From Nash Equilibria to Chain Recurrent Sets: Solution Concepts and Topology

Christos Papadimitriou
University of California, Berkeley
christos@berkeley.edu

Georgios Piliouras
Singapore University of Technology and Design
Simons Institute for the Theory of Computing
University of California, Berkeley
georgios@sutd.edu.sg

ABSTRACT

Nash's universal existence theorem for his notion of equilibria was essentially an ingenious application of fixed point theorems, the most sophisticated result in his era's topology — in fact, recent algorithmic work has established that Nash equilibria are in fact computationally equivalent to fixed points. Here, we shift focus to universal non-equilibrium solution concepts that arise from an important theorem in the topology of dynamical systems that was unavailable to Nash. This approach takes as input both a game and a learning dynamic, defined over mixed strategies. Nash equilibria are guaranteed to be fixed points of such dynamics; however, the system behavior is captured by a more general object that is known in dynamical systems theory as *chain recurrent set.* Informally, once we focus on this solution concept, every game behaves like a potential game with the dynamic converging to these states. We characterize this solution for simple benchmark games under replicator dynamics, arguably the best known evolutionary dynamic in game theory. For potential games it coincides with the notion of equilibrium; however, in simple zero sum games, it can cover the whole state space. We discuss numerous novel computational as well as structural, combinatorial questions that chain recurrence raises.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

Keywords

Algorithmic Game Theory; Replicator Dynamics; Invariant; Entropy

1. INTRODUCTION

Game theory has enjoyed a close relationship with topology of dynamical systems from its very beginnings. Nash's brilliant proof of the universality of equilibria is based on fixed point theorems, the most sophisticated class of topological theorems of his time. Recent work in algorithmic game theory has coupled these concepts even further arguing that that these ideas are, in fact, in a formal sense computationally equivalent [11]. However, natural dynamics do not always converge to a Nash equilibrium. *To which object do they converge?*

To answer, one must apply the modern pinnacle of dynamical systems theory, the fundamental theorem of dynamical systems, which was introduced in the seminal work of Conley [8] in 1978. The theorem states that, given an arbitrary initial condition, a dynamical system converges to a set of states with a natural recurrence property, an elegant concept that should appeal to a computer scientist, especially a student of cryptography.

Imagine that Alice is trying to simulate the trajectory of a given system on a powerful computer. Every time she computes a single iteration of the dynamical process, there is a rounding error ϵ. Furthermore, imagine that inside the machine there is an infinitely powerful demon, Bob who, before the next computational step is taken, rounds the result in arbitrary fashion of his own choosing (but within distance ϵ of the actual outcome). If, no matter how high the accuracy of Alice's computer is, Bob can always fool her into believing that a specific point is periodic, then this point is a chain recurrent point.

Naturally, periodic points are chain recurrent, since Bob does not have to intervene. Therefore, equilibria are also chain recurrent since they are trivially periodic. On the other hand, there exist chain recurrent point that are not periodic. Imagine for example a system that is defined as the product of two points moving at constant angular speed along circles of equal length, but where the ratio of the speed of the two points is an irrational number, *e.g.*, $\sqrt{2}$. This system will, after some time, get arbitrarily close to its initial position and a slight bump of the state by Bob would convince Alice that the initial condition is periodic.

In a sense, chain recurrent points are the natural generalization of periodic points in a world where measurements work like a PTAS/FPTAS algorithm, with arbitrarily high accuracy. The surprising implication of the fundamental theorem of dynamical system is that this generalization is not just necessary when arguing about computational systems but also sufficient. It captures all possible limit points of the system.

How does this connect back to game theory? We can apply the fundamental theorem of dynamical systems on the system (G, f) that emerges by coupling any game G and a

ITCS'16, January 14–16, 2016, Cambridge, MA, USA.

ACM 978-1-4503-4057-1/16/01 ...$15.00.
DOI: http://dx.doi.org/10.1145/2840728.2840757.

learning dynamic f. (For most of this paper, we take f to be the replicator dynamics, that is, the continuous-time version of Multiplicative Weights Update.) This defines a new solution concept, the set of chain recurrent points of the system, which we denote R(G,f).

Naturally, we wish to understand how this concept compares against the standard game theoretic solution of Nash equilibria. Since Nash equilibria are fixed points for any reasonable learning dynamics the set of chain recurrent points will be a superset of the set of Nash. However, can it be a strict superset and by how much? Also, potential games comprise the one class of games where learning dynamics indeed work well and converge to equilibria. Does this new solution concept reflect this?

We exemplify this solution concept in two simple and well-understood classes of games, in which it is particularly well behaved, namely 2×2 zero-sum games, and potential games.

In 2×2 zero-sum games we show that once a fully mixed Nash equilibrium exists (like in Matching Pennies) then the set of chain recurrent states is the whole state space. This is an extremal example, in the sense that this is a zero-sum game with a unique Nash equilibrium, where the dominant belief is that any natural, self-interested dynamic will "solve" the game. In fact, it is conventional wisdom amongst game theorists that all no-regret dynamics "converge" in a weak sense to Nash equilibria. By that it is meant that the time-averages of the strategies of both agents will converge to the Nash equilibrium. This is true for the replicator dynamic, which indeed is a continuous time analogue of well known regret minimizing dynamics (*i.e.* Multiplicative Weights Update Algorithm). Nevertheless, time averages fail to fully capture the actual system behavior.

The phase space is partitioned into cyclic trajectories that are essentially "co-centric" orbits around the maxmin [12]. Rounding errors can cause the system to jump from one trajectory to a nearby one. No matter how small these jumps are allowed to be one can fully migrate from any point of the phase space to another making the system completely unpredictable in the presence of any arbitrary small and any arbitrary infrequent perturbation. Chain recurrent sets capture perfectly this uncertainty by declaring the system may lie anywhere in its state space. This comes to a direct contrast with the common knowledge belief that zero-sum games are easy. Just because a Nash equilibrium is unique and efficiently computable does not mean that it fully captures the behavior of self-interested agents, especially when we target nonlinear properties of such systems, *i.e.*, the variance of the utilities of the agents. Technically, this analysis requires combining insights from the fundamental theorem of dynamical systems, Poincaré-Bendixson theorem, a powerful theorem for planar dynamical systems along with known connections between replicator dynamics and information theory and entropy (e.g., [12, 23]).

On the contrary, for potential games, it is already well known that replicator (and many other dynamics) converge to equilibria (e.g., [12, 16]). As the name suggests, these games have a potential function that strictly decreases as long as we are not at equilibrium. One could naturally come to hypothesize that in such systems, which are called gradient-like, the chain recurrent set must be equal to the set of equilibria. However, this is not the case.

Here is a counterexample that is due to Conley [8]. Imagine a continuous dynamical system on the $[0, 1]^2$ where all points on the boundary of this square are equilibria and all other trajectories flow straight downwards. That is, the limit behavior given any point (x, y) with $0 < x, y < 1$ is $(x, 0)$. This is a gradient-like system since the height y is always strictly decreasing unless we are at an equilibrium. Nevertheless, by allowing hops of size ϵ, for any $\epsilon > 0$, we can return from every point to itself. Starting from any (x, y) we move downwards along the flow, and when we get close enough to the boundary we hop on it on the point $(x, 0)$. Afterwards, we use these hops to traverse the boundary till we reach point $(x, 1)$ then one last last hop to $(x, 1 - \epsilon)$ places us on a trajectory that will revisit our starting point. Hence, once again, the whole state space is chain recurrent despite the fact that the system has a potential function. A memorable analogy for these systems is the game of snakes and ladders. The non-equilibrium states corresponds to the snakes where you move downwards alongside it, whereas the equilibrium states given these ϵ perturbations work as ladders that you can traverse upwards. Unlike the standard game of snakes and ladders there is no ending state in this game and you can keep going in circles indefinitely.

In the case of potential games we show that such contrived counterexamples cannot arise. The key components to this proof are techniques developed by Hurley [13], along with a beautiful, simplified and equivalent definition of chain recurrence, connecting chain recurrence of continuous flows and their discrete time maps, as well as tools from analysis (Sard's theorem) to establish that set of all possible potential values amongst all system equilibria is of zero measure. Combining these two, we establish that as we decrease the size of allowable perturbations eventually at the limit these perturbations do not allow for any new recurrent states to emerge and thus the set of equilibria and chain recurrent sets do coincide.

2. RELATED WORK

Studying dynamics in game theoretic settings is an area of research that is effectively as old as game theory itself. However, dating all the way back to the work of Brown and Robinson [6, 25] on dynamics in zero-sum games, the role of dynamics in game theory played a second fiddle to the main tune of Nash equilibria. The role of dynamics in game theory has predominantly been viewed as a way to provide validation and computational tools for equilibria.

At the same time, over the last fifty years and starting with Shapley [29] there has been an ever increasing trickle of counterexamples to the dominant Nash equilibrium paradigm that has been building up to a steady stream within the algorithmic game theory community ([10, 15, 17, 23, 19, 20] and references therein). Several analogous results are well known within evolutionary game theory [12, 26]. Such observations do not fit the current theoretical paradigm. The recent intractability results in terms of computing equilibria [11, 7] have provided an alternate, more formal tone to this growing discontent with the Nash solution concept, however, the key part is still missing. We need a general theory that fully encompasses all these special cases.

The definition of chain recurrent sets as well as a reference to the fundamental theorem of dynamical systems have been actually introduced in what is currently the definitive textbook reference of evolutionary game theory [26], however, the treatment is rather cursory, limited to abridged definitions and references to the dynamical systems literature. No intuition is built and no specific examples are discussed.

One recent paper [3] is quite related to the issues examined here. [3] establishes stability properties for chain recurrence sets in general dynamical systems. Applications to game theory are discussed. It is shown that there exist games where chain recurrent sets for learning dynamics can be more inclusive than Nash. However, no explicit characterization for these sets is established and all interior initial conditions actually converge to Nash. Connections between chain recurrence sets in potential games and equilibria are discussed without formal proofs. Our results provide complete characterizations of chain recurrent sets and reflect the realized behavior given generic initial states. We believe that chain recurrent sets is a central object of interest whose properties need to be carefully analyzed and we set up a series of analytical goals along these lines in the future work section.

3. PRELIMINARIES

3.1 Game Theory

We denote an n-agent game as $(n, \times_i S_i, \times u_i)$. Each agent chooses a strategy s_i from its set of available strategies S_i. Given a strategy profile $s = (s_1, \ldots, s_n)$, the payoff to each agent i is defined via its utility function $u_i : \times_i S_i \to \mathbb{R}$. Every potential game has a potential function $\Phi : \times_i S_i \to \mathbb{R}$, such that at any strategy profile s and for each agent i possible deviation from strategy s_i to s_i': $\Phi(s_i', s_{-i}) - \Phi(s_i, s_{-i}) = u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})$. Naturally the definitions of strategy and utility can be extended in the usual multilinear fashion to allow for randomized strategies. In that case, we usually overload notation in the following manner: if x_i a mixed strategy for each agent i, then we denote by $u_i(x)$ the expected utility of agent i, $\mathbf{E}_{s\sim x}[u_i(s)]$. We denote by x_{is_i} the probability that agent i assigns to strategy $s_i \in S_i$ in (mixed) strategy profile x. To simplify notation, sometimes instead of x_{is_i} we merely write x_{s_i}.

3.2 Replicator Dynamics

The replicator equation [30, 27] is described by:

$$\frac{dp_i(t)}{dt} = \dot{p}_i = p_i[u_i(p) - \hat{u}(p)], \quad \hat{u}(p) = \sum_{i=1}^{n} p_i u_i(p)$$

where p_i is the proportion of type i in the population, $p = (p_1, \ldots, p_m)$ is the vector of the distribution of types in the population, $u_i(p)$ is the fitness of type i, and $\hat{u}(p)$ is the average population fitness. The state vector p can also be interpreted as a randomized strategy. The replicator dynamic enjoys connections to classic models of ecological growth (*e.g.*, Lotka-Volterra equations [12]), as well as discrete time dynamics (*e.g.*, Multiplicative Weights algorithm) [16, 2, 18].

Remarks: In the context of game theory p_i will replaced with $x_{is_i}(t)$, *i.e.*, the probability that agent i plays strategy $s_i \in S_i$ at time t. Also, many times we will drop the explicit reference to time and just use x_{is_i}, or even just x_{s_i}.

3.3 Topology of dynamical systems

Our treatment follows that of [32], the standard text in evolutionary game theory, which itself borrows material from the classic book by Bhatia and Szegö [5]. Our chain recurrent set approach follows from [1].

Definition 1. A flow on a topological space X is a continuous function $\phi : \mathbb{R} \times X \to X$ such that

(i) $\phi(t, \cdot) : X \to X$ is a homeomorphism for each $t \in \mathbb{R}$.

(ii) $\phi(s+t, \mathbf{x}) = \phi(s, (\phi(t, \mathbf{x})))$ for all $s, t \in \mathbb{R}$ and all $x \in X$.

The second property is known as the group property of the flows. The topological space X is called the phase (or state) space of the flow.

Definition 2. Let X be a set. A map (or discrete dynamical system) is a function $f : X \to X$.

Typically, we write $\phi^t(x)$ for $\phi(t, x)$ and denote a flow $\phi : \mathbb{R} \times X \to X$ by $\phi^t : X \to X$, where the group property appears as $\phi^{t+s}(x) = \phi^s(\phi^t(x))$ for all $x \in X$ and $s, t \in \mathbb{R}$. Sometimes, depending on context, we use the notation ϕ^t to also signify the map $\phi(t, \cdot)$ for a fixed real number t. The map ϕ^1 is useful to relate the behavior of a flow to the behavior of a map.

Definition 3. If ϕ^t is a flow on a topological space X, then the function ϕ^1 defines the time-one map of ϕ^t.

Since our state space is compact and the replicator vector field is Lipschitz-continuous, we can present the unique solution of our ordinary differential equation by a flow $\phi : \mathbb{R} \times X \to X$, where X denotes the set of all mixed strategy profiles. Fixing starting point $x \in X$ defines a function of time which captures the trajectory (orbit, solution path) of the system with the given starting point. This corresponds to the graph of $\phi(\cdot, x) : \mathbb{R} \to X$, *i.e.*, the set $\{(t, y) : y = \phi(t, x) \text{ for some } t \in \mathbb{R}\}$.

If the starting point x does not correspond to an equilibrium then we wish to capture the asymptotic behavior of the system (informally the limit of $\phi(t, x)$ when t goes to infinity). Typically, however, such functions do not exhibit a unique limit point so instead we study the set of limits of all possible convergent subsequences. Formally, given a dynamical system $(\mathbb{R}, X, \phi)$ with flow $\phi : \mathbb{R} \times X \to X$ and a starting point $x \in X$, we call point $y \in X$ an ω-limit point of the orbit through x if there exists a sequence $(t_n)_{n\in\mathbb{N}} \in \mathbb{R}$ such that $\lim_{n\to\infty} t_n = \infty$, $\lim_{n\to\infty} \phi(t_n, x) = y$. Alternatively the ω-limit set can be defined as: $\omega_\Phi(x) = \cap_t \overline{\cup_{\tau \geq t} \phi(\tau, x)}$.

We denote the boundary of a set X as bd(X) and the interior of S as $\text{int}(X)$. In the case of replicator dynamics where the state space X corresponds to a product of agent (mixed) strategies we will denote by $\phi_i(t, x)$ the projection of the state on the simplex of mixed strategies of agent i. In our replicator system we embed our state space with the standard topology and the Euclidean distance metric.

3.4 Poincaré-Bendixson theorem

The Poincaré-Bendixson theorem is useful in proving the existence of periodic orbits and limit cycles[1] in two dimensional systems. The main idea is to find a trapping region, *i.e.*, a region from which trajectories cannot escape. If a trajectory enters and does not leave such a closed and bounded region of the state space that contains no equilibria then this trajectory must approach a periodic orbit as time goes to infinity. Formally, we have:

[1]A periodic orbit is called a limit cycle if it is the ω-limit set of some point not on the periodic orbit.

THEOREM 1. *[4, 31] Given a differentiable real dynamical system defined on an open subset of the plane, then every non-empty compact ω-limit set of an orbit, which contains only finitely many fixed points, is either a fixed point, a periodic orbit, or a connected set composed of a finite number of fixed points together with homoclinic and heteroclinic orbits connecting these.*

Homeomorphisms and Conjugacy of Flows

A function f between two topological spaces is called a *homeomorphism* if it has the following properties: f is a bijection, f is continuous, and f has a continuous inverse. A function f between two topological spaces is called a *diffeomorphism* if it has the following properties: f is a bijection, f is continuously differentiable, and f has a continuously differentiable inverse. Two flows $\Phi^t : A \to A$ and $\Psi^t : B \to B$ are conjugate if there exists a homeomorphism $g : A \to B$ such that for each $x \in A$ and $t \in \mathbb{R}$: $g(\Phi^t(x)) = \Psi^t(g(x))$. Furthermore, two flows $\Phi^t : A \to A$ and $\Psi^t : B \to B$ are *diffeomorhpic* if there exists a diffeomorphism $g : A \to B$ such that for each $x \in A$ and $t \in \mathbb{R}$ $g(\Phi^t(x)) = \Psi^t(g(x))$. If two flows are diffeomorphic then their vector fields are related by the derivative of the conjugacy. That is, we get precisely the same result that we would have obtained if we simply transformed the coordinates in their differential equations [21].

3.5 The fundamental theorem of dynamical systems

The standard formulation of the fundamental theorem of dynamical systems is built on the following set of definitions, based primarily on the work of Conley [8].

Definition 4. Let ϕ^t be a flow on a metric space (X, d). Given $\epsilon > 0$, $T > 0$, and $x, y \in X$, an (ϵ, T)-chain from x to y with respect to ϕ^t and d is a pair of finite sequences $x = x_0, x_1, \ldots, x_{n-1}, x_n = y$ in X and $t_0, \ldots, t_{n-1}$ in $[T, \infty)$, denoted together by $(x_0, \ldots, x_n; t_0, \ldots, t_{n-1})$ such that

$$d(\phi_i^t(x_i), x_{i+1}) < \epsilon$$

for $i = 0, 1, 2, \ldots, n-1$.

Definition 5. Let ϕ^t be a flow on a metric space (X, d). The forward chain limit set of $x \in X$ with respect to ϕ^t and d is the set

$$\Omega^+(x) = \bigcap_{\epsilon, T > 0} \{y \in X \mid \exists \text{ an } (\epsilon, T)\text{-chain from } x \text{ to } y\}.$$

Definition 6. Let ϕ^t be a flow on a metric space (X, d). Two points $x, y \in X$ are chain equivalent with respect to ϕ^t and d if $y \in \Omega^+(x)$ and $x \in \Omega^+(y)$.

Definition 7. Let ϕ^t be a flow on a metric space (X, d). A point $x \in X$ is chain recurrent with respect to ϕ^t and d if x is chain equivalent to itself. The set of all chain recurrent points of ϕ^t, denoted $\mathcal{R}(\phi)$, is the chain recurrent set of ϕ^t.

One key definition is the notion of a complete Lyapunov function. The game theoretic analogue of this idea is the notion of a potential function in potential games. In a potential game, as long as we are not at an equilibrium, the potential is strictly decreasing guiding the dynamics towards the standard game theoretic solution concept, *i.e.*, equilibria. The notion of a complete Lyapunov function switches the target solution concept from equilibria to chain recurrent points. More formally:

Definition 8. Let ϕ^t be a flow on a metric space X. A complete Lyapunov function for ϕ^t is a continuous function $\gamma : X \to \mathbb{R}$ such that

(i) $\gamma(\phi^t(x))$ is a strictly decreasing function of t for all $x \in X \setminus \mathcal{R}(\phi^t)$,

(ii) for all $x, y \in \mathcal{R}(\phi^t)$ the points x, y are chain equivalent with respect to ϕ^t if and only if $\gamma(x) = \gamma(y)$,

(iii) $\gamma(\mathcal{R}(\phi^t))$ is nowhere dense.

The powerful implication of the fundamental theorem of dynamical systems is that complete Lyapunov functions always exist. In game theoretic terms, every game is a "potential" game, if only we change our solution concept from equilibria to chain recurrent sets.

THEOREM 2. *[8] Every flow on a compact metric space has a complete Lyapunov function.*

Alternative, equivalent formulations of chain equivalence

For the purpose of our investigation, it will be useful to apply the following alternative definitions of chain equivalence, which are due to Hurley [13].

Definition 9. Let (X, d) be a metric space, and let $f : X \to X$. Given $\epsilon > 0$ and $x, y \in X$, an ϵ-chain from x to y is a finite sequence

$$x = x_0, x_1, \ldots, x_{n-1}, x_n = y$$

in X such that $d(f(x_i), x_{i+1}) < \epsilon$ for $i = 0, 1, 2, \ldots, n-1$.

Definition 10. Let X be a metric space, and let $f : X \to X$. Two points $x, y \in X$ are called chain equivalent if for every $\epsilon > 0$ there exists an ϵ-chain from x to y and there exists an ϵ-chain from y to x.

Next, we provide three alternative formulations of chain equivalence which are equivalent with our original definition for a flow on a compact metric space as shown in [13].

THEOREM 3. *[13] If ϕ^t is a flow on a compact metric space (X, d) and $x, y \in X$, then the following statements are equivalent.*

1. *(i) The points x and y are chain equivalent with respect to ϕ^t.*
2. *(ii) For every $\epsilon > 0$ and $T > 0$ there exists an $(\epsilon, 1)$-chain*

$$(x_0, \ldots, x_n; t_0, \ldots, t_n - 1)$$

from x to y such that

$$t_0 + \cdots + t_{n-1} \geq T$$

and there exists an $(\epsilon, 1)$-chain

$$(y_0, \ldots, y_m; s_0, \ldots, s_m - 1)$$

from y to x such that

$$s_0 + \cdots + s_{m-1} \geq T.$$

3. (iii) *For every $\epsilon > 0$ there exists an $(\epsilon, 1)$-chain from x to y and an $(\epsilon, 1)$-chain from y to x.*

4. (iv) *The points x and y are chain equivalent with respect to ϕ^1.*

4. CHAIN RECURRENT SETS FOR MATCHING PENNIES

Zero-sum games are amongst the most well studied class of games within game theory. Equilibria here are classically considered to completely "solve" the setting. This is due to the fact that the equilibrium prediction is essentially unique, Nash computation is tractable, and many natural classes of learning dynamics are known to "converge weakly" to the set of Nash equilibria.

The notion of weak convergence encodes that the time average of the dynamics converge to the equilibrium set. However, this linguistic overloading of the notion of convergence is unnatural and arguably can lead to a misleading sense of certainty about the complexity that learning dynamics may exhibit in this setting. For example, would it be meaningful to state that the moon "converges weakly" to the earth instead of stating that e.g. the moon follows a trajectory that has earth at its center?

The complexity and unpredictability of the actual behavior of dynamics becomes apparent when we characterize the set of chain recurrent points even for the simplest zero-sum games, Matching Pennies. Despite the uniqueness and symmetry of the Nash equilibrium it is shown to not capture fully the actual dynamics. The set of chain recurrent points is the whole strategy space. This means that in the presence of arbitrary small noise the replicator dynamics can become completely unpredictable. Even in an idealized implementation without noise there exist absolutely no initial conditions that converge to a Nash equilibrium. To argue this we will use the following known lemma, whose proof we provide for completeness in the appendix.

LEMMA 1. *[12, 24] Let ϕ denote the flow of the replicator dynamic when applied to a zero sum game with a fully mixed Nash equilibrium $q = (q_1, q_2)$ then given any starting point $x(0) = (x_1(0), x_2(0))$ then the sum of the KL-divergences between each agent's mixed Nash equilibrium q_i and his evolving strategy $x_i(t)$ is time invariant. Equivalently, $D_{\mathrm{KL}}(q_1\|x_1(t)) + D_{\mathrm{KL}}(q_2\|x_2(t)) = D_{\mathrm{KL}}(q_1\|x_1(0)) + D_{\mathrm{KL}}(q_2\|x_2(0))$ for all t.*

The main theorem of this section is the following:

THEOREM 4. *Let ϕ denote the flow of the replicator dynamic when applied to Matching Pennies then the set of chain recurrent points is the whole state space. This characterization holds for all 2x2 zero-sum games that have a fully mixed Nash equilibrium.*

PROOF. The proof proceeds in two steps. In the first step, which follows along the lines of [23], we establish that every interior point of the state space, i.e., any fully mixed strategy profile, lies on a periodic orbit. The second part of the proof establishes that the geometry of the trajectories implies that for any point in the state space and for any $\epsilon > 0$ we can find an ϵ-orbit connecting this point to itself.

In order to establish periodicity of all orbits starting from a fully mixed initial condition we need to apply the Poincaré-Bendixson theorem. The Poincaré-Bendixson theorem can only be applied for two dimensional systems, whereas our system has four variables $(x_{1Heads}, x_{1Tails}, x_{2Heads}, x_{2Tails})$. Since $x_{1\mathrm{Heads}} + x_{1Tails} = 1$ and $x_{2\mathrm{Heads}} + x_{2Tails} = 1$ there exists a natural projection to the plane (x_{1Heads}, x_{2Heads}). The two flows are conjugate and periodicity is a topological invariant property [1], so we can identify all periodic points of all four variable system by lifting up all the periodic points of the flow on the (x_{1Heads}, x_{2Heads}) plane.

We know from lemma 2 that in zero-sum games with fully mixed equilibria, and hence in Matching Pennies, the sum of the KL-divergences between the Nash equilibrium strategies and the time evolving strategies of each agent remains constant as we move along the trajectories of the replicator. KL-divergence is a (pseudo)-metric implying the existence of trapping regions in the interior of our state space. Specifically, as long as we start from an interior point other than the unique Nash then the trajectory stays bounded away from the boundary (KL-divergence becomes infinite) and from the unique equilibrium (KL-divergence becomes zero). Naturally, these restrictions also apply to the conjugate flow via the homeomorphic projection. By the Poincaré-Bendixson theorem we have that starting from any point (other than the Nash) the resulting limit set is a periodic orbit. It is straightforward to check that the KL-divergence invariance condition when projected to our subspace translates to an invariance of the quantity $x_{1Head}(1 - x_{1Head})x_{2Head}(1 - x_{2Head})$. This defines a closed continuous curve on our subspace that is symmetric along the axis $x_{1Head} = 1/2$ and $x_{2Head} = 1/2$. In each of the four regions defined by $x_{1Head} = 1/2$, $x_{2Head} = 1/2$ the signs of $\frac{dx_{1Head}}{dt}$, $\frac{dx_{2Head}}{dt}$ are fixed and define a clockwise direction. The periodicity of the interior points now follows from the uniqueness of the solution of the projected replicator flow.

The case of non-interior points is simpler. Naturally each of the four "pure" initial conditions $(H, H)(H, T)(T, H)(T, T)$ are fixed points for the replicator dynamics. Finally, initial conditions where exactly one agent applies a pure strategy lie on an orbit connection two pure states. For example, if the first agent chooses $Heads$ with probability one, then as the second agent moves forward in time he will continuously increase his probability of playing strategy $Tails$. Thus, all these points lie on an orbit connecting (H, H) to (H, T). The geometry of these orbits is depicted in Figure 1.

The fact that all states are chain recurrent follows easily. All interior states are chain recurrent since they are periodic. All pure states are chain recurrent since they are equilibria. Let x be any of the remaining states with exactly one randomizing agent. For any $\epsilon > 0$ we can create a $(\epsilon, 1)$-orbit which travels exclusively along the boundary of the state space and as it approaches a fixed point, uses a jump of size at most ϵ to hop onto the next best response ray. Thus, $x \in \Omega^+(x)$ and these states are chain recurrent as well. □

5. CHAIN RECURRENT SETS FOR POTENTIAL GAMES

For gradient systems, it is understood that under sufficient smoothness conditions the set of chain recurrent points coincides with the set of equilibria. Many learning dynamics define gradient-like behavior in potential games (i.e., the potential function always increases), but without being formally a gradient (moving in the direction of maximum increase). We argue that under replicator dynamics the chain

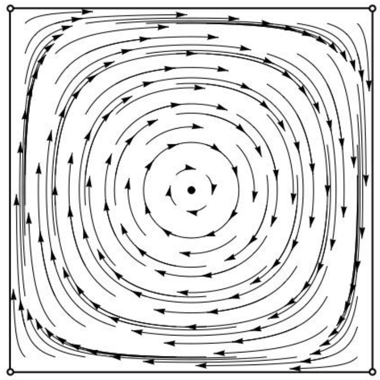

Figure 1: Replicator trajectories in Matching Pennies game. Each point encodes the probability assigned by the agents to their first strategy.

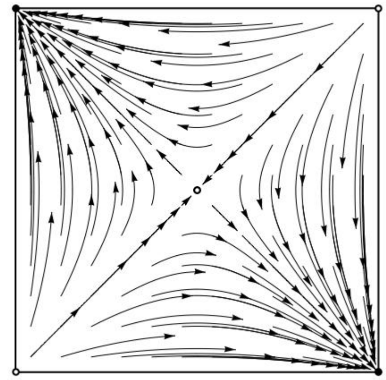

Figure 2: Replicator trajectories in a partnership, coordination game. Each point encodes the probability assigned by the agents to their first strategy.

recurrent sets for any potential games coincides with the set of system equilibria. The set of chain recurrent points of gradient-like systems can be rather complicated and Conley [8] constructs a specific example of a gradient-like system with a zero measurable set of equilibria where the set of chain recurrent points is the whole state space. Nevertheless, we establish that such bad examples do not arise in potential games. For the proof of this characterization we will apply the following theorem due to Hurley that we have already discussed in the preliminaries:

THEOREM 5. *[13] The chain recurrent set of a continuous (semi)flow on an arbitrary metric space is the same as the chain recurrent set of its time-one map.*

THEOREM 6. *Let ϕ denote the flow of the replicator dynamic when applied to a potential game then the set of chain recurrent points coincides with its set of equilibria.*

PROOF. Replicator dynamics defines a gradient-like system, where the (expected) value of the potential function always increases unless we are at a fixed point. Specifically, it is well known that in any potential game the utility of any agent at a state s, $u_i(s)$ can be expressed as a summation of the potential Φ and a dummy term $D_i(s_{-i})$ that depends on the strategies of all agents other than i. Indeed, by the definition of the potential game for any possible deviation of agent i from strategy s_i to s_i': $\Phi(s_i', s_{-i}) - \Phi(s_i, s_{-i}) = u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})$ and hence for each $s_{-i} = \times_{j \neq i} s_j$ and any two possible strategies s_i, s_i' of agent i we have that $u_i(s_i', s_{-i}) - \Phi(s_i', s_{-i}) = u_i(s_i, s_{-i}) - \Phi(s_i, s_{-i})$. Hence these differences are independent of the choice of strategy of agent i and can be expressed as $D_i(s_{-i})$, a function of the choices of all other agents. We can now express $u_i(s) = \Phi(s) + D_i(s_{-i})$ and similarly $u_i(x) = \Phi(x) + D_i(x_{-i})$ for mixed strategy profiles. Furthermore, we have that since $\Phi(x) = \sum_{s_i \in S_i} x_{s_i} \Phi(s_i, x_{-i})$, we have that for each $s_i \in S_i$ $\frac{\partial \Phi(x)}{\partial x_{s_i}} = \Phi(s_i, x_{-i})$. Therefore, we have that:

$$
\begin{aligned}
\dot{\Phi}(x) &= \sum_i \sum_{s_i \in S_i} \frac{\partial \Phi(x)}{\partial x_{s_i}} \dot{x}_{s_i} = \sum_i \sum_{s_i \in S_i} \Phi(s_i, x_{-i}) \dot{x}_{s_i} \\
&= \sum_i \sum_{s_i \in S_i} \Phi(s_i, x_{-i}) x_{s_i} [u_i(s_i, x_{-i}) - u_i(\mathbf{x})] \\
&= \sum_i \sum_{s_i \in S_i} \Phi(s_i, x_{-i}) x_{s_i} [\Phi(s_i, x_{-i}) - \Phi(\mathbf{x})] \\
&= \sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} [\Phi(s_i, x_{-i})^2 - \Phi(s_i, x_{-i}) \Phi(s_i', x_{-i})] \\
&= \frac{1}{2} \sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} [\Phi(s_i, x_{-i}) - \Phi(s_i', x_{-i})]^2 \qquad (1) \\
&= \frac{1}{2} \sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} \sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} [\Phi(s_i, x_{-i}) - \Phi(s_i', x_{-i})]^2 \\
&\geq \frac{1}{2} \Big[\sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} |\Phi(s_i, x_{-i}) - \Phi(s_i', x_{-i})| \Big]^2 \\
&= \frac{1}{2} \Big[\sum_{i, s_i, s_i' \in S_i} x_{s_i} x_{s_i'} |u_i(s_i, x_{-i}) - u_i(s_i', x_{-i})| \Big]^2 \\
&\geq \frac{1}{2} \Big[\sum_{i, s_i, s_i' \in S_i} x_{s_i} |u_i(s_i, x_{-i}) - u_i(x)| \Big]^2 = \frac{1}{2} \|\xi\|_1^2
\end{aligned}
$$

where ξ expresses the replicator vector field. Although the convergence of the replicator dynamics in congestion and hence potential games can be derived by an analogous construction in [16] here we will also use this proof to establish that the set of all potential values attained at equilibrium points, *i.e.*, $\mathcal{V} = \{\Phi(x), x \text{ is an equilibrium}\}$ is of measure zero.

We will argue this by showing that it can be expressed as the finite union of zero measure sets. By the above derivation we have that the potential is strictly decreasing unless we are at a system equilibrium. Furthermore, its establishes that in the places where the potential does not increase, *i.e.*, at equilibrium, we have that for all agents i if $x_{s_i}, x_{s_i'} > 0$, then $\Phi(s_i, x_{-i}) - \Phi(s_i', x_{-i}) = 0$. However, this immediately implies that $\frac{\partial\Phi(x)}{\partial x_{s_i}} - \frac{\partial\Phi(x)}{\partial x_{s'_i}} = 0$. Any equilibrium, either corresponds to a pure state, in which cases the union of their potential values is trivially of zero measure or its corresponds to a mixed state where one or more agents is randomizing. In order to account for the possibility of continuums of equilibria, we will use Sard's theorem that implies that the set of critical values (that is, the image of the set of critical points[2]) of a smooth function from one Euclidean space to another is a null set, i.e., it has Lebesgue measure 0. We define an arbitrary fixed ordering over the strategy set of each agent. Given any mixed system equilibrium x, the expected value of the potential $\Phi(x)$ can be written as a multi-variate polynomial over all strategy variables x_{s_i} played with strictly positive probability. Since the x_{s_i}'s represent probabilities we can replace the lexicographically smaller variable x_{s_i} as one minus the summation of all other variables in the support of the current mixed strategy of agent i. Now, however, all partial derivatives of this polynomial at the equilibrium are equal to zero. Hence, each equilibrium can be expressed as a critical point of a smooth function from some space $\mathbb{R}^k$ to $\mathbb{R}$ and hence its image (i.e. its set of potential values) is a zero measure subset of $\mathbb{R}$. It is clear that the set of polynomials needed to capture all equilibria depends only on the choice of strategies for each agent that are played with positive probability and hence although they are exponential many they are finite. Putting everything together the set of all potential values attained at equilibria is of zero measure.

Naturally, however, the complement of equilibrium values, which we denote $\mathcal{C}$ is dense in the set[3] $[\min_s \Phi(s), \max_s \Phi(s)]$. Indeed, if that is not the case, *i.e.*, if $\mathcal{C}$ is not dense in $[\min_s \Phi(s), \max_s \Phi(s)]$ then there exists a point $y \in [\min_s \Phi(s), \max_s \Phi(s)]$ such that $y \notin \mathcal{C}$ and at the same time y is not an accumulation point of $\mathcal{C}$. This implies that there exists a neighborhood of y that contains no points of $\mathcal{C}$. We reach contradiction since $[\min_s \Phi(s), \max_s \Phi(s)] \setminus \mathcal{C}$ is of zero measure.

Next, we will apply the fact that the complement of equilibrium values of the potential is dense in the set $[\min_s \Phi(s), \max_s \Phi(s)]$ to establish that the chain recurrent points of the time one map ϕ^1 of the flow coincide with the set of equilibria. As we stated above, Hurley [13] has shown that the chain recurrent points of the flow coincide with those of its time one map and hence the theorem follows.

We have that $\Phi(\phi^1(x)) < \Phi(x)$ for all x, with equality if and only if we are at equilibrium. Suppose that we choose a regular value r of the potential, *i.e.* a value that does not correspond to a fixed point. Let's consider the sets $K_r = \Phi^{-1}((-\infty, r])$, and $U_r = \Phi^{-1}((-\infty, r))$. Note that K_r is closed while U_r is open (in the topology defined by the set of strategy profiles) and contained in K_r. If $\Phi(p) = r$, then $\Phi(\phi^1(p)) < \Phi(p) = r$. This means that $\phi^1(K_r) \subset U_r$. However, since ϕ^1, the time one map of the flow is a homeomorphism, the fact that K_r is closed yields that $\overline{\phi^1(U_r)} \subset \overline{\phi^1(K_r)} = \phi^1(K_r) \subset U_r$.

Any chain recurrent point whose forward orbits meets U_r is furthermore contained in $\overline{\phi^1(U_r)}$ [13, 8]. Let q be a non-equilibrium point, then we have that $\Phi(\phi^1(q)) < \Phi(q)$. However, due to the fact that the images of the potential values of non-equilibrium points are dense in $(\min_s \Phi(s), \max_s \Phi(s))$ we can choose such a value r such that $\Phi(\phi^1(q)) < r < \Phi(q)$. Then $\phi^1(q) \in U_r$ but $q \notin U_r$ and thus q is not a chain recurrent point of ϕ^1. As such it cannot be a chain recurrent point for the replicator flow as well. The theorem follows immediately since all equilibria are trivially chain recurrent points. $\square$

[2]For a differentiable function of several real variables, a critical point is a value in its domain where all partial derivatives are zero.

[3]For simplicity we disregard the trivial potential games where the potential is everywhere constant $\min_s \Phi(s) = \max_s \Phi(s)$. After all in those games are states are equilibria and the theorem is trivially true.

6. FUTURE WORK

The purpose of this paper is to introduce the fundamental theorem of dynamical systems, and to sketch and motivate its use towards developing a principled, rigorous, and informative discourse on game-theoretic solution concepts. We focused on two simple and evocative examples, namely zero-sum games with fully mixed equilibria and potential games. Naturally, there is much that needs to be done, and below we sample a few research goals that are immediate, important, and each open ended in its own way.

- **The structure of Chain Recurrent Sets (CRSs).** A game may have many CRSs (for example, the coordination game in Figure 2 has five). It is not hard to see that chain equivalence defines an equivalence relation that partitions the set of CRSs to equivalence classes. These classes are called chain components and can be arranged as vertices of a directed acyclic graph, where directed edges signify possible transitions after an infinitesimal jump; for the coordination game in Figure 2 this DAG has two sinks (the pure Nash equilibria), two sources (the other two pure profiles), and a node of degree 4 (the mixed Nash equilibrium). Identifying this DAG is tantamount to analyzing the game, the generalization of finding its Nash equilibria. Understanding the fundamental structure in games of interest is an important challenge.

- **Price of Anarchy through CRSs.** We can define a natural distribution over the sink CRSs of a game, namely, assign to each sink CRS the probability that a trajectory started at a (say, uniformly) random point of the state space will end up, perhaps after infinitesimal jumps, at the CRS. This distribution, together with the CRS's expected utility, yield a new and productive definition of the average price of anarchy in a game, as well as a methodology for calculating it, see for example [22].

- **Inside a CRS.** Equilibria and limit cycles are the simplest forms of a chain recurrent component, in the sense that no "jumps" are necessary for going from one state in the component to another. In Matching Pennies, in contrast, $O(\frac{1}{\epsilon})$ many ϵ-jumps are needed to reach the Nash equilibrium, starting from a pure strategy profile. What is the possible range of this form of complexity of a CRS?

- **Complexity.** There are several intriguing complexity questions posed by this concept. What is the complexity of determining, given a game and two strategy profiles, whether they belong to the same component? What is the complexity of finding a point in a sink chain component? What is the speed of convergence to a CRS?

7. ACKNOWLEDGMENTS

This work was supported by NSF grant CCF1408635 and SUTD grant SRG ESD 2015 097.

8. REFERENCES

[1] J. M. Alongi and G. S. Nelson. *Recurrence and topology*, volume 85. American Mathematical Soc., 2007.

[2] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta algorithm and applications. Technical report, 2005.

[3] M. Benaïm, J. Hofbauer, and S. Sorin. Perturbations of set-valued dynamical systems, with applications to game theory. *Dynamic Games and Applications*, 2(2):195–205, 2012.

[4] I. Bendixson. Sur les courbes définies par des équations différentielles. *Acta Mathematica*, 24(1):1–88, 1901.

[5] N. P. Bhatia and G. P. Szegö. *Stability theory of dynamical systems*. Berliin: Springer Verlag, 1970.

[6] G. Brown. Iterative solutions of games by fictitious play. *In Activity Analysis of Production and Allocation, T.C. Koopmans (Ed.), New York: Wiley.*, 1951.

[7] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3):14:1–14:57, May 2009.

[8] C. Conley. Isolated invariant sets and the morse index. cbms regional conference series in mathematics, 38. *American Mathematical Society, Providence, RI*, 16, 1978.

[9] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.

[10] C. Daskalakis, R. Frongillo, C. Papadimitriou, G. Pierrakos, and G. Valiant. On learning algorithms for Nash equilibria. *Algorithmic Game Theory*, pages 114–125, 2010.

[11] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. pages 71–78. ACM Press, 2006.

[12] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, 1998.

[13] M. Hurley. Chain recurrence, semiflows, and gradients. *Journal of Dynamics and Differential Equations*, 7(3):437–456, 1995.

[14] G. P. Karev. Replicator equations and the principle of minimal production of information. *Bull. Math. Biol.*, 72(1124), 2010.

[15] R. Kleinberg, K. Ligett, G. Piliouras, and É. Tardos. Beyond the Nash equilibrium barrier. In *Symposium on Innovations in Computer Science (ICS)*, 2011.

[16] R. Kleinberg, G. Piliouras, and É. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *ACM Symposium on Theory of Computing (STOC)*, 2009.

[17] K. Ligett and G. Piliouras. Beating the best Nash without regret. *SIGecom Exchanges 10*, 2011.

[18] V. Losert and E. Akin. Dynamics of games and genes: Discrete versus continuous time. *Journal of Mathematical Biology*, 1983.

[19] R. Mehta, I. Panageas, and G. Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics [working paper abstract]. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, ITCS '15, pages 73–73, New York, NY, USA, 2015. ACM.

[20] R. Mehta, I. Panageas, G. Piliouras, and S. Yazdanbod. The Complexity of Genetic Diversity: Sex with Two Chromosomes is Advantageous but Unpredictable. *ArXiv e-prints*, Nov. 2014.

[21] J. Meiss. *Differential Dynamical Systems*. SIAM, 2007.

[22] I. Panageas and G. Piliouras. Approximating the geometry of dynamics in potential games: Point-wise convergence, regions of attraction, average case performance analysis and system invariants. *http://arxiv.org/abs/1403.3885*, 2014.

[23] G. Piliouras, C. Nieto-Granda, H. I. Christensen, and J. S. Shamma. Persistent patterns: Multi-agent learning beyond equilibrium and utility. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 181–188, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[24] G. Piliouras and J. S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence. In *SODA*, 2014.

[25] J. Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54:296–301, 1951.

[26] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.

[27] P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100(3):533 – 538, 1983.

[28] C. R. Shalizi. Dynamics of bayesian updating with dependent data and misspecified models. *Electronic Journal of Statistics*, 3:1039–1074, 2009.

[29] L. Shapley. Some topics in two-person games. *Advances in game theory, edited by M. Dresher, R. J. Aumann, L. S. Shapley, A. W. Tucker*, 1964.

[30] P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(12):145–156, 1978.

[31] G. Teschl. *Ordinary differential equations and dynamical systems*, volume 140. American Mathematical Soc., 2012.

[32] J. Weibull. *Evolutionary Game Theory*. MIT Press; Cambridge, MA: Cambridge University Press., 1995.

APPENDIX

A. INFORMATION THEORY

Entropy is a measure of the uncertainty of a random variable and captures the expected information value from a measurement of the random variable. The entropy H of a discrete random variable X with possible values $\{1, \ldots, n\}$ and probability mass function $p(X)$ is defined as $H(X) = -\sum_{i=1}^{n} p(i) \ln p(i)$.

Given two probability distributions p and q of a discrete random variable their *K-L divergence* (relative entropy) is defined as $D_{\mathrm{KL}}(p\|q) = \sum_i \ln\left(\frac{p(i)}{q(i)}\right) p(i)$. It is the average of the logarithmic difference between the probabilities p and q, where the average is taken using the probabilities p. The K-L divergence is only defined if $q(i) = 0$ implies $p(i) = 0$ for all i[4]. K-L divergence is a "pseudo-metric" in the sense that for it is always non-negative and is equal to zero if and only if the two distributions are equal (almost everywhere). Other useful properties of the K-L divergence is that it is additive for independent distributions and that it is jointly convex in both of its arguments; that is, if (p_1, q_1) and (p_2, q_2) are two pairs of distributions then for any $0 \leq \lambda \leq 1$: $D_{\mathrm{KL}}(\lambda p_1 + (1-\lambda)p_2 \| \lambda q_1 + (1-\lambda)q_2) \leq \lambda D_{\mathrm{KL}}(p_1\|q_1) + (1-\lambda) D_{\mathrm{KL}}(p_2\|q_2)$.

A closely related concept is that of the *cross entropy* between two probability distributions, which measures the average number of bits needed to identify an event from a set of possibilities, if a coding scheme is used based on a given probability distribution q, rather than the "true" distribution p. Formally, the cross entropy for two distributions p and q over the same probability space is defined as follows: $H(p, q) = -\sum_{i=1}^{n} p(i) \ln q(i) = \mathrm{H}(p) + D_{\mathrm{KL}}(p\|q)$. For more details and proofs of these basic facts the reader should refer to the classic text by Cover and Thomas [9].

[4]The quantity $0 \ln 0$ is interpreted as zero because $\lim_{x \to 0} x \ln(x) = 0$.

B. MISSING PROOFS OF SECTION 4

LEMMA 2. *[12, 24] Let ϕ denote the flow of the replicator dynamic when applied to a zero sum game with a fully mixed Nash equilibrium $q = (q_1, q_2)$. Given any starting point $x(0) = (x_1(0), x_2(0))$ then the sum of the KL-divergences between each agent's mixed Nash equilibrium q_i and his evolving strategy $x_i(t)$ is time invariant. Equivalently, $D_{\mathrm{KL}}(q_1\|x_1(t)) + D_{\mathrm{KL}}(q_2\|x_2(t)) = D_{\mathrm{KL}}(q_1\|x_1(0)) + D_{\mathrm{KL}}(q_2\|x_2(0))$ for all t.*

PROOF. It suffices to establish that the time derivative of $D_{\mathrm{KL}}(q_1\|x_1(t)) + D_{\mathrm{KL}}(q_2\|x_2(t))$ is everywhere equal to zero. By properties of the KL-divergence, it suffices to show that the time derivative of the quantity $\sum_i H(q_i, \phi_i(t, x_0)) = -\sum_i \sum_{s_i \in S_i} q_{s_i} \cdot \ln(x_{s_i})$ is everywhere zero, where $i \in \{1, 2\}$ and S_i the available strategies of agent i. We denote by $A^{1,2}$ the payoff matrix of agent 1 and $A^{2,1}$ the payoff matrix of agent 2. Since this is a zero-sum game: $A^{1,2} + (A^{2,1})^T = \mathbf{0}$ where $(A^{2,1})^T$ the transpose of $A^{2,1}$.

$$
\begin{aligned}
&\sum_i \sum_{s_i \in S_i} q_{s_i} \frac{d \ln(x_{s_i})}{dt} = \sum_i \sum_{s_i \in S_i} q_{s_i} \frac{\dot{x}_{s_i}}{x_{s_i}} = \\
&= \sum_i \big(\sum_{s_i \in S_i} q_{s_i} u_i(s_i, x_{-i}) - \sum_{s_i \in S_i} x_{s_i} u_i(s_i, x_{-i}) \big) = \\
&= \big(q_1^{\mathrm{T}} A^{1,2} x_2 - x_1^{\mathrm{T}} A^{1,2} x_2\big) + \big(q_2^{\mathrm{T}} A^{2,1} x_1 - x_2^{\mathrm{T}} A^{2,1} x_1\big) = \\
&= \big(q_1^{\mathrm{T}} - x_1^{\mathrm{T}}\big) A^{1,2} (x_2 - q_2) + \big(q_2^{\mathrm{T}} - x_2^{\mathrm{T}}\big) A^{2,1} (x_1 - q_1) = \\
&= -\big(q_1^{\mathrm{T}} - x_1^{\mathrm{T}}\big) \Big[A^{1,2} + (A^{2,1})^T\Big] (q_2 - x_2) = 0
\end{aligned}
$$

□