

On Robustness Properties in Empirical Centroid Fictitious Play

BRIAN SWENSON^{†*}, SOUMMYA KAR[†] AND JOÃO XAVIER^{*}

Abstract

Empirical Centroid Fictitious Play (ECFP) is a generalization of the well-known Fictitious Play (FP) algorithm designed for implementation in large-scale games. In ECFP, the set of players is subdivided into equivalence classes with players in the same class possessing similar properties. Players choose a next-stage action by tracking and responding to aggregate statistics related to each equivalence class. This setup alleviates the difficult task of tracking and responding to the statistical behavior of every individual player, as is the case in traditional FP. Aside from ECFP, many useful modifications have been proposed to classical FP, e.g., rules allowing for network-based implementation, increased computational efficiency, and stronger forms of learning. Such modifications tend to be of great practical value; however, their effectiveness relies heavily on two fundamental properties of FP: robustness to alterations in the empirical distribution step size process, and robustness to best-response perturbations. The main contribution of the paper is to show that similar robustness properties also hold for the ECFP algorithm. This result serves as a first step in enabling practical modifications to ECFP, similar to those already developed for FP.

I. INTRODUCTION

The field of learning in games is concerned with the study of systems of interacting agents, and in particular, the question of how simple behavior rules applied at the level of individual agents can lead to desirable global behavior. Fictitious Play (FP) [1] is one of the best studied game-theoretic learning algorithms. While attractive for its intuitive simplicity and proven convergence results, certain practical issues make FP prohibitively difficult to implement in games with a large number of players [2]–[5].

Empirical Centroid FP (ECFP) [4], [5] is a recently proposed generalization of FP designed for implementation in large games. In ECFP, the set of players is subdivided into sets of “equivalence classes” of players sharing similar properties. In this formulation, players only track and respond to an aggregate statistic (the empirical centroid) for each class of players, rather than tracking and responding to statistical properties of every individual player, as in classical FP. ECFP has been shown to learn elements of the set of symmetric Nash equilibria for the class of multi-player games known as potential games.

The main focus of this paper will be to study ECFP and show that certain desirable properties possessed by classical FP also hold for the more general ECFP. In particular, the work [6] studied classical FP and proved that the fundamental learning properties of FP can be retained in the following scenarios:

- (i) The step size sequence of the empirical distribution process takes on a form other than $\{1/t\}_{t \geq 1}$.
- (ii) Players are permitted to make suboptimal choices when choosing a next-stage action so long as the degree of suboptimality decays asymptotically to zero with time.

We say a FP-type algorithm is *step-size robust* if it retains its fundamental learning properties in the first scenario, and we say an algorithm is *best-response robust* if it retains its fundamental learning properties in the second scenario.

The work was partially supported by the FCT project FCT [UID/EEA/50009/2013] through the Carnegie Mellon/Portugal Program managed by ICTI from FCT and by FCT Grant CMU-PT/SIA/0026/2009, and was partially supported by NSF grant ECCS-1306128.

[†]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA (brianswe@andrew.cmu.edu and soumyak@andrew.cmu.edu).

^{*}Institute for Systems and Robotics (ISR), Instituto Superior Tecnico (IST), Technical University of Lisbon, Portugal (jxavier@isr.ist.utl.pt).

The notion of step-size robustness generalizes the concept of the empirical distribution of classical FP. A player’s empirical distribution in classical FP is taken to be the time-averaged histogram of the player’s action history; implicitly, this has an incremental step size of $1/t$. Scenario (i) allows players to choose alternate step-size sequences. Of particular interest is that it allows for construction of an empirical distribution that places more emphasis on recent observations while discounting observations from the distant past.

The notion of best-response robustness generalizes FP by relaxing the traditional assumption that players are always perfect optimizers. In particular, in classical FP, it is assumed that players are capable of choosing their next-stage action as a (precise) best response to the empirical action history of opposing players. In practice, this is a stringent assumption, requiring that players have perfect knowledge of the empirical distribution of all opposing players at all times, and are capable of precisely solving a (non-trivial) optimization problem each iteration of the algorithm. By relaxing this implicit assumption slightly (as in scenario (ii)), one is able to consider many useful extensions of FP of both practical and theoretical value.

In [6], the best-response robustness of FP was used to show convergence to the set of Nash equilibria of stochastic FP with vanishing smoothing, and to prove convergence of an FP-inspired actor-critic learning algorithm. In [3], best-response robustness of FP was used to show convergence of sampled FP—a variant of FP in which computational complexity is mitigated by approximating the expected utility using a Monte-Carlo method—and used again in [7] to ensure convergence of an even more computationally efficient version of sampled FP. In [8], the best-response robustness of FP is used to construct a variant of FP achieving a strong form of learning in which the player’s period-by-period strategies are guaranteed to converge to equilibrium (rather than only convergence in terms of the empirical frequencies, as is typical in FP). The best-response robustness of FP is also useful in that it allows for practical network-based implementations of FP; e.g., [5].

The main contribution of this paper is to demonstrate that ECFP is both step-size robust and best-response robust; i.e., ECFP retains its fundamental learning properties under scenarios (i) and (ii) above. This result is a necessary first step in order to develop practical modifications for ECFP similar in spirit to those already developed for FP; e.g., improved computational efficiency, network-based implementation rules, and strongly convergent variants of the algorithm, as mentioned above.¹ We prove the result following a similar line of reasoning to [6], [9]; we first study a continuous-time version of ECFP, and then use results from the theory of stochastic approximations to prove our main result regarding convergence of discrete-time ECFP based on properties of the continuous-time counterpart.

The remainder of the paper is organized as follows. Section II sets up the notation to be used in the subsequent development and reviews the classical FP algorithm. Section III presents discrete-time ECFP and states the main result. Section IV reviews relevant results in differential inclusions and stochastic approximations to be used in the proof of the main result. Section V presents continuous-time ECFP. Section VI proves convergence of discrete-time ECFP using properties of continuous-time ECFP. Section VII provides concluding remarks.

II. PRELIMINARIES

A. Game Theoretic Preliminaries

A review of game-theoretic learning algorithms—including classical FP—can be found in [10], [11].

¹The results of this paper are directly applied in [8] to prove a strong learning result for a variant of ECFP. We also note that one possible network-based implementation of ECFP has been presented in [5]. This implementation—which considers a fixed communication graph topologies and synchronous communication rules—relies on a weak form of best-response robustness (see [5], A.3). In order to consider ECFP in more general distributed scenarios (e.g., random communication graph topology and asynchronous communication rules) it is necessary to have the full robustness property derived in this paper.

A normal form game is given by the triple $\Gamma = (N, (Y_i)_{i \in N}, (u_i(\cdot))_{i \in N})$, where $N = \{1, \dots, n\}$ represents the set of players, Y_i —a finite set of cardinality m_i —denotes the action space of player i and $u_i(\cdot) : \prod_{i=1}^n Y_i \rightarrow \mathbb{R}$ represents the utility function of player i .

Throughout this paper we assume:

A. 1. *All players use identical utility functions.*

Under this assumption we drop the subscript i and denote by $u(\cdot)$ the utility function used by all players. The set of mixed strategies for player i is given by $\Delta_i = \{p \in \mathbb{R}^{m_i} : \sum_{k=1}^{m_i} p(k) = 1, p(k) \geq 0 \forall k = 1, \dots, m_i\}$, the m_i -simplex. A mixed strategy $p_i \in \Delta_i$ may be thought of as a probability distribution from which player i samples to choose an action. The set of joint mixed strategies is given by $\Delta^n = \prod_{i=1}^n \Delta_i$. A joint mixed strategy is represented by the n -tuple (p_1, \dots, p_n) , where $p_i \in \Delta_i$ represents the marginal strategy of player i , and it is implicitly assumed that players' strategies are independent.

The mixed utility function is given by $U(\cdot) : \Delta^n \rightarrow \mathbb{R}$, where,

$$U(p_1, \dots, p_n) := \sum_{y \in Y} u(y) p_1(y_1) \dots p_n(y_n).$$

Note that $U(\cdot)$ may be interpreted as the expected value of $u(y)$ given that the players' mixed strategies are statistically independent. For convenience, the notation $U(p)$ will often be written as $U(p_i, p_{-i})$, where $p_i \in \Delta_i$ is the mixed strategy for player i , and p_{-i} indicates the joint mixed strategy for all players other than i .

For $\epsilon \geq 0$, $i \in N$ and $p_{-i} \in \Delta_{-i}$, define the ϵ -best response set for player i as

$$BR_i^\epsilon(p_{-i}) := \{p_i \in \Delta_i : U(p_i, p_{-i}) \geq \max_{\alpha_i \in \Delta_i} U(\alpha_i, p_{-i}) - \epsilon\}$$

and for $p \in \Delta$ define

$$BR^\epsilon(p) := (BR_1^\epsilon(p_{-1}), \dots, BR_n^\epsilon(p_{-n})).$$

The set of Nash equilibria is given by

$$NE := \{p \in \Delta^n : U(p_i, p_{-i}) \geq U(p'_i, p_{-i}), \forall p'_i \in \Delta_i, \forall i\}.$$

As a matter of convention, all equalities and inequalities involving random objects are to be interpreted almost surely (a.s.) with respect to the underlying probability measure, unless otherwise stated.

B. Repeated Play

The learning algorithms considered in this paper assume the following format of repeated play.

Let a normal form game Γ be fixed. Let players repeatedly face off in the game Γ , and for $t \in \{1, 2, \dots\}$, let $a_i(t) \in \Delta_i$ denote the action played by player i in round t .² Let the n -tuple $a(t) = (a_1(t), \dots, a_n(t))$ denote the joint action at time t .

Denote by $q_i(t) \in \Delta_i$, the *empirical distribution*³ of player i . The precise manner in which the empirical distribution is formed will depend on the algorithm at hand. In general, $q_i(t)$ is formed as a function of the action history $\{a_i(s)\}_{s=1}^t$ and

²An action is usually assumed to be pure strategy, or a vertex of the simplex Δ_i . In this work, an action is permitted to be an arbitrary mixed strategy (cf. [6], for the case of FP). Since the results hold for any actions of this form, they also hold for the typical case where actions are restricted to be pure strategies.

³The term *empirical distribution* is often used to refer explicitly to the time-averaged histogram of the action choices of some player i ; i.e., $q_i(t) = \frac{1}{t} \sum_{s=1}^t a_i(s)$. However, using a broader definition as considered here, allows for interesting algorithmic generalizations; e.g., learning processes that place greater emphasis on observations of more recent actions. See [6] for further discussion.

serves as a compact representation of the action history of player i up to and including the round t . The joint empirical distribution is given by $q(t) := (q_1(t), \dots, q_n(t))$.

C. Classical Fictitious Play

FP may be intuitively described as follows. Players repeatedly face off in a stage game Γ . In any given stage of the game, players choose a next-stage action by assuming (perhaps incorrectly) that opponents are using stationary and independent strategies. In particular, let the empirical distribution be given by the time-averaged histogram

$$q_i(t) := \frac{1}{t} \sum_{s=1}^t a_i(s); \quad (1)$$

in FP, players use the empirical distribution of each opponent's past play as a prediction of the opponent's behavior in the upcoming round and choose a next-round strategy that is optimal (i.e., a best response) given this prediction.

A sequence of actions $\{a(t)\}_{t \geq 1}$ such that⁴

$$a_i(t+1) \in BR_i(q_{-i}(t)), \quad \forall i,$$

for all $t \geq 1$, is referred to as a *fictitious play process*. It has been shown that FP achieves Nash equilibrium learning in the sense that $d(q(t), NE) \rightarrow 0$ as $t \rightarrow \infty$ for select classes of games including two-player zero-sum games [12], two-player two-move games [13], and multi-player potential games [14], [15].

D. Empirical Centroid FP Setup

A presentation of ECFP in its most elementary form (i.e., all players are grouped into a single equivalence class) is given in [5]; the elementary formulation is less notationally involved, and can serve as a useful means of conveying the basic ideas of the approach in a straightforward manner. In this paper we focus on the general formulation of the ECFP algorithm.

In ECFP, players are grouped into sets of equivalence classes, or “permutation invariant” classes. Such grouping allows players to analyze collective behavior by tracking only the statistics of each equivalence class, rather than tracking the statistics of every individual player.

Let $m \leq n$, denote the number of classes, let $I = \{1, \dots, m\}$ be an index set, and let $\mathcal{C} = \{C_1, \dots, C_m\}$ be a collection of subsets of N ; i.e. $C_k \subseteq N$, $\forall k \in I$. A collection \mathcal{C} is said to be a *permutation-invariant partition* of N if,

- (i) $C_k \cap C_\ell = \emptyset$, for $k, \ell \in I$, $k \neq \ell$,
- (ii) $\bigcup_{k \in I} C_k = N$,
- (iii) for $k \in I$, $i, j \in C_k$, $Y_i = Y_j$,
- (iv) for $k \in I$, $i, j \in C_k$, there holds for any strategy profile $y = (y_i, y_j, y_{-(i,j)}) \in Y$,

$$u(y_i, y_j, y_{-(i,j)}) = u([y_j]_i, [y_i]_j, y_{-(i,j)}),$$

where the notation $([y_i]_j, [y_j]_i, y_{-(i,j)})$ indicates a permutation of (only) the strategies of players i and j in the strategy profile $y = (y_i, y_j, y_{-(i,j)})$.

For a collection \mathcal{C} , define $\phi(\cdot) : N \rightarrow I$ to be the unique mapping such that $\phi(i) = k$ if and only if $i \in C_k$.

⁴In all learning algorithms discussed in this paper, the initial action $a_i(1)$ may be chosen arbitrarily for all i .

For $k \in I$, and $p \in \Delta^n$, and permutation-invariant partition \mathcal{C} , define

$$\bar{p}^k := |C_k|^{-1} \sum_{i \in C_k} p_i \quad (2)$$

to be the k -th centroid with respect to \mathcal{C} , where $|C_k|$ denotes the cardinality of the set C_k . Likewise for $p \in \Delta^n$ define

$$\bar{p} := (\bar{p}_1, \bar{p}_2, \dots, \bar{p}_n), \quad (3)$$

where $\bar{p}_i := \bar{p}^{\phi(i)}$, to be the centroid distribution with respect to \mathcal{C} .

Given a permutation-invariant partition \mathcal{C} , let the set of symmetric Nash equilibria (relative to \mathcal{C}) be given by,

$$SNE := \{p \in NE : p_i = p_j \forall i, j \in C_k, \forall k \in I\},$$

and let the set of mean-centric equilibria (relative to \mathcal{C}) be given by,

$$MCE := \{p \in \Delta^n : U(p_i, \bar{p}_{-i}) \geq U(p'_i, \bar{p}_{-i}), \forall p'_i \in \Delta_i, \forall i\}.$$

The set of MCE is neither a strict superset nor subset of the NE—rather, it is a set of natural equilibrium points tailored to the ECFP dynamics [16]. The set of SNE however, is contained in the set of MCE.

The sets of SNE and MCE relative to a partition \mathcal{C} can be shown to be non-empty under A.1 using fixed point arguments similar to [16], [17].

III. EMPIRICAL CENTROID FICTITIOUS PLAY

Let the game Γ be played repeatedly as in Section II-B. Let the empirical distribution for player i be formed recursively with $q_i(1) = a_i(1)$ and for $t \geq 1$,

$$q_i(t+1) = q_i(t) + \gamma_t (a_i(t+1) - q_i(t)), \quad (4)$$

where we assume:

A. 2. The sequence $\{\gamma_t\}_{t \geq 1}$ in (4) satisfies $\gamma_t \geq 0$, $\forall t$, $\sum_{t \geq 1} \gamma_t = \infty$, and $\lim_{t \rightarrow \infty} \gamma_t = 0$.

Let the joint empirical distribution be given by $q(t) := (q_1(t), \dots, q_n(t))$.

Typical FP-type learning algorithms⁵ consider the empirical distribution to be a time-averaged histogram that places equal weight on all rounds; this corresponds to a step size of form $\gamma_t = \frac{1}{t+1}$, $\forall t$ (e.g., (1)). If a FP-type algorithm retains its fundamental learning properties under the more general assumption A.2, then we say the algorithm is *step-size robust*.

In ECFP [5], players do not track the empirical distribution of each individual player. Instead, they track only the centroid $\bar{q}^k(t)$ for each $k \in I$ (see (2)). Intuitively speaking, in ECFP each player i assumes (perhaps incorrectly) that for each class $C_k \in \mathcal{C}$ the centroid $\bar{q}^k(t)$ accurately represents the mixed strategy for all players $j \in C_k$. Each player i chooses her next-stage action as a myopic best response given this assumption.

Formally, the joint action at time $(t+1)$ is chosen according to the rule⁶

$$a(t+1) \in BR^{\epsilon_t}(\bar{q}(t)), \quad (5)$$

⁵We use the term FP-type learning algorithm to refer to an algorithm in which players choose their next-stage action as a myopic best response to some forecast rule based on the current time-averaged empirical distribution of play; cf. the learning framework considered in [18].

⁶The action $a(1)$ may be chosen arbitrarily.

where $\bar{q}(t)$ is the centroid distribution associated with $q(t)$ (see (3)), and where it is assumed that,

A. 3. *The sequence $\{\epsilon_t\}_{t \geq 1}$ in (5) satisfies $\lim_{t \rightarrow \infty} \epsilon_t = 0$.*

Typical FP-type learning algorithms assume that players are always perfect optimizers; i.e., $\epsilon_t = 0$, $\forall t$. If a FP-type learning algorithm retains its fundamental learning properties under A.3, we say the algorithm is *best-response robust*. The work [6] first considered generalizations to FP of the forms indicated in A.2–A.3 and showed that classical FP is both step-size robust and best-response robust. In this work we show that ECFP is also robust in both these senses.

Combining (4) with (5), gives the following difference inclusion governing the behavior of $\{q(t)\}_{t \geq 1}$,

$$q(t+1) \in (1 - \gamma_t) q(t) + \gamma_t BR^{\epsilon_t}(\bar{q}(t)). \quad (6)$$

Likewise, Lemma 5 (see appendix) shows that the sequence of centroid distributions $\{\bar{q}(t)\}_{t \geq 1}$ follows the difference inclusion,

$$\bar{q}(t+1) \in (1 - \gamma_t) \bar{q}(t) + \gamma_t BR^{\epsilon_t}(\bar{q}(t)). \quad (7)$$

We refer to the sequence $\{q(t), \bar{q}(t)\}_{t \geq 1}$ as a discrete-time ECFP (DT-ECFP) process with respect to (Γ, \mathcal{C}) .

The following theorem is the main result of the paper—it states that, if Γ is an identical interests game, then under the relatively weak assumptions A.2–A.3, players engaged in ECFP asymptotically learn elements of sets of SNE and MCE. Learning of MCE occurs in the sense that $d(q(t), MCE) \rightarrow 0$ —this form of learning corresponds to the typical notion of setwise *convergence in empirical distribution* typical in classical FP (see Section II-C and [19], [11]). Learning of SNE occurs in the sense that $d(\bar{q}(t), SNE) \rightarrow 0$. This notion of learning, while similar in spirit to the typical notion of convergence in empirical distribution, differs in that it is the empirical centroid distribution (3) that is converging to the set of SNE, rather than the empirical distribution itself.

Theorem 1. *Assume A.1–A.3 hold. Let \mathcal{C} be a permutation-invariant partition of the player set N . Let $\{q(t), \bar{q}(t)\}_{t \geq 1}$ be an ECFP process with respect to (Γ, \mathcal{C}) . Then,*

- (i) *players learn a subset of the MCE in the sense that $\lim_{t \rightarrow \infty} d(q(t), MCE) = 0$,*
- (ii) *players learn a subset of the SNE in the sense that $\lim_{t \rightarrow \infty} d(\bar{q}(t), SNE) = 0$.*

We note that if $\epsilon_t = 0$ and $\gamma_t = \frac{1}{t+1}$, then convergence of ECFP in the sense of Theorem 1 was established in our prior work [5].

In order to prove Theorem 1 in its full generality we follow the approach of [6], [9]—we first study the set of continuous-time differential inclusions associated with ECFP, and then derive Theorem 1 from the continuous-time results via tools from the theory of stochastic approximations.

In particular, Section IV discusses the notion of a perturbed solution of a differential inclusion, introduces the notion of a chain transitive set, and presents key results that allow one to relate the limit sets of perturbed solutions to internally chain transitive sets of the associated differential inclusion. Section V then presents continuous-time ECFP (CT-ECFP) and shows convergence of CT-ECFP to the sets of SNE and MCE using Lyapunov arguments.

Section VI presents Lemmas 1 and 2 that relate the limit sets of DT-ECFP to the limit sets of CT-ECFP. Lemma 1 shows that the limit sets of DT-ECFP are contained in the internally chain transitive sets of the corresponding CT-ECFP process. This is accomplished by first showing that DT-ECFP processes may be considered to be perturbed solutions of the associated CT-ECFP differential inclusion, and then invoking Theorem 2 to clinch the result. Lemma 2 then shows that the internally

chain transitive sets of CT-ECFP are contained in the sets of MCE and SNE. This is accomplished by invoking Proposition 1 together with the Lyapunov arguments derived for CT-ECFP processes in Section V.

The proof of Theorem 1 then follows by combining Lemmas 1 and 2, as noted in Section VI-A.

IV. CHAIN TRANSIENT SETS

We study the limiting behavior of DT-ECFP by first studying the behavior of a continuous-time version of ECFP, and then relating the limit sets of DT-ECFP to the limit sets of its continuous-time counterpart. In particular, we will relate the limit sets of DT-ECFP to the chain transitive sets of CT-ECFP. Following the approach of [9], [6], let F denote a set-valued function mapping each point $\xi \in \mathbb{R}^m$ to a set $F(\xi) \in \mathbb{R}^m$. We assume:

A. 4. (i) F is a closed set-valued map.⁷

(ii) $F(\xi)$ is a nonempty compact convex subset of \mathbb{R}^m for all $\xi \in \mathbb{R}^m$.

(iii) For some norm $\|\cdot\|$ on \mathbb{R}^m , there exists $c > 0$ such that for all $\xi \in \mathbb{R}^m$, $\sup_{\eta \in F(\xi)} \|\eta\| \leq c(1 + \|\xi\|)$.

Definition 1. A solution for the differential inclusion $\frac{dx}{dt} \in F(x)$ with initial point $\xi \in \mathbb{R}^m$ is an absolutely continuous mapping $x : \mathbb{R} \rightarrow \mathbb{R}^m$ such that $x(0) = \xi$ and $\frac{dx(t)}{dt} \in F(x(t))$ for almost every $t \in \mathbb{R}$.

Definition 2. Let $\|\cdot\|$ be a norm on \mathbb{R}^m , and let $F : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a set valued function satisfying A.4. Consider the differential inclusion

$$\frac{dx}{dt} \in F(x). \quad (8)$$

(a) Given a set $X \subset \mathbb{R}^m$ and points ξ and η , we write $\xi \leftrightarrow \eta$ if for every $\epsilon > 0$ and $T > 0$ there exist an integer $n \geq 1$, solutions x_1, \dots, x_n to the differential inclusion (8), and real numbers t_1, \dots, t_n greater than T such that

(i) $x_i(s) \in X$, for all $0 \leq s \leq t_i$ and for all $i = 1, \dots, n$,

(ii) $\|x_i(t_i) - x_{i+1}(0)\| \leq \epsilon$ for all $i = 1, \dots, n-1$,

(iii) $\|x_1(0) - \xi\| \leq \epsilon$ and $\|x_n(t_n) - \eta\| \leq \epsilon$.

(b) X is said to be internally chain transient if X is compact and $\xi \leftrightarrow \xi$ for all $\xi \in X$.

The following theorem from [9] allows one to relate the set of limit points of certain discrete-time processes to the internally chain transient sets of their continuous-time counterparts.

Theorem 2. Assume $F : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a set valued function satisfying A.4. Let $\{x(t)\}_{t \geq 1}$ be a process satisfying

$$x(t+1) - x(t) - \alpha_{t+1}M(n+1) \in \alpha_{t+1}F(x(t)), \quad (9)$$

where $\{\alpha_t\}_{t \geq 1}$ is a sequence of non-negative numbers such that

$$\sum_{t \geq 1} \alpha_t = \infty \quad \text{and} \quad \lim_{t \rightarrow \infty} \alpha_t = 0,$$

⁷I.e., $\text{Graph}(F) := \{(\xi, \eta) : \eta \in F(\xi)\}$ is a closed subset of $\mathbb{R}^m \times \mathbb{R}^m$.

and $\{M(t)\}_{t \geq 1}$ is a sequence of deterministic or random perturbations. If

$$(a) \text{ for all } T > 0,$$

$$\lim_{t \rightarrow \infty} \sup_k \left(\left\| \sum_{i=t}^{k-1} \alpha_{i+1} M(i+1) \right\| : \sum_{i=t}^{k-1} \alpha_{i+1} \leq T \right) = 0,$$

$$(b) \sup_{t \geq 1} \|x(t)\| < \infty,$$

then the set of limit points of $\{x(t)\}_{t \geq 1}$ is a connected internally chain transitive set of the differential inclusion

$$\frac{d}{dt}x(t) \in F(x(t)).$$

In an abuse of terminology,⁸ we sometimes refer to a discrete-time process $\{x(t)\}_{t \geq 1}$ verifying the recursion (9) as a perturbed solution of (8).

The differential inclusion (8) induces a set-valued dynamical system $\{\Phi_t\}_{t \in \mathbb{R}}$ defined by

$$\Phi_t(x_0) := \{x(t) : x \text{ is a solution to (8) with } x(0) = x_0\}.$$

Let Λ be any subset of \mathbb{R}^m . A continuous function $V : \mathbb{R}^m \rightarrow \mathbb{R}$ is called a Lyapunov function for Λ if $V(y) < V(x_0)$ for all $x_0 \in \mathbb{R}^m \setminus \Lambda$, $y \in \Phi_t(x_0)$, $t > 0$, and $V(y) \leq V(x_0)$ for all $x_0 \in \Lambda$, $y \in \Phi_t(x_0)$ and $t \geq 0$. The following proposition ([9], Proposition 3.27) allows one to relate the chain transitive sets of a differential inclusion to Lyapunov attracting sets.

Proposition 1. *Suppose that V is a Lyapunov function for Λ . Assume that $V(\Lambda)$, the image of Λ under V , has empty interior. Then every internally chain transitive set L is contained in Λ and $V|_L$, the restriction of V to the set L , is constant.*

V. CONTINUOUS-TIME ECFP

In this section we consider a continuous-time version of ECFP. Let Γ satisfy A.1 and let \mathcal{C} be a permutation-invariant partition of N . In analogy to⁹ (6), for $t > 0$ let

$$\dot{q}^c(t) \in BR(\bar{q}^c(t)) - q^c(t), \tag{10}$$

where we use the superscript $q^c(t)$ to indicate a continuous-time analog of the empirical distribution, and where, for $p \in \Delta^n$, we let $BR(p) := BR^\epsilon(p)$ with $\epsilon = 0$, and $\bar{q}^c(t)$ is the centroid distribution associated with $q^c(t)$ (see (3)). We refer to the process $\{q^c(t), \bar{q}^c(t)\}_{t \geq 0}$ as a continuous-time ECFP (CT-ECFP) process relative to (Γ, \mathcal{C}) .

As our end goal involves studying the limiting behavior of $\{\bar{q}^c(t)\}_{t \geq 1}$, note that for $k \in I$, and $\bar{q}^{c,k}(t)$ defined similar to

⁸A perturbed solution of (8) typically refers to a continuous time process that is associated with (8) by means of an integrable perturbation. Under appropriate conditions, processes of the form (9) may be transformed via an interpolation procedure into a continuous-time process satisfying the typical definition of a perturbed solution. See [9], Section I for more details.

⁹Note that (6) may be written as $q(t+1) - q(t) \in \frac{1}{t+1} (BR^{\epsilon t}(\bar{q}(t)) - q(t))$.

(2), there holds

$$\begin{aligned}\dot{\bar{q}}^{c,k}(t) &= \frac{d}{dt}|C_k|^{-1} \sum_{j \in C_k} q_j^c(t) \\ &= |C_k|^{-1} \sum_{j \in C_k} \frac{d}{dt} q_j^c(t) \\ &= |C_k|^{-1} \sum_{j \in C_k} \dot{q}_j^c(t),\end{aligned}$$

Let $p(t) = \dot{q}^c(t) + q^c(t)$, so that $\bar{p}(t) = (\bar{p}_1(t), \dots, \bar{p}_n(t))$ with $\bar{p}_i(t) = \bar{p}^{\phi(i)}(t)$ for $i \in N$ (see (3)). By the above, and the linearity of differentiation, $\bar{p}(t) = \dot{\bar{q}}^c(t) + \bar{q}^c(t)$. Thus, by Lemma 4, (10) implies that $\bar{p}(t) \in BR(\bar{q}^c(t))$, or equivalently,

$$\dot{\bar{q}}^c(t) \in BR(\bar{q}^c(t)) - \bar{q}^c(t). \quad (11)$$

A. Convergence in Continuous Time

This section studies the convergence of continuous-time ECFP to the sets of SNE and MCE.

For any solution $q^c(t)$ of (10) and associated centroid process $\bar{q}^c(t)$, let $w(t) := U(\bar{q}^c(t))$ and let $v(t) := \frac{1}{n} \sum_{i=1}^n U(q_i^c(t), \bar{q}_{-i}^c(t))$.

There holds,

$$\begin{aligned}\dot{w}(t) &= \sum_{i=1}^n \frac{\partial}{\partial \bar{q}_i^c} U(\bar{q}^c(t)) \dot{\bar{q}}_i^c(t) \\ &\geq \sum_{i=1}^n \left[U(\dot{\bar{q}}_i^c(t) + \bar{q}_i^c(t), \bar{q}_{-i}^c(t)) - U(\bar{q}^c(t)) \right] \\ &= \sum_{i=1}^n \left[\max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - U(\bar{q}^c(t)) \right] \geq 0,\end{aligned} \quad (12)$$

where the second line follows from the concavity of U in p_i , and the third follows from (11).

By Lemma 3 there holds

$$\frac{1}{n} \sum_{i=1}^n U(q_i^c(t), \bar{q}_{-i}^c(t)) = U(\bar{q}^c(t)). \quad (13)$$

Hence $v(t) = w(t)$, there holds $\dot{v}(t) \geq 0$. Moreover, the following expansion is useful in order to study v as a Lyapunov function for the set of MCE:

$$\begin{aligned}\dot{v}(t) = \dot{w}(t) &\geq \sum_{i=1}^n \left[\max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - U(\bar{q}^c(t)) \right] \\ &= \sum_{i=1}^n \max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - nU(\bar{q}^c(t)) \\ &= \sum_{i=1}^n \max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - \sum_{i=1}^n U(q_i^c(t), \bar{q}_{-i}^c(t)) \\ &= \sum_{i=1}^n \left[\max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - U(q_i^c(t), \bar{q}_{-i}^c(t)) \right] \geq 0,\end{aligned} \quad (14)$$

where the inequality follows from (12), and the third line follows again from (13).

By (12), $w(t)$ is weakly increasing, and is constant in a time interval T if and only if $\max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - U(\bar{q}^c(t)) = 0$, $\forall i$; i.e., if and only if $\bar{q}^c(t) \in SNE$ for all $t \in T$.

By (14), $v(t)$ is weakly increasing, and $\dot{v}(t) = 0$ in some interval T implies $\max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}^c(t)) - U(q_i^c(t), \bar{q}_{-i}^c(t)) = 0$, $\forall i \in N, t \in T$; i.e., $q^c(t) \in MCE$ for all $t \in T$. Moreover, by Lemma 4, $q^c(t) \in MCE, \forall t \in T \implies \bar{q}^c(t) \in SNE \forall t \in T$,

which by the above comments implies $\dot{w}(t) = 0$ in T , or equivalently $\dot{v}(t) = 0$ in T . Thus, $v(t)$ is constant in a time interval T if and only if $q^c(t) \in MCE$ for all $t \in T$.

Proposition 2. *Assume A.1 holds. Then,*

- (i) *The limit set of every solution of (11) is a connected subset of SNE along which U is constant;*
- (ii) *For $p \in \Delta^n$, let $V(p) := \frac{1}{n} \sum_{i=1}^n U(p_i, \bar{q}_{-i}^c)$. The limit set of every solution of (10) is a connected subset of MCE along which V is constant.*

The proof of this proposition follows from the above comments.

VI. LIMIT SETS OF DISCRETE-TIME ECFP

In this section we study the limit sets of DT-ECFP by relating them to the internally chain transitive sets of CT-ECFP.

The following lemma relates the limit sets of DT-ECFP to the internally chain transient sets of the CT-ECFP differential inclusions (10) and (11).

Lemma 1. *Assume A.1–A.3 hold. Let $\{q(t), \bar{q}(t)\}_{t \geq 1}$ be a discrete-time ECFP process. Then,*

- (i) *The set of limit points of $\{\bar{q}(t)\}_{t \geq 1}$ is a connected internally chain transient set of (11),*
- (ii) *The set of limit points of $\{q(t)\}_{t \geq 1}$ is a connected internally chain transient set of (10).*

Proof: Proof of (ii): Observe that adding and subtracting the set $BR(\bar{q}(t))$ and rearranging terms in (6) gives,

$$\begin{aligned} q(t+1) - q(t) + \gamma_t \{BR(\bar{q}(t)) - BR^{\epsilon_t}(\bar{q}(t))\} \\ \in \gamma_t \{BR(\bar{q}(t)) - q(t)\}. \end{aligned} \quad (15)$$

Thus, the process $\{q(t)\}_{t \geq 1}$ fits the template of Theorem 2 with $x(t) := q(t)$, $F(x) := BR(\bar{x}) - x$, and $M(t) := BR(\bar{x}) - BR^{\epsilon_t}(\bar{x})$. It is straightforward to verify that F satisfies A.4.

It suffices to show that the process (15) satisfies conditions (a) and (b) of Theorem 2. Condition (b) is trivially satisfied since $q(t) \in \Delta^n$ for all t .

If k is such that $\sum_{i=n}^{k-1} \gamma_i \leq T$, then

$$\begin{aligned} \sup_k \left\| \sum_{i=n}^{k-1} \gamma_i \{BR^{\epsilon_i}(\bar{q}(t)) - BR(\bar{q}(t))\} \right\| \\ \leq T \sup_k \|BR^{\epsilon_i}(\bar{q}(t)) - BR(\bar{q}(t))\|. \end{aligned}$$

Since BR is upper semicontinuous, $BR^\epsilon(p) \rightarrow BR(p)$ uniformly as $\epsilon \rightarrow 0$. Thus condition (a) holds, and (ii) is proved.

Proof of (i): Adding and subtracting $BR(\bar{q}(t))$ and rearranging terms, the difference inclusion (7) may be written as

$$\begin{aligned} \bar{q}(t+1) - \bar{q}(t) + \gamma_t \{BR(\bar{q}(t)) - BR^{\epsilon_t}(\bar{q}(t))\} \\ \in \gamma_t \{BR(\bar{q}(t)) - \bar{q}(t)\}. \end{aligned} \quad (16)$$

Thus, the process $\{\bar{q}(t)\}$ fits the template of Theorem 2 with $x(t) := \bar{q}(t)$, $M(t) := BR(\bar{q}(t)) - BR^{\epsilon_t}(\bar{q}(t))$, and $F(x) := BR(x) - x$. It was shown in [9] that F satisfies A.4. It is sufficient to show that the process (16) satisfies conditions (a) and

(b) of Theorem 2. Condition (b) is trivially satisfied since $\bar{q}(t) \in \Delta^n$ for all t . Since $M(t)$ is defined the same manner as in case (i), the proof that condition (a) is satisfied follows the same reasoning as in the proof of assertion (ii). ■

The following lemma relates internally chain transitive sets of (10) to the set of MCE, and the internally chain transitive sets of (11) to the set of SNE; combined with Lemma 1 this will prove Theorem 1.

Lemma 2. *Let Γ be an identical interests game. Let \mathcal{C} be a permutation-invariant partition on Γ . Then,*

(i) *Every internally chain transitive set of (11) is contained in the set of SNE.*

(ii) *Every internally chain transitive set of (10) is contained in the set of MCE.*

Proof: Proof of (i): Let $W := -U$. By Section V-A (in particular, see (12)), W is a Lyapunov function for the set of SNE with $x(t) := \bar{q}(t)$. Note that W is multilinear and hence continuously differentiable.

For a differentiable function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, we say $x \in \mathbb{R}^m$ is a critical point of f if for $i = 1, \dots, m$, the partial derivative at x is zero, i.e., $\frac{\partial}{\partial x_j} f(x) = 0$. By Sard's Theorem ([20], p. 69), if CP is the critical points set of W , then $W(CP)$ contains no intervals. By definition, the set of NE is contained in the critical points set of U , and hence also contained in the critical points of W . Furthermore, by definition, $SNE \subset NE$, and hence the set SNE is contained in the critical points set of W . Thus, by Proposition 1, every internally chain transitive set of (11) is contained in the set SNE .

Proof of (ii): Note that, by Lemma 4, $p \in BR(\bar{p}) \implies \bar{p} \in BR(p)$. Thus, $p \in MCE$ implies that $\bar{p} \in SNE$. Let $V : \Delta^n \rightarrow \mathbb{R}$, with $V(p) := \frac{1}{n} \sum_{i=1}^n U(p_i, \bar{p}_{-i})$, and note that by Lemma 3, $V(p) = U(\bar{p})$. Invoking again Sard's Theorem, $U(NE)$ contains no intervals, and hence $U(SNE) \subset U(NE)$ contains no intervals. Since $U(SNE)$ contains no intervals, $V(MCE)$ also contains no intervals.

By Section V (in particular, see (14)) the function V is a Lyapunov function for the set of MCE with $x(t) := q(t)$. It follows from Proposition 1 that every chain transitive set of (10) is contained in MCE . ■

A. Proof of Theorem 1

Theorem 1 follows directly from Lemmas 1 and 2.

VII. CONCLUSIONS

Classical Fictitious Play (FP) is robust to alterations in the empirical distribution step-size process and robust to best-response perturbations. These robustness properties allow for interesting modifications to FP which can be of great practical value. Empirical Centroid Fictitious Play (ECFP) is a generalization of FP designed for large games. The paper showed that ECFP is also robust to step-size alterations and best-response perturbations. This result enables future research to consider practical modifications to ECFP, similar to those already developed for FP.

APPENDIX

Lemma 3. *Let \mathcal{C} be a partition of N , and for $p \in \Delta^n$ let \bar{p} be as defined in (3). Then $\frac{1}{n} \sum_{i=1}^n U(p_i, \bar{p}_{-i}) = U(\bar{p})$.*

Proof: Let I be an index set for \mathcal{C} and let m be the cardinality of I . For $k \in I$, and $j \in C_k$ note that

$$\begin{aligned} |C_k|U(\bar{p}) &= |C_k|U(\bar{p}_j, \bar{p}_{-j}) = |C_k|U\left(\left[|C_k|^{-1} \sum_{i \in C_k} p_i\right]_j, \bar{p}_{-j}\right) \\ &= \sum_{i \in C_k} U([p_i]_j, \bar{p}_{-j}) = \sum_{i \in C_k} U(p_i, \bar{p}_{-i}), \end{aligned}$$

where the second line follow from the definition of \bar{p}_j (see (3)), the third by multilinearity of U , and the fourth by permutation invariance of elements in C_k . Thus,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n U(p_i, \bar{p}_{-i}) &= \frac{1}{n} \sum_{k \in I} \sum_{i \in C_k} U(p_i, \bar{p}_{-i}) \\ &= \frac{1}{n} \sum_{k \in I} |C_k| U(\bar{p}) = \frac{1}{n} \sum_{i=1}^n U(\bar{p}) = U(\bar{p}). \end{aligned}$$

■

Lemma 4. *Let $q \in \Delta^n$, let \bar{q} be as defined in (3), and let $\epsilon \geq 0$. If $p \in BR^\epsilon(\bar{q})$, then $\bar{p} \in BR^\epsilon(\bar{q})$.*

Proof: Let $i \in N$. Recall that $\bar{p} := (p_1, \dots, p_n)$ with $p_i = p^{\phi(i)}$. There holds

$$\begin{aligned} U(\bar{p}_i, \bar{q}_{-i}) &= U\left(\left[|C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} p_j\right]_i, \bar{q}_{-i}\right) \\ &= |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} U([p_j]_i, \bar{q}_{-i}(t)) \\ &= |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} U([p_j]_i, [\bar{q}^{\phi(i)}]_j, \bar{q}_{-(i,j)}) \\ &= |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} U([p_j]_j, [\bar{q}^{\phi(i)}]_i, \bar{q}_{-(i,j)}) \\ &= |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} U([p_j]_j, \bar{q}_{-j}(t)) \\ &\geq |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} \max_{p'_j \in \Delta_j} (U(\alpha_j, \bar{q}_{-j}) - \epsilon) \\ &= |C_{\phi(i)}|^{-1} \sum_{j \in C_{\phi(i)}} \max_{p'_i \in \Delta_i} (U(\alpha_i, \bar{q}_{-i}) - \epsilon) \\ &= \max_{\alpha_i \in \Delta_i} U(\alpha_i, \bar{q}_{-i}) - \epsilon, \end{aligned}$$

where the first line follows by the definition of \bar{p}_i (see (3)), the second from the multilinearity of U , the fourth by permutation invariance of elements of $C_{\phi(i)}$, the sixth by the fact that, by hypothesis, $p_j \in BR_j^\epsilon(\bar{q}_{-j})$, and the seventh by permutation invariance of elements of $C_{\phi(i)}$. Since this holds for all $i \in N$, it follows that $\bar{p} \in BR^\epsilon(\bar{q})$. ■

Lemma 5. *Assume A.1 holds and suppose the action sequence $\{a(t)\}_{t \geq 1}$ is chosen according to (5). Then centroid process $\{\bar{q}(t)\}_{t \geq 1}$ follows the differential inclusion (7).*

Proof: Note that $\bar{q}(t+1)$ may be written recursively as

$$\bar{q}(t+1) = \bar{q}(t) + \frac{1}{t+1} (\bar{a}(t+1) - \bar{q}(t)).$$

By Lemma 4, $a(t+1) \in BR^{\epsilon t}(\bar{q}(t))$ implies $\bar{a}(t+1) \in BR^{\epsilon t}(\bar{q}(t))$. Substituting this into the above recursion and rearranging terms shows that $\{\bar{q}(t)\}$ follows the difference inclusion (7). ■

REFERENCES

- [1] G. W. Brown, "Iterative Solutions of Games by Fictitious Play" In *Activity Analysis of Production and Allocation*, T. Coopmans, Ed. New York: Wiley, 1951.

- [2] J. R. Marden, G. Arslan, and J. S. Shamma, "Joint strategy fictitious play with inertia for potential games," *IEEE Transactions on Automatic Control*, vol. 54, no. 2, pp. 208–220, 2009.
- [3] T. J. Lambert, M. A. Epelman, and R. L. Smith, "A fictitious play approach to large-scale optimization," *Operations Research*, vol. 53, no. 3, pp. 477–489, 2005.
- [4] B. Swenson, S. Kar, and J. Xavier, "Distributed learning in large-scale multi-agent games: A modified fictitious play approach," in *46th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, USA, 2012, pp. 1490 – 1495.
- [5] —, "Empirical centroid fictitious play: an approach for distributed learning in multi-agent games," 2012, submitted for journal publication, <http://arxiv.org/abs/1304.4577>.
- [6] D. S. Leslie and E. Collins, "Generalised weakened fictitious play," *Games and Economic Behavior*, vol. 56, no. 2, pp. 285–298, 2006.
- [7] B. Swenson, S. Kar, and J. Xavier, "A computationally efficient implementation of fictitious play for large-scale games," 2015, submitted for conference publication.
- [8] —, "From weak learning to strong learning in fictitious play-type algorithms," 2015, submitted for journal publication.
- [9] M. Benaïm, J. Hofbauer, and S. Sorin, "Stochastic approximations and differential inclusions," *SIAM Journal on Control and Optimization*, vol. 44, no. 1, pp. 328–348, 2005.
- [10] J. Marden and J. Shamma, "Game theory and distributed control," in *Handbook of Game Theory Vol. 4*, H. Young and S. Zamir, Eds. Elsevier Science, 2013, forthcoming.
- [11] H. P. Young, *Strategic learning and its limits*. Oxford University Press on Demand, 2004, vol. 2002.
- [12] J. Robinson, "An iterative method of solving a game," *The Annals of Mathematics*, vol. 54, no. 2, pp. 296–301, Sep. 1951.
- [13] K. Miyasawa, "On the convergence of the learning process in a 2 x 2 non-zero-sum two-person game," DTIC Document, Tech. Rep., 1961.
- [14] D. Monderer and L. S. Shapley, "Fictitious play property for games with identical interests," *Journal of Economic Theory*, vol. 68, no. 1, pp. 258–265, Jan. 1996.
- [15] —, "Potential Games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, May 1996.
- [16] B. Swenson, S. Kar, and J. Xavier, "Mean-centric equilibrium: an equilibrium concept for learning in large-scale games," in *Proceedings of IEEE Global Conference on Signal and Information Processing (GlobalSIP), 2013*, 2013, pp. 571–574.
- [17] S.-F. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman, "Notes on equilibria in symmetric games," in *Proceedings of Workshop on Game Theory and Decision Theory*, 2004.
- [18] J. S. Jordan, "Three problems in learning mixed-strategy Nash equilibria," *Games and Economic Behavior*, vol. 5, no. 3, pp. 368–386, 1993.
- [19] D. Fudenberg and D. K. Levine, *The theory of learning in games*. MIT press, 1998, vol. 2.
- [20] M. W. Hirsch, *Differential topology*. Springer, 1976.