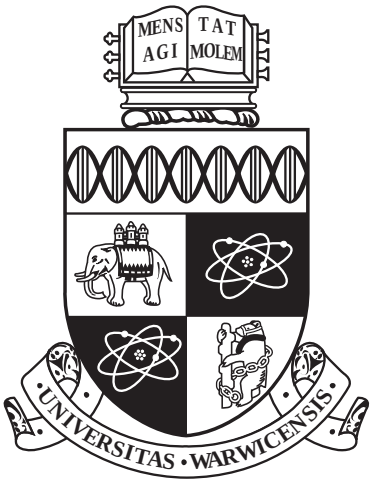

Topics Arising from Fictitious Play Dynamics

by

Georg Ostrovski

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Mathematics

October 2013

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments **iii**

Declarations **iv**

Abstract **v**

Introduction **1**

Chapter 1 Game theory and fictitious play **6**

1.1 Bimatrix games 7

1.1.1 Best response and Nash equilibrium 8

1.1.2 Game equivalences 10

1.2 Fictitious play dynamics 13

1.2.1 Discrete-time fictitious play 13

1.2.2 Continuous-time fictitious play 15

1.2.3 Uniqueness of the fictitious play flow 18

1.2.4 Fictitious play in 2×2 games 19

1.2.5 Convergence of fictitious play 20

1.2.6 Fictitious play as a Hamiltonian system 27

1.3 Why fictitious play? 29

Chapter 2 Combinatorics of fictitious play **34**

2.1 Combinatorial description of fictitious play 34

2.1.1 Coding of fictitious play 35

2.1.2 Transition diagrams 36

2.1.3 Realisable diagrams 37

2.1.4 Admissible and realisable sequences 39

2.1.5 Combinatorics of zero-sum games 44

2.1.6 Periodic and quasi-periodic orbits 54

2.2 Numerical investigation of ergodic properties 57

Chapter 3 Payoff performance of fictitious play **69**

3.1 Limit set for fictitious play 70

3.2 Fictitious play vs. Nash equilibrium payoff 73

3.2.1 An example family of games revisited 78

3.2.2 Fictitious play can be worse than Nash equilibrium play 80

3.3 Concluding remarks on fictitious play performance 83

Chapter 4 Non-self maps of the plane **86**

4.1 Modern results on Brouwer homeomorphisms 86

4.2 Non-self maps of Jordan domains 87

4.3 Non-self maps for compact simply connected planar sets 92

4.4 Proofs of Lemmas 4.7 and 4.14 96

4.5 Discussion of assumptions and possible extensions 98

Chapter 5 Piecewise affine model maps **100**

5.1 Construction of the family of maps 101

5.2 Properties of the maps 102

5.3 Invariant circles 104

5.4 Special parameter values 108

5.5 General parameter values and discussion 112

Bibliography **119**

Acknowledgments

First and foremost, I would like to warmly thank my supervisor, Sebastian van Strien, for his invaluable help and constant support throughout the years of my doctorate. Without his inspiration, ideas, guidance, and most importantly, unceasing encouragement, I would not have been able to finish this work in its present form. His readiness to help with any problem along the way, the countless fruitful discussions on all the topics in this thesis and many others, and his always optimistic attitude never failed to make me leave his office in a more inspired, positive, and constructive mood, than when I entered it. For this most valuable contribution to my PhD I am very grateful to Sebastian.

I thank the Mathematics Department at Warwick University for providing the great working environment and financial support. In particular, I would like to thank the members of the Ergodic Theory and Dynamical Systems group at Warwick: Italo Cipriano, Adam Epstein, Andrew Ferguson, David Franco, José Pedro Gaivão, Vassili Gelfreich, Peter Hazard, Dintle Kagiso, Tom Kempton, Oleg Kozlovski, Robert MacKay, Anthony Manning, Mark Pollicott, Dayal Strub, Sofia Trejo-Abad, Sebastian van Strien, Polina Vytnova, and many others. The seminars, meetings, and many personal discussions contributed much to my research. I am grateful to the Mathematics Department of the Imperial College London and its Dynamical Systems group for its hospitality during the last year of my doctorate.

Thanks to Colin Sparrow for his comments and suggestions at different stages of my work throughout the years. I am grateful to Robbert Fokkink, Saul Schleimer and an anonymous referee for their comments and ideas regarding the paper underlying Chapter 4 of this thesis, to Shaun Bullett for his suggestions on the work presented in Chapter 5, to Martin Rasmussen, Janosch Rieger and Dmitry Turaev for their suggestions on my work in the last year of my PhD.

I warmly thank the many friends who helped me in my research through critical questions, valuable discussions, and proofreading my manuscripts: Sebastian Helmensdorfer, Vlad Moraru, Heather Reeve-Black, Julia Slipantschuk, Jorge Vitória. Thanks to the numerous friends I made during these years in the UK, without whom that time would not have been as pleasant.

Last but not least, I would like to thank my family and friends back home and elsewhere, for their incredible support and encouragement throughout these years: my parents Alexander Ostrovski and Tatjana Pelech and my brother Daniel Ostrovski, my friends Sascha Adler, Dennis Kost, Aileen Ong, Daria Yakovleva.

My PhD has been supported financially by the Warwick Postgraduate Research Scholarship and the EPSRC.

Declarations

I declare that, to the best of my knowledge, the material contained in this thesis is original and my own work except where otherwise indicated. This thesis has not been submitted for a degree at any other university. This thesis was typeset using LaTeX and a modified version of the style package *warwickthesis*.

Abstract

In this thesis, we present a few different topics arising in the study of the learning dynamics called fictitious play. We investigate the combinatorial properties of this dynamical system describing the strategy sequences of the players, and in particular deduce a combinatorial classification of zero-sum games with three strategies per player. We further obtain results about the limit sets and asymptotic payoff performance of fictitious play as a learning algorithm.

In order to study coexistence of regular (periodic and quasi-periodic) and chaotic behaviour in fictitious play and a related continuous, piecewise affine flow on the three-sphere, we look at its planar first return maps and investigate several model problems for such maps. We prove a non-recurrence result for non-self maps of regions in the plane, similar to Brouwer's classical result for planar homeomorphisms. Finally, we consider a family of piecewise affine maps of the square, which is very similar to the first return maps of fictitious play, but simple enough for explicit calculations, and prove several results about its dynamics, particularly its invariant circles and regions.

Introduction

Modern game theory began in the first half of the 20th century, in its formative years most notably driven by John von Neumann's paper *Zur Theorie der Gesellschaftsspiele* (1928, [94]) and his book with Oskar Morgenstern *Theory of Games and Economic Behavior* (1944, [69]). It provides a framework to mathematically formalise decision making in situations of conflict and cooperation, in other words, situations in which choice under uncertainty governs the outcome for the participants.

Since its ascent, game theory has proven to be an immensely powerful tool to describe and analyse interactions of individuals involved in situations of mutual influence and interdependence with each other or an inanimate 'nature', in a highly diverse range of fields, such as economics, biology and genetics, psychology and sociology, political science, and more recently, computer science. It provides a means of analysing the interactions of market participants; it can be used to model the ecological and evolutionary interplay of species or individuals, as well as human behaviour on individual or societal level, ranging up to the scale of international political decision making; finally, it is successfully applied in the modelling and design of computer networks with complex interaction patterns between large numbers of (electronic) agents. Examples include predator-prey models in evolutionary biology, the analysis of international conflicts in political science, the analysis and design of automatic electronic trading systems, and the design and operation of wireless networks consisting of large numbers of nodes which share common resources (such as bandwidth).

From a theoretical point of view, mathematical analysis of game-theoretic models helps to *understand and describe* such pre-existing, observed situations of conflict and cooperation. For instance, after the realisation of the importance of notions of equilibrium in games (most prominently introduced by John Forbes Nash, Jr. in his PhD thesis[1]), the analysis of certain predator-prey models and other evolutionary settings revealed hidden patterns in large-scale animal behaviour. In particular, it facilitated the understanding of natural balances, such as proportions of predator and prey animals, or males and females of a given species, as 'Nash equilibria' of the underlying 'games'.

[1] Published in 1951 as the paper *Non-cooperative games* [72].

More practically, surpassing this descriptive approach, the tools developed in game theory help to *design or affect* such situations, by better understanding complex interaction networks, attempting to model and predict individual or collective behaviour, and set up regulations and policies (in the case of electronic networks, protocols) according to certain outcome goals. In this context the term 'decision theory' is often used synonymously.

Almost from the very beginning, the development of (modern) game theory has been linked with a dynamical view of 'games' being 'played' repeatedly or continuously in time, for instance, in ongoing evolutionary or economic processes. This led to the study of a large number of mathematical dynamical systems modelling such processes, often referred to as *evolutionary dynamics* or *learning dynamics*[2].

An early example of such dynamical systems is the famous Lotka-Volterra predator-prey model, proposed in the 1920s as an explanation of the fluctuations in different Adriatic fish populations in response to a changing fishing behaviour during the first world war. The invention of this model preceded a game-theoretic interpretation of the setting, but was later recognized as a special case of the famous 'replicator dynamics' in evolutionary game theory. Other examples include learning processes studied in psychology and behavioural economy, such as 'reinforcement learning', a process capturing the dynamics of an agent learning by responding accordingly to positive and negative feedback from its past actions.

This thesis is mainly concerned with one of the earliest learning algorithms, known in its different forms as *best response* or *fictitious play dynamics*. The basic idea of fictitious play is that at each stage of a repeatedly or continuously played game, each of the involved players assumes the probability distribution of her opponents' strategies to be stationary. She then chooses to play a strategy which maximises her expected payoff against the empirical average play of her opponents, as observed through their past play. This is a so-called *myopic* learning algorithm, as it does not incorporate any notion of predictive or strategic behaviour: the players do not attempt to influence their opponents' future behaviour through strategic considerations, but only try to maximise their respective next-round payoff.

Introduced in 1951 by George W. Brown [22], the original purpose of fictitious play was not to model actually observed behaviour, but rather to serve as a computational device to study the underlying game. At that time, when Nash's recently proposed equilibrium concept began to be recognised as crucial to understanding games, it was a computational challenge to determine Nash equilibria of a given game. It was observed that in two-player zero-sum games[3], fictitious play dynamics converges to the set of Nash equilibria, and

[2]The term *learning dynamics* is usually employed when 'players' are assumed to be rationally and strategically decision-making agents, whereas *evolutionary dynamics* is more often used to describe the interplay of merely mechanistically responding agents; however, this distinction is quite vague and not always appropriate.

[3]A zero-sum game is one in which the interests of the two players are diametrically opposed, the gain of one being exactly the loss of the other.

hence provides a means of numerical approximation of equilibria. The name 'fictitious play' derives from the following intuition: rather than expecting rational players to follow the fictitious play rule, one can assume that both players 'fictitiously' play the game according to the rule to pre-compute the Nash equilibrium, and then start the actual game by immediately using this equilibrium strategy.

Indeed, the earliest results on fictitious play confirmed that it converges to the set of Nash equilibria in the most important class of games at the time, two-player zero-sum games. Later on, similar convergence properties have been proven for various other relevant classes of games (for example, two-player games in which one of the players has at most two strategies, identical interest games, or potential games). However, in 1964 Lloyd S. Shapley constructed a famous example [87] demonstrating that, in general, fictitious play for two-player games with more than two strategies for each player need not converge.

Shapley's example of a game with two players, each with three strategies, can be seen as a variant of the 'Rock-Paper-Scissors' game: it carries a cyclical structure with no single 'best strategy' for either player. For this game, fictitious play (which effectively describes the dynamics of the players' empirical strategy averages) converges to a stable limit cycle for typical initial conditions. This limit cycle can be recognised as the cyclic behaviour of the two players continuously trying to choose a strategy beating their opponent's current average play (for instance, choosing to play 'rock' to beat an average which is predominantly 'scissors'), soon triggering an analogous response from the other player (who starts to play 'paper' to beat the just chosen 'rock'), and so on. Hence the strategies and, crucially, their time averages are caught in an infinite loop of cyclical strategy switching.

Besides showing that fictitious play does not necessarily converge to the set of Nash equilibria, Shapley's example provided a dynamical system that sparked mathematical interest in its own right. More recently, Colin Sparrow, Sebastian van Strien and Christopher Harris [90, 93] investigated fictitious play dynamics for a one-parameter family of two-player games closely related to Shapley's example and with a similar cyclic Rock-Paper-Scissors-like structure. They demonstrated that this system has rather remarkable dynamical behaviour, in particular being chaotic in a very strong sense, containing elements of random walk dynamics, and exhibiting features which cannot occur in smooth dynamical systems[4]. Furthermore, in [92] van Strien showed that fictitious play can be viewed as a particular example of a more general class of piecewise affine, not differentiable Hamiltonian systems. In analogy to how circle diffeomorphisms can be modelled by circle rotations and how the Lozi map can be regarded as a simplified piecewise affine version of the Hénon map, these piecewise linear Hamiltonian flows (in this case, on the three-sphere) could pro-

[4]Continuous-time fictitious play gives rise to a continuous and piecewise affine but typically not differentiable dynamical system.

vide simplified models for smooth Hamiltonian flows, retaining the dynamical complexity but often allowing more explicit study. This motivates fictitious play dynamics from a purely mathematical point of view, even after stripping off all of its numerous applications. The work by Sparrow and van Strien on the complex behaviour observed in fictitious play forms a starting point for the investigations presented in this thesis.

To set the stage for the main part of this thesis, in Chapter 1 we first introduce the necessary notions and definitions of game theory and fictitious play dynamics. We then give an overview of classical results on fictitious play in general, and the zero-sum case in particular, with a focus on its convergence properties. Finally, we present the recent results by Sparrow and van Strien on the chaotic dynamics emerging from fictitious play and a related class of Hamiltonian systems.

In the main chapter of this thesis, Chapter 2, we investigate fictitious play from a combinatorial point of view, by studying the natural coding of its dynamics in terms of the players' strategy itineraries (the sequence of strategies taken by the players along fictitious play trajectories). This leads to a combinatorial classification of zero-sum games with three strategies per player[5] and a numerical investigation of their ergodic properties. This chapter partially consists of joint work with Sebastian van Strien, published in 2011 as [76].

Chapter 3 is available as a preprint [74] and is submitted for publication at the time of submission of this thesis. It compares the performance of fictitious play as a learning algorithm (measured as the time-averaged payoff to the players) and Nash equilibrium play. We show that Nash equilibrium need not generally be better, and is indeed Pareto dominated by fictitious play in many games. We also observe that while fictitious play does not necessarily converge to Nash equilibrium, it does converge to the larger set of 'coarse correlated equilibria'. This shows that, in the limit, the performance of this very simple learning rule is comparable to some more sophisticated algorithms, such as 'regret-based learning'.

Chapters 4 and 5 can be regarded as separate investigations, motivated by questions arising in the study of certain maps related to the fictitious play flow in zero-sum games with three strategies per player. This flow in four-space is decomposable into a (radial) motion converging to Nash equilibrium and a (spherical) volume-preserving flow on the three-sphere (the 'induced flow'). The latter contains all the information and complexity of the combinatorial behaviour of the fictitious play dynamics; it admits certain topological disks as first return sections, with (planar) first return maps whose study is therefore central to understanding fictitious play dynamics.

[5]The case of 3×3 games, that is, two-player games with three strategies for each player, provides the simplest case in which fictitious play dynamics can have non-trivial, and in fact highly complicated, dynamical behaviour. The intuitive reason for this is that 3×3 games have the lowest dimension, such that both cooperative and competitive elements can be embedded in the game (as 'subgames'). Two-player games where one of the players has less than three strategies are known to converge in a simple cyclical way (see Chapter 1).

Chapter 4 is published in 2013 as [75]. It tackles a primarily topological question about the dynamics of iteration of planar homeomorphisms, very close in flavour to the classical non-recurrence results of Brouwer[6], tailored to the special setting of non-self maps of planar continua. The main result of this chapter states that under a few technical topological conditions, a fixed point free orientation-preserving homeomorphism from one planar continuum to another has no recurrent orbits within any simply connected component of the intersection of its domain and image.

Chapter 5 is available as a preprint [73] and is submitted for publication at the time of submission of this thesis. It studies the dynamics of a one-parameter family of piecewise affine homeomorphisms of the square. This forms an attempt to study phenomenologically the first return maps of the induced flow of fictitious play in 3×3 zero-sum games. For this, we consider the simplest possible (non-trivial) model maps of a compact disk in the plane, with the same formal properties as the first return maps: continuous, area-preserving, piecewise affine and fixing the disk's boundary pointwise. In this chapter we present an investigation of invariant circles and regions for these maps, determining their dynamics completely for certain parameter values. We then make observations about their ergodic properties, and compare them to twist maps, to which they are closely related.

The range of phenomena observed through computational experiments turns out to be quite similar to the one in the numerical investigation of fictitious play dynamics in Chapter 2, confirming the validity of this family of maps as models for the fictitious play first return maps. Although these first return maps are typically much richer (for instance, having many more pieces), a better understanding of this or similar families of piecewise affine planar maps might provide valuable insight into at least qualitative properties of fictitious play, and more generally, piecewise linear Hamiltonian flows.

Publications by the author:

- Piecewise linear Hamiltonian flows associated to zero-sum games: transition combinatorics and questions on ergodicity (with S. van Strien). *Regul. Chaotic Dyn.*, 16(1-2):129-154, 2011.
- Fixed point theorem for non-self maps of regions in the plane. *Topology Appl.*, 160(7):915-923, 2013.
- Dynamics of a continuous piecewise affine map of the square. arXiv:1305.4282, 2013. Submitted for publication to *Phys. D*.
- Payoff performance of fictitious play (with S. van Strien). arXiv:1308.4049, 2013.

[6] Brouwer's classical plane translation theorem and its modern versions due to John Franks and others state that an orientation-preserving fixed point free homeomorphism of the plane has no recurrent points.

Chapter 1

Game theory and fictitious play

In this chapter we introduce basic notions from game theory and the learning dynamics called fictitious play. Our focus will be on two-player games given by a bimatrix, for which we will discuss a geometric representation of the strategy space. We will then present an overview of classical results on fictitious play dynamics and discuss its properties in the special setting of zero-sum games. The main purpose of this chapter is to set the stage for the investigation of fictitious play dynamics in zero-sum games in Chapter 2.

Some of the earliest and most prominent works that started game theory in its modern form are the paper *Zur Theorie der Gesellschaftsspiele* from 1928 by von Neumann [94] and the book *Theory of Games and Economic Behavior* from 1944 by Morgenstern and von Neumann [69], introducing many of the notions used in game theory to this day, followed by Nash's PhD thesis (in particular, its part on equilibria in games, published in 1951 as the paper *Non-cooperative games* [72]).

For further reference on game theory in general, and learning dynamics in particular, the reader is referred to one of the books by Fudenberg and Tirole [34], Fudenberg and Levine [32], Hofbauer and Sigmund [50, 51], Weibull [95] or Sandholm [84]. For an excellent overview of the dynamics of the most popular learning algorithms, see the introductory book by Young [99] or the review articles by Hart [44] and Fudenberg and Levine [33].

We begin by giving the notion of a game a mathematical definition.

Definition 1.1. A *finite game in normal form* is a tuple $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$, where

- $\mathcal{I} = \{1, \dots, N\}, N \in \mathbb{N}$, is a (finite) collection of *players*;
- $S^i = \{1, \dots, n_i\}, n_i \in \mathbb{N}$, is the (finite) collection of *pure strategies* of player $i \in \mathcal{I}$;
- $u^i \colon S^1 \times \cdots \times S^N \to \mathbb{R}$ is the *payoff function* of player $i \in \mathcal{I}$.

The space of strategy tuples is denoted by $S = S^1 \times \cdots \times S^N$ and an element $s \in S$ is called a *(pure) strategy profile*.

The interpretation of this definition is as follows. To play the game Γ, each player $i \in \mathcal{I}$ chooses one of her strategies $s_i \in S^i$, independently and without knowing the other players' choices. Then, each player i receives a payoff $u^i(s_1, \dots, s_N)$, which depends both on her and her opponents' strategies. Each player strives to maximise her own payoff.[1]

Instead of restricting to the discrete sets of pure strategies as above, we further consider the so-called *mixed strategies*. For $i \in \mathcal{I}$, we write

$$\Sigma^i := \left\{ x \in \mathbb{R}^{n_i} : x_k \geq 0, \sum_{k=1}^{n_i} x_k = 1 \right\}$$

for the (geometric representation of) probability distributions over a player's pure strategies. Here, implicitly each strategy $k \in S^i$ (or more precisely, the distribution assigning full probability to this strategy) is identified with the kth standard unit vector in $\mathbb{R}^{n_i}$. Note that Σ^i is geometrically the $(n_i - 1)$-dimensional simplex in $\mathbb{R}^{n_i}$ spanned by these n_i unit vectors. We write $\Sigma = \Sigma^1 \times \cdots \times \Sigma^N$ and say that $\sigma \in \Sigma$ is a *(mixed) strategy profile*.

Since mixed strategies are probability distributions over pure strategies, we can linearly extend the payoff functions to $\tilde{u}^i \colon \Sigma \to \mathbb{R}$. For this, let $\sigma = (\sigma_1, \dots, \sigma_N) \in \Sigma$, where each $\sigma_i \in \Sigma^i$ is a probability distribution over the pure strategies of player i, $\sigma_i = (\sigma_i^1, \dots, \sigma_i^{n_i})$. Then we define

$$\tilde{u}^i(\sigma) = \tilde{u}^i(\sigma_1, \dots, \sigma_N) := \sum_{k_1=1}^{n_1} \cdots \sum_{k_N=1}^{n_N} \sigma_1^{k_1} \cdots \sigma_N^{k_N} u^i(k_1, \dots, k_N).$$

By linearity, this can be interpreted as the expected payoff to player i, if each of the players is randomising over her strategies according to her mixed strategy σ_i. Henceforth, we will not distinguish between u^i and $\tilde{u}^i$, and refer to both payoff functions as u^i.

Throughout this entire thesis, we will focus our attention on two-player games. This setting allows for a simple representation of the two players' payoff functions as matrices, hence encoding the entire information about a game in a 'bimatrix'.

1.1 Bimatrix games

Let a two-player game Γ be defined as above, with players 1 and 2 having m and n (pure) strategies, respectively. Then, the two payoff functions $u^1, u^2 \colon \Sigma \to \mathbb{R}$ can be represented by a *bimatrix*, that is, a tuple of matrices (A, B), where $A, B \in \mathbb{R}^{m \times n}$.

For this, we identify the standard unit vector $e_i \in \Sigma^1$ with the first player's strategy

[1] Note that in general, no assumption is made on any connections between the players' payoff functions u^i. In particular, no notion of 'winning' or 'losing' the game is implied automatically; in fact, the payoff functions might be such that all players benefit from the same strategy tuples.

$i \in \{1, \dots, m\}$ and $e_j \in \Sigma^2$ with the second player's strategy $j \in \{1, \dots, n\}$. Then, let the (i, j) entry of A be given by $a_{ij} = e_i^\top A e_j = u^1(i, j)$, and similarly, $b_{ij} = e_i^\top B e_j = u^2(i, j)$. In other words, the (i, j) entries of the matrices A and B are the respective payoffs of the (pure) strategy profile (i, j) to the players 1 and 2. By linearity, it follows immediately that the (expected) payoffs of the (mixed) strategy profile $(p, q) \in \Sigma^1 \times \Sigma^2 = \Sigma$ for players 1 and 2 are given by $u^1(p, q) = p^\top A q$ and $u^2(p, q) = p^\top B q$, respectively. A finite two-player game given in this form (A, B) is called a *bimatrix game*.

For convenience, in the case of bimatrix games we will use the notation $\Sigma^A = \Sigma^1$, $\Sigma^B = \Sigma^2$, $u^A = u^1$, $u^B = u^2$, and we will interpret vectors $p \in \Sigma^A$ as row vectors in $\mathbb{R}^{1\times m}$, so that for $(p, q) \in \Sigma$ we can write

$$u^A(p, q) = pAq \quad \text{and} \quad u^B(p, q) = pBq.$$

Now that we defined a setting for two-player games, we will introduce some strategic notions, expressing how players can reason about a game and evaluate their different strategy options when pursuing their goal of payoff maximisation.

1.1.1 Best response and Nash equilibrium

We define the so-called *best response correspondences*, which assign payoff-maximising response strategies to any given strategy of a player's opponent. Let $(p, q) \in \Sigma$, then

$$\mathcal{BR}_A(q) \coloneqq \arg\max_{\bar{p}\in\Sigma^A} \bar{p}Aq \quad \text{and} \quad \mathcal{BR}_B(p) \coloneqq \arg\max_{\bar{q}\in\Sigma^B} pB\bar{q}.$$

We further denote the maximal-payoff functions

$$\bar{A}(q) \coloneqq \max_{\bar{p}\in\Sigma^A} \bar{p}Aq \quad \text{and} \quad \bar{B}(p) \coloneqq \max_{\bar{q}\in\Sigma^B} pB\bar{q},$$

so that $\bar{A}(q) = u^A(\bar{p}, q)$ for $\bar{p} \in \mathcal{BR}_A(q)$ and $\bar{B}(p) = u^B(p, \bar{q})$ for $\bar{q} \in \mathcal{BR}_B(p)$. Observe that $\bar{A}(q) = \max_i (Aq)_i$ and $\bar{B}(p) = \max_j (pB)_j$: the maximal payoff to player A given player B's strategy q is equal to the maximal entry of the vector Aq, and similarly for player B.

For generic bimatrix games, the best response correspondences $\mathcal{BR}_A \colon \Sigma^B \to \Sigma^A$ and $\mathcal{BR}_B \colon \Sigma^A \to \Sigma^B$ are almost everywhere single-valued, with the exception of a finite number of hyperplanes. The singleton value taken by $\mathcal{BR}_A$ whenever it is single-valued is always one of the standard unit vectors e_i, $i = 1, \dots, m$, corresponding to a pure strategy of player A. When $\mathcal{BR}_A(p)$ is not a singleton, it is the set of convex combinations of a subset of $\{e_i \colon i = 1, \dots, m\}$, that is, a face of the simplex Σ^A, or possibly all of Σ^A. The analogous statement holds for $\mathcal{BR}_B$.

It follows that Σ^A and Σ^B can be divided into respectively n and m regions (in fact, convex polytopes):

$$R_j^B := \mathcal{BR}_B{}^{-1}(e_j) \subseteq \Sigma^A \qquad \text{for } j = 1, \dots, n,$$
$$R_i^A := \mathcal{BR}_A{}^{-1}(e_i) \subseteq \Sigma^B \qquad \text{for } i = 1, \dots, m.$$

We will call R_i^A the *preference region* of strategy i for player A, as it is the subset of the second player's strategies against which player A expects the highest payoff by playing strategy i; similarly, for R_j^B. We also use the following notation for the subsets of Σ on which both players have a fixed strategy preference:

$$R_{ij} := R_j^B \times R_i^A \qquad \text{for } i = 1, \dots, m \text{ and } j = 1, \dots, n.$$

Next, note that for a generic game (A, B), the subset of Σ^B on which $\mathcal{BR}_A$ contains two distinct pure strategies e_i and $e_{i'}$ (and hence, automatically, all their convex combinations) is a codimension-one hyperplane of Σ^B:

$$Z_{ii'}^A := \{q \in \Sigma^B : (Aq)_i = (Aq)_{i'} \geq (Aq)_k \ \forall k = 1, \dots, m\} = R_i^A \cap R_{i'}^A \subseteq \Sigma^B.$$

Analogously, for $j, j' \in \{1, \dots, n\}$,

$$Z_{jj'}^B := \{p \in \Sigma^A : (pB)_j = (pB)_{j'} \geq (pB)_l \ \forall l = 1, \dots, n\} = R_j^B \cap R_{j'}^B \subseteq \Sigma^A.$$

These hyperplanes are subsets of linear codimension-one subspaces of Σ^B and Σ^A, respectively. See Figure 1.1 for an illustration in the case $n = m = 3$. We call these sets the *indifference sets* of players A and B. As we will see later, they turn out to be of crucial importance to understanding learning dynamics on the space Σ, when this dynamics is defined by means of the best response correspondences.

The following notion was first introduced by John Nash in his PhD thesis (see [72]).

Definition 1.2. A (mixed) strategy profile $(\bar{p}, \bar{q}) \in \Sigma$ is called a *Nash equilibrium*, if

$$\bar{p} \in \mathcal{BR}_A(\bar{q}) \quad \text{and} \quad \bar{q} \in \mathcal{BR}_B(\bar{p}).$$

If a Nash equilibrium lies in the interior of Σ, it is called *completely mixed*.

The key idea is that in a Nash equilibrium, neither player benefits from *unilateral* deviation. Note that this does not necessarily imply that both players are receiving the highest possible payoff, but rather that they are both receiving the highest possible payoff *given their opponent's strategy*.

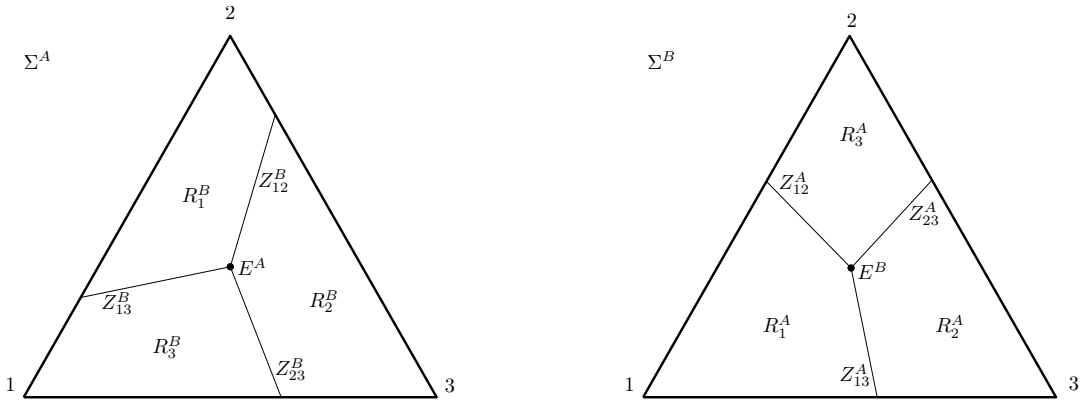

Figure 1.1: Geometry of a 3×3 bimatrix game. The spaces of mixed strategies Σ^A and Σ^B are each a simplex spanned by three vertices (the pure strategies). Note the convex preference regions $R_j^B \subset \Sigma^A$ and $R_i^A \subset \Sigma^B$, their intersections as indifference sets $Z_{jj'}^B$ and $Z_{ii'}^A$, and the projections to Σ^A and Σ^B of the (in this case, unique) Nash equilibrium (E^A, E^B) at the intersection of all these sets.

The concept of Nash equilibrium is central to most of modern game theory; the study of learning algorithms for games is to a large extent concerned with questions about whether and how convergence to Nash equilibrium can be achieved. It has been proved by Nash [72] that the set of Nash equilibria is non-empty for every bimatrix game (for several alternative proofs, see also Hofbauer [49]). On the other hand, generally, Nash equilibrium need not be unique; the set of Nash equilibria can be a discrete set of points in Σ or even contain a continuum of points.

The following lemma is a standard fact and easy to check.

Lemma 1.3. *The point $(E^A, E^B) \in \mathrm{int}(\Sigma)$ is a (completely mixed) Nash equilibrium of an $m \times n$ bimatrix game (A, B) if and only if, for all $i, i' = 1, \ldots, m$ and $j, j' = 1, \ldots, n$,*

$$(AE^B)_i = (AE^B)_{i'} \quad \text{and} \quad (E^A B)_j = (E^A B)_{j'}.$$

Note that this implies that $E^A \in R_j^B$ and $E^B \in R_i^A$, for all i, j. It follows that $(E^A, E^B) \in R_{ij}$, and also $E^A \in Z_{ii'}^B$, $E^B \in Z_{jj'}^A$, for all i, i', j, j'.

1.1.2 Game equivalences

For use in the context of learning dynamics, it will be convenient to introduce a notion of equivalence of games. As we will see in Section 1.2, fictitious play dynamics is defined in terms of the best response correspondences $\mathcal{BR}_A$ and $\mathcal{BR}_B$, so that different bimatrix games which have the same best response structures produce identical fictitious play systems.

This gives rise to the notion of 'best response equivalence'.

Definition 1.4. Two $m \times n$ bimatrix games (A, B) and $(\tilde{A}, \tilde{B})$ are *best response equivalent*, if for all $(p, q) \in \Sigma$,

$$\mathcal{BR}_A(q) = \mathcal{BR}_{\tilde{A}}(q) \quad \text{and} \quad \mathcal{BR}_B(p) = \mathcal{BR}_{\tilde{B}}(p).$$

Another, stronger, notion of game equivalence is 'better response equivalence', requiring the same preference ordering of pure strategies.

Definition 1.5. Two $m \times n$ bimatrix games (A, B) and $(\tilde{A}, \tilde{B})$ are *better response equivalent*, if for any mixed strategy $q \in \Sigma^B$ and any pure strategies $x, x' \in S^A$

$$u^A(x, q) \geq u^A(x', q) \Leftrightarrow u^{\tilde{A}}(x, q) \geq u^{\tilde{A}}(x', q),$$

and for any $p \in \Sigma^A$, $y, y' \in S^B$,

$$u^B(p, y) \geq u^B(p, y') \Leftrightarrow u^{\tilde{B}}(p, y) \geq u^{\tilde{B}}(p, y').$$

Clearly, better response equivalent games are also best response equivalent.

Although, in general, best and better response equivalence are very well suited to describe dynamically equivalent bimatrix games, for technical convenience we will mostly be using the following stronger form of equivalence.

Definition 1.6. We say that two $m \times n$ bimatrix games (A, B) and $(\tilde{A}, \tilde{B})$ are *(linearly) equivalent*, $(A, B) \sim (\tilde{A}, \tilde{B})$, if the matrix $\tilde{A}$ can be obtained by multiplying A with a positive constant $c > 0$ and adding constants $c_1, \ldots, c_n \in \mathbb{R}$ to the matrix columns, and $\tilde{B}$ can be obtained from B by multiplication with $d > 0$ and addition of $d_1, \ldots, d_m \in \mathbb{R}$ to its rows:

$$\tilde{a}_{ij} = c \cdot a_{ij} + c_j \quad \text{and} \quad \tilde{b}_{ij} = d \cdot b_{ij} + d_i \qquad \text{for } i = 1, \ldots, m \text{ and } j = 1, \ldots, n.$$

The following lemma follows by direct computation.

Lemma 1.7. *Let (A, B) and $(\tilde{A}, \tilde{B})$ be two $m \times n$ bimatrix games. If (A, B) and $(\tilde{A}, \tilde{B})$ are linearly equivalent, then they are also better and best response equivalent. In particular, their best response correspondences coincide, $\mathcal{BR}_A \equiv \mathcal{BR}_{\tilde{A}}$ and $\mathcal{BR}_B \equiv \mathcal{BR}_{\tilde{B}}$.*

Remark 1.8. Linear equivalence is stronger than better response equivalence, which in turn is stronger than best response equivalence, and in general, neither of these equivalence notions coincide[2]. However, for the class of games that we will mostly be considering in

[2]This is due to the fact that both best and better response equivalences only take into account payoff comparisons for strategy vectors in Σ, whereas linear equivalence enforces a definite relation between the payoff matrices, irrespective of what vectors are plugged in as strategy vectors.

this thesis (namely, $n \times n$ bimatrix games with unique, completely mixed Nash equilibrium), all three notions of equivalence can be regarded as equal. An extensive treatment of these and other game equivalences, as well as their relations to each other, can be found in [70]; see also [71].

The following statement follows immediately from Lemma 1.7 and the definitions of preference regions, indifference sets, and Nash equilibrium.

Corollary 1.9. *If two $m \times n$ bimatrix games (A, B) and $(\tilde{A}, \tilde{B})$ are (linearly) equivalent, then all their preference regions and indifference sets coincide, that is,*

$$R_i^{\tilde{A}} = R_i^A,\ R_j^{\tilde{B}} = R_j^B,\ Z_{ii'}^{\tilde{A}} = Z_{ii'}^A,\ Z_{jj'}^{\tilde{B}} = Z_{jj'}^B \qquad \text{for } i, i' = 1, \ldots, m \text{ and } j, j' = 1, \ldots, n.$$

Moreover, the Nash equilibrium sets of (A, B) and $(\tilde{A}, \tilde{B})$ coincide.

Throughout most of this thesis, we will be dealing with games that have a unique, completely mixed Nash equilibrium. We now provide a lemma which tells us that only $n \times n$ bimatrix games (with equally many strategies for each player) can satisfy this condition.

Lemma 1.10. *Let (A, B) be an $m \times n$ bimatrix game with an isolated, completely mixed Nash equilibrium $(E^A, E^B) \in \operatorname{int}(\Sigma)$. Then $m = n$.*

Proof. Assume for a contradiction that $m \neq n$. Without loss of generality take $m < n$; the case $m > n$ is similar after swapping the roles of the players (replacing (A, B) by $(B^\top, A^\top)$).

We denote $\mathbb{1} = (1, \ldots, 1)^\top \in \mathbb{R}^n$ and $U = \operatorname{span}(\mathbb{1})$. By Lemma 1.3, $AE^B \in U$, and E^B is the only point in Σ^B with this property. This implies that $V = \{v \in \mathbb{R}^n : Av \in U\}$ is a one-dimensional subspace of $\mathbb{R}^n$ with $\{E^B\} = V \cap \Sigma^B$. Now, since $A \in \mathbb{R}^{m \times n}$ with $m < n$, we have that $\dim(\ker(A)) \geq 1$, where A is viewed as a linear map from $\mathbb{R}^n$ to $\mathbb{R}^m$. We differentiate two cases.

Case 1: $V \not\subseteq \ker(A)$. Then $W = V \oplus \ker(A)$ has dimension greater than or equal to 2, and therefore $\dim(W \cap \Sigma^B) \geq 1$, since $\dim(\Sigma^B) = n - 1$ and $E^B \in W \cap \operatorname{int}(\Sigma^B)$. But $Aw \in U$ for every $w \in W$, which contradicts the fact that E^B is the only such point in Σ^B.

Case 2: $V \subseteq \ker(A)$. In this case, let $c \neq 0$ be some constant and let $\tilde{A} \in \mathbb{R}^{m \times n}$ be the matrix obtained from A by adding c to each of its entries, that is, $\tilde{a}_{ij} = a_{ij} + c$ for all i, j. Clearly, $(\tilde{A}, B)$ is linearly equivalent to (A, B) and by the previous corollary, both games have the same sets of Nash equilibria. Therefore, with $\tilde{V} = \{v \in \mathbb{R}^n : \tilde{A}v \in U\}$ we have $\{E^B\} = \tilde{V} \cap \Sigma^B$, and it follows that $\tilde{V} = V$. On the other hand, $\tilde{A}E^B = c \cdot \mathbb{1}$, so that $\tilde{V} \not\subseteq \ker(A)$, and we can apply case 1 to obtain a contradiction.

This finishes the proof of the lemma. □

1.2 Fictitious play dynamics

We now proceed to the main goal of this chapter, defining a dynamical system modelling the repeated or continuous play of a game. For this, we will look at fictitious play, a simple learning algorithm, which captures the following straightforward intuition:

> At each stage of the game, each player determines the (mixed) average strategy of her opponent's play up to this time and plays a (pure) best response to it.

The empirical average strategy of a player up to a certain time could be interpreted as the *belief* of the opponent about this player's (internal and unobserved) strategy distribution, making the assumption that this player follows a stationary strategy distribution throughout the entire course of play. Fictitious play is defined as the dynamics of these beliefs, which is a dynamical system on $\Sigma = \Sigma^A \times \Sigma^B$, the space of mixed strategy profiles.

The algorithm is a so-called *myopic* learning rule, since players only aim at maximising next-round payoff based on the play history, and do not attempt to make any further strategic considerations to influence their opponent's future behaviour. It is rather naive in that the assumption of stationary strategy distributions is not realistic for even just moderately adaptive opponents. However, as we will see in Chapter 3, its performance in terms of time-average payoff is not necessarily worse than, say, constantly playing Nash equilibrium. Moreover, it turns out that fictitious play converges to the same limit sets as some much more sophisticated learning dynamics, which provides further motivation to the study of this simple algorithm.

In 1951, Brown introduced both a discrete- and continuous-time version of the algorithm [22]; see also the unpublished report [21]. The discrete-time version was initially more popular, intended to be an algorithm for numerical approximation of Nash equilibrium in zero-sum games, which was a computational challenge at the time. This motivation was justified by the earliest work on discrete-time fictitious play in 1951 by Robinson [82], showing that indeed the process converges to Nash equilibrium in zero-sum games.

Our focus in this thesis will be on continuous-time fictitious play, but we begin by defining the discrete-time process, as from a historical and intuitive point of view it precedes the continuous-time version.

1.2.1 Discrete-time fictitious play

Let (A, B) be an $m \times n$ bimatrix games with strategy sets S^A and S^B, played repeatedly at times $k \in \mathbb{N}_0$. Let $(x_k, y_k) \in S^A \times S^B$ denote the (pure) strategies played by the players at time $k \in \mathbb{N}$, with some (mixed) initial condition $(x_0, y_0) \in \Sigma$. For $k \in \mathbb{N}$, denote the

empirical average play through time $k-1$ by

$$p_k \coloneqq \frac{1}{k}\sum_{i=0}^{k-1} x_i \in \Sigma^A \quad \text{and} \quad q_k \coloneqq \frac{1}{k}\sum_{i=0}^{k-1} y_i \in \Sigma^B.$$

Then the fictitious play rule requires that

$$x_k \in \mathcal{BR}_A(q_k) \cap S^A \quad \text{and} \quad y_k \in \mathcal{BR}_B(p_k) \cap S^B \tag{1.1}$$

for all times $k \in \mathbb{N}$.

Remark 1.11. Recall that both sets $\mathcal{BR}_A(q_k)$ and $\mathcal{BR}_B(p_k)$ each contain at least one pure strategy. Here we do not specify any so-called 'tie-breaking rule', that is, a choice of pure strategy from $\mathcal{BR}_A(q_k)$ and $\mathcal{BR}_B(p_k)$ in (1.1), whenever one of these sets is not a singleton. All of the results to follow hold independently of the choice of such tie-breaking rule.

The p_k and q_k are the beliefs of the players at any time $k \in \mathbb{N}$ about the strategy distribution of their respective opponent. We calculate

$$\begin{aligned} p_{k+1} &= \frac{1}{k+1}\sum_{i=0}^{k} x_i \\ &= \frac{1}{k+1}x_k + \frac{k}{k+1}\frac{1}{k}\sum_{i=0}^{k-1} x_i \\ &= \frac{1}{k+1}x_k + \frac{k}{k+1}p_k \\ &\in \frac{1}{k+1}\mathcal{BR}_A(q_k) + \frac{k}{k+1}p_k, \end{aligned}$$

and similarly

$$q_{k+1} \in \frac{1}{k+1}\mathcal{BR}_B(p_k) + \frac{k}{k+1}q_k.$$

We get the following definition.

Definition 1.12. For a bimatrix game (A, B) with mixed strategy space Σ, *discrete-time fictitious play* is the process $(p_k, q_k) \in \Sigma$, $k \geq 1$, given by the initial condition $(p_1, q_1) \in \Sigma$ and for $k \geq 1$,

$$p_{k+1} \in \frac{1}{k+1}(\mathcal{BR}_A(q_k) + k \cdot p_k), \qquad q_{k+1} \in \frac{1}{k+1}(\mathcal{BR}_B(p_k) + k \cdot q_k), \tag{DFP}$$

where some tie-breaking rule for $\mathcal{BR}_A$ and $\mathcal{BR}_B$ is applied whenever these are multi-valued.

Remark 1.13. (1) Geometrically, discrete-time fictitious play follows a very simple intuition. When the players have beliefs $(p_k, q_k) \in R_{ij} \subset \Sigma$, that is, their best responses

to each other's empirical average strategies are i and j, respectively, then p_{k+1} lies on the line segment between p_k and $e_i \in \Sigma^A$, and q_{k+1} on the line segment between q_k and $e_j \in \Sigma^B$. That is, both players' beliefs p_k and q_k always move towards their currently preferred pure strategy, with step size decreasing with time k.

(2) Note that

$$p_{k+1} - p_k \in \frac{1}{k+1}(\mathcal{BR}_A(q_k) - p_k), \quad q_{k+1} - q_k \in \frac{1}{k+1}(\mathcal{BR}_A(p_k) - q_k), \tag{1.2}$$

so that $\|p_{k+1} - p_k\|$ and $\|q_{k+1} - q_k\|$ are bounded by $\sqrt{2}/(k+1)$. In other words, the step size of fictitious play decreases like $1/k$, reflecting the intuitive fact that in this learning dynamics, new incoming data about the opponent has decreasing influence as k grows, since it is being compared to an increasingly long play history, and no discounting of the past is applied in this dynamics.

(3) As Berger points out in [12], Brown's original algorithm differs from what we (and most other authors) call fictitious play in a subtle yet significant detail: in his version of the algorithm, the beliefs p_k and q_k of the players are updated *alternatingly* instead of *simultaneously*, which allows for certain convergence proofs to be carried out rather easily.

1.2.2 Continuous-time fictitious play

Instead of considering a game being played repeatedly, we can look at a continuous-time process. *Continuous-time fictitious play* is a dynamical system in which both players are assumed to continuously play a given bimatrix game by playing a best response to the average of their respective opponent's past play at each time $t > 0$. In analogy to the discrete-time case, we give the following definition (see, for example, Hofbauer [48]).

Definition 1.14. For a bimatrix game (A, B) with mixed strategy space Σ, *continuous-time fictitious play* is the process $(p(t), q(t)) \in \Sigma$, $t \geq t_0 > 0$, given by the differential inclusion

$$\dot{p}(t) \in \frac{1}{t}(\mathcal{BR}_A(q(t)) - p(t)), \qquad \dot{q}(t) \in \frac{1}{t}(\mathcal{BR}_A(p(t)) - q(t)), \tag{FP}$$

with some initial condition $(p(t_0), q(t_0)) = (p_0, q_0) \in \Sigma$.

Remark 1.15. (1) Note that (FP) is a differential inclusion. Uniqueness of solutions cannot generally be guaranteed, but the fact that $\mathcal{BR}_A$ and $\mathcal{BR}_B$ are upper semi-continuous correspondences with closed and convex values (faces of Σ^A and Σ^B) implies by general theory that solutions exist for all initial conditions (see Aubin and Cellina [4,

Chapter 4, Section 2, Theorem 1]). Later on, we are going to exploit the fact that these correspondences are single-valued except on a finite number of hyperplanes (the indifference planes of the players) to deduce that in the games we consider, solutions are indeed unique.

(2) It is clear from the definition that Nash equilibria are precisely the equilibria of fictitious play dynamics (FP). In particular, if an orbit of (FP) converges to a point, this point is necessarily a Nash equilibrium.

(3) As in the discrete-time case, the fact that $\mathcal{BR}_A$ and $\mathcal{BR}_B$ are piecewise constant on the convex sets $R_j^B \subseteq \Sigma^A$ and $R_i^A \subseteq \Sigma^B$ respectively leads to a simple geometric intuition for the (local) dynamics of (FP). Orbits are locally straight line segments heading for vertices of Σ, which only change direction upon hitting an indifference set, or equivalently, passing into a different block R_{ij}. See Figures 1.4 and 1.2 for examples in 2×2 and 3×3 games.

(4) Definition 1.14 can be obtained (informally) from discrete-time fictitious play by noting that in the difference inclusion (1.2) step sizes go to zero as $n \to \infty$ (see [48]). An alternative, equivalent definition (see Harris [43]) is somewhat more constructive and we introduce it here since its terminology will help to give a shorter proof of a result in Section 3.1 of Chapter 3.

Mimicking notation from the discrete-time case, we denote by $x(t)$ and $y(t)$ the strategies played by the two players at time $t \geq 0$, where $x\colon [0,\infty) \to \Sigma^A$ and $y\colon [0,\infty) \to \Sigma^B$ are assumed to be measurable functions. We write the average (empirical) past play of the respective players from time 0 through t as

$$p(t) \coloneqq \frac{1}{t}\int_0^t x(s)\,ds \quad \text{and} \quad q(t) \coloneqq \frac{1}{t}\int_0^t y(s)\,ds.$$

Then, analogously to (1.1), continuous-time fictitious play is given by the rule expressed in the following integral inclusions:

$$x(t) \in \mathcal{BR}_A(q(t)), \quad y(t) \in \mathcal{BR}_B(p(t)) \qquad \text{for } t \geq 1$$

and $(x(t), y(t)) \in \Sigma$ arbitrary for $0 \leq t < 1$. Note that with an appropriate tie-breaking rule, $x(t)$ and $y(t)$ can be chosen to be pure strategies for any time $t \geq 1$, since $\mathcal{BR}_A(q)$ and $\mathcal{BR}_B(p)$ each contain at least one pure strategy for any $(p, q) \in \Sigma$.

Defined this way, $(p(t), q(t))$, $t \geq 1$, is a solution of the differential inclusion (FP) with $t_0 = 1$ and initial condition $p(1) = \int_0^1 x(s)\,ds$ and $q(1) = \int_0^1 y(s)\,ds$.

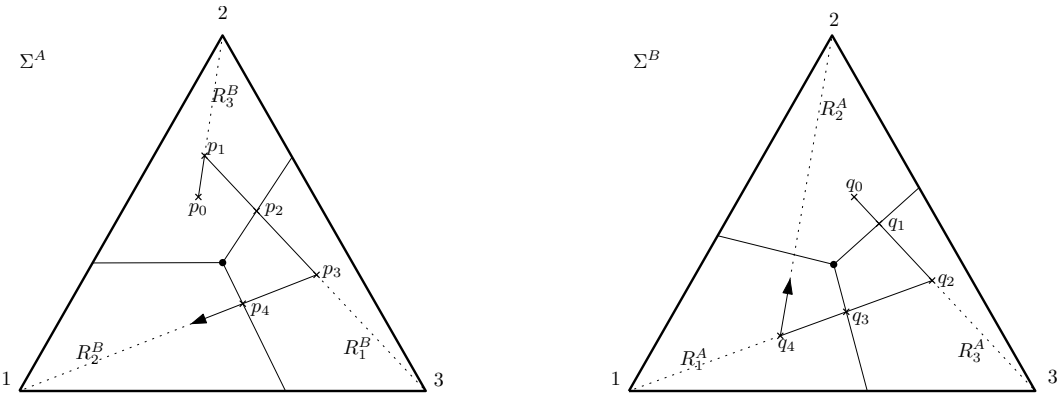

Figure 1.2: A fictitious play orbit in a 3×3 game with unique, completely mixed Nash equilibrium, similar to Shapley's Rock-Paper-Scissor-like game (1.5) or (1.7). The initial condition is $(p_0, q_0) \in \Sigma^A \times \Sigma^B$. Note that the players' trajectories are piecewise straight line segments. Player A changes direction whenever player B crosses one of her indifference lines, and vice versa. These changes of direction are marked as the points (p_i, q_i), indicating that $(p(t), q(t)) = (p_i, q_i)$ for times $t = t_i$, $i = 0, 1, 2, \ldots$.

An important observation is that although (FP) is formally time-dependent, a simple time-reparametrisation $t = e^s$ turns it into an autonomous system. That is, the only time dependence in fictitious play dynamics is the slowing down of the motion as time progresses, but not the actual shape of the trajectories. Thinking of fictitious play as an autonomous system will sometimes simplify certain arguments, so we give a formal definition. This system is sometimes referred to as 'best response dynamics', it has first been considered by Gilboa and Matsui [37] (see also Matsui [61], or for a wider overview, Hofbauer [48] and Hofbauer and Sorin [52]).

Definition 1.16. For a bimatrix game (A, B) with mixed strategy space Σ, *best response dynamics* is the process $(p(t), q(t)) \in \Sigma$, $t \geq 0$, given by the differential inclusion

$$\dot{p}(t) \in \mathcal{BR}_A(q(t)) - p(t), \qquad \dot{q}(t) \in \mathcal{BR}_A(p(t)) - q(t), \tag{BR}$$

with some initial condition $(p(0), q(0)) = (p, q) \in \Sigma$.

Remark 1.17. Due to its motivation as a computational device, discrete-time fictitious play has been preferred and continuous-time fictitious play (also introduced by Brown [21]) was almost forgotten until it reappeared in 1971 in a paper by Rosenmüller [83]. However, the issue of 'overshooting' arising in discrete-time fictitious play makes many analytic proofs much easier for the continuous-time dynamics. Consequently, most modern literature deals with the continuous-time case, sometimes deducing certain statements for the discrete-time case from analogous continuous-time results via approximation techniques. See also the

paragraphs following Theorem 1.40.

For the rest of this and the next two chapters, we will be mainly concerned with the dynamics of (FP) and (BR). We now proceed to properties of the flow defined by these dynamical systems, before presenting classical results about fictitious play and best response dynamics in the subsequent sections of this chapter.

1.2.3 Uniqueness of the fictitious play flow

As discussed in Remark 1.15(1), the differential inclusion (FP) has solutions for any initial conditions $(p_0, q_0) \in \Sigma$. Uniqueness and continuity of the flow induced by this system turn out to be less certain, and in fact only hold for generic (but not for all) bimatrix games.

Clearly, in the interior of each of the convex blocks $R_{ij} \subseteq \Sigma$, solutions of (FP) are (locally) unique and continuous (and, in fact, locally consist of straight line segments). Hence the problem only arises when these solution curves meet any of the indifference sets $Z^A = \bigcup_{i,i'} Z^A_{ii'}$ or $Z^B = \bigcup_{j,j'} Z^B_{jj'}$. On these sets, the right-hand sides of (FP) are multi-valued, giving rise to potential non-uniqueness of solutions. To avoid this, a possible requirement is that the game is such that all of $Z^A_{ii'}$ and $Z^B_{jj'}$ are codimension-one planes, and that the flow crosses these transversally; this would guarantee that solutions leave such sets instantaneously and cannot get trapped inside sets with non-unique flow direction. The following proposition formalises this intuition.

Proposition 1.18 (Sparrow et al. [90, Proposition 2.1])**.** *Let (A, B) be an $m \times n$ bimatrix game. Denote $Z^* = Z^B \times Z^A$ the set where each of the players is indifferent between at least two strategies. Assume that for all $(p, q) \in \Sigma \setminus Z^*$, if $p \in Z^B$ and $q \notin Z^A$ so that, say, $\mathcal{BR}_A(q) = e_k$, then e_k is not parallel to the plane $Z^B \subset \Sigma^A$ at the point p, and similarly for the roles of p and q reversed. Then* (FP) *defines a continuous flow on $\Sigma \setminus Z^*$.*

Using that Z^* is the finite union of codimension-two planes, one can immediately deduce the following statement.

Corollary 1.19 (Sparrow et al. [90, Corollary 2.1])**.** *Under the hypotheses of Proposition 1.18, there exists a set $X \subseteq \Sigma$ which is open and dense in Σ and has full Lebesgue measure, such that for all $(p_0, q_0) \in X$, the solution of* (FP) *with initial condition (p_0, q_0) is unique and continuous for all times $t \geq 1$.*

Finally, a detailed analysis of the phenomenon of non-uniqueness on the 'bad' set Z^* allows an even stronger statement.

Proposition 1.20 (Sparrow et al. [90, Proposition 2.2])**.** *Assume the bimatrix game (A, B) satisfies the hypotheses of Proposition 1.18, and additionally assume that A and B have*

maximal rank. Then the flow on the interior of the sets R_{ij} has a unique continuous extension everywhere (on Σ), except possibly to points in the subset of Z^* where one of the players is indifferent between at least three of her strategies (and the other between at least two). This remaining set has codimension three.

For zero-sum games (see Definition 1.24) and under suitable non-degeneracy conditions, a unique and continuous extension of the fictitious play flow to an even bigger set can be proved. In [92], van Strien showed that the vector field defined by (FP) in the case of a zero-sum bimatrix game $(M, -M)$ can be viewed as a particular case of a more general family of Hamiltonian vector fields. Among other things, his results imply conditions for (FP) to define a unique and continuous flow on all of $\Sigma \setminus \{(E^A, E^B)\}$ in a zero-sum game. We have adjusted the formulation of the following proposition by taking into account our Lemma 1.10 (so that by assuming a unique, completely mixed Nash equilibrium, only $n \times n$ games need to be considered).

Proposition 1.21 (van Strien [92])**.** *Let M be the payoff matrix of the first player in an $n \times n$ zero-sum game $(M, -M)$ with unique, completely mixed Nash equilibrium E.*

Let M^* denote some $(n-1) \times n$ or $n \times (n-1)$ matrix obtained by removing one row or column from M and subtracting it from each of the other rows or columns. Assume that for every such matrix M^* and every $r \geq 1$, each $r \times r$ minor[3] of M^* is non-zero, and for every $r \geq 2$, each $r \times r$ minor of M is non-zero. Then (FP) *defines a unique continuous flow on $\Sigma \setminus \{E\}$.*

Remark 1.22. Conditions such as the hypotheses of Propositions 1.18 or 1.21 can be found in the literature in various forms. They are usually transversality or non-degeneracy conditions on the entries of the matrices A and B, and hold for generic bimatrices (open and dense set of bimatrices, whose complement, as a subset of $\mathbb{R}^{2mn}$, has Lebesgue measure zero). For several older examples, see [80, 83, 88].

1.2.4 Fictitious play in 2×2 games

The case of 2×2 bimatrix games (A, B) forms the simplest possible 'building block' of fictitious play dynamics. It has been among the earliest results in this field, that in any 2×2 bimatrix game, all solutions of (FP) (and hence also of (BR)) converge to a Nash equilibrium point. A first proof of this has been given by Miyasawa [64], but a later proof by Metrick and Polak [62] is more geometric and conceptual, and we will briefly sketch its idea here, as it reveals the entire geometric structure of fictitious play in 2×2 games.

[3] An $r \times r$ minor of an $m \times n$ matrix M with $r \leq \min\{m, n\}$ is the determinant of the $r \times r$ matrix obtained by removing $m - r$ rows and $n - r$ columns from M.

Theorem 1.23 (Miyasawa [64], Metrick and Polak [62])**.** *For any* 2×2 *bimatrix game* (A, B)*, any orbit* $\{(p(t), q(t)), t \geq t_0\}$ *of* (FP) *with initial conditions* $(p(t_0), q(t_0)) = (p_0, q_0) \in \Sigma$ *for some* $t_0 \geq 0$ *converges to a Nash equilibrium* $(p^*, q^*) \in \Sigma$ *as* $t \to \infty$*.*

Sketch of proof. The basic idea is that there are only a few combinatorially distinct types of 2×2 games. Each of Σ^A and Σ^B can be simply viewed as a line segment or closed interval $[0, 1]$ with 0 corresponding to the first and 1 to the second of the respective player's strategies. Then in these coordinates, $\Sigma = \Sigma^A \times \Sigma^B$ is just the unit square $[0, 1] \times [0, 1]$.

A simple calculation shows that there is an $a \in \mathbb{R}$, such that $0 \in \mathcal{BR}_A(y)$ for $y \leq a$ and $1 \in \mathcal{BR}_A(y)$ for $y \geq a$, or vice versa. Note however, that a may not lie in $[0, 1]$, in which case one of R_0^A, R_1^A might be empty, while the other one is the entire Σ^B. Similarly, Σ^A consists of two closed intervals $[0, b] = R_0^B$, $[b, 1] = R_1^B$, or vice versa (one of which might be empty or a singleton $\{0\}$ or $\{1\}$). It follows that Σ consists of at most four rectangles R_{ij}, $i, j \in \{0, 1\}$, on each of which orbits of (FP) are straight line segments heading towards one of the vertices $(i, j) \in \Sigma$, see Figure 1.3. Hence, up to swapping strategies and players, there are only a few possible combinatorial types of 2×2 games. Most of these (Figure 1.3(a,b,d)) trivially give rise to convergent fictitious play dynamics, converging either to one of the vertices (i, j) or the interior point (a, b).

The only slightly more subtle case is when combinatorially, fictitious play leads to a 'spiralling' motion circling the point (a, b) (Figures 1.3(c) and 1.4). Convergence to (a, b) in this case can be proven by direct computation, showing that the motion has to be inward spiralling for geometric reasons. Alternatively and more conceptually, it can be shown that any such game is linearly equivalent to a zero-sum game (see Definitions 1.6 and 1.24), and then apply Theorem 1.27 from the next section. □

1.2.5 Convergence of fictitious play

The question of whether or not fictitious play dynamics converges has been among the most important and most studied questions in its investigation. In this section we present some of the classical results on classes of games for which convergence can be guaranteed.

We begin with the class of games that motivated fictitious play dynamics in the first place, the zero-sum games. Classically, a bimatrix game (A, B) is called *zero-sum*, if $A + B = 0$, that is, if the payoff to the first player is precisely the negative of the payoff to the second player, independently of the played strategy profile. As noted above, from a dynamical point of view, linearly equivalent games are not distinguishable, so we make the following definition.

Definition 1.24. An $m \times n$ bimatrix game (A, B) is *zero-sum*, if there exists a linearly equivalent $m \times n$ bimatrix game $(\tilde{A}, \tilde{B}) \sim (A, B)$ such that $\tilde{A} + \tilde{B} = 0$.

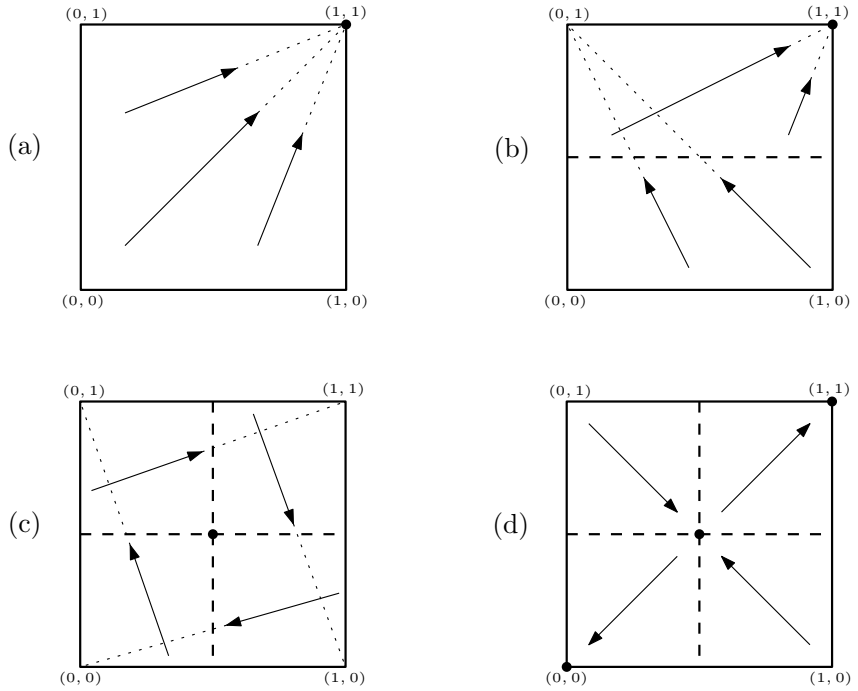

Figure 1.3: The different types of 2×2 bimatrix games. The spaces of mixed strategies Σ^A and Σ^B are each represented by a line segment $[0, 1]$. The space $\Sigma = \Sigma^A \times \Sigma^B$ consists of one, two, or four regions, on each of which all fictitious play trajectories run towards a single target (one of the vertices of Σ), indicated by arrows. In (a) and (b), all trajectories converge to the unique (pure) Nash equilibrium. The game in (c) is equivalent to a zero-sum game; trajectories spiral towards the unique, completely mixed Nash equilibrium. In (d), there are two pure (stable) and one completely mixed (saddle) Nash equilibrium.

The very first convergence result for discrete-time fictitious play (DFP) was proved in 1951 by Robinson[4].

Theorem 1.25 (Robinson [82])**.** *For any zero-sum bimatrix game, any orbit* $\{(p_k, q_k), k \geq 1\}$ *of discrete-time fictitious play* (DFP) *converges to the set of Nash equilibria as* $k \to \infty$.

The original statement of the theorem actually states that the payoffs of the players converge to the Nash equilibrium payoffs (the minmax payoff of the game), and this is only shown for bimatrix games (A, B) with $A + B = 0$, but the stronger statement of the above theorem can be deduced from that.

For continuous-time fictitious play, we use the following terminology.

Definition 1.26. A bimatrix game is said to have the *fictitious play property*, if every solution of (FP) (and hence (BR)) in this game converges to the set of Nash equilibria.

[4]In fact, fictitious play is sometimes called the Brown-Robinson process, due to its introduction by Brown [21, 22] and the first convergence result by Robinson [82].

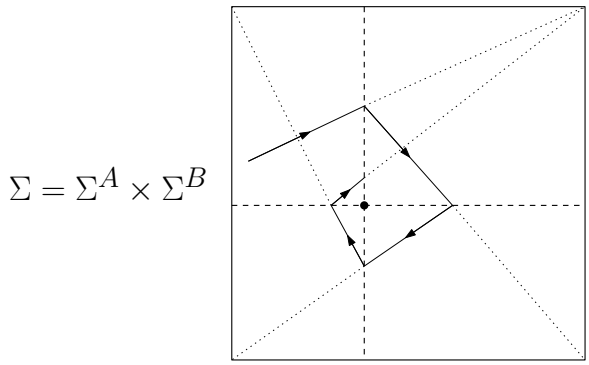

Figure 1.4: A fictitious play orbit in a 2×2 game with unique, completely mixed Nash equilibrium (the same type as in Figure 1.3(c)). Any orbit in this system spirals towards the equilibrium.

Convergence of continuous-time fictitious play in zero-sum games is actually easier to prove than in the discrete-time case, and we provide a sketch proof of the result. The argument is based on an explicitly constructed Lyapunov function for the fictitious play flow. It seems to go back to Brown and Robinson, just as the discrete-time case, but full formal proofs seem to have appeared in the literature only later, see, for example, Hofbauer [48, 52].

Theorem 1.27. *Every zero-sum game has the fictitious play property.*

Sketch of proof. We only discuss the case which is of greatest importance to us, where (A, B) is a game with a unique, completely mixed Nash equilibrium (E^A, E^B).

Without loss of generality, assume $A + B = 0$; convergence in this case implies convergence for all linearly equivalent games, since the dynamics and the Nash equilibria remain unchanged under linear equivalence. The key to a simple proof of this theorem is the following important function on Σ:

$$H(p, q) \coloneqq \mathcal{BR}_A(q)Aq + pB\mathcal{BR}_B(p) = \bar{A}(q) + \bar{B}(p), \qquad (p, q) \in \Sigma. \tag{1.3}$$

Since $A = -B$, we get that

$$\bar{A}(q) = \mathcal{BR}_A(q)Aq \geq pAq = -pBq \geq -pB\mathcal{BR}_B(p) = -\bar{B}(q), \tag{1.4}$$

so that $H(p, q) \geq 0$ for all $(p, q) \in \Sigma$. By Lemma 1.3, equality in (1.4) holds if and only if $(p, q) = (E^A, E^B)$. Additionally, $H \colon \Sigma \to \mathbb{R}^{\geq 0}$ is continuous on Σ and affine on each of the convex blocks R_{ij}. It follows that H is a function expressing a notion of distance to (E^A, E^B), with level sets topological spheres, which are made up of polytopes.

We will now show that H is strictly decreasing along non-equilibrium solutions

$(p(t), q(t))$ of (BR) (and hence (FP)). We skip certain technicalities involving almost-everywhere differentiability of fictitious play solution curves (see [43, 48, 52]), as well as the use of the so-called envelope theorem (see, for example, [91]) and present only the main idea of the calculation. Since $\mathcal{BR}_A(q)$ and $\mathcal{BR}_B(p)$ are locally constant (constant on each of the regions R_{ij}), by (BR) we get that

$$\begin{aligned}\frac{d}{dt}H(p,q) &= \frac{d}{dt}(\mathcal{BR}_A(q)Aq - pA\mathcal{BR}_B(p)) \\ &= \mathcal{BR}_A(q)A\dot{q} - \dot{p}A\mathcal{BR}_B(p) \\ &= \mathcal{BR}_A(q)A(\mathcal{BR}_B(p) - q) - (\mathcal{BR}_A(q) - p)A\mathcal{BR}_B(p) \\ &= -\mathcal{BR}_A(q)Aq + pA\mathcal{BR}_B(p) = -H(p,q) \leq 0,\end{aligned}$$

almost everywhere along a trajectory of (BR). Hence $\dot{H} = -H$, and $H(p(t), q(t)) = c \cdot e^{-t}$ for some $c > 0$, so that H is a Lyapunov function for (BR) and trajectories of the best response dynamics converge to (E^A, E^B) like e^{-t}. The same calculation for (FP) gives $\dot{H} = -H/t$ and $H(p(t), q(t)) = c/t$, so that fictitious play trajectories converge to (E^A, E^B) like $1/t$. □

Remark 1.28. In fact, the proof of Theorem 1.27 shows that if the unique Nash equilibrium of a zero-sum game is a single isolated point $(E^A, E^B) \in \text{int}(\Sigma)$ (that is, completely mixed), then it is an asymptotically stable fixed point of (FP) and (BR).

The following partially converse conjecture by Hofbauer remains open to this day.

Conjecture 1.29 (Hofbauer [48])**.** *A bimatrix game with a unique Nash equilibrium point in* int(Σ) *that is stable under the dynamics of* (FP) *(and hence* (BR)*) must be a zero-sum game.*

Another important question, both for practical computational reasons and from a theoretical point of view, is about the convergence rate of fictitious play dynamics in zero-sum games. As the proof of Theorem 1.27 shows, non-equilibrium solutions of (FP) converge to Nash equilibrium like $1/t$, and accordingly, solutions of (BR) converge like e^{-t} (see also [43]).

However, the same question for discrete-time fictitious play is more subtle. It is easy to see that a solution (p_k, q_k) of (DFP) in a zero-sum game converges to Nash equilibrium fastest, if it follows precisely a trajectory of (FP), and in this case the convergence rate is $1/k$. However, generically the phenomenon of 'overshooting' occurs, so that a solution of (DFP) 'falls behind' the corresponding continuous-time trajectory. To the best of the author's knowledge, the following result from 1958 by Shapiro is the strongest statement available to date.

Theorem 1.30 (Shapiro [86])**.** *Let* (A, B) *be an* $m \times n$ *zero-sum bimatrix game. Then solutions* (p_k, q_k) *of* (DFP) *converge to Nash equilibrium at a rate* $O(k^{-1/(m+n-2)})$.

It is known that Shapiro's estimate is sharp for 2×2 games, where convergence typically takes place at a rate $k^{-1/2}$, see [38]. On the other hand, its sharpness has not been demonstrated in games of higher dimension, and it is conjectured that the rate $O(k^{-1/2})$ is universal (independent of the dimension $m \times n$). The conjecture seems to go back to Karlin [53] (see also [43]).

Conjecture 1.31. *For any zero-sum bimatrix game* (A, B)*, solutions* (p_k, q_k) *of discrete-time fictitious play* (DFP) *converge to Nash equilibrium (uniformly) at a rate* $O(k^{-1/2})$.

At the time of submission of this thesis, the author is collaborating with Sebastian van Strien on a proof of this conjecture.

While zero-sum games are the most notorious class of games possessing the fictitious play property, there are several other classes for which convergence to Nash equilibrium is guaranteed: $2 \times n$ games, 'weighted potential games', 'quasi-supermodular' games of dimension $3 \times n$ and 4×4, 'quasi-supermodular' games with 'diminishing returns', and several others. We only state some of the results without giving proofs. To avoid certain degenerate cases, we will often make the assumption of non-degenerate games, according to the following definition.

Definition 1.32. An $m \times n$ bimatrix game (A, B) is called *non-degenerate*, if there is a unique best response to every pure strategy of either player. Equivalently, (A, B) is non-degenerate if the entries of A and B satisfy $a_{ij} \neq a_{i'j}$ and $b_{ij} \neq b_{ij'}$ for all $1 \leq i, i' \leq m$, $1 \leq j, j' \leq n$ with $i \neq i'$ and $j \neq j'$.

Theorem 1.33 (Berger [11])**.** *Any non-degenerate* $2 \times n$ *bimatrix game has the fictitious play property.*

Remark 1.34. For historical reasons it should be mentioned that for 2×3 games, this result has been proved earlier by Monderer and Sela [66, 85].

Berger's proof of Theorem 1.33 is rather geometrical: essentially, he shows that the formally n-dimensional fictitious play dynamics of a non-degenerate $2 \times n$ game can be reduced to a two-dimensional system, geometrically very much like a 2×2 game, and the only possibly non-trivial behaviour is an inward spiralling motion which necessarily converges to Nash equilibrium.

The following classes of games, whose definition for the first time involves an ordering of the players' strategies, are meaningful in certain economic contexts.

Definition 1.35. An $m \times n$ bimatrix game (A, B) is *supermodular*, if for $i < i'$ and $j < j'$,

$$a_{i'j'} - a_{ij'} > a_{i'j} - a_{ij} \quad \text{and} \quad b_{i'j'} - b_{i'j} > b_{ij'} - b_{ij}.$$

It is *quasi-supermodular*, if for $i < i'$ and $j < j'$,

$$a_{i'j} > a_{ij} \Rightarrow a_{i'j'} > a_{ij'} \quad \text{and} \quad b_{ij'} > b_{ij} \Rightarrow b_{i'j'} > b_{i'j}.$$

It has *diminishing returns*, if

$$\begin{aligned} a_{i+1,j} - a_{i,j} < a_{i,j} - a_{i-1,j} \qquad & \text{for all } i = 2, \dots, m-1 \text{ and all } j, \\ b_{i,j+1} - b_{i,j} < b_{i,j} - b_{i,j-1} \qquad & \text{for all } i \text{ and all } j = 2, \dots, n-1. \end{aligned}$$

Supermodular games are also known as games with 'strategic complementarities', and quasi-supermodular games as games with 'ordinal complementarities'. The following results have been proved in several weaker forms previously, we only state their most recent and strongest forms.

Theorem 1.36 (Berger [13, 14])**.** *Every non-degenerate quasi-supermodular game with either diminishing returns, or dimension* $3 \times n$ *or* 4×4*, has the fictitious play property.*

Theorem 1.37 (Hahn [41])**.** *Every supermodular* 3×3 *game has the fictitious play property.*

The fictitious play property for supermodular games of arbitrary dimension is an open question [15]:

Conjecture 1.38. *Every supermodular game has the fictitious play property.*

The last class of games with the fictitious play property that we are presenting are the 'weighted potential games'. These games generalise the so-called 'identical interest games': games in which all players have the same payoff function, so that their incentives are entirely aimed at cooperation and coordination. Identical interest games are complementary to zero-sum games, which can be seen as games of pure conflict, where all incentives for the players are entirely opposed to each other. The fictitious play property for identical interest games has been proved by Monderer and Shapley [67].

In a weighted potential game, payoffs are not identical, but payoff differences expressing the players' incentives are assumed to have a fixed positive ratio between players.

Definition 1.39. An $m \times n$ bimatrix game (A, B) is a *weighted potential game*, if there exists a matrix $P = (p_{ij}) \in \mathbb{R}^{m \times n}$ and positive weights $w_1, w_2 > 0$ such that for all i, i', j, j',

$$a_{ij} - a_{i'j} = w_1(p_{ij} - p_{i'j}) \quad \text{and} \quad b_{ij} - b_{ij'} = w_2(p_{ij} - p_{ij'}).$$

Theorem 1.40 (Monderer and Shapley [68])**.** *Every weighted potential bimatrix game has the fictitious play property.*

Most of the above convergence results hold for discrete-time fictitious play as well as for continuous-time fictitious play. For example, analogously to Theorem 1.33, discrete-time fictitious play converges in any non-degenerate $2 \times n$ game [11]. One of the available tools to transfer such results from continuous to discrete time is the following theorem due to Hofbauer. For more general theory on discrete-time processes related to differential inclusions and applications to dynamics in game theory, see Benaïm et al. [9, 10].

Theorem 1.41 (Hofbauer [48])**.** *For any bimatrix game, the set of limit points for any solution of* (DFP) *is invariant for the dynamics* (FP) *or* (BR)*. Hence the global attractor of these continuous-time processes contains the limit points of the discrete-time process.*

In general, the occurrence of 'overshooting' in discrete-time fictitious play makes it easier to prove convergence (and many other properties) for the continuous-time version, and indeed most of the more recent works on fictitious play focus exclusively on the latter.

Although its convergence to Nash equilibrium in certain games was one of the driving forces and main motivation for fictitious play, it turned out quite early that convergence need not hold for all types of games.

The first example of a bimatrix game without the fictitious play property was constructed in 1964 by Shapley [87]. His 3×3 example can be seen as a version of the well-known 'Rock-Paper-Scissors' game:

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}. \tag{1.5}$$

This game has a unique Nash equilibrium $(E^A, E^B) \in \operatorname{int}(\Sigma)$, with $E^A = (E^B)^\top = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Non-equilibrium trajectories of (FP) converge to a stable periodic orbit which can be interpreted as the two players trying to 'catch up' with their respective opponent's currently prevailing strategy. For instance, if one player empirically plays predominantly 'rock', the other player eventually switches to 'paper', which in turn causes the first player eventually to switch to 'scissors', and so on, until the pattern repeats after six such switches. The periodic orbit is a hexagon in Σ, which projects to a triangle in each of Σ^A and Σ^B. See Figure 1.2 for a similar orbit.

Remark 1.42. This type of periodic orbit is indeed typical for learning dynamics in bimatrix games; following this first example, such orbits are known as *Shapley polygons*, see [36].

1.2.6 Fictitious play as a Hamiltonian system

In this section, we give a somewhat different perspective on the dynamics of (FP) in the case of $n \times n$ zero-sum games with a unique, completely mixed Nash equilibrium. We briefly sketch how the fictitious play flow in this case gives rise to a Hamiltonian flow on S^{2n-3}, which will be important for our combinatorial investigations in Chapter 2, as well as motivate much of the work presented in Chapters 4 and 5. Fictitious play of zero-sum games has been studied extensively in the Hamiltonian context in [92]; our exposition here follows closely the introduction of the author's paper with van Strien [76].

Recall from Theorem 1.27, that every solution of (FP) in an $n \times n$ zero-sum game (A, B) with unique, completely mixed Nash equilibrium (E^A, E^B) converges to (E^A, E^B), and that a Lyapunov function for this is given by $H\colon \Sigma \to \mathbb{R}$ as in (1.3),

$$H(p,q) = \mathcal{BR}_A(q)Aq + pB\mathcal{BR}_B(p) = \max_i (Mq)_i - \min_j (pM)_j, \qquad (p,q) \in \Sigma,$$

where we denote $M = A = -B$. (Alternatively, we can think of H as a function $\mathbb{R}^{2n} \to \mathbb{R}$ with the same expression as above.)

It was shown in [92] that for an open set of full Lebesgue measure of $n \times n$ matrices M, each level set $H^{-1}(r)$, $r > 0$, is a topological $(2n-3)$-sphere bounding a convex ball, and $H^{-1}(0) = \{(E^A, E^B)\}$ (note that the dimension of $\Sigma = \Sigma^A \times \Sigma^B$ is $2n-2$).

The function H is continuous and piecewise affine, and $(\partial H/\partial q, \partial H/\partial p)$ is piecewise constant outside the indifference hyperplanes $Z^B_{ii'} \times \Sigma^B$ and $\Sigma^A \times Z^A_{jj'}$. On these hyperplanes, the derivatives can be thought of as multi-valued.

Now, consider the Hamiltonian differential inclusion

$$\frac{dp}{dt} \in \frac{\partial H}{\partial q}, \qquad \frac{dq}{dt} \in -\frac{\partial H}{\partial p}.$$

More generally, take the Hamiltonian vector field X_H associated to H and the symplectic 2-form $\omega = \sum_{ij} \omega_{ij} dp_i \wedge dq_j$, where (ω_{ij}) are the (real) coefficients of some constant non-singular $n \times n$ matrix Ω, and the corresponding differential inclusion is

$$\left(\frac{dp}{dt}, \frac{dq}{dt}\right) \in X_H(p,q). \tag{1.6}$$

Here X_H is defined by requiring that $\omega(X_H, Y) = dH(Y)$ for each vector field Y.

For an open set of full Lebesgue measure of $n \times n$ matrices M and Ω, the above differential inclusions have unique solutions for all initial conditions. The corresponding flow $(p, q, t) \mapsto \psi_t(p, q)$ is continuous and piecewise a translation flow, so that first return maps to hyperplanes in $H^{-1}(r)$ are piecewise affine maps.

While the above construction at first glance looks somewhat arbitrary, it turns out to be closely related to the fictitious play flow in zero-sum games with unique, completely mixed Nash equilibrium. Namely, let $(p, q) \in \Sigma \setminus \{(E^A, E^B)\}$ and let $l(p, q)$ denote the half-ray from (E^A, E^B) through (p, q). It can be shown that $l(p, q) \cap H^{-1}(1)$ defines a unique point in the topological $(2n-3)$-sphere $H^{-1}(1)$, so that we can define the projection

$$\pi \colon \Sigma \setminus \{(E^A, E^B)\} \to H^{-1}(1), \quad \pi(p, q) = l(p, q) \cap H^{-1}(1).$$

Crucially, the results in [92] imply the following proposition.

Proposition 1.43. *Let (A, B) be an $n \times n$ zero-sum bimatrix game with unique, completely mixed Nash equilibrium. Let ϕ_t be the flow of the corresponding continuous-time fictitious play dynamics* (FP)*, and let $H \colon \mathbb{R}^{2n} \to \mathbb{R}$ be the corresponding Lyapunov function. Then the projection under π of the flow ϕ_t to $H^{-1}(1)$ corresponds to a solution of a Hamiltonian system* (1.6)*.*

The formulation of the proposition is kept somewhat sketchy, for technical details see [76, 92]. The important conclusion is that solution curves of (FP) are projected to solution curves of (1.6). In other words, we can think of the Hamiltonian dynamics (1.6) on $H^{-1}(1)$ as a system induced by the fictitious play flow.

This *induced system* describes the 'spherical component' of the fictitious play dynamics, by 'projecting out' the converging 'radial' motion. The importance of this lies in the projective nature of the geometry of bimatrix games and their fictitious play dynamics. The regions R_{ij} can be seen as (convex) cones with apex (E^A, E^B), so that the strategies chosen by the players at the point (p, q) do not depend on the radial component, but only on $\pi(p, q)$. Therefore, the entire combinatorial complexity of the fictitious play dynamics is preserved in the induced system.

On the other hand, this system has certain nice properties and therefore allows to study the combinatorics of fictitious play dynamics in a sometimes more convenient setting than the original fictitious play. For instance, while the induced flow is still piecewise a translation flow, it is additionally volume-preserving (with an appropriately chosen volume form) and has no stationary points. In Chapter 2 we will exploit some of the advantages of the induced flow to study the combinatorial and ergodic properties of fictitious play.

In our main example of 3×3 games, first return maps of the induced flow to certain sections are continuous area-preserving piecewise affine planar maps, which can sometimes be visualised and studied more easily than the higher-dimensional flows. In Chapter 4 we will prove certain recurrence results relevant to such maps of compact planar sets, and in Chapter 5, we will study a concrete family of maps very similar to these first return maps of the induced flow.

1.3 Why fictitious play?

In this last section of the chapter, we would like to explain the motivation behind investigating fictitious play, beyond its original purpose as a computational tool to approximate Nash equilibrium in zero-sum games[5].

As a dynamical system, continuous-time fictitious play gives rise to a piecewise linear flow. A closer look at this flow in examples such as Shapley's Rock-Paper-Scissors system (1.5) reveals that its dynamics is far from trivial, and in fact can have highly complicated chaotic behaviour, with features that are unusual or even impossible in smooth dynamical systems.

In [90, 93], Sparrow, van Strien and Harris construct a one-parameter family of bimatrix games of a similar Rock-Paper-Scissors-like structure:

$$A_\beta = \begin{pmatrix} 1 & 0 & \beta \\ \beta & 1 & 0 \\ 0 & \beta & 1 \end{pmatrix}, \qquad B_\beta = \begin{pmatrix} -\beta & 1 & 0 \\ 0 & -\beta & 1 \\ 1 & 0 & -\beta \end{pmatrix}, \qquad \beta \in (0,1). \tag{1.7}$$

For $\beta = 0$, this corresponds to Shapley's example (1.5). Their investigation shows that fictitious play dynamics in these bimatrix games gives rise to a rich variety of different dynamical behaviours. Here, we summarise their results.

Note first, that (A_β, B_β) has a unique Nash equilibrium $(E^A, E^B) \in \operatorname{int}(\Sigma)$ given by $E^A = (E^B)^\top = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. By Proposition 1.18 and an analysis of the exceptional set Z^*, the differential inclusion (FP) defines a unique and continuous flow on all of $\Sigma \setminus \{(E^A, E^B)\}$.

If we denote the golden mean by $\sigma := (\sqrt{5} - 1)/2 \approx 0.618$, one can calculate that (A_β, B_β) is (linearly equivalent to) a zero-sum game if and only if $\beta = \sigma$. By Theorem 1.27, in that case all solutions of (FP) converge to the unique Nash equilibrium (E^A, E^B). For $\beta \neq \sigma$, the situation is more complicated.

Theorem 1.44 (Sparrow et al. [90, Theorem 3.3])**.** *For $\beta \in (0, \sigma)$, the flow defined by* (FP) *has a (locally) attracting periodic orbit, along which the players follow a 6-periodic cycle of strategy profiles:*

$$(1,2) \to (2,2) \to (2,3) \to (3,3) \to (3,1) \to (1,1) \to (1,2).$$

This orbit is a hexagon in Σ, which is a continuous deformation of the one in Shapley's

[5] In fact, the value of fictitious play as an algorithm for approximating Nash equilibrium is limited. Its convergence is relatively slow (see Theorems 1.27, 1.30 and Conjecture 1.31), and even its originator Brown wrote in [22] that *"it may be possible to use this method to get near, and some other method to finish, the calculation"*. Many more sophisticated algorithms with better convergence properties are known. On the other hand, its simplicity is also a strong argument in favour of fictitious play. Its computational efficiency makes it a valuable algorithmic tool even today, for example, in the area of artificial intelligence.

example for $\beta = 0$. It shrinks in diameter to 0 and converges to (E^A, E^B) as $\beta \to \sigma$, and it is not present for $\beta > \sigma$. This orbit is called the Shapley orbit.

For $\beta \in (\sigma, 1)$, there exists a periodic orbit, called the anti-Shapley orbit[6]*, with the following cycle of strategy profiles:*

$$(1,3) \to (1,2) \to (3,2) \to (3,1) \to (2,1) \to (2,3) \to (1,3).$$

There exists some $\tau \approx 0.915$ (the root of some expression), such that the anti-Shapley orbit is of saddle type when $\beta \in (\sigma, \tau)$ and attracting when $\beta \in (\tau, 1)$. For $\beta < \sigma$, this orbit does not exist.

Recall from Section 1.2.3 that $Z^* \subset \Sigma$ denotes the set of points where each of the players is indifferent between at least two of her strategies. The flow on this set is not defined a priori by (FP), but only by continuous extension from the sets R_{ij}. Moreover, since *both* players are indifferent between multiple strategies on Z^*, it cannot be guaranteed that trajectories move off this set instantaneously (as is the case when only one of the players is indifferent). So, in principle, orbits can remain within Z^* for a positive amount of time, or even be contained within it for all times.

Continuing their analysis, the authors look at the dynamics on and near a certain two-dimensional topological manifold $J \subset Z^*$ (the 'jitter set'). This manifold J is the union of six sets of the form $Z^B_{ii'} \times Z^A_{jj'}$ and is fully invariant for (FP). The authors show that the dynamics on and near J gives the most significant contribution to the richness of the entire system's behaviour. Since on J the players are undecided between multiple strategies, in its vicinity rapid switching between strategies can occur, and fictitious play orbits spend long times near J, as they 'spiral' along this set. The following theorem describes some of the fictitious play orbits on J.

Theorem 1.45 (Sparrow et al. [90, Theorem 3.4])**.** *For $\beta \in (0, \sigma]$, trajectories on J spiral towards the Nash equilibrium (E^A, E^B). For $\beta = \sigma$, the system (on J) undergoes a Hopf-like bifurcation, and for $\beta \in (\sigma, 1)$, there is a periodic orbit $\Gamma \subset J$, to which all orbits on J (except for the stationary Nash equilibrium orbit) converge. The diameter of Γ shrinks to 0 as $\beta \to \sigma$. Unlike in a usual Hopf-bifurcation, the decrease of the diameter is roughly linear for $\beta - \sigma > 0$ small.*

The trajectories on J spiralling in or out of (E^A, E^B) (for $\beta > \sigma$ and $\beta < \sigma$, respectively) do so in finite time, which emphasises the genuine non-uniqueness of the flow at (E^A, E^B) for $\beta \neq 0$.

[6]The name 'anti-Shapley orbit' derives from the fact that in the projections to Σ^A and Σ^B this orbit runs in the opposite direction (counter-clockwise instead of clockwise) to the Shapley orbit.

Furthermore, their analysis shows that the periodic orbit Γ on J 'organises' nearby dynamics. It is neither attracting or repelling, nor of saddle type; in [93], the authors call it 'jitter type': nearby orbits can remain close for very long times and jitter towards it and away from it in a strongly chaotic and random-walk-like fashion. Such orbits near Γ roughly follow the same sequence of six strategy profiles as Γ, but 'dither' around it, which gives rise to essentially periodic sequences of strategy profiles interspersed with bursts of 'undecided' back-and-forth strategy jumping between essential steps.

The following theorem summarises the abundance of periodic and quasi-periodic[7] orbits resulting from this. The term *essential period* refers to the actual underlying period of a sequence of strategy profiles, after removing the 'dithering' moves.

Theorem 1.46 (van Strien and Sparrow [93, Theorem 1.1])**.** *For each $\beta \in (0,1)$ and each $n \geq 1$, there are infinitely many fictitious play orbits γ_s, $s \in \mathbb{N}$, with periodic strategy profile sequences of period $N_s \to \infty$ as $s \to \infty$, but essential period $6n$, such that*

- *for $\beta \in (0, \sigma)$, these orbits reach (E^A, E^B) in finite time,*
- *for $\beta \in (\sigma, 1)$, these orbits are genuine periodic orbits of* (FP)*.*

Remark 1.47. Note that for $\beta \in (0, \sigma)$, this shows that not all orbits are attracted to the (locally) attracting Shapley orbit from Theorem 1.44.

Considering the special periodic orbit $\Gamma \subset J$ from Theorem 1.45 and taking some $x \in \Gamma$ and a section Z through x transversal to Γ, with corresponding first return map F, the authors show that:

- for each $k \in \mathbb{N}$, the map F has a sequence $x_n \in Z$ of periodic points of period k, converging to x as $n \to \infty$;
- F has infinite topological entropy;
- the dynamics acts like a random walk: there are annuli $A_i \subset Z$ centred at x and partitioning a neighbourhood of x (pairwise disjoint and geometrically shrinking to x), so that for each sequence $n(i) \geq 0$ with $|n(i+1) - n(i)| \leq 1$ there exists $z \in Z$, such that $F^i(z) \in A_{n(i)}$ for all $i \geq 1$.

Especially the last point displays a striking difference to any smooth dynamical system. It has the following consequence. Denote the flow of (FP) by ϕ_t, let $\epsilon > 0$ and define the local stable set corresponding to rate τ as

$$W_\epsilon^{s,\tau} := \left\{ x : \operatorname{dist}(\phi_t(x), \Gamma) \leq \epsilon \ \forall t \geq 0 \text{ and } \lim_{t\to\infty} \frac{1}{t} \log(\operatorname{dist}(\phi_t(x), \Gamma)) = \tau \right\}.$$

[7] We call an orbit of fictitious play *quasi-periodic*, if the strategy profiles played along this orbit repeat in a periodic fashion. See Section 2.1.6 for more details.

Then for each $\epsilon > 0$, $\tau > 0$ sufficiently small, $W_\epsilon^{s,\tau}$ is non-empty in any neighbourhood of Z. The same holds for the corresponding unstable sets $W_\epsilon^{u,\tau}$. This is in stark contrast to any smooth dynamical system, where one has that all orbits converging to a periodic orbit (in forward or backward time) do so at a fixed exponential rate and form a (stable or unstable) manifold. Here on the other hand, the stable and unstable sets are highly complicated and contain orbits of all (sufficiently small) exponential convergence rates.

We finish the discussion with a final remark on the robustness of these results.

Remark 1.48. All the above results are robust to small perturbations of the matrices A_β and B_β, that is, there is some $\epsilon > 0$, such that all of the above still holds for any bimatrix game (A, B) with $\|A - A_\beta\| < \epsilon$ and $\|B - B_\beta\| < \epsilon$ (see [93, Theorem 1.4]).

The above results are meant to motivate fictitious play from a mathematical point of view, as a dynamical system with a rather unique set of features. Another interesting side of fictitious play dynamics is revealed by its induced Hamiltonian flow, presented in Section 1.2.6. For an $n\times n$ zero-sum game with unique, completely mixed Nash equilibrium (like the game in (1.7) with $\beta = \sigma$), using the Lyapunov function H from (1.3) one can show that the motion projects to a spherical motion on $H^{-1}(1)$, which is a (topological) $(2n-3)$-sphere. In this case fictitious play is (almost, namely up to time parametrisation) the product of the radial motion towards (E^A, E^B) and an induced flow on S^{2n-3} (see Proposition 1.43).

This induced flow and a whole class of systems containing it are studied extensively in [92]. It is a continuous piecewise linear Hamiltonian flow, whose properties largely shadow the ones described above for the fictitious play flow. It has no stationary points, meaning that convergence to (E^A, E^B) in the fictitious play flow never happens along a well-defined direction, and orbits of (FP) keep spiralling in various ways around (E^A, E^B) as they converge to it. It admits certain first return sections, with Poincaré maps that are continuous, piecewise affine and volume-preserving.

In the particular example (1.7) with $\beta = \sigma$, it was shown that the induced flow on S^3 has a global first return section D, which is a topological disk made up of four triangles, with a Poincaré map R of a very special form. This first return map is continuous, piecewise affine, area-preserving, and extends continuously to the closure of D, by $R|_{\partial D} = \mathrm{id}$. (In fact, the boundary ∂D corresponds to a periodic orbit of the flow.) Model maps of this type are studied in Chapter 5.

The class of flows arising from fictitious play and its induced flow on S^{2n-3} is of interest as a model for Hamiltonian flows on the level-set of a given function $H\colon \mathbb{R}^{2n} \to \mathbb{R}$. Very little is generally known about such systems, and a possible approach is to study piecewise linear versions, as these often allow for much more explicit analytic computations, even without employing perturbation arguments as are often needed in smooth systems.

A similar paradigm of approaching general systems by looking at piecewise linear models can be traced to existing analogies such as circle diffeomorphisms and circle rotations, the quadratic map and the tent map, or the Hénon map and the Lozi map. Exploring the analogy between dynamics arising from smooth and piecewise affine Hamiltonians, as well as smooth and piecewise affine area-preserving maps, is perhaps the strongest mathematical motivation for studying fictitious play and various systems arising from it, and forms one of the main motivations for the investigations presented in this thesis.

Chapter 2

Combinatorics of fictitious play

This chapter is largely based on the publication [76] with Sebastian van Strien. Its focus lies on the combinatorics of continuous-time fictitious play dynamics, and its ergodic properties.

As we discussed in Section 1.3, in games of dimension 3×3 and higher, fictitious play dynamics can display highly complex behaviour, whereas in dimensions 2×2 (or, more generally, $2 \times n$), it is almost completely understood and classifies into a few simple dynamical types, with all its trajectories converging to the set of Nash equilibria (see Section 1.2.4 and Theorem 1.33 of the previous chapter). Therefore, in this chapter we focus our attention on 3×3 bimatrix games, and moreover we will be mostly dealing with games that have a unique, completely mixed Nash equilibrium.

In Section 2.1 we introduce a combinatorial (symbolic) description of fictitious play dynamics, show some of its basic properties, investigate admissible and realisable sequences of strategy profiles, and prove a classification result about the combinatorics of 3×3 zero-sum games. Then, in Section 2.2 we investigate the ergodic properties of fictitious play dynamics in 3×3 zero-sum games, by numerically studying certain first return maps of the fictitious play flow. We finish the chapter by posing a number of open questions and conjectures relating to both the combinatorial and ergodic properties of fictitious play.

2.1 Combinatorial description of fictitious play

Our first aim in this section is to obtain a combinatorial description of bimatrix games and an appropriate symbolic representation of their fictitious play dynamics. Next, we explore the possible combinatorial configurations for the classes of games we are interested in, in particular for zero-sum games.

For this, we begin by recalling some simple observations on the geometry of an $m \times n$ bimatrix game (A, B) and its implications for the shape of the orbits of continuous-time

fictitious play dynamics, that is, the solutions of the differential inclusion (FP) (or (BR)).

As discussed in the previous chapter, the first player's mixed strategy space Σ^A can be divided into n convex regions (polytopes) $R_j^B = \mathcal{BR}_B{}^{-1}(j)$, $j \in \{1, \dots, n\}$, and analogously Σ^B comprises of the regions $R_i^A = \mathcal{BR}_A{}^{-1}(i)$, $i \in \{1, \dots, m\}$.

Since $\mathcal{BR}_B \times \mathcal{BR}_A$ is constant on $R_{ij} = R_j^B \times R_i^A$, the orbits of (FP) and (BR) are continuous curves $(p(t), q(t)) \in \Sigma = \Sigma^A \times \Sigma^B$, which consist piecewise of straight line segments heading for vertices $(e_k, e_l) \in \Sigma$ whenever $(p(t), q(t)) \in R_{kl}$. The orbits only change direction when crossing one of a finite number of hyperplanes, the indifference sets $Z_{ii'}^B \subset \Sigma^A$ and $Z_{jj'}^A \subset \Sigma^B$, that is, whenever $\mathcal{BR}_A$ or $\mathcal{BR}_B$ (or both) become multi-valued. More precisely, $p(t)$ changes direction whenever $q(t)$ passes from R_i^A into $R_{i'}^A$ for some $i \neq i'$, and $q(t)$ changes direction whenever $p(t)$ crosses from R_j^B into $R_{j'}^B$, $j \neq j'$. See Figure 1.2 for an example with $n = m = 3$.

2.1.1 Coding of fictitious play

The partition of Σ into the convex blocks R_{ij} quite naturally gives rise to a coding of fictitious play. We codify an orbit $(p(t), q(t))$ by a (finite or infinite, one-sided) *itinerary*

$$(i_0, j_0) \to (i_1, j_1) \to \cdots \to (i_k, j_k) \to \cdots$$

indicating that there exists a strictly increasing sequence $(t_k)_{k\in\mathbb{N}}$, such that $(p(t), q(t)) \in R_{i_k j_k}$ for $t_k < t < t_{k+1}$, and $(i_k, j_k) \neq (i_{k+1}, j_{k+1})$ for $k \geq 0$. In other words, along the orbit $(p(t), q(t))$, the players choose the strategy profiles $(i_0, j_0), (i_1, j_1), \dots$, with switches between strategy profiles occurring at certain discrete times t_k, $k \in \mathbb{N}$. The following observation is essentially the 'improvement principle' of Monderer and Sela [66].

Lemma 2.1. *The itinerary of an orbit of* (FP) *for a fixed bimatrix game* (A, B) *can contain a step* $(i, j) \to (i', j')$ *only if* $a_{i'j} \geq a_{ij}$ *and* $b_{ij'} \geq b_{ij}$.

Next observe that, generically, only transitions of the form $(i, j) \to (i', j)$ and $(i, j) \to (i, j')$ can occur.

Lemma 2.2. *Let* (A, B) *be a non-degenerate bimatrix game. Then for almost all initial conditions* $(p_0, q_0) \in \Sigma$*, the components* $p(t)$ *and* $q(t)$ *of the corresponding solution of* (FP) *never switch directions simultaneously, that is, the itinerary of the orbit* $(p(t), q(t))$ *only contains transitions of the form* $(i, j) \to (i', j)$ *and* $(i, j) \to (i, j')$*,* $i \neq i'$*,* $j \neq j'$.

Proof. First, note that if the orbit $(p(t), q(t))$ undergoes a transition from (i, j) to (i', j') with $i \neq i'$ and $j \neq j'$, then there exists a $t > 0$, such that $(p(t), q(t)) \in Z_{ii'}^B \times Z_{jj'}^A$, that is, the orbit crosses the set where both players are indifferent between at least two strategies each.

Now, $p \in Z^B_{jj'} \subset \Sigma^A$ if and only if $(pB)_j = (pB)_{j'}$, that is,

$$p \cdot (\vec{b}_j - \vec{b}_{j'}) = 0,$$

where $\vec{b}_k$ denotes the kth column of the matrix B. From this one can see that $Z^B_{jj'}$ is all of Σ^A if $\vec{b}_j = \vec{b}_{j'}$, and has at least codimension 1, otherwise. By the non-degeneracy of the game, $b_{ij} \neq b_{ij'}$ for all i and j, j', so that only the second possibility can occur.

By the same reasoning, $Z^A_{ii'}$ has at least codimension 1, and it follows that the set $Z^B_{ii'} \times Z^A_{jj'}$ is of codimension 2 or more. Therefore, the set of initial conditions (p_0, q_0) whose fictitious play orbits eventually cross this set has at least codimension 1 in Σ, and hence zero Lebesgue measure. □

Using the above lemmas together with the non-degeneracy condition introduced in Definition 1.32 of Chapter 1, we can now make sure that for any $(i, j) \neq (i', j')$ there is only one possible transition direction between (i, j) and (i', j') which is realised by fictitious play orbits of an open set of initial conditions. Recall that a bimatrix game (A, B) is called non-degenerate, if $a_{ij} \neq a_{i'j}$ and $b_{ij} \neq b_{ij'}$ for all i, i', j, j' with $i \neq i'$ and $j \neq j'$; the non-degenerate bimatrices form an open dense subset of full Lebesgue measure in the space of $m \times n$ bimatrices.

Corollary 2.3. *Let (A, B) be a non-degenerate bimatrix game. Then for any $(i, j) \neq (i', j')$, at most one of the transitions $(i, j) \to (i', j')$ or $(i', j') \to (i, j)$ is realised by fictitious play orbits of an open set of initial conditions.*

2.1.2 Transition diagrams

To conveniently represent the collection of possible transitions in a given $m \times n$ bimatrix game (A, B), we can consider a graph (the 'transition graph') of $m \cdot n$ vertices (i, j), $i \in \{1, \ldots, m\}$, $j \in \{1, \ldots, n\}$, with a directed edge from (i, j) to (i', j') indicating that there exists a fictitious play orbit crossing (directly) from R_{ij} to $R_{i'j'}$. Here we only consider those transitions that occur along the orbits of an open set of initial conditions.

Notation 2.4. For a given non-degenerate game (A, B), we write $(i, j) \Rightarrow (i', j')$ to denote that the step $(i, j) \to (i', j')$ can occur along fictitious play orbits corresponding to an open set of initial conditions.

By Corollary 2.3 and Lemma 2.2, this leaves us with transitions of the form $(i, j) \Rightarrow (i', j)$ and $(i, j) \Rightarrow (i, j')$, with at most one transition direction between any two regions R_{ij} and $R_{i'j'}$.

For the case of non-degenerate 3×3 bimatrix games, we get a particularly simple visual representation of the transition graph, expressed in a so-called *transition diagram*

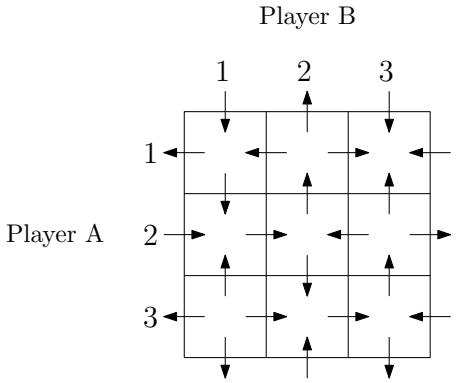

Figure 2.1: Example transition diagram of a 3×3 game. A block in row i and column j corresponds to the region R_{ij}, where i and j are the players' respective best response strategies in this subset of Σ. The arrows indicate which transitions between regions can occur along (open sets of) fictitious play trajectories in this game. Note that the left and right sides of the diagram should be though of as identified, as well as the top and bottom sides. Note also that the diagram only illustrates admissible itineraries: a formally admissible sequence of steps through the diagram cannot necessarily be realised by an actual fictitious play orbit.

as in Figure 2.1. The three rows and three columns of the diagram represent the regions R_i^A, $i = 1, 2, 3$, and R_j^B, $j = 1, 2, 3$, respectively. The arrows, only horizontal or vertical by Lemma 2.2, indicate the possible transitions between the regions, which by Corollary 2.3 always have a unique direction. For example, $(1, 2) \Rightarrow (1, 3)$ for a given bimatrix game if and only if in the first row of its transition diagram an arrow points from the second into the third column. Opposite sides of the diagram should be thought of as identified, so that possible transitions between the first and third rows and columns are indicated by arrows on the boundary of the diagram.

2.1.3 Realisable diagrams

One may now ask whether a given transition diagram can occur as the transition diagram of fictitious play dynamics for a bimatrix game (A, B), and how properties of a game relate to the combinatorial information given by its transition diagram.

A simple first observation is that no row of a transition diagram of a non-degenerate game can have all its horizontal arrows pointing in the same direction, as this would imply

$$a_{i,1} > a_{i,2} > \cdots > a_{i,n} > a_{i,1}$$

if the arrows points left, or

$$a_{i,1} < a_{i,2} < \cdots < a_{i,n} < a_{i,1}$$

if they point right. Analogously, no column of such diagram can have all its vertical arrows pointing in the same direction. More generally, the horizontal arrows in any row of a transition diagram have to correspond to a particular (strict) ordering of the entries of the matrix A in the corresponding row, and the vertical arrows in any column have to correspond to a (strict) ordering of the entries of the matrix B in the corresponding column.

It is not difficult to see that apart from this restriction, any transition diagram can be realised by appropriate choice of (A, B). However, our interest lies in bimatrix games whose fictitious play dynamics gives rise to non-trivial behaviour. In particular, the game should have no 'irrelevant' strategies, which are never a best response and hence never chosen by fictitious play. For this, we introduce the following classical notion from game theory.

Definition 2.5. Let (A, B) be an $m \times n$ bimatrix game. A strategy $i \in \{1, \dots, m\}$ of the first player is *dominated by strategy* $i' \in \{1, \dots, m\}$, if

$$e_i A q < e_{i'} A q$$

for all $q \in \Sigma^B$. We say that the strategy i is *dominated*, if there exists such a strategy i' dominating i. We define a dominated strategy of the second player analogously.

Remark 2.6. In classical game theory, such strategy is known as a 'strongly dominated strategy', as opposed to a 'weakly dominated strategy', which only requires $e_i A \tilde{q} < e_{i'} A \tilde{q}$ for a single $\tilde{q} \in \Sigma^B$, and $e_i A q \leq e_{i'} A q$ for general $q \in \Sigma^B$. For our purposes, the notion of strongly dominated strategy is sufficient, and we are referring to it simply as 'dominated'.

Lemma 2.7. *The first player's strategy i is dominated by her strategy i' if and only if $a_{ij} < a_{i'j}$ for all $j \in \{1, \dots, n\}$. The second player's strategy j is dominated by her strategy j' if and only if $b_{ij} < b_{ij'}$ for all $i \in \{1, \dots, m\}$.*

Proof. Clearly, if $a_{ij} < a_{i'j}$ for all j, then $e_i A e_j = a_{ij} < a_{i'j} = e_{i'} A e_j$. Now every $q \in \Sigma^B$ can be written as $q = \sum_j q_j e_j$ with $q_j \geq 0$ and $\sum_j q_j = 1$, so that

$$e_i A q = \sum_j q_j \cdot e_i A e_j < \sum_j q_j \cdot e_{i'} A e_j = e_{i'} A q.$$

Conversely, if $e_i A q < e_{i'} A q$ for every $q \in \Sigma^B$, then in particular, $a_{ij} = e_i A e_j < e_{i'} A e_j = a_{i'j}$ for every $j \in \{1, \dots, n\}$. The proof for the second player is similar. □

The lemma implies that the first player's strategy i is dominated by i' in the game (A, B) if and only if in the corresponding transition diagram, all (vertical) arrows between the rows i and i' point from i to i'. Similarly, the second player's strategy j is dominated by j' if and only if all (horizontal) arrows between the columns j and j' point from j to j'. See Figure 2.6(a)-(b).

Lemma 2.8. *If a bimatrix game (A, B) has a unique, completely mixed Nash equilibrium, then neither player has a dominated strategy.*

Proof. Suppose that one of the player has a dominated strategy. Without loss of generality we can assume that the first player's strategy 1 is dominated, so that $R_1^A = \emptyset$. But $E^B \in (\bigcap_i R_i^A) \cap \text{int}(\Sigma^B)$, whenever E^B is an isolated point in Σ^B, which is a contradiction. □

Consequently, in our combinatorial investigation of non-degenerate bimatrix games, transition diagrams of interest will be those that satisfy the following rules:

- no row has all its horizontal arrows pointing in the same direction;
- no column has all its vertical arrows pointing in the same direction;
- not all of the vertical arrows between any two rows point in the same direction;
- not all of the horizontal arrows between any two columns point in the same direction.

2.1.4 Admissible and realisable sequences

It is important to note that the partition of Σ given by $\{R_{ij} = R_j^B \times R_i^A : 1 \le i \le m,\ 1 \le j \le n\}$ does not have the nice properties of a Markov partition: there is no claim that every itinerary that can be obtained from the transition diagram of a bimatrix game can actually be realised by an orbit of the fictitious play dynamics. In fact, even a *finite* (formally admissible) sequence $(i_1, j_1) \to \cdots \to (i_K, j_K)$ cannot necessarily be realised as (part of) the itinerary of an actual fictitious play orbit.

Definition 2.9. Let (A, B) be a non-degenerate $m \times n$ bimatrix game. We say that a sequence $(i_1, j_1) \to \cdots \to (i_K, j_K)$, $K > 1$, with $i_k \in \{1, \ldots, m\}$ and $j_k \in \{1, \ldots, n\}$ for $1 \le k \le K$, is *admissible* for (A, B), if $(i_k, j_k) \Rightarrow (i_{k+1}, j_{k+1})$ for $k = 1, \ldots, K-1$.

An admissible sequence $(i_1, j_1) \Rightarrow \cdots \Rightarrow (i_K, j_K)$ is said to be *realisable*, if there exists $(p_0, q_0) \in \Sigma$, such that the itinerary of a corresponding fictitious play orbit $(p(t), q(t))$ with initial conditions (p_0, q_0) has $(i_1, j_1), \ldots, (i_K, j_K)$ as its first K entries.

Recall that in a bimatrix game with a unique, completely mixed Nash equilibrium (E^A, E^B), all the indifference sets $Z_{ii'}^A = R_i^A \cap R_{i'}^A \subset \Sigma^B$ and $Z_{jj'}^B = R_j^B \cap R_{j'}^B \subset \Sigma^A$ are

codimension-one hyperplanes that meet at E^B and E^A, respectively. It follows immediately that every admissible sequence $(i_1, j_1) \to (i_2, j_2)$ (length $K = 2$) is realisable. For non-degenerate 3×3 games, we can make the analogous statement for $K = 3$.

Proposition 2.10. *In a non-degenerate* 3×3 *bimatrix game with a unique, completely mixed Nash equilibrium, every admissible sequence of length* $K \leq 3$ *is realisable.*

Proof. We have already proved the statement for $K \leq 2$. For $K = 3$, let

$$(i_1, j_1) \to (i_2, j_2) \to (i_3, j_3)$$

be any admissible sequence, and assume without loss of generality $i_1 \neq i_2$. We consider two cases.

Case 1: $i_1 \neq i_2 \neq i_3$. It follows that $j_1 = j_2 = j_3$ and by non-degeneracy and Corollary 2.3, $i_1 \neq i_3$. So, up to relabelling strategies, we can assume $i_k = k$ and $j_k = 1$ for $k = 1, 2, 3$, so that the sequence to be realised is

$$(1, 1) \to (2, 1) \to (3, 1).$$

Since (A, B) has a unique Nash equilibrium $(E^A, E^B) \in \text{int}(\Sigma)$, every region $R^B_j \subset \Sigma^A$ is a convex region with non-empty interior. Let p_0 be an interior point of R^B_1. Solutions $(p(t), q(t))$ of (FP) satisfy $\dot{p} = \mathcal{BR}_A(q) - p$, therefore in particular, $\|\dot{p}\|$ is bounded and we can choose $T > 0$ such that $p(t) \in R^B_1$ for $0 \leq t \leq T$ for any such orbit with $p(0) = p_0$.

In Σ^B, the admissibility of $(1, 1) \to (2, 1) \to (3, 1)$ implies that $e_1 \in R^A_3$ (here we use the fact that there are only three strategies for each player) and that there are rays through e_1 that cross both indifference lines Z^A_{12} and Z^A_{23} transversally. These two lines meet at E^B. We can choose a ray through e_1 crossing Z^A_{12} and Z^A_{23} arbitrarily close to E^B. In particular, we can choose a ray r and a point $q_0 \in r \cap R^A_1$ such that a solution of $\dot{q} = e_1 - q$ with $q(0) = q_0$ satisfies $q_1 = q(t_1) \in r \cap Z^A_{12}$ and $q_2 = q(t_2) \in Z^A_{23}$ with $0 < t_1 < t_2 < T$.

By construction, the fictitious play orbit with initial conditions (p_0, q_0) has the desired itinerary, see Figure 2.2.

Case 2: $i_1 \neq i_2 = i_3$. Then $j_1 = j_2 \neq j_3$ and without loss of generality we can assume $i_1 = j_1 = j_2 = 1$, $i_2 = i_3 = j_3 = 2$, so that the sequence to be realised in this case is

$$(1, 1) \to (2, 1) \to (2, 2).$$

See Figure 2.3 for the construction to follow. Pick a point p_2 in the interior of the line segment Z^B_{12} and q_1 in the interior of Z^A_{12}. Let r_1 be the ray from $e_2 \in \Sigma^A$ through p_2 and let r_2 be the ray from $e_1 \in \Sigma^B$ through q_1. Since R^B_1 and R^A_2 have non-empty interior,

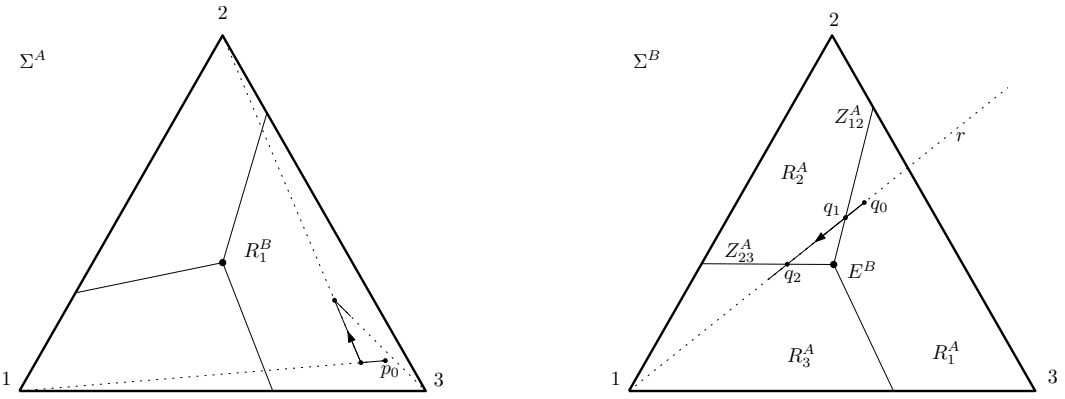

Figure 2.2: Orbit with itinerary $(1,1) \to (2,1) \to (3,1)$ constructed in case 1 of the proof of Proposition 2.10.

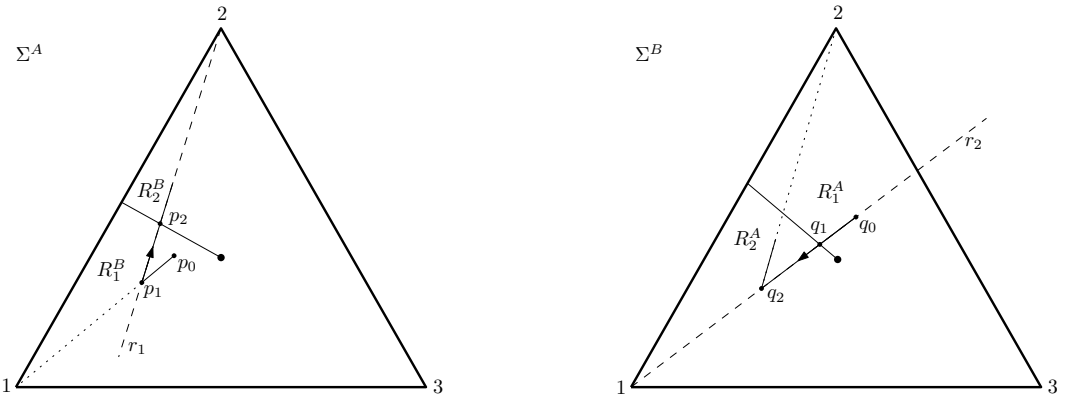

Figure 2.3: Orbit with itinerary $(1,1) \to (2,1) \to (2,2)$ constructed in case 2 of the proof of Proposition 2.10.

$r_1 \cap R_1^B$ and $r_2 \cap R_2^A$ are line segments of positive lengths. One can then choose points $p_1 \in r_1 \cap \text{int}(R_1^B)$ and $q_2 \in r_2 \cap \text{int}(R_2^A)$, such that the solutions of $\dot{p} = e_2 - p$ with $p(0) = p_1$ and $\dot{q} = e_2 - q$ with $q(0) = q_1$ satisfy $p(t_1) = p_2$ and $q(t_1) = q_2$ for some $t_1 > 0$. Since $p_1 \in \text{int}(R_1^B)$, we can now choose some (small) $t_2 > 0$ and points $p_0 \in R_1^B$, $q_0 \in R_1^A$, such that the solutions of $\dot{p} = e_1 - p$ with $p(0) = p_0$ and $\dot{q} = e_1 - q$ with $q(0) = q_0$ satisfy $p(t_2) = p_1$ and $q(t_2) = q_1$. Analogously, since $q_2 \in \text{int}(R_2^A)$, we can find $t_3 > 0$ such that the solution to $\dot{p} = e_2 - p$ with $p(0) = p_2$ and $\dot{q} = e_2 - q$ with $q(0) = q_2$ satisfy $p(t) \in R_2^B$ and $q(t) \in R_2^A$ for $0 < t < t_3$.

We conclude that the itinerary of the fictitious play orbit with initial conditions (p_0, q_0) begins with $(1,1) \to (1,2) \to (2,2)$, as desired (see Figure 2.3). □

One is tempted to think that the previous result can be extended to admissible sequences of length $n \geq 4$. However, as the next result shows, there are indeed admissible

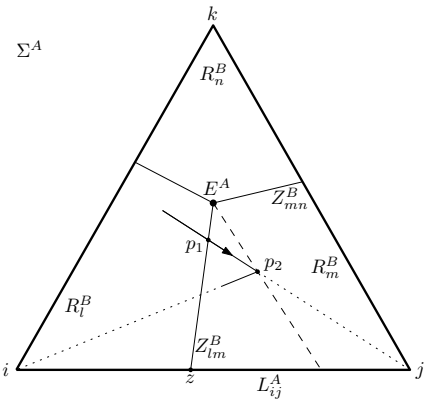

Figure 2.4: (Proof of Lemma 2.11) Only the projection to Σ^A of the orbit to be constructed is shown. To realise the sequence $(j,l) \to (j,m) \to (i,m) \to (i,n)$, the orbit needs to next go through a point $p_3 \in Z_{mn}^B$, which is impossible since $[p_2, p_3] = [p_2, e_i] \cap R_m^B$ lies within the cone with apex E^A spanned by Z_{lm}^B and the ray from E^A through $p_2 \in R_m^B$ (dashed line).

sequences of length 4 which cannot be realised by an actual fictitious play orbit. Moreover, these 3-step sequences can be shown not to be realisable under entirely combinatorial conditions on the transition diagram and without any considerations on the finer geometric structure of a given game. We start with a geometric lemma.

Lemma 2.11. *Let (A, B) be a non-degenerate 3×3 bimatrix game with a unique, completely mixed Nash equilibrium. Let $\{i,j,k\} = \{l,m,n\} = \{1,2,3\}$. Denote by S_{ij}^A the edge of Σ^A connecting the vertices e_i and e_j and assume that $Z_{lm}^B \cap S_{ij}^A \neq \emptyset$ in such a way that $(j,l) \Rightarrow (j,m)$ and $(i,m) \Rightarrow (i,l)$ (that is, such that e_i lies in the same half plane as R_l^B and e_j in the same half plane as R_m^B). Then the sequence*

$$(j,l) \to (j,m) \to (i,m) \to (i,n)$$

is not realisable (irrespective of whether it is admissible).

Proof. We only need to check the case that the given sequence is admissible. Assume some fictitious play orbit $(p(t), q(t))$ realises the given sequence as the first steps of its itinerary. Let $0 < t_1$ be the minimal time such that $p_1 = p(t_1) \in Z_{lm}^B$, $t_2 > t_1$ minimal such that $q(t_2) \in Z_{ij}^A$, and $t_3 > t_2$ minimal such that $p_3 = p(t_3) \in Z_{mn}^B$. Denote the point $Z_{lm}^B \cap S_{ij}^A$ by z and $p_2 = p(t_2)$. Then for $t_1 \le t \le t_2$, $p(t)$ lies on the line segment $[p_1, e_j]$ and for $t_2 \le t \le t_3$, on the line segment $[p_2, e_i]$. Observe that $[p_2, p_3] = [p_2, e_i] \cap R_m^B$ lies in the cone with apex E^A spanned by the two rays from E^A through $p_1 \in Z_{lm}^B$ and $p_2 \in R_m^B$, respectively. See Figure 2.4. This cone is contained in R_m^B and only intersects R_n^B in E^A, so that $p_3 \in [p_2, e_i]$ cannot lie on Z_{mn}^B, which is a contradiction. This finishes the proof. □

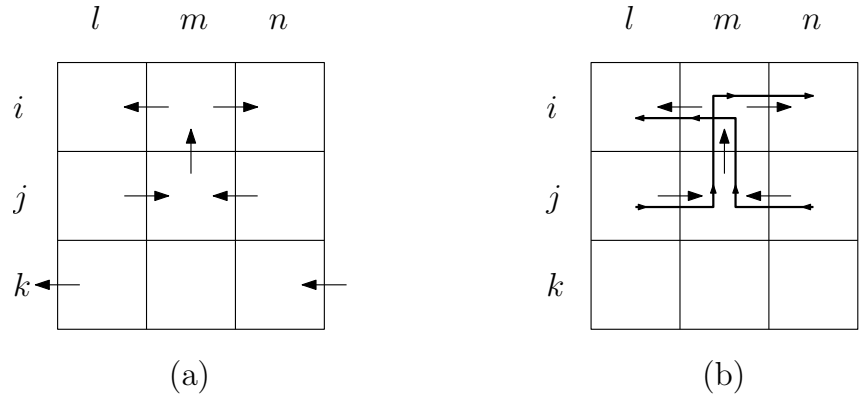

Figure 2.5:
(a) The prescribed combinatorial data in Corollary 2.12. Other arrows can be chosen freely.
(b) Proposition 2.13: Only one of the two shown itinerary paths can be realized.

Corollary 2.12. *Let (A, B) be a non-degenerate 3×3 bimatrix game with a unique, completely mixed Nash equilibrium. Let $\{i, j, k\} = \{l, m, n\} = \{1, 2, 3\}$, and let the following combinatorial data be given (see Figure 2.5(a)):*

- *$(j, l) \Rightarrow (j, m)$ and $(j, n) \Rightarrow (j, m)$ (that is, $b_{jm} > b_{jl}$ and $b_{jm} > b_{jn}$);*
- *$(i, m) \Rightarrow (i, l)$ and $(i, m) \Rightarrow (i, n)$ (that is, $b_{im} < b_{il}$ and $b_{im} < b_{in}$);*
- *$(j, m) \Rightarrow (i, m)$ (that is, $a_{im} > a_{jm}$);*
- *$(k, l) \Rightarrow (k, n)$ (that is, $b_{kn} > b_{kl}$).*

Then the sequence $(j, l) \to (j, m) \to (i, m) \to (i, n)$ is not realisable.

Proof. Note that by Lemma 2.8, there are no dominated strategies, and all the combinatorial restrictions given at the end of Section 2.1.3 apply. Together with the hypotheses, this leaves a certain number of possible configurations for the transition diagram of (A, B).

It can be checked that each of these gives rise to a geometric configuration (that is, a configuration of the players' indifference lines in Σ^A and Σ^B) satisfying the hypotheses of Lemma 2.11, which yields the conclusion. □

Proposition 2.13. *Let (A, B) be a non-degenerate 3×3 bimatrix game with a unique, completely mixed Nash equilibrium. Let $\{i, j, k\} = \{l, m, n\} = \{1, 2, 3\}$. Then at most one of the sequences*

$$(j, l) \to (j, m) \to (i, m) \to (i, n) \quad \textit{and} \quad (j, n) \to (j, m) \to (i, m) \to (i, l)$$

is realisable (see Figure 2.5(b)).

Proof. Assume for a contradiction that both sequences are realisable, and in particular, both are admissible. This implies that the first three combinatorial assumptions of Corollary 2.12 hold. If $b_{kn} > b_{kl}$, this corollary implies that the first sequence is not realisable. On the other hand, if $b_{kn} < b_{kl}$, then after swapping the roles of l and n, the same result implies that the second sequence is not realisable. Since the game is non-degenerate, $b_{kn} \neq b_{kl}$, and the proof is finished. □

More combinatorial rules such as the above can be deduced for non-degenerate 3×3 games (or particular types of such games), but we restrict ourselves to the few mentioned above to illustrate the principle. In the next section, we present a somewhat bigger combinatorial result of this type. Rather than dealing with single itineraries, it classifies all types of transition diagram occurring for 3×3 zero-sum games.

2.1.5 Combinatorics of zero-sum games

Recall from Definition 1.24 that we call a bimatrix game (A, B) zero-sum, if it is linearly equivalent (see Definition 1.6) to a game $(\tilde{A}, \tilde{B})$ which satisfies $\tilde{A} + \tilde{B} = 0$. This class of games is special in that it gives rise to a fictitious play flow which necessarily converges to the set of Nash equilibria (Theorem 1.27), and moreover, conjecturally is the only class of games with this property when the Nash equilibrium is an isolated point in the interior of Σ (Conjecture 1.29). Furthermore, as discussed in Section 1.2.6, by projecting the flow of a 3×3 zero-sum game to a level-set of a certain Lyapunov function (1.3), one can obtain a Hamiltonian 'induced flow' on S^3, which is an exciting object of study on it own, as it can provide insights into more general piecewise affine Hamiltonian flows on the three-sphere (see Section 1.3).

In this section we classify all 3×3 zero-sum games combinatorially, that is, we provide a list of all the possible transition diagrams that can occur for such games. It turns out that the zero-sum property imposes strong restrictions on the possible configurations of the transition diagram, and, up to relabelling strategies and swapping players, leaves only 23 possible transition diagrams.

We start by introducing a few combinatorial notions, where we always assume a non-degenerate bimatrix game (A, B) to be given.

Definition 2.14. An admissible sequence $(i_1, j_1) \to (i_2, j_2) \to \cdots \to (i_K, j_K) = (i_0, j_0)$, $K > 1$, is called *alternating cycle*, if after reversing the direction of all 'vertical' transitions along the sequence (that is, transitions $(i_k, j_k) \to (i_{k+1}, j_{k+1})$ with $i_k \neq i_{k+1}$), it forms a directed loop in the transition diagram of (A, B). More formally, the sequence is alternating, if the transition diagram satisfies one of the following:

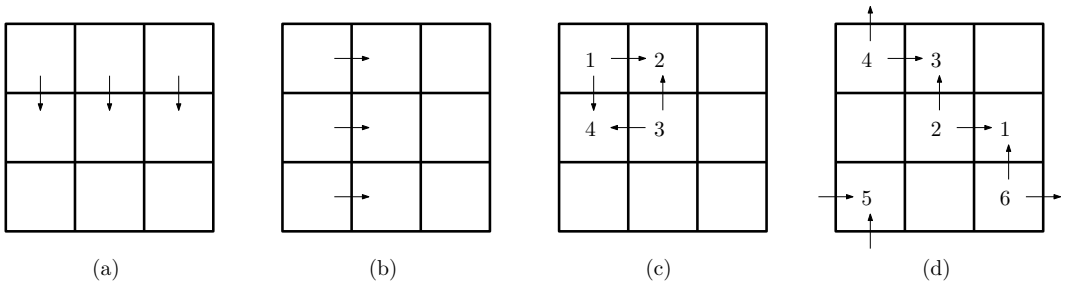

Figure 2.6: (a) Row 2 dominates row 1. (b) Column 2 dominates column 1.
(c),(d) Examples of alternating cycles.

- $(i_k, j_k) \Rightarrow (i_{k+1}, j_{k+1})$ whenever $i_k \neq i_{k+1}$, $(i_{k+1}, j_{k+1}) \Rightarrow (i_k, j_k)$ whenever $j_k \neq j_{k+1}$; or
- $(i_{k+1}, j_{k+1}) \Rightarrow (i_k, j_k)$ whenever $i_k \neq i_{k+1}$, $(i_k, j_k) \Rightarrow (i_{k+1}, j_{k+1})$ whenever $j_k \neq j_{k+1}$.

Examples of alternating cycles are shown in Figures 2.6(c) and (d).

Definition 2.15. We call (i, j) a *sink*, if it can be entered but not left by an admissible sequence (that is, $(i', j) \Rightarrow (i, j)$ and $(i, j') \Rightarrow (i, j)$ for $i' \neq i$ and $j' \neq j$). Conversely, we call it a *source*, if it can be left but not entered.

We can now formulate several facts for the transition diagram of a zero-sum game.

Lemma 2.16. *Let (A, B) be a non-degenerate zero-sum game. Then its transition diagram does not have alternating cycles.*

Remark 2.17. Without loss of generality we can only consider cycles in which the i- and j-component change alternatingly, which justifies the notion of alternating cycle. In fact, in 3×3 games this lemma reduces to saying that there are no alternating cycles of the two kinds depicted in Figures 2.6(c) and (d).

Proof. Assume that $A + B = 0$ (otherwise choose linearly equivalent matrices such that this holds and note that this does not change the arrows in the transition diagram). Recall that $(i, j) \Rightarrow (i', j)$ if and only if $a_{i'j} > a_{ij}$. Further, $(i, j) \Rightarrow (i, j')$ if and only if $b_{ij'} > b_{ij}$, which in a zero-sum game is equivalent to $a_{ij} > a_{ij'}$. It follows that an alternating cycle leads to a chain of inequalities

$$a_{i_0 j_0} > a_{i_1 j_1} > \cdots > a_{i_n j_n} = a_{i_0 j_0},$$

which is impossible. □

Lemma 2.18. *Let (A, B) be a non-degenerate 3×3 zero-sum game with a unique, completely mixed Nash equilibrium. Then its transition diagram does not have sinks or sources.*

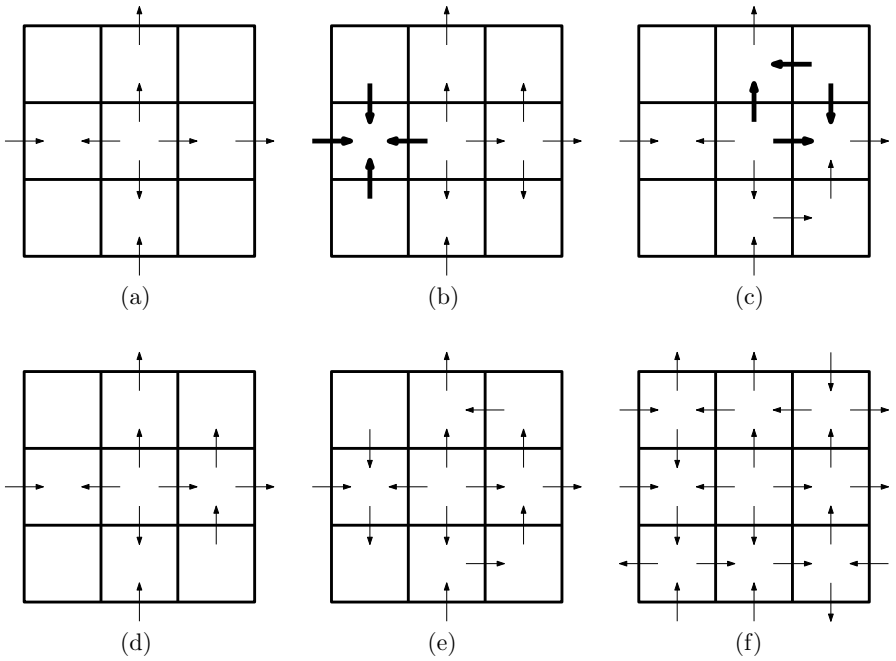

Figure 2.7: (Proof of of Lemma 2.18)
(a) General case for a diagram with a source.
(b) Case 1: sink in (2,1) contradicts zero-sum property.
(c) Case 2: alternating cycle contradicts zero-sum property.
(d) Case 3.
(e),(f) Case 3: necessarily following configuration.

Proof. To show that such game has no sinks, we use Theorem 1.27, that is, every orbit of fictitious play in a zero-sum game converges to the unique Nash equilibrium $(E^A, E^B) \in \text{int}(\Sigma)$. A sink in the transition diagram however would imply that all orbits of (FP) that start in R_{ij} for some i, j remain in it for all times. Since within R_{ij} all orbits are straight line segments with target point $(e_i, e_j) \in \partial\Sigma$, this can only be the case if they converge along straight line segments toward this point, contradicting the uniqueness of the Nash equilibrium in $\text{int}(\Sigma)$.

To rule out the existence of sources, suppose for a contradiction that such a zero-sum game with a source in its transition diagram exists. After possibly permuting rows and columns and swapping the roles of the two players, we can assume that the source is $(2, 2)$ and we have the (incomplete) diagram as seen in Figure 2.7(a). Let us now consider all four possible cases for the vertical arrows in $(2, 3)$:

- Case 1: $(2, 3) \Rightarrow (i, 3)$, $i = 1, 2$, that is, both arrows pointing *out of* $(2, 3)$.
  Either row 2 dominates row 1 or 3, or $(2, 1)$ is a sink, see Figure 2.7(b). While the former is impossible by Lemma 2.8, the latter contradicts the first part of this lemma.

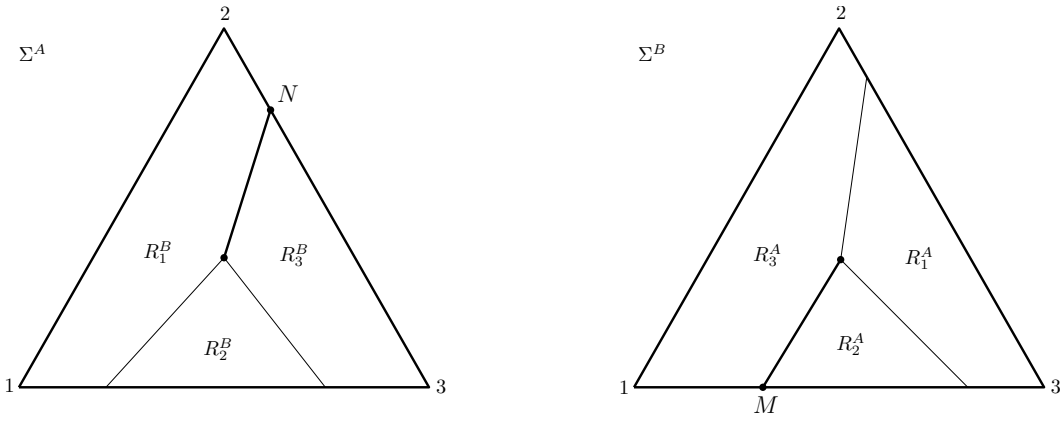

Figure 2.8: (Proof of Lemma 2.18) This configuration necessarily follows from the transition diagram in case 3 of the proof. Then (N, M) is a Nash equilibrium, contradicting the uniqueness of the completely mixed Nash equilibrium.

- Case 2: $(i, 3) \Rightarrow (2, 3)$, $i = 1, 2$, that is, both arrows pointing *into* $(2, 3)$.
  Then either column 2 dominates column 3, contradicting Lemma 2.8, or there is an alternating cycle, contradicting Lemma 2.16. See Figure 2.7(c).

- Case 3: $(3, 3) \Rightarrow (2, 3)$, $(2, 3) \Rightarrow (1, 3)$, both arrows point *upward* (Figure 2.7(d)).
  To avoid an alternating cycle and a dominated column, one necessarily has $(3, 2) \Rightarrow (3, 3)$ and $(1, 3) \Rightarrow (1, 2)$. Further, since row 2 may not dominate row 1 one gets $(1, 1) \Rightarrow (2, 1)$. Then $(2, 1) \Rightarrow (3, 1)$ is necessary to avoid a source in $(2, 1)$, see Figure 2.7(e). With some further deductions of the same kind one can show that the only possible transition diagram is the one shown in Figure 2.7(f).

  We can now deduce that Σ^A and Σ^B are partitioned into the regions R_j^B and R_i^A as shown in Figure 2.8. Consider the point $(Z_{13}^B \cap S_{23}^A) \times (Z_{23}^A \cap S_{13}^B) \subset \partial\Sigma$ denoted by (N, M) and note that $\mathcal{BR}_A(M)$ contains e_2 and e_3, hence all their convex combinations. Therefore, $N \in \mathcal{BR}_A(M)$. Analogously, $M \in \mathcal{BR}_B(N)$. Hence (N, M) is a Nash equilibrium, contradicting our assumption that the completely mixed Nash equilibrium is unique. (In fact, it follows from this configuration that there exist initial conditions arbitrarily close to the Nash equilibrium whose trajectories spiral off toward (N, M), and therefore the completely mixed Nash equilibrium cannot be stable for the dynamics.)

- Case 4: $(1, 3) \Rightarrow (2, 3)$ and $(2, 3) \Rightarrow (3, 3)$, that is, both arrows point *downward*.
  This case can be treated analogously to the previous one.

We have shown that a source in the diagram contradicts our assumption of a zero-sum game with unique, completely mixed Nash equilibrium, which finishes the proof. □

Remark 2.19. The above result remains true for general zero-sum games of dimension $n \times n$ with unique, completely mixed Nash equilibrium (recall from Lemma 1.10 that in a game with unique, completely mixed Nash equilibrium, both players necessarily have the same number of strategies). In this thesis we only need the statement for 3×3 games, and therefore restricted ourselves to the more elementary proof of this case.

Note however, that the above proof does not make use of the dimension to show that such a zero-sum game does not admit sinks in its transition graph. To prove that the transition graph of an $n \times n$ zero-sum game with unique, completely mixed Nash equilibrium also cannot have any sources, one could, for instance, use the induced Hamiltonian flow of fictitious play presented in Section 1.2.6. As was shown in [92], this induced flow is volume-preserving and has no stationary points. Both properties can be shown to contradict the existence of sources in the transition graph. One argument is that the block structure of fictitious play is preserved in the induced flow, and the existence of a convex block of positive volume with an entirely outward pointing vector field on its boundary contradicts the volume preservation of the flow.

Our aim is now to get a full characterisation of all combinatorial configurations that can occur in zero-sum games. This can be done after defining a suitable notion of combinatorially equivalent games.

Definition 2.20. We call two non-degenerate $m \times n$ bimatrix games (A, B) and (C, D) *combinatorially identical*, if they induce the same transition relation, that is, $(i, j) \Rightarrow (i', j')$ for (A, B) if and only if $(i, j) \Rightarrow (i', j')$ for (C, D).

We call two bimatrix games (A, B) and (C, D) *combinatorially equivalent*, if there exist permutation matrices P and Q such that (A, B) and (PCQ, PDQ) are combinatorially identical or $(B^\top, A^\top)$ and (PCQ, PDQ) are combinatorially identical.

Remark 2.21. Note that the bimatrix game $(B^\top, A^\top)$ is the game obtained from (A, B) by swapping the two players' roles.

The definition expresses the idea that games are combinatorially equivalent if they have the same transition diagram up to permutation of rows and columns. We can now state and proof our main result in this section.

Theorem 2.22. *The types of transition diagram (combinatorial equivalence classes) that can occur for a non-degenerate* 3×3 *zero-sum game with unique, completely mixed Nash equilibrium are precisely all those that satisfy the following (combinatorial) conditions:*

(1) No row of the diagram has three horizontal arrows pointing in the same direction and no column has three vertical arrows pointing in the same direction.

(2) No three horizontal arrows between two columns point in the same direction and no three vertical arrows between two rows point in the same direction.

(3) The diagram has no sinks.

(4) The diagram has no sources.

(5) The diagram has no alternating cycles.

This gives precisely 23 *distinct types of transition diagrams, up to combinatorial equivalence. These are listed in Figure 2.11*[1].

Throughout the proof of the theorem we will make use of the following notion.

Definition 2.23. We call an oriented loop of length four formed by the arrows in a transition diagram a *short loop*, see Figure 2.9(a). A short loop always has the form

$$(i, j) \Rightarrow (i', j) \Rightarrow (i', j') \Rightarrow (i, j') \Rightarrow (i, j)$$

and we indicate the vertex in the diagram encircled by such loop by a $\bullet$.

Remark 2.24. If we consider the 2×2 bimatrix game (A', B') formed by removing from A and B all rows except for i, i' and all columns except for j, j', then it can be checked that $(i, j) \Rightarrow (i', j) \Rightarrow (i', j') \Rightarrow (i, j') \Rightarrow (i, j)$ is a short loop if and only if (A', B') is linearly equivalent to a 2×2 zero-sum game with unique, completely mixed Nash equilibrium (see Figure 1.3(c) and Figure 1.4).

Proof of Theorem 2.22. We already know that (1) is true for any transition diagram of a game and (2) is equivalent to the game not having any dominated strategies, which is implied by the existence of a unique, completely mixed Nash equilibrium and Lemma 2.8. So the only conditions that are left to check are (3)-(5). By Lemmas 2.16 and 2.18, we already know that (3)-(5) are necessary conditions for a transition diagram to be realisable by a zero-sum game.

To show that (1)-(5) are also sufficient, we will proceed in two steps: we will show that combinatorially these conditions give rise to precisely 23 types of diagrams (up to permutation of rows and columns and transposition), and then we will give examples of zero-sum games realising each of these types. Because of the initially large number of possible transition diagrams, we will group them by the number of short loops contained in them.

[1]Coincidentally (or not?) the number 23 is the most sacred number for the religion called 'Discordianism'. In this religion 23 is the number of the highest deity, Eris, who is the Greek goddess of Chaos.

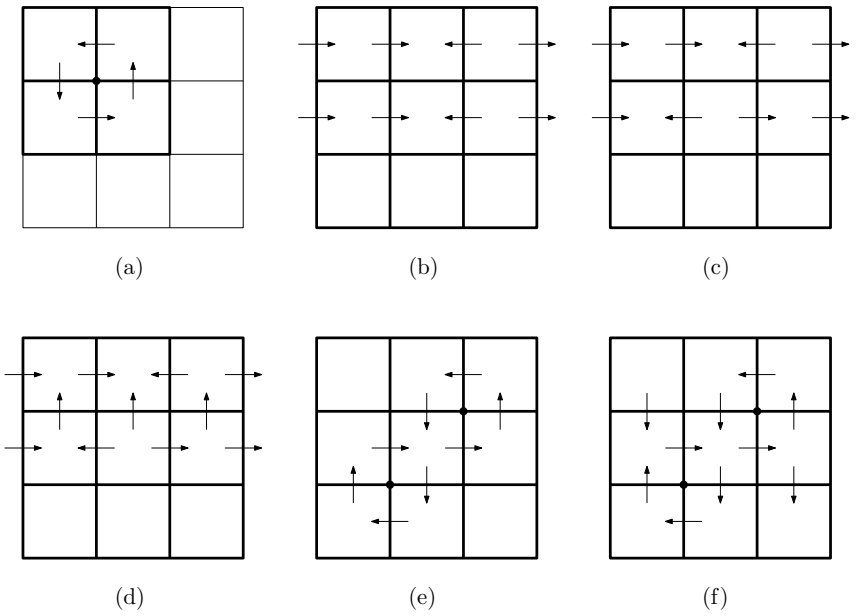

Figure 2.9: (Proof of Theorem 2.22)
(a) A short loop.
(b) Two rows coinciding at all positions.
(c) Two rows coinciding at one and differing at two positions.
(d) Lemma 2.26(1): if two rows differ at two positions, then either there is a short loop or one row dominates the other.
(e) Lemma 2.27: two short loops not between the same rows/columns.
(f) Lemma 2.26(2) applied to columns 1,2 and 2,3 in previous diagram.

Let us introduce the notion of rows (or columns) coinciding or differing at a position. We say that two rows i and i' *coincide between columns* j *and* j', if

$$(i, j) \Rightarrow (i, j') \quad \Leftrightarrow \quad (i', j) \Rightarrow (i', j'),$$

and they *differ at this position* (between columns j and j') otherwise. For example, in Figure 2.9(b) rows 1 and 2 coincide at all positions, whereas in Figure 2.9(c) they coincide at one and differ at two positions.

We now need a few technical lemmas about transition diagrams satisfying (1)-(5) of Theorem 2.22. The first lemma is obvious, and we omit a formal proof.

Lemma 2.25. *Two columns (or rows) of a* 3×3 *transition diagram satisfying (1)-(5) can have at most two short loops between them.*

Lemma 2.26. *In a* 3×3 *transition diagram satisfying (1)-(5), the following statements hold.*

(1) If two rows (columns) differ at two positions, then there is a short loop between these

rows (columns).

(2) If two rows (columns) differ at all three positions, then there are precisely two short loops between them.

(3) If two rows (columns) coincide at two positions, then there is a short loop between each of these rows (columns) and the third row (column). In particular, the diagram has at least two short loops.

(4) If two rows (columns) coincide at all three positions, then there are precisely two short loops between each of these rows (columns) and the third row (column). Then the diagram has precisely four short loops.

Proof. For (1), note that every 2×2 block obtained by deleting one row and one column from a 3×3 transition diagram either has two arrows pointing in the same direction or contains an alternating cycle or a short loop. Assume that two rows (columns) differ at two positions (Figure 2.9(c)). Since by hypothesis (5) of the theorem we don't allow alternating cycles, the only way a short loop between the two rows (columns) can be avoided is by having all arrows between them pointing in the same direction (Figure 2.9(d)). But this case is ruled out by hypothesis (2). Hence there is a short loop between the two rows (columns).

Essentially the same argument shows that statement (2) of the lemma holds.

If two rows (columns) coincide at two positions, then since no column (row) is allowed to be dominated, each of these rows (columns) differs at two positions from the third row (column). Statement (3) then follows from statement (1).

The same argument proves statement (4). The fact that the diagram then has precisely four short loops follows from Lemma 2.25 and the fact that there cannot be any short loops between the two rows (columns) that coincide at all positions. □

We now proceed to grouping all possible transition diagrams by the number of short loops contained in them.

Lemma 2.27. *A 3×3 transition diagram satisfying (1)-(5) has between three and six short loops.*

Proof. Note first that for any pair of rows (columns) of a transition diagram, at least one of the cases of Lemma 2.26 applies and we can make the following list of cases for two rows (columns), say i and i':

- i and i' coincide at 0 positions, then they have precisely 2 short loops between them.
- i and i' coincide at 1 position, then they have 1 or 2 short loops between them.

- i and i' coincide at 2 positions, then there is at most 1 short loop between them, and there is at least 1 short loop between each of them and the third row (column).
- i and i' coincide at 3 positions, then they have no short loops between them, and there are precisely 2 short loops between each of them and the third row (column).

It is clear from these that there cannot be such transition diagram without any short loops: pick any two rows, and in each of the above cases it follows that the diagram has at least one short loop.

Similarly, the transition diagram cannot have precisely one short loop. Suppose for a contradiction that without loss of generality there is a single short loop between rows 1 and 2. Now consider rows 2 and 3, which do not have a short loop between them. Then they must have at least 2 coinciding positions. But having 2 or more coinciding positions implies that there is also a short loop between rows 3 and 1, which is a contradiction.

We now show that there is no transition diagram with precisely two short loops. Assume first that such diagram exists and both short loops are between the same two rows (or columns), say rows 1 and 2. Applying the above rules to rows 2 and 3 we see that either there have to be more short loops between rows 2 and 3 or between rows 3 and 1, which is in either case a contradiction. So the only possibility left is that the two short loops are not between the same two rows or columns.

Here there are two cases to check: either both short loops run clockwise or one runs clockwise and one runs anti-clockwise (any other configuration leads to a combinatorially equivalent diagram). Assume the short loops have different orientation. Without loss of generality, we have the configuration shown in Figure 2.9(e). By Lemma 2.26(2) applied to columns $1, 2$ and $2, 3$ we get that $(1, 1) \Rightarrow (2, 1)$ and $(2, 3) \Rightarrow (3, 3)$ (Figure 2.9(f)). But now Lemma 2.26(1) applied to columns 1 and 3 implies that there is a third short loop. A similar chain of deductions shows that the case with both short loops having the same orientation also cannot happen. It follows that a transition diagram with two short loops is not possible.

Finally, the upper bound of six short loops follows directly from Lemma 2.25, which finishes the proof of the lemma. □

We can now state the last lemma needed for Theorem 2.22. The proof of the lemma consists of easy (but somewhat tedious) deductions of the only possible combinatorial configurations for the transition diagrams and we do not provide complete details. Lemma 2.26 is very useful to reduce the number of diagrams that have to be checked.

Lemma 2.28. *Up to combinatorial equivalence, there are precisely*

- 2 *non-equivalent transition diagrams with precisely* 3 *short loops,*

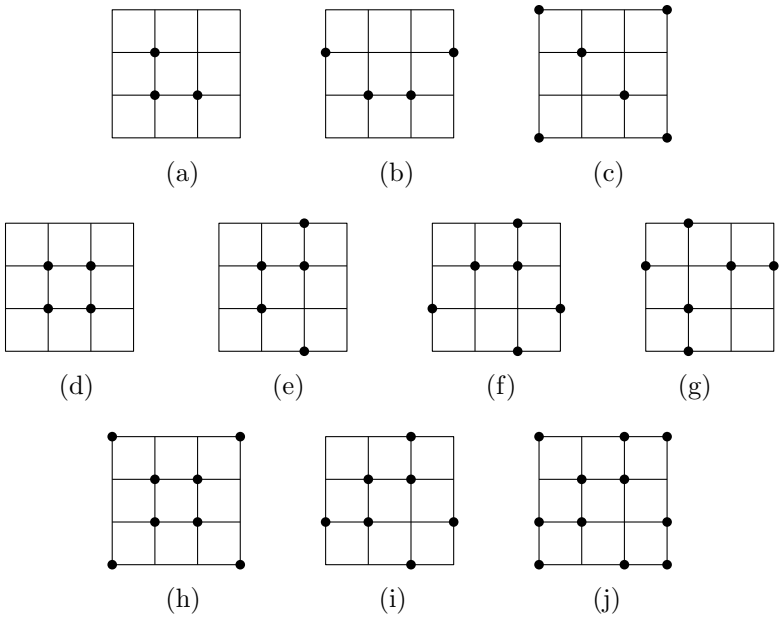

Figure 2.10: The possible (non-equivalent) configurations for a transition diagram containing 3 (a-c), 4 (d-g), 5 (h-i) or 6 (j) short loops.

- 15 *non-equivalent transition diagrams with precisely* 4 *short loops,*
- 5 *non-equivalent transition diagrams with precisely* 5 *short loops,*
- 1 *transition diagram with precisely* 6 *short loops,*

that satisfy conditions (1)-(5) of Theorem 2.22.

Proof. Up to combinatorial equivalence, there are three ways in which three short loops can be positioned, see Figure 2.10(a)-(c). It can be checked that only the first of these can give a transition diagram that satisfies (1)-(5), and there are two non-equivalent such diagrams.

Further, there are four ways to position four short loops (Figure 2.10(d)-(g)), the first three of which admit five non-equivalent transition diagrams each, whereas the last one contradicts (1)-(5).

The two ways to position five short loops (Figure 2.10(h)-(i)) admit two and three non-equivalent transition diagrams satisfying (1)-(5), respectively.

At last, by Lemma 2.25 it is obvious that up to combinatorial equivalence the only way to position six short loops is the one shown in Figure 2.10(j), and it is straightforward to check that there is only one possible transition diagram of this type. □

Together with Lemma 2.27, this shows that there are precisely 23 transition diagram types satisfying (1)-(5). We already know that (1)-(5) are necessary for a transition diagram

to correspond to a non-degenerate 3×3 bimatrix game with unique, completely mixed Nash equilibrium. A list of the 23 transition diagram types and example zero-sum game bimatrices realising them is given in Figure 2.11, showing that (1)-(5) are also sufficient.

This finishes the proof of Theorem 2.22. □

2.1.6 Periodic and quasi-periodic orbits

As a last item of the combinatorial description of fictitious play we discuss its periodic and quasi-periodic orbits. These will also be of importance in Section 2.2, where we numerically observe and categorise different types of fictitious play orbits.

We first introduce the game-theoretic notion of quasi-periodicity and investigate the relation to its usual mathematical definition. This notion was first introduced in [83].

Definition 2.29. We say that a solution of (FP) is *quasi-periodic* (in the game-theoretic sense), if its itinerary is periodic.

Quasi-periodic orbits of fictitious play have been studied by Rosenmüller [83] and Krishna and Sjöström [57]. Their main result shows that in almost all games of dimension greater than 2×2, cyclic convergence to a completely mixed Nash equilibrium (that is, along a game-theoretically quasi-periodic orbit) cannot occur for an open set of initial conditions.

Theorem 2.30 (Krishna and Sjöström [57])**.** *For (Lebesgue) almost all bimatrix games, if there is an open set of initial conditions whose orbits converge to a Nash equilibrium* (E^A, E^B) *along the same periodic itinerary, then* $\#\operatorname{supp}(E^A) = \#\operatorname{supp}(E^B) \leq 2$*, where* $\operatorname{supp}(x)$ *denotes the number of non-zero entries of the vector* $x \in \Sigma^A$ *or* Σ^B*.*

Remark 2.31. Note how this result relates to zero-sum games. We know from Theorem 1.27 that in a zero-sum game with unique, completely mixed Nash equilibrium (E^A, E^B) *all* initial conditions have orbits converging to (E^A, E^B), and in many examples there are open sets of such initial conditions converging with the same periodic itinerary (see, for instance, the examples in Section 2.2). This however does not contradict Theorem 2.30, since zero-sum games form a null set among all bimatrix games.

On the other hand, this theorem seems somewhat related to Hofbauer's converse conjecture (Conjecture 1.29) that in a non-zero-sum game, a completely mixed Nash equilibrium cannot be stable for fictitious play dynamics, as it shows that at least cyclic convergence to such equilibrium is exceptional in games of dimension greater than 2×2.

A priori the game-theoretic notion of quasi-periodicity is different from the usual mathematical definition (of an orbit which is dense in an invariant torus). In particular, a game-theoretically quasi-periodic orbit might be an actual periodic or quasi-periodic orbit,

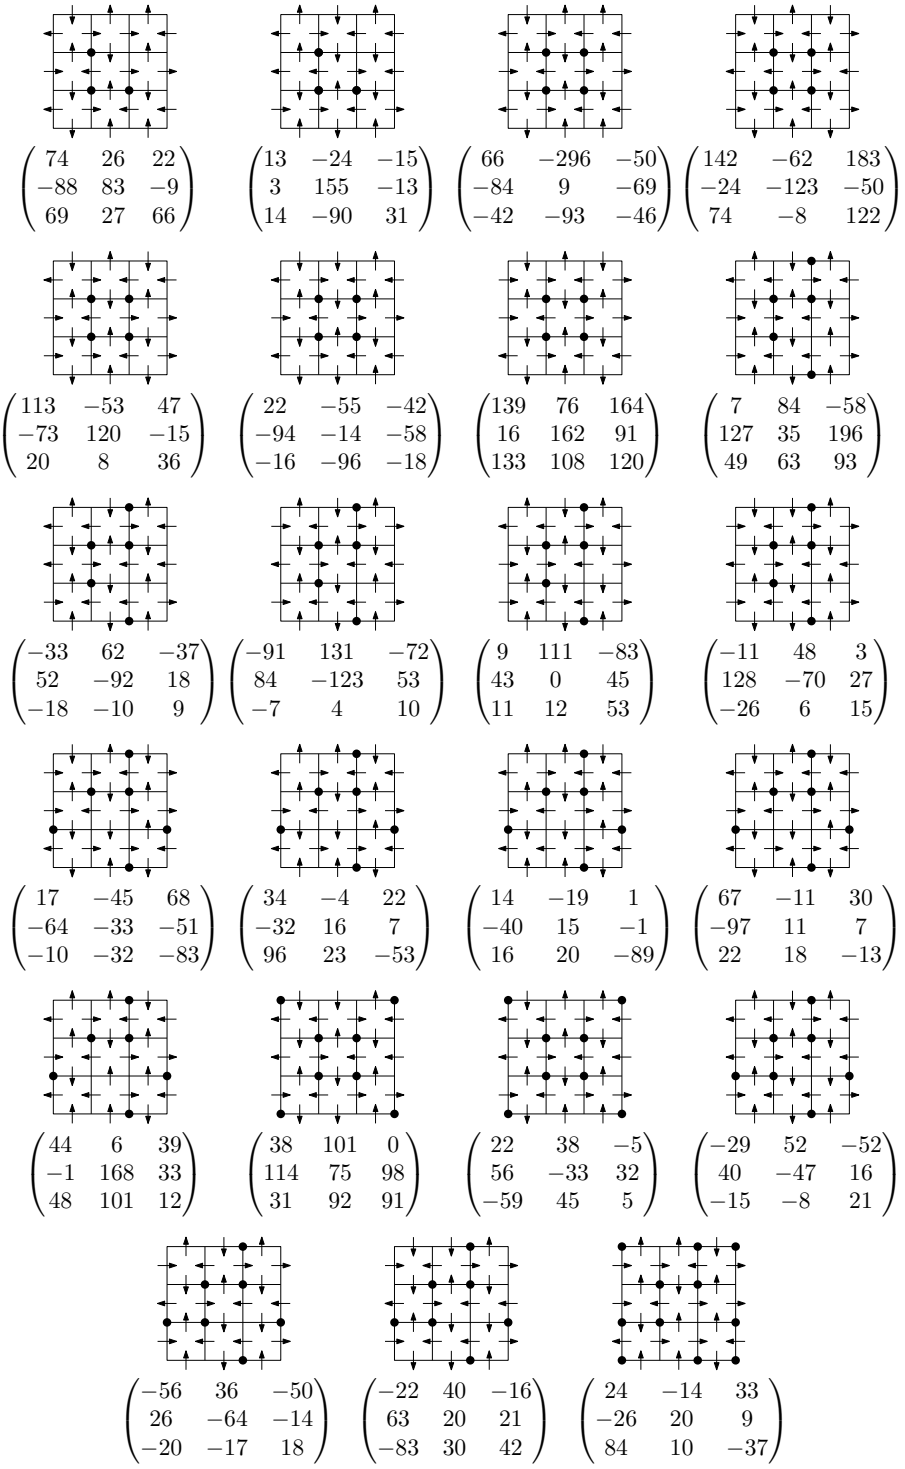

Figure 2.11: The 23 transition diagram types as in Theorem 2.22, sorted by the number of short loops contained in them, together with the respective example matrices A, such that the bimatrix games $(A, -A)$ realise the given diagram types.

but it might also be an orbit which converges in a cyclic way, periodically switching strategies as it approaches an equilibrium point. However, we will show that the two notions are closely related.

As before, we restrict our attention to the case of 3×3 zero-sum games with unique, completely mixed Nash equilibrium, although the arguments are similar for any higher-dimensional $n \times n$ zero-sum game. Every orbit of fictitious play in such a game converges to Nash equilibrium, so neither periodic nor quasi-periodic orbits (in the classical sense) can occur. Recall from Section 1.2.6 that fictitious play in this case induces a Hamiltonian flow on S^3 given by the differential inclusion (1.6), which describes its 'spherical' motion obtained by a certain central projection from the Nash equilibrium and ignores its (converging) radial direction. This induced flow still contains the entire symbolic information of the original fictitious play flow. In particular, an orbit of this Hamiltonian system has the same game-theoretic itinerary as any fictitious play orbit of which it is a projection.

This makes the induced Hamiltonian flow (1.6) perfectly suitable to look at game-theoretic quasi-periodicity, as it reduces the dimension of the studied dynamical system without removing any relevant information. In fact, like the fictitious play flow, the induced flow is continuous and a piecewise translation flow, while additionally it is volume-preserving and has no stationary points. To tackle the question of quasi-periodicity, we will consider its first return maps to the (two-dimensional) planes on which its trajectories change direction, that is, where one of the best response correspondences is multi-valued. These are precisely the projections of the hyperplanes $Z^B_{ii'} \times \Sigma^B$ and $\Sigma^A \times Z^A_{jj'}$ to the topological three-sphere $H^{-1}(1)$ (see Section 1.2.6).

Let S be such a plane, let $x \in S$ have a quasi-periodic orbit under the induced flow with (infinite) periodic itinerary $I = I(x)$, and let $\hat{T}$ be the first return map to S (defined on the non-empty subset of S of points whose orbits return to S). Note that $\hat{T}$ acts as a shift by a finite number of symbols on the itinerary of x. In particular, there exists $n \geq 1$, such that $\hat{T}^n(x)$ has the same itinerary as x. Let us denote $T = \hat{T}^n$ so that each point in the T-orbit of x has the same periodic itinerary I, and denote $U = U_I = \{z \in S : I(z) = I(x) = I\}$.

Lemma 2.32. *The set U is convex, $T(U) \subseteq U$, and $T|_U : U \to U$ is affine.*

Proof. Convexity follows from the convexity of the indifference planes and the fact that the flow in each region R_{ij} follows the 'rays' of a central projection. By the definition of T we have that $T(U) \subseteq U$. It is proved in [92] that T is area-preserving and piecewise affine on S. Since all points in U have the same itinerary, it follows that $T|_U$ is indeed affine. □

Theorem 2.33. *Let $x \in S$ correspond to a (game-theoretically) quasi-periodic orbit of the induced flow of a 3×3 zero-sum game with unique, completely mixed Nash equilibrium,*

where S is an indifference plane and T is the return map to S, such that $I(T(x)) = I(x)$. Let $U \subseteq S$ be the set of points with the same itinerary as x. Then one of the following holds:

(1) The orbit of x is periodic and $T^n(x) = x$ for some $n \geq 1$.

(2) There is a point $x_0 \in U$ whose orbit is a saddle periodic orbit; that is, for T, it is a hyperbolic fixed point with a contracting and an expanding direction. The T-orbit of x lies on its stable manifold and converges to x_0, that is, $T^n(x) \to x_0$ as $n \to \infty$.

(3) The T-orbit of x lies on a T-invariant circle (or more generally, an ellipse) and x corresponds to a quasi-periodic orbit of the Hamiltonian dynamics (in the usual sense), that is, its orbit under the flow is dense in an invariant torus.

In the third case, U is a disk and $T|_U \colon U \to U$ is a rotation by an irrational angle (in suitable linear coordinates). In this case (3) *holds for every $x' \in U$.*

Proof. The result follows immediately from the previous lemma, that is, the fact that $T|_U$ is a planar affine transformation with $T(U) \subseteq U$. □

The theorem shows that every quasi-periodic orbit (in the game-theoretic sense) is actually either periodic or quasi-periodic in the usual sense, or converges to a periodic orbit. Conversely, every quasi-periodic orbit in the usual sense, which lies on a torus that only intersects the indifference surfaces along whole circles (and never just partially along an arc) is clearly quasi-periodic in the game theoretic sense. Throughout the rest of this chapter, we will always refer to the game-theoretic definition, when using the notion of quasi-periodicity. We conclude with the following corollary.

Corollary 2.34. *If the flow induced by fictitious play dynamics in a 3×3 zero-sum bimatrix game has an orbit with periodic itinerary, then it also has an actual periodic orbit.*

2.2 Numerical investigation of ergodic properties

In this section we present some numerical observations on the behaviour of fictitious play dynamics for zero-sum games. We investigate a few different aspects:

- the time fraction that different orbits spend in each of the regions R_{ij},
- the frequencies with which different orbits visit the regions R_{ij} and the transition probabilities for transitions between regions,
- the different types of orbits that can occur and their itineraries (periodic, quasi-periodic, space-filling).

The systems we consider are randomly generated examples of zero-sum games of different combinatorial types (see Section 2.1.5), for which we look at the induced Hamiltonian dynamics (see Section 1.2.6). For randomly chosen initial points we compute the orbits of the best response dynamics (more precisely, its induced analogue) and study the time fractions spent in each region R_{ij} and the frequencies with which the orbits visit the regions. Note that here for the first time we make use of the autonomous version of fictitious play called best response dynamics, given by the differential inclusion (BR). Other than fictitious play orbits, best response orbits do not slow down over time, and the time they spend in each region relates more straightforwardly to the arc length of the orbit passing through that region. Regarding the presented types of orbit we do not claim to give an exhaustive account of occurring types but rather a list of examples illustrating a few key concepts.

Formally, for an orbit of the best response dynamics $(p(t), q(t))$, $t > 0$, with itinerary $(i_0, j_0) \to (i_1, j_1) \to \cdots \to (i_k, j_k) \to \cdots$ and switching times (t_n) we define

$$P_{ij}(n) = \frac{1}{t_n} \int_0^{t_n} \chi_{ij}(p(s), q(s))\, ds\,,$$

where χ_{ij} is the characteristic function of the region R_{ij}. Alternatively, we record the number of times that each region is being visited by an orbit and compute the frequencies:

$$Q_{ij}(n) = \frac{1}{n} \sum_{k=0}^{n-1} I_{ij}(i_k, j_k)\,, \quad \text{where } I_{ij}(i_k, j_k) = \begin{cases} 1 & \text{if } (i_k, j_k) = (i, j), \\ 0 & \text{otherwise.} \end{cases}$$

We write $P = (P_{ij})$ and $Q = (Q_{ij})$ for the matrices containing all the above frequencies.

Throughout the following examples we look at orbits of the first return maps for the best response dynamics to certain surfaces of section. The most convenient choice for such a surface is an indifference plane (either $Z^B_{ii'} \times \Sigma^B$ or $\Sigma^A \times Z^A_{jj'}$). We will mostly use the indifference planes of the form $Z^B_{ii'} \times \Sigma^B$, denoted by

$$S_{ii'} = \{(p, q) \in \Sigma : \{i, i'\} \subset \mathcal{BR}_B(p)\}\,.$$

Example 2.35 (Uniquely ergodic case)**.** Let the zero-sum bimatrix game (A, B) be given by

$$A = \begin{pmatrix} 22 & 34 & -4 \\ 7 & -32 & 16 \\ -53 & 96 & 23 \end{pmatrix}, \qquad B = -A.$$

We numerically calculate orbits with itineraries of 10^4 transitions for several hundreds of randomly chosen initial conditions. For all of these orbits, the evolution of $P(n)$

and $Q(n)$ indicates convergence to

$$P \approx 10^{-2} \cdot \begin{pmatrix} 13 & 5 & 27 \\ 14 & 5 & 27 \\ 3 & 1 & 5 \end{pmatrix}, \qquad Q \approx 10^{-2} \cdot \begin{pmatrix} 12 & 9 & 19 \\ 9 & 13 & 15 \\ 10 & 5 & 8 \end{pmatrix}.$$

In Figure 2.12, the evolution of some of the $P_{ij}(n)$ and $Q_{ij}(n)$ along an orbit is shown. This or very similar statistical behaviour is observed for all sampled initial conditions. It seems to suggest that initial conditions with quasi-periodic orbits have zero or very small Lebesgue measure in the phase space of best response dynamics for this bimatrix game, as quasi-periodicity in all our experiments leads to very rapid convergence to certain frequencies. Most of the space seems to be filled with orbits that statistically resemble each other in the sense that they all visit certain portions of the space (the regions R_{ij}) with asymptotically equal (or very close) frequencies. The same seems to hold for the fraction of time spent in each region by the orbits.

Figure 2.13 shows the intersections of a best response orbit for this game with S_{12}, S_{23} and S_{31}, that is, an orbit of the first return map of these surfaces. Each $S_{ii'}$ consists of three triangular pieces, corresponding to the three pieces of the indifference set between regions R_{ij} and $R_{i'j}$ for $j = 1, 2, 3$. Inside each of these triangles, the orbit seems to fill the space rather uniformly, suggesting ergodicity (of Lebesgue measure). If the best response dynamics had invariant tori, these would appear on all or some of these sections as invariant circles whose interior cannot be entered by orbits starting outside. Judging from the above observations, in this example they either don't exist or have very small radius.

Example 2.36 (Space decomposed into ergodic and elliptic regions)**.** In this example we consider a bimatrix game from the family of games studied in [90, 93], which we discussed in Section 1.3. Let the zero-sum bimatrix game (A, B) be given by

$$A = \begin{pmatrix} 1 & 0 & \sigma \\ \sigma & 1 & 0 \\ 0 & \sigma & 1 \end{pmatrix}, \qquad B = -A,$$

where $\sigma = (\sqrt{5} - 1)/2 \approx 0.618$ is the golden mean.

Two types of orbit can be (numerically) observed for the best response dynamics of this game. The first type resembles the orbits in the previous example. The empirical frequencies $P_{ij}(n)$ and $Q_{ij}(n)$ along such orbits initially behave erratically but seem to converge to certain values or narrow ranges of values, which are the same for all such orbits:

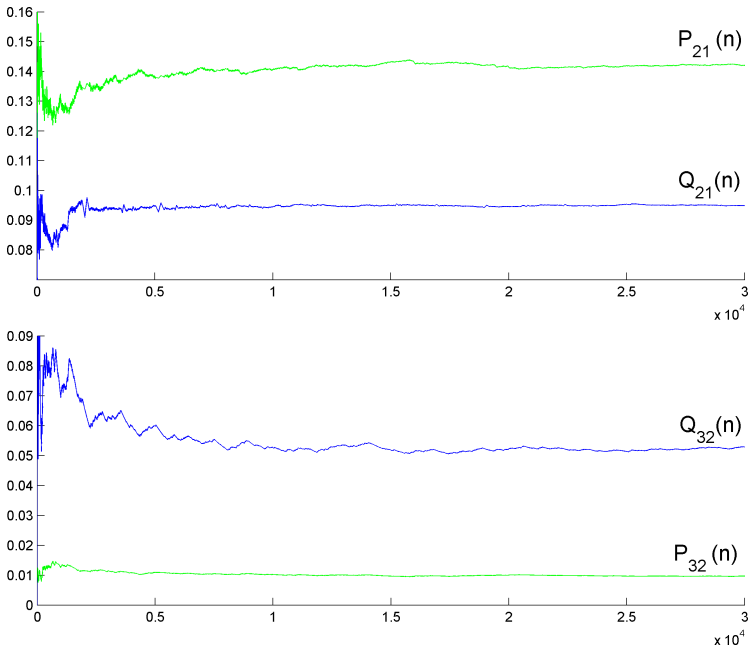

Figure 2.12: (Example 2.35) The first plot shows the evolution of $P_{21}(n)$ and $Q_{21}(n)$ along an orbit with itinerary of length $3 \cdot 10^4$. A certain 'stabilisation' and convergence to the values $P_{21} \approx 0.14$ and $Q_{21} \approx 0.09$ or small intervals containing these values can be observed. The second plot shows the evolution of $P_{32}(n)$ and $Q_{32}(n)$ along the same orbit. Here the observed limits (or limit intervals) are near the values $P_{32} \approx 0.01$ and $Q_{32} \approx 0.05$.

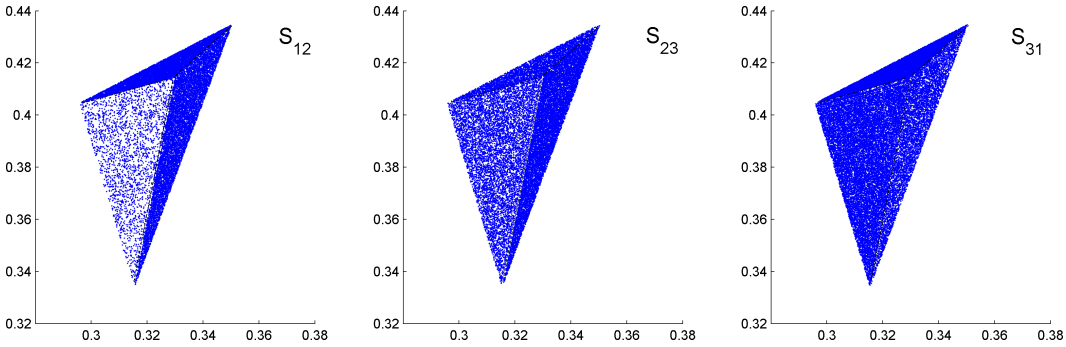

Figure 2.13: (Example 2.35) Typical orbit of induced flow intersected with the three surfaces S_{12}, S_{23} and S_{31}. The original orbit has an itinerary of length $2 \cdot 10^6$. The three triangular regions in each of the $S_{ii'}$ correspond to the different possible transitions between regions R_{ij} and $R_{i'j}$ for $j = 1, 2, 3$ and are indicated by dashed lines. In other words, the images show an orbit of the first return map to the surfaces of section $S_{ii'}$. The different visiting frequencies of the regions can clearly be seen from the different densities of orbit points. The orbit points inside each triangular region seem to be uniformly distributed.

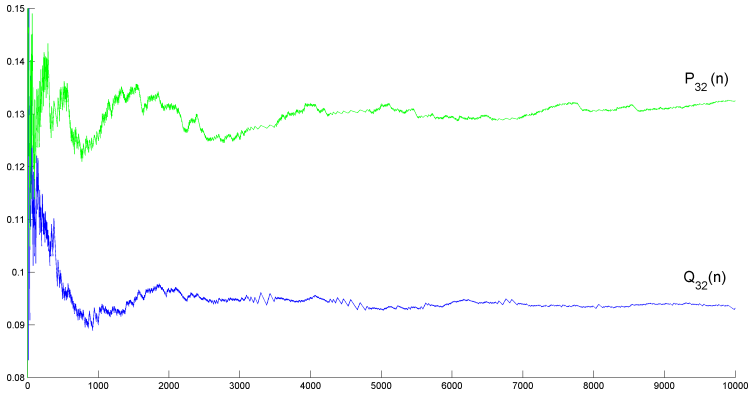

Figure 2.14: (Example 2.36) The evolution of $P_{32}(n)$ and $Q_{32}(n)$ along a typical orbit outside of the invariant torus.

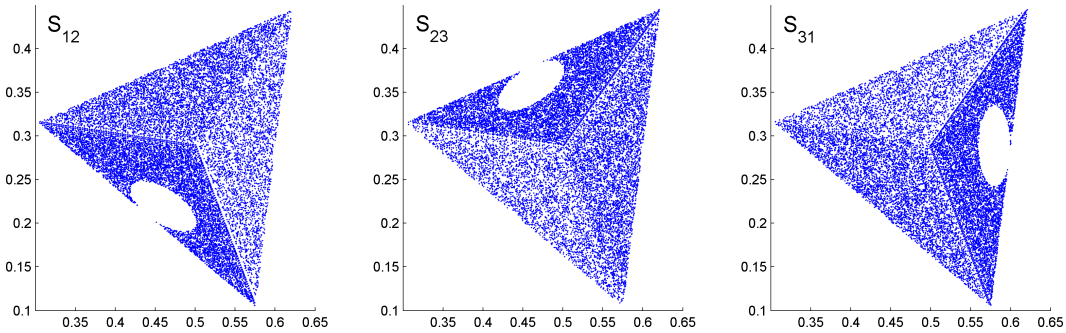

Figure 2.15: (Example 2.36) Typical orbit of induced flow intersected with the three indifference planes of a player. The orbit has an itinerary of length 10^6. The different visiting frequencies of the regions can be seen by the different densities of orbit points. The orbit points inside each triangular region seem to be uniformly distributed but leave out elliptical regions in some of the regions. These contain quasi-periodic orbits on invariant circles.

$$P \approx 10^{-2} \cdot \begin{pmatrix} 9 & 11 & 13 \\ 13 & 9 & 11 \\ 11 & 13 & 9 \end{pmatrix}, \qquad Q \approx 10^{-2} \cdot \begin{pmatrix} 11 & 13 & 9 \\ 9 & 11 & 13 \\ 13 & 9 & 11 \end{pmatrix}.$$

Seemingly, the values of $Q(n)$ are less erratic, and in most of our experiments they seem to converge faster than those of $P(n)$. As an example, the evolution of $P_{32}(n)$ and $Q_{32}(n)$ along a typical orbit can be seen in Figure 2.14.

As in the previous example, in Figure 2.15 we show the intersection of one such orbit with the surfaces $S_{ii'}$. Once again the orbit points have a certain seemingly uniform density inside each region, but here they leave out an elliptical region on each of the indifference planes. This invariant region consists of invariant circles, formed by quasi-periodic orbits of the system (the second type of observed orbits). The centre of the circles corre-

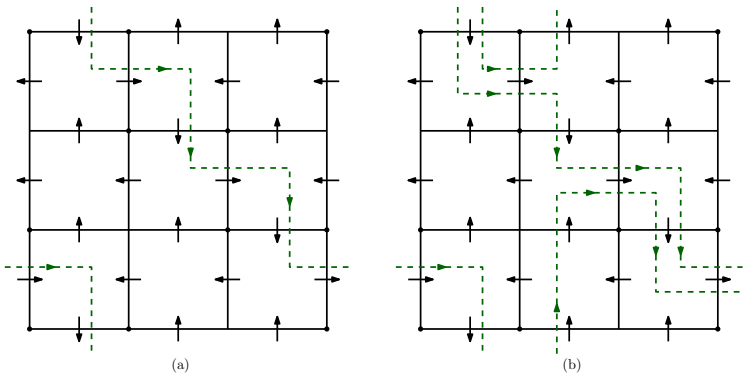

Figure 2.16: (a) (Example 2.36) The transition diagram of the bimatrix game. The periodic itinerary of the quasi-periodic orbits forming an invariant torus is indicated by a dashed line. (b) (Example 2.37) The transition diagram is the same as in the previous example. However, there is an invariant torus of quasi-periodic orbits with a more complicated periodic itinerary.

sponds to an actual periodic orbit, that is, an elliptic fixed point of the return map to one of the indifference planes. See [93] for an explicit analytic investigation of this (which is made possible by the high symmetry of this particular bimatrix game).

The invariant circles in the elliptical region correspond to invariant tori of the best response dynamics. Their itinerary is periodic with period 6:

$$(1,1) \to (1,2) \to (2,2) \to (2,3) \to (3,3) \to (3,1) \to (1,1) \to \cdots .$$

Figure 2.16(a) shows the transition diagram for this bimatrix game. The periodic itinerary is indicated by a dashed line as a loop in the transition diagram. The empirical frequencies along such quasi-periodic orbits converge to

$$P = Q = \begin{pmatrix} 1/6 & 1/6 & 0 \\ 0 & 1/6 & 1/6 \\ 1/6 & 0 & 1/6 \end{pmatrix}.$$

Figure 2.15 indicates that there are no other invariant tori for this system, that is, no open set of initial conditions outside of the visible elliptical regions, whose orbits are all quasi-periodic.

The following questions arise naturally from the above example:

- Does an invariant torus of quasi-periodic orbits always have a 'simple' periodic itinerary?

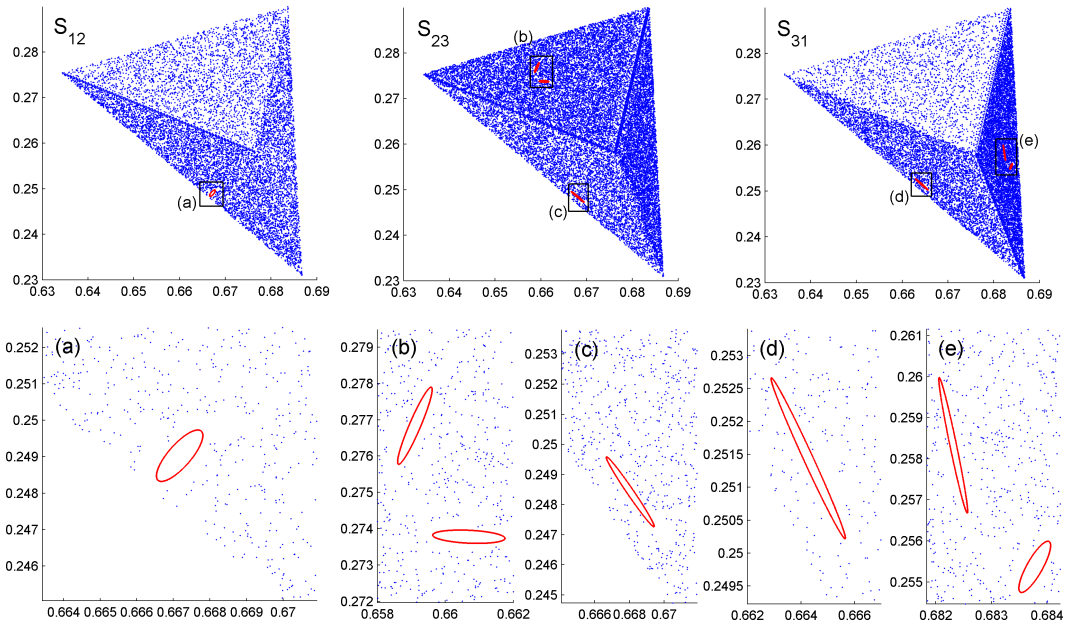

Figure 2.17: (Example 2.37) Two orbits (their intersections with the indifference planes $S_{ii'}$) are shown: one which stochastically fills most of the space and one which lies in an invariant torus. The latter intersects the first surface once and the other two three times each. The intersections of the (very thin) torus with the planes are marked with rectangles and shown magnified in the bottom row.

- Are the periods of such elliptic islands necessarily equal to 6?

As the next example shows, the situation can indeed be more complicated and less simple paths through the transition diagram are possible candidates for the periodic itinerary of quasi-periodic orbits in an invariant torus.

Example 2.37 (Quasi-periodic behaviour with itineraries of higher period)**.** Consider the bimatrix game (A, B) with

$$A = \begin{pmatrix} 84 & -37 & 10 \\ 24 & 33 & -14 \\ -26 & 9 & 20 \end{pmatrix}, \qquad B = -A.$$

Generally, the observations here coincide with Example 2.35. However, one can detect a (quite thin) invariant torus. Figure 2.17 shows a typical orbit stochastically filling most of the space. In the bottom row of the same figure, the regions marked by rectangles are enlarged to reveal a thin invariant torus. These quasi-periodic orbits intersect S_{12} once, S_{23} and S_{31} three times each. The orbits look essentially like the quasi-periodic orbits in the previous example, but with an extra loop added. The itineraries are periodic with period

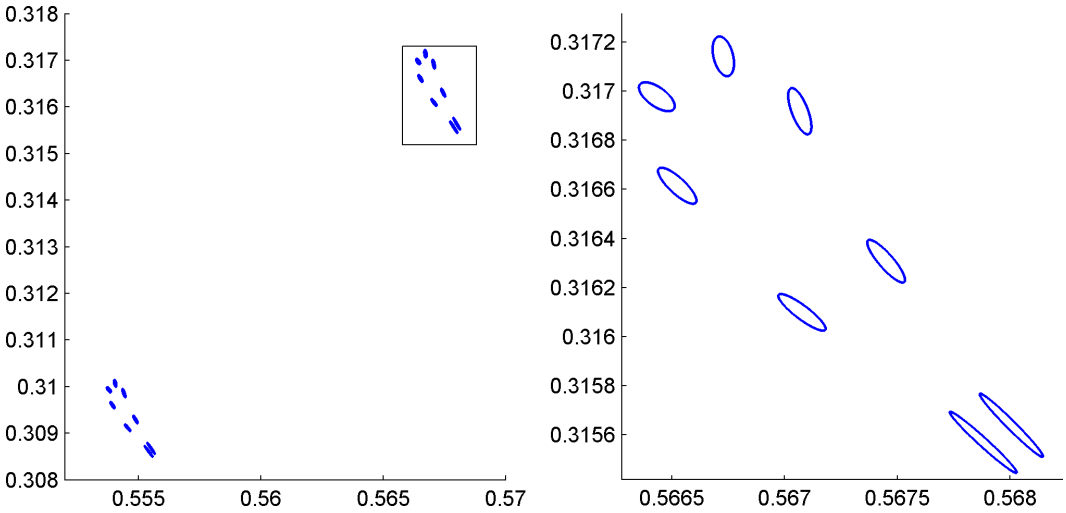

Figure 2.18: (Example 2.38) The invariant torus (corresponding to quasi-periodic orbits) intersects two of the indifference planes of the second player, while not intersecting the third one at all. The right image is a magnification of the region marked by a rectangle.

13, where each period is of the form

$$\begin{aligned}&(1,1) \to (1,2) \qquad\quad \to (2,2) \to (2,3) \to (3,3) \to (3,1) \to \\ &(1,1) \to (1,2) \to (3,2) \to (2,2) \to (2,3) \to (3,3) \to (3,1) \to (1,1).\end{aligned}$$

In Figure 2.16(b), this itinerary is shown as a loop in the transition diagram. The example demonstrates that combinatorially more complicated quasi-periodic orbits are possible for open sets of initial conditions in the best response dynamics of zero-sum games.

The next example shows an even more complex quasi-periodic structure and gives numerical evidence for more subtle and involved effects than those observed above.

Example 2.38 (Coexistence of different elliptic behaviour)**.** Let us now consider the bimatrix game (A, B) with

$$A = \begin{pmatrix} -92 & 18 & 52 \\ 62 & -37 & -33 \\ -10 & 9 & -18 \end{pmatrix}, \qquad B = -A.$$

As in all the previous examples, the largest part of the phase space of the best response dynamics seems to be filled with orbits which stochastically fill most of the space and along which the frequency distributions $P(n)$ and $Q(n)$ seem to converge to certain (orbit-independent) values. Again, an invariant torus can be found. It is more complicated than those observed in the other examples (see Figure 2.18). The orbits forming this invariant torus are quasi-periodic and have an itinerary of period 60. Its structure suggests a

generalisation of the type of itinerary observed in Example 2.37. It consists of a sequence of blocks of the following two forms:

$$a = (1,1) \to (3,1) \to (2,1) \to (2,2) \to (3,2) \to (3,3) \to (1,3) \to (1,1),$$
$$b = (1,1) \to (3,1) \to (2,1) \to (2,2) \to (3,2) \to (3,3) \to (1,3) \to (1,2) \to (1,1).$$

The two blocks are shown as paths in the transition diagram in Figure 2.21. Each period of the itinerary of orbits in the invariant torus then has the form

$$a \to a \to b \to a \to b \to b \to a \to b.$$

As in the previous example, the two blocks are the same except for one element (in the previous example the itinerary consisted of two blocks concatenated alternatingly).

In Figure 2.19 the intersection of an orbit outside of the invariant torus with one of the indifference planes is shown together with a quasi-periodic orbit. The orbit seems to have essentially the same property of filling the space outside the invariant torus, as in the previous examples. However, a closer look at a neighbourhood of the invariant circles (see the right part of Figure 2.19) reveals that the orbit not only misses out the invariant circles but also a certain 'heart-shaped' region surrounding these. The investigation of orbits with initial conditions in this set reveals a range of phenomena not observed in any of the previous examples.

Several different orbits with initial conditions in this region can be seen in Figure 2.20. The orbit points show complicated structures, revealing a large number of regions of 'stochastic' behaviour, as well as invariant regions of periodic orbits of high periods and corresponding quasi-periodic orbits (Figure 2.22 shows some examples of such quasi-periodic orbits of different higher periods). Some of these orbits spend very long times (itineraries of length 10^6 and more) in the heart-shaped region before diffusing into the much larger stochastic portion of the space. On the other hand, we observe orbits that stochastically fill (heart-shaped) annuli leaving out islands of quasi-periodic orbits. These annuli seem to be invariant for the dynamics (see Figure 2.23).

Altogether, the observations described above strongly indicate the occurrence of 'Arnol'd diffusion': the coexistence of a family of invariant annuli, which contain regions of stochastic (space-filling) motion and islands of further periodic orbits and invariant circles (quasi-periodic orbits).

We would like to end this section by proposing some open questions for further investigation:

(1) Does the Hamiltonian system induced by a 3×3 zero-sum bimatrix game always

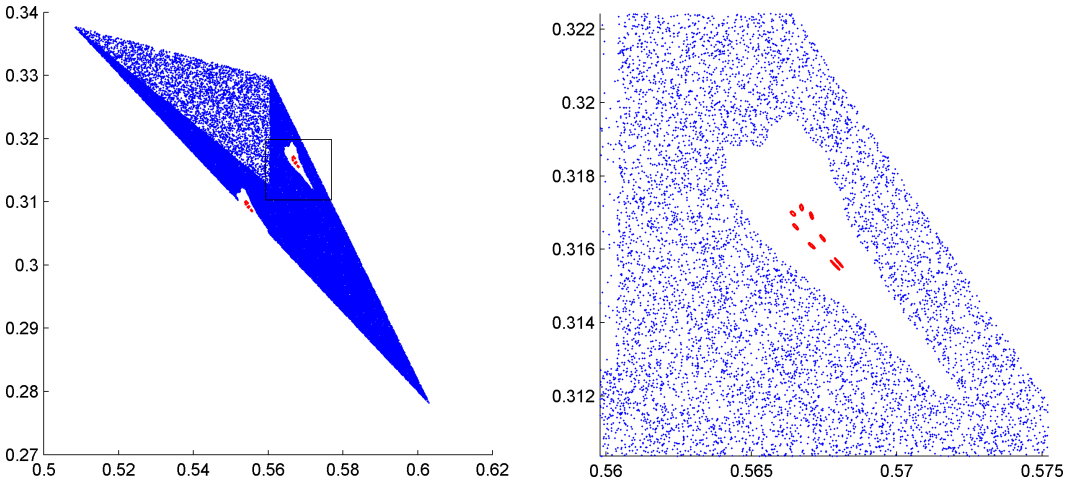

Figure 2.19: (Example 2.38) The left image shows a typical orbit that seems to fill almost all of the space. The right image shows a magnified view of the 'heart-shaped' region spared out by this orbit. Also both images show the invariant circles of a quasi-periodic orbit inside the 'heart-shaped' region.

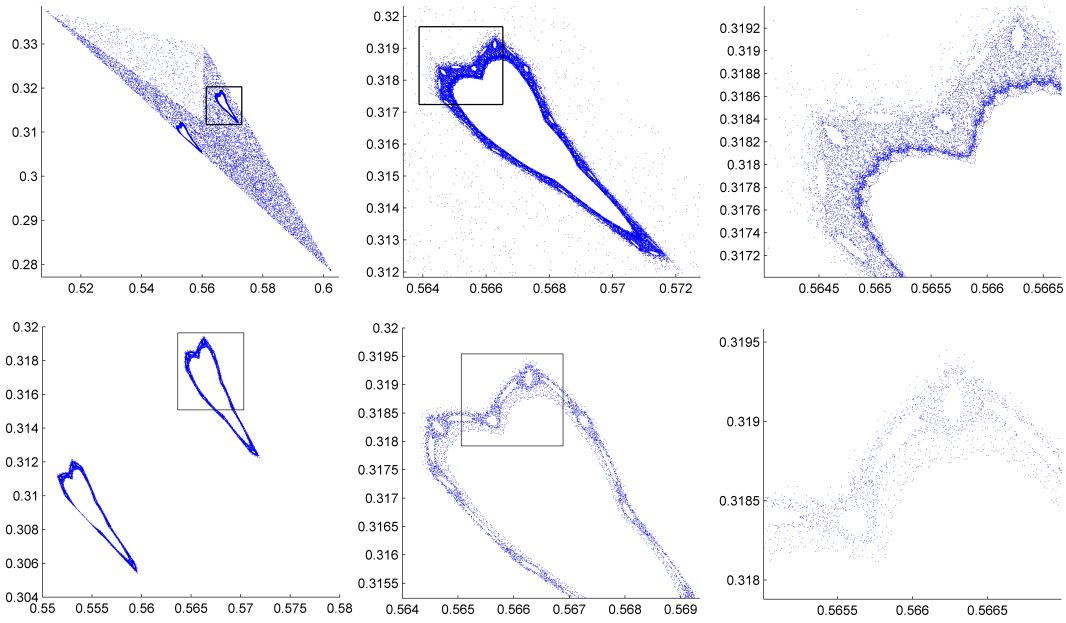

Figure 2.20: (Example 2.38) Two different orbits with initial conditions in the 'heart-shaped' region. Regions of stochastic motion as well as invariant islands of periodic and quasi-periodic orbits are clearly visible. The first orbit spends a long time in the 'heart-shaped' region before it diffuses into the larger 'ergodic' part of the space.

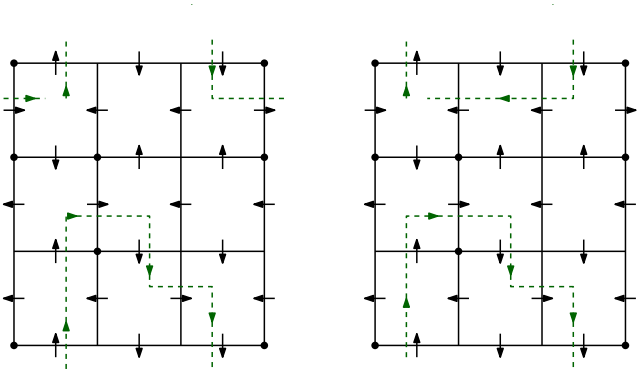

Figure 2.21: (Example 2.38) The two block types a and b of the itinerary for the quasi-periodic orbits in this system. Note that they differ by one step only.

have quasi-periodic orbits / invariant tori? Example 2.35 suggests that it is possible to have topological mixing and that the Lebesgue measure is ergodic. However, this might be due to the limited resolution of our numerical simulations and images.

(2) Are there orbits which are dense outside of the elliptic regions? Are almost all orbits outside of the elliptic regions dense?

(3) Does Example 2.38 (and similar ones) have infinitely many quasi-periodic orbits of different periods? The pictures of orbits (for instance, Figures 2.20 and 2.23) show many regions that could potentially contain the corresponding elliptic islands of different periods. All regions that we investigated for this property actually revealed quasi-periodic orbits.

(4) Given a specific bimatrix example, are there a finite number of blocks, so that the itinerary of any orbit on an elliptic island is periodic with each period being a (finite) concatenation of these blocks? The examples we looked at suggest the answer to be positive.

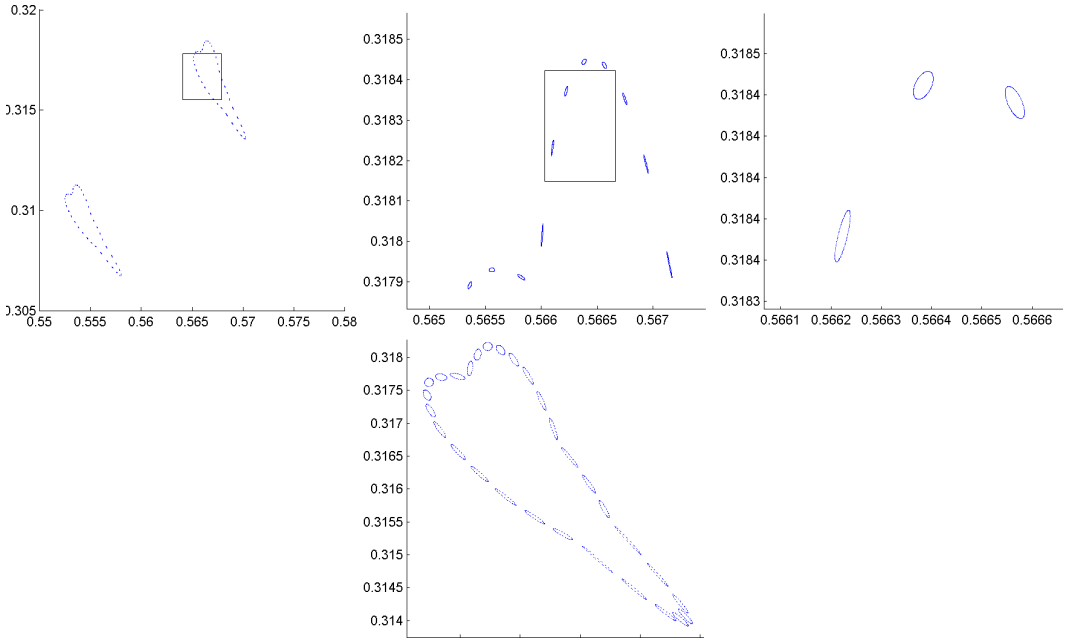

Figure 2.22: (Example 2.38) Two different quasi-periodic orbits (invariant tori for the best response dynamics) of different periods.

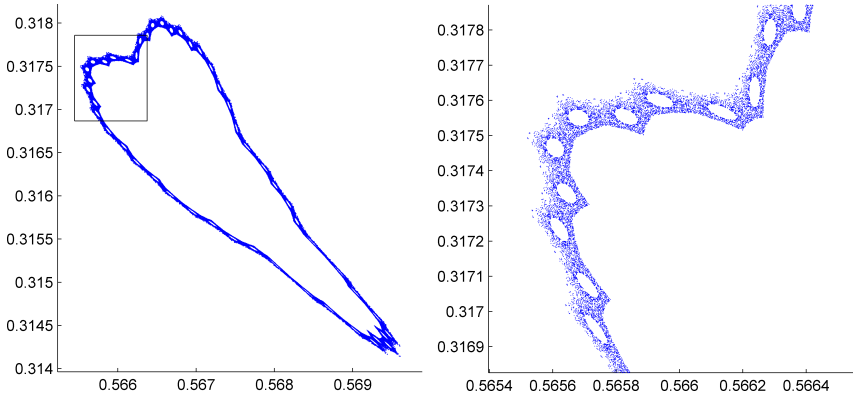

Figure 2.23: (Example 2.38) An orbit restricted to an invariant heart-shaped annulus.

Chapter 3

Payoff performance of fictitious play

In this chapter we investigate how well fictitious play performs in terms of average payoff, particularly compared to the payoff of two players who permanently play a Nash equilibrium strategy. The latter is often silently assumed to be optimal by classical rationality theory (at least when the Nash equilibrium is unique). In contrast, we show that in many games, two players involved in fictitious play may in the long run both earn a higher payoff than their Nash equilibrium payoff, either on average, or even at all times.

In Section 3.1 we analyse the limiting behaviour of fictitious play dynamics, and in Section 3.2 we use this to compare the payoff along the limit sets with the Nash equilibrium payoffs. Ultimately, this allows us to show that every bimatrix game is linearly equivalent to one in which fictitious play Pareto dominates Nash equilibrium and we discuss the conditions governing the payoff comparison of these two.

In Section 3.2.1 we apply these results to the family of Rock-Paper-Scissor-like games from Section 1.3 and show that in these games, fictitious play yields better average payoff to both players than Nash equilibrium.

Conversely, in Section 3.2.2 we investigate the possibility of games in which Nash equilibrium play dominates fictitious play. We also deduce conditions for this and numerically determine examples in which this is the case. The discussion shows that these examples are relatively 'rare'.

Finally, in Section 3.3 we discuss the implications of the results in this chapter for the notions of equilibrium (in the context of payoff performance of learning algorithms) and game equivalence.

The contents of this chapter are contained in the preprint [74], which at the time of submission of this thesis is submitted for publication.

3.1 Limit set for fictitious play

In this section we study the long term behaviour of (continuous-time) fictitious play. It has been known since Shapley's famous version of the Rock-Paper-Scissors game (1.5) that fictitious play does not necessarily converge to a Nash equilibrium even when the latter is unique, and can converge to a limit cycle instead (see Section 1.2.5). In fact, convergence to a unique Nash equilibrium in the interior of Σ seems to be rather the exception than the rule: It is a standing conjecture that such Nash equilibrium can only be stable for fictitious play dynamics, if the game is equivalent to a zero-sum game (Conjecture 1.29).

We will show that every fictitious play orbit converges to (a subset of) the set of so-called 'coarse correlated equilibria', sometimes also referred to as the 'Hannan set' (see Hannan [42], Young [99], Hart [44]). In fact, this result follows directly from the 'belief affirming' property of fictitious play, shown in Monderer et al. [65]. However, to the best of our knowledge, the conclusion that fictitious play has its limit set contained in the Hannan set, and therefore asymptotically the players have 'zero regret', has not been mentioned in the literature. We also provide a slightly different proof of this fact.

Let us fix some notation. We denote by $S^A = \{1, \dots, m\}$ and $S^B = \{1, \dots, n\}$ the sets of pure strategies of the two players in an $m \times n$ bimatrix game (A, B). We call $S = S^A \times S^B$ the *joint strategy space*, and we call a probability distribution over S a *joint probability distribution*. Note that $\Sigma = \Sigma^A \times \Sigma^B$ can be seen as a (proper) subset of the set of joint probability distributions.

The following definition can be found in Moulin and Vial [71].

Definition 3.1. A joint probability distribution $P = (p_{ij})$ over S is a *coarse correlated equilibrium (CCE)* for the bimatrix game (A, B) if

$$\sum_{i,j} a_{i'j} p_{ij} \leq \sum_{i,j} a_{ij} p_{ij}$$

and

$$\sum_{i,j} b_{ij'} p_{ij} \leq \sum_{i,j} b_{ij} p_{ij}$$

for all $(i', j') \in S$. The set of CCE is also called the *Hannan set*.

One way of viewing the concept of CCE is in terms of the notion of *regret*. Let us assume that two players are (repeatedly or continuously) playing a bimatrix game (A, B), and let $P(t) = (p_{ij}(t))$ be the empirical joint distribution of their past play through time t, that is, $p_{ij}(t)$ represents the frequency (or fraction of time) of the strategy profile (i, j) along their play through time t. For $x \in \mathbb{R}$, let $[x]_+$ denote the positive part of x: $[x]_+ = x$ if $x > 0$,

and $[x]_+ = 0$ otherwise. Then the expression

$$\left[\sum_{i,j} a_{i'j} p_{ij}(t) - \sum_{i,j} a_{ij} p_{ij}(t)\right]_+$$

can be interpreted as a measure for the regret of the first player from not having played action $i' \in S^A$ throughout the entire past history of play. It is (the positive part of) the difference between player A's actual past payoff[1] and the payoff she would have received if she always played i', given that player B would have played the same way as she did. Similarly, $[\sum_{i,j} b_{ij'} p_{ij}(t) - \sum_{i,j} b_{ij} p_{ij}(t)]_+$ is the regret of the second player from not having played $j' \in S^B$. This regret notion is sometimes called *unconditional* or *external regret* to distinguish it from the *internal* or *conditional regret*[2]. In this context the set of CCE can be interpreted as the set of joint probability distributions with non-positive regret.

It has been shown that there are learning algorithms with no regret, that is, such that asymptotically the regret of players playing according to such algorithm is non-positive for all their actions. Dynamically this means that if both players in a two-player game use a no-regret learning algorithm, the empirical joint probability distribution of actions taken by the players converges to (a subset of) the set of CCE (*not* necessarily to a certain point in this set).

The concept of no-regret learning (also known as *universal consistency*, see Fudenberg and Levine [31]) and the first such learning algorithms have been introduced by Blackwell [17] and Hannan [42]. More such algorithms have been found later on and moreover algorithms with asymptotically non-positive *conditional* regrets have been found (see, for example, Foster and Vohra [29], Hart and Mas-Colell [45, 46]; for good surveys see Young [99] or Hart [44]).

We now show that continuous-time fictitious play converges to the set of CCE.

Theorem 3.2. *Every trajectory of fictitious play dynamics* (FP) *in a bimatrix game* (A, B) *converges to the boundary of the set of CCE, that is, to the set of joint probability distributions* $P = (p_{ij})$ *over* $S^A \times S^B$ *such that for all* $(i', j') \in S^A \times S^B$

$$\sum_{i,j} a_{i'j} p_{ij} \le \sum_{i,j} a_{ij} p_{ij} \quad \text{and} \quad \sum_{i,j} b_{ij'} p_{ij} \le \sum_{i,j} b_{ij} p_{ij},$$

with equality for at least one $(i', j') \in S^A \times S^B$. *In other words, fictitious play dynamics asymptotically leads to non-positive (unconditional) regret for both players.*

[1] Note that $\sum_{i,j} a_{ij} p_{ij}(t)$ and $\sum_{i,j} b_{ij} p_{ij}(t)$ are the two players' respective average payoffs in their play through time t.

[2] Conditional regret is the regret from not having played an action i' whenever a certain action i has been played, that is, $[\sum_j a_{i'j} p_{ij} - \sum_j a_{ij} p_{ij}]_+$ for some fixed $i \in S^A$.

Remark 3.3. (1) Note that an orbit of fictitious play $(p(t), q(t))$, $t \geq t_0$, gives rise to a joint probability distribution $P(t) = (p_{ij}(t))$ via $p_{ij}(t) = p_i(t) \cdot q_j(t)$. So, when we say that fictitious play dynamics converges to a certain set of joint probability distributions, we mean that $P(t)$ obtained this way converges to this set.

(2) Monderer et al. [65] proved a stronger result, that continuous-time fictitious play is 'belief affirming' or 'Hannan-consistent'. This means that it leads to asymptotically non-positive unconditional regret for the player following it, *irrespective of her opponent's play* (even if the opponent is playing according to some very different algorithm). This however is not the case for discrete-time fictitious play (DFP), whereas the conclusion of Theorem 3.2 also holds in the discrete-time case by application of Hofbauer's limit set theorem for discrete- and continuous-time fictitious play (Theorem 1.41).

Proof of Theorem 3.2. Let us briefly recall notation from Chapter 1: The payoff function of player A is $u^A(p, q) = pAq$, her best response correspondence is $\mathcal{BR}_A(q) = \arg\max_{\bar{p} \in \Sigma^A} \bar{p}Aq$, and the maximal-payoff function is $\bar{A}(q) = \max_{\bar{p} \in \Sigma^A} \bar{p}Aq$; similarly, for player B.

We assume here that we have an orbit of (FP), $(p(t), q(t))$, $t \geq 0$. Recall from Remark 1.15(4) that $p(t) = \frac{1}{t}\int_0^t x(s)\,ds$ and $q(t) = \frac{1}{t}\int_0^t y(s)\,ds$, where $x\colon [0, \infty) \to \Sigma^A$ and $y\colon [0, \infty) \to \Sigma^B$ are measurable functions representing the players' strategies at any time $t \geq 0$, so that $x(t) \in \mathcal{BR}_A(q(t))$ and $y(t) \in \mathcal{BR}_B(p(t))$ for $t \geq 1$.

By the envelope theorem (see, for example, [91]), for $\bar{p} \in \mathcal{BR}_A(q)$ we have that

$$\frac{d\bar{A}(q)}{dq} = \frac{\partial u^A(p, q)}{\partial q}\bigg|_{p=\bar{p}} = \bar{p}A.$$

Therefore, since $x(t) \in \mathcal{BR}_A(q(t))$ for $t \geq 1$,

$$\frac{d}{dt}\left(t\bar{A}(q(t))\right) = \bar{A}(q(t)) + t\frac{d}{dt}\left(\bar{A}(q(t))\right) = \bar{A}(q(t)) + t \cdot x(t) \cdot A \cdot \frac{dq(t)}{dt}.$$

Using the definition of fictitious play (FP) and $\bar{A}(q(t)) = x(t) \cdot A \cdot q(t)$, it follows that

$$\frac{d}{dt}\left(t\bar{A}(q(t))\right) = \bar{A}(q(t)) + x(t) \cdot A \cdot (y(t) - q(t)) = x(t) \cdot A \cdot y(t)$$

for $t \geq 1$. We conclude that for $T > 1$,

$$\int_1^T x(t) \cdot A \cdot y(t)\,dt = T\bar{A}(q(T)) - \bar{A}(q(1)),$$

and therefore

$$\lim_{T\to\infty}\left(\frac{1}{T}\left(\int_0^T x(t) \cdot A \cdot y(t)\,dt\right) - \bar{A}(q(T))\right) = 0.$$

Note that

$$\frac{1}{T}\int_0^T x(t)\cdot A\cdot y(t)\,dt = \sum_{i,j} a_{ij}p_{ij}(T),$$

where $P(T) = (p_{ij}(T))$ is the empirical joint distribution of the two players' play through time T. On the other hand,

$$\bar{A}(q(T)) = \max_{i'}\sum_j a_{i'j}q_j(T) = \max_{i'}\sum_{i,j} a_{i'j}p_{ij}(T).$$

Hence,

$$\lim_{T\to\infty}\left(\sum_{i,j} a_{ij}p_{ij}(T) - \max_{i'}\sum_{i,j} a_{i'j}p_{ij}(T)\right) = 0.$$

By a similar calculation for B, we obtain

$$\lim_{T\to\infty}\left(\sum_{i,j} b_{ij}p_{ij}(T) - \max_{j'}\sum_{i,j} b_{ij'}p_{ij}(T)\right) = 0.$$

This shows that any orbit of fictitious play converges to the boundary of the set of CCE. □

Let us denote the average payoffs through time T along an orbit of fictitious play as

$$\hat{u}^A(T) = \frac{1}{T}\int_0^T x(t)\cdot A\cdot y(t)\,dt \quad\text{and}\quad \hat{u}^B(T) = \frac{1}{T}\int_0^T x(t)\cdot B\cdot y(t)\,dt.$$

As a corollary to the proof of the previous theorem we get the following useful result.

Theorem 3.4. *In any bimatrix game, along every orbit of fictitious play dynamics we have*

$$\lim_{T\to\infty}\left(\hat{u}^A(T) - \bar{A}(q(T))\right) = \lim_{T\to\infty}\left(\hat{u}^B(T) - \bar{B}(p(T))\right) = 0.$$

Remark 3.5. This formulation of the result shows why Monderer et al. [65] call this property 'belief affirming'. Since $\bar{A}(q(T))$ and $\bar{B}(p(T))$ can be interpreted as the players' expected payoffs given their respective opponent's play $q(T)$ and $p(T)$, the above theorem says that the difference between expected and actual average payoff of each player vanish, so that asymptotically their 'beliefs' are 'confirmed' when playing fictitious play.

3.2 Fictitious play vs. Nash equilibrium payoff

In this section we investigate the average payoff to players in a two-player game along the orbits of fictitious play dynamics and compare it to the Nash equilibrium payoff (in particular, in games with a unique, completely mixed Nash equilibrium).

Since fictitious play does not generally converge, the question of average payoff along fictitious play orbits in comparison to Nash equilibrium payoff is non-trivial and of interest both theoretically and for practical assessment of fictitious play as a sensible learning algorithm.

We show that in contrast to the usual assumption that players should primarily attempt to play Nash equilibria and that learning algorithms converging to Nash equilibria are desirable, the payoff along orbits of fictitious play can in some games be better on average or even Pareto dominate[3] the Nash equilibrium payoff. Moreover, we demonstrate that to every bimatrix game with unique, completely mixed Nash equilibrium, there is a dynamically equivalent game for which this superiority of fictitious play over Nash equilibrium holds.

Throughout the rest of this section we will assume that all the games under consideration have a unique, completely mixed Nash equilibrium point (E^A, E^B) (recall from Lemma 1.10 that in such a game, both players necessarily have the same number of strategies). A first simple situation in which fictitious play can improve upon such a Nash equilibrium is given by the following direct consequence of Theorem 3.4.

Proposition 3.6. *Let (A, B) be a bimatrix game with unique, completely mixed Nash equilibrium (E^A, E^B). If $\bar{A}(q) \geq \bar{A}(E^B)$ and $\bar{B}(p) \geq \bar{B}(E^A)$ for all $(p, q) \in \Sigma$, then asymptotically the average payoff along fictitious play orbits is greater than or equal to the Nash equilibrium payoff (for both players).*

Remark 3.7. The hypothesis of this proposition, $\bar{A}(q) \geq \bar{A}(E^B)$ and $\bar{B}(p) \geq \bar{B}(E^A)$ for all $(p, q) \in \Sigma$, means that

$$u^A(E^A, E^B) = \min_{q \in \Sigma^B} \max_{p \in \Sigma^A} pAq \quad \text{and} \quad u^B(E^A, E^B) = \min_{p \in \Sigma^A} \max_{q \in \Sigma^B} pBq,$$

that is, the Nash equilibrium payoff equals the minmax payoff of the players. For a non-zero-sum game this is a rather strong assumption, suggesting an unusually bad Nash equilibrium in terms of payoff. However, as we will show in the next result, at least from a dynamical point of view, the situation is not at all exceptional.

Theorem 3.8. *Let (A, B) be an $n \times n$ bimatrix game with unique, completely mixed Nash equilibrium (E^A, E^B). Then there exists a linearly equivalent game (A', B'), for which $\bar{A}'(q) > \bar{A}'(E^B)$ and $\bar{B}'(p) > \bar{B}'(E^A)$ for all $p \in \Sigma^A \setminus \{E^A\}$ and $q \in \Sigma^B \setminus \{E^B\}$.*

This result states that every bimatrix game with unique, completely mixed Nash equilibrium is linearly equivalent (see Definition 1.6) to one in which players are better

[3] By saying that fictitious play *Pareto dominates* Nash equilibrium play, we mean that it yields better payoffs to *both* players for *all times* $t \geq t_0$ for some sufficiently large $t_0 > 0$.

off playing according to fictitious play rather than playing the (unique) Nash equilibrium strategy. In the proof of the theorem we will make use of the following lemma.

Lemma 3.9. *Let (A, B) be an $n \times n$ bimatrix game with unique, completely mixed Nash equilibrium (E^A, E^B). Then for each of the first player's strategies k, $L_k^A := \left(\bigcap_{i\neq k} R_i^A\right) \setminus R_k^A$ is non-empty. More precisely, L_k^A is a ray from E^B in the direction v^k, such that any $(n-1)$ of the n vectors $v^1, \ldots, v^n$ form a basis for the space $\{v \in \mathbb{R}^n : \sum_i v_i = 0\}$. The analogous statement applies to $L_l^B := \left(\bigcap_{j\neq l} R_j^B\right) \setminus R_l^B$, $l = 1, \ldots, n$.*

Proof. Define the projection

$$\pi\colon \left\{x \in \mathbb{R}^n : \sum_i x_i = 1\right\} \to \mathbb{R}^{n-1}, \quad \pi(x_1, \ldots, x_n) = (x_1, \ldots, x_{n-1}),$$

and note that π is invertible with inverse

$$\pi^{-1}(y) = (y_1, \ldots, y_{n-1}, 1 - \sum_{k=1}^{n-1} y_k).$$

For $q \in \Sigma^B$ we have that $\sum_{k=1}^n q_k = 1$ and therefore

$$(Aq)_i - (Aq)_j = \sum_{k=1}^{n}(a_{ik} - a_{jk})q_k = \sum_{k=1}^{n-1}(a_{ik} - a_{jk} - a_{in} + a_{jn})q_k + (a_{in} - a_{jn}),$$

and we define the affine map $P\colon \mathbb{R}^{n-1} \to \mathbb{R}^{n-1}$ by

$$P_l(x) = \sum_{k=1}^{n-1}(a_{l,k} - a_{l+1,k} - a_{l,n} + a_{l+1,n})x_k + (a_{l,n} - a_{l+1,n}),$$

for $l = 1, \ldots, n-1$ (that is, $P_l(x) = A(\pi^{-1}(x))_l - A(\pi^{-1}(x))_{l+1}$).

Recall from Lemma 1.3 that for $(p, q) \in \Sigma$, $q = E^B$ if and only if $(Aq)_i = (Aq)_j$ for all i, j, and $p = E^A$ if and only if $(pB)_i = (pB)_j$ for all i, j. It follows that

$$P(x) = 0 \quad \text{if and only if} \quad x = \pi(E^B).$$

In particular, the affine map P is invertible and there is a unique vector $v^1 \in \{v \in \mathbb{R}^n : \sum_i v_i = 0\}$, such that $P(\pi(E^B + v^1)) = w^1 := (-1, 0, \ldots, 0)^\top$. Since E^B is in the interior of Σ^B, $x^1 = E^B + s\cdot v^1 \in \Sigma^B$ for sufficiently small $s > 0$, and we have that $P(\pi(x^1)) = (-s, 0, \ldots, 0)^\top$. By the definition of P, this means that

$$(Ax^1)_1 < (Ax^1)_2 = (Ax^1)_3 = \cdots = (Ax^1)_n.$$

Hence $x^1 \in L_1^A = \left(\bigcap_{k\neq 1} R_k^A\right) \setminus R_1^A$. Note also that every $x \in L_1^A$ is of the form $E^B + s \cdot v^1$ for some $s > 0$, that is, L_1^A is a ray from the point E^B.

For $1 < k < n$, let w^k be the vector in $\mathbb{R}^n$ with $(k-1)$th and kth entries equal to 1 and -1 respectively, and all other entries equal to 0. Then choose v^k such that $P(\pi(E^B + v^k)) = w^k$. Again for sufficiently small $s > 0$, we get $x^k = E^B + s \cdot v^k \in L_k^A$. Finally, for $k = n$, let $w^k = (0, \ldots, 0, 1)$ and proceed as above to get v^n and $x^n = E^B + v^n \in L_n^A$.

Writing the affine map P as $P(x) = Mx + b$ for some invertible matrix $M \in \mathbb{R}^{(n-1)\times(n-1)}$ and $b \in \mathbb{R}^{n-1}$, we get

$$w^k = P(\pi(E^B + v^k)) = P(\pi(E^B)) + M(v_1^k, \ldots, v_{n-1}^k)^\top = M(v_1^k, \ldots, v_{n-1}^k)^\top, \quad k = 1, \ldots, n.$$

Since any $n-1$ of the vectors

$$w^1 = \begin{pmatrix} -1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad w^2 = \begin{pmatrix} 1 \\ -1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \ldots, \quad w^{n-1} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ -1 \end{pmatrix}, \quad w^n = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

are linearly independent and M is invertible, it follows that any $n-1$ of the vectors $v^1, \ldots, v^n$ are linearly independent, as claimed.

The same argument applied to the matrix $B^\top$ shows the analogous result for L_l^B, $l = 1, \ldots, n$, which finishes the proof. □

Proof of Theorem 3.8. Let $A' \in \mathbb{R}^{n\times n}$, such that $a'_{ij} = a_{ij} + c_j$ for some $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$. Then for any $q \in \Sigma^B$,

$$\bar{A}'(q) = \max_i (A'q)_i = \max_i \left(Aq + \sum_{j=1}^n c_j q_j \cdot \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \right)_i = \bar{A}(q) + c \cdot q. \tag{3.1}$$

Observe that, restricted to R_k^A, level sets of $\bar{A}$ are precisely the $(n-2)$-dimensional hyperplane pieces in Σ^B orthogonal to $\underline{a}_k$, the kth row vector of A:

$$q - \tilde{q} \perp \underline{a}_k \Leftrightarrow \underline{a}_k \cdot q = \underline{a}_k \cdot \tilde{q} \Leftrightarrow \max_j (Aq)_j = \max_j (A\tilde{q})_j \qquad \text{for } q, \tilde{q} \in R_k^A.$$

So all level sets of $\bar{A}$ restricted to R_k^A are parallel hyperplane pieces. Figure 3.1 illustrates this situation for the case $n = 3$.

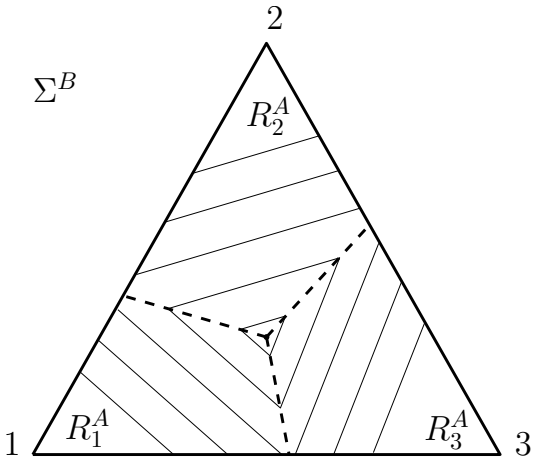

Figure 3.1: (Proof of Theorem 3.8) Level sets for $\bar{A}$ restricted to each region R_i^A are parallel line segments in Σ^B (in a 3×3 game).

By Lemma 3.9 we can choose n points $P_1, \ldots, P_n \in \Sigma^B$ such that

$$P_k \in L_k^A = \left(\bigcap_{i \neq k} R_i^A \right) \setminus R_k^A.$$

Each point P_k is in the relative interior of the line segment $L_k^A \subset \Sigma^B$. This line segment has endpoint E^B and is adjacent to all of the regions R_i^A, $i \neq k$. By the same lemma, $P_1 - E^B, \ldots, P_{n-1} - E^B$ form a basis for $\{v \in \mathbb{R}^n \colon \sum_k v_k = 0\}$. Therefore, the vectors $P_1, \ldots, P_n$ form a basis for $\mathbb{R}^n$.

It follows that one can choose $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$, such that

$$c \cdot P_1 + \bar{A}(P_1) = \cdots = c \cdot P_n + \bar{A}(P_n),$$

and hence by (3.1),

$$\bar{A}'(P_1) = \cdots = \bar{A}'(P_n).$$

Then level sets of $\bar{A}'$ are boundaries of $(n-1)$-dimensional simplices centred at E^B (each similar to the simplex with vertices $P_1, \ldots, P_n$).

Now we show that E^B is a minimum for $\bar{A}'$. The uniqueness of the completely mixed Nash equilibrium implies that A is not the zero matrix. Therefore, there exists a vector $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ with $\sum_k v_k = 0$, such that at least one of the entries of Av is positive. Let $r(t) = E^B + t \cdot v$, $t \geq 0$, be a ray from E^B in Σ^B. Then for $t_2 > t_1$ we get

$$\bar{A}'(r(t_2)) - \bar{A}'(r(t_1)) = \max_j (AE^B + t_2 Av)_j - \max_j (AE^B + t_1 Av)_j = (t_2 - t_1) \max_j (Av)_j > 0.$$

So, along some ray from E^B, $\bar{A}'$ is increasing. By the spherical structure of the level sets, this implies that $\bar{A}'$ is increasing along every ray from E^B. So $\bar{A}'(E^B) \leq \bar{A}'(q)$ for every $q \in \Sigma^B$ with equality only for $q = E^B$.

The same reasoning shows that one can choose $d_1, \ldots, d_n \in \mathbb{R}$ and $B' \in \mathbb{R}^{n \times n}$, $b'_{ij} = b_{ij} + d_i$, such that $\bar{B}'(E^A) \leq \bar{B}'(p)$ for every $p \in \Sigma^A$ with equality only for $p = E^A$. □

The previous results, Theorem 3.8 and Proposition 3.6, assert that every game possesses a dynamically equivalent version, in which fictitious play payoff Pareto dominates Nash equilibrium payoff. This shows that dynamical equivalence does not in general preserve the global payoff structure of a game, since there are clearly games for which Pareto dominance of fictitious play over Nash equilibrium does not hold a priori.

In the famous Shapley game (see (1.5) in Section 1.2.5) or variants of it (see (1.7) in Section 1.3) fictitious play typically converges to a limit cycle, known as a Shapley polygon, and usually the payoff along this polygon is greater than the Nash equilibrium payoff in some parts of the cycle, and less in others. On average, this can be still preferable for both players compared to playing Nash equilibrium, if they aim to maximise their time-average payoffs. In a similar setting, this has been previously observed by Gaunersdorfer and Hofbauer in [36]. We will investigate this for the family of games (1.7) in Section 3.2.1.

In fact, the proof of Theorem 3.8 shows that the unique, completely mixed Nash equilibrium (E^A, E^B) can never be an isolated payoff-maximum, since there are always directions from E^B in Σ^B and from E^A in Σ^A along which $\bar{A}$ and $\bar{B}$ are non-decreasing. Heuristically one would therefore expect that fictitious play typically improves upon Nash equilibrium in at least parts of any limit cycle. In Section 3.2.2 we will demonstrate that this need not always be the case: there are games in which fictitious play typically produces a lower average payoff than Nash equilibrium.

3.2.1 An example family of games revisited

Here we consider the one-parameter family of 3×3 bimatrix games from (1.7), which can be viewed as a generalisation of the Shapley game. In Section 1.3 we presented some of the most important results from [90, 93] on the remarkable properties of its fictitious play dynamics. Recall that the family is given by

$$A_\beta = \begin{pmatrix} 1 & 0 & \beta \\ \beta & 1 & 0 \\ 0 & \beta & 1 \end{pmatrix}, \qquad B_\beta = \begin{pmatrix} -\beta & 1 & 0 \\ 0 & -\beta & 1 \\ 1 & 0 & -\beta \end{pmatrix}, \qquad \beta \in (0, 1). \tag{3.2}$$

The unique Nash equilibrium of this game is $E^A = (E^B)^\top = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, which yields the respective payoffs

$$u^A(E^A, E^B) = \frac{1+\beta}{3} \quad \text{and} \quad u^B(E^A, E^B) = \frac{1-\beta}{3}$$

for players A and B. To check the hypothesis of Proposition 3.6, let $q = (q_1, q_2, q_3)^\top \in \Sigma^B$, then

$$\begin{aligned}
\bar{A}(q) &= \max\{q_1 + \beta q_3, q_2 + \beta q_1, q_3 + \beta q_2\} \\
&\geq \frac{1}{3}((q_1 + \beta q_3) + (q_2 + \beta q_1) + (q_3 + \beta q_2)) \\
&= \frac{1}{3}(q_1 + q_2 + q_3)(1 + \beta) \\
&= \frac{1+\beta}{3} \\
&= u^A(E^A, E^B) = \bar{A}(E^B).
\end{aligned}$$

Moreover, equality holds if and only if

$$q_1 + \beta q_3 = q_2 + \beta q_1 = q_3 + \beta q_2,$$

which is equivalent to $q_1 = q_2 = q_3$, that is, $q = E^B$. We conclude that $\bar{A}(q) > \bar{A}(E^B)$ for all $q \in \Sigma^B \setminus \{E^B\}$, and by a similar calculation, $\bar{B}(p) > \bar{B}(E^A)$ for all $p \in \Sigma^A \setminus \{E^A\}$. As a corollary to Proposition 3.6 we get the following result.

Theorem 3.10. *Consider the one-parameter family of bimatrix games* (A_β, B_β) *in* (3.2) *for* $\beta \in (0, 1)$. *Then any (non-stationary) orbit of fictitious play Pareto dominates constant Nash equilibrium play in the long run, that is, for large times t we have*

$$\hat{u}^A(t) > u^A(E^A, E^B) \quad \textit{and} \quad \hat{u}^B(t) > u^B(E^A, E^B).$$

In fact, one can say more: there is a $\beta \in (0, 1)$ such that fictitious play has an attracting closed orbit (the so-called 'anti-Shapley orbit', see Theorem 1.44) along which payoffs Pareto dominate the Nash equilibrium payoff *at all times*. In other words, both players are receiving a higher payoff than at Nash equilibrium at any time along this orbit. We omit the details of the proof: techniques developed by Krishna and Sjöström [57], building on earlier work by Rosenmüller [83], can be used to analyse fictitious play along this orbit, whose existence is guaranteed by Theorem 1.44. In particular, the times spent in each region R_{ij} along the orbit can be worked out explicitly, which can be directly applied to obtain average payoffs.

Remark 3.11. An important equilibrium notion in game theory is that of a *correlated equilibrium*, first introduced by Aumann [5, 6], which is defined as follows. A joint probability distribution $P = (p_{ij})$ over the joint strategy space $S = S^A \times S^B$ is a *correlated equilibrium* for the bimatrix game (A, B) if

$$\sum_k a_{i'k} p_{ik} \leq \sum_k a_{ik} p_{ik} \quad \text{and} \quad \sum_l b_{lj'} p_{lj} \leq \sum_k b_{lj} p_{lj}$$

for all $i, i' \in S^A$ and $j, j' \in S^B$. One interpretation of this notion is similar to that of the coarse correlated equilibrium (see paragraph after Definition 3.1), with the notion of '(unconditional) regret' replaced by the finer notion of 'conditional regret'. If we think of P as the empirical distribution of play up to a certain time for two players involved in repeatedly or continuously playing a given game, then P is a correlated equilibrium if neither player regrets not having played a strategy i' (or j') whenever she actually played i (or j). In other words, the average payoff to player A would not be higher, if she would have played i' at all times when she actually played i throughout the history of play (assuming her opponent's behaviour unchanged), and the same for player B.

One can check that the set of Nash equilibria is always contained in the set of correlated equilibria, which in turn is always contained in the set of coarse correlated equilibria. In the game (A_β, B_β) in Theorem 3.10, the Nash equilibrium (E^A, E^B) is also the unique correlated equilibrium, which can be checked by direct computation. Hence our result shows that in this case, fictitious play also improves upon correlated equilibria in the long run.

3.2.2 Fictitious play can be worse than Nash equilibrium play

We have seen that fictitious play can (and often does) improve upon Nash equilibrium in terms of payoff. Moreover, we have shown that for any bimatrix game with unique, completely mixed Nash equilibrium, linear equivalence can be used to obtain dynamically equivalent examples in which fictitious play payoff Pareto dominates Nash equilibrium payoff. In this section we investigate the converse possibility of fictitious play having lower payoff than Nash equilibrium. Again we restrict our attention to $n \times n$ games with unique, completely mixed Nash equilibrium.

Let us define the *sub-Nash payoff cones*, the set of those mixed strategies of player A, for which the best possible payoff to player B is not greater than Nash equilibrium payoff,

$$P_A^- = \{p \in \Sigma^A : \max_i (pB)_i \leq \max_i (E^A B)_i\},$$

and similarly

$$P_B^- = \{q \in \Sigma^B : \max_j (Aq)_j \leq \max_j (AE^B)_j\}.$$

By adding suitable constants to the player's payoff matrices we can assume without loss of generality that $u^A(E^A, E^B) = u^B(E^A, E^B) = 0$. Then one can see that

$$P_A^- = (B^\top)^{-1}(\mathbb{R}^n_-) \cap \Sigma^A \quad \text{and} \quad P_B^- = A^{-1}(\mathbb{R}^n_-) \cap \Sigma^B,$$

where $\mathbb{R}^n_-$ denotes the quadrant of $\mathbb{R}^n$ with all coordinates non-positive, and by $(B^\top)^{-1}$ and A^{-1} we mean the pre-images under the linear maps $B^\top, A \colon \mathbb{R}^n \to \mathbb{R}^n$. Therefore, P_A^- and P_B^- are convex cones in Σ^A and Σ^B with apexes E^A and E^B respectively.

Now a fictitious play orbit is Pareto worse than Nash equilibrium if and only if it (or its part for $t \geq t_0$ for some t_0) is contained in the interior of $P_A^- \times P_B^-$. This shows that a result like Theorem 3.8 with the roles of Nash equilibrium and fictitious play reversed cannot hold: if a game has a fictitious play orbit whose projections to Σ^A and Σ^B are not both contained in some convex cones with apexes E^A and E^B, then for any linearly equivalent game, along this orbit there are times at which one of the players enjoys higher payoff than Nash equilibrium payoff. In order to find fictitious play orbits along which payoffs are permanently worse than Nash equilibrium payoff, one therefore needs to find orbits contained in a halfspace (whose boundary plane contains the Nash equilibrium). The following lemma ensures that one can then obtain a linearly equivalent game with $P_A^- \times P_B^-$ containing this orbit.

Lemma 3.12. *Let (A, B) be any $n \times n$ bimatrix game with unique, completely mixed Nash equilibrium (E^A, E^B). Then for any convex cones $C_A \subset \Sigma^A$, $C_B \subset \Sigma^B$ with apexes E^A and E^B respectively, and opening angles in $[0, \pi)$, there exists a linearly equivalent game $(\tilde{A}, \tilde{B})$, such that $P_A^- = C_1$ and $P_B^- = C_2$.*

The proof of this lemma follows from the proof of Theorem 3.8. Note that in the proof of Theorem 3.8, to any given game we constructed a linearly equivalent game with $P_A^- = P_B^- = \emptyset$.

By Lemma 3.12, to find an example of a game with an orbit which is Pareto worse than Nash equilibrium, it suffices to find a game with an orbit whose projections to Σ^A and Σ^B are completely contained in convex cones with apexes E^A and E^B respectively. One can then construct a linearly equivalent game, for which this orbit is actually contained in the sub-Nash payoff cones. We will demonstrate one such example in the 3×3 case, which we obtained by numerically randomly generating 3×3 games and testing large numbers of initial conditions to detect orbits of the desired type.

Observe that by convexity of the preference regions R_i^A, a halfspace in Σ^B whose boundary line contains the (unique, completely mixed) Nash equilibrium contains at most two of the three rays L_i^A, $i = 1, 2, 3$. The same holds for a halfspace in Σ^A and the rays L_j^B, $j = 1, 2, 3$. Hence an orbit entirely contained in such halfspace never crosses at least one of these lines for each player.

Example 3.13. Let the bimatrix game (A, B) be given by

$$A = \begin{pmatrix} -1.353259 & -1.268538 & 2.572738 \\ 0.162237 & -1.800824 & 1.584291 \\ -0.499026 & -1.544578 & 1.992332 \end{pmatrix}, \; B = \begin{pmatrix} -1.839111 & -2.876997 & -3.366031 \\ -4.801713 & -3.854987 & -3.758662 \\ 6.740060 & 6.590451 & 6.898102 \end{pmatrix}.$$

This bimatrix game has a unique Nash equilibrium (E^A, E^B) with

$$E^A \approx (0.288, 0.370, 0.342), \quad E^B \approx (0.335, 0.327, 0.338)^\top.$$

The matrices A and B are chosen in such a way that the Nash equilibrium payoffs are both normalised to zero: $u^A(E^A, E^B) = u^B(E^A, E^B) = 0$. Numerical simulations suggest that fictitious play has a periodic orbit as a stable limit cycle, which attracts almost all initial conditions. This trajectory forms an octagon in the four-dimensional space $\Sigma = \Sigma^A \times \Sigma^B$, it is depicted in Figure 3.2. The orbit follows an 8-periodic itinerary of the form

$$(2,1) \to (2,2) \to (3,2) \to (3,3) \to (1,3) \to (1,2) \to (1,1) \to (3,1) \to (2,1).$$

Note that the first player's best response never changes from 1 to 2, or vice versa. Similarly, for player B the best response never directly changes between 1 and 3 without an intermediate step through 2. Moreover, it can be seen from Figure 3.2 that the projections of the periodic orbit to Σ^A and Σ^B lie in half planes whose boundaries contain the points E^A and E^B respectively. By Lemma 3.12 this allows us to choose the matrices A and B such that this orbit lies completely in $P_A^- \times P_B^-$, that is, such that the payoffs to both players are permanently worse than Nash equilibrium payoff. Figure 3.3 shows the (negative) payoffs to both players along several periods of the orbit and the higher (zero) Nash equilibrium payoff.

This example has been obtained through numerical experimentation. The difficulty in finding an example of a periodic orbit with the key property of lying in a convex cone with apex at the unique, completely mixed Nash equilibrium seems to suggest that such examples are relatively rare. For most games with unique, completely mixed Nash equilibrium, payoff along typical fictitious play orbits either Pareto dominates equilibrium payoff or at least improves upon it along parts of the orbit. We formulate the following two conjectures.

Conjecture 3.14. *Bimatrix games with unique, completely mixed Nash equilibrium, where Nash equilibrium Pareto dominates typical fictitious play orbits are 'rare'. To be precise, within the space of all $n \times n$ games with entries in* $[0, 1]$*, those where typical fictitious play orbits are Pareto dominated by Nash equilibrium (in terms of payoff) form a set with at most Lebesgue measure* 0.01.

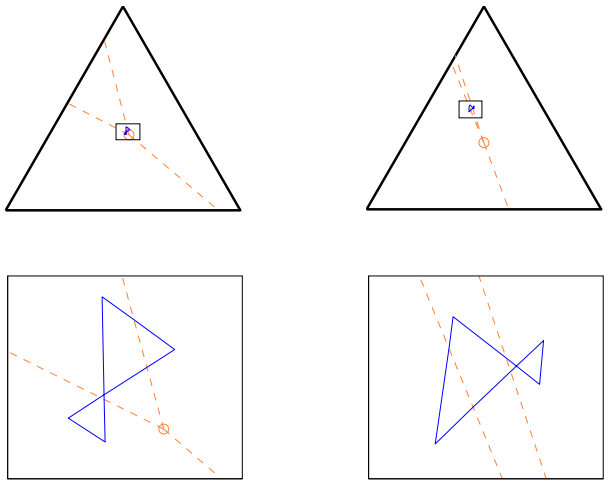

Figure 3.2: Periodic orbit whose projections to Σ^A (left) and Σ^B (right) are contained in convex cones with apexes E^A and E^B respectively. The dashed lines indicate the indifference lines of the players. Their intersections are the projections of the unique Nash equilibrium, E^A and E^B. For better visibility, the bottom row shows a zoomed version of the periodic orbit.

Conjecture 3.15. *For bimatrix games with unique, completely mixed Nash equilibrium and certain transition combinatorics (see Chapter 2), Nash equilibrium does not Pareto dominate typical fictitious play orbits. In particular, this is the case if* $\mathcal{BR}_A(e_j) \neq \mathcal{BR}_A(e_{j'})$ *for all* $j \neq j'$ *and* $\mathcal{BR}_B(e_i) \neq \mathcal{BR}_B(e_{i'})$ *for all* $i \neq i'$.

3.3 Concluding remarks on fictitious play performance

Conceptually, the overall observation is that playing Nash equilibrium might not be an advantage over playing according to some learning algorithm (such as fictitious play) in a wide range of games, in particular in many common examples of games occurring in the literature. Even in cases where playing fictitious play does not dominate Nash equilibrium at all times, it might still be preferable in terms of time-averaged payoff. In contrast, the previous section shows that there are examples in which Nash equilibrium indeed Pareto dominates fictitious play orbits, but the restrictive nature of the example suggests that this situation is quite rare.

Conversely, the discussion also shows that certain notions of game equivalence (for instance, linear, best response or better response equivalence, see Section 1.1.2), which are popular in the literature on learning dynamics, often do not preserve essential features of the payoff structure of games, even though they preserve Nash equilibria (and other notions

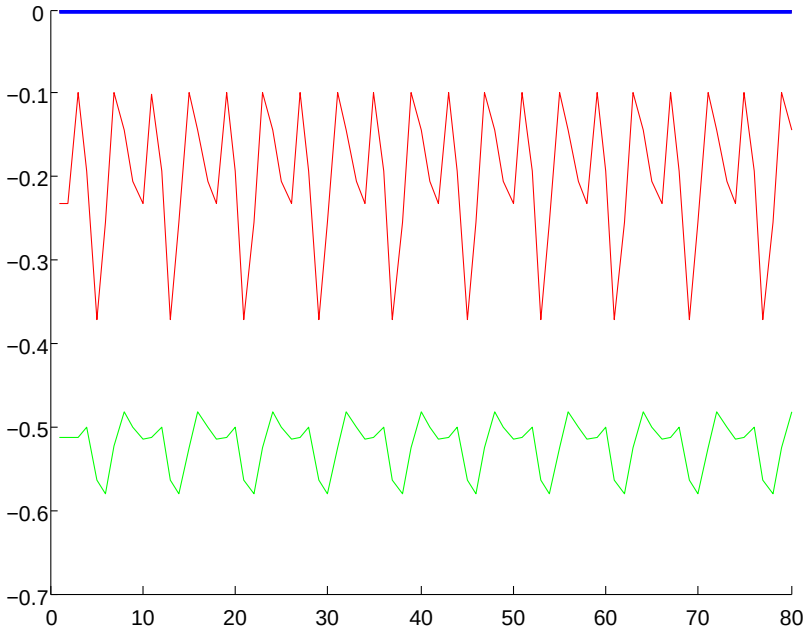

Figure 3.3: Payoff along 10 periods of the periodic orbit contained in $P_A^- \times P_B^-$. Player A's payoff oscillates around -0.5, player B's payoff around -0.25. Nash equilibrium payoff is zero to both players.

of equilibrium) and conditional preferences of the players. While some dynamics (in particular, best response or fictitious play dynamics) are invariant under all of these equivalence relations, the actual payoffs along their orbits and the average payoff comparison of different orbits can strongly depend on the chosen representative bimatrix, as becomes apparent from Theorem 3.8. This is to some extent analogous to the situation in the classical example of the 'prisoner's dilemma' given by the bimatrix

$$A = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 3 & 5 \\ 0 & 1 \end{pmatrix}.$$

Under linear equivalence, this corresponds to the bimatrix game

$$\tilde{A} = \begin{pmatrix} 0 & 0 \\ 2 & 1 \end{pmatrix}, \quad \tilde{B} = \begin{pmatrix} 0 & 2 \\ 0 & 1 \end{pmatrix},$$

which shares all essential features such as equilibria, best response structures, etc with the prisoner's dilemma. Both games are dynamically identical, with all fictitious play orbits converging along straight lines to the unique pure Nash equilibrium $(2, 2)$. However, the

second game does not constitute a prisoner's dilemma in the classical sense: whereas in the prisoner's dilemma the Nash equilibrium is Pareto dominated by the (dynamically irrelevant) strategy profile $(1, 1)$, in the second game this is not the case and no 'dilemma' occurs.

Theorem 3.8 can be interpreted in a similar vain: linear equivalence turns out to be sufficiently coarse, so that by changing the representative bimatrix inside an equivalence class, one can create certain regions in Σ in which payoff is arbitrarily high in comparison to the payoff at the unique Nash equilibrium. Since orbits of fictitious play remain unchanged, this can be done in such a way that a given periodic orbit lies completely or predominantly in these desired 'high payoff portions' of Σ. On the other hand, it can be seen from the proof that the conditions for this to happen are not at all exceptional. Consequently, it could be argued that in many games of interest the assumption that Nash equilibrium play is the most desirable outcome might not hold and a more dynamic view of 'optimal play' might be reasonable.

Chapter 4

Non-self maps of the plane

In this chapter we study planar homeomorphisms, a field in which many important investigations and results were started by Brouwer at the beginning of the twentieth century. In the context of this thesis, the importance of this class of maps lies in the fact that, as we have seen in Chapters 1 and 2, certain Poincaré maps of the fictitious play flow turn out to be planar homeomorphisms[1] and understanding of their fixed, periodic and recurrent points is crucial to the study of the global dynamics.

The main result of this chapter is a fixed point theorem deduced from classical results of Brouwer and, more recently, Franks. The crucial difference between our result and the classical versions is that in our result we assume a non-self map of a compact subset of the plane instead of a homeomorphism of the whole plane.

The contents of this chapter are published as [75].

4.1 Modern results on Brouwer homeomorphisms

Since Brouwer's proof of his plane translation theorem [20], many alternative proofs of the theorem and its key ingredient, the translation arc lemma, have been given (for several more recent ones, see Brown [23], Fathi [28], Franks [30]). The following is a concise formulation of its main statement.

Theorem 4.1 (Barge and Franks [7])**.** *Suppose $f\colon \mathbb{R}^2 \to \mathbb{R}^2$ is an orientation-preserving homeomorphism of the plane. If f has a periodic point then it has a fixed point.*

Slightly stronger versions assume a weaker form of recurrence, for example the existence of periodic disk chains (Barge and Franks [7]), to obtain the existence of fixed points.

[1] More precisely, the first return maps of the Hamiltonian flow induced by fictitious play in 3×3 zero-sum games with unique, completely mixed Nash equilibrium have this property, see Sections 1.2.6, 2.1.6 and 2.2.

The point of this chapter is to show that an analogue of Theorem 4.1 also holds for a homeomorphism which is merely defined on a compact subset of a surface. This situation could arise when considering a restriction of a self-map of a surface. Our main theorem is the following.

Theorem 4.2. *Let $X \subset \mathbb{R}^2$ be a compact, simply connected, locally connected subset of the real plane and let $f\colon X \to Y \subset \mathbb{R}^2$ be a homeomorphism isotopic to the identity on X. Let C be a connected component of $X \cap Y$. If f has a periodic orbit in C then it also has a fixed point in C.*

In fact, as a corollary we will obtain a slightly stronger result.

Corollary 4.3. *Let $X \subset \mathbb{R}^2$ and $f\colon X \to Y$ be as in Theorem 4.2 and let C be a connected component of $X \cap Y$. If f has no fixed point in C, then the orbit of every point $x \in C$ eventually leaves C, that is, there exists $n = n(x) \in \mathbb{N}$ such that $f^n(x) \notin C$.*

In particular, if $X \cap Y$ is connected and f has no fixed points, then the non-escaping set *of f is empty:*

$$\{x \in X : f^n(x) \in C\ \forall n \in \mathbb{N}\} = \emptyset.$$

To prove these results, in Section 4.2 we will first consider the case of an orientation-preserving homeomorphism $f\colon D \to E \subset \mathbb{R}^2$ of a Jordan domain D into the plane. In this special case, the statement will be first proved for $D \cap E$ connected, by suitably extending f to the real plane and applying Theorem 4.1. We will then proceed to show that the connectedness assumption can be removed if one formulates the result more precisely, taking into account the connected components of $D \cap E$ individually. Finally, in Section 4.3, we will deduce the general case of Theorem 4.2 by reducing the problem to the Jordan domain case. In Section 4.5 we will discuss the assumptions of our results and questions about possible extensions.

4.2 Non-self maps of Jordan domains

A set $D \subset \mathbb{R}^2$ is a *Jordan domain*, if it is a compact set with boundary ∂D a simple closed curve (*Jordan curve*). By Schoenflies' theorem, Jordan domains are precisely the planar regions homeomorphic to the (closed) disk. Here we assume that all Jordan curves are endowed with the counter-clockwise orientation. For a Jordan curve C and $x, y \in C$, we denote by $(x, y)_C$ (respectively $[x, y]_C$) the open (respectively closed) arc in C from x to y according to this orientation.

For $X \subset \mathbb{R}^2$, we say that $x \in X$ is a *fixed point* for the map $f\colon X \to \mathbb{R}^2$ if $f(x) = x$. We say that $x \in X$ is a *periodic point* for f if there exists $n \in \mathbb{N}$ such that $f^n(x) = x$. This of

course requires that the entire *orbit* of x, $\mathrm{orb}(x) = \{x = f^n(x), f(x), \ldots, f^{n-1}(x)\}$, is included in X.

In analogy to Theorem 4.1, we will prove the following result.

Theorem 4.4. *Let $D \subset \mathbb{R}^2$ be a Jordan domain and let $f\colon D \to E \subset \mathbb{R}^2$ be an orientation-preserving homeomorphism. Assume that $D \cap E$ is connected. If f has a periodic point then it also has a fixed point.*

Remark 4.5. All periodic or fixed points of $f\colon D \to E$ necessarily lie in $D \cap E$.

The strategy of our proof is to show that a homeomorphism $f\colon D \to E$ as in Theorem 4.4 and without fixed points can be extended to an orientation-preserving fixed point free homeomorphism $F\colon \mathbb{R}^2 \to \mathbb{R}^2$ (a *Brouwer homeomorphism*). The result then follows from classical results such as Theorem 4.1.

Our first step is to consider the structure of the set $D \cap E$. It is known that each connected component of the intersection of two Jordan domains is again a Jordan domain (see, for example, Kerékjártó [55, p.87]). We need the following more detailed statement (Bonino [19, Proposition 3.1], proved in Le Calvez and Yoccoz [60, Part 1]).

Proposition 4.6. *Let U, U' be two Jordan domains containing a point $o \in U \cap U'$ such that $U \not\subset U'$ and $U' \not\subset U$. Denote the connected component of $U \cap U'$ containing o by $U \wedge U'$.*

(1) There is a partition

$$\partial(U \wedge U') = (\partial(U \wedge U') \cap \partial U \cap \partial U') \cup \bigcup_{i \in I} \alpha_i \cup \bigcup_{j \in J} \beta_j, \text{ where}$$

- *I, J are non-empty, at most countable sets,*
- *for every $i \in I$, $\alpha_i = (a_i, b_i)_{\partial U}$ is a connected component of $\partial U \cap U'$,*
- *for every $j \in J$, $\beta_j = (c_j, d_j)_{\partial U'}$ is a connected component of $\partial U' \cap U$.*

(2) For $j \in J$, $U \wedge U'$ is contained in the Jordan domain bounded by $\beta_j \cup [d_j, c_j]_{\partial U}$.

(3) $\partial(U \wedge U')$ is homeomorphic to ∂U, and hence is itself a Jordan curve.

(4) Three points $a, b, c \in \partial(U \wedge U') \cap \partial U$ (respectively $\partial(U \wedge U') \cap \partial U'$) are met in this order on ∂U (respectively $\partial U'$) if and only if they are met in the same order on $\partial(U \wedge U')$.

We consider a Jordan domain $D \subset \mathbb{R}^2$ and a homeomorphism $f\colon D \to E = f(D) \subset \mathbb{R}^2$. If $D \cap E$ is connected, it follows that both $D \cap E$ and $D \cup E$ are Jordan domains (see Figure 4.1). With Proposition 4.6 in mind, we obtain the following lemma; we delay its proof until Section 4.4.

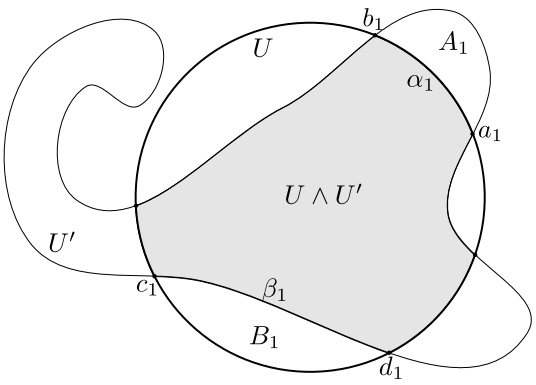

Figure 4.1: (Proposition 4.6) Partition of $\partial(U \wedge U')$ in the case when $U \cap U'$ is connected.

Lemma 4.7. *Let $D, E \subset \mathbb{R}^2$ be Jordan domains such that $D \not\subset E$, $E \not\subset D$ and $D \cap E$ is non-empty and connected. Then there exists a partition of $\mathbb{R}^2 \setminus \mathrm{int}(D \cap E)$ into arcs, each of which connects a point on $\partial(D \cap E)$ to ∞ and intersects each of ∂D and ∂E in precisely one point.*

Furthermore, there exist continuous functions $\lambda^D \colon \mathbb{R}^2 \setminus \mathrm{int}(D) \to \mathbb{R}_{\geq 0}$ and $\lambda^E \colon \mathbb{R}^2 \setminus \mathrm{int}(E) \to \mathbb{R}_{\geq 0}$, which are strictly monotonically increasing along the arcs of the partition and such that $\lambda^D|_{\partial D} \equiv 0$ and $\lambda^E|_{\partial E} \equiv 0$.

Remark 4.8. For a partition arc γ connecting some point $p \in \partial D$ (or ∂E) to ∞, the function λ^D (λ^E) can be viewed as providing a notion of arc length (assigning to each $x \in \gamma$ the length of the subarc $[p, x] \subset \gamma$). It is finite for every $x \in \gamma$, even if the usual (Euclidean) arc length $\ell([p, x])$ is not.

The partition of $\mathbb{R}^2 \setminus \mathrm{int}(D \cap E)$ obtained this way allows us to prove our key proposition, from which the main result of this section will follow as a corollary.

Proposition 4.9. *Let $D \subset \mathbb{R}^2$ be a Jordan domain and $f \colon D \to E \subset \mathbb{R}^2$ an orientation-preserving homeomorphism with no fixed points and such that $D \cap E$ is non-empty and connected. Then there exists a fixed point free orientation-preserving homeomorphism $F \colon \mathbb{R}^2 \to \mathbb{R}^2$ extending f (that is, $F|_D \equiv f$).*

Proof. Since f has no fixed points, by Brouwer's fixed point theorem we get that $E \not\subset D$ and $D \not\subset E$. Therefore, we can apply Lemma 4.7 to obtain a partition $\mathcal{P}$ of $\mathbb{R}^2 \setminus \mathrm{int}(D \cap E)$ into arcs, each connecting a point in $\partial(D \cap E)$ with ∞ and intersecting each of ∂D and ∂E in exactly one point.

For $x \in \mathbb{R}^2 \setminus \mathrm{int}(D \cap E)$, we denote by $L^x \in \mathcal{P}$ the partition element containing the point x, and let $\pi^D(x) = L^x \cap \partial D$ and $\pi^E(x) = L^x \cap \partial E$. From the construction of $\mathcal{P}$ in the proof of Lemma 4.7 it is clear that $\pi^D \colon \mathbb{R}^2 \setminus \mathrm{int}(D \cap E) \to \partial D$ and $\pi^E \colon \mathbb{R}^2 \setminus \mathrm{int}(D \cap E) \to \partial E$ are continuous.

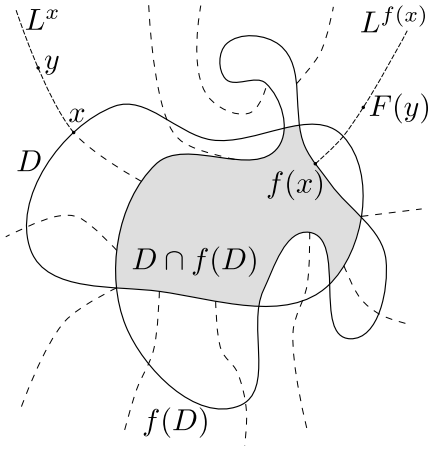

Figure 4.2: (Proof of Proposition 4.9) Extension $F\colon \mathbb{R}^2 \to \mathbb{R}^2$ of $f\colon D \to f(D)$ maps each arc L^x, $x \in \partial D$, into the arc $L^{f(x)}$.

We now construct $F\colon \mathbb{R}^2 \to \mathbb{R}^2$ as an extension of f in such a way, that each arc $L^x \setminus D$, $x \in \partial D$, is mapped into the arc $L^{f(x)}$. For $x \in D$ we define $F(x) = f(x)$, and for $x \notin D$ we define $F(x)$ to be the unique point $y \in L^{f(\pi^D(x))} \setminus E$ such that $\lambda^E(y) = \lambda^D(x)$ (see Figure 4.2). Then, by the construction of $\mathcal{P}$ and continuity of π^D and π^E, F is a continuous extension of f.

The map F is an orientation-preserving homeomorphism of $\mathbb{R}^2$; its inverse can be obtained the same way by swapping the roles of D, f and E, f^{-1}. Moreover, F has no fixed points in D since $F|_D \equiv f$. On the other hand, suppose $p \in \mathbb{R}^2 \setminus D$ with $F(p) = p$. Then $F(L^p \setminus D) \subset L^p$. Since $\lambda^D(p) = \lambda^E(F(p)) = \lambda^E(p)$, we get $\pi^D(p) = \pi^E(p)$, and therefore $f(\pi^D(p)) = \pi^D(p)$, a contradiction. Hence F is fixed point free. □

Combining Proposition 4.9 and Theorem 4.1, we now obtain Theorem 4.4.

Next, we formulate a somewhat more general form of Theorem 4.4, by removing the restriction of $D \cap f(D)$ to be connected, and instead making an assertion for each individual connected component of this set. This version of the result turns out to be much more useful in its application to more general compact sets in Section 4.3.

Theorem 4.10. *Let $D \subset \mathbb{R}^2$ be a Jordan domain and $f\colon D \to E \subset \mathbb{R}^2$ an orientation-preserving homeomorphism. Let C be a connected component of $D \cap E$. If f has a periodic orbit in C then it also has a fixed point in C.*

Proof. We show that this more general setting can be reduced to the one in Theorem 4.4. Let $\{C_i\}$ be the collection of connected components of $\overline{E \setminus C}$ and let $\gamma_i = C \cap C_i$, which is a closed subarc of ∂D (a connected component of $C \cap \partial D$).

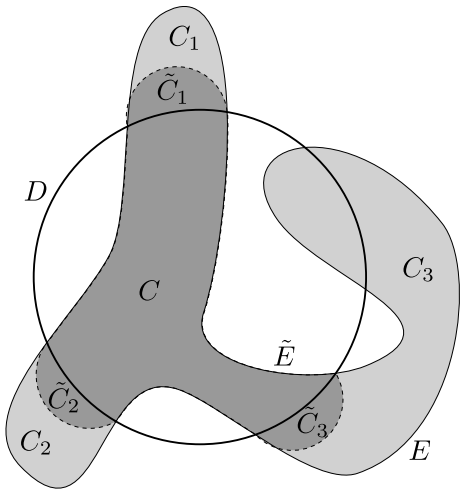

Figure 4.3: (Proof of Theorem 4.10) The Jordan domains D (white), $E = \left(\bigcup_i C_i\right) \cup C$ (light grey) and $\tilde{E} = \left(\bigcup_i \tilde{C}_i\right) \cup C$ (dark grey). The $\tilde{C}_i$ are pairwise disjoint, the homeomorphism $g\colon E \to \tilde{E}$ maps each C_i homeomorphically to $\tilde{C}_i$ and fixes C pointwise. Moreover, $\tilde{E} \cap D = C$ is connected.

Each C_i is a Jordan domain bounded by the union of γ_i and a closed subarc of ∂E. Let $\{\tilde{C}_i\}$ be another collection of disks $\tilde{C}_i \subset \overline{\mathbb{R}^2 \setminus D}$ such that $\tilde{C}_i \cap D = \gamma_i$ and the $\tilde{C}_i$ are pairwise disjoint. (Note that by Schoenflies' theorem we can assume without loss of generality that D is the closed standard unit disk $\overline{\mathbb{D}} = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. Since the γ_i are pairwise disjoint, letting $\tilde{C}_i$ be the round halfdisk over $\partial\mathbb{D}$ centred at the midpoint of γ_i gives a collection of pairwise disjoint disks as required, see Figure 4.3.)

By Schoenflies' theorem, for each i there exists a homeomorphism $g_i\colon C_i \to \tilde{C}_i$ and we can assume that g_i is the identity on γ_i. Noting that $E = \left(\bigcup_i C_i\right) \cup C$, let $\tilde{E} = \left(\bigcup_i \tilde{C}_i\right) \cup C$ and define a homeomorphism $g\colon E \to \tilde{E}$ by gluing together the identity on C and the homeomorphisms g_i:

$$g(y) = \begin{cases} y & \text{if } y \in C, \\ g_i(y) & \text{if } y \in C_i. \end{cases}$$

Finally, define the homeomorphism $\tilde{f} = g \circ f\colon D \to \tilde{E}$. By construction, $\tilde{f}$ and f coincide on $f^{-1}(C)$, and if f has a periodic orbit in C, so does $\tilde{f}$. Since $\tilde{f}(D) \cap D = \tilde{E} \cap D = C$ is connected, we can apply Theorem 4.4 to get that $\tilde{f}$ has a fixed point in C. Hence f has a fixed point in C. □

Corollary 4.11. *Let $f\colon D \to E$ be as in Theorem 4.10 and C a connected component of $D \cap E$. If f has no fixed point in C, then the* non-escaping set *of f in C is empty:*

$$\{x \in C : f^n(x) \in C \ \forall n \in \mathbb{N}\} = \emptyset.$$

Proof. By repeating verbatim the argument in the proof of Theorem 4.10, we can construct a homeomorphism $\tilde{f}\colon D \to \tilde{E}$, which coincides with f on C and such that $\tilde{E} \cap D = C$.

Observe that $\tilde{f}$ is fixed point free, so we can apply Proposition 4.9 to get an extension of $\tilde{f}$ to a fixed point free orientation-preserving homeomorphism $F\colon \mathbb{R}^2 \to \mathbb{R}^2$. By [30, Corollary 1.3], F is a so-called *free* homeomorphism, that is, for any topological disk $U \subset \mathbb{R}^2$ we have

$$F(U) \cap U = \emptyset \quad \Rightarrow \quad F^i(U) \cap F^j(U) = \emptyset \text{ whenever } i \neq j. \tag{4.1}$$

Suppose f has a point $x \in C$ with $f^n(x) \in C$ for all $n \in \mathbb{N}$. Then $f^n(x) = F^n(x)$ for all $n \in \mathbb{N}$ and the forward orbit of x has an accumulation point, $F^{n_k}(x) \to x_0 \in C$ as $k \to \infty$. A sufficiently small neighbourhood U of x_0 then satisfies $F(U) \cap U = \emptyset$ but clearly $F^{n_{k+1}-n_k}(U) \cap U \neq \emptyset$ for sufficiently large k, contradicting (4.1). Hence such point never escaping C does not exist. □

4.3 Non-self maps for compact simply connected planar sets

In this section we will generalise the results of the previous section to non-self maps of compact, simply connected, locally connected sets in the real plane (also known as *non-separating Peano continua*). Note that in our terminology simply connected always implies connected.

We denote the Riemann sphere by $\hat{\mathbb{C}}$. If $X \subset \hat{\mathbb{C}}$ is compact, connected and non-separating, then its complement $U = \hat{\mathbb{C}} \setminus X$ is simply connected (an open set $U \subset \hat{\mathbb{C}}$ is simply connected if and only if both U and $\hat{\mathbb{C}} \setminus U$ are connected).

The classical Riemann mapping theorem states that if $U \subsetneq \mathbb{C}$ is non-empty, simply connected and open, then there exists a biholomorphic map from U onto the open unit disk $\mathbb{D} = \{z \in \mathbb{C} : \|z\| < 1\}$, known as the *Riemann map*. We will make use of the following stronger result by Carathéodory (see [63, Theorem 17.14]).

Theorem 4.12 (Carathéodory's theorem)**.** *If $U \subsetneq \hat{\mathbb{C}}$ is non-empty, simply connected and open, $\hat{\mathbb{C}} \setminus U$ has at least two points, and additionally ∂U (or $\hat{\mathbb{C}} \setminus U$) is locally connected, then the inverse of the Riemann map, $\phi\colon \mathbb{D} \to U$, extends continuously to a map from the closed unit disk $\overline{\mathbb{D}}$ onto $\overline{U}$.*

As in the case of Jordan domains, we will need a requirement on the homeomorphisms to be orientation-preserving. To make sense of this notion for more general subsets of the plane, we cite the following result, simplified by our assumption of a non-separating and locally connected set (see also Oversteegen and Tymchatyn [77]).

Theorem 4.13 (Oversteegen and Valkenburg [78])**.** *Let $X \subset \hat{\mathbb{C}}$ be compact, simply connected and locally connected, and $f\colon X \to Y \subset \hat{\mathbb{C}}$ a homeomorphism. Then the following are equivalent:*

(1) f is isotopic to the identity on X;

(2) there exists an isotopy $F\colon \hat{\mathbb{C}} \times [0,1] \to \hat{\mathbb{C}}$ such that $F^0 = \mathrm{id}_{\hat{\mathbb{C}}}$ and $F^1|_X = f$;

(3) f extends to an orientation-preserving homeomorphism of $\hat{\mathbb{C}}$;

(4) if $U = \hat{\mathbb{C}} \setminus X$ and $V = \hat{\mathbb{C}} \setminus Y$, then f induces a homeomorphism $\hat{f}$ from the prime end circle of U to the prime end circle of V which preserves the circular order.

To explain assertion (4) in the above theorem, let $\phi\colon \overline{\mathbb{D}} \to \overline{U}$ and $\psi\colon \overline{\mathbb{D}} \to \overline{V}$ be the extended inverse Riemann maps given by Carathéodory's theorem, and denote $S^1 = \partial\mathbb{D}$. Then the homeomorphism $\hat{f}\colon S^1 \to S^1$ is said to be *induced by* f, if

$$\psi|_{S^1} \circ \hat{f} = f|_{\partial U} \circ \phi|_{S^1}.$$

More details can be found in [78], for an introduction to Carathéodory's theory of prime ends the reader is referred to Milnor's book [63, Chapter 17].

Furthermore, we will make use of the following notation: For $\varepsilon > 0$, we denote the closed ε-neighbourhood of a set $X \subset \mathbb{R}^2$ by

$$X_\varepsilon = \{z \in \mathbb{R}^2 : \inf_{x \in X} \|x - z\| \leq \varepsilon\}.$$

In the proof of our main theorem, we will consider ε-neighbourhoods of disconnected compact sets. A key fact we need is that any two given connected components of such a set are separated by the ε-neighbourhood of the set, if $\varepsilon > 0$ is chosen sufficiently small. This is the following technical lemma, whose proof can be found in Section 4.4.

Lemma 4.14. *Let X be a compact subset of $\mathbb{R}^n, n \in \mathbb{N}$. Let $\mathcal{C} = \{C_i\}_{i \in I}$ be the collection of connected components of X and assume $|\mathcal{C}| \geq 2$. Let $C, C' \in \mathcal{C}$ be two distinct connected components. Then for $\varepsilon > 0$ sufficiently small, C and C' lie in two distinct connected components of X_ε.*

We are now ready to prove our main results.

Proof of Theorem 4.2. Let C be a connected component of $X \cap Y$ which contains a periodic orbit of f. We will prove that f has a fixed point in C. The strategy of the proof is to construct Jordan-domain neighbourhoods of X and Y and to extend f to a homeomorphism between these neighbourhoods, so that Theorem 4.10 can be applied to the extension. The

existence of a fixed point for the extension will then be shown to imply the existence of a fixed point for f in C.

For fixed $\varepsilon > 0$, we first construct a closed neighbourhood N^ε of X with boundary $\Gamma = \partial N^\varepsilon$ a Jordan curve such that $\overline{N^\varepsilon \setminus X}$ has a foliation $\{\gamma_z\}_{z\in\Gamma}$ with the following properties:

- each γ_z is an arc connecting the point $z \in \Gamma$ with a point $x(z) \in \partial X$;
- $\gamma_z \cap \partial X = \{x(z)\}$ and $\gamma_z \cap \Gamma = \{z\}$ for each $z \in \Gamma$;
- $\bigcup_{z\in\Gamma} \gamma_z = \overline{N^\varepsilon \setminus X}$;
- if $z \neq z'$, then either $\gamma_z \cap \gamma_{z'} = \emptyset$, or $x(z) = x(z')$ and $\gamma_z \cap \gamma_{z'} = \{x(z)\} \subset \partial X$;
- each arc γ_z lies within an ε-ball centred at its basepoint $x(z) \in \partial X$: $\gamma_z \subset B_\varepsilon(x(z))$.

Observe that by the last property, N^ε is included in the ε-neighbourhood X_ε of X.

We identify $\mathbb{R}^2 = \mathbb{C}$ and $\mathbb{C} \subset \hat{\mathbb{C}}$ in the usual way and denote $U = \hat{\mathbb{C}} \setminus X$. Then U satisfies the hypotheses of Theorem 4.12, so there exists a continuous map $\phi\colon \overline{\mathbb{D}} \to \overline{U}$ (the extended inverse Riemann map), whose restriction to the open unit disk $\mathbb{D}$ is a conformal homeomorphism from $\mathbb{D}$ to U and $\phi(\partial\mathbb{D}) = \partial U$. The map ϕ can be chosen such that $\phi(0) = \infty$. Since $\overline{\mathbb{D}}$ is compact, ϕ is uniformly continuous, so we can find $\delta = \delta(\varepsilon) > 0$ such that for all $x, y \in \overline{\mathbb{D}}$ with $\|x - y\| < \delta$, $\|\phi(x) - \phi(y)\| < \varepsilon$.

Assume without loss of generality that $\delta < 1$ and set $A_\delta = \{x \in \mathbb{C} : 1 - \delta \leq \|x\| \leq 1\}$ and $N^\varepsilon = \phi(A_\delta) \cup X$. Then N^ε is a closed neighbourhood of X, whose boundary $\Gamma = \partial N^\varepsilon = \phi(\{x \in \mathbb{D} : \|x\| = 1 - \delta\})$ is a Jordan curve. We can then construct the foliation $\{\gamma_z\}_{z\in\Gamma}$ of $\overline{N^\varepsilon \setminus X}$ by taking the images of radial lines in A_δ (see Figure 4.4):

$$\gamma_z = \phi(\{r \cdot v : 1 - \delta \leq r \leq 1\}), \qquad \text{where } v = \frac{\phi^{-1}z}{\|\phi^{-1}z\|}.$$

All the required properties of N^ε and its foliation $\{\gamma_z\}$ then follow because ϕ maps $A_\delta \setminus \partial\mathbb{D}$ bijectively and uniformly continuously onto $N^\varepsilon \setminus X$; in particular, the last required property follows by the choice of $\delta = \delta(\varepsilon)$.

Next, since f is a homeomorphism, $Y = f(X)$ is simply connected, compact and also locally connected (see [63, Theorem 17.15]). Hence, with $V = \hat{\mathbb{C}} \setminus Y$ and $\phi\colon \overline{\mathbb{D}} \to \overline{V}$ the corresponding extended inverse Riemann mapping, we can repeat the same construction as before to obtain a closed neighbourhood M^ε of Y, such that $\Theta = \partial M^\varepsilon$ is a Jordan curve and $\{\theta_z\}$ is a foliation of $\overline{M^\varepsilon \setminus X}$ with the same properties as $\{\gamma_z\}$.

By choosing $\delta > 0$ small enough such that both constructions work with this given value, we get that $N^\varepsilon = X \cup \phi(A_\delta)$ and $M^\varepsilon = Y \cup \psi(A_\delta)$. Denote by $\{l_s : s \in S^1\}$ the

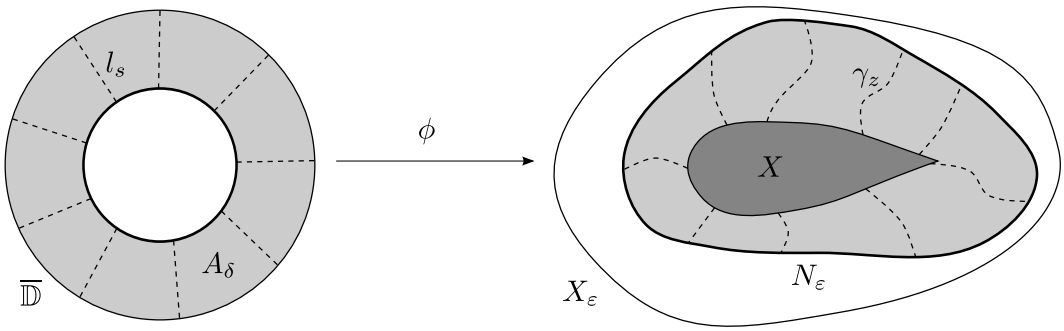

Figure 4.4: (Proof of Theorem 4.2) The extended inverse Riemann map $\phi\colon \overline{\mathbb{D}} \to \hat{C} \setminus \mathrm{int}(X)$ maps the annulus $A_\delta = \{x \in \mathbb{C} : 1 - \delta \le \|x\| \le 1\}$ onto $N^\varepsilon \setminus \mathrm{int}(X)$ such that $\partial\mathbb{D}$ is mapped onto ∂X. Each leaf γ_z of the foliation of $N^\varepsilon \setminus \mathrm{int}(X)$ is the image under ϕ of a radial line segment $l_s \subset A_\delta$.

foliation of A_δ by radial lines such that $\gamma_z = \phi(l_s)$ for $z = \phi((1-\delta)\cdot s)$ and $\theta_{z'} = \psi(l_s)$ for $z' = \psi((1-\delta)\cdot s)$.

From Theorem 4.13 we get that f induces a homeomorphism $\hat{f}\colon S^1 \to S^1$ such that

$$\psi|_{S^1} \circ \hat{f} = f|_{\partial U} \circ \phi|_{S^1}.$$

Hence we can extend $f\colon X \to Y$ to $F_\varepsilon\colon N^\varepsilon \to M^\varepsilon$ by mapping each arc $\gamma_z \subset \overline{N^\varepsilon \setminus X}$ to the corresponding arc $\theta_{z'} \subset \overline{M^\varepsilon \setminus X}$. More precisely, let $H\colon A_\delta \to A_\delta$ be the homeomorphism which maps the radial line segment l_s linearly to the radial line segment $l_{\hat{f}(s)}$ (that is, in polar coordinates $H = \mathrm{id}\,|_{[1-\delta,1]} \times \hat{f}$). Then define the homeomorphic extension of f to N^ε by setting

$$F_\varepsilon(x) = \begin{cases} f(x) & \text{if } x \in X, \\ \psi \circ H \circ \phi^{-1}(x) & \text{if } x \in N^\varepsilon \setminus X. \end{cases}$$

Note that for $0 < \varepsilon' < \varepsilon$, we can repeat the above construction to obtain Jordan domains $N^{\varepsilon'}$ and $M^{\varepsilon'}$ and a homeomorphism $F_{\varepsilon'}\colon N^{\varepsilon'} \to M^{\varepsilon'}$, and moreover $F_\varepsilon(x) = F_{\varepsilon'}(x)$ for all $x \in N^{\varepsilon'} \subset N^\varepsilon$.

Every connected component of $X \cap Y$ is a subset of a connected component of $N^\varepsilon \cap M^\varepsilon$. Let K^ε be the connected component of $N^\varepsilon \cap M^\varepsilon$ which contains C.

Claim 1. *$F_{\varepsilon'}$ has a fixed point in $K^{\varepsilon'}$ for every $\varepsilon' \in (0, \varepsilon]$.*

This follows from the main result of the previous section: the periodic point for f in C is also a periodic point for $F_{\varepsilon'}$ in $K^{\varepsilon'}$. By Theorem 4.10, $F_{\varepsilon'}$ has a fixed point in $K^{\varepsilon'}$.

Claim 2. *F_ε has a fixed point in $K^\varepsilon \cap (X \cap Y)$.*

By Claim 1 and the fact that for $0 < \varepsilon' < \varepsilon$, $F_{\varepsilon'}$ is the restriction of F_ε, F_ε has a

sequence of fixed points $(p_k)_{k\in\mathbb{N}}$ such that $p_k \in N^{\varepsilon/k} \cap M^{\varepsilon/k}$. By passing to a convergent subsequence, we get that $p_k \to p$ as $k \to \infty$, with $p \in K^\varepsilon \cap (X \cap Y)$ and $F_\varepsilon(p_k) = F_{\varepsilon/k}(p_k) = p_k$. By continuity $F_\varepsilon(p) = p$, which proves the claim.

Claim 3. *F_ε has a fixed point in C.*

By Claim 2 we know that F_ε (and hence $F_{\varepsilon'}$ for $\varepsilon' \in (0, \varepsilon]$) possesses a fixed point $p \in K^\varepsilon \cap (X \cap Y)$. If $p \notin C$, then p lies in another connected component of $K^\varepsilon \cap (X \cap Y)$, say C'. We can now apply Lemma 4.14 to $K^\varepsilon \cap (X \cap Y)$ and its connected components. Since $N^{\varepsilon'} \subseteq X_{\varepsilon'}$ and $M^{\varepsilon'} \subseteq Y_{\varepsilon'}$, we obtain that for $\varepsilon' > 0$ sufficiently small, $p \notin K^{\varepsilon'}$. By again applying Theorem 4.10 and Claim 2, we get that $F_{\varepsilon'}$ has a fixed point $p' \in K^{\varepsilon'} \cap (X \cap Y)$, which is also a fixed point for F_ε.

We can iterate this argument and obtain a sequence of F_ε-fixed points $(p_k)_{k\in\mathbb{N}}$ such that $p_k \in K^{1/k} \cap (X \cap Y)$. As in the proof of the previous claim, by passing to a convergent subsequence, $p_k \to p$ as $k \to \infty$, we get an F_ε-fixed point $p \in K^\varepsilon \cap (X \cap Y)$. If p lies in a connected component of $K^\varepsilon \cap (X \cap Y)$ distinct from C, then, by Lemma 4.14, $p \notin K^{\varepsilon'}$ for sufficiently small ε'. But by construction, $p \in K^{\varepsilon'}$ for all $\varepsilon' \in (0, \varepsilon]$, so $p \in C$, finishing the proof of Claim 3.

Claim 3 implies that f has a fixed point in C, and Theorem 4.2 is proved. □

Proof of Corollary 4.3. If C is a connected component of $X \cap Y$ which contains a point x such that $f^n(x) \in C$ for all $n \in \mathbb{N}$, then exactly as in the proof of Theorem 4.2, we can construct an extension F of f to a closed neighbourhood N^ε of X, for small $\varepsilon > 0$. Using Corollary 4.11 instead of Theorem 4.10, we get that the extension has a fixed point. The same argument then shows that one such fixed point lies in C, as required. □

4.4 Proofs of Lemmas 4.7 and 4.14

Proof of Lemma 4.7. We apply Proposition 4.6 with $U = D$ and $U' = E$. The set $D \cap E$ is a Jordan domain and by (1), its boundary curve consists of the (pairwise disjoint) sets

- $\partial D \cap \partial E$ (isolated points or closed arcs),
- $\alpha_i \subset \partial D$, $i \in I$ (open arcs), and
- $\beta_j \subset \partial E$, $j \in J$ (open arcs).

It also follows from Proposition 4.6, that $\overline{E \setminus D}$ is the union of Jordan domains A_i bounded by $\alpha_i \cup [a_i, b_i]_{\partial E}$ and $\overline{D \setminus E}$ is the union of Jordan domains B_j bounded by $\beta_j \cup [c_j, d_j]_{\partial D}$ (see Figure 4.1).

By Schoenflies' theorem, each Jordan domain A_i is homeomorphic to the closed unit disk $\overline{\mathbb{D}} = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. Let $h_i \colon A_i \to \overline{\mathbb{D}}$ be such homeomorphism. We can assume $h_i(a_i) = (-1, 0)$ and $h_i(b_i) = (1, 0)$. Then the partition of $\overline{\mathbb{D}}$ into vertical line segments $l_s := \{(x, y) \in \overline{\mathbb{D}} : x = s\}$, $s \in (-1, 1)$, gives rise to a partition $\{h_i^{-1}(l_s) : s \in (-1, 1)\}$ of A_i into closed arcs, each connecting a point of α_i to a point of $(a_i, b_i)_{\partial E}$. Similarly, one can obtain a partition of B_j into closed arcs connecting points on β_j to points on $(c_j, d_j)_{\partial D}$.

Note that $D \cup E$ is also a Jordan domain and that the $(a_i, b_i)_{\partial E}$, $i \in I$, and $(c_j, d_j)_{\partial D}$, $j \in J$, together with $\partial D \cap \partial E$ form a partition of its boundary. Thus we obtained a collection of (pairwise disjoint) closed arcs, each connecting precisely one point of $\partial(D \cap E)$ to precisely one point of $\partial(D \cup E)$. Denote the arc corresponding to $z \in \partial(D \cup E)$ by l^z, and let $l^z = \{z\}$ whenever $z \in \partial D \cap \partial E$.

Since $D \cup E$ is a Jordan domain, we can again apply Schoenflies' theorem to obtain a homeomorphism $h \colon \mathbb{R}^2 \setminus \operatorname{int}(D \cup E) \to \mathbb{R}^2 \setminus \mathbb{D}$. Let r_θ be the radial line segment in $\mathbb{R}^2 \setminus \mathbb{D}$ expressed in polar coordinates as $r_\theta = \{(r, \phi) : r \geq 1, \phi = \theta\}$. Then $\{r_\theta : \theta \in [0, 2\pi)\}$ forms a partition of $\mathbb{R}^2 \setminus \mathbb{D}$, which can be pulled back to a partition $\{h^{-1}(r_\theta) : \theta \in [0, 2\pi)\}$ of $\mathbb{R}^2 \setminus \operatorname{int}(D \cup E)$. This partition consists of (pairwise disjoint) arcs m^z, each connecting a point on $z \in \partial(D \cup E)$ to ∞.

Combining the above, we get that each point $z \in \partial(D \cup E)$ is the endpoint of two uniquely defined arcs l^z and m^z (l^z possibly being the trivial arc $\{z\}$). Let $L^z := l^z \cup m^z$, which is an arc connecting a point on $\partial(D \cap E)$ to ∞, for every $z \in \partial(D \cup E)$. Then $\{L^z : z \in \partial(D \cup E)\}$ is a partition of $\mathbb{R}^2 \setminus \operatorname{int}(D \cap E)$ with the desired properties.

At last, we construct a continuous function $\lambda^D \colon \mathbb{R}^2 \setminus \operatorname{int}(D) \to \mathbb{R}_{\geq 0}$ strictly monotonically increasing along each arc L^z and such that $\lambda^D|_{\partial D} \equiv 0$; the function $\lambda^E \colon \mathbb{R}^2 \setminus \operatorname{int}(D) \to \mathbb{R}_{\geq 0}$ can be constructed similarly.

For $x \in \mathbb{R}^2 \setminus \operatorname{int}(D)$, let z be the unique point in $\partial(D \cup E)$ such that $x \in L^z$ and let $\{p\} = L^z \cap \partial D$. If $p \notin E$, then $p = z$ and the arc $[p, x] \subset L^z$ is contained in $\mathbb{R}^2 \setminus \operatorname{int}(D \cup E)$. Then $h([p, x])$ is a straight line from $h(p)$ to $h(x)$ and we set $\lambda^D(x) = \ell(h([p, x]))$, where ℓ denotes the usual (Euclidean) arc length on arcs in $\mathbb{R}^2$. If $p \in E$ and $x \in E$, $[p, x] \subset L^z$ lies in a component A_i of $\overline{E \setminus D}$, and we set $\lambda^D(x) = \ell(h_i([p, x]))$. Finally, if $p \in E$ and $x \notin E$, $[p, x]$ is the union of the arcs $[p, z] \subset L^z$ and $[z, x] \subset L^z$, and we set $\lambda^D(x) = \ell(h_i([p, z])) + \ell(h([z, x]))$. □

Proof of Lemma 4.14. Suppose C and C' lie in the same component C_k of $X_{1/k}$ for all $k \in \mathbb{N}$. The space of subcontinua (non-empty, compact, connected subsets) of a compact planar set equipped with the Hausdorff metric is a compact complete metric space. By passing to a subsequence which we also call C_k, we have $C_k \to C_0$ as $k \to \infty$, where C_0 is itself non-empty, compact and connected. Convergence in the Hausdorff metric implies that $C \cup C' \subseteq C_0 \subseteq X$, contradicting that C and C' are distinct connected components of X. □

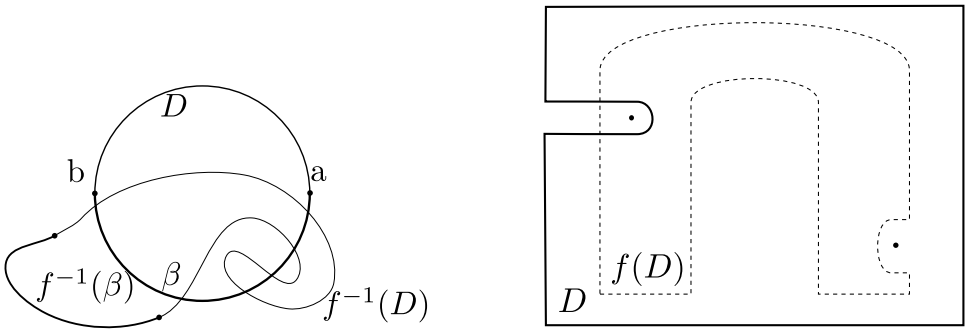

(a) (b)

Figure 4.5: (a) Jordan domain satisfying Bonino's [18] condition.
(b) Horseshoe map with neighbourhood of a period-two point removed from its domain: The map has period-three points of non-trivial braid type, but no period-two points, showing that a simple generalisation of period-forcing results to non-self maps of planar domains does not work.

4.5 Discussion of assumptions and possible extensions

Easy examples show that in our main results we cannot omit the hypothesis of a periodic orbit being completely contained in a given connected component of $X \cap f(X)$: take X to be a closed ε-neighbourhood of the straight line segment $S = [-1, 1] \times \{0\}$ and let $f\colon X \to Y$ be a homeomorphism which maps S to the semicircle $\{(x, y) : x^2 + y^2 = 1, y \geq 0\}$ with $f(-1, 0) = (1, 0)$ and $f(1, 0) = (-1, 0)$. For $\varepsilon > 0$ small, f has a period-two orbit (spread over two different connected components of $X \cap f(X)$) but no fixed point.

Furthermore, obvious counterexamples show that our main result, Corollary 4.3, fails when the domain $X \subset \mathbb{R}^2$ of the homeomorphism f is not compact. Also, we are not aware of generalisations to higher dimensions; in fact, simple counterexamples to Theorem 4.2 in its current form can be constructed, when $X \subset \mathbb{R}^n$ is an n-dimensional ball with $n \geq 3$.

Bonino [18] showed that if an orientation-preserving homeomorphism $f\colon \mathbb{R}^2 \to \mathbb{R}^2$ possesses non-escaping points in a topological disk D bounded by a simple closed curve C, and C can be split into two arcs $\alpha = [a, b]_C$, $\beta = [b, a]_C$ such that $D \cap f^{-1}(\beta) = \emptyset$ and $f^{-1}(D) \cap \alpha = \emptyset$ (see Figure 4.5(a)), then f has a fixed point *in* D. Bonino's theorem is similarly topological, but does not imply (and is not implied by) our results.

Another result related to Theorem 4.2 is due to Brown [24, Theorem 5]. Brown assumed that for the n-periodic point $x \in X$ there is a connected neighbourhood W of $\{x\} \cup \{f(x)\}$ with $f^i(W) \subset X$ for $i = 1, \ldots, n$, and showed that f then has a fixed point in X.

A somewhat different class of results relates to the question of period forcing. Gam-

baudo et al. [35] showed that a C^1 orientation-preserving embedding of the disk into itself has periodic orbits of all periods, if it has a period-three orbit of braid type different than the rotation by angle $2\pi/3$ (see also Kolev [56] for a topological version of this result). For $f(X) \not\subset X$, this statement is false: a counterexample is a version of the Smale horseshoe map, where a period-two point is removed from the disk together with a narrow strip such that the remaining domain X is still simply connected (see Figure 4.5(b)). Then f, the restriction of the horseshoe map to X, is an orientation-preserving homeomorphism with $X \cap f(X)$ connected, and one can choose the removed strip narrow enough such that f still has period-three orbits but no period-two orbit in X. We conclude with the following question:

Question 4.15. *Can one find conditions on X and $f(X)$, under which* every extension *of f to a homeomorphism of $\mathbb{R}^2$ has periodic orbits of every period* passing through *X, whenever f has a period-three point (of non-trivial braid type) in X?*

The method of extending a non-self map of $X \subset \mathbb{R}^2$ to a self-map of $\mathbb{R}^2$ (or $\mathbb{D}$) without adding fixed (or periodic) points could possibly also be applied to this problem.

Chapter 5

Piecewise affine model maps

Piecewise affine maps and piecewise isometries have received a lot of attention, either as simple, computationally accessible models for complicated dynamical behaviour, or as a class of systems with their own unique range of dynamical phenomena. For a list of examples, see [1–3, 8, 25, 26, 39, 40, 58, 76, 79, 81, 96–98] and references therein.

As we have seen in Section 1.2.6, fictitious play dynamics of 3×3 zero-sum games with unique, completely mixed Nash equilibrium can be represented by a Hamiltonian flow on S^3. In certain cases (for instance, the game (1.7) with $\beta = \sigma$; see the end of Section 1.3) this induced flow has a global first return section, which is a topological disk D and whose first return map R is continuous, piecewise affine, area-preserving and extends continuously to the boundary of the disk by $R|_{\partial D} = \mathrm{id}$.

In Chapter 2, we investigated the itinerary structure of fictitious play and the induced flow, and undertook numerical simulations of the qualitative behaviour of its first return maps. Then in Chapter 4, we looked at more general planar homeomorphisms, in particular non-self maps of compact planar regions, which is a frequently occurring setting when studying the behaviour of such first return maps restricted to certain subsets of their domains.

In this chapter, we take a somewhat different approach and study the phenomenology of such maps by considering instead a formally similar family of planar piecewise affine maps. These serve as models for the first return maps of the induced fictitious play flow by sharing their formal properties, while at the same time being simple enough to allow for explicit calculations. More precisely, each map in this one-parameter family is a homeomorphism of the standard unit square to itself, which is

- piecewise affine,
- (everywhere) continuous,

- area-preserving, and
- fixes the boundary of the square pointwise.

The nine-piece construction considered here seems to be the simplest possible (non-trivial) example satisfying all of these formal properties. The qualitative behaviour of these maps resembles that seen in Section 2.2 for the first return maps of fictitious play. In particular, the ways in which stochastic and (quasi-)periodic behaviour coexist seem to be of similar type, giving rise to similar phenomena.

Apart from these connections with fictitious play, the family of maps studied in this chapter is of independent mathematical interest, as it raises a number of questions related to classical problems (such as, for example, the ergodic hypothesis). These are often very hard to answer generally, but might be accessible for simple models such as this one, where explicit calculations and numerical simulations are significantly easier.

The chapter is organised as follows. In Sections 5.1 and 5.2 we present a geometric construction of our family of maps and describe its basic formal properties. Then, interested in the long-term behaviour of iterates of the maps, our next goal is to establish the existence of certain invariant regions. For that, in Section 5.3 we develop some technical results about periodic orbits and invariant curves, and in Section 5.4 prove their existence for certain parameter values. Finally, in Section 5.5 we discuss the dynamics for more general parameter values, present numerical observations and discuss open questions.

The contents of this chapter are contained in the preprint [73], which at the time of submission of this thesis is submitted for publication.

5.1 Construction of the family of maps

Let us denote the unit square by $\mathcal{S} = [0,1] \times [0,1]$. We construct a one-parameter family of continuous, piecewise affine maps $F_\theta \colon \mathcal{S} \to \mathcal{S}$, $\theta \in (0, \frac{\pi}{4})$, as follows (see Figure 5.1 for an illustration).

Denote the four vertices of $\mathcal{S}$ by $E_1 = (0,0)$, $E_2 = (1,0)$, $E_3 = (1,1)$ and $E_4 = (0,1)$. In the following we will use indices $i \in \mathcal{I} = \{1,2,3,4\}$ with cyclic order, that is, with the understanding that index $i+1$ is 1 for $i = 4$ and index $i-1$ is 4 for $i = 1$.

Let $\theta \in (0, \frac{\pi}{4})$, and for $i \in \mathcal{I}$ let L_i be the ray through E_i, such that the angle between the segment $\overline{E_i E_{i+1}}$ and L_i is θ. Let $P_i \in \operatorname{int}(\mathcal{S})$ be the point $L_{i-1} \cap L_i$, then the P_i, $i \in \mathcal{I}$, form a smaller square inside $\mathcal{S}$. Now we divide $\mathcal{S}$ into the following nine regions (see Figure 5.1(left)):

- four triangles $\mathcal{A}_i = \Delta(E_i, E_{i+1}, P_i)$, $i \in \mathcal{I}$, each adjacent to one of the sides of $\mathcal{S}$;

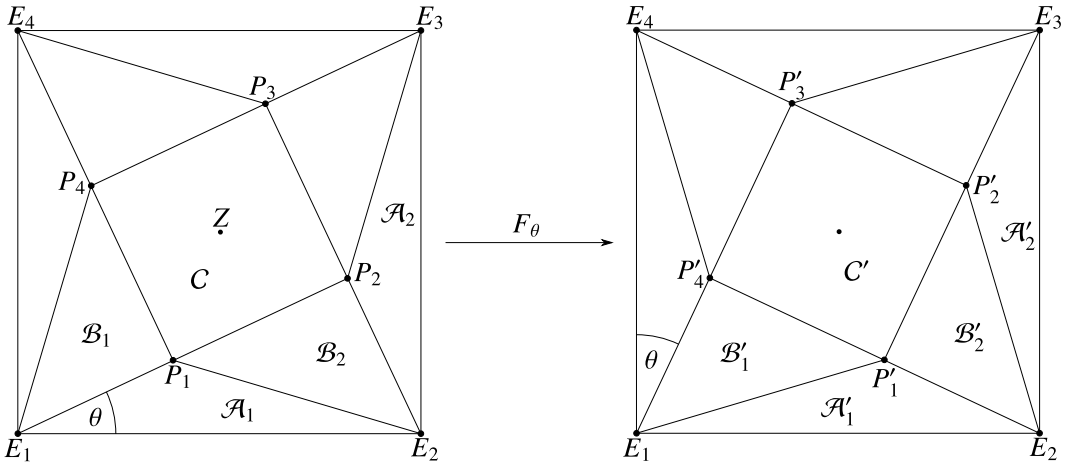

Figure 5.1: (Construction of the map $F = F_\theta$) $F(E_i) = E_i$, $F(P_i) = P'_i$; F is affine on each of $\mathcal{A}_i$, $\mathcal{B}_i$ and C, such that $F(\mathcal{A}_i) = \mathcal{A}'_i$, $F(\mathcal{B}_i) = \mathcal{B}'_i$ and $F(C) = C'$.

- four triangles $\mathcal{B}_i = \Delta(E_i, P_i, P_{i-1})$, $i \in \mathcal{I}$, each sharing one side with $\mathcal{A}_{i-1}$ and $\mathcal{A}_i$;
- a square $C = \square(P_1, P_2, P_3, P_4)$, each side of which is adjacent to one of the $\mathcal{B}_i$.

Now we repeat the same construction 'in reverse orientation', to obtain a second, very similar partition of $\mathcal{S}$, as shown in Figure 5.1(right). Here we denote the vertex of the inner square which has the same y-coordinate as P_1 by P'_1, and the other vertices of the inner square by P'_2, P'_3, P'_4, in counter-clockwise order. For $i \in \mathcal{I}$ we denote the triangles $\mathcal{A}'_i = \Delta(E_i, E_{i+1}, P'_i)$, $\mathcal{B}'_i = \Delta(E_i, P'_i, P'_{i-1})$ and the square $C' = \square(P'_1, P'_2, P'_3, P'_4)$.

Finally, the map $F = F_\theta \colon \mathcal{S} \to \mathcal{S}$ is uniquely defined by the data

- $F(E_i) = E_i$, $i \in \mathcal{I}$;
- $F(P_i) = P'_i$, $i \in \mathcal{I}$;
- F affine on each of the pieces $\mathcal{A}_i$, $\mathcal{B}_i$, and C.

5.2 Properties of the maps

It is easy to see that F is a piecewise affine homeomorphism of $\mathcal{S}$, with $F(\mathcal{A}_i) = \mathcal{A}'_i$ and $F(\mathcal{B}_i) = \mathcal{B}'_i$ for each $i \in \mathcal{I}$, and $F(C) = C'$. Moreover, F is area- and orientation-preserving, since dF is constant on each of the pieces, with $\det dF = 1$ everywhere. Note that while F is continuous, its derivative dF has discontinuity lines along the boundaries of the pieces; we call these the *break lines*. Note also that $F|_{\partial\mathcal{S}} = \mathrm{id}$.

We denote $P_1 = (s, t)$, where $t \in (0, \frac{1}{2})$ and $s \in (0, \frac{1}{2})$. The coordinates of the other points P_i and P_i' are then given by symmetry. Simple geometry gives that s, t and θ satisfy

$$t - t^2 = s^2, \quad t = \sin^2\theta, \quad s = \sin\theta\cos\theta. \tag{5.1}$$

The map F is given by three types of affine maps A, B and C:

- $A\colon \mathcal{A}_1 \to \mathcal{A}_1'$ is a shear fixing $\overline{E_1E_2} = [0,1] \times \{0\}$ and mapping P_1 to P_1':

$$A(x, y) = \begin{pmatrix} 1 & \frac{1-2s}{t} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

  It leaves invariant horizontal lines and moves points in $\mathrm{int}(\mathcal{A}_1)$ to the right (since $(1 - 2s)/t > 0$).

- $B\colon \mathcal{B}_1 \to \mathcal{B}_1'$ is a linear scaling map with a contracting and an expanding direction, defined by $B(E_1) = E_1$, $B(P_1) = P_1'$ and $B(P_4) = P_4'$:

$$B(x, y) = \frac{1}{t - s} \begin{pmatrix} t^2 - (1 - s)^2 & (1 - 2s)t \\ (2s - 1)t & (2t - 1)t \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

  It can be checked that the contracting direction of B lies in the sector between $\overline{E_1P_4}$ and $\overline{E_1E_4}$, and the expanding direction in the sector betwen $\overline{E_1E_2}$ and $\overline{E_1P_1}$. From general theory we also have that B preserves the quadratic form given by

$$Q_B(x, y) = t(x^2 + y^2) - xy. \tag{5.2}$$

- $C\colon \mathcal{C} \to \mathcal{C}'$ is the rotation about $Z = (\frac{1}{2}, \frac{1}{2})$, mapping P_i to P_i', $i \in \mathcal{I}$:

$$C(x, y) = \begin{pmatrix} 2s & 2t - 1 \\ 1 - 2t & 2s \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 1 - t - s \\ t - s \end{pmatrix}.$$

  The rotation angle is $\alpha = \frac{\pi}{2} - 2\theta$, where θ is the parameter angle in the construction of the map F.

All other pieces of F are analogous, by symmetry of the construction. To capture this high degree of symmetry, we make the following observations which follow straight from the definition of F.

Lemma 5.1. *Let R denote the rotation about Z by the angle $\frac{\pi}{2}$. Then F and R commute:*

$$F \circ R = R \circ F.$$

Lemma 5.2. *Let $S(x, y) = (y, x)$. Then F is S-reversible, that is, S conjugates F and F^{-1}:*

$$S \circ F \circ S = F^{-1}.$$

Further, F is T_1- and T_2-reversible for the reflections $T_1(x, y) = (1 - x, y)$ and $T_2(x, y) = (x, 1 - y)$.

Heuristically, F acts similarly to a twist map: The iterates $F^n(X)$ of any point $X \in \text{int}(\mathcal{S})$ rotate counterclockwise about Z as $n \to \infty$. The 'rotation angle' (the angle between $\overline{ZX}$ and $\overline{ZF(X)}$) is not constant, but it is bounded away from zero as long as X is bounded away from $\partial\mathcal{S}$; in particular, every point whose orbit stays bounded away from the boundary runs infinitely many times around the centre.

Note also that the rotation angle is monotonically decreasing along any ray emanating from the centre Z. However, it is not strictly decreasing, as all points in C rotate by the same angle; this sets the map F apart from a classical twist map, for which strict monotonicity (the 'twist condition') is usually required.

5.3 Invariant circles

Clearly, the circle inscribed to the inner square C and all concentric circles in it centred at Z are invariant under F, which acts as a rotation on these circles. When θ is a rational multiple of π, the rotation C is a rational (periodic) rotation, and a whole regular n-gon inscribed to C is F-invariant, see Figure 5.4.

We are interested in other invariant circles encircling Z, as these form barriers to the motion of points under F and provide a partitioning of $\mathcal{S}$ into F-invariant annuli. Numerical simulations indicate that such curves exist for many parameter values θ and create invariant annuli, on which the motion is predominantly stochastic.

This section follows closely the arguments of Bullett [25], where in a similar way, invariant circles are studied for a piecewise linear version of the standard map. The idea is to study the orbits of the points where the invariant circles intersect break lines and to prove that these follow a strict symmetry pattern, forming so-called cancellation orbits.

We consider invariant circles Γ on which F preserves the S^1-order of points. This, for example, is the case if all rays from Z intersect the circle Γ in precisely one point. For an invariant circle Γ, we denote the rotation number of $F|_\Gamma \colon \Gamma \to \Gamma$ by $\rho_\Gamma = \rho(F|_\Gamma)$. By simple geometric considerations, we get the following lemma.

Lemma 5.3. *Let Γ_1, Γ_2 be two invariant circles for F encircling Z. If Γ_1 is contained in the component of $\mathcal{S} \setminus \Gamma_2$ containing Z, then $\rho_{\Gamma_1} \geq \rho_{\Gamma_2}$.*

In other words, if there is a family of such nested invariant circles, their rotation number is monotonically decreasing as the circles approach $\partial\mathcal{S}$. It also follows that the rotation number ρ of any orbit is bounded above by the rotation number on the centre piece C, that is,

$$0 \leq \rho \leq \frac{\alpha}{2\pi} = \frac{1}{4} - \frac{\theta}{\pi}.$$

We now consider F-invariant circles near the boundary $\partial\mathcal{S}$, which do not intersect the centre piece C. Any such curve Γ intersects exactly two types of break line segments: the segments $\mathcal{B}_i \cap \mathcal{A}_i = \overline{E_i P_i}$ and $\mathcal{A}_i \cap \mathcal{B}_{i+1} = \overline{E_{i+1} P_{i+1}}$, $i \in \mathcal{I}$. Let us call these intersection points $U_i = \Gamma \cap (\mathcal{B}_i \cap \mathcal{A}_i)$ and $V_i = \Gamma \cap (\mathcal{A}_i \cap \mathcal{B}_{i+1})$.

We say that an invariant curve which encircles Z and on which F preserves the S^1-order is *rotationally symmetric*, if it is invariant under the rotation R (cf. Lemma 5.1): $R(\Gamma) = \Gamma$.

In the remainder of this section we will use the fact that any invariant circle Γ with rational rotation number is of one of the following two types [54]:

- pointwise periodic, that is, $F|_\Gamma$ is conjugate to a rotation of the circle;
- non-periodic, that is, $F|_\Gamma$ is not conjugate to a rotation; however, in this case, $F|_\Gamma$ still has at least one periodic orbit.

Following the ideas in [25], we now prove a number of results illustrating the importance of the orbits of U_i and V_i for the invariant circle containing them.

Lemma 5.4. *Let Γ be a rotationally symmetric invariant circle disjoint from C. Assume that $F|_\Gamma$ has rational rotation number $\rho_\Gamma = \frac{p}{q} \in \mathbb{Q}$, and that $F|_\Gamma$ is periodic. Then for $i \in \mathcal{I}$ the orbit* $\mathrm{orb}(U_i) = \{F^n(U_i) : n \in \mathbb{Z}\}$ *contains some V_j, $j \in \mathcal{I}$, and vice versa.*

Proof. By symmetry it is sufficient to show the result for U_1. Suppose for a contradiction that $V_j \notin \mathrm{orb}(U_1)$ for all $j \in \mathcal{I}$. Let $O = (\bigcup_i \mathrm{orb}(U_i)) \cup (\bigcup_j \mathrm{orb}(V_j))$, which by periodicity of $F|_\Gamma$ is finite. Let $X, Y \in O$ be the two points closest to U_1 on either side along Γ. Consider the line segment $\overline{XY}$ which crosses the break line going through U_1. Then $F(\overline{XY})$ is a 'bent' line, consisting of two straight line segments, so that F maps the triangle $\Delta(X, U_1, Y)$ to a quadrangle.

Now, by assumption and periodicity of $F|_\Gamma$, there exists $k > 1$ such that $F^k(U_1) = U_j$ for some $j \in \mathcal{I}$ and $F^l(U_1) \notin \{U_i : i \in \mathcal{I}\} \cup \{V_j : j \in \mathcal{I}\}$ for $1 \leq l < k$. This implies that the triangle $\Delta(F(X), F(U_1), F(Y))$ is mapped by F^{k-1} to the triangle $\Delta(F^k(X), F^k(U_1), F^k(Y))$ without bending any of its sides (since X and Y are the points in O closest to U_1).

By symmetry we have that $\Delta(F^k(X), F^k(U_1), F^k(Y)) = \tilde{R}(\Delta(X, U_1, Y))$, where $\tilde{R}$ is

the rotation about Z by one of the angles $0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}$, and hence

$$\operatorname{area}(\Delta(F^k(X), F^k(U_1), F^k(Y))) = \operatorname{area}(\Delta(X, U_1, Y)).$$

But F is area-preserving, so

$$\operatorname{area}(\Delta(F^k(X), F^k(U_1), F^k(Y))) = \operatorname{area}(\Delta(F(X), F(U_1), F(Y))),$$

which is a contradiction because F maps $\Delta(X, U_1, Y)$ to a quadrangle which either properly contains or is properly contained in $\Delta(F(X), F(U_1), F(Y))$. □

With slightly bigger effort, we can extend the result to the case of non-periodic $F|_\Gamma$.

Lemma 5.5. *Let Γ be a rotationally symmetric invariant circle disjoint from C. Assume that $F|_\Gamma$ has rational rotation number $\rho_\Gamma = \frac{p}{q} \in \mathbb{Q}$, and that $F|_\Gamma$ is not periodic. Then for $i \in \mathcal{I}$, the orbit of U_i contains some V_j, $j \in \mathcal{I}$, and vice versa.*

Proof. As in the previous lemma, we give a proof for U_1, the other cases following by symmetry. We distinguish two cases.

Case 1: U_1 non-periodic for F. Let us write $z_k = F^k(U_1)$ for $k \in \mathbb{Z}$. Since $\rho_\Gamma = \frac{p}{q} \in \mathbb{Q}$, there exist points Q and Q' in Γ, each periodic of period q, such that $z_{nq} \to Q$ and $z_{-nq} \to Q'$ as $n \to \infty$. Note that F^q is affine in a sufficiently small neighbourhood on either side of Q. Then for sufficiently large N, the points z_{nq}, $n > N$, lie on the straight line segment $\overline{z_{Nq}Q}$ (the contracting direction at Q). Hence for large n, Γ contains the straight line segment $\overline{z_{nq}Q}$. Analogously, for large n, the straight line segment $\overline{Q'z_{-nq}}$ is contained in Γ. In particular, $\ell = \overline{z_{-(m+2)q}z_{-mq}}$ is contained in Γ for large m. But $U_1 \in F^{(m+1)q}(\ell)$, so $F^{nq}(\ell)$ has a kink for large n unless $z_{Nq} = V_j$ for some N and j (note that since U_1 is non-periodic, $U_i \notin \operatorname{orb}(U_1)$, $i \in \mathcal{I}$). Since $F^{nq}(\ell)$ is near Q for large n, it has to be straight, and it follows that $V_j \in \operatorname{orb}(U_1)$ for some j.

Case 2: U_1 periodic for F. In this case $F^q(U_1) = U_1$ and the argument is similar to the proof of Lemma 5.4. Assume for a contradiction that $V_j \notin \operatorname{orb}(U_1)$ for all j. By symmetry, this implies $V_j \notin \bigcup_i \operatorname{orb}(U_i)$. Pick $X, Y \in \bigcup_i \operatorname{orb}(U_i)$ nearest to V_1 from each side and denote by S the segment of Γ between X and Y. Since the straight line segment $\overline{XY}$ crosses the break line which contains V_j, its image $F(\overline{XY})$ has a kink. Therefore, since F is area-preserving, the area between $F(S)$ and $\overline{F(X)F(Y)}$ is either greater or less than the area between S and $\overline{XY}$. For $0 \le k < q$, the area between $F^k(S)$ and $\overline{F^k(X)F^k(Y)}$ is equal to the area between $F^{k+1}(S)$ and $\overline{F^{k+1}(X)F^{k+1}(Y)}$, unless $F^k(S)$ contains one of the V_j. Whenever $V_j \in F^k(S)$ for some j and k (which can happen at most four times for $0 \le k < q$), this area decreases or increases. By symmetry, these up to four changes have the same form, so the

area either always decreases or always increases. Note on the other hand that $F^q(S) = S$, $F^q(X) = X$ and $F^q(Y) = Y$ (since X and Y are in the q-periodic orbit of U_1). So if we denote the region between S and $\overline{XY}$ by Ω, area$(F^q(\Omega)) \neq$ area(Ω), which contradicts the fact that F is area-preserving. This finishes the proof. □

Combinining the above lemmas, we get the following result.

Proposition 5.6. *Let Γ be a rotationally symmetric invariant circle disjoint from C with rotation number $\rho_\Gamma = \frac{p}{q} \in \mathbb{Q}$. Then for every $i \in \mathcal{I}$, the F-orbit of U_i contains some V_j, $j \in \mathcal{I}$, and vice versa. Moreover, every such orbit contains an equal number n of the U_i and V_j, which are traversed in alternating order. If $n \geq 2$, then any such orbit is periodic.*

Remark 5.7. In [25], Bullett coined the term 'cancellation orbits' for these orbits of break points on an invariant circle, reflecting the insight that each 'kink' introduced by the discontinuity of dF at one such point needs to be 'cancelled out' by an appropriate 'reverse kink' at another discontinuity point of dF, if the invariant circle has rational rotation number.

However, cancellation orbits can also occur on invariant circles with $\rho \notin \mathbb{Q}$. In that case these orbits are not periodic, and each cancellation orbit would only contain one of the U_i and one of the V_j. We will see examples for this in Section 5.4.

The next result shows that cancellation orbits in fact determine the behaviour of the whole map $F|_\Gamma$.

Proposition 5.8. *Let Γ be a rotationally symmetric invariant circle disjoint from C. Then $F|_\Gamma$ is periodic if and only if the U_i- and V_i-orbits are periodic.*

Proof. Of course, if the rotation number ρ_Γ is irrational, neither $F|_\Gamma$ nor the break point orbits can be periodic, so we only need to consider rational rotation number.

Further, if $F|_\Gamma$ is periodic, so are all cancellation orbits. For the converse, suppose for a contradiction that U_i and V_i, $i \in \mathcal{I}$, are periodic, but $F|_\Gamma$ is not. We repeat an argument already familiar from the proof of Lemma 5.5.

Pick any non-periodic point $P \in \Gamma$, then there exists $Q \in \Gamma$, such that $F^{nq}(P) \to Q$ as $n \to \infty$. Note that F^q is affine in a sufficiently small neighbourhood on either side of Q. Then for $n > N$ sufficiently large, the points $F^{nq}(P)$ lie on a straight line segment from $F^{Nq}(P)$ to Q (the contracting direction at Q). So Γ contains a straight line segment S which expands under F^{-q}. This expansion cannot continue indefinitely, so $F^{-mq}(I)$ must meet some U_i (or V_i) for some m. But then this U_i (or V_i) cannot be periodic, which contradicts the assumption. □

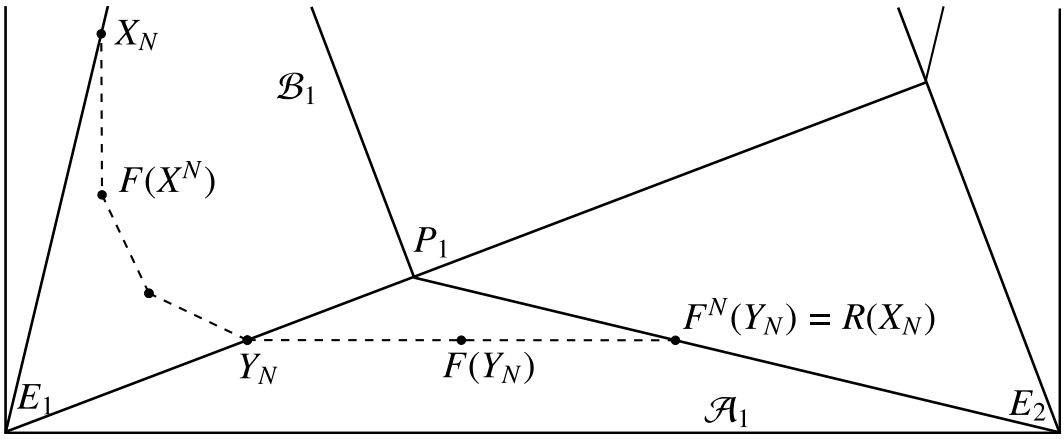

Figure 5.2: (Proof of Theorem 5.9) Construction of a periodic cancellation orbit for $K = 3$, $N = 2$. The points $X_N \in \overline{E_1P_4}$ and $Y_N \in \overline{E_1P_1}$ are chosen such that $F^K(X_N) = Y_N$ (Lemma 5.10) and $F^N(Y_N) = R(X_N) \in \overline{E_2P_1}$ (Lemma 5.11). The dots are the F-iterates of X_N, the dashed lines indicate the line segments making up (part of) the invariant circle Γ_K^N.

5.4 Special parameter values

In this section, we will show that for a certain countable subset of parameter values $\theta \in (0, \frac{\pi}{4})$, the map $F = F_\theta$ has invariant circles of the form described in the previous section.

Theorem 5.9. *There exists a sequence of parameter angles $\theta_3, \theta_4, \ldots \in (0, \frac{\pi}{4})$, $\theta_K \to \frac{\pi}{4}$ as $K \to \infty$, such that for each K, F_{θ_K} has a countable collection of invariant circles $\{\Gamma_K^N : N \geq 0\}$, each of rational rotation number $\rho(\Gamma_K^N) = 1/(4(K+N))$. The curves Γ_K^N consist of straight line segments, are rotationally symmetric and converge to the boundary $\partial\mathcal{S}$ as $N \to \infty$.*

Proof. We will prove the result by explicitly finding periodic orbits in $\bigcup_i(\mathcal{A}_i \cup \mathcal{B}_i)$, which hit the break lines whenever passing from one of the pieces to another. More precisely, we will show that for $K \geq 3$ and $N \geq 0$, there is a parameter value $\theta = \theta_K$ and a point $X_N \in \overline{E_1P_4}$, such that $F^n(X_N) \in \mathcal{B}_1$ for $0 \leq n < K$, $F^K(X_N) \in \overline{E_1P_1}$, $F^n(X_N) \in \mathcal{A}_1$ for $K \leq n < K+N$, and $F^{N+K}(X_N) = R(X_N) \in \overline{E_2P_1}$, where R is the counter-clockwise rotation by $\frac{\pi}{2}$ about the centre of the square (see Figure 5.2). By symmetry this clearly gives a periodic cancellation orbit, and an invariant circle is then given by

$$\Gamma_K^N = \bigcup_{i=0}^{4(K+N)} \overline{F^i(X_N)F^{i+1}(X_N)}.$$

We use the following two lemmas, whose proofs are delayed until after the proof of the theorem.

Lemma 5.10. *For every $K \in \mathbb{N}$, $K \geq 3$, there exists a parameter $\theta_K \in (0, \frac{\pi}{4})$, such that for $F = F_{\theta_K}$, $F^k(\overline{E_1P_4}) \subset \mathcal{B}_1$ for $0 \leq k < K$, and $F^K(\overline{E_1P_4}) = \overline{E_1P_1} = \mathcal{B}_1 \cap \mathcal{A}_1$. For $K \to \infty$, the angle θ_K tends to $\frac{\pi}{4}$.*

Lemma 5.11. *For every $\theta \in (0, \frac{\pi}{4})$ and $N \geq 0$ there exists a point $Y_N \in \overline{E_1P_1} = \mathcal{B}_1 \cap \mathcal{A}_1$ such that $F^n(Y_N) \in \mathcal{A}_1$ for $0 \leq n < N$ and $F^N(Y_N) \in \overline{E_2P_1} = \mathcal{A}_1 \cap \mathcal{B}_2$. For $N \to \infty$, the points Y_N converge to E_1.*

By Lemma 5.10, for every $K \geq 3$ we can find θ_K such that F^K maps $\overline{E_1P_4}$ to $\overline{E_1P_1}$ (in $\mathcal{B}_1$). Then by Lemma 5.11, for any $N \geq 0$ there exists $X_N \in \overline{E_1P_4}$ such that $F^K(X_N) = Y_N \in \overline{E_1P_1}$ and $F^{K+N}(X_N) \in \overline{E_2P_1}$, and it can be seen that each open line segment $(F^k(X_N), F^{k+1}(X_N))$, $k = 0, \ldots, K+N-1$, lies in the interior of either $\mathcal{B}_1$ ($0 \leq k < K$) or $\mathcal{A}_1$ ($K \leq k < K+N$), see Figure 5.2.

Further, B preserves the quadratic form $Q_B(x, y) = t(x^2 + y^2) - xy$. With $X_N = (x_1, y_1) \in \overline{E_1P_4}$ and $F^K(X_N) = (x_2, y_2) \in \overline{E_1P_1}$ as above, a simple calculation shows that $Q_B(X_N) = Q_B(F^K(X_N))$ implies that $x_1 = y_2$. Then with $F^{K+N}(X_N) = (x_3, y_3) \in \overline{E_2P_1}$, one has $y_3 = y_2$, as A preserves the y-coordinate. We have that $R(\overline{E_1P_4}) = \overline{E_2P_1}$, and since $x_1 = y_3$, it follows that $F^{K+N}(X_N) = R(X_N)$. Rotational symmetry (Lemma 5.1) then implies that X_N is a periodic point of period $4(K+N)$ for F, and the line segments connecting the successive F-iterates of X_N form a rotationally symmetric invariant circle for F.

Hence we obtain, for each $K \geq 3$, a parameter θ_K such that F_{θ_K} has a sequence of rotationally invariant circles Γ_K^N, $N \geq 0$, each consisting of straight line segments. Moreover, Lemma 5.11 implies that $\Gamma_K^N \to \partial\mathcal{S}$ (in the Hausdorff metric) and $\rho_{\Gamma_K^N} = 1/(4(K+N)) \to 0$ as $N \to \infty$. □

Proof of Lemma 5.10. First, recall that the map $B\colon \mathcal{B}_1 \to \mathcal{B}_1 \cup \mathcal{A}_1 \cup C$ leaves invariant the quadratic form (5.2). Then, using $P_4 = (1-s, t)$ and $P_1 = (s, t)$, we calculate

$$Q_B(P_4) - Q_B(P_1) = \left((1-s)^2 + t^2 - \frac{(1-s)t}{t}\right) - \left(s^2 + t^2 - \frac{st}{t}\right) = 0,$$

and hence $Q_B(P_4) = Q_B(P_1) =: c$. So the line $\{x \in \mathcal{B}_1 : Q_B(x) = c\}$ is a segment of a hyperbola connecting P_4 and P_1. Therefore, with $V = \{x \in \mathcal{B}_1 : Q_B(x) \leq c\}$, we get $F(V) = B(V) \subset (V \cup \mathcal{A}_1)$.

Further, note that B maps rays through E_1 to other rays through E_1. This implies that all points on the straight line segment $\overline{E_1P_4} = \mathcal{A}_4 \cap \mathcal{B}_1$ remain in the piece $\mathcal{B}_1$ for equally many iterations of the map F, before being mapped into $\mathcal{A}_1$. In particular, if for $X \in \overline{E_1P_4}$ we have that

$$F^k(X) \in \begin{cases} \mathcal{B}_1 & \text{if } 0 \leq k < K, \\ \overline{E_1P_1} & \text{if } k = K, \end{cases}$$

then the same holds for every other point $X' \in \overline{E_1P_4}$, and $F^K(\overline{E_1P_4}) = \overline{E_1P_1}$.

We will now show that there exists a sequence of parameter values θ_K, $K \geq 3$, such that for $F = F_{\theta_K}$, $F^k(P_4) = B^k(P_4) \in \mathcal{B}_1$ for $0 \leq k \leq K$ and $F^K(P_4) = P_1$. For this, we need a few elementary facts about the map B, which follow from straightforward (but rather tedious) calculations:

- Let $f(t) = \sqrt{1-4t^2}$. Then the hyperbolic map B has two eigendirections

$$v_1 = \begin{pmatrix} 1+f(t) \\ 2t \end{pmatrix}, \quad v_2 = \begin{pmatrix} 2t \\ 1+f(t) \end{pmatrix}$$

  with corresponding eigenvalues

$$\lambda_1 = \lambda = \frac{4t^2 - 2t + (2s-1)(1+f(t))}{2(t-s)} > 1, \quad \lambda_2 = \lambda^{-1} < 1. \tag{5.3}$$

- By a linear change of coordinates $\Phi\colon \mathbb{R}^2 \to \mathbb{R}^2$ mapping v_1 and v_2 to $(1,0)$ and $(0,1)$, respectively, one gets a conjugate linear map

$$\tilde{B} = \Phi \circ B \circ \Phi^{-1} = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}.$$

  Setting $\Phi(P_4) = \Phi(t, 1-s) =: (x_1, y_1)$ and $\Phi(P_1) = \Phi(s,t) =: (x_2, y_2)$, a somewhat tedious calculation gives

$$\frac{x_2}{x_1} = \frac{2ts + t(f(t)-1)}{2t^2 + (1-s)(f(t)-1)}. \tag{5.4}$$

Now, since $\tilde{B}^K = \Phi \circ B^K \circ \Phi^{-1}$, it follows that $B^K(P_4) = P_1$ is equivalent to $\tilde{B}^K\Phi(P_4) = \Phi(P_1)$. By the simple form of $\tilde{B}^K$, this is equivalent to $x_2 = \lambda^K x_1$, that is,

$$K = \frac{\log(x_2/x_1)}{\log(\lambda)}. \tag{5.5}$$

Substituting (5.3) and (5.4) into (5.5), and using $s = \sqrt{t-t^2}$ from (5.1), we get an expression $K = K(t)$ for $0 < t < \frac{1}{2}$. Then K is differentiable and strictly monotonically decreasing as a function of t, and (by application of L'Hôpital's rule)

$$K \to \begin{cases} 2 & \text{as } t \to 0, \\ \infty & \text{as } t \to \frac{1}{2}. \end{cases}$$

Since $\sin^2\theta = t$ with $t \in (0, \frac{1}{2})$, $\theta \in (0, \frac{\pi}{4})$, we get that for each $K \geq 3$ there exists $\theta_K \in (0, \frac{\pi}{4})$

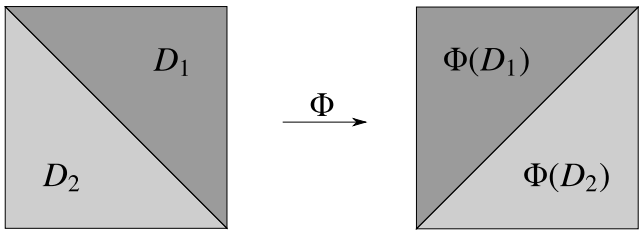

Figure 5.3: The form of the map $\Phi\colon \mathcal{R} \to R(\mathcal{R})$, where $\mathcal{R}$ is the quadrangle formed by two consecutive cancellation orbit points on each of two adjacent invariant circles Γ_K^N, Γ_K^{N+1}. The first return map of F to $\mathcal{R}$ is (by symmetry) the fourth iterate of Φ.

such that $B^K(P_4) = P_1$, hence $B^K(\overline{E_1P_4}) = \overline{E_1P_1}$, as claimed. □

Proof of Lemma 5.11. First, the shear map $A\colon \mathcal{A}_1 \to \mathcal{A}_1 \cup \mathcal{B}_2$ is such that any point $(x, y) \in \mathcal{A}_1$ is mapped to $F(x, y) = A(x, y) = (x + \tilde{c}y, y)$, $\tilde{c} > 0$. By continuity, it follows that for any $N \geq 0$ there exists a point $Y_N \in \overline{E_1P_1} = \mathcal{B}_1 \cap \mathcal{A}_1$ such that $F^n(Y_N) \in \mathcal{A}_1$ for $0 \leq n < N$ and $F^N(Y_N) \in \overline{E_2P_1} = \mathcal{A}_1 \cap \mathcal{B}_2$. Clearly, $Y_0 = P_1 = (s, t)$, and one can calculate that

$$Y_N = \frac{1}{1 + N(1 - 2s)}(s, t), \quad N \geq 0,$$

and $Y_N \to (0, 0) = E_1$ as $N \to \infty$. □

In the proof of the theorem, for a sequence of special parameter values θ_K we constructed periodic orbits hitting the break lines and invariant circles made up of line segments connecting the points of the periodic orbits. A closer look at the behaviour of F *between* any two consecutive invariant circles in this construction reveals that in fact the dynamics of F in these regions is very simple for these parameter values.

To see this, let $\theta = \theta_K$, $K \geq 3$, $N \geq 0$, and take $X \in \Gamma_K^N \cap \overline{E_1P_1}$, $Y \in \Gamma_K^{N+1} \cap \overline{E_1P_1}$. Then X and Y give rise to periodic cancellation orbits of periods $4(K + N)$ and $4(K + N + 1)$ on the respective invariant circles. Let $\mathcal{R}$ be the quadrangle with vertices $X, Y, F(Y), F(X)$. Then $F^{K+N}(X) = R(X)$ and $F^{K+N+1}(Y) = R(Y)$ lie in $\overline{E_2P_2} = R(\overline{E_2P_2})$, and F^{K+N} maps the triangle $D_1 = \Delta(X, F(Y), F(X))$ affinely to the triangle $R(\Delta(X, Y, F(X)))$ and F^{K+N+1} maps $D_2 = \Delta(X, Y, F(Y))$ affinely to $R(\Delta(Y, F(Y), F(X)))$. This gives a piecewise affine map $\Phi\colon D_1 \cup D_2 = \mathcal{R} \to R(\mathcal{R})$ of the simple form shown in Figure 5.3. By symmetry, the first return map of F to the quadrangle $\mathcal{R}$ is then the fourth iterate of such map, and is easily seen to preserve the y-coordinate.

It follows immediately that in fact all points in the annulus between Γ_K^N and Γ_K^{N+1} lie on invariant circles. These invariant circles are of the same form as the Γ_K^N, that is, they consist of line segments parallel to those explicitly constructed in the proof of Theorem 5.9. The rotation number of the invariant circle through the point $W \in \overline{E_1P_1}$ changes continu-

ously (more precisely, linearly) from $1/(4(K+N))$ to $1/(4(K+N+1))$, as W goes from X to Y.

These invariant circles take on both rational and irrational rotation numbers, and their intersections with the break lines are not necessarily periodic as those of the Γ_K^N, but their orbits still form cancellation orbits, since $F^K(\overline{E_1P_4}) = \overline{E_1P_1}$ (see Remark 5.7). We get the following corollary.

Corollary 5.12. *For $\theta = \theta_K$, $K \geq 3$, as in Theorem 5.9, the annulus between $\partial\mathcal{S}$ and Γ_K^0 (the invariant circle containing P_i and P_i', $i \in \mathcal{I}$) is completely foliated by rotationally symmetric invariant circles with rotation numbers continuously and monotonically varying from 0 on $\partial\mathcal{S}$ to $1/(4K)$ on Γ_K^0.*

Remark 5.13. One can check that $K = 3$ in the proof of Theorem 5.9 is obtained by setting $\theta = \frac{\pi}{8}$, corresponding to $t = (2 - \sqrt{2})/4$ and $s = \sqrt{2}/4$. In this case C is the rotation by $\frac{\pi}{4}$ on C. For $K \geq 4$, exact values for θ are less easy to determine explicitly.

In the special case $\theta = \frac{\pi}{8}$, the map F turns out to be of a very simple form, allowing a complete description of the dynamics on all of $\mathcal{S}$. By a similar argument applied to the region inside Γ_3^0 (containing the rotational part C), the statement of Corollary 5.12 can then be strengthened, stating that in this case the whole space $\mathcal{S}$ is foliated by invariant circles, with rotation numbers varying continuously and monotonically from 0 on $\partial\mathcal{S}$ to $\frac{1}{8}$ on the invariant octagon O inscribed in C. The invariant circles between Γ_3^0 and O each consist of twelve straight line segments, parallel to the twelve segments of Γ_3^0, see Figure 5.4.

5.5 General parameter values and discussion

For other parameter values than the ones considered in the previous section, we generally cannot prove the existence of any invariant circles. Let us briefly mention more general 'higher order' cancellation orbits and invariant circles.

Recall that in Theorem 5.9, we constructed a family of invariant circles consisting of line segments connecting successive orbit points of periodic cancellation orbits. The chosen cancellation orbits were of the simplest possible kind, where a point on any break line is mapped by a certain number of iterations to the next possible break line.

It is possible to construct invariant circles from one or several more complicated periodic cancellation orbits[1] (for other values of θ than the ones in Theorem 5.9). On such periodic cancellation orbit, the iterates of a point on a break line would cross several break

[1]Due to the rotational symmetry of the system, a periodic cancellation orbit could contain one, two or four pairs of break points. In the first two cases, the union of all rotated copies of the cancellation orbit would need to be considered, to form an invariant circle by adding straight line segments between successive points in this union.

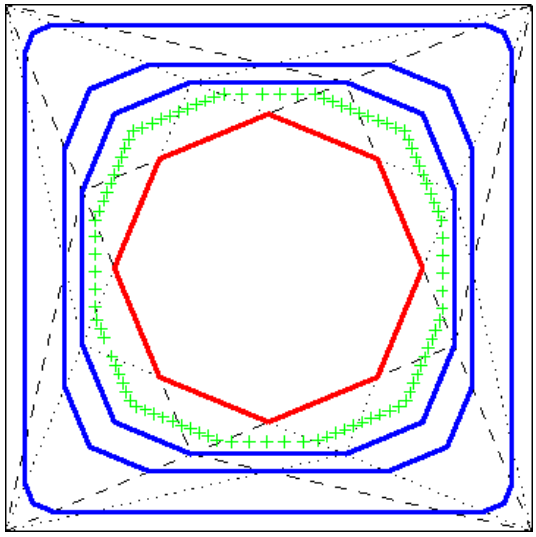

Figure 5.4: (The case $\theta = \theta_3 = \frac{\pi}{8}$ from Remark 5.13) The solid piecewise straight line circles are the invariant octagon O inscribed in C, as well as the invariant circles Γ_3^0, Γ_3^1 and Γ_3^{10}. In the annulus between O and Γ_3^0, the first 100 iterates of an orbit are depicted, all lying on an invariant circle which consists of twelve straight line segments.

lines before hitting one. The resulting invariant circle would still consist of straight line segments, but in such a way that a given segment and its image under F are not adjacent on the circle.

Doing this 'higher order' construction is conceptually not much more difficult, but certainly more tedious than the 'first order' construction of Theorem 5.9. The effort to construct even just a single higher order periodic cancellation orbit seems futile, unless a more general scheme to construct all or many of them at once can be found.

It is unclear whether invariant circles of this type exist for all θ, and whether for typical θ there exist piecewise line segment invariant circles (of rational or irrational rotation number) containing non-periodic cancellation orbits, as the ones seen in Corollary 5.12.

Question 5.14. *For which parameter values $\theta \in (0, \frac{\pi}{4})$ does F_θ have invariant circles with periodic cancellation orbits (and therefore rational rotation number)? For which θ are there piecewise line segment invariant circles with non-periodic cancellation orbits?*

Moreover, we can at this point not rule out the existence of invariant circles (outside of C) of an entirely different kind, not consisting of straight line segments. Indeed, some numerical experiments seem to indicate the occurrence of invariant regions with smooth boundaries, but it is unclear whether this is due to the limited resolution (see, for example, left pictures in Figures 5.5 and 5.6).

By Proposition 5.6, the intersections of any invariant circle of rational rotation num-

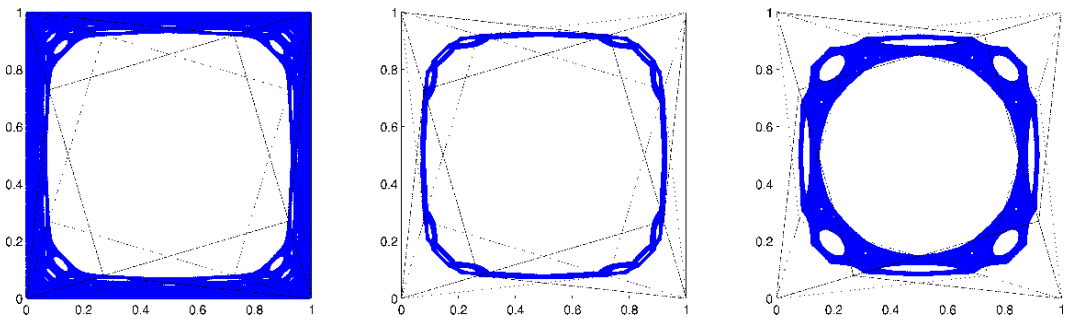

Figure 5.5: The first $5 \cdot 10^5$ iterates of three initial points for $\theta = \frac{\pi}{11}$. Each of the orbits seems to be confined to an invariant annulus and fill this annulus densely, except for a number of elliptic islands. The left picture seems to show an invariant region bounded by a *smooth* invariant circle. However, this might be an actual piecewise line segment curve, seemingly smooth because of the limited resolution and orbit length.

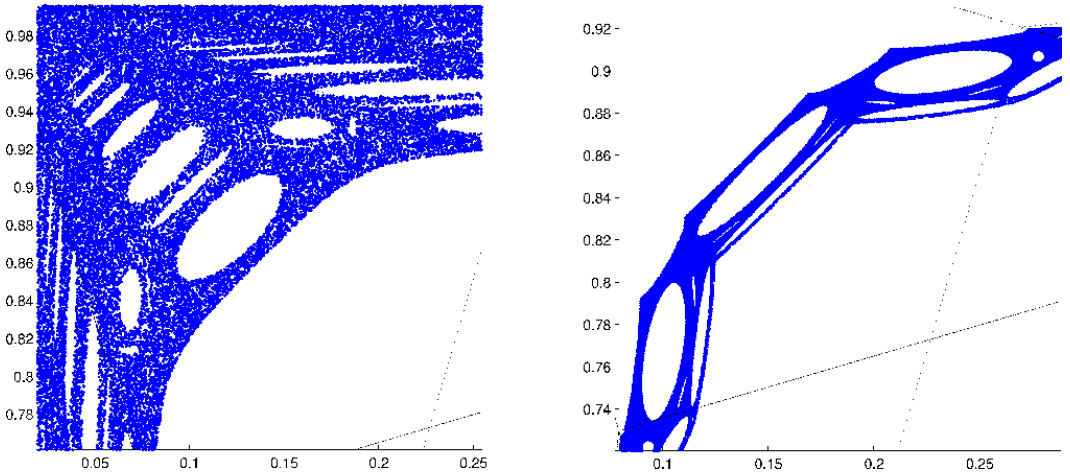

Figure 5.6: Zoom-in of the first two orbits from Figure 5.5, showing that the invariant annuli contain large numbers of smaller and smaller elliptic island chains. Further zoom-in (and longer orbits) reveal increasingly intricate patterns of such quasi-periodic elliptic regions.

ber with the break lines form cancellation orbits. For irrational rotation number, this need not be the case. Note however, that under irrational rotation, the 'kink' introduced to an invariant circle at its intersection with a break line propagates densely to the entire circle, unless it is cancelled by eventually being mapped to another break line intersection. Hence, an invariant circle of irrational rotation number would either contain a cancellation orbit for each pair of break lines, or otherwise would be geometrically complicated, namely nowhere differentiable.

Question 5.15. *Are there invariant circles for F (outside of C) which are not comprised of a finite number of straight line segments? Are there invariant circles whose intersections with the break lines do not form cancellation orbits?*

As for the dynamics of F between invariant circles, we can only point to numerical evidence that annuli between consecutive invariant circles form ergodic components interspersed with 'elliptic islands'. An elliptic island consists of a periodic point of, say, period p, surrounded by a family of ellipses which are invariant under F^p, such that F^p acts as an irrational rotation on each of these ellipses; this is referred to as 'quasi-periodic' behaviour. In the case when F^p is a rational rotation, these quasi-periodic invariant circles would take on the shape of polygons, consisting entirely of p-periodic points. The rest of the annuli seems to be filled with what is often referred to as 'stochastic sea', that is, the dynamics seems to be ergodic and typical orbits seem to fill these regions densely.

As in many similar systems (for instance, perturbations of the standard map), numerical observations seem to indicate that these 'ergodic regions' have positive Lebesgue measure (see Figures 5.5-5.9). This is related to questions surrounding the famous 'quasi-ergodic hypothesis', going back to Ehrenfest [27] and Birkhoff [16], conjecturing that typical Hamiltonian dynamical systems have dense orbits on typical energy surfaces (see also [47]).

For a piecewise linear version of the standard map, Bullett [25] established a number of results on cancellation orbits and invariant circles of both rational and irrational rotation numbers. Wojtkowski [97, 98] showed that the map is almost hyperbolic and deduced that it has ergodic components of positive Lebesgue measure. Almost hyperbolicity here means the almost everywhere existence of invariant foliations of the space (or an invariant subset) by transversal local contracting and expanding fibres. Equivalently, this can be expressed through the existence of invariant contracting and expanding cone fields. By classical theory (see, for example, [89]) almost hyperbolicity implies certain mixing properties of the map, and in particular, the existence of an at most countable collection of ergodic components of positive Lebesgue measure.

A similar kind of cone construction as in [97, 98] seems to be more difficult for the map studied here. One important difference is that the piecewise linear standard map

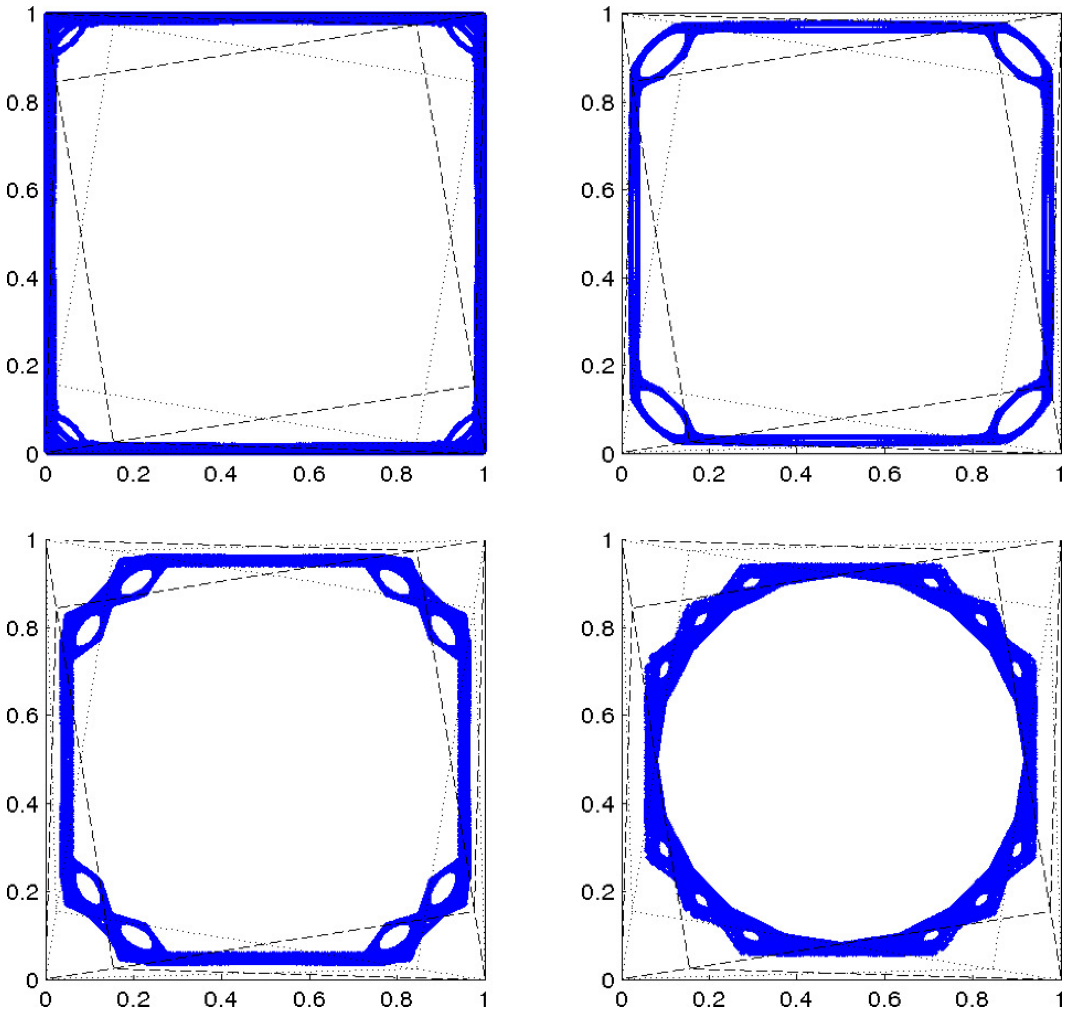

Figure 5.7: The first 10^5 iterates of four initial points for $\theta = \frac{\pi}{20}$. Each seems to be a dense orbit in an invariant annulus. The four annuli (without a number of elliptic islands and the invariant periodic regular 20-gon inscribed in C) seem to be the ergodic components of F.

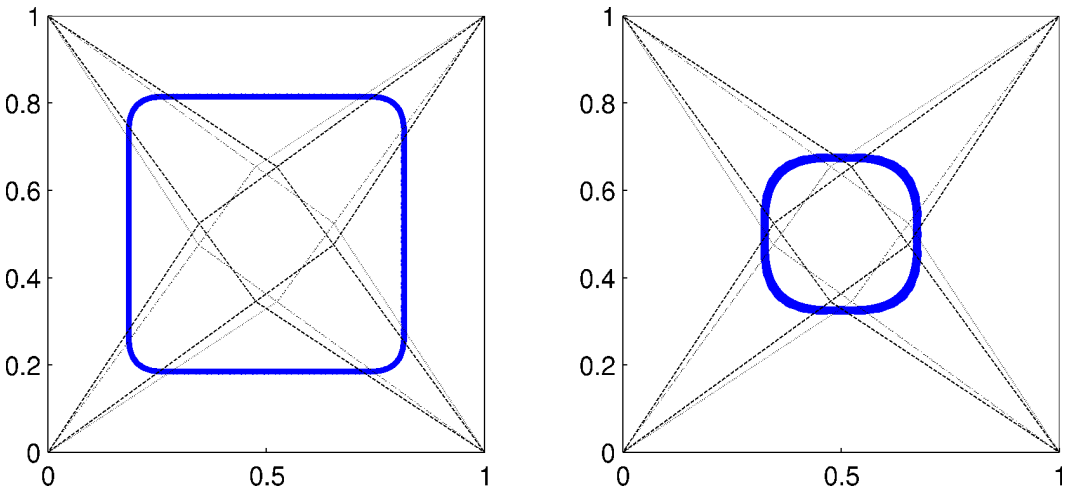

Figure 5.8: The first 10^5 iterates of two initial points for $\theta = \frac{\pi}{5}$. Each of the orbits seems to densely fill a thin invariant annulus. The rectangle seems to be partitioned into finitely many such invariant annuli (which get thinner and more numerous as $\theta \to 0$).

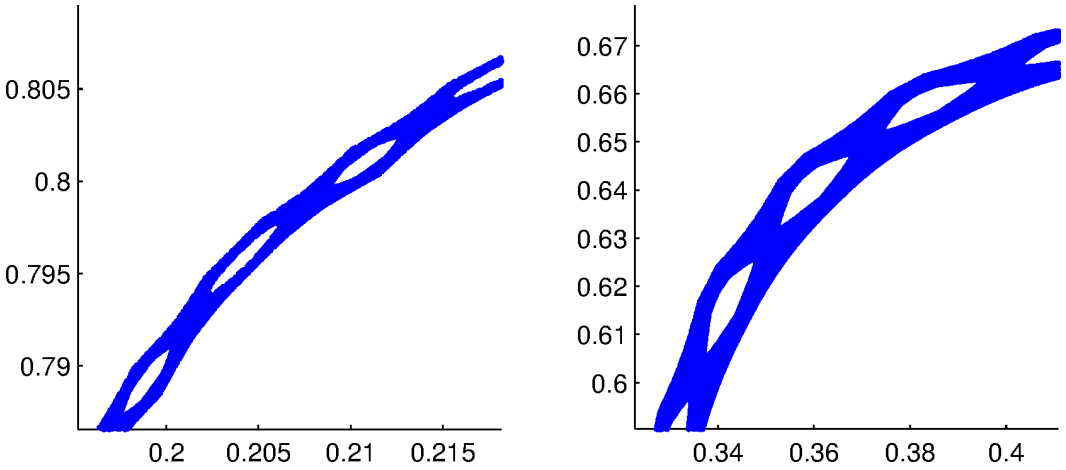

Figure 5.9: Zoom-in of the orbits from Figure 5.8. The thin invariant annuli contain periodic island chains, which under even stronger magnification could be seen to be surrounded by further, finer, islands of quasi-periodic motion.

in these papers is a twist map, which is not strictly the case for the map F studied here (see [59] and references therein for an overview over the numerous classical results for twist maps, mostly based on Birkhoff and Aubry-Mather theory). The additional property that F equals the identity on the boundary of the square also sets it apart. In particular, the motion of points under F close to the boundary can be arbitrarily slow, that is, take arbitrarily many iterations to pass through the piece $\mathcal{A}_i$ (while the number of iterations for a passage through $\mathcal{B}_i$ remains bounded). This seems to make it more difficult to explicitly construct an invariant contracting or expanding cone field, as was done for the piecewise linear standard map. Moreover, such an invariant cone field construction can not be carried out uniformly for all θ. In fact, as can be seen from Corollary 5.12, there are parameter values θ_K, $K = 3, 4, \ldots$, for which almost hyperbolicity cannot hold on large parts of $\mathcal{S}$, as the dynamics is completely integrable on the annulus between the invariant circle Γ_K^0 and $\partial\mathcal{S}$ (and even on all of $\mathcal{S}$ for $\theta = \theta_3 = \frac{\pi}{8}$, see Remark 5.13).

We are led to leave the following as a question.

Question 5.16. *Are there parameter values θ for which the map F_θ is almost hyperbolic on some invariant subset of $\mathcal{S}$? How large is the set of parameters θ for which this is the case?*

While it does not seem likely that almost hyperbolicity can be shown for almost all θ, numerical evidence suggests that for typical θ, the map F_θ has a finite number of ergodic components of positive Lebesgue measure, in which typical points have dense orbits.

Conjecture 5.17. *For Lebesgue almost all $\theta \in (0, \frac{\pi}{4})$, there is a finite number of F-invariant sets $A_1, \ldots, A_m$, each of positive Lebesgue measure, such that $F|_{A_i} \colon A_i \to A_i$ is ergodic for every $i = 1, \ldots, m$. Each A_i is a topological annulus with a certain number of elliptic islands removed from it, and, together with the elliptic islands and the invariant disk inscribed in C, the A_i form a partition of $\mathcal{S}$.*

Numerical experiments also seem to indicate that for many parameter values, the way in which chaotic and (quasi-)periodic behaviour coexist, that is, the structure of invariant annuli containing families of quasi-periodic elliptic islands, can be quite rich, see Figures 5.5 and 5.6. Besides the total measure of such quasi-periodic elliptic islands, it would be also interesting to know whether a general scheme for their itineraries, periods and rotation numbers can be found.

Bibliography

[1] R. L. Adler, B. Kitchens, and C. Tresser. Dynamics of non-ergodic piecewise affine maps of the torus. *Ergod. Th. & Dynam. Sys.*, 21(04):959–999, 2001.

[2] D. Aharonov, R. L. Devaney, and U. Elias. The dynamics of a piecewise linear map and its smooth approximation. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 7(2):351–372, 1997.

[3] P. Ashwin and A. Goetz. Polygonal invariant curves for a planar piecewise isometry. *Trans. Amer. Math. Soc.*, 358(1):373–390, 2005.

[4] J.-P. Aubin and A. Cellina. *Differential Inclusions*. Springer, Berlin, 1984.

[5] R. J. Aumann. Subjectivity and correlation in randomized strategies. *J. Math. Econom.*, 1:67–96, 1974.

[6] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.

[7] M. Barge and J. Franks. Recurrent sets for planar homeomorphisms. In: *From Topology to Computation: Proceedings of the Smalefest (Berkeley, CA, 1990)*, pages 186–195. Springer, New York, 1993.

[8] A. F. Beardon, S. R. Bullett, and P. J. Rippon. Periodic orbits of difference equations. *Proc. Roy. Soc. Edinburgh Sect. A*, 125(4):657–674, 1995.

[9] M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions. *SIAM J. Control Optim.*, 44(1):328, 2005.

[10] M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions, part II: applications. *Math. Oper. Res.*, 31(4):673–695, 2006.

[11] U. Berger. Fictitious play in $2 \times n$ games. *J. Econ. Theory*, 120(2):139–154, 2005.

[12] U. Berger. Brown's original fictitious play. *J. Econ. Theory*, 135(1):572–578, 2007.

[13] U. Berger. Two more classes of games with the continuous-time fictitious play property. *Game. Econ. Behav.*, 60(2):247–261, 2007.

[14] U. Berger. Learning in games with strategic complementarities revisited. *J. Econ. Theory*, 143(1):292–301, 2008.

[15] U. Berger. The convergence of fictitious play in games with strategic complementarities: a comment. Preprint, 2009, http://mpra.ub.uni-muenchen.de/id/eprint/20241.

[16] G. D. Birkhoff and B. O. Koopman. Recent contributions to the ergodic theory. *Proc. Natl. Acad. Sci. USA*, 18(3):279–282, 1932.

[17] D. Blackwell. Controlled random walks. In: *Proceedings of the International Congress of Mathematicians*, 1954, vol. III, pages 336–338.

[18] M. Bonino. A dynamical property for planar homeomorphisms and an application to the problem of canonical position around an isolated fixed point. *Topology*, 40(6):1241–1257, 2001.

[19] M. Bonino. Lefschetz index for orientation reversing planar homeomorphisms. *Proc. Amer. Math. Soc.*, 130(7):2173–2177, 2002.

[20] L. E. J. Brouwer. Beweis des ebenen Translationssatzes. *Math. Ann.*, 72(1):37–54, 1912.

[21] G. W. Brown. Some notes on computation of games solutions. Report P-78, The Rand Corporation, 1949.

[22] G. W. Brown. Iterative solution of games by fictitious play. In: *Activity Analysis of Production and Allocation*, Cowles Commission Monograph No. 13, pages 374–376. John Wiley & Sons, New York, 1951.

[23] M. Brown. A new proof of Brouwer's lemma on translation arcs. *Houston J. Math.*, 10(1):35–41, 1984.

[24] M. Brown. On the fixed point index of iterates of planar homeomorphisms. *Proc. Amer. Math. Soc.*, 108(4):1109–1114, 1990.

[25] S. R. Bullett. Invariant circles for the piecewise linear standard map. *Comm. Math. Phys.*, 107(2):241–262, 1986.

[26] R. L. Devaney. A piecewise linear model for the zones of instability of an area-preserving map. *Phys. D*, 10(3):387–393, 1984.

[27] P. Ehrenfest and T. Ehrenfest. *The Conceptual Foundations of the Statistical Approach in Mechanics*. Cornell University Press, Ithaca, 1959.

[28] A. Fathi. An orbit closing proof of Brouwer's lemma on translation arcs. *Enseign. Math.*, 33(3-4):315–322, 1987.

[29] D. P. Foster and R. V. Vohra. Calibrated learning and correlated equilibrium. *Game. Econ. Behav.*, 21(1-2):40–55, 1997.

[30] J. Franks. A new proof of the Brouwer plane translation theorem. *Ergod. Th. & Dynam. Sys.*, 12(02):217–226, 2008.

[31] D. Fudenberg and D. K. Levine. Consistency and cautious fictitious play. *J. Econ. Dyn. Control*, 19(5-7):1065–1089, 1995.

[32] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press Series on Economic Learning and Social Evolution, 2. MIT Press, Cambridge, MA, 1998.

[33] D. Fudenberg and D. K. Levine. Learning and equilibrium. *Annu. Rev. Econ.*, 1(1):385–420, 2009.

[34] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

[35] J.-M. Gambaudo, S. van Strien, and C. Tresser. Vers un ordre de Sarkovskiĭ pour les plongements du disque préservant l'orientation. *C. R. Math. Acad. Sci. Paris*, 310(5):291–294, 1990.

[36] A. Gaunersdorfer and J. Hofbauer. Fictitious play, Shapley polygons, and the replicator equation. *Game. Econ. Behav.*, 11(2):279–303, 1995.

[37] I. Gilboa and A. Matsui. Social stability and equilibrium. *Econometrica*, 59(3):859–867, 1991.

[38] S. Gjerstad. The rate of convergence of continuous fictitious play. *Econom. Theory*, 7:161–178, 1996.

[39] A. Goetz. Dynamics of piecewise isometries. *Illinois J. Math.*, 44(3):465–478, 2000.

[40] E. Gutkin and N. Haydn. Topological entropy of generalized polygon exchanges. *Bull. Amer. Math. Soc.*, 32(1):50–57, 1995.

[41] S. Hahn. The convergence of fictitious play in 3×3 games with strategic complementarities. *Econom. Lett.*, 64(1):57–60, 1999.

[42] J. Hannan. Approximation to Bayes risk in repeated play. In: *Contributions to the Theory of Games*, vol. 3 (Annals of Mathematics Studies No. 39), pages 97–139, Princeton University Press, Princeton, NJ, 1957.

[43] C. Harris. On the rate of convergence of continuous-time fictitious play. *Game. Econ. Behav.*, 22(2):238–259, 1998.

[44] S. Hart. Adaptive heuristics. *Econometrica*, 73(5):1401–1430, 2005.

[45] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[46] S. Hart and A. Mas-Colell. A general class of adaptive strategies. *J. Econ. Theory*, 98(1):26–54, 2001.

[47] M. Herman. Some open problems in dynamical systems. In: *Proceedings of the International Congress of Mathematicians*, vol. II, Berlin, 1998.

[48] J. Hofbauer. Stability for the best response dynamics. Preprint, 1995, http://homepage.univie.ac.at/josef.hofbauer/br.ps.

[49] J. Hofbauer. From Nash and Brown to Maynard Smith: equilibria, dynamics and ESS. *Selection*, 1(1):81–88, 2001.

[50] J. Hofbauer and K. Sigmund. *The Theory of Evolution and Dynamical Systems: Mathematical Aspects of Selection (London Mathematical Society Student Texts, 7)*. Cambridge University Press, Cambridge, 1988.

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

[52] J. Hofbauer and S. Sorin. Best response dynamics for continuous zerosum games. *Discrete Contin. Dyn. Syst. Ser. B*, 6(1):215–224, 2005.

[53] S. Karlin. *Mathematical Methods and Theory in Games, Programming and Economics, Vols. I and II*. Addison-Wesley Publishing Co., Inc., Reading, MA, 1959.

[54] A. Katok. Periodic and quasiperiodic orbits for twist maps. In: *Dynamical Systems and Chaos (Lecture Notes in Physics, 179)*, pages 47–65. Springer, Berlin, 1983.

[55] B. Kerékjártó. *Vorlesungen über Topologie*. Springer, Berlin, 1923.

[56] B. Kolev. Periodic orbits of period 3 in the disc. *Nonlinearity*, 7(3):1067–1071, 1994.

[57] V. Krishna and T. Sjöström. On the convergence of fictitious play. *Math. Oper. Res.*, 23(2):479 – 511, 1998.

[58] J. C. Lagarias and E. Rains. Dynamics of a family of piecewise-linear area-preserving plane maps II. Invariant circles. *J. Difference Equ. Appl.*, 11(13):1137–1163, 2005.

[59] P. Le Calvez. *Dynamical Properties of Diffeomorphisms of the Annulus and of the Torus*. American Mathematical Society, Providence, RI, 2000.

[60] P. Le Calvez and J.-C. Yoccoz. Un théorème d'indice pour les homéomorphismes du plan au voisinage d'un point fixe. *Ann. of Math.*, 146(2):241–293, 1997.

[61] A. Matsui. Best response dynamics and socially stable strategies. *J. Econ. Theory*, 57(2):343–362, 1992.

[62] A. Metrick and B. Polak. Fictitious play in 2 x 2 games: a geometric proof of convergence. *Econom. Theory*, 4(6):923–933, 1994.

[63] J. Milnor. *Dynamics in One Complex Variable, third edition. (Annals of Mathematics Studies, 160)*. Princeton University Press, Princeton, NJ, 2006.

[64] K. Miyasawa. On the convergence of the learning process in a 2 x 2 non-zero-sum two-person game. *Econometric Research Program, Research Memorandum No. 33*, 1961.

[65] D. Monderer, D. Samet, and A. Sela. Belief affirming in learning processes. *J. Econ. Theory*, 73(2):438–452, 1997.

[66] D. Monderer and A. Sela. Fictitious play and no-cycling conditions. Technical report, Sonderforschungsbereich 504, University of Mannheim, 1997.

[67] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *J. Econ. Theory*, 68(1):258–265, 1996.

[68] D. Monderer and L. S. Shapley. Potential games. *Game. Econ. Behav.*, 14(1):124–143, 1996.

[69] O. Morgenstern and J. von Neumann. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

[70] S. Morris and T. Ui. Best response equivalence. *Game. Econ. Behav.*, 49(2):260–287, 2004.

[71] H. J. Moulin and J.-P. Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *Internat. J. Game Theory*, 7(3-4):201–221, 1978.

[72] J. F. Nash Jr. Non-cooperative games. *Ann. of Math.*, 54(2):286–295, 1951.

[73] G. Ostrovski. Dynamics of a continuous piecewise affine map of the square. Preprint, 2013, arXiv:1305.4282.

[74] G. Ostrovski and S. van Strien. Payoff performance of fictitious play. Preprint, 2013, arXiv:1308.4049.

[75] G. Ostrovski. Fixed point theorem for non-self maps of regions in the plane. *Topology Appl.*, 160(7):915–923, 2013.

[76] G. Ostrovski and S. van Strien. Piecewise linear Hamiltonian flows associated to zero-sum games: transition combinatorics and questions on ergodicity. *Regul. Chaotic Dyn.*, 16(1-2):129–154, 2011.

[77] L. G. Oversteegen and E. D. Tymchatyn. Extending isotopies of planar continua. *Ann. of Math.*, 172(3):2105–2133, 2010.

[78] L. G. Oversteegen and K. I. S. Valkenburg. Characterizing isotopic continua in the sphere. *Proc. Amer. Math. Soc.*, 139(04):1495–1495, 2011.

[79] F. Przytycki. Ergodicity of toral linked twist mappings. *Ann. Sci. Éc. Norm. Supér. (4)*, 16(3):345–354, 1983.

[80] T. E. S. Raghavan. Zero-sum two-person games. In: *Handbook of Game Theory with Economic Applications Vol. II (Handbooks in Economics Vol. 11)*, pages 735–768. North-Holland, Amsterdam, 1994.

[81] H. Reeve-Black and F. Vivaldi. Near-integrable behaviour in a family of discretized rotations. *Nonlinearity*, 26(5):1227–1270, 2013.

[82] J. Robinson. An iterative method of solving a game. *Ann. of Math.*, 54(2):296–301, 1951.

[83] J. Rosenmüller. Über Periodizitätseigenschaften spieltheoretischer Lernprozesse. *Z. Wahrscheinlichkeit.*, 17(4):259–308, 1971.

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

[85] A. Sela. Fictitious play in 2x3 games. *Game. Econ. Behav.*, 31(1):152–162, 2000.

[86] H. N. Shapiro. Note on a computation method in the theory of games. *Comm. Pure Appl. Math.*, 11:587–593, 1958.

[87] L. S. Shapley. Some topics in two-person games. *Advances in Game Theory*, 52:1–29, 1964.

[88] L. S. Shapley and R. N. Snow. Basic solutions of discrete games. In: *Contributions to the Theory of Games (Annals of Mathematics Studies No. 24)*, pages 27–35. Princeton University Press, Princeton, NJ, 1950.

[89] J. G. Sinai. Classical dynamic systems with countably-multiple Lebesgue spectrum. II. (Russian). *Izv. Akad. Nauk SSSR Ser. Mat.*, 30:15–68, 1966.

[90] C. Sparrow, S. van Strien, and C. Harris. Fictitious play in 3×3 games: the transition between periodic and chaotic behaviour. *Game. Econ. Behav.*, 63:259–291, 2008.

[91] K. Sydsaeter and P. Hammond. *Essential Mathematics for Economic Analysis, third edition*. Prentice Hall, 2008.

[92] S. van Strien. Hamiltonian flows with random-walk behaviour originating from zero-sum games and fictitious play. *Nonlinearity*, 24(6):1715–1742, 2011.

[93] S. van Strien and C. Sparrow. Fictitious play in 3×3 games: chaos and dithering behaviour. *Game. Econ. Behav.*, 73:262–286, 2011.

[94] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Math. Ann.*, 100(1):295–320, 1928.

[95] J. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1997.

[96] M. Wojtkowski. Linked twist mappings have the K-property. *Ann. N. Y. Acad. Sci.*, 357(1):65–76, 1980.

[97] M. Wojtkowski. A model problem with the coexistence of stochastic and integrable behaviour. *Comm. Math. Phys.*, 80(4):453–464, 1981.

[98] M. Wojtkowski. On the ergodic properties of piecewise linear perturbations of the twist map. *Ergod. Th. & Dynam. Sys.*, 2(3-4):525–542, 1982.

[99] H. P. Young. *Strategic Learning and Its Limits (Arne Ryde Memorial Lectures Series)*. Oxford University Press, 2005.