

NO-REGRET LEARNING AND MIXED NASH EQUILIBRIA: THEY DO NOT MIX

LAMPROS FLOKAS*, EMMANOUIL V. VLATAKIS-GKARAGKOUNIS*, THANASIS LIANEAS†,
PANAYOTIS MERTIKOPOULOS**,‡, AND GEORGIOS PILIOURAS‡

ABSTRACT. Understanding the behavior of no-regret dynamics in general N -player games is a fundamental question in online learning and game theory. A folk result in the field states that, in finite games, the empirical frequency of play under no-regret learning converges to the game’s set of coarse correlated equilibria. By contrast, our understanding of how the day-to-day behavior of the dynamics correlates to the game’s Nash equilibria is much more limited, and only *partial* results are known for *certain* classes of games (such as zero-sum or congestion games). In this paper, we study the dynamics of *follow the regularized leader* (FTRL), arguably the most well-studied class of no-regret dynamics, and we establish a sweeping negative result showing that *the notion of mixed Nash equilibrium is antithetical to no-regret learning*. Specifically, we show that any Nash equilibrium which is not *strict* (in that every player has a unique best response) cannot be stable and attracting under the dynamics of FTRL. This result has significant implications for predicting the outcome of a learning process as it shows unequivocally that only strict (and hence, *pure*) Nash equilibria can emerge as stable limit points thereof.

1. INTRODUCTION

Regret minimization is one of the most fundamental requirements for online learning and decision-making in the presence of uncertainty and unpredictability [13]. Defined as the difference between the cumulative performance of an adaptive policy and that of the best fixed action in hindsight, the regret of an agent provides a concise and meaningful benchmark for quantifying the ability of an online algorithm to adapt to an otherwise unknown and unpredictable environment.

Arguably, the most widely studied class of no-regret algorithms is the general algorithmic scheme known as *follow the regularized leader* (FTRL) [60, 61]. This umbrella learning framework includes as special cases the multiplicative weights update (MWU) [3, 4, 36, 66] and online gradient descent (OGD) algorithms [68], both of which achieve a min-max optimal $\mathcal{O}(T^{1/2})$ regret guarantee. For obvious reasons, the ability of FTRL to adapt optimally to an unpredictable environment makes them ideal for applying them in multi-agent environments – i.e., games. In this case, if all agents adhere to a no-regret learning process based on FTRL (or one of its variants), as the sequence of play becomes more predictable, stronger regret guarantees are achievable, possibly down to constant regret, see

*DEPARTMENT OF COMPUTER SCIENCE, COLUMBIA UNIVERSITY, NEW YORK, NY10025.

†SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING, NATIONAL TECHNICAL UNIVERSITY OF ATHENS, ATHENS, GREECE.

**UNIV. GRENOBLE ALPES, CNRS, INRIA, LIG, 38000, GRENOBLE, FRANCE.

‡CRITEO AI LAB.

‡SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN, SINGAPORE.

E-mail addresses: lamflokas@cs.columbia.edu, emvlatakis@cs.columbia.edu,
lianeas@corelab.ntua.gr, panayotis.mertikopoulos@imag.fr, georgios@sutd.edu.sg.

2020 *Mathematics Subject Classification*. Primary 91A26, 37N40; Secondary 91A68, 68Q32, 68T05.

Key words and phrases. Regret; follow the regularized leader; game theory; stability of equilibria.

Acknowledgments of financial support are given in p. 11.

e.g., [6, 7, 23, 34, 41, 42, 54, 62] and references therein. As such, several crucial questions arise:

What are the *game-theoretic implications* of the no-regret guarantees of FTRL?
Do the dynamics of FTRL converge to an equilibrium of the underlying game?

A folk answer to this question is that “*no-regret learning converges to equilibrium in all games*” [47], suggesting in this way that no-regret dynamics inherently gravitate towards game-theoretically meaningful states. However, at this level of abstraction, both the *type of convergence* as well as the specific *notion of equilibrium* that go in this statement are not as strong as one would have hoped for. Formally, the only precise conclusion that can be drawn is as follows: *under a no-regret learning procedure, the empirical frequency of play converges to the game’s set of coarse correlated equilibria* [26, 27].

This leads to an important disconnect with standard game-theoretic solution concepts on several grounds. First, even in 2-player games, coarse correlated equilibria may be exclusively supported on *strictly* dominated strategies [64], so they fail even the most basic requirements of rationalizability [22, 25]. Second, the archetypal game-theoretic solution concept is that of *Nash equilibrium* (NE), and convergence to a Nash equilibrium is a much more tenuous affair: since no-regret dynamics are, by construction, uncoupled (in the sense that a player’s update rule does not *explicitly* depend on the payoffs of other players), the impossibility result of Hart & Mas-Colell [28] precludes the convergence of no-regret learning to Nash equilibrium in *all* games. This is consistent with the numerous negative complexity results for finding a Nash equilibrium [21, 58]: an incremental method like FTRL simply cannot have enough power to overcome PPAD completeness and converge to Nash equilibrium given adversarially chosen initial conditions.

In view of the above, a natural test of whether the dynamics of FTRL favor convergence to a Nash equilibrium is to see whether they eventually stabilize and converge to it when initialized nearby. In more precise language, *are Nash equilibria asymptotically stable in the dynamics of FTRL?* And, perhaps more importantly, *are all Nash equilibria created equal in this regard?*

Our contributions. We establish a stark and robust dichotomy between how the dynamics of FTRL treat Nash equilibria in *mixed* (i.e., randomized) vs. *pure* strategies. For the case of mixed Nash equilibria we establish a sweeping negative result to the effect that *the notion of mixed Nash equilibrium is antithetical to no-regret learning*. More precisely, we show that any Nash equilibrium which is not *strict* (in the sense that every player has a unique best response) cannot be stable and attracting under the dynamics of FTRL. Schematically:

Informal Theorem: Asymptotically stable point for FTRL \implies Pure Nash equilibrium
Equivalently: Mixed Nash equilibrium \implies Not asymptotically stable under FTRL

The linchpin of our analysis is the following striking property of the FTRL dynamics: when viewed in the space of “payoffs” (their natural state space), *they preserve volume irrespective of the underlying game*. More precisely, the Lebesgue measure of any open set of initial conditions in the space of payoffs remains invariant as it is carried along the flow of the FTRL dynamics (cf. Fig. 2). Importantly, this result is *not* true in the problem’s “primal” space, i.e., the space of the player’s mixed strategies: here, sets of initial conditions can expand or contract indefinitely under the standard Euclidean volume form.

This duality between payoffs and strategies is the leitmotif of our approach and has a number of important consequences. First, exploiting the volume-preservation property of FTRL, we show that no interior Nash equilibrium (and, furthermore, no closed set in the interior of the strategy space) can be asymptotically stable under the dynamics of FTRL, as this effectively would necessitate volume contraction in the interior of the space (Theorem 1).

To move beyond this result and disqualify *all* non-strict Nash equilibria (not just interior ones) more intricate arguments are required. In this case, a fundamental distinction arises between classes of dynamics that may attain the boundary of the players’ strategy space in finite time versus those that do not. The first case concerns FTRL dynamics with an everywhere-differentiable regularizer, like the Euclidean regularizer that gives rise to OGD and the associated projection dynamics. The second concerns dynamics where the regularizer becomes *steep* at the boundary of the strategy simplex, e.g., like the Shannon-Gibbs entropy that gives rise to the multiplicative weights update (MWU) algorithm and the replicator dynamics. While the interior of the strategy simplex is invariant for the second class of dynamics, this is not the case for the former: in Euclidean-like cases, the support of the mixed strategy of an agent may change over time. This leads to an essential dichotomy in the boundary behavior of different classes of FTRL dynamics. Nonetheless, despite the qualitatively distinct long-run behavior of the dynamics, a unified message emerges: *under the dynamics of FTRL, only strict Nash equilibria survive* (Theorem 2).

Finally, for the case of steep, entropy-like regularizers we prove that not only their asymptotically stable points but much more generally *any* asymptotically stable set must contain at least one pure strategy profile (Theorem 3).

Related work. The regret properties of FTRL have given rise to a vast corpus of literature which we cannot hope to review here; for an appetizer, we refer the reader to [12, 60] and references therein. On the other hand, the long-run behavior of FTRL in games (even finite ones) is nowhere near as well understood. A notable exception to this is the case of the replicator dynamics which have been studied extensively due to their origins and connection with evolutionary game theory, cf. [29, 59, 63, 67] for a review. For the replicator dynamics, a special instance of the volume preservation principle was first discovered by Akin [1] and ultimately gave rise to the so-called “folk theorem” of evolutionary game theory:¹ in population games, the notions of strict Nash equilibrium and asymptotic stability coincide [30]. This instability of mixed Nash equilibria plays a major role in the theory of population games as it shows that even the weakest form of mixing cannot be stable in an evolutionary sense. The volume preservation result that we establish here can be seen as a much more general “learning analogue” of this biological principle and provides an important link between population dynamics and the theory of online learning in games.

Recent work has examined the non-convergence of FTRL dynamics in more specialized settings. Coucheney et al. [20] established a version of the folk theorem of evolutionary game theory for a subclass of “decomposable”, steep FTRL dynamics. By contrast, Mertikopoulos et al. [42] focused on two-player *zero-sum* games (and networked versions thereof), and showed that almost all trajectories of FTRL orbit interior equilibria at a fixed distance without ever converging to equilibrium, generalizing the previous analysis for replicator dynamics by Piliouras & Shamma [53]. This is an interior equilibrium avoidance result, but one that *uniquely concerns zero-sum games*. Although the above results apply for continuous-time dynamics, in discrete-time non-convergence results only become stronger. Bailey & Piliouras [5] proved that discrete-time FTRL diverges away from the Nash equilibrium in zero-sum games, whereas Cheung & Piliouras [15] established Lyapunov chaos (volume-expansion, butterfly effects). Understanding the detailed geometry of non-equilibrating FTRL dynamics, e.g., periodicity/chaos, is an interesting direction where volume analysis has found application [7, 10, 16, 43, 44, 52]. Non-convergence, recurrence results have recently been established for FTRL dynamics via volume analysis even outside normal form games, e.g., in non-convex non-concave min-max differential games [65] and imperfect information zero-sum games [51]. Finally, such instability, non-convergence

¹Interestingly, Akin’s result was established under a special *non-Euclidean* volume form on the game’s strategy space, a fact which made any attempts at generalization particularly elusive.

results have inspired new, dynamics-based, solution concepts for games that generalize strict Nash while allowing cyclic, recurrent behavior [31, 48–50, 57].

In the converse direction, a complementary research thread has shown strict Nash equilibria are asymptotically stable under several incarnations of the FTRL dynamics [11, 18, 20, 38–41]. Our paper establishes the *converse* to this stability result, thus leading to the the following overarching principle (which covers all generic N -player games):

$$\textit{Asymptotic stability under FTRL} \iff \textit{Strict Nash equilibrium}$$

This result has significant implications for predicting the outcome of a learning process as it shows unequivocally that its pointwise stable outcomes are *precisely* the strict (and hence, *pure*) Nash equilibria of the underlying game.

2. PRELIMINARIES

Notation. If f is a function of a single variable, we will abuse notation slightly and extend it to vector variables $x \in \mathbb{R}^n$ by letting $f(x) \leftarrow (f(x_1), \dots, f(x_n))$. We will also understand inequalities involving vectors component-wise, i.e., $(x_1, \dots, x_n) > 0$ means that $x_i > 0$ for all $i = 1, \dots, n$.

The game. Throughout the sequel, we will focus on finite games. Formally, a *finite game in normal form* is defined as a tuple $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ consisting of (i) a finite set of *players* $i \in \mathcal{N} = \{1, \dots, N\}$; (ii) a finite set of *actions* (or *pure strategies*) $\mathcal{A}_i = \{\alpha_1, \dots, \alpha_{n_i}\}$ per player $i \in \mathcal{N}$; and (iii) each player's payoff function $u_i: \mathcal{A} \rightarrow \mathbb{R}$, where $\mathcal{A} := \prod_i \mathcal{A}_i$ denotes the ensemble of all possible *action profiles* $\alpha = (\alpha_1, \dots, \alpha_N)$. In this general context, players can also play *mixed strategies*, i.e., probability distributions $x_i = (x_{i\alpha_i})_{\alpha_i \in \mathcal{A}_i} \in \Delta(\mathcal{A}_i)$ over their pure strategies $\alpha_i \in \mathcal{A}_i$. Collectively, we will write $\mathcal{X}_i := \Delta(\mathcal{A}_i)$ for the mixed strategy space of player i and $\mathcal{X} := \prod_i \mathcal{X}_i$ for the space of all mixed strategy profiles $x = (x_1, \dots, x_N)$.

Given a mixed profile $x \in \mathcal{X}$, the corresponding expected payoff of player i will be

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

To keep track of the payoffs of each individual action, we will also write

$$v_{i\alpha_i}(x) := u_i(\alpha_i; x_{-i}) \quad (2)$$

for the payoff of the pure strategy $\alpha_i \in \mathcal{A}_i$ in the mixed profile $x = (x_i; x_{-i}) \in \mathcal{X}$.² Hence, writing $v_i(x) := (v_{i\alpha_i}(x))_{\alpha_i \in \mathcal{A}_i} \in \mathbb{R}^{\mathcal{A}_i}$ for the *payoff vector* of player i , we get the compact expression

$$u_i(x) = \langle v_i(x), x_i \rangle = \sum_{\alpha_i \in \mathcal{A}_i} x_{i\alpha_i} v_{i\alpha_i}(x) \quad (3)$$

where, in standard notation, $\langle v, x \rangle = v^\top x$ denotes the ordinary pairing between v and x .

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

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

Writing $\text{supp}(x_i^*) = \{\alpha_i \in \mathcal{A}_i : x_{i\alpha_i}^* > 0\}$ for the support of x_i^* , Nash equilibria can be equivalently characterized via the variational inequality

$$v_{i\alpha_i^*}(x^*) \geq v_{i\alpha_i}(x^*) \quad \text{for all } \alpha_i^* \in \text{supp}(x_i^*) \text{ and all } \alpha_i \in \mathcal{A}_i, i \in \mathcal{N}. \quad (4)$$

In turn, this characterization leads to the following taxonomy:

- (1) x^* is called *pure* if $\text{supp}(x^*) = \prod_i \text{supp}(x_i^*)$ is a singleton.

²We are using here the standard game-theoretic shorthand $(x_i; x_{-i}) := (x_1, \dots, x_i, \dots, x_N)$ to highlight the strategic choice of a given player $i \in \mathcal{N}$ versus that of the player's opponents $\mathcal{N}_{-i} := \mathcal{N} \setminus \{i\}$.

- (2) If x^* is not pure, we say that it is *mixed*; and if $\text{supp}(x^*) = \mathcal{A}$, we say that it is *fully mixed*.

By definition, pure Nash equilibria are themselves pure strategies and correspond to vertices of \mathcal{X} ; at the other end of the spectrum, fully mixed equilibria belong to the relative interior $\text{ri}(\mathcal{X})$ of \mathcal{X} , so they are often referred to as *interior* equilibria.

Another key distinction between Nash equilibria concerns the defining inequality (NE): if this inequality is strict for all $x_i \neq x_i^*$, $i \in \mathcal{N}$, x^* is called itself *strict*. Strict Nash equilibria are pure a fortiori, and they play a key role in game theory because any unilateral deviation incurs a strict loss to the deviating player; put differently, if x^* is strict, *every player has a unique best response*. Taking this idea further, x^* is called *quasi-strict* if (4) is strict for all $\alpha_i \in \mathcal{A}_i \setminus \text{supp}(x_i^*)$, i.e., if all best responses of player i are contained in $\text{supp}(x_i^*)$. By a deep result of Ritzberger [55], all Nash equilibria are quasi-strict in almost all games;³ in view of this, we will tacitly assume in the sequel that all equilibria considered are quasi-strict, a property known as “genericity” [17, 25, 35].

Remark. We should stress here that quasi-strict equilibria *need not be pure*: they could be partially or even fully mixed, e.g., as in the case of Stag Hunt, Rock-Paper-Scissors, Matching Pennies, the Battle of the Sexes, etc. We provide a series of illustrative examples in the supplement.

Regret. A key requirement in online learning is the minimization of the players’ *regret*, i.e., the cumulative payoff difference between a player’s mixed strategy at a given time and the player’s best possible strategy in hindsight. In more detail, assuming that play evolves in continuous time $t \geq 0$, the (external) regret of a player $i \in \mathcal{N}$ relative to a sequence of play $x(t) \in \mathcal{X}$ is defined as

$$\text{Reg}_i(T) = \max_{p_i \in \mathcal{X}_i} \int_0^T [u_i(p_i; x_{-i}(t)) - u_i(x(t))] dt, \quad (5)$$

and we say that player i has *no regret* under $x(t)$ if $\text{Reg}_i(T) = o(T)$.

No-regret learning via regularization. The most widely used method to achieve no-regret is the class of policies known as *follow the regularized leader* (FTRL) [60, 61]. Heuristically, at each $t \geq 0$, FTRL prescribes a mixed strategy that maximizes the players’ cumulative payoff up to time t minus a regularization penalty which incentivizes exploration. Formally, this is represented by the dynamics

$$\begin{aligned} y_{i\alpha_i}(t) &= y_{i\alpha_i}(0) + \int_0^t v_{i\alpha_i}(x(s)) ds && \{\text{aggregate payoffs}\} \\ x_{i\alpha_i}(t) &= Q_{i\alpha_i}(y_i(t)) && \{\text{choice of strategy}\} \end{aligned}$$

or, in more compact notation:

$$\dot{y}(t) = v(Q(y(t))). \quad (\text{FTRL})$$

In the above, each $y_{i\alpha_i}$ plays the role of an auxiliary “score variable” which measures the aggregate performance of the pure strategy $\alpha_i \in \mathcal{A}_i$ over time. These scores are subsequently transformed to mixed strategies by means of a player-specific *choice map* $y_i \mapsto x_i = Q_i(y_i)$ which is defined as

$$Q_i(y_i) = \arg \max_{x_i \in \mathcal{X}_i} \{ \langle y_i, x_i \rangle - h_i(x_i) \} \quad \text{for all } y_i \in \mathcal{Y}_i := \mathbb{R}^{n_i}. \quad (6)$$

In other words, $Q_i: \mathcal{Y}_i \rightarrow \mathcal{X}_i$ essentially acts as a “soft” version of the best-response correspondence $y_i \mapsto \arg \max_{x_i \in \mathcal{X}_i} \langle y_i, x_i \rangle$, suitably regularized by a convex penalty term $h_i(x_i)$. The precise assumptions regarding the *regularizer function* $h_i: \mathcal{X}_i \rightarrow \mathbb{R}$ will be

³Specifically, on a set which is open and dense (and hence of full measure) in the space of all games.

discussed in detail later; for now, we provide two prototypical examples of (FTRL) that will play a major role in the sequel:

Example 2.1 (Entropic regularization and exponential weights). One of the most widely used regularizers in online learning is the (negative) Gibbs-Shannon entropy $h_i(x_i) = \sum_{\alpha_i} x_{i\alpha_i} \log x_{i\alpha_i}$. A standard calculation then yields the so-called *logit choice map*, written in vectorized form as $\Lambda_i(y_i) = \exp(y_i) / \sum_{\alpha_i \in \mathcal{A}_i} \exp(y_{i\alpha_i})$. In turn, this leads to the *exponential weights dynamics*:

$$\begin{aligned} \dot{y}_i(t) &= v_i(x(t)), \\ x_i(t) &= \Lambda_i(y_i(t)). \end{aligned} \tag{EW}$$

The system (EW) describes the mean dynamics of the so-called *multiplicative weights update* (MWU) algorithm (or ‘‘Hedge’’); for an (incomplete) account of its long history, see [3, 4, 13, 24, 32, 36, 37, 66] and references therein.

Example 2.2 (L^2 regularization). Another popular choice of regularizer is the quadratic penalty $h_i(x_i) = (1/2)\|x_i\|^2$. In this case, the associated choice map is the Euclidean projector on the simplex, $\Pi_i(y_i) = \arg \min_{x_i \in \mathcal{X}_i} \|y_i - x_i\|$, which gives rise to the *Euclidean regularization dynamics*

$$\begin{aligned} \dot{y}_i(t) &= v_i(x(t)), \\ x_i(t) &= \Pi_i(y_i(t)). \end{aligned} \tag{ERD}$$

Beyond the two prototypical examples discussed above, the origin of the dynamics (FTRL) can be traced to Shalev-Shwartz & Singer [61], Nesterov [46], and, via their link to online mirror descent (OMD), all the way back to Nemirovski & Yudin [45]. Describing the history and literature surrounding these dynamics would take us too far afield, so we do not attempt it.

3. THE FUNDAMENTAL DICHOTOMY OF FTRL DYNAMICS

To connect the long-run behavior of (FTRL) to the Nash equilibria of the underlying game, we must first understand how the players’ mixed strategies evolve under (FTRL). Our goal in this section is to provide some background to this question as a precursor to our analysis in Section 4. To lighten notation, we will drop in what follows the player index i , writing for example x_α instead of the more cumbersome $x_{i\alpha_i}$; we will only reinstate the index i if absolutely necessary to avoid confusion.

3.1. Scores vs. strategies. To begin, we note that (FTRL) exhibits a unique duality: on the one hand, the variables of interest are the players’ mixed strategies $x(t) \in \mathcal{X}$; on the other, the dynamics (FTRL) evolve in the space \mathcal{Y} of the players’ score variables $y(t)$. Mixed strategies are determined by the corresponding scores via the players’ choice maps $y \mapsto x = Q(y)$, but this is not a two-way street: as we explain below, the map $Q: \mathcal{Y} \rightarrow \mathcal{X}$ is *not invertible*, so obtaining an autonomous dynamical system on the strategy space \mathcal{X} is a delicate affair. In the general case, invoking standard arguments from convex analysis [8, 56] we have $y(t) \in \nabla h(x(t)) + \text{PC}(x(t))$, where

$$\text{PC}(x) = \{y \in \mathcal{Y} : y_\alpha \geq y_\beta \text{ for all } \alpha, \beta \in \text{supp}(x)\} \tag{7}$$

denotes the polar cone to \mathcal{X} at x .⁴

In the entropic case of Example 2.1, the logit choice map $Q = \Lambda$ only returns *fully mixed* strategies since $\exp(y) > 0$. In the relative interior $\text{ri}(\mathcal{X})$ of \mathcal{X} , we have by Equation (7)

⁴In particular, for all $y \in \text{PC}(x)$, we have $y_\alpha = y_\beta$ whenever $\alpha, \beta \in \text{supp}(x)$. The similarity of this condition to the characterization (4) of Nash equilibria is not a coincidence: x^* is a Nash equilibrium of Γ if and only if $v(x^*) \in \text{PC}(x^*)$ [17, 39].




Figure 1: The inverse images of neighborhoods of different points in \mathcal{X} under the Euclidean choice map $Q = \Pi$.

that $\text{PC}(x) = \{(t, \dots, t) : t \in \mathbb{R}\}$. As a result, Λ is not surjective; however, up to a multiple of $(1, \dots, 1)$, it is *injective*. On the other hand, in the Euclidean framework of [Example 2.2](#), the choice map $Q = \Pi$ can also return non-fully mixed strategies. Both [Equation \(7\)](#) and [Fig. 1](#) show that on the boundary $\text{PC}(x)$ is strictly larger compared to the interior. Thus Π is surjective but not injective, even modulo a subspace of \mathcal{Y} .

The key obstacle to mapping the dynamics ([FTRL](#)) to \mathcal{X} is the lack of injectivity of Q . In turn, this allows us to make two key observations: (i) there is an important split in behavior between boundary and interior states; and (ii) this split is linked to whether the underlying choice map is surjective or not. We elaborate on this below.

3.2. The steep/non-steep dichotomy. The lack of injectivity of Λ on $\text{ri}(\mathcal{X})$ is a technical artifact of the sum-to-one constraints of the strategy probabilities: knowing all but one of the strategy probabilities we can easily recover the remaining one. Thus the \mathcal{Y} space, having the same number of coordinates as the \mathcal{X} space, also contains redundant information. With an appropriate projection we can remove this redundancy and restore injectivity in the interior, deriving the dynamics of $x(t)$ on \mathcal{X} . Making this argument precise for the entropic case of [Example 2.1](#), we obtain the *replicator dynamics*:

$$\dot{x}_\alpha = x_\alpha [v_\alpha(x) - u(x)]. \quad (\text{RD})$$

On the other hand, this is not enough for the Euclidean framework of [Example 2.2](#). When trajectories approach $\text{bd}(\mathcal{X})$, the positivity constraints $x_i \geq 0$ kick in finite time. Unlike the sum-to-one constraints of the previous case, these cannot be resolved with a dimensionality reduction so we cannot obtain a well-posed dynamical system on \mathcal{X} as above. This problem can only be temporarily avoided for time intervals where $\text{supp}(x(t))$ remains constant. For these intervals $x(t)$ can be shown to satisfy the *projection dynamics* [\[38\]](#)

$$\dot{x}_\alpha = v_\alpha(x) - |\text{supp}(x)|^{-1} \sum_{\beta \in \text{supp}(x)} v_\beta(x) \quad \text{if } \alpha \in \text{supp}(x). \quad (\text{PD})$$

In contrast to the replicator dynamics, different trajectories of [\(PD\)](#) can merge or split any number of times, and they may transit from one face of \mathcal{X} to another in finite time [\[38, 39\]](#).

The two cases above are not just conveniently chosen examples, but archetypes of the fundamentally different behaviors that can be observed under ([FTRL](#)) for different regularizers. As we discuss in the supplement, this polar split is intimately tied to the behavior of the derivatives of h at the boundary of \mathcal{X} . To formalize this, we say that h is *steep* if $\|\nabla h(x)\| \rightarrow \infty$ whenever $x \rightarrow \text{bd}(\mathcal{X})$; by contrast, if $\sup_{x \in \mathcal{X}} \|\nabla h(x)\| < \infty$, we say that h is *non-steep*. Thus, in terms of our examples, the negentropy function of [Example 2.1](#) is the archetype for steep regularizers, while the L^2 penalty of [Example 2.2](#) is the non-steep one. The split between steep and non-steep dynamics may then be stated as follows:

- (1) If h is steep, the mixed-strategy trajectories $x(t) = Q(y(t))$ carry all the information required to predict the evolution of the system; in particular, $x(0)$ fully determines $x(t)$ for all $t \geq 0$, and $x(0)$ remains fully mixed for all time.

- (2) If h is non-steep, the trajectories $x(t) = Q(y(t))$ do not fully capture the state of the system: $x(0)$ *does not* determine $x(t)$ for all $t \geq 0$, and even the times when $x(t)$ changes support cannot be anticipated by knowing $x(0)$ alone. For concision, we defer the precise statement and proof of this dichotomy to the paper’s supplement.

4. CONVERGENCE ANALYSIS AND RESULTS

We now turn to the equilibrium convergence properties of (FTRL). The central question that we seek to address here is the following: *Which Nash equilibria can be stable and attracting under (FTRL)? Are all equilibria created equal in that regard?*

4.1. Notions of stability. At a high level, a point is (a) *stable* when every trajectory that starts nearby remains nearby; and (b) *attracting* when it attracts all trajectories that start close enough. Already, this heuristic shows that defining these notions for (FTRL) is not straightforward: the target points are strategy profiles in \mathcal{X} , while the dynamics (FTRL) evolve in the dual space \mathcal{Y} . When h is *steep*, we can define an equivalent presentation of (FTRL) on \mathcal{X} , so this problem can be circumvented by working solely with mixed strategies; however, when h is *non-steep*, this is no longer possible and we need to navigate carefully between \mathcal{X} and \mathcal{Y} . In view of this, we have the following definitions:

- $x^* \in \mathcal{X}$ is *stable* if, for every neighborhood U of x^* in \mathcal{X} , there exists a neighborhood U' of x^* such that $x(t) = Q(y(t)) \in U$ for all $t \geq 0$ whenever $x(0) = Q(y(0)) \in U'$.
- $x^* \in \mathcal{X}$ is *attracting* if there exists a neighborhood U of x^* in \mathcal{X} such that $x(t) = Q(y(t)) \rightarrow x^*$ whenever $x(0) = Q(y(0)) \in U$.
- $x^* \in \mathcal{X}$ is *asymptotically stable* if it is both stable and attracting.

For obvious reasons, asymptotic stability is the “gold standard” for questions pertaining to equilibrium convergence and it will be our litmus test for the appropriateness of an equilibrium $x^* \in \mathcal{X}$ as an outcome of play. Specifically, if a Nash equilibrium is not asymptotically stable under (FTRL), it is not reasonable to expect a no-regret learner to converge to it, meaning in turn that it cannot be justified as an end-state of the players’ learning process. We expound on this below.

4.2. Volume preservation. A key observation regarding asymptotic stability is that neighborhoods of initial conditions near an asymptotically stable point should “contract” over time, eventually shrinking down to the point in question. Our first result below provides an apparent contradiction to this principle: it shows that volume is preserved under (FTRL), *irrespective of the underlying game*.

Proposition 1. *Let $\mathcal{R}_0 \subseteq \mathcal{Y}$ be a set of initial conditions for (FTRL) and let $\mathcal{R}_t = \{y(t) : y(0) \in \mathcal{R}_0\}$ denote its evolution under (FTRL) after time $t \geq 0$. Then, $\text{vol}(\mathcal{R}_t) = \text{vol}(\mathcal{R}_0)$.*

Proposition 1 (which we prove in the supplement through an application of Liouville’s formula) is surprising in its universality as it holds for *all games* and *all instances* of (FTRL). As such, it provides a blanket generalization of the well-known volume-preserving property for the replicator dynamics established by Akin [1], as well as subsequent results for zero-sum games [42].

4.3. Instability of fully mixed equilibria. As stated above, the volume-preserving property of (FTRL) would seem to suggest that no strategy can be asymptotically stable. However, this is a figment of the duality between strategy and score variables: a mixed strategy orbit $x(t) = Q(y(t))$ could converge in \mathcal{X} , even though the corresponding dual orbit $y(t)$ diverges in \mathcal{Y} (for an illustration, see Fig. 2 above). This again brings into sharp contrast the behavior of (FTRL) at the boundary of \mathcal{X} versus its behavior at the interior. Our first instability result below shows that the volume-preserving property of (FTRL) rules out the stability of *any* fully mixed equilibrium, in *any* game:




Figure 2: The duality between scores and strategies under (FTRL): the dynamics are volume-preserving in \mathcal{Y} , but a volume of initial conditions could either collapse in finite time (in the Euclidean case, left), or shrink asymptotically (in the logit case, right). This is due to the vastly different geometric properties of each system.

Theorem 1. *A fully mixed Nash equilibrium cannot be asymptotically stable under (FTRL).*

The main idea of the proof of [Theorem 1](#) relies on a tandem application of [Proposition 1](#) together with the dimensionality reduction idea we discussed for the entropic case in [Section 3](#). In the resulting quotient space, the inverse image of an interior point $x^* \in \text{ri}(\mathcal{X})$ is a single point and the induced dynamics remain volume-preserving. If x^* is asymptotically stable, a limit point argument rules out the possibility of a trajectory entering and exiting a small neighborhood of its preimage infinitely many times. At the same time, Lyapunov stability and volume preservation imply that the dynamics are locally recurrent. This contradicts the transient property established above and proves that x^* cannot be asymptotically stable; the details involved in making these arguments precise are fairly intricate, so we defer the proof of [Theorem 1](#) to the supplement.

This universal instability result has significant implications as it provides a dynamic justification of the fragility of fully mixed Nash equilibria. [Theorem 1](#) illustrates this principle through the lens of regret minimization: any deviation from a fully mixed equilibrium invariably creates an opportunity that can be exploited by a no-regret learner. When every player adheres to such a policy, this creates a vicious cycle which destroys any chance of stability for fully mixed equilibria.

4.4. The case of partially mixed equilibria. Taking this premise to its logical extreme, a natural question that arises is whether this instability persists as long as even a *single* player employs a mixed strategy at equilibrium. In the previous case, after the dimensionality reduction argument we described in [Section 3](#), neighborhoods of fully mixed equilibria in the space of strategies (\mathcal{X}) correspond to sets of finite volume in the space of payoffs (\mathcal{Y}). On the contrary, the case of *partially* mixed equilibria is much more complex because neighborhoods of points on the boundary of \mathcal{X} correspond to sets of *infinite* volume in the space of payoffs – and this, even after dimensionality reduction (cf. [Fig. 1](#)). Because of this, volume preservation arguments cannot rule out asymptotic stability of Nash equilibria lying at the boundary of the strategy space: indeed, pure Nash equilibria also lie on the boundary but they *can* be asymptotically stable [17, 20, 38, 39].

In view of the above, it is not a priori clear whether partially mixed equilibria would behave more like pure or fully mixed ones – or if no conclusion can be drawn whatsoever. Our next result shows that the dynamics of FTRL represent a very sharp selection mechanism in this regard:

Theorem 2. *Only strict Nash equilibria can be asymptotically stable under (FTRL).*

Corollary 1. *If x^* is partially mixed, it cannot be asymptotically stable under (FTRL).*

Viewed in isolation, [Theorem 1](#) would seem to be subsumed by [Theorem 2](#), but this is not so: the former plays an integral role in the proof of the latter, so it cannot be viewed as a special case. In more detail, the proof of [Theorem 2](#) builds on [Theorem 1](#) along two separate axes, depending on whether the underlying regularizer is steep or not:

- (1) In the steep case, as we discussed in [Section 3](#) there is a well-posed dynamical system on \mathcal{X} . As we show in the supplement, each face of \mathcal{X} is forward-invariant in this system, so x^* must also be asymptotically stable when constrained to the face \mathcal{X}^* of \mathcal{X} spanned by $\text{supp}(x^*)$. The conclusion of [Theorem 2](#) then follows by noting that x^* is interior in \mathcal{X}^* and applying [Theorem 1](#) to the restriction of the underlying game to \mathcal{X}^* .
- (2) The non-steep case is considerably more difficult because (FTRL) no longer induces a well-posed system on \mathcal{X} . In lieu of this, by examining the finer structure of the inverse image of x^* , it is possible to show the following: for every small enough compact neighborhood \mathcal{K} of x^* in \mathcal{X} , there exists a finite time $\tau_{\mathcal{K}} \geq 0$ such that $\text{supp}(x(t)) = \text{supp}(x^*)$ for all $t \geq \tau_{\mathcal{K}}$ whenever $x(0) \in \mathcal{K}$. As it turns out, the dynamics after $t \geq \tau_{\mathcal{K}}$ locally coincide with the mixed strategy dynamics of (FTRL) applied to the restriction of the underlying game to the face \mathcal{X}^* of \mathcal{X} spanned by x^* . Since x^* is a fully mixed equilibrium in this restricted game, it cannot be asymptotically stable.

4.5. Stable limit sets. We conclude our analysis with a result concerning more general behaviors whereby the dynamics of FTRL do not converge to a point, but to a more general *invariant set* – such as a chain of stationary points interconnected by solution orbits, a structure known as a *heteroclinic cycle* [see e.g., [29](#), [59](#), and references therein]. As an example, in the case of two-player zero-sum games with a fully mixed equilibrium, it is known that the trajectories of (FTRL) form periodic orbits (cycles). However, these orbits are *not* asymptotically stable: if the initialization of the FTRL dynamics is slightly perturbed, the resulting trajectory will be a different periodic orbit, which does not converge to the first (in the language of dynamical systems, the cycles observed in zero-sum games are not *limit cycles*). We are thus led to the following natural question:

What type of invariant structures can arise as stable limits of (FTRL)?

To state this question formally, we will require the setwise version of asymptotic stability: a set \mathcal{S} is called *asymptotically stable* under (FTRL) if *a*) all orbits $x(t) = Q(y(t))$ of (FTRL) that start sufficiently close to \mathcal{S} remain close; and *b*) all orbits that start nearby eventually converge to \mathcal{S} . Then, focusing on the case of steep dynamics to avoid more complicated statements, we have:

Theorem 3. *Every asymptotically stable set of steep (FTRL) contains a pure strategy.*

The proof of [Theorem 3](#) relies on an “infinite descent” argument whereby the faces of \mathcal{X} that intersect with \mathcal{S} are eliminated one-by-one, until only pure strategies remain as candidate elements of \mathcal{S} with minimal support; we provide the details in the supplement.

The importance of [Theorem 3](#) lies in that it provides a succinct criterion for identifying possible attracting sets of (FTRL). Indeed, by Conley’s decomposition theorem (also known as the “fundamental theorem of dynamical systems”) [[19](#)], the flow of (FTRL) in an arbitrary game decomposes into a chain recurrent part and an attracting part (see [[49](#), [50](#)] for several examples/discussion in the case of replicator dynamics). The recurrent part is exemplified by the periodic orbits that arise in zero-sum games with an interior equilibrium (there are no attractors in this case) [[42](#)]. [Theorem 3](#) goes a long way to showing that the attracting part of (FTRL) always intersects the *extremes* of the game’s strategy space – i.e., the players’ set of *pure* strategies. A special case of [Theorem 3](#), in the case of replicator dynamics, was employed in [[48](#)] as a step in the definition of new, dynamics/decomposition-based solution concepts. Formalizing the exact form of this

decomposition in arbitrary games is an open direction for future research with far-reaching implications for the theory of online learning in games.

5. CONCLUDING REMARKS

The well known universal existence theorem for (mixed) Nash equilibria in general games has been very influential not only from a mathematics perspective but also from a public policy one as it seems to suggest that there is no inherent tension in any societal setting between the single-minded pursuit of individual profits and societal stability. Nash equilibria satisfy both desiderata simultaneously. Thus, there is in principle no need for centralized intervention and guidance as market forces will converge upon such a solution.

Our results present an argument in the opposite direction. Unless the game has a pure Nash equilibrium, which is definitely not satisfied in numerous strategic interactions, then societal systems do not self-stabilize, even if they are driven by our most effective payoff seeking dynamics, i.e., gradient learning and its follow-the-regularizer-leader variants. Exploring the tradeoffs between individual optimality and societal stability is thus a much more subtle issue than it first meets the eye, and we hope that we inspire follow-up work that can elucidate these questions further.

ACKNOWLEDGMENTS

This research was partially supported by the COST Action CA16228 “European Network for Game Theory” (GAMENET), the French National Research Agency (ANR) under grant ALIAS, and the Onassis Foundation under Scholarship ID: F ZN 010-1/2017-2018.

E.V. Vlatakis-Gkaragkounis is grateful to be supported by NSF grants CCF-1703925, CCF-1763970, CCF-1814873, CCF-1563155, and by the Simons Collaboration on Algorithms and Geometry.

T. Lianas is supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “First Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant”, project BALSAM, HFRI-FM17-1424.

P. Mertikopoulos is grateful for financial support by the French National Research Agency (ANR) in the framework of the “Investissements d’avenir” program (ANR-15-IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), and MIAI@Grenoble Alpes (ANR-19-P3IA-0003).

G. Piliouras gratefully acknowledges AcRF Tier-2 grant (Ministry of Education – Singapore) 2016-T2-1-170, grant PIE-SGP-AI-2018-01, NRF2019-NRF-ANR095 ALIAS grant and NRF 2018 Fellowship NRF-NRFF2018-07 (National Research Foundation Singapore).

APPENDIX A. AN ONTOLOGY OF NASH EQUILIBRIA: REPRESENTATIVE EXAMPLES

In the archetypal game of Prisoner’s Dilemma (left), it is easy to check that the unique Nash equilibrium is the mutual betrayal which is strict (and hence pure). On the other hand, Matching Pennies (right) is an example of a zero-sum game whose unique Nash equilibrium is fully mixed but still quasi-strict (since all strategies present in its support are unilateral best responses to it). We mention the above to clarify that quasi-strict *does not mean* pure equilibria and includes also the fully mixed Nash equilibrium; the terminology is, perhaps, unfortunate, but otherwise deeply entrenched in the game-theoretic literature [25].

APPENDIX B. BASIC PROPERTIES OF THE FTRL DYNAMICS

B.1. Definitions from dynamical systems. In this appendix, we provide some general preliminaries from general topology and the theory of dynamical systems that we will use freely in the sequel.

A key notion in our analysis is that of (*Poincaré*) *recurrence*. Intuitively, a dynamical system is recurrent if, after a sufficiently long (but *finite*) time, almost every state returns

		Player Y		
		B	S	
Player X	B	(3, 3)	(0, 5)	
	S	(5, 0)	(1, 1)	

		Player Y	
		H	T
Player X	H	(1, -1)	(-1, 1)
	T	(-1, 1)	(1, -1)

Table 1: Prisoner’s Dilemma (left) & Matching Pennies (right).

arbitrarily close to the system’s initial state.⁵ More formally, given a dynamical system on \mathcal{X} that is defined by means of a *semiflow* $\Phi: \mathcal{X} \times [0, \infty) \rightarrow \mathcal{X}$, we have:⁶

Definition 1. A point $x \in \mathcal{X}$ is said to be *recurrent* under Φ if, for every neighborhood U of x in \mathcal{X} , there exists an increasing sequence of times $t_n \uparrow \infty$ such that $\Phi_{t_n}(x) \in U$ for all n . Moreover, the flow Φ is called (*Poincaré*) *recurrent* if, for every measurable subset A of \mathcal{X} , the set of recurrent points in A has full measure.

The above definition directly implies that the flow $\Phi_t(x)$ from a recurrent point x cannot converge to any $x' \neq x$. Poincaré’s recurrence theorem gives sufficient condition for the existence of such points.

Theorem B.1 (Poincaré Recurrence Theorem). *If a flow Φ preserves volume and its orbits are bounded, then almost every point is recurrent under Φ .*

The key notion in the above formulation of the theorem is that of volume preservation: formally, a flow Φ is *volume-preserving* if $\text{vol}(\Phi_t(\mathcal{R})) = \text{vol}(\mathcal{R})$ for any set of initial conditions $\mathcal{R} \subseteq \mathcal{X}$. A useful condition to establish this property is via *Liouville’s formula*, as stated below:

Theorem B.2 (Liouville’s formula). *Let Φ be the flow of a dynamical system with infinitesimal generator V , i.e., $\Phi_t(x)$ is the solution trajectory of the ordinary differential equation*

$$\frac{d}{dt}x(t) = V(x(t)) \tag{B.1}$$

with initial condition $x(0) = x$. Then, letting $\mathcal{R}_t = \Phi_t(\mathcal{R})$ for an arbitrary measurable set \mathcal{R} , we have

$$\frac{d}{dt} \text{vol}[\mathcal{R}_t] = \int_{\mathcal{R}_t} \text{div}[V(x)] dx \tag{B.2}$$

Corollary 2. *If V is incompressible over \mathbb{R}^n (i.e., $\text{div} V(x) = 0$ for all $x \in \mathbb{R}^n$), the induced flow Φ is volume-preserving.*

B.2. Structural properties of the FTRL dynamics: the steep/non-steep dichotomy. To proceed with our analysis, we will need to clarify the precise technical requirements for the dynamics’ regularizer function h . These are as follows: (i) $h \in C_0(\mathbb{R}_+^n) \cap C_2(\mathbb{R}_{++}^n)$, i.e., h is continuous on \mathbb{R}_+^n and two times continuously differentiable on \mathbb{R}_{++}^n ; (ii) h is strongly convex on \mathcal{X} ; and (iii) the inverse Hessian $H(x) = \text{Hess}(h(x))^{-1}$ of h admits a Lipschitz extension to all of \mathcal{X} such that $z^\top H(x)z > 0$ whenever $\text{supp}(z) \supseteq \text{supp}(x)$. These conditions are purely technical in nature and they are satisfied by all the regularizers used in practice, cf. [2, 9, 14, 33, 38, 60] and references therein. As we discussed in the main body of our paper, there is an important distinction to be made depending on whether h is *steep* or *non-steep*. Formally, we say that h is *steep* if $\|\nabla h(x)\| \rightarrow \infty$ whenever $x \rightarrow \text{bd}(\mathcal{X})$ and

⁵Here, “almost” means that the set of such states has full Lebesgue measure.

⁶A smooth map $\Phi: \mathcal{X} \times [0, \infty) \rightarrow \mathcal{X}$ is called a *semiflow* if $\Phi_0(x) = x$ and $\Phi_{t+s}(x) = \Phi_s(\Phi_t(x))$ for all $x \in \mathcal{X}$ and all $t, s \geq 0$. Heuristically, $\Phi_t(x) \equiv \Phi_t(x)$ describes the trajectory of the dynamical system starting at x .




Figure 3: Steep vs. non-steep regularizers (note in particular the singular behavior of the gradient at the boundary in the case of steep regularizers).

$\text{rank}(H(x)) = |\text{supp}(x)|$ for all x ; by contrast, if $\sup_{x \in \mathcal{X}} \|\nabla h(x)\| < \infty$ and $\text{rank}(H(x)) = n$, we say that h is *non-steep*. The qualitative difference in behavior between these cases is illustrated in the figure below (which shows the very different behavior of the derivatives of h near the boundary of the state space).

The following lemma illustrates the relation between mixed strategies and score vectors and the mirror map (6) that defines the dynamics (FTRL). We focus on the perspective of an arbitrary player, say i , and for ease of notation we write Q , x and y instead of Q_i , x_i and y_i respectively. The lemma begins to illustrate the gulf between the steep and non-steep cases.

Lemma B.1. $x = Q(y)$ if and only if there exist $\mu \in \mathbb{R}$ and $\nu_\alpha \in \mathbb{R}_+$ such that, for all $\alpha \in \mathcal{A}$, we have: a) $y_\alpha = \frac{\partial h}{\partial x_\alpha} + \mu - \nu_\alpha$; and b) $x_\alpha \nu_\alpha = 0$. In particular, if h is steep, we have $\nu_\alpha = 0$ for all $\alpha \in \mathcal{A}$.

Proof. Recall that

$$\begin{aligned} Q(y) &= \arg \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\} \\ &= \arg \max \left\{ \sum_{\alpha \in \mathcal{A}} y_\alpha x_\alpha - h(x) : \sum_{\alpha \in \mathcal{A}} x_\alpha = 1 \text{ and } \forall \alpha \in \mathcal{A} : x_\alpha \geq 0 \right\} \end{aligned}$$

The result follows by applying the Karush–Kuhn–Tucker (KKT) conditions to this optimization problem and noting that, since the constraints are affine, the KKT conditions are sufficient for optimality. Our Lagrangian is

$$\mathcal{L}(x, \mu, \nu) = \left(\sum_{\alpha \in \mathcal{A}} y_\alpha x_\alpha - h(x) \right) - \mu \left(\sum_{\alpha \in \mathcal{A}} x_\alpha - 1 \right) + \sum_{\alpha \in \mathcal{A}} \nu_\alpha x_\alpha$$

where the set of constraints (i) of the statement of the lemma are the stationarity constraints, which in our case are $\nabla \mathcal{L}(x, \mu, \nu) = 0 \Leftrightarrow \nabla \left(\sum_{\alpha \in \mathcal{A}} y_\alpha x_\alpha - h(x) \right) = \mu \nabla \left(\sum_{\alpha \in \mathcal{A}} x_\alpha - 1 \right) - \sum_{\alpha \in \mathcal{A}} \nu_\alpha \nabla x_\alpha$, while the set of constraints (ii) of the statement of the lemmas are the complementary slackness constraints. Note that complementary slackness implies that

whenever $\nu_\alpha > 0$ whenever $\alpha \notin \text{supp}(x)$. Finally, if h is steep, we have $|\partial_\alpha h(x)| \rightarrow \infty$ as $x \rightarrow \text{bd}(\mathcal{X})$, which implies that the KKT conditions admit a solution with $\nu_\alpha = 0$. \square

The following lemma shows that if the support of $x(t)$ does not change over a given interval of time, then the evolution of the players' mixed strategies under (FTRL) follows a certain differential equation that can be calculated explicitly. The lemma below also shows that the trajectory of play coincides with the trajectory that would have resulted if the game were constrained to the strategies present in the support of $x(t)$. Again, for ease of notation we focus on player i and omit i from all subscripts.

Proposition B.1. *Let $x(t) = Q(y(t))$ be a mixed strategy orbit of (FTRL), and let \mathcal{T} be an interval over which $\text{supp}(x(t))$ is constant. Then, for all $t \in \mathcal{T}$, $x(t)$ satisfies the mixed strategy dynamics:*

$$\dot{x}_\alpha = \sum_{\beta \in \text{supp}(x)} [H_{\alpha\beta}(x) - H_\alpha(x)H_\beta(x)] v_\alpha(x), \quad (\text{FTRL-s})$$

where $H_\alpha(x) = \left[\sum_{\beta, \beta' \in \text{supp}(x)} H_{\beta\beta'}(x) \right]^{-1/2} \sum_{\beta \in \text{supp}(x)} H_{\alpha\beta}(x)$. In particular, we have the following dichotomy:

- (1) If h is steep, the dynamics (FTRL-s) are well-posed, i.e., they admit unique global solutions from any initial condition $x \in \mathcal{X}$ (including the boundary). Moreover, the faces of \mathcal{X} are forward-invariant under (FTRL-s): the support of $x(t)$ remains constant for all $t \geq 0$.
- (2) If h is non-steep, the dynamics (FTRL-s) are not well-posed: solutions $x(t)$ to (FTRL-s) exist only up to a finite time, after which the support of $x(t)$ may change.

Proof. For the first part of the lemma, we follow a line of reasoning due to [38]. Specifically, letting $g_\alpha(x) = \partial_\alpha h(x)$, Lemma B.1 yields

$$y_\alpha(t) = g_\alpha(t) + \mu(t), \quad \forall t \in I, \forall \alpha \in \mathcal{A}^* \quad (\text{B.3})$$

Since y_α and g_α are both smooth, so is $\mu(t)$. Thus, differentiating with respect to t we get

$$\begin{aligned} \dot{y}_\alpha(t) &= \sum_{\beta \in \mathcal{A}} \frac{\partial^2 h}{\partial x_\beta \partial x_\alpha} \dot{x}_\beta(t) + \dot{\mu}(t) \\ &= \sum_{\beta \in \mathcal{A}^*} \frac{\partial^2 h}{\partial x_\beta \partial x_\alpha} \dot{x}_\beta(t) + \dot{\mu}(t) \end{aligned}$$

since for all $t \in I$ and $\beta \in \mathcal{A} \setminus \mathcal{A}^*$, $x_\beta(t) = 0$, and thus $\dot{x}_\alpha(t) = 0$. Multiplying with the inverse of the Hessian, and omitting t for brevity, we get

$$\dot{x}_\alpha = \sum_{\beta \in \mathcal{A}^*} H_{\alpha\beta} \dot{y}_\beta + \sum_{\beta \in \mathcal{A}^*} H_{\alpha\beta} \dot{\mu} \quad (\text{B.4})$$

By the definition of the dynamics, $\dot{y}_\beta = v_\beta$ and since the support remains constant $\sum_{\alpha \in \mathcal{A}^*} \dot{x}_\alpha = 0$. Summing up Equation (B.4) for $\alpha \in \mathcal{A}^*$ we get

$$\sum_{\alpha, \beta \in \mathcal{A}^*} H_{\alpha\beta} v_\beta + G \dot{\mu} = 0 \Leftrightarrow \dot{\mu} = -\frac{\sum_{\beta \in \mathcal{A}^*} H_\beta v_\beta}{G} \quad (\text{B.5})$$

where $G = \sum_{\beta, \beta' \in \text{supp}(x)} H_{\beta\beta'}(x)$. Substituting the latter and $\dot{y}_\beta = v_\beta$ to Equation (B.4) we get the desired result.

For the second part of the lemma, the well-posedness of (FTRL-s) follows from the fact that $H(x)$ admits a Lipschitz continuous extension to all of x ; moreover, by the rank

assumption, the field $H_\alpha(x)$ is also Lipschitz continuous (since the denominator does not vanish; recall that $\text{im } H(x) = \mathbb{R}^{\text{supp}(x)}$). Finally, forward invariance follows by noting that, for every initial condition $x \in \mathcal{X}$, the quantity $\sum_{\alpha \in \text{supp}(x)} x_\alpha(t)$ is a constant of motion (identically equal to 1), and that $\dot{x}_\alpha = 0$ whenever $x_\alpha = 0$. \square

A key take-away from the above result is that, in the steep case, the dynamics (FTRL-s) can be seen as a patchwork of dynamical systems, each evolving on a specific face of x , and each a continuous extension of the other at the points where they come into contact. Neither of the above is true for the non-steep case.

B.3. Volume preservation in \mathcal{Y} and $\text{ri}(\mathcal{X})$.

After restating it, we proceed to show [Proposition 1](#) by simply applying Liouville's formula ([Theorem B.2](#)).

Proposition 1. *Let $\mathcal{R}_0 \subseteq \mathcal{Y}$ be a set of initial conditions for (FTRL) and let $\mathcal{R}_t = \{y(t) : y(0) \in \mathcal{R}_0\}$ denote its evolution under (FTRL) after time $t \geq 0$. Then, $\text{vol}(\mathcal{R}_t) = \text{vol}(\mathcal{R}_0)$.*

Proof. We have to show that the dynamics of (FTRL) (i.e., $\dot{y}(t) = v(Q(y(t)))$) are incompressible. For any player i and any $\alpha \in \mathcal{A}_i$ we have

$$\frac{\partial v_{i\alpha}}{\partial y_{i\alpha}} = \sum_{\beta \in \mathcal{A}_i} \frac{\partial v_{i\alpha_i}}{\partial x_{i\beta}} \frac{\partial x_{i\beta}}{\partial y_{i\alpha}} = 0, \quad (\text{B.6})$$

because v_i does not depend on x_i . We thus obtain $\text{div}_y v(y) = 0$, i.e., the dynamics (FTRL) are incompressible. The result then follows from Liouville's formula \square

In the following lemma, we show that the flow defined by FTRL in the interior $\text{ri}(\mathcal{X})$ of \mathcal{X} is incompressible under a suitably defined measure. For that, using Liouville's formula, we first show that in the so-called z -space, i.e., a ‘‘slice’’ of the payoff space, the respective flow is incompressible. Using that, we can easily get a diffeomorphism from $\text{ri}(\mathcal{X})$ to the z -space, we define the volume of a set in $\text{ri}(\mathcal{X})$ to be the volume of the corresponding set in the z -space, and thus incompressibility in the interior comes for free.

Lemma B.2. *There exists a measure μ_x for which the flow in the interior of \mathcal{X} is incompressible, i.e., for any subset $U \subset \text{ri}(\mathcal{X})$ of initial conditions, and any $t_0 \geq 0$ so that for any $0 \leq t \leq t_0$: $\Phi(U, t) \subset \text{ri}(\mathcal{X})$, it is $\mu_x(U) = \mu_x(\Phi(U, t))$.*

Proof. First we go on to define the z -space. The intuition for defining and using the z -space can be based on [Lemma B.1](#) which implies that for any $x \in \text{ri}(\mathcal{X})$, any two corresponding points y, y' in the y -space differ by a constant, since for all i and $\alpha_j \in \mathcal{A}_i$, $y_{i\alpha_j} = \frac{\partial h}{\partial x_{i\alpha_j}} + \mu_i$ and $y'_{i\alpha_j} = \frac{\partial h}{\partial x_{i\alpha_j}} + \mu'_i$ for some μ_i and μ'_i (recall $x \in \text{ri}(\mathcal{X})$ implies $\nu_{i\alpha_j} = 0$). Thus, all y 's that correspond to an $x \in \text{ri}(\mathcal{X})$ form an equivalent class. For each class, we pick as representative the y in the class that has 0 in some specific coordinate $\hat{\alpha}_i$, for every player i . The set of representatives form the z -space and there is a one to one correspondence of points of $\text{ri}(\mathcal{X})$ to points in the z -space which moreover can be used to define an incompressible flow in $\text{ri}(\mathcal{X})$.




Figure 4: From payoffs to strategies and back, the steep case: deformation of neighborhoods under the logit map $Q = \Lambda$.




Figure 5: The three different representation spaces of (FTRL): primal (strategies, left); dual (payoffs/scores, center); dual quotient (payoff/score differences, right).

So, for a *benchmark* strategy $\hat{\alpha}_i \in \mathcal{A}_i$ for every player $i \in \mathcal{N}$ and for all $\alpha \in \mathcal{A}_i \setminus \{\hat{\alpha}_i\} \equiv \hat{\mathcal{A}}_i$ consider the corresponding score differences

$$z_{i\alpha} = y_{i\alpha} - y_{i\hat{\alpha}_i}. \quad (\text{B.7})$$

Obviously, $z_{i\alpha} = y_{i\alpha} - y_{i\hat{\alpha}_i}$ is identically zero so we can ignore it in the above definition. In so doing, we obtain a linear map $\Pi_i: \mathbb{R}^{\mathcal{A}_i} \rightarrow \mathbb{R}^{\hat{\mathcal{A}}_i}$ sending $y_i \mapsto z_i$; aggregating over all players, we also write Π for the product map $\Pi = (\Pi_1, \dots, \Pi_N)$ sending $y \mapsto z$. For posterity, note that this map is surjective but *not* injective,⁷ so it does not allow us to recover the score vector y from the score difference vector z .

Now, under FTRL, the score differences (B.7) evolve as

$$\dot{z}_{i\alpha} = v_{i\alpha}(x(t)) - v_{i\hat{\alpha}_i}(x(t)). \quad (\text{B.8})$$

⁷Specifically, $\Pi_i(y_i) = \Pi_i(y'_i)$ if and only if $y'_{i\alpha_i} = y_{i\alpha_i} + c$ for some $c \in \mathbb{R}$ and all $\alpha_i \in \mathcal{A}_i$.

Our first step below is to show that (B.8) constitutes a well-defined dynamical system on z as long as the corresponding x 's remain in $\text{ri}(\mathcal{X})$.

To do so, consider the reduced mirror map $\hat{Q}_i: \mathbb{R}^{\mathcal{A}_i} \rightarrow \mathcal{X}_i$ defined as

$$\hat{Q}_i(z_i) = Q_i(y_i) \quad (\text{B.9})$$

for some $y_i \in \mathbb{R}^{\mathcal{A}_i}$ such that $\Pi_i(y_i) = z_i$. That such a y_i exists is a consequence of Π_i being surjective; furthermore, that $\hat{Q}_i(z_i)$ is well-defined is a consequence of the fact that Q_i is invariant on the fibers of Π_i . Indeed, by construction, and as long as the corresponding x 's remain in $\text{ri}(\mathcal{X})$ we have $\Pi_i(y_i) = \Pi_i(y'_i)$ if and only if $y'_{i\alpha} = y_{i\alpha} + c$ for some $c \in \mathbb{R}$ and all $\alpha \in \mathcal{A}_i$. Hence, by the definition of Q_i , we get

$$\begin{aligned} Q_i(y'_i) &= \arg \max_{x_i \in \mathcal{X}_i} \{ \langle y_i, x_i \rangle + c \sum_{\alpha \in \mathcal{A}_i} x_{i\alpha} - h_i(x_i) \} \\ &= \arg \max_{x_i \in \mathcal{X}_i} \{ \langle y_i, x_i \rangle - h_i(x_i) \} = Q_i(y_i), \end{aligned} \quad (\text{B.10})$$

where we used the fact that $\sum_{\alpha \in \mathcal{A}_i} x_{i\alpha} = 1$. The above shows that $Q_i(y'_i) = Q_i(y_i)$ if and only if $\Pi_i(y_i) = \Pi_i(y'_i)$, so \hat{Q}_i is well-defined. Letting $\hat{Q} \equiv (\hat{Q}_1, \dots, \hat{Q}_N)$ denote the aggregation of the players' individual mirror maps \hat{Q}_i , it follows immediately that $Q(y) = \hat{Q}(\Pi(y)) = \hat{Q}(z)$ by construction.

Hence, the dynamics (B.8) may be written as

$$\dot{z} = V(z), \quad (\text{B.11})$$

where

$$\nu_{i\alpha}(z) = v_{i\alpha}(\hat{Q}_i(z)) - v_{i\hat{\alpha}_i}(\hat{Q}_i(z)). \quad (\text{B.12})$$

These dynamics obviously constitute an autonomous system.

Next, we show incompressibility of the z -space. Indeed, for all $\alpha \in \mathcal{A}$ we have

$$\frac{\partial \nu_{i\alpha}}{\partial z_{i\alpha}} = \sum_{\beta \in \mathcal{A}_i} \frac{\partial \nu_{i\alpha_i}}{\partial x_{i\beta}} \frac{\partial x_{i\beta}}{\partial z_{i\alpha}} = 0, \quad (\text{B.13})$$

because v_i does not depend on x_i . We thus obtain $\nabla_z \cdot V(z) = 0$, i.e., the dynamics (B.11) are incompressible.

For the last step, for a set $A \subset \text{ri}(\mathcal{X})$ define $\mu_x(A) := \mu_z(Q^{-1}(A))$, where μ_z is the Lebesgue measure in the z -space. Then for any $U \subset \text{ri}(\mathcal{X})$, as long as $\Phi(U, t)$ remains in $\text{ri}(\mathcal{X})$, it is

$$\mu_x(U) = \mu_z(Q^{-1}(U)) = \mu_z(Q^{-1}(\Phi(U, t))) = \mu_x(\Phi(U, t))$$

as needed. \square

APPENDIX C. PROOF OF THEOREM 1

Below we show that there are no asymptotically stable sets (or points) in $\text{ri}(\mathcal{X})$. Indeed, if this were the case, there would be a full-measure set of initial conditions outside the asymptotically stable set A^* that converges to A^* , while at the same time its trajectories are bounded (by stability). This contradicts Poincaré's recurrence theorem, because the flow in $\text{ri}(\mathcal{X})$ is volume-preserving by Lemma B.2.

Theorem C.1. *Let A^* be a closed set of $\text{ri}(\mathcal{X})$. Then A^* is not asymptotically stable under FTRL.*




Figure 6: The various sets in the proof of [Theorem C.1](#).

Proof. To reach a contradiction, let A^* be an asymptotically stable set, i.e., attracting and Lyapunov stable, belonging in $\text{ri}(\mathcal{X})$. Since A^* is attracting, there exists a neighborhood U of A^* all points of which converge to A^* . Without loss of generality, since A^* is closed, we may assume that U lies in $\text{ri}(\mathcal{X})$, and its closure is disjoint from the boundary of \mathcal{X} .

Now, since A^* is Lyapunov stable, there exists some neighborhood U_0 of A^* so that whenever $x(0) \in U_0$, $x(t) \in U$. Pick some $x_0 \in U_0 \setminus A^*$. Since A^* is closed and U_0 is open, there is a small enough neighborhood E of x_0 so that all points of E lie inside $U_0 \setminus A^*$. By Lyapunov stability, for all $t \geq 0$, $\Phi(E, t) \subseteq U$. But then the set $E_\infty = \cup_{t \geq 0} \Phi(E, t)$ is bounded, having positive measure that does not change over time ([Lemma B.2](#)). Therefore it is a Poincaré recurrent set. But this means that all but a measure zero set of initializations in E lead to recurrent trajectories that return infinitely often to E . Picking E to be bounded away from A^* (which is a closed set) we conclude that there are points in E (and thus U) that do not converge to A^* , a contradiction. \square

Now, given that any singleton set $\{x\}$, $x \in \mathcal{X}$, is closed, the above yields:

Theorem C.2. *There are no asymptotically stable points in $\text{ri}(\mathcal{X})$*

[Theorem 1](#) (restated below) then follows as a corollary.

Theorem 1. *A fully mixed Nash equilibrium cannot be asymptotically stable under (FTRL).*

APPENDIX D. PROOF OF [THEOREM 2](#): THE NON-STEEP CASE

Our goal in this appendix is to provide the proof of [Theorem 2](#), which we restate below for convenience:

Theorem 2. *Only strict Nash equilibria can be asymptotically stable under (FTRL).*

Because of the fundamental dichotomy between steep and non-steep FTRL dynamics, we will break the proof in two cases, treating here the non-steep regime; the steep case will be proved in [Appendix E](#) as a consequence of a more general result. The fundamental distinction between the two cases is that, in the non-steep regime, the mixed-strategy dynamics of (FTRL) could change support infinitely many times, which means that the type of volume-preservation arguments employed in the previous section cannot work (because the corresponding preimages in the z -space could have infinite volume; see below for a graphical illustration). However, as we show below, this “change of support” is a

blessing in disguise: if x^* is asymptotically stable, nearby trajectories will end up employing only those strategies present in x^* in finite time.




Figure 7: From payoffs to strategies and back, the non-steep case: deformation of neighborhoods under the Euclidean projection map $Q = \Pi$.

The following lemma shows that, for generic games, if the underlying regularizer is non-steep, all trajectories starting near an asymptotically stable point x^* attain the face of x^* in some uniform, finite time. The intuition for this is that, generically, for any player i , the coordinates of y_i that correspond to the support of x_i^* increase with a “speed” that is uniformly higher than those strategies not supported in x^* . The regularizer of player i could possibly act in favor of the coordinates that do not belong to the support, but in a bounded way, since it is non-steep. Thus, there is a time after which the coordinates of y_i corresponding to the support are bigger enough than the other coordinates, so that the mirror map Q_i keeps returning a point with support equal to the support of x_i^* .

Lemma D.1. *Let x^* be an asymptotically stable equilibrium of a generic finite game Γ , with the regularizers used, being non-steep. For any neighborhood U of x^* , there exists a neighborhood U_0 of x^* and a finite time T_0 such that if $x(t) = Q(y(t))$ is an orbit of FTRL starting at $x(0) \in U_0$, then $\text{supp}(x(t)) = \text{supp}(x^*)$ for all $t \geq T_0$.*

Proof. By the genericity assumption, all Nash equilibria are quasi-strict. Clearly we have that for any player i , $u_{i\alpha}(x^*) > u_{i\beta}(x^*)$ for all $\alpha \in \text{supp}(x_i^*) \equiv \mathcal{A}_i^*$ and $\beta_i \notin \text{supp}(x_i^*)$. Thus, by continuity there exists some neighborhood U of x^* and a $c > 0$ so that for any $x \in U$ and any player i , $u_{i\alpha}(x) > u_{i\beta}(x) + c$ for all $\alpha \in \mathcal{A}_i^*, \beta \in \mathcal{A}_i \setminus \mathcal{A}_i^*$. Additionally, we can choose U small enough so that for all $x \in U$, $\text{supp}(x^*) \subseteq \text{supp}(x)$. Since x^* is asymptotically stable there exists a neighborhood U_0 of x^* so that $x(t) \in U$ for all t whenever $x(0) \in U_0$, and $\lim_{t \rightarrow \infty} x(t) = x^*$.

Consider some $x(0) \in U_0$ and let $\alpha \in \mathcal{A}^*$, $\beta \in \mathcal{A} \setminus \mathcal{A}^*$. For ease of notation in the following we focus on the perspective of an arbitrary player, say i , and omit i from the subscripts.

By Lemma B.1, for any $t \geq 0$ there exist a $\mu(t)$ and non negative $v_\alpha(t)$'s so that

$$\begin{aligned} y_\alpha(t) &= g_\alpha(x(t)) + \mu(t) & \forall \alpha \in \mathcal{A}^* \\ y_\alpha(t) &= g_\alpha(x(t)) + \mu(t) - v_\beta(t) & \forall \beta \in \mathcal{A} \setminus \mathcal{A}^* \end{aligned}$$

since, by complementary slackness, $v_\alpha(t) = 0$, whenever $x_\alpha(t) > 0$. Subtracting we get

$$y_\alpha(t) - y_\beta(t) = g_\alpha(x(t)) - g_\beta(x(t)) + v_\beta(t) \leq v_\beta(t) + G \quad (\text{D.1})$$

with the inequality following, for some constant G , by h being non-steep.

On the other hand by the definition of the dynamics, using Equation (D.1) and that $u_\alpha(x(t)) > u_\beta(x(t)) + c$, for all t (since $x(0) \in U_0$), we get

$$\begin{aligned} y_\alpha(t) - y_\beta(t) &= y_\alpha(0) - y_\beta(0) + \int_0^t [u_\alpha(x(s)) - u_\beta(x(s))] ds \\ &> g_\alpha(x(0)) - g_\beta(x(0)) + v_\beta(0) + ct \\ &\geq ct + v_\beta(0) - G \end{aligned}$$

with the last inequality following again by h being non-steep. Combining the latter with Equation (D.1), and since $v_\beta(0) \geq 0$, we get

$$v_\beta(t) + G > ct + v_\beta(0) - G \Rightarrow v_\beta(t) > ct - 2G$$

which implies that for $t \geq \frac{2G}{c}$ it is $v_\beta(t) > 0$. This in turn, by complementary slackness, yields $x_\beta(t) = 0$ for all $t \geq \frac{2G}{c}$, implying $\text{supp}(x(t)) \subseteq \text{supp}(x^*)$. By the choice of U and since $\forall t : x(t) \in U$ we have $\text{supp}(x(t)) \supseteq \text{supp}(x^*)$ and thus setting $T_0 = \frac{2G}{c}$ proves the claim, since the above holds for any player i . \square

The main result of this section is the following theorem that covers the non-steep case, stating that for generic games at an asymptotically stable point under non-steep regularizers, every player plays a pure strategy. The proof combines results presented above. It reaches a contradiction by showing that points in a small enough neighborhood of the asymptotically stable point x^* , instead of converging to it as they ought to, they follow recurrent trajectories. In a first step it finds points that after a finite time T_0 reach and stay forever at the simplex formed by the support of x^* (using Lemma D.1), which moreover have non-zero volume in that simplex. But then, these points follow FTRL trajectories in the restricted simplex (Proposition B.1) and x^* belongs in the interior of this simplex. However, we already know this cannot be the case in the interior (Theorem C.2) and we follow a similar reasoning.

Theorem D.1. *If $x^* \in \mathcal{X}$ is an asymptotically stable point under non-steep regularizers of a generic game Γ , then it consists of only pure strategies.*

Proof. Let $x^* \in \mathcal{X}$ be asymptotically stable, $\mathcal{A}^* = \text{supp}(x^*)$, with $|\mathcal{A}_i^*| = |\text{supp}(x_i^*)| \geq 2$ for some player i , and \mathcal{X}^* be its respective simplex. Since x^* is attracting, there exists some (bounded) neighborhood U of x^* for which if $x(0) \in U$, then $\lim_{t \rightarrow \infty} x(t) = x^*$. By Lemma D.1, there exists a neighborhood U_0 of x^* and a finite time T_0 such that if $x(t) = Q(y(t))$ is an orbit of FTRL starting at $x(0) \in U$, then $\text{supp}(x(t)) = \mathcal{A}^*$ for all $t \geq T_0$.

By Proposition B.1, for $t \geq T_0$ all trajectories satisfy Equation (FTRL-s) and these trajectories coincide with the trajectories of a generic game Γ' played on \mathcal{A}^* , with the

restricted simplex being \mathcal{X}^* . At the same time, similar to the proof of [Theorem C.1](#), $\Phi(U_0, T_0)$ is a bounded (as a subset of $U \cap \mathcal{X}^*$), positive measure for \mathcal{X}^* (as the evolution of an open set after a finite time (recall $|\mathcal{A}_i^*| \geq 2$ for some i)), and invariant ([Lemma B.2](#)) set of \mathcal{X}^* and, therefore, it is also Poincaré recurrent, contradicting that for $x(0) \in \Phi(U_0, T_0) \subseteq U$, $\lim_{t \rightarrow \infty} x(t) = x^*$, as implied by the asymptotic stability of x^* . Thus, if $x^* \in \mathcal{X}$ is asymptotically stable then $|\text{supp}(x_i^*)| = 1$ for all i . \square

As we stated in the beginning of this appendix, the steep case of [Theorem 2](#) comes as a corollary ([Corollary 3](#)) of a more general result on asymptotically stable sets that we prove in the next section ([Theorem 3](#)).

APPENDIX E. PROOF OF [THEOREM 3](#) AND [THEOREM 2](#): THE STEEP CASE

The following theorem shows that any asymptotically stable set A cannot be contained in the interior of any non-singleton face \mathcal{X}' . This comes as a consequence of [Theorem C.1](#). When the regularizers are steep, any point starting in $\text{ri}(\mathcal{X}')$ stays in $\text{ri}(\mathcal{X}')$ over time and A being asymptotically stable implies that $A \cap \mathcal{X}'$ is an asymptotically stable set under the FTRL dynamics of the restricted game played on \mathcal{X}' . But for the restricted game [Theorem C.1](#) applies excluding the possibility of $A \cap \mathcal{X}'$ being an asymptotically stable set inside $\text{ri}(\mathcal{X}')$.

Theorem E.1. *Let $A \subseteq \mathcal{X}$ be an asymptotically stable set intersecting a non-singleton face \mathcal{X}' of \mathcal{X} . Then, $A \cap \mathcal{X}'$ cannot be contained in the relative interior of \mathcal{X}' .*

Proof. To reach a contradiction let A intersect a non-singleton face \mathcal{X}' of \mathcal{X} and $A' = A \cap \mathcal{X}'$ be a subset of the relative interior of \mathcal{X}' . We will show that if this is the case then A' is an asymptotically stable set under the dynamics of FTRL restricted to \mathcal{X}' , that lies in the relative interior of \mathcal{X}' . This contradicts [Theorem C.1](#).

To reach the contradiction, we go on to prove that A' is an asymptotically stable set under FTRL dynamics in \mathcal{X}' . We will crucially use that with steep regularizers for any $x(0) \in \text{ri}(\mathcal{X}')$, $x(t) \in \text{ri}(\mathcal{X}')$, for all $t \geq 0$, i.e., \mathcal{X}' is forward invariant under FTRL ([Proposition B.1](#)).

To show Lyapunov stability of A' in \mathcal{X}' pick any neighborhood U' of A' in \mathcal{X}' . It can be written as $U' = U \cap \mathcal{X}'$ for some neighborhood U of A in \mathcal{X} . Since A is Lyapunov stable, there exists a neighborhood U_0 of A in \mathcal{X} such that for any $x(0) \in U_0$, it is $x(t) \in U$ for all $t \geq 0$. Let $U'_0 = U_0 \cap \mathcal{X}'$. Using that \mathcal{X}' is forward invariant, the latter implies that for any $x(0) \in U'_0 = U_0 \cap \mathcal{X}'$, it is $x(t) \in U \cap \mathcal{X}' = U'$ for all $t \geq 0$, as needed.

We use similar ideas to show that A' is attracting in \mathcal{X}' . Since A is attracting in \mathcal{X} there exist a neighborhood U of A in \mathcal{X} such that for any $x(0) \in U$, $x(t) \rightarrow A$. Let $U' = U \cap \mathcal{X}'$. The latter combined with the forward invariance of \mathcal{X}' implies that for any $x(0) \in U \cap \mathcal{X}' = U'$, $x(t) \rightarrow A \cap \mathcal{X}' = A'$, as needed. \square

As a corollary of the above theorem we may get [Theorem 3](#), restated below. Whenever an asymptotically stable set intersects a non-singleton face it must intersect its respective boundary, and thus a face of smaller dimension. Consequently, “in the long run”, it must intersect a singleton face. Put differently, it should contain a point consisting of only pure strategies.

Theorem 3. *Every asymptotically stable set of steep (FTRL) contains a pure strategy.*

Proof. Let A be asymptotically stable and \mathcal{X}_{\min} be a face of minimal dimension intersected by A which is not a singleton. By [Theorem E.1](#), A cannot be contained in the relative

interior of \mathcal{X}_{min} , so it must intersect the boundary of \mathcal{X}_{min} . However, this means that A intersects a face of dimension strictly smaller than that of \mathcal{X}_{min} , a contradiction. Thus, \mathcal{X}_{min} is a singleton and A must contain a vertex of \mathcal{X} . \square

Corollary 3. *If x is an asymptotically stable point, then it consists of only pure strategies.*

REFERENCES

- [1] Ethan Akin. Domination or equilibrium. *Mathematical Biosciences*, 50(3-4):239–250, 1980.
- [2] Felipe Alvarez, Jérôme Bolte, and Olivier Brahic. Hessian Riemannian gradient flows in convex programming. *SIAM Journal on Control and Optimization*, 43(2):477–501, 2004.
- [3] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [4] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995.
- [5] James P. Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In *EC '18: Proceedings of the 19th ACM Conference on Economics and Computation*, 2018.
- [6] James P. Bailey and Georgios Piliouras. Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. In *NeurIPS '19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [7] James P. Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *COLT '20: Proceedings of the 33rd Annual Conference on Learning Theory*, 2020.
- [8] Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, New York, NY, USA, 2 edition, 2017.
- [9] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [10] Victor Boone and Georgios Piliouras. From Darwin to Poincaré and von Neumann: Recurrence and cycles in evolutionary and algorithmic game theory. In *WINE 2019: Proceedings of the 15th Conference on Web and Internet Economics*, 2019.
- [11] Mario Bravo and Panayotis Mertikopoulos. On the robustness of learning in games with stochastically perturbed payoff observations. *Games and Economic Behavior*, 103, John Nash Memorial issue:41–66, May 2017.
- [12] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [13] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [14] Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, August 1993.
- [15] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *COLT '19: Proceedings of the 32nd Annual Conference on Learning Theory*, 2019.
- [16] Yun Kuen Cheung and Georgios Piliouras. Chaos, extremism and optimism: Volume analysis of learning in games. <https://arxiv.org/abs/2005.13996>, 2020.
- [17] Johanne Cohen, Amélie Héliou, and Panayotis Mertikopoulos. Learning with bandit feedback in potential games. In *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [18] Johanne Cohen, Amélie Héliou, and Panayotis Mertikopoulos. Hedging under uncertainty: Regret minimization meets exponentially fast convergence. In *SAGT '17: Proceedings of the 10th International Symposium on Algorithmic Game Theory*, 2017.
- [19] Charles Cameron Conley. *Isolated Invariant Set and the Morse Index*. American Mathematical Society, Providence, RI, 1978.
- [20] Pierre Coucheney, Bruno Gaujal, and Panayotis Mertikopoulos. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, August 2015.

- [21] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC '06: Proceedings of the 38th annual ACM symposium on the Theory of Computing*, 2006.
- [22] Eddie Dekel and Drew Fudenberg. Rational behavior with payoff uncertainty. *Journal of Economic Theory*, 52:243–267, 1990.
- [23] Dylan J. Foster, Thodoris Lykouris, Kathrik Sridharan, and Éva Tardos. Learning in games: Robustness of fast convergence. In *NIPS '16: Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4727–4735, 2016.
- [24] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.
- [25] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.
- [26] James Hannan. Approximation to Bayes risk in repeated play. In Melvin Dresher, Albert William Tucker, and P. Wolfe (eds.), *Contributions to the Theory of Games, Volume III*, volume 39 of *Annals of Mathematics Studies*, pp. 97–139. Princeton University Press, Princeton, NJ, 1957.
- [27] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, September 2000.
- [28] Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- [29] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.
- [30] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.
- [31] Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. <https://arxiv.org/abs/2006.09065>, 2020.
- [32] Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, 1997.
- [33] Krzysztof C. Kiwiel. Free-steering relaxation methods for problems with strictly convex costs and linear constraints. *Mathematics of Operations Research*, 22(2):326–349, 1997.
- [34] Joon Kwon and Panayotis Mertikopoulos. A continuous-time approach to online optimization. *Journal of Dynamics and Games*, 4(2):125–148, April 2017.
- [35] Rida Laraki, Jérôme Renault, and Sylvain Sorin. *Mathematical Foundations of Game Theory*. Universitext. Springer, 2019.
- [36] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- [37] Panayotis Mertikopoulos and Aris L. Moustakas. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4):1359–1388, July 2010.
- [38] Panayotis Mertikopoulos and William H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, November 2016.
- [39] Panayotis Mertikopoulos and William H. Sandholm. Riemannian game dynamics. *Journal of Economic Theory*, 177:315–364, September 2018.
- [40] Panayotis Mertikopoulos and Yannick Viossat. Imitation dynamics with payoff shocks. *International Journal of Game Theory*, 45(1-2):291–320, March 2016.
- [41] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.
- [42] Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *SODA '18: Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.
- [43] Sai Ganesh Nagarajan, Sameh Mohamed, and Georgios Piliouras. Three body problems in evolutionary game dynamics: Convergence, periodicity and limit cycles. In *AAMAS '18: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2018.
- [44] Sai Ganesh Nagarajan, David Balduzzi, and Georgios Piliouras. From chaos to order: Symmetry and conservation laws in game dynamics. In *ICML '20: Proceedings of the 37th International Conference on Machine Learning*, 2020.
- [45] Arkadi Semen Nemirovski and David Berkovich Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley, New York, NY, 1983.
- [46] Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.

- [47] Noam Nisan, Tim Roughgarden, Éva Tardos, and V. V. Vazirani (eds.). *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [48] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. α -rank: Multi-agent evaluation by evolution. *Scientific reports*, 9(1):1–29, 2019.
- [49] Christos H. Papadimitriou and Georgios Piliouras. From Nash equilibria to chain recurrent sets: An algorithmic solution concept for game theory. *Entropy*, 20(10):782, 2018.
- [50] Christos H. Papadimitriou and Georgios Piliouras. Game dynamics as the meaning of a game. *ACM SIGecom Exchanges*, 16(2):53–63, 2019.
- [51] Julien Perolat, Remi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, Mark Rowland, Pedro Ortega, Neil Burch, Thomas Anthony, David Balduzzi, Bart De Vylder, Georgios Piliouras, Marc Lanctot, and Karl Tuyls. From Poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. <https://arxiv.org/abs/2002.08456>, 2020.
- [52] Georgios Piliouras and Leonard J. Schulman. Learning dynamics and the co-evolution of competing sexual species. <https://arxiv.org/abs/1711.06879>, 2017.
- [53] Georgios Piliouras and Jeff S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence. In *SODA '14: Proceedings of the 25th annual ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [54] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *COLT '13: Proceedings of the 26th Annual Conference on Learning Theory*, 2013.
- [55] Klaus Ritzberger. The theory of normal form games from the differentiable viewpoint. *International Journal of Game Theory*, 23:207–236, September 1994.
- [56] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.
- [57] Mark Rowland, Shayegan Omidshafiei, Karl Tuyls, Julien Perolat, Michal Valko, Georgios Piliouras, and Remi Munos. Multiagent evaluation under incomplete information. In *NeurIPS '19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [58] Aviad Rubinstein. Inapproximability of Nash equilibrium. *SIAM Journal on Computing*, 47(3): 917–959, 2018.
- [59] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.
- [60] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [61] Shai Shalev-Shwartz and Yoram Singer. Convex repeated games and Fenchel duality. In *NIPS' 06: Proceedings of the 19th Annual Conference on Neural Information Processing Systems*, pp. 1265–1272. MIT Press, 2006.
- [62] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *NIPS '15: Proceedings of the 29th International Conference on Neural Information Processing Systems*, 2015.
- [63] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.
- [64] Yannick Viossat and Andriy Zapechelnjuk. No-regret dynamics and fictitious play. *Journal of Economic Theory*, 148(2):825–842, March 2013.
- [65] Emmanouil Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, and Georgios Piliouras. Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games. In *NeurIPS '19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [66] Vladimir G. Vovk. Aggregating strategies. In *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, pp. 371–383, 1990.
- [67] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.
- [68] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML '03: Proceedings of the 20th International Conference on Machine Learning*, pp. 928–936, 2003.