

The Attractor of the Replicator Dynamic in Zero-Sum Games

Oliver Biggar

Iman Shames

CIICADA Lab, Australian National University, Acton, ACT, Australia

OLIVER.BIGGAR@ANU.EDU.AU

IMAN.SHAMES@ANU.EDU.AU

Editors: Claire Vernade and Daniel Hsu

Abstract

In this paper we characterise the long-run behaviour of the replicator dynamic in zero-sum games (symmetric or non-symmetric). Specifically, we prove that every zero-sum game possesses a unique global replicator attractor, which we then characterise. Most surprisingly, this attractor depends *only* on each player’s preference order over their own strategies and not on the cardinal payoff values, defined by a finite directed graph we call the game’s *preference graph*. When the game is symmetric, this graph is a tournament whose nodes are strategies; when the game is not symmetric, this graph is the game’s response graph. We discuss the consequences of our results on chain recurrence and Nash equilibria.

Keywords: Replicator dynamic, attractors, learning in games, zero-sum games

1. Introduction

Learning in the presence of other learning agents is an increasingly fundamental topic in modern machine learning, motivated by its role at the core of cutting-edge techniques like *learning from self-play* (Silver et al., 2016, 2018) and *Generative Adversarial Networks* (Goodfellow et al., 2020). The challenge of these systems is analysing their collective behavior, which is where learning theory intersects with game theory. To quote Hofbauer and Sigmund (1998), “a major task of game theory [is] to describe the dynamical outcome of model games described by strategies, payoffs and adaptive mechanisms.” That is, when agents learn collectively, what do they learn?

In online learning, the best-known approaches use variants of the Multiplicative Weights Update algorithm (MWU) (Arora et al., 2012). To achieve the ‘no-regret’ property, these algorithms typically decrease the step size as more samples are observed. In the long-run, as the step size becomes small, the behavior of a collection of MWU-playing agents converges to the flow of a famous differential equation: the *replicator dynamic* (Taylor and Jonker, 1978). This model was originally inspired by biological models of evolution (Smith and Price, 1973), and is a central object of study in *evolutionary game theory*, the subfield of game theory which focuses on dynamic processes. Since its discovery, the replicator has been extensively analysed by biologists, mathematicians, economists and computer scientists (Hofbauer and Sigmund, 2003; Sandholm, 2010; Kleinberg et al., 2011). Indeed, just as MWU is the flagship algorithm in online learning, the replicator is the flagship dynamic in evolutionary game theory (Sandholm, 2010).

Describing the ‘dynamical outcome’ of games under the replicator dynamic (and hence MWU) involves answering a basic question: *to which strategy profiles do we converge over time?* In dynamical systems, a system’s long-run behavior is defined by its *attractors* (Strogatz, 2018). Attractors are sets of points which are *invariant* (points inside the set remain there for all time), *asymptotically stable* (points in some neighbourhood converge to the set) and *minimal* (they do not




Figure 1: The preference graphs (Definition 3.2) of two zero-sum games: (left) Rock-Paper-Scissors (symmetric), and (right) Matching Pennies (non-symmetric).

contain a smaller set with the first two properties). Understandably, identifying the attractors of the replicator is one of the core questions of evolutionary game theory, and increasingly also algorithmic game theory (Zeeman, 1980; Hofbauer and Sigmund, 2003; Papadimitriou and Piliouras, 2019; Omidshafiei et al., 2019; Vlatakis-Gkaragkounis et al., 2020, see Section 2). The broadest takeaway of this line of work is that a strategy profile is an attractor *if and only if* it is a *strict (pure) Nash equilibrium* (Sandholm, 2010). However, this result only describes the simplest attractors—those which contain only a single point. Most games—especially zero-sum games—don’t have any pure Nash equilibria, and most attractors contain more than one point! Instead, the trajectories of the replicator in zero-sum games are typically periodic (Mertikopoulos et al., 2018), and under MWU they are often *chaotic* (Cheung and Piliouras, 2019, 2020).

We conclude that, despite four decades of research, **the attractors of the replicator remain largely unknown**, even in zero-sum games, arguably the best-studied special case (Akin and Losert, 1984; Hofbauer, 1996; Hofbauer and Sigmund, 1998, 2003; Piliouras and Shamma, 2014, see Section 2). This is what we achieve in this paper: we characterise the attractors of the replicator dynamic in every zero-sum game. Beyond this result, our concepts and techniques shed new light on equilibria, the graph structure of games (Biggar and Shames, 2023a,b) and the modelling of payoffs/losses in economics and machine learning.

Characterising the replicator attractors of zero-sum games is a valuable development, but we believe the most remarkable aspect of this result is the form that this attractor takes. Specifically, the attractor depends *only* on players’ discrete preferences over their strategies, and not on the cardinal payoff values. These ‘preferences’ are captured in a directed graph called the *preference graph* of the game. In non-symmetric games this graph coincides with the game’s *response graph* (Papadimitriou and Piliouras, 2019)¹, an object which has been of increasing interest in algorithmic game theory, particularly in relation to the replicator (Candogan et al., 2011; Omidshafiei et al., 2019; Biggar and Shames, 2023a,b). The nodes of the preference graph are the profiles of the game, and the arcs represent which strategies players prefer, given the strategies of the other player. As an example, consider Figure 1, which shows the preference graph of the Matching Pennies game. In this game, player 1 prefers to match the choice of player 2, and player 2 prefers to mismatch player 1. The arc $(T, H) \longrightarrow (H, H)$, for example, captures the fact that, given player 2 plays Heads, player 1 ‘prefers’ Heads over Tails. In symmetric zero-sum games, like Rock-Paper-Scissors (Figure 1), the preference graph has an even simpler form where each arc represents the preferred option between some pair of strategies. For example, given Rock ‘beats’ Scissors, in a match-up of Rock and Scissors, one prefers to play Rock, hence the arc $\text{Scissors} \longrightarrow \text{Rock}$. **Conceptually, the**

1. We use the name *preference graph* to unify the symmetric and non-symmetric cases, and because we find the name “response graph” can be confusing. The word ‘response’ suggests a *repeated* or *extensive-form* game, when actually the name is just a contraction of “better-response relation”, an ordering which defines player’s *preferences*.




Figure 2: A zero-sum game (left) and its associated preference graph (right). The sink component of the graph consists of all profiles other than (a, a) . By Theorem 4.3, this game’s unique replicator attractor is the *content* (Definition 4.1) of the sink component, which is the union of the strategy spaces of the subgames $\{a, b, c\} \times \{b, c\}$ and $\{b, c\} \times \{a, b, c\}$, represented by the shaded region on the graph. Note that the strategy space of the game is 4-dimensional, with the attractor a 3-dimensional region on the boundary.

preference graph stores the underlying combinatorial structure of the game. Most game-theoretic concepts—including the replicator dynamic and the Nash equilibrium—are defined by cardinal payoffs, which serve as a numerical instantiation of the underlying preference structure. Our result shows that the choice of representation of preferences by numbers has a transient effect: two games with different payoffs but the same preferences have the same long-run behavior, in that their attractors are identical. One practical consequence is that computing the attractor is easy (we can do it by traversing the preference graph). More fundamentally, this lends our prediction stability in the face of uncertainty in our model, an important and rare property in game theory (Von Neumann and Morgenstern, 1944).

1.1. Contributions

The main result of the paper is Theorem 4.3, which characterises the attractors of zero-sum games. For each zero-sum game, we prove that an attractor exists, is unique and attracts all points on the interior of the game. This attractor is precisely the *content* (Definition 4.1, see Biggar and Shames, 2023b) of the preference graph’s unique² *sink connected component*, which is a strongly connected component with no arcs from a node inside the component to a node outside. The sink component is a set of pure profiles; its content is the set of *mixed* profiles whose support contains only profiles in this component, which is always an invariant set under the replicator. See Figure 2. A recent result (Biggar and Shames, 2023b) demonstrated that every replicator attractor contains the content of some sink component. The challenging part of our proof is showing that the content of the sink component is asymptotically stable in any zero-sum game, and so is itself an attractor. We demonstrate stability using a potential function argument. Our choice of function derives from the preference graph: specifically, we use the total probability mass over all sink component profiles. The proof then separates into two cases, reflecting two standard types of zero-sum game:

2. Uniqueness of the sink connected component is a property of zero-sum games (Biggar and Shames, 2023a). Non-zero-sum games, such as the 2×2 Coordination game, may have preference graphs with multiple sink components, and thus can have multiple attractors under the replicator dynamic (Biggar and Shames, 2023a,b).

symmetric and non-symmetric. In evolutionary game theory these are often called *single*³ and *multiple-population* games. The replicator has different properties in each case. In symmetric zero-sum games the preference graph possesses a simplified form, which makes the remainder of the proof straightforward. The non-symmetric case is much more complex. Here we prove a novel transformation of the replicator dynamic (Theorem 4.8): the flow of the two-population replicator on a zero-sum game embeds in the flow of the single-population replicator on a larger symmetric zero-sum game, known as its *von Neumann symmetrisation* (Definition 4.6). We believe this result (which also works for MWU) to be of independent interest. While symmetrisation sets up the proof, a further step is required because the sink component of the von Neumann symmetrisation may be larger than that of the original game. Lemma 4.9 completes the proof using the fact that the dynamics are embedded on the subspace of product distributions. In Section 5 we discuss some consequences of our result. Lemma 5.4 shows an interesting connection between the preference graph and Nash equilibria in zero-sum games: the support of the equilibrium must be strongly connected as a subgraph of the preference graph and wholly contained within the unique sink connected component. This has important consequences for characterising *sink chain components* (see Section 2).

2. Related Work

Long-run stability of strategy profiles under the replicator is a fundamental topic in evolutionary game theory, especially with regard to classical solution concepts, such as Nash equilibria and *evolutionarily stable strategies*. See Sandholm (2010) for a summary. A particularly relevant work is that of Eshel et al. (1983), who studied zero-sum games under the replicator, proving a crucial *volume conservation* property, analysed in depth by Hofbauer (1996). Many papers since (such as Piliouras and Shamma, 2014; Vlatakis-Gkaragkounis et al., 2020; Biggar and Shames, 2023b, and this paper) have used this property to bound asymptotically stable sets of the replicator. Zeeman (1980) performed an early study of replicator attractors, suggesting that the qualitative behaviour of the replicator can be split into a finite number of classes; we extend this by showing that in zero-sum games the qualitative behaviour is defined by the preference graph alone. Ritzberger and Weibull (1995) showed that when a *subgame* is closed under “weakly better responses”, then it is asymptotically stable under the replicator. The preference graph is defined by the weakly better responses, and so this follows as a special case of our result when a connected component of the preference graph is a subgame. Finally, one of the most famous of these classical results proves that the *time-average* of the replicator (and MWU) converges to the Nash equilibrium in zero-sum games (Freund and Schapire, 1999; Hofbauer et al., 2009). However, the time-average behavior is distinct from the day-to-day or *last-iterate* behavior (Papadimitriou and Piliouras, 2019).

Despite these efforts, a general negative conclusion of this line of work is that the replicator doesn’t converge to mixed equilibria (Sandholm, 2010), and moreover no dynamic can converge to equilibria in all games (Hart and Mas-Colell, 2003; Benaïm et al., 2012). Instead, its behavior can be chaotic (Sato et al., 2002). Further, finding equilibria is generally computationally intractable (Daskalakis et al., 2009). Consequently, the algorithmic game theory and learning communities have increasingly shifted towards new notions of dynamical outcome which can predict the day-to-day behavior of computational agents in games (Kleinberg et al., 2011; Papadimitriou and Piliouras, 2019).

3. Single-population games are also called *population games* (Sandholm, 2010) or *matrix games* (Andrade et al., 2021).

To this end, *sink chain components* were recently proposed (Papadimitriou and Piliouras, 2019) as the outcome of dynamic games, with the replicator used as the motivating example. Sink chain components are built on a concept called *chain recurrence* (Definition 5.1), a generalisation of periodicity which forms the foundation of the Fundamental Theorem of Dynamical Systems (Conley, 1978). Crucially, chain components are grounded in computational considerations. To quote Papadimitriou and Piliouras, informally, a “point x is chain recurrent if, whenever Alice starts at x , Bob can convince her that she is on a cycle by manipulating the round off error of her computer—no matter how much precision Alice brings to bear.” Interestingly, this computational solution concept brings us back to the classical notion of an attractor: when a replicator attractor exists, it is a sink chain component (Lemma 5.2). Thus, finding the attractors of the replicator dynamic is motivated not only by dynamical systems but also computer science: the attractors give us the strongest prediction of long-run behavior which is consistent with reliable computation.

This chain recurrence approach has inspired a number of new results on games and the replicator dynamic (Omidshafiei et al., 2019; Biggar and Shames, 2023a,b). In zero-sum games, when a *fully-mixed* Nash equilibria exists, the behavior is essentially unpredictable: the sink chain component is the whole game (Papadimitriou and Piliouras, 2016, 2018). Under MWU, we observe chaotic behavior in these games (Cheung and Piliouras, 2019, 2020). Further, when a fully-mixed NE does not exist, Piliouras and Shamma (2014) showed that all fully-mixed strategy profiles converge to the subgame containing the equilibrium, called the *essential subgame*, and within this subgame all profiles are chain recurrent (Papadimitriou and Piliouras, 2016). Surprisingly, however, these results do not characterise the chain components of zero-sum games—the essential subgame is generally *not* a sink chain component. The reason is that the convergence to the essential subgame is not *uniform*, and so the essential subgame is typically *not* asymptotically stable (see Section 5), one of the defining properties of attractors and sink chain components (Alongi and Nelson, 2007) and a key property for predicting learning processes (Vlatakis-Gkaragkounis et al., 2020; Omidshafiei et al., 2019). Instead, interior profiles which are *arbitrarily close* to the essential subgame may move far away before returning. The question remains unanswered: what are the replicator sink chain components/attractors of zero-sum games?

The story of predicting the replicator seems generally negative: the replicator may be chaotic, may not converge, or may only converge in time-average or non-uniformly. Our results tell a different, more positive story: we characterise the attractor/sink chain component of the replicator, which, while larger than the essential subgame, is the smallest outcome which is consistent with computation, in the sense of chain recurrence. What’s more, being defined by discrete preferences, the solution is natural and robust. As an example, in Figure 2, shifting the payoff for the profile (a, c) from $(3, -3)$ to $(4, -4)$ moves the essential subgame from $\{b, c\} \times \{b, c\}$ to $\{a, c\} \times \{b, c\}$. However, all preferences remain unchanged, so the graph and hence the attractor do not change. Shifting focus from the Nash equilibrium also gives a new perspective on chain recurrence. Prior approaches (Papadimitriou and Piliouras, 2016; Mertikopoulos et al., 2018) suggested a connection between equilibria and chain recurrence in zero-sum games. We find instead that chain recurrence in zero-sum games is *entirely* defined by the preference graph (Lemma 5.3). The previous findings are now explained by a non-trivial connection between equilibria and the preference graph: the existence of a fully-mixed equilibrium implies strong connectedness of the preference graph (Lemma 5.4). See Section 5.

3. Preliminaries

In game theory, a game is defined by a triple consisting of the *players*, *strategy sets* for each player, and *payoffs*. A combination of strategies for each player is called a *strategy profile* or simply a profile, and for each profile there is a real-valued payoff to each player. In this paper we focus on two-player games, where we denote the players by the integers 1 and 2 and their strategy sets by S_1 and S_2 . The strategy names are simply labels, so we assume $S_1 = [n] := \{i \in \mathbb{N}_0 \mid i < n\}$ and $S_2 = [m] := \{i \in \mathbb{N}_0 \mid i < m\}$. The profiles are the pairs $S_1 \times S_2$. We call this an $n \times m$ game. An $n \times m$ game is defined by a pair (A, B) of matrices, $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{m \times n}$, representing the payoffs to players 1 and 2 respectively. A game is *symmetric* if $A = B$. Intuitively, a game is symmetric if the payoff is determined solely by the choice of strategies, and not identity of the player. We focus on *zero-sum games*, which we represent by a single matrix $M \in \mathbb{R}^{n \times m}$, implicitly assumed to be the payoffs for the first player, which defines the game $(M, -M^T)$. That is, in a profile (s_1, s_2) , player 1 receives M_{s_1, s_2} and player 2 receives $(-M^T)_{s_2, s_1} = -M_{s_1, s_2}$. A zero-sum game is also a symmetric game if $M = -M^T$, that is, M is *anti-symmetric*. Thus, there is a natural one-to-one correspondence between anti-symmetric real matrices and symmetric zero-sum games. A *subgame* of a game is formed by choosing subsets $T_1 \subseteq S_1$ and $T_2 \subseteq S_2$ of each player's strategy sets and restricting the game to the profiles in $T_1 \times T_2$. We typically represent a subgame by its product set of profiles $T_1 \times T_2$.

A *mixed strategy* is a distribution over a player's 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. If x is a mixed strategy, and s a strategy, we write x_s for the s -entry of x . Like profiles, we denote mixed profiles by pairs (x, y) where x and y are mixed strategies for the first and second player, respectively. The *support* of a mixed strategy x , written $\text{support}(x)$, is the set of strategies s where x_s is non-zero. The *support* of a mixed profile (x, y) is the Cartesian product $\text{support}(x) \times \text{support}(y)$, the set of profiles whose strategies are in the support of x and y respectively. As distributions over a finite set, mixed strategies can be naturally embedded in the standard probability simplex in Euclidean space, by choosing some arbitrary ordering of the strategies in S_1 and S_2 . We denote these spaces by $\Delta(S_1)$ and $\Delta(S_2)$. The set of mixed profiles is the product $\Delta(S_1) \times \Delta(S_2)$, which we call the *strategy space* of the game, and we refer to mixed profiles as 'points' in strategy space. The strategy space is also naturally embedded in Euclidean space, so we can talk about geometric properties of sets of mixed profiles. The payoffs extend to mixed profiles using expectation. In a zero-sum game $(M, -M^T)$, the *expected payoff* of a mixed profile (x, y) is $y^T M x$ to player 1 and $x^T (-M^T) y = -(y^T M x)$ to player 2.

In evolutionary game theory, we think of an $n \times m$ game as a pair of 'populations', with the strategies as 'types' within each population. A mixed strategy x represents the distribution of types in the population. Symmetric games, called *single-population games* in this context, are viewed as having a single underlying population, with types from the strategy set S . See [Sandholm \(2010\)](#). Viewed as a two-player game, all mixed profiles in a single-population game are of the form (x, x) , because there is only one underlying population. This leads to an important terminology convention: *in symmetric games, (mixed) profiles are the same as (mixed) strategies*. Intuitively, this is because a profile is a choice of strategy for each player, and symmetric games have 'only one player'. The strategy space of the game consists of the symmetric profiles (x, x) in $\Delta(S) \times \Delta(S)$. We write the mixed profile (x, x) as the mixed strategy x , and write simply $\Delta(S)$ for the strategy space of the symmetric game.

3.1. Preference graphs

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*. In symmetric games, profiles and strategies are the same, and we define all profiles to be comparable. Up to strategic equivalence, the game is defined by the *payoff differences* between comparable profiles (Candogan et al., 2011). We store this in a matrix we call the *weight matrix* W of the game.

Definition 3.1 *Let M be a zero-sum game, and let p and q be comparable profiles. If M is a symmetric game, then profiles and strategies are the same, and we define $W_{p,q}$ to be the same as $M_{p,q}$. If M is non-symmetric, then $p = (p_1, p_2)$ and $q = (q_1, q_2)$ and*

$$W_{p,q} = \begin{cases} M_{p_1,p_2} - M_{q_1,q_2} & p_2 = q_2 \text{ (i.e., } p \text{ and } q \text{ are 1-comparable)} \\ M_{q_1,q_2} - M_{p_1,p_2} & p_1 = q_1 \text{ (i.e., } p \text{ and } q \text{ are 2-comparable)} \end{cases}$$

If p and q are not comparable, then $W_{p,q}$ is undefined.

We deliberately leave the payoff differences between incomparable profiles undefined, so that it is clear to the reader that we will only reference W when the associated profiles are comparable. Note that for any comparable profiles p and q , $W_{p,q} = -W_{q,p}$.

Definition 3.2 *Let M be a zero-sum game. The preference graph of M is the graph whose nodes are the profiles of the game and where there is an arc $p \longrightarrow q$ between profiles p and q if and only if they are comparable and $W_{p,q} \leq 0$ (equivalently, $W_{q,p} \geq 0$).*

While the definition is the same for symmetric and non-symmetric games, the resultant graphs are not the same, because the weight matrix is defined differently. In symmetric zero-sum games all profiles are comparable, so the preference graph is a *tournament*⁴ (a directed graph with an arc between every pair of nodes). In non-symmetric games the preference graph is the game's response graph (Biggar and Shames, 2023a), which is never a tournament because not all profiles are comparable. We think of the entries in the weight matrix as being weights on the associated arc, as in $p \xrightarrow{|W_{p,q}|} q$.

3.2. Dynamical Systems and the Replicator

The replicator dynamic is a continuous-time dynamical system (Sandholm, 2010), defined by an ordinary differential equation. Let $x \in \Delta(S_1)$ and $y \in \Delta(S_2)$ be mixed strategies, and let $s \in S_1$ and $t \in S_2$ be pure strategies. Then, for a (non-symmetric) zero-sum game M we have

Definition 3.3 (Non-Symmetric Zero-Sum Replicator Equation)

$$\begin{aligned} \dot{x}_s &= x_s((My)_s - x^T My) \\ \dot{y}_t &= -y_t((M^T x)_t - x^T My) \end{aligned}$$

4. In the degenerate case where $W_{p,q} = 0$, there is a pair of arcs 0-weighted arcs between nodes p and q , which makes the preference graph not, strictly speaking, a tournament, but this will cause no problems.

In a symmetric game M , the replicator is defined a similar way:

$$\dot{x}_s = x_s((Mx)_s - x^T Mx)$$

If M is also zero-sum, then $x^T Mx = 0$ (by anti-symmetry), and this reduces to

Definition 3.4 (Symmetric Zero-Sum Replicator Equation)

$$\dot{x}_s = x_s(Mx)_s$$

Note that the symmetric and non-symmetric replicator, while similar, are distinct equations with different properties. The solutions to these equations define a *flow* (Sandholm, 2010) on the strategy space of the game, which is a function $\phi : X \times \mathbb{R} \rightarrow X$ that is a continuous group action of the reals on X . We call these the symmetric or non-symmetric *replicator flow*, respectively. The forward orbit of the flow from a given point is called a *trajectory* of the system. A set of points Y is called *invariant* under ϕ if $\phi(Y, t) = Y$ for any $t \in \mathbb{R}$.

Definition 3.5 *Let X be a compact space, ϕ a flow on X and A a compact subspace of X . If there is a neighbourhood U of A such that*

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

under any metric \mathbf{d} , then we say A is asymptotically stable. If A is also invariant under ϕ , we call it an attracting set. An attracting set which contains no smaller attracting sets is called an attractor.

There are some differences in terminology in the literature to be wary of. What we call attracting sets are sometimes called attractors (Conley, 1978; Sandholm, 2010; Biggar and Shames, 2023b), in which case what we call an attractor is a *minimal attractor*. Otherwise, our definition is the same as Sandholm (2010); Biggar and Shames (2023b). The set of points which approach an attractor A in the limit $t \rightarrow \infty$ is called the attractor's *basin of attraction*. We call an attractor *global* if its basin of attraction includes all points in $\text{int}(X)$ (Hofbauer and Sigmund, 2003).

4. The Attractor of the Replicator

In this section we prove Theorem 4.3, which characterises the attractor of the replicator in zero-sum games. We begin by noting that any mixed profile x naturally defines a distribution over profiles, with x_p denoting the mass on a profile p . If the game is symmetric, profiles and strategies are the same and so this is trivial. In a non-symmetric game, $x = (x_1, x_2)$ is a pair of mixed strategies, and the distribution is defined by the *product*, with the mass on a profile $p = (p_1, p_2)$ defined by

$$x_p := x_{1p_1} x_{2p_2} \tag{1}$$

This distribution over profiles is used to define an important concept: the *content* of a set of profiles.

Definition 4.1 (Biggar and Shames) *Let H be a set of profiles in a game. The content of H , denoted $\text{content}(H)$, is the set of all mixed profiles x where all profiles in the support of x are in H .*

Equivalently, $x \in \text{content}(H)$ if and only if $x_H := \sum_{h \in H} x_h = 1$, that is, x defines a distribution whose mass is entirely distributed over profiles in H . The content is a union of subgames, and so is an invariant set under the replicator [Sandholm \(2010\)](#). An example is shown in [Figure 2](#). We will show that the unique global attractor is the content of the unique sink component of the preference graph. Uniqueness follows from graph structure, established originally in [Biggar and Shames \(2023a\)](#).

Lemma 4.2 (Uniqueness) *The preference graph of a zero-sum game has a unique sink component.*

The proof of [Lemma 4.2](#) can be found in the appendix. Now we can prove our main theorem.

Theorem 4.3 (The Attractor of the Replicator) *In a (symmetric or non-symmetric) zero-sum game M , the content of the unique sink component H of its preference graph is the unique global attractor of the (respectively symmetric or non-symmetric) replicator dynamic.*

Proof Proving [Theorem 4.3](#) requires showing (i) $\text{content}(H)$ is an invariant set (*invariance*), (ii) every attracting set contains $\text{content}(H)$ (*minimality*) and (iii) $\text{content}(H)$ is asymptotically stable and its basin of attraction contains $\text{int}(X)$ (*global asymptotic stability*). We establish (i) and (ii) in [Lemma 4.4](#).

Lemma 4.4 (Invariance and Minimality) *If H is the sink component of the preference graph of a (symmetric or non-symmetric) zero-sum game M , then $\text{content}(H)$ is invariant under the replicator. Further, for any attracting set A , $\text{content}(H) \subseteq A$.*

[Lemma 4.4](#) is a streamlined presentation of existing results (particularly [Theorem 5.2 of Biggar and Shames, 2023b](#)), so we defer its proof to the appendix. The challenge and main contribution of [Theorem 4.3](#) lies in (iii): showing global asymptotic stability of $\text{content}(H)$. We do this by demonstrating that x_H , the total mass on the sink component H , increases over time ([Lemmas 4.5 and 4.9](#)). That is, $\dot{x}_H = \frac{d}{dt}x_H = \sum_{h \in H} \frac{d}{dt}x_h = \sum_{h \in H} \dot{x}_h > 0$, for any $x \in \text{int}(X)$. The function x_H is a natural choice, because x_H is uniformly continuous, bounded in $[0, 1]$ and $x_H = 1$ if and only if $x \in \text{content}(H)$, so x_H can be thought of as a metric for the distance between x and the content. Showing $\dot{x}_H > 0$ requires different arguments for the symmetric and non-symmetric cases. The symmetric case is straightforward, and we complete it in [Lemma 4.5](#) below.

Lemma 4.5 *In a symmetric zero-sum game M , under the symmetric replicator ([Definition 3.4](#)), if $x_H \in (0, 1)$ then $\dot{x}_H > 0$.*

Proof First, $\dot{x}_H = \sum_{h \in H} x_h(Mx)_h$. As M is anti-symmetric, for any profiles h and q we have $x_h x_q M_{h,q} + x_q x_h M_{q,h} = 0$. Hence,

$$\dot{x}_H = \sum_{q \in S} \sum_{h \in H} x_q x_h M_{h,q} = \sum_{q \in H} \sum_{h \in H} x_q x_h M_{h,q} + \sum_{q \notin H} \sum_{h \in H} x_q x_h M_{h,q} = \sum_{q \notin H} \sum_{h \in H} x_q x_h M_{h,q}$$

For any $q \notin H$ and $h \in H$, as H is a sink component, there is an arc $q \longrightarrow h$ in the preference graph with strictly positive weight. That is, $M_{h,q} > 0$. Hence the summands are nonnegative. Finally, because $x_H \in (0, 1)$, there must exist a $q \notin H$ and an $h \in H$ such that $x_q x_h > 0$. Thus, $\dot{x}_H > 0$. ■

The more challenging case of the proof is showing that $\dot{x}_H > 0$ in non-symmetric games (Lemma 4.9). We solve this using Theorem 4.8, which establishes that we can embed the flow of the non-symmetric replicator on our non-symmetric game M into the flow of the symmetric replicator on a larger symmetric game, M 's *von Neumann symmetrisation*.

Definition 4.6 (Von Neumann Symmetrisation) *Let M be a zero-sum game, with $M \in \mathbb{R}^{n \times m}$. The von Neumann symmetrisation of M , written \mathcal{S}_M , is defined as the following $\mathbb{R}^{(nm) \times (nm)}$ matrix, which we index by profiles $p = (p_1, p_2)$ and $q = (q_1, q_2)$:*

$$(\mathcal{S}_M)_{p,q} = M_{p_1,q_2} - M_{q_1,p_2}.$$

The strategy space of \mathcal{S}_M is $S_1 \times S_2$. That is, the *strategy profiles* of the original game M become *strategies* of the von Neumann symmetrisation. Mixed profiles likewise become *mixed strategies*, using the production distribution as in equation (1): $x_p := x_{1p_1}x_{2p_2}$.

This construction appeared first in Gale et al., who attributed it to John von Neumann, hence the name. It possesses an intuitive explanation: it results from holding two plays of the original game M , with each player taking the role of the row player 1 in one play and player 2 in the other (Gale et al., 1950). The von Neumann symmetrisation has two important properties: (1) it is anti-symmetric, and so is a symmetric zero-sum game, and (2) it can be viewed as an extension of the weight matrix W , previously only defined on comparable profiles, to a relation over all profiles. That is, if p and q are comparable then $(\mathcal{S}_M)_{p,q} = W_{p,q}$. More generally,

Lemma 4.7 *Let $p = (p_1, p_2)$ and $q = (q_1, q_2)$ be profiles. Then:*

$$(\mathcal{S}_M)_{p,q} = W_{p,(p_1,q_2)} + W_{p,(q_1,p_2)} = W_{(p_1,q_2),q} + W_{(q_1,p_2),q}.$$

In Theorem 4.8, we establish a novel link between the replicator flow on a zero-sum game and its von Neumann symmetrisation.

Theorem 4.8 (Symmetrising the Replicator Dynamic) *Let M be a non-symmetric zero-sum game. Let $x = (x_1, x_2)$ be a mixed profile and $p = (p_1, p_2)$ a pure profile. Write $x_p := x_{1p_1}x_{2p_2}$ as in equation (1). Then, under the non-symmetric replicator (Definition 3.3),*

$$\dot{x}_p = x_p(\mathcal{S}_M x)_p.$$

In English, Theorem 4.8 states that the flow of the non-symmetric replicator (Definition 3.3) on M embeds (on the subspace of product distributions) in the flow of the symmetric replicator (Definition 3.4) on \mathcal{S}_M . The proofs of Theorem 4.8 and Lemma 4.7 are straightforward and can be found in the appendix. Using Theorem 4.8, we can complete the proof of Theorem 4.3 on non-symmetric games by showing that x_H is increasing in $x_H \in (0, 1)$.

Lemma 4.9 *In a non-symmetric zero-sum game M , under the non-symmetric replicator (Definition 3.3), if $x_H \in (0, 1)$ then $\dot{x}_H > 0$.*

Proof If $h = (h_1, h_2)$ is a pure profile, then by Theorem 4.8, $\dot{x}_h = \frac{d}{dt}(x_{1h_1}x_{2h_2}) = x_h(\mathcal{S}_M x)_h$. Because \mathcal{S}_M is symmetric, by the same argument as in Lemma 4.5, we can show that

$$\dot{x}_H = \sum_{q \notin H} \sum_{h \in H} x_q x_h (\mathcal{S}_M)_{h,q}.$$

Now pick some profiles $q \notin H$ and $h \in H$ with $x_q x_h > 0$. As in Lemma 4.5, because $x_H \in (0, 1)$, at least one such pair exist. We will show that the sum above is strictly positive. Firstly, observe that if q and h are comparable, the arc $q \longrightarrow h$ of the preference graph of M goes from q to h with strictly positive weight, because $q \notin H$ and $h \in H$, and so $(\mathcal{S}_M)_{h,q} = W_{h,q} > 0$. Otherwise, suppose that $q = (q_1, q_2)$ and $h = (h_1, h_1)$ are not comparable, and let $a = (q_1, h_2)$ and $b = (h_1, q_2)$. We have the following three cases:

1. $a, b \in H$. Then the arcs $q \longrightarrow b$ and $q \longrightarrow a$ are directed towards a and b because q is outside the sink component H . By Lemma 4.7, $(\mathcal{S}_M)_{h,q} = W_{a,q} + W_{b,q} > 0$.
2. $a, b \notin H$. Then the arcs $a \longrightarrow h$ and $b \longrightarrow h$ are directed towards h because a and b are outside the sink component. By Lemma 4.7, $(\mathcal{S}_M)_{h,q} = W_{h,a} + W_{h,b} > 0$.
3. $a \in H, b \notin H$ (the case $b \in H, a \notin H$ is identical). The sum $\sum_{q \notin H} \sum_{h \in H} x_q x_h (\mathcal{S}_M)_{h,q}$ includes the terms $x_q x_h (\mathcal{S}_M)_{h,q}$ and $x_b x_a (\mathcal{S}_M)_{a,b}$. However, $x_q x_h = (x_{1q_1} x_{2q_2})(x_{1h_1} x_{2h_2}) = (x_{1q_1} x_{2h_2})(x_{1h_1} x_{2q_2}) = x_a x_b$, and so $x_q x_h (\mathcal{S}_M)_{h,q} + x_b x_a (\mathcal{S}_M)_{a,b} = x_q x_h ((\mathcal{S}_M)_{h,q} + (\mathcal{S}_M)_{a,b}) = x_q x_h (W_{a,q} + W_{b,q} + W_{q,b} + W_{h,b}) = x_q x_h (W_{a,q} + W_{h,b})$ by Lemma 4.7. Since a and h are inside the sink component and q and b are outside, and the arcs $q \longrightarrow a$ and $b \longrightarrow h$ must be directed into the component, so $W_{a,q} > 0$ and $W_{h,b} > 0$ and thus $x_q x_h (\mathcal{S}_M)_{h,q} + x_b x_a (\mathcal{S}_M)_{a,b} > 0$.

Overall, we conclude that $\dot{x}_H > 0$. ■

Finally, asymptotic stability of $\text{content}(H)$ follows easily from the fact that x_H is uniformly continuous. Pick any $0 < \alpha < \beta < 1$. Then for any x with $x_H \in [\alpha, \beta]$, $\dot{x}_H > \epsilon$ for some ϵ (uniform continuity), so after some finite time $x_H > \beta$. Repeating this argument for any α, β shows that $\text{content}(H)$ is asymptotically stable. ■

5. Chain Recurrence and Nash Equilibria

In this section we discuss some consequences of Theorem 4.3. The first concerns chain recurrence, which is defined by (ϵ, T) -chains.

Definition 5.1 ((ϵ, T) -Chains, Alongi and Nelson) *Let ϕ be a flow on a compact metric space X , with x and y in X . An (ϵ, T) -chain from x to y is a finite sequence of points x_1, x_2, \dots, x_n with $x = x_1$ and $y = x_n$, and times $t_1, \dots, t_n \in [T, \infty)$ such that $\mathbf{d}(\phi(x_i, t_i), x_{i+1}) < \epsilon$. 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 .*

A point is called *chain recurrent* if it has a pseudo-orbit to itself. Two points are *chain equivalent* if there are pseudo-orbits between them in both directions, and equivalent chain recurrent points are grouped in topologically connected components called *chain components* (Alongi and Nelson, 2007). Reachability under pseudo-orbits provides an ordering on the chain components, and sink chain components are those which are minimal in this order. Sink chain components have been increasingly studied in algorithmic game theory (Papadimitriou and Piliouras, 2016, 2018, 2019). A connection between the preference graph and chain components of the replicator was demonstrated by Biggar and Shames (2023b), who proved that replicator sink chain components always exist. However, sink

chain components have not generally been characterised. Attractors, when they exist, are sink chain components, so our results present the first characterisation of sink chain components of zero-sum games.

Lemma 5.2 (Folklore, see the appendix) *In any flow, every attractor is a sink chain component.*

Lemma 5.3 *The content of the sink component of the preference graph is the unique sink chain component of a zero-sum game.*

Lemma 5.3 follows directly from Theorem 4.3 combined with Lemma 5.2. The proof of both lemmas can be found in the appendix. In zero-sum games, previous state-of-the-art results on chain recurrence made use of the Nash equilibrium (Piliouras and Shamma, 2014; Papadimitriou and Piliouras, 2016; Mertikopoulos et al., 2018). This line of inquiry established first (1) that all points in the essential subgame are contained within the sink chain component. Secondly (2), if the essential subgame is not the whole game, all interior starting points converge to the essential subgame. This suggests a connection between chain recurrence and equilibria in zero-sum games. Lemma 5.3 comes to a different conclusion; it proves that sink chain components are characterised solely by the preference graph. Lemma 5.4 resolves this seeming discrepancy: chain components are determined by the preference graph, but the presence of an equilibrium in a subgame forces some structure on the induced preference graph of that subgame—in particular, it must be strongly connected and contained within the sink component of the whole game’s preference graph.

Lemma 5.4 *In a zero-sum game, any subgame with a Nash equilibrium in its interior must (i) be contained within the sink component of the preference graph and (ii) the subgraph which this subgame induces in the preference graph must be strongly connected.*

Proof (i) Assume the attractor is not the whole game, in which case it is on the boundary. By Theorem 4.3, all interior points are in the basin of attraction, and by Theorem 3.4 of Piliouras and Shamma (2014), trajectories starting from interior points converge to the essential subgame in the limit, and so the essential subgame must be within the attractor. (ii) Each subgame is independent, so we can assume we are working with the whole game and the equilibrium x is fully-mixed. For contradiction, assume the preference graph is not strongly connected. Consequently, there is an attractor A on the boundary of the strategy space (Theorem 4.3). The point x is in the basin of attraction of A , so x converges to A , which contradicts the fact that x is a Nash equilibrium, which are fixed points under the replicator (Sandholm, 2010). ■

This Lemma rephrases our understanding of equilibria and chain recurrence, highlighting the key role of the preference graph. As an example, Lemma 5.3 implies that sink chain component is the whole game *if and only if* the preference graph is strongly connected. When a fully-mixed equilibrium exists, the preference graph must be connected, and so the sink chain component is the whole game. Most interestingly, even though we prove Lemma 5.4 using the replicator dynamic, this lemma is a *purely game-theoretic result* which only relates equilibria and the preference graph. This connection between the preference graph and equilibria is useful for analysing games. For instance, examining Figure 2 we find the Nash equilibrium is $((0, 0.5, 0.5), (0, 0.5, 0.5))$ (both players play b and c half the time), which has support $\{b, c\} \times \{b, c\}$. As Lemma 5.4 predicts, this is within the attractor and its induced preference graph (a 4-cycle) is strongly connected.

6. Conclusions and Future Work

In this paper we gave the first characterisation of the unique attractor of the replicator in zero-sum games, thereby describing the long-term behaviour of the dynamic in these games. As a secondary result, we have demonstrated the importance of the preference graph as a tool for analysing game dynamics. In particular, in the long run, only the *preferences* dictates the outcome of the game. This surprising conclusion has potentially significant consequences for modelling strategic interactions. In game theory—especially evolutionary game theory—precise knowledge of utilities is difficult to achieve; our results show that modelling only the preferences for each player is sufficient to characterise the attractor of the replicator dynamic.

An important goal for future research is to characterise the attractors of the replicator in *all* games, not just zero-sum ones. Currently, the attractors of the replicator have been characterised in only a few classes of games, such as potential games and $2 \times n$ games without dominated strategies (Biggar and Shames, 2023b), as well as some individual games, like the Asymmetric Cyclic Matching Pennies game from (Kleinberg et al., 2011). Our paper adds zero-sum games to this list of solved classes.

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.
- 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. *Scientific Reports*, 13(1):1833, 2023a.
- Oliver Biggar and Iman Shames. The replicator dynamic, chain components and the response graph. In *International Conference on Algorithmic Learning Theory*, pages 237–258. PMLR, 2023b.
- 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.
- Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- David Gale, Harold W Kuhn, and Albert W Tucker. On symmetric games. In *Contributions to the Theory of Games (AM-24), Volume I*, pages 81–88. Princeton University Press, 1950.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- 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 games and population dynamics*. Cambridge university press, 1998.
- Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American mathematical society*, 40(4):479–519, 2003.
- Josef Hofbauer, Sylvain Sorin, and Yannick Viossat. Time average replicator and best-reply dynamics. *Mathematics of Operations Research*, 34(2):263–269, 2009.
- 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.
- 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.
- 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.
- 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.
- Klaus Ritzberger and Jörgen W Weibull. Evolutionary selection in normal-form games. *Econometrica: Journal of the Econometric Society*, pages 1371–1399, 1995.
- William H Sandholm. *Population games and evolutionary dynamics*. MIT press, 2010.
- Yuzuru Sato, Eizo Akiyama, and J Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002.
- 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, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan 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.
- J Maynard 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.
- Peter D Taylor and Leo B Jonker. Evolutionary stable strategies and game dynamics. *Mathematical biosciences*, 40(1-2):145–156, 1978.
- Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianas, 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.
- E Christopher Zeeman. Population dynamics from game theory. In *Global theory of dynamical systems*, pages 471–497. Springer, 1980.

Appendix A. Proofs

Lemma A.1 (Lemma 4.2) *The preference graph of a zero-sum game has a unique sink component.*

Proof In the non-symmetric case, the preference graph is the response graph, and the result follows from Theorem 4.10 of Biggar and Shames (2023a). In the symmetric case, the preference graph is a tournament, and all tournaments have one sink component, as they are orientations of complete graphs. ■

Lemma A.2 (Lemma 4.4) *If H is the sink component of the preference graph of a (symmetric or non-symmetric) zero-sum game M , then $\text{content}(H)$ is invariant under the replicator. Further, for any attracting set A , $\text{content}(H) \subseteq A$.*

Proof This proof is largely the same as Theorem 5.2 of Biggar and Shames (2023b), with the addition of the symmetric case. (*Invariance:*) Observe that if $x \in \text{content}(H)$, then the support of x is contained in H , and because all mixed profiles in the subgame $\Delta(\text{support}(x))$ have the same support, $\Delta(\text{support}(x)) \subseteq \text{content}(H)$. It follows that $\text{content}(H)$ is a union of subgames. By Theorem 5.4.7 of Sandholm (2010), all subgames are invariant sets under the replicator, and unions of invariant sets are invariant.

($\text{content}(H) \subseteq A$.) By Theorems 9.1.2 and 9.1.6 of Sandholm (2010), no asymptotically stable set can exist in the interior of the strategy space of a symmetric or non-symmetric zero-sum game. Subgames have the same properties as the whole game under the replicator, so the same is true of all subgames. Dually, no repelling set can exist in the interior of any subgame.

(*Claim:* every attracting set contains a profile.) This follows by induction, using the fact that the replicator dynamic on a subgame has the same properties as on the whole game. In the whole game, an asymptotically stable set intersects the boundary. This intersection with the boundary must also be asymptotically stable in any subgame it intersects on the boundary, and so it intersects the boundary of this smaller subgame, and so on. We conclude that such a set contains a pure profile, the smallest possible subgame. This claim generalises to other dynamics—see Vlatakis-Gkaragkounis et al. (2020), Theorem 4.5.

(*Claim:* every attracting set contains all profiles in H .) An arc $p \longrightarrow q$ of the preference graph is also a subgame, where only the profiles p and q are in the support. The (symmetric or non-symmetric) replicator reduces to $\dot{x}_p = x_p(1 - x_p)W_{q,p}$ on this subgame, where $W_{q,p} \geq 0$ (for the player for which these profiles are comparable). If p is contained in an asymptotically stable set, then q must also be contained in this set, because points near p move to q along this arc. We know that asymptotically stable sets contain a pure profile—by this argument we deduce that they contain all pure profiles reachable from that one in the preference graph. Such a set of profiles always contains the sink component H .

(*Claim:* every attracting set contains $\text{content}(H)$.) Let Y be a subgame, where the pure profiles in Y are in an attracting set A . If Y is a pure profile, then all mixed profiles in Y are in the set, trivially. Now suppose for induction that all points on the boundary of Y are in A . Suppose for contradiction that there is a point $x \in \text{int}(Y)$ that is not in A . The set $A \cap Y$ is attracting in Y , and the boundary is contained in $A \cap Y$, but this means that the dual repelling set of $A \cap Y$ is contained in the interior of Y , but no such sets can be contained in the interior. Hence all of Y is within A . By induction on subgames, we find that all of $\text{content}(H)$ is within every attracting set. ■

Lemma A.3 \mathcal{S}_M is anti-symmetric.

Proof For $p = (p_1, p_2)$ and $q = (q_1, q_2)$, $(\mathcal{S}_M)_{p,q} = M_{p_1,q_2} - M_{q_1,p_2} = -(M_{q_1,p_2} - M_{p_1,q_2}) = -(\mathcal{S}_M)_{q,p}$. ■

Lemma A.4 (Lemma 4.7) Let $p = (p_1, p_2)$ and $q = (q_1, q_2)$ be profiles. Then:

$$(\mathcal{S}_M)_{p,q} = W_{p,(p_1,q_2)} + W_{p,(q_1,p_2)} = W_{(p_1,q_2),q} + W_{(q_1,p_2),q}.$$

Proof

$$\begin{aligned} W_{p,(p_1,q_2)} + W_{p,(q_1,p_2)} &= (M_{p_1,q_2} - M_{p_1,p_2}) + (M_{p_1,p_2} - M_{q_1,p_2}) \\ &= M_{p_1,q_2} - M_{q_1,p_2} = (\mathcal{S}_M)_{p,q} \quad \text{and} \\ W_{(p_1,q_2),q} + W_{(q_1,p_2),q} &= (M_{p_1,q_2} - M_{q_1,q_2}) + (M_{q_1,q_2} - M_{q_1,p_2}) \\ &= M_{p_1,q_2} - M_{q_1,p_2} = (\mathcal{S}_M)_{p,q} \end{aligned}$$

■

Theorem A.5 (Theorem 4.8) Let M be a non-symmetric zero-sum game. Let $x = (x_1, x_2)$ be a mixed profile and $p = (p_1, p_2)$ a pure profile. Write $x_p := x_{1p_1}x_{2p_2}$ as in equation (1). Then, under the non-symmetric replicator (Definition 3.3),

$$\dot{x}_p = x_p(\mathcal{S}_M x)_p.$$

Proof The two-population replicator dynamic (written for player 1, the player 2 case is similar) is equivalent to

$$\begin{aligned} \dot{x}_{1s} &= x_{1s}((Mx_2)_s - x_1^T Mx_2) && \text{(Definition 3.3)} \\ \dot{x}_{1s}/x_{1s} &= \sum_{t \in S_1} x_{1t}((Mx_2)_s - (Mx_2)_t) && \text{(as } \sum_{t \in S_1} x_{1t} = 1) \\ &= \sum_{t \in S_1} x_{1t} \sum_{r \in S_2} x_{2r} (M_{s,r} - M_{t,r}) \\ &= \sum_{t \in S_1} \sum_{r \in S_2} x_{1t}x_{2r} (M_{s,r} - M_{t,r}) \\ &= \sum_{p=(p_1,p_2) \in S_1 \times S_2} x_p (M_{s,p_2} - M_{p_1,p_2}) && \text{(relabelling)} \end{aligned}$$

Now we observe that for $p = (p_1, p_2)$,

$$\begin{aligned}
 \dot{x}_p &= \frac{d}{dt}(x_{1p_1} x_{2p_2}) \\
 &= (x_{1p_1} x_{2p_2}) \left(\frac{\dot{x}_{1p_1}}{x_{1p_1}} + \frac{\dot{x}_{2p_2}}{x_{2p_2}} \right) && \text{(product rule)} \\
 &= x_p \left(\sum_{q \in S_1 \times S_2} x_q (M_{p_1, q_2} - M_{q_1, q_2}) \right. \\
 &\quad \left. + \sum_{q \in S_1 \times S_2} x_q (M_{q_1, q_2} - M_{q_1, p_2}) \right) && \text{(by above)} \\
 &= x_p \sum_{q \in S_1 \times S_2} x_q (M_{p_1, q_2} - M_{q_1, q_2} + M_{q_1, q_2} - M_{q_1, p_2}) \\
 &= x_p \sum_{q \in S_1 \times S_2} x_q (\mathcal{S}_M)_{p, q} && \text{(Definition 4.6)} \\
 &= x_p (\mathcal{S}_M x)_p.
 \end{aligned}$$

■

Lemma A.6 (Lemma 5.2) *In any flow, every attractor is a sink chain component.*

Proof We first show that all points in the attractor are chain recurrent. Attracting sets are closed under intersection (Kalies et al., 2021; Biggar and Shames, 2023b), so an attractor cannot overlap any other attracting set—that would contradict the minimality of the attractor. Conley (1978) showed that points are chain recurrent if, for each attracting set A , the point is contained in either A or its dual repelling set A^* . The attractor is compact, invariant, and no attracting set overlaps it, so an attracting or repelling set must contain all points in the attractor. Hence all points are chain recurrent.

Pseudo-orbits cannot leave attracting sets (Akin and Losert, 1984). Consequently, no point outside the attractor is chain equivalent to a point inside it, and all points in the attractor are chain equivalent, so it is a chain component. It is a *sink* chain component because no pseudo-orbits leave the set. ■

Corollary A.7 *The content of the sink component of the preference graph is the unique sink chain component of a zero-sum game.*

Proof By Lemma 4.2 and Theorem 3.3, the content of the sink component is a sink chain component. Uniqueness follows for the same reason as in Theorem 3.3: distinct sink chain components are disjoint, but every sink chain component contains the content (Biggar and Shames, 2023b). ■