Proceedings of Machine Learning Research vol 201:1–22, 2023 34th International Conference on Algorithmic Learning Theory

The Replicator Dynamic, Chain Components and the Response Graph

Oliver Biggar OLIVER.BIGGAR@ANU.EDU.AU

Iman Shames IMAN.SHAMES@ANU.EDU.AU

CIICADA Lab, Australian National University, Canberra, 2601, Australia

Editors: Shipra Agrawal and Francesco Orabona

Abstract

In this paper we examine the relationship between the flow of the replicator dynamic, the continuum limit of Multiplicative Weights Update, and a game's *response graph*. We settle an open problem establishing that under the replicator, *sink chain components*—a topological notion of long-run outcome of a dynamical system—always exist and are approximated by the *sink connected components* of the game's response graph. More specifically, each sink chain component contains a sink connected component of the response graph, as well as all mixed strategy profiles whose support consists of pure profiles in the same connected component, a set we call the *content* of the connected component. As a corollary, all profiles are chain recurrent in games with strongly connected response graphs. In any two-player game sharing a response graph with a zero-sum game, the sink chain component is unique. In two-player zero-sum and potential games the sink chain components and sink connected components are in a one-to-one correspondence, and we conjecture that this holds in all games.

1. Introduction

When a collection of players simultaneously learn a game, what strategies do they eventually learn to play? This is the fundamental question of evolutionary game theory, and by extension is of critical importance in economics (Sandholm, 2010), biology (Smith and Price, 1973), and computer science (Roughgarden, 2010). In particular, this question lies at the heart of *multi-agent learning* (Yang and Wang, 2020) which itself is at the core of recent breakthroughs in AI (Silver et al., 2016, 2017, 2018).

Unpacking this question in more depth, we find three key observations. The first: the question assumes a fixed learning or evolution algorithm, which in evolutionary game theory is called a *dynamic*. When paired with the strategy space, we obtain a *dynamical system* (Strogatz, 2018). We expect dynamical systems theory to play a role in the solution. The second: as the word 'eventually' suggests, we are interested in *long-run* or asymptotic analysis of the dynamical system. The third: learning is a computational process, and so our analysis must be compatible with computational feasibility.

Despite our game-theoretic intuition, *Nash equilibria* (Nash, 1951) are generally not the answer to our question. A series of results in evolutionary game theory (Kleinberg et al., 2011; Sandholm, 2010; Papadimitriou and Piliouras, 2019) have shown that standard game dynamics generally do not converge to (mixed) Nash equilibria (Vlatakis-Gkaragkounis et al., 2020)[1], even in zero-sum games (Mertikopoulos et al., 2018). In fact, no choice of

1. By contrast, many dynamics do converge to *pure* Nash equilibria.

dynamic can converge to Nash equilibria in all games (Hart and Mas-Colell, 2003; Benaïm et al., 2012; Milionis et al., 2022). From a computational perspective, there are also problems: Nash equilibria are PPAD-complete to compute (Daskalakis et al., 2009), and so our players cannot be expected to feasibly learn them.

If the Nash equilibrium is not the right answer, we must determine what *is*. Using our key ideas, we set out three criteria for a solution concept, inspired by Fabrikant and Papadimitriou (2008):

1. (**Convergence**) Almost all points should converge to the solution concept, and should remain there once reached.

2. (**Existence**) The solution concept should exist in all games.

3. (**Computability**) The solution concept should be efficiently computable.

Finally, to avoid trivial solutions[2], we desire the *minimal* set of strategy profiles satisfying these properties—**Existence**, in particular. A priori, these criteria—particularly the first!—seem very optimistic. We restrict our attention to the best-studied continuous-time dynamic, the *replicator* (Sandholm, 2010). The replicator emerged originally from population models in biology, and is the continuum limit of the *Multiplicative Weights Update* algorithm (Arora et al., 2012). Yet even this well-studied and comparatively well-behaved dynamic's convergence property is notoriously difficult to understand, with general results only known for a few classes of games (such as zero-sum and potential games, as in Mertikopoulos et al.).

But hope remains, if we use the right mathematical tools. Papadimitriou and Piliouras (2019, 2018, 2019) recently proposed a new solution concept meeting the **Convergence** criterion, inspired by the concept of *chain recurrence*, a dynamical systems concept underlying the Fundamental Theorem of Dynamical Systems (Conley, 1978). To understand this concept, we first note that our **Convergence** criterion breaks into two parts, (1) converging to the solution concept, and (2) remaining there once reached. An example of a point satisfying the latter is a *stationary point*, a point which remains in place under the dynamic. A more general notion of 'staying in place' is a *periodic* point, one which returns to itself repeatedly. A yet more general concept is a *recurrent* point, one which returns arbitrarily close to itself infinitely often. A major historical challenge of dynamical systems was to find the appropriate generalisation of a recurrent point, such that all points 'end up at' these stationary ones (Alongi and Nelson, 2007). This was achieved by Conley (1978), who introduced the notion of *chain recurrent* points.

Conley's insight was to use a 'noisy' generalisation of an orbit of the system. He introduced the concept of an (ϵ, T)-chain, which is an orbit which allows for a finite number of tiny 'jumps' of size at most ϵ. To ensure we cannot jump 'too often', each jump must be separated by a time of at least $T > 0$. If there exists an (ϵ, T)-chain from x to y for *any* $\epsilon > 0$ and $T > 0$, we say there is a *pseudo-orbit* from x to y. A point x is *chain recurrent* if there is a pseudo-orbit from x to itself. Chain recurrent points form connected components of the space, called *chain components*. The Fundamental Theorem of Dynamical

2. Taking the entire strategy space of the game as the solution concept satisfies these criteria, but it is not a particularly enlightening solution.

Systems (Conley, 1978) establishes that all points converge to chain components. When we view the dynamical system from the perspective of pseudo-orbits rather than 'true' orbits, we obtain a partial ordering on chain components. The minimal elements in this order are called *sink chain components*. Points in these components satisfy both parts of the **Convergence** criterion: they are chain recurrent, so they remain in the component once reached, and because they are *sink* chain components other points end up there under pseudo-orbits.

The definition of chain recurrence is fundamentally *computational*. Any computational device with arbitrarily large but finite precision cannot distinguish between a 'true' orbit of the dynamical system and a pseudo-orbit. In fact, many common game dynamics exhibit chaotic behaviour, demonstrating the difficulty in tracking 'true' orbits (Cheung and Piliouras, 2019, 2020; Andrade et al., 2021). Consequently, any solution concept which is 'computational' in this sense cannot distinguish between points in the same chain component, making sink chain components the minimal set satisfying our **Convergence** and **Computability** criteria.

Unfortunately, despite their computational inspiration, actually computing the sink chain components is not obviously feasible (failing **Computability**). There is a second problem: the order defined by pseudo-orbits is infinite in general, and so sink chain components need not exist (failing **Existence**). Papadimitriou and Piliouras (2019) solve the second problem by assumption: they conjecture that *in game dynamics, sink chain components always exist*, which we prove for the replicator in this paper. They address the first problem by suggesting a computable surrogate of sink chain components: the *sink connected components* of the game's *response graph*.

The response graph is a directed graph defined on the pure profiles of the game. There is an arc between profiles if they differ in the strategy of a single player, with the arcs directed toward the preferred profile for that player. The response graph can be thought of as underlying combinatorial structure of the game, capturing precisely the order in which each player prefers their strategies given fixed choices of strategy for each other player. The response graph is *structurally stable*; changing the payoffs of the game in a small way generally does not affect it. Response graphs also underlie *ordinal games* (Mertens, 2004) and *strategic games* (Candogan et al., 2011), and so are a unifying model of the game structure. In spite of their generality, response graphs store many structural features of a game (Biggar and Shames, 2022). As a solution concept, sink connected components can be easily computed by traversing the response graph (**Computability**), and always exist because the connected components of a graph is a *finite* partial order (**Existence**). Restricted to the edges of the graph, the replicator dynamic flows in the direction of higher payoff, suggesting that this discrete structure will be a useful analogue of the replicator flow over the whole game. If the sink connected components also capture the **Convergence** properties of the sink chain components, then they meet all of the desired criteria for a solution concept.

Sink connected components are a simple combinatorial object which depend only on each player's preference order. By contrast, the sink chain components are complex topological objects which depend on payoffs and the dynamic. Yet, remarkably, we show that the sink connected components capture key properties of the sink chain components of the replicator. This is the topic of this paper: **the connection between sink chain components of**

games under the replicator dynamic and the sink connected components of the response graph.

1.1. Contributions

We show, firstly, that *sink chain components always exist under the replicator dynamic* (Theorem 4.1), as conjectured by Papadimitriou and Piliouras (2019). The proof uses the structure of the response graph; specifically, each attractor contains an attracting (in the sense of paths) set of nodes in the response graph. In the same theorem, we prove that *sink chain components of the replicator contain sink connected components*. A weaker result is proved in Omidshafiei et al. (2019), which establishes that *asymptotically stable* sink chain components, if they exist, contain sink connected components of the replicator. We believe our proof is the first to establish the general result for the replicator (without the asymptotic stability requirement)[3]

To extend this result, we define the *content* of a sink connected component as the set of all mixed strategy profiles where all pure profiles in their support are in the sink connected component. We then prove that *the content of a connected component is always contained in the associated chain component* (Theorem 5.2). The content of a strongly connected response graph is all mixed profiles; as a corollary, we obtain the surprising result that all profiles are chain recurrent in any game with a strongly connected response graph (such as Matching Pennies or Rock-Paper-Scissors). Augmenting this result, we prove that *if a sink connected component is a subgame, then the associated sink chain component is precisely that subgame* (Corollary 5.5).

We then analyse the influence of the zero-sum and potential properties on the chain components. Because we are interested in the sink connected components, and thus the graph structure, we study all games whose response graphs are isomorphic to a zero-sum or potential game respectively; these are called *preference-zero-sum* and *preference-potential* games (Biggar and Shames, 2022). *The sink chain components of preference-potential games are exactly the pure Nash equilibria* (Theorem 5.6). Preference-zero-sum games have a unique sink connected component, and consequently have exactly one sink chain component (Lemma 5.7). This result is surprisingly analogous to the well-known result that two-player zero-sum games have a convex set of Nash equilibria. However, preference-zero-sum games are a much more general set of games (Biggar and Shames, 2022), defined entirely by their graph structure.

Using our results, we show in Section 6 that the content of sink connected components completely characterises, and thus allows us to compute, the replicator sink chain components of *all* $2 \times n$ strict games, all-but-one 3×3 strict preference-zero-sum games, and all games where every sink connected component is a subgame, such as preference-potential and weakly acyclic games (Young, 1993).

These results suggest that the response graph has a much more significant impact on the outcome of the game than we might expect. We conjecture that there is always a one-to-one correspondence between sink chain components and sink connected components, suggesting

3. A more general result is stated in (Papadimitriou and Piliouras, 2019, Theorem 4.1), but no proof is given and the statement is not true in general.

that the long-run behaviour of the replicator dynamic is fundamentally governed by the graph structure of the game.

Wherever a result has no reference, a proof can be found in the Appendix.

2. Related Work

Though the idea of sink chain components as a tool to analyse games has appeared historically in the game theory literature (Akin and Losert, 1984), the movement towards dynamic solution concepts has accelerated in the modern algorithmic game theory community (Sandholm, 2010). Our solution concept criteria are derived from those in Fabrikant and Papadimitriou (2008). Kleinberg et al. (2011) demonstrate that the replicator generically does not converge to Nash, and give an example where the game converges to a cycle with strictly higher social welfare than the unique Nash equilibrium. Papadimitriou and Piliouras (2016, 2018) present the argument for using chain recurrence to analyse games, and study the replicator dynamic on zero-sum and weighted potential games as a demonstration. This was generalised to Follow-The-Regularised-Leader (FTRL) dynamics and network zero-sum games in Mertikopoulos et al. (2018). Our work most closely follows Papadimitriou and Piliouras (2019), where the authors argue for sink chain components and suggest a connection with sink connected components of the response graph. The authors also set out the importance of further understanding of sink chain components. Since then, Vlatakis-Gkaragkounis et al. (2020) demonstrated that mixed Nash are never attracting under FTRL dynamics and argued for greater understanding the long-run behaviour of these dynamics.

Following Papadimitriou and Piliouras (2019), Omidshafiei et al. (2019) use the replicator, sink chain components and sink connected components to develop a method, called *α-rank*, to evaluate the strength of algorithms in multi-agent learning settings, such as AlphaGo (Silver et al., 2016). The validity of α-rank as a ranking method for algorithms is *predicated on the premise* that sink connected components are a good surrogate for sink chain components of the replicator. In Omidshafiei et al. (2019), this premise rests on a proof that *asymptotically stable* sink chain components of the replicator are finite in number and contain sink connected components. It is not established that any sink chain components (asymptotically stable or not) exist. In this paper, we prove that sink chain components of the replicator always exist and always contain sink connected components (dropping the requirement that they be asymptotically stable) . We also conjecture there is a one-to-one correspondence between sink chain and connected components, which if true would greatly strengthen the motivation for α-rank.

The response graph is a concept of increasing interest, particularly in algorithmic game theory (Fabrikant and Papadimitriou, 2008; Goemans et al., 2005; Kleinberg et al., 2011). A labelled form of the response graph is a key component of the decomposition results of Candogan et al. (2011). Biggar and Shames (2022) explicitly the study the response graph and its sink connected components, and examine the influence of the zero-sum and potential properties on the graph. We use these results extensively in Section 6. More recently, the structure of the response graph is used to describe the 'landscape' of games, for the purposing of analysing multi-agent learning (Omidshafiei et al., 2020).

3. Preliminaries

A game is a triple consisting of N players, strategy sets $S_1, S_2, \ldots, S_N$ and a utility function $u : \prod_{i=1}^N S_i \to \mathbb{R}^N$. We assume each strategy set S_i is finite. An element of $\prod_{i=1}^N S_i$ (an assignment of strategies to players) we call a *profile*, and we denote the set of profiles by Z. We use the notation $\bar{p}_{-i}$ to denote an assignment of strategies to all players other than i, which we call an *antiprofile*, and we denote the set of all i-antiprofiles by $\bar{Z}_{-i}$. If we insert a strategy $s \in S_i$ in the ith index of $\bar{p}_{-i}$ we obtain a profile, and we denote this operation by $(s; \bar{p}_{-i})$. A *subgame* of a game $([N], \{S_1, \ldots, S_n\}, u)$ is a game $([N], \{T_1, \ldots, T_n\}, u')$ where for each i, $T_i \subseteq S_i$, and u' is u restricted to $\prod_{i=1}^N T_i$.

Two profiles are *i-comparable* if they differ only in the strategy of player i; they are *comparable* if they are i-comparable for some player i. If two profiles are comparable, then there is exactly one i such that they are i-comparable. We say a game is *strict* if the payoffs to player i in two i-comparable profiles are never equal. The *response graph* of a game $u : Z \to \mathbb{R}^N$ is the graph $\mathcal{G}_u = (Z, A)$ where there is an arc $p \longrightarrow q \in A$ between profiles p and q if and only if they are i-comparable and $u_i(p) \leq u_i(q)$. A subgraph of the response graph is *attracting* if there are no paths out of it. The sink connected components are minimal attracting subgraphs.

A *mixed strategy* is a distribution over a player's pure strategies, and a *mixed profile* is an assignment of a mixed strategy to each player. We sometimes refer to a profile as a *pure profile* to distinguish it from a mixed profile. For a mixed profile x, we write x^i for the distribution over player i's strategies, and x^i_s for the s-entry of player i's distribution, where $s \in S_i$. The set of mixed profiles on a game is given by $\prod_{i=1}^N \Delta_{|S_i|}$ where $\Delta_{|S_i|}$ are the simplices in $\mathbb{R}^{|S_i|}$. We denote $\prod_{i=1}^N \Delta_{|S_i|}$ simply by X, and call X the *strategy space* of the game. The utility function u of a game extends naturally to mixed profiles. The *expected utility function* of u is $\mathbb{U} : X \to \mathbb{R}^N$, where

$$\mathbb{U}(x) = \sum_{q \in Z} u(q) \prod_{i=1}^N x^i_{q_i}$$

3.1. Dynamical Systems

We study the *replicator dynamic*, which is a continuous-time dynamical system (Sandholm, 2010; Hofbauer and Sigmund, 2003) defined by the following ordinary differential equation, where N is the number of players, $\mathbb{U}$ is the expected utility function, and x is a mixed profile.

$$\dot{x}^p_s = x^p_s \left(\mathbb{U}_p(s; \bar{x}_{-p}) - \sum_{t \in S_p} x^p_t \mathbb{U}_p(t; \bar{x}_{-p}) \right)$$

The solutions to this equation define a *flow* (Sandholm, 2010; Vlatakis-Gkaragkounis et al., 2020) on the strategy space of a game, which is a function $\phi : X \times \mathbb{R} \to X$ which is a continuous group action of the reals on X. We call this the *replicator flow*. The forward orbit of the flow from a given point is called a *trajectory* of the system. Flows are *invertible*, that is, ϕ^{-1} is also a flow, called the *time-reversed* flow. We make use of two special properties of the replicator dynamic. The first is that the replicator is subgame-independent: the support of a point is invariant along an orbit, and the trajectory is only defined by the payoffs in that subgame (Theorem 3.1). The second property of the replicator is *volume preservation*: after a differentiable change of variables, the replicator preserves the volume of all sets on the interior of a game (Akin and Losert, 1984; Hofbauer, 1996; Eshel et al., 1983; Selten, 1988; Sandholm, 2010; Vlatakis-Gkaragkounis et al., 2020). Consequently, no attractor or repellor (Definition 3.2) can exist in the interior of the state space (Theorem 3.4).

Theorem 3.1 (Subgame-independence of the replicator) *Let X be the strategy space of a game u, and Y be the strategy space of a subgame u' of u. The flow $\phi_u|_Y$ of the replicator on u restricted to Y is identical to $\phi_{u'}$, the replicator flow on u'.*

The fact that replicator trajectories have constant support is well-known (Sandholm, 2010). The fact that the flow is defined by the payoffs for strategies in the support follows easily from the fact that all other terms in the differential equation vanish. This result allows us to analyse the flow of the replicator on a subgame of a game using induction on subgames, as we do in the proof of Theorem 5.2. Another important dynamical systems concept are *attractors*.

Definition 3.2 (Sandholm (2010)) *Let A be a compact, non-empty invariant set under a flow ϕ on a compact space X. If there is a neighbourhood U of A such that*

$$\lim_{t\to\infty} \sup_{x\in U} \inf_{y\in A} \mathbf{d}(\phi(x,t), y) = 0$$

where $\mathbf{d}$ is a metric, then we call A an attractor*. An attractor of the time-reversed flow ϕ^{-1} we call a* repellor*.*

There are many equivalent ways of defining an attractor (Sandholm, 2010). In particular, a compact set A being an attractor is equivalent to requiring that there is an open forward-invariant set B with $\phi(\mathrm{cl}(B), t) \subset B$ for all times $t \geq T > 0$, and $A = \bigcap_{t\geq 0} \phi(B, t)$. Such a B is called a *trapping region* for A. Each attractor has a *dual repellor*, defined by trapping regions.

Lemma 3.3 (Sandholm (2010)) *Let A be an attractor, with B a trapping region for A. Then $A^* := \bigcap_{t\leq 0} \phi(X \setminus B, t)$ is a repellor, which we call the* dual repellor *of A.*

Attractors and repellors are dual in the sense that A^* is an attractor of the time-reversed flow ϕ^{-1}, and in this flow A is its dual repellor. On a compact space, the dual repellor is non-empty. Now we can formally establish that the replicator has no attractors in its interior.

Theorem 3.4 (The replicator has no interior attractors, Eshel et al. (1983)) *Let G be a game, with X the strategy space. Under the replicator, there are no attractors in the interior of X.*

This does not mean that there are no attractors at all—only that each must contain at least some points on the boundary of X. This theorem also applies identically to repellors. Inductive repetition of this theorem combined with the flow along edges gives us:

Lemma 3.5 (Attractors contain pure profiles, Vlatakis-Gkaragkounis et al. (2020)) *Every attractor of the replicator flow ϕ contains an attracting set of nodes in the response graph.*

Lemma 3.5 establishes a connection between attractors and the response graph. In particular, because sink connected components are minimal attracting sets, it follows that every attractor of the replicator contains a sink connected component.

3.2. Chains

Definition 3.6 (Alongi and Nelson (2007)) *Let ϕ be the flow of a dynamical system on a compact metric space X. Let x and y be points in X. There is an (ϵ, T)-*chain *from x to y if there is a finite sequence of points $x_1, x_2, \ldots, x_n$ with $x = x_1$ and $y = x_n$, and times $t_1, \ldots, t_n \in [T, \infty)$ such that $\mathbf{d}(\phi(x_i, t_i), x_{i+1}) < \epsilon$.*

Definition 3.7 *If there is an (ϵ, T)-chain from x to y for* all *$\epsilon > 0$ and $T > 0$ we say there is a* pseudo-orbit *from x to y, and write $x \dashrightarrow y$. We say x and y are* chain equivalent *if $x \dashrightarrow y$ and $y \dashrightarrow x$ and write $x \text{ ------ } y$. If $x \text{ ------ } x$ we say x is* chain recurrent.

The relation defined by pseudo-orbits is transitive, though not reflexive, as in general not all points are chain recurrent. However, restricting our attention to the chain recurrent points gives a preorder. Every preorder has an associated partial order given by grouping points into equivalence classes. Each equivalence class under chain equivalence is called a *chain component*, and they are connected components of the topological space (Alongi and Nelson, 2007). Given a chain recurrent point x, we denote its chain component by $[x]$. The chain recurrent points can be characterised by attractors.

Theorem 3.8 (Alongi and Nelson (2007)) *A point x is chain recurrent if and only if, for every attractor A, either $x \in A$ or $x \in A^*$.*

We write $\mathcal{A}_Y$ for the set of attractors of ϕ wholly containing a set Y of points. If Y is a singleton $\{y\}$ then we simplify notation by writing $\mathcal{A}_y$ rather than $\mathcal{A}_{\{y\}}$. Likewise, we write $\mathcal{R}_Y$ for the set of repellors containing Y. The following lemma is critical.

Lemma 3.9 *If $x \dashrightarrow y$ then $\mathcal{A}_x \subseteq \mathcal{A}_y$ and $\mathcal{R}_y \subseteq \mathcal{R}_x$. Conversely, if x is chain recurrent, then $\mathcal{A}_x \subseteq \mathcal{A}_y$ implies $x \dashrightarrow y$, and $\mathcal{R}_x \subseteq \mathcal{R}_y$ implies $y \dashrightarrow x$. Consequently, if x is chain recurrent, $x \text{ ------ } y$ if and only if $\mathcal{A}_x = \mathcal{A}_y$ and $\mathcal{R}_x = \mathcal{R}_y$. If y is also chain recurrent then one of these equalities is sufficient.*

This states that pseudo-orbits never leave (but possibly enter) attractors. Dually, a pseudo-orbit never enters repellors, but may leave them. The following result connects the response graph with pseudo-orbits.

Lemma 3.10 (Omidshafiei et al. (2019)) *If v and w are pure profiles, and there is a path $v, v_1, v_2, \ldots, w$ in the response graph, then $v \dashrightarrow w$ under the replicator.*

The idea of this proof is that the points on an arc move toward in the direction of the arc under the replicator. Upon nearing the head of the arc, an ϵ-jump will take us to the subsequent arc of the path. It follows that if H is a connected component of the response graph, then all profiles in H are contained in exactly one chain component, which we denote $[H]$.

4. The Existence Theorem for Sink Chain Components

In this section we establish that sink chain components exist under the replicator, as conjectured by Papadimitriou and Piliouras. Specifically we prove that all chain recurrent points in the game can reach a sink chain component via a pseudo-orbit. Combining this with the Fundamental Theorem of Dynamical Systems, which states that all points converge to chain recurrent ones, we find that under the replicator there are pseudo-orbits from all points to sink chain components.

Theorem 4.1 (Existence of sink chain components) *Let $x \in X$ be a chain recurrent point under the replicator. Then there exists a pure profile y and pseudo-orbit $x \dashrightarrow y$ where y is in a sink connected component H contained in a sink chain component $[H]$.*

The proof proceeds in two steps. The first step is to show that there are pseudo-orbits from all chain recurrent points to a pure profile. Lemmas 3.5 and 3.10 establish that attractors of the replicator contain sink connected components. Lemma 3.9 demonstrates that the chain components are defined by the structure of the attractors. If we knew that there were only finitely many attractors then we would be done—any minimal attractor is a sink chain component. To establish existence without assuming there are finitely many attractors we use the relationship with the sink connected components, which we know are finite in number. The key observation is that from any chain recurrent point x, there must exist a pure profile p contained in all attractors containing x, which implies the existence of a pseudo-orbit from x to p. Finally, we use the attractors of the system to order the sink connected components, and so find the sink chain components.

5. Bounding the Sink Chain Components using Connected Components

Theorem 4.1 establishes that sink chain components exist and contain sink connected components. Our goal in this section is to strengthen these results. That is, given a sink connected component H, what can we say about the unique sink chain component $[H]$ which contains it[4]?

We first prove a lower bound: $[H]$ always contains the *content* of H.

Definition 5.1 (Content) *Let W be a collection of pure profiles in a game u. The* content *of W, denoted* content(W)*, is the set of all mixed strategy profiles x where all pure profiles in the support of x are in W.*

4. Uniqueness here follows from the fact that chain components are disjoint. Each connected component H is chain recurrent (Lemma 3.10), and so is contained in a chain component.

The content is a topological subspace of the strategy space. If W is a connected component of the response graph, then the content of W is topologically connected. We now show that if W is strongly connected in the graph, the content of W is always a subset of $[W]$.

Theorem 5.2 *Let H be a strongly connected component of the response graph. Then the content of H is contained in $[H]$, the chain component of the replicator containing H.*

The proof uses induction on subgames, following Theorem 3.1. Specifically, if the entire boundary of a subgame is contained in the same chain component, then there can be no attractors in that subgame, by Lemma 3.9 and Theorem 3.4. As a corollary, whenever the response graph is strongly connected the *entire strategy space is chain recurrent.* An example is the Matching Pennies game (Figure 1).

Corollary 5.3 *In any game with a strongly connected response graph, every mixed profile is chain recurrent.*

Now we can obtain an upper bound in some cases: when the sink connected component H is a subgame, then $[H] = \text{content}(H)$. This is because attracting subgames are attractors.

Lemma 5.4 *If Y is an attracting subgame, then the set of mixed profiles in Y is an attractor under the replicator.*

Pure NEs are attracting *singleton subgames*—by this lemma, they are always attractors under the replicator, as we know from existing results (Vlatakis-Gkaragkounis et al., 2020). Combining this result with Theorem 5.2 gives us:

Corollary 5.5 *If a sink component H is a subgame, then $[H] = \text{content}(H)$.*

Proof Sink connected components are attracting, so H is an attracting subgame, and thus is an attractor under the replicator (Lemma 5.4). The content of H is exactly the mixed profiles in that subgame, and by Theorem 5.2, all points in the content of H are in $[H]$, and so are chain equivalent. However, as an attractor, no pseudo-orbit can leave this subgame by Lemma 3.9 and so $[H] = \text{content}(H)$. ■

The sink connected component is generally not a subgame, but when this is the case our upper and lower bounds meet, and the sink chain component can be precisely characterised.

5.1. Zero-sum and Potential games

A number of special properties are known for potential (Monderer and Shapley, 1996) and zero-sum (Von Neumann and Morgenstern, 1944) games. The current understanding of the chain components of such games comes from Piliouras and Shamma (2014); Papadimitriou and Piliouras (2016, 2018); Mertikopoulos et al. (2018) where it is proved that the chain components of a weighted potential game are its pure Nash equilibria, and when a two-player zero-sum game has an interior equilibrium, all profiles are chain recurrent. We generalise the former result in Theorem 5.6 as a straightforward consequence of Corollary 5.5.

The case of two-player zero-sum games is more intriguing. All two-player zero-sum games with interior equilibria that we know of have *strongly connected response graphs*. That is, in all known cases, this result is subsumed by Corollary 5.3. Indeed, if Conjecture 7.2 is true, then this would become a theorem: any two-player zero-sum game with an interior equilibrium would have a strongly connected response graph. Corollary 5.3 is much broader. Many games are strongly connected but not two-player zero-sum with an interior equilibrium. Focusing on two-player games, we find that the response graphs of all $2 \times n$ zero-sum games are strongly connected (Lemma A.8). On 3×3 strict games, there are only two response graphs of zero-sum games which are *not* strongly connected (Biggar and Shames, 2022). These are the Inner and Outer Diamond graphs (Biggar and Shames, 2022), depicted in Figure 2. No zero-sum game with either of these response graphs ever has an interior equilibrium (Lemma A.9).

Following Biggar and Shames (2022), we call a game *preference-zero-sum* if its response graph is isomorphic to that of a two-player zero-sum game. We call it *preference-potential* if its response graph is isomorphic to that of a potential game. This equivalence relation is quite broad. In particular, any weighted potential game is preference-potential, but not vice versa.

Theorem 5.6 *In a strict preference-potential game under the replicator, the sink chain components are exactly* $\{p_1\}, \ldots, \{p_m\}$ *where* $p_1, \ldots, p_m$ *are the pure Nash equilibria.*

Proof Each pure Nash equilibrium p in a strict game is an attracting subgame, so by Corollary 5.5 the associated sink chain component is precisely $\{p\}$. Since every sink chain component contains a sink connected component, and thus a pure Nash equilibrium, these are precisely the sink chain components. ∎

Recalling the Coordination game (Figure 1), preference-potential games can have mixed Nash equilibria, but these are *never in the sink chain components*, while the pure Nash equilibria are always precisely the sink chain components. Beyond preference-potential games, this argument applies to any game where every sink connected component is a subgame, such as *weakly acyclic games* (Young, 1993).

Lemma 5.7 *Every preference-zero-sum game has exactly one sink connected component and thus precisely one sink chain component.*

Proof A sink chain component exists by Theorem 4.1. Every preference-zero-sum game has exactly one sink connected component by (Biggar and Shames, 2022, Theorem 4.9). Sink chain components are disjoint and contain sink connected components, so there must be precisely one. ∎

We conclude that, at least in preference-zero-sum and preference-potential games, there is a one-to-one correspondence between sink chain components and sink connected components.

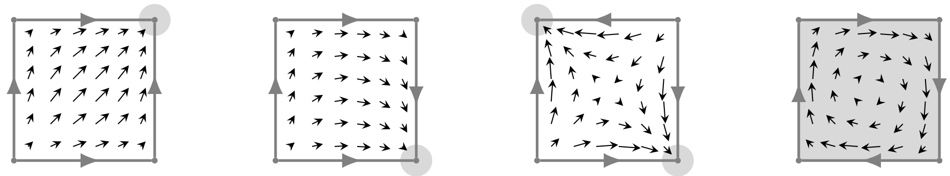

Figure 1: Sink chain components (highlighted in grey) of the replicator dynamic on the strict 2x2 games. These response graphs are called Double-Dominance, Single-Dominance, Coordination and Matching Pennies (Biggar and Shames, 2022). We have overlaid a typical example of the replicator vector field for clarity in each case.

6. Applications

In this section we identify the sink chain components of the replicator on a collection of small games, demonstrating how our findings extend the existing literature.

- **(Any game equivalent to)** 2×2 **Coordination**: This game has been well-studied (Papadimitriou and Piliouras, 2016, 2018, 2019). It is known that any game sharing a response graph with 2×2 Coordination (Figure 1) has two sink chain components which are the two pure Nash equilibria. All such games are preference-potential, and so this is an immediate corollary of Theorem 5.6.
- **(Any game equivalent to) Matching Pennies**: Again, this game is well-studied (Papadimitriou and Piliouras, 2016, 2018, 2019; Balduzzi et al., 2018), and it is known that all profiles are chain recurrent. However, it is assumed in these papers that the game is zero-sum. Because the response graph of Matching Pennies is a cycle, it is strongly connected and so by Corollary 5.3 any game sharing a response graph with Matching Pennies is chain recurrent. Almost all such games are not zero-sum.
- ***Any*** 2×2 **strict game**: There are only four response graphs of strict 2×2 games (Biggar and Shames, 2022), of which we have already discussed two (Matching Pennies and 2×2 Coordination). The remaining two (Single- and Double-dominance) are preference-potential, so Theorem 5.6 applies. See Figure 1.
- ***Any*** 2×3 **strict game**: We can focus on games without dominated strategies, because the sink chain component will never contain such strategies (Sandholm, 2010). By Biggar and Shames (2022), there are three response graphs of 2×3 games without dominated strategies. One is strongly connected, and so all profiles are chain recurrent (Corollary 5.3); another is preference-potential; in the final graph the unique sink connected component is a pure Nash equilibrium, so in both of the latter cases the sink chain components are exactly the pure Nash equilibria (Theorem 5.6).
- ***Any*** $2 \times n$ **strict game**: Generalising the previous cases, it turns out (Lemma A.8) that sink connected components of strict $2 \times n$ games are always subgames, so Corollary 5.5 applies. Additionally, every preference-zero-sum $2 \times n$ game without dominated strategies is strongly connected.

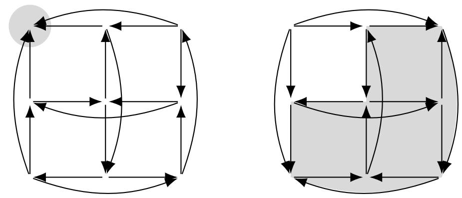

Figure 2: The Inner and Outer Diamond graphs and the content of their sink connected components.

- ***Any*** 3×3 **strict preference-zero-sum game**: By Biggar and Shames (2022), all but two of the response graphs of such games are strongly connected, so Corollary 5.3 applies. The two exceptions are known as the Inner and Outer Diamond graphs (Figure 2). The unique (Lemma 5.7) sink connected component in the Inner Diamond game is a pure Nash equilibrium, so this is the unique sink chain component by Theorem 5.6. The Outer Diamond game is also preference-zero-sum, so has a unique sink chain component containing the content of its sink connected component. We conjecture (Conjecture 7.2) that the sink chain component is always precisely the content of the sink connected component.

7. Conclusions and Open Problems

In this paper we demonstrated the influence of the response graph on the chain components of the replicator dynamic. In particular, we showed that sink chain components exist and contain the content of sink connected components of the response graph.

The close relationship between connected components and chain components is interesting from the perspective of both game theory and dynamical systems, and poses several questions that are deserving of further investigation. In this section we outline some of the most important. In all the cases we discussed above where we could identify the sink chain components, each not only contained a sink connected component but in fact exactly one, and each sink connected component was contained in a sink chain component. We conjecture that this relationship is generally true.

Conjecture 7.1 *There is a one-to-one correspondence between sink chain components and sink connected components.*

Using the construction in the proof of Theorem 4.1, we can cast this question as one about attractors: Conjecture 7.1 is true if and only if each sink connected component has at least one attractor containing it and no others. If true, it would greatly strengthen the premise of Papadimitriou and Piliouras that sink connected components are the *right* combinatorial surrogate of sink chain components. However, even if such a one-to-one correspondence existed, it wouldn't necessarily specify precisely which points are in sink chain components. The following stronger conjecture does do this:

Conjecture 7.2 *The sink chain components of a game are exactly the content of the sink connected components.*

Again, we can rephrase this as a statement about attractors: the conjecture holds if and only if the content of a sink chain component is always an attractor. Conjecture 7.2 would be particularly shocking, because it would imply that the 'long-run' outcome of the replicator dynamic is *entirely dictated* by the response graph. That is, from the perspective of the replicator, only the preference orders of each player affect the chain components. We do not currently know of counterexamples to either conjecture.

There is another important question: to which other dynamics do these results generalise? In particular, the steep FTRL dynamics (Vlatakis-Gkaragkounis et al., 2020) are known to satisfy Theorem 3.4, which is necessary for these results, so it seems plausible that these results may apply to this class.

References

Ethan Akin and Viktor Losert. Evolutionary dynamics of zero-sum games. *Journal of Mathematical Biology*, 20(3):231–258, 1984.

John M Alongi and Gail Susan Nelson. *Recurrence and topology*, volume 85. American Mathematical Soc., 2007.

Gabriel P Andrade, Rafael Frongillo, and Georgios Piliouras. Learning in matrix games can be arbitrarily complex. In *Conference on Learning Theory*, pages 159–185. PMLR, 2021.

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.

David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *International Conference on Machine Learning*, pages 354–363. PMLR, 2018.

Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Perturbations of set-valued dynamical systems, with applications to game theory. *Dynamic Games and Applications*, 2(2):195–205, 2012.

Oliver Biggar and Iman Shames. The graph structure of two-player games. *arXiv preprint arXiv:2209.10182*, 2022.

Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3):474–503, 2011.

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 *Conference on Learning Theory*, pages 807–834. PMLR, 2019.

Yun Kuen Cheung and Georgios Piliouras. Chaos, extremism and optimism: Volume analysis of learning in games. *Advances in Neural Information Processing Systems*, 33:9039–9049, 2020.

Charles C Conley. *Isolated invariant sets and the Morse index*. Number 38. American Mathematical Soc., 1978.

Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

Ilan Eshel, Ethan Akin, et al. Coevolutionary instability of mixed nash solutions. *Journal of Mathematical Biology*, 18(2):123–133, 1983.

Alex Fabrikant and Christos H Papadimitriou. The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond. In *SODA*, volume 8, pages 844–853. Citeseer, 2008.

Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–151. IEEE, 2005.

Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.

Josef Hofbauer. Evolutionary dynamics for bimatrix games: A hamiltonian system? *Journal of Mathematical Biology*, 34(5):675–688, 1996.

Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American mathematical society*, 40(4):479–519, 2003.

William D Kalies, Konstantin Mischaikow, and Robert CAM Vandervorst. Lattice structures for attractors iii. *Journal of Dynamics and Differential Equations*, pages 1–40, 2021.

Robert D Kleinberg, Katrina Ligett, Georgios Piliouras, and Éva Tardos. Beyond the nash equilibrium barrier. In *ICS*, pages 125–140, 2011.

Jean-François Mertens. Ordinality in non cooperative games. *International Journal of Game Theory*, 32(3):387–430, 2004.

Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2703–2717. SIAM, 2018.

Jason Milionis, Christos Papadimitriou, Georgios Piliouras, and Kelly Spendlove. Nash, conley, and computation: Impossibility and incompleteness in game dynamics. *arXiv preprint arXiv:2203.14129*, 2022.

Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14 (1):124–143, 1996.

John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

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.

Shayegan Omidshafiei, Karl Tuyls, Wojciech M Czarnecki, Francisco C Santos, Mark Rowland, Jerome Connor, Daniel Hennes, Paul Muller, Julien Pérolat, Bart De Vylder, et al. Navigating the landscape of multiplayer games. *Nature communications*, 11(1):1–17, 2020.

Christos Papadimitriou and Georgios Piliouras. From nash equilibria to chain recurrent sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 227–235, 2016.

Christos Papadimitriou and Georgios Piliouras. From nash equilibria to chain recurrent sets: An algorithmic solution concept for game theory. *Entropy*, 20(10):782, 2018.

Christos Papadimitriou and Georgios Piliouras. Game dynamics as the meaning of a game. *ACM SIGecom Exchanges*, 16(2):53–63, 2019.

Georgios Piliouras and Jeff S Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In *Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 861–873. SIAM, 2014.

Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.

William H Sandholm. *Population games and evolutionary dynamics*. MIT press, 2010.

Reinhard Selten. A note on evolutionarily stable strategies in asymmetric animal conflicts. In *Models of Strategic Rationality*, pages 67–75. Springer, 1988.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

JMPGR Smith and George R Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

Steven H Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. CRC press, 2018.

Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed nash equilibria: They do not mix. *Advances in Neural Information Processing Systems*, 33:1380–1391, 2020.

John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior.* Princeton university press, 1944.

Yaodong Yang and Jun Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583*, 2020.

H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

Appendix A. Proofs

In addition to the notation introduced in Section 3, we will use some additional notation in the proofs in this appendix. If Y is a subgame, we denote the strategies used by player p in Y by Y_p. Similarly, we denote the p-antiprofiles where all other players play strategies in Y by $\bar{Y}_{-p}$. Given a p-antiprofile $\bar{q}$, we write $x_{\bar{q}}$ as shorthand for $\prod_{j \neq p} x^j_{\bar{q}_j}$.

We will use a rearranged form of the replicator equation in several proofs.

Lemma A.1 *The replicator equation is equivalent to*

$$\dot{x}^p_s = x^p_s \sum_{t \in S_p} x^p_t \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(u_p(s; \bar{q}) - u_p(t; \bar{q}) \right)$$

Proof

$$\begin{aligned}
\dot{x}^p_s &= x^p_s \left(\mathbb{U}_p(s; \bar{x}_{-p}) - \sum_{t \in S_p} x^p_t \mathbb{U}_p(t; \bar{x}_{-p}) \right) \\
&= x^p_s \left(\sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} u_p(s; \bar{q}) - \sum_{t \in S_p} x^p_t \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} u_p(t; \bar{q}) \right) \\
&= x^p_s \sum_{t \in S_p} x^p_t \left(\sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} u_p(s; \bar{q}) - \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} u_p(t; \bar{q}) \right) \\
&= x^p_s \sum_{t \in S_p} x^p_t \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(u_p(s; \bar{q}) - u_p(t; \bar{q}) \right)
\end{aligned}$$

where in the third equality we have used the fact that $\sum_{t \in S_p} x^p_t = 1$. ■

Theorem A.2 (Theorem 3.1) *Let X be the strategy space of a game u, and Y be the strategy space of a subgame u' of u. The flow $\phi_u|_Y$ of the replicator on u restricted to Y is identical to $\phi_{u'}$, the replicator flow on u'.*

Proof Let x be a mixed profile in the subgame Y. Observe from the definition of the replicator that for any player p with strategy t outside of Y, $x_t^p = 0$ and so $\dot{x}_t^p = 0$. Now we show that if $x_t^p > 0$ then $\dot{x}_t^p$ depends on the payoffs of profiles in Y. Then note that $x_{\bar{q}} = 0$ if $\bar{q} \notin \bar{Y}_{-p}$. Then, for $s \in Y_p$, by Lemma A.1,

$$\begin{aligned}
\dot{x}_s^p &= x_s^p \sum_{t \in S_p} x_t^p \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(u_p(s; \bar{q}) - u_p(t; \bar{q}) \right) \\
&= x_s^p \sum_{t \in Y_p} x_t^p \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(u_p(s; \bar{q}) - u_p(t; \bar{q}) \right) \\
&= x_s^p \sum_{t \in Y_p} x_t^p \sum_{\bar{q} \in \bar{Y}_{-p}} x_{\bar{q}} \left(u_p(s; \bar{q}) - u_p(t; \bar{q}) \right)
\end{aligned}$$

and this is the definition of the replicator on the restricted game Y. ■

Lemma A.3 *If $x \dashrightarrow y$ then $\mathcal{A}_x \subseteq \mathcal{A}_y$ and $\mathcal{R}_y \subseteq \mathcal{R}_x$. Conversely, if x is chain recurrent, then $\mathcal{A}_x \subseteq \mathcal{A}_y$ implies $x \dashrightarrow y$, and $\mathcal{R}_x \subseteq \mathcal{R}_y$ implies $y \dashrightarrow x$. Consequently, if x is chain recurrent, x ------ y if and only if $\mathcal{A}_x = \mathcal{A}_y$ and $\mathcal{R}_x = \mathcal{R}_y$. If y is also chain recurrent then one of these equalities is sufficient.*

Proof The first claim is well-known, with a proof in (Akin and Losert, 1984).

For the second claim, assume x is chain recurrent and $\mathcal{A}_x \subseteq \mathcal{A}_y$. Now fix some $\epsilon > 0$ and $T > 0$, and let $C_{\epsilon,T}(x)$ be the set of points reachable by (ϵ, T)-chains from x. Note that $C_{\epsilon,T}(x)$ is open. If $C_{\epsilon,T}(x)$ is all of X, then $y \in C_{\epsilon,T}(x)$. Otherwise, consider $V = N_{\epsilon/2}(C_{\epsilon,T}(x))$, that is the set of points within $\epsilon/2$ distance of $C_{\epsilon,T}(x)$. This is open. Now let z be some point within ϵ of a point $\chi \in C_{\epsilon,T}(x)$. By definition, $\phi(z,t) \in C_{\epsilon,T}(x)$ for all $t \geq T$, as we can obtain a (ϵ, T)-chain to it by adding a step to the chain from x to χ. Thus, we find that $\phi(\mathrm{cl}(V), t) \subseteq \phi(N_\epsilon(C_{\epsilon,T}(x)), t) \subseteq C_{\epsilon,T}(x) \subset V$. Hence V is a trapping region for some attractor $A \subseteq C_{\epsilon,T}(x) \subseteq V$. Let A^* be the dual repellor of A. By Theorem 3.8, $x \in A \cup A^*$, and since $x \in V$ and $A^* \subseteq X \setminus V$ we have $x \in A$. But then $y \in A$, since y is contained in all attractors containing x, and so $y \in A \subseteq C_{\epsilon,T}(x)$. Since ϵ and T were arbitrary, $x \dashrightarrow y$. Now observing that in the time-reversed flow ϕ^{-1}, repellors become attractors and vice versa, and pseudo-orbits $x \dashrightarrow y$ become pseudo-orbits $y \dashrightarrow x$, establishing the claim for repellors.

If x ------ y, then $x \dashrightarrow y$ and $y \dashrightarrow x$, so $\mathcal{A}_x = \mathcal{A}_y$ and $\mathcal{R}_y = \mathcal{R}_x$ by the first claim. By the second claim, if x is chain recurrent and $\mathcal{A}_x = \mathcal{A}_y$ and $\mathcal{R}_y = \mathcal{R}_x$ then x ------ y, and so y is necessarily also chain recurrent. If we know in advance that y is chain recurrent, then one of the equalities is sufficient, because by Theorem 3.8 if two chain recurrent points are contained in all the same attractors then they are also contained in all the same repellors. ■

We will need the following result.

Lemma A.4 (Kalies et al. (2021)) *If ϕ is a flow, then any finite non-empty intersection of attractors is an attractor.*

Theorem A.5 (Theorem 4.1) *Let $x \in X$ be a chain recurrent point under the replicator. Then there exists a pure profile y and pseudo-orbit $x \dashrightarrow y$ where y is in a sink connected component H contained in a sink chain component $[H]$.*

Proof First, we define a preorder on the set of sink connected components of a game, defined by the attractors of the replicator. If K and H are sink connected components, we say $K \preceq H$ if $\mathcal{A}_K \supseteq \mathcal{A}_H$. This is a preorder because it inherits reflexivity and transitivity from the sets $\mathcal{A}_K$ and $\mathcal{A}_H$. By finiteness, this order has at least one minimal element, and we denote the minimal elements by $M_1, \ldots, M_m$. We will show that the sink chain components are exactly the chain components $[M_1], \ldots, [M_m]$ (which are not necessarily distinct).

(*Claim:* Every chain recurrent point has a pseudo-orbit to some M_i.)

Let x be chain recurrent. Now define

$$\mathcal{M}_x := \bigcap_{A \in \mathcal{A}_{\{x\}}} \{\text{sink connected components contained in } A\}$$

Recall that each such A contains a sink connected component (Lemma 3.5). Suppose for contradiction that $\mathcal{M}_x$ is empty. Then for each sink connected component $H_1, \ldots, H_n$ in the game there exist attractors $B_1, \ldots, B_n$, each containing x, where B_i does not contain H_i. By Lemma A.4, $\bigcap_{i=1}^n B_i$ is an attractor, which by Lemma 3.5 contains a sink connected component H_i. This is a contradiction, because H_i is not contained in B_i, by construction. Hence $\mathcal{M}_x$ is not empty. Let K be a sink connected component in $\mathcal{M}_x$. By definition of $\mathcal{M}_x$, $\mathcal{A}_{\{x\}} \subseteq \mathcal{A}_K$.

There exists some sink connected component M_i with $M_i \preceq K$ and so $\mathcal{A}_{\{x\}} \subseteq \mathcal{A}_K \subseteq \mathcal{A}_{M_i}$, and by Lemma 3.9 there is a pseudo-orbit $x \dashrightarrow y$ from x to every point y in M_i.

(*Claim:* The $[M_i]$s are sink chain components.)

Let z be some chain recurrent point, and suppose there is a pseudo-orbit $x \dashrightarrow z$, with $x \in [M_i]$. By the above argument there is some M_j with $z \dashrightarrow y$, $y \in M_j$. By transitivity there is a pseudo-orbit $x \dashrightarrow y$, which by Lemma 3.9 implies that $\mathcal{A}_{M_i} \subseteq \mathcal{A}_{M_j}$, but since each M_i is minimal in this order we must have $\mathcal{A}_{M_i} = \mathcal{A}_{M_j}$. Lemma 3.9 establishes that $[M_i] = [M_j]$. Hence $z \in [M_i]$ and so $[M_i]$ is a sink chain component. ■

Theorem A.6 (Theorem 5.2) *Let H be a connected component of the response graph. Then the content of H is contained in $[H]$, the chain component of the replicator containing H.*

Proof We show that if all of the pure profiles in a game Y are in the same chain component K, then all profiles in the game are in K. We will use induction, by Theorem 3.1. First, observe that the base case where Y is a pure profile is trivial.

Now suppose for induction that all mixed profiles in all proper subgames of Y are contained in K. The union of all proper subgames is the boundary of Y, so the boundary of Y is contained in K. Suppose there is an attractor A in Y. By Theorem 3.4, A intersects the boundary, but then by Lemma 3.9, A must contain all of the boundary, since all points on the boundary are chain equivalent. But then the dual repellor A^* is contained in the interior of Y, which is impossible by Theorem 3.4. We conclude that there are no attractors or repellors in Y, and so all points are chain equivalent by Lemma 3.9.

Finally, we know by Lemma 3.10 that all profiles in a connected component H are chain recurrent and in the same chain component $[H]$. By the above argument, in any subgame where all pure profiles are in H, all mixed profiles are chain equivalent, and so in $[H]$. The content of H is the union of all of these subgames. ■

Lemma A.7 (Lemma 5.4) *If Y is an attracting subgame, then the set of mixed profiles in Y is an attractor under the replicator.*

Proof Given a number $0 < M < 1$ we define $\mathcal{Y}_M$ to be the M-neighbourhood of Y in the infinity norm; that is, the set of mixed profiles x where for any player p and strategy $s \in S_p \setminus Y_p$, $x^p_s < M$.

We know that the points in Y form a compact non-empty invariant set, by Theorem 3.1, and we assume that Y is not the whole game.

(*Claim:* there exists an $M > 0$ such that in $\mathcal{Y}_M$, $\dot{x}^p_r < 0$ for any strategy $s \in S_p \setminus Y_p$)

Fix a player p, and let $s \in S_p \setminus Y_p$. By Lemma A.1,

$$\dot{x}^p_s = x^p_s \sum_{t \in S_p} x^p_t \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(u_p(s;\bar{q}) - u_p(t;\bar{q})\right)$$

Now, if $t \in S_p$, we define $\alpha^p_{s,t} = \max_{\bar{q} \in \bar{Y}_{-p}}(u_p(s;\bar{q}) - u_p(t;\bar{q}))$ and $\beta^p_{s,t} = \max_{\bar{q} \in \bar{Z}_{-p} \setminus \bar{Y}_{-p}}(u_p(s;\bar{q}) - u_p(t;\bar{q}))$. As Y is attracting and $t \in S_p$, $\alpha^p_{s,t}$ must be negative. Finally, define $\gamma^p_s = \max_{t \in S_p \setminus Y_p} \max_{\bar{q} \in \bar{Z}_{-p}}(u_p(s;\bar{q}) - u_p(t;\bar{q}))$.

Then

$$\begin{aligned}
\dot{x}^p_s &= x^p_s \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \sum_{t \in S_p} x^p_t \left(u_p(s;\bar{q})_p - u_p(t;\bar{q})\right) \\
&= x^p_s \sum_{\bar{q} \in \bar{Z}_{-p}} x_{\bar{q}} \left(\sum_{t \in Y_p} x^p_t \left(u_p(s;\bar{q}) - u_p(t;\bar{q})\right) + \sum_{t \in S_p \setminus Y_p} x^p_t \left(u_p(s;\bar{q}) - u_p(t;\bar{q})\right) \right) \\
&\leq x^p_s \left(\sum_{t \in Y_p} x^p_t \left(\alpha^p_{s,t} \sum_{\bar{q} \in \bar{Y}_{-p}} x_{\bar{q}} + \beta^p_{s,t} \sum_{\bar{q} \in \bar{Z}_{-p} \setminus \bar{Y}_{-p}} x_{\bar{q}} \right) + \gamma^p_s I(\gamma^p_s > 0) \sum_{t \in S_p \setminus Y_p} x^p_t \right)
\end{aligned}$$

where $I(\cdot)$ is the indicator function. Now observe that every antiprofile $\bar{q} \in \bar{Z}_{-p} \setminus \bar{Y}_{-p}$ contains at least one strategy not in Y, and so for $x \in \mathcal{Y}_M$, $x_{\bar{q}} \leq M$. Thus $\sum_{\bar{q} \in \bar{Z}_{-p} \setminus \bar{Y}_{-p}} x_{\bar{q}} \leq MQ$, where Q is the cardinality of $Z_{-p} \setminus \bar{Y}_{-p}$. Hence $\sum_{\bar{q} \in \bar{Y}_{-p}} x_{\bar{q}} \geq 1 - MQ$. For $M < 1/Q$, $1 - MQ > 0$. Given $\alpha^p_{s,t} < 0$, we have $\alpha^p_{s,t} \sum_{\bar{q} \in \bar{Y}_{-p}} x_{\bar{q}} \leq \alpha^p_{s,t}(1 - MQ)$ and $\beta^p_{s,t} I(\beta^p_{s,t} >$

$0) \sum_{\bar{q}\in \bar{Z}_{-p}\setminus \bar{Y}_{-p}} x_{\bar{q}} \le MQ$. Also, for $x \in \mathcal{Y}_M$ we have $\sum_{t\in S_p\setminus Y_p} x^p_t \le MN$. Hence

$$\dot{x}^p_s \le x^p_s \left(\sum_{t\in Y_p} x^p_t \left(\alpha^p_{s,t}(1-MQ) + \beta^p_{s,t}MQ\right) + \gamma^p_s I(\gamma^p_s > 0)NM \right)$$
$$\le x^p_s \left(\sum_{t\in Y_p} x^p_t \alpha^p_{s,t} - MQ\sum_{t\in Y_p} \alpha^p_{s,t} + MQ\sum_{t\in Y_p} \beta^p_{s,t} + \gamma^p_s I(\gamma^p_s > 0)NM \right)$$

Finally, let $\alpha^p_s = \max_{t\in Y_p} \alpha^p_{s,t}$. Then $\sum_{t\in Y_p} x^p_t \alpha^p_{s,t} \le \alpha^p_s \sum_{t\in Y_p} x^p_t \le \alpha^p_s(1-ML_p)$, where $L_p = |S_p \setminus Y_p|$, and recalling that $\alpha^p_s < 0$. Thus we obtain the bound:

$$\dot{x}^p_s \le x^p_s \left(\alpha^p_s(1-ML_p) - MQ\sum_{t\in Y_p} \alpha^p_{s,t} + MQ\sum_{t\in Y_p} \beta^p_{s,t} + \gamma^p_s I(\gamma^p_s > 0)NM \right)$$
$$\le x^p_s \left(\alpha^p_s + M\left(-L_p\alpha^p_s - Q\sum_{t\in Y_p} \alpha^p_{s,t} + Q\sum_{t\in Y_p} \beta^p_{s,t} + \gamma^p_s I(\gamma^p_s > 0)N \right)\right)$$

The sign of this equation is determined by the sign of the bracketed term, which is a linear equation in M that is independent of x. Given $\alpha^p_s < 0$, there is a positive value M^p_s ($< 1/Q$)) such that for any $M \le M^p_s$ we have $\dot{x}^p_s$ always negative. Finally, define $M := \min_p \min_{s\in S_p\setminus Y_p} M^p_s$, which is positive as there are only finitely many strategies. This establishes the claim.

Since x^p_s is strictly decreasing inside $\mathcal{Y}_M$, this set is forward-invariant and so $\phi^t(x)^p_s \to 0$ as $t \to \infty$. The result then follows from

$$\lim_{t\to\infty} \sup_{x\in\mathcal{Y}_M} \min_{y\in Y} \mathbf{d}(\phi^t(x), y) = \lim_{t\to\infty} \sup_{x\in\mathcal{Y}_M} \min_{y\in Y} \|\phi^t(x) - y\|_\infty$$
$$= \lim_{t\to\infty} \sup_{x\in\mathcal{Y}_M} \min_{y\in Y} \max_p \max_{s\in S_p} |\phi^t(x)^p_s - y^p_s|$$
$$= \lim_{t\to\infty} \sup_{x\in\mathcal{Y}_M} \max_p \max_{s\in S_p} \min_{y\in Y} |\phi^t(x)^p_s - y^p_s|$$

The value $\min_{y\in Y} |\phi^t(x)^p_s - y^p_s|$ is zero if the strategy $s \in S_p$ is inside Y. Consequently, since $\phi^t(x)$ is not in Y, the strategy s which maximises

$$\max_p \max_{s\in S_p} \min_{y\in Y} |\phi^t(x)^p_s - y^p_s|$$

must be outside Y. Thus $\phi^t(x)^p_s \to 0$ as $t \to \infty$, for any $x \in \mathcal{Y}_M$, and so

$$\lim_{t\to\infty} \sup_{x\in\mathcal{Y}_M} \max_p \max_{s\in S_p} \min_{y\in Y} |\phi^t(x)^p_s - y^p_s| = 0$$

■

Lemma A.8 *The sink connected components of strict $2 \times n$ games are subgames. If the game is preference-zero-sum game without dominated strategies, then the response graph is strongly connected.*

Proof We can assume there are no dominated strategies, because the sink connected components are always contained within the subgame surviving iterated dominance. We call the strategy sets $\{A, B\}$ and $\{s_1, \dots, s_n\}$ respectively. Let $s_1, s_2, \dots, s_n$ be the order in which player 2 prefers their strategies when player 1 plays A. Because no strategy dominates any other, when player 1 plays B player 2 must prefer their own strategies in the reverse order, which is $s_n, s_{n-1}, \dots, s_1$. Thus there are arcs $(A, s_i) \longrightarrow (A, s_{i+1})$ and $(B, s_i) \longrightarrow (B, s_{i-1})$ in the response graph for each s_i.

Any non-singleton sink connected component must contain profiles with both of player 1's strategies A and B. Without loss of generality suppose that (A, s_i) and (B, s_j) are in a sink connected component. By the structure of the graph, there are nodes (B, s_k) and (A, s_ℓ) with $k \geq n$, $\ell \leq i$ and paths from (A, s_i) and (B, s_j) respectively. All nodes reachable from (B, s_k) and (A, s_ℓ) are also in the sink connected component, and so in particular (A, s_j) and (B, s_i) are both in the component. This is true for any pair of profiles, so the component is a subgame.

Now we consider the second claim. There are two cases: (1) either one of (B, s_1) or (A, s_n) is a sink, or (2) there are arcs $(B, s_1) \longrightarrow (A, s_1)$ and $(A, s_n) \longrightarrow (B, s_n)$ in the graph. In case (2) there is a Hamiltonian cycle

$$(A, s_1), (A, s_2), \dots, (A, s_n), (B, s_n), (B, s_{n-1}), \dots, (B, s_1), (A, s_1)$$

containing all nodes so the graph is strongly connected. We show case (1) is impossible in a preference-zero-sum game. Assume without loss of generality that (A, s_n) is a sink, so there is an arc $(B, s_n) \longrightarrow (A, s_n)$. As A does not dominate B, there is an s_i with $(A, s_i) \longrightarrow (B, s_i)$. But then the 2×2 subgame $\{A, B\} \times \{s_i, s_n\}$ has the response graph of 2×2 Coordination, which is impossible in a preference-zero-sum game, by (Biggar and Shames, 2022, Corollary 4.7). ■

Lemma A.9 *No zero-sum game with an interior equilibrium has a response graph isomorphic to the Inner or Outer Diamond graphs.*

Proof Suppose for contradiction that p is an interior Nash equilibrium of a game $u = (u_1, u_2)$ with the response graph isomorphic to the Inner Diamond. Then by Papadimitriou and Piliouras (2016), all profiles must be chain recurrent. However by Lemma 5.4 the sink of the Inner Diamond graph is an attractor, so cannot be chain equivalent to all other points by Lemma 3.9. For the Outer Diamond graph, observe similarly that the source is a repellor, which again cannot be chain equivalent to all other points, by Lemma 3.9. ■