $t > t_0$, $tR_{i,c}(t) = t_0R_{i,c}(t_0)$, hence $R_{i,c}(t) \rightarrow 0$. It follows that $R_{i,\max}(t) \rightarrow 0$ and that for all $t > t_0$, player i plays an $\varepsilon(t)$ -best reply with $\varepsilon(t) := \max_{k \in A_i} u_i(k, z_{-i}(t)) - u_i(c, z_{-i}(t)) = R_{i,\max}(t) - R_{i,c}(t) \rightarrow 0$. For a discussion of other possible assumptions, see Hart and Mas-Colell [26, Appendix A].

Note that there are no-regret dynamics that do not satisfy (i'). For instance, stochastic fictitious play with a noise parameter that declines with time at an appropriate rate (see, e.g., Benaïm and Faure [3]). This process is not potential based in our sense due to the time inhomogeneity, but this is not the crucial point, since (i')–(ii') would suffice.

Define the *Hannan set* H of the stage game Γ as the set of all correlated actions of the players where each player has no regret:

$$H = \left\{ z \in \Delta(A) \mid \max_{k \in A_i} u_i(k, z_{-i}) \leq u_i(z) \text{ for each } i = 1, 2 \right\}.$$

The *reduced Hannan set* H_R is the subset of H in which at least one regret is exactly zero for each player:

$$H_R = \left\{ z \in \Delta(A) \mid \max_{k \in A_i} u_i(k, z_{-i}) = u_i(z) \text{ for each } i = 1, 2 \right\}.$$

The next property of no-regret dynamics is straightforward by the definition of no-regret strategies and Remark 2 (see, e.g., Hart and Mas-Colell [26, Corollary 3.2]).

⁵ This additional property is needed for Corollary 1 below, but for our main results (ii') suffices.

Corollary 1. *For every no-regret dynamics in class \mathcal{R} , the trajectories converge almost surely to the reduced Hannan set.*

Convergence of the average play $z(t)$ to set H_R does not imply its convergence to any particular point in H_R . Moreover, even if $z(t)$ converges to a point, this point need not be a Nash equilibrium.

3. Fictitious play and no-regret dynamics

3.1. Fictitious play

In *discrete fictitious play*, in every period t after the initial one, player i plays a pure best reply $a_i(t)$ to the average past play of her opponent $x_{-i}(t - 1) := \frac{1}{t-1} \sum_{\tau=1}^{t-1} a_{-i}(\tau)$ (here $a_{-i}(\tau)$ is a vertex of $\Delta(A_{-i})$). The latter is called the *belief* of player i on her opponent’s next move. Formally, for any $x = (x_1, x_2)$ in $\Delta(A_1) \times \Delta(A_2)$, denote by $BR_i(x_{-i})$ player i ’s set of best replies to x_{-i} :

$$BR_i(x_{-i}) := \left\{ x_i \in \Delta(A_i) \mid u_i(x_i, x_{-i}) = \max_{k \in A_i} u_i(k, x_{-i}) \right\}, \quad i = 1, 2.$$

Let $BR(x) = BR_1(x_2) \times BR_2(x_1)$. A discrete-time trajectory $(x(t))_{t=1}^\infty$ on $\Delta(A_1) \times \Delta(A_2)$ is a solution of *discrete fictitious play* (DFP) if for every $t > 1$

$$x(t) - x(t - 1) = \frac{1}{t} (a(t) - x(t - 1)) \tag{3}$$

where $a(t) = (a_1(t), a_2(t))$ and $a_i(t) \in BR_i(x_{-i}(t - 1))$ is a vertex of $\Delta(A_i)$ associated with some pure best-reply action, $i = 1, 2$.

Analogously, an absolutely continuous function $x : [1, \infty) \rightarrow \Delta(A_1) \times \Delta(A_2)$ is a solution of *continuous fictitious play* (CFP) if for almost all $t \geq 1$, $x(t)$ is differentiable and

$$\dot{x}(t) = \frac{1}{t} (q(t) - x(t)),$$

where $q(t) \in BR(x(t))$ is now a profile of *mixed* actions. This may be written as the differential inclusion:

$$\dot{x}(t) \in \frac{1}{t} (BR(x(t)) - x(t)). \tag{4}$$

The average correlated play satisfies $z(t) := \frac{1}{t} (z(1) + \int_1^t q(\tau) d\tau)$ for some initial condition $z(1)$ such that $z_i(1) = x_i(1)$, $i = 1, 2$. Thus, for almost all t , $z(t)$ is differentiable and

$$\dot{z}(t) = \frac{1}{t} (\bar{q}(t) - z(t)), \tag{5}$$

where $\bar{q} = q_1 \otimes q_2 \in \Delta(A)$ is the product distribution corresponding to the mixed strategy profile $q = (q_1, q_2) \in \Delta(A_1) \times \Delta(A_2)$, and q_i is a best reply to z_{-i} .⁶

⁶ This definition of CFP guarantees that solutions exist in all games and for all initial conditions, and that by the change of time scale $y(t) = x(e^t)$, CFP corresponds to the *best-reply dynamics* [22,37] defined by $\dot{y} \in BR(y) - y$. Another definition of CFP (e.g., Monderer et al. [38, p. 445] and Berger [7, pp. 252–253]) considers only trajectories that are piecewise linear, such that $q_i(t)$ is always a pure action (technically, a vertex of $\Delta(A_i)$), and that the times at which $q(t)$ changes have no finite accumulation point. This restricted definition is easier to handle, but in many games there do not exist such trajectories from every initial condition.

In discrete or continuous fictitious play, the marginals $z_1(t), z_2(t)$ of the average past play are equal to the beliefs $x_1(t), x_2(t)$. By analogy, if $z(t)$ is the average past play generated by a no-regret dynamics, it is convenient to call $z_{-i}(t)$ the belief of player i about her opponent’s next move. This illuminates a crucial difference between fictitious play and no-regret dynamics in class \mathcal{R} : under fictitious play, a player chooses a *best reply* to her belief, whereas under no-regret dynamics, she chooses a *better reply* (“better” with respect to her average realized payoff).

3.2. Continuous fictitious play leads to no regret

It is well known that discrete fictitious play does not lead to *no regret* [47,25]. Since the continuous fictitious play process is a continuous-time version of DFP, intuitively, it should not lead to *no regret* either. The following result – a generalization of Theorem D of Monderer et al. [38] – shows that this intuition is misleading.

Proposition 2. *Under any solution of continuous fictitious play, the average correlated play converges to the reduced Hannan set.*

This discrepancy between DFP and CFP may be explained as follows. Playing an action with positive regret decreases the regret for this action. In CFP, roughly, when an action is played it remains a best reply, hence it is associated with maximal regret for a small time increment. Precisely, the derivative of the regret for the action played is equal to the derivative of the maximal regret. Since the regret for this action decreases, so does the maximal regret. In contrast, in DFP, an action played at stage t has maximal regret at stage t , but not necessarily at stage $t + 1$. Thus the fact that the regret for this action decreases does not entail that the maximal regret does.

Proof of Proposition 2. For comparison with Hart and Mas-Colell [26, Theorem 3.1], rescale time (let $\tilde{t} = \exp t$) so that (5) becomes $\dot{z} = \bar{q} - z$. For any mixed action $\sigma_i \in \Delta(A_i)$, let

$$R_{i,\sigma_i}(t) := \sum_{k \in A_i} \sigma_i(k) R_{i,k}(t) = u_i(\sigma_i, z_{-i}(t)) - u_i(z(t)).$$

Let $v_i(t) = R_{i,\max}(t)$. Note that $R_{i,k}$ is Lipschitz continuous for all k in A_i . Thus it follows from Theorem A.4 of Hofbauer and Sandholm [27] that, for almost all t , v_i and $R_{i,k}$ are differentiable, and for all k such that $q_{i,k}(t) > 0$, we have $\dot{v}_i(t) = \dot{R}_{i,k}(t)$. It follows that $\dot{v}_i = \sum_k q_{i,k} \dot{R}_{i,k} = \dot{R}_{i,q_i}$. Furthermore:

$$\begin{aligned} \dot{R}_{i,q_i} &= u_i(q_i, \dot{z}_{-i}) - u_i(\dot{z}) = u_i(q_i, q_{-i} - z_{-i}) - u_i(\bar{q} - z) \\ &= -[u_i(q_i, z_{-i}) - u_i(z)] = -R_{i,q_i} = -v_i. \end{aligned}$$

Thus, $\dot{v}_i = -v_i$. Therefore, $v_i(t)$ converges to zero for all $i = 1, 2$, hence $z(t) \rightarrow H_r$. \square

Remark 3. In the proof, we did not use that q_{-i} is a best reply to z_i . This shows that the fact that CFP leads to no regret is a unilateral property. That is, if a player’s behavior evolves according to CFP, then she has no asymptotic regret, independently of her opponent’s behavior (see also Monderer et al. [38, p. 445]).

Remark 4. CFP and the best-reply dynamics converge to the set of Nash equilibria in finite zero-sum games [29]. The usual proof is to show that the “duality gap” $W(x) = \max_{k \in A_1} u_1(k, x_2) -$

$\min_{s \in A_2} u_1(x_1, s)$ converges to zero. This follows from the above proof, since in a two-player zero-sum game $W(x(t)) = R_{1,\max}(z(t)) + R_{2,\max}(z(t))$, where x is a solution of CFP and z the associated correlated play.

3.3. No-regret dynamics is perturbed CFP

In the previous subsection we showed that CFP leads to no regret. Conversely, we now show that any no-regret dynamics in class \mathcal{R} (as defined in Section 2) is closely related to CFP. We first explain the intuition. Denote by $BR_i^\varepsilon(x_{-i})$ the set of ε -best replies of player i to the mixed action x_{-i} of her opponent:

$$BR_i^\varepsilon(x_{-i}) = \left\{ x_i \in \Delta(A_i) \mid u_i(x_i, x_{-i}) \geq \max_{k \in A_i} u_i(k, x_{-i}) - \varepsilon \right\}, \quad i = 1, 2.$$

The crucial observation is the following.

Lemma 1. *Assume that the maximal regret is less than ε . Then any action with positive regret is an ε -best reply to the average play of the opponent.*

Proof. If player i has positive regret for action a_i at some $z \in \Delta(A)$, then $u_i(z) - u_i(a_i, z_{-i}) < 0$. But by assumption $\max_{k \in A_i} u_i(k, z_{-i}) - u_i(z) \leq \varepsilon$. Therefore, $\max_{k \in A_i} u_i(k, z_{-i}) - u_i(a_i, z_{-i}) < \varepsilon$, and a_i is an ε -best reply to z_{-i} . \square

Since no-regret dynamics in class \mathcal{R} only pick actions with positive regret, they only pick ε -best replies to the average play of the others, where ε is the maximal regret. Since this maximal regret approaches zero almost surely, eventually only almost-exact best replies are picked. This provides the intuition why no-regret dynamics and fictitious play may exhibit similar asymptotic behavior. Finding a precise link, however, is not obvious. For instance, there could exist actions that are ε_t -best replies in each period t , with $\varepsilon_t \rightarrow 0$, but never exact best replies. Thus a limit play of no-regret dynamics may include such actions, but this cannot happen under fictitious play. Yet, as we show below, the asymptotic behavior of no-regret dynamics and CFP can be formally related using the theory of perturbed differential inclusions [4,5].

Before stating a precise result, we need some definitions. A set $L \subset \Delta(A_1) \times \Delta(A_2)$ is *invariant* under CFP if for every initial point $x \in L$ there exists a solution $x(\cdot)$ of CFP, defined for all $t > 0$ (not only $t \geq 1$) and such that $x(1) = x$ and $x(t) \in L$ for all $t > 0$. A nonempty compact invariant set is an *attractor* if it attracts uniformly all trajectories starting in its neighborhood. An invariant set L is *attractor-free* if no proper subset of L is an attractor for the dynamics restricted to L . A nonempty compact set L is *internally chain transitive* (ICT) for continuous fictitious play if every pair of points in L can be connected by finitely many arbitrarily long pieces of orbits of CFP lying completely within L with arbitrarily small jumps between them.⁷ Every ICT set is invariant and attractor-free [5, Property 2]. The *limit set of the beliefs* of a trajectory $z(t)$ on $\Delta(A_1 \times A_2)$ is the set of all accumulation points of its marginals $(z_1(t), z_2(t)) \in \Delta(A_1) \times \Delta(A_2)$ as $t \rightarrow \infty$.

⁷ For the formal definitions of *attractor* and *attractor-free set* see Benaïm et al. [5, p. 675]; for the definition of *ICT* see Benaïm et al. [4, p. 337]. Note that the definition of invariance in Benaïm et al. [4,5] applies to the best-reply dynamics, so an appropriate time rescaling must be used to apply it to CFP (see footnote 6). This explains that their definition considers solutions defined for all $t \in \mathbb{R}$ while ours considers solutions defined for all $t > 0$.

Theorem 1. For every no-regret dynamics in class \mathcal{R} , the limit set of the beliefs is almost surely internally chain transitive for continuous fictitious play.⁸

We give here a sketch of the proof. The details are given in [Appendix A](#). A discrete-time trajectory $(x_1(t), x_2(t))_{t=1}^{\infty}$ on $\Delta(A_1) \times \Delta(A_2)$ is a *payoff-perturbed* DFP trajectory if there exists a positive sequence (ε_t) converging to zero such that (3) holds and $a_i(t)$ is a vertex of $\Delta(A_i)$ associated with a pure ε_t -best reply to $x_{-i}(t-1)$, for all $i = 1, 2$ and all $t > 1$. A no-regret dynamics in class \mathcal{R} generates a trajectory $(z(t))_{t=1}^{\infty}$ on $\Delta(A)$ and an associated *sequence of beliefs* $(z_1(t), z_2(t))$ on $\Delta(A_1) \times \Delta(A_2)$. Building on [Lemma 1](#), we show that this sequence of beliefs is almost surely a payoff-perturbed DFP trajectory. By an auxiliary lemma, this implies that this is almost surely a *graph-perturbed* DFP trajectory: a notion similar to payoff-perturbed trajectory, but for another definition of perturbed best reply, the one used in the theory of perturbed differential inclusions [4,5]. It follows that the continuous-time interpolation of this sequence of beliefs is almost surely a perturbed solution of CFP, in the sense of Benaïm et al. [4]. [Theorem 1](#) then follows from [Theorem 3.6](#) of Benaïm et al. [4].

Since ICT sets are invariant, a consequence of [Theorem 1](#) is the following:

Corollary 2. Let A be the global attractor of CFP (i.e., its maximal invariant set, see Benaïm et al. [4]). For any no-regret dynamics in class \mathcal{R} , the limit set of the beliefs is almost surely a subset of A .

Note the similarity with [Propositions 5.1 and 5.2](#) of Hofbauer et al. [30], who study the links between the time average of the replicator dynamics and CFP.

3.4. Applications of [Theorem 1](#) and comments

[Theorem 1](#) allows for alternative and sometimes much shorter proofs of most known convergence properties of no-regret dynamics. Below, we write that no-regret dynamics converge to some set E if the limit set of the beliefs is almost surely a subset of E .⁹

(a) For any game which is best-reply equivalent to a two-person zero-sum game, the global attractor of CFP is the set of Nash equilibria [29]. Hence all no-regret dynamics in class \mathcal{R} converge to the set of Nash equilibria. Actually, in zero-sum games, if the correlated action z is in the Hannan set (recall that this is the set of correlated actions that satisfy *no regret* for all players), then (z_1, z_2) is a Nash equilibrium. Consequently, in zero-sum games all dynamics that lead to no regret (not only those in class \mathcal{R}) converge to the set of Nash equilibria. This holds more generally for *stable bimatrix games* [27], because these are rescaled zero-sum games in the sense of Hofbauer and Sigmund [28], as is easily shown and was known to Josef Hofbauer (private communication).

(b) For games with strictly dominated strategies, the global attractor of CFP is contained in the face of the simplex with no weight on these strategies. Hence all no-regret dynamics in class

⁸ In the statement of [Theorem 1](#), CFP can be replaced by the best-reply dynamics since they clearly have the same ICT sets (see footnote 6).

⁹ Note that some applications of [Theorem 1](#) (points (a), (b) and (c) below) lead to the same conclusions about no-regret dynamics as those about the time average of the replicator dynamics described in Hofbauer et al. [30, p. 267, points (2), (3) and (4)].

	A	B	C
A	2	1	-4
B	1	0	-1
C	-4	-1	-2

(i)

	A	A ⁻	B	B ⁻
A	1	1	0	0
A ⁻	1 - ε	1 - ε	-ε	-ε
B	0	0	1	1
B ⁻	-ε	-ε	1 - ε	1 - ε

(ii)

Fig. 1. Example 1.

	A	B	C
A	0, 0	1, 0	0, 1
B	0, 1	0, 0	1, 0
C	1, 0	0, 1	0, 0

Fig. 2. Example 2.

\mathcal{R} converge to this face. Similarly, these dynamics converge to the unique Nash equilibrium in strictly dominance solvable games.

Contrary to (a), this need not be true for all dynamics that lead to no regret. Indeed, convergence to the Hannan set or even to the reduced Hannan set does not guarantee elimination of strictly dominated strategies. Consider, for instance, the games shown on Fig. 1. Both games are symmetric, so we indicate only the payoffs of the row player. Game (i) is an identical interest game which is strictly dominance solvable; yet the correlated action putting probabilities 1/3 on each diagonal square is in the reduced Hannan set. For $\varepsilon = 0$, game (ii) is a coordination game with duplicate strategies. For $\varepsilon > 0$, the duplicates A^- , B^- are penalized and become strictly dominated. Thus, the correlated action putting probability 1/2 on (A^-, A^-) and 1/2 on (B^-, B^-) puts only weight on strictly dominated actions. Yet, for $\varepsilon \leq 1/2$, it belongs to the Hannan set.¹⁰

(c) In weighted potential games, all internally chain transitive sets of CFP are (subsets of) connected components of Nash equilibria on which the payoffs are constant, see [4, Theorem 5.5 and Remark 5.6]. Hence by Theorem 1, all no-regret dynamics in class \mathcal{R} converge to such components. Note that the original proof is much longer [26, Appendix A].

(d) If the beliefs $(z_1(t), z_2(t))$ of a no-regret dynamics converge to the set of Nash equilibria, then the average realized payoff converges to the set of Nash equilibrium payoffs. To see why this is true, let $\hat{z} \in \Delta(A)$ be a limit point of $\{z(t)\}$ and let the marginals $(\hat{z}_1, \hat{z}_2) \in \Delta(A_1) \times \Delta(A_2)$ constitute a Nash equilibrium. By Corollary 1 the maximal regret converges to zero, so for every $i = 1, 2$

$$u_i(\hat{z}) = \max_{k \in A_i} u_i(k, \hat{z}_{-i}) = u_i(\hat{z}_i, \hat{z}_{-i}).$$

This result illuminates an important difference between no-regret dynamics and discrete fictitious play. It is well known that under DFP, if the beliefs of the players converge to a Nash equilibrium, their average realized payoffs need not approach the set of Nash equilibrium payoffs, whereas under no-regret dynamics it is always the case.

(e) The game of Fig. 2 has a unique equilibrium, but almost all solutions of CFP converge to a hexagon [21,41,42]. It may be shown that the only ICT sets are the Nash equilibrium and this

¹⁰ See also the game of Moulin and Vial [39, p. 205], where the third strategy of player 1 is strictly dominated but has a positive marginal probability under some correlated actions in the Hannan set.

hexagon. Consequently, the limit set of any no-regret dynamics in class \mathcal{R} is almost surely one of these two sets.

(f) In a number of classes of games, convergence of discrete fictitious play to the set of Nash equilibria has been established, but analogous results for continuous fictitious play are lacking. Thus we cannot use [Theorem 1](#). These classes of games include generic $2 \times n$ games [6], generic ordinal potential games, quasi-supermodular games¹¹ with diminishing returns [7], and some other special classes (see, e.g., Sparrow et al. [42, p. 260]). For ordinal potential games and quasi-supermodular games with diminishing returns, Berger [7] proves convergence to the set of Nash equilibria of *some* solutions of continuous fictitious play as defined by (4) (see our footnote 6). This is not enough to apply the results of Benaïm et al. [4]. The same problem arises in Krishna and Sjöström [32]. Actually, as explained below, convergence of CFP to the set of Nash equilibria would not suffice to use [Theorem 1](#): we would need some additional structure, such as a Lyapunov function, to get more information on the ICT sets.

(g) Consider a bimatrix game in which all solutions of CFP converge to the set of Nash equilibria. Because the definition of attractor requires uniform attraction, this does not imply that the set of Nash equilibria is an attractor. Neither does it imply that all ICT sets are contained in the set of Nash equilibria, as shown in Appendix A.2 of the working paper version of this note [45]. Therefore, we cannot deduce from [Theorem 1](#) that no-regret dynamics in class \mathcal{R} converge to the set of Nash equilibria; whether this is always the case remains an open question.

(h) We show in Section 5 that [Theorem 1](#) also applies, and under weaker assumptions, to the continuous-time version and to the expected version of no-regret dynamics in class \mathcal{R} . As apparent from the proof, the existence of a potential is not essential: for a good choice of behavior when there are no regrets, [Theorem 1](#) holds for any no-regret dynamics such that a player always chooses an action with positive regret whenever he has one. It also applies to certain no-regret dynamics that do not have this property, such as the exponential weight algorithm (see [Remark 6](#) at the end of [Appendix A](#)).

(i) [Theorem 1](#) extends to n -player games with linear incentives [40], also known as polymatrix games [46], but not to general n -player games. The main issue is technical: for n -player games, no-regret dynamics are related to a correlated version of fictitious play. This correlated version is defined through a best-reply correspondence which is not convex valued. For this reason, we cannot use the current theory of perturbed differential inclusions [4], which applies only to convex valued correspondences. More details are given in Viossat and Zapechelnyuk [45].

4. Curb sets

[Theorem 1](#) does not answer whether *attracting sets* of CFP have an analogous property under no-regret dynamics.

A set $\mathcal{C} \subset \Delta(A)$ is *eventually attracting* under a no-regret dynamic process if with any given probability it captures all no-regret trajectories originating from a small enough neighborhood of \mathcal{C} at all distant enough periods. Formally, \mathcal{C} is eventually attracting if for every $\pi > 0$ there exists $\varepsilon > 0$ and a period T such that: for every $t_0 \geq T$, if $z(t_0)$ is in an ε -neighborhood of \mathcal{C} , then $z(t)$ converges to set \mathcal{C} with probability at least $1 - \pi$.¹²

¹¹ Also known as games of strategic complementarities (e.g., Tirole [44]).

¹² We say that $z(t)$ converges to \mathcal{C} if $\inf_{c \in \mathcal{C}} \|z(t) - c\| \rightarrow 0$ as $t \rightarrow \infty$.

For this section it is convenient to replace assumption (Q2) by the following one:

If a player's maximal regret is nonpositive, then she plays a best reply to the empirical distribution of her opponent. (Q2')

This is not essential, since the interesting histories are those where both players have positive regrets, in which case (Q2) plays no role.¹³

A strict Nash equilibrium is eventually attracting. Indeed, if $z(t_0)$ is close enough to a vertex of $\Delta(A)$ corresponding to a strict Nash equilibrium $a = (a_1, a_2)$, then for each player i , action a_i is the unique best reply and there is a negative regret for any action other than a_i . Since by (R3) only actions with positive regret can be chosen, and by (Q2') only best-reply actions can be chosen if all regrets are nonpositive, action a_i will be played by each player i in the following period, and so on.

Let us now consider a standard generalization of strict Nash equilibria. For each $i = 1, 2$, let $B_i \subset A_i$. With a slight abuse of notation, denote by $\Delta(B_i)$ the set of probability measures on A_i with support on B_i only. The product set $B = B_1 \times B_2$ is closed under rational behavior (curb) (Basu and Weibull [2]) if

$$BR_i(x_{-i}) \subset \Delta(B_i) \quad \text{whenever } x_{-i} \in \Delta(B_{-i}), \quad i = 1, 2.$$

That is, the set B is curb if the players' pure best-reply profiles are contained in B whenever they believe that no actions outside of B should be played.

Curb sets are known to be attracting under CFP (e.g., Balkenborg et al. [1, Lemma 7]). However, they need not be attracting under no-regret dynamics in class \mathcal{R} . Indeed, even if the support of $z(t_0)$ is contained in some curb set B , there may be positive regrets for actions outside of B , since B need not be closed under better replies. However, we show that the intersection of the Hannan set and the set of correlated actions with support on a curb set is eventually attracting.

Formally, let $B = B_1 \times B_2$ be a curb set. Let $\Delta_B(A)$ denote the set of correlated actions with support on B only. Let $H_B = H \cap \Delta_B(A)$.

Proposition 3. *For every curb set B , the set H_B is eventually attracting under every no-regret dynamics in \mathcal{R} .*

The proof is based on the following observations. For every curb set B , if the average play is close enough to H_B , then regrets for all actions outside of B are negative (since B is curb). Hence, by condition (R3), only actions in B will be played in the immediate future. On the other hand, almost sure convergence of maximum regret to zero suggests that, so long as the players choose only actions in B , the average play will approach H_B , thus reinforcing the former observation. To prove the result, however, we need to establish bounds on the maximal future regret conditional on certain histories (namely, conditional on being close to H_B) that Hart and Mas-Colell [25] do not provide. For the complete proof we refer the interested reader to Viossat and Zapechelnyuk [45, Appendix A.3].

¹³ Recall that by Remark 1, if a player has positive maximal regret, then it remains positive forever. So we can consider histories from a distant enough period t_0 where both players have positive regrets and (Q2) plays no role. If t_0 does not exist, i.e., some player always has nonpositive maximal regret, then Proposition 1 and (Q2) imply that her play is constant, whereas her opponent's play must approach a best reply to it, leading to Nash equilibrium. By replacing (Q2) by (Q2') we avoid dealing with this issue.

5. Continuous-time and expected no-regret dynamics

We now prove an analog of [Theorem 1](#) for continuous-time dynamics [[26](#)] and the expected version of discrete-time dynamics. Both describe trajectories of average *intended (mixed)* play, rather than average realized (pure) play. For this reason, condition (R3) is not needed. Indeed, the interest of (R3) is that, together with (Q1), it requires every realized action to be a better reply to the opponents empirical distribution of play (whenever such actions exist). But now we only need every mixed (expected) action to be a better reply, and this follows already from conditions (R1)–(R2) and (Q1). Besides, these dynamics are deterministic, hence the results we obtain hold *surely* (not just *almost surely*). The proofs are based on [Appendix A](#) and are best understood after reading it.

Consider a continuous-time dynamics

$$\dot{z}(t) = \frac{1}{t}(\bar{q}(t) - z(t)), \tag{6}$$

where $\bar{q}(t) = q_1(t) \otimes q_2(t) \in \Delta(A)$ is the (independent) joint play at time t and $z(t)$ the average correlated play. There are two differences with (1): time is now continuous, and, more importantly, realized play $a(t)$ has been replaced by intended mixed play $\bar{q}(t)$. As in CFP, start at time 1 with some initial condition $z(1) \in \Delta(A)$. Assume that whenever $R_{i,\max}(t) > 0$

$$q_{i,k}(t) = \frac{\nabla_k P_i(R_i(t))}{\sum_{s \in A_i} \nabla_s P_i(R_i(t))}, \quad k \in A_i, \tag{7}$$

where P_i is a C^1 potential function satisfying (R1), (R2) and the technical condition:

(P4') There exists $0 < \rho_2 < \infty$ such that $\nabla P_i(x) \cdot x \leq \rho_2 P_i(x)$ for all $x \in \mathbb{R}_-^{A_i}$.

This is a part of condition (P4) in Hart and Mas-Colell [[26](#)].

Proposition 4. *Let $z(t)$ be a solution of (6) and (7) with P_i satisfying conditions (R1), (R2) and (P4') for all $i = 1, 2$. Assume that the initial condition $z(1)$ is such that both players have some positive regrets: $R_{i,\max}(1) > 0$ for all $i = 1, 2$. Then the limit set of the beliefs is internally chain transitive for continuous fictitious play.*

Proof. Let $\varepsilon_i(t) := R_{i,\max}(t)$. Hart and Mas-Colell [[26](#), Theorem 3.1 and Lemma 3.3¹⁴] show that if $\varepsilon_i(1) > 0$, then $\varepsilon_i(t) > 0$ for all t , and $\varepsilon_i(t) \rightarrow 0$ as $t \rightarrow +\infty$. Moreover, by (R2) applied to $x = R_i(t)$ and definition of q_i , we have: $u_i(q_i, z_{-i}) - u_i(z) = q_i \cdot R_i > 0$ (this is Eq. (3.3) in [[25](#)]). Thus by [Lemma 1](#), $q_i \in BR_i^{\varepsilon_i(t)}(z_{-i})$. Together with [Lemma 3](#) in [Appendix A](#), this implies that $(z_1(\cdot), z_2(\cdot))$ is a perturbed solution of CFP in the sense of Benaïm et al. [[4](#)]. The result then follows from Theorem 3.6 of Benaïm et al. [[4](#)]. \square

Remark 5. Assume that if all initial regrets of a player are nonpositive then the dynamics is defined as in Hart and Mas-Colell [[26](#), Eq. (4.9)]. Then it is easily seen that the result of [Proposition 4](#) holds for any initial condition $z(1)$.

¹⁴ Note a typo in the proof of Lemma 3.3 in Hart and Mas-Colell [[26](#)]: (P3) should be replaced by (P4). Moreover, only our condition (P4') is used in the proof of Lemma 3.3 in [[26](#)].

	L	R
T	1	0
C	0	1
B	$\frac{1}{2} - \eta$	$\frac{1}{2} - \eta$

Fig. 3. Example 3.

Expected discrete-time dynamics. The expected motion in (1) is described by

$$z(t) - z(t - 1) = \frac{1}{t}(\bar{q}(t) - z(t - 1)),$$

where $\bar{q}(t) = q_1(t) \otimes q_2(t)$ is the expectation of $a(t)$. Assume that q_i is derived by (Q1) from a potential function satisfying (R1)–(R2). Let $\varepsilon_i(t) := R_{i,\max}(t)$. It is easily seen that, as for continuous-time dynamics, $\varepsilon_i(t) \rightarrow 0$ as $t \rightarrow +\infty$, and if $\varepsilon_i(1) > 0$, then for all t , $\varepsilon_i(t) > 0$ and $q_i \in BR_i^{\varepsilon_i(t)}(z_{-i})$. Due to Lemmata 3 to 5 of Appendix A and to Theorem 3.6 of Benaïm et al. [4], it follows that for a good choice of behavior when all regrets are initially nonpositive, the limit set of the beliefs is internally chain transitive for CFP.

Appendix A. Proof of Theorem 1

Denote by $\widehat{BR}_i^\varepsilon(x)$ the correspondence whose graph is the ε -neighborhood of the graph of BR_i :

$$\widehat{BR}_i^\varepsilon(x_{-i}) = \left\{ x_i \in \Delta(A_i) \mid \begin{array}{l} \exists (x_i^*, x_{-i}^*) \in \Delta(A_1) \times \Delta(A_2) \text{ s.t.} \\ x_i^* \in BR_i(x_{-i}^*), \text{ and } \|(x_i^*, x_{-i}^*) - (x_i, x_{-i})\|_\infty \leq \varepsilon \end{array} \right\}.$$

Let $\widehat{BR}^\varepsilon(x) = \widehat{BR}_1^\varepsilon(x_2) \times \widehat{BR}_2^\varepsilon(x_1)$. In words, action x_i is an ε -graph perturbed best reply to x_{-i} if there is an action ε -close to x_i which is an exact best reply to an action ε -close to x_{-i} . This is the notion of perturbation used in the theory of perturbed differential inclusions (Benaïm et al. [4,5]). As illustrated by the example below, it is different from the notion of perturbation of payoffs in the ε -best reply correspondence, i.e. $BR^\varepsilon(x) = BR_1^\varepsilon(x_2) \times BR_2^\varepsilon(x_1)$ with

$$BR_i^\varepsilon(x_{-i}) = \left\{ x_i \in \Delta(A_i) \mid u_i(x_i, x_{-i}) \geq \max_{k \in A_i} u_i(k, x_{-i}) - \varepsilon \right\}, \quad i = 1, 2.$$

Consider a game where the payoffs of player 1 are given by Fig. 3. Let $\varepsilon \in (0, 1/2)$ and let $x_2^\varepsilon = (\frac{1}{2} + \varepsilon)L + (\frac{1}{2} - \varepsilon)R$. The pure action C is a 2ε -best reply to x_2^ε . Using the sup norm, it is at distance 1 from pure action T , the unique exact best reply to x_2^ε . Nevertheless, C is an ε -graph perturbed best reply, because it is an exact best reply to x_2^0 , which is ε -close (in sup norm) to x_2^ε . By contrast, for all $\eta > 0$, action B is an $(\varepsilon + \eta)$ -best reply, but only a 1-graph perturbed best reply to x_2^ε .

A discrete-time trajectory $(x_1(t), x_2(t))_{t=1}^\infty$ on $\Delta(A_1) \times \Delta(A_2)$ is a *payoff-perturbed* fictitious play trajectory if there exists a positive sequence (ε_t) converging to zero such that

$$x(t) - x(t - 1) = \frac{1}{t}(q(t) - x(t - 1))$$

with $q(t) = (q_1(t), q_2(t))$ and $q_i(t) \in BR_i^{\varepsilon_t}(x_{-i}(t - 1))$ for all $i = 1, 2$ and all $t > 1$. It is a *graph-perturbed* fictitious play trajectory if the same holds but replacing $BR_i^{\varepsilon_t}$ with $\widehat{BR}_i^{\varepsilon_t}$. A trajectory $(z(t))_{t=1}^\infty$ on $\Delta(A)$ generates a *sequence of beliefs* $(z_1(t), z_2(t))$ in $\Delta(A_1) \times \Delta(A_2)$.

The proof goes as follows. **Lemma 2** shows that the sequence of beliefs generated by a no-regret dynamics is a payoff-perturbed FP trajectory. Together with **Lemma 3**, this implies that it is a graph-perturbed FP trajectory (**Lemma 4**). It follows that the interpolated process of a no-regret dynamics trajectory is a perturbed solution of CFP (**Lemma 5**). The result then follows from Benaïm et al. [4].

Lemma 2. *The sequence of beliefs of a solution of a no-regret dynamics in class \mathcal{R} is almost surely a payoff-perturbed DFP trajectory.*

Proof. If $R_{i,\max}(t) \leq 0$ for all t , then by **Remark 2**, player i plays an $\varepsilon(t)$ -best reply for some $\varepsilon(t)$ converging to zero. Otherwise, $R_{i,\max}(t_0) > 0$ for some $t_0 \in \mathbb{N}^*$. Then for all times $t > t_0$, $R_{i,\max}(t) > 0$ (by **Remark 1**) and player i plays an $R_{i,\max}(t)$ -best reply by **Lemma 1** and conditions (R3) and (Q1). Since $R_{i,\max}(t) \rightarrow 0$ almost surely, the result follows. \square

Lemma 3. *Let X be a compact subset of \mathbb{R}^m and F a correspondence from X to itself. For any $\delta \geq 0$, let $\hat{F}_\delta : X \rightrightarrows X$ denote the correspondence whose graph is the δ -neighborhood of the graph of F :*

$$\hat{F}_\delta(x) = \{y \in X \mid \exists(x^*, y^*) \in X^2 \text{ s.t. } y^* \in F(x^*) \text{ and } \|(x^*, y^*) - (x, y)\|_\infty \leq \delta\}.$$

For any $\alpha > 0$, let G_α be a u.s.c. correspondence from X to itself. Assume that for each x in X :

- (i) $\alpha < \alpha' \Rightarrow G_\alpha(x) \subset G_{\alpha'}(x)$ (that is, $(G_\alpha)_{\alpha>0}$ is increasing w.r.t. inclusion);
- (ii) $\bigcap_{\alpha>0} G_\alpha(x) \subset F(x)$.

Then for every $\delta > 0$ there exists $\alpha > 0$ such that for each x in X , $G_\alpha(x) \subset \hat{F}_\delta(x)$.

Proof. By contradiction, assume that there exists $\delta > 0$, a decreasing sequence (α_n) converging to zero, and sequences (x_n) and (y_n) of points in X such that $y_n \in G_{\alpha_n}(x_n) \setminus \hat{F}_\delta(x_n)$ for all n . By compactness of X , we can assume that (x_n) and (y_n) converge respectively to x^* and y^* . Fix $k \in \mathbb{N}$. For all $n \geq k$, $y_n \in G_{\alpha_n}(x_n) \subset G_{\alpha_k}(x_n)$ by (i). Since G_{α_k} is u.s.c., it follows that $y^* \in G_{\alpha_k}(x^*)$. Therefore, by (i) and (ii)

$$y^* \in \bigcap_{k \in \mathbb{N}} G_{\alpha_k}(x^*) = \bigcap_{\alpha > 0} G_\alpha(x^*) \subset F(x^*).$$

But for n large enough, $\|(x^*, y^*) - (x_n, y_n)\|_\infty < \delta$, hence $y_n \in \hat{F}_\delta(x_n)$, a contradiction. \square

Applied to the best-reply correspondence, **Lemma 3** implies that for any $\delta > 0$, an ε -perturbed best reply is a δ -graph perturbed best reply, provided ε is small enough. Thus we have the next result.

Lemma 4. *Any payoff-perturbed DFP trajectory is a graph-perturbed DFP trajectory.*

Proof. Let $\varepsilon_t \rightarrow 0$. Let

$$\delta_t = \min\{\delta \geq 0 \mid \forall i = 1, 2, \forall x \in \Delta(A_1) \times \Delta(A_2), BR_i^{\varepsilon_t}(x_{-i}) \subset \widehat{BR}_i^\delta(x_{-i})\}.$$

Applying **Lemma 3** with $X = \Delta(A_1) \times \Delta(A_2)$, $G_\varepsilon = BR^\varepsilon$ and $F = BR$, we obtain that $\delta_t \rightarrow 0$. The result follows. \square

Given a discrete-time trajectory $x(n) = (x_1(n), x_2(n))$ on $\Delta(A_1) \times \Delta(A_2)$, with $n \in \mathbb{N}^*$, define its interpolated process $x : [1, +\infty) \rightarrow \Delta(A_1) \times \Delta(A_2)$ as follows. For all $t \in [n, n + 1)$ let $tx(t) = nx(n) + (t - n)q(n)$, where $q_i(n) = (n + 1)x_i(n + 1) - nx_i(n)$, $i = 1, 2$. This is equivalent to

$$x_i(t) - x_i(n) = \frac{t - n}{t}(q_i(n) - x_i(t)), \quad i = 1, 2.$$

Hence for all $t \in (n, n + 1)$ we have $\|x(t) - x(n)\|_\infty \leq \frac{1}{n+1}$ and

$$\dot{x}(t) = \frac{1}{t}(q(n) - x(t)). \tag{8}$$

An absolutely continuous function $x : [1, +\infty) \rightarrow \Delta(A_1) \times \Delta(A_2)$ is a perturbed solution of CFP if there exists a vanishing function $\varepsilon : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for almost all t ,

$$\dot{x} \in \frac{1}{t}(\widehat{BR}^{\varepsilon(t)}(x) - x) \quad \text{where } x = x(t). \tag{9}$$

Lemma 5. *The interpolated process of a graph-perturbed DFP trajectory is a perturbed solution of CFP.*

Proof. Consider a discrete-time trajectory $(x_1(n), x_2(n))_{n \in \mathbb{N}}$ such that

$$x_i(n) - x_i(n - 1) = \frac{1}{n}(q_i(n) - x_i(n)), \quad i = 1, 2,$$

with $q_i(n) \in \widehat{BR}_i^{\varepsilon_n}(x_{-i}(n - 1))$ and $\varepsilon_n \rightarrow 0$. For all n and all $t \in [n, n + 1)$, let $\varepsilon(t) = \varepsilon_n + 2/n$. Obviously, $\varepsilon(t) \rightarrow 0$ as $t \rightarrow \infty$. Moreover, for all $t \in (n, n + 1)$, the interpolated process satisfies $\|x_{-i}(t) - x_{-i}(n - 1)\|_\infty \leq \frac{1}{n+1} + \frac{1}{n} < 2/n$, so $q_i(n) \in \widehat{BR}_i^{\varepsilon(t)}(x_{-i}(t))$. Therefore (8) implies (9) (see also Faure and Roth [16, Proposition 2.2]). \square

We can now prove Theorem 1. By Lemmata 2 and 4, the sequence of beliefs of a solution of a no-regret dynamics in class \mathcal{R} is almost surely a graph-perturbed DFP trajectory. Hence, by Lemma 5, its interpolated process $x(t)$ is a perturbed solution of CFP. This implies that $x(e^t)$ is almost surely a perturbed solution of the best-reply dynamics, in the sense of Benaïm et al. [4, Definition II]. Theorem 1 now follows from Theorem 3.6 of Benaïm et al. [4].¹⁵

Remark 6. Assume that at stage t , for each $i = 1, 2$, player i chooses a pure action according to a mixed action $q_i(t)$ that depends on the previous history $h(t - 1)$. Do not assume conditions (R1)–(R3) and (Q1), but assume that there exists a vanishing sequence (ε_t) such that for all $t > 1$ and any previous history $h(t - 1)$, $q_i(t) \in BR_i^{\varepsilon_t}(z_{-i}(t - 1))$, $i = 1, 2$. Then it follows from Lemma 3, the above proof and Benaïm et al. [4, Proposition 1.4 and a variant of Proposition 1.3] that Theorem 1 applies. As is well known, this is the case for the exponential weights algorithm [19,35] that corresponds to

$$q_{i,k}(t) := \frac{\exp \beta_t u_i(k, z_{-i})}{\sum_{s \in A_i} \exp(\beta_t u_i(s, z_{-i}))}$$

¹⁵ The definition of perturbed solution in Benaïm et al. [4] is different from ours but equivalent.

with $z_{-i} = z_{-i}(t-1)$, $\beta_t \rightarrow +\infty$ as $t \rightarrow \infty$, and $\beta_t < t^\alpha$ for some $\alpha \in (0, 1)$ to ensure that this is a no-regret dynamics (see, e.g., Benaïm and Faure [3]). The above assumptions are not (or not trivially) satisfied by no-regret dynamics in class \mathcal{R} . Indeed, the rate at which the maximal regret vanishes, hence the value ε_t such that $q_i(t) \in BR_i^{\varepsilon_t}(z_{-i}(t))$, may depend on the trajectory.

References

- [1] D. Balkenborg, J. Hofbauer, C. Kuzmics, Refined best reply correspondence and dynamics, *Theoretical Econ.*, forthcoming, available at <http://econtheory.org/ojs/index.php/te/article/viewForthcomingFile/652/6253/1>.
- [2] K. Basu, J.W. Weibull, Strategy subsets closed under rational behavior, *Econ. Letters* 36 (1991) 141–146.
- [3] M. Benaïm, M. Faure, Consistency of vanishingly smooth fictitious play, *Math. Operations Res.*, forthcoming, <http://dx.doi.org/10.1287/moor.1120.0568>.
- [4] M. Benaïm, J. Hofbauer, S. Sorin, Stochastic approximations and differential inclusions, *SIAM J. Control Optim.* 44 (2005) 328–348.
- [5] M. Benaïm, J. Hofbauer, S. Sorin, Stochastic approximations and differential inclusions. Part II: Applications, *Math. Operations Res.* 31 (2006) 673–695.
- [6] U. Berger, Fictitious play in $2 \times n$ games, *J. Econ. Theory* 120 (2005) 139–154.
- [7] U. Berger, Two more classes of games with the continuous-time fictitious play property, *Games Econ. Behav.* 60 (2007) 247–261.
- [8] D. Blackwell, An analog of the minmax theorem for vector payoffs, *Pacific J. Math.* 6 (1956) 1–8.
- [9] A. Blum, E. Even-Dar, K. Ligett, Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games, in: *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing*, 2006, pp. 45–52.
- [10] G. Brown, Iterative solutions of games by fictitious play, in: T. Koopmans (Ed.), *Activity Analysis of Production and Allocation*, in: *Cowles Commission Monogr.*, vol. 13, Wiley, New York, 1951, pp. 374–376.
- [11] N. Cesa-Bianchi, G. Lugosi, Potential-based algorithms in on-line prediction and game theory, *Mach. Learn.* 51 (2003) 239–261.
- [12] N. Cesa-Bianchi, G. Lugosi, *Prediction, Learning, and Games*, Cambridge Univ. Press, 2006.
- [13] Y. Chen, J.W. Vaughan, A new understanding of prediction markets via no-regret learning, in: *Proc. 11th ACM Conference on Electronic Commerce*, 2010, pp. 189–198.
- [14] R.T. Clemen, R.L. Winkler, Aggregating probability distributions, in: W. Edwards, R. Miles, D. von Winterfeldt (Eds.), *Advances in Decision Analysis*, Cambridge Univ. Press, 2007, pp. 154–176.
- [15] P. DeMarzo, I. Kremer, Y. Mansour, Online trading algorithms and robust option pricing, in: *Proc. 38th Annual ACM Symposium on Theory of Computing*, 2006, pp. 477–486.
- [16] M. Faure, G. Roth, Stochastic approximations of set-valued dynamical systems: Convergence with positive probability to an attractor, *Math. Operations Res.* 35 (2010) 624–640.
- [17] D. Foster, R. Vohra, A randomization rule for selecting forecasts, *Operations Res.* 41 (1993) 704–709.
- [18] D. Foster, R. Vohra, Regret in the online decision problem, *Games Econ. Behav.* 29 (1999) 7–35.
- [19] Y. Freund, R. Schapire, Adaptive game playing using multiplicative weights, *Games Econ. Behav.* 29 (1999) 79–103.
- [20] D. Fudenberg, D. Levine, Consistency and cautious fictitious play, *J. Econ. Dynam. Control* 19 (1995) 1065–1089.
- [21] A. Gaunersdorfer, J. Hofbauer, Fictitious play, Shapley polygons, and the replicator equation, *Games Econ. Behav.* 11 (1995) 279–303.
- [22] I. Gilboa, A. Matsui, Social stability and equilibrium, *Econometrica* 59 (1991) 859–867.
- [23] J. Hannan, Approximation to Bayes risk in repeated play, in: M. Dresher, A.W. Tucker, P. Wolfe (Eds.), *Contributions to the Theory of Games*, vol. III, in: *Ann. of Math. Stud.*, vol. 39, Princeton Univ. Press, 1957, pp. 97–139.
- [24] S. Hart, A. Mas-Colell, A simple adaptive procedure leading to correlated equilibrium, *Econometrica* 68 (2000) 1127–1150.
- [25] S. Hart, A. Mas-Colell, A general class of adaptive procedures, *J. Econ. Theory* 98 (2001) 26–54.
- [26] S. Hart, A. Mas-Colell, Continuous-time regret-based dynamics, *Games Econ. Behav.* 45 (2003) 375–394.
- [27] J. Hofbauer, W.H. Sandholm, Stable games and their dynamics, *J. Econ. Theory* 144 (2009) 1665–1693.
- [28] J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge Univ. Press, 1998.
- [29] J. Hofbauer, S. Sorin, Best response dynamics for continuous zero-sum games, *Discrete Contin. Dyn. Syst. Ser. B* 6 (2006) 215–224.

- [30] J. Hofbauer, S. Sorin, Y. Viossat, Time average replicator and best reply dynamics, *Math. Operations Res.* 34 (2009) 263–269.
- [31] S. Irani, A. Karlin, On-line computation, in: D. Hochbaum (Ed.), *Approximation Algorithms for NP-Hard Problems*, PWS-Kent, Boston, 1996, pp. 521–564.
- [32] V. Krishna, T. Sjöström, On the convergence of fictitious play, *Math. Operations Res.* 23 (1998) 479–511.
- [33] R.P. Larrick, J.B. Soll, Intuitions about combining opinions: Misappreciation of the averaging principle, *Manag. Sci.* 52 (2006) 111–127.
- [34] E. Lehrer, A wide range no-regret theorem, *Games Econ. Behav.* 42 (2003) 101–115.
- [35] N. Littlestone, M. Warmuth, The weighted majority algorithm, *Info. Computat.* 108 (1994) 212–261.
- [36] Y. Mansour, Regret minimization and job scheduling, in: *Proc. 36th Conference on Current Trends in Theory and Practice of Computer Science*, Springer, 2010, pp. 71–76.
- [37] A. Matsui, Best response dynamics and socially stable strategies, *J. Econ. Theory* 57 (1992) 343–362.
- [38] D. Monderer, D. Samet, A. Sela, Belief affirming in learning processes, *J. Econ. Theory* 73 (1997) 438–452.
- [39] H. Moulin, J.P. Vial, Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon, *Int. J. Game Theory* 7 (1978) 201–221.
- [40] R. Selten, An axiomatic theory of a risk dominance measure for bipolar games with linear incentives, *Games Econ. Behav.* 8 (1995) 213–263.
- [41] L.S. Shapley, Some topics in two-person games, in: M. Dresher, L.S. Shapley, A.W. Tucker (Eds.), *Advances in Game Theory*, Princeton Univ. Press, 1964, pp. 1–28.
- [42] C. Sparrow, S. van Strien, C. Harris, Fictitious play in 3×3 games: The transition between periodic and chaotic behaviour, *Games Econ. Behav.* 63 (2008) 259–291.
- [43] A. Timmermann, Forecast combinations, in: G. Elliott, C.W. Granger, A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Elsevier, 2006.
- [44] J. Tirole, *The Theory of Industrial Organization*, MIT Press, 1988.
- [45] Y. Viossat, A. Zapechelnyuk, No-regret dynamics and fictitious play, Working paper, 2012, available at: <http://arxiv.org/abs/1207.0660>.
- [46] E. Yanovskaya, Equilibrium points in polymatrix games, *Lit. Mat. Sb.* 8 (1968) 381–384 (in Russian).
- [47] H.P. Young, The evolution of conventions, *Econometrica* 61 (1993) 57–84.
- [48] H.P. Young, *Strategic Learning and Its Limits*, Oxford Univ. Press, 2004.