Rates of mixing in models of fluid devices with discontinuities

Hannah Elizabeth Kreczak

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

University of Leeds
Centre for Doctoral Training in Fluid Dynamics

September 2019

I confirm that the work submitted is my own, except where work which has formed part of jointly authored publications has been included. No material is included which has been accepted for the award of any other degree or diploma at any university or equivalent institution.

The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The work in Chapter 5 of the thesis has appeared in publication as follows:

Deceleration of one-dimensional mixing by discontinuous mappings, H. Kreczak, R. Sturman, and M. C. T. Wilson, Phys. Rev. E 96, 053112 (2017).

Conception of the model that is investigated, computational simulations, data analysis of the results and writing of the paper is attributable to myself.

Supervision of the research with suggestions on lines of enquiry, manuscript revision and suggested edits of the paper were carried out by my supervisor team Dr Rob Sturman and Dr Mark Wilson.

"That continued mixing results in uniformity, and that uniformity is only to be obtained by mixing, will be generally acknowledged, but how deeply and universally this enters into all the arts can but rarely have been apprehended." Osborne Reynolds, 1894

This thesis is dedicated to my parents, Stephen and Lynn.

Acknowledgements

First of all I would like to express my immense gratitude to my supervisors, Rob Sturman and Mark Wilson. I am thankful for the guidance they have both shown me throughout this supervisory period, not only relating to my academic studies, but also in my personal wellbeing. I am appreciative of Mark's engineering eye, keeping the research grounded in reality, and patience during the discussions which spiralled into abstract dynamical theory. I extend additional gratitude to Rob, for the mentorship during my attendance at Snowbird, and most importantly the initial encouragement to purse a PhD in the first place.

I would like to thank Jean-Luc Thiffeualt, and the extended network of academics at the University of Madison-Wisconsin, who warmly received me during my academic placement. I would also like to thank the teams at Northwestern University, and the Universities of Exeter and Manchester, for inviting me to speak about my work and engaging me in insightful discussions, especially those over a few pints.

I would like to thank EPSRC for providing the funding allowing me to undertake this research as part of the Leeds Centre of Doctoral training in Fluid Dynamics. I have benefited massively from the financial support, providing me with the opportunities to visit national and international conferences and academic institutions. I am particularly grateful to Claire Savy and the CDT management team, for arranging extra professional development to accompany my academic studies, preparing me for my future career.

On a personal note, I am extremely indebted to my friends and colleagues for the necessary distractions over the last five years. Special mentions to Ollie, Inna, Georgie, Jacob, Jenny, and Andy, for the lunch time chats and cryptic crosswords, laughs over after work wines, and encouragements during climbs. I would especially like to thank

my house mate Caitlin, for being both a fun and supportive friend over the years.

Last but not least, I would like to thank my parents and brothers for the unending support I have received during my education to date. Although the financial support my parents have given me has helped dispel worries that would have made it difficult to complete this thesis, it is their unwavering love for which I am most grateful, and which has allowed me to pursue my passions and achieve my many successes.

Abstract

In the simplest sense, mixing acts on an initially heterogeneous system, transforming it to a homogeneous state through the actions of stirring and diffusion. The theory of dynamical systems has been successful in improving understanding of underlying features in fluid mixing, and how smooth stirring fields, coherent structures and boundaries affect mixing rates. The main stirring mechanism in fluids at low Reynolds number is the stretching and folding of fluid elements, although this is not the only mechanism to achieve complicated dynamics.

Mixing by cutting and shuffling occurs in many situations, for example in micro–fluidic split and recombine flows, through the closing and re-orientation of values in sink–source flows, and within the bulk flow of granular material. The dynamics of this mixing mechanism are subtle and not well understood. Here, mixing rates arising from fundamental models capturing the essence of discontinuous, chaotic stirring with diffusion are investigated.

In purely cutting and shuffling flows it is found that the number of cuts introduced iteratively is the most important mechanism driving the approach to uniformity. A balance between cutting, shuffling and diffusion achieves a long-time exponential mixing rate, but similar mechanisms dominate the finite time mixing observed through the interaction of many slowly decaying eigenfunctions. The time to achieve a mixed condition varies polynomially with diffusivity rate κ, obeying $t \propto \kappa^{-\eta}$. For the transformations meeting good stirring criteria, $\eta < 1$. Considering the time to achieve a mixed condition to be governed by a balance between cutting, shuffling, and diffusion derives $\eta \sim 1/2$, which shows good agreement with numerical results.

In stirring fields which are predominantly chaotic and exponentially mixing, it is observed that the addition of discontinuous transformations contaminates mixing when the stretching rates are uniform, or

close to uniform. The contamination comes from an increase in scales of the concentration field by the reassembly of striations when cut and shuffled. Mixing stemming from this process is unpredictable, and the discontinuities destroy the possibility to approximate early mixing rates from stretching histories. A speed up in mixing rate can be achieved if the discontinuity aids particle transport into islands of the original transformation, or chops and rearranges large striations generated from highly non-uniform stretching. The long-time mixing rates and time to achieve a mixed condition are shown to behave counter-intuitively when varying the diffusivity rate. A deceleration of mixing with increasing diffusion coefficient is observed, sometimes overshooting analytically derived bounds.

Contents

List of Figures

Chapter 1

Introduction

Mixing is ubiquitous in a variety of natural, industrial and technological applications, therefore it is essential to understand the underlying mechanisms of the process. In the simplest sense, the process of mixing can be divided into two areas: stirring, which is the kinematic transport of two or more different materials resulting in decreased segregation, and diffusion, which evens out gradients of high and low concentrations on the molecular level, achieving equal distribution in the long time.

The key stirring mechanism in fluid mixing is the stretching and folding of fluid elements, where in high Reynolds number flows this is efficiently achieved via turbulent eddies, or can be attained in laminar flows through specially prescribed velocity fields with chaotic characteristics [Aref (1984)]. The limit of zero Reynolds number allows kinematic transport in fluid flows to be modelled as a dynamical system, for which there is a wealth of literature aiding the understanding of observed mixing phenomena arising from particular underlying advective dynamics. As well as highlighting which dynamical features result in efficient stirring [Ottino (1989); Wiggins (2013)], or those which hinder particle transport [Springham & Sturman (2014); Sturman & Springham (2013)], models of chaotic advection have unearthed the driving mechanisms for rates of mixing to uniformity from molecular diffusion [Antonsen Jr *et al.* (1996); Gouillart *et al.* (2008); Haynes & Vanneste (2005)].

It is important to distinguish the dynamical mechanisms contributing to stirring and mixing rates, since dominant characteristics leading to rapid or contaminated mixing may be easily identifiable in complex systems, which are difficult to model as a whole. For example, identifying coherent structures and basins of attraction in coarse grain models of the oceans informs the potential spread or
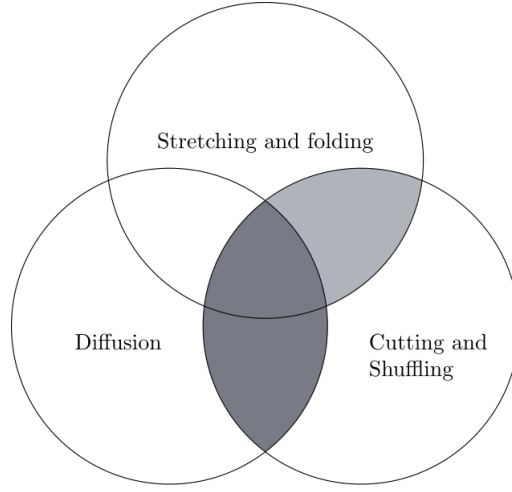
Figure 1.1: Venn Diagram of mixing mechanisms. The highlighted regions show emerging topics of interest in the literature. The dark grey regions indicate the mixing mechanisms of main interest in this thesis.

confinement of contaminants [Froyland *et al.* (2007)]. Additionally, knowledge of stirring mechanisms can enhance mixing protocol designs without building and testing many configurations. For example, boundaries have been shown to contaminate the approach to uniformity across the whole mixing domain, where slow stretching caused by no-slip boundaries is injected into the bulk flow [Gouillart *et al.* (2008); Sturman & Springham (2013)]. Rotating the walls and generating a homoclinic orbit causes a barrier from contamination in the bulk flow and re-achieves fast mixing rates [Gouillart *et al.* (2010); Thiffeault *et al.* (2011)]. Similar knowledge of dynamical systems can be utilised to increase understanding and innovation across a range of engineering and environmental mixing applications saving considerable resources. Designing models which capture dominant mechanisms of mixing, help isolate specific dynamical behaviours and aid in identifying their contribution to overall particle transport organisation and mixing rates.

While stretching and folding is the dominant stirring mechanism in most fluid flows, there are many stirring protocols where fluid material may also be cut, shifted and shuffled. Employing moving stirrers may be infeasible in certain applications, due to the energetic cost in highly laminar flows, or if the material being stirred consists of long, polymer chains which may be damaged in the process [Priye *et al.* (2013)]. Split-and-recombine micro-mixers are configured to cut and rearrange viscous fluid, increasing the number of striations and encouraging the approach to uniformity [Schönfeld *et al.* (2004)]. Likewise, chaotic trajecto-

ries can be implemented via pressure differences caused from the subtraction and re-injection of fluid using syringes [Cola (2004)]. Reorientation of the syringes creates time dependence, essential to establishing chaos, and discontinuities introduced on the closing and opening of valves. Similarly, granular material in tumblers possess both stirring mechanisms of stretching and folding, and cutting and shuffling [Sturman *et al.* (2008)]. Shearing occurs in a flowing layer at the surface, while cutting and shuffling stirs the bulk material, which is in solid body rotation with the tumbler walls. While discontinuous mixing arises in many systems, there is little understanding of the impact discontinuities have on fluid transport and the resulting homogenization in fluid mixing.

A fairly new area of research in dynamical systems is the study of piecewise isometries (PWIs), associated with transformations constructed entirely from translations, rotations and re-orientations [Goetz (1998, 2000)]. Although lacking exponential separation of near-by trajectories, the hallmark of chaos, PWIs can generate complex structures such as fractal coverings of coding partitions [Goetz (2003)]. Although PWIs do not contain the main mechanism of fluid stirring, stretching and folding, cutting and shuffling of finite pieces can decrease segregation and increase material interface. This was observed in PWI models of idealised granular tumbler flow with a vanishing shearing layer [Juarez *et al.* (2010)].

More recently, interest has grown in understanding particle transport organisation in combined systems of stretching and folding, and cutting and shuffling. Models investigating the cut-dominated limit in reorientated dipole flows with slip deformations identified novel possibilities for particle transport, not attainable in smoothly deforming Hamiltonian systems [Smith *et al.* (2016)]. When measuring the rate of decrease in the scale of segregation [Smith *et al.* (2017b)] it was observed that typically a decrease in the speed of stirring occurs from the presence of cutting and shuffling, compared to uniform stretching acting alone. This was also seen in uniformly expanding maps composed with permutations [Byott *et al.* (2013, 2016)]

However, the impact of discontinuities on the complete composition of mixing mechanisms has not been investigated, since in these studies the approach to uniformity from the action of diffusion has been neglected. A handful of studies have considered mixing rates arising in one-dimensional PWIs with a diffusive step, and find that discontinuous transformations mix faster than pure diffusion

alone [Ashwin *et al.* (2002)], and universality in mixing rates emerges under time-rescaling [Wang & Christov (2018)].

However, none of these investigations attempt to quantify the effect different stirring parameters have on the rates of mixing. Additionally, completely overlooked in the literature is the combined mechanisms of stretching and folding, cutting and shuffling, and diffusion. Figure 1.1 illustrates a Venn diagram of these three, distinct mechanism of mixing, and highlights the overlaps which have received little or no attention. The dark grey regions indicate the overlaps which are the focus of this thesis.

The aim of this thesis is to illuminate the impact discontinuous particle transport has on the rates of mixing of passive scalar fields in their approach to homogenization. Simplified models capturing the essence of the three mixing mechanisms are devised, in which the stirring dynamics are wholly comprehensive, aiding understanding on the mixing dynamics observed on the inclusion of diffusion. The models take the form of one and two-dimensional discrete time maps, which present themselves as ideal models to address the aim since they possess the following advantages.

1. Dynamical maps are a populated area of research, with many analytical and computational studies to draw upon.

2. The devised maps can be reformulated in terms of their transfer operators, which describe the evolution of typical densities and concentration fields. This allows comment on the typical mixing behaviours which would be expected.

3. Derivation of transfer operators allows diffusion effects to be captured via numerically accurate computational methods.

The points (1-3) are employed in each results chapter, which individually investigate the rates of mixing arising from a specific composition of discontinuous stirring dynamics with diffusion.

Chapter 2 reviews the relevant literature and introduces the computational methodologies used throughout the thesis. In Chapter 3 there is a short presentation on mixing rates from a range of dynamics in smooth stirring velocity fields with diffusion. This is presented to allow a direct comparison with the results reported for discontinuously stirred velocity fields. The main results chapters

begin with Chapter 4, which investigates the rates of mixing in cutting and shuffling transformations with a diffusive step. The effect of stirring parameters is discussed, and a scaling law between the time to achieve a mixed condition and the rate of diffusivity found. Next, advection governed by the composition of uniform stretching with cutting, shuffling and diffusion in explored in Chapter 5. A deceleration of mixing rates is observed, both via the addition of discontinuous transformations, and counter intuitively, increasing the rate of diffusion. The latter result is informed by analytical mixing rates for the diffusion-less limit. Whether discontinuous stirring protocols can be used to speed up, or hinder mixing rates in non-uniform hyperbolic maps is discussed in Chapter 6. A direct comparison to mixing phenomena arising from non-uniform stretching rates increases understanding of the driving mechanisms governing mixing is discontinuous transformations. In Chapters 4, 5, and 6, the models are one-dimensional, therefore an extension investigating mixing in two-dimensional discontinuous models is introduced in Chapter 7. It is revealed that the driving mechanisms of mixing observed in the one-dimensional models operate in higher dimensions. The key conclusions of the thesis are summarised in Chapter 8, with a discussion of future research directions.

Chapter 2

Literature Review

Mixing occurs in a variety of natural, industrial and technological applications: from something as simple as stirring milk into tea, to the transfer of heat through the Earth's atmosphere via jet streams. In addition to fluid mixing, mixing can occur in particle laden flows, such as granular material in tumblers. In a tumbler, granular material will rotate in solid body rotation with the tumbler walls while a flowing layer exhibits shearing, intrinsic to fluid mixing. The diversity of applications means that there are many approaches which are used to study such flows, however we can ask questions which are central to all problems to help quantify and predict the amount of mixing that occurs in each scenario.

Dynamical systems captures the kinematic behaviour of mixing, describing fluid particle transport from the action of stirring and has uncovered fundamental mechanisms which control rates of mixing. However, stirring is unlikely to occur in isolation and processes such as molecular diffusion, chemical reactions, and particle-particle interaction, assist homogenization and uniformity in the long time. Understanding the evolution of typical initial concentration fields under the combination of these processes will lead to a more comprehensive understanding of mixing.

In this chapter literature is reviewed with covers a large breadth of methods used to investigate the driving mechanisms behind rates of mixing in models of fluid and granular mixing. In Section 2.1, mixing is described in the simplest sense, as a process leading to homogenization in the long-time limit from the actions of advection and diffusion. Non-dimensional scaling numbers are introduced characterising fluid flow properties. In Section 2.2 the concept of chaotic advection is introduced, a term originated by Aref (1984), who discovered efficient stirring could be generated in unsteady laminar flows. By analysing the

velocity field as a dynamical system with chaotic trajectories, the relevant theory is introduced to determine if a chaotic system has good mixing properties via the ergodic hierarchy. A reduction of continuous time flows to discrete time maps is a simplification which captures time periodic mixing protocols.

The limitations of including diffusion using standard computational approaches is discussed in Section 2.3, accompanied by dynamical methods to include diffusion via discrete time maps. The relevant theory of transfer operators to predict typical asymptotic mixing rates is introduced. In Section 2.4 methods to measure mixing when there is stirring alone, alongside standard functional norms used in models with molecular diffusion are introduced.

Section 2.5 reviews key results from studies covering a combination of chaotic advection and diffusion. The impact of non-uniformity in the stretching rates of the flow field, non-mixing regions, and the effect of parabolic points due to boundaries are discussed. Section 2.6 introduces discontinuous transformations and mixing applications in which discontinuities arise. Analytical results from piecewise isometric transformations are given, which do not include material deformation by shearing, a hallmark of chaotic advection. A discussion on the handful of studies concerned with pair-wise combinations of the mixing mechanisms, stretching and folding, cutting and shuffling, and molecular diffusion, concludes the review.

2.1 Mixing and Applications

2.1.1 Advection-diffusion equation

Mixing is a mechanical or natural processes acting on an initially heterogeneous system, transforming it to a homogeneous state. A requirement for homegenization are mechanisms such as molecular diffusion or chemical reaction, encouraged by the kinematic action of stirring from an applied or incidental velocity field. Ignoring chemical reactions and additional effects on the velocity, such as heat transport, the simplest way to capture mixing of a passive concentration field is via the advection-diffusion equation

$$\frac{\partial c}{\partial t} + \boldsymbol{u}(\boldsymbol{x}, t) \cdot \nabla c(\boldsymbol{x}, t) = \kappa \nabla^2 c(\boldsymbol{x}, t). \tag{2.1}$$

The constant κ represents the diffusion coefficient and $\boldsymbol{u}$ the underlying velocity field. If the flow is incompressible then the velocity field $\boldsymbol{u}$ satisfies $\nabla \cdot \boldsymbol{u} = 0$.

Throughout this work pure advection ($\kappa = 0$) will be referred to as stirring, while the term mixing is reserved for the process of homogenization from the dual actions of stirring and diffusion ($\kappa \neq 0$).

Defining the average $\langle c \rangle$ as the average of $c(\boldsymbol{x})$ over a fixed domain of interest Ω,

$$\langle c \rangle = \frac{1}{\Omega} \int_\Omega c(\boldsymbol{x}) \mathrm{d}\Omega \tag{2.2}$$

then if the domain Ω has periodic boundary conditions, it can be shown from (2.1) that

$$\partial_t \langle c \rangle = 0, \tag{2.3}$$

and average of c is conserved. The variance, defined $\psi = \langle c^2 \rangle - \langle c \rangle^2$, evolves in time as

$$\partial_t \psi = -2\kappa \langle |\nabla c|^2 \rangle, \tag{2.4}$$

which implies $\psi \to 0$ when κ is non-zero. Note that the velocity field $\boldsymbol{u}$ is no longer represented, instead the gradients $|\nabla c|$ of the concentration field govern decrease in ψ. However it is precisely the velocity field which generates these necessary gradients for faster decay. Shearing in a velocity field filaments initially concentrated regions, with continued stirring reducing the lengths scales across widths of these filaments, increasing the surface area of material interface, and consequently increasing gradients. A good outline to these ideas is given by Thiffeault (2008).

Therefore the primary interest when wanting to achieve good mixing is to understand how certain flow features of an underlying velocity field lead to efficient growth in gradients in a concentration field. Eliminating barriers to mixing or optimizing mixing with fixed energetic cost is the motivation behind many studies. However, idealised systems can be devised where specific kinematic dynamics of the velocity field are isolated with the potential for analytic results to inform observations of mixing rates. The applicability of results to a wide range of applications can be quantified by dimensionless numbers capturing the ratios of fluid properties.

2.1.2 Scaling numbers

Flow dynamics and their effect on mixing can be generalised across a range of applications, comparing length and time scales with fluid properties via dimensionless ratios. These ratios characterise the importance of competing mechanisms

from advection and diffusion. Two important dimensionless numbers to consider when interested in (2.1) are the Reynolds number and the Péclet number.

The Reynolds number (Re) is a measure on inertial forces to viscous forces, defined

$$Re = \frac{UL}{\nu},\tag{2.5}$$

with U the typical velocity scale, L a characteristic length scale and ν the kinematic viscosity. Low Re number implies a dominance in viscous forces, a laminar regime, while large Re number corresponds to turbulent flows. The limit of zero Reynolds number is called Stokes flow, which demonstrates time-reversibility.

The Péclet number (Pe) is the ratio of advective transport rate to diffusive transport rate, defined

$$Pe = \frac{UL}{\kappa},\tag{2.6}$$

with U the typical velocity scale, L a characteristic length scale and κ diffusivity coefficient. In idealised models with normalised length and velocity scales it is conventional to take $\kappa = 1/Pe$. For a fixed advective rate, a larger Pe number implies the effect of diffusivity is weak. Therefore for large Pe, (2.4) can only decay quickly if large gradients are introduced into the concentration field. Vortical eddies are inherent to turbulent flow and are extremely efficient in shearing fluid, inducing large gradients which ensure fast mixing. In laminar flow, specialised velocity fields need to be applied for effective mixing, discussed with further detail in Section 2.2.

Flows with low Re are common across mixing applications. On geophysical scales, heat transfer in the Earth's mantle has length scales on the order of 10^6 metres but slow flow speeds, resulting in low Reynolds numbers ($Re \sim 10^{-20}$) [Ottino (1989)]. In the context of pharmaceutical, food, polymer and biochemical manufacturing, mixing is frequently carried out at low speeds or involves highly viscous fluids [Ottino (1990)].

Micro-fluidic devices are employed in typical mixing tasks, such as gas absorption, emulsification, foaming, and blending, or specialised applications such as sample preparation in biomedical and chemical analysis [Hessel *et al.* (2005)]. They operate at length scales as small as 10^{-6} metres to macroscales on the order of centimetres, with times scales on the order of milliseconds to seconds [Sturman *et al.* (2006)]. For example, water has a kinematic viscosity of $\nu \sim 0.01cm^2/s$, so at these length and times scales, flows are typically in a laminar regime.

Similarly, in many of these applications large Péclet numbers are plausible. Molecular dyes in water or glycerol solutions in micro-fluidic devices have Pe in

the range 10^3 to 10^5 [Stroock *et al.* (2002)]. Results from experimental granular flow suggest analogies can be drawn between mechanical and thermal behaviour for both laminar and turbulent flows of simple fluids [Wang & Campbell (1992)]. Macro-scale granular flows in rotating tumblers have $Pe \sim 10^5$ [Christov *et al.* (2011)]. In turbulent reactive flows, Pe can be as high as $Pe \sim 10^{10}$ [Muzzio & Liu (1996)], meaning a slow approach to uniformity is possible even in large Reynolds number flows. Therefore, advancing expertise in implementing good stirring protocols in laminar, high Péclet number flows is essential.

2.2 Fluid transport as a dynamical system

One of the main questions when studying fluid mixing is whether the mixing performance can be determined given a prescribed velocity field. The basic mechanisms for stirring were first understood and advanced by Reynolds (1894), who noted that stretching and folding (SF) were key elements in the deformation of fluid parcels. Others also highlighted this important mechanism in mixing [Eckart (1948); Welander (1955)]. Turbulent flow fields are by definition unsteady and contain greatly varying velocity fields in which fluid shearing causes deformation and stretching, leading to efficient stirring through the exponential lengthening of material lines [Batchelor (1952)]. In laminar, low Reynolds number flow, turbulent eddy structures are not present and specially designed stirring protocols need to be implemented to ensure good mixing occurs.

Flow behaviour in low Reynolds number regimes can be well approximated by Stokes flow, a considerable simplification of the Navier-Stokes equations governing the motion of fluids. When varying the Reynolds number in experiments of rod stirring flow and comparing with numerical Stokes flow approximations, Clifford *et al.* (2004) concluded that the zero Reynolds number model was a good indicator for the evolution of material lines up to $Re \sim 3$. The advantages of using such an approximation are outlined in this section.

2.2.1 Chaotic advection

The trajectories for non-inertial particles are simply the integration along the underlying velocity field

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{u}. \tag{2.7}$$

This collection of ordinary differential equations forms a dynamical system capturing the fluid transport. For a point initially located at $\boldsymbol{x}_0$, its trajectory after a time t is the solution to the integration of (2.7), and can be denoted as $\boldsymbol{x}(t) = \Phi_t(\boldsymbol{x}_0)$, $\Phi_t : \mathbb{R}^d \to \mathbb{R}^d$, known as the flow.

A seminal investigation by Aref (1984) found that in a model of Stokes flow as a dynamical system, pulsed, counter rotating vortices generated complex dye deformations after just a handful of iterations. From this study, the term *chaotic advection* was coined to describe the phenomenon, since the complex patterns were generated by exponential separation of nearby fluid particles, a hallmark of chaotic systems. Before Aref's study it was not known that such striation structures could be generated in laminar flow using a deterministic velocity field. Following this paper, the interest in studying chaotic advection, and dynamical systems in general, as a foundation for understanding mixing in low Reynolds number flows has grown across a number of multidisciplinary areas [Aref (2002)].

Many stirring protocols are periodic in time, allowing the assumption that $\boldsymbol{u}(\boldsymbol{x}, t + T) = \boldsymbol{u}(\boldsymbol{x}, t)$, where T is the time period. A flow with time periodicity is easily represented as an iterative map $M : \Omega \to \Omega$ acting along the streamlines from $t \to t + T$. The domain $\Omega \subset \mathbb{R}^d$ represents the full domain in which the mixing takes place. If the map M can be found, it is numerically easier to analyse than solving for the full time-integration along the streamlines of (2.7). The evolution of a given point $x \in \Omega$ by the mapping M is counted iteratively by iteration number $j \in \mathbb{Z}$ such that $x^{j+1} = M(x^j)$. If the flow represented by (2.7) is incompressible, $\nabla \cdot \boldsymbol{u} = 0$, then the map M is area-preserving as a consequence which implies that the determinant $|\det \boldsymbol{D}M| = \pm 1$, where $\boldsymbol{D}M$ is the Jacobian matrix

$$\boldsymbol{D}M = \begin{bmatrix} \dfrac{\partial M}{\partial x_1} & \cdots & \dfrac{\partial M}{\partial x_d} \end{bmatrix}, \tag{2.8}$$

defining the localised deformation by the map. Maps can be thought of as a snapshot of the full flow dynamics and exhibit chaotic behaviour similar to their chaotic flow counterparts.

Viewing the advection of fluid particles as a dynamical systems is advantageous since there is a vast amount of theory supporting the topic [Katok & Hasselblatt (1995); Ott (2002)]. For example, two-dimensional flows require time dependence to generate chaos by the Poincaré-Bendixson theorem. This is evidenced by the study of Aref (1984), where the vortices where required to alternate their rotation, to introduce time dependence and chaotic trajectories. In maps, only one-dimension is sufficient to generate chaos if specific criteria are met.

Fundamental features of dynamical systems are periodic points, where their stability determines how nearby trajectories evolve in time. In periodic flows of time period T, periodic points are defined as fixed points of $\Phi_{pT}(\boldsymbol{x}_f) = \boldsymbol{x}_f$, where p is the periodicity of the point. In a discrete map they are simply described by $M^p(\boldsymbol{x}_f) = \boldsymbol{x}_f$, fixed points of the pth iteration. The stability of periodic points are classified by the eigenvalues of the Jacobian determinant $\boldsymbol{D}\Phi$ (2.8) evaluated at the point $\boldsymbol{x} = \boldsymbol{x}_f$. Elliptic points have complex eigenvalues and surrounding points are stable and rotate about the periodic point, generating orbits which act as barriers to particle transport. Hyperbolic points are classified by real eigenvalues and are unstable, causing a contraction and expansion of nearby trajectories. Hyperbolic points are a hallmark of chaotic behaviour. Parabolic points on the other hand are characterised by degenerate eigenvalues, $\lambda_1 = \cdots = \lambda_d = \pm 1$, and as such are neither stable and rotating, nor unstable and expanding locally. The books by Ottino (1989) and Wiggins (2003) provide a comprehensive discussion on the application of dynamical systems, and the effect of periodic points in mixing.

Dye tracer experiments have been effective in determining if properties of dynamical systems predict mixing qualities of a prescribed flow. Macroscopic structures of fluorescent dye evolved in flow between two rotating, non-concentric cylinders qualitatively agreed with the position of periodic points, with stable and unstable manifolds indicating new barriers to mixing, providing an outline to striation arrangement [Chaiken et al. (1986); Swanson & Ottino (1990)]. Similar conclusions were drawn from cavity driven flows generated from pulsed sides [Leong & Ottino (1989)]. The dye trace experiments of [Jana et al. (1994) presented that in almost all cases a few or one of the lower-order hyperbolic periodic points contribute to striation arrangement, while other hyperbolic points with $p \geq 1$ add to the fine structure.

Other dynamical behaviours have been observed experimentally outside of comparisons with Stokes flow. Structures analogous to horseshoes, core examples of chaotic stretching and folding from a square domain into itself, were identified within flow visualisation of cavity flow [Chien et al. (1986); Ottino et al. (1994)]. Similarly, the study of a more topological approach, braid theory, has informed the best way to interweave multiple stirrers, and has been the topic of experimental [Boyland et al. (2000)], computational and analytic studies [Finn et al. (2003); Thiffeault et al. (2008)]. Mixing protocols inspired by braid theory do not require much tuning of stirrer parameters in comparison to one stirrer protocols.

2.2.2 Ergodic hierarchy

The additional benefit of studying stirring as a dynamical system is the potential for analytic results describing the transport and mixing of typical initial conditions via ergodic theory. Flows and maps can be constructed to capture the essence of mixing protocols, for example link-twist maps imitate the perpendicular shearing of egg beater flows [Sturman *et al.* (2006)]. To address whether a transformation is stirring a blob of dye evenly thought a domain, a precise and quantitative approach is required. The ergodic hierarchy provides a mathematical framework, as an ordering of behaviours of increasing complexity, capturing the kinematic properties essential to good stirring. To begin a discussion in ergodic theory, some basic and necessary terminologies on quantifying the size of sets and describing the behaviour of functions is required.

A transformation or map M requires a tuple $(\Omega, \mathcal{A}, M, \mu)$, consisting of a domain Ω, a σ-algebra $\mathcal{A}$ over Ω, the transformation itself M, and an invariant measure μ. A measure is a function which assigns a number to a given set, and can be thought of as size, volume or probability of the set. The measure μ is a non-negative, real-valued function defined on the σ-algebra, a collection of subsets on Ω. The transformations of interest in this current work are to represent incompressible, one or two-dimensional flows, therefore the measure μ will be taken as Lebesgue measure with the appropriate σ-algebra; the collection of half open intervals and products thereof. Lebesgue measure can be considered as the mathematical formalisation of Euclidean length, area or volume. The formalisation of incompressibility in ergodic theory is to describe a transformation representing an incompressible velocity field as a measure preserving transformation, defined such that for any set $A \subset \Omega$,

$$(M^{-1}(A)) = \mu(A) \text{ for all } A \subset \mathcal{A}. \tag{2.9}$$

The starting point to the ergodic hierarchy is the notion of in-decomposability of a domain, known as ergodicity. Within a domain of interest, elliptic island and invariant surfaces can generate barriers to mixing since they do not allow typical trajectories to traverse the entirety of the domain. Typical here refers to the trajectories of almost all initial starting positions, since many of the points will have some form of regularity in their behaviour. Typical trajectories of an ergodic transformation traverse the whole domain Ω, motivating the following definition in terms of sets.

Definition 2.2.1. (Ergodicity) *A measure-preserving dynamical system* $(\Omega, \mathcal{A}, M, \mu)$ *is ergodic if* $\mu(A) = 0$ *or* $\mu(A) = 1$ *for all* $A \in \mathcal{A}$ *such that* $M(A) = A$.

Sometimes, f is called an ergodic transformation, or that μ is an ergodic, invariant measure. The domain of an ergodic transformation can not be broken down into two or more sets of positive measure on which the transformation can be studied separately. Following from the definition, regular points in ergodic transformations will have a total volume within the domain equivalent to zero.

Ergodicity is not strong enough to determine if good mixing is possible. An irrational rotation on the unit circle is ergodic, but simply rotating a concentration field does not allow for a decrease in the scale of segregation. Instead, the definition of (strong) mixing is required to characterise such a desirable property of stirring.

Definition 2.2.2. ((Strong) Mixing*) A measure preserving, invertible transformation* $(\Omega, \mathcal{A}, M, \mu)$ *is (strong) mixing if for any two non-empty, measurable sets* A, B

$$\mu(M^j(A) \cap B) - \mu(A)\mu(B) \to 0 \ as \ j \to \infty. \tag{2.10}$$

This fundamentally says that in the infinite time limit, all arbitrary sets B of domain Ω contain the same proportion of the set A, as the proportion of A in Ω. For a non-invertible transformation, M^j is replaced with the backward time evolution of the map M^{-j} given in the definition. In many transformations strong mixing does not hold, and the notion of weak mixing applies instead.

Definition 2.2.3. (Weak mixing) *The measure-preserving transformation* M : $(\Omega, \mu) \to (\Omega, \mu)$ *is said to be weak mixing if for any two measurable sets* A, $B \subset \Omega$ *we have:*

$$\lim_{j \to \infty} \frac{1}{j} \sum_{k=0}^{n-1} |\mu(M^{-k}(A) \cap B) - \mu(A)\mu(B)| = 0. \tag{2.11}$$

Strong mixing implies weak mixing, which implies ergodicity, however the converse statements are not true. Of course, these definitions of strong and weak mixing are for infinite time, and a system may be mixing but with no understanding of "how quickly". In many applications, the *rate of decay of correlations* is used as a measure to quantify the rate of mixing. As a starting point, the definition of strong mixing can be reformulated in terms of a correlation function between arbitrary functions from a class of interest, such that

$$C_j(\phi, \psi) = \left| \int (\phi \cdot M^{-j})\psi d\mu - \int \phi d\mu \int \psi d\mu \right| \to 0 \ \text{as} \ j \to \infty, \tag{2.12}$$

for any pair of bounded measurable functions ϕ and ψ. For a given rate function $r(j)$, if there exists some constant C such that

$$C_j(\phi, \psi) \leq Cr(j), \tag{2.13}$$

then the rate of decay of correlations can be determined, and a mixing rate assigned to the transformation. Analytic derivation of the decay rate of correlations is not general and as such finding $r(j)$ via rigorous approaches is limited to a number of ergodic maps and the choice of observation functions ϕ and ψ.

2.3 Mixing via chaotic advection

While dynamical systems model well the kinematic properties of stirring in fluids, they can not wholly capture the mixing dynamics when diffusion is introduced. The advection-diffusion equation (2.1) is a partial differential equation in which there is a multitude of approaches to solve for initial value problems, particularly via numerical schemes such as finite difference methods. These computational methods employ either structured of variable grids to approximate the time evolution of desired functional field.

Problems arise when using finite sized grids to approximate mixing in turbulent or laminar flows driven by chaotic advection, since the exponential expansion of near by particles, and the consequential compression of striations from incompressibility, quickly generates variations in the concentration field with widths smaller than the grid cells. Refining a grid to accurately approximate the evolution of a concentration field in chaotic advective flows with large Péclet number would result in infeasible computational times to complete comprehensive parameter studies. For example, in the cavity flows of Leong & Ottino (1989), finite element and finite difference schemes were employed alongside a dynamical systems model, and were able to approximate the coarse structures of the flow but predicting the exact arrangement of striations were beyond the capabilities of the numerical schemes.

As was noted in Section 2.2.2, simple systems can be developed which capture the essence of stirring dynamics which are open to analytic investigation on stirring qualities and expected rates using ergodic theory. An extension to these ideas can be made to develop iterative advection-diffusion maps, which are open to analytic study or more efficient computational methods, to investigate the driving mechanisms behind mixing rates. These approaches are outlined in detail in this section.

2.3.1 Iterative advection-diffusion equation

Although (2.7) can be used to study fluid particle evolution it tells us little about the evolution of a passive scalar. If the velocity field is captured by a mapping M over the time period T, then the evolution of a concentration field c by this velocity field is also an iterative process $c(\boldsymbol{x}, t) \rightarrow c(\boldsymbol{x}, t + T)$. This is represented by a linear operator acting on the space of functions $P_M : \mathcal{X} \rightarrow \mathcal{X}$, $c \in \mathcal{X}$. This operator is known as the Frobenius-Perron operator in the dynamics literature, capturing the evolution of densities by a map M. However, the reformulation of stirring as a map means the integration of diffusional effects can not be incorporated for the full-time evolution in the interval $[t, t + T)$.

Similar to considering the advection as an iterative step, an idealisation of the advection-diffusion equation is to consider it iteratively via an operator splitting. First the concentration field is evolved according to (2.1) with $\kappa = 0$ for a time T. The advective step is then followed by a diffusive step in which (2.1) is solved for the time T with $\boldsymbol{u} = 0$. This is equivalent to applying a functional operator $P_D : \mathcal{X} \rightarrow \mathcal{X}$, with $P_D = \exp(t\kappa\nabla^2)$ for a time $t = T$. The operator for the full advective-diffusive step over time T is thus the composition of the operators $P = P_D \circ P_M$ and $c^{(j+1)}(x) = P(c^{(j)})(x)$. The operator P is linear due to the linearity of the advection-diffusion equation (2.1). When the maps used to investigate mixing are not descriptions of real velocity fields for some real periodic time interval T, the factor κT in the diffusive operator P_D can be treated with $T = 1$ normalised, such that the coefficient κ is needed only to describe the diffusional effects. This normalisation of the time-step T is taken throughout this work. The advantage of studying the evolution of a scalar field by analysing transfer operators acting on functional spaces is introduced in the next section.

2.3.2 Properties of Transfer Operators

Common approaches in studying dynamical systems is to study the long time behaviour of a specific trajectory or looking at local, Lagrangian dynamics near periodic points. However when discussing mixing we are generally concerned with how a function evolves in time, such as the concentration of particles, dyes, contaminants etc. This approach can be re-imagined as the evolution of *densities* by considering the statistical evolution of typical trajectories.

A measurable function $\phi : \Omega \rightarrow \mathcal{X}$, called an *observable*, will have a time series of functions $\{\phi \circ M^n\}_{k \geq 0}$, following a repeated action on points in Ω of

the mapping M. The composition $\phi \circ M^j$ represents the observation ϕ made at time j. For example, the concentration of particles could represent a particular observation. For a given map $M : \Omega \to \Omega$, its transfer operator P acting on a suitable Banach space of functions $\mathcal{X} \to \mathcal{X}$ is given by

$$(P\phi)(x) = \sum_{M(y)=x} \frac{\phi(y)}{|DM(y)|}, \quad \phi \in \mathcal{X}, \tag{2.14}$$

where $|DM(y)|$ is the Jacobian determinant evaluated at the point y. If M preserves an ergodic measure μ with density $\rho(x) \in M$, then $(P\rho)(x) = \rho(x)$ and ρ is called in invariant density of P.

For example, the dyadic transformation is defined as the map $M : I \to I$, $M(x) = 2x \mod 1$, with $I = [0,1]$. Under this mapping, each x has exactly two pre-images under M, so the transfer operator for functions ϕ is

$$P\phi(x) = \frac{1}{2}\left(\phi\left(\frac{x}{2}\right) + \phi\left(\frac{1}{2} + \frac{x}{2}\right)\right). \tag{2.15}$$

An invariant density of (2.15) is $\rho(x) = 1$ since

$$P\rho(x) = \frac{1}{2}(1+1) = 1,$$

which implies that a uniform density in the interval $I = [0,1]$ remains uniform under the dyadic transformation.

Recall that for a measure preserving system $(\Omega, \mathcal{A}, M, \mu)$ the rate of mixing can be quantified by the rate function $r(j)$ governing the decay of correlations (2.13). Provided functions are chosen from an appropriate functional space, the properties of the transfer operator give a tool for analysing the mixing properties of the system. For a complete Banach space B of functions the rate function may not be specified, i.e. there might exists $\phi \in \mathcal{B}$ for which $r(j)$ decays arbitrarily slowly. For example, expanding maps on the unit interval require $\mathcal{X}_S \in B$ is taken as function of bounded variation for the decay of correlations to decay exponentially [Byott *et al.* (2013)]. The operator restricted to the given functional space can be defined as $P_M|_{\mathcal{X}}$.

Provided a suitable space of functions is chosen, the spectrum of a transfer operator $Spec(P_M|_{\mathcal{X}})$ yields information on the asymptotic mixing rates of the transformation. The spectrum of the operator $P : \mathcal{X} \to \mathcal{X}$ is the set

$$\sigma(P) = \{\lambda \in \mathbb{C} | P - \lambda I \text{ is not an invertible operator on } \mathcal{X}\}. \tag{2.16}$$
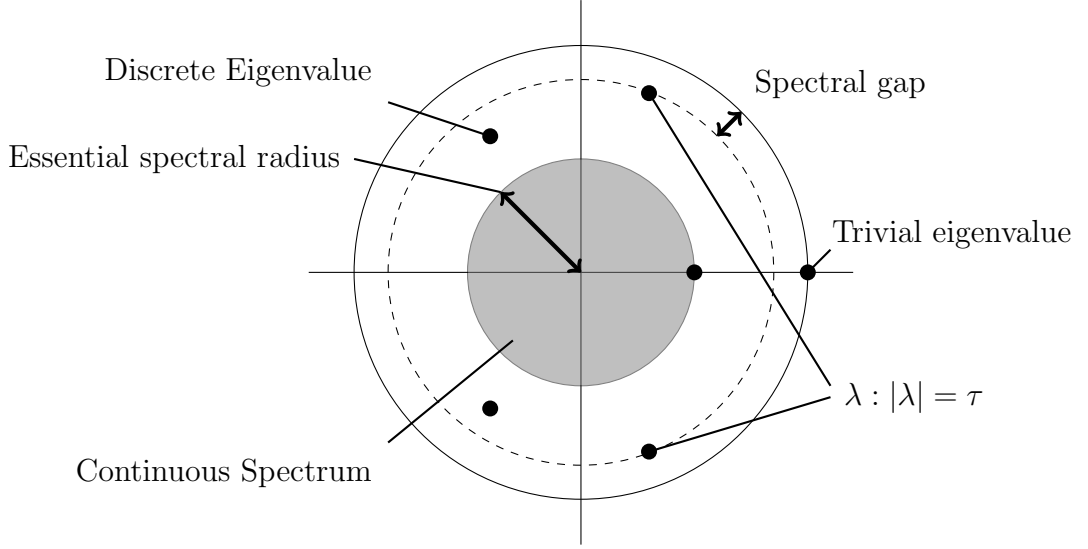
Figure 2.1: A representative spectrum for a transfer operator P is illustrated, which has a continuous spectrum bounded above by r_{ess}, and a point spectrum with a spectral gap since only the trivial eigenvalue $\lambda_1 = 1$ lies on the unit circle.

The spectrum σ will contain, but is not necessarily equal to, the set of eigenvalues of P called the point spectrum. Eigenvalues in the point spectrum are called discrete, and these eigenvalues and their respective eigenfunctions satisfy

$$P_M v_i = \lambda_i v_i. \tag{2.17}$$

There will exist a trivial eigenvalue equal to $\lambda_1 = 1$ which corresponds to the invariant density ρ. There may exist a *continuous spectrum* bounded by a closed disc of radius r_{ess}. The *essential spectral radius* $r_{ess} = r_{ess}(P_M|_x)$, is defined by

$$r_{ess} = \inf\{r \geq 0 : \lambda \in Spec(P_M|_x), \lambda > r \implies \lambda \text{ isolated of finite multiplicity}\}. \tag{2.18}$$

Figure 2.1 illustrates an example spectrum for an operator with both a continuous and point spectrum. To aid in discussion, the point spectrum (2.17) can be organised according to the modulus of their values, such that $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq 0$.

Bounds on mixing rates can be obtained from the spectrum and we say that $P_M|_x$ has a *spectral gap* if

$$\tau = \sup\{|\lambda| : \lambda \in Spec(P_M|_x)\backslash\{1\} < 1\}. \tag{2.19}$$

If an operator has a spectral gap, then the following lemma is important in determining the rate of mixing via the decay of correlations. An alternative formulation

of (2.13) in terms of the Frobenius-Perron operator P_M and associated spectrum is captured in the following Lemma.

Lemma 2.3.1. *Let $\mathfrak{X}_S$ be a class of real-valued functions preserved by P. Let $\sigma(P)$ denote the spectrum of P when considered as an operator of $\mathfrak{X}_S$, and set $\tau = \{|z| : z \in \sigma(P)\backslash\{1\}\}$. Then there is a constant $C < \infty$ such that $C_{\phi,\psi}(j) \leq C\tau^j$ if $\psi \in \mathfrak{X}_S$ and $\phi \in L^\infty$ [Froyland (2001)].*

Therefore τ, the largest non-unit eigenvalue, determines the rate of decay of correlations. As long as $P_M|_{\mathfrak{X}}$ has a spectral gap, then this mixing rate is exponential with rate τ. Smaller values of τ imply a faster decay of correlations while a global mixing rate close to 1 implies that there are eigenfunctions for which convergence is very slow. A transformation will only have a spectral gap if it is strong mixing, otherwise non-trivial eigenvalues will reside on the unit circle. The essential spectral radius essentially gives a lower bound for the value of τ. Note however that estimates on the constant C are in general difficult to obtain, and the value of the smallest such constant C will depend on both the mixing protocol and the initial concentration.

The above describes how spectra of transfer operators help determine strong mixing rates of advective maps, however the discussion can be extend to consider advective-diffusive systems. An operator $P = P_D \circ P_M$ for the iterative advection-diffusion equation will also have an associated spectrum which can be used to predict the asymptotic mixing rate. Even in the absence of stirring, diffusion ensures an exponential approach to uniformity. All systems in which diffusion is included will have a spectral gap, even those in which the underlying advective map itself is not strong mixing. Therefore all eigenvalues λ_k will have a modulus less than one except for the trivial eigenvalue $\lambda_1 = 1$. The eigenfunction of the trivial eigenvalue $\lambda_1 = 1$ is the mean field of the scalar concentration $\bar{c}$, and all other eigenfunctions v_k will decay to $\bar{c}$ with the rate given by λ_k. The eigenfunction associated with λ_2 will be the slowest decaying from the action of the operator. A distinction is made throughout this thesis with τ denoting the largest eigenvalue of the point spectrum for P_M, and λ_2 the eigenvalue with the largest modulus from the spectrum of $P_D \circ P_M$ for $\kappa \neq 0$.

Eigenvalue spectra alone do not give definite information of finite time behaviours, where instead eigenfunctions v can aid understanding of the evolution of densities in early stages.

For two functions $f, g \in \mathcal{X}$ for the operator P, the inner product is defined as

$$\langle f, g \rangle = \int_\Omega f(x)\overline{g(x)}\mathrm{d}x. \tag{2.20}$$

Suppose the functional space has an orthonormal basis given by the set of functions $\{u_1(x), u_2(x), \ldots, u_n(x)\}$ where n may be infinite. For the orthonormal basis

$$\langle u_k, u_l \rangle = \int_\Omega u_k(x)\overline{u_l(x)}\mathrm{d}x = \delta_{kl} \tag{2.21}$$

and functions can be written as a linear combination of the basis functions,

$$f(x) = \sum_{k=1}^n b_k u_k(x). \tag{2.22}$$

The eigenfunctions v_k of P form a set of basis functions for the transfer operator, therefore functions $c(x, j)$ representing the concentration field can be written as a linear combination of the eigenfunctions

$$c(x, j) = \sum_{k=1}^n b_k(j)v_k(x), \tag{2.23}$$

and the coefficients $b_k(j)$ found by evaluating the inner product (2.20) $\langle c, v_k \rangle$. The evolution in time of the coefficients $b_k(j)$ for a given $c(x, 0)$ reveal information on the finite time behaviour of the operator P.

Typically the methods used in determining the spectra of transformations, the values of τ and λ_2 in particular, and analysis of their eigenfunctions are not general and are far from trivial.

2.3.3 Numerical approximations of transfer operators

In special cases the transfer operator (2.14) can be written down explicitly for all functions ϕ, and the point spectrum of P determined from (2.17). However, in general this approach might be too complex. An easier approach is to define a set of basis functions which are dense in the domain of interest, and define the operator as a transfer matrix $\boldsymbol{P}$, describing the transfer in functional space as the change in amplitude in the elements of the basis set. This approach informs the numerical method to investigate the properties of transfer operators. The transfer matrix will satisfy an eigenvalue equation $\boldsymbol{P}\boldsymbol{v} = \lambda\boldsymbol{v}$, and the eigenfunctions $\boldsymbol{v}$ given as a summation of the amplitudes for the basis functions, represented by the elements of the eigenvectors $\boldsymbol{v}$.

When employing transfer operators to study mixing, the most common approach is to consider a transformation on $\mathbb{T}^d$, where d denotes the dimension, and apply periodic boundary conditions. Then, a discrete Fourier series can be used as the set of basis functions. This approach is advantageous since time evolution of the diffusive operator P_D is easily solved is discrete Fourier space. This method has been frequently used in the literature [Ashwin *et al.* (2002); Fereday *et al.* (2002); Popovych *et al.* (2007); Thiffeault & Childress (2003)]. How to explicitly find the advective and diffusive transfer matrices for this set of basis functions is described in detail in Section 2.3.4. When other boundary conditions are taken a different set of basis functions can be used. For example Gilbert (2006) outlines the appropriate functions to solve the heat kernel for zero and insulating boundaries when applying the non-uniform baker's transformation to represent fluid stretching and folding.

In many cases, finding a suitable set of basis functions in which the transfer matrix elements can be found analytically, is impossible. A numerically simple way to construct a transfer matrix for the advective operator is to form a basis of indicator functions over a finite partition of disjoint elements within the domain. Conventionally, evenly distributed squares define the partition, and the transfer matrix element $\boldsymbol{P}_{ij}$ is the proportion of square i carried into square j during one periodic iteration of the velocity field. Diffusion is either included via an operator splitting, such as the Marchuk-Yaneko method, or approximated by numerical diffusion. Such methods were inspired by the work of Spencer & Wiley (1951), who addressed their potential for understanding mixing in viscous fluids, and Ulam who proposed the method to approximate Frobenius-Perron operators representing the evolution of densities [Ding & Zhou (2010); Ulam (1960)]. This approach is sometimes referred to as the Mapping method.

This numerical method has been suggested as a tool box to optimize micromixer configurations, used to detect eigenmodes in granular flows, and study autocatalytic reactions into chaotic advective flows [Christov *et al.* (2011); Schlick *et al.* (2014); Singh *et al.* (2008b)].

The advantage of the Mapping method is its simplicity in derivation, in which matrix elements can be computed by tracking the proportion of many initialised points from one cell into others. Rather than having to evolve single initial conditions with finite difference methods, eigenfunctions of the advectice-diffusive operator can be well approximated from the transfer matrix alone. However, similar to other grid based methods, the computational time and memory can be

particularly costly to ensure finer details of the concentration field are captured. Additionally an effective Péclet number needs to be determined by taking into consideration the numerical diffusion introduced by approximating the transfer operator, which may mean that diffusive effects are not accurate.

2.3.4 Numerical approach using Discrete Fourier expansion

Depending on the map M and the boundary conditions an advective-diffusive transfer operator $P = P_D \circ P_M$ can sometimes be computed with ease using the following method. A specific advantage of this numerical method is that the effect of numerical diffusion is largely eliminated, which can occur for velocity fields with chaotic trajectories when using numerical schemes based on grids.

The numerical problem is to evolve an initial concentration field $c^{(0)}(\boldsymbol{x}) \in \mathcal{X}$ with a domain Ω as an iterative operator problem $c^{(j)} = P^j(c^{(0)})(\boldsymbol{x})$. The transfer operator P can be constructed as a matrix, denoting the transfer of concentration as a change in amplitude for a suitable set of basis functions, which are dense for any $f \in \mathcal{X}$ over the domain Ω. When choosing periodic boundary conditions on the unit circle $\Omega \in [0, 1)$, an integrable function f can be represented by a infinite series of sines and cosines via a discrete Fourier expansion. The set of functions $\cos(2\pi kx)$ and $\sin(2\pi kx)$ for all $k \in \mathbb{Z}$ form a dense set of basis functions and are suitable to construct a transfer matrix acting on the Fourier coefficients. The method is easily expendable to higher dimensions, such as the unit torus $\mathbb{T}^2 = [0, 1) \times [0, 1)$.

For the one-dimensional case on the unit circle $\Omega = \mathbb{T} = [0, 1)$, consider the periodic function $c \in \mathcal{X}$, then its discrete, exponential Fourier expansion is,

$$c^{(j)}(x) = \sum_{q=-\infty}^{\infty} \hat{c}_q^{(j)} e^{2\pi iqx} \tag{2.24}$$

where the complex coefficients $\hat{c}$ are given by

$$\hat{c}_q^{(j)} = \int_0^1 c^{(j)}(x) e^{-2\pi iqx} \mathrm{d}x. \tag{2.25}$$

The basis functions $e^{2\pi ikx}$ are dense in $\mathcal{X}$. Iterating the concentration field by a mapping $M(x)$ produces a new concentration field with a Fourier expansion and its own respective Fourier coefficients $c^{(j+1)}(x) = \sum_{k=-\infty}^{\infty} \hat{c}_k^{(j+1)} e^{2\pi ikx}$.

If the mapping is considered on the Fourier coefficients, then the new coefficients after one iteration will be of the form

$$\hat{c}_k^{(j+1)} = \int_0^1 c^{(j)}(M^{-1}(x))e^{-2\pi ikx}\mathrm{d}x. \tag{2.26}$$

$$= \sum_{q=-\infty}^{\infty} \hat{c}_q^{(j)} \int_0^1 e^{2\pi iqM^{-1}(x)-2\pi ikx}\mathrm{d}x. \tag{2.27}$$

The transformation between Fourier coefficients can be represented by a linear transformation in the form of a transfer matrix

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \boldsymbol{M}_{kq}\hat{c}_q^{(j)} \implies \boldsymbol{M}_{kq} = \int_0^1 e^{2\pi iqM^{-1}(x)-2\pi ikx}\mathrm{d}x. \tag{2.28}$$

Similarly, if the transfer operator acting on the concentration field was purely diffusive P_D with diffusivity coefficient κ then a transfer matrix representing this is

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \boldsymbol{D}_{kq}\hat{c}_q^{(j)}, \qquad \boldsymbol{D}_{kq} = \rho^{k^2}\boldsymbol{\delta}_{kq}, \qquad \rho = \exp(-4\pi^2\kappa T). \tag{2.29}$$

where the Kronecker delta $\boldsymbol{\delta}_{kq} = 1$ if $k = q$ and 0 otherwise. Since the advection component is considered as an iterative map, the time period T only appears in the diffusive contribution $\exp(-4\pi^2\kappa T)$, and so can be normalised to $T = 1$, allowing the diffusive step to be parametrized by the κ alone. The full iterative advection-diffusion mapping for one time step is represented by the composition of the transfer matrices

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \boldsymbol{D}_{km}\boldsymbol{M}_{mq}\hat{c}_q^{(j)}. \tag{2.30}$$

If an analytic expression for the entries of the transfer matrix $\boldsymbol{M}$ can be found, then numerical calculation in discrete Fourier space can be accomplished easily without much concern for numerical diffusion. The only consideration is that the number of modes used in the computation can not be infinite, and a truncation on the number of Fourier modes is required, so that $-Q \le q, k \le Q$. Wonhas & Vassilicos (2002) note that the diffusive contribution to the transfer between Fourier modes is $\exp(-4\pi^2\kappa k^2)$ so modes in which $|k| \gg k_d$ are negligible, where

$$k_d = \frac{1}{2\pi\sqrt{\kappa}}. \tag{2.31}$$

The effect of the truncation Q will be investigated to ensure the correct evolution and decay of the concentration field occurs irrespective of truncation value.

The method is easily extended to N dimensions by Fourier expanding $c^{(j)}(\boldsymbol{x})$

$$c^{(j)}(\boldsymbol{x}) = \sum_{\boldsymbol{k}} \hat{c}_{\boldsymbol{k}}^{(j)} e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}} \tag{2.32}$$

then

$$\hat{c}_{\boldsymbol{k}}^{(j+1)} = \sum_{\boldsymbol{q}} \boldsymbol{M}_{\boldsymbol{k}\boldsymbol{q}} \hat{c}_{\boldsymbol{q}}^{(j)} \tag{2.33}$$

with

$$\boldsymbol{M}_{\boldsymbol{k}\boldsymbol{q}} = \int_{\mathbb{T}^N} \exp\left(2\pi i (\boldsymbol{q} \cdot M^{-1}(\boldsymbol{x}) - \boldsymbol{k} \cdot \boldsymbol{x})\right) \mathrm{d}\boldsymbol{x}, \tag{2.34}$$

or for the full advective-diffusive step

$$\boldsymbol{DM}_{\boldsymbol{k}\boldsymbol{q}} = \int_{\mathbb{T}^N} \exp(2\pi i (\boldsymbol{q} \cdot M^{-1}(\boldsymbol{x}) - \boldsymbol{k} \cdot \boldsymbol{x}) - 4\pi^2 \kappa \boldsymbol{k} \cdot \boldsymbol{k}) \mathrm{d}\boldsymbol{x}, \tag{2.35}$$

where $\boldsymbol{x} = (x_1, x_2, \ldots, x_N)$ and similarly for $\boldsymbol{k}$ and $\boldsymbol{q}$. The summation is over all possible $\boldsymbol{q} = (q_1, \ldots, q_N)$ and $\boldsymbol{k} = (k_1, \ldots, k_N)$ with $-Q \leq k_n, q_n \leq Q$ for the chosen truncation Q.

The leading, non-trivial eigenvalue λ_2 of the transfer matrices $\boldsymbol{DM}$ give a single value for the asymptotic mixing rate and the eigenvectors $\boldsymbol{v}$ contain the coefficients $\hat{c}_k$ for the discrete Fourier expansion for the eigenfunctions v_k. The inner product (2.20) can be rewritten in terms of discrete Fourier coefficients by the following derivation.

Given the functions $f(\boldsymbol{x}) = \sum_{k=-\infty}^{\infty} a_k e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}}$ and $g(\boldsymbol{x}) = \sum_{l=\infty}^{\infty} b_l e^{2\pi i \boldsymbol{l} \cdot \boldsymbol{x}}$ the inner product

$$\langle f, g \rangle = \int_0^1 f(\boldsymbol{x}) \overline{g(\boldsymbol{x})} \mathrm{d}\boldsymbol{x}, \tag{2.36}$$

can be written as

$$= \int_0^1 \sum_{k=-\infty}^{\infty} a_k e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}} \sum_{l=\infty}^{\infty} \overline{b_l} e^{-2\pi i \boldsymbol{l} \cdot \boldsymbol{x}} \mathrm{d}\boldsymbol{x},$$

$$= \int_0^1 \sum_{k=-\infty}^{\infty} \sum_{l=\infty}^{\infty} a_k \overline{b_l} e^{2\pi i (\boldsymbol{k} - \boldsymbol{l}) \cdot \boldsymbol{x}} \mathrm{d}\boldsymbol{x}.$$

Evaluating the integral gives

$$\langle f, g \rangle = \sum_{k=-\infty}^{\infty} \sum_{l=\infty}^{\infty} a_k \overline{b_l} \delta_{kl} = \sum_{k=-\infty}^{\infty} a_k \overline{b_k}. \tag{2.37}$$

Therefore, the vectors $\boldsymbol{v}$ can be used to determine the influence the eigenfunctions v have on the finite time behaviour of a given initial condition $c(x, j)$, acted on by the transfer matrix $\boldsymbol{DM}$ by evaluating (2.37).

Truncating the matrix makes it possible to compute the discrete spectrum of the transfer operator. Eigenvalues and eigenvectors of the transfer matrix can be computed easily using preprogrammed algorithms, however when increasing Q the computation time quickly becomes infeasible. Even constructing the transfer matrix for large Q involves large amounts of computing memory, especially for a two-dimensional system where memory restrictions may mean this is not achievable in practice. In this work, construction of the transfer matrix is restricted to one-dimensional systems. In two-dimensional systems the numerical approach used is to iterate over the sums of wavemodes instead. The Python package Scipy is used to solve for the eigenvalues and eigenvectors of the composition matrix $\boldsymbol{DM}$.

2.4 Measures of mixing

The ergodic hierarchy provides a mathematical formalisation encapsulating ideas of good stirring within an advective transformation, however solving analytically for the properties is complicated in general and rarely possible in realistic flows. In addition they say nothing about the mixing quality or rates of mixing for observable functions, such as a scalar concentration fields decaying via the action of molecular diffusion. Transfer operators provide a means to study the evolution of typical concentrations fields in both advective and advective-diffusive systems, but properties of their spectra only relate to mixing rates in the asymptotic limit. Understanding mixing throughout the full finite time evolution is essential in real mixing scenarios.

As such an array of techniques quantifying the size of blobs, their occupation throughout a domain, and their homogenization with the surrounding fluid, has been developed throughout the literature to evaluate mixing in a range of applications. Although a measure on the quality of mixing is essential, how to quantify such a process is not straightforward. A handful of measures are introduced which are important to the current work.

2.4.1 Functional norms and the mix norm

An approach to uniformity from diffusive effects is easily quantified from functional norms, where a norm is defined measuring the distance between the concentration field being mixed and its mean field.

The family of norms called the L^q norms are frequently used, and are defined such that given a concentration field $c(\boldsymbol{x})$ with mean field $\bar{c}$,

$$||c(\boldsymbol{x},t) - \bar{c}||_q = \left(\int_\Omega |c(\boldsymbol{x},t) - \bar{c}|^q d\Omega \right)^{1/q}. \tag{2.38}$$

The L^∞ norm is specifically defined as,

$$||c(\boldsymbol{x},t) - \bar{c}||_\infty = \inf\{R : |c(\boldsymbol{x},t) - \bar{c}| \leq R \text{ a.e. } \boldsymbol{x} \in \Omega\}. \tag{2.39}$$

Determining which L norm is better suited to quantify mixing is dependent on the application at hand. For example Ashwin *et al.* (2002) highlight that when disinfecting in food processing, L^∞ would be more desirable since it is important that all locations are smoothly covered during the mixing process. The most common norm in the study of mixing is the L^2 norm, the variance of the concentration field from its mean field. While L^q norms are good measures in systems which are mixing to uniform, in systems that are purely stirring, functional norms such as those defined by (2.38) and (2.39) would not decay. As such in purely stirring systems other measures of mixing are required.

One of the first quantifications of stirring quality was devised by Gibbs (1902), who formulated the idea of coarse-grained density. Grid cells are assigned to the mixing domain of interest and values computed on the average number of grid cells occupied by a component of the stirred fluid. Problems arise with such methods since they are sensitive to the size of the cells used in computation. Later Danckwerts (1952) defined the scale of segregation, which encapsulated the average of the size of clumps in a mixed component. Mathew *et al.* (2005) developed a functional norm to capture the idea of the scale of segregation which does not need diffusion to decay, called the mix norm, which is now described.

Consider the scalar function $c : \mathbb{T}^d \to \mathbb{R}$, where $\mathbb{T}^d = [0,1]^d$ is an d-dimensional torus and μ is the Lebesgue measure. For a given length scale $s \in (0,1)$ and point $\boldsymbol{p} \in \mathbb{T}^n$, then define the sphere

$$B(\boldsymbol{x},s) = \{\boldsymbol{y} \in \Omega \text{ such that } ||\boldsymbol{y} - \boldsymbol{x}||_2 \leq s/2\}, \tag{2.40}$$

that is the ball is defined as all the points within the d-dimensional sphere about the point $\boldsymbol{x}$ of radius $s/2$, and volume of the ball noted $\text{Vol}(B(s))$.

To define the mix norm let

$$b(c, \boldsymbol{p}, s) = \frac{\int_{x \in B(\boldsymbol{p},s)} c(\boldsymbol{x})\mu(\mathrm{d}\boldsymbol{x})}{\mathrm{Vol}B(s)}, \tag{2.41}$$

for all length scales $s \in (0,1)$ and $\boldsymbol{p} \in \mathbb{T}^d$. Taking the L^2 norm of b

$$\phi(c, s) = \left(\int_{\Omega} b^2(c, \boldsymbol{x}, s)\mu(d\boldsymbol{x}) \right)^{1/2} \tag{2.42}$$

and finally integrating over all possible values of the scale s gives the mix norm for the concentration field c as

$$\Phi(c) = \left(\int_0^1 \phi^2(c, s)\mu(ds) \right)^{1/2}. \tag{2.43}$$

This norm was shown to capture the exponential and algebraic decay rates expected in specific mixing problems [Mathew *et al.* (2005)]. Numerically integrating over spatial scales is time consuming, therefore Smith (2016) presents a recursive algorithm to reduce computational cost. If the problem is computed, or accurately transformed into discrete Fourier space, the computation is less intensive taking advantage of the spatial scales represented by the wavemode number. Similarly, negative Sobolev norms will decay in the absence of diffusion and are easily computed in discrete Fourier space. These such norms are collectively called multiscale norms, and Thiffeault (2012) provides a review of their use in quantifying mixing and transport.

Profiles of norms decaying to the mean field, plotted against iteration number, give a full finite time analysis for a given stirring protocol and diffusivity coefficient. When comparing across a range of parameters however, the full profiles can be difficult to compare. A single quantifying value for finite time mixing can be found by choosing a mixed condition and computing the iteration number when this condition is achieved; called the *time to a mixed condition.* In line with Ashwin *et al.* (2002), if interest lies in finding how many iterations it takes to stir to within 5% of uniform, or the time to 95% mixed state t_{95}, then define t_{95} as the smallest $t_{95} > 0$ such that

$$||c^{t_{95}}(\boldsymbol{x}) - \bar{c}|| < 0.05 \tag{2.44}$$

for all $j > t_{95}$ and the chosen norm $||\cdot||$. Other *time to $t_{\%}$* can be defined similarly as $t_{\%} > 0$ such that

$$||c^{t_{\%}}(\boldsymbol{x}) - \bar{c}|| < 1 - \%. \tag{2.45}$$

The *time to a mixed condition* $t_\%$ gives a single value which can be used to compare across varying parameters and initial conditions for mixing systems, capturing how they effect mixing in the initial stages without analysing the full time of a norm decay. Throughout the thesis, eigenvalues τ and $|\lambda_2|$ will be computed where possible to understand the asymptotic behaviour of each mixing system, while t_{95} and $t_\%$ will be used within parameter studies as a simple method to quantify changes within the finite time mixing behaviour, where a full spectra and eigenfunction analysis would be demanding.

2.4.2 Measures of mixing in Fourier space

The variance of a concentration field from the mean-field is easily computed from Fourier coefficients via Parseval's Theorem. Recall the variance is taken as

$$\psi(j) = \int_0^1 |c(\boldsymbol{x}) - \hat{c}|^2 \mathrm{d}\boldsymbol{x}, \tag{2.46}$$

and substituting in $c(x) = \sum_{\boldsymbol{k}=-\infty}^{\infty} \hat{c}_{\boldsymbol{k}} e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}}$ gives

$$\psi(j) = \int_0^1 \left| \sum_{\boldsymbol{k}:\boldsymbol{k} \neq \boldsymbol{0}} \hat{c}_{\boldsymbol{k}} e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}} \right|^2 \mathrm{d}\boldsymbol{x}, \tag{2.47}$$

where the amplitude of the wavemode $\boldsymbol{0} = (0, \ldots, 0)$ is ignored since $\hat{c}_{\boldsymbol{0}} = \bar{c}$. Using the fact that when $c(\boldsymbol{x})$ is complex, $|c(\boldsymbol{x})|^2 = c(\boldsymbol{x})\overline{c(\boldsymbol{x})}$, then

$$\left| \sum_{\boldsymbol{k}:\boldsymbol{k} \neq \boldsymbol{0}} \hat{c}_{\boldsymbol{k}} e^{2\pi i \boldsymbol{k} \cdot \boldsymbol{x}} \right|^2 = \sum_{\boldsymbol{k}:\boldsymbol{k} \neq \boldsymbol{0}} \sum_{\boldsymbol{l}:\boldsymbol{l} \neq \boldsymbol{0}} \hat{c}_{\boldsymbol{k}} \overline{\hat{c}_{\boldsymbol{l}}} e^{2\pi i (\boldsymbol{k}-\boldsymbol{l}) \cdot \boldsymbol{x}}. \tag{2.48}$$

Substituting (2.48) into the integral (2.47) gives

$$\psi(j) = \sum_{\boldsymbol{k}:\boldsymbol{k} \neq \boldsymbol{0}} \sum_{\boldsymbol{l} \neq \boldsymbol{0}} \hat{c}_{\boldsymbol{k}} \overline{\hat{c}_{\boldsymbol{l}}} \delta_{kl} = \sum_{\boldsymbol{k}:\boldsymbol{k} \neq \boldsymbol{0}} |\hat{c}_{\boldsymbol{k}}|^2 \tag{2.49}$$

Therefore the variance can be computed as the sum of the modulus squared of the discrete Fourier coefficients of the concentration field. It is from (2.49) that it can be easily understood, that if the long-time mixing is dominated by the slowest decaying eigenfunction v_2, the coefficients $\hat{c}_k$ of the eigenfunction v_2 will vary iteratively as $\lambda_2 \hat{c}_k$, and the variance decay on average given by

$$\psi(j) \sim |\lambda_2|^{2j}. \tag{2.50}$$

For the L^∞ norm, computation from Fourier coefficients alone is not possible and the concentration field $c(\boldsymbol{x})$ would need to found from $\hat{c}_{\boldsymbol{k}}$ via an inverse Fast Fourier Transform (IFFT). Many computational programmes provide packages which compute IFFTs, in this work the Python package Numpy is used.

2.4.3 Additional measures of mixing

A number of other measures exist in the literature but have limitations. Two of the more commonly used are the growth in length of the material interface, and the decrease in the largest or average lamellar width.

The length of interface between two materials is a measure on the amount of stretching occurred on a fluid element. Used in one-dimensional and two-dimensional studies of lamellar formations, increased interface is desirable in problems without diffusion [Krasnopolskaya *et al.* (1999); Krotter *et al.* (2012); Meleshko & Aref (1996)]. Usually in two-dimensions the interface of an initial dye parcel is represented by a collection of points and the mapping these point forward in time reveals how the interface changes and extends in the fluid flow. However it can be difficult to compute efficiently in such studies, since tracking an interface which stretches exponentially results in resolution problems and gaps in the quickly appear. Algorithms have been developed which introduce points when the separation becomes large [Krasnopolskaya *et al.* (1999)], however the exponential stretching leads to a limit on precision. In one-dimensional models, tracking the growth in interface is easy since it is simply represented by a countable set of finite points. For example, the increased interface via discontinuous mapping has been measured by the increasing number of cuts C_j at each iteration [Krotter *et al.* (2012); Yu *et al.* (2016)].

In one-dimensional models of strips of fluid, the average lamellar width or largest lamellar width gives an indication on how much the fluid is contracted and folded, however it has been shown that the decay rate in diffusion and reaction problems is sensitive to lamellar arrangement [Clifford *et al.* (1999)]. An additional problem with such measures as that they are not easily extended to two-dimensional studies due to the variation in the orientation of striations [Finn *et al.* (2004)]. Such measures have been employed when cutting and shuffling a one-dimensional concentration field [Krotter *et al.* (2012); Smith *et al.* (2018); Yu *et al.* (2016)].

The main problem with both of these measures of mixing is that they lack knowledge of the distribution of mixing throughout the full domain, and the average scales of segregation. For example, an interface within a small chaotic region may stretch exponentially without visiting much of the domain [Finn *et al.* (2004)]. Although computationally more intensive to compute, the mix norm eliminates such problems and gives a quantification to uniformity from stirring.

The measures described so far are measures which quantify a degree of mixing at snapshots during the mixing time for a particular initial concentration field, however other methods in the literature, such as homogenisation theory, have been successful at predicting finite time mixing behaviours of a passive scalar in simple flows by describing the evolution of desired statistics but these methods can be difficult to apply in general [Majda & Kramer (1999)].

2.5 Mixing rates in chaotic advection

Now that methods to capture and quantify mixing have been outlined, a review on results from the literature concerning mixing rates arising from a range of stirring dynamics are presented. Molecular diffusion acting alone will cause a concentration field to tend to uniformity at an exponential rate, although this rate is generally very slow. Stirring will increase the rate to uniformity.

In fully chaotic flows exponential stretching and compression of fluid parcels produce an exponential rate in the reduction of length scales and an increase in the material interface. However, elliptic islands and invariant tori, boundaries surrounding fully chaotic domains, or parabolic points can contaminate this exponential mixing rate [Horita *et al.* (1990); Springham & Sturman (2014); Sturman & Springham (2013)]. However, these mixing processes are unlikely to arise in isolation in real life applications.

There is an extensive literature studying the combination of stretching and folding from chaotic advection, with diffusion, and how the underlying stirring dynamics produce the observed mixing rates. The key results are outlined in the following sections.

2.5.1 Non-uniformity and strange-eigenmodes

Recall (2.4) which describes the evolution of the variance ψ from the diffusivity coefficient κ and the average gradients in the concentration field $|\nabla c|$. If κ is small, in the initial stages of stirring the gradients in the concentration field are small and the variance will remain almost constant. It is only when the gradients $|\nabla c|$ are large that ψ will begin to decay significantly. It is suggested from (2.4) that good stirring occurs when the system becomes independent of diffusivity, such that $\kappa \langle |\nabla c^2| \rangle$ balances and

$$\nabla c \sim \kappa^{-1/2}, \tag{2.51}$$

31

and the smallest scales visible in the concentration field have size $\sim \sqrt{\kappa}$. If the flow is assumed to be chaotic with average local strain rate h, the gradients in the concentration field would vary as $|\nabla c| \sim e^{-ht}$. The time T when the concentration striations stabilise due to the balance between the effect of diffusion and strain is given by

$$e^{-ht} \sim \sqrt{\kappa}, \implies T \sim h^{-1}\log\kappa. \tag{2.52}$$

The length scale of the striations where this occurs is

$$l_B = \sqrt{\frac{\kappa}{h}}, \tag{2.53}$$

called the Batchelor length [Balkovsky & Fouxon (1999); Thiffeault (2008)]. Once the spatial scale of the concentration field reaches the Batchelor length, the concentration field will deplete quickly. From (2.52), this initial stage where ψ is almost constant varies weakly with diffusivity, only logarithmically with κ. The decay rates of the concentration field after this stage, quantified by norms such as the variance, will vary depending on the kinematic dynamics of the underlying flow.

In chaotic systems which have a uniform rate of strain throughout the domain, the addition of diffusion predicts an unrealistic super-exponential mixing rate due to the combined action of exponential diffusive rates alongside exponential reduction in concentration scales [Thiffeault & Childress (2003); Wonhas & Vassilicos (2002)]. Such rates have never been observed outside of perfectly realised mathematical models, since real fluid flow contains large amounts of non-uniformity in the stretching and deformation from shearing in the velocity field.

Introducing non-uniformity in the strain rates of the underlying flow field produces exponential mixing rates overall. This was first observed by Pierrehumbert (1994), where a non-uniformly deforming chaotic flow produced a persistent pattern in the concentration field after several iterations, and the variance decay achieved an exponential long-time mixing rate. It was noted that on decreasing the rate of diffusivity, the initial stages where not much mixing took place extended, but the exponential mixing rate became independent of the diffusivity. The spatial structure on the concentration field remained time independent but the amplitude decreased at each iteration. The pattern appeared to have arbitrary fractal scales, limited only by the presence of diffusion, and as such was described as a "strange eigenmode", where strange was chosen to describe the

fractal spatial scales observed. A strange eigenmode was observed experimentally by Rothstein *et al.* (1999), where a spatially persistent pattern emerged in an electromagnetically driven, two-dimensional fluid layer with florescent dye.

Local Lagrangian dynamics were used by Antonsen Jr *et al.* (1996) in an attempt to explain the exponential decay in a unifying theory. It was proposed that the exponential decay is determined by the small number of orbits which are aligned perpendicular to the contracting direction of the flow. Thiffeault (2008) considered probability distributions of randomised stretching histories, where again a slower decay rate was produced, however concern was raised over the limitations of local theory if blobs of fluid were to interact with each other or boundaries during evolution. Indeed, examples were found in which the local theory could not account for an exponential decay of tracer variance.

Fereday *et al.* (2002) posed a one-dimensional baker's transformation with a diffusional step, which exhibited long-time exponential decay when the stretching factors were non-homogeneous within the interval. In this one-dimensional map, perpendicular alignment of stretching histories do not account for the observed decay rate, since the initial condition is directly aligned with the stretching direction. In such a system the local theory predicts a super-exponential mixing rate.

A similar study in two-dimensional discrete Fourier space, used the cat map with non-uniformity added via a small-wave perturbation [Thiffeault & Childress (2003)]. This perturbed system was a rare example of a two-dimensional map in which the variance decay could be investigated analytically. It concluded that the exponential mixing rate which emerged was not governed by the alignment of stretching histories, since the parameters ensured the map was close to uniform stretching, but the dispersion of variance into wavemodes of longer wavelength via the non-uniformity. Both Fereday *et al.* (2002) and Thiffeault & Childress (2003) showed the decay of variance was dominated by the variance trapped in the gravest Fourier modes.

This theory in which the gravest wavemodes dominate the decay of variance, resulting in the emergence of eigenfunctions of the concentration field, was described as a "global" theory of mixing, in contrast to "local" Lagragian description from stretching histories. The global mechanism arises from the arrangement of striations with varying widths, due to the non-uniformity in the stretching rates. This latter theory is supported by the long-time mixing rate being sensitive to boundary conditions which the local description would not predict, shown for

the one-dimensional non-uniform baker's transformation by Wonhas & Vassilicos (2002) and Gilbert (2006).

It has been argued however that the local mechanism describes well the initial transient stage of mixing, before the eigenmode regime begins [Fereday & Haynes (2004); Hu & Pierrehumbert (2002); Wonhas & Vassilicos (2002)]. Studies on the one-dimensional non-uniform baker's transformation map were extended by Wonhas & Vassilicos (2002) who showed that the probability density function (PDF) of stretching histories capture well the evolution of concentration in the first few iterations [Wonhas & Vassilicos (2002)]. Additionally Tsang *et al.* (2005) and Haynes & Vanneste (2005) found the local mechansim can capture the asymptotic stages of mixing depending on the flow and velocity scales in the vanishing diffusivity limit, confirmed by Sukhatme & Pierrehumbert (2002). The local argument provides a lower bound on the exponential mixing rate, associated with the essential spectral radius of the continuous spectrum, while the global mechanism is associated with discrete isolated eigenvalues of the advective transfer operator [Haynes & Vanneste (2005)]. The global mechanism describes the regime where the domain scale is significantly larger than the flow scale, and scalar transport is dominated by large scale transport and the domain geometry.

Although in the global mechanism, Lypaynov exponent statistics fail to capture the asymptotic mixing rate of tracer decay, qualitatively the spatial distributions of finite time stretching histories seem linked to the spatial arrangement of the strange eigenmodes. Gilbert (2006) used boundary layer theory to predict the eigenfunctions in the presence of small diffusion and showed eigenmode alignment to regions of low stretch were evident, similarly reported in other numerical studies [Gouillart *et al.* (2008)].

The emergence of eigenfunctions in the concentration field via the advective-diffusive transfer operator with $\kappa > 0$ is easily understood, since the eigenvalues all satisfy $|\lambda_k| < 1$. The eigenfunctions v_k satisfying (2.17) decay at the rate of their respective eigenvalues under each application of the transfer operator $P = P_D \circ P_M$. For a given initial concentration field under the action of P, after some time the eigenfunctions with the slowest decaying rate will begin to emerge in the concentration field and dominate the long-time evolution. The average long-time exponential mixing rate is given by the modulus second leading eigenvalue of the transfer operator for the discrete time evolution, $|\lambda_2|$. Therefore the long-time mixing rate is independent of initial condition. The variance decays proportionally to the leading eigenvalue as $\psi_j \propto |\lambda_2|^{2j}$, however its possible for

oscillations about this average mixing rate arise due to the complex nature of the eigenvalue λ_2 [Toussaint *et al.* (1995)].

The independence in the long-time mixing rate which emerges in the limit of zero-diffusivity is due to the leading eigenvalue tending to the strong mixing rate of the advective transfer operator, $|\lambda_2| \to \tau$, however this approach is non general. In the local regime, Haynes & Vanneste (2005) reported that if the spectrum has no isolated eigenvalues, then the leading eigenvalue λ_2 of P tends to r_{ess}, the upper limit of the continuous spectrum of P_M, like $1/\log^2 \kappa$. In the global regime the diffusive correction to τ goes like κ^σ for some $0 < \sigma < 1$.

For many of the numerical studies in the global regime, the approach $\lambda_2 \to \tau$ is monotonic with κ [Thiffeault & Childress (2003); Wonhas & Vassilicos (2002)]. Non-monotonic convergence in the diffusive limit was reported by Gilbert (2006) for the non-uniform baker's transformation with a flux boundary condition, where regular oscillations in the correction $|\tau - \lambda_2|$ where observed with scaling $\log(\kappa)$, but in general the approach exhibited a power law in line with Haynes & Vanneste (2005). Similar non-monotonic convergence was reported by Eckhardt *et al.* (2003) for a 3 branched, folded baker's transformation with a diffusive step.

The pre-mentioned studies on non-uniformity in stretching rates gave some insight into the complicated behaviour of laminar advection-diffusion, but were idealised cases concerning fully chaotic domains with mainly periodic boundaries. In many real life mixing systems, barriers to mixing, such as islands, and no-slip boundary conditions would be present.

2.5.2   Boundaries and walls

No-slip boundaries have also been shown to inhibit mixing. Numerical studies considering the effect of slip and no-slip wall boundaries reported that the variance was shown to decay via a power-law during intermediate times of mixing for both Neumann and Dirichlet boundary conditions [Salman & Haynes (2007); Zaggout & Gilbert (2012)]. Additional studies such as an experiment of a figure of 8 stirrer in a circular vessel will walls [Gouillart *et al.* (2007)], a numerical adaptation of the baker's map with a parabolic point at the origin [Gouillart *et al.* (2008)], and an abstract two-dimensional ergodic map of an egg beater flow with parabolic points along the boundaries [Sturman & Springham (2013)] also revealed power-law mixing rates in both diffusive or non-diffusive systems. In these latter studies, analytical results or numerical Poincaré sections showed that the chaotic regions extended to the boundaries and there were no barriers to mixing in all cases.

They concluded that the fluid experiencing slow, algebraic stretching rates at the boundaries injected into the bulk flow via repeated stretching and folding contaminates mixing rates across the whole domain.

In the figure of 8 stirring vessel, exponential decay rate was successfully re-achieved in the bulk through rotation of the outer walls generating a homoclinic orbit, protecting the inner mixing from boundary contamination Gouillart *et al.* (2010); Thiffeault *et al.* (2011). The central region at long-times recovered the formation of an eigenmode, however there remained a thin layer of unmixed fluid in the region at the boundaries.

A question can be posed however, across a number of protocols, can a stirring strategy with an algebraic decay rate beat one of exponential rate to a sufficient mixing condition? Rotating an exterior wall will provide additional energy costs in practice and a least cost effective protocol may be just as effective. Mixing protocols are generally studied in isolation and only a handful compare a large number of parameters [Finn *et al.* (2004)].

2.5.3 Barriers to mixing

A transformation with elliptic islands, or several chaotic domains separated by invariant tori, are not ergodic and as such limit particle transport. Elliptic islands have been shown to cause trajectories which pass near to be slow and stick to them for some time, posing problems to mixing similar to boundary behaviour discussed in Section 2.5.2 [Horita *et al.* (1990); Pierrehumbert (1991)].

However, the main contamination from elliptic orbits and invariant tori in the approach to uniformity, is the restriction of concentration transfer to only diffusion across the barriers in the long-time limit. The inhibition of material transfer from stirring means that large elliptic islands will prevent an approach to the Batchelor length in regions of the flow, preventing a rapid approach to uniformity. In the limit of small diffusivity, the transfer of material by diffusion through the interfaces is slow.

Popovych *et al.* (2007) investigated the evolution of a scalar concentration field via advection-diffusion in a phase space with weakly connected chaotic regions, and a phase space consisting of a chaotic sea with large elliptic islands. In both cases the transfer operator describing the iterative evolution was computed and it was found that the finite time variance decay could be well approximated as a summation of the two slowest decaying eigenfunctions of the transfer operator. The latter phase space resulted in a polynomial decay rate. Plotting the spatial

distribution of stretching histories via finite time Lyapunov exponents, it was observed that the slowest decay eigenfunctions aligned with regions of low stretch, near islands or boundaries to chaotic regions.

2.6 Discontinuous mixing

The studies discussed so far have been mainly concerned with material deformations arising via stretching and folding alone, however systems with discontinuous stirring also exists. The stirring mechanisms arising from discontinuities in such systems has previously been described as cutting and shuffling (CS) [Smith *et al.* (2016); Sturman (2012); Wang & Christov (2018)].

2.6.1 Dynamics of Piecewise Isometries

Most discontinuous systems studied analytically to date have been those composed of CS dynamics alone. Transformations composed of translations, rotations and re-orientations, are called Piecewise Isometries, which was largely instigated by Goetz [Goetz (1998, 2000, 2003)]. Such systems can give rise to complex structures, where coding partitions can have fractal self similar patterns [Goetz (2003)], and non-smooth invariant curves forming barriers to ergodicity [Ashwin & Goetz (2005)]. Positive Lyapunov exponents and exponential mixing are the hallmark of chaotic advection, and although complex behaviour can emerge in piecewise isometries, Lyapunov exponents have be shown to be zero [Fu & Duan (2008)], and transformations having zero entropy [Buzzi (2001)].

Although stirring from CS lacks essential chaotic properties and does not cause material deformation from shearing, it can still create interface and decrease the scale of segregation. A handful of analytical results outline the potential mixing properties in one-dimensional piecewise isometries.

An Interval Exchange Transformation (IET) is a one-dimensional subset of generalised piecewise exchanges and translations in higher dimensions. Mathematical studies of IETs have proven that they increase the number of interfaces linearly [Novak (2009)] and are minimal and ergodic if the permutation on the intervals is irreducible and the subinterval lengths rationally independent [Keane (1975)]. Katok (1980) asserts IETs are never strong mixing and it was later proved that almost all minimal permutations which are not a rotation are weakly mixing [Avila & Forni (2007)], at a rate which is at most polynomial. The consequence of a transformation that is not strong mixing is that the transfer operator

describing the evolution of observables contains no spectral gap. Therefore transfer operators describing Interval Exchange transformations have a point spectrum with non-trivial eigenvalues on the unit-circle.

Weak mixing is an asymptotic mixing rate, and more recent numerical studies have been concerned with the finite time segregation properties of Interval Exchange transformations [Krotter *et al.* (2012); Yu *et al.* (2016)] and higher dimensional Piecewise isometries constructed on a hemi-spherical shell [Juarez *et al.* (2012); Park *et al.* (2016, 2017); Smith *et al.* (2017a)]. All studies show that significant complexity can arise in such systems after a finite number of iterations. Both Krotter *et al.* (2012) and Yu *et al.* (2016) found that the parameters which result in weak mixing inform the parameters for good finite time stirring, even if they are not strictly met i.e., rational dependence between subintervals can be relaxed to some extent. Introducing cut variation, mimicking uncertainty in mechanical stirring, also improves stirring [Yu *et al.* (2016)]. Smith *et al.* (2018) proposed an IET with variable parameters at each iteration which optimized decreasing the scale of segregation.

The hemi-spherical shell model was first proposed by Sturman *et al.* (2008) as an idealisation of a spherical granular tumbler. In the limit of a vanishing flowing layer, where shearing like mixing takes place, the mixing is captured purely by cutting and shuffling in the bulk from solid body rotation. A half filled tumbler is rotated first in one-direction, followed by a second rotation in the perpendicular direction, represented as a piecewise isometry with four pieces whose sizes are determined by the rotation angles. In such systems, streamline jumping governed by solid body rotation and a limiting shear layer leads to complexity in the absence of chaos [Christov *et al.* (2010)]. There has been extensive investigation into the rotation angle parameter space and resulting coverage of cutting lines within the domain [Juarez *et al.* (2012); Park *et al.* (2017)]. By identifying resonances in the parameter space which generate large regions untouched by cutting lines, Smith *et al.* (2017a) where able to suggest a broad collection of rotation angles likely to result in a successful stirring in finite time.

IETs, the hemi-spherical shell, and many other piecewise isometric transformations informed a recent review on their possibilities in mixing [Smith *et al.* (2019)]. However, these mixing processes are unlikely to arise in isolation in real life applications. The next section outlines the range of applications involving combined dynamics of stretching and folding with discontinuities, followed by a

review of studies considering the combined effects of stretching and folding (SF), cutting and shuffling (CS) and diffusion (D).

2.6.2 Examples of discontinuous mixing

There are many instances in mechanical mixing in which discontinuities arise, although there is comparatively little understanding in their implication on mixing. For example, if the underlying fluid being stirred is a polymer solution, paste, or is highly viscous, stirring by conventional methods such as using paddles or rods may require too much energetic costs to be effective. Split–and–recombine mixers adopt the action of cutting and shuffling to increase the number of lamellae between substances and can have complicated geometries. They are commonly employed in micro-mixing flows only millimetres across [Hardt *et al.* (2008); Hobbs & Muzzio (1997); Hossain & Kim (2015); Schönfeld *et al.* (2004)]. Notably, the Kenics mixer cuts and rotates fluid to mimic the stretching and folding action of the baker's transformation, a paradigm transformation which is strong mixing [Hobbs & Muzzio (1997)].

Sink–source flows stir via pressure differences arising from fluid subtraction and reinjection, which may be configured to generate chaotic velocity profiles [Cola (2004); Lester *et al.* (2009)]. Discontinuities can be introduced via the closing and opening of valves during syringe reorientation, required to ensure streamline crossing and chaotic trajectories produced [Hertzsch *et al.* (2007); Jones & Aref (1988)]. Streamline jumping has been known to destroy dynamical features [Lester *et al.* (2009)], or in the presence of no-slip boundary condition create pseudo-elliptic and pseudo-hyperbolic periodic points in the cut-dominated limit [Smith *et al.* (2016)].

As briefly mentioned in Section 2.6.1, granular materials also exhibit stirring mechanisms of both stretching and folding and cutting and shuffling. Piecewise isometries have been shown to capture the underlying structure in spherical tumbler flows [Juarez *et al.* (2012); Sturman *et al.* (2008); Yu *et al.* (2019)], with deviations in experimental models occurring due to the material passing through the flowing layer or diffusive-like effects from particle–particle interaction. Work has been done to understand the bifurcations that occur in the models when introducing shearing from the flowing layer into the piecewise isometry. In a model with reduced granular material height, stretching and folding is present even in the no shearing limit, and new bifurcations have been observed when cutting lines meet elliptic islands [Smith *et al.* (2017c)].

Another mechanism which could introduce discontinuities in fluid mixing is from the underlying properties of the material being stirred. High strain in polymeric, plastic or metallic material may cause slip deformations due to shear banding [Louzguine-Luzgin *et al.* (2012)].

Even though the applications are numerous, there a few studies on the interplay between stretching, folding, cutting and shuffling and diffusion. The next two sections cover in detail analytical and numerical studies investigating the rates of stirring and mixing in systems with a discontinuous mechanism to particle transport.

2.6.3 Dynamics of mixed SF and CS systems

There are many applications in which a combination of SF and CS stirring mechanisms arise but comparatively little analytical or computational studies investigating the impact of discontinuous fluid transport on mixing rates. Sturman (2012) provides a substantial review on the necessary background and applications in which discontinuities play a role in fluid mixing.

Vaienti (1992) rigorously proved the sawtooth map is ergodic, a discontinuous area-preserving automorphism of the unit torus onto itself. However, the proof is far from trivial and employs the work of Katok *et al.* (1985), who generalised Pesin theory to maps with singularities. Byott *et al.* (2013) showed that for one-dimensional expanding maps composed with permutations of equal sized cells, permutations do not improve the rate of decay of correlations, but typically make it worse. For a given stretch rate m, whether permutations from S_N have the potential to cause a decomposition in the composed map is proven by combinatorial methods. The strong mixing rates were explicitly computed from probability transition matrices between Markov states, which are the isolated eigenvalues of the maps transfer operators. Byott *et al.* (2016) extended the study to include all combinations of expanding maps with inverted branches.

Jones & Aref (1988) recognised that discontinuous fluid transport arises in Pulsed sink–sources flows from fluid extraction and re-injection. There have been numerous other studied on sink–source based flows considering the particle transport invoked by configuring the systems to generate chaotic advection, however none of the studies commented on the effect of the discontinuities [Cola (2004); Cola *et al.* (2006); Hertzsch *et al.* (2007)].

Smith *et al.* (2016) recently investigated the dynamics of a Reoriented Potential Mixing (RPM) flow, where a free-slip boundaries created discontinuities in

SF & D	SF & CS	CS & D
Pierrehumbert (1994)	Viana (2006)	Ashwin *et al.* (2002)
Antonsen Jr *et al.* (1996)	Byott *et al.* (2013, 2016)	Wang & Christov (2018)
Fereday *et al.* (2002)	Smith *et al.* (2016)	
Sukhatme & Pierrehumbert (2002)	Smith *et al.* (2017b)	
Wonhas & Vassilicos (2002)	Smith *et al.* (2017c)	
Gilbert (2006)		
Popovych *et al.* (2007)		
Schlick *et al.* (2013)		

Table 2.1: Key literature is listed which investigates fluid transport and mixing rates in systems with combined mixing dynamics of stretching and folding (SF), cutting and shuffling (CS) and molecular diffusion (D).

fluid transport. Analysis of a simplified map, which they called the Cut-Shear-Shear (CSS) map, revealed how changing the direction of the slip deformation and shears caused variation in the transport dynamics, all of which featured in different regions within the full RPM flow model. Novel periodic points and bifurcation transitions were identified which would be impossible in standard Hamiltonian systems of chaotic advection. Tracing the images and pre-images of cutting lines assisted in identifying these periodic points named pseudo-elliptic and pseudo-hyperbolic. The impact on the rates of mixing in the CSS map were investigated by Smith *et al.* (2017b) measuring the decay of the mix norm, and found that rate of decrease in material segregation was slowed from the uniform hyperbolic case when the slip deformation was present. Additionally, exotic bifurcations were identified in idealised granular flows which consisted of combined dynamics from SF and CS [Smith *et al.* (2017c)], where contact from cutting lines eliminates invariant tori surrounding elliptic periodic points.

2.6.4 Dynamics of mixed CS and D systems

In addition to systems of combined SF and CS, there has been a handful of studies considering the combined dynamics of CS and D. Recall the decay of variance in (2.4), which states that an increase in average gradients will generate a quicker decay of variance. Although CS will not increase gradients through shearing, sharp interfaces can be introduced into the concentration field from discontinuous transformations, and so increase the speed of mixing. This was confirmed by Ashwin *et al.* (2002), who showed that permutations of equal sized cells on the unit interval, a subset of Interval exchange transformations (IET),

although themselves are not strong or weak mixing, accelerate the asymptotic mixing rate of diffusion acting alone. It was observed that the ordering of the permutations which achieve the quickest mixing was dependent on the rate of diffusivity. IETs with diffusion have been used as toy models to investigate mixing rates and optimizing diffusional schemes [Froyland *et al.* (2016); Sturman (2012)], but there was no investigation of mixing rates in a larger parameter space.

Wang & Christov (2018) studied the average mixing rates across groups of IETs with a diffusional step, and concluded a universal stretched exponential in the decay of variance emerges on rescaling with a mixed time. The study was inspired by the work of Liang & West (2009), who looked for cut-offs in chaotic advective mixing, a phenomenon seen in card shuffling [Aldous & Diaconis (1986)]. The variation distance is a measure used in probability which can be used to specify how far away a deck of cards is from a perfectly randomised deck during a shuffling process. A cut-off describes a sharp transition in variation distance and as an such could predict a potential stopping time. Wang & Christov (2018) questioned whether a similar phenomenon is observed in mixing by IETs and could be used as a proposed stopping time on stirring. They conclude that the system presents evidence of cut-offs, however this is not clearly evident from their results. A cut-off is described as a relative quicker transition to a well mixed state when decreasing the diffusivity rate. The trend in the rescaled variance profiles of Wang & Christov (2018) vary qualitatively from both Aldous & Diaconis (1986) and Liang & West (2009), and do not show the behaviour expected of cut-offs. The main problem is that the rescaling of the profiles does not see an agreement between the respective stopping times, which is needed to assess whether a steeping in the transition between a not well mixed state, and well mixed state occurs. A more rigorous discussion on Wang & Christov (2018) is carried out in Chapter 4.

Table 2.1 lists the key literature investigating the combination of mixed dynamics in chaotic advective systems. The combination of chaotic advection and diffusion (SF & D) has received much attention, as outlined in Section 2.5, and as such a wealth of literature is omitted and the named studies are those which are particularly relevant to the current body of work. The combined dynamics of SF & CS has a handful of studies associated to understanding the dynamic and ergodic properties in such mixed systems, however CS & D, and additionally SF, CS & D has received little or no attention. This latter fact is the main motivation of the work carried out herein, in which the role of molecular diffusion on the

rates of mixing to uniformity are investigated in discontinuous transformations.

Chapter 3

Fundamental models of mixing by chaotic advection

This Chapter outlines the use of dynamical maps to study mixing in chaotic advective systems where the velocity field is continuous. Discrete time-periodic maps can be devised which capture the essence of fluid mixing, where stretching and folding lead to good mixing while boundaries and separated domains contaminate and increase the time to homogenization. Discrete maps allow for intuitive understanding of the effect specific dynamics have on mixing due to their simplicity. In more complicated models deciphering which flow field dynamics contribute to the observed behaviours is more difficult. Methodologies for constructing maps, including diffusion, and analysing mixing dynamics are introduced which will be repeatedly used in the current work.

3.1 Discrete maps as models for mixing

3.1.1 Toral Automorphisms

A selection of toral automorphisms are presented which capture a range of dynamics in which stretching and folding are the key mechanism for chaotic advection. Methods of analysis are introduced throughout which will be used within the chapters which follow.

Toral automorphisms are paradigm examples of two-dimensional maps which can exhibit chaotic dynamics, namely exponential separation of nearby particles. Non-uniformity in the stretching dynamics is added via a small wave-perturbation. Since discrete time intervals describe the stirring of fluid parti-
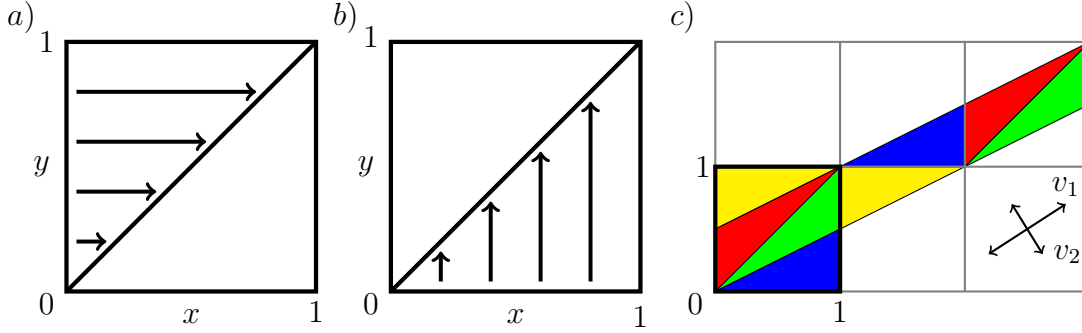
Figure 3.1:   The action of two shear maps are shown in $a)$ and $b)$, independent of y and x respectively. The composition of the two shears, applying $b)$ followed by $a)$, results in the cat map. The action on the unit square of the cat map is shown in $c)$ with a representation of reassembly onto the unit torus.

cles, diffusion is introduced iteratively by solving the heat equation after each time step. Implementing periodic boundary conditions means that the evolution of a concentration field is computed using discrete Fourier coefficients without numerical diffusion. Dynamical metrics such as Poincaré maps and Lyapunov exponents, and how they can be used to determine observations in mixing systems, are discussed.

The unit torus is defined by the projection $\pi : \mathbb{R}^2 \to \mathbb{T}^2$ such that

$$\pi(x + m, y + n) = \pi(x, y) \text{ for all } (m, n) \in \mathbb{Z}^2. \tag{3.1}$$

Let A be an integer matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \quad \text{and let } M \begin{pmatrix} x \\ y \end{pmatrix} = A \begin{pmatrix} x \\ y \end{pmatrix} \quad \text{mod } 1. \tag{3.2}$$

then $M : \mathbb{T}^2 \to \mathbb{T}^2$. Provided $\det A = \pm 1$, A^{-1} will also be an integer matrix. Maps $M : \mathbb{T}^2 \to \mathbb{T}^2$ in which A and A^{-1} are both integer matrices are called *Toral Automorphisms*. All toral automorphism preserve Lebesgue measure since they are invertible and have $\det A = 1$. A toral automorphism will have characteristic polynomial for the eigenvalues of the matrix, given by $\chi(\lambda) = \lambda^2 - t_A \lambda + 1$, $t_A = \text{trace} A$, so the eigenvalues are given by $\lambda = (t_A \pm \sqrt{t_A^2 - 4})/2$ with t_A an integer. The eigenvalues determine the behaviour of the toral automorphism. For $t_A \geq 3$, or $t_A \leq -3$ A has two real eigenvalues λ_1 and λ_2 with $|\lambda_1| > 1$ and $|\lambda_2| < 1$, resulting in hyperbolic behaviour, with a stretching and contracting direction.
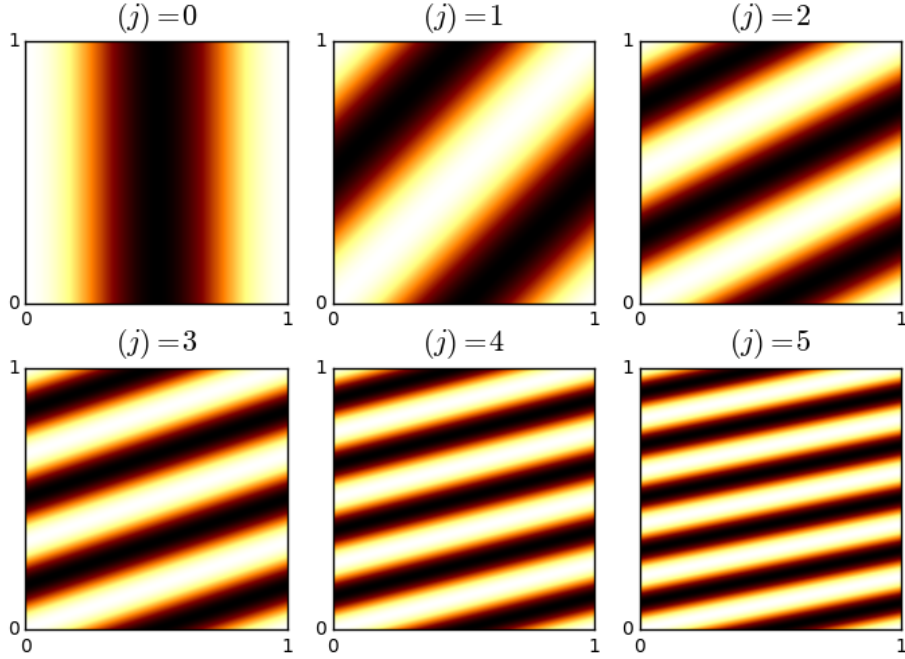
Figure 3.2: The evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ by the shear map (3.3). Striations are compressed uniformly but at a rate which is subexponential.

The simplest toral automorphism is the shear map $M_S : \mathbb{T}^2 \to \mathbb{T}^2$, which has

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \text{ so that } \quad M_S \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + y \\ y \end{pmatrix}. \tag{3.3}$$

Figure 3.1 a) depicts the action of the map on the points x and y, where for all points the y value remains fixed and they move in lines parallel to the x axis on repeated iteration. The shearing action of the map results in stretching of fluid elements. This is seen when showing the evolution of a concentration field.

A simple initial concentration field $c^{(0)}(x, y)$ can be used as an example to show how a map M might evolve typical concentration fields. First, represent $c^{(0)}(x, y)$ on a grid of evenly spaced points x_i, y_i. If x_i^j, y_i^j denotes the values at the jth iteration, then the concentration field at iteration (j) can be plotted by working out $c^{(j)}(x_i, y_i) = c^{(0)}(x_i^{-j}, y_i^{-j})$, where x_i^{-j}, y_i^{-j} are worked out from the inverse inverse map $M^{-1}(x, y)$. Figure 3.2 shows the evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$, which is sheared and stretched by the map. Although this map is stretching the concentration field, it is doing so at a rate which is subexponential. The shear map is neither mixing or even chaotic, reflected in the eigenvalues of the shear map which are both $\lambda_1 = \lambda_2 = 1$, and no hyperbolicity is present.
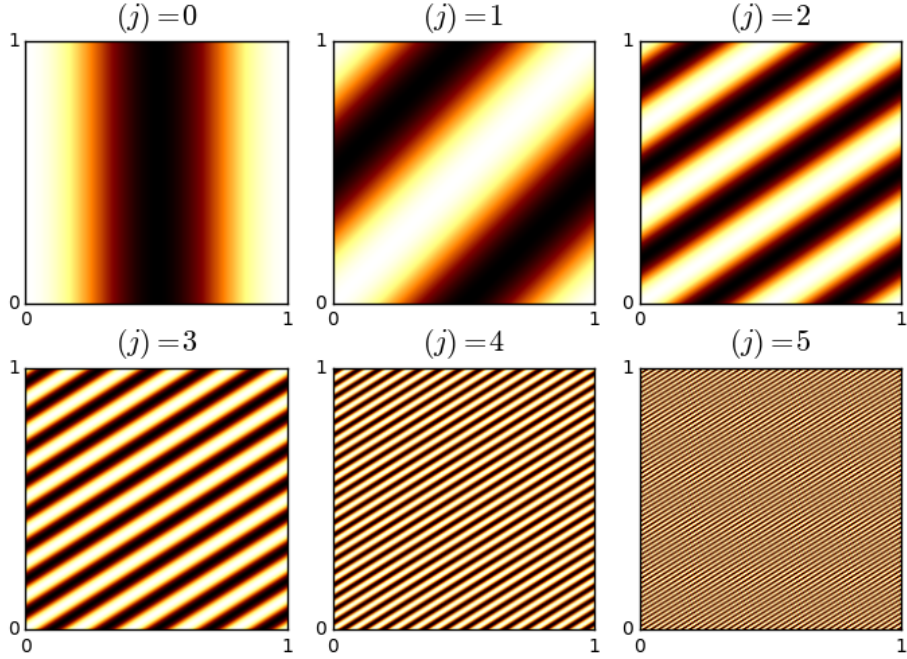
Figure 3.3: The evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ by the cat map (3.4), a composition of two shears. Striations are compressed uniformly with an exponential rate of compression.

The composition of the two shears depicted in Figure 3.1 $a)$ and $b)$, is called Arnold's cat map. The cat map M_C is an example of a toral automorphism which is hyperbolic and as such captures the essence of good mixing, defined by

$$M_C \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \mod 1 = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \mod 1. \quad (3.4)$$

The action of the cat map on the unit torus is shown in Figure 3.1 $c)$, before it is reassembled by periodic boundary conditions. The evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ by the action of the cat map is depicted in Figure 3.3, and in comparison to the shear map, stretches the initial condition at a much faster rate. This is because the cat map is hyperbolic; in that it has exponential separation of infinitesimally nearby trajectories. This is due to the hyperbolicity of the cat map, which has eigenvalues $\lambda_1 = (3 + \sqrt{5})/2$ and $\lambda_2 = (3 - \sqrt{5})/2$, and respective eigenvectors v_1 and v_2 shown in Figure 3.1 $c)$ showing the directions of stretch and contraction.

Compositions of shears have been considered as paradigm models of stretching and folding dynamics in mixing. Sturman $et\ al.$ (2006) consider how co-rotating or counter-rotating effects of shears describe the dynamics on co-twisting or counter-
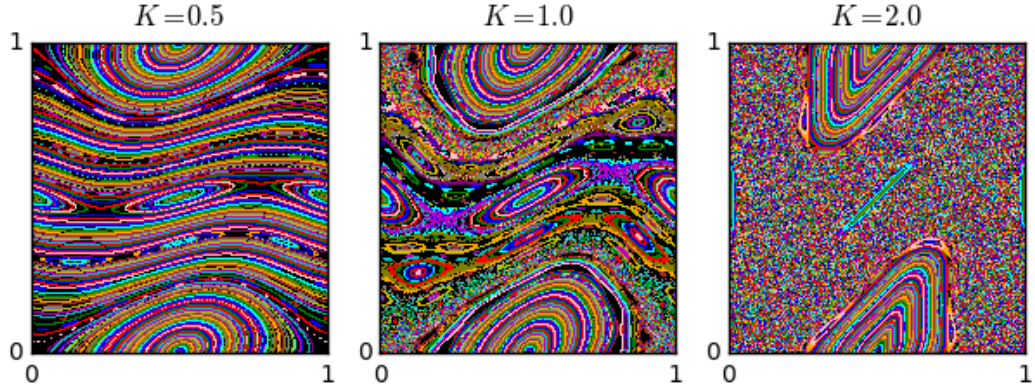
Figure 3.4: Poincaré sections for the Standard map where varying colours distinguish the trajectories of multiple initial conditions along the line $(x = 1/2, y)$ for 1000 time steps. On increasing the effect of the wave perturbation, for large K a chaotic sea emerges with elliptic islands.

twisting link-twist maps. The action of counter-twists can lead to un-stirring dynamics. Additionally, Sturman & Springham (2013) use a composition of shears in a link twist map with parabolic boundaries to prove the decay of correlations is polynomial at best in the presence of walls. D'Alessandro *et al.* (1999) investigate the optimization of entropy as a potential measure for mixing, and conclude that simple periodic application of counter-rotating shears gives the maximum measure of entropy over randomised applications. This periodic composition is simply repeated applications of the cat map.

The failing of these maps is the uniformity in the stretching rates, which is not applicable to realised physical mixing scenarios. Non-uniformity can be introduced by adding a wave perturbation such that

$$
M \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \frac{K}{2\pi} \begin{pmatrix} \sin(2\pi x) \\ \sin(2\pi x) \end{pmatrix} \quad \text{mod } 1. \tag{3.5}
$$

For the shear map, this results in the standard map, also known as the Chirikov-Taylor map, defined

$$
M_{PS} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + y + \frac{K}{2\pi} \sin(2\pi x) \\ y + \frac{K}{2\pi} \sin(2\pi x) \end{pmatrix} = M_S(\boldsymbol{x}) + \omega(\boldsymbol{x}), \quad \omega(\boldsymbol{x}) = \frac{K}{2\pi} \begin{pmatrix} \sin(2\pi x) \\ \sin(2\pi x) \end{pmatrix}. \tag{3.6}
$$

Similarly the cat map with a wave perturbation can be defined as $M_{PC}(\boldsymbol{x}) = M_C(\boldsymbol{x}) + \omega(\boldsymbol{x})$. This latter map was used to investigate mixing in non-uniformly stretching systems by Thiffeault & Childress (2003).

49
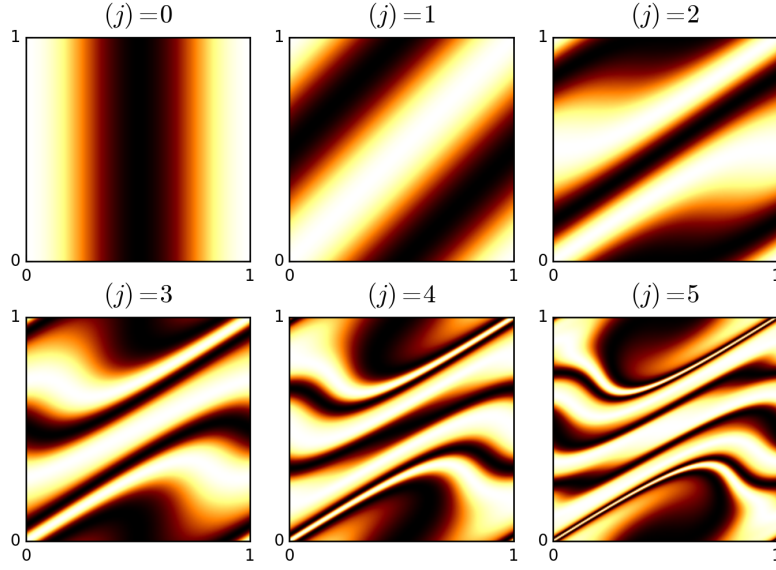
Figure 3.5: The initial concentration field $c^{(0)}(x, y) = \cos(2\pi x)$ is evolved via the standard map with $K = 1.0$. The wave perturbation introduces non-uniformity in the width of striations, and within regions aligning with elliptic islands in the Poincar'e maps, limited filamentation occurs.

3.1.2 Poincaré sections

Plotting long trajectories of points under a mapping M will highlight structures in the underlying flow. These plots are known as Poincaré sections, and they are useful in studying mixing systems since non-mixing regions will appear as islands or decompositions of the domain, while a chaotic sea, in which trajectories appear to fill the state space, suggest good mixing may be present.

Figure 3.4 shows multiple trajectories initialised along the line $(x = 1/2, y)$ for the standard map (3.6). For $K = 0$ the trajectories lie along straight lines parallel to the x-axis. Adding the wave perturbation causes the lines to bend, in some instances generating chains of elliptic islands, apparent in the Poincaré sections when $K = 0.5$. As K increases, regions of chaotic trajectories between elliptical chains emerges, shown for $K = 1.0$, until finally a chaotic sea appears with several elliptic islands, depicted for $K = 2.0$. Figure 3.5 shows the evolution of $c^{(0)}(x, y) = \cos(2\pi x)$ by the standard map with $K = 1.0$, where elliptic islands have restricted regions of the concentration field from filamenting, while the chaotic sea between elliptic chains have resulted in a large increase in gradients in the concentration field due to exponential stirring dynamics there.

Poincaré section for the cat map with added wave perturbation are shown in
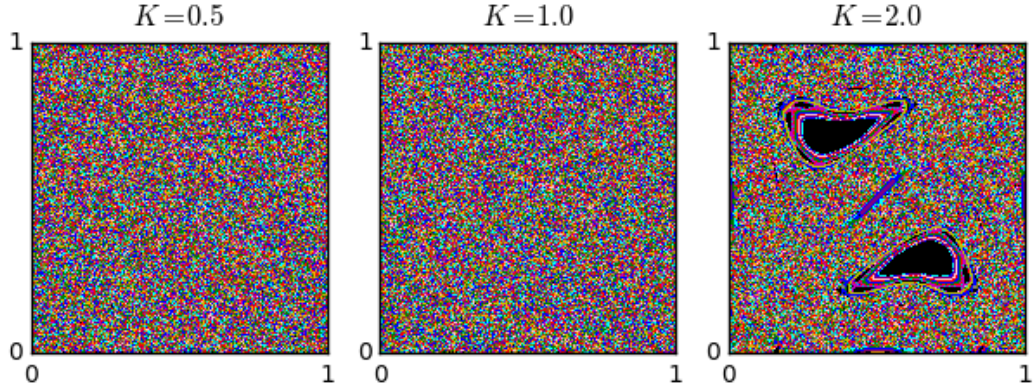
Figure 3.6: Poincaré sections for the wave perturbed cat map where varying colours distinguish the trajectories of multiple initial conditions along the line $(x = 1/2, y)$ for 1000 time steps. For small values of K the cat map remains chaotic however for $K > 1.5$ elliptic island emerge within the chaotic sea.

Figure 3.6. When $K = 0$ the map is fully chaotic and a chaotic sea fills the unit torus. As K is increased the transformation remains fully chaotic, observed when $K = 0.5$ and 1.0, until finally when the parameter K is large enough, elliptic islands emerge within the chaotic sea, shown for $K = 2.0$.

In Figure 3.6 there is no qualitative difference between the Poincaré sections for $K = 0.5$ and 1.0, with trajectories filling the whole of the domain $\mathbb{T}^2$, but looking at the evolution of $c^{(0)}(x, y) = \cos(2\pi x)$ in Figure 3.7 when $K = 1.0$, the concentration field shows variations in the striation widths as it is being stretched.

Although Poincaré sections show that trajectories can densely cover the domain, they reveal only long-time dynamics of trajectories and give no quantitative information about the transformation. Lyapunov exponents, and the distribution of finite time Lyapunov exponents (FTLEs), provided quantitative and qualitative measures of stretching rates in a dynamical system and are described in the next section.

3.1.3   Lyapunov exponents

Lyapunov exponents quantify the exponential separation of infinitesimally nearby trajectories within a dynamical system and are typically used to determine if a dynamical system is chaotic. Stretching histories and short-time evolution of nearby trajectories are captured by finite-time Lyapunov exponents (FTLEs).
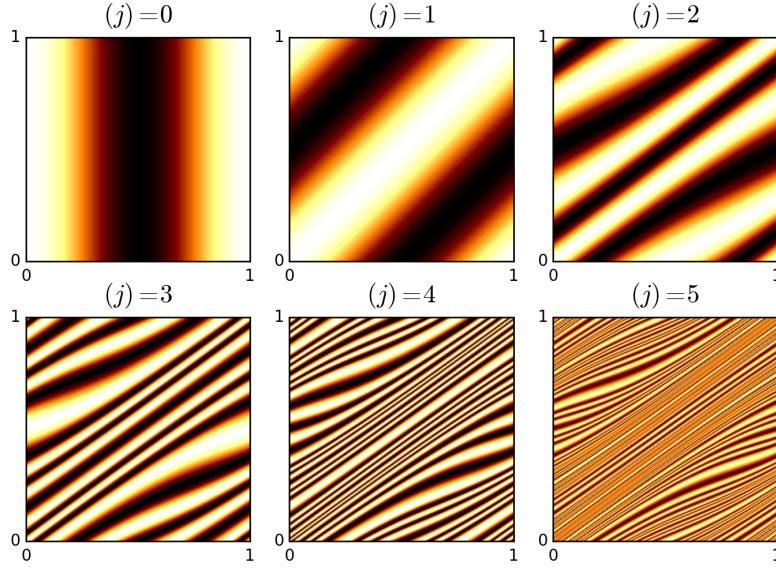
Figure 3.7: The initial concentration field $c^{(0)}(x, y) = \cos(2\pi x)$ is evolved via the wave perturbed cat map with $K = 1.0$. The hyperbolicity of the cat map is retained under perturbation, with efficient filamentation after just a handful of iterations. Non-uniformity in the stretching rates produces uneven compression of lamellar throughout the domain.

For an initial location $\boldsymbol{x}_0$ and an initial perturbation $\delta\boldsymbol{x}_0$ the evolution the Lyapunov exponent is defined as

$$h(\boldsymbol{x}_0, \delta\boldsymbol{x}_0) = \lim_{t\to\infty} \lim_{\delta\boldsymbol{x_0}\to\boldsymbol{0}} \frac{1}{t} \ln\left(\frac{||\delta\boldsymbol{x}_t||}{||\delta\boldsymbol{x_0}||}\right) \tag{3.7}$$

where $\delta\boldsymbol{x}_t$ denote the deformation of the perturbation after time t.

A similar definition applies for discrete maps M

$$h(\boldsymbol{x}_0, \delta\boldsymbol{x}_0) = \lim_{n\to\infty} \frac{1}{n} \lim_{\delta\boldsymbol{x_0}\to\boldsymbol{0}} \sum_{k=0}^{n-1} \ln\frac{||\delta\boldsymbol{x}_k||}{||\delta\boldsymbol{x}_0||} \tag{3.8}$$

The deformation of the element $\delta\boldsymbol{x}$ can be thought of as the local deformation. The infinitesimal limit of the local deformation is the Jacobian determinant $\boldsymbol{D}f$ (2.8) at the point along the trajectory.

For a N-dimensional system there are N Lypaunov exponents defined for each point, typically however interest lies with the Lyapunov exponent value for the fastest stretching direction, called the maximal Lyapunov exponent. A chaotic dynamical system will have a positive maximal Lyapunov exponent. For an area or volume preserving system, the Lyapunov exponents of a trajectory will sum to 0. For $\boldsymbol{x}_0$ which belongs to an ergodic partition $\mathcal{B}$ say, almost all $\boldsymbol{x}_0 \in \mathcal{B}$ will
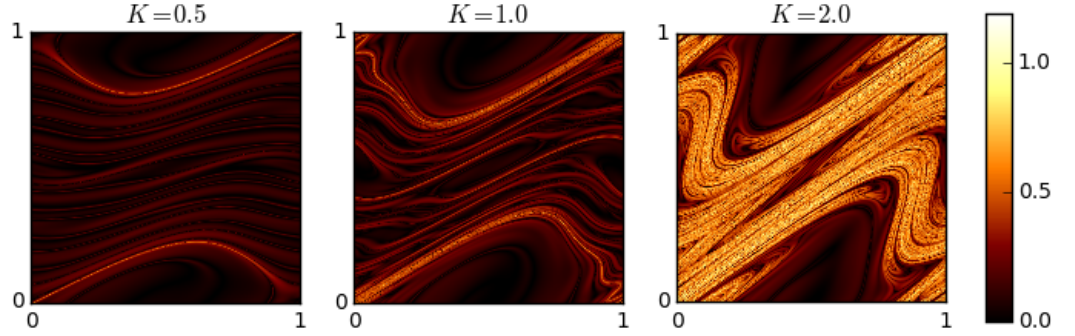
Figure 3.8: Stretching histories along element trajectories are shown by finite time Lyapunov exponents (FTLEs) averaged over 10 iterates for the Standard map with varying K. Elliptic islands appear as dark islands while the chaotic sea shows regions with low and high stretch.
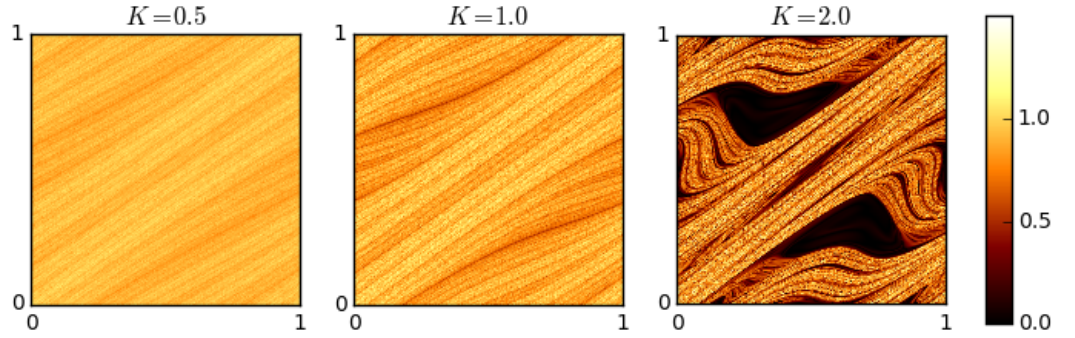
Figure 3.9: Stretching histories along element trajectories are shown by finite time Lyapunov exponents averaged over 10 iterates for the cat map with added wave perturbation with varying K.

have $\sigma_1, \ldots \sigma_N$ equivalent. This is due to Birkhoff's ergodic theory and which states that time averages are equivalent to spatial averages which results in

$$h(\boldsymbol{x}) = \lim_{n \to \infty} \frac{1}{n} \sum_{k=0}^{n-1} \ln ||\boldsymbol{D}f(\boldsymbol{x})|| = \int \ln ||\boldsymbol{D}f(\boldsymbol{x})|| d\mu(\boldsymbol{x}). \tag{3.9}$$

where $|| \cdot ||$ denotes the matrix norm, and $\boldsymbol{D}f$ denotes the Jacobian determinant matrix (2.8). Thus within an ergodic partition, finding the Lyapunov exponents along a single trajectory is enough to calculate the Lyapunov exponent for that ergodic partition. Lyapunov exponents can be computed numerically via the method of Parker & Chua (2012). This method is described in detail in Appendix A.

Since for almost all points within an ergodic partition the Lyapunov exponents will be the same they do not highlight regions which may experience low or high stretch locally. Finite time Lyapunov exponents can be calculated for individual points within an ergodic partition by computing (3.8) can be iterated for a specified iterations number T. However, to observe the specific stretching history of a given trajectory we need the maximal finite time Lyapunov exponent. This can be calculated as

$$h^T(\boldsymbol{x}_0, \delta\boldsymbol{x}_{-T}) = \frac{1}{T}\sum_{k=1}^{T}\ln\frac{||\delta\boldsymbol{x}_{-k}^1||}{||\delta\boldsymbol{x}_0||}, \qquad (3.10)$$

where computing in backward time finds the average of the finite time stretching histories for the point of interest $\boldsymbol{x}_0$. For area-preserving maps, h_1 and h_1^T are always zero or positive in value. For a grid of chosen points (x_i, y_i) in the domain $\mathbb{T}^2$, Figures 3.8 and 3.9 plot the maximal finite time Lyapunov exponents for each point, averaged over $T = 10$ backwards iterations for the standard map and the cat map with added wave perturbation respectively. The same values of K are used as those of the Poincaré sections in Figure 3.4 and 3.6 for comparison.

For the standard map in Figure 3.8, when K is small, only small regions between chains of elliptic islands experience large values of h_1^T. As K is increased, there is an increase in regions which experience high values of stretch in the trajectory history, until $K = 2.0$ where the chaotic sea shows a good distribution of $h_1^T > 0$. The stretching histories surrounding the elliptic islands are slow in comparison to regions far from the elliptic islands in the chaotic sea.

For the cat map, qualitative differences in the stretching behaviours is seen between $K = 0.5$ and $K = 1.0$ which are not seen in the fully chaotic Poincaré sections. As K is increased, regions of lower stretch emerge as bands throughout the domain, aligned with the fastest eigenvector of the cat map with some small perturbation about it. Comparing with Figure 3.7 which shows evolution of $c^{(j)}(x) = \cos(2\pi x)$, striations are larger where the finite time Lyapunov exponents are the lowest, highlighting regions of low stretch and compression which are reflected in the evolution of striation widths in the concentration field.

Now that the dynamics of the maps in the absence of diffusion have been established, attention is now directed at including diffusion and observing how the dynamics effect mixing rates to the uniform distribution when full advection-diffusion iterative operator applied for $\kappa > 0$.

54

3.2 Discrete chaotic maps with diffusion

The advantage of the maps presented in Section 3.1.1 is that the discrete advective-diffusive transfer operator can be derived easily using the methods of Section 2.3.4 via discrete Fourier expansion of the concentration field for periodic boundary conditions.

3.2.1 Analytic solutions of idealised chaotic mixing

Shear map

Solving for the transfer matrix (2.34) for the shear map (with no wave perturbation) (3.3) gives

$$\boldsymbol{DM_{kq}} = e^{-4\pi^2 \kappa \boldsymbol{k} \cdot \boldsymbol{k}} \int_{[0,1)} \exp(2\pi i Q_1 x)\mathrm{d}x \cdot \int_{[0,1)} \exp(2\pi i Q_2 y)\mathrm{d}y \qquad (3.11)$$

where $Q_1 = k_1 - q_1$ and $Q_2 = k_1 + k_2 - q_2$. Since $Q_1, Q_2 \in \mathbb{Z}$ the integral is periodic and equates to zero unless $Q_1 = Q_2 = 0$, then

$$\boldsymbol{DM_{kq}} = e^{-4\pi^2 \kappa \boldsymbol{k} \cdot \boldsymbol{k}} \delta_{0,Q1} \delta_{0,Q2}. \qquad (3.12)$$

The wavemodes which exchange concentration can be written as an iterative map

$$\begin{aligned} q_1^{(j+1)} &= q_1^{(j)} \\ q_2^{(j+1)} &= q_2^{(j)} - q_1^{(j)}. \end{aligned} \qquad (3.13)$$

From this iterative relation on wave modes, the decay of variance can be worked out directly. If for a given $\boldsymbol{q}^{(0)} = (q_1^{(0)}, q_2^{(0)})$, $\psi_{\boldsymbol{q}^{(j)}} = |\hat{c}_{\boldsymbol{q}}^{(j)}|^2$, then

$$\psi_{\boldsymbol{q}^{(1)}} = \psi_{\boldsymbol{q}^{(0)}} \times \exp(-4\pi^2 \kappa \boldsymbol{q}^{(1)} \cdot \boldsymbol{q}^{(1)})^2. \qquad (3.14)$$

Then it follows that

$$\psi_{\boldsymbol{q}^{(j)}} = \psi_{\boldsymbol{q}^{(0)}} \times \prod_{i=1}^{j} \exp\left(-4\pi^2 \kappa \boldsymbol{q}^{(i)} \cdot \boldsymbol{q}^{(i)}\right)^2 = \psi_{\boldsymbol{q}^{(0)}} \times \exp\left(-8\pi^2 \kappa \sum_{i=1}^{j} \boldsymbol{q}^{(j=i)} \cdot \boldsymbol{q}^{(i)}\right), \qquad (3.15)$$

and the total variance at each iteration is

$$\psi(j) = \sum_{\boldsymbol{q}} \psi_{\boldsymbol{q}^{(j)}}. \qquad (3.16)$$

For a y independent initial condition $c^{(0)}(x, y) = \cos(2q_1\pi x)$, the iterative equation (3.13) on the wavemodes reduces to $q_1^{(j)} = q_1^0$ and $q_2^{(j)} = -j \cdot q_1^0$ and $\psi(j)$ and is written analytically as

$$\psi(j) = \psi(0) \exp\left(-8\pi^2\kappa \sum_{n=1}^{j} ((q_1^0)^2 + (-n \cdot q_1^0)^2)\right) \tag{3.17}$$

$$\psi(j) = \psi(0) \exp\left(-8\pi^2\kappa |q_1^0|^2 \sum_{n=1}^{j} (1 + n^2)\right). \tag{3.18}$$

Faulhaber's formula expresses the sum of squares of the first n positive integers as

$$\sum_{k=1}^{n} k^2 = \frac{n^3}{3} + \frac{n^2}{2} + \frac{n}{6}. \tag{3.19}$$

Hence the variance decay (3.18) can be written as,

$$\psi(j) = \psi(0) \exp\left(-8\pi^2\kappa |q_1^0|^2 \mathcal{P}(j)\right), \tag{3.20}$$

where the polynomial $\mathcal{P}(j) = \frac{j^3}{3} + \frac{j^2}{2} + \frac{7j}{6}$. Therefore for the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ when $q_1 = 1$, the decay of variance can be computed explicitly. For this initial condition, the decay of variance profiles from numerical computations for different κ are plotted in Figure 3.10, and (3.20) agrees exactly. As κ goes to zero, the polynomial growth in the transfer of wavemodes means that the rate of mixing becomes very slow. Although the map is stretching, the shear map is slow at reducing the scale of the concentration field to a scale in which diffusion is effective and the map is not an effective stirring protocol.

Cat map

Solving for the transfer matrix for the cat map results in

$$\boldsymbol{DM_{kq}} = e^{-4\pi^2\kappa \boldsymbol{k}\cdot\boldsymbol{k}} \delta_{0,Q1}\delta_{0,Q2}, \tag{3.21}$$

with $Q_1 = 2k_1 + k_2 - q_1$ and $Q_2 = k_1 + k_2 - q_2$, $Q_1, Q_2 \in \mathbb{Z}$. An iterative equation for the wavemodes can be written as

$$\begin{pmatrix} q_1^{(j+1)} \\ q_2^{(j+1)} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} q_1^{(j)} \\ q_2^{(j)} \end{pmatrix} \tag{3.22}$$

As the map is iterated forward the wavemodes align with the eigenvector corresponding to the stretching direction of the matrix in (3.22) multiplied by the maximal eigenvalue σ_{max} at each iteration such that for a given $\boldsymbol{q}^{(0)}$

$$||\boldsymbol{q}^{(j)}|| \approx ||\boldsymbol{q}^{(0)}|| \cdot \sigma_{max}^j. \tag{3.23}$$
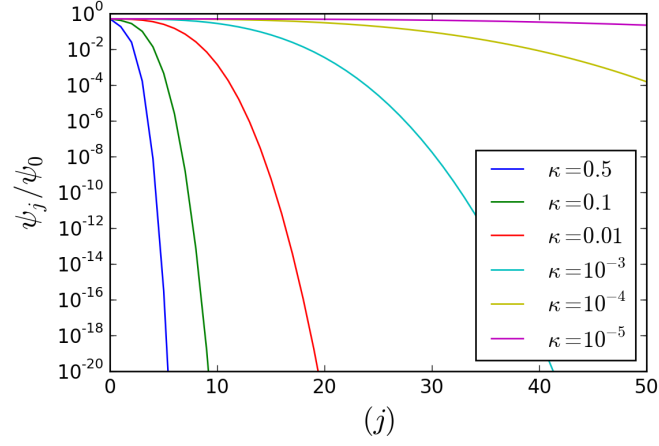
Figure 3.10: Profiles for the decay of variance in the shear map with no wave perturbation and a range of diffusivity coefficients κ. The initial condition for all profiles was $c^{(0)}(x, y) = \cos(2\pi x)$. The decay has a stretched exponential profile, the rate of which slows down as $\kappa \to 0$.
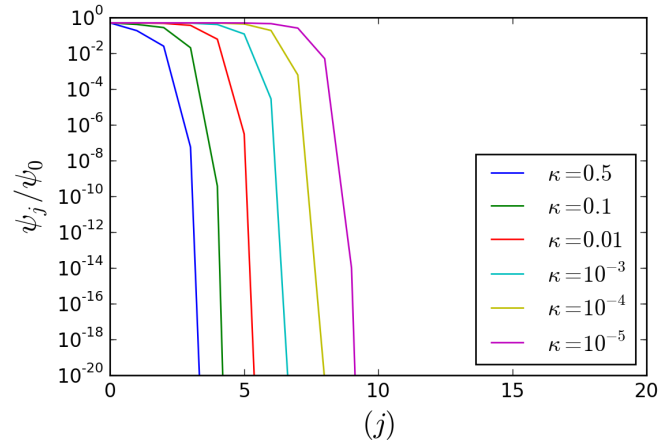
Figure 3.11: Profiles for the decay of variance in the cat map with no wave perturbation and a range of diffusivity coefficients κ. The initial condition for all profiles was $c^{(0)}(x, y) = \cos(2\pi x)$. Super-exponential decay emerges for all values of κ.

This means the variance can be approximated by

$$\psi_{\boldsymbol{q}^{(j)}} \sim \psi_{\boldsymbol{q}^{(0)}} \cdot \exp\left(-8\pi^2 \kappa ||q_0||^2 \sigma_{max}^{2(j-1)}\right) \tag{3.24}$$

and the total variance given by (3.16). This leads to super-exponential decay in variance, seen in Figure 3.11 for the initial condition $c^{(0)}(x,y) = \cos(2\pi x)$. When κ is small, the variance does not deplete for the first few iterations until the concentration profile reaches a scale in which diffusion becomes more effective, at which the decay begins. The main effect of changing the diffusion coefficient is changing the scale at which diffusion begins to effectively reduce gradients in the concentration field. In Figure 3.11 the onset of super-exponential decay changes linearly with the values of κ, which vary by orders of magnitude, implying $j \propto -\log(\kappa)$, as discussed at the start of Section 2.5.1 and given by equation (2.52).
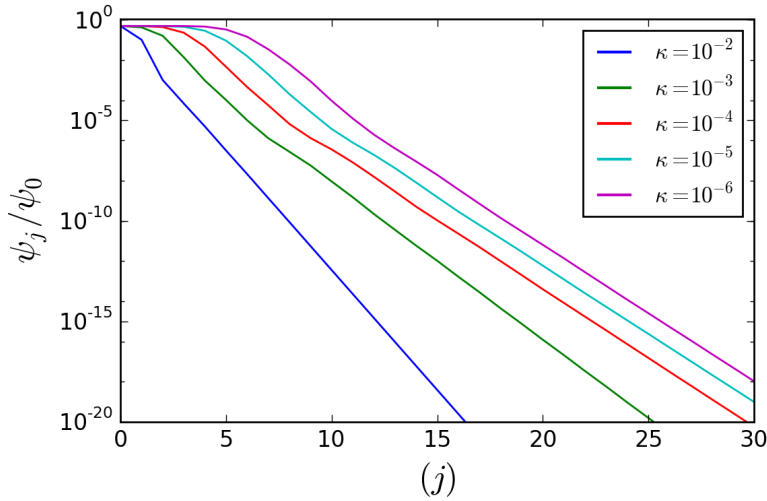
Figure 3.12: Profiles for the decay of variance in the cat map with wave-perturbation with strength $K = 1.0$ and a range of diffusivity coefficients κ. The initial condition for all profiles was $c^{(0)}(x,y) = \cos(2\pi x)$. The long-time mixing is exponential and as $\kappa \to 0$ this mixing rate becomes independent of κ.

3.2.2 Numerical solutions of non-uniform chaotic mixing

Wave perturbed shear and cat maps

When $K > 0$, the evolution of a concentration field, and the respective variance decay, is no longer simple enough to solve analytically. Instead, when adding the
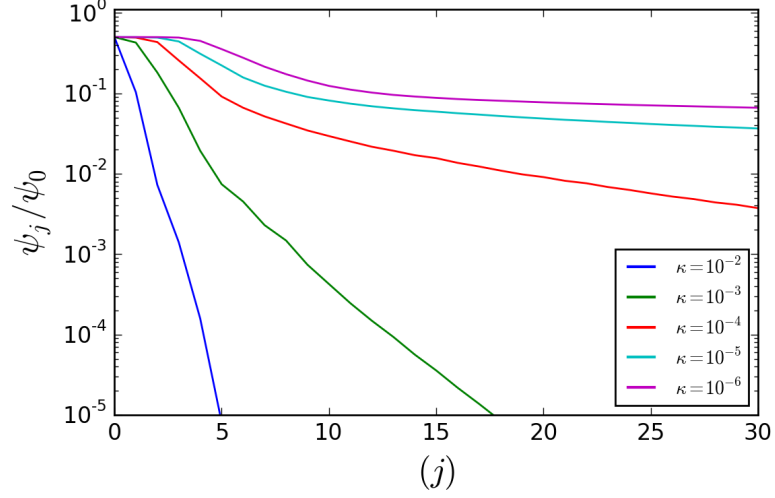
Figure 3.13: Profiles for the decay of variance in the cat map with wave-perturbation with strength $K = 2.0$ and a range of diffusivity coefficients κ. The initial condition for all profiles was $c^{(0)}(x, y) = \cos(2\pi x)$. Since islands are present in the flow, in the long-time limit the transport is restricted to diffusion across boundaries and as such the exponential mixing rate is not independent of κ as $\kappa \to 0$.

wave perturbation the transfer matrix $\boldsymbol{M}$ can be derived analytically for both the wave-perturbed shear and cat maps. This was done for the cat map in Thiffeault & Childress (2003). The transfer matrix relations between Fourier modes for all toral automorphism maps with the added wave-perturbation (3.5) and diffusion coefficient κ are given by

$$\boldsymbol{DM_{kq}} = \exp(-4\pi^2 \kappa \boldsymbol{k} \cdot \boldsymbol{k}) \delta_{0,Q_2} (-1)^{Q_1} J_{Q_1}((k_1 + k_2)K), \qquad \boldsymbol{Q} = \boldsymbol{k} \cdot A - \boldsymbol{q} \quad (3.25)$$

where A denotes the integer matrix for the toral automorphism. The derivation of the transfer matrix is in Appendix B.

The non-uniformity in the stretching rates causes a dispersion in the amplitude of a given wavemode to several wavemodes at each iteration, rather than a direct cascade, such as (3.13) and (3.22). This introduces complexity in the transfer of concentration through the wavemodes and the decay of variance can not be written down analytically. The evolution of a concentration field and its approach to the mean field is studied numerically only.

Figure 3.12 shows the decay of variance in the cat map with added wave perturbation for varying diffusivity rates. The strength of the perturbation is $K = 1.0$, which varies the stretching histories of trajectories significantly but does not introduce regular islands or barriers to mixing to the flow field, observed in Figure

3.6. The initial condition was $c^{(0)}(x,y) = \cos(2\pi x)$. The decay of variance ψ_j is shown of a linear-log axis, revealing long-time exponential mixing for all values of κ. Exponential mixing occurs due to the dispersion between the Fourier modes, resulting in the emergence of a slowly decaying eigenfunction in the asymptotic limit. Changing the initial condition does not change the exponential mixing rates. Although in two-dimensions it is not feasible to compute the eigenvalues and eigenfunctions for the operators $P_D \circ P_M$ when κ is small, the mechanism for the emergence of exponential mixing due to the emergence of the slowest decaying eigenfunction in the flow field remains the same.

On decreasing the effect of diffusion, the initial stages in which the variance does not decay much is extended, similar to the super-exponential mixing case when $K = 0$, however as $\kappa \to 0$ the asymptotic rate of mixing becomes independent of diffusivity.

When the strength of the perturbation is increased such that $K = 2.0$, elliptic islands appear and form barriers to mixing through stirring alone. This is reflected in Figure 3.13, where the long-time mixing rate does not become independent of κ, and instead the exponential mixing rate becomes incredibly slow as $\kappa \to 0$, due to the fact that $|\lambda_2| \to 1$ since the transfer operator P_M is no longer strong mixing. Although computation of λ_2 is computationally infeasible from the computational method employed herein for these two-dimensional systems, and approximation of $|\lambda_2|$ can be computed from the variance profile in these late stages.

The mixing in the long-time limit is restricted to diffusive transport across the edge of the islands. In the wave perturbed shear map, barriers to mixing are present for all values $K > 1.0$ and as such the long-time mixing behaviour is similar in nature and as such plots of the variance decay are not included for brevity.

There has been much discussion that the long-time mixing can be predicted from stretching histories alone [Antonsen Jr *et al.* (1996); Sukhatme & Pierrehumbert (2002)], however it is widely accepted that situations where this would be correct are rare in occurrence and the exponential mixing arises from the global mechanism of slowly decaying eigenfunctions in most cases. However comparisons of the distributions of stretching histories, captured in FTLEs, to the concentration field show that stretching histories can be used to predict the appearance of emerging eigenfunctions. Figures 3.14 and 3.15 show both the concentration field after 15 iterations when $\kappa = 10^{-5}$ and the FTLEs for the wave perturbed
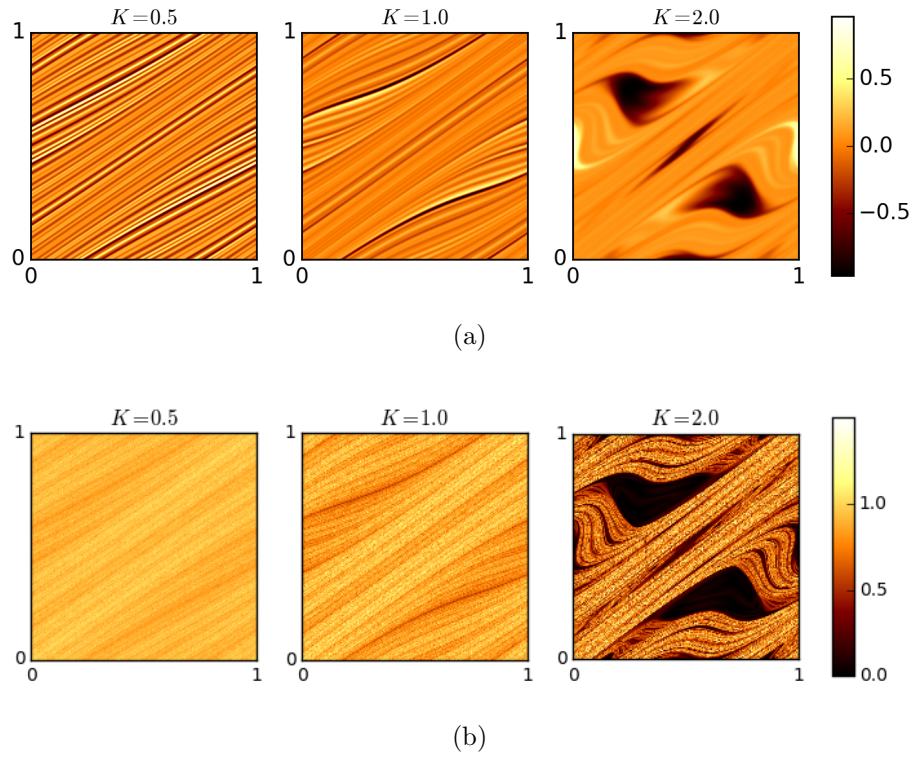
Figure 3.14: (*a*) Concentration fields after 15 iterations for the wave perturbed cat map with different wave perturbation strengths K. The initial condition was $c^{(0)}(x, y) = \cos(2\pi x)$ and $\kappa = 10^{-5}$. The concentration field can be compared with (*b*) the distribution of FTLEs in the domain, showing alignment of regions where little mixing has occurred, with areas of little or no stretch experienced by trajectories.

(a)

(b)
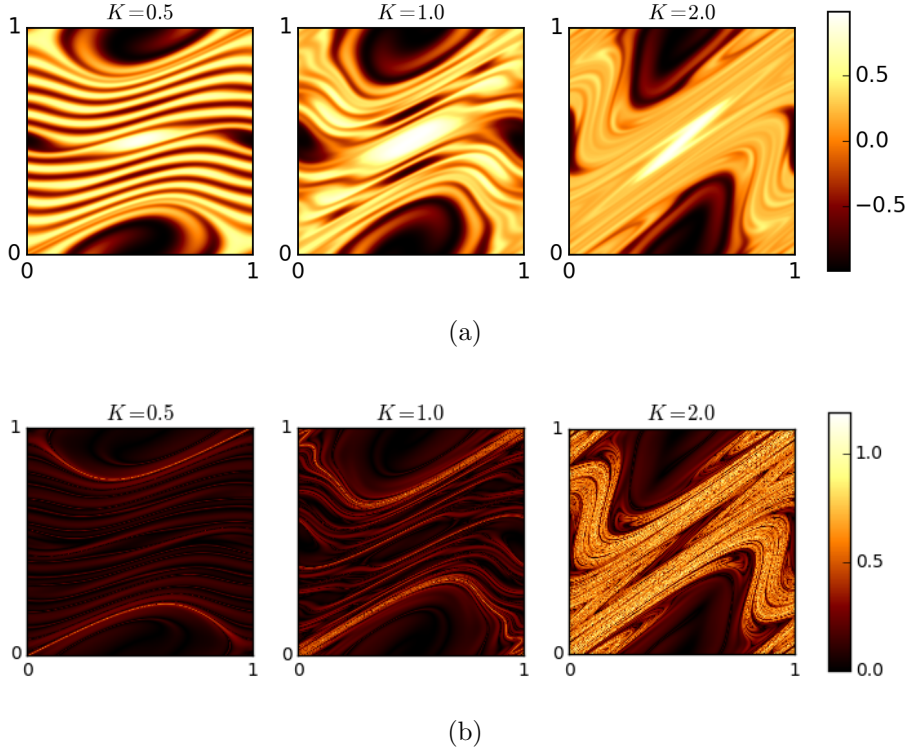
Figure 3.15: (*a*) Concentration fields after 15 iterations for the wave perturbed shear map with different wave perturbation strengths K. The initial condition was $c^{(0)}(x, y) = \cos(2\pi x)$ and $\kappa = 10^{-5}$. The concentration field can be compared with (*b*) the distribution of FTLEs in the domain, showing alignment of regions where little mixing has occurred, with areas of little or no stretch experienced by trajectories.

cat map and shear map respectively. Iteration 15 is chosen because for this diffusivity coefficient and initial condition $c^{(0)}(x, y) = \cos(2\pi x)$, the evolution of the concentration field is in the eigenfunction regime of the advective-diffusive operator for the wave perturbed cat map. The concentration fields have been rescaled and normalised with the maximum absolute value of the concentration field at the iteration, revealing well mixed regions as zeros in the concentration field, and peaks far from the concentration field revealing areas where mixing is not as good.

In Figure 3.14 there is a correlation between well stretched regions and those that have experienced a significant approach the to the mean field. When $K = 0.5$ and $K = 1.0$ bands which have experienced the least amount of average stretching history are reflected by white and black bands in the concentration field, representing the peaks and troughs of the concentration field furtherest from the mean field. When $K = 2.0$, islands in the flow exhibit the least amount of
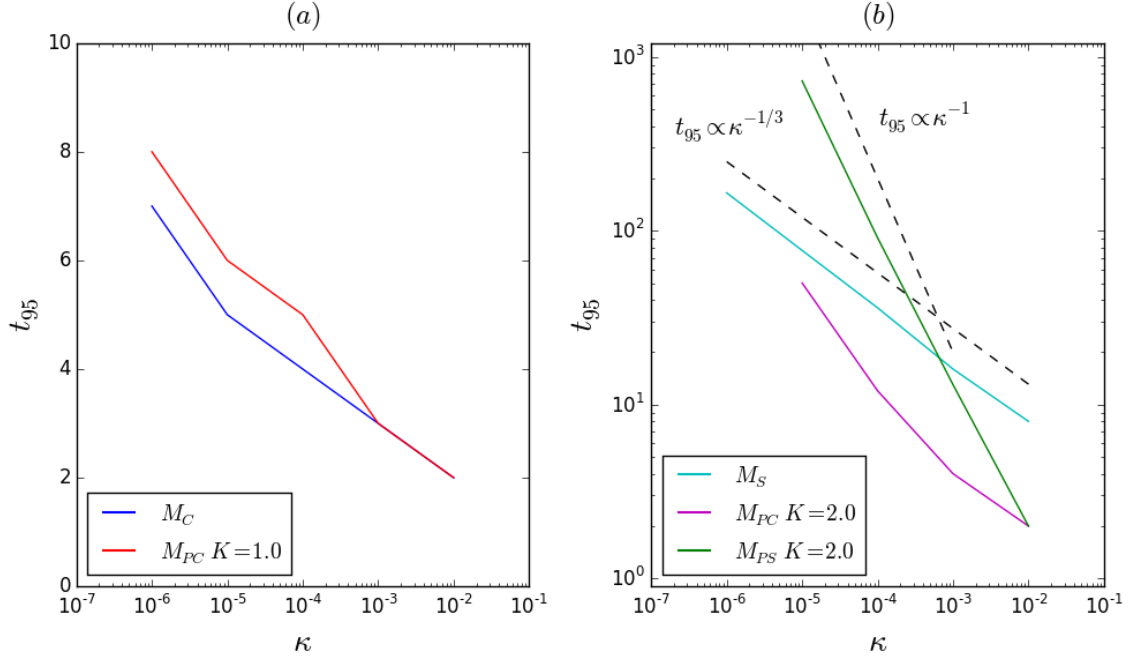
Figure 3.16: Number of iterations t_{95} is plotted for the mixed condition; within 5% of uniform, with varying rates of diffusivity coefficient κ for different stirring transformations on the unit torus. (a) plots t_{95} for the cat map M_C and the wave perturbed cat map M_{PC} showing and almost linear relation with a logarithmic change in κ, while (b) shows t_{95} varies polynomially with κ for the shear map M_S and perturbed shear map M_{PS}, the exponent of which is dependent on the dynamics.

mixing and dominate the concentration field while within the chaotic sea where some positive average stretching rates exist, the concentration field is close to approaching zero. Again, there is a qualitative agreement between bands of least stretching and regions where there are some peaks in the concentration field away from these islands.

Similar results are shown in Figure 3.15 for the wave perturbed shear map, where the fastest mixing has occurred close to the islands where the fastest average stretch has been experienced for $K = 0.5$ and $K = 1.0$. For larger $K = 2.0$ when a chaotic sea exists, good mixing occurs with regions of peaks and troughs aligning with regions of low stretch, surrounding the non-mixing islands. So although FTLEs and distributions thereof, can not be used to quantify long-time exponential mixing rates, they give qualitative understanding of eigenfunction appearance.

3.2.3 Time to a mixed condition

As an additional measurement, the time to a mixed condition is used to compare across example toral automorphism maps with and without wave-perturbation, and varying diffusivity coefficient κ. Figure 3.16 (a) plots how the time changes with κ when the underlying mapping is fully chaotic, comparing the cat map alone M_C and the wave perturbed cat map M_{PC} for $K = 1.0$ when there are no islands in the flow. The log-linear axis shows that the approximation that $t_{95} \propto \log(\kappa)$ is a reasonable approximation to find how long in takes to achieve a well-mixed concentration field.

When the mixing is not chaotic but either slowing stretching, like the shear map M_S, or has islands, like the perturbed shear map M_{PS} or perturbed cat map with $K = 2.0$, the time to a mixed condition changes more abruptly with decreasing κ. For the shear map, when plotted on a log-log axis it is observed that $t_{95} \propto \kappa^{-1/3}$. Recall from (3.20) that for the shear map the variance can be computed analytically as

$$\psi(j) = \psi(0) \exp\left(-8\pi^2 \kappa C \left(\frac{j^3}{3} + \frac{j^2}{2} + \frac{7j}{6}\right)\right) \tag{3.26}$$

where C is a constant relating to the initial condition. As such the iterate j when the concentration field is within 5% of uniform, that is $\psi(j)/\psi(0) \leq 0.05$, is approximated by

$$\exp\left(-8\pi^2 C \kappa \left(\frac{j^3}{3} + \frac{j^2}{2} + \frac{7j}{6}\right)\right) \approx 0.05. \tag{3.27}$$

To find how this scales with diffusion, some rearrangement gives

$$\frac{j^3}{3} + \frac{j^2}{2} + \frac{7j}{6} \approx \frac{1}{-8\pi^2 \kappa} \log(0.05), \tag{3.28}$$

which to leading order of the polynomial in j predicts the scaling $j \propto \kappa^{-1/3}$, which agrees with the line in Figure 3.16 (b) for the shear map.

Flow with islands predict long-time mixing rates dominated purely by diffusive effects across boundaries. Large islands in the flow will produce a slowly decaying eigenfunction with concentration of amplitude in small wavemodes, irrespective of the diffusivity coefficient. Say the dominant wavemode had wavenumber k , and since the main mechanism of mixing is purely diffusive, there is not transfer of concentration to other mode with stirring and dispersion, then the only way

amplitude can decay in this mode is via diffusive effects. Therefore the variance decay contribution from this mode will be

$$\psi(j) \approx \psi(0) \prod_{n=0}^{j} \exp\left(-4\pi^2 C\kappa\boldsymbol{k} \cdot \boldsymbol{k}\right) = \psi(0)\exp\left(-4\pi^2 C\kappa j \boldsymbol{k} \cdot \boldsymbol{k}\right). \qquad (3.29)$$

By similar arguments to the shear map, this means that the time to a mixed condition scales with diffusion coefficient κ as $j \propto \kappa^{-1}$. Indeed, for the wave perturbed shear map with $K = 2.0$ this is a good approximation, shown in Figure 3.16 (b). Although the wave perturbed cat map with $K = 2.0$ has islands in the transformation, significant mixing occurs in the initial stages, seen in the decay of variance profiles in Figure 3.13. This results in a diversion from the scaling $\sim \kappa^{-1}$ for t_{95}. However, looking at Figure 3.16 it seems reasonable to assume that the profile of t_{95} as $\kappa \to 0$ may approach this limit. This deviation from the expected scaling of $\sim \kappa^{-1}$ occurs because when the t_{95} condition is satisfied the system is not always in the asymptotic regime, and the transition to the asymptotic regime in finite time may vary depending on the initial conditions and parameters of the flow. Therefore, t_{95} and other $t_\%$ relations are used throughout the thesis to compare deviations in the mixing behaviour in the initial, finite-time stages of mixing.

3.3 Conclusions

Methodologies where introduced to study advective-diffusive systems iteratively. The advantage to this approach is the potential in studying the long-time mixing dynamics via the spectra of transfer operators; where eigenvalues of the point spectrum can be used to determine if a transformation M is strong mixing and determine the rate of decay of correlations for purely advective systems, or the rate of approach to the mean-field in advective-diffusive systems. In the advective-diffusive systems it is the emergence of slowly decaying eigenfunctions towards the mean field which produce an exponential long-time mixing rate.

Toral automorphism maps, and wave perturbed versions thereof, were used as simple examples to show how varying underlying dynamics contribute to finite time mixing rates. This was akin to Thiffeault & Childress (2003). The maps were chosen specifically since periodic boundary conditions allows for a direct computation of a transfer matrix using a dense collection of discrete Fourier waveforms. This method is accurate in computing the evolution of a concentration

field via advection and diffusion as long as a reasonably large number of Fourier modes are kept in construction of a finite sized transfer matrix.

Standard dynamical techniques, such as Poincaré sections and the distribution of FTLEs, were shown to highlight features of the flow likely to contaminate mixing, such as barriers to stirring by islands, and regions which experience slow stretching histories. The latter also qualitatively suggested the appearance of peaks in the concentration field away from the mean field in the long-time mixing of advective-diffusive systems, providing a tool to predict the appearance of eigenfunctions. Similar observations on the alignment of concentration peaks with regions of low compression have been made before [Gilbert (2006); Gouillart *et al.* (2008); Popovych *et al.* (2007)]

Although variance decay profiles are helpful in picturing the full finite time evolution of a concentration field by a particular transformation and diffusion coefficient, the time to a mixed condition provides a single value, beneficial to compare across stirring and diffusive parameters. It was shown that fully chaotic flows are likely to have weak-relation to varying diffusivity, only logarithmic with κ, while non-chaotic shear flow and islands have a polynomial relation $\propto \kappa^{-\eta}$, where η was different depending on the underlying dynamics of the non-chaotic stirring.

Chapter 4

Mixing by cutting and shuffling

As discussed in Chapter 3, smoothly varying transformations with chaotic trajectories result in fast mixing in the presence of diffusion. The stretching caused by the exponential separation of nearby particles is the dominating mixing mechanism of these systems. There are many instances in mechanical mixing in which the main underlying transport mechanism is from cutting and shuffling, which in and of itself is non-exponential, as discussed in Section 2.6.2. Examples are split-and-recombine micro-mixers and the mixing of granular material, where particle interlocking in the bulk of the flow restricts shearing.

Transformations in which cutting and shuffling form the basic mechanism of increasing surface area and decreasing segregation are explored in this chapter, in particular the mixing rates which emerge when combined with the homogenization effect of molecular diffusion. Mathematical results of such transformations so far have been concerned with asymptotic ergodic properties, and they can be at most weak mixing. Little is known of finite time mixing behaviour in such transformations.

Krotter *et al.* (2012) and Yu *et al.* (2016) studied the potential of Interval Exchange Transformations (IETs) to mix well in finite time, investigating which parameters resulted in efficient decreases in the scale of segregation. Here, their defined parameter space is utilised to investigate the rate of mixing from diffusion in these one-dimensional cutting and shuffling transformations. Ashwin *et al.* (2002) investigated the asymptotic rates of mixing by permutations of equally sized cells composed with diffusion and found a speed-up over the action of pure diffusion, however only permutations from the groups S_3, S_4 and S_5 were considered, which have low periodic orders. In the absence of diffusion, reassembly of any initial condition would occur after just a handful of iterations. A diffusive

step in the IETs of Krotter *et al.* (2012) and Yu *et al.* (2016) has already been investigated by Wang & Christov (2018), attempting to find evidence of cut-offs, a phenomenon in card shuffling. A function is found which described the averaged finite mixing behaviour, however there is no parameter space exploration or discussion on long-time mixing rates. Additionally, the diffusional numerical method used by Wang & Christov (2018) would not be independent of the discretization resolution, and a detailed review of the numerical method and results of this paper are contained within this chapter.

Here a superior computational model is used to investigate the evolution of concentration fields to the mean field in asymptotic and finite time, for a range of parameter choices defining an Interval Exchange transformation. The effect of permutation choice, interval lengths, number of intervals and the diffusivity coefficient κ are all considered and compared to results of previous literature. This computational study shows that previous suggestions on parameters which result in good stirring in the absence of diffusion, hold to some extent when there is diffusion but not necessarily in the way previously thought in the literature.

The chapter concludes with a potentially universal scaling on the effect of diffusion to achieve a desired mixed condition in IETs. When the parameters are chosen to satisfy the conditions of good stirring, this scaling is shown to be faster than diffusion acting alone and numerical evidence suggests it holds well for all such IETs. The scaling also appears in the modulus of the leading, non-trivial eigenvalue which defines the asymptotic mixing rate. This leads the discussion to hypotheses on a mechanism which would derive such a scaling in both finite and asymptotic time; a balance between cutting, shuffling and diffusion.
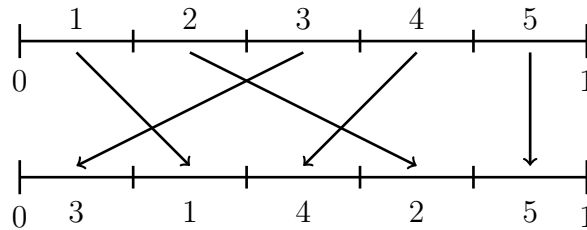
Figure 4.1: A permutation of equal sized cells on the unit interval by an interleaving permutation from S_5.

4.1 Interval exchange transformations

4.1.1 Formal definitions

The cutting and rearranging of a unit interval has been explored in many dynamical problems. The simplest formation is the rearrangement of equal sized pieces which is defined as a map as follows.

Definition 4.1.1. (Permutation of equal sized cells) *Let $\Omega = [0, 1)$ and consider a permutation $\Pi \in S_N$ where S_N is the group of all permutations on the set of symbols $\{1, 2, \ldots, N\}$. The action of the map M_Π can be defined such that for a point $x \in [(l-1)/N, l/N)$,*

$$M_\Pi(x) = x + \frac{\Pi(l) - l}{N}. \tag{4.1}$$

The definition can be extended to any interval $\Omega = [a, b)$, but in the discussion which follows, Ω is always taken to be the unit interval.

The notation defining a permutation Π is non-standardised. For example Ashwin *et al.* (2002) adopted disjoint cyclic notation which is helpful in understanding the periodic order of a given permutation. Figure 4.1 shows an interleaving permutation of equal sized cells from S_5 and can be represented by either (1243)(5), meaning $1 \to 2$, while $2 \to 4$ and $4 \to 8 \mod 5 = 3$ and $3 \to 1$, to complete the cycle while 5 appears as its own disjoint cycle. An alternative cyclic notation is (1243), where elements which do not move are omitted. The disadvantage of the alternative notation is that difficulty can arise from distinguishing which permutation group S_N a permutation belongs to if it is not stated. The order of a permutation Π is the minimal common multiple of the lengths of the disjoint cycles.

Krotter *et al.* (2012) and Yu *et al.* (2016) choose to represent a permutation as the final rearrangement order of the sub-intervals. For example if the initial configuration of sub-intervals from S_5 is represented by [12345], then the rearrangement of these intervals by one iteration of the interleaving permutation of Figure 4.1 would be [31425]. The latter notation is adopted throughout the work herein, since is avoids unnecessary confusion and is more compact than the former cyclic notation.

An extension rearranging sub-intervals of varying length is captured in an Interval Exchange Transformation (IET), described by a tuple of a permutation and a set of interval lengths. The following definition is taken from Avila & Forni (2007).

Definition 4.1.2. (Interval Exchange Transformation (IET)) *Let $N \geq 2$ be a natural number and let Π be a permutation $\Pi \in S_N$. Given $L \in \mathbb{R}_+^N$, we define an Interval Exchange Transformation $f(L, \Pi)$ by considering the interval*

$$I = I(L) = \left[0, \sum_{i=1}^{N} L_i \right) \tag{4.2}$$

broken into sub-intervals

$$I_i = I_i(L) = \left[\sum_{j<i} L_j, \sum_{j\leq i} L_j \right), \quad 1 \leq i \leq N. \tag{4.3}$$

Then the action of the IET is to rearrange the I_i according to Π (the ith interval is mapped to the $\Pi(i)$th interval) such that $f(L, \Pi) : I \to I$ is given by

$$f(x) = x + \sum_{\Pi(j)<\Pi(i)} L_j - \sum_{j<i} L_j, \quad x \in I_i. \tag{4.4}$$

IETs are one-to-one and continuous except for a finite set of points, preserve Lebesgue measure and, as defined above, preserve orientation. For the remainder of the chapter the action of IETs on the unit interval $I = \Omega = [0, 1)$ is considered without loss of generality.

Finally two additional definitions are introduced which are important to the discussion. Permutations of S_N can be divided into subgroups by the following definitions.

Definition 4.1.3. *A* **rotation permutation** *is defined as $\Pi \in S_N$ which satisfies $\Pi(i+1) = \Pi(i) + 1 \mod N$ for all $i \in \{1, 2, \ldots N\}$.*

Definition 4.1.4. *A* **reducible permutation** *is defined as Π where there exists $k \in \{1, \ldots N-1\}$ such that*

$$\Pi_1(\{1, \ldots, k\}) \in S_k \tag{4.5}$$

then for any choice of L, the sub-interval

$$J = \bigcup_{i \leq k} I_i = \bigcup_{\Pi(i) \leq k} I_i \tag{4.6}$$

is invariant under the transformation f.

A permutation that is not reducible is defined as **irreducible**. For example, [32154] is a reducible permutation, since the first 3 elements are a permutation of only themselves, and similarly the last two elements. An example of an irreducible permutation [31524], which can not be decomposed into two or more permutations. Note that all rotation permutations are irreducible, except an identity permutation such as [12345]. Ergodic properties and dynamics of IETs were discussed in Section 2.6.1.

4.1.2 Model parameters for investigation

In line with past literature, the parameter space of interval exchange transformations conceived by Krotter *et al.* (2012) is chosen for investigation herein. Consider the set of sub-intervals $\{L_1, \ldots, L_N\}$ where the length $L_i = |I_i|$ of adjacent sub-intervals satisfies a constant ratio $r = L_i / L_{i-1}$. Since it is required that the total length of the interval I to be 1, then the primary interval length L_1 is found to be;

$$\sum_{i=1}^{N} r^{i-1} L_1 = 1 \implies L_1 = \frac{r-1}{r^N - 1}. \tag{4.7}$$

The IET can then be redefined as a tuple with respect to r such that the map is $f(r, \Pi) : [0, 1) \to [0, 1)$. Figure 4.2 depicts an example IET under this construction. The action of the IET $f(1.5, [3142])$ is shown as a space-time plot of incurred decrease in segregation from repeat application on a given initial condition, here a *half black, half white* unit interval. The new concentration field across the interval is shown at each iteration to aid understanding of the full finite time evolution up to iteration j. Space-time plots are similar to panelled concentration plots shown for a number of iterations in Chapter 3, and are used numerously in this and further Chapters.

Computational investigations on the finite time stirring of IETs $f(r, \Pi)$ are restricted to only rational values of r due to limited precision in computing. This implies that weak mixing IETs can not be computed, but since weak mixing is an asymptotic result and gives no information on finite time mixing properties,
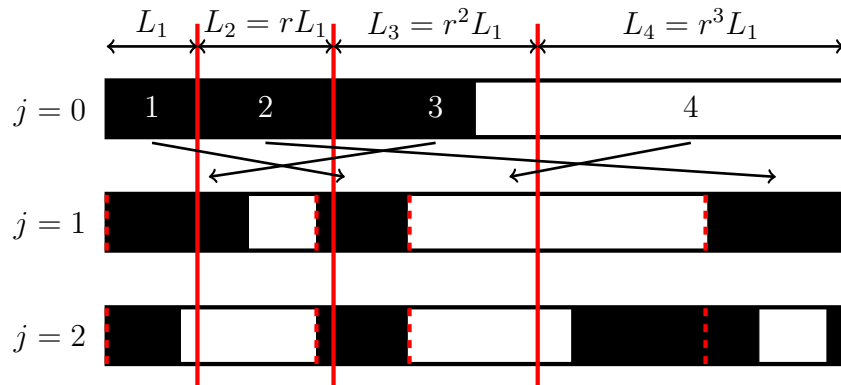
Figure 4.2: Construction of an IET using the parameters $f(r, \Pi)$, where $r = 1.5$ and $\Pi = [3142]$. Repeat application of the IET decreases the scale of segregation of the initial condition *half black, half white*, shown qualitatively by a space-time plot. Cut locations (dashed red) are tracked by the rearrangement (solid red) of intervals

it is not considered a problem for this discussion. Due to the choice of rational r, the IET $f(r, \Pi)$ can be redefined as a map M_{Π_f} (4.1), a permutation of equal sized cells of length N_f. The length N_f and permutation Π_f are found as follows.

Since r is rational, the fractions representing the rational numbers $\{L_1, L_2, \ldots, L_N\}$ will have a largest common denominator. This can be found from

$$N_f = lcd = \frac{(r_n^N - r_d^N)}{(r_n - r_d)} \tag{4.8}$$

when r is converted to a fraction $r = r_n/r_d$. The permutation Π_f acting on cells of equal length $1/N_f$ can be constructed by the permutation arrangement

$$\Pi_{N_f}(i) = i + \sum_{\pi(j) < \pi(i)} L_j N_f - \sum_{j < i} L_j N_f, \tag{4.9}$$

since all values $L_i N_f$ are integers. The rearrangement of this larger permutation formed the basis of the computational algorithm of Krotter *et al.* (2012). All IETs with rationally dependent intervals will be periodic and all points $x \in I$ will have a periodic order given by the order of the disjoint cycles Π_f or the minimal common multiple of all disjoint cycle orders. The order of the permutation Π_f will give the total number of iterations needed to observe a full reassembly of an initial concentration field.

The results of Krotter *et al.* (2012) and Yu *et al.* (2016) found that good reduction in the scale of segregation can be achieved in finite time by establishing the following criteria on the IET $f(r, \Pi)$;

1. Π should be an irreducible, non-rotational permutation,

2. the ratio of adjacent sub-interval lengths should not be large; closer to 1 but not equal to 1,

3. and the ratio of adjacent sub-intervals should be "closer-to-irrational", or there should be slight random variation in the cut locations which define the boundaries of the sub-intervals.

In Krotter *et al.* (2012) an r is described as "closer-to-irrational" if it has a longer continued fraction, that if a is a rational number, then it can be represented by

$$a = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots + \cfrac{1}{a_k}}}}, \tag{4.10}$$

a finite continued fraction of k partial denominators. Real numbers have a unique continued fraction representation where a_k are all integers, and irrational numbers correspond to infinite continued fractions. If b has a longer continued fraction than a, this means that more a_k are needed to write b in its continued fraction form. Krotter *et al.* (2012) contribute needing more a_k elements to represent r results in the permutation Π_f having a larger N_f (number of equal-sized cells to capture the sub-intervals).

This definition is mathematically ambiguous and whether this holds for all rational r is not known. To compare with the results and terminology of Krotter *et al.* (2012) and Yu *et al.* (2016), for a fixed N, a "closer-to-irrational" r will result in a Π_f with a larger N_f. Krotter *et al.* (2012) also state an observed diminishing returns of increasing the number of sub-intervals and state that $N = 6$ is the most practical, while Yu *et al.* (2016) state that as long as the final criterion of the IET good stirring criteria, listed on the previous page, is met then $N = 4$ is enough to produce good stirring.

Two choices in the methodologies of Krotter *et al.* (2012), Yu *et al.* (2016) and Wang & Christov (2018) which will not be used here are the measures of mixing diagnostics and the initial concentration fields. The two mixing measures in the diffusion-less investigations are the increase in the number of interfaces C_j, and the largest percent unmixed U_j; both of which were briefly described in Section 2.4.3. They give no spatial information on the total arrangement across the interval, are not easily computable in higher dimensions, and are not computable in systems with molecular diffusion. Wang & Christov (2018) investigated the latter point. Similarly, the main initial condition considered in these studies is constructed by assigning different colours, or values, to the IET sub-intervals and then tracking how these are cut and shuffled. This choice of initial condition is again specific to the configuration at hand and unrealistic compared to real mixing processes. These initial configurations were likely chosen due to the limitations of the computational schemes; shuffling the larger permutation [Krotter *et al.* (2012)] and tracking interfaces [Yu *et al.* (2016)]. However, a result of Yu *et al.* (2016) shows that the conclusions on the value r hold when the initial condition is composed of two components; a *half-black, half white* interval.

The next section outlines the numerical method to include diffusion as an iterative step in Interval Exchange transformations.

4.2 Iterative Advection-Diffusion Transformation

Investigating stirring from cutting and shuffling alone is interesting mathematically but in real-life mixing problems is unlikely to occur in isolation. Yu *et al.* (2016) introduced random, but small, variations about sub-interval boundaries at each iteration, mimicking mechanical variation as one approach to a more realistic mixing scenario. Similarly, stretching and folding may be present as part of a stirring mechanism, or diffusion may blur the edges between neighbouring sub-segments of varying colour, be it fluid via molecular diffusion or particle-particle interaction in granular material. For the remainder of this chapter, the mixing rates of Interval Exchange transformations with molecular diffusion are the investigated.

4.2.1 Transfer operators for permutations of equal sized cells

The incorporation of diffusion can be achieved by previous computational methods introduced in Chapter 3, solving the advection-diffusion equation as a discrete time evolution via operator splitting. By assuming periodic boundary conditions, discrete Fourier modes are once again used as the basis functions to construct a finite sized transfer matrix. Ashwin *et al.* (2002) derived the transfer matrix for the permutation of equally sized cells by solving

$$\hat{c}_k^{(j+1)} = \hat{c}_q^{(j)} \int_0^1 e^{2\pi i q x - 2\pi i k M_\Pi(x)} \mathrm{d}x \tag{4.11}$$

where q is regarded as the initial wavenumber, and k as the final wavenumber at each iteration. This results in the transfer matrix $\boldsymbol{M}_{kq}$ denoting a transfer of concentration from q to k of

$$\boldsymbol{M}_{kq} = \frac{1 - \omega^{(q-k)}}{2\pi i (q-k)} \sum_{\ell=1}^{N} \omega^{k\Pi(\ell) - q\ell} \tag{4.12a}$$

when $k \neq q$, where the primitive Nth root of unity is defined as $\omega = e^{-2\pi i/N}$. When $k = q$,

$$\boldsymbol{M}_{kq} = \frac{1}{N} \sum_{\ell=1}^{N} \omega^{k(\Pi(\ell) - \ell)}. \tag{4.12b}$$

The derivation of transfer matrices are within Appendix B. Taking the notation $\bar{c}$ to mean the complex conjugate of c, the matrix elements $\boldsymbol{M}_{kq}$ have the symmetry

$\boldsymbol{M}_{-k-q} = \overline{\boldsymbol{M}_{kq}}$, ensuring the reality of concentration fields are preserved by the transfer matrix.

Extension of the matrix derivation to IETs $f(r, \Pi)$ is possible due to the map M_{Π_f} (4.9) where there is the construction of a larger permutation of equal sized cells. This is extendible to all IETs $f(L, \Pi)$ with the set L of rational sub-intervals L_i. The matrix defining the diffusive operator P_D is simply the diagonal matrix

$$\boldsymbol{d}_{kq} = \rho^{k^2} \boldsymbol{\delta}_{kq}, \quad \rho = \exp(-4\pi^2 \kappa), \tag{4.13}$$

and the matrix describing the full iterative step is then

$$\hat{c}_k^{(j+1)} = \sum_{q=-Q}^{Q} \boldsymbol{d}_{km} \boldsymbol{M}_{mq} \hat{c}_q^{(j)}, \tag{4.14}$$

for a given truncation value Q. A resolution test can be performed to find an appropriate value Q required for accurate representation of the evolution when κ is non-zero, which is addressed in the next section. Since the system is one-dimensional, the size of the transfer matrix for certain values of Q will be of a feasible size to directly compute the eigenvalues and eigenvectors of the transfer matrix.

Note that computation of the transfer matrix (4.12) for the IET $f(r, \Pi)$ in terms of M_{Π_f} involves computing $2Q \times 2Q$ entries, all of which require a summation over the length N_f of the extend permutation. IETs which result in large N_f produce significant computation time for each entry of the transfer matrix, hence there are some restrictions on the reasonable choices for r and N with feasible computation time. Therefore the choice of IETs with "closer-to-irrational" r is limited. Additionally, extension to investigating IETs with cut variation is not intuitively straightforward and is excluded.

The main measure of mixing used throughout the discussion is the decay of variance to the mean field. However, when a comparison is needed to the stirring, non-diffusive dynamics of the IETs, the mix norm of Mathew *et al.* (2005), which gives indication of spatial scales within the interval, is also used. The computation of the mix norm in the diffusion-less case is taken from Smith (2016).

4.2.2 Initial conditions of interest

Krotter *et al.* (2012) and Yu *et al.* (2016) both consider initial conditions in which a different colour, or numerical value, is assigned uniformly across each initial sub-interval L_i to investigate stirring. Most recently the assigned-component initial

condition was also employed by Wang & Christov (2018) to study diffusion in IETs.

This study discards this initial condition due to several Reasons. Firstly, such an initial condition depends purely on the parameters of the underlying stirring transformations, meaning comparison across varying parameter spaces do not yield the same starting conditions. Secondly, it is an unlikely configuration assigned in a real-life mixing scenario. Thirdly, on the inclusion of diffusion the reasoning for choosing such a condition is made completely invalid; which was purely visual rather than a numerical consideration of an approach to a mean field. Finally, a unique problem to this study, the computation of such an initial condition in discrete Fourier space may be complicated, with no simple way to analytically compute the Fourier coefficients. Using a fast Fourier transform of the initial configuration is possible, but the errors introduced when doing so would be unknown. Instead the main initial conditions considered herein are: *half black, half white* and variations thereof, and in line with the previous chapter, configurations of cosine and sine waves.

The *half black, half white* initial condition, represented periodically across the interval as the square wave, has a Fourier representation which is easy to compute. The advantage of this initial condition is its utility in investigating both mixing when there is diffusion, and segregation when the system is purely stirring. Other initial conditions composed of sine and cosine waves are not directly transferable to computational methods of the past literature on mixing in Interval Exchange transformations.

The square wave on the unit torus is given by

$$c^{sq}(x) = \begin{cases} -1 \text{ if } x \in [0, 1/2) \\ +1 \text{ if } x \in [1/2, 1). \end{cases} \tag{4.15}$$

The discrete Fourier coefficients for the square wave are calculated from

$$\hat{c}_k = \int_0^{1/2} 1 \cdot e^{-2\pi i x k} dx + \int_{1/2}^1 -1 \cdot e^{-2\pi i x k} dx \tag{4.16}$$

such that

$$\hat{c}_k = \begin{cases} 0 \text{ when } k \text{ even,} \\ \frac{2i}{k\pi} \text{ when } k \text{ odd.} \end{cases} \tag{4.17}$$

For the square wave, a simple rotation of α can be applied as $c^\alpha(x) = c^{sq}(x + \alpha)$; in which the coefficients are evaluated as

$$\hat{c}_k^\alpha = \begin{cases} 0 \text{ when } k \text{ even,} \\ \frac{2i}{k\pi} \cdot e^{2\pi i \alpha k} \text{ when } k \text{ odd.} \end{cases} \tag{4.18}$$
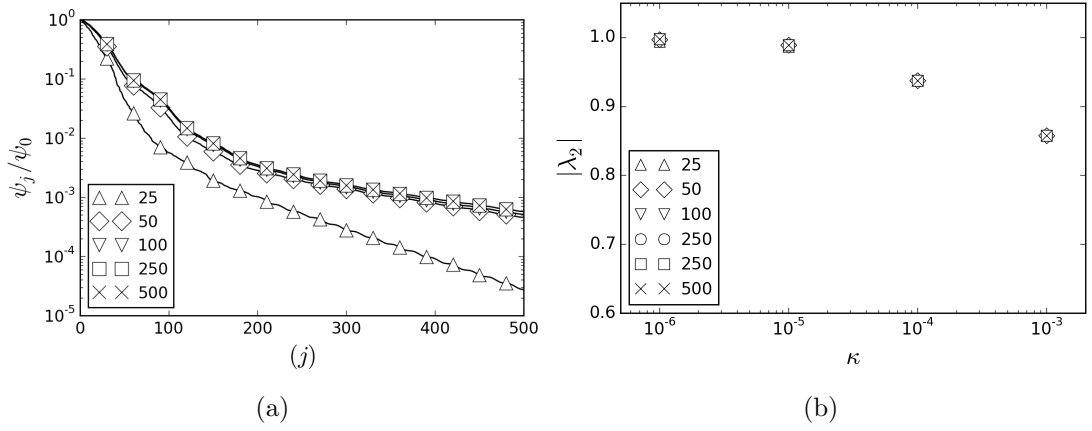
Figure 4.3: Results from a resolution test for Q using the IET $f(r, \Pi) = f(1.5, [4321])$ as an example transformation. Both (a) the decay of variance ψ_j/ψ_0, for one diffusive rate $\kappa = 10^{-6}$, and (b) the modulus of the leading eigenvalue $|\lambda_2|$ for several diffusivity rates show a convergence in values for $Q = 250$ and 500.

Initial conditions composed of cosine and sine waves can be created similarly to Toussaint *et al.* (1995) in which the amplitude of the first few Fourier coefficients are randomised to provide a simple random way to test the effect of the initial condition. Taking $\bar{c}$ to represent the complex conjugate of c, providing $\hat{c}_{-k} = \overline{\hat{c}_k}$, the initial condition is real.

4.2.3 Resolution of Fourier Modes

A resolution test is required to ensure enough wave modes are used to correctly capture the evolution of the concentration field, otherwise there may be significant contribution from numerical diffusion. The decay of variance ψ_j/ψ_0, and the modulus of the non-trivial leading eigenvalue of the matrix $\boldsymbol{dM}$, are used as metrics to determine the resolution required. The results for both metrics are shown in Figure 4.3 using a representative IET with $\Pi = [4321]$ and $r = 1.5$, and the values $Q = 25, 50, 100, 250, 500$ tested. The initial condition for the variance decay was $c^{(0)}(x) = \cos(2\pi x)$. Profiles for both ψ_j/ψ_0 and $|\lambda_2|$ show no variation between the values $Q = 250$ and 500, presented by the perfect alignment of the symbols on both plots. Additionally, $Q = 1000$ was computed and showed no variation. For lower values of Q variance loss in the initial stages is significant due to numerical diffusion.

It is surprising that $Q = 250$ accurately captures the evolution correctly for values when the diffusivity coefficient is as low as $\kappa = 10^{-6}$. For the one-

dimensional approximation of the non-uniform baker's transformation with diffusion, Wonhas & Vassilicos (2002) suggest the diffusive contribution $\exp(-4\pi^2 \kappa k^2)$ makes modes $|k| \gg Q_d$ negligible where

$$Q_d = \frac{1}{2\pi\sqrt{\kappa}}, \tag{4.19}$$

which for $\kappa = 10^{-6}$, $Q_d \approx 160$, however it is arguable whether $250 \gg 160$. Compare the value of Q here to that of the cat map in Section 3.2.2, in which $Q \approx 1800$ was required to accurately capture the evolution of the concentration field. However, there is a contrast here, since from cutting and shuffling alone there is an inefficient transfer of concentration to higher modes, in contrast to the cat map in which there is an exponential cascade to large wavemodes. From stirring by the IET, the transfer of concentration to higher modes is due primarily to the discontinuities introduced at each iteration. In this case the value of the coefficient $\hat{c}_k$ from a discontinuity is likely to be small, as is seen in the complex coefficients of the square wave where to capture the discontinuity, coefficients decrease as $|\hat{c}_k| = 2/k\pi$, which for $Q = 250$ is $|\hat{c}_{250}| \approx 0.002$. Therefore, with the contribution from diffusion, the value of coefficients at these wavemodes are likely to be very small, and thus negligible in evolution of the concentration field. For the remainder of the study, the truncation $Q = 250$ is used with diffusivity coefficients no smaller than $\kappa = 10^{-6}$ considered.

To ensure that the resolution results are not dependent on the chosen parameters $f([4321], 1.5)$, other permutations of larger N and different r were also tested and the same conclusions reached.

4.2.4 An initial example of mixing by Interval Exchange transformations with diffusion

As an introductory example, the IET depicted in Figure 4.2 is used to show the incorporation of diffusion alongside a second IET for comparison. Figure 4.4 (a) shows a space-time plot plotting the first 200 iterations of the initial condition *half white, half black* under the evolution of the IET $f_1 = f([4321], 1.5)$. Small sub-segments of white and black are dispersed between larger segments, and qualitatively there appears to be some order to the pattern, where a banded structure is visible in many of the rows, rotated across the domain. Figure 4.4 (b) shows the space-time plot for the evolution of the same initial condition under the action of a second IET $f_2(r_2, \Pi_2)$ with parameters $\Pi_2 = [653241]$ and $r_2 = 1.25$.
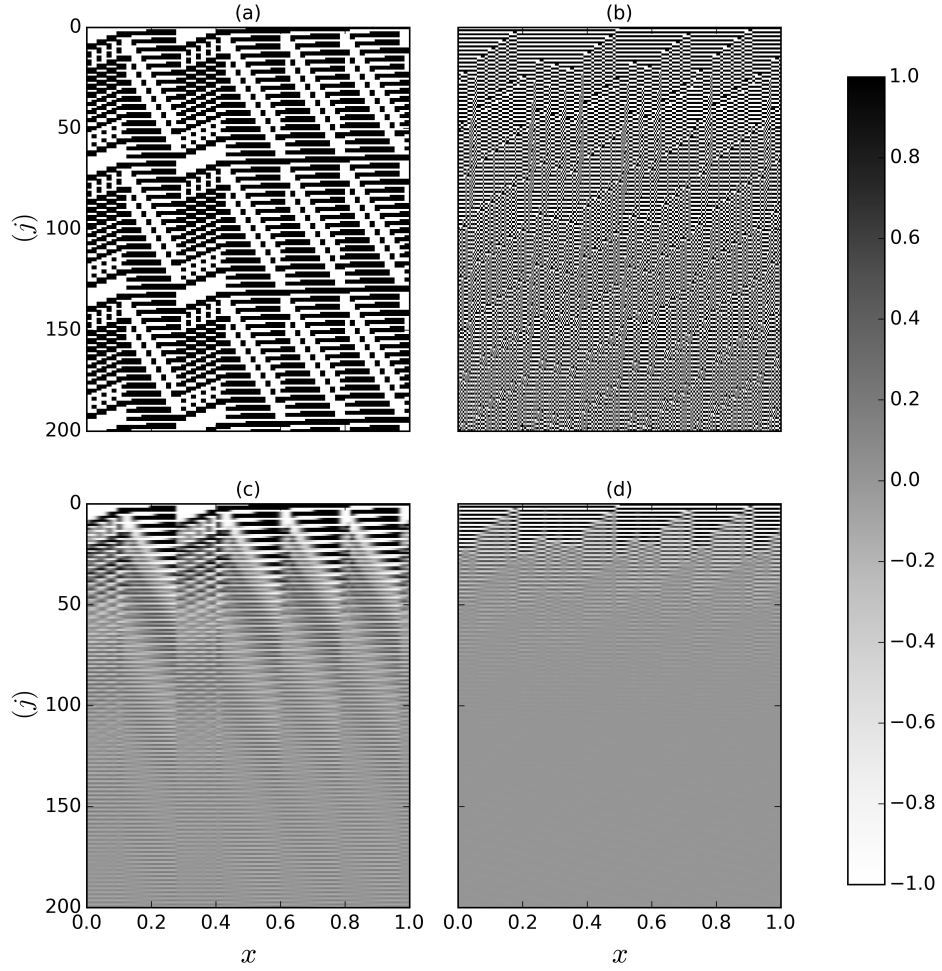
Figure 4.4: The mixing of two example IETs is compared. In (a),(b) a qualitative comparison using space-time plots for stirring is plotted in, while in (c),(d) an approach to uniformity with diffusivity coefficient $\kappa = 10^{-5}$ is plotted. In (a),(c) $f_1(r_1, \Pi_1)$ with $r_1 = 1.5$ and $\Pi_1 = [3142]$, in (b),(d) $f_2(r_2, \Pi_2)$ with $r_2 = 1.25$, $\Pi_2 = [653241]$. The initial condition in all plots was the square wave $c^{(0}(x) = c^{sq}$.

79
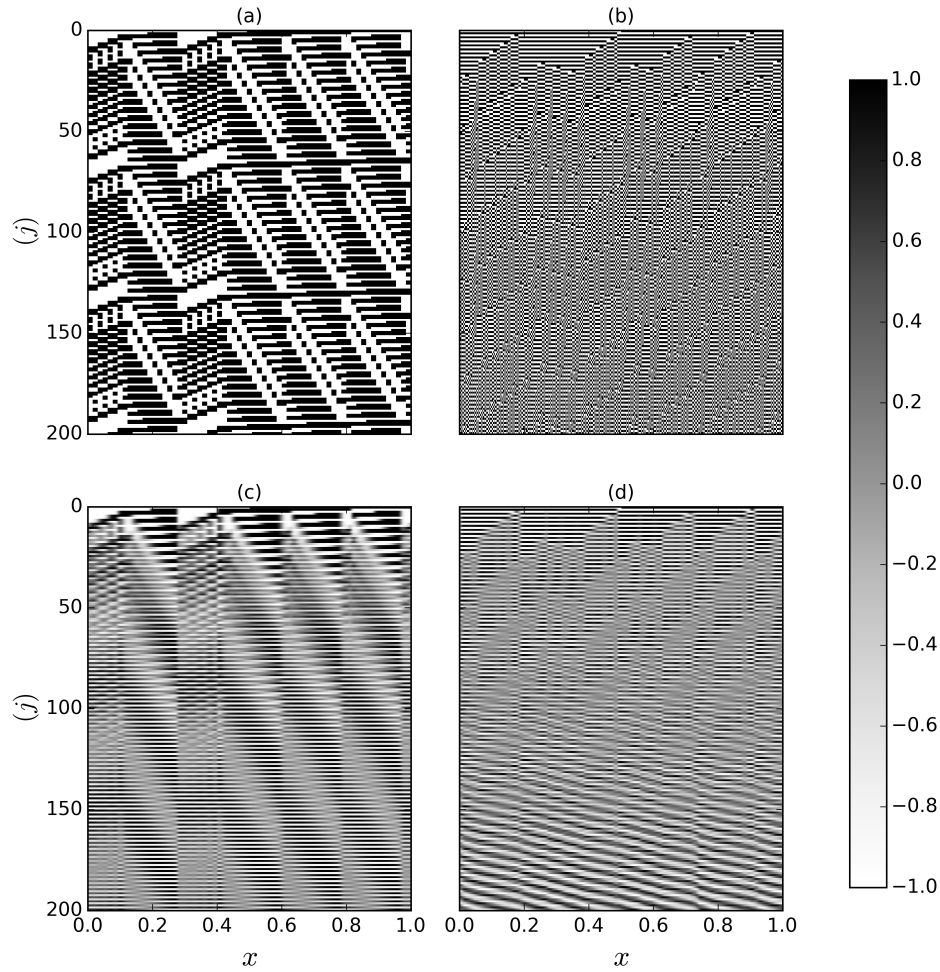
Figure 4.5:   The mixing of two example IETs from Figure 4.4 is compared again but in the plots (c),(d) an approach to uniformity with diffusivity coefficient $\kappa = 10^{-5}$ is plotted with the concentration field normalised at each iteration. It can be seen that the inclusion of diffusion significantly changes the scale and arrangement of the concentration field as it approaches the mean field.

These two IETs were compared in Yu *et al.* (2016) to see the effects of varying the parameters Π and r. In that study, f_2 improved stirring of a coloured component initial condition, however because the permutations varied in size, so too did the initial conditions. Here, qualitatively it is seen that both IETs acting on the same initial condition preserves the conclusions that f_2 is superior at stirring over f_1.

The space-time plots of Figures 4.4 (c) and (d) show the time evolution for the iterative advective-diffusive equation (4.14) with diffusion coefficient $\kappa = 10^{-5}$ and advection provided by the same IETs, f_1 and f_2 respectively. A transition to grey, the mean field of the initial condition, is observed in both however the qualitative comparison between the two space-time plots appears to show a faster transition for f_2.

Figure 4.5 re-plots the space-time evolutions of Figure 4.4 but with the concentration field normalised at each iteration. The inclusion of diffusion cause a depletion of a concentration field to its mean field and, from these renormalised space-time plots, it is also observed that diffusion encourages an increase in spatial scale horizontally within the domain. Particularly for the IET f_2 in Figure 4.5 (d), the final concentration field appears to have a dominant wave mode which undergoes rotational symmetry under repeat iteration of the transfer operator $P = P_D \circ P_M$ which is much larger in scale when compared to the concentration field at the same time in the diffusionless system Figure 4.5 (c). It is also seen that the IET f_2 undergoes two regimes during the finite time evolution of the initial condition $c^{(0)} = c_{sq}$ with diffusion, where a repeat pattern in the renormalised concentration field seems visible between iterations $\sim 25 - 125$, before there is a transition into a different, later stage pattern as $j \to 200$. The presence of these two exponential regimes are reflected in the decay of variance.

Computing the variance at each iteration gives a quantitative comparison between the two IETs, plotted in Figure 4.6 on two different axes. In the linear-linear plot in Figure 4.6 (*a*) the variance depletion in the initial stages decays similarly for both IETs. In these early stages not much stirring has taken place, however sharp gradients introduced at each iteration have initiated an approach to the mean field. Around iteration 20, there is a divergence in the variance decay profiles and IET f_2 depletes the variance quicker than f_1. This agrees with the superior stirring properties of f_2.

In both cases there appears to be a distinct transition to a later stage of mixing, around iteration 40 for f_1 and iteration 25 for f_2. On the linear-log axis in Figure 4.6 (*b*) this transition leads to exponential mixing. However, for f_2,
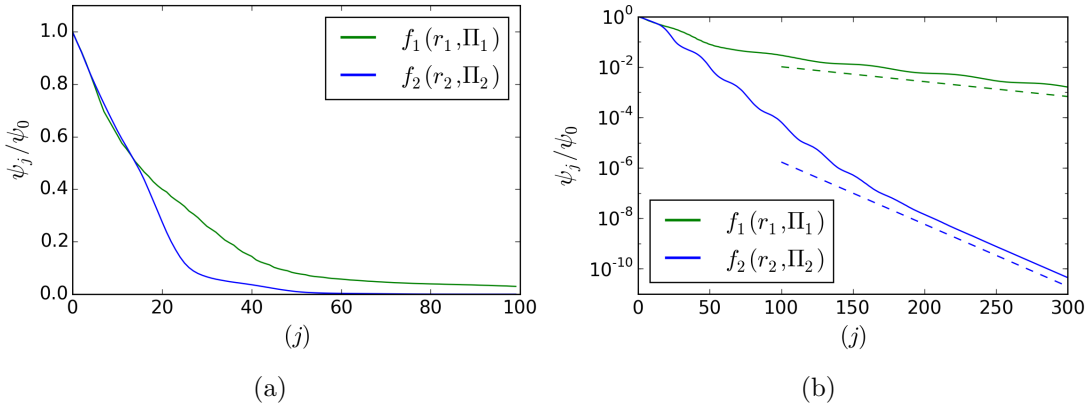
Figure 4.6: The decay of variance in iterative diffusive IETs are compared between $f_1(r_1, \Pi_1)$ and $f_2(r_2, \Pi_2)$ with $r_1 = 1.5$, $\Pi_1 = [3142]$ and $r_2 = 1.25$, $\Pi_2 = [653241]$, and diffusion coefficient $\kappa = 10^{-5}$. The initial condition was the square wave $c^{(0)}(x) = c^{sq}$. (a) Linear-linear plot shows clearly the initial stages of mixing, while a (b) linear-log plot shows the long-time exponential mixing rate.

there appears to be two exponential decay rates in the variance profile, with a long-time exponential mixing rate emerging around iteration 150, after significant depletion of variance has occurred. Assuming the long-time decay of variance is dominated by the slowest decaying eigenfunction (2.50), plotted for comparison are the variance decay profiles predicted from the second leading eigenvalue λ_2,

$$\psi_j \propto |\lambda_2|^{2j}. \tag{4.20}$$

computed from the respective diffusive transfer matrices $\boldsymbol{dM}$ for each IET. There is good agreement to both profiles confirming this is indeed the long-time mixing rate of the transfer operator. Note that the long-time mixing rate is also quicker for IET f_2, suggesting that good stirring properties may impact all stages of mixing. In both cases this long-time mixing rate emerges after significant mixing has occurred, especially for f_2 where ψ_j/ψ_0 is $O(10^{-6})$ before the final stage begins. Following this illustrative example of diffusive IET transformations, the next section considers each of the mixing stages separately. A full parameter study is left for later in the chapter.

4.3 Stages of mixing in Diffusive Interval Exchange Transformations

Thiffeault & Childress (2003) noted that chaotic mixing systems had three dis-

tinct stages of mixing. First there is an initial stage where there is hardly any depletion of variance. This lasts until the concentration field is stirred and the scales of segregation reach that on the order of which diffusion is effective. Then there is the possibility of a super-exponential stage of mixing due to the dual exponential mixing rates from stirring combined with the exponential effect of diffusion at small scales. This stage is rarely seen. The final stage of mixing is when there is a balance between stretching to smaller scales and diffusion, seeing the emergence of eigenfunctions of the concentration field and the corresponding long-time exponential rate. If the underlying system is strong mixing, the exponential decay rate becomes independent of diffusion in the diffusion-less limit. Similarly in systems with contamination from slow regions of mixing, such as physical or island boundary behaviour, an additional stage of polynomial decay rate is observed before the final exponential stage [Gouillart *et al.* (2008); Popovych *et al.* (2007)].

IETs lack hyperbolic behaviour leading to exponential decay in the scale of segregation, and instead discontinuous transformations introduce sharp interfaces into a concentration fields. These sharp gradients will then be washed out by diffusion. Due to these differences with chaotic systems it can not be assumed that similar stages of mixing occur. In the next two sections, the evolution of the concentration field is considered in two parts. First the asymptotic, long-time exponential mixing rate is discussed, computed from the leading eigenvalues of the transfer operator, and compared with how it varies with diffusion coefficient κ. Secondly, the full finite time mixing profile is considered where the variance decay profiles deviate from an exponential mixing rate. A function fit for this mixing rate is investigated and a mechanism for the mixing rate proposed.

4.3.1 Asymptotic mixing rates

In the transfer matrix (4.12), for each given mode k there is a dispersion of amplitude to lots of modes rather than a direct cascade. As would be expected, this dispersion results in long-time exponential decay. This was reported for permutations of equal size cells by Ashwin *et al.* (2002). This is due to the fact that the dispersion creates eigenfunctions for the transfer operator with eigenvalues which predict the exponential mixing rate. This dispersion is caused from the cutting and rearranging of the concentration field at each iteration. Therefore, although the underlying mechanism producing dispersion in the Fourier modes differs from that that continuously deforming transformations, where it is caused
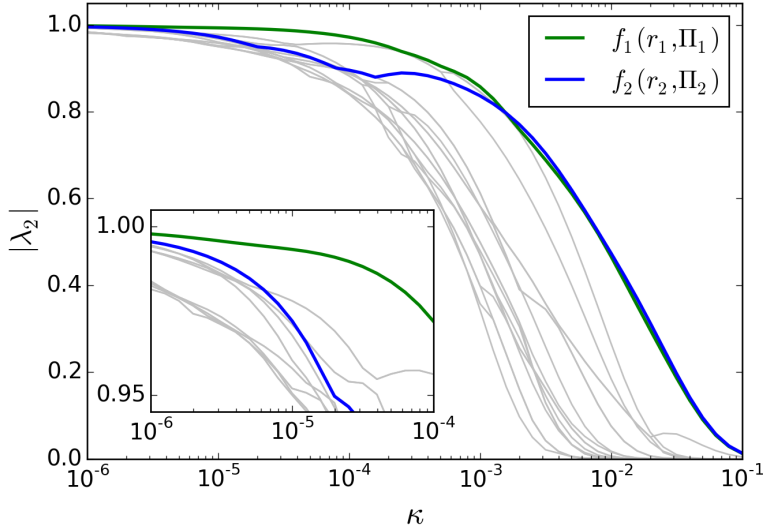
Figure 4.7: The modulus of the leading non-trivial eigenvalue $|\lambda_2|$ for both $f_1(r_1, \Pi_1)$ and $f_2(r_2, \Pi_2)$ tends to 1 as $\kappa \to 0$. Some non-monotonicity in the approach to the diffusion-less limit is seen for (r_2, Π_2). The values of $|\lambda_k|$ are plotted as light-grey lines suggesting $|\lambda_k| \to 1$ as $\kappa \to 0$ for all k.

by non-uniformity in stretching rates, the resulting asymptotic mixing dynamics are somewhat similar.

In the two-dimensional examples discussed in Chapter 3, the transfer matrices constructed for the tuple of wavenumbers (k_1, k_2) are extensive and solving for eigenvalues computationally impractical. However, one-dimensional transformations produce a less extensive matrix of size $2Q \times 2Q$, and as such it is efficient to compute a spectrum of eigenvalues. Most importantly, the second leading eigenvalue and the long-time mixing rate $|\lambda_2|$, can be compared across a range of parameters.

Figure 4.7 plots how $|\lambda_2|$ changes with the diffusivity coefficient κ for the two example IETs, f_1 and f_2 of Section 4.2.4. For large values of κ the asymptotic mixing rate for both IETs is similar until around $\kappa = 10^{-3}$ where there is a divergence as κ is decreased. The profile of $|\lambda_2|$ for f_2 shows some monotonic behaviour, decreasing slightly in value before tending to 1 as $\kappa \to 0$. For all values of $\kappa < 10^{-3}$ the value of $|\lambda_2|$ for f_1 is greater than that for f_2. This agrees with the finite time mixing results of Figure 4.6, and suggests that the value of $|\lambda_2|$ may depend on the stirring properties of the IET.

For chaotic maps it was observed that the rate $|\lambda_2|$ becomes independent of diffusivity κ as $\kappa \to 0$, however since Interval Exchange Transformations are

(a) $f(r, \Pi) = f(1.5, [4312])$ (b) $f(r, \Pi) = f(1.5, [3421])$
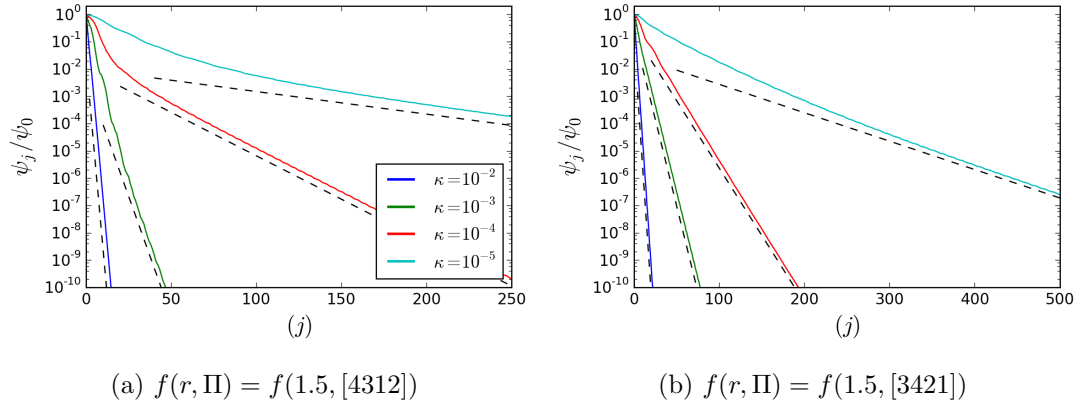
Figure 4.8: The decay of variance ψ_j/ψ_0 in iterative diffusive IETs is plotted for a range of diffusivity coefficients κ for two different IETs shown separately in (a) and (b). Depending on the IET and diffusivity coefficient, either the variance decay can be approximated well by the long time exponential decay for nearly all of the finite time mixing or large deviations occur in the initial iterations.

at most only weak mixing and not strong mixing, no spectral gap will exist in the diffusion-less limit. This implies that for the advective-diffusive operator, as $\kappa \to 0$, $|\lambda_k| \to 1$ for all k. Indeed, for the example permutations, f_1 and f_2, the values of $|\lambda_k|$ are plotted as light-grey lines alongside the values of $|\lambda_2|$ in Figure 4.7, and as $\kappa \to 0$ all profiles tend to 1 for all k with no sign of becoming independent with decreased diffusivity. As with eigenmodes resulting from smooth stirring, this long-time mixing rate is independent of almost all initial conditions. Since the stirring action of IETs is complicated it is unlikely that an initial condition would be prescribed that would not fall into this slowest decaying eigenfunction under evolution.

The effect of IET parameter choices, r, N and Π on the long-time exponential mixing rate will be discussed in a later section. The next two sections discuss the contributing factor for the rate of mixing in finite-time and the effect that the initial condition has on this mixing.

4.3.2 Finite time mixing: interaction of slowly decaying eigenfunctions

The asymptotic mixing rates for diffusive IETs with rational sub-intervals can be found easily enough and compared for changing parameters, but they tell us little of finite time mixing behaviour. In Figure 4.6 for both illustrative IETs it is

observed that although an exponential mixing rate is achieved in the long time, significant decay of variance occurs within the initial stages where the mixing rate varies from the final exponential rate.

Figure 4.8 plots the decay of variance for two new IETs. These IETs vary from Section 4.2.4 and are chosen to clearly show that deviations from the long-time mixing rate can exist for many iterations, particularly with changing diffusion. For varying rates of diffusivity coefficient κ, profiles are plotted for (a) $f(r, \Pi) = f(1.5, [4312])$ and (b) $f(r, \Pi) = f(1.5, [3421])$. The expected decay of variance from the second leading eigenvalue is plotted as a comparison. Note that the scales between both plots differ to best represent the overall decay of variance, and the initial condition for all was $c^{(0)}(x) = c^{sq}$.

In both plots when κ is large, the mixing rate achieves the final exponential mixing rate after a small number of iterations, shown by the dark blue and green lines. This similarly holds for $f(1.5, [3421])$ in Figure 4.8 (b) when $\kappa = 10^{-5}$, where there is little deviation from the final mixing rate across the full profile. For $f(1.5, [4312])$ in Figure 4.8 (a), the initial stages of mixing show a steeper descent in variance before the long-time mixing rate is achieved when $\kappa = 10^{-5}$. For a smaller diffusivity rate $\kappa = 10^{-6}$, both profiles have early, non-exponential like mixing stages where significant depletion of variance occurs before tending to the final exponential mixing rate approximated by $\psi_j = C|\lambda_2|^{2j}$.

Previous work of Wang & Christov (2018), investigating finite time mixing in IETs with diffusion, suggests a stretched exponential function approximates the averaged profiles of variance decay across permutation groups and fixed r. They propose an appropriate fit is given by

$$||c(x, j) - c(x, 0)||_p = C \cdot e^{-(j/\bar{T})^\gamma}, \tag{4.21}$$

and find good agreement in the initial stages. Fitting to numerical results, values for the parameters $\bar{T}$ and γ are found, but the authors state that both have non-trivial dependencies on the length ratio r and diffusivity κ. No mechanism was suggested for why a stretched exponential mixing rate emerges.

The numerical method used in Wang & Christov (2018) models the evolution of concentration limited to the equal sized cells in the construction of the large permutation (4.9) for given IET parameters r and Π. Diffusion is applied as a simple finite-difference step between consecutive cells. This means that comparing mixing between different IETs for the same diffusivity coefficient requires a rescaling in time to take into account the spatial resolution problems of the

$f(r, \Pi)$	λ_1	$|\lambda_2|$	$|\lambda_3|$	$|\lambda_4|$	$|\lambda_5|$	$|\lambda_6|$
$f(1.25, 6)$	1.0	0.972014	0.972014	0.957913	0.957913	0.951927

$f(\Pi, r)$	$|\lambda_7|$	$|\lambda_8|$	$|\lambda_9|$	$|\lambda_{10}|$	$|\lambda_{11}|$	$|\lambda_{12}|$
$f(1.25, 6)$	0.951927	0.950489	0.950489	0.949988	0.949988	0.948377

Table 4.1: Leading Eigenvalues in the spectrum for IET $f_2(r_2, \Pi_2)$ with $\kappa = 10^{-5}$
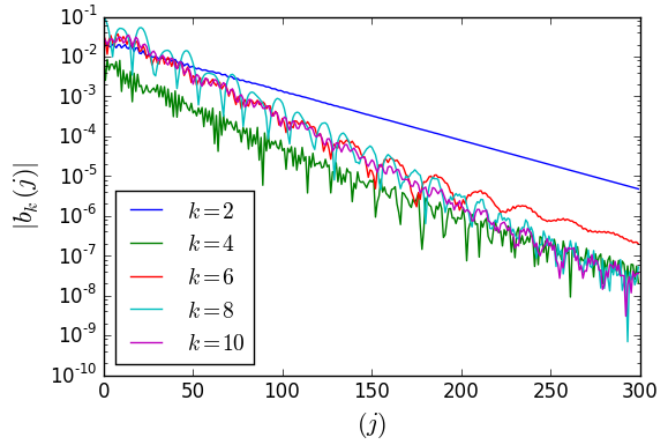
Figure 4.9: The evolution of the modulus of the coefficients $b_k(j)$ with iteration number j for the IET $f_2(r_2, \Pi_2)$. An influence of eigenfunctions with $k \geq 6$ occurs in the initial iterations until $k \sim 100$ when $b_2(j)$ emerges as the dominant coefficient for all future times. The initial condition was $c^{(0)}(x) = c^{sq}$ and $\kappa = 10^{-5}$.

numerical method. Therefore a direct comparison between the model presented here and the results of Wang & Christov (2018) is not easily achieved due to a discrepancy in the diffusion coefficient κ, and their parameter D defining the sharing factor of concentration between adjoining cells. However comments on their conclusions, and comparison to the results presented herein, will be made throughout the discussion in this and following sections.

Recall the initial example of Section 4.2.4 where the variance decay for the diffusive IET f_2 was plotted with $\kappa = 10^{-5}$. Two exponential mixing rates with oscillations were observed. The modulus for the first leading eigenvalues are listed in Table 4.1, revealing that the eigenvalues are particularly close in value. All eigenvalues appear in complex conjugate pairs for those listed, which is why $|\lambda_{2k}| = |\lambda_{2k+1}|$ for $k = 1, 2, \ldots$ across the table.

The eigenvalues listed in Table 4.1 are very close in value, which occurs for all eigenvalues in the spectrum of $P = P_D \circ P_M$ for IETs as $\kappa \to 0$. If eigenfunctions

v_k of the operator are decaying on average at a rate given by $|\lambda_k|$, an eigenfunction with only a slightly faster mixing rate may emerge as a dominant feature in the concentration field for several stirring periods before the slowest decaying eigenfunction emerges as the long-time dominant persistent pattern. In line with past literature [Popovych *et al.* (2007)], eigenvalues close in modulus value will be referred to as *close-to-degenerate*, since degenerate eigenvalues are those which are equal in value.

Popovych *et al.* (2007) investigated if close-to-degenerate eigenvalues contributed to non-exponential, finite time mixing rates in smoothly stirred systems, such as the polynomial mixing rates in flows with islands or boundaries. Here, the interaction of slowly decaying eigenfunctions are investigated as the potential mechanism for the finite time mixing behaviour of IETs with weak diffusivity. If the concentration field evolves via the summation of the underlying eigenfunctions, then

$$c(x, j) = \sum_k b_k(j) v_k(x). \tag{4.22}$$

Since the eigenfunctions are decaying as a product with their respective eigenvalues, the concentration field can be approximate as

$$c(x, j) = \sum_k a_k \lambda_k^j v_k(x), \tag{4.23}$$

for some coefficients a_k. After some time we would expect lower order terms which have smaller values of $|\lambda_k|$ can be ignored and the decay of variance decay approximated by

$$\psi(c(j)) = \langle (c(j) - \bar{c})^2 \rangle \approx d_2 \lambda_2^{2j} + d_3 \lambda_3^{2j} + \dots, \tag{4.24}$$

which is the precise approximations taken by Popovych *et al.* (2007). However, in contrast to that work here the eigenvalues in the limit of weak diffusivity result in many close-to-degenerate eigenvalues. An eigenfunction analysis can be performed to compute the coefficients $b_k(j)$ and observe the competing interaction of eigenfunctions and how they contribute to the finite time mixing behaviour. The coefficients $b_k(j)$ from (4.22) can be found by evaluating the inner product (2.37) from Section 2.3.4, and since the coefficients $b_k(j)$ are complex, the modulus $|b_k(j)|$ plotted to see competing influence the eigenfunctions have on the concentration field at each iteration.

Figure 4.9 plots the modules $|b_k(j)|$ for $k = 2, 4, 6, 8, 10$ against iteration number j for the IET $f_2 = f(1.25, 6)$ from Section 4.2.4. In the initial iterations,
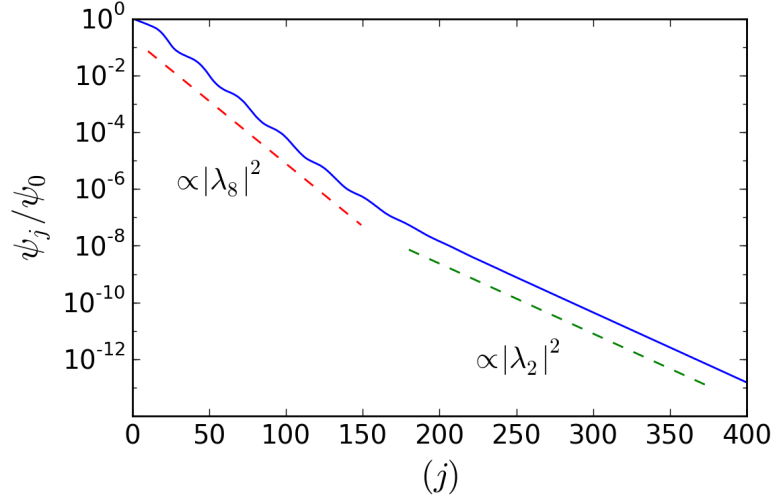
Figure 4.10: Variance decay for $f_2(r_2, \Pi_2)$ with $\kappa = 10^{-5}$ and $c^{(0)}(x) = c^{sq}$. Expected variance decays $\psi = C|\lambda_k|^2$ match well the exponential decay rates which occur.

the dominance in the values of $|b_k(j)|$ fluctuate in value between $k = 2$, 6, 8 and 10, with no single eigenfunction a dominating influence on the concentration field. For many iterations, the eigenfunction v_8 dominates the concentration field with the largest values in $|b_8(j)|$, until around iteration 100 when $|b_2(j)|$ dominates for all future iterates, associated with the eigenfunction v_2, the slowest decaying eigenfunction. Figure 4.10 replots the variance for f_2 with decay rates $\psi \propto |\lambda_8|^{2j}$ and $\psi \propto |\lambda_2|^{2j}$ showing good agreement for each respective stages.

The examples in Figure 4.8 can be used to see if close-to-degenerate eigenvalues are the mechanism contributing to the finite time mixing profiles observed which deviate from exponential in the initial stages. Figure 4.11 plots the coefficients $|b_k(j)|$ for the first five slowest decaying eigenfunctions, where the complex conjugate eigenfunctions have been neglected since they result in the same values of $|b(j)|$. In both cases, in the early stages of finite time evolution the modulus of the coefficients $|b_k(j)|$ are greater for a couple of values $k > 2$, than $|b_2(j)|$ relating to the dominant eigenfunction. This suggests that in the earlier iterations, the concentration field is influenced by the interaction of many eigenfunctions decaying with similar decay rates.

An approximation of the variance decay can be computed from the evolution of the leading eigenfunctions and coefficients $|b_k(j)|$ as

$$\psi_{approx}^K = \sum_{k=2}^{K} |b_k(j)v_k|^2 \tag{4.25}$$

89

(a) $f(r, \Pi) = f(1.5, [4312])$

(b) $f(r, \Pi) = f(1.5, [3421])$
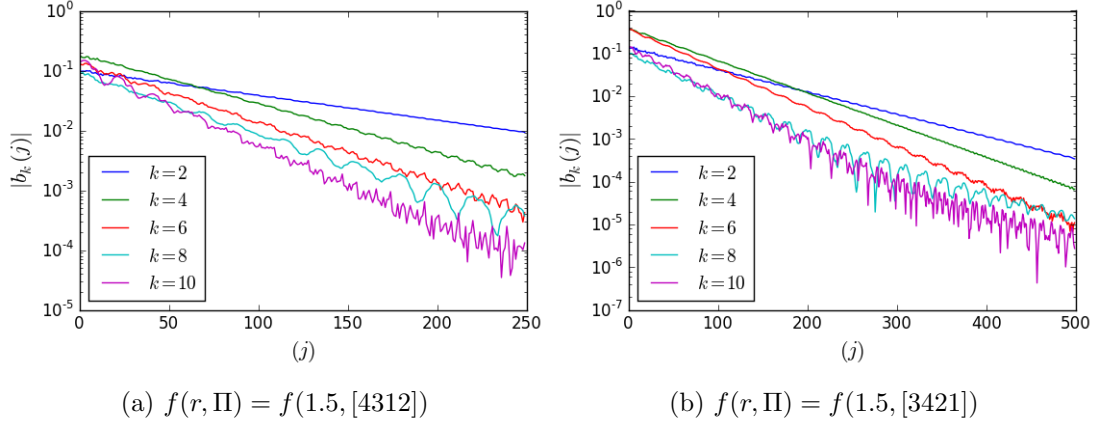
Figure 4.11: The evolution of the modulus of the coefficients $b_k(j)$ with iteration number j are plotted for two IETs. The initial condition was $c^{(0)}(x) = c^{sq}$ and $\kappa = 10^{-5}$.

(a) $f(r, \Pi) = f(1.5, [4312])$

(b) $f(r, \Pi) = f(1.5, [3421])$
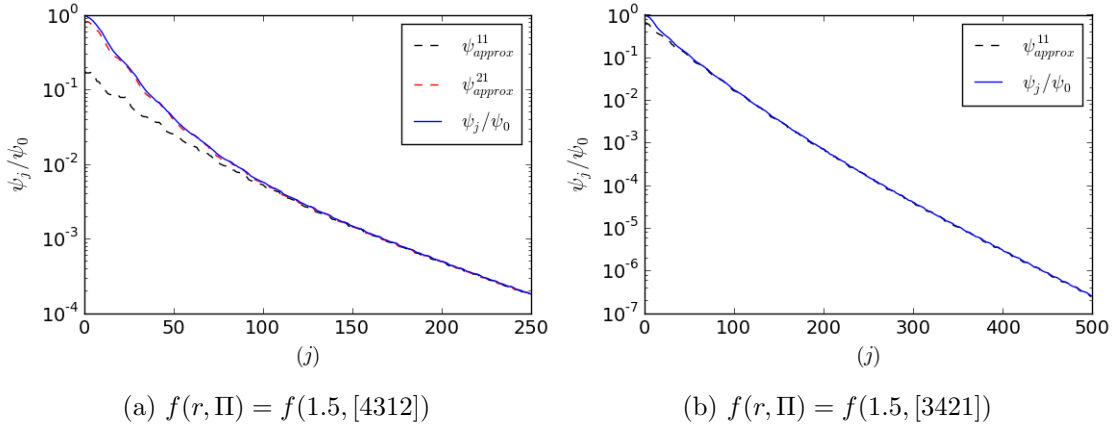
Figure 4.12: The approximation of the variance decay from the interaction of slowly decaying eigenfunctions by (4.25) are plotted alongside the full variance decay for the initial condition $c^{(0)}(x) = c^{sq}$ with $\kappa = 10^{-5}$.

and plotted alongside the full variance evolution of from the initial condition $c^{(0)}(x) = c^{sq}$, as is shown in Figure 4.12. In Figure 4.12 (a) $K = 11$ under approximates the variance decay and $K = 21$ is required for a more accurate approximation of the variance decay profile, while in Figure 4.12 (b) $K = 11$ approximates well the variance profile. The former is not surprising due to the values of $|b_k(j)|$ being close in value for the initial iterations in Figure 4.11. These examples reveal that the interaction of the slowly decaying eigenfunctions indeed contribute to the finite time evolution of IETs with a diffusive step, however since all eigenvalues have $|\lambda_k| \to 0$ for all k as $\kappa \to 0$, for weak diffusivity several slowly decaying eigenfunctions interact and contribute to the observed mixing behaviour.

4.3.3 Effect of initial condition

The asymptotic mixing rate of an IET is independent of almost all initial conditions, as discussed in Section 4.3.1. In the previous section, the initial stages of mixing were well approximated by the combination of several eigenfunctions with similar decay rates, however how long each of the eigenfunctions dominates the concentration field, and how this depends on the initial condition, was not discussed. Popovych *et al.* (2007) note that in their systems of interest, the time interval over which the non-exponential decay of concentration is observed depends strongly on initial conditions. This reflects a strong variation in the coefficients d_k in (4.24). Is this also observed for the non-chaotic transformations investigated here?

A return to the time to achieve a mixed condition, specifically t_{95}, allows a clear understanding of the effect of initial conditions on finite time mixing rates. For the example IET $f(1.5, [4312])$, Figure 4.13 (a) plots the decay of variance for different initial conditions when the diffusion coefficient is taken as $\kappa = 10^{-6}$. The changes in the initial condition considered were five rotations of the square wave, such that $c^{\alpha}(x) = c^{sq}(x + \alpha)$ where $\alpha \in \{0.0, 0.05, 0.1, 0.15, 0.2\}$. It is seen that changing the initial condition in such a simple way has a considerable effect on the decay of variance. The time to achieve a variance of decay below 5% of the mean field varies over 200 iterations. For $\alpha \in \{0.25, 0.3, 0.35, 0.4, 0.45, 0.5\}$ the plot is similar to Figure 4.13 (a), and due to the rotational symmetry of the square wave, for $\alpha > 0.5$ the results are repeated.

The IET $f(1.5, [4312])$ is an irreducible, non-rotational permutation and has a short periodic order of 65. Figure 4.13 (b) shows how the mix norm changes
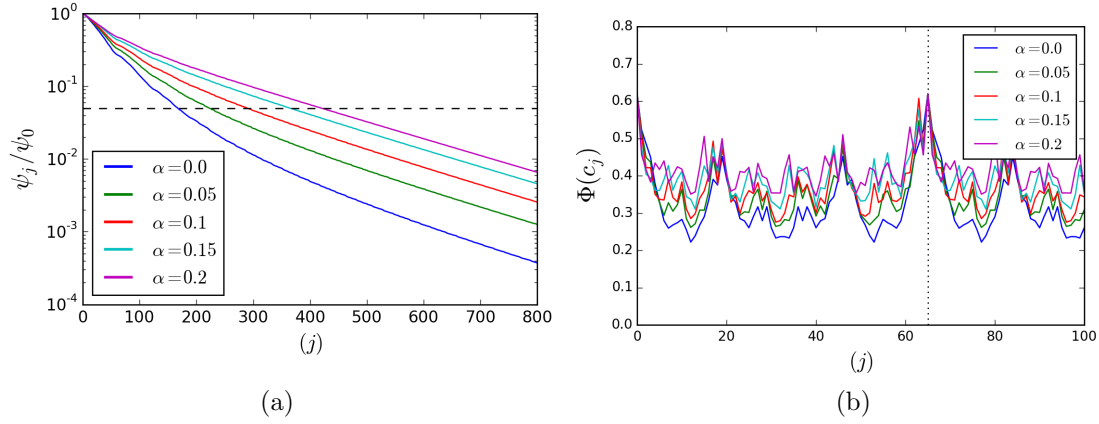
Figure 4.13: (a) The decay of variance for the IET $r = 1.5$, $\Pi = [4312]$ with $\kappa = 10^{-6}$ shown for various rotations of the initial condition the square wave, representing half white, half black. Dashed line represents when ψ_j/ψ_0 is within 5% of the mean field. (b) Mix norm measurement for the segregation by the IET $f(1.5, [4312])$ in the diffusion-less limit for rotated initial conditions of the square wave. For all ICs rejoining occurs are 65 periods, however between these periods the mixing at each iteration depends on the initial condition.

for the rotated square wave initial conditions under evolution of the IET in the diffusion-less case. The mix norm is calculated using the recursive computational method of Smith (2016) on an extended lattice with 10,000 points sampling the concentration field. Reassembly of the initial condition occurs after just 65 iterations, however the value of the mix norm between these peaks varies across the range of α. For the case $\alpha = 0.2$, which takes the longest time to achieve the mixed condition t_{95} in Figure 4.13 (a), the stirring is the least effective in decreasing the value of the mix norm..

The effect of changing the initial condition was considered by Yu et al. (2016), but only in comparing across the parameter value r and averaging across permutation groups. They did not report any variation in mixing properties of IETs due to varying initial conditions. Of course, this is a specific example and with a small periodic order of 65 stirring periods, does not satisfy the parameter specifications for good finite time mixing by Krotter et al. (2012) and Yu et al. (2016). There is no mention of the effect of initial conditions in the diffusive IETs of Wang & Christov (2018).

It would be hoped that an IET that mixes well for one initial condition should mix well for nearly all initial conditions. Therefore, in addition to the time to achieve a mixed condition $t_{\%}$, an additional quantitative measure of the ability

92

to mix well in finite time will be $t_\%^{max} - t_\%^{min}$. For a fixed IET with parameters defined by r and Π, and the rotated square wave initial condition parameterized by α, $t_\%^{max}$ denotes the largest number of iterations taken to achieve the mixed condition while $t_\%^{min}$ will be the smallest number of iterations required, when varying α. Therefore a small value of $t_\%^{max} - t_\%^{min}$ suggests that the IET mixing properties do not vary much across the range of initial conditions, while a large $t_\%^{max} - t_\%^{min}$ implies the mixing properties are sensitive to the initial condition.

4.4 Dependence on parameters

So far a handful of example permutations have been used to explain in detail the mixing behaviour and resulting mixing rates that emerge in IETs with a diffusive step. The parameter r defining the ratio between sub-interval lengths has only had the values $r = 1.5$ or $r = 1.25$, and similarly the rearrangement in the shuffling has only been given by irreducible, non-rotational permutations. Now an aim to quantify the effect of parameters defining an IET $f(r, \Pi)$ is carried out. These are the permutation arrangement Π, the sub-interval length scaling parameter r, and the number of sub-intervals N.

4.4.1 Effect of permutation rearrangement

The simplest division of permutation groups into subgroups depending on their rearrangement order are given by rotation permutations, reducible permutations or irreducible non-rotational permutations. The asymptotic and finite time mixing rates for these three groups are investigated for fixed r.

Rotation permutations

In the case of periodic boundary conditions, the mixing rates by rotation permutations can be explicitly computed. An IET $f(r, \Pi)$ where Π is a rotation will result in the permutation Π_f which is also a rotation, such that

$$\Pi_f(i) = i + m \mod N_l \tag{4.26}$$

for all $j \in \{1, 2, \ldots, N_f\}$, where m is fixed and dependent on the rotation. This results in the map

$$M_{\Pi_f}(x) = x + \frac{m}{N_f} \mod 1. \tag{4.27}$$

(a) Space-time plot for $f(1.5, [1432])$ (b) Space-time plot for $(f1.5, [2143])$
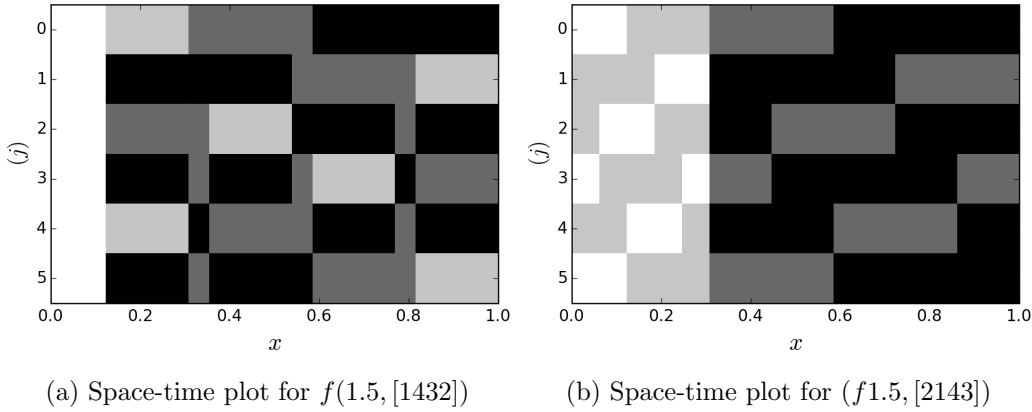
Figure 4.14: Space-time plots for the evolution of a concentration field under the action of two reducible permutations. A component initial condition [Krotter *et al.* (2012); Yu *et al.* (2016)] shows that a separation of the domain occurs preventing further stirring and reduction in the scale of segregation.

The transfer matrix for the IET then has elements $\boldsymbol{M}_{kk} = \omega^m$ where $\omega = e^{2\pi i/N_f}$, and with $|\boldsymbol{M}_{kk}| = 1$ while $\boldsymbol{M}_{kq} = 0$ for $k \neq q$. Since the transfer matrix is a diagonal matrix, representing only a shift in phase for each of the Fourier modes, there is no dispersion between other modes. Therefore mixing to uniformity can only occur through the action of diffusion. The eigenvalues can be explicitly written down since $\boldsymbol{M}_{nq}$ and $\boldsymbol{d}_{kn}$ are diagonal matrices and

$$\boldsymbol{M}\boldsymbol{d} - \lambda\boldsymbol{I} = 0 \implies \lambda_k = \boldsymbol{M}\boldsymbol{d}_{kk} \tag{4.28}$$

where $\boldsymbol{M}\boldsymbol{d}_{kk} = \omega^m \rho^{k^2}$, such that $|\lambda_2| = \rho$ for all rotations. In the initial stages, the finite time mixing rate depends only on the initial condition and will be the same for all rotation permutations, including the identity permutation.

Reducible and irreducible permutations

As an illustrative example, the space-time evolution of two IETs with $r = 1.5$ and reducible permutations from S_4 are shown in Figure 4.14. The first permutation $\Pi = [1432]$ plotted in Figure 4.14 (a) is one of the simplest permutations to satisfy the reducible definition; in which the first element is fixed but the others are free to be shuffled within the rest of the domain. Assigning the component coloured initial condition from Krotter *et al.* (2012); Yu *et al.* (2016), it is easily seen that this reducible permutation allows no mixing of the first sub-interval into the rest of the domain, while cutting and shuffling generates some smaller

(a) Space-time plot for $f(1.5, [3241])$ (b) Space-time plot for $f(1.5, [4231])$
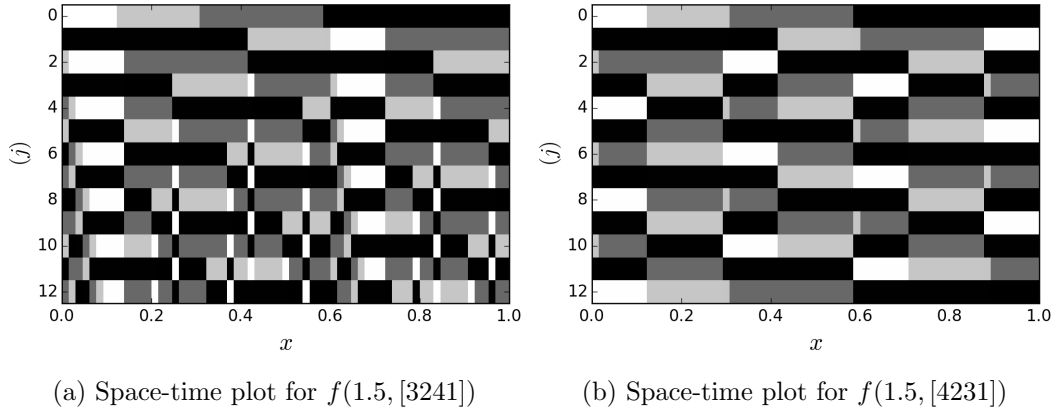
Figure 4.15: Space-time plots for the evolution of a concentration field under the action of two irreducible permutations. A component initial condition [Krotter *et al.* (2012); Yu *et al.* (2016)] shows the respective mixing from each. Although (a) suggests irreducibles reduce the scale of segregation well, a counter example permutation in (b) shows in some cases they result in bad mixing and low periodic order, resulting in reassembly of the initial condition after just 12 iterations.

sub-segments of colour in the remainder of the domain. Figure 4.14 (b) shows a second reducible permutation $\Pi = [2143]$ in which the first two, and last two sub-intervals are swapped, again resulting in a decomposition of the domain into two sub-intervals.

These examples reveal that reducible permutations as defined by Definition 4.1.4 separate the domain in which the IET is acting into two or more distinct intervals. These distinct sub-intervals may have their own segregating dynamics, however between them no transfer of concentration can occur via stirring. This implies that uniformity can be only achieved via diffusion across the boundary between the intervals.

In contrast, Figure 4.15 plots the space time evolution for two irreducible permutations from S_4 with $r = 1.5$. The first in Figure 4.15 (a) is good at reducing the scale of segregation, generating smaller and smaller sub-segments with each cut and shuffle. The second in Figure 4.15 (b) however does not, and reassembles the initial condition after just 12 iterations. These examples highlight some of the subtleties of mixing in IETs. Yu *et al.* (2016) put the bad mixing properties of $f(1.5, [4231])$, and like permutations, where the first and last elements are simply swapped, down to consecutive sub-intervals contaminating the potential for mixing at each iteration. They neglect the fact that the additional rational scaling r between all sub-intervals contributes too, and produces short periodic

95
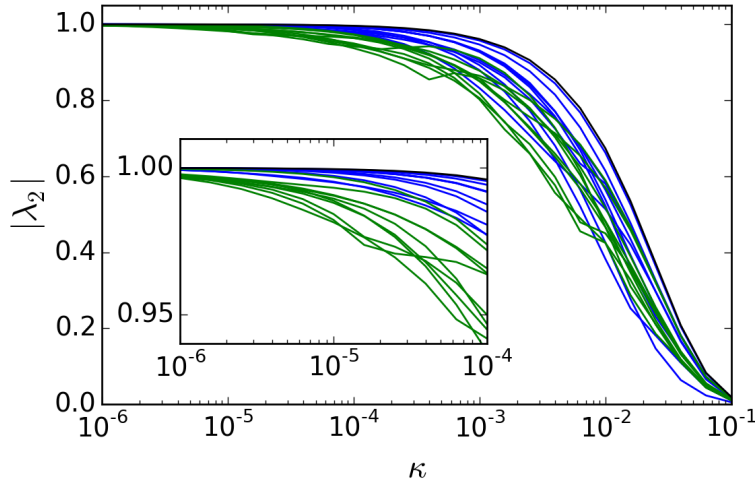
Figure 4.16: Asymptotic mixing rate $|\lambda_2|$ varies with κ for all $\Pi \in S_4$ with $r = 1.5$. Profiles for rotations (black), reducible (blue) and irreducbile, non-rotational (green) permutations are distinguished by colour.

orders for permutations of this type. Varying the lengths of the sub-intervals so that there is no common ratio between them all produces better reduction in the scale of segregation and a longer periodic order for these permutations. However consecutive sub-intervals do restrict the potential decrease in segregation scale at each iteration. This is discussed in more detail in a later section. The IET $f(1.5, [3241])$ has a periodic order of 177, while all other irreducible, non-rotational permutations from S_4 with $r = 1.5$ have a periodic order of 65, equivalent in value to the length N_f of the larger permutation Π_f for these parameters. The relevance of N_f and the periodic order of an IET on mixing is more important when discussing the parameters r and N and as such is left for the discussion of those sections.

Now that the segregation properties of different permutations groups are understood, the resulting mixing rates from including diffusion iteratively are investigated. Unlike rotations, it is not possible to analytically solve for the leading eigenvalues of reducible and irreducible, non-rotational permutations. Instead numerical results on the asymptotic and finite time mixing rates are quantitatively compared.

Figure 4.16 shows for the IETs $f(1.5, \Pi \in S_4)$ how the second leading eigenvalue varies with diffusivity rate. The lines are coloured dependent on the permutation properties, black for the identity or rotational permutations, blue for

	Mean	Median	Max	Min	Mean	Median	Max	Min
$\kappa = 10^{-3}$	13.58	13.05	25.6	7.1	8.14	8.05	10.8	5.9
$\kappa = 10^{-4}$	102.86	89.05	237.8	48.5	26.76	23.95	37.3	19.1
$\kappa = 10^{-5}$	905.53	813.7	2333.6	245.4	74.31	67.6	140.5	44.2
$\kappa = 10^{-6}$	8639.9	7932.65	23202.7	1158.8	256.43	243.25	498.4	139
		(a)					(b)	

Table 4.2: Statistics of t_{95} averaged over varying initial condition for (a) reducible permutations and (b) non-rotational, irreducible permutations from S_4 with $r = 1.5$. The initial conditions were rotations of the square wave with 10 evenly distributed values of α in $[0, 0.5)$.

the reducible permutations, and green for the irreducible, non-rotational permutations.

For fast diffusivity, $\kappa > 10^{-4}$ the profiles of $|\lambda_2|$ are complicated, varying non-monotonically in some cases, with overlapping profiles regardless of whether the permutation is reducible or irreducible. This was similarly reported by Ashwin *et al.* (2002) for permutations of equal sized cells, the order of permutations by the value of $|\lambda_2|$ changes with the value of κ, thus there is not one permutation defining the IET which mixes better for all diffusion coefficients. The rotations and identity permutation always have the slowest long-time mixing rate, determined by the slowest decaying mode from diffusion alone.

As κ is decreased, on average it appears as though the irreducible, non-rotational permutations have a quicker long-time mixing rate, seen in the inlay plot. This result implies that irreducible, non-rotational permutations are better at mixing, which is not surprising since Krotter *et al.* (2012) and Yu *et al.* (2016) report that IETs with irreducible, non-rotational permutations are also better at reducing the scale of segregation in finite time. The slow long-time mixing rates for reducible permutations with weak diffusivity is due to the large unstirred regions of the domain and limited mixing across their boundaries by diffusion only. Since the scale of this unmixed region will be much larger than that which can be achieved by an irreducible permutation, the rate of decay of the largest mode capturing this diffusion-limited unmixed region will be much slower.

As discussed in Section 4.3, the asymptotic time does not reveal the whole finite time mixing behaviour, which is highly sensitive to properties of the IET and the initial condition. For $r = 1.5$ and each $\Pi \in S_4$ the time to achieve

a 95% mixed state, t_{95}, was computed for 10 different initial condition and the mean value taken. The different initial condition were rotations of the square-wave (4.18) with α evenly distributed in the interval [0,0.5). The average of t_{95} is computed since, as previously discussed, large variations in finite time mixing can occur for simple variations in initial condition.

Table 4.2 lists the statistics for averaged values of t_{95} across the permutation group S_4, split for the reducible permutations (of which there are 10) and irreducible non-rotational permutations (10). In each row the value of the diffusivity coefficient is changed to show the drastic differences in how many iterations are required to achieve the mixed condition depending on Π. For all diffusivity values the irreducible, non-rotational permutations have lower values across all statistical measures. In line with the asymptotic mixing rate, irreducible, non-rotational permutations are significantly faster at mixing to a desired condition than reducible permutations and confirms that choosing such a permutation is essential to achieve good mixing; in both diffusive and diffusion-less transformations.

	N=4	N=5	N=6	N=7
r=1.5	65	211	665	2059
r=1.2	671	4651	31031	-
r=1.1	4641	61051	-	-

Table 4.3: N_f for the parameter choices of $\Pi \in S_N$ and different values of r. On average the periodic order of IETs from this parameter space have $order = N_f$ with a few outliers, more often larger than N_f.

4.4.2 Effect of sub-interval ratio r

Krotter *et al.* (2012) and Yu *et al.* (2016) concluded that better reduction in the scale of segregation can be achieved when the ratio r between sub-interval lengths is "closer-to-irrational". This is because such values of r parametrising an IET essentially generates a permutation of equal sized cells with large length N_f, and subsequently the possibility to achieve a small scale of segregation of width $1/N_f$. Since at each iteration a diffusive step will begin to blur sharp gradients in the concentration field between sub-segments of varying colour, this preciseness in the possible concentration scale will be lost. A numerical parameter study on the effect of r highlights whether the conclusions of previous stirring studies carry over into IETs with a diffusive step. Additionally r closer to, but not equal to, 1 achieves the best stirring results, therefore only $r \in [1.1, 1.5]$ is considered here.

Table 4.3 lists N_f for a variety of parameter pairings $f(r, \Pi \in S_N)$. Three values of r are listed; $r = 1.5$, 1.2, and 1.1. As the value of r decreases for these selected values, the length of the extended permutation Π_f increases roughly by one order of magnitude. For fixed r increasing the number of sub-intervals N in the permutations also increases the value of N_f. Since increasing N also decreases the possible scale of segregation, it would be difficult to determine whether it is a larger N_f, or more intervals N, which contribute to better mixing. However these values allow for a comparative study between the two since if it is larger N_f that contributes mainly to better mixing overall, it would be expected that $f(1.2, \Pi \in S_4)$ and $f(1.5, \Pi \in S_6)$ would have similar mixing properties due to having comparatively similar values of N_f. Similarly $f(1.1, \Pi \in S_4)$ and $f(1.2, \Pi \in S_5)$. Since rotational and reducible permutations result in bad mixing, only the subset of irreducible, non-rotational permutations are taken from the groups S_N; of which there are 10 in S_4, 67 in S_5 and 457 in S_6.

First the effect of r on the mixing rates is investigated. The long-time mixing rates, finite time mixing rates and the spread of finite time mixing rates dependent on the initial conditions are compared across the values of r, given by $|\lambda_2|$, t_{95} and $t_{95}^{max} - t_{95}^{min}$ respectively. Since t_{95} is variable dependent on the initial condition, as was done when comparing reducible and irreducible permutations, an averaged value will be taken for each permutation across 10 initial conditions; rotations of the square-wave parametrised by α. Box plots are used to graphically depict the spread in these measures across the permutation groups S_N. Like all methods to show a spread in values, there is a loss in information by not including all data points. However, box plots were chosen since they capture the median (red line),
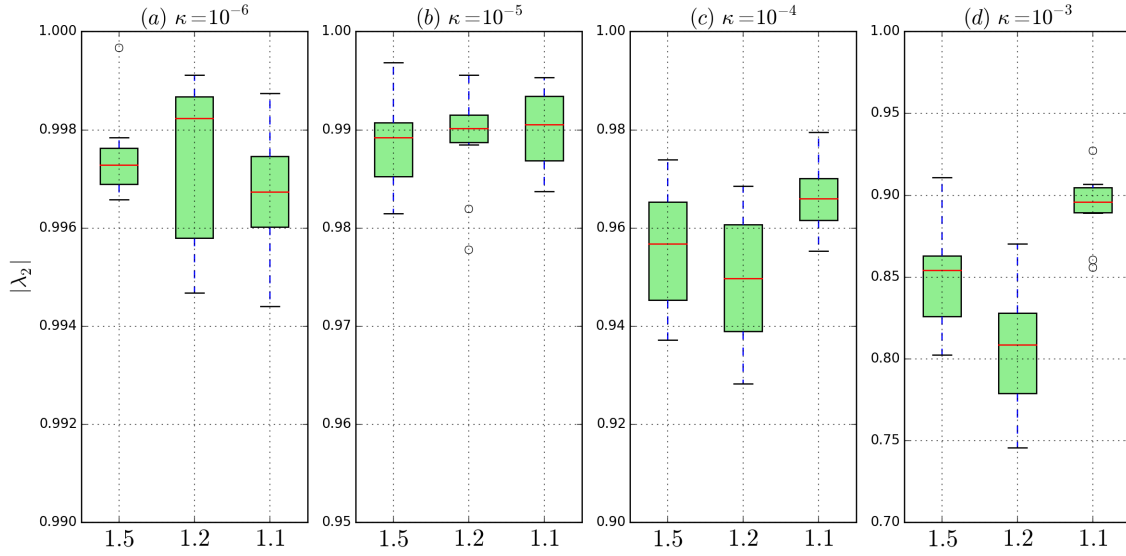
Figure 4.17: Boxplots showing the spread of the modulus of the second leading eigenvalue $|\lambda_2|$ for non-rotational irreducible permutations from S_4 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-6}$, c) $\kappa = 10^{-6}$, and d) $\kappa = 10^{-6}$,
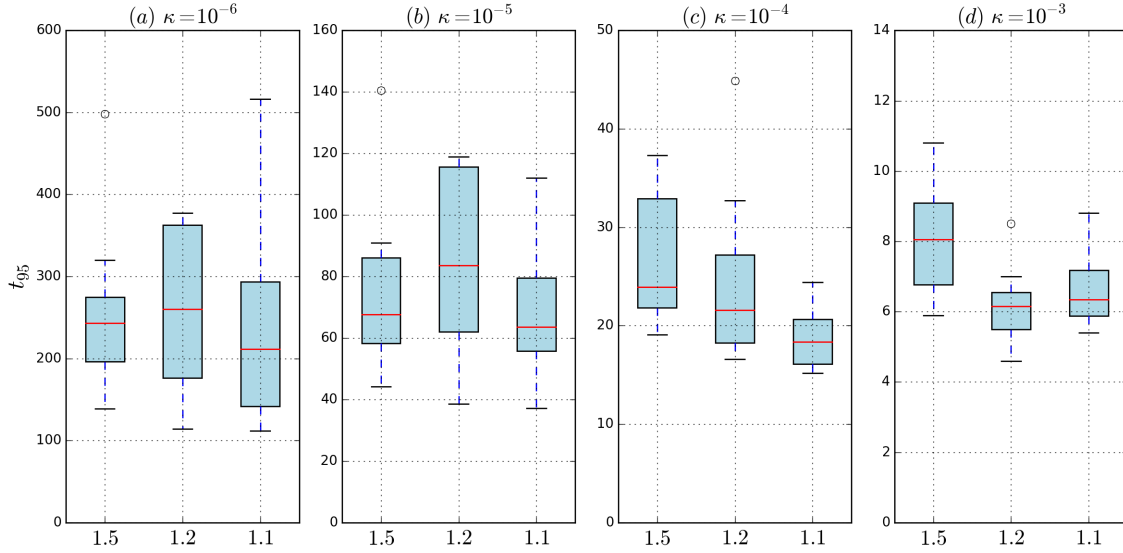
Figure 4.18: Boxplots showing the spread of t_{95} for non-rotational irreducible permutations from S_4 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-5}$, c) $\kappa = 10^{-4}$, and d) $\kappa = 10^{-3}$,

upper and lower quartiles (top and bottom of boxes respectively), and maximum and minimum values in a data set excluding anomalies (whiskers) presenting a sizeable amount of information on the spread, as opposed to say the mean and standard deviation. Similarly, box plots seem more appropriate since each data point in the spread is a unique mixing measure for a single IET ($r\Pi$) for $\Pi \in S_N$, rather than one realisation of the action of (r, Π) with some small variation in the conditions, such as a different initial concentration field.

Figure 4.17 shows box plots representing the statistics on the spread of $|\lambda_2|$ across each of the 10 irreducible non-rotational permutations $\Pi \in S_4$. The subfigures represent different values of the diffusion coefficient κ, and note that each has a varying y-axis scale to better capture the spread of values. Box plots are shown to represent a selection of the full $|\lambda_2|$ profiles for easy comparison across parameters. Additionally, the spread in values is captured better via the median and upper and lower quartiles, than from the mean value and standard deviation alone.

Across the range of diffusivity coefficients, no value of r emerges which has a faster on average long-time mixing rated for all values of κ. When κ is large, it appears as though $r = 1.1$ has the slowest overall long-time exponential decay, while for the smallest value of $\kappa = 10^{-6}$, the median and lower quartiles are the lowest in value compared to $r = 1.5$ and $r = 1.2$. Although $r = 1.1$ for $\Pi \in S_4$ has the largest N_f, these results suggest that the long-time mixing rate determined by the decay rate of the dominant eigenfunction is not necessarily governed by the possibility to reduce the scale of segregation in finite time.

The finite time mixing behaviour similarly shows little advantage in increasing the value of N_f via choosing a "more irrational" r. The spread of t_{95} averaged over varying initial conditions is plotted in Figure 4.18, again as box plots showing the median, upper and lower quartiles and general spread of all values. In Figure 4.18 (c) when $\kappa = 10^{-4}$ a trend is seen in which choosing a "more irrational" r, resulting in larger N_f, on average gives a quicker finite time mixing rate, however, this trend does not hold in any of the other sub-figures (a), (b) and (d). Figure 4.19 plots the variation in the time to achieve a mixed condition with varying initial condition, $t_{95}^{max} - t_{95}^{min}$, for small diffusivity coefficients and shows that although the median values for all three parameters are very similar, a number of permutations $r = 1.5$ and $r = 1.2$ result in larger variations depending on the initial condition in contrast to $r = 1.1$ which shows less variation overall.
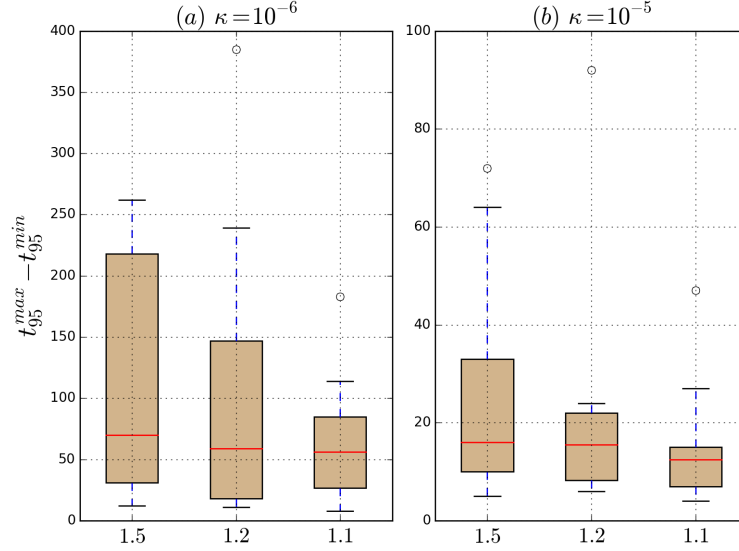
Figure 4.19: Boxplots showing the spread $t_{95}^{max} - t_{95}^{min}$ for non-rotational irreducible permutations from S_4 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$ and b) $\kappa = 10^{-5}$.
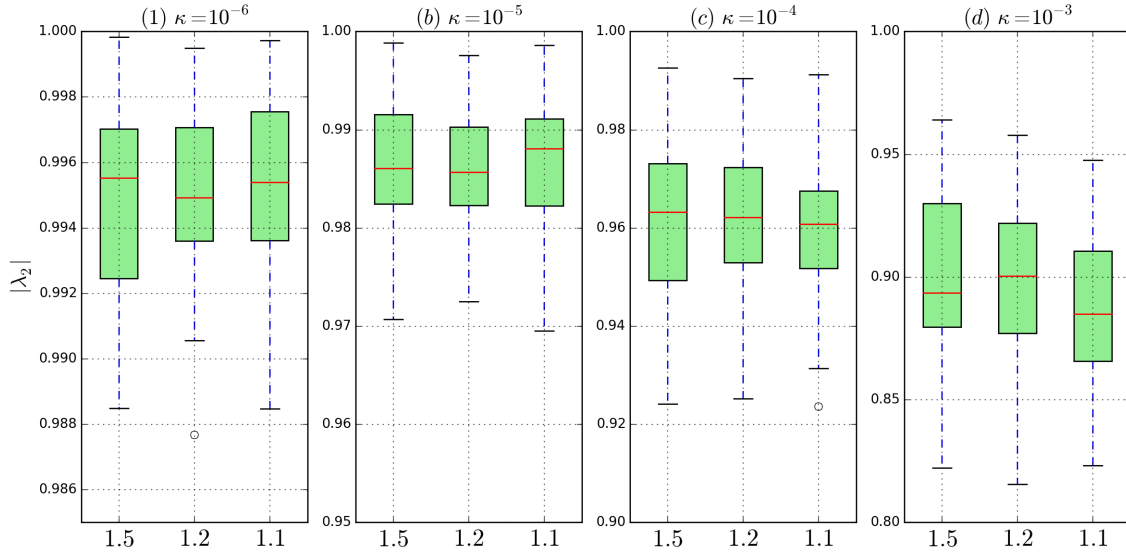
Figure 4.20: Boxplots showing the spread of the modulus of the second leading eigenvalue $|\lambda_2|$ for non-rotational irreducible permutations from S_5 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-5}$, c) $\kappa = 10^{-4}$, and d) $\kappa = 10^{-3}$.
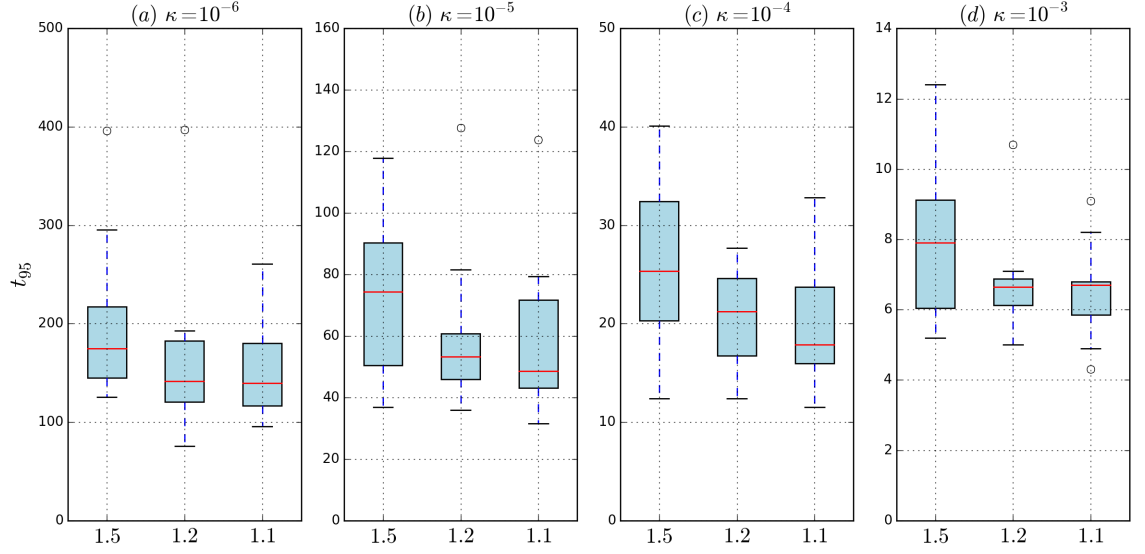
Figure 4.21: Boxplots showing the spread of t_{95} for non-rotational irreducible permutations from S_5 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-5}$, c) $\kappa = 10^{-4}$, and d) $\kappa = 10^{-3}$.
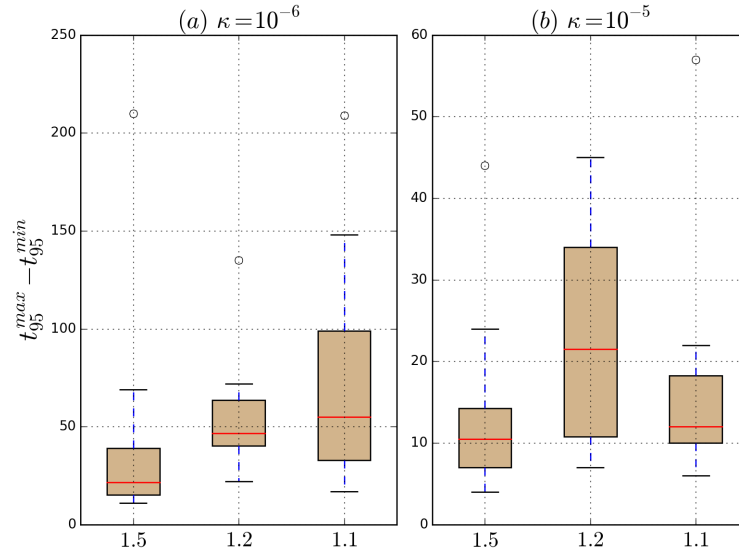
Figure 4.22: Boxplots showing the spread of $t_{95}^{max} - t_{95}^{min}$ for non-rotational irreducible permutations from S_5 dependent on the sub-interval ratio $r = 1.5$, 1.2 and 1.1. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$ and b) $\kappa = 10^{-5}$.

A repeated study using irreducible, non-rotational permutations $\Pi \in S_5$ was performed and showed similar results. Since S_4 only has 10 suitable permutations, S_5 with 67 provides a better number of IETs to compare across. In this case the variations in the value of N_f are larger with $N_f = 211$ for $r = 1.5$, $N_f = 4651$ for $r = 1.2$ and $N_f = 61051$ for $r = 1.1$. Even though the size of the smallest scale of segregation changes drastically for varying r in this case, Figure 4.20 shows for all values of κ the spread in the long-time exponential mixing rate given by $|\lambda_2|$ has little correlation to this fact.

In contrast to S_4, in Figure 4.21 the box plots for the average value of t_{95} appear to suggest $r = 1.5$ has the slowest finite time mixing rates on average when compared to $r = 1.2$ and $r = 1.1$ across all diffusivity values, except for one or two outliers. Looking again at $t_{95}^{max} - t_{95}^{min}$ in Figure 4.22, the trend in behaviour is now opposite to that reported in Figure 4.19, with $r = 1.5$ showing the least variation in the time to achieve a mixed condition from changing the initial condition.

These numerical results are inconclusive in suggesting having a larger N_f from the choice of r results in better mixing in diffusive IET transformations. In the next section the effect of including more sub-intervals N is investigated and the conclusions compared to the results thus far.

4.4.3   Effect of the number of sub-intervals N

The effect of N on reducing the scale of segregation in cutting and shuffling IETs was discussed by Krotter *et al.* (2012) and Yu *et al.* (2016) in addition to the parameter r. Krotter *et al.* (2012) noted diminishing returns on the effect of N to achieve on average a smaller, largest unmixed sub-segment U_j over the initial first 50 iterations. When including cut variation, Yu *et al.* (2016) found that the profiles for the average of U_j with r aligned for all values $N = 4$, 5 and 6, if the underlying permutations are irreducible and non-rotational. They concluded that 4 sub-intervals in an IET is enough to achieve good stirring in finite time. Whether this hypothesis can be carried over into IETs with a diffusive step is considered here.

The previous mixing metrics used in Section 4.4.2 to compare across values of r are used to compare across changing N. Again, only irreducible, non-rotational permutations are used. For fixed $r = 1.5$, Figure 4.23 plots the spread of the second leading eigenvalue $|\lambda_2|$, contributing to the long-time exponential mixing rate, for varying permutation groups S_N. In addition to S_4, S_5, and S_6, the
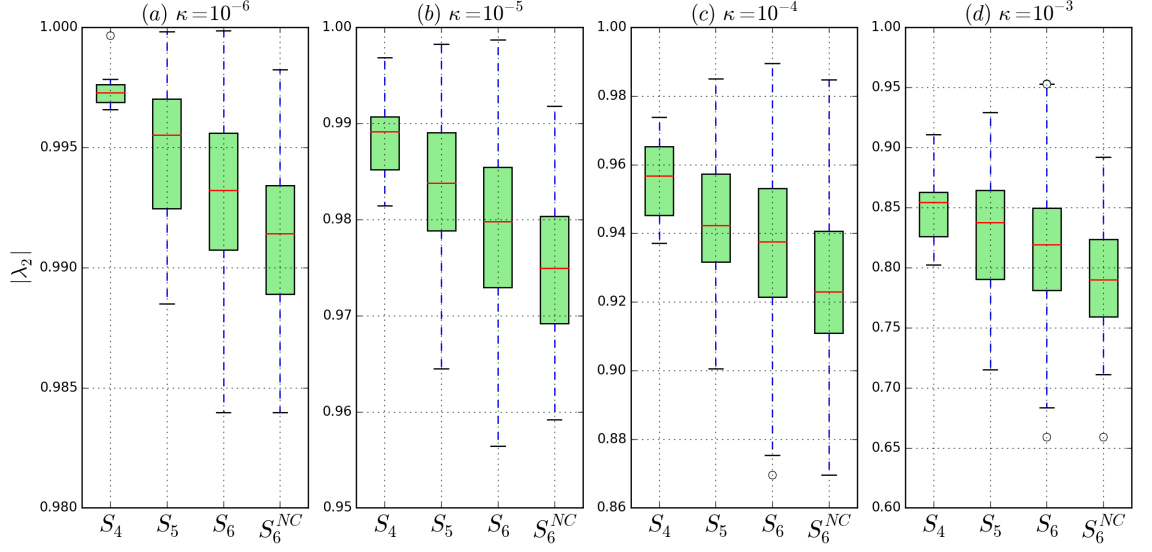
Figure 4.23: Boxplots showing the spread of the modulus of the second leading eigenvalue dependent on the number of sub-intervals in the permutation Π. All non-rotational irreducible permutations from S_4, S_5 and S_6 are considered with $r = 1.5$. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-5}$, c) $\kappa = 10^{-4}$, and d) $\kappa = 10^{-3}$.

group S_6^{NC} is included, denoting the group of non-consecutive, irreducible, non-rotational permutations of S_6, defined such that after permutation, sub-intervals which were consecutive in the initial arrangement [123456], are consecutive following rearrangement. For example, $\Pi = [635214]$ is non-consecutive.

In contrast to changing r, here an obvious trend emerges with the median value of $|\lambda_2|$ being the lowest for S_6 for all values of κ. The variation is more significant as κ is decreased. Note however that when increasing κ, the upper tails on the spread is the largest in value for S_6, showing that some of the slowest mixing IETs belong to $f(1.5, \Pi \in S_6)$. This is likely due to the fact that S_6 has many permutations in which a couple or more of the sub-intervals are consecutive in the permutation rearrangement. Removing these and focusing on only the non-consecutive permutations indeed lowers the upper whiskers and/or reduces the median of $|\lambda_2|$ across all diffusivity coefficients, shown by plotting the spread of $|\lambda_2|$ for S_6^{NC}.

Looking at the spread of values in Figure 4.24 for the average t_{95} across different initial conditions, and Figure 4.25 for the variation in the time to achieve the mixed condition $t_{95}^{max} - t_{95}^{min}$, the trend persists that increasing the number of sub-intervals has a significant effect on increasing the speed of mixing, the
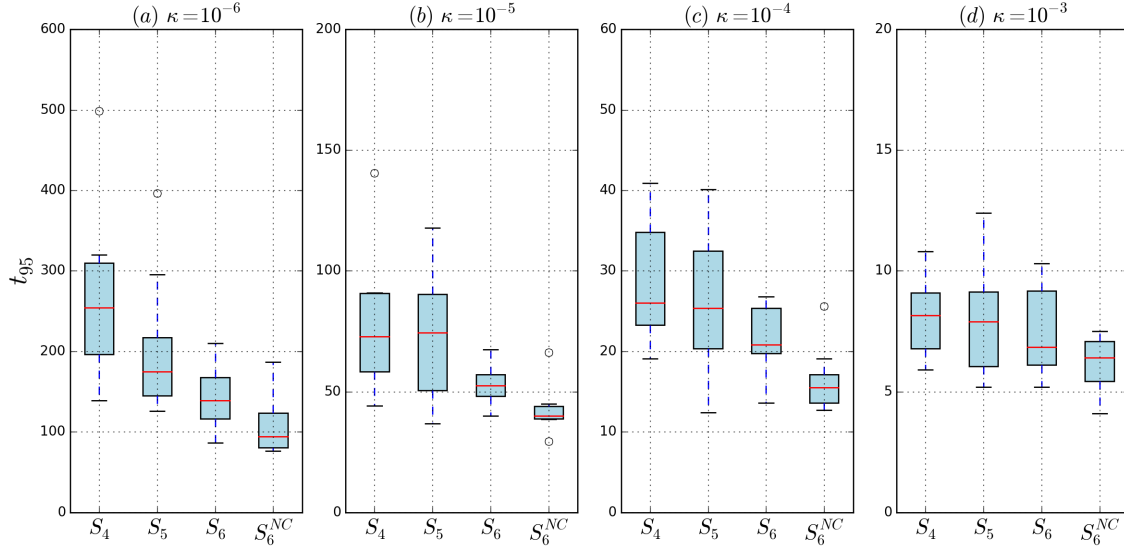
Figure 4.24: Box plots showing t_{95} for non-rotational, irreducible permutations from S_4, S_5, and S_6, alongside the non-consecutive irreducible permutation of S_6, for varying rates of diffusivity coefficient κ and fixed $r = 1.5$. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$, b) $\kappa = 10^{-5}$, c) $\kappa = 10^{-4}$, and d) $\kappa = 10^{-3}$ Eliminating the consecutive permutations from $\Pi \subset S_6$ in the group S_6^{NC} removes most of the outliers.
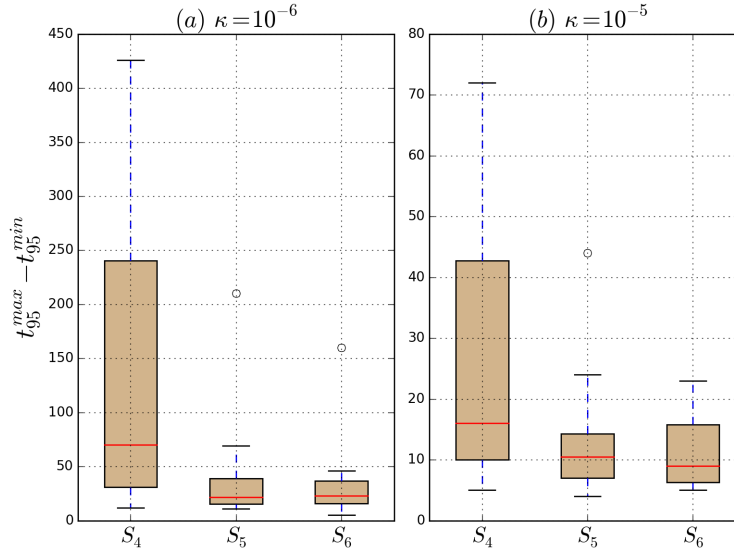
Figure 4.25: Box plots showing $t_{95}^{max} - t_{95}^{min}$ for non-rotational irreducible permutations from S_4, S_5 and S_6 with fixed sub-interval ratio $r = 1.5$. In each of the sub-figures the diffusion coefficient is a) $\kappa = 10^{-6}$ and b) $\kappa = 10^{-5}$.

latter suggesting that good mixing holds across varying initial conditions. When plotting the same figures for $r = 1.2$, the same conclusions can be drawn (not included).

From Table 4.3, the difference in N_f between $N = 4$ and $N = 5$ is a factor of 3, similarly between $N = 5$ and $N = 6$, when $r = 1.5$. Recall from the results of fixed $N = 4$ with different r, there was little effect in the finite and long-time mixing rates from just large increases in N_f. Therefore as a comparative study, although N_f increases with N, a more considerable improvement occurs due to the increase in N, rather than N_f alone. The implication of more cuts at each iteration is that more interfaces are introduced across which diffusion is effective. Additionally more cuts increases the potential to reduce the scale of segregation at each iteration following rearrangement. This implies that it is not the potential to achieve the smallest scales of segregation over a finite time alone which gives better mixing in diffusive IET transformations, but the ability to introduce more new cuts at each iteration, which larger N allows.

4.5 Polynomial dependence on the time to achieve a mixed condition with diffusivity rate

Previous work by Wang & Christov (2018) on IETs with a diffusive step investigated the effect of varying the Péclet number. This is equivalent to the current work by changing the diffusivity coefficient κ. The mixing metric used was the variance $||\cdot||^2$ with $t_{\sim 36\%}$ as a mixed condition, which they call the e-folding time. One could argue t_{36} is not a practical mixed condition since depending on the application, other mixing metrics such as $||\cdot||^\infty$ can still be significantly large in value for this condition on the variance, and they do not state any physical reasoning for applying this condition except it allows an analytical approximation for a mixing time. Earlier drafts of their work used t_{50}, inspired by problems from shuffling cards [Aldous & Diaconis (1986),Diaconis (1996)]. Famously, card shuffling problems ask how many shuffles are required to achieve a desired variation distance from a perfectly randomised permutation of cards. In card shuffling a variation distance around 50% makes sense practically because for this condition, the pack is essentially random by the definition of variation distance, and hence well mixed. For the variance of a concentration field, this is not necessarily the case. From numerical results they did not find any definitive scaling relation between varying the Péclet number and the time to achieve a mixed condition.
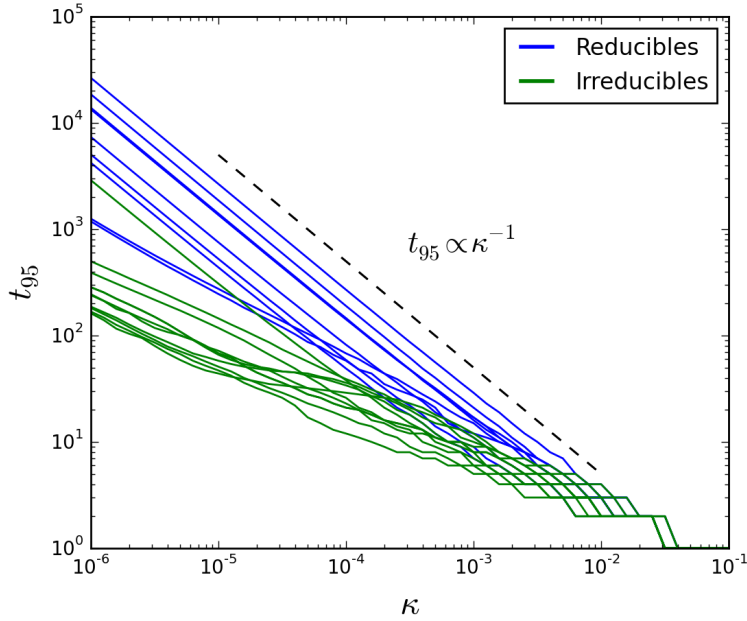
Figure 4.26: t_{95} with κ from numerical simulations for all non-rotational $\Pi \in S_4$, $r = 1.5$. Initial condition $c^{(0)}(x) = c_{sq}$.

In the previous section, the time to achieve a mixed condition t_{95} was used as a diagnostic to investigate the effect of IET parameters on rates of mixing. Only a handful of diffusion coefficients were considered. This section addresses in detail how the time to achieve a mixed condition varies with the rate of diffusivity. A scaling $t_{\%}$ with κ is found and related to the eigenvalues of the transfer operator. A mechanism is introduced which predicts a similar scaling rate.

4.5.1 Numerical results for t_{95}

Figure 4.26 plots t_{95}, the iteration at which the variance achieves a 95% mixed state, for all permutations $\Pi \in S_4$ when $r = 1.5$. The profiles are coloured according to the rearrangement properties of the permutation, whether reducible (blue), or irreducible and non-rotational (green, denoted only as irreducible in the legend). The rotation permutations are omitted. For each the initial condition was $c^{(0)}(x) = c_{sq}$ and 51 values of κ tested, uniformly distributed in log-scale for $\kappa \in [10^{-6}, 10^{-1}]$. For large κ the profiles appear to plateau, with stepped profiles, which is caused by the integer values in t_{95} from discrete time iteration in j.

A log-log scaling on the axes reveals a roughly polynomial relation for all

permutations of

$$t_{95} = C\kappa^{-\eta}, \tag{4.29}$$

where η is dependent on the underlying permutation used in the IET. For most of the reducible permutations $\eta = 1.0$, while for nearly all of the irreducible permutations η is smaller in value. The line $t_{95} \propto \kappa^{-1}$ is included as a comparison.

Using (4.29) for the profiles of t_{95}, the curve-fit function from the Scipy package in Python is used to find approximations for the exponent η. Table 4.4 lists approximations of η from the numerically computed profiles for the irreducible permutations alone. For each IET $f(1.5, \Pi)$, the profile of t_{95} with κ was computed for 5 rotated initial conditions of the square wave, η computed for each, and then the mean value taken. For nearly all initial conditions the values of η had small deviations from those listed in the table, varying by roughly 0.01 absolute error. $\Pi = [2413]$ is omitted due to the profile being highly variable and not accurately described by (4.29). Note also that $\Pi = [4231]$ has $\eta = 1.0$, the permutation which was discussed in Section 4.4.1 for having bad stirring properties. The latter result is due to the fact that for this permutation the IET has a periodic order of just 12 iterations and is not good at reducing the scale of segregation. Similarly, reducible permutations which seem to have an exponent $\eta < 1$ are those in which the first interval I_1 remains fixed, resulting in only a small static island. This is quite surprising since it would be expected that when $\sim 12.3\%$ of the domain is static, diffusion-limited mixing would be expected. This may suggest that the large values of κ, scales in the concentration field under repeated cutting, shuffling and diffusion are larger than the static interval I_1 for these reducible cases. However as $\kappa \to 0$, for both of these permutations the profile of t_{95} starts to curve upwards, and diffusion limited mixing predicting $t_{95} \propto \kappa^{-1}$ dominates the mixing dynamics.

All values of η for the irreducible permutations fall within the interval $\eta \in [0.5, 0.62]$. Re-plotting this figure for other mixed conditions, such as t_{99}, t_{90} and

Π	[3142]	[4132]	[2413]	[4231]	[4213]	[4312]	[2431]	[3241]	[3241]	[3421]
η	0.5	0.54	-	1.0	0.56	0.57	0.59	0.54	0.56	0.62

Table 4.4: Approximation for η from curve fitting $t_{95} = C\kappa^{-\eta}$ to numerically computed profiles for irreducible, non-rotational permutations $\Pi \in S_4$. Profiles were computed for 5 different initial conditions of the square wave rotated by α and the exponent of each averaged.
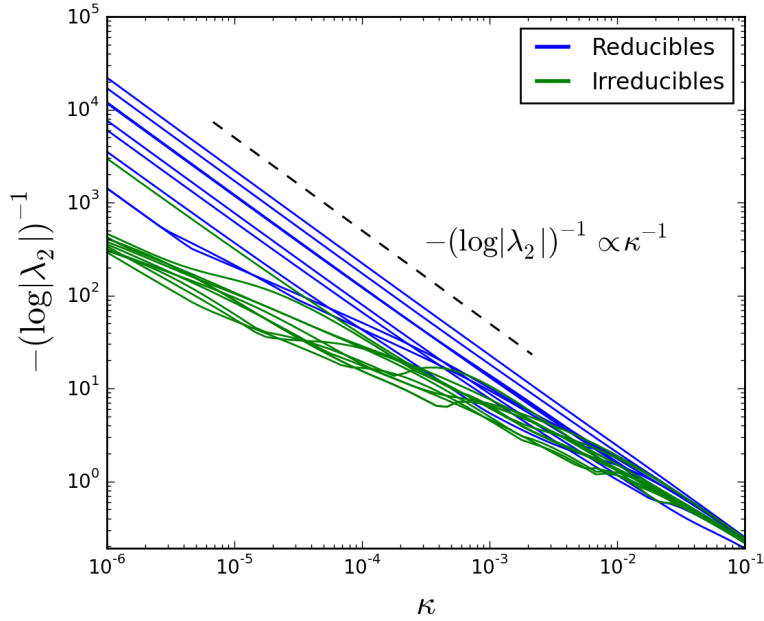
Figure 4.27: The variation in the second leading eigenvalue $|\lambda_2|$ with diffusivity coefficient κ is plotted for all non-rotational $\Pi \in S_4$, $r = 1.5$. The values of $|\lambda_2|$ where computed from the respective transfer matrices dM.

t_{80} retains the polynomial scaling (4.29), where the only main difference is the changing value of C. Also for other values of r, such as $r = 1.1$ and $r = 1.2$, similar conclusions can be made.

Since κ can vary over orders of magnitude in certain mixing problems, the polynomial scaling (4.29) predicts orders of magnitude increases in the time to achieve a mixed condition with decreasing diffusivity rates. This contrasts chaotic stirring systems where $t_\% \sim \log(\kappa)$, which is weakly related to orders of magnitude change in κ.

4.5.2 Comparison to leading eigenvalue

In the last section numerical results for a range of $t_\%$ revealed a polynomial scaling with the diffusion coefficient κ. An attempt to generalise the scaling observed can be done following the ideas of Ashwin *et al.* (2002), who stated that the time to achieved a mixed condition is bounded by the mixing rate of the leading eigenvalue λ_2. This is reminiscent of Lemma 2.3.1 for diffusion-less maps M which are strong mixing, where the decay of correlations is given by $C_{\phi,\psi}(j) \leq C\tau^j$, for

a class of suitable, real-valued functions ϕ, ψ and some constant $C < \infty$, where τ is the leading eigenvalue of the spectrum $\sigma(P)$.

Consider that for an advective-diffusive transfer operator P the asymptotic mixing rate is given by the second leading eigenvalue $|\lambda_2|$ of P, then there exists a constant $C < \infty$ such that

$$|c(x,j) - \bar{c}| < |c(x,0) - \bar{c}|C|\lambda_2|^j \tag{4.30}$$

where $c(x,j) = P^j c(x,0)$ and assume that for a given initial concentration field, C is constant over mixing protocols. Then, a time $t_\%$ can be found in terms of C and the eigenvalue λ_2 in which the concentration field achieves a % mixed state;

$$|c(x,j) - \bar{c}| < |c(x,0) - \bar{c}|C|\lambda_2|^{t_\%} = R, \tag{4.31}$$

$$|\lambda_2|^{t_\%} = \left(R^{-1}C|c(x,0) - \bar{c}|\right)^{-1}, \tag{4.32}$$

$$\implies t_\% = \frac{\log(R^{-1}C|c(x,0) - \bar{c}|)}{-\log|\lambda_2|}, \tag{4.33}$$

where R represents the ratio of percent mixed, for example 0.05 for t_{95}. Hence for a given percentage %, $t_\%$ will approximately decrease by a factor given by the ratio of the logarithm of the leading eigenvalues for the mixing protocol. Since all $|\lambda_2| < 1$, a smaller $|\lambda_2|$ decreases $t_\%$ faster as expected.

The relation (4.33) leads to $t_\% \propto -(\log|\lambda_2|)^{-1}$, so if the leading eigenvalue strongly effects the time to achieve a mixed condition, then it is expected that

$$-(\log|\lambda_2|)^{-1} \propto \kappa^{-\eta}. \tag{4.34}$$

Plotting $-(\log|\lambda_2|)^{-1}$ in Figure 4.27 for the same parameters as in Figure 4.26, indeed both scaling relations for the reducible permutations, $\eta = 1$, and for the non-rotational irreducible permutations a slope of $\eta < 1$ are retained. The figures are qualitatively similar except for the eigenvalue profiles when $\kappa < 10^{-5}$, where a slight up turn in $-(\log|\lambda_2|)^{-1}$ occurs.

Figure 4.28 (a) plots the profiles in more detail for the irreducible permutations, minus $\Pi = [4231]$. Different colours in the profiles show complicated variations as κ decreases. It is surprising that the scaling relation is visible in both t_{95} and $-(\log|\lambda_2|)^{-1}$ since it was shown in Section 4.3.2 that the finite time mixing is not dependent on the long-time mixing rate when κ is small. Figure 4.28 (b) plots both profiles of t_{95} and $-(\log|\lambda_2|)^{-1}$ for an example permutation, alongside how $-(\log|\lambda_k|)^{-1}$ varies for a handful of k representing the largest eigenvalues in the spectrum of $P_D \circ P_M$ (grey). This permutation is chosen as it
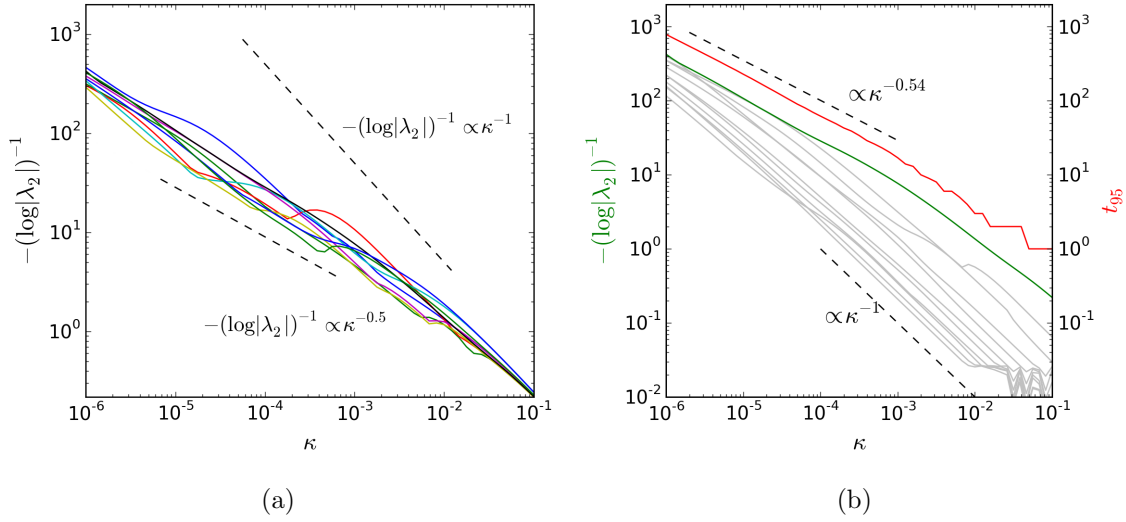
Figure 4.28: (a) The variation in the leading eigenvalues, plotted as $-(\log|\lambda_2|)^{-1}$ is plotted again diffusivity, for only the irreducible, non-rotational permutations in S_4, with $r = 1.5$. Dashed line represents the expected polynomial relation from diffusion alone. In (b) for a single permutation $\Pi = [3241]$ profiles of $-\log(|\lambda_k|)$ for many k are plotted in grey, alongside the $|\lambda_2|$ profile and the time to achieve a 95% mixed state t_{95}. The latter profiles are well approximated by $\propto \kappa^{-0.54}$.

is one of the smoothest profiles. For large values of κ all eigenvalue profiles appear to change $\propto \kappa^{-1}$. As κ decreases below 10^{-3}, the relation $-(\log|\lambda_2|)^{-1} \propto \kappa^{-0.54}$ holds. The other eigenvalues continue to have $-(\log|\lambda_k|)^{-1} \propto \kappa^{-1}$ until they are bounded by the leading eigenvalue. As κ decreases, the eigenvalues become close-to-degenerate and predict the finite time mixing, all of which go like $-(\log|\lambda_k|)^{-1} \propto \kappa^{-0.54}$.

Since the finite time mixing was shown to be well approximated by close-to-degenerate eigenvalues, this similarity in the variation of $|\lambda_k|$ with $|\lambda_2|$ with small diffusivity could explain why the scaling is observed in both long-time mixing rates $|\lambda_2|$ and finite time mixing rates $t_\%$. Of course this is only a single example and the other profiles do not provide good comparisons to make a conclusive hypothesis when κ is small. In the next section a potential mechanism for the mixing rates is explored and predicts the scaling (4.29) with $\eta < 1$, more precisely $\eta \sim 0.5$.

4.5.3 Mechanism for polynomial scaling relations

The mechanism for the $\eta = 1$ scaling when the underlying IET has a reducible permutation is easy to understand. A reducible permutation will leave a large continuous colour segment unmixed from having a single, or several, static subintervals. Mixing between this segment and the domain outside is only possible due to diffusion across the boundaries, since no scalar concentration can be stirred into the segment from the transformation. Even if some stirring via cutting and shuffling occurs within the segment, eventually the final approach to a uniform mean field will be via diffusion. Therefore the long-time mixing will be dominated by the slow decay of the eigenfunction capturing the diffusion-limited mixing across the unstirred segment. For fixed scales in a concentration field, the effect of varying the diffusivity coefficient is proportional to $\exp(-\kappa)$. Therefore eigenvalues will change with diffusivity value as $\lambda \propto \exp(-\kappa)$ and (4.33) predicts

$$t_{\%} = \frac{K}{\kappa}, \tag{4.35}$$

with $\eta = 1$. This was similarly discussed in Chapter 3.2.3 for the shear map with a strong wave perturbation and resulting island, where it was shown that the time to achieve a mixed condition fin variance decay was $\propto \kappa^{-1}$. Now to address the scaling when the IET has an irreducible, non-rotational permutation, which satisfy good stirring conditions.

In chaotic systems an eigenfunction emerges in the concentration field due to a balance between the repeated strain from stretching, and leading to the eradication of large gradients in the concentration. This occurs when the smallest length scales in the scalar concentration field reaches the Batchelor length scale

$$w_B = \sqrt{\kappa/\lambda} \tag{4.36}$$

where the effects of compression and diffusion balance. In purely cutting and shuffling transformations there is no local stretching to result in a balance between compression and diffusion. Whether a scale such as the Batchelor length scale exists, presenting a minimum scale where there is a balance between cutting and shuffling and diffusion, is not known.

Wang & Christov (2018) hypothesised a method for predicting a stopping time T in cutting and shuffling systems, similar to card shuffling systems. They attempted to find a time T_{Pe} in which the concentration field reaches a scale where diffusion is effective, and fast depletion results from then on with no further

cutting and shuffling required. For this they employed the analysis of Schlick *et al.* (2013) for the rate of depletion of two sub-segments of unequal colour via diffusion. It is worth recalling the discussion of Schlick *et al.* (2013) and Wang & Christov (2018) as it is beneficial in the suggestion of a mechanism predicting a similar scaling to (4.29) observed for the irreducible, non-rotational permutations in Figure 4.26. Some mistakes were made in their arguments and these will be addressed throughout.

Schlick *et al.* (2013) asked us to consider two subsegments of opposing colour. Let them be represented by values -1 and 1 respectively, both of equal length l and periodic boundary conditions. Then the concentration profile under the action of diffusion with diffusivity coefficient κ is

$$c(x,t) = \sum_{k=0}^{\infty} \frac{2}{\pi(2k+1)} \sin\left(\frac{n\pi x}{l}\right) \exp\left(-(2k+1)^2 \pi^2 \frac{t\kappa}{l^2}\right). \qquad (4.37)$$

For the dominant mode to be *washed out* by a significant factor, say about $e^{-2} \approx 13\%$, this gives a diffusional length scale l^*, such that

$$l^* = l^*(T) \approx \pi\sqrt{\frac{T\kappa}{2}}, \qquad (4.38)$$

where T is the time period over which diffusion is applied for the given diffusivity coefficient κ.

Wang & Christov (2018) then ask: when will the *average* continuous color subsegment length, denoted l_m, in the cutting and shuffling process *without* diffusion reach l^*? They state the importance of this question, since if $l_m \approx l^*$ then T iterations of cutting and shuffling *with* diffusion later, the concentration of the striation will be damped out by $e^{-2} = 13\%$, so T is the e-folding time; their specified mixed condition.

It is in this statement that the confusion lies. The statement defines T as though it is the iteration number, implying that l^* increases with increasing iterations of the IET. However, in the analysis of Schlick *et al.* (2013), T is the time taken during the diffusive intervals in the iterative process, where significant depletion begins to occur at every iteration, once the concentration field reaches this scale. If the question and the accompanying statement were to be linked, l^* should be a constant for T constant, instead of l^* varying with time which they have in their study. Their argument then continues with finding numerical values for l_m and rescaling variance profiles depending on the iteration number in which $l^* = l_m$, which they denote $\tilde{T}_{Pe}$. They find no discernible pattern in $\tilde{T}_{Pe}$ with

Peclet number and state that numerical evidence suggests cut-off phenomena is observed, similar to card shuffling, however this remaining fact is also arguable due the problems of definitions already discussed.

Other problems lie in this suggested mechanism of mixing in comparison to the model presented here. Numerical results of this study suggest extremely slow, long time mixing rates attributed to eigenfunctions with eigenvalues ~ 1, extremely slow compared to a fast depletion time of 13% at each iteration. Additionally this approach completely neglects the effect of diffusion in the early stages of cutting and shuffling. However, the idea of comparing l_m and l^* can be re-imagined in terms of defining a balance between cutting, shuffling and diffusion, and how this balance would vary with diffusion.

The results presented in this current work show that an eigenfunction regime emerges in IET systems with diffusion, dominating both the finite and long-time mixing dynamics. An eigenfunction regime implies there is a balance between cutting and shuffling and diffusion since, on average at every iteration, the concentration field is decaying with a fixed factor given by the respective eigenvalue of the eigenfunction. Therefore at each iteration a decay rate ϕ of an eigenfunction can be defined as

$$\phi = |\lambda_2|. \tag{4.39}$$

Comparing this to the analysis idea of Schlick *et al.* (2013), the mixing rate is also related to the average scale of the concentration field l, that is

$$\phi = \exp(-4\pi^2\kappa/l^2). \tag{4.40}$$

Now if the decay rate was taken over a number of iterations with a fixed rate with unchanging scale l, ϕ could be written as a function of j,

$$\phi(j) = \exp(-4\pi^2\kappa j/l^2). \tag{4.41}$$

However, in cutting and shuffling systems the largest scale of the concentration field l is not fixed and is instead a function of iteration number $l(j)$. Then, interest lies in working out how the decay rate $\phi(j)$ of the concentration field achieves on average a constant decay rate over j iterations, say $\phi(j) = e^{-C}$, and how many more iterations it takes if the diffusion coefficient κ was decreased. This would imply that

$$\phi(j) = \exp(-4\pi^2\kappa/l(j)^2) \implies l^2(j) = \frac{4\pi^2\kappa}{C}. \tag{4.42}$$

All that is left now is to approximate the average variation of $l(j)$ over j iterations. However, following the assumption that a balance occurs between cutting,

shuffling and diffusion an approximation from l is not simple to conceive. One simple assumption is between j iterations of an decay rate by ϕ, the mid-way point of changing $l(j)$ captures an average length scale l^*. This is the assumption carried forward, however other claims can be made on what the lengthscale l^* should be capturing the balance between cutting, shuffling and diffusion. For all cutting and shuffling IETs, at most new interface can be generated linearly, so that if N new cuts are made at each iteration, then a minimum function of the value of $l(j)$ is

$$l(j) = \frac{1}{jN}. \tag{4.43}$$

Then the length scale at the mid-way iterations $j/2$ would be

$$l^* = \frac{2}{jN}. \tag{4.44}$$

Substituting this into (4.42) gives

$$\left(\frac{2}{jN}\right)^2 = \frac{4\pi^2 \kappa}{C}, \implies j^2 = \frac{C}{2\pi^2 \kappa N^2}. \tag{4.45}$$

Therefore if the diffusivity coefficient is varied, the number of iterations needed to ensure the same decay factor $\phi = \exp(-C)$ is achieved goes like

$$j \propto \kappa^{-1/2}. \tag{4.46}$$

This mechanism of attempting to derive a simplistic representation of a balance between cutting and shuffling has derived a polynomial scaling between the time to achieve a fixed decay rate and the diffusivity coefficient κ, which is remarkably similar to the numerical results presented in Sections 4.5.1 and 4.5.2. Figure 4.29 re-plots Figure 4.26 with (4.46) relation included for comparison, showing good agreement. The scaling (4.45) also captures that increasing N, the number of sub-intervals in the IET, would see an improvement on the time to achieve a mixed condition by reducing the number of iterations needed.

Of course, this argument ignores the diffusion across sharp interfaces in the concentration field, however it is expected in the eigenfunction stages that these have been balanced and are not as effective across the eigenfunction as they are in the initial stages of mixing. Additionally, choosing l^* to be (4.44) may not be appropriate and other functions could be derived and investigated in future work. The numerical results suggest $\eta = 0.5$ as a potential lower bound on how t_{95} varies with diffusivity, with values of η greater than this in general, listed in Table 4.4. More development is needed in the intuition that mixing rates arise from a balance of cutting, shuffling and diffusion.
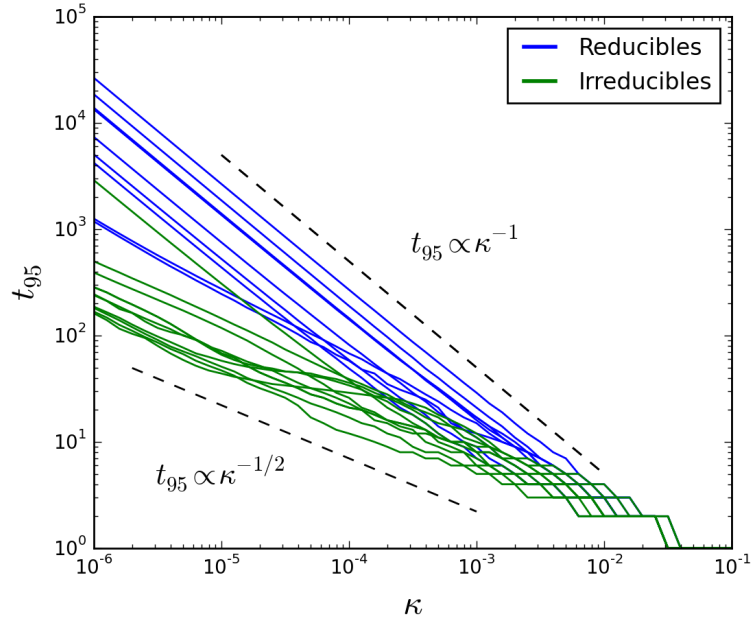
Figure 4.29: t_{95} with κ from numerical simulations for all non-rotational $\Pi \in S_4$, $r = 1.5$. The line $t_{95} \propto \kappa^{-1/2}$ is included for comparison from (4.46). The initial condition for all profiles was $c^{(0)}(x) = c_{sq}$.

4.6 Conclusions

The work presented in this chapter extended the studies of Ashwin *et al.* (2002), Krotter *et al.* (2012), and Yu *et al.* (2016) to study rates of mixing in one-dimensional IETs with a diffusive step. The construction of the transfer operator as a transfer matrix describing the advective step was possible due to the rational ratio r taken between the sub-interval lengths. This allowed for a computational study with results independent of the resolution of wavemodes if Q was taken large enough, which can not be said for a previous study on diffusive IETs by Wang & Christov (2018). Derivation of a transfer matrix suggests a long-time exponential mixing rate, as is observed, and as $\kappa \to 0$, $|\lambda_k| \to 1$ for all k. This is due to the fact that the underlying advective transformation is not strong mixing by definition.

For a considerable number of iterations when the diffusion coefficient is small, a deviation from this final exponential mixing rate is observed for the IETs presented here, and the profiles of variance decay follow a stretched exponential first suggested as an appropriate fit by Wang & Christov (2018). Close-to-degenerate eigenvalues and corresponding eigenfunctions are presented as a mechanism to

explain the mixing rate in this stage where a summation of multiple exponential decay rates also approximated the variance profile well. The application of this mechanism was inspired by Popovych *et al.* (2007) who showed that close-to-degenerate eigenvalues contributed to the mixing for weakly connected chaotic domains or flows with islands, both of which are not strong mixing in the diffusion-less limit. However, where they showed the variance profile could predict the two leading eigenvalues, here more than two eigenvalues were needed for an accurate approximation of the decay of variance profile. This is likely due to the fact that many close-to-degenerate eigenvalues emerge in the IETs systems with diffusion as $\kappa \to 0$. These initial stages where the decay of concentration is dominated by the superposition of many slowly decaying eigenfunctions is sensitive to the initial condition, shown herein by the time to achieve a mixed condition varying over hundreds of iterations for initial conditions with fixed scales, but are simply rotated in the domain. This implies that predicting when the transition to the asymptotic regime, governed by the slowest decaying eigenfunction, is difficult for knowledge of the initial condition and flow parameters alone.

An illustrative example of two IETs suggested that the decay of a concentration field to its mean value occurs faster for an IET which satisfies the good stirring conditions of Krotter *et al.* (2012) and Yu *et al.* (2016), than for an IET which does not meet all the criteria. A parameter study was carried out to find which of the criteria were essential, and which were less important, for an IET with diffusion. The choice of Π as an irreducible, non-rotational permutation was found to be essential for fast mixing.

However it was found that the choice or r leading to a larger N_f was not as important as having more sub-intervals in the permutation rearrangement. In Krotter *et al.* (2012) and Yu *et al.* (2016) an IET with larger N_f suggested better decrease in the scale of segregation due to the fact that this decreases the smallest scale that could be achieved from cutting and shuffling ($1/N_f$). Although increasing N for a fixed r also increases N_f, it was observed that faster long-time and short-time mixing rates are on average better with increasing N, compared to changing r, even if the latter caused a much more significant increase of N_f. This suggests that in the presence of diffusion it is not the reduction to smaller scales over finite time in the initial stages which will achieve a better long-time mixing rate, but rather introducing new cuts and decreasing segregation at each iteration which will mix the concentration field faster. In the diffusion-less limit, a choice of r resulting in a larger N_f would become important, since mixing would

be restricted by only diffusion across the edges of the equal sized cells of size $1/N_f$.

Studies on cutting and shuffling in higher dimensional systems, such as those on the hemi-spherical shell [Park *et al.* (2016, 2017); Smith *et al.* (2017a)], and new studies in two-dimensional planar systems [Smith *et al.* (2019)] have been concerned with fine tuning rotation parameters within the piecewise isometries, to achieve the best coverage of cutting lines and as such small scales of segregation across finite times. In all of the examples the number of pieces within the piecewise isometry remains fixed, implying the total length of cuts introduced at each iteration is also fixed. The results of the current work herein, suggest that such parameter tuning may not be of high importance when diffusion is present. Similarly, numerical results presented here on the finite-time mixing rates reveal that roughly 100 iterations is required to achieve a 95% mixed condition when $\kappa = 10^{-6}$, so IETs do not present themselves as efficient options as practical mixing protocols when the diffusivity coefficient is small.

This is because the relationship between the time to achieve a mixed condition, $t_\%$ and diffusivity coefficient κ, has a polynomial relation such that $t_\% \propto \kappa^{-\eta}$. For IETs which satisfy the good stirring conditions, $\eta \sim 0.5$, which is also derived from an argument considering a balance between cutting, shuffling and diffusion. This exponent is smaller in value than for diffusive restricted mixing, which gives $\eta = 1.0$, as seen for nearly all reducible IETs with unstirred sub-segments. For rational r, a likely assumption is that as $\kappa \to 0$ for all IETs, $\eta \to 1$ since eventually all mixing would be restricted across the boundaries of the cells of equal size. This is seen briefly in the eigenvalue profiles for $r = 1.5$, $\Pi \in S_4$ in Figure 4.27, however this is as an open question for future work.

A recent study was interested in investigating decreasing scales of segregation via IETs with variable parameters. Smith *et al.* (2018) devised an optimization procedure to find IETs with variable cut locations to achieve the best stirred condition. The goal was to achieve a final concentration field which would have a scale of segregation which, as well as small in scale, is evenly distributed within the domain. However studies such as Miles & Doering (2018) find that homogenization from diffusion, and filamentiation via advection, are sometimes in conflict and produce a negative impact on mixing in optimized flows. Therefore it is likely that the optimized cutting method of Smith *et al.* (2018) is not directly transferable to diffusive IETs. They also address the limitation that knowing the initial condition is required for their optimization scheme to work. As evidenced

herein, in IET transformations better segregation does not necessarily suggest fast mixing with diffusion, and varying the initial condition can have a drastic effect on the time to achieve a mixed condition. One may be able to construct an optimization procedure for an IET with varying cut location when there is diffusion, where at each iteration the best rearrangement is chosen for an assigned number of cuts. However whether this is of practical interest or computationally feasible is debatable.

Chapter 5

Mixing by piecewise, uniformly expanding maps

The last chapter investigated the rates of mixing in cutting and shuffling transformations with diffusion, however, it is unlikely that these mechanisms will occur in isolation. Generally stretching and folding is the key stirring mechanism. This chapter investigates the effect discontinuous transformations have on mixing rates in chaotic advective systems.

Hyperbolic maps with discontinuities have been at the heart of a number of studies investigating mixing. It has been observed that strong mixing can be preserved in expanding maps composed with permutations [Byott *et al.* (2013)], although it was reported that typically the mixing rate is slowed down. In special cases a discontinuous slip deformation can lead to exotic pseudo-elliptic points where mixing can be hindered completely [Smith *et al.* (2016)]. However, there has been no investigation into the mixing rates for the composition of the three mixing dynamics in the literature; stretching and folding, cutting and shuffling and diffusion.

This is addressed in the current work where a simple map on the unit interval is derived representing an idealised time-periodic flow with the three mixing dynamics. Comparison to previously published results by Byott *et al.* (2013) allows for computation of strong mixing rates in the diffusion-less limit. Numerical results on the mixing rates when there is diffusion are presented and the implication of results on finite time mixing are discussed.

It is observed that the addition of cutting and shuffling behaves similarly to non-uniformity in stretching rates, discussed in Section 2.5.1, causing dispersion of the concentration field between wavemodes in discrete Fourier space.
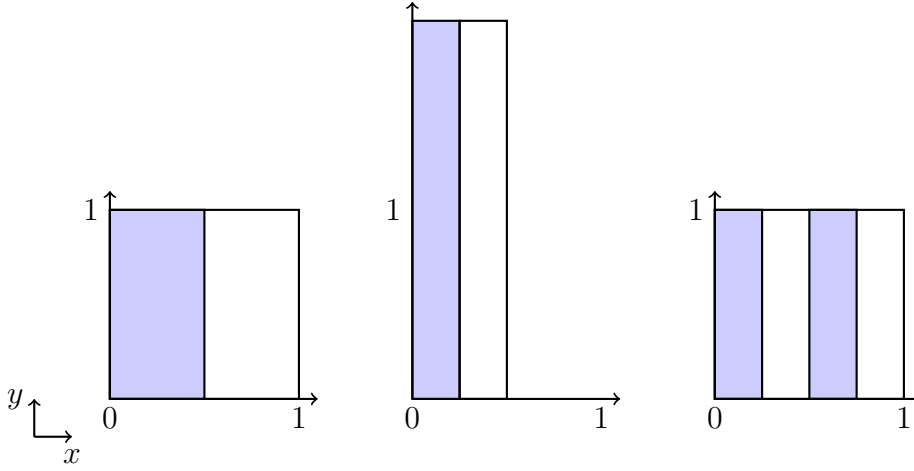
Figure 5.1: The action of the baker's transformation (5.1) is to contract the x direction, stretch the y direction and reassemble onto the unit torus. Taking a y independent initial condition, here $c(x, y) = blue$ for $x < 1/2$ and $c(x, y) = white$ for $x \geq 1/2$, reduces the baker's transformation to a one-dimensional map on the unit interval.

This results in eigenfunctions with long-time exponential decay rates. The decay rates of the eigenfunctions are shown to be dependent on the choice of permutation, although counter-intuitively for many of the permutations these mixing rates change non-monotonically with the diffusion coefficient. The strong mixing rates of the underlying maps, the limit of vanishing diffusivity, are found using Markov partitions and bounds on the rates presented. The non-monotonicity of the mixing rates results in the bounds no longer being applicable when diffusion is included.

5.1 Formulation of the problem

5.1.1 Baker's transformation with permutations

The two-dimensional incompressible baker's transformation is a paradigm example of a hyperbolic chaotic map which is mixing. In two-dimensions the map is defined as follows

Definition 5.1.1. (The baker's transformation) *Let Ω be the unit square such that $\Omega = [0, 1) \times [0, 1)$. We define the transformation $T : \Omega \to \Omega$ by*

$$M_B(x, y) = \begin{cases} (x/2, 2y) & y \in [0, 1/2) \\ ((x+1)/2, 2y - 1) & y \in [1/2, 1) \end{cases} \qquad (5.1)$$

Figure 5.1 depicts the action of the baker's transformation, which shows that there is a compression by $1/2$ in the x direction whilst a stretch of 2 in the y direction, followed by horizontally dividing the rectangle for reassembly onto the unit square.

Considering a y independent initial condition, the baker's transformation reduces to a one-dimensional map which is one-to-two and represented acting on a concentration field by

$$c^{(j+1)}(x) = \begin{cases} c^{(j)}(2x) & x \in [0, 1/2), \\ c^{(j)}(2(x - 1/2)) & x \in [1/2, 1). \end{cases} \tag{5.2}$$

The y independent initial condition remains y independent in Figure 5.1. The baker's transformation is a paradigm example of a hyperbolic, chaotic transformation which is strong mixing. The inverse of the one-dimensional baker's transformation, is captured in the following transformation;

Definition 5.1.2. (m-**adic transformation**) *Consider the transformation T : $[0, 1) \to [0, 1)$ given by*

$$f(x) = mx \qquad \mod 1, \tag{5.3}$$

where m is an integer. For $m = 2$ it is customarily called the doubling map.

The one-dimensional approximation of the baker's transformation is the preimage of the doubling map, (5.3) with $m = 2$. The doubling map is a one-dimensional map which is also chaotic, ergodic, mixing and exact. The one-dimensional baker's transformation (5.2) will form the stretching and folding component of the dynamics.

For the discontinuous transformation, permutations of equal sized cells (4.1) are used acting in the x direction. Interval exchange transformations with varying interval lengths are withheld until later in the discussion. For the sake of conciseness, only the permutation groups S_3, S_4 and S_5 are investigated. S_3 has 6 permutations, 3 rotations and 3 interleaving permutations, while S_4 and S_5 have 24 and 120 permutations respectively. This gives a total number of permutations which is feasible to compute, while enough to make a varied and comparative study. The composition maps will be denoted $\Pi \circ M_B$.

Although this model is one-dimensional, and thus highly idealised, the action of stretching and folding creates long lamellar structures which align with the unstable manifolds. In high dimensional systems, the action of cutting and shuffling will introduce discontinuities, but it is likely that these will propagate with the

flow and also align with the unstable manifolds from the stretching and folding dynamics. Thus this one-dimensional model can be considered as a cross-section approximation of the dynamics.

The one-dimensional baker's transformation has successfully been used many times as an initial investigation into the mixing rates of chaotic maps with non-uniform stretching rates [Fereday *et al.* (2002),Wonhas & Vassilicos (2002)] and the effect of no-slip boundary conditions, sufficiently capturing experimental behaviour [Gouillart *et al.* (2007),Gouillart *et al.* (2008)]. These studies were reviewed in more detail in Section 2.5.1, alongside other relevant literature on mixing arising from non-uniform stretching.
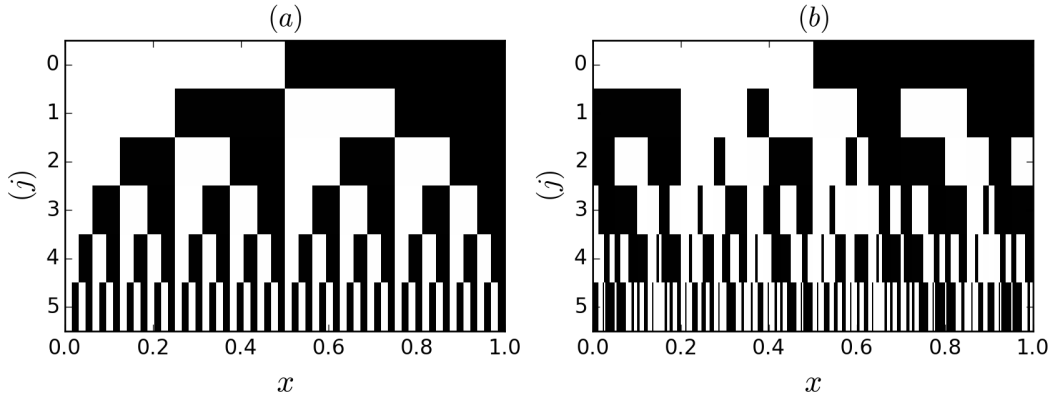
Figure 5.2: Space-time plots showing the evolution of striations of a half-white, half-black initial condition. In (a) the action of the baker's transformation (5.1) decreases striation width exponentially, while in (b) the transformation $\Pi \circ M_B$ with $\Pi = [35421]$ appears to retain the exponential creation of striations but introduces non-uniformity in the striation widths.

5.2 Dynamics in the absence of diffusion

5.2.1 Space-time plots of segregation

Observing how the concentration field evolves in time gives a qualitative explanation of how including permutations can perturb the mixing of the baker's transformation. For this the use of space-time plots is employed, similar to Chapter 4 for IETs. The simplest initial condition *half-white, half-black*, and slight variations thereof, are enough to present the range of behaviour. Example permutations from S_3, S_4 and S_5 are used without a full parameter study presented.
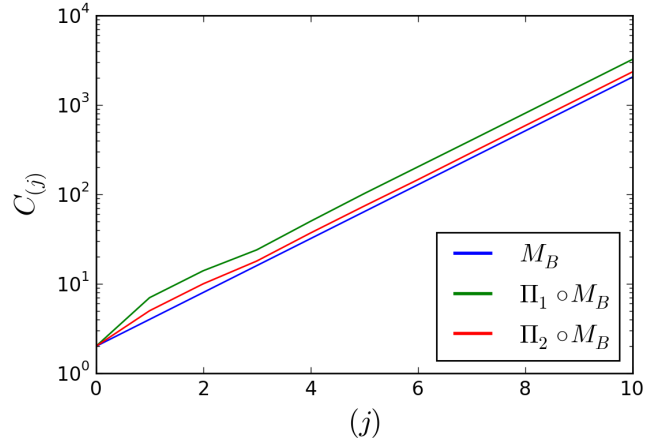
Figure 5.3: Increase in the number of interfaces $C_{(j)}$ is plotted against iteration number, showing that exponential growth in striations is preserved for the mappings M_B, $\Pi_1 \circ M_B$ and $\Pi_2 \circ M_B$ where $\Pi_1 = [35421]$ and $\Pi_2 = [213]$. The permutations also increase the total amount of interface. The initial condition was *half-white, half-black*.

Figure 5.2 (*a*) shows how a *half-white, half black* initial condition evolves under the action of the one-dimensional baker's transformation. At each iteration the white and black sub-segments are halved in size and total number doubled within the domain, giving an exponential decay of correlations. The equal distribution of white and black in the initial condition is preserved uniformly in the striation widths at each iteration.

In Figure 5.2 (*b*), the evolution of the same initial condition is shown under the transformation $\Pi_1 \circ M_B$ with $\Pi_1 = [35421]$, a random permutation chosen from S_5. After the first iteration, parts of the concentration field are swapped by the permutation and uniformity in the striation widths is no longer retained. This swapped striation arrangement in then halved in size and doubled in number, before being rearranged by the cutting and shuffling once more. After several iterations, a highly non-uniform concentration field of many striation widths is observed. It appears as though the exponential increase in striations is preserved by the composition $\Pi_1 \circ M_B$. This is confirmed in Figure 5.3 where exponential growth in the striations is inferred from plotting the number of interfaces between sub-segments of the concentration field of varying colour $C_{(j)}$. The growth in interfaces is also plotted for the baker's transformation M_B alone, and the composition map $\Pi_2 \circ M_B$ with the interleaving permutation $\Pi_2 = [213]$. The results show that alongside retaining exponential growth in the number of interfaces, the
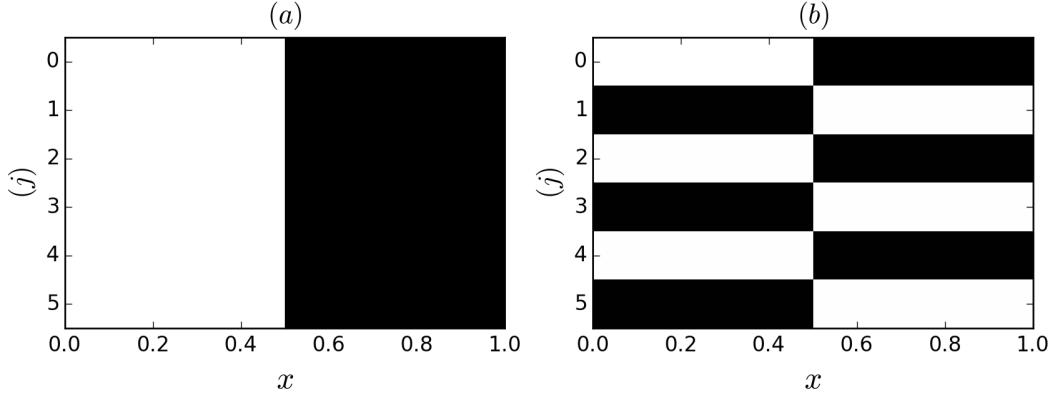
Figure 5.4: Space-time plots show the evolution of striations of a half-white, half-black initial condition by a composition $\Pi \circ M_B$. In (a) the permutation $\Pi = [1324]$ unmixes the new striations formed by the baker's transformation, reassembling the initial condition, while in (b) the permutation $\Pi = [4231]$ behaves similarly reassembling the black and white sub-segments after they are swapped within the domain.

number of interfaces in the concentration field is increased by the permutations. However this is not always the case.

For some of the permutations in S_4, different behaviour is observed. Take for example the permutation $\Pi = [1324]$ which swaps the middle two sub-intervals. After one iteration by the baker's transformation, the concentration field consists of four sub-segments of width $1/4$, as seen in Figure 5.4 (a). Applying the permutation reassembles the white sub-intervals and the black sub-intervals, so that only two sub-intervals remain with width $1/2$. Repeated application of the map $\Pi \circ M_B$ repeats this behaviour and no increased segregation can occur. Similarly, the permutation $\Pi = [4231]$ reassembles the striations but interchanges their position in the domain at each iteration, with no increased segregation. This is because permutations $\Pi \in S_4$ align with the point $x = 1/2$. An initial condition constructed from four colours of equal size reveals what is happening, shown in Figure 5.5. For $\Pi = [1324]$ although no exchange between the sub-intervals $[0, 1/2)$ and $[1/2, 1)$ can occur from the transformation $\Pi \circ M_B$, exponential mixing occurs within the two halves independently of each other. The permutation creates two ergodic partitions.

Finally, Figure 5.6 (a) shows the evolution of the *half-white, half-black* initial condition from the baker's transformation M_B compared with (b) when a rotation permutation $\Pi = [45123]$ is applied. The rotation preserves the uniformity
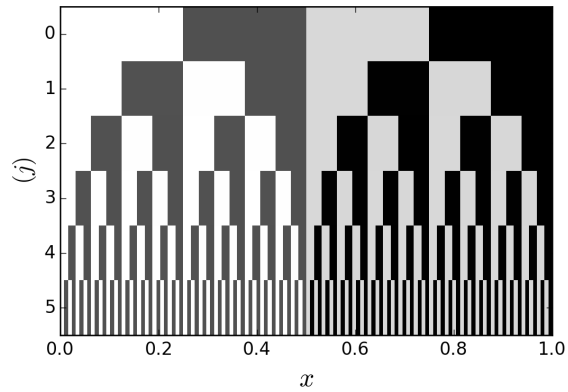
Figure 5.5: Space-time plot for IC with 4 colours, evolving under the map $\Pi \circ M_B$ with $\Pi = [1324]$. Exponential mixing occurs on the partitions $x \leq 1/2$ and $x > 1/2$ but no exchange of colour between the regions.
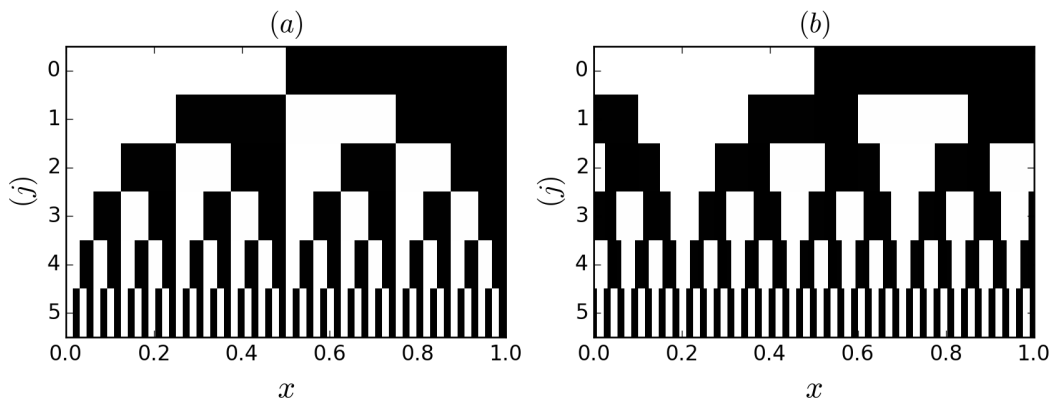
Figure 5.6: Space-time plots for the evolution of a half-white, half-black initial concentration is compared for (a) the baker's transformation and (b) when a rotation permutation $\Pi = [45123]$ is applied for $\Pi \circ M_B$. The rotation preserves the uniformity of the striation widths.
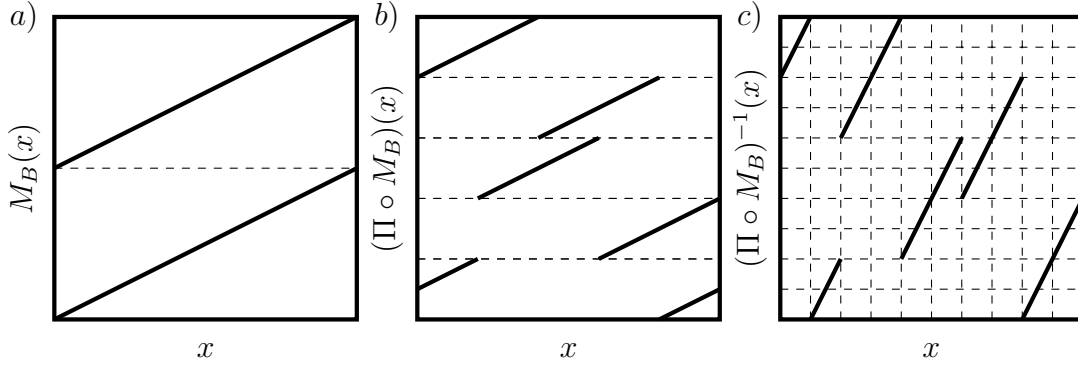
Figure 5.7: Graphical representations of a) the one-dimensional approximation of the baker's transformation, the halving map, b) the halving map composed with the permutation $\Pi = [35421]$, denoted $(\Pi \circ M_B)$, and c) the preimage map $(\Pi \circ M_B)^{-1}$ with an appropriate Markov Partition.

of the striation widths at each iteration, but rotates the concentration field periodically within the domain $[0, 1)$. In the next section the effect of permutations on segregation is quantified by the rates of decay of correlations.

5.2.2 Rate of decay of correlations

Figure 5.7 a) shows graphically the map M_B representing the evolution of points by the one-dimensional baker's transformation (5.2), which can be considered as a *halving map*. In this configuration, it is easy to see that the map is contracting to the stable fixed point $x = 0$ and is not chaotic, nor mixing. Figure 5.7 b) shows the map $\Pi \circ M_B$ for the permutation $\Pi = [35421]$, the same permutation used in the space-time plot in Figure 5.2 (b). The map $\Pi \circ M_B$ is no longer continuous but is clearly piecewise smooth with multiple pieces. The branches all have the same slope as the halving map, suggesting contraction and stability of points, however, the evolution of striations in Figure 5.2 (b) clearly show good mixing properties for the halving map and $\Pi \circ M_B$ with exponential reduction in striation widths and increase in interface. This is because the action of sets under the map $\Pi \circ M_B$ is related to ergodic, and strong mixing properties of the pre-image map $(\Pi \circ M_B)^{-1}$.

Recall from Definition 2.2.2 for strong mixing, then the action of the map M on the set A is taken to be the backward time iteration of the transformation. Figure 5.7 c) shows the inverse map $(\Pi \circ M_B)^{-1}$ which is not invertible, and as such to determine if the map $(\Pi \circ M_B)^{-1}$ is strong mixing, then it is precisely the evolution of sets by the mapping $\Pi \circ M_B$ which needs to be considered. The
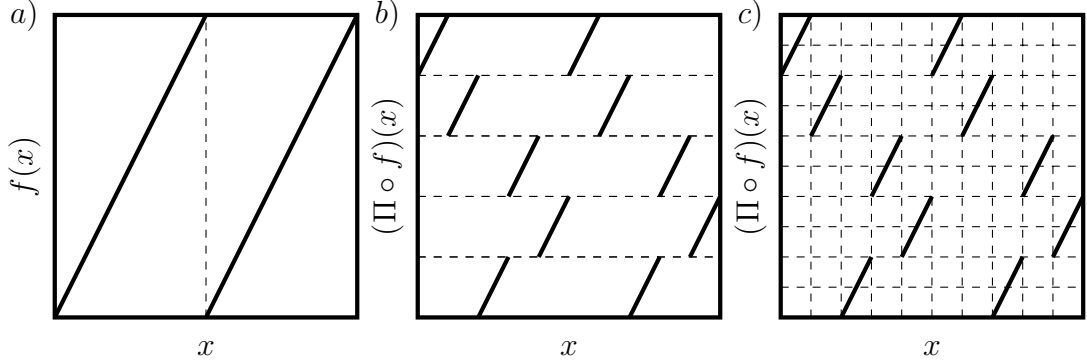
Figure 5.8: Graphical representations of a) the one-dimensional approximation of the doubling map b) the doubling map composted with permutation $\Pi = [35421]$, denoted $(\Pi \circ f)$, and c) $(\Pi \circ f)$ with an appropriate Markov Partition.

definition of strong mixing for $(\Pi \circ M_B)^{-1}$ can be thought of heuristically as the infinite time limit of the *half white, half black* intervals iterated by $\Pi \circ M_B$. If the set A is the half-open interval $[0, 1/2) \subset \mathcal{A}$, identified by colouring it white, then $\mu(A) = 1/2$. For any other interval $B = [a, b) \subset [0, 1)$, if under repeated images of A by $\Pi \circ M_B$, $\lim_{j \to \infty} B \cap (\Pi \circ M_B)^j(A) = \mu(B)/2$, then the map $(\Pi \circ M_B)^{-1}$ is strong mixing. That is, in every subinterval B, in the infinite limit half of the interval is occupied by white striations.

Therefore, the problem lies in determining if the pre-image maps $(\Pi \circ M_B)^{-1}$ are strong mixing. Fortunately, the pre-image maps have a special property that a Markov partition can be determined. An example partition is shown in Figure 5.7 c). Constructing a probability transition matrix between the Markov states gives a way for finding the mixing rate τ (2.19). The method for finding the mixing rates τ follows closely the work of Byott *et al.* (2013), who investigate the mixing properties of expanding maps f (5.3) composed with permutations, $\Pi \circ f$. Their application and method is now briefly explained, followed by how the method and results are extended to find the mixing rates for the maps $(\Pi \circ M_B)^{-1}$.

Mixing rates for $\Pi \circ f$

Byott *et al.* (2013) investigate the mixing rates of expanding maps (5.3) composed with permutations of equal sized cells (4.1), $\Pi \circ f$. A graphical representation of the doubling map f when $m = 2$ is shown in Figure 5.8 a), and in Figure 5.8 b) $\Pi \circ f$ with $\Pi = [35421]$. The map $\Pi \circ f$ is similarly piecewise smooth with multiple pieces, however there are differences between this map and $(\Pi \circ M_B)^{-1}$ with the same permutation. The relevant methods for defining the Markov partition,

Markov transition matrices and theory for transfer operators for the maps $\Pi \circ f$ from Byott *et al.* (2013) are now described.

A piecewise linear Markov map $M : \Omega \to \Omega$ has a finite partition $\mathcal{B}_q = \{I_i\}_{i=1}^q$ and a representative transition matrix B. Here B is a $q \times q$ matrix with $B_{ij} = 1$ if $I_j \subset M(I_i)$, and $B_{ij} = 0$ if $M(I_i) \cap I_j = \emptyset$, where it is assumed that M is differentiable on the interior of each element. From Figure 5.7 c) and Figure 5.8 c), the maps $\Pi \circ f$ with $m = 2$, and $(\Pi \circ M_B)^{-1}$ are piecewise linear Markov maps, with a suitable partition $\mathcal{B}$ shown. The partition $\mathcal{B}$ can be determined systematically for $\Pi \circ f$ for a given m and $\Pi \in S_N$.

Consider a partition B_q of constructed for $q \geq 1$ by

$$\mathcal{B}_q = \{[j/q, (j+1)/q) \quad : \quad 0 \leq j \leq q-1\}, \tag{5.4}$$

of $\Omega = [0, 1)$ into q equal subintervals. The map f with stretch factor m, is Markov with respect to the partition $\mathcal{B}_m$, while the permutation map Π, is Markov with respect to $\mathcal{B}_N$. The composition map $\Pi \circ f$ however, is Markov with respect to $\mathcal{B}_{mN}$, shown in Figure 5.8 c).

The probability transition matrix for $\Pi \circ f$ is constructed in the following way. Index the subintervals of the partition $\mathcal{B}_q$ by $\{0, 1, \ldots, q\}$ and number the rows and columns of the associated transition matrix accordingly. Define the matrix $B(m, N)$ to represent the transition matrix for the expanding map f with respect to the partition $\mathcal{B}_{mN}$. That is for $0 \leq i, j \leq mN - 1$

$$B(m, N)_{ij} = \begin{cases} 1 \text{ if } j = mi + d \mod Nm \text{ with } 0 \leq d \leq m - 1. \\ 0 \text{ otherwise.} \end{cases} \tag{5.5}$$

Additionally, define the matrix $A(m, N)$ with respect to $\mathcal{B}_N$ such that $0 \leq i, j \leq N - 1$ such that

$$A(m, N)_{ij} = \begin{cases} 1 \text{ if } j = mi + d \mod N \text{ with } 0 \leq d \leq m - 1. \\ 0 \text{ otherwise.} \end{cases} \tag{5.6}$$

For example, taking $m = 2$ and $N = 3$, the matrices are

$$B(2, 3) = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}, \qquad A(2, 3) = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}. \tag{5.7}$$

Although A is not defined on the Markov partition $\mathcal{B}_{mN}$, it has importance in the proof of the theorems of Byott *et al.* (2013), and is essential in establishing bounds on the mixing rates of $(\Pi \circ M_B)^{-1}$ in the next section.

To find the corresponding state transition matrices for $\Pi \circ f$, the columns of matrices $B(m, N)$ and $A(m, N)$ are permuted. Given a permutation Π, the matrix $P(\Pi)$, an $N \times N$ matrix representing the permutation is given by

$$P(\Pi)_{ij} = \begin{cases} 1 \text{ if } j = \Pi(i) \\ 0 \text{ otherwise.} \end{cases} \tag{5.8}$$

and similarly the matrix $Q(\Pi)$ can be defined as a $mN \times mN$ matrix derived from $P(\Pi)$, in which every 1 is replaced with a $m \times m$ identity matrix and 0 with a $m \times m$ zero matrix. The state transition matrix for the composition map $\Pi \circ f$ is then the product $B(m, n)Q(\Pi)$. The action of this product is that $Q(\Pi)$ causes a permutation on the columns of $B(m, N)$. For example, take the $m = 2$ and $N = 3$ and $\Pi = [213]$, then

$$P(\Pi) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad Q(\Pi) = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{5.9}$$

The state transition matrices for the transformation $\Pi \circ f$ with $\Pi = [213]$ are derived by permuting the columns and are

$$A(2,3)P(\Pi) = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}, \qquad B(2,3)Q(\Pi) = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}. \tag{5.10}$$

Note that $P(\Pi)A(m, N)$ is the matrix obtained by applying the inverse permutation Π^{-1} to the rows of $A(m, N)$. Byott *et al.* (2013) prove that the nonzero eigenvalues of $B(m, N)Q(\Pi)$ are equivalent to the non-zero eigenvalues of $A(m, N)P(\Pi)$. This is covered in more detail in the next section. The main results and theorems from Byott *et al.* (2013) are now discussed.

It is proven combinatorially that $\Pi \circ f$ is topologically mixing dependent on the parameters m and N and Byott *et al.* (2013) arrive at the following theorem.

Theorem 5.2.1. *[Byott* et al. *(2013)] Let f be as described in* (5.3)

1. *if N is not a multiple of m, then $\Pi \circ f$ is mixing for all $\Pi \in S_N$.*

2. *if $N > m$ and N is a multiple of m, say $N = ml$, there there will be some $\Pi \in S_N$ for which $\Pi \circ f$ is not mixing. As $l \to \infty$ (with fixed m), however, the proportion of permutations $\Pi \circ f$ mixing tends to 1.*

This theorem gives a condition on the parameters m and Π which ensures mixing. When determining the mixing rates of the maps, i.e. the rate of decay of correlations, a suitable functional space needs to be defined. Since the maps $\Pi \circ f$ are no longer continuous, the transformation does not preserve the space of Hölder continuous functions. Instead $\mathfrak{X}_s$ is taken to be the Banach space of functions with bounded variation BV. Given a function $\phi : [a, b] \to \mathbb{R}$, the total variation of ϕ is defined

$$TV(\phi) = \sup \left\{ \sum_{k=1}^{n} |\phi(x_k) - \phi(x_{k-1})| : a = x_0 \leq \cdots \leq x_n = n \right\} \qquad (5.11)$$

where the supremum is taken over all partitions of $[a, b]$. The function ϕ has bounded variation if $TV(\phi) < \infty$.

Now a suitable Banach space has been chosen, the spectrum of the transfer operator for the map acting on this space of functions $P_f|_{BV} : BV \to BV$ will give a means to determine the mixing rate. The spectrum is contained in the unit disk of the complex plane and has a trivial eigenvalue 1, corresponding to the equilibrium state. Since f is piecewise expanding, with the expansion factor uniformly bounded away from 1, the essential spectral radius r_{ess} can be interpreted as the slowest local mixing rate of the system. The spectrum may also contain isolated points of modulus greater than r_{ess}. The *global* mixing rate τ of the system is therefore determined by

$$\tau = \sup\{|\lambda| : \lambda \in Spec(P_f|_{BV}) \backslash \{1\}\}. \qquad (5.12)$$

For $P_f|_{BV}$ the essential spectal radius is $r_{ess} = 1/m$ and there are no isolated eigenvalues with $\lambda > 1/m$ except for 1, hence $\tau = 1/m$. For the composition maps $\Pi \circ f$ this is not the case. When $\Pi \circ f$ is mixing, the mixing rate is defined as

$$\tau_\Pi = sup \left\{ |\lambda| : \lambda \in Spec(P_{\Pi \circ f|BV} \backslash \{1\} \right\} \qquad (5.13)$$

For the choice of two integers m and N, Byott *et al.* (2013) prove the following theorem for the mixing rates of $P_{\Pi \circ f}$;

Theorem 5.2.2. *Fix m, $N \geq 2$ and consider the transformations $\Pi \circ f$ for $\Pi \in S_N$. Then the following hold;*

1. *For all $\Pi \in S_N$, the essential spectral radius is given by $r_{ess}(P_{\Pi \circ f}|_{BV}) = 1/m$.*

2. *If $N > m$ and $gcd(m, N) = 1$, then for each $\Pi \in S_N$, we have*

$$\tau_{\Pi \circ f} \leq \tau_{max} = \frac{\sin(m\pi/N)}{m \sin(\pi/N)}. \tag{5.14}$$

The proof of the theorem is arrived at by a detailed study of the Fredholm matrices, and the related Markov transition matrices associated with $\Pi \circ f$. It is not necessary to discuss the full details of the proof, only that the mixing rate τ_Π is determined from the leading non-unit eigenvalue λ of the matrices $B(m, N)Q(\Pi)$ and $A(m, N)P(\Pi)$ as $\tau_\Pi = (1/m) * \lambda$. The values r_{ess} and τ_{max} essentially provide lower and upper bounds on the mixing rates. The main conclusion of the study is that permutations preserve the mixing rate or typically make it worse.

For a given m and N and $\Pi \in S_N$, Byott *et al.* (2013) define

$$\Lambda_\Pi = \left\{ \lambda \in Spec(P_{\Pi \circ f}|_{BV}) \text{ such that } |\lambda| = \tau_\Pi \right\}, \tag{5.15}$$

τ_Π as defined in 5.13. This is because there may be multiple eigenvalues with $|\lambda| = \tau_\Pi$. The geometric arrangement of the set $\cup_{\Pi \in S_N} \Lambda_\Pi$ in the complex plane was of interest for Byott *et al.* (2013).

Mixing rates for $(\Pi \circ M_B)^{-1}$

Similar to Byott *et al.* (2013), the construction of the Markov matrices for the maps $(\Pi \circ M_B)^{-1}$ can be generalised. Compare the differences between the map $\Pi \circ f$ in Figure 5.8 and $(\Pi \circ M_B)^{-1}$ in Figure 5.7. The map could be rewritten in terms of a permutation Γ such that $(\Pi \circ M_B)^{-1} = f \circ \Gamma$. In terms of a Markov transition matrix, this implies that we take a permutation on the rows with Γ, in contrast for the maps $\Pi \circ f$ in which a permutation occurs on the columns.

For a given permutation Π, $Q(\Pi^{-1})B(m, N)$ produces the appropriate Markov transition matrices for the map $(\Pi \circ M_B)^{-1}$. The eigenvalues of these transition matrices can be computed and would give the mixing rates of the map $(\Pi \circ M_B)^{-1}$. However no assumptions can be made that the Theorems of Byott *et al.* (2013) carry over to this system of maps, or that $P(\Pi^{-1})A(m, N)$ has the same eigenvalue spectrum as $Q(\Pi^{-1})B(m, N)$. This is addressed in the following two lemmas.
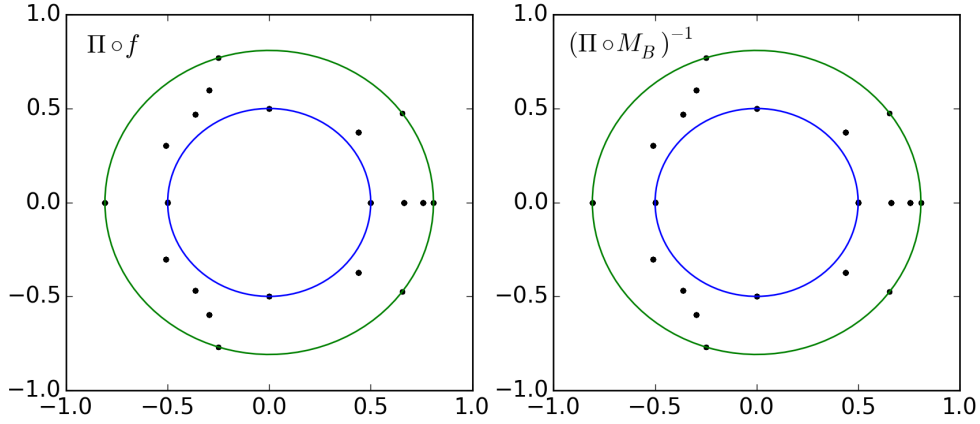
Figure 5.9: Λ_Π is plotted for $\Pi \circ f$ and $(\Pi \circ M_B)^{-1}$ and shows that the distributions of τ in the unit plane agree, along with the lower (r_{ess}) and upper (τ_{max}) bounds shown in blue and green respectively.

Lemma 5.2.3. *For all m, N and $\Pi \in S_N$, the non-zero eigenvalues of $Q(\Pi)B(m, N)$ are the same as those of $P(\Pi)A(m, N)$.*

Proof. The proof follows the method of Byott *et al.* (2013) for the proof that the matrices $B(m, N)Q(\Pi)$ and $A(m, N)P(\Pi)$ have the same none zero eigenvalues. View the matrix $Q(\Pi)B(m, N)$ as determining a linear endomorphism θ on the space $V = \mathbb{C}^{mN}$ of column vectors. Clearly $B(m, N)$ has rank N, since the first N rows are linearly independent and the remaining rows merely repeat these. Therefore since $Q(\Pi)B(m, N)$ is just a permutation of these rows, then $Q(\Pi)B(m, N)$ also has rank N. The kernel W of θ therefore has dimension $N(m - 1)$, and θ induces an endomorphism $\bar{\theta}$ on the quotient space V/W of dimension N. The eigenvalues of θ, (that is of QB) are therefore the eigenvalues of $\bar{\theta}$ together with the eigenvalue 0 with multiplicity $N(m - 1)$ coming from W. The result follows if the matrix PA represents $\bar{\theta}$, which it does following from the construction of the basis vectors in Byott *et al.* (2013). This is omitted for brevity. $\square$

Lemma 5.2.4. *Fix m, $N \geq 2$ and consider the transformations $(\Pi \circ M_B)^{-1}$ where M_B is the one-dimensional baker's transformation and $\Pi \in S_N$. Then*

1. *For all $\Pi \in S_N$, the essential spectral radius is given by $r_{ess} = 1/m$.*

2. *If $N > m$ and $\gcd(m, N) = 1$, then, for each $\Pi \in S_N$, we have*

$$\tau_\Pi \leq \tau_{max} = \frac{\sin(m\pi/N)}{m\sin(\pi/N)}.$$

Proof. Note that for a given $\Pi \in S_N$, a $\Gamma \in S_N$ can be found such that

$$P(\Pi^{-1})A(m, N) = A(m, N)P(\Gamma).$$

Since it follows that the non-zero eigenvalues of $P(\Pi^{-1})A(n, M)$ and $Q(\Pi^{-1})B(n, M)$ are equivalent, then

$$\cup_{\Pi \in S_N} \Lambda_\Pi = \cup_{\Gamma \in S_N} \Lambda_\Gamma, \tag{5.16}$$

and all results and bounds on the mixing rate found by for $\Pi \circ f$ hold for $(\Pi \circ M_B)^{-1}$. $\qquad\square$

In Figure 5.9 the distribution of the leading eigenvalues $\cup_{\Pi \in S_N} \Lambda_\Pi$ are plotted in the unit-circle for the maps $\Pi \circ f$ and $(\Pi \circ M_B)^{-1}$ and find agreement for $m = 2$ and $N = 5$. For all other choices m and N the distributions agree.

In conclusion, the mixing rates of the baker's transformation composed with permutations can be found using Markov partitions. This method follows closely the work of Byott *et al.* (2013), who found bounds on the mixing rate dependent on the stretch factor m and the number of equal sized cells in the permutation. Although slightly different in construction, both $\Pi \circ f$ and $(\Pi \circ M_B)^{-1}$ have the same distribution of eigenvalues for a given m and N.

5.3 Introducing diffusion

5.3.1 Deriving the transfer operator

The inclusion of diffusion is computed through the action of operator splitting, as was previously performed in Chapters 3 and 4. In this case, the advective step can also be considered as split operators, where the operator evolving the concentration field by the baker's transformation P_{M_B}, is followed by an operator representing the evolution by the permutation P_Π, such that one full advective step is

$$c^{(j+1)}(x) = P_D \circ P_\Pi \circ P_{M_B} c^{(j)}(x). \tag{5.17}$$

As for the maps considered in Chapter 3 and 4, the evolution of the concentration field $c(x)$ is captured easily by a transfer matrix representing the transfer of amplitude between complex coefficients of a discrete Fourier expansion of $c(x)$ when periodic boundary conditions are applied.

The action of the one-dimensional baker's transformation is given by

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \hat{c}_q \int_0^1 e^{2\pi i M_B^{-1}(x) - 2\pi i k x} = \sum_{q=-\infty}^{\infty} \hat{c}_q \int_0^1 e^{2\pi i q 2x - 2\pi i k x} \mathrm{d}x \tag{5.18}$$

which results in the transfer matrix

$$\boldsymbol{T}_{kq} = \boldsymbol{\delta}_{0,(k-2q)}. \tag{5.19}$$

The transfer matrix $\boldsymbol{M}_{kq}$ for the operator P_Π is the same as in 4.12. However, deriving the single advective transfer matrix via $\boldsymbol{MT}$ is not computationally efficient since for a truncation of Q modes, the matrix for the baker's transformation produces a matrix of size $2Q \times 2Q$, with Q columns that are all zeros. It follows that the product matrix $\boldsymbol{MT}$ similarly has Q zero columns.

It is more computationally efficient to derive the transfer matrix from the action of both maps simultaneously. The application of M_B followed by M_Π is such that the Fourier coefficients evolve according to

$$\hat{c}_k^{(j+1)} = \sum \hat{c}_q \int_0^1 e^{2\pi i q (M_\Pi \circ M_B)^{-1}(x) - 2\pi i k x} \mathrm{d}x \tag{5.20}$$

and the transfer matrix defined

$$\boldsymbol{M}_{kq} = \int_0^1 e^{2\pi i q 2x - 2\pi i k M_\Pi(x)} \mathrm{d}x. \tag{5.21}$$

When $k \neq 2q$,

$$\boldsymbol{M}_{kq} = \frac{1 - \omega^{(2q-k)}}{2\pi i (2q - k)} \sum_\ell^N \omega^{k\Pi(\ell) - 2q\ell}, \tag{5.22a}$$

while for $k = 2q$

$$\boldsymbol{M}_{kq} = \frac{1}{N} \sum_{\ell=1}^N \omega^{k(\Pi(\ell) - \ell)}. \tag{5.22b}$$

Since $f^{-1} = M_B^x$ when M_B is the one-dimensional map presented Figure 5.7 when $m = 2$, the transfer matrix $\boldsymbol{M}_{kq}$ for $\Pi \circ f^{-1}$ can be easily derived for any stretch factor m via

$$\boldsymbol{M}_{kq} = \frac{1 - \omega^{(mq-k)}}{2\pi i (mq - k)} \sum_\ell^N \omega^{k\Pi(\ell) - mq\ell} \tag{5.23a}$$

when $k \neq mq$, while for $k = mq$

$$\boldsymbol{M}_{kq} = \frac{1}{N} \sum_{\ell=1}^N \omega^{k(\Pi(\ell) - \ell)}. \tag{5.23b}$$

Matrices $\boldsymbol{M}_{kq}$ defined in (5.22) and (5.23) share the symmetry $\boldsymbol{M}_{-k-q} = \overline{\boldsymbol{M}_{kq}}$, which implies the reality of concentration fields are preserved on the application of the transfer operator. Additionally, if λ is an eigenvalue with eigenvector v, then $\overline{\lambda}$ is also an eigenvalue with eigenvector $\overline{v}$.
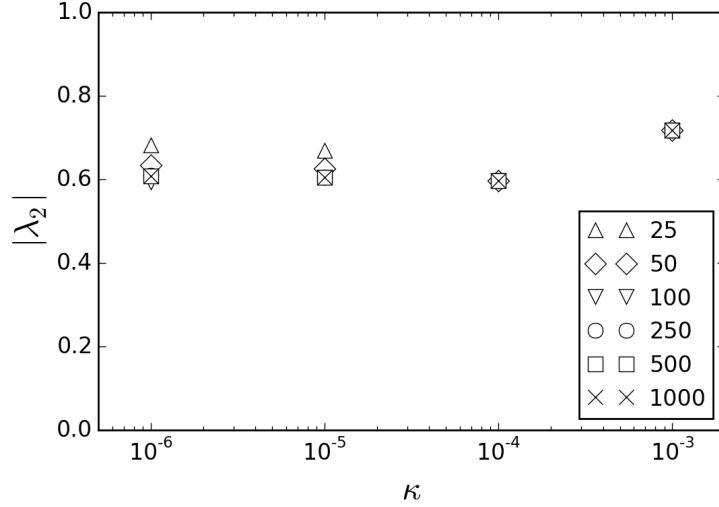
Figure 5.10: Resolution testing showing convergence in the leading eigenvalue $|\lambda_2|$ on increasing the number of wavemodes Q for $\Pi \circ M_B$ with $\Pi = [213]$. Different diffusivity values κ are test and convergence occurs around $Q = 500$ for all considered.

The full advective-diffusive step is then to iterate the concentration field according to

$$\hat{c}_k^{(j+1)} = \boldsymbol{d}_{kp}\boldsymbol{M}_{pq}\hat{c}_q^{(j)} \tag{5.24}$$

where $\boldsymbol{d}_{kp}$ is defined as in (4.13).

As in Section 4.2.2, the initial conditions considered will be variations of the square wave, sine or cosine profiles, or randomised concentration fields constructed by randomising the amplitude $\hat{c}_k$ of the first handful of discrete Fourier modes. The decay of variance forms the main measure of mixing.

5.3.2 Resolution of Fourier modes

For the computations, it is required that the number of Fourier modes are truncated to deal with finite matrices. A resolution test is performed to ensure a sufficient number of Fourier modes are included for the diffusivity values of interest, without introducing significant numerical diffusion. In the previous Chapter, a surprisingly small truncation was needed, $Q \sim 250$, however since the model of this Chapter involves chaotic behaviour it is expected that a larger truncation is required.

The leading eigenvalue of the full advective-diffusive matrix $\boldsymbol{dM}$ is easily computed and the values compared for varying Q. As a test case, the transformation
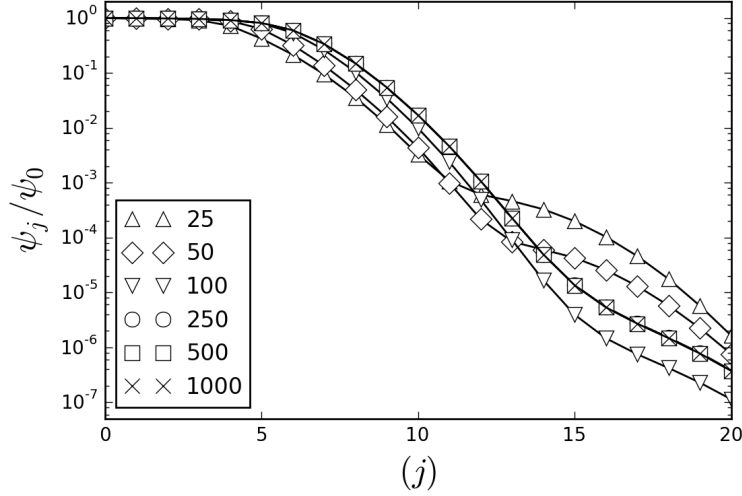
Figure 5.11: The decay of variance is plotted for the evolution of $c^0(x) = \cos(2\pi x)$ under the mapping $\Pi \circ M_B$ with $\Pi = [213]$ and diffusivity coefficient $\kappa = 10^{-6}$. Increasing the number of wavemodes used in computation shows convergence occurs around $Q = 500$.

$\Pi \circ M_B$ with $\Pi = [213]$ is used. Figure 5.10 plots the modulus of the leading eigenvalue $|\lambda_2|$, comparing 6 values of the truncation number Q of Fourier mode and 4 diffusivity values κ. The permutation perturbs the uniformity of the baker's transformation and results in leading eigenvalues with $|\lambda_2| \neq 0$. Although it can be noted that when $\kappa = 10^{-6}$, the approximation of $|\lambda_2|$ is quite accurate when only a small number of modes are taken with slight over approximations when low truncation values are used. When Q is taken to be $Q = 500$ or $Q = 1000$ no difference in $|\lambda_2|$ is observable. Convergence was also observed in the values $|\hat{c}_k|$ for all $k \in [-Q, Q]$ of the eigenfunction v_2 for $Q = 500$ and $Q = 1000$, but the results are omitted for brevity.

Additionally, the decay of variance can be used to observe the effects of the truncation value Q. Figure 5.11 shows the results for evolving the initial condition $c^{(0)}(x) = \cos(2\pi x)$ with $\kappa = 10^{-6}$ for different values of Q. In this case the variation in the finite time mixing behaviour is readily seen when Q is not larger enough and again agreement occurs well for $Q \geq 250$, where the $Q = 250$ marker is hidden behind the markers for higher values. Furthermore $Q = 2500$ was tested which showed no improvement over the chosen value, hence $Q = 1000$ is sufficient to capture the dynamics of this transformation with diffusion for $\kappa \geq 10^{-6}$ and is taken throughout the computations whose results follow.

5.4 Mixing with diffusion

5.4.1 Mixing in the initial transient

As a leading example, a comparison between mixing in the initial stages can be made between the baker's transformation acting alone, or composed with rotations and interleaving permutations from S_3.

It was observed in Section 5.2.1 that the rotation permutation $\Pi = [45123]$, while rotating the concentration field after the action of the baker's transformation preserved the uniformity of the striation widths. This is observed for all rotation permutations, which will be denoted as the group S_N^R for brevity. Given $\Pi \in S_N^R$ the map for the full advective step $\Pi \circ M_B$ reduces to

$$M(x) = 2x + m/N \quad \mod 1, \tag{5.25}$$

and the transfer matrix (5.21) reduces to

$$\boldsymbol{M}_{kq} = \boldsymbol{\delta}_{0,(k-2q)}\omega^{kq}, \tag{5.26}$$

with $|\omega^{kq}| = 1$ for all k,q. This represents explicitly that the amplitude of concentration at each wavemode is directly mapped to the wavemode twice its value and rotated by the factor ω^{kq}. For the initial condition, $c^{(0)}(x) = \cos(2\pi x)$, which has $\hat{c}_{-1,1} = 0.5$ and $\hat{c}_{|k|>1} = 0$ the evolution of the variance for M_B and $\Pi \circ M_B$ for $\Pi \in S_N^R$ can be explicitly computed as

$$\psi(j) = \psi(0)\exp\left(-\frac{32}{3}\pi^2\kappa(4^j - 1)\right), \tag{5.27}$$

in which details on the derivation are in Appendix B.3. There is an exponential cascade to large wavemodes where the action of diffusion is extremely effective, resulting in super-exponential decay of variance, expressed by the 4^j term in the exponential of (5.27). This result is the same as discussed in Fereday *et al.* (2002) for the baker's transformation, who investigate the effect of adding non-uniformity in stretching rates. The extension here is that it applies to all composition maps $\Pi \circ M_B$ where $\Pi \in S_N^R$.

Figure 5.12 (*a*) shows three initial conditions, and in Figure 5.12 (*b*) the respective decay of variance under the action of the baker's transformation M_B and the composition map $\Pi \circ M_B$ with $\Pi = [213]$. The diffusivity coefficient was $\kappa = 10^{-5}$ for all. The initial evolution of $c^{(0)}(x) = \cos(2\pi x)$ is plotted as the solid blue line and follows (5.27). For this initial condition evolved by the baker's
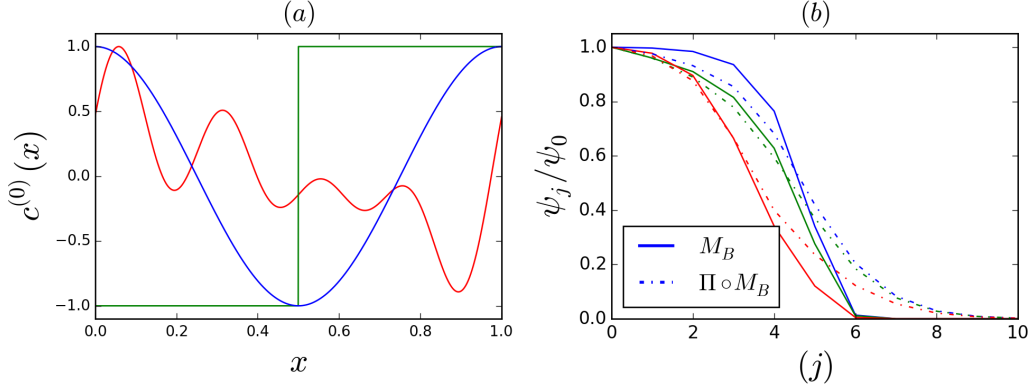
Figure 5.12: The initial transient of variance decay is shown for $\kappa = 10^{-5}$ and 3 different initial conditions, plotted in (a). The mixing of transformations M_B (solid) and $\Pi \circ M_B$ with $\Pi = [213]$ (dot-dashed) are compared in (b). In the initial stages composition with a permutation increases the decay of variance over the baker's transformation acting alone, however eventually the non-perturbed map is faster.
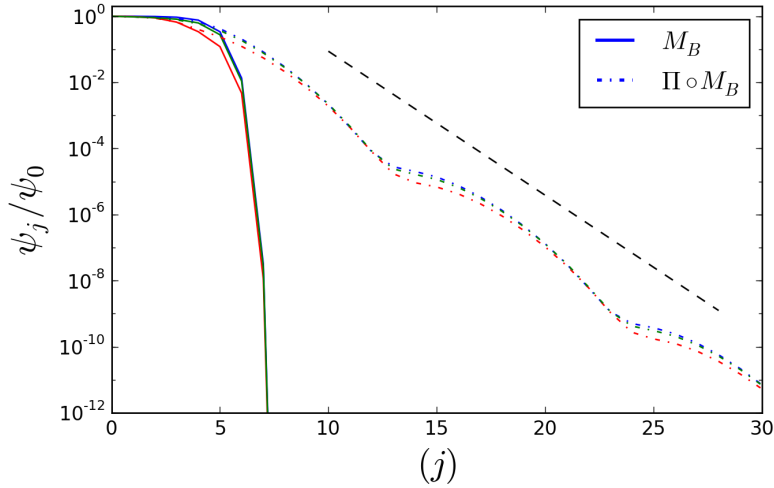
Figure 5.13: The decay of variance is shown for $\kappa = 10^{-5}$ and 3 different initial conditions of Figure 5.12 on a linear-log plot revealing the long-time mixing. M_B acting alone has super-exponential decay of variance for all initial conditions, while $\Pi \circ M_B$, $\Pi_1 = [213]$ has long time exponential decay with average rate given by $|\lambda|^{2j}$.

transformation M_B, the first few iterations sees little depletion in the variance which remains nearly constant. This is due to the fact that for small rates of diffusivity, the scales on the concentration field are too large in the initial stages for diffusion to be effective. It is only after a number of iterations, in which the segregation of the concentration field becomes significant, when diffusion is effective. The variance in the first few iterations is depleted quicker for these other initial conditions compared to $c^{(0)}(x) = \cos(2\pi x)$ due to increased gradients caused by the discontinuity and dominant amplitude in higher wavemodes. The conditions are the square wave and a concentration field composed by randomising the amplitudes of the first four modes, plotted as solid lines in green and red respectively. Figure 5.13, plots the decay of variance on a linear-log scale revealing clearly that eventual super-exponential occurs for M_B for all initial conditions.

For all three initial conditions, the addition of the permutation $\Pi = [213]$ increases depletion of variance in the initial iterations, represented by the dot-dashed lines in Figures 5.12 (a) and (b). Recall from Section 2.1.1 that the decrease in variance from diffusive mixing contributed entirely to an increase in gradients in the concentration field. In smoothly deforming transformations this can only happen through increasing the number of striations and decreasing their width by repeated stretching and folding. In discrete Fourier space this is the cascade to large wavemodes, such as the action presented by the baker's transformation. If the stirring transformation is discontinuous, large, sharp interfaces can be introduced into the concentration field at each iteration, effectively increasing the homogenization by diffusion in the early stages of mixing.

For the composite maps $\Pi \circ M_B$ where Π is not a rotation, there is dispersion of amplitude in the complex coefficients to large wavemodes due to the introduction of discontinuities, and the diffusive step reduces the amplitude more effectively at the resulting sharp gradients in the concentration field. This explains the enhancement of mixing in the first handful of iterations in Figure 5.12 when the diffusivity coefficient κ is small. This improvement in mixing over the baker's transformation holds for nearly all Π investigated in the maps $\Pi \circ M_B$ which are not rotations.

However, this enhancement of the mixing by $\Pi \circ M_B$ when $\Pi \notin S_N^R$, does not hold in the later stages. From Figure 5.13 it is seen that the long-time mixing rate is exponential and slower than the baker's transformation acting alone. This long-time mixing rate is explored in more detail in the next section.
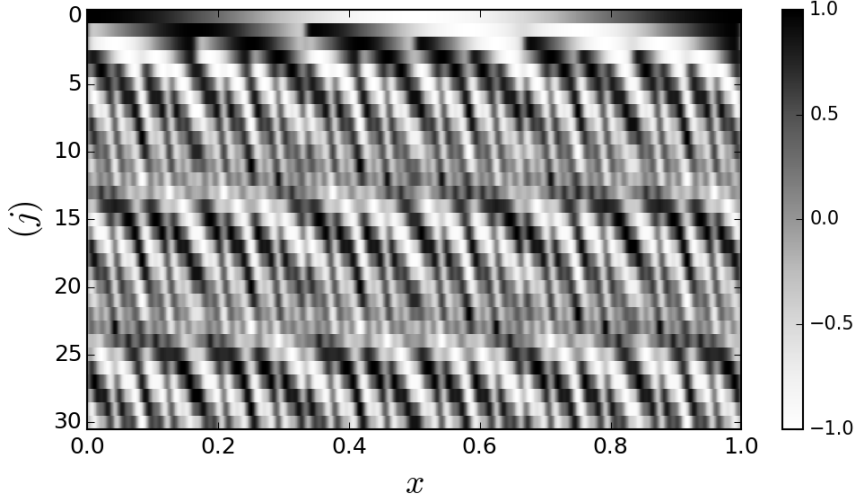
Figure 5.14: Eigenfunction can be seen emerging as a periodic pattern in the evolution of the concentration field. There is an attempt to renormalise the concentration field at each iteration to highlight the self-similar pattern on the eigenmode. $c^{(0)}(x) = \cos(2\pi x)$ and $\kappa = 10^{-5}$

5.4.2 Long-time exponential mixing

In Figure 5.13 the decay of variance for $\Pi \circ M_B$ is exponential in the long time limit and is the same for all initial conditions considered. This exponential decay rate is predicted from the transfer matrix $\boldsymbol{dM}$ of the composed advection-diffusion iterative operator. After a number of initial iteration the eigenfunctions v_k decay at the rate of their respective eigenvalues until the slowest decaying eigenfunction with the slowest decay rate, the second leading eigenvalue λ_2, dominates the evolution of the concentration field and the long-time exponential decay rate of the variance is given by $\psi(j) \sim |\lambda_2|^{2j}$ as expected. This expected decay of variance is plotted on Figure 5.13 as the black dashed line. Although the actual decay of variance shows some variation in the decay profile, the dashed line accurately predicts the average exponential decay rate. Oscillations arise due to the complex value of the eigenvalues λ_k of the transfer matrix, and the respective eigenfunctions evolve spatially on further applications of the transfer operator. Figure 5.14 plots the space-time evolution for $c^{(0)}(x) = \cos(2\pi x)$ with $\kappa = 10^{-5}$, where $c^{(j)}(x)$ is renormalised at each iteration to show spatial comparison of the field between iterations. It is observed that the concentration field has a spatially periodic pattern, aligning with the periodic oscillations in the variance profile of Figure 5.13. This is the eigenfunction v_2 associated with λ_2. Calculating the

Permutations	eigenvalue	$\kappa = 10^{-3}$	$\kappa = 10^{-4}$	$\kappa = 10^{-5}$	$\kappa = 10^{-6}$
[132],[213],[321]	$|\lambda_2|$	0.7177	0.5962	0.6054	0.6091

Table 5.1: Decay rates for permutations $\Pi \in S_3$ for $\Pi \notin \Pi \in S_N^R$ for different values of κ. The modulus $|\lambda_2|$ agrees for all values of κ due to a rotation symmetry of the map $M_\Pi \circ M_B$ for these permutations.

eigenvalues for other two interleaving permutations in S_3 gives the same leading eigenvalue, listed in Table 5.1, and plots of the variance match those seen in Figure 5.13.

The long-time eigenfunction can be thought of a persistant patterns with a decay rate irrespective of the initial condition, similar to those seen in chaotic advective systems with non-uniformity in the stretching rates of the flow field. This is a well know global mechanism for mixing in smoothly deforming systems (Section 2.5). The contrast here however, is that the dispersion between Fourier modes occurs due to permutation composition, rather than the non-uniformity in stretching rates for non-uniform maps. The decay rate is not governed by any Local Lagrangian behaviour since the stretching rates are the same almost everywhere except for a countable number of discontinuities which form a set of measure zero, introduced by the cutting and shuffling transformation. The non-uniformity arises in the rearrangement of striations by the interval exchange, observed in the segregation plots of Figure 5.2 for the permutation $\Pi = [35421]$. Non-uniformity in striation arrangement of the concentration field occurs for all $\Pi \notin S_N^R$, such as $\Pi = [213]$ here.

However, unlike the strange eigenfunctions of non-uniform smooth systems which align with regions of low stretch, the eigenfunction for the composition $\Pi \circ M_B$, $\Pi = [213]$, does not appear to correlate with physical characteristics of the underlying advective dynamics, such as periodic points or where the discontinuities are introduced. The same can be said when considering a range of diffusivity coefficients κ.

5.4.3 Effect of diffusion

Figure 5.15 plots the decay of variance for $\Pi \circ M_B$ with $\Pi = [213]$ for a range of diffusivity coefficients κ. For all values of κ considered, the initial condition is $c^{(0)}(x) = \cos(2\pi x)$. The effect of diffusion on the early stages of mixing is comparable to smoothly deforming systems. The transition from early stages, where
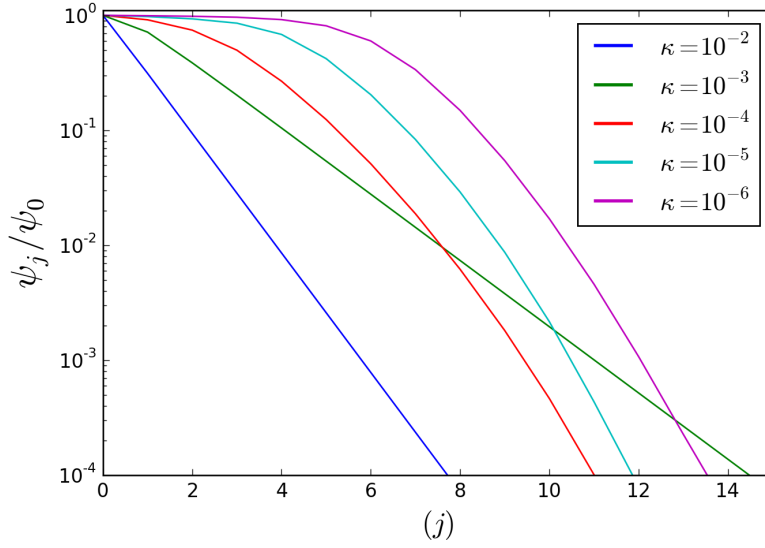
Figure 5.15: Decay of variance for $\Pi \circ M_B$ with $\Pi = [213]$ with the initial condition $c^{(0)}(x) = \cos(2\pi x)$. The diffusivity coefficient κ is varied showing a $\sim \log(\kappa)$ dependency on the long-time mixing behaviour emerging.

variance depletion is weak, to long-time exponential or super-exponential mixing stages is only a handful of iterations. It is weakly related by $\log(\kappa)$, similarly reported for non-uniform smoothly deforming systems [Thiffeault (2008); Wonhas & Vassilicos (2002)]. This contrasts the results for the purely cutting and shuffling, Interval Exchange Transformations of Chapter 4, where the relation on finite time mixing was polynomial with diffusion coefficient.

Once variance depletion becomes significant however, in Figure 5.15 it appears as though long-time mixing when $\kappa = 10^{-3}$ is slower than all $\kappa < 10^{-3}$ computed. To look at the effect of diffusion in the later, asymptotic mixing stage, the change in the modulus of the leading eigenvalues need only be considered. The permutations fall into two groups; $\Pi \in S_N^R$ or the interleaving permutations with $|\lambda_2| \neq 0$. For $\Pi \in S_N^R$ only a direct cascade between wavemodes $k = 2q$ occurs which results in $|\lambda_k| = 0$ and a super-exponential mixing rate for all initial conditions, seen in Figure 5.13.

For the interleaving permutations $|\lambda_2| \neq 0$ and are listed in Table 5.1 for a handful of diffusivity coefficients. The modulus of the leading eigenvalue is equivalent for all interleaving Π due to a rotational symmetry on the transformation $M_\Pi \circ M_B$, and $\lambda_3 = \overline{\lambda_2}$ due to the symmetry $\boldsymbol{M}_{-k-q} = \overline{\boldsymbol{M}_{kq}}$. Therefore the leading eigenvalue $|\lambda_2|$ need only be computed for one of the permutations $\Pi \notin S_3^R$ to see the changes of $|\lambda_2|$ with finer values of κ. Figure 5.16 plots the profile of $|\lambda_2|$

144
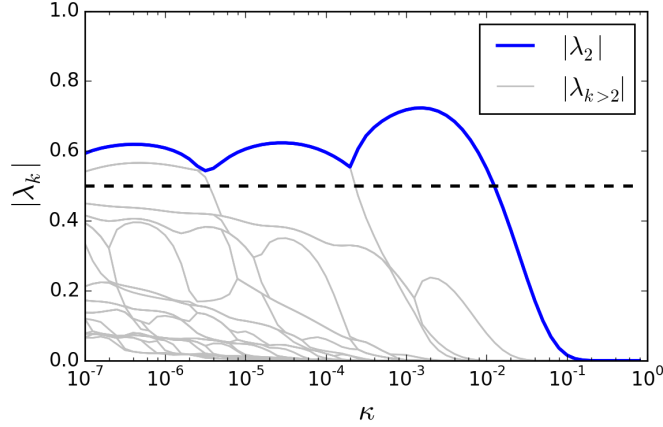
Figure 5.16: The leading eigenvalue $|\lambda_2|$ is plotted for the interleaving permutation and rotations thereof. Non-monotonicity in the value can be observed with periodic like behaviour in value occurring logarithmically with κ. $|\lambda_k|$ is plotted for $2 < k < 50$ also and shows the non-monotonicity arises with collision.

for $\Pi \circ M_B$ with $\Pi = [213]$ against a log scaling for the diffusion coefficient κ.

It is observed that the leading eigenvalue changes non-monotonically with diffusion, the highest peak occurring for large κ, and oscillatory-like peaks continuing on further decreasing values of κ. It appears as though there is some periodicity in growth and decay of these peaks which is logarithmic in κ. The largest peak achieves a value of $|\lambda_\kappa| = 0.7231$ around $\kappa = 0.00158$. The maximum values of the following peaks in the decease of κ are 0.6231 and 0.62 respectively, hence a slight decrease in their value. The profile for $\kappa = 10^{-3}$ in Figure 5.15 which has the slowest long-time mixing rate coincides which the largest bump of $|\lambda_2|$ plotted against κ in Figure 5.16. Plotted in grey are the modulus of the other eigenvalues $|\lambda_k|$, showing that the troughs occur through collisions with other eigenvalues in the spectrum of the composition operator $P_D \circ P_M$. Non-monotonicity persists in all eigenvalue profiles.

The eigenvalue collision which occurs in the region $\kappa \in [10^{-4}, 10^{-3}]$ can be understood by looking at the eigenfunctions. Figure 5.17 plots the eigenfunctions v_2 (blue) and v_4 (green) for the two values $\kappa = 10^{-3}$ and $\kappa = 10^{-4}$. Note that $\lambda_3 = \overline{\lambda_2}$ and $\lambda_5 = \overline{\lambda_4}$, and similarly $v_3 = \overline{v_2}$ and $v_5 = \overline{v_4}$, therefore the functions v_3 and v_5 are not shown. When $\kappa = 10^{-3}$ the eigenfunctions have distinct spatial scales in both the real ($\Re(v)$) and imaginary ($\Im(v)$) parts. As κ is decreased the spatial scales become more similar until the eigenvalue collision, which is observed in Figure 5.16, occurs and the spatial scales of the eigenfunctions agree. After
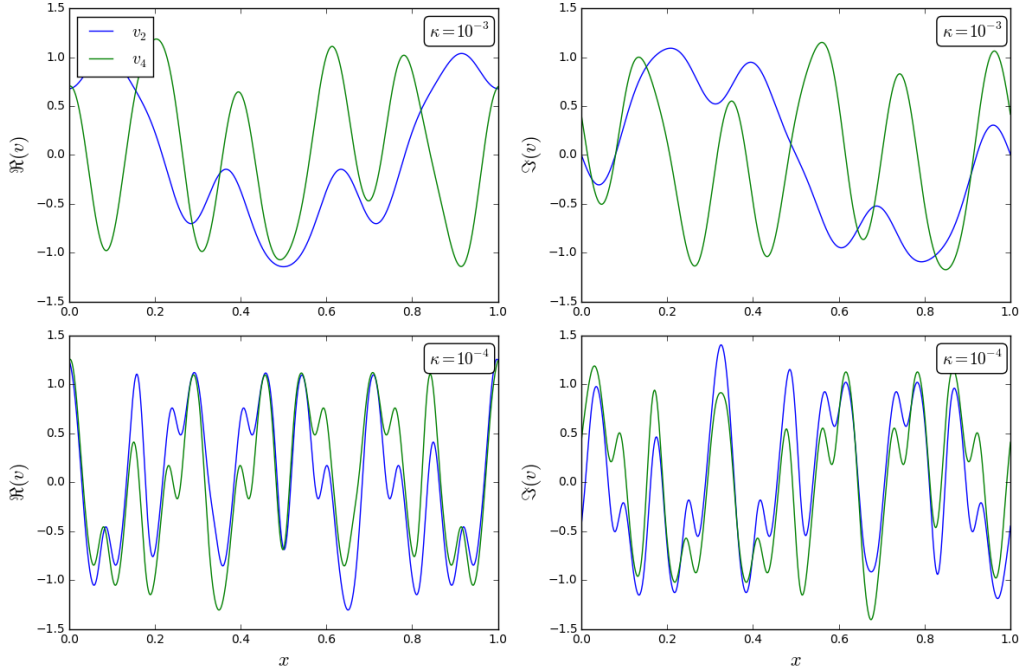
Figure 5.17: The real ($\Re(v)$) and imaginary ($\Im(v)$) parts of the eigenfunctions v_2 and v_4 are plotted for different values of κ, $\kappa = 10^{-3}$ and $\kappa = 10^{-4}$ on the top and bottom rows respectively. When $\kappa = 10^{-3}$, the eigenfunctions v_2 and v_4 have distinct spatial scales. When $\kappa = 10^{-4}$, where v_2 and v_4 share a symmetry in both $\Re(v)$ and $\Im(v)$ about $x = 0.5$.

collision the eigenfunctions v_2 and v_4 share a spatial symmetry, a reflection about $x = 0.5$. This is seen in Figure 5.17 when $\kappa = 10^{-4}$, where the reflected symmetry about $x = 0.5$ is seen in both the real and imaginary parts of the eigenfunctions. A similar eigenvalue collision occurs within the window $\kappa = [10^{-6}, 10^{-5}]$. When $\kappa = 10^{-5}$, v_2 and v_4 maintain the shared spatial scales. On decreasing κ eventually there is a division so that $|\lambda_2| = |\lambda_3| \neq |\lambda_4| = |\lambda_5|$, and a collision occurs in which $|\lambda_4| = |\lambda_6| = |\lambda_5| = |\lambda_7|$ instead. After this collision the spatial scales between v_2 and v_4 no longer agree, and instead the spatial scales between v_4 and v_6 now agree with symmetries in the eigenfunctions.

The dotted black line represents the diffusion-less, asymptotic mixing rate τ for the map $(\Pi \circ M_B)^{-1}$ for $\Pi = [213]$, the strong mixing rate. Values of $|\lambda_2|$ are seen to overshoot τ for most of the diffusivity coefficients considered, with no suggestion that convergence might occur as $\kappa \to 0$. Note that the parameters $N = 3$ and $m = 2$ satisfy Theorem 5.2.1, and predict lower and upper bounds

$$\tau_{min} = \frac{1}{2}, \quad \tau_{max} = \frac{\sin(2\pi/3)}{2\sin(\pi/3)} \approx 0.5, \tag{5.28}$$

146

Group	Permutations	eigenvalue	$\kappa = 10^{-3}$	$\kappa = 10^{-4}$	$\kappa = 10^{-5}$	$\kappa = 10^{-6}$
S_4^1	[1243],[1423],[2134],[4132], [2314],[4312],[3241],[3421]	$\|\lambda_2\| = \|\lambda_3\|$	0.4083	0.4563	0.485	0.4927
S_4^2	[1324],[1342],[3124],[3142], [2413],[4213],[2431],[4231]	λ_2	0.8923	0.9707	0.9913	0.9973
		λ_3	0.4461	0.4853	0.4957	0.4987

Table 5.2: Eigenvalues for permutations $\Pi \in S_4$, with $\Pi \notin S_4^R$, grouped according to their leading non-unit eigenvalues.

derived from Lemma 5.2.4. Therefore the overshoot of τ by the profile of $|\lambda_2|$, also overshoots the upper bound on the diffusion-less mixing rate, τ_{max}.

This deceleration of mixing rate with increasing diffusivity coefficient is counter-intuitive, as it would be expected that fast mixing from diffusion should encourage faster mixing over all. However, the formation of a long-time eigenfunction allows this complicated behaviour to emerge. Since there is no spatial relation between the eigenfuction and the dynamics of the underlying map, the local effect of diffusion on the formation of the eigenfunction is more difficult to understand and could result in unexpected spatial arrangement in the eigenfunction, and as such a persistent pattern in the concentration field with a slow decay rate. The spatial properties of the eigenfunction are discussed further in the next chapter.

In this section S_3 was used as a starting example. The next two sections briefly discuss the results from other permutation groups of interest, S_4 and S_5.

5.4.4 Results for S_4

For the strong mixing results in the diffusion-less limit, S_4 with $m = 2$ does not coincide with the rules from Theorem 5.2.1, and as such it is possible for permutations in S_4 for the composition maps $(\Pi \circ M_B)^{-1}$ to not be topologically mixing.

Table 5.2 lists the modulus of the non-trivial leading eigenvalue λ_2 and λ_3 and shows that similar to S_3, the values of the eigenvalues suggest permutations of S_4 can be grouped according to their long-term mixing behaviour. For $\Pi \in S_4^1$ $\lambda_3 = \overline{\lambda_3}$, while for $\Pi \in S_4^2$ their eigenvalues λ_2 and λ_3 are real and vary in modulus. Permutations S_4^R are neglected since it is understood that their eigenvalues satisfy $|\lambda_k| = 0$ for all k and values of κ.

The first of the two groups, denoted S_4^1, has a leading eigenvalues which appears to converge as as $\kappa \to 0$, to $|\lambda_2| \sim= 0.5$. This coincides with the
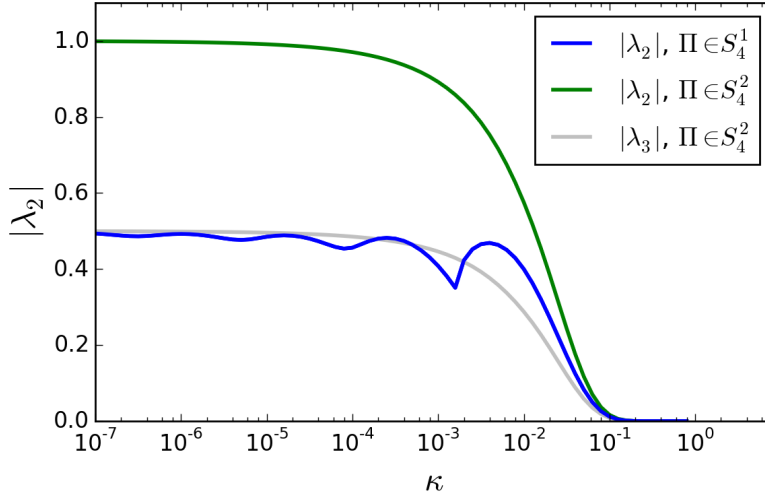
Figure 5.18: The leading eigenvalue $|\lambda_2|$ is plotted for the two groups $\Pi \in S_4^1$ and $\Pi \in S_4^2$ identified from Table 5.2, and additionally $|\lambda_3|$ for S_4^2.

computation of τ in the diffusionless limit, which is found by constructing the Markov transition matrix as described in Section 5.2.2, giving $\tau = 0.5$ for $\Pi \in S_4^1$. This is confirmed in Figure 5.18 where the profile of $|\lambda_2|$ is plotted on a finer resolution of diffusion coefficients κ for a selected permutation from S_4^1. As with the profile of $|\lambda_2|$ when $\Pi \circ M_B$ has permutation $\Pi = [213]$, oscillations occur in the profile, however it never overshoots τ. Oscillations in the non-monotonicity appear to be on the order $\propto \log(\kappa)$. The short and long-time mixing behaviour of the maps $\Pi \circ M_B$ with $\Pi \in S_4^1$ is similar to the interleaving permutations of S_3, where rearrangement of the concentration field by cutting and shuffling causes long-time exponential mixing, independent of initial condition. This is reflected in the evolution of variance profiles for the three different initial conditions in Figure 5.19, represented by the solid lines for $\Pi = [1243] \in S_4^1$.

The permutations $\Pi \in S_2^4$, the final group, have varying behaviour. As $\kappa \to 0$, $|\lambda_2| \to 1$ monotonically, which corresponds to $\tau_{max} = 1$ for S_4 as the diffusion-less strong mixing rate. This is due to the separated ergodic partitions which were shown in Figure 5.4 (a), where no mixing occurred between the two partitions. Instead, only diffusion can act across the boundaries at $x = 0$ and $x = 1/2$. However, note that the second-leading, non-trivial eigenvalues $|\lambda_3| \neq 1$, and instead appears to converge $|\lambda_3| = 0.5$ as $\kappa \to 0$. In Figure 5.18 the profile of $|\lambda_3|$ plotted in grey, converges monotonically to 0.5.

In Figure 5.19 the implication of this result is observed. For the square wave
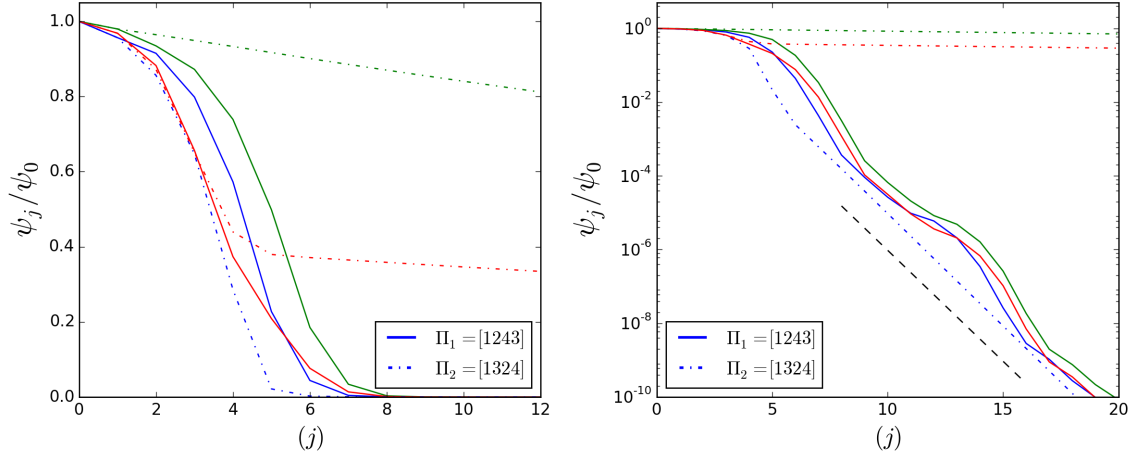
Figure 5.19: The initial transient of variance decay is shown for $\kappa = 10^{-5}$ and the 3 different initial conditions of Figure 5.12 (a) where colours are respective of initial conditions. Different permutations $\Pi_1 \in S_4^1$, and $\Pi_2 \in S_4^2$ result in different long-time behaviour. For $\Pi \in S_4^2$, $\Pi \circ M_B$ is not strong mixing and dramatic change between mixing stages is observed, dependent on the initial condition. Black dashed line plots $\psi_j \propto 0.5^{2j}$ for comparison.

initial condition the map $\Pi \circ M_B$ with $\Pi = [1324]$ does not encourage increased segregation and the long time mixing rate is diffusive only across the boundaries at the edges of the two partitions. Slow mixing occurs from the initial iterations, plotted as the green dotted-dashed line. For the initial condition $c^{(0)}(x) = \cos(2\pi x)$, after the action of the baker's transformation during the first iteration, the swapping of the middle two sub-intervals does not reassemble the initial condition. Instead mixing occurs within the two ergodic partitions and the long-time mixing is exponential, faster than predicted by $|\lambda_2|$. The variance profile decays according to $\psi(j) \propto |\lambda_3|^{2j}$, represented by the dashed black line in Figure 5.19 aligning with the variance profile in dotted-dashed blue. It is important that at each iteration, the mean field within each partition has the value $\bar{c} = 0$, resulting in this long-time mixing behaviour.

For the randomised initial condition, the mixing rate is similar to that of $c^{(0)}(x) = \cos(2\pi x)$ for the first few iterations, where significant mixing occurs within each ergodic partition. This is plotted as the dashed-dotted red line in Figure 5.19. However $\bar{c} \neq 0$ on the partitions and eventually, once significant strong mixing has encouraged the concentration field to approach the mean value within each partition, only diffusion across the partitions' boundaries can cause $c(x) \to \bar{c}$ for the whole domain. There is a sharp transition between the strong

mixing in the initial stages and the slow mixing dominated by diffusion in the later stages.

Table 5.3: Decay rates for permutations $\Pi \in S_5$, grouped according to their leading non-unit eigenvalues

Permutations	$|\lambda_k|$	$\kappa = 10^{-3}$	$\kappa = 10^{-4}$	$\kappa = 10^{-5}$	$\kappa = 10^{-6}$
[12354],[15234],[21345],[51324],[35124],	$|\lambda_2|$	0.4270	0.5589	0.5872	0.5977
[23415],[24513],[45213],[35412],[24351]	$|\lambda_3|$	0.4270	0.5589	0.5872	0.5977
[12435],[13245],[13452],[41235],[51243],	$|\lambda_2|$	0.3459	0.5445	0.5978	0.6318
[41523],[54123],[34152],[32451],[34521]	$|\lambda_3|$	0.3459	0.509	0.56	0.5569
[12534],[51342],[23145],[53124],[35142],	$|\lambda_2|$	0.6486	0.6760	0.6707	0.6856
[42513],[53412],[24531], [42351], [45231]	$|\lambda_3|$	0.6486	0.6760	0.6707	0.6856
[12453],[14235],[13425],[14523],[31245],	$|\lambda_2|$	0.7189	0.6880	0.6079	0.6255
[41253],[51423],[31452],[32154],[34251]	$|\lambda_3|$	0.7189	0.6880	0.6079	0.6255
[12543],[14532],[31542],[51432],[32145],	$|\lambda_2|$	0.6093	0.6707	0.6426	0.5993
[42153],[43125],[43251],[54231],[53421],	$|\lambda_3|$	0.6093	0.6707	0.6426	0.575
[13254],[15243],[14352],[21435],[31524],	$|\lambda_2|$	0.5821	0.6810	0.7223	0.7420
[41325],[24153],[32415],[54213],[35421],	$|\lambda_3|$	0.5821	0.5571	0.5611	0.5393
[14253],[31425],[53142],[25314],[42531]	$|\lambda_2|$	0.8853	0.8720	0.8238	0.7955
	$|\lambda_3|$	0.8853	0.8720	0.8238	0.7955
[13524],[41352],[24135],[52413],[35241]	$|\lambda_2|$	0.6256	0.7422	0.7834	0.7987
	$|\lambda_3|$	0.6256	0.7422	0.7834	0.7987

Table 5.3 – continued from previous page

Permutations	$|\lambda_k|$	$\kappa = 10^{-3}$	$\kappa = 10^{-4}$	$\kappa = 10^{-5}$	$\kappa = 10^{-6}$
[14325],[15423],[21453],[31254],[34215]	$|\lambda_2|$	0.5459	0.5727	0.5878	0.5813
	$|\lambda_3|$	0.5459	0.5727	0.5878	0.5813
[15324],[21354],[24315],[25413],[35214]	$|\lambda_2|$	0.6563	0.7406	0.7867	0.7984
	$|\lambda_3|$	0.4833	0.6041	0.5742	0.5836
[13542],[41532],[42135],[52143],[43152],	$|\lambda_2|$	0.5436	0.5728	0.6084	0.6305
[54132],[32541],[52431],[53241],[43521]	$|\lambda_3|$	0.4513	0.5729	0.5774	0.5395
[15342],[21534],[23154],[25143],[32514],	$|\lambda_2|$	0.6115	0.5928	0.6381	0.6148
[42315],[53214],[54312],[25431],[45321]	$|\lambda_3|$	0.6115	0.5928	0.6381	0.6148
[15432],[21543],[32154],[43215],[54321]	$|\lambda_2|$	0.3694	0.5099	0.4807	0.5124
	$|\lambda_3|$	0.3694	0.5099	0.4807	0.5124
[25134],[23514],[52314],[45312],[25341]	$|\lambda_2|$	0.7719	0.7628	0.7762	0.8023
	$|\lambda_3|$	0.3241	0.4368	0.4484	0.4733
[52134],[45132],[43512],[23541],[52341]	$|\lambda_2|$	0.6522	0.7027	0.6435	0.5911
	$|\lambda_3|$	0.3558	0.5107	0.5518	0.5911
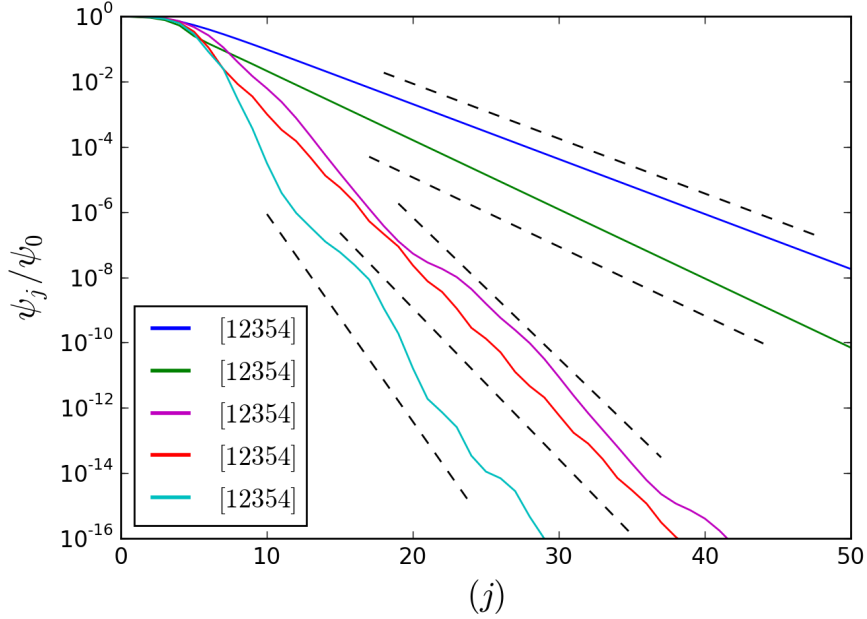
Figure 5.20: Exponential decay of variance is shown for $\kappa = 10^{-5}$ for a number of permutations representing the groupings shown in Table 5.3.

5.4.5 Results for S_5

As a final example, variance profiles and eigenvalues for the composition maps $\Pi \circ M_B$ with permutations from S_5 were computed. In Figure 5.20, profiles for the variance decay are plotted for the same initial condition, evolved according to $\Pi \circ M_B$, for a handful of permutations $\Pi \in S_5$. There is significant variation in the exponential rates of mixing depending on the choice of Π. Oscillations occur in a few of the profiles but average mixing rates are predicted well by the eigenvalues of each transformation.

For the maps $\Pi \circ M_B$ with diffusion, permutations from S_3 and S_4 fell into two and three groups respectively. For the permutation group S_5, with has 120 elements, Table 5.3 lists the modulus of the leading and second non-trivial eigenvalues $|\lambda_2|$ and $|\lambda_3|$ for 4 different diffusion coefficients. In total 16 groups emerge consisting of 5 or 10 permutations giving the same asymptotic mixing rates. The rotation permutations are neglected in the table since $\lambda_k = 0$ for all κ. Similar to S_4, in a number of groups there is a discrepancy in value between $|\lambda_2|$ and $|\lambda_3|$, which occurs when λ_2 is real, while $|\lambda_2| = |\lambda_3|$ corresponds to $\lambda_2 = \overline{\lambda_3}$. None of
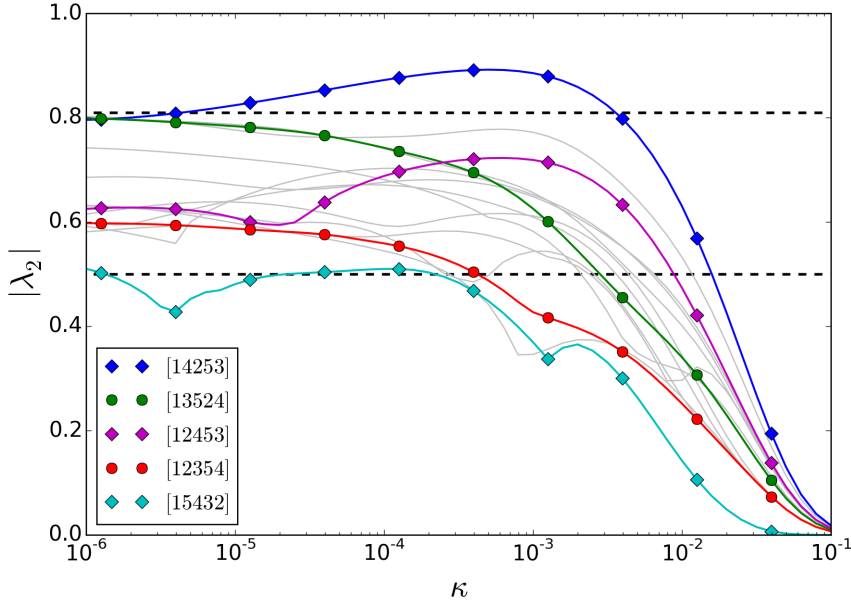
Figure 5.21: Profiles of the exponential mixing rate for $\Pi \in S_5$ given by the leading eigenvalue $|\lambda_2|$. All profiles are shown in *lightgrey* showing a range of complicated behaviour. A number of representative permutations are highlighted in *red* and *blue* as examples to varying behaviour.

the eigenvalues tend to 1 because for $m = 2$ and $N = 5$ the conditions of Theorem 5.2.1 are satisfied and all maps $(\Pi \circ M_B)^{-1}$ are strong mixing.

It is not intuitive from the permutations alone what the groups should be and they are dependent on the full composition of the transformation. When computing the eigenvalues for another inverse-expanding map, say $\Pi \circ f^{-1}$ with $m = 3$, again 16 groups emerge but there is disagreement between the groupings of $m = 2$, S_5 and $m = 3$, S_5.

Figure 5.21 plots the profiles of $|\lambda_2|$ with κ for a selected permutation from each group. Five profiles have been highlighted for discussion, however all are shown in *light-grey* revealing the complicated behaviour that emerges across the 16 profiles. Many of the profiles change monotonically with diffusion, and tend to a constant mixing rate in the limit of $\kappa \to 0$. Example profiles of this behaviour are shown with circled lines. Additionally, non-monotonic like profiles emerge, highlighted by the diamond lines. For the permutation $\Pi = [14253]$, $|\lambda_2|$ appears to convergence as $\kappa \to 0$, however there is an overshoot of this value when κ is large. The other two highlighted profiles have non-monotonic mixing rates, permutations $\Pi = [15432]$ and $\Pi = [12453]$. Repeated decreases and increases in

154

	[12354]	[12434]	[12534]	[12453]	[12543]	[13254]	[14253]	[13524]
$\|\lambda_2^{10^{-7}}\|$	0.6042	0.6444	0.673	0.5878	0.6013	0.7493	0.7996	0.8046
$\|\lambda_2^{10^{-8}}\|$	0.6061	0.6518	0.6641	0.6235	0.5904	0.753	0.8075	0.807
Markov τ	0.5919	0.6624	0.6677	0.5755	0.5755	0.7564	0.809	0.809

	[14325]	[15324]	[13542]	[15342]	[15432]	[25134]	[52134]	
$\|\lambda_2^{10^{-7}}\|$	0.5751	0.8048	0.6443	0.6086	0.5089	0.8075	0.6156	
$\|\lambda_2^{10^{-8}}\|$	0.561	0.8072	0.6531	0.5926	0.4788	0.807	0.5808	
Markov τ	0.5	0.809	0.6624	0.5919	0.5	0.809	0.5	

Table 5.4:   For each of the subgroups, the modulus of the second leading eigenvalues for low values of the diffusion coefficient, denoted $|\lambda_2^\kappa|$, are computed. A comparison with τ, the mixing rate from the respective Markov transition matrices of the pre-image mapping $(\Pi \circ M_B)^{-1}$, shows good agreement.
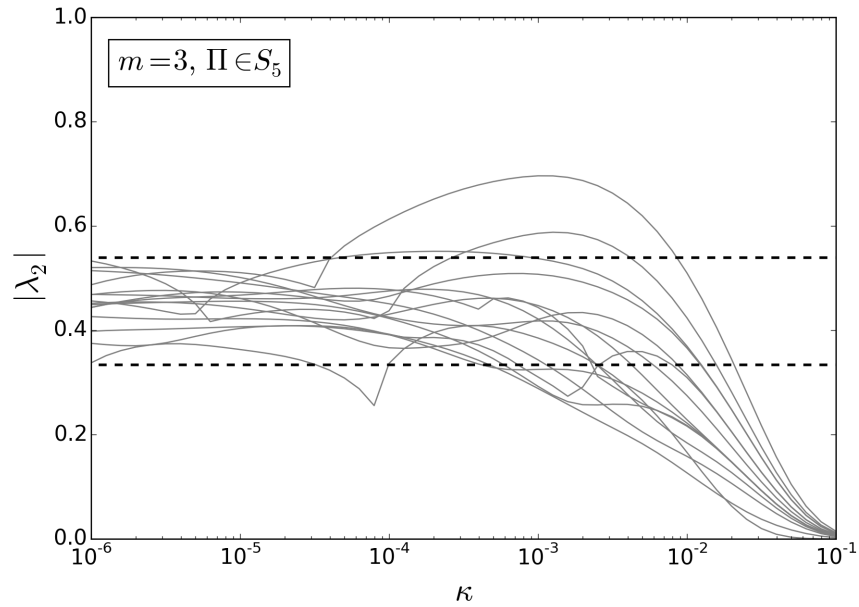
Figure 5.22: Profiles of $|\lambda_2|$ for the transformations $\Pi \circ f^{-1}$ with $m = 3$ and $\Pi \in S_5$. Non-monotonicity is observed in many of the profiles with overshooting of the upper bound on the value of τ achieved by a few.

value occur with no obvious convergence to a fixed mixing rate as κ decreases.

Table 5.4 compares τ for the pre-image maps $(\Pi \circ M_B)^{-1}$ with $|\lambda_2|$ for the respective diffusive transfer matrices in the limit of small diffusivity. The values show good agreement for many of the permutations, however this is not the case for all. The permutations in which there is not good agreement correspond to the profiles of $|\lambda_2|$ with show significant non-monotonic variation with κ. It is not known whether the profiles will converge or continue oscillating as $\kappa \to 0$, however investigating smaller values of κ is computationally infeasible.

The lower and upper bounds on the value of τ for $(\Pi \circ M_B)^{-1}$ when $\Pi \in S_5$ are

$$\tau_{min} = \frac{1}{2}, \quad \tau_{max} = \frac{\sin(2\pi/5)}{2\sin(\pi/5)} \approx 0.809, \tag{5.29}$$

respectively. These are plotted as the black dashed-line in Figure 5.21, and seem to be reasonable bounds on the asymptotic mixing rates when the diffusivity coefficient is small. However, at large values of κ the upper bound τ_{max} does not agree as a bound for all $|\lambda_2|$, due to the non-monotonicity with κ. Table 5.3 also confirms that for many of the permutations, $|\lambda_2^\kappa| > \tau_\Pi$. These results persist for $\Pi \circ f^{-1}$ with $m = 3$ and $\Pi \in S_5$, shown in Figure 5.22, where non-monotonicity leads to values of $|\lambda_2| > \tau_{max}$ for non-zero κ.

If the subgroups were listed by the value of $|\lambda_2|$, then it is apparent from Figure 5.21 and Figure 5.22 that the ordering of permutations would be dependent on the value of κ. This suggests that for the maps $\Pi \circ M_B$ where $\Pi \notin S_N^R$, the best mixing protocol, that which achieves a mixed condition in the fastest time, may be dependent on the rate of diffusion. This has been similarly reported for the mixing of permutations of equal sized cells Ashwin *et al.* (2002). However, the values of $|\lambda_2|$ give the long-time mixing rate of the system and may have little affect in the early stages where significant mixing can occur. This is addressed in the next section.

5.4.6 Time to achieve a mixed condition

For all of the permutations groups considered, many presented non-monotonic variations with the asymptotic rate of mixing when varying the diffusivity coefficient. In several of the cases, for $\kappa > 0$, $|\lambda_2| > \tau$, where τ is the diffusion-less, strong mixing rate of the map. This result is interesting itself as it is counter-intuitive in the expectations of increasing the rate of diffusivity at each iteration. However, this result is for the long-time mixing. Mixing of a concentration field
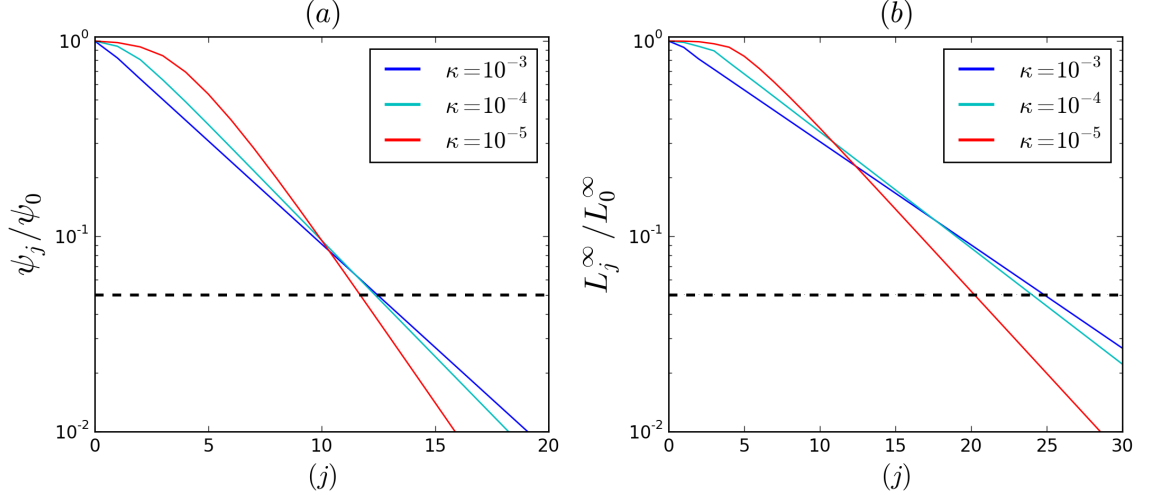
Figure 5.23: (a) Variance decay ψ_j/ψ_0 and (b) L^∞ norm L_j^∞/L_0^∞ are plotted for the map $\Pi \circ M_B$ with $\Pi = [14253]$ evolving the $c^{(0)} = \cos(2\pi x)$ for three different diffusivity rates. 95% mixed condition is plotted as a dashed black line, revealing that the slowest diffusivity rate $\kappa = 10^{-5}$ crosses the threshold at the earliest time.

has several stages which contribute to the finite time mixing. In practical situations one would want to consider the time needed to mix to a desired condition. Hence, a return to studying the time to achieve a mixed condition is necessary to understand the full implications of these results.

In addition to taking the decay of variance as a quantifier for the amount of mixing, the definition 2.44 is generalised to consider any norm. In addition to the variance (L^2 norm), the L^∞-norm (2.39) is used as an additional measure of mixing to determine the time to achieve a mixed condition.

In smoothly deforming chaotic flows, during the initial stages of advection and diffusion, when the effect of diffusivity is small, an L^q norm will remain constant for some time until the gradients or length scales in the concentration field are on the order of the Batchelor length scale $\sqrt{\kappa/h_{ls}}$, the balance between the diffusive rate and local strain rate h_{ls}. In the limit of small diffusivity the exponential mixing rate becomes independent of the diffusivity coefficient, thus the main consequence of decreasing the effect of diffusion is to extend the initial transient where a measured norm is close to constant. Recall that this initial stage is weakly dependent on diffusivity, on the order of $\log(\kappa)$.

The systems $\Pi \circ M_B$ presented in this chapter are discontinuously deforming with chaotic stretching rates. It has already been observed that the asymptotic, exponential mixing rate does not always become independent of the diffusivity

rate as $\kappa \to 0$, and deceleration with increasing κ is common. The numerical results for the mixing of $c^{(0)}(x) = \cos(2\pi x)$, by the mapping $\Pi \circ M_B$ with $\Pi = [14253]$ are plotted in Figure 5.23 for three different diffusivity coefficients. Recall that this permutation from S_5 had $|\lambda_2| > \tau_{max}$ for non-zero κ in Figure 5.21. For both (a) the decay of variance and (b) the decay of the L^∞ norm, the time to achieve a mixed condition is first achieved when the diffusivity coefficient is $\kappa = 10^{-5}$, the smallest diffusivity value considered. For the decay of variance, the condition is achieved about 1 iteration before $\kappa = 10^{-3}$ and 10^{-4}. For the L^∞ norm the condition is reached approximately five iterations earlier. In applications of fluid mixing devices this would equate to five additional stirring periods or channel segments to achieve the desired result, which could be overlooked if approximating mixing time from advective properties only and thus the correct mixing criteria not achieved in the predicted time.

As with chaotic mixing systems, it is observed in Figure 5.23 that changing κ prolongs any significant depletion an L^q norm in the first initial iterates, which appears as though it may follows a similar $\log(\kappa)$ relation. It is not surprising that this is the case, since the dominant mechanism of stirring in the system is still from an exponential reduction in lengths scale, therefore a logarithm variation in the value of κ has only a linear effect on the time needed to achieve the length scales where significant depletion takes place. However, this example illustrates that even with such a weak relation from the diffusivity coefficient on the initial stages of mixing, the counter-intuitive deceleration of the asymptotic mixing rate with diffusivity coefficient could still have a notable effect on the time to achieve a practical mixing condition, such as within 5% of uniform.

5.5   Conclusions

The results from a one-dimensional model which captures a mixture of stretching and discontinuous advective dynamics with a diffusive step, are presented. Three permutation groups S_3, S_4 and S_5, were investigated for the discontinuous step and similarities in the results persisted across the groups.

In the initial transient, permutations in which sub-intervals are swapped (all but rotations), introduced sharp interfaces into the concentration field and these increases in gradients sped up the decay of variance over the purely chaotic advection of the baker's transformation. However, once the Batchelor scale is reached and variance begins to deplete significantly, a phase of exponential decay emerges

due to non-uniformity caused by the permutation of the striation arrangement. Diffusive and reactive systems have been shown previously to be sensitive to striation arrangement Clifford *et al.* (1999). Lagrangian arguments fail to explain the varying decay rates across the range of permutations since stretching rates are the same almost everywhere except where they are undefined on the discontinuities which form a set of measure zero. The mechanism for emergence of an exponential asymptotic mixing rate is global.

Rather than a direct cascade to large wavenumbers which occurs for the baker's transformation alone, the amplitude in a wavemode of the discrete Fourier expansion is dispersed to many wavemodes at each iteration. This is comparable to dispersion due to non-uniformity in stretching rates reported by Fereday *et al.* (2002) and Wonhas & Vassilicos (2002). However, the mechanism causing the dispersion is different. The slowest decaying eigenfunction of the transfer operator, and respective eigenvalue, are computed from the transfer matrices representing the linear operator and the eigenvalue predicts average long-time mixing rate of the system.

As $\kappa \to 0$, it is well reported for smoothly deforming chaotic maps that $|\lambda_2| \to \tau$, where τ is an isolated discrete eigenvalue for the advective transfer operator. For systems where the asymptotic mixing mechanism is global, this convergence is

$$|\tau - \lambda_2| \propto \kappa^\eta \tag{5.30}$$

with $0 < \eta < 1$ Haynes & Vanneste (2005). For the discontinuous systems $\Pi \circ M_B$ the mixing rate in the diffusion-less limit is well approximated by the strong mixing rate τ, predicted from Markov partitions of the pre-image map.

However, this approach to the mixing rate in the zero-diffusivity limit is non-monotonic in many cases, predicting a deceleration with increasing diffusivity coefficient, which is counter-intuitive. Non-monotonicity in profiles of $|\lambda_2|$ with κ have been observed in one-dimensional maps before; for an expanding map with three branches where one branch is inverted Eckhardt *et al.* (2003), and a non-uniform inverted baker's transformations with a no-flux boundary condition Gilbert (2006). However the oscillations in both of these maps occur as $|\lambda_2|$ converges to τ, the latter shown to have a polynomial power-law on average, in line with (5.30). Oscillations in $|\lambda_2|$ occur for the maps $\Pi \circ M_B$, but where convergence to τ happens for some, this is not always the case for others.

All maps which report non-monotonicity in the mixing rate with diffusivity have a common property in that they contain points which are non-differentiable.

A sensible hypothesis would be that this deceleration of mixing rate is a feature of non-continuous mappings, in which discontinuous transformations are a subset. Explaining the non-monotonicity is beyond the scope of this work.

In the dynamical systems and ergodic theory literature there is great interest in finding global mixing rates of maps, but these computational results suggest that in studying fluid mixing systems in which there is a mixture of stretching and folding and cutting and shuffling, diffusion may have to be taken into account for accurate mixing rate predictions and comparisons across mixing protocols.

For example, designing parameter configurations for periodic and non-periodic stirring protocols has been successfully completed using Ulam's method as a crude approximation for the advective transfer operator [Kang *et al.* (2008); Schlick *et al.* (2015); Singh *et al.* (2008a)]. A transfer matrix is constructed using grid-based partitions and computing the ratio of the volume in partition element mapped to all others when advected by the flow. Numerical diffusion is significant in these grid based methods, but it has been argued by some to be an a possible approximation for an applied diffusion equation [Schlick *et al.* (2013)], although the rate of this numerical diffusion is uncontrolled and will be variable throughout the domain. The counter-intuitive deceleration with diffusion observed herein suggests that in models with discontinuous stirring, under-approximation of stirring times to a desired condition might occur if the rate of diffusivity is too large. This is likely to occur in these models where Ulam's methods are used to approximate the transfer of concentration where numerical diffusion increases the diffusion of a concentration field at each step. The significance of this is shown to effect mixing rates in finite time considerations to achieve physical mixing conditions, chosen arbitrarily to be a 95% mixed state.

Additionally, behaviour was presented when the transformation was not strong mixing across the whole domain, but instead created two ergodic partitions in which the dynamics on each are strong mixing. For two of the reported initial conditions, the long-time mixing rate was slow, limited to diffusion across the boundaries between the partitions. The initial decay of variance could be fast, and no transition to parabolic like mixing rates emerged between the two stages of mixing. This contrasts what is observed in chaotic flows with islands or weakly-connected, chaotic partitions where non-exponential, parabolic decay of variance described the main stage of mixing Popovych *et al.* (2007). The difference here is because the transformation does not contain parabolic points, which are known

to contaminate mixing rates [Gouillart et $al.$ (2007, 2008)], and instead the dynamics are non-Hamiltonian in the stability arrangement of periodic points in the flow. These ergodic partitions are comparable to the pseudo-elliptic points which contaminated mixing in cut-shear-shear flow of Smith et $al.$ (2016). If the partitions are denoted as B_1 and B_2, for specially chosen initial conditions in which $\bar{c}|_{B_i} = \bar{c}|_{[0,1)]}$, $i = 1, 2$, the long-time mixing is exponential given by $|\lambda_3|$. However, for almost all initial conditions diffusion-limited mixing between partitions would dominate the asymptotic mixing rate.

There are many extensions which could be considered to this model. One of the most obvious would be to consider extending the permutations of equal size to Interval Exchange Transformations of the last Chapter. Recall that IETs with rational sub-intervals can be approximated as larger permutations of equal sized cells. Consider the upper bound function on non-diffusive mixing rates, which has the property that

$$\lim_{N \to \infty} \frac{\sin(m\pi/N)}{m\sin(\pi/N)} \to 1. \qquad (5.31)$$

Hence, one implication of considering IETs is that they are more likely to lead to slower mixing rates as an upper bound. However, in Zhang (2012), the distribution of τ_Π was investigated in the limit as $N \to \infty$ and probability distribution functions on the spread of τ appeared to converge to an average mixing rate. They contemplated whether convergence on an average mixing rate would emerge;

$$\lim_{N \to \infty} \frac{1}{N!} \sum_{\Pi \in S_N} \tau_\Pi. \qquad (5.32)$$

Bordenave et $al.$ (2018) found that although a permutation can arbitrarily slow down the mixing rate in the limit as $N \to \infty$, a typical permutation is more likely to have $\tau = 1/\sqrt{m}$, which for $m = 2$ has $\tau_{typical} = 2/\sqrt{2} \approx 0.7071$. In light of these results, it can not be readily seen why extending the study to interval exchange transformations will reveal anything of interest than is already reported herein. This is a similar conclusion to extending the study to tent maps with various inversions of uniform expanding branches where Byott et $al.$ (2016) has already investigated values of τ and distributions in the unit-circle.

Another obvious extension would be to consider the non-uniform baker's transformation [Fereday et $al.$ (2002); Gilbert (2006); Wonhas & Vassilicos (2002)] composed with cutting and shuffling transformations. Since the non-uniformity in stretching rates itself produces eigenfunctions of the transfer operators with a

long-time exponential mixing rate in the presence of diffusion, it is worth considering if permutations help or hinder the mixing rates. This is the topic of the next Chapter.

Chapter 6

Mixing by piecewise, non-uniformly expanding maps

In the previous chapter, introducing discontinuous transformations to uniformly stretching transformations resulted in a deceleration of the asymptotic mixing rate. Rather than super-exponential mixing as expected in the baker's transformation with diffusion, exponential mixing emerged due to the formation of eigenfunctions. These arose due to the non-uniformity in the concentration field evolution, caused by the cutting and rearranging of the concentration field from the permutation transformations. Other studies considering uniform stretching with discontinuous transformations concluded that they generally stir worse than uniform stretching alone [Byott *et al.* (2013); Smith *et al.* (2016, 2017b)]. However, it is unlikely in real life mixing situations for such strict uniformity in stretching rates to exist. Non-uniformly stretching systems exhibit the same long-time mixing phenomena where persistent patterns with exponential decay rates emerge in the concentration field evolution. These eigenfunctions arise from a different mechanism in generating non-uniformity in striation arrangement, from the varying compression rates of non-uniform stretching and folding.

In this chapter, transformations which are a composition of non-uniform stretching and discontinuous transformations are discussed. In line with a number of studies interested in non-uniform stirring [Fereday *et al.* (2002),Wonhas & Vassilicos (2002)], the non-uniform baker's transformation is used to investigate the composition of mixing dynamics in one-dimension. The discontinuous transformations used are permutations of equal sized cells.

It is found that little improvement on asymptotic mixing rates can be achieved with permutation transformations which cut and shuffle. Improvement in finite
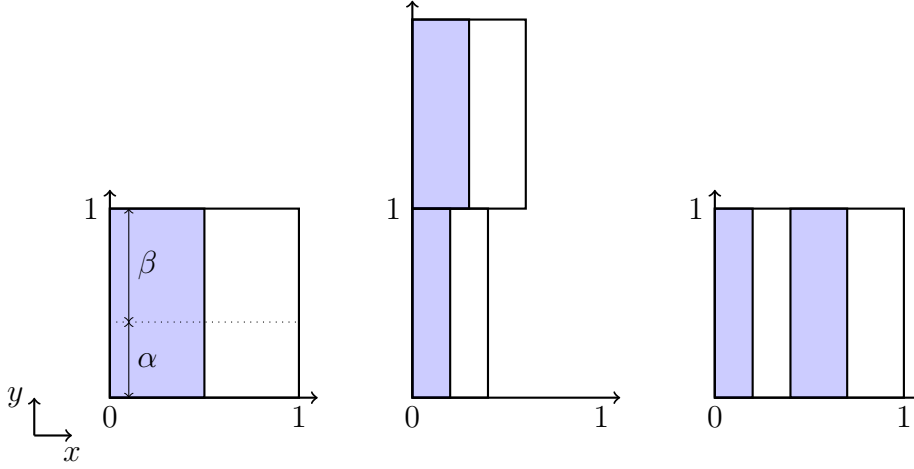
Figure 6.1: Illustration of the non-uniform baker's transformation. The unit square is divided horizontally into two rectangles with areas α and β. These are stretched in the y direction by factors α^{-1} and β^{-1} respectively, and then reassembled onto the unit square by bringing the top rectangle to the right of the lower rectangle.

time mixing is only observed in the initial stages, where sharp gradients introduced by the discontinuities increase depletion of variance. In the later stages, the global effects of rearrangement contaminate the mixing with the possibility for significant reduction in mixing rates. As well as the potential to cut the concentration field into smaller striations, the permutations allow for the possibility of striation rejoining, increasing the scales in the concentration field at each iteration. This is the same mechanism which contributes to the contamination of mixing in Chapter 5, however the comparison between the non-uniformity arising from stretching alone, highlights the differences between the mechanisms.

When the non-uniformity in the stretching rates is particularly large, discontinuous rearrangement does improve the asymptotic mixing rate when diffusivity is particularly fast, but shows little improvement in the diffusion-less limit. Rotation permutations on the other hand, which are not discontinuous but rearrange the stretching histories of points in the domain, can improve the mixing rate. The qualitative differences in the mixing behaviours between continuous non-uniform stirring and discontinuous non-uniform stirring are discussed. By looking in detail at the evolution of the concentration field and resulting eigenfunctions, the reason for the slow long-time mixing becomes clear. The predictability in the appearance of the eigenfunctions in relation to the underlying dynamics of the composition transformations is investigated.

6.1 Non-uniform baker's transformation

6.1.1 Formulation of the map

The non-uniform, area-preserving baker's transformation is a paradigm model of a non-uniformly mixing transformation. The unit square is divided into two rectangles of area α and β by a cut parallel to the x-axis. Choosing $\alpha + \beta = 1$ ensures incompressibility and without loss of generality the convention $\alpha < \beta$ is taken. The two rectangles are stretched in the y direction by factors α^{-1} and β^{-1} respectively, and then reassembled on the unit square. The action on the points $(x, y) \in \mathbb{T}^2$ is

$$M_B^\alpha(x, y) = \begin{cases} (\alpha x, y/\alpha) & y \in [0, \alpha) \\ (\beta x + \alpha, [y - \alpha]/\beta) & y \in [\alpha, 1). \end{cases} \tag{6.1}$$

Figure 6.1 illustrates one iteration of the transformation. Taking a y-independent initial concentration field, the non-uniform baker's transformation reduces to a one-dimensional map which is one-to-two, and the action on the concentration field given by

$$c(x) = \begin{cases} c\left(\frac{x}{\alpha}\right), & x \in [0, \alpha), \\ c\left(\frac{x - \alpha}{\beta}\right), & x \in [\alpha, 1). \end{cases} \tag{6.2}$$

This produces a simple one-dimensional mapping in which on repeated iteration the number of striations increases exponentially, akin to the uniform baker's transformation, but now the contraction rate of striations is no longer uniform. The action of the map on points x is illustrated in Figure 6.2, where the fixed point $x_f = 1$ is highlighted.

As in previous chapters, a diffusive step can be easily incorporated by applying periodic boundary conditions and representing the concentration field by its discrete Fourier expansion. A transfer matrix $\boldsymbol{B}^\alpha$ is derived which describes the changes to the Fourier coefficients by the action of the map. Due to the non-uniformity of the map, the transfer of concentration to the Fourier modes is more complicated than the uniform case in Section 5.3. For one-dimensional non-uniform baker's transformation of (6.2), the transfer matrix is computed for each pair of modes kq by the following. If $q = \alpha k$, $\boldsymbol{B}_{kq}^\alpha = \alpha$ and similarly, if $q = \beta k$, $\boldsymbol{B}_{kq}^\alpha = \beta$. These represent resonances in the map. Otherwise,

$$\boldsymbol{B}_{kq}^\alpha = \frac{\sin(k\pi\alpha)}{\pi} \frac{(\beta - \alpha)q e^{-ki\pi\alpha}}{(q - \alpha k)(q - \beta k)}. \tag{6.3}$$
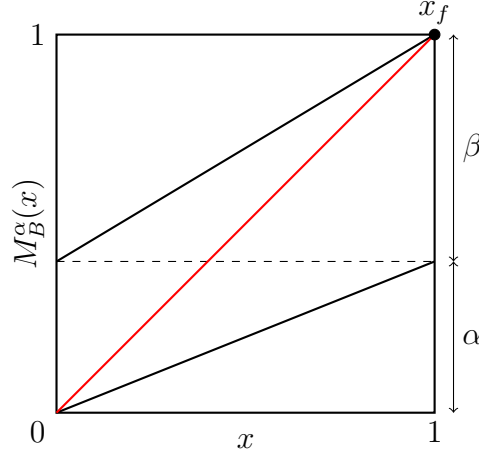
Figure 6.2:   Illustration of the one-dimensional mapping for the non-uniform baker's transformation as shown in Figure 6.1 when taking a y-independent initial condition. The red line represents $M_B^\alpha(x) = x$ and x_f marks the non-trivial fixed point $x = 1$, which lies on the β branch of least compression.

As with the transfer matrices of earlier chapters, it is true that $\boldsymbol{B}^\alpha_{-k-q} = \overline{\boldsymbol{B}^\alpha}_{kq}$, and the matrix preserves reality an the initial condition. The matrix (6.3) was given by Fereday *et al.* (2002) and Wonhas & Vassilicos (2002), but the full derivation is listed in Appendix B. Note the special cases when $\alpha = 1/2$, which derives just the uniform baker's transformation, while $\alpha = 0$ results in no stretching and the action of pure diffusion only. The intermediate stages of interest are where $\alpha \in (0, 1/2)$.

Before the main investigation into the composition of non-uniform stretching with permutations, a short review is presented, with references to the relevant literature, on the results of mixing in the non-uniform baker's transformation important to the discussion.

6.1.2   Mixing rates with diffusion

Figure 6.3 a) shows the decay of variance for the non-uniform baker's transformation with $\alpha = 0.2$ and a range of diffusivity rates. On decreasing the value of the diffusion coefficient κ, independence of the exponential mixing rate emerges. Although the profiles feature some oscillation, the mixing rate given by $|\lambda_2|^{2j}$ approximates the average exponential decay of variance. The relation $\psi_j/\psi_0 \propto |\lambda_2|^{2j}$ is plotted as the dashed black line for $\kappa = 10^{-5}$ showing good agreement for the variance decay when $\kappa = 10^{-6}$ also. This is because as $\kappa \to 0$, $|\lambda_2|$ becomes independent of diffusivity. The values of $|\lambda_2|$ varying with diffusion coefficient κ

(a) Decay of variance for $\alpha = 0.2$ (b) Leading eigenvalue profiles for different α.
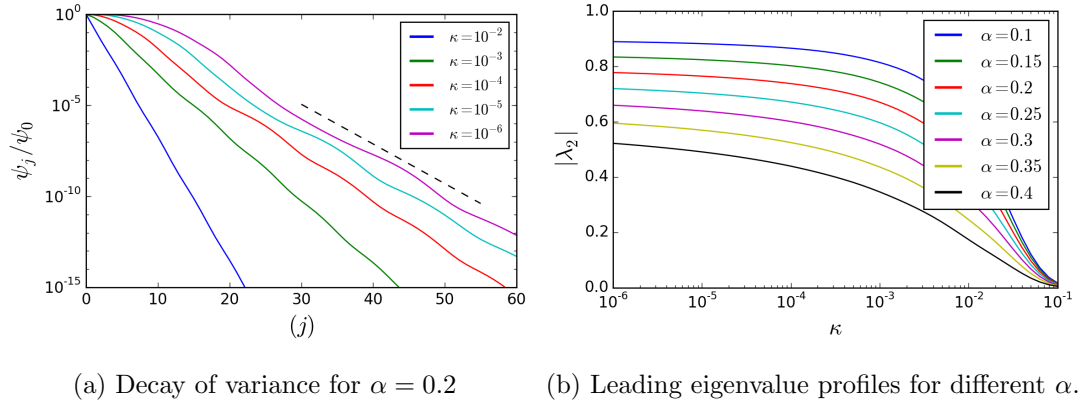
Figure 6.3: Mixing properties of the non-uniform baker's transformation are shown from a) decay of variance when $\alpha = 0.2$ and b) the leading eigenvalues $|\lambda_2|$ for a selection of α values with varying diffusivities κ. In b) the profiles of $|\lambda_2|$ increase monotonically with decreasing κ and become independent in the zero limit. This is reflected in a) where the dashed line representing $\psi \propto |\lambda_2|^{2j}$, with $|\lambda_2|$ when $\kappa = 10^{-5}$, shows agreement with ψ_j/ψ_0 with $\kappa = 10^{-6}$. The initial condition was $c(x) = \cos(2\pi x)$.

is shown in Figure 6.3 b). For all values of α the profiles change monotonically, and each becomes independent of diffusivity coefficient in the zero-limit. In the diffusion-less limit, Gilbert (2006) gives the leading eigenvalues to be complex conjugate pairs with

$$|\lambda| = \max(\alpha, \beta) + o(\kappa), \tag{6.4}$$

and states that as $\kappa \to 0$, convergence to $|\lambda| = \max(\alpha, \beta)$ is monotonic and slow, only logarithmic in κ.

Figure 6.4 shows the resulting eigenfunctions for $\kappa = 10^{-5}$ and different values of α. Each of the eigenfunctions is normalised but plotted without scale, and the value of the mean field $c(x) = \bar{c}$ included as a dashed black line for reference. For large α, when the non-uniformity in the compression rates of the map does not differ too much, there are large fluctuations in the eigenfunction across the domain. Peaks occur in the eigenfunction and align with images of the fixed point $x_f = 1$ where the least compression takes place at each iteration, plotted as red dashed lines for $\alpha = 0.4$. As $\alpha \to 0.1$, where the non-uniformity in the stretching rates is more diverse, images of $x = 1$ are more pronounced in the concentration field. The regions between these peaks are seemingly better mixed. For $\alpha = 0.1$ the eigenfunction has little variation in values across the domain except for the distinctive peak around $x \sim 1$. All the eigenfunctions vary in time on further iteration of the map, due to the complex value of the second leading eigenvalue
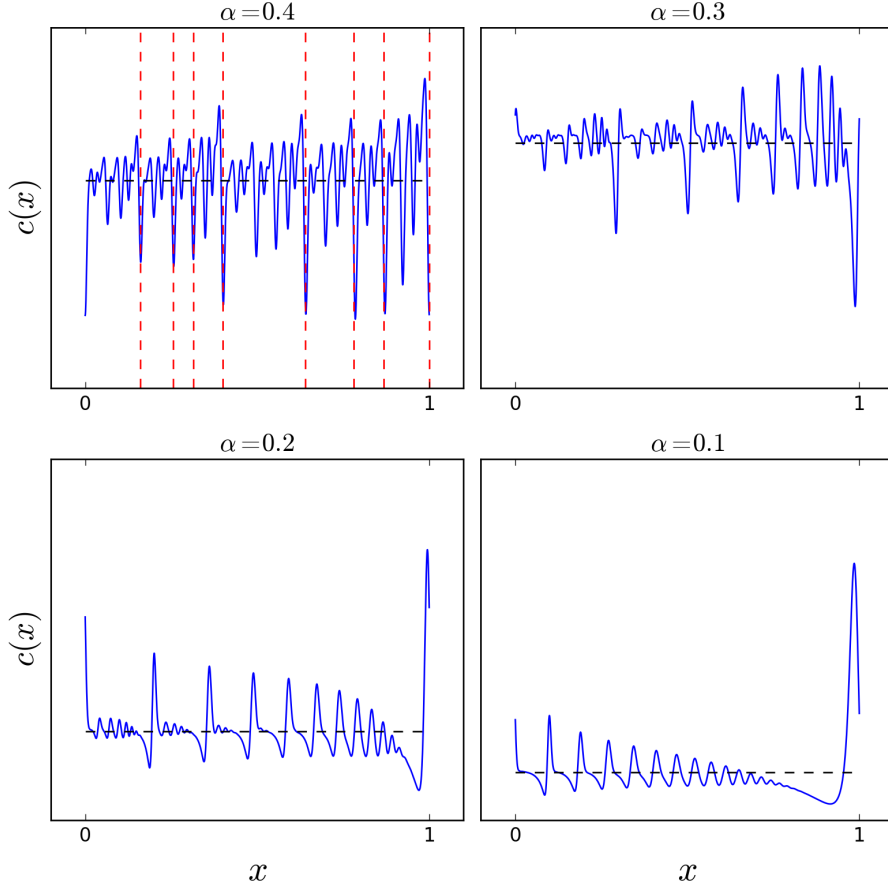
Figure 6.4: The dominant eigenfunction v_2, corresponding to the eigenvalue λ_2 of the diffusive-advective transfer matrix $\boldsymbol{dB}^\alpha$, plotted as a scalar field solution for different values of α and $\kappa = 10^{-5}$. The eigenfunctions were obtained in Fourier space and transformed to real space using a numerical inverse Fourier transform and the dashed black line in each represents $\bar{c}$. When $\alpha = 0.4$, the images of the fixed point $x_f = 1$ are plotted as dashed red lines. The fixed point corresponds to the place where least compression takes place, and its images align with peaks in the concentration field.

which causes a rotational shift in the Fourier coefficients of the concentration field.

The non-uniform baker's transformation of the form (6.1) was used as a paradigm example to prove that global effects can contribute to the long-time exponential mixing rate in chaotic-diffusive maps. The respective literature reviewed in Section 2.5.1.

6.2   Composition with permutations

In line with the previous chapter, the simplest discontinuous transformation in one-dimension are permutations of equal sized cells, such that a complete stirring step is the composition $\Pi \circ M_B^\alpha$. The transfer matrix for the transformations $\Pi \circ M_B^\alpha$ in discrete Fourier space is

$$\boldsymbol{T} = \boldsymbol{M}\boldsymbol{B}^\alpha \tag{6.5}$$

where $\boldsymbol{M}$ is given by (4.12) and $\boldsymbol{B}^\alpha$ given by (6.3), such that one full advective-diffusive step is

$$\hat{c}_k^{j+1} = \sum_{q=-\infty}^{\infty} \boldsymbol{d}_{km}\boldsymbol{T}_{mq}\hat{c}_q^{(j)} \tag{6.6}$$

with d_{km} as defined in (4.13). As in Section 5.3.2, finite truncation to $Q = 1000$ suffices for diffusivity values $\kappa \geq 10^{-6}$. A resolution test was carried out for confirmation, however it is omitted for brevity. The decay of variance ψ_j/ψ_0 is used to measure the mixing for given initial conditions $c^{(0)}(x)$, and eigenvalues and eigenfunctions of the transfer matrices $\boldsymbol{dT}$ computed using the Numpy and Scipy packages in Python. The initial condition taken in every computation was $c(x,y) = \cos(2\pi x)$, however the main results discussed refer to the asymptotic mixing rates and behaviours, which are independent of the initial condition.

First the stirring dynamics in the absence of diffusion are discussed, since for certain parameters a decomposition of the unit interval is possible resulting in separate ergodic partitions, similar to those observed for S_4 is Section 5.4.4.

6.2.1   Decomposition from cutting and shuffling

Depending on the choice of parameters, Π and α, a decomposition into two or more ergodic components can occur. Take for example $\alpha = 1/3$ and $\Pi = [213]$, where Figure 6.5 a) illustrates the action of the one-dimensional non-uniform
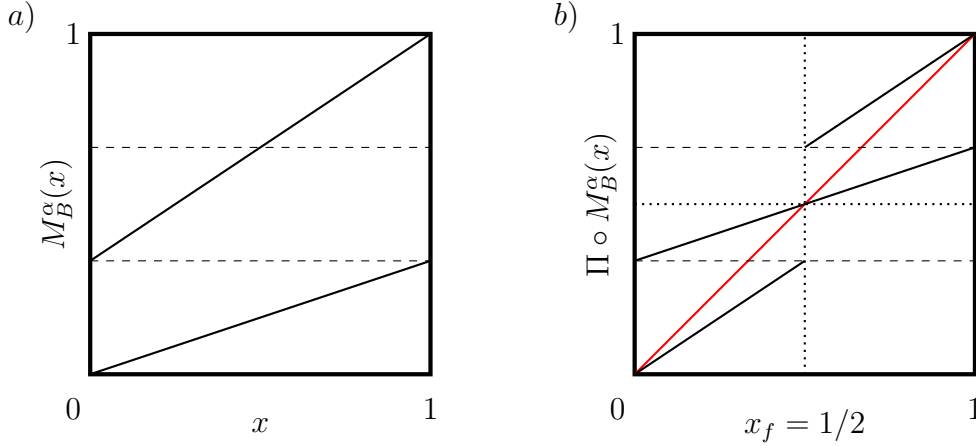
Figure 6.5:   Illustrations of a) the one-dimensional mapping of the non-uniform baker's transformation with $\alpha = 1/3$ and b) the mapping under the composition $\Pi \circ M_B^\alpha$ with $\Pi = [213]$.   The red line represents $\Pi \circ M_B^\alpha(x) = x$, and a fixed point emerges at $x_f = 1/2$.   Note that the pre-image of the cut locations $x = 1/3$ and $x = 2/3$, represented by the dashed black lines, is $x = 1/2$.

baker's transformation and b) the accompanying transformation when composed with the permutation. It is highlighted that for these parameters a decomposition of the map into two halves, $[0, 1/2)$ and $[1/2, 1)$, occurs. Note that the fixed point $x_f = 1/2$ aligns with the pre-image of the cut locations $x_c = 1/3$ and $x_c = 2/3$.

Unfortunately the analytical proof of Byott *et al.* (2013) used to derive Theorem 5.2.1 that states which parameters in expanding maps composed with permutations result in a decomposition, can not be directly applied. There, combinatorial arguments were used to prove which permutations preserve topological mixing. The non-uniform baker's transformation introduces additional complexity by having two branches with different expansion and compression rates. Instead, herein a computational study reveals which parameters might result in a decomposition and a conjecture made on the attributing criteria. Similarly, an analytical approach to find the strong mixing rates in the diffusion-less limit would be non-trivial and as such is not investigated here.

An easy approach to determine if a decomposition is likely to occur is to numerically approximate the transfer matrix $\boldsymbol{T}$ (6.5) for the advective step. A decomposition will produce a transfer operator with a leading eigenvalue $|\lambda_2| \sim 1$. Numerical diffusion will be present due to the finite truncation of the transfer matrix, however the effect will be negligible in determining if $|\lambda_2| \sim 1$. If the decomposition is large in scale, such as the halving for the domain for $\alpha = 1/3$,

(a) Eigenvalues for $\Pi \in S_4$ (b) Map $\Pi \circ M_B^\alpha$ with $\Pi = [3214]$
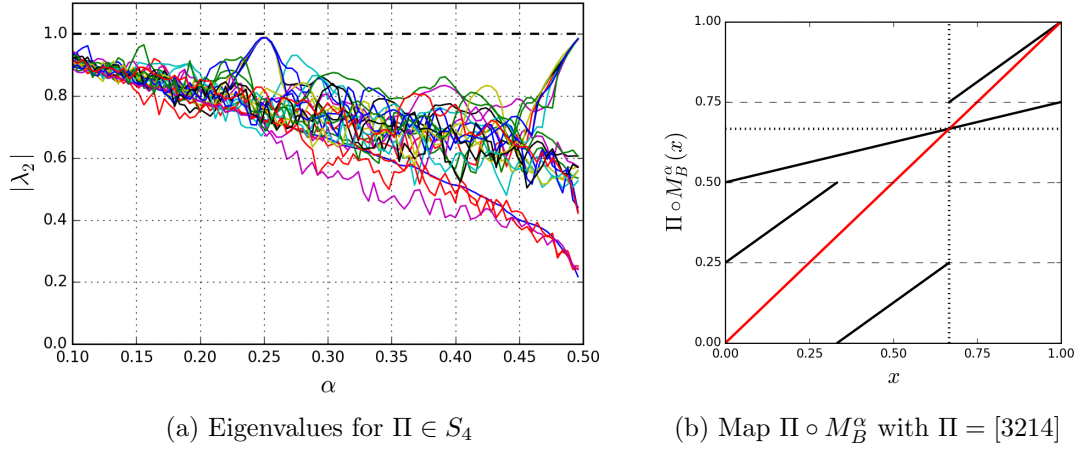
Figure 6.6: Profiles of approximated eigenvalues $|\lambda_2|$ with α for different Π are computed from the transfer matrices T_{nk} using a low resolution of $Q = 50$. Peaks in the profiles towards $|\lambda_2| = 1$ highlight the likely parameters which results in a decomposition. Confirmation of a decomposition can be found by plotting the one-dimensional map, shown in b) for $\alpha = 0.25$ and $\Pi = [3214]$, where again the images of the cut locations align with the fixed point x_f of the map.

$\Pi = [132]$, even small truncation values K will highlight transformations which produce a decomposition. This is because small wavemodes will be dominant features of the eigenfunction predicting the decomposition, so the eigenfunctions would be well approximated if enough wavemodes are included, and the respective eigenvalue ~ 1.

To investigate which values of α produce a decomposition for $\Pi \in S_N$, the transfer matrices $\boldsymbol{T}$ (6.5) are constructed for all permutations $\Pi \in S_N$ for a given N and 100 values of $\alpha \in [0.1, 0.5)$. Smaller value of $\alpha \in [0.0, 0.1)$ are ignored because as $\alpha \to 0$, all eigenvalues of the transfer operator will have $|\lambda_k| \to 1$ since $\alpha = 0$ represents the diffusion only limit. The resolution $Q = 50$ is chosen as it is small enough to ensure quick computation but large enough to avoid significant numerical diffusion. The values of $|\lambda_2|$ are then plotted against α to see where peaks towards 1 occur. Whether a decomposition has occurred can be checked by plotting the one-dimensional map for which it seems $|\lambda_2| \sim 1$, similar to those shown in Figure 6.5 for the $\alpha = 1/3$ example.

For example, in Figure 6.6 (a) the profiles of $|\lambda_2|$ approximated for all $\Pi \in S_4$ are shown, with peaks occurring around $\alpha = 1/4$ and $1/2$ for a handful of permutations. In Figure 6.6 (b) the map $\Pi \circ M_B^\alpha$ for an example permutation, $\Pi = [3214]$, which resulted in a value of $|\lambda_2| \sim 1$ in the low resolution approximation

171

(a) Eigenvalues for $\Pi \in S_3$ (b) Eigenvalues for $\Pi \in S_5$
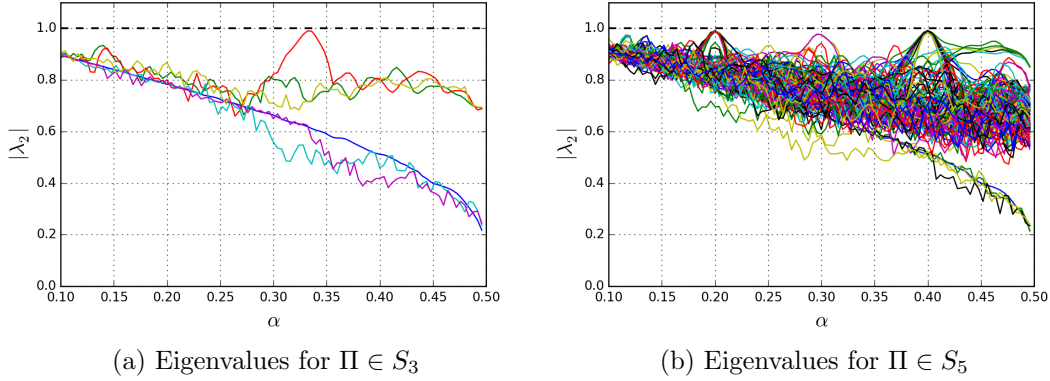
Figure 6.7: Profiles of approximated eigenvalues $|\lambda_2|$ with α for different Π . In a) for $\Pi \in S_3$, only one profile approximates a decomposition around $\alpha = 1/3$, while for b) $\Pi \in S_5$, two peaks occur for many permutations around $\alpha = 0.2$ and $\alpha = 0.3$. Further investigation into the peak at $\alpha = 0.3$ reveals a decomposition does not occur.

is illustrated. When $\alpha = 0.25$, indeed a decomposition of the domain into two components has occurred, $[0, 2/3)$ and $[2/3, 1)$. Note that the point $x_f = 2/3$ is a fixed point of the map but is also a pre-image of the cut-locations $x = 1/4$ and $x = 3/4$.

In Figure 6.7 the eigenvalue results are plotted for the permutation groups S_3 and S_5. When $N = 3$, the only problematic point is around $\alpha = 1/3$ with only one permutation causing a decomposition which was mentioned previously. When $N = 5$, the value of α resulting in $|\lambda_2| \sim 1$ are $\alpha = 0.2$ and $= 0.4$. These were confirmed by plotting the one-dimensional maps. Around $\alpha \sim 0.3$ there seems to be another peak of slow mixing, however it does not seem to reach $|\lambda| = 1$. The one-dimensional map reveals that a decomposition does not occur however the period 2 mapping shows a near decomposition dominates the dynamics. Not shown here, but for a representative subset of permutations from S_6, $\alpha = 1/6$, $1/3$ and $1/2$ resulted in decompositions. These computational results suggest the following conjecture.

Conjecture 6.2.1. *Given a permutation group S_N, and transformation $\Pi \circ M_B^\alpha :$ $[0, 1) \to [0, 1]$, where $\Pi \in S_N$ and $M_B^\alpha(x)$, where $M_B^\alpha : x \to M_1(x) \cup M_2(x)$,*

$$M_1(x) = \alpha x, \qquad M_2(x) = \beta x + \alpha, \qquad (6.7)$$

for $\alpha, \beta \in [0, 1]$, satisfying $\beta = 1 - \alpha$. A decomposition of the interval $[0, 1)$ into two or more disjoint intervals occurs if an interval $[a, b)$ exists such that

172

Permutation group	α	Π
S_3	1/3	[213]
S_4	1/4	[2134],[2143],[2314],[3214]
	1/2	[1324],[1342],[3124],[3142],[2413],[4213],[2431]
S_5	1/5	[21345],[21354],[21435],[21534],[21453][21543],[23145], [23154],[32145],[32154],[23415],[24315],[32415] [42315],[34215],[43215]
	2/5	[21435],[21453], [13425],[13425],[15423],[13452],[31245], [31254],[31425],31524],[41325],[51423],[31452],[31542], [51432],[34125][43125],[35421],[53421]

Table 6.1: Permutations which create a decomposition are listed under their respective S_N group and the parameter α for which the decomposition occurs. In all cases, pre-images of the cut locations align with fixed points, a necessity for a splitting of the interval $[0, 1)$ into two or more pieces.

$(\Pi \circ M_B^\alpha)([a, b)) = [a, b)$ *and* $(\Pi \circ M_B^\alpha)([a, b)) \cup [0, 1) \backslash [a, b) = \emptyset$. *The map* $\Pi \circ M_B^\alpha$, *does not have a decomposition if* $\alpha, \beta \notin \{1/N, 2/N \ldots, (N-1)/N\}$.

In all of the maps investigated, a decomposition occurred when a pre-image of a cut location aligned with a fixed point x_f of the map $\Pi \circ M_B^\alpha$ and could inform a starting point for the proof of Conjecture 6.2.1. However, whether a decomposition does not occur if the alignment does not occur is not known and so this fact is not sufficient to prove the Conjecture alone. Table 6.1 for each α and permutation group S_N lists the permutations which result in a decomposition.

For the remainder of the chapter, interest lies in the rates of mixing in maps which do not results in a decomposition of the domain. Therefore given a permutation group S_N, values of α which result in decompositions, and values of α in close proximity, are avoided.

6.2.2 Finite-time mixing behaviours

To begin examining the mixing behaviour of the maps $\Pi \circ M_B^\alpha$ with diffusion, initial examples are presented with $\alpha = 0.4$ and $\Pi \in S_3$, since S_3 has a manageable 6 permutations. Figure 6.8 plots the decay of variance ψ_j/ψ_0 for all 6 permutations when $\alpha = 0.4$, $\kappa = 10^{-5}$, and the initial condition was $\cos(2\pi x)$. Figure 6.8
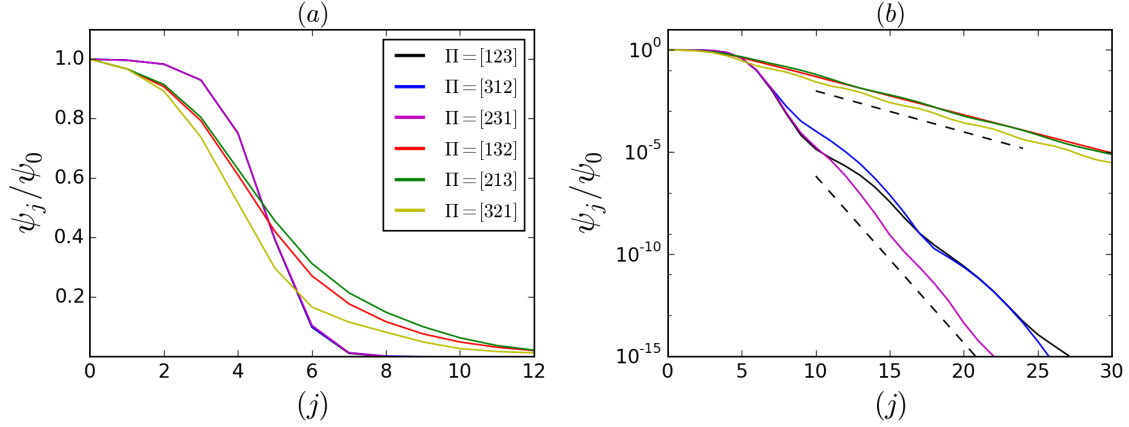
Figure 6.8: Profiles of variance decay for the composition maps $\Pi \circ M_B^\alpha$ for $\Pi \in S_3$ and $\alpha = 0.4$ are plotted on a) a linear axis and b) a semilogrithmic axis. The diffusivity coefficient is $\kappa = 10^{-5}$ and the initial condition $c^0(x) = \cos(2\pi x)$. The interleaving permutations deplete variance quicker in the first initial iterates, however the identity and rotation permutations are superior at depleting variance in the long-time limit. Example expected mixing rates $|\lambda_2|^{2j}$ are plotted as black dashed lines in b) for two permutations.

(a) plots the variance for the initial 12 iterations on a linear-linear plot, while Figure 6.8 (b) uses a linear-log plot to reveal the long-time mixing behaviour. All permutations $\Pi \in S_3$ have unique profile of variance, unlike the transformations in Chapter 5 where subgroups of the permutation groups S_N emerged. This is due to the fact that when non-uniformity in the stretching rates is included, symmetries between the maps under the action of rotations and reflections of the unit interval no longer occur.

In Figure 6.8 (a) it is seen that for the identity and rotation permutations $\Pi \in S_3^R$ the initial decay of variance is the same and all profiles are overlaid by the variance decay for $\Pi = [231]$. There is little depletion of variance in the first few iterations until a fast decay emerges thereafter. Note however that in Figure 6.8 (b), it is observed that once significant decay of variance has occurred for the rotation permutations, exponential decay emerges in the long-time with varying rates across $\Pi \in S_3^R$.

The interleaving permutations, $\Pi \notin S_3^R$ have contrasting mixing behaviour. They are named as such since one sub-interval in each is swapped and interleaved between the remaining two. In the initial stages the depletion of variance is faster than the rotation permutations. This is due to the discontinuities introduced on rearranging the concentration field via cutting and shuffling at each iteration.

174

These results are similar to the initial stages of mixing seen in Chapter 5, where the sharp interfaces also assisted the approach to uniformity. In the final mixing stages, shown in Figure 6.8 (b), the interleaving permutations have a long-time exponential mixing rate much slower than that of the rotation permutations.

This is reminiscent of the mixing behaviour observed in Chapter 5, where all permutations which cut and shuffled mixed slower than the rotation permutations when composed with the uniform baker's transformation, particularly notable for the long-time exponential mixing rates listed in Table 5.3 for $\Pi \in S_5$. However there, the rotation permutations had long-time super-exponential mixing. This is not the case for the systems discussed here. It is not immediately clear why, for this non-uniformly stretching system, the interleaving permutations are inferior at mixing in the long-time limit to the rotation permutations.

In line with other strong mixing transformations with diffusion, the main effect of decreasing the diffusivity coefficient κ on finite time mixing, is to extend the early stages before exponential mixing is achieved. Therefore, the finite time mixing is of no great interest, and so the next section looks in detail at how the asymptotic mixing rate changes depending on the parameters α, κ and Π. Profiles of $|\lambda_2|$ computed from the transfer matrices are compared across parameters since they approximate the average exponential decay of variance in the long-time limit, shown for the maps $\Pi \circ M_B^{\alpha}$ for two of the permutations in Figure 6.8 (b).

6.2.3 Asymptotic mixing rates

In the last section, for a specific example, it was observed that the initial stages of mixing are sped up by cutting and shuffling through the introduction of discontinuities, however the final stages, dominated by the emergence of eigenfunctions, were slowed by the discontinuous transformations. Figure 6.9 plots the profiles of $|\lambda_2|$ against the diffusivity coefficient κ for $\Pi \in S_3$. Each plot shows the profiles for different values of α. The profiles for the identity permutation, i.e. the map M_B^{α} acting alone, is plotted as a solid black line.

When $\alpha = 0.4$, for all values of κ, the identity and rotation permutations have a faster long-time mixing rate than the interleaving permutations, with the rotations also having lower values of $|\lambda_2|$ compared to the identity permutation for most values of κ. The value of $|\lambda_2|$ quickly becomes independent of diffusivity as $\kappa \to 0$ for the interleaving permutations.

On decreasing α, the discrepancy in the mixing rates across the profiles between $\Pi \in S_3^R$ and $\Pi \notin S_3^R$ decreases, where for $\alpha = 0.2$ the rotations no longer
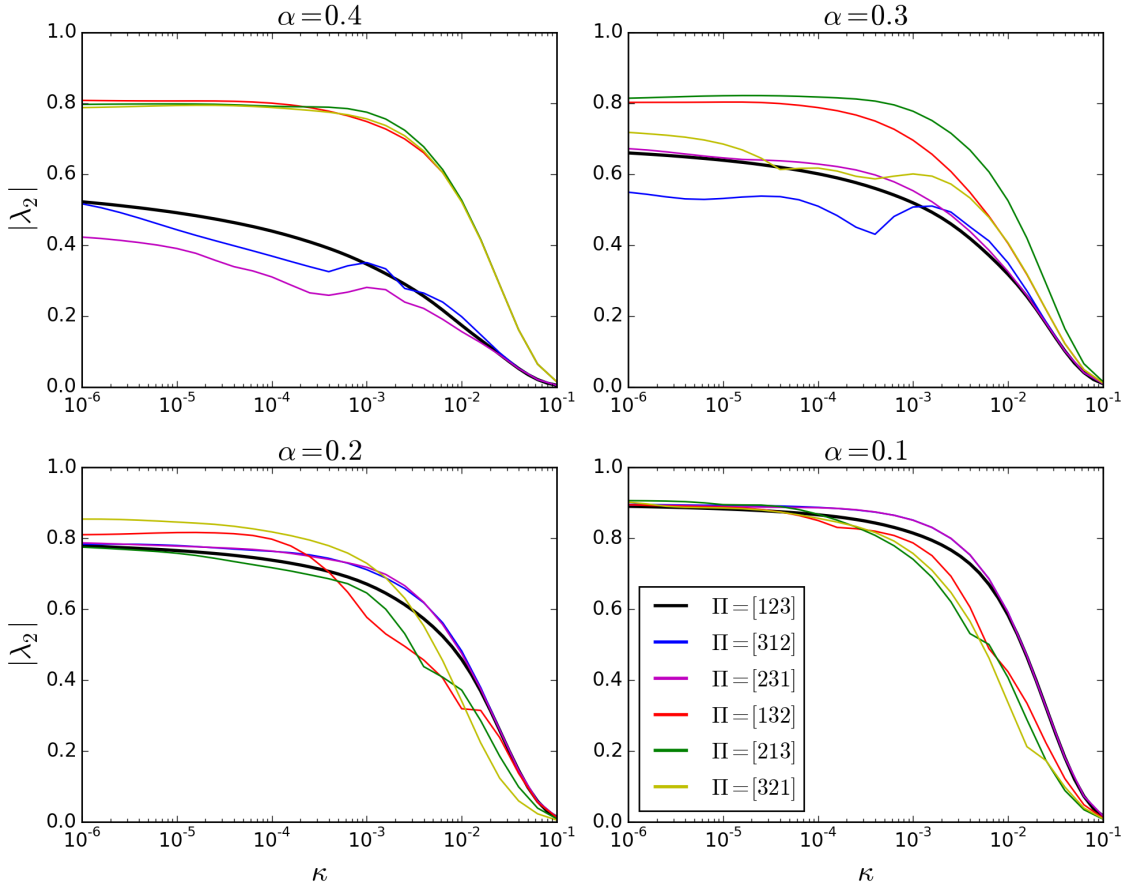
Figure 6.9: The modulus of the leading eigenvalues $|\lambda_2|$ are plotted against the diffusivity coefficient κ for $\Pi \circ M_B^\alpha$ with $\Pi \in S_3$. Each plot represents a different α value. Two behaviours emerge categorised by the rotation permutations and interleaving permutations. The interleaving permutations only mix faster than the rotation permutations on decreasing α and increasing κ.
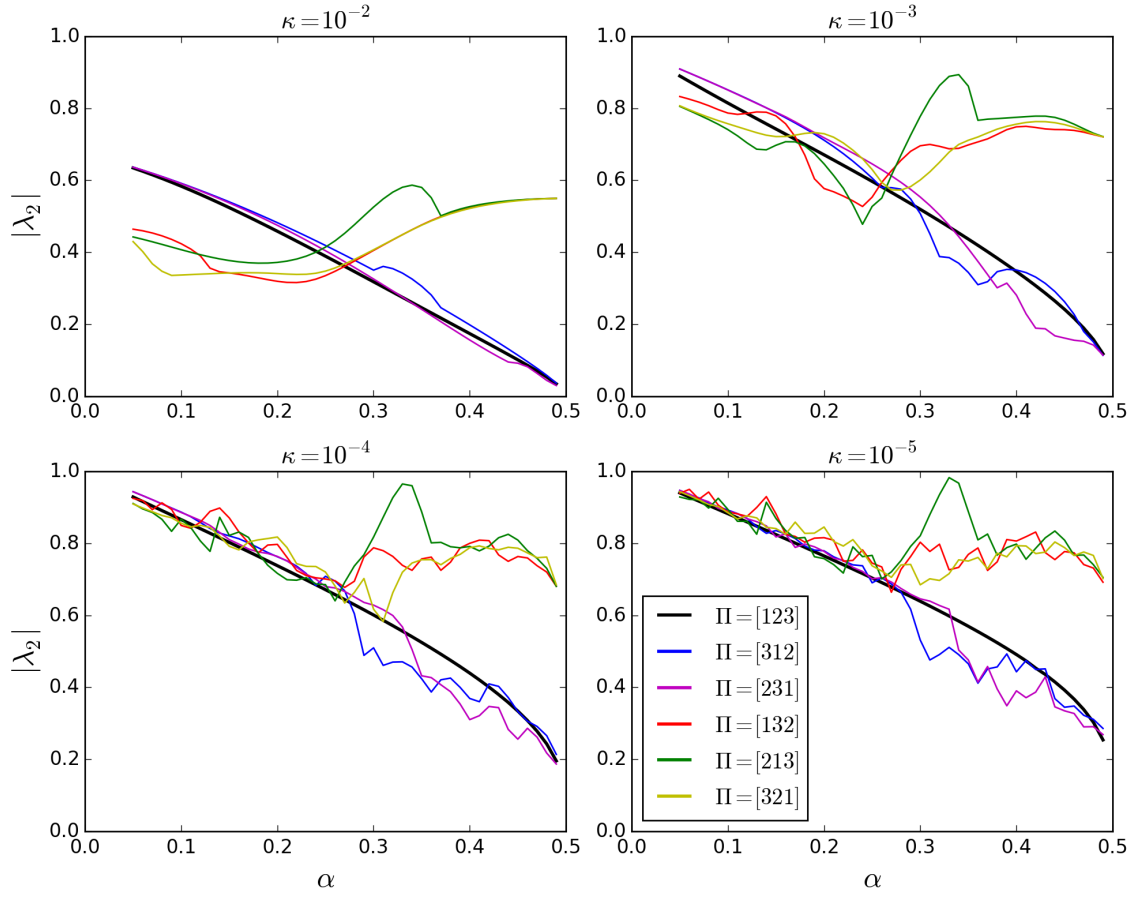
Figure 6.10: The modulus of the leading eigenvalues $|\lambda_2|$ are plotted against the non-uniformity parameter α for $\Pi \circ M_B^\alpha$ with $\Pi \in S_3$. Each plot represents a different diffusivity coefficient κ. The finer values of α confirm the observations of Figure 6.9

177

have faster mixing rates across the whole range of κ. By $\alpha = 0.1$, for large values of κ, the interleaving permutations have a faster mixing rate than the rotation permutations. However, as $\kappa \to 0$, the profiles converge to a similar mixing rate, ~ 0.9, the diffusion-less strong mixing rate for M_B^α with $\alpha = 0.1$. Many of the profiles varying non-monotonically across the range of κ, although overshooting the diffusion-less mixing rate does not appear to happen. This can not be fully understood since the strong mixing rates in the diffusion-less limit are not known.

Note the same trends appear in Figures 6.10, where each plot represents a different value of κ, and the profiles of $|\lambda_2|$ plotted for finer values in α. Note for $\Pi = [123]$, shown again as a solid black line, the profile changes smoothly with α and monotonically. For the other permutations, the profiles do not vary smoothly or monotonically and as diffusion coefficient κ becomes very small, the profiles vary erratically with α. Note the peak in the profile of $\Pi = [213]$ corresponds to $\alpha = 1/3$ where the transformation is no longer mixing or ergodic as discussed in Section 6.2.1. Decreasing the value of diffusion further does not qualitatively change the plot observed for $\kappa = 10^{-5}$.

The trends in the profiles of $|\lambda_2|$ for $\Pi \in S_3$ with varying α values are seen again for the permutations from S_4, where Figure 6.11 plots $|\lambda_2|$ for every $\Pi \in S_4$ with κ for 4 different values of α. The identity permutation is plotted in black again, while the three rotation permutations are plotted as red lines, and the remaining permutations which swap two or more elements are all plotted in grey. For $\alpha = 0.4$, the permutations $\Pi \in S_4^R$ have profiles of $|\lambda_2|$ with the lowest values for almost all values of κ, while for $\Pi \notin S_4^R$ the profiles are either similar in value to the interleaving permutations of S_3 in Figure 6.9, or fall at intermediate between these and the profiles for $\Pi \in S_4^R$. As $\alpha \to 0$ and the non-uniformity in stretching rates is larger, the permutations $\Pi \notin S_4^R$ which cut and shuffle parts of the concentration field improve the mixing rates on average when the diffusivity is fast (large κ). As $\kappa \to 0$ a transition occurs and most of values of $|\lambda_2|$ for $\Pi \notin S_4^R$ are larger than for $\Pi \in S_4^R$. When $\alpha = 0.1$, only a handful of permutations improve the mixing rate for the smallest values of κ, with the majority having a larger value of $|\lambda_2|$ than the identity or rotation permutations.

These results suggest that it is only for fast diffusivity rates and large non-uniformity in the stretching rates that improvement in the long-time mixing is generally achieved from cutting and rearranging the concentration field. The latter observation makes sense, since the addition of sharp interfaces introduced at each iteration combined with a fast diffusivity rate will speed up the average
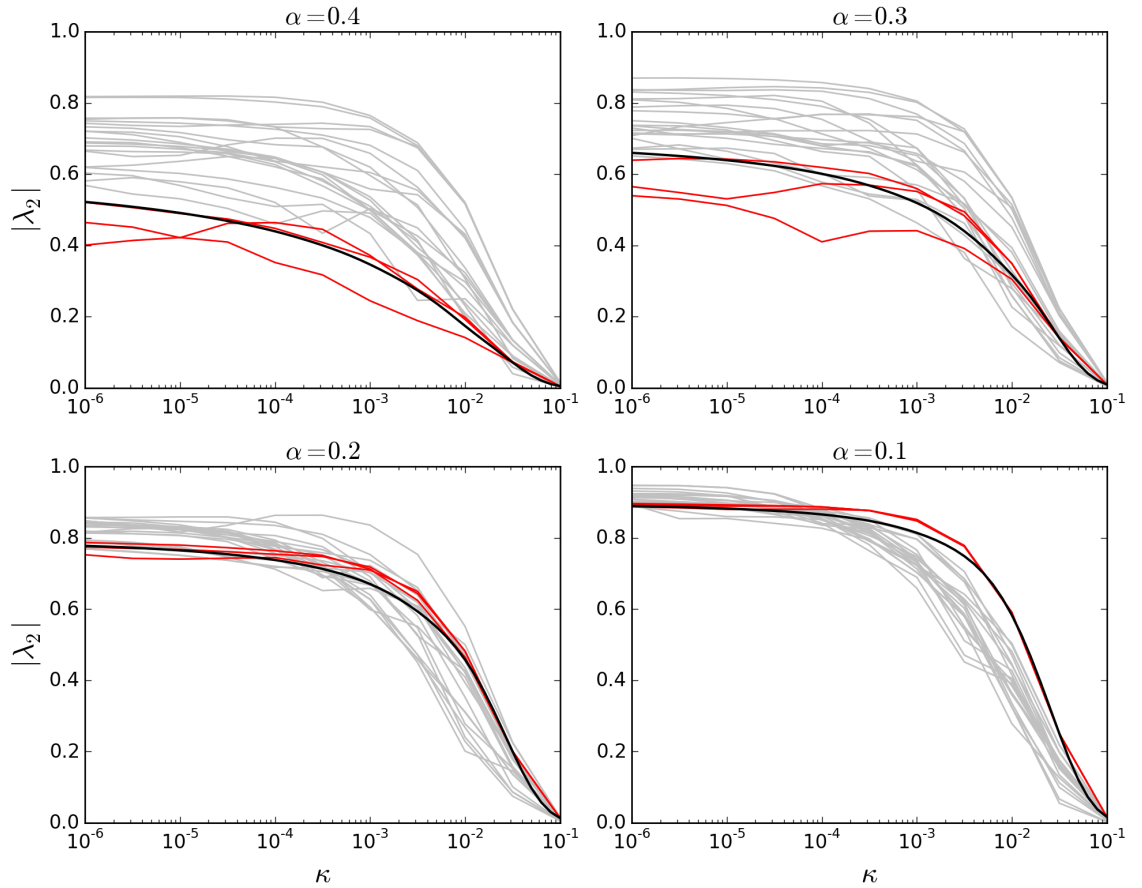
Figure 6.11: The modulus of the leading eigenvalues $|\lambda_2|$ are plotted against the diffusivity coefficient κ for M_B^α with $\Pi \in S_4$. The identity permutation (black) and rotation permutations (red) are highlighted as specific profiles against all other permutations (grey). Trends in the profiles of S_3 persist in S_4, such that for small non-uniformity, cutting and shuffling hinders the long-time mixing. When the non-uniformity is larger, the interleaving permutations improve mixing for large κ, but as κ is decreased, improvement in the rate of mixing does not happen on average
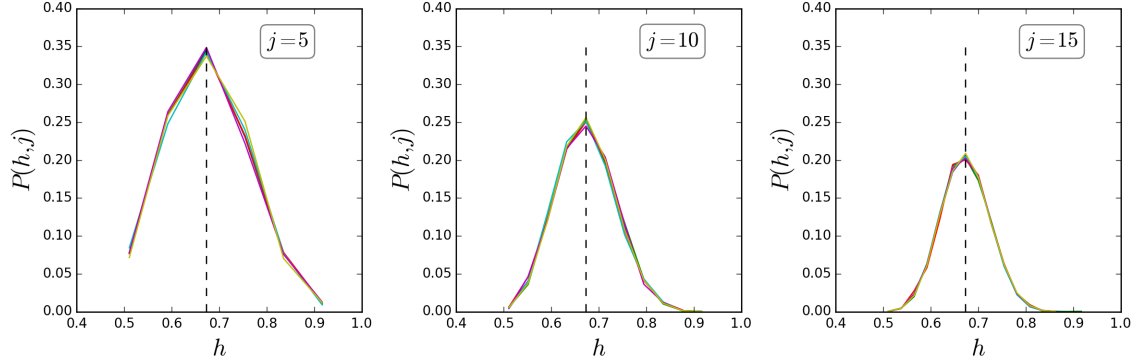
Figure 6.12: Probability density functions (PDFs) of FTLEs for the composite maps $\Pi \circ M_B^\alpha$ for $\Pi \in S_3$ and $\alpha = 0.4$. Each Π is represented by a different coloured profile as in previous plots. Different iteration values j are shown in each plot. For all permutations the PDFs are the same, except for small deviations arising from numerical errors in computation.

mixing rate. The rest of the discussion focuses on the cases with $\alpha \geq 0.3$ and $\kappa \leq 10^{-5}$ where the mixing rates are hindered by the interleaving permutations. In the next section a closer look at the concentration field evolution and the eigenfunctions in the long-time limit help explore the reasons for trends observed in the values of $|\lambda_2|$.

6.3 Stretching histories and eigenfunctions

6.3.1 Distribution of FTLEs and mixing rates

Previously, literature has tried to understand the mixing rates of passive scalars in chaotic flows from the underlying chaotic dynamics. Notably Antonsen Jr *et al.* (1996) approximated the decay of scalar variance by relating it to the growth of particle pair separations, implying a Lagrangian analysis of stretching histories is enough to determine the amount of mixing. However, the one-dimensional non-uniform baker's transformation was used by Fereday *et al.* (2002) to show this isn't always the case, although Wonhas & Vassilicos (2002) showed that the probability distribution function (PDF) of FTLEs could be used to approximate well the decay of variance in its initial stages.

Figure 6.12 shows PDFs of FTLEs for the 6 maps $\Pi \circ M_B^\alpha$ for $\Pi \in S_3$ when $\alpha = 0.4$. The distributions were computed as the ratios of stretching histories from 1000 evenly distributed points within the domain $[0, 1)$. Three different j
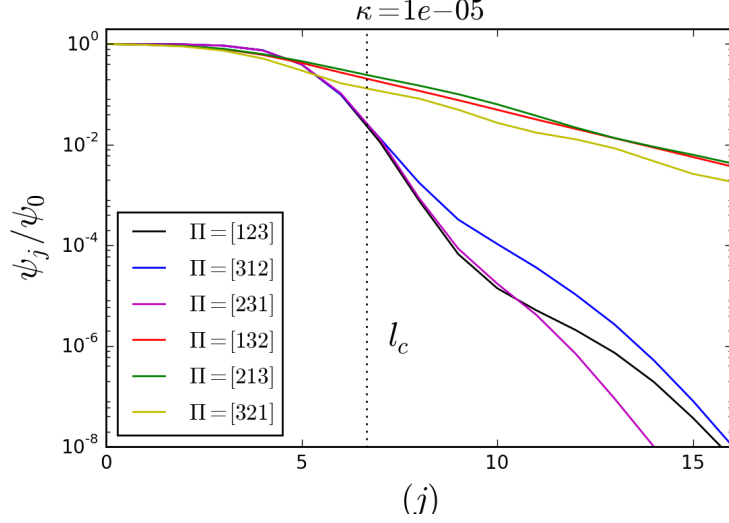
Figure 6.13: The initial transient of variance decay from Figure 6.8 is replotted with l_c marked, the Batchelor length scale for the parameters α and $\kappa = 10^{-5}$. For the identity and rotation permutations, the variance decay profiles agree until divergence when the Batchelor lengthscale is reached and the eigenfunction regime emerges. For the interleaving permutations divergence from FTLE prediction happens from the first iteration.

iterations of the averages are shown, and in all the PDFs for the composite maps are exactly the same. Small variations arise purely from numerical computation, since the PDFs can be computed analytically.

At each iteration, every point representing the concentration field will experience two compression rates, one of value α and one of value β, and stretch factors $1/\alpha$ and $1/\beta$ respectively. The addition of a permutation at each iteration does not change this fact; there will always be two images with two different compression rates. Therefore, for each composite map $\Pi \circ M_B^\alpha$ the distribution of stretching histories is represented by

$$P(h, j) = \sum_{l=0}^{j} \binom{j}{l} \alpha^{j-l} \beta^l \delta \left(h - h_\alpha + \frac{j}{l}(h_\alpha - h_\beta) \right) \tag{6.8}$$

where $h_\alpha = -ln\alpha$ and $h_\beta = -\ln\beta$, corresponding to the Lyapunov exponent responding to uniform strain rate by $1/\alpha$ and $1/\beta$ respectively. This was given for the non-uniform baker's transformation by Wonhas & Vassilicos (2002). In the limit $j \to \infty$ (6.8) tends to a normal distribution. The fact that the PDFs of FTLEs are the same for all the composition transformations is specific to the one-dimensional mapping under consideration. It is not assumed to hold for

higher-dimensional flows which are a composition of stretching and folding, and cutting and shuffling.

The fact that the PDFs of stretching histories approximate well the mixing in the early stages of smoothly deforming transformations is seen here in the profiles for the decay of variance for $\Pi \in S_3^R$, the rotation permutations. Figure 6.13 replots the variance decay in the initial stages for the maps $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$, $\Pi \in S_3$ and $\kappa = 10^{-5}$. Wonhas & Vassilicos (2002) approximate the iterate l_c the Batchelor's length scale is reached for the non-uniform baker's transformation as

$$l_c = \frac{1}{2h_\beta} \log \left(\frac{h_\alpha + h_\beta}{16\pi^2\kappa} \right), \tag{6.9}$$

The iteration l_c is plotted as a dotted black line in Figure 6.13. For the 3 rotation permutations, which are smoothly deforming, agreement in the variance decay occurs up until l_c is reached and the eigenfunction regime begins. After this time the profiles diverge in long-time mixing rate, although by this time the variance has already reached around $\psi_j/\psi_0 \sim 10^{-2}$. Therefore since the stretching histories approximated well the mixing for M_B^α, they do for all $\Pi \circ M_B^\alpha$ with $\Pi \in S_3^R$. Approximating the appearance of the Batchelor length scale from (6.9) does not appear to hold for the three interleaving permutations, which each depleted the variance quicker in the initial stages. This suggests where discontinuous transformations are concerned, mixing rates will not be well approximated by stretching histories across all stages of mixing.

6.3.2 Eigenfunctions and asymptotic mixing rates

For the examples discussed thus far, the large discrepancies in the asymptotic mixing rates between the rotation permutations and the interleaving permutations in the composition maps can be understood from the qualitative appearances of the dominant eigenfunctions. This further gives insight to the mechanisms contributing to the values of $|\lambda_2|$ across varying parameter values.

For $\alpha = 0.4$ and $\kappa = 10^{-5}$, for which the finite time mixing behaviour was shown in Figure 6.8, Figure 6.14 plots the eigenfunctions v_2 computed from the transfer matrices via their real (blue) and imaginary (grey) parts. The dominant eigenfunctions v_2 were computed from the transfer matrices and transformed into real space via a inverse Fast Fourier transform algorithm in Python. Since the transfer matrices (6.3) share the symmetry $\boldsymbol{B}_{-k-q} = \overline{\boldsymbol{B}_{kq}}$, $\lambda_3 = \overline{\lambda_2}$ is also a dominant eigenvalue with corresponding eigenfunction $v_3 = \overline{v_2}$, which is omitted
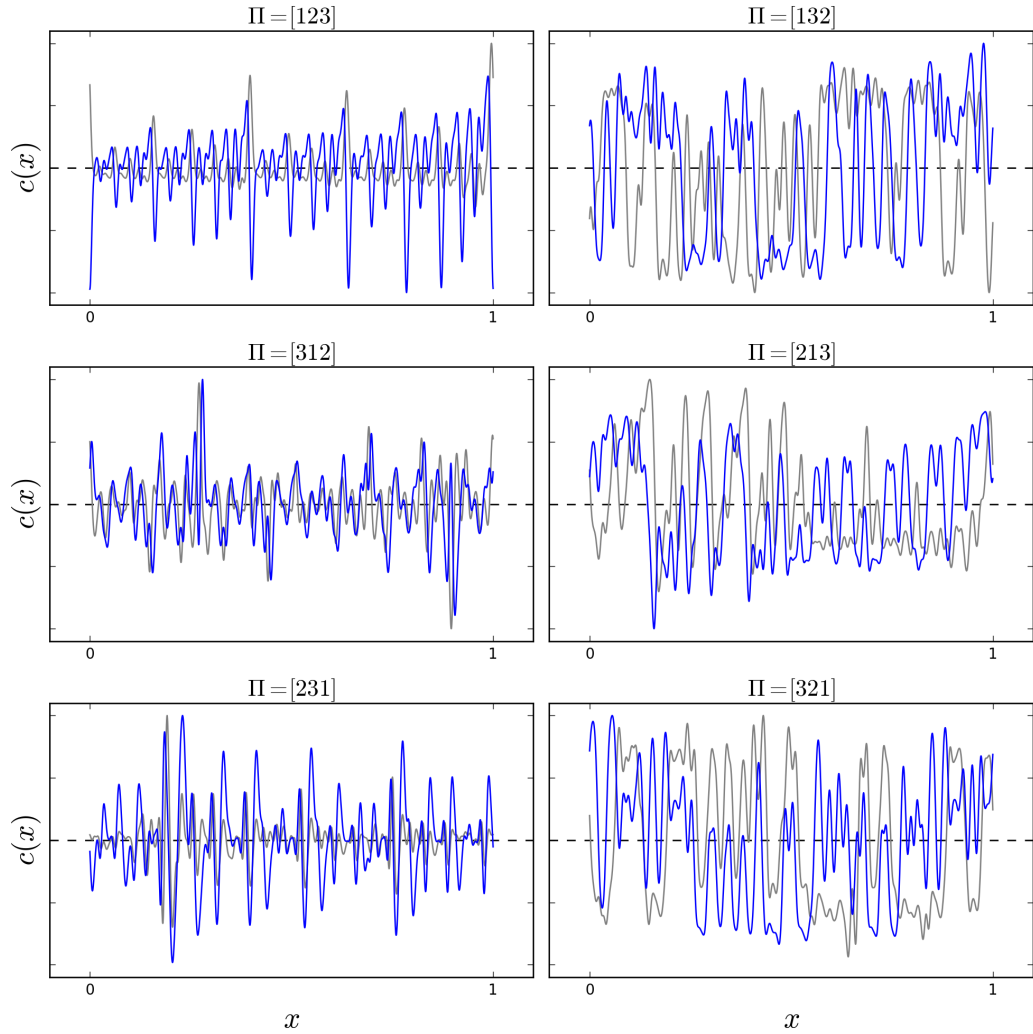
Figure 6.14: Real (blue) and imaginary (grey) parts of the dominant eigenfunctions v_2 are plotted as scalar fields for composition transformations $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$ and $\kappa = 10^{-5}$. Each plot represented a different permutation $\Pi \in S_3$. The dashed black line in each represents $\bar{c}$. Qualitative differences in the eigenfunctions appearance are visible between the identity and rotation permutations (left), and the interleaving permutations (right).
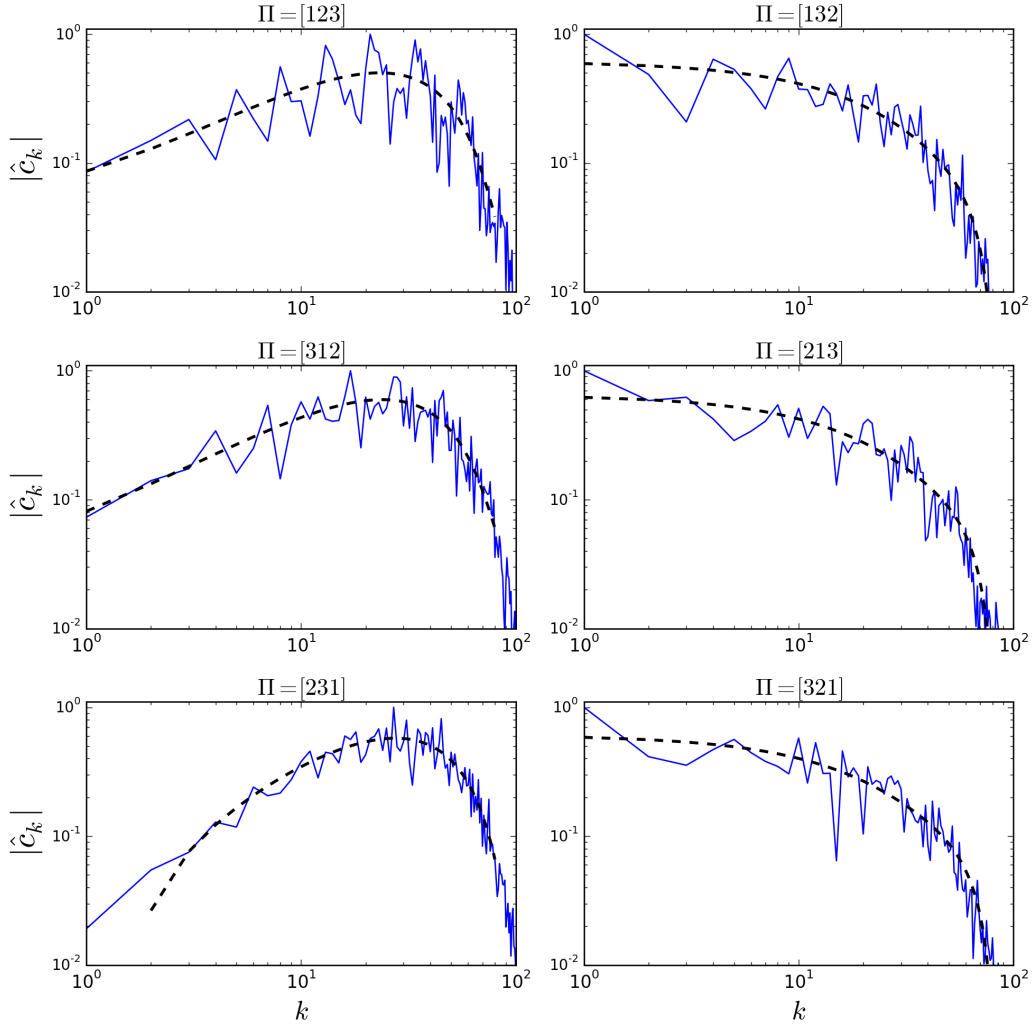
Figure 6.15: The modulus of the coefficients $|\hat{c}_k|$ are plotted for the spectrum of the dominant eigenfunctions v_2 shown in Figure 6.14 for the composition transformations $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$ and $\kappa = 10^{-5}$. Each plot represented a different permutation $\Pi \in S_3$. A four degree polynomial line of best fit is plotted, showing that dominating scales in the eigenfunctions of the rotation permutations (left) are at larger wavemodes than the interleaving permutations (right).

from the plots since there is only a symmetry in the imaginary part due to the conjugacy. Each plot represents a different permutation $\Pi \in S_3$. The black dashed line represents the mean field of the concentration, which will be regarded as $c(x) = \bar{c} = 0$ for the discussion.

The eigenfunctions for the identity and rotation permutations (left) and the interleaving permutations (right) are qualitatively different. The rotation permutations have a number of large peaks in the concentration field, which are small in width, and frequent transitions between positive and negative values occur across the whole unit interval. This contrasts the appearance of the eigenfunctions for the interleaving permutations, where large peaks occur across the whole domain, many large in width and large regions dominated either by all positive or all negative values.

Figure 6.15 plots the wave spectrum for each of the eigenfunctions, where the respective coefficients $|\hat{c}_k|$ are plotted against wave-number k. A line of best fit is plotted using a polynomial of 4th degree found using polyfit function in Numpy, Python. Although the spectrum varies non-smoothly between each consecutive wave number, for the rotation permutations a peak in the fitted function occurs around $k \sim 25$. However, for the interleaving permutations, the peak in the spectrum occurs for the smallest value $k = 1$ and as such the dominant features in the concentration field are large in scale. The spectrum provides a quantitative description of the features seen in the eigenfunctions of Figure 6.14. It is this dominant scale in the eigenfunctions of the interleaving permutations which contribute to the overall slow mixing rates.

In the eigenfunction regime, the dominant mechanism stirring the concentration field is the compression of v_2 by M_B^α at each iteration. The permutation step could cut and rearrange the concentration field, reducing the scale of one or more striations and introduce large gradients via a discontinuity, however the overall effect of this is the concentration field will be negligible compared to the action of the stretching. Conflictingly, it could also rejoin striations of similar colour, increasing the average scale in the concentration field. The diffusion step will then blur any large gradients. Since the scales in the eigenfunctions are larger for the interleaving permutations in comparison to the rotation permutations, they will remain larger after the action of the stirring, hence less chance of striations being washed out by diffusion, and a slow decay rate overall.

Although the large scales in the eigenfunction contribute to the slow asymptotic mixing rates, how the scales emerge is not initially clear. Figure 6.16 plots
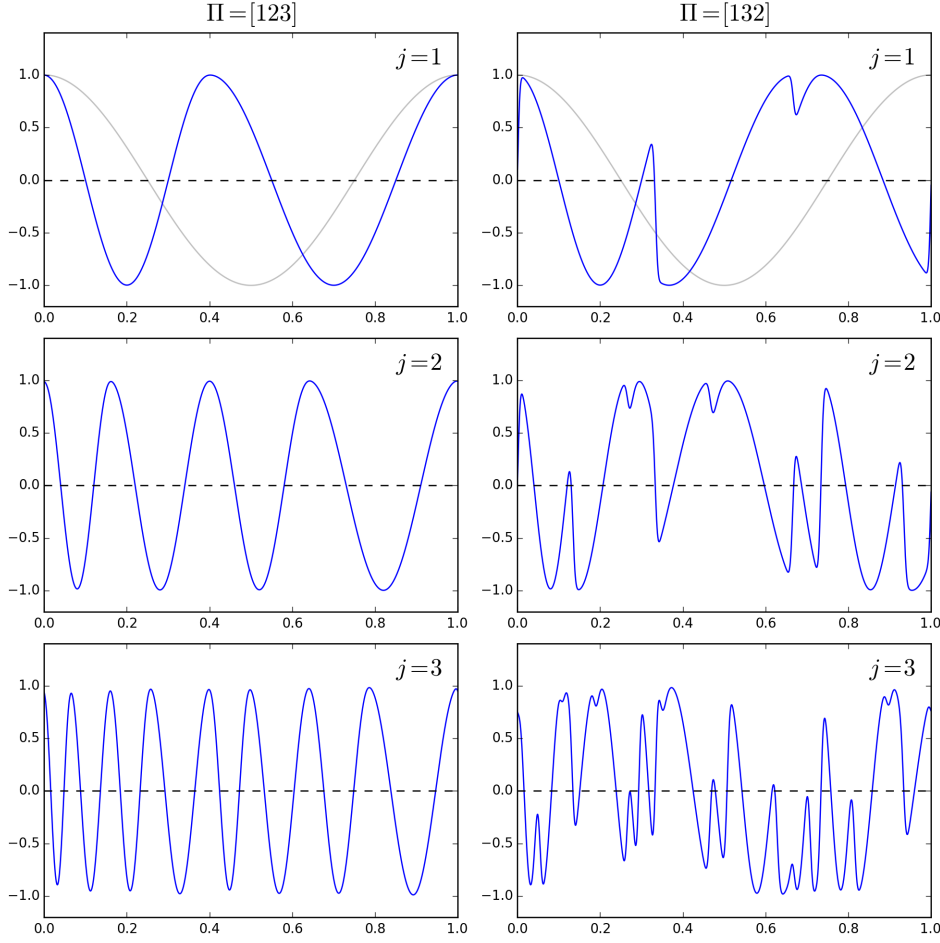
Figure 6.16: The initial iterations of the concentration field evolution for two composite maps $\Pi \circ M_B^\alpha$ are plotted for $\Pi = [123]$ (left) and $\Pi = [132]$ (right). The initial condition was $c^{(0)}(x) = \cos(2\pi x)$, plotted in grey in the top images. While variation in striation width arises from the non-uniformity in stretching rates for the identity permutation $\Pi = [123]$, the interleaving permutation $\Pi = [132]$ produces large striations from reassembly by cutting and shuffling.

the evolution of the concentration field under the action of M_B^α and $\Pi \circ M_B^\alpha$ with $\Pi = [132]$ for $\alpha = 0.4$ and $\kappa = 10^{-5}$. The iteration number of each concentration field in plotted in the top, right hand corner. The initial condition for both was $c^{(0)}(x) = \cos(2\pi x)$, plotted as the grey profile in both plots with $j = 1$. For $\Pi = [123]$, the differences in the compression rates at each iteration produce slight variations across the widths of the striations, which the largest striations eventually becoming dominant features in the eigenfunction. For $\Pi = [132]$ at iteration $j = 1$, after the initial compression by the map M_B^α, the rearrangement of the concentration field introduces two sharp discontinuities at $x = 0, 1/3$ and $2/3$, however the rearrangement has rejoined striations of similar value, generating large scale features in the concentration field. On further iteration, these large features are carried into the future concentration fields and further rejoining of striations occurs. By the third iteration, large regions of all positive or all negative values already dominate the concentration field.

This example only considers a single initial concentration field, but since eigenfunction regimes and mixing rates are independent of almost all initial conditions, it would be unsurprising that the contamination to mixing from the rejoining of striations is a global mechanism in cutting and shuffling transformations, which would contaminate for almost all initial conditions. In smoothly deforming systems where the stirring mechanism is from stretching and folding alone, the only contamination to the mixing is from areas which experience slow compression rates, while for the composite transformations, both mechanisms of contamination are possible; slow compression and striation reassembly. The next section considers the dominant eigenfunctions in the long-time limits and whether any of the distinctive features can be approximated by the dynamics of the underlying transformations.

6.3.3 Eigenfunction features and their relation to underlying dynamical properties

To briefly summarise the results so far on the mixing rates of the composite maps $\Pi \circ M_B^\alpha$; permutations which cut and rearrange the concentration field can lead to slow mixing rates arising from the reassembly of striations, and a loss in the predictability of short time mixing rates approximated by stretching histories. In this final section, whether any features of the concentration field in
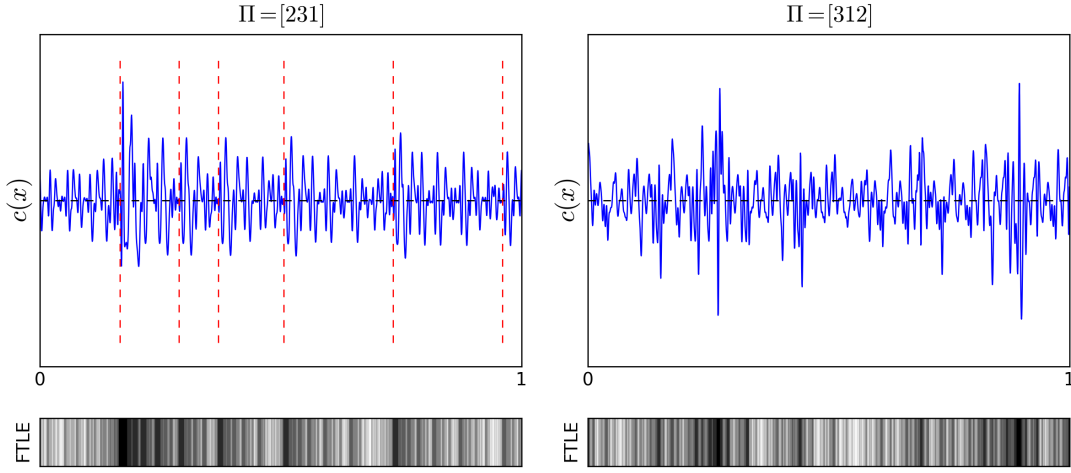
(a) v_2 for $\Pi \circ M_B^\alpha$ with $\Pi = [231]$ (b) v_2 for $\Pi \circ M_B^\alpha$ with $\Pi = [312]$

Figure 6.17: Dominant eigenfunctions v_2 are plotted for the composite maps $\Pi \circ M_B^\alpha$ with $\Pi = [231]$ and $[312]$, the rotation permutations, with $\alpha = 0.4$ and $\kappa = 10^{-6}$. Additionally, the distributions of the FTLEs are plotted for $j = 8$ iterations below the plots. Black, representing areas of least compression, align with peaks in the concentration field. For $\Pi = [231]$, which has a fixed point on a branch of least compression, images of this fixed point also align with the peaks in the concentration field.

the eigenfunction regime can be pre-determined from the dynamical features of the underlying stirring flow is investigated.

Figure 6.17 plots of the eigenfunctions for $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$ and $\Pi \in S_3^R$, the rotation permutations. Underneath each of the plots, a colour map representing the distribution of FTLEs across the unit interval is shown for $j = 8$ iterations. Light regions correspond to areas which have experienced fast stretching histories while dark regions correspond to the areas which have seen the least stretching. The 8th iteration of FTLEs was chosen since it captured enough of the details in the FTLE field to compare with the eigenfunctions for the diffusivity coefficient $\kappa = 10^{-6}$. A smaller diffusivity coefficient was applied in comparison to Figure 6.14, which showed eigenfunctions for $\kappa = 10^{-5}$, to show the finer details which emerge when decreasing the diffusivity rate. For $\Pi = [231]$, the non-trivial fixed point x_f which falls on the permuted β branch of least compression, is plotted alongside its images under the mapping as dashed red lines, and align with the regions of least stretch in the FTLE distribution. The mean field, $\bar{c}$ is included as a dashed black line.

The eigenfunctions for the composition maps with rotations are well mixed, with fluctuations about the mean value small in general. Peaks in the concentration field exist as distinctive features in the eigenfunction, focused near areas which have experienced the least compression. This is seen by comparing the peaks to the position of the dashed red lines, or the dark areas in the FTLE distributions. Since the rotations do not chop and rearrange the concentration field at each iteration, the stirring dynamics are still smooth and the eigenfunctions behave similarly to past literature of smoothly deforming chaotic systems, where eigenfunctions have been shown to align with regions of low stretch [Gouillart et al. (2008); Popovych et al. (2007); Wonhas & Vassilicos (2002)].

The eigenfunctions for the interleaving permutations in S_3 depicted in Figure 6.14 are not well mixed, instead composed of regions of large width, composed of either all positive or all negative values. To see whether dynamical features of the underlying map $\Pi \circ M_B^\alpha$ help predict the appearance of the eigenfunctions, $\Pi = [213]$ is used as an initial example.

Figure 6.18 shows the illustrative map for $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$ and (a) $\Pi = [213]$, and a second permutation (b) $\Pi = [4321]$. These permutations cut and shuffle the concentration field at each iteration, but also have fixed points which fall on the β branches of least compression. This is not true for the other interleaving permutations of S_3 or many other permutations in general. Since these fixed points and their pre-images align with the slowest stretching histories of the FTLE distributions, they should provide accurate point comparisons to see if the eigenfunctions in the composite maps which are cut and shuffled, also align with points which have experienced the least stretching rates.

Figure 6.19 (a) plots the dominant eigenfunctions v_2 for $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$, $\Pi = [213]$ and $\kappa = 10^{-6}$. Red dashed lines plots the images of the two fixed points which fall on the β branch, $x_f = 1/6$ and 1, and again the black dashed line plots the mean-field value $\bar{c}$. The eigenfunction of this smaller diffusivity coefficient is also not well mixed, with large width fluctuations between positive and negative values. Across much of the domain, the images of the fixed points align with many sharp transitions between positive and negative values in the concentration field. However, the eigenfunction is not static and the concentration field oscillates with further iteration of the transfer operator.

As an additional diagnostic, maximum and minimum values in the evolution of the eigenfunction can be found by evolving v_2 by the transfer operator, $P^j v_2$, and rescaling by the decay factor of the eigenvalue $|\lambda_2|$ at each iteration,

(a) $\alpha = 0.4$, $\Pi = [213]$ (b) $\alpha = 0.4$, $\Pi = [4321]$
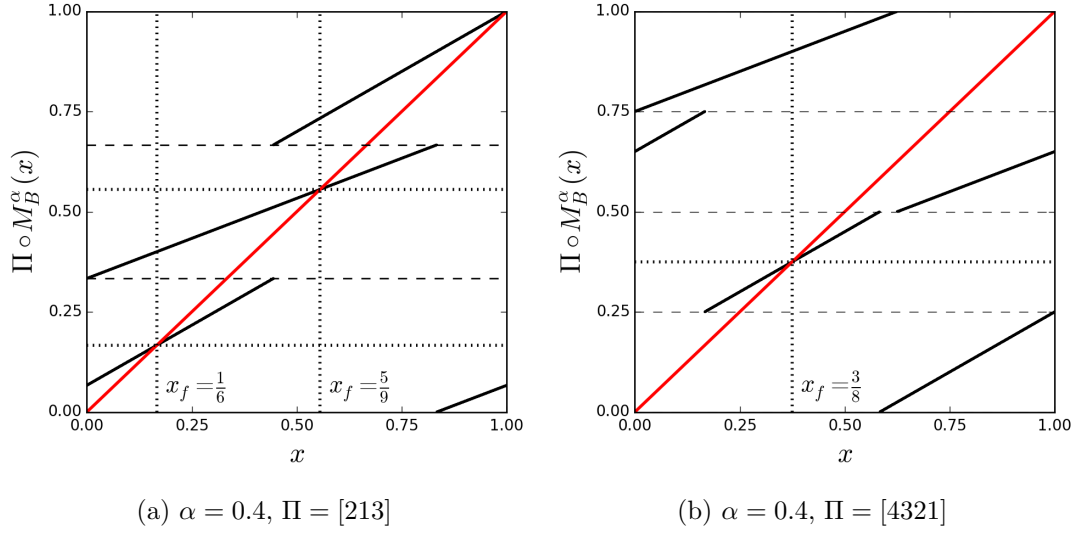
Figure 6.18: Maps for $\Pi \circ M_B^\alpha$ shown with periodic points x_f marked as they cross the line $\Pi \circ M_B^\alpha(x) = x$ (red). In a) the fixed point $x_f = 5/9$ lines on the faster compression branch while $x_f = 1/6$ and 1 correspond to points of least compression. In b) the one fixed point for $\Pi = [4321]$ lines on the branch of least compression.

$P^j v_2/|\lambda_2|$, therefore renormalising the scale of the concentration field to highlight distinctive features. Figure 6.19 (b) plots the maximum and minimum values of the concentration field during the evolution of the eigenfunction. The profiles of maximum and minimum values (blue) are seen to be symmetric abound the mean field $\bar{c}$ (dotted black). The top plot shows again the locations of the fixed points and their images, showing alignment with many distinctive peaks which now emerge as maximums and minimums of the eigenfunction. The bottom plot also shows the locations of the cuts and their images as an additional dynamical feature. Again, alignment with some of the peaks occurs, but this is due to the fact that the cut locations align images of fixed points. The cut locations do not highlight any other distinctive aspects of the eigenfunctions appearance.

As a second example, Figure 6.20 (a) plots the dominant eigenfunctions v_2 for $\Pi \circ M_B^\alpha$ with $\alpha = 0.4$, $\Pi = [4321]$ and $\kappa = 10^{-6}$. Red dashed lines plots the images of the single fixed point which fall on the β branch, $x_f = 3/8$. Analogous to the eigenfunctions of the interleaving permutations in S_3, the eigenfunction is not well mixed in general, with large fluctuations about the mean field across the whole domain. Unlike the previous example however, images of the fixed point do not align with the dominant peaks in the eigenfunction. Many of the peaks indeed align with dark regions in the FTLE distribution, included below the plot,
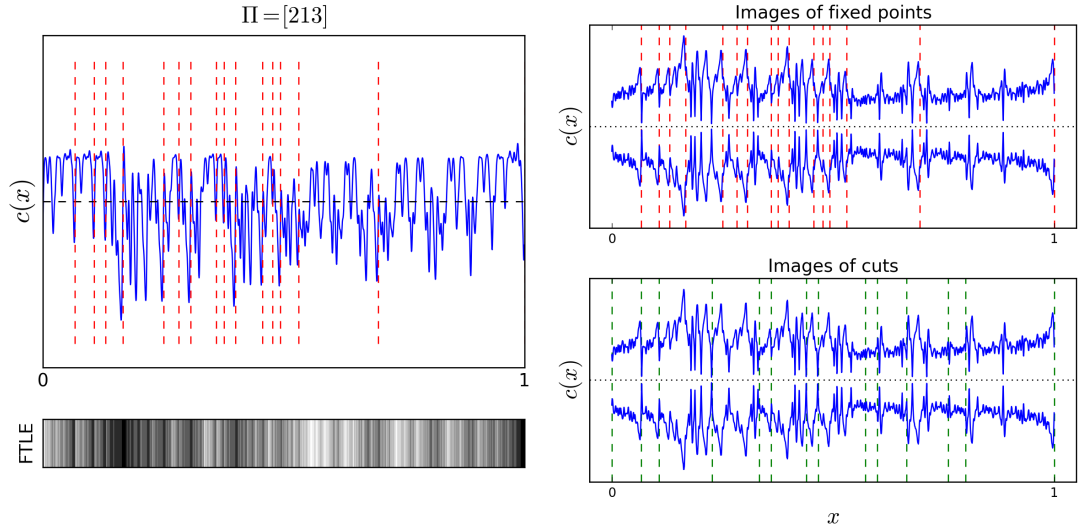
(a) v_2 plotted with FTLE distribution.

(b) Maximum and minimum values of v_2 plotted with images of x_f and images of cut locations

Figure 6.19: In a) dominant eigenfunction v_2 is plotted for the composite maps $\Pi \circ M_B^\alpha$ with $\Pi = [213]$, with $\alpha = 0.4$ and $\kappa = 10^{-6}$. Additionally, the distribution of FTLEs is plotted for $j = 8$ iterations below the plots. Black, representing areas of least compression, align with peaks in the concentration field. When $\Pi = [231]$, there is a fixed point on a branch of least compression, and images of this fixed point agree with the peaks in the concentration field. In b) maximum and minimum values of the concentration field under further iteration of the eigenfunction show that many of the dominant peaks coincide with areas of least compression, while images of the cuts show some agreement with peaks but show no obvious alignment with any features of the eigenfunction.
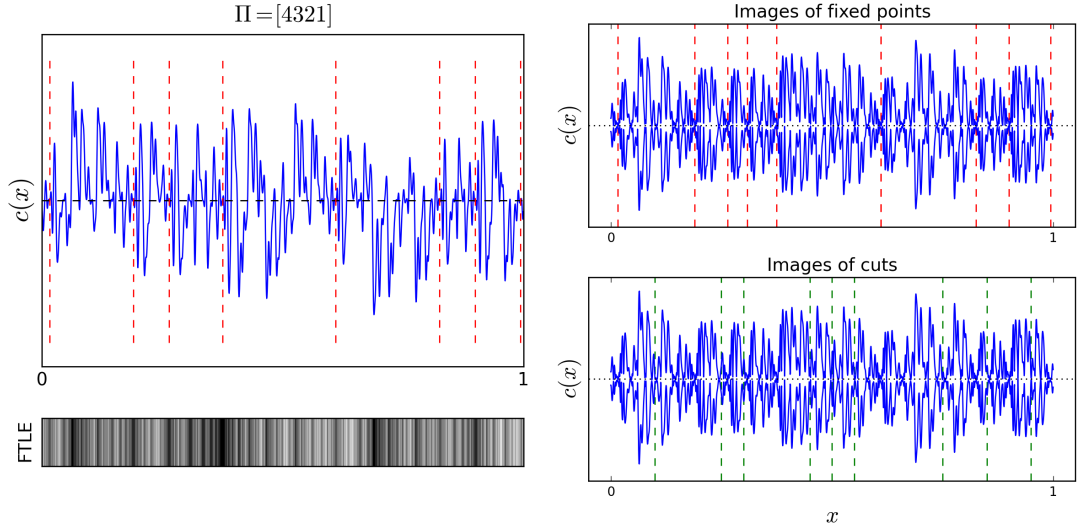
(a) v_2 plotted with FTLE distribution.

(b) Maximum and minimum values of v_2 plotted with images of x_f and images of cut locations

Figure 6.20: In a) dominant eigenfunction v_2 is plotted for the composite maps $\Pi \circ M_B^\alpha$ with $\Pi = [4321]$, with $\alpha = 0.4$ and $\kappa = 10^{-6}$. Additionally, the distribution of FTLEs is plotted for $j = 8$ iterations below the plots. Black, representing areas of least compression, align with peaks in the concentration field. When $\Pi = [4321]$ the images of the fixed point on a branch of least compression do not align with peaks on the eigenfunction, although there is some agreement with dark regions of the FTLE distribution. In b) maximum and minimum values of the iterated eigenfunction show neither images of fixed points or cut locations predict any dominating features of the concentration field.

which in this example do not directly correspond to the fixed point.

Figure 6.20 (b) again plots the maximum and minimum values of the evolving eigenfunction under further iterations of the transfer operator. For this case, $P^j v_2$ is period two, therefore the maximum and minimums of v_2 and $Pv_2/|\lambda_2|$ need only be plotted. Neither the first few images of the fixed point (top) or the images of the cut locations (bottom) predict any of the distinctive features of the eigenfunction. Although not plotted, further images of the fixed point did not align with the most dominant maximum and minimum peaks in the eigenfunction, even though other regions of least compression in the FTLE distribution do appear to predict some of the largest peaks. Similar to the last example, this result suggests that distinctive features in eigenfunctions of mixed systems of stretching and folding, cutting and shuffling, may be predictable from stretching histories of the underlying system, however numerically accurate predictors for some of least compressed areas, such as images of least compressed fixed points, may not be wholly representable.

6.4 Conclusions

In Chapter 5 it was demonstrated that transformations composed of stretching and folding, and cutting and shuffling could result in slower mixing rates by introducing non-uniformity in striation arrangement from the rearranging of the concentration field. Similar to non-uniformity introduced through non-uniform stretching rates, this is a mechanism which results in long-time exponentially decaying eigenfunctions. In this chapter, the differences in the two mechanisms are explicitly compared. The rearrangement of a concentration field from cutting and shuffling can result in the rejoining of striations of similar species, and has the ability to drastically slow the mixing rates in the small diffusive limit. Only in the fast diffusive limit, improvements in mixing rates were observed on including a discontinuous transformation. This is when the non-uniformity in the mixing rates was particularly diverse and the diffusion coefficient sufficiently large.

It is found that employing the distribution of stretching histories to approximate mixing in the early stages would prove unsuccessful in discontinuous transformations, since the sharp interfaces introduced at each iteration speed up the mixing and are not captured by the stretching histories. This was shown herein where for continuous maps with the same stretching histories distribution the

decay of variance agreed up until the Batchelor scale was reached and the eigen-function regime began, while in contrast the discontinuous compositions from cutting and shuffling led to a diversion in finite time mixing from the first iteration.

Finally, whether any features of the mixing behaviour could be predicted from distributions of stretching histories was addressed by comparing such to the appearance of eigenfunctions which emerge in the long-time limit. Although the eigenfunctions for the cut and shuffle transformations are qualitatively different from their smoothly deforming counterparts, regions of low stretch in the FTLE distributions appeared to approximate well many distinctive maximum and minimum peaks in the concentration field. This contrasts the eigenfunctions in the systems of Chapter 5, where the transformations lacked dynamical features for comparison since FTLE distributions in those systems with uniform stretching rates did not vary across the domain. Since the transformations in this chapter are more realistic, it is of greater interest to understand the contributing factors to the mixing behaviours. More work is needed for a complete understanding of these mixed systems.

Briefly it was shown that decompositions are possible in these system, with initial numerical results suggesting a hypothesis on the parameters which contribute to such. An extensive study to determine analytically the parameters which restrict mixing, and deriving the strong mixing rates in the diffusion-less limit such as that of Byott *et al.* (2013, 2016), would further assist the understanding of the numerical results presented herein. For example, whether a deceleration of mixing rate with increased diffusivity is observed only for the idealised uniform baker's transformation or whether it is trait intrinsic to discontinuous transformations. While results for the non-uniform baker's transformation thus far show non-monotonicity in the mixing rate persists, they do not strongly suggest a slower than diffusion-less mixing rate is likely.

Chapter 7

Mixing by discontinuous transformations in 2D

In previous chapters extensive investigation into mixing properties of one-dimensional discontinuous transformations, and their dependency on the underlying parameters, was carried out. The one-dimensional models were advantageous due to several factors, such as their ease in computation, the availability of analytical results from past literature for comparison, and the simplicity of the concentration field evolution aiding understanding in the mechanisms which drive the observed behaviours. In this chapter the study of mixing in discontinuous mappings is briefly extended to two-dimensional maps as an initial investigation into which phenomena persists in higher dimensions.

Firstly, simple models of planar piecewise isometries with a diffusive step are investigated, namely two-dimensional Interval and Rectangle Exchange Transformations. For two IETs acting in perpendicular directions on the unit torus, one-dimensional mixing dominates the long-time mixing, suggesting the results of Chapter 4 hold in higher dimensions. Similarly, Rectangle Exchange transformations (RETs) which have good stirring properties see a polynomial relation between diffusion and the time to achieve a mixed condition.

Secondly, toral automorphisms, and wave perturbed versions thereof, are used as the underlying chaotic maps to investigate the combination of the three mixing dynamics; stretching and folding, cutting and shuffling and diffusion. When the underlying map is strong mixing, permuting the concentration field at each iteration results in a slower mixing rate. This is again due to the potential to reassemble striations of like colour, leading to increased scales observed in the eigenfunctions. An example is shown in which an improvement in mixing can be
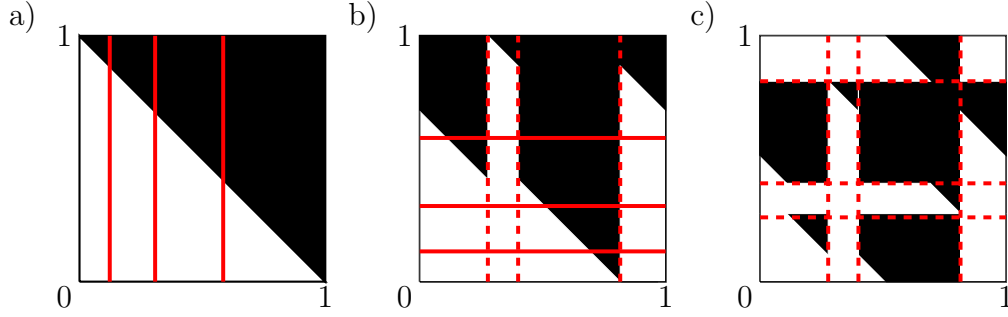
Figure 7.1: A complete iteration of a Double IET with $f^x = f^y = IET_1$ with parameters $r = 1.5$ and $\Pi = [3142]$. Panel a) shows the diagonalised *half-white, half black* initial condition, and panel b) the concentration field after the IET is applied in the x directions, followed by panel c) where the IET is applied in the y direction. The solid red lines represent cuts about to be made, while dashed red are the images of the discontinuities. A decrease in the scale of segregation occurs through the action of cutting and shuffling.

achieved from permutations through the destruction of large elliptic islands from enhanced particle transport.

7.1 Two-dimensional Piecewise Isometries

7.1.1 Double Interval Exchange Transformations

The simplest model extending cutting and shuffling on the unit interval to cutting and shuffling on the unit torus is to apply an IET in one direction, followed by an IET in the other direction. Yu *et al.* (2016) touched on a toy model of such a transformation, looking at the resulting stirring in the initial iterations. They hypothesise that similar properties from the one-dimensional case will hold for two-dimensions. More recently Smith *et al.* (2019) define these transformations as *Double IETs*, discuss their limitations in practical applications, and note that material transport is restricted to the horizontal and vertical slices of each respective point. This latter fact is a contributor to the results which follow.

Explicitly, a Double Interval Exchange Transformation can be defined as a mapping $S : \mathbb{T}^2 \to \mathbb{T}^2$ such that

$$S(x, y) = (f^x(x), f^y(y)) \tag{7.1}$$

where f^x and f^y are IETs with parameters (L_x, Π_x) and (L_y, Π_y) respectively. The superscript denotes the direction in which the interval exchange transformation is applied. In continuation of the parameter space of Chapter 4, for the investigation herein the set of sub-intervals, L_x and L_y, are defined such that adjacent sub-intervals are scaled by the ratio r.

Figure 7.1 illustrates the stirring effect of a Double IET when applying the same IET in both directions with $r = 1.5$ and $\Pi = [3142]$. The initial condition is such that the unit torus is *half-white, half-black*, symmetric about the line $y = x$. Even at the intermediate stage, the scale of segregation is decreased across parts of the domain. Rather than extensively investigating the decrease in the scale of segregation by Double IETs, the results presented here focus instead on the resulting mixing which occurs from including a diffusional step, and reveal why such a study on segregation may not be necessary. The numerical approach to include diffusion is briefly outlined, followed by a handful of results which epitomise the mixing behaviour in such systems.

Iterative Advection-Diffusion Transformation

Applying periodic boundary conditions on the unit torus, the concentration field $c(x, y)$ can be represented by a two-dimensional discrete Fourier expansion with coefficients defined as

$$\hat{c}_{k_1,k_2} = \int_0^1 \int_0^1 c(x,y) e^{-2\pi i k_1 x} e^{-2\pi i k_2 y} \mathrm{d}x \mathrm{d}y. \tag{7.2}$$

A transfer matrix can then be constructed, representing the evolution of $c(x, y)$ by a Double IET through the action of the transformation acting on values of the Fourier coefficients. Since each Double IET is composed of two IETs acting independently in the vertical and horizontal directions, the transfer in discrete Fourier space can be applied first on the action of the wavemodes representing the horizontal shuffle followed by the action on the wavemodes representing the vertical shuffle. Therefore, rather than deriving the transfer matrix acting on each pair of wavemodes (k_1, k_2), a simpler computational method can be devised.

Firstly, represent the values of the Fourier coefficients as a two-dimensional matrix $\boldsymbol{C}_{k_1,k_2}$, such that at each iteration the entries are

$$\boldsymbol{C}_{k_1,k_2}^{(j)} = \hat{c}_{k_1,k_2}^{(j)}. \tag{7.3}$$

For each of the IETs f^x and f^y derive the transfer matrix between Fourier coefficients for one dimension as in (4.12), say $\boldsymbol{M}^x$ and $\boldsymbol{M}^y$ respectively. Then, when
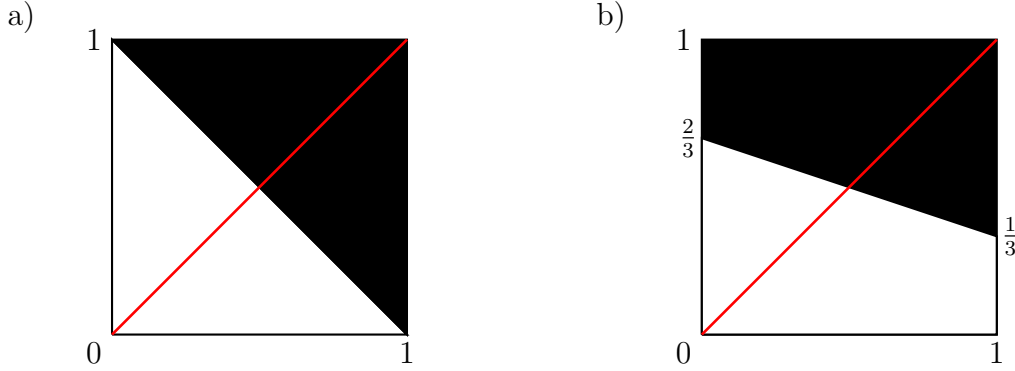
Figure 7.2: Two initial conditions are illustrated a) the diagonalised *half-white half-black*, symmetric about the line $y = x$ (red), and b) *half-white half-black* but non-symmetric about $y = x$.

applying the IET f^x, fix each vertical wavenumber, say $k_2 = n$ and define the vector

$$C_{k_1}^{(j)}(n) = \boldsymbol{C}_{k_1,k_2}^{(j)}|_{k_2=n}, \tag{7.4}$$

and perform

$$C_{k_1}^{(j+1)}(n) = \boldsymbol{M}^x C_{k_1}^{(j)}(n). \tag{7.5}$$

The row is then updated in the original matrix, such that $\boldsymbol{C}_{k_1,n} = C_{k_1}(n)$. One complete application of the IET f^x is then to loop over all values $-Q \leq n \leq Q$ defining (7.4), applying (7.5) and then updating the matrix of Fourier coefficients, where Q denotes the chosen truncation value.

Similarly, for the IET acting in the horizontal direction, fix $k_1 = n$, define $C_{k_2}^{(j)}(n) = \boldsymbol{C}_{k_1,k_2}^{(j)}|_{k_1=n}$, and apply $C_{k_2}^{(j+1)}(n) = \boldsymbol{M}^y C_{k_2}^{(j)}(n)$, looping over all wavenumbers and updating the original matrix. The diffusional step is achieved by multiplying each entry of $\boldsymbol{C}_{(k_1,k_2)}$ by the factor $\exp(-4\pi^2\kappa(k_1^2+k_2^2))$ following the complete stirring iteration.

Numerical results

Mixing results are presented for Double IETs $S = f^x \circ f^y$ with f^x and f^y taken from two IETs; IET_1 with parameters $(r_1, \Pi_1) = (1.5, [3142])$ and IET_2 with $(r_2, \Pi_2) = (1.5, [4213])$. Three combinations are presented; $IET_1^x \circ IET_1^y$, $IET_1^x \circ IET_2^y$, and $IET_2^x \circ IET_2^y$. For all of the examples, the decay of variance was measured as a quantitative degree of mixing and the initial conditions considered shown in Figure 7.2. The initial conditions were plotted on a grid in real-space
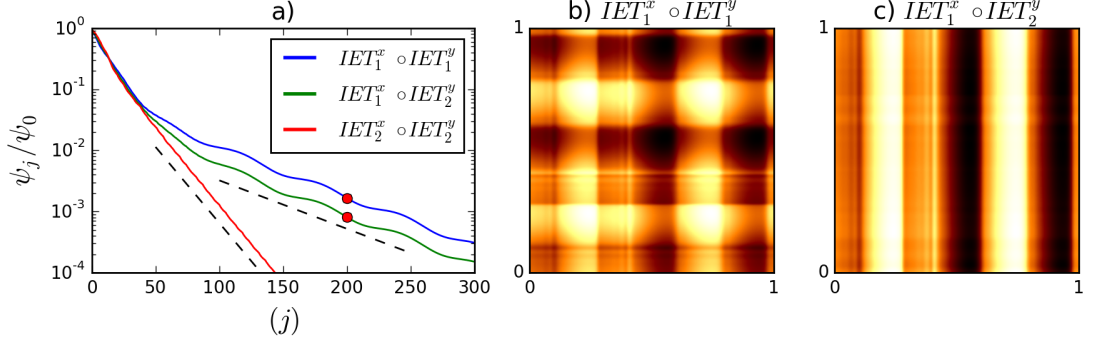
Figure 7.3: Mixing properties for double IET transformations shown from a) the decay of variance when $\kappa = 10^{-5}$ and b) the concentration fields after 200 iterations for two different transformations. The initial condition was Figure 7.2 a) symmetric *half-white half-black*. The red dots on the variance decay profiles indicate the concentration fields shown are representative of the eigenfunction regime. $IET_1^x \circ IET_1^y$ and $IET_1^x \circ IET_2^y$ both have the same decay rate once the eigenfunction regime is reached.
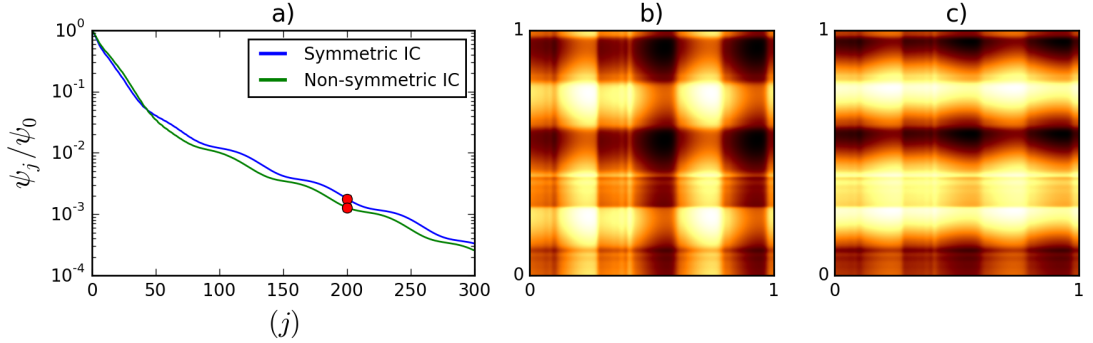
Figure 7.4: Mixing properties for double IET transformations shown from a) the decay of variance when $\kappa = 10^{-5}$ and b) the concentration fields after 200 iterations for the same transformation with a symmetric initial condition and c) a non-symmetric initial condition, shown respectively in Figure 7.2 a) and b). The red dots on the variance decay profiles indicate the concentration fields shown are representative of the eigenfunction regime. Even though the concentration fields look qualitatively different the decay rates are the same. This is due to the perpendicular independence in the mixing of the IETs.

and then a Fast Fourier Transform algorithm used to find approximations for the Fourier coefficients.

Figure 7.3 a) plots the decay of variance for the three combinations of Double IETs when the symmetric initial condition was applied, Figure 7.2 a). Although the variance decay is similar in the first 50 iterations, the mixing for $IET_2^x \circ IET_2^y$ is superior at depleting the variance in the long time. For the combination $IET_1^x \circ IET_1^y$ and the mixed $IET_1^x \circ IET_2^y$ the long-time mixing has the same exponential mixing rate.

Figures 7.3 b) and c) display the renormalised concentration field at iteration 200 for these final two cases. For $IET_1^x \circ IET_2^y$ the concentration field is symmetric about $y = x$, while for $IET_1^x \circ IET_2^y$ the concentration field has reduced to varying mainly in one-dimension, showing little dependency in the y coordinate. These results arise from the orthogonality of the transformations, where the eigenfunctions of each transfer matrix $\boldsymbol{M}^x$ and $\boldsymbol{M}^y$ are also orthogonal to each other. For the case where the IET parameters differ, the dominant eigenfunction v_2^y varying in the y direction, decays at a fast rate leaving the slowest decaying eigenfunction v_2^x, orientated in the x direction, to dominate the concentration field and the long-time mixing rate.

Computing the mixing of IET_1 and IET_2 as individual stirring transformations on the unit interval indeed confirms IET_2 to be superior at depleting variance. The average asymptotic mixing rate for the Double IETs is predicted by the largest of the leading eigenvalues, $|\lambda_2|^x$ or $|\lambda_2|^y$, from the transfer matrices $\boldsymbol{M}^x$ and $\boldsymbol{M}^y$ respectively. The slowest decaying eigenfunction for IET_2 decays quicker than that of IET_1, and essentially produces a one-dimensional mixing regime governed by the long-time dynamics of IET_1. For the case where the stirring is given by $IET_1^x \circ IET_1^y$ the orthogonal eigenfunctions are decaying with the same rate and the resulting concentration field in the long-time is a composition of the eigenfunction in both directions. This is due to the fact that the initial condition was symmetric about $y = x$.

Figure 7.4 compares the variance decay for $IET_1^x \circ IET_1^y$ between having the symmetric initial condition and the non-symmetric initial condition from Figure 7.2. The decay of variance in both cases in nearly identical with the same asymptotic exponential mixing rate. However, the concentration field at iterate 200 shows the composition of the two orthogonal eigenfunctions. Having a non-symmetric initial condition has resulted in one of the orthogonal eigen-
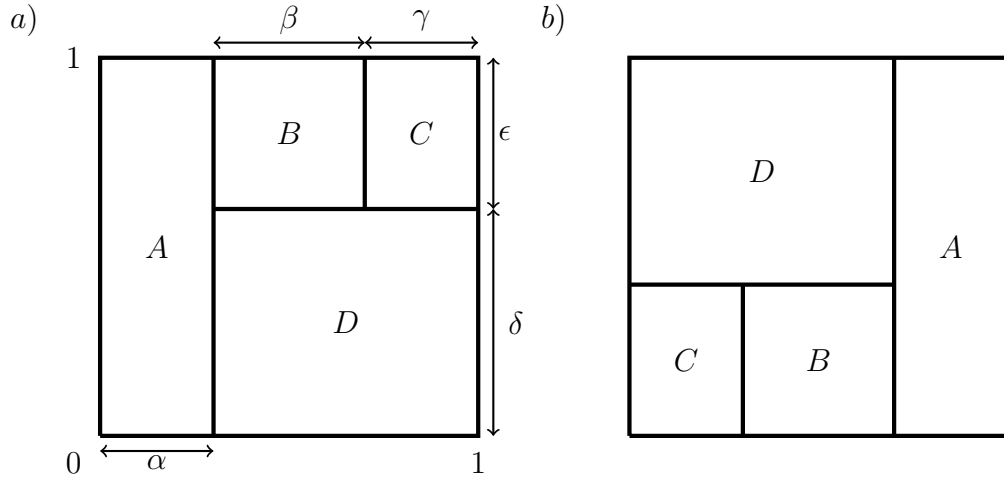
Figure 7.5: One configuration of Haller's rectangle transformation which is potentially weak mixing from initial computational investigations of Haller (1981). The parameter tuple (α, β, δ) generates a full scaling on the transformation acting on the unit torus

functions being less dominant in the concentration field, and the preference for one-dimensionality is observed slightly in its appearance.

This one-dimensionality due to orthogonality suggests that the results on asymptotic mixing rates for interval exchange transformations from Chapter 4 carry over to higher dimensional transformations composed of orthogonal Interval Exchange transformations. Finite time behaviours may vary but investigation into these has been left for future work. In these transformations, two one-dimensional IETs are applied orthogonal to each other with perpendicular mixing dynamics, therefore these results are unlikely to reveal anything about generalised planar piecewise isometries.

7.1.2   Rectangle transformations on the unit torus

The generalised extension of interval exchange transformations to the unit torus are Rectangle Exchange Transformations (RETs). They are obtained by cutting the unit square up into a finite number of rectangular pieces and rearranging the pieces. Unlike interval exchange transformations, the unit torus can not be arbitrarily cut into rectangular pieces which will reassemble under permutation to construct the original domain. Haller (1981) investigated the admissible partitions for 4 subrectangles and determined when they are ergodic. Figure 7.5 illustrates one of the five classes of admissible partitions found by Haller (1981)

with the general partition defined

$$P_A = [0, \alpha) \times [0, 1), \qquad P_B = [\alpha, 1 - \gamma) \times [\delta, 1)$$
$$P_C = [1 - \gamma, 1) \times [\delta, 1), \qquad P_D = [\alpha, 1) \times [0, \delta). \tag{7.6}$$

This rectangle transformation is ergodic if permuted as shown in Figure 7.5, and the parameters α, $\alpha \cdot \beta / (1 - \alpha)$ and 1 are independent over $\mathbb{Q}$; i.e the scaling parameters for the sub-rectangle dimensions are rationally independent. Additionally, following a computational study on the stirring properties of RETs by Hetmanski (2012), it was hypothesised that this particular case could also be weak mixing under the conditions specified, however no rigorous analytical results support this claim. This RET case informs the rest of the discussion.

To investigate the mixing properties of rectangle transformations, parameter tuples (α, β, δ) are chosen for the case illustrated in Figure 7.5. These will define the RET which will be repeatedly applied as the stirring step in the transformation, followed by a diffusional step after each iteration. Rational dependence on the sub-rectangle dimensions α, β, and γ is taken due to computational restrictions. Unfortunately, it is not possible to define a transfer matrix between the coefficients of the discrete Fourier expansion for rectangle transformations, as was done for Interval Exchange transformations. Instead, rational dependence between the sub-rectangle dimensions allows construction of the transformation on a structured grid and a finite difference scheme to be applied for the diffusional step. The full computational approach is described in Appendix C, in which periodic boundary conditions are taken in line with the computational studies of previous Chapters. Although rationality of sub-rectangle lengths means the underlying transformation is not ergodic, and as such no potential to be weak mixing, this idealisation is not a problem herein.

Rather than present results from a full parameter study on mixing in rectangle exchange transformations, which would be a large undertaking for such a system with little research literature associated with it, focus instead lies with the effect of diffusion on the rate of mixing. Appendix C presents results on a small stirring study looking at the effect of changing the parameters α, β, and δ, and the qualitative decrease in the scale of segregation. Figure 7.6 shows three examples of the rectangle transformations which have sufficient stirring properties to ensure good mixing is likely to occur, i.e. not restricted by large stirred regions. The initial configuration of the 4 rectangles are accompanied by the respectively concentration fields after 100 shuffles. The initial condition was a *half black, half white* domain, segregated on $x \in [0, 1/2)$ and $x \in [1/2, 1)$ respectively.
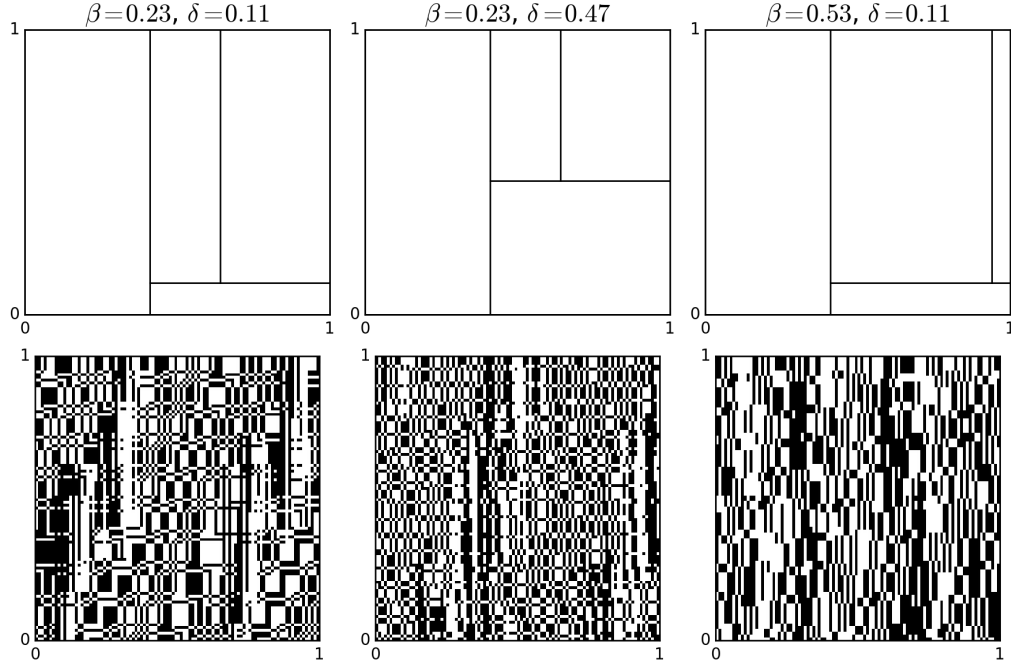
Figure 7.6: Three parameter tuples (α, β, δ) and the resulting decrease in the scale of segregation via the Haller rearrangement shown in Figure 7.5 after 100 iterations. For fixed α, β and δ are specified at the top of each rectangle. The initial condition was black for $x \in [0, 1/2)$ and white for $x \in [1/2, 1)$.
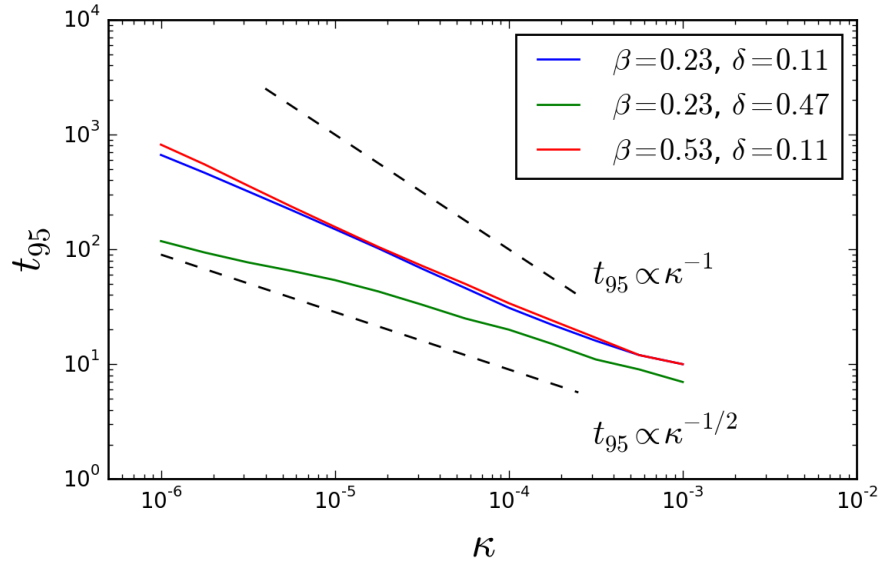
Figure 7.7: Time to achieve a mixed condition t_{95} with varying diffusivity coefficient κ for the three example rectangle transformations shown in Figure 7.6. The profiles have a polynomial scaling with κ shallower than pure diffusion, however $t_{95} \propto \kappa^{-1/2}$ does not appear to approximate the relation well.

For the three examples the time to achieve the mixed condition t_{95} was computed for varying rates of diffusivity coefficient κ. Figure 7.7 shows the results for t_{95} in which all three profiles follow the trend $t_{95} \propto \kappa^{-\beta}$ for $\beta \in (0,1)$. The diffusion-driven trend $\propto \kappa^{-1}$ is included for comparison, and reveals that equivalent to the results in Chapter 4, the polynomial relation has an exponent $\beta < 1$. However, the suggestion that $\beta = 1/2$ may be a suitable scaling does not provide a good approximation for any of the profiles. Profiles for the mixed conditions t_{80}, t_{90}, and t_{99}, were computed and also showed the same polynomial trends for all three examples.

This result suggests that the polynomial scaling relation found in Section 4.5 persists in higher dimensions, but correctly approximating of the exponent β depending on the stirring behaviour is not immediately obvious. An extension of studying diffusion in other planar diffusional piecewise isometries is essential to better the mechanism driving this scaling relation.

7.2 Toral automorphisms composed with permutations

Chapters 5 and 6 investigated the effect of cutting and shuffling on the rates of mixing in stretching and folding transformations with diffusion. The main conclusion from the studies, was that after a short time in which the added sharp gradients assisted the mixing of a concentration field to the mean field, usually discontinuous stirring hindered mixing in the long-time limit. However the one-dimensionality of the systems allowed for unrealistic idealisations, the largest being the alignment of discontinuities with the unstable manifolds of the stretching direction. Whether a similar contamination to mixing rates arises in higher dimensional models is not immediately obvious.

Smith *et al.* (2016) investigated novel particle transport and the resulting stirring from slip deformations in the two-dimensional Re-orientated Potential Mixing flow, and a simplified model capturing the key dynamics, the Cut-Shear-Shear (CSS) map. When the slip deformation counteracted the direction of the shear, the stirring was hindered with island like regions emerging in the flow about periodic points, which were named *pseudo-elliptic* islands. These non-Hamiltonian structures where shown to be identifiable from tracking the evolution of discontinuities as a new diagnostic in mixed systems with discontinuities. The

rates of strong mixing in the CSS maps were investigated in Smith *et al.* (2017b), concluding that cutting could be enhanced or impeded by the presence of the cut.

Rather than deriving a numerical scheme to accurately include diffusion in the Cut-Shear-Shear model, the simplest approach to begin investigating mixing in two-dimensional models is to return to the toral automorphism of Chapter 3 and compose them with permutations, where the mixing can be accurately computed in discrete Fourier space. In these systems the idealisation of alignment between the unstable and stable manifolds with the permutations no longer persists. Similarly, non-uniformity can be easily introduced via the wave-perturbed versions of the maps. An initial example of the wave-perturbed cat map composed with a single permutation concludes the results of the Chapter.

7.2.1 Mathematical description and formulation

The toral automorphisms of interest are the cat map, and the double application of the cat map,

$$M_C = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, \qquad M_{CC} = \begin{pmatrix} 5 & 3 \\ 3 & 2 \end{pmatrix}, \tag{7.7}$$

respectively. Both maps are strong mixing and have eigenvalues $\lambda_{M_C} = (3 \pm \sqrt{5})/2$ and $\lambda_{M_{CC}} = (7 \pm \sqrt{45})/2$. The rate of stretch by the map M_{CC} is much larger than that of the cat map, over 2 fold. As such M_{CC} is used as a comparison map to see the general effect of increasing the rate of stretch when composed with permutations.

The computational approach is to solve the transfer operator for the toral automorphism first, as in Section 2.3.4, and apply the permutation and diffusional step by the approach of Section 7.1.1. The permutation transformation is only performed in the x direction, and chosen from $\Pi \in S_N$ with $N = 3$, 4 and 5 for permutations of equal sized cells. A permutation can be easily applied in the y direction, or two IETs applied in both directions, at each iterative step but for simplicity of discussion only the x direction is presented here. None of the chosen permutations composed with the toral automorphisms produced a decomposition of the domain.

A handful of numerical results are presented which capture the overall trends in the transformations considered. The discussion evolves around the results present here, and the similarities and differences to main results of Chapter 5, in which uniformly stretching transformations were composed with permutations. Namely, whether non-uniformity of concentration rearrangement results in a slow
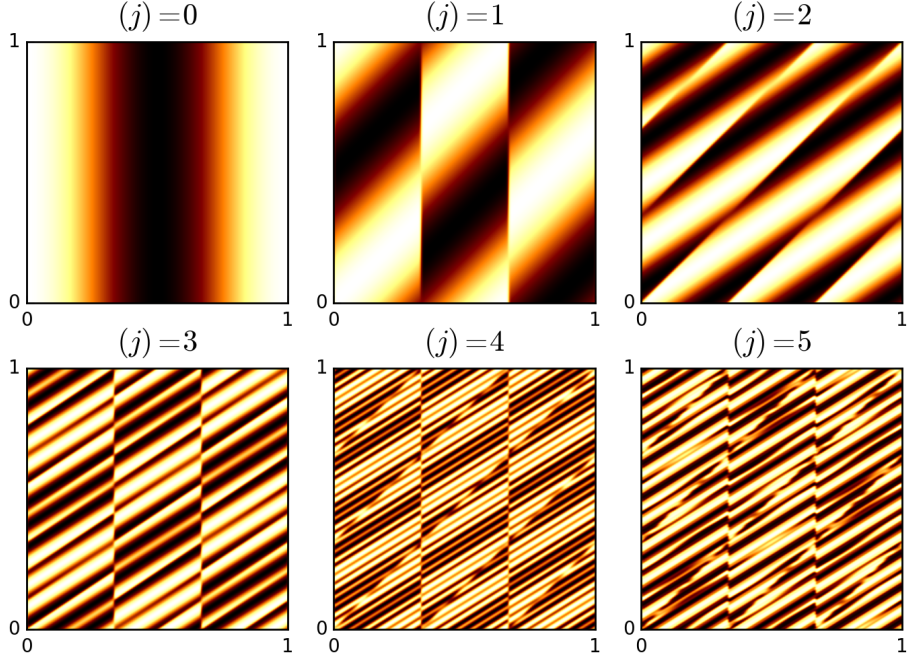
Figure 7.8: Evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ by transformation $\Pi \circ M_C$ with $\Pi = [132]$ and diffusion coefficient $\kappa = 10^{-5}$. The concentration field is renormalised at each iteration to highlight the spatial features.

down of mixing rate, and whether this mixing rate changes monotonically or non-monotonically with diffusion. Additionally the effect of increased stretch factor is compared, comparative to choosing the value of m in the extension of the uniform baker's transformation. No analytical results on the mixing rates were derived from the cat map, or double cat map, composed with permutations.

7.2.2 Mixing in uniformly stretching toral automorphisms composed with permutations

Snap-shots of the concentration field evolution for a single example provide a qualitative means to understand the resultant mixing. Figure 7.8 shows the evolution of the initial condition $c^{(0)}(x, y) = \cos(2\pi x)$ under the action of the cat map M_C and the permutation $\Pi = [132]$ acting in the x-direction. The diffusion coefficient was $\kappa = 10^{-5}$ and the concentration field renormalised and shown without scale at each iteration, to highlight the changing scales across the concentration field striations, rather than the reduction in the height of the concentration field as it approaches the mean field.

In line with the observations of Chapter 5, where the cutting and rearranging
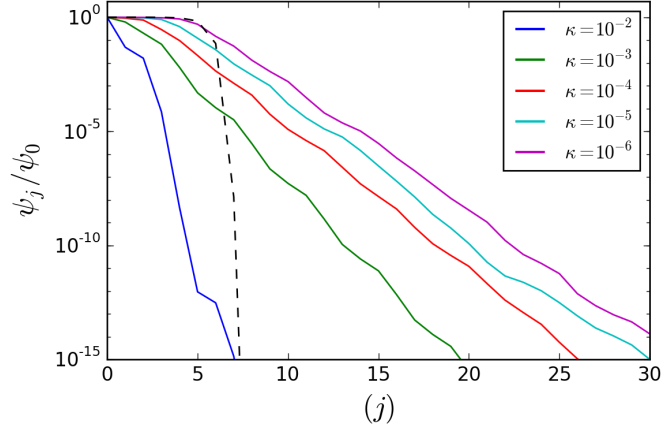
Figure 7.9: Decay of variance $\Pi \circ M_C$ with $\Pi = [132]$ for varying diffusivity values κ. The initial condition $c^{(0)}(x) = \cos(2\pi x)$. Black dashed line represents the decay of variance for the cat map M_C alone with $\kappa = 10^{-6}$.

of the concentration field rejoined striations of similar colour, a similar mechanism is in mixing contamination can be seen in Figure 7.8. This is most clearly shown by iteration $j = 3$, where regions of mostly white, and regions of most black are easily distinguishable throughout the domain. By iteration $j = 5$, an eigenfunction forms, where smaller striations are starting to be washed out by the balance between shearing and diffusion.

Figure 7.9 plots the decay of variance for the underlying transformation shown in Figure 7.8, $\Pi \circ M_C$ with $\Pi = [132]$, for varying diffusivity coefficients κ. Also included as a dashed black line is the decay of variance for the cat map M_C mixing alone with $\kappa = 10^{-6}$. The addition of the permutation results in long-time exponential mixing, due to the non-uniformity in the striation widths from the rearrangement of the concentration field at each iteration. Recall that for the maps $\Pi \circ M_B$ with $\Pi \in S_3$, non-monotonicity in the long-time mixing rate $|\lambda_2|$ with κ was reported, additionally observed in the variance decay profiles. Here it appears as though $|\lambda_2|$ may vary monotonically.

The values $|\lambda_2|$ can be approximated from the average decay rate of the variance profiles as a comparison across finer values of κ. Figure 7.10 shows the profiles $|\lambda_2|$ computed for $\Pi \circ M_C$ with the interleaving permutations $\Pi \in S_3$. All produce the same long-time mixing rates and non-monotonicity in the values of $|\lambda_2|$ are observed. There appears to be a convergence to a mixing rate with $\kappa \to 0$, the maximum value across the profile being when $\kappa = 10^{-6}$.

Similarly, Figure 7.11 plots the approximated values of $|\lambda_2|$ for $\Pi \circ M_C$ with (a)

207
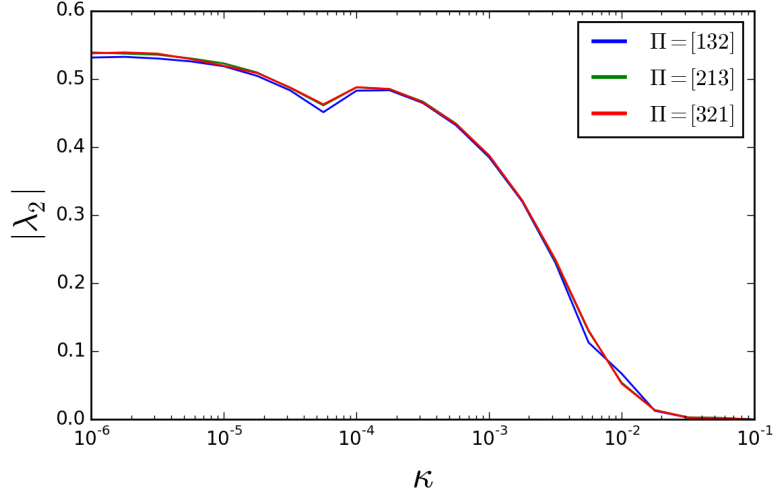
Figure 7.10: The modulus of the leading eigenvalues $|\lambda_2|$ against diffusivity coefficient computed from the average decay rate of the variance profiles. The three interleaving permutations $\Pi \in S_3$ have the same mixing rates.

(a) $|\lambda_2|$ approximation for $\Pi \in S_4$

(b) $|\lambda_2|$ approximation for $\Pi \in S_5$
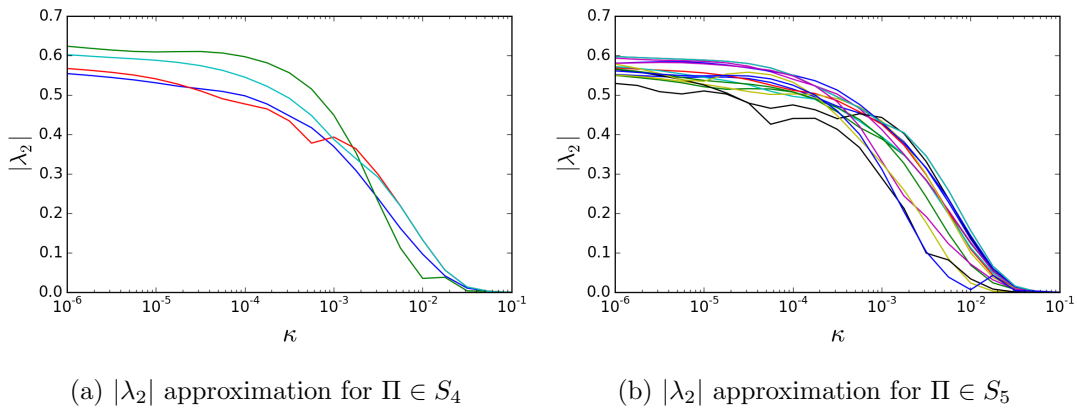
Figure 7.11: The modulus of the leading eigenvalues $|\lambda_2|$ against diffusivity coefficient κ, computed from the average decay rate of the variance profiles for a) $\Pi \in S_4$ and b) $\Pi \in S_5$.

(a) Comparison of $|\lambda_2|$ for $\Pi \in S_3$

(b) Comparison of $|\lambda_2|$ for $\Pi \in S_4$
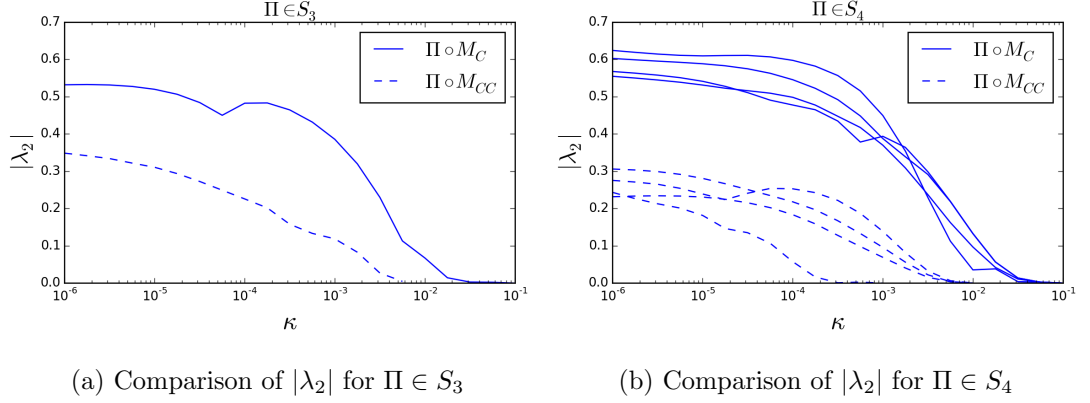
Figure 7.12: The modulus of the leading eigenvalues $|\lambda_2|$ against diffusivity coefficient κ, computed from the average decay rate of the variance profiles for a) $\Pi \in S_3$ and b) $\Pi \in S_4$ for both the cat map composed with permutations $\Pi \circ M_C$ (solid) and the cat map applied twice, $\Pi \circ M_{CC}$ (dashed). The additional stretch from applying the cat map twice always produces faster mixing rates.

$\Pi \in S_4$ and (b) $\Pi \in S_5$. As with the interleaving permutations of S_3, symmetry in the advective transformations results in subgroups emerging from the permutation groups S_N with the same long-time mixing rates. An example profile is plotted for each subgroup, showing the range of mixing rates across κ. As with S_3, non-monotonicity in the value of $|\lambda_2|$ occurs in some of the profiles as $\kappa \to 0$, and for nearly all of the profiles, there appears to be a convergence in the value of $|\lambda_2|$ as $\kappa \to 0$.

In Chapter 5 the most surprising result was that non-monotonicity in the value of $|\lambda_2|$ sometimes achieved $|\lambda_2| > \tau$ for non-zero κ, where τ is the diffusion-less mixing rate. Since no analytical results on the value of τ were found for the maps $\Pi \circ M_C$, it is assumed that $|\lambda_2|$ when $\kappa = 10^{-6}$ gives a good approximation since convergence seems to be happening for small values of κ. In contrast to the results of Chapter 5, either no overshoot of τ for non-zero κ happens for these systems, or if it is occurring the effect is less dramatic.

As a final comparison study to the results of Chapter 5, the profiles of $|\lambda_2|$ with κ can be compared for $\Pi \circ M_C$ and $\Pi \circ M_{CC}$, to see whether having a larger stretch is more likely to result in faster long-time mixing. Figure 7.12 plots the profiles of the long-time mixing rate for each of the subgroups of $\Pi \in S_3$ and $\Pi \in S_4$, for the composition maps $\Pi \circ M_C$ (solid) and $\Pi \circ M_{CC}$ (dashed). The subgroups between $\Pi \circ M_C$ and $\Pi \circ M_{CC}$ differed, showing that again it is the action of both stretch and folding, and cutting and shuffling, which determines

209

the resultant long-time mixing rate. The subgroups are included as one-profile for each of the maps.

All of the profiles for $\Pi \circ M_{CC}$ have values of $|\lambda_2|$ smaller than $\Pi \circ M_C$ across all values of κ. The additional stretch at each iteration arising from applying the cat map twice will reduce the scales in the concentration field of the eigenfunction significantly, allowing fast depletion from diffusion. Although the cutting and re-arranging of the concentration field by the permutation has the potential to rejoin regions of similar concentration, stretching is the dominant mixing mechanism in the transformation. This outcome mirrors that of the uniform baker's transformation with additional branches, where a larger m speeds up the mixing rate on average; more easily understood by the $1/m$ factor in the bounds of mixing from Byott *et al.* (2013).

7.2.3 Mixing in the wave perturbed cat map composed with permutations

Chapter 6 presented a one-dimensional study on the combined stirring dynamics of non-uniform stretching with cutting and shuffling. Having a combination of both mixing mechanisms which produce non-uniformity in the widths of striation evolution allowed for a direct comparison between the mechanisms of contamination; where cutting and shuffling can significantly slow down mixing rates by generating large scale non-uniformities in the concentration field from striation reassembly, in comparison to contamination arising from slow compression of striations alone.

For the toral automorphisms, introducing the wave-perturbation increases the complexity in the computation of the transfer operator, directly increasing the computation time. Therefore a parameter study in such systems is left for future work. Instead the wave perturbed cat map composed with a single permutation $\Pi = [132]$ is used as an initial investigation into the combined stirring dynamics in two-dimensions. As before, the computations are carried out in discrete Fourier space, applying the wave perturbed cat map first, via the numerical method of Section 2.3.4, followed by the permutation in the x direction and diffusional step of Section 7.1.1.

Figure 7.13 plots the decay of variance for 6 different maps; the wave perturbed cat map M_{PC}, and the composition map $\Pi \circ M_{PC}$ with $\Pi = [132]$, for three values of K. The decay of variance for the maps M_{PC} are plotted as dashed lines, while

(a) (b)
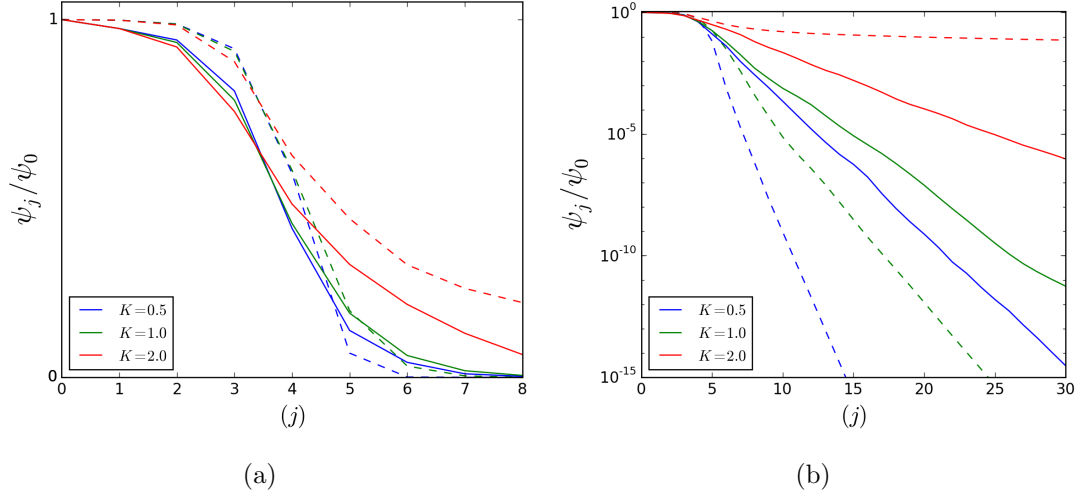
Figure 7.13: Decay of variance for the cat map with an added wave perturbation with three values of K (dashed) and when the permutation $\Pi = [132]$ is applied as a composition transformation $\Pi \circ M_{PC}$ (solid). (a) shows the initial iterations on a linear-linear axis, while the long-time exponential mixing behaviour is more easily observed in (b) on a linear-log axis. For smaller values of K the permutation slows down the decay of variance, while when $K = 2.0$ the permutation improves the mixing rate.

(a) M_{PC} (b) $\Pi \circ M_{PC}$, $\Pi = [132]$
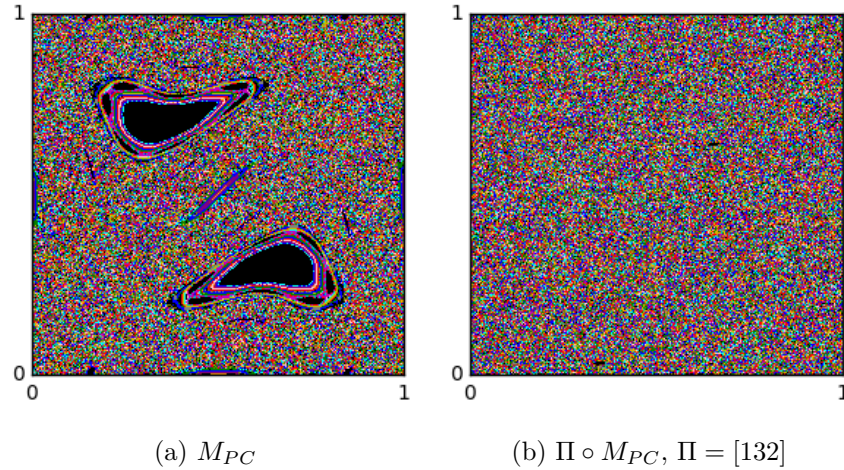
Figure 7.14: Poincarè sections for (a) the cat map wave with $K = 2.0$, a number of islands are seen to decompose the domain, while in (b) when the permutation is applied the islands are destroyed leaving only potentially small decompositions in the domain. 1000 iterations of 200 points are plotted, initialised along the line $x = 1/2$.
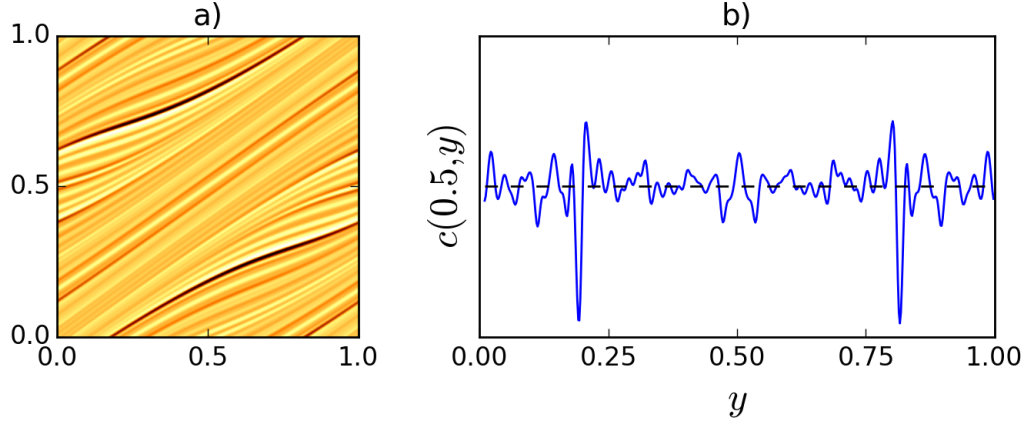
Figure 7.15: The concentration field for the wave-perturbed cat map M_{PC} with $K = 1.0$ and diffusivity coefficient $\kappa = 10^{-5}$ after 15 iterations shows distinctive peaks where the least compression occurs just this stirring stage. At this iteration the mixing is in the eigenfunction regime. In b) a slice of the concentration field at $x = 0.5$ confirms that dominant peaks of small width occur across the domain. The initial condition was $c(x, y) = \cos(2\pi x)$.

the composition maps are plotted as solid lines. The assigned colours correspond to the values of K in each. For all composition maps, the mixing in the first handful of iterations is faster than for the maps M_{PC}. When $K = 0.5$ and 1.0, the addition of the permutation slows down the speed of mixing in the long-time limit. It is only for $K = 2.0$ that the mixing rate is improved by the addition of the permutation throughout all times.

The latter can be easily understood, since for $K = 2.0$ the map M_{PC} has two large islands which restrict transport throughout the domain. These are observed in Figure 7.14 (a) from a Poincaré section plotting 1000 iterations of 200 points initialised along the line $x = 1/2$. Long-time mixing is restricted to diffusion across the boundaries of the islands. When including the permutation $\Pi = [132]$, the large islands are destroyed and no other sizeable islands visible, shown in Figure 7.14 (b) for $\Pi \circ M_{PC}$. Small islands may remain within the domain, not visible on the scale of this plot. Since large decompositions of the domain no longer exist, fast mixing can occur via all mechanisms of stretching, folding, cutting, shuffling and diffusion. These initial results suggest that an improvement to mixing from cutting and shuffling can only be achieved when the underlying transformation has bad stirring properties, reminiscent of the conclusions from Chapter 6 on one-dimensional non-uniform stirring.
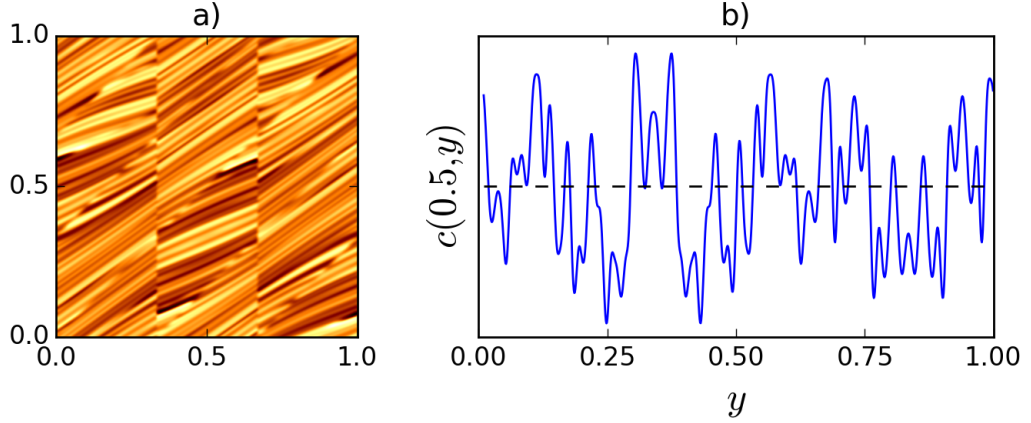
Figure 7.16: The concentration field for the wave-perturbed cat map and permutation $\Pi \circ M_{PC}$ with $K = 1.0$, $\Pi = [132]$ and diffusivity coefficient $\kappa = 10^{-5}$ after 15 iterations. At this iteration the mixing is in the eigenfunction regime. In b) a slice of the concentration field at $x = 0.5$ reveals that the concentration field has regions dominated either by all positive or all negative values, with sharp transitions between these striations with large width. The initial condition was $c(x, y) = \cos(2\pi x)$.

Finally, the appearance of the eigenfunctions and comparisons to the stretching histories are addressed for this example. Figure 7.15 a) plots the full concentration field at iteration 15 under the action of the map M_{PC} with $K = 1.0$ for the initial condition $c(x, y) = \cos(2\pi x)$ and $\kappa = 10^{-5}$. At iteration 15 the concentration field has reached the eigenfunction regime, where on average the decay rate is exponential but oscillations in the modulus of the leading eigenvalue and the concentration field occur. The eigenfunction is shown renormalised, using the same colour map as in Chapter 3, with white and black representing maximum and minimum values respectively, and orange representing the mean field. Peaks and troughs in the concentration field exist as thin bands throughout the domain, shown in Chapter 3 to align regions which have experienced the smallest stretching histories. A slice of the eigenfunction is taken through $c(0.5, y)$, and plotted in Figure 7.15 b), where distinctive peaks that are small in width, are the dominant features of the concentration field. Across most of the domain, the concentration field is well mixed and close to the mean field, with only extreme values found in the peaks.

Figure 7.16 a) shows the eigenfunction observed in the map $\Pi \circ M_{PC}$ with $K = 1.0$, $\Pi = [132]$, for the initial condition $c(x, y) = \cos(2\pi x)$ and $\kappa = 10^{-5}$, and Figure 7.16 b) the slice taken at $c(0.5, y)$. Large regions of either all positive,

(a) FTLEs

(b) Eigenfunctions
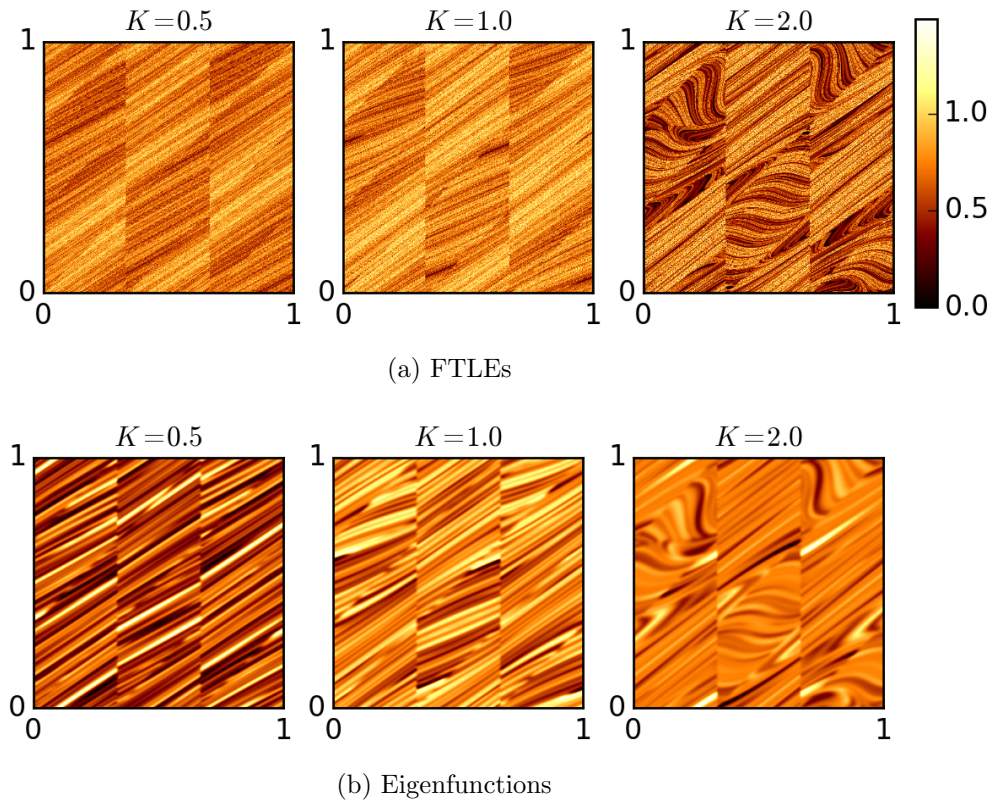
Figure 7.17: Comparison between a) the distribution of FTLEs within the domain under the transformations $\Pi \circ M_{PC}$ with different K and $\Pi = [132]$, and b) the respective concentration fields after 15 iterations. Regions of least compression highlight areas which exhibit peaks in the concentration field, particularly when $K = 2.0$ and the non-uniformity in the stretching is the largest.

or all negative values are seen across the slice where values in the concentration field are comparable in size across the striations, without a single peak or trough emerging as a dominant feature. This coincides with the observations from the one-dimensional case, where contamination from striation reassembly dominates the long-time mixing.

Figure 7.17 (a) plots the FTLE fields for the maps $\Pi \circ M_{PC}$ for different K and (b) a comparison to the concentration fields at iteration 15. The regions of low stretch align well with the darkest and lightest regions of the concentration field, while well mixed regions are well approximated by areas which have experienced the largest stretching histories. When $K = 2.0$ and the non-uniformity in the stretching rates is the largest, the concentration field is more well mixed across the domain in comparison to $K = 1.0$. The largest peaks in the eigenfunction align well with the darkest regions in the FTLE plot. This example suggests that while cutting and shuffling are likely to slow down mixing from increasing scales in the concentration field via striation reassembly, potentially when there is large variation in the stretching histories of the flow, areas of least compression dominate the contamination to mixing. As with smoothly deforming systems, the distribution of FTLEs approximate well the appearance of the eigenfunction.

7.3   Conclusions

An initial investigation into mixing in two-dimensional, discontinuous transformations was presented as an accompaniment to the one-dimensional models of earlier chapters. While the one-dimensional models were open to rigorous analytical and numerical investigation, they are highly idealised in comparison to real laminar mixing flows.

Two-dimensional interval exchange transformations were observed to reduce to one-dimensional systems in the presence of diffusion, due to the composition of two transformations which are stirring in perpendicular directions only. Therefore it is assumed that the results of Chapter 4 on the long-time mixing rates can be readily applied in two-dimensional or higher dimensional IETs, as defined by Smith *et al.* (2019).

Mixing in another planar piecewise isometry was investigated, an RET of Haller (1981). A single case of rectangle rearrangement was chosen which could result in qualitatively well stirred concentration fields if the scaling parameters defining the rectangular shapes satisfied certain conditions, outlined in Appendix

C. The relation $t_{95} \propto \kappa^{-\eta}$ with $\eta < 1$ was observed for all three parameter tuples (α, β, δ) considered, but $\eta \sim 0.5$ did not approximate the scaling well. Future work should focus firstly on a thorough parameter space investigation of the stirring resulting from different values of α, β and δ, akin to the work of Krotter *et al.* (2012) and Yu *et al.* (2016) for IETs. This should be followed by an in depth study on finite and long-time mixing rates in such systems, and whether there is sensitivity to the initial condition as was observed for IET stirring systems in Chapter 4 but not explored here for RETs.

Toral automorphisms with uniform stretching, composed with permutations mix slower than the non-composed maps in the long-time limit. Similar to their one-dimensional counterpart, the baker's transformation composed with permutations, rearrangement of the concentration field by the discontinuous transformation generates non-uniformity in striation arrangement by rejoining regions of like colour. In the long-time limit, eigenfunctions form with exponential decay rates. The asymptotic mixing rates similarly vary non-monotonically with diffusivity coefficient. However lacking knowledge of the diffusion-less mixing rates, it can not be concluded whether a counter-intuitive deceleration of mixing rate with diffusivity is observed.

Composing the wave perturbed cat map with an example permutation results in contamination to mixing, expect for when the composition assists fluid transport, destroying islands. The scales in the concentration field are increased from cutting and shuffling, resulting in eigenfunctions which are larger in scale than the eigenfunctions arising from purely non-uniform stretching and folding. The observations in this two-dimensional model are analogous to those of Chapter 6, when the stretching and folding mechanism is provided by one-dimensional non-uniform baker's transformation. However, only a limited parameter space was studied due to the computational cost of solving the iterative advection-diffusion problem for these transformations.

There has been only a handful of studies investigating the dynamics in discontinuous, chaotic advective transformations in two or more dimensions [Smith *et al.* (2016, 2017b); Vaienti (1992)]. Derivations of analytical results on the ergodic properties in such systems are likely to be non-trivial. This chapter presented a brief numerical investigation into mixing arising in two dimensional systems where discontinuities where introduced via chopping and rearranging the concentration field following a continuous stirring transformation. Rigorous proofs of ergodic properties, rates of stirring and mixing, and the role discontinuities play

in these results should be the focus of future work. A range of discontinuous maps should be the focus of these investigations, such as those where a discontinuity arises via slip deformation, which have not been addressed in the current work.

Chapter 8

Conclusions

The primary aim of this thesis was to illuminate the impact discontinuous particle transport has on the rates of mixing of passive scalar fields. Where previous studies have mainly investigated the effect discontinuities have on organisation of particle trajectories, and rates of decrease in the scale of segregation, the novelty of this work was to include diffusion and study the rates of mixing to uniformity in a range of discontinuous models, with and without chaotic advection.

In fulfilment of this aim, models were devised using discrete time maps which allowed for the derivation of transfer operators capturing the evolution of concentration fields, and a diffusive step easily introduced. Computational methods and analysis techniques were introduced in Chapter 3, alongside an introduction of discrete time maps with smooth velocity fields, and results presented on their mixing rates in relation to the underlying dynamics of the flow.

In Chapter 4 the action of cutting and shuffling transformations to achieve homogenization was investigated using a one-dimensional computational model. The model utilised the parameter space of Krotter *et al.* (2012) for Interval Exchange transformations, and transfer operators derived from transfer matrices computed for permutations of equal sized cells from Ashwin *et al.* (2002). The long-time mixing rates are exponential, however since IETs lack strong mixing properties, the interaction of slowly decaying eigenfunctions dominate the finite time mixing profiles in the limit of small diffusivity. This fact was shown to contribute to large variations in the time to achieve a desired mixed state from small variations in the initial condition. A parameter study highlighted that the dynamics which on average achieve the fastest rates of mixing, are those which introduce more cuts at each iteration, and ensure there is no significant decomposition of the domain. This has the dual effect of increasing gradients in the

concentration field via sharp discontinuities, as well as decreasing the scale of segregation at each iteration. The time to a mixed condition was shown to be polynomial with $t_\% \propto \kappa^{-\eta}$, where $\eta < 1$ when IET parameters satisfied those with good stirring criteria. This relation was reflected in the changing profiles of the leading eigenvalues with diffusion, suggesting that the balance between cutting, shuffling and diffusion could be the main driving mechanism for the result. If the mixing rate is governed by the scale of the concentration field balanced by the mixing mechanisms, this predicts $\eta \sim 0.5$, which shows good agreement with numerical results.

Chapter 5 began the investigation into the impact of discontinuous transport when the main stirring mechanism is from exponential stretching, the hallmark of chaotic advection. This initial study implemented the one-dimensional approximation of the baker's transformation with uniform stretching rates composed with permutations. When the permutations do not generate a decomposition of the domain, exponential stirring is maintained, found by extending the analytical results of Byott *et al.* (2013). The exponential rates of stirring are either preserved, or slowed by the permutations. On the inclusion of a diffusive step, the asymptotic mixing rates are exponential due to the emergence of eigenfunctions, generated by the non-uniformity in striation arrangement arising from the chopping and rearranging of the concentration field following the exponential stretching and folding. Mixing rates are sensitive to striation arrangement [Clifford *et al.* (1999)], therefore the mixing rates vary significantly across the choice of permutations. Only in the initial stages do the discontinuous transformations assist mixing over the uniform baker's transformation, by introducing sharp interfaces which help the depletion of variance before the Batchelor scale is reached from shearing. A counter-intuitive observation is that the asymptotic mixing rates vary non-monotonically with diffusivity rate for many of the composition transformations, with some mixing slower than their diffusion-less limit.

The emergence of eigenfunctions via the variation of striation widths also occurs in smoothly varying velocity fields with non-uniform stretching rates, and can be described as a global mechanism of mixing. Chapter 6 investigated the mixing behaviour in the non-uniform baker's transformation composed with permutations, creating a system in which both stirring mechanisms resulting in the global regime could be compared: non-uniform stretching and cutting and shuffling. As was previously observed, the depletion of variance is encouraged in the initial iterations from the sharp interfaces introduced by the discontinuous

transformations. However, when the non-uniformity in the stretching rates is not significant, permuting the concentration field at each iteration leads to significant reduction in the asymptotic mixing. Qualitative assessment of the resulting eigenfunctions and evolution of the concentration in finite time, reveals that cutting and shuffling the concentration field can lead to reassembly of striations of like concentration, contaminating the mixing by significantly increasing scales of the concentration field. While it is commonly understood that stretching histories are not able to predict asymptotic mixing rates arising from a global mechanism, in non-uniform smooth stretching flows, they have been reported to approximate well the early stages of mixing, prior to the Batchelor length scale being reached [Wonhas & Vassilicos (2002)]. It is shown herein that in discontinuous transformations, stretching histories can no longer be used to approximate variance depletion at any stage of mixing, however there is some evidence to suggest that areas of least compression may dominate the long-time concentration field evolution in the eigenfunction regime.

In Chapter 7, a handful of introductory, two-dimensional models with discontinuous material transport were presented. The two-dimensional extension of IETs were shown to reduce to one-dimensional mixing in the long-time limit, due to the perpendicular action of cutting and shuffling. Therefore, a Rectangle Exchange transformation was employed to initiate investigation into mixing from planar PWIs without this idealisation. A polynomial relation on the time to achieve a mixed condition with diffusivity rate is observed when the scaling parameters of the transformation led to good stirring in finite time. When modelling the composition of the three mixing mechanism in toral automorphisms composed with permutations, striation reassembly leads to mixing contamination and the emergence of eigenfunctions. This is similarly observed on including a wave perturbation to generate non-uniformity in the stretching and folding element of the transformation. Observations mirror those of Chapter 6, where scales in the concentration field are larger compared to the smooth velocity field counterparts, but regions of low compression still qualitatively capture dominant features of the concentration field. An improvement in mixing is achieved when the composition of a discontinuous transformation aid particle transport, extending stirring throughout the domain.

Discrete time modelling using maps satisfied the requirement of periodic dynamics to allow the derivation of transfer operators, however, modelling mixing via operator splitting is an idealisation. Similarly, the velocity fields lacked a

representation of walls and boundaries, and only cut and rearrange discontinuous transformations were considered due to their ease in deriving transfer matrices, which was essential for reliable computational modelling. In spite of these limitations, the use of discrete time maps has provided substantial insight into the driving mechanisms of mixing rates governed by discontinuous stirring fields.

Previously, the decay of variance from IETs with a diffusive step was well approximated by a stretched exponential function [Wang & Christov (2018)]. The results herein suggests that the mechanism leading to this observed finite time mixing rate is due to the interaction of slowly decaying eigenfunctions. This is a product of the fact IETs are not strong mixing, and as such the transfer operators lack a spectral gap, with the modulus of all isolated eigenvalues tending to 1 in the diffusion-less limit.

While Krotter *et al.* (2012) and Yu *et al.* (2016) suggested criteria for good finite time stirring, a parameter study revealed that some of the criteria could be relaxed when a diffusive step was included, and in contrast to the suggestions of Yu *et al.* (2016), that 4 sub-intervals is enough for good stirring, the number of cuts introduced at each iteration was found to be one of the best contributors to faster mixing rates. However, the criteria that rearrangement of intervals should follow an irreducible, non-rotational permutation is essential, which is physically analogous to reducing large decompositions in the system. This latter condition has been the main consideration when investigating the parameter space of rotation angles in spherical tumblers, where a larger coverage of cutting lines is ascribed to the potential for better mixing [Park *et al.* (2016, 2017); Smith *et al.* (2017a)]. However, the conclusions of the study herein suggest that these other studies may be fruitless if diffusion is present, since their aim is to reduce almost all visible decompositions, which it is found here can be relaxed.

In contrast to chaotic systems which are weakly related to Péclet number, the mixing rates of PWIs both in one and two dimensions were shown herein to have a polynomial relation to the diffusivity rate. Also it was shown that in the low diffusivity limit, the time to a mixed condition is sensitive to initial concentration arrangement. Therefore when investigating parameter optimization in cut and shuffle systems a wide range of initial conditions would be required. For small diffusivity rates, cutting and shuffling does not present itself as an efficient method of stirring alone. The most successful split-and-recombine micro-mixers employed configurations which mimic exponential stretching [Hobbs & Muzzio

(1997); Schönfeld *et al.* (2004)], and the results of this thesis suggests that indeed this is the best approach.

In discontinuous systems where the dominant mechanism driving the decrease in the scale of segregation is from chaotic stirring, results of this thesis suggest that discontinuities can both improve and contaminate mixing. If islands are present in the flow, discontinuous permutations can be applied to enhance particle transport, and potentially achieve a chaotic sea with almost full measure. In the literature, slip deformations were similarly shown to increase the possibilities of particle transport [Smith *et al.* (2016)], and cutting lines shown to destroy boundaries of islands on contact [Smith *et al.* (2017c)]. In those flows, the slip deformations are a consequence of the stretching dynamics, where here the results suggest that adding a discontinuous step into a stirring protocol could assist mixing, similar to how rotating walls can prevent contamination for no-slip boundaries and increase mixing in the bulk flow [Gouillart *et al.* (2010); Thiffeault *et al.* (2011)]. In all results presented, even when the chaotic stirring did not result in a decomposition, discontinuities increased mixing in the early iterations from the introduction of sharp interfaces into the concentration field. Therefore if the majority of mixing occurs before the Batchelor scale is reached, then a discontinuous transformation in a predominantly stretching velocity field could speed up the homogenization process to a mixed condition.

However, the results of the models developed in this work showed that the mixing rate can be severely contaminated when striations of like concentration are reassembled by cutting and shuffling. Although the reassembled striations are stretched and folded by the chaotic advection, they were shown to also contaminate long-time mixing rates. This is reminiscent of the deceleration in exponential rate of decay of correlations reported for expanding maps composed with permutations [Byott *et al.* (2013, 2016)]. The contamination is also driven by the inclusion of diffusion, which blurs the reassembly of like striations generating increased scales in the emerging eigenfunctions. Present in both one and two-dimensional systems, the mechanism is essentially dependent on the initial condition and unpredictable from stretching histories.

This additional contamination from blurring striations together with diffusion, and increasing scales of the concentration field, could be the reason for the counter-intuitive observation of deceleration in asymptotic mixing rate with increased diffusivity. Although non-monotonicity has been reported in other maps in the approach of the leading eigenvalue to the diffusion-less strong mixing rate

τ, herein lies a novel observation of significant overshoot in $|\lambda_2|$ in models where τ could be analytically derived. Whether similar observations would be made in other systems, and whether a defining feature of these systems is non-continuous or non-differentiable velocity fields, could be of interest to the wider dynamical systems literature. For example, would systems driven by stochastic diffusion, such as Brownian motion of particles, see similar scale increases in density profile evolution when particles can jump discontinuously?

An additional impact of this results lies in the approach many studies in the literature take, where numerical approximations of transfer operators are commonly computed when analytical derivation may be too complex. The mapping method, an extension of Ulam's method, is a common computational approach outlined in Section 2.3.3. Including diffusion is problematic in such methods and in some cases the action of numerical diffusion is relied upon as a basic approximation [Schlick *et al.* (2015); Singh *et al.* (2008b)]. In applications of optimizing mixing, such as in micro-fluidic flows, if comparisons across discontinuous stirring protocols were computed neglecting the true diffusivity properties of the fluid to be stirred, an under or over approximation of the time to achieve the desired result may occur. This could result in the correct mixing criteria not achieved in the predicted time in real application.

8.1 Suggestions for future work

Preliminary results on mixing rates from Double IETs and RETs were presented in this thesis, however these PWIs are idealised in some sense. The cutting lines in the transformations are parallel to the x and y axes, and they could be thought of as some of the simplest manifestations of planar piecewise isometries. The advantage of these transformations were that either a transfer matrix could be derived analytically for a set of basis functions, or a computational method easily applied on a structured grids. The literature on piecewise isometries presents exotic examples of transformations which are formed from translating, rotating and reorienting pieces of the domain, generating cutting lines in a variety of directions [Goetz (2000, 2003)]. The difficulty of including diffusion in these models is the derivation of a numerical scheme which preserves the concentration field during the advective step without numerical diffusion. Novel computational methods to study diffusion in such piecewise isometries could inform the basis of future work. Additionally, advancement could be made by devising experiments which exhibit

discontinuous stirring mechanisms. A recent overview on segregation by cutting and shuffling transformations by Smith *et al.* (2019) listed a number of piece-wise isometries which could be realised, most notably they showed that rotating overlapping circles, a toy model for duct flow [Shearer (1973)], can generate complicated dynamics. As such experiments designed around the systems discussed in Smith *et al.* (2019) could form a comparison to the results presented in Chapters 4 and 7 on purely cutting and shuffling stirring.

Advancing numerical models for the inclusion of diffusion in PWIs could then extended investigation into systems composed with the three mixing mechanisms of stretching and folding, cutting and shuffling and diffusion. In the studies herein, the addition of discontinuous particle transport was only performed via the cutting and shuffling of permutations since they allowed analytic derivation of transfer operators and easy implementation of diffusion. It would be advantageous to study the effect of diffusion in velocity fields with slip deformations, such as those of Smith *et al.* (2016) and Smith *et al.* (2017b), and see if mixing phenomena reported here also emerges. Similarly, not all stirring dynamics were investigated. Whether discontinuous transformations can be used to limit the impact of parabolic points to contaminate mixing could be another avenue of potential investigation. Again, experimental configurations of a dye tracer diffusing in systems which exhibit discontinuous, chaotically stirred configurations could advance understanding of mixing rates in these mixed systems with an adaptation to the Re-orientated Dipole Mixer of Smith *et al.* (2016) an obvious contender.

In general, the impact of discontinuities in particle transport has many open questions, even in the absence of diffusion. For example, if a transformation which is strong mixing is composed with a piecewise isometric map at each step, is ergodicity and strong mixing preserved in the composed transformation if a decomposition does not occur? Byott *et al.* (2013) showed that in expanding maps, the exponential strong mixing rate could be made arbitrarily slow when composed with a permutation if the number of intervals in the permutation tended to infinity. Viana (2006) proved that in the sawtooth map, ergodicity is preserved, but no results were found on the rates of decay of correlations. Future work could focus on developing analytical techniques to prove ergodic properties of a range of discontinuous maps, which would prove advantageous in addressing this question. However, in comparison to the necessary developments of Katok *et al.* (1986) to extend Pesin theory to maps with singularities, it is expected that such developments would be non-trivial.

Appendix A

Computation of Lyapunov exponents

This method to compute the Lyapunov exponents for 2 dimensions is taken from Parker & Chua (2012).

First a point $\boldsymbol{x}_0$ and an initial perturbation $\delta\boldsymbol{x}_0$ are chosen. The point is then mapped forward to $\boldsymbol{x}_1 = M(\boldsymbol{x}_0)$ and the Jacobian matrix $\boldsymbol{D}M$ computed at the point $\boldsymbol{x}_1$ such that the deformation of the perturbation evolves as

$$\delta\boldsymbol{x}_1 = \boldsymbol{D}M(\boldsymbol{x}_1) \cdot \delta\boldsymbol{x}_0 \qquad\qquad (A.1)$$

Continued integration of (A.1) would give the rate of growth of $\delta\boldsymbol{x}_n$, and the maximal Lyapunov exponent calculated from

$$h(\boldsymbol{x}_0, \delta\boldsymbol{x}_0) = \frac{1}{N} \ln \prod_{n=1}^{N} \frac{||\delta\boldsymbol{x}_n||}{||\delta\boldsymbol{x}_0||}. \qquad\qquad (A.2)$$

However when considering chaotic systems the length $||\delta\boldsymbol{x}_n||$ will grow unbounded with n and numerical problems will arise. Similarly all $\delta\boldsymbol{x}$ will align with the eigenvector for the fastest stretching direction and computation of other eigenvalues, bar the maximal Lyapunov exponent, would be impossible.

To reliably calculate the spectrum of eigenvalues in an N dimensional system, the perturbation at each iteration can be renormalised and an orthogonalisation performed for the set of perturbation vectors $\delta\boldsymbol{x}$ which will prevent the alignment of all vectors with the fastest growing eigenvector. The maps considered herein have $N = 2$, so the method is described in detail for calculating h_1 and h_2, however the method is easily extend to higher N using Gram-Schmidt orthogonalisation. The only constraint is that the orthogonalised vectors must span the same subspace as the original vectors.

First chose an initial state $\boldsymbol{x}_0$ and two initial perturbations $\delta\boldsymbol{x}^1$ and $\delta\boldsymbol{x}^2$ as two linearly independent vectors, setting $\delta\boldsymbol{x}_0^1 = \delta\boldsymbol{x}^1$ and $\delta\boldsymbol{x}_0^2 = \delta\boldsymbol{x}^2$. The normalised vectors are then $\boldsymbol{u}_0^1 = \frac{\delta\boldsymbol{x}^1}{||\delta\boldsymbol{x}^1||}$ and $\boldsymbol{u}_0^2 = \frac{\delta\boldsymbol{x}^2}{||\delta\boldsymbol{x}^2||}$. Iterating the initial state $\boldsymbol{x}_0$ is denoted as $\boldsymbol{x}_k = M^k(\boldsymbol{x}_0)$ and the perturbation equations from (A.1)

$$\delta\boldsymbol{x}_k^1 = \boldsymbol{D}M(\boldsymbol{x}_{k-1}) \cdot \boldsymbol{u}_{k-1}^1, \qquad \delta\boldsymbol{x}_k^2 = \boldsymbol{D}M(\boldsymbol{x}_{k-1}) \cdot \boldsymbol{u}_{k-1}^2. \tag{A.3}$$

The final step it to ensure that the deformation vectors remain orthonormal at each iteration. This can be done by letting $\boldsymbol{v}_k^1 = \delta\boldsymbol{x}_k^1$, $\boldsymbol{u}_k^1 = \boldsymbol{v}_k^1/||\boldsymbol{v}_k^1||$ and then

$$\boldsymbol{v}_k^2 = \delta\boldsymbol{x}_2^k - \langle \delta\boldsymbol{x}_k^2, \boldsymbol{u}_k^1 \rangle \boldsymbol{u}_k^1$$

and finally $\boldsymbol{u}_k^2 = \boldsymbol{v}_k^2/||\boldsymbol{v}_k^2||$, which is orthogonal to $\boldsymbol{u}_k^1$. Then on repeated iteration of (A.1) and the orthonormalisation process that follows, the Lyapunov exponents h_1 and h_2 can be found as

$$h_1 = \frac{1}{K} \sum_{k=1}^{K} \ln ||\boldsymbol{v}_k^1||, \quad h_2 = \frac{1}{K} \sum_{k=1}^{K} \ln ||\boldsymbol{v}_k^2||. \tag{A.4}$$

for K large enough.

Appendix B

Derivation of transfer matrices

B.1 Derivation of discrete Fourier transfer matrix for wave-perturbed toral automorphisms

Define $\mathcal{M} = A \cdot \boldsymbol{x} + \phi(\boldsymbol{x})$ where A is a non-singular matrix with integer coefficients which has unit determinant and $\phi(\boldsymbol{x})$ is periodic in both directions with unit period. The transfer matrix defining transfer between Fourier coefficients is

$$\boldsymbol{M}_{\boldsymbol{kq}} = \int_{\mathcal{T}^2} \exp(2\pi i(\boldsymbol{q} \cdot \mathcal{M}^{-1}(\boldsymbol{x}) - \boldsymbol{k} \cdot \boldsymbol{x}) - 4\pi^2 \kappa \boldsymbol{k}^2) \mathrm{d}\boldsymbol{x}, \tag{B.1}$$

but since $\mathcal{M}$ has unit determinant this can be written as

$$\boldsymbol{M}_{\boldsymbol{kq}} = \int_{\mathcal{T}^2} \exp(2\pi i(\boldsymbol{q} \cdot \boldsymbol{x} - \boldsymbol{k} \cdot \mathcal{M}(\boldsymbol{x})) - 4\pi^2 \kappa \boldsymbol{k}^2) \mathrm{d}\boldsymbol{x} \tag{B.2}$$

with a change of variables. In line with Thiffeault & Childress (2003), taking

$$\phi(\boldsymbol{x}) = \frac{K}{2\pi} \begin{pmatrix} \sin(2\pi x_1) \\ \sin(2\pi x_1) \end{pmatrix} \tag{B.3}$$

and expanding out the map $\mathcal{M}$ gives

$$\boldsymbol{M}_{\boldsymbol{kq}} = \exp(-4\pi^2 \kappa \boldsymbol{k}^2) \int_{\mathcal{T}^2} \exp[2\pi i(\boldsymbol{q} - \boldsymbol{k} \cdot A)\boldsymbol{x} - i(k_1 + k_2)K \sin(2\pi x_1)] d\boldsymbol{x} \tag{B.4}$$

where $\boldsymbol{q}$ is the initial wave number and $\boldsymbol{k}$ is the final wave number after one advective iteration. The expression $\boldsymbol{q} - \boldsymbol{k} \cdot A$ can be written as

$$\boldsymbol{Q} = \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} = \begin{pmatrix} q_1 - k_1 A_{00} - k_2 A_{10} \\ q_2 - k_1 A_{01} - k_2 A_{11} \end{pmatrix}. \tag{B.5}$$

The integral over x_2 gives a Kronecker delta, thus the matrix reduces to

$$\boldsymbol{M}_{\boldsymbol{kq}} = e^{-4\pi^2 \kappa \boldsymbol{k}^2} \delta_{0,Q_2} \int_{\mathcal{T}} \exp[-2\pi i Q_1 x_1 - i(k_1 + k_2)K \sin(2\pi x_1)] \mathrm{d}x_1. \tag{B.6}$$

The integral over x_1 can be written in terms of a Bessel function of the first kind via the following manipulation; first shift the integrals limits with $x = x_1 + 1/2$;

$$\int_{\mathcal{I}} \ldots \mathrm{d}x_1 = \int_{-1/2}^{1/2} \exp(-2\pi i Q_1 x - \pi i Q_1 - i\chi \sin(2\pi x + \pi)) \mathrm{d}x \tag{B.7}$$

where $\chi = (k_1 + k_2)K$ and then defining $\eta = 2\pi x$, such that $\mathrm{d}\eta = 2\pi \mathrm{d}x$

$$= e^{-i\pi Q_1} \frac{1}{2\pi} \int_{-\pi}^{\pi} \exp(-i[Q_1 \eta + \sin(\eta)]) \mathrm{d}\eta = (-1)^{Q_1} J_{Q_1}(\chi). \tag{B.8}$$

The full transfer matrix is then written as

$$\boldsymbol{M_{kq}} = \exp(-4\pi^2 \kappa \boldsymbol{k}^2) \delta_{0,Q_2} (-1)^{Q_1} J_{Q_1}((k_1 + k_2)K). \tag{B.9}$$

B.2 Fourier transfer matrix for permutations of equal sized cells

The mapping $T_\sigma(x)$ for a given permutation σ on symbols $\{1, 2, \ldots, N\}$ is given by

$$T_\sigma(x) = x + \frac{\sigma(i) - i}{N} \quad \text{when } x \in [(i-1)/N, i/N). \tag{B.10}$$

The transfer matrix between complex Fourier coefficient is

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \boldsymbol{M}_{kq} \hat{c}_q^{(j)}$$

where the transfer matrix can be computed as

$$\boldsymbol{M}_{kq} = \int_0^1 e^{2\pi i q T_\sigma^{-1}(x) - 2\pi i k x} \mathrm{d}x \tag{B.11}$$

Since $M'(x) = 1$ for all x, from a change of variables the matrix can be found as

$$\boldsymbol{M}_{kq} = \int_0^1 e^{2\pi i q x - 2\pi i k T_\sigma(x)} \mathrm{d}x$$

$$\boldsymbol{M}_{kq} = \sum_{l=1}^{N} \int_{(l-1)/N}^{l/N} e^{2\pi i q x - 2\pi i k x - (2\pi i k/N)(\sigma(l) - l)} \mathrm{d}x$$

$$= \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l) - l)} \int_{(l-1)/N}^{l/N} e^{2\pi i (q-k)x} \mathrm{d}x$$

$$= \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l) - l)} \frac{1}{2\pi i (q - k)} \left[e^{2\pi i (q-k)x} \right]_{(l-1)/N}^{l/N}$$

$$= \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l)-l)} \frac{1}{2\pi i(q-k)} \left[e^{2\pi i(q-k)(l/N)} - e^{2\pi i(q-k)(l-1)/N)} \right]$$

$$= \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l)-l)} \frac{1}{2\pi i(q-k)} \left[e^{2\pi i(q-k)(l/N)} - e^{2\pi i(q-k)(l/N)-2\pi i(q-k)/N} \right].$$

Define the primitive Nth root of unity $\omega = e^{-2\pi i/N}$ then,

$$= \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l)-l)} \frac{(1-\omega^{(q-k)})}{2\pi i(q-k)} e^{2\pi i(q-k)(l/N)}$$

$$= \frac{(1-\omega^{(q-k)})}{2\pi i(q-k)} \sum_{l=1}^{N} e^{(-2\pi i k/N)(\sigma(l)-l)} e^{2\pi i(q-k)(l/N)}$$

$$= \frac{(1-\omega^{(q-k)})}{2\pi i(q-k)} \sum_{l=1}^{N} e^{-2\pi i k\sigma(l)/N + 2\pi i k l/N + 2\pi i q l/N - 2\pi i k l/N}$$

$$= \frac{(1-\omega^{(q-k)})}{2\pi i(q-k)} \sum_{l=1}^{N} e^{-2\pi i k\sigma(l)/N + 2\pi i q l/N}$$

$$= \frac{(1-\omega^{(q-k)})}{2\pi i(q-k)} \sum_{l=1}^{N} \omega^{k\sigma(l)-ql}.$$

B.3 Derivation of the uniform baker's transformation variance decay

Wonhas & Vassilicos (2002) present the solution that given the initial condition $c(x) = \cos(2\pi x)$, that the variance is given by the equation

$$\psi(j) = \psi(0) \exp\left(-\frac{32}{3}\pi\kappa(4^j - 1)\right). \tag{B.12}$$

Here it is briefly noted how this is derived.

First, for the initial condition $\cos(2\pi x)$ would have $\hat{c}_{\pm 1} = 0.5$ such that $\psi(0) = 2||\hat{c}_1||^2 = 0.5$. The action of the baker's transformation is to map $\hat{c}_{\pm k} \to \hat{c}_{\pm 2k}$, so in the absence of diffusion and starting with the wave number $k = 1$ the wavenumber at iteration (j) is $\hat{c}_{\pm 2^j}$.

Considering this, when there is diffusion the variance changes by a factor at each iteration according to the diffusion coefficient and the corresponding wavenumber resulting in

$$\psi(j) = \psi(0) \prod_{l=1}^{j} \exp(-8\pi^2\kappa 2^{2l}) \tag{B.13}$$

$$\psi(j) = \psi(0) \exp(-8\pi^2 \kappa \sum_{l=1}^{j} 4^l) \tag{B.14}$$

Now what is left to do is to evaluate the summation of the powers of 4. Have that

$$\sum_{l=1}^{j} 4^{l+1} - \sum_{l=1}^{j} 4^l = 4^{j+1} - 4$$

$$\implies (4-1) \sum_{l=1}^{j} 4^l = 4(4^j - 1)$$

$$\implies \sum_{l=1}^{j} 4^l = \frac{4}{3}(4^j - 1)$$

Hence, substituting this into (B.14) derives the solution of Wonhas & Vassilicos (2002).

B.4 Derivation of the Fourier transfer matrix for non-uniform baker's transformation

Assume the baker's transformation is uneven such that for $0 < \alpha < 1$, have $\beta = 1 - \alpha$ and the mapping is such that

$$c(x) = \begin{cases} c(\alpha^{-1}x) \text{ for } 0 \le x < \alpha, \\ c(\beta^{-1}(x - \alpha)) \text{ for } \alpha \le x < 1. \end{cases} \tag{B.15}$$

This mapping is given in terms $c(x) = c(T^{-1}(x))$. The transfer matrix between Fourier modes is given from

$$\hat{c}_k^{(j+1)} = \sum_{q=-\infty}^{\infty} \hat{c}_q^{(j)} \int_0^1 e^{2\pi i q T^{-1}(x) - 2\pi i k x} \mathrm{d}x \tag{B.16}$$

Since the map is $1 : 2$ and not area preserving we derive the transfer matrix directly from (B.16) as

$$\boldsymbol{M}_{qk}^{\alpha} = \int_0^{\alpha} \exp(2\pi i q x / \alpha - 2\pi i k x) dx + \int_{\alpha}^1 \exp(2\pi i q (x - \alpha)/\beta - 2\pi i k x) dx. \tag{B.17}$$

If $(q/\alpha - k) = 0$ then the first integral simply becomes

$$[x]_0^{\alpha} .$$

The second integral becomes

$$\int_{\alpha}^{1} \exp(-2\pi i(q\alpha/\beta) + 2\pi ix(q/\beta - k))dx$$

Substituting $\alpha = 1 - \beta$ into the integral derives

$$\int_{\alpha}^{1} \exp(-2\pi iq/\beta + 2\pi ix(q/\beta - k))dx$$

Equate the integral and again, substitute $\alpha = 1 - \beta$ for α, the integral equates to zero. Thus the first bracket equates to α and the second to zero, so the resonance $q/\alpha = k$ has $\boldsymbol{M}_{qk}^{\alpha} = \alpha$. Similarly, $q/\beta = k$, $\boldsymbol{M}_{qk}^{\alpha} = \beta$. Otherwise

$$\boldsymbol{M}_{qk}^{\alpha} = \left[\frac{\exp(2\pi ix(q/\alpha - k))}{2\pi i(q/\alpha - k)}\right]_{0}^{\alpha} + \left[\frac{\exp(-2\pi iq\alpha/\beta + 2\pi ix(q/\beta - k))}{2\pi i(q/\beta - k)}\right]_{\alpha}^{1} \quad \text{(B.18)}$$

$$= \frac{\alpha}{2\pi i(q - \alpha k)}\left[\exp(-2\pi i\alpha k) - 1\right] + \frac{\beta}{2\pi i(q - \beta k)}\left[\exp(-2\pi iq(1-\alpha)/\beta) - \exp(-2\pi i\alpha k)\right]$$

$$= \frac{\alpha(q - \beta k)}{2\pi i(q - \alpha k)(q - \beta k)}\left[\exp(-2\pi i\alpha k) - 1\right] + \frac{\beta(q - \alpha k)}{2\pi i(q - \alpha k)(q - \beta k)}\left[1 - \exp(-2\pi i\alpha k)\right]$$

$$\boldsymbol{M}_{qk}^{\alpha} = \frac{1}{2\pi i(q - \alpha k)(q - \beta k)}\left[(\beta - \alpha)q\left(1 - e^{-2\pi i\alpha k}\right)\right].$$

Taking out a factor of $e^{-\pi i\alpha k}$ in the square bracket results and using the fact that

$$\sin(x) = \frac{e^{ix} - e^{-ix}}{2i} \quad \text{(B.19)}$$

arrives at the transfer matrix of for the non-uniform baker's transformation

$$\boldsymbol{M}_{qk}^{\alpha} = \frac{\sin(k\pi\alpha)}{\pi}\frac{(\beta - \alpha)qe^{-ki\pi\alpha}}{(q - \alpha k)(q - \beta k)}. \quad \text{(B.20)}$$

Appendix C

Rectangle Exchange transformation

C.1 Computational scheme for Haller's rectangle transformations with diffusion

In this section the computational method is described to compute the iterative evolution of a scalar-field $c(x, y, t)$ by Haller's rectangle transformation with a diffusional step.

For the advective step, the concentration field is shuffled by one application of the rectangle transformation from Chapter 7 for a given tuple (α, β, γ). Similar to an Interval Exchange transformation with rational dependence between sub-intervals, if α, β and δ are rationally dependent, then a permutation between sub-rectangles of equal size can be constructed on the unit torus. By converting the scaling parameters to fractions, for example $\alpha = \alpha_N / \alpha_D$, then M can be defined as the least common multiple between α_D, β_D and δ_D. A sub-rectangle with sides of equal length $1/M$ is the smallest piece that can be created from cutting and shuffling by the transformation. By assigning a reference number i to each of these small sub-rectangles and an associated colour C to the face of the rectangle, segregation by the permutation can be investigated in a similar fashion to that of Krotter et. al. by constructing a permutation Π between the sub-rectangles and where they are shuffled to by $C^{(j)}(i) = C^{(j-1)}(\Pi^{-1}(i))$. This method is restrictive to the types of concentration field which can be considered and does not necessarily allow a high resolution spatially if the effect of diffusion in the system is of interest.

Rather, a grid can be constructed of uniform size in both the x and y direction, such that $\Delta x = \Delta y = h$ and $h = 1/MN$, where $N \in \mathbb{N}_{>0}$ and is chosen to be sufficiently large. The concentration field $c(x, y, t)$ is then approximated numerically as a vector $\boldsymbol{C}_{i \cdot MN+j} = c(x_i, y_j, t)$. A transfer matrix $\boldsymbol{\phi}$ for the rectangle transformation S can be constructed by

$$\phi_{i' \cdot MN+j', i \cdot MN+j} = \begin{cases} 1 \text{ if } (x_i, y_j) = S(x_{i'}, y_{j'}) \\ 0 \text{ otherwise.} \end{cases} \tag{C.1}$$

By choosing $h = 1/MN$ the transformation is captured perfectly on the uniform grid with no numerical diffusion in the advective step.

For the diffusive step the heat equation

$$\frac{\partial c}{\partial t} = \kappa \left(\frac{\partial^2 c}{\partial x^2} + \frac{\partial^2 c}{\partial y^2} \right). \tag{C.2}$$

is solved with periodic boundary conditions over the interval $t = 0$ to $t = 1$. Finite difference methods (FDMs) are simple and effective solvers for partial differential equations such as the heat equation. They approximate the derivatives using finite differences on a discretization in space and time and reduce the problem to solving a system of algebraic equations. Taking the spatial discretisation already established for the advective step, the time interval is discretised so that it has uniform increments $\Delta t = t_{k+1} - t_k$ for all k. One-dimensional FDMs are relatively quick and many unconditionally stable methods exist, such as Crank-Nicolson, however extension to two-dimensions using such methods produces a complicated system of equations.

An Alternating Direction Implicit (ADI) scheme splits the finite difference scheme into two steps, for each derivative, and can be computed relatively quickly if designed to solve systems of tridiagonal matrices. At each time step the finite difference is solved first for the x derivative implicitly and and then the y derivative implicitly, such that

$$\frac{c_{ij}^{k+1/2} - c_{ij}^k}{\Delta t/2} = \kappa \frac{\delta_x^2 c_{ij}^{k+1/2} + \delta_y^2 c_{ij}^k}{h^2} \tag{C.3}$$

$$\frac{c_{ij}^{k+1} - c_{ij}^{k+1/2}}{\Delta t/2} = \kappa \frac{\delta_x^2 c_{ij}^{k+1/2} + \delta_y^2 c_{ij}^{k+1}}{h^2} \tag{C.4}$$

where δ_p^2 is the central second difference operator for the $p - th$ coordinate

$$\delta_x^2 c_{ij} = c_{i+1,j} - 2c_{ij} + c_{i-1,j}, \quad \delta_y^2 c_{ij} = c_{i,j-1} - 2c_{ij} + c_{i,j-1}. \tag{C.5}$$

The systems of equations to solve these can be made tridiagonal by changing the referencing on the coordinates (i, j) between each of the half steps. The Thomas Algorithm is particularly fast at solving tridiagonal systems and is of $\mathcal{O}(N)$ for N systems of equations, while Gaussian elimination is (N^3).

Complications arise by choosing periodic boundary conditions for the system which means that the systems of unknowns are no longer tridiagonal but have elements $A_{1,MN} \neq 0$ and $A_{MN,1} \neq 0$, however the Sherman-Morrison algorithm reduces the problem to solving two tridiagonal systems via the Thomas algorithm at each implicit derivative step retaining $\mathcal{O}(N)$.

For the computations, α, β and δ were chosen such that $M = 100$ and N was taken to be 10 such that $\Delta x = \Delta y = 1/1000$ in the spatial directions and $\Delta t = 0.01$ was taken. For diffusion coefficient values $\kappa \in [10^{-6}, 10^{-4}]$ this time step was sufficient.
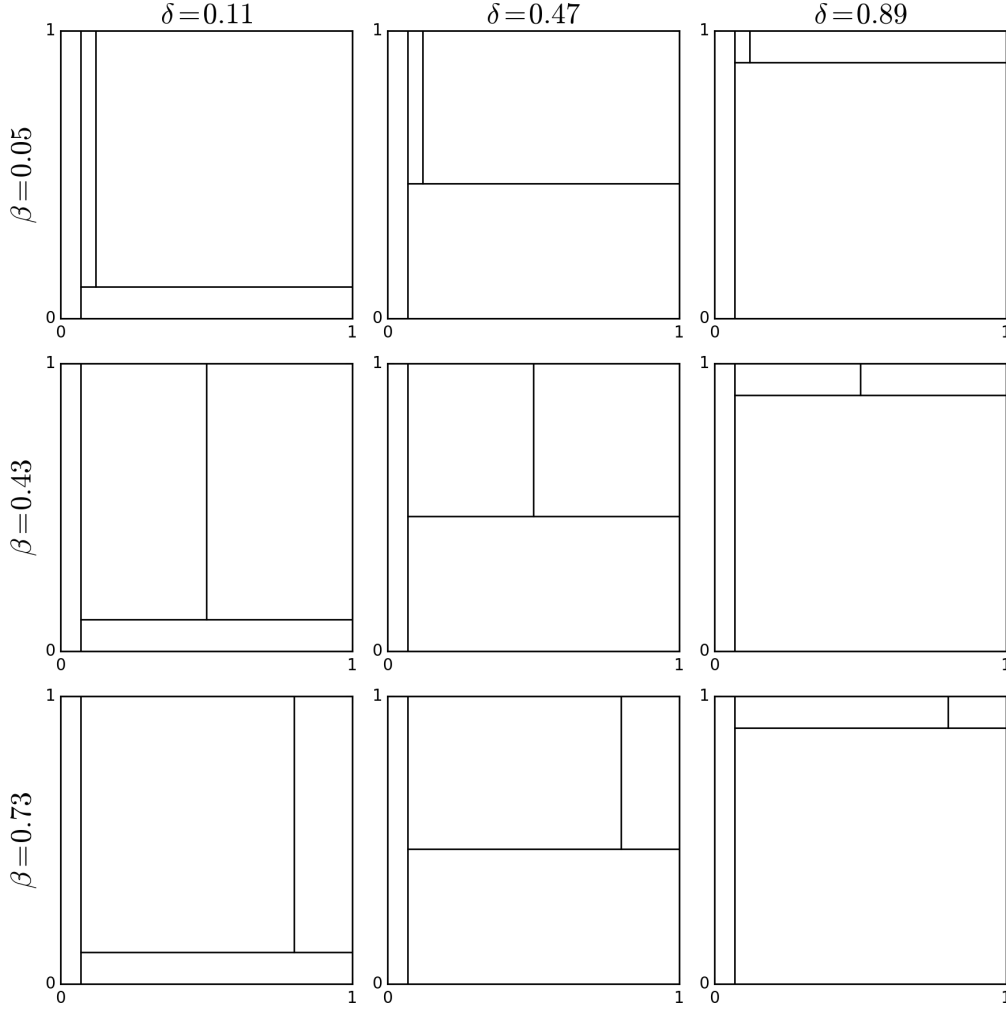
Figure C.1: $\alpha = 0.07$ in all images. β and δ are changed according to labels.

C.2 Segregation study

To see if the scaling $t_\% \propto \kappa^{-\eta}$ with $\eta < 1$ holds for piecewise isometries in higher dimensions, the parameters α, β and γ need to be chosen such that they produce good segregation in finite time. This appendix presents a small study on the choice of parameters (α, β, δ), more specifically the respective scaling sized between the sub-rectangles, and the resulting decrease in the scale of segregation after a number of iterations.

The parameters are characterised as being either small, medium or large in width, relative to the other parameters. For example, Figure C.1 gives 9 examples where the value of α is small, $\alpha = 0.07$ on the unit torus, while β and δ vary between three values. The values $\beta = 0.11$, 0.43 and 0.73, and $\delta = 0.11$, 0.47,
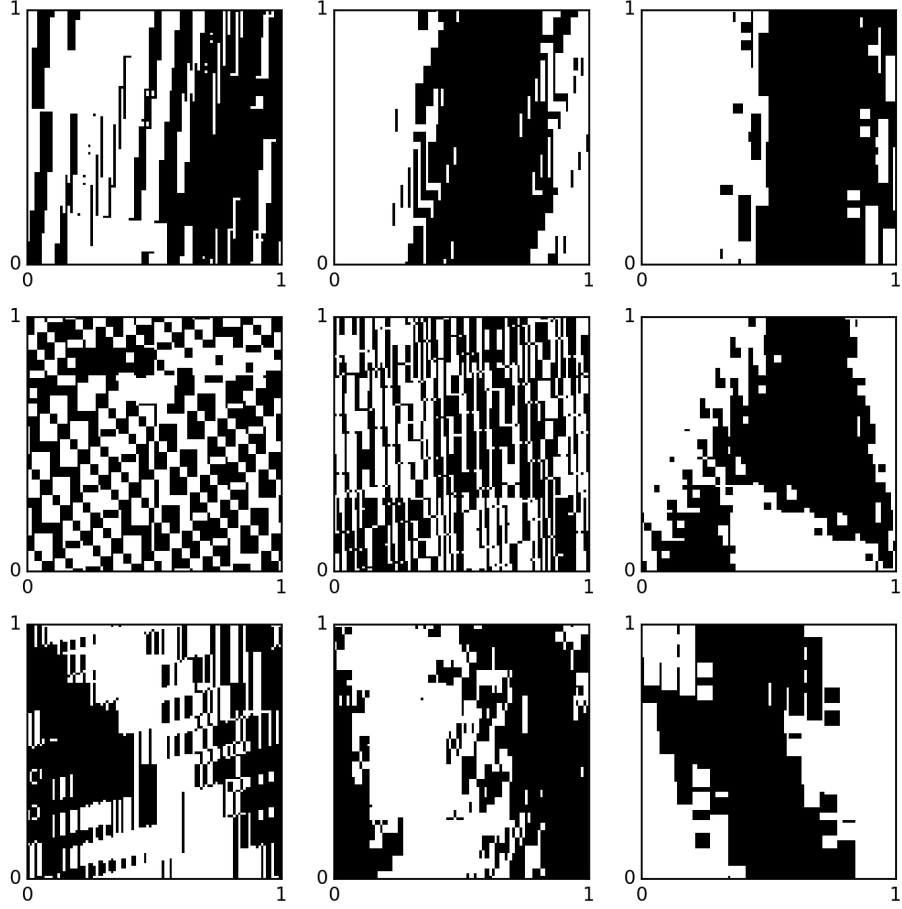
Figure C.2: Segregation after 100 iterations for the initial condition *half-black, half white*. $\alpha = 0.07$ in all images. β and δ are changed according to Figure C.1.

0.89, are chosen to represent small, medium or large values of these parameters in comparison to small α, but also all result in long-periodic dynamics of the RETs, i.e. reassembly of the initial condition only occurs for an extremely large number of iterations. This is because no smaller value than $M = 100$ can represent the RET for these parameters.

Figure C.2 plots the resulting stirring after 100 iterations for the parameter fields in Figure C.1. For nearly all of the tuples with small α, large regions which are poorly stirred exist within the domain. Only when β is medium sized, a more even dispersion of black and white are visible throughout the domain.

Figure C.3 repeats a similar scaling plot with α fixed at a medium value and the same δ values explored. The scales of β vary however, taking into account the updated value of α, were now the small, medium and large values of β are associated with the remainder of the domain once α is fixed. Figure C.4 plots
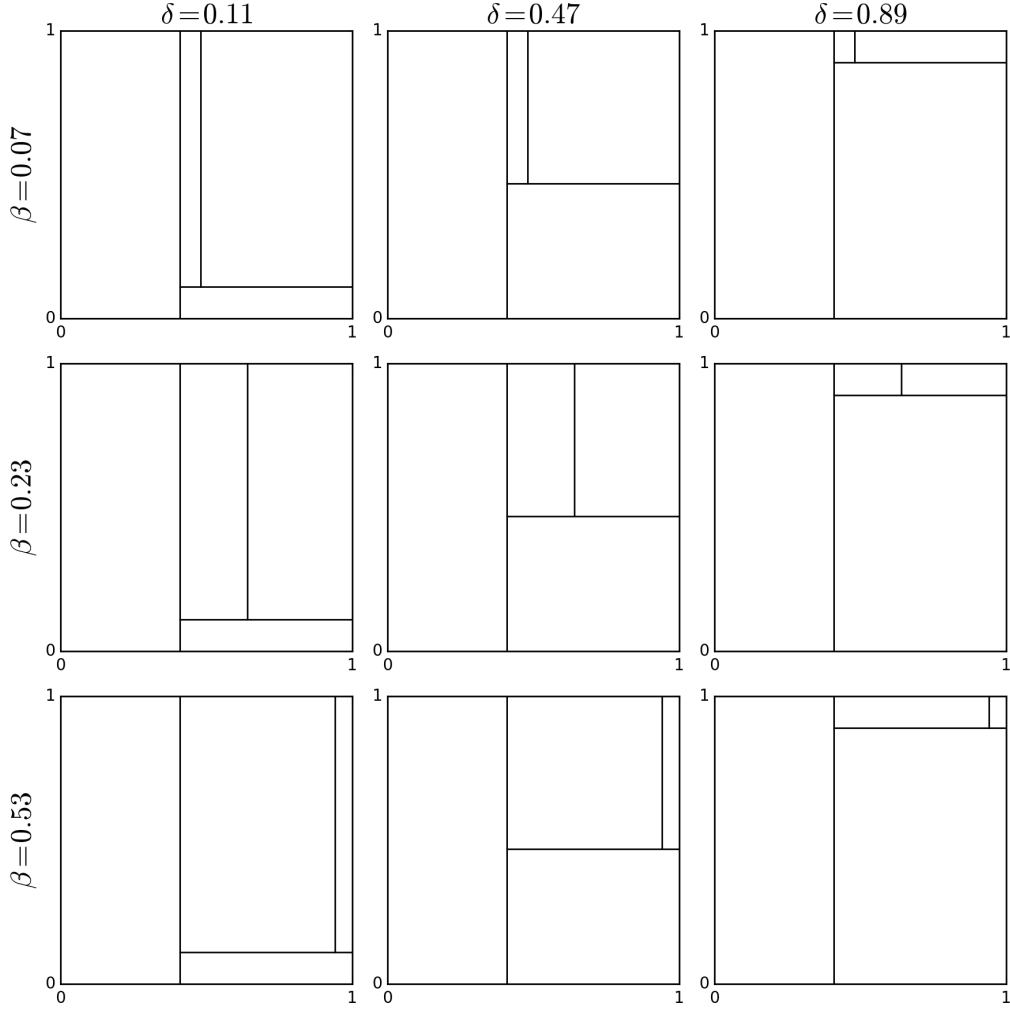
Figure C.3: $\alpha = 0.07$ in all images. β and δ are changed according to labels.

the resulting stirring after 100 iterations for the parameter fields in Figure C.3. Better mixing is observed throughout the 9 parameter tuples explored. The best decreases in the scale of segregation appear when $\beta = 0.23$ and $\delta = 0.11$ or 0.47. Qualitatively, it appears as though a medium scaling of all parameters result in the best stirred concentration field after 100 iterations.

As a final study, for large α, fixing small β, Figure C.5 shows three configurations of the RET with varying δ. The segregation plots of Figure C.6 reveal these to the be worse parameter configurations to achieve good stirring in a short finite time.
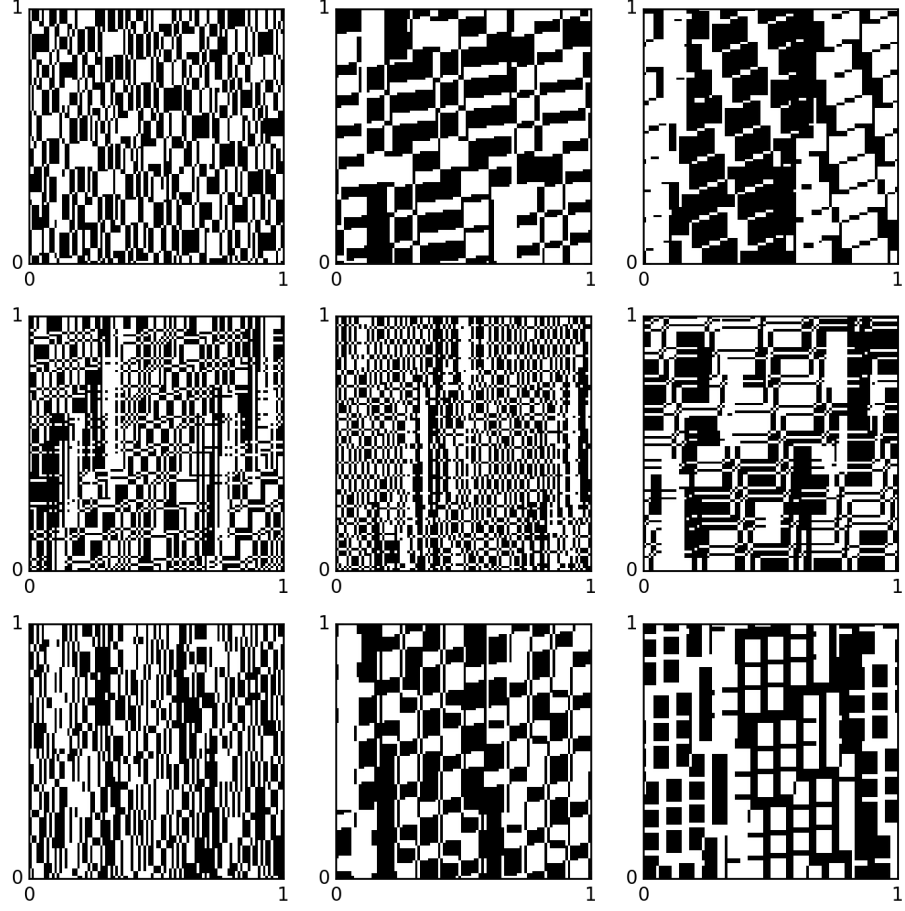
Figure C.4: Segregation after 100 iterations for the initial condition *half-black, half white*. $\alpha = 0.07$ in all images. β and δ are changed according to Figure C.3.

Figure C.5: $\alpha = 0.07$ in all images. β and δ are changed according to labels.

241

Figure C.6: Segregation after 100 iterations for the initial condition *half-black, half white*. $\alpha = 0.07$ in all images. β and δ are changed according to Figure C.5.
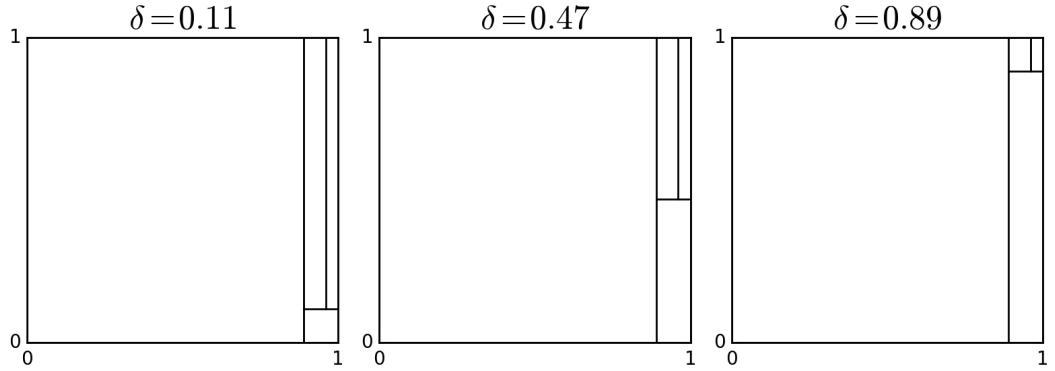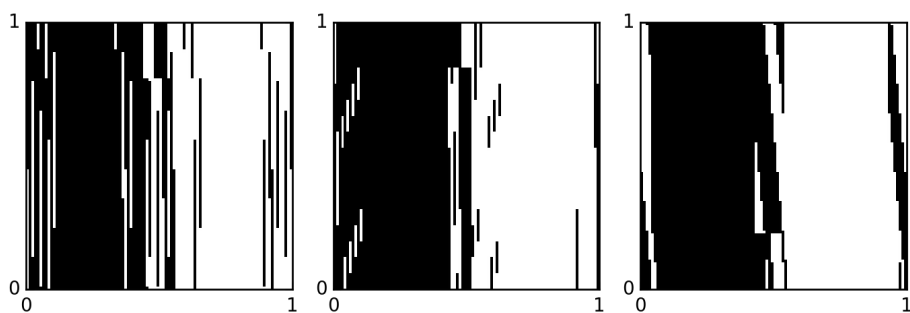
References

ALDOUS, D. & DIACONIS, P. (1986). Shuffling cards and stopping times. *The American Mathematical Monthly*, **93**, 333–348. 42, 107

ANTONSEN JR, T.M., FAN, Z., OTT, E. & GARCIA-LOPEZ, E. (1996). The role of chaotic orbits in the determination of power spectra of passive scalars. *Phys. Fluids*, **8**, 3094–3104. 1, 33, 41, 60, 180

AREF, H. (1984). Stirring by chaotic advection. *J. Fluid Mech*, **143**, 1–21. 1, 7, 12

AREF, H. (2002). The development of chaotic advection. *Physics of Fluids (1994-present)*, **14**, 1315–1325. 12

ASHWIN, P. & GOETZ, A. (2005). Invariant curves and explosion of periodic islands in systems of piecewise rotations. *SIAM Journal on Applied Dynamical Systems*, **4**, 437–458. 37

ASHWIN, P., NICOL, M. & KIRKBY, N. (2002). Acceleration of one-dimensional mixing by discontinuous mappings. *Physica A*, **310**, 347–363. 4, 22, 27, 28, 41, 67, 69, 74, 83, 97, 110, 117, 156, 219

AVILA, A. & FORNI, G. (2007). Weak mixing for interval exchange transformations and translation flows. *Ann. Math.*, 637–664. 37, 69

BALKOVSKY, E. & FOUXON, A. (1999). Universal long-time properties of lagrangian statistics in the Batchelor regime and their application to the passive scalar problem. *Physical Review E*, **60**, 4164. 32

BATCHELOR, G.K. (1952). The effect of homogeneous turbulence on material lines and surfaces. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 213, 349–366, The Royal Society. 11

BORDENAVE, C., QIU, Y. & ZHANG, Y. (2018). Spectral gap of sparse bistochastic matrices with exchangeable rows with application to shuffle-and-fold maps. *arXiv preprint arXiv:1805.06205*. 161

BOYLAND, P.L., AREF, H. & STREMLER, M.A. (2000). Topological fluid mechanics of stirring. *Journal of Fluid Mechanics*, **403**, 277–304. 13

BUZZI, J. (2001). Piecewise isometries have zero topological entropy. *Ergodic Theory and Dynamical Systems*, **21**, 1371–1377. 37

BYOTT, N., HOLLAND, M. & ZHANG, Y. (2013). On the mixing properties of piecewise expanding maps under composition with permutations. *Discrete Contin. Dyn. Syst.*, **33**, 3365– 3390. 3, 18, 40, 41, 121, 129, 130, 131, 132, 133, 134, 135, 163, 170, 194, 210, 220, 223, 225

BYOTT, N.P., LIN, C. & ZHANG, Y. (2016). On the mixing properties of piecewise expanding maps under composition with permutations, ii: Maps of non-constant orientation. *Stochastics and Dynamics*, **16**, 1660013. 3, 40, 41, 161, 194, 223

CHAIKEN, J., CHEVRAY, R., TABOR, M. & TAN, Q. (1986). Experimental study of lagrangian turbulence in a Stokes flow. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 408, 165–174, The Royal Society. 13

CHIEN, W.L., RISING, H. & OTTINO, J.M. (1986). Laminar mixing and chaotic mixing in several cavity flows. *Journal of Fluid Mechanics*, **170**, 355–377. 13

CHRISTOV, I.C., OTTINO, J.M. & LUEPTOW, R.M. (2010). Streamline jumping: A mixing mechanism. *Physical Review E*, **81**, 046307. 38

CHRISTOV, I.C., OTTINO, J.M. & LUEPTOW, R.M. (2011). From streamline jumping to strange eigenmodes: Bridging the lagrangian and eulerian pictures of the kinematics of mixing in granular flows. *Physics of Fluids*, **23**, 103302. 11, 22

CLIFFORD, M., COX, S. & ROBERTS, E. (1999). Reaction and diffusion in a lamellar structure: the effect of the lamellar arrangement upon yield. *Physica A*, **262**, 294–306. 30, 159, 220

CLIFFORD, M.J., COX, S.M. & FINN, M.D. (2004). Reynolds number effects in a simple planetary mixer. *Chemical engineering science*, **59**, 3371–3379. 11

COLA, B.A. (2004). *Optimization of a Pulsed Source-sink Microscale Mixing Device*. Ph.D. thesis, Vanderbilt University. 3, 39, 40

COLA, B.A., SCHAFFER, D.K., FISHER, T.S. & STREMLER, M.A. (2006). A pulsed source-sink fluid mixing device. *Journal of microelectromechanical systems*, **15**, 259–266. 40

D'ALESSANDRO, D., DAHLEH, M. & MEZIC, I. (1999). Control of mixing in fluid flow: A maximum entropy approach. *IEEE Transactions on Automatic Control*, **44**, 1852–1863. 49

DANCKWERTS, P. (1952). The definition and measurement of some characteristics of mixtures. *Applied Scientific Research, Section A*, **3**, 279–296. 27

DIACONIS, P. (1996). The cutoff phenomenon in finite Markov chains. *Proceedings of the National Academy of Sciences*, **93**, 1659–1664. 107

DING, J. & ZHOU, A. (2010). *Statistical properties of deterministic systems*. Springer Science & Business Media. 22

ECKART, C. (1948). An analysis of the stirring and mixing processes in incompressible fluids. *Jounral of Marine Research*, **7**, 265–275. 11

ECKHARDT, B., HASCOËT, E. & BRAUN, W. (2003). Passive fields and particles in chaotic flows. In *IUTAM Symposium on Nonlinear Stochastic Dynamics*, 415–424, Springer. 35, 159

FEREDAY, D. & HAYNES, P. (2004). Scalar decay in two-dimensional chaotic advection and Batchelor-regime turbulence. *Physics of Fluids*, **16**, 4359–4370. 34

FEREDAY, D.R., HAYNES, P.H., WONHAS, A. & VASSILICOS, J.C. (2002). Scalar variance decay in chaotic advection and Batchelor-regime turbulence. *Phys. Rev. E*, **65**, 035301. 22, 33, 41, 124, 139, 159, 161, 163, 166, 180

FINN, M.D., COX, S.M. & BYRNE, H.M. (2003). Topological chaos in inviscid and viscous mixers. *Journal of Fluid mechanics*, **493**, 345–361. 13

FINN, M.D., COX, S.M. & BYRNE, H.M. (2004). Mixing measures for a two-dimensional chaotic Stokes flow. *Journal of Engineering Mathematics*, **48**, 129–155. 30, 36

FROYLAND, G. (2001). *Extracting dynamical behavior via Markov models*, 281. Birkhäuser Verlag AG, Berlin. 20

FROYLAND, G., PADBERG, K., ENGLAND, M.H. & TREGUIER, A.M. (2007). Detection of coherent oceanic structures via transfer operators. *Physical review letters*, **98**, 224503. 2

FROYLAND, G., GONZÁLEZ-TOKMAN, C. & WATSON, T.M. (2016). Optimal mixing enhancement by local perturbation. *SIAM Rev.*, **58**, 494–513. 42

FU, X.C. & DUAN, J. (2008). On global attractors for a class of non-hyperbolic piecewise affine maps. *Physica D: Nonlinear Phenomena*, **237**, 3369–3376. 37

GIBBS, J.W. (1902). Elementary principles in statistical physics. *The Collected Works of JW Gibbs (Yale University, New Haven, CT, 1957)*, **2**. 27

GILBERT, A.D. (2006). Advected fields in maps—iii. passive scalar decay in baker's maps. *Dynam. Syst.*, **21**, 25–71. 22, 34, 35, 41, 66, 159, 161, 167

GOETZ, A. (1998). Dynamics of a piecewise rotation. *Discrete and Continuous Dynamical Systems*, **4**, 593–608. 3, 37

GOETZ, A. (2000). Dynamics of piecewise isometries. *Illinois Journal of Mathematics*, **44**, 465–478. 3, 37, 224

GOETZ, A. (2003). Piecewise isometries—an emerging area of dynamical systems. In *Fractals in Graz 2001*, 135–144, Springer. 3, 37, 224

GOUILLART, E., KUNCIO, N., DAUCHOT, O., DUBRULLE, B., ROUX, S. & THIFFEAULT, J.L. (2007). Walls inhibit chaotic mixing. *Phys. Rev. Lett.*, **99**, 114501. 35, 124, 161

GOUILLART, E., DAUCHOT, O., DUBRULLE, B., ROUX, S. & THIFFEAULT, J.L. (2008). Slow decay of concentration variance due to no-slip walls in chaotic mixing. *Phys. Rev. E*, **78**, 026211. 1, 2, 34, 35, 66, 83, 124, 161, 189

GOUILLART, E., THIFFEAULT, J.L. & DAUCHOT, O. (2010). Rotation shields chaotic mixing regions from no-slip walls. *Phys. Rev. Lett.*, **104**, 204502. 2, 36, 223

HALLER, H. (1981). Rectangle exchange transformations. *Monatshefte für Mathematik*, **91**, 215–232. xx, 201, 215

HARDT, S., DRESE, K., HESSEL, V. *et al.* (2008). Passive micro mixers for applications in the micro reactor and μtas field. In *ASME 2004 2nd International Conference on Microchannels and Minichannels*, 45–55, American Society of Mechanical Engineers Digital Collection. 39

HAYNES, P. & VANNESTE, J. (2005). What controls the decay of passive scalars in smooth flows? *Phys. Fluids*, **17**, 097103. 1, 34, 35, 159

HERTZSCH, J.M., STURMAN, R. & WIGGINS, S. (2007). Dna microarrays: design principles for maximizing ergodic, chaotic mixing. *Small*, **3**, 202–218. 39, 40

HESSEL, V., LÖWE, H. & SCHÖNFELD, F. (2005). Micromixers—a review on passive and active mixing principles. *Chemical Engineering Science*, **60**, 2479–2501. 10

HETMANSKI, J. (2012). Mixing by cutting and shuffling, unpublished. 202

HOBBS, D. & MUZZIO, F. (1997). The Kenics static mixer: a three-dimensional chaotic flow. *Chem. Eng. J.*, **67**, 153–166. 39, 222

HORITA, T., HATA, H., ISHIZAKI, R. & MORI, H. (1990). Long-time correlations and expansion-rate spectra of chaos in Hamiltonian systems. *Progress of Theoretical Physics*, **83**, 1065–1070. 31, 36

HOSSAIN, S. & KIM, K.Y. (2015). Mixing analysis in a three-dimensional serpentine split-and-recombine micromixer. *Chemical Engineering Research and Design*, **100**, 95–103. 39

HU, Y. & PIERREHUMBERT, R.T. (2002). The advection–diffusion problem for stratospheric flow. part ii: Probability distribution function of tracer gradients. *Journal of the atmospheric sciences*, **59**, 2830–2845. 34

JANA, S.C., METCALFE, G. & OTTINO, J.M. (1994). Experimental and computational studies of mixing in complex Stokes flows: the vortex mixing flow and multicellular cavity flows. *Journal of Fluid Mechanics*, **269**, 199–246. 13

JONES, S.W. & AREF, H. (1988). Chaotic advection in pulsed source–sink systems. *Phys. Fluids*, **31**, 469–485. 39, 40

JUAREZ, G., LUEPTOW, R.M., OTTINO, J.M., STURMAN, R. & WIGGINS, S. (2010). Mixing by cutting and shuffling. *EPL (Europhysics Letters)*, **91**, 20003. 3

JUAREZ, G., CHRISTOV, I.C., OTTINO, J.M. & LUEPTOW, R.M. (2012). Mixing by cutting and shuffling 3d granular flow in spherical tumblers. *Chemical engineering science*, **73**, 195–207. 38, 39

KANG, T.G., SINGH, M.K., KWON, T.H. & ANDERSON, P.D. (2008). Chaotic mixing using periodic and aperiodic sequences of mixing protocols in a micromixer. *Microfluidics and nanofluidics*, **4**, 589–599. 160

KATOK, A. (1980). Interval exchange transformations and some special flows are not mixing. *Israel J. Math.*, **35**, 301–310. 37

KATOK, A. & HASSELBLATT, B. (1995). *Introduction to the modern theory of dynamical systems*, vol. 54. Cambridge university press. 12

KATOK, A., STRELCYN, J.M., LEDRAPPIER, F. & PRZYTYCKI, F. (1985). *Smooth maps with singularities: invariant manifolds, entropy and billiards*. Université Paris-Nord. Centre Scientifique et Polytechnique [CSP]. 40

KATOK, A., STRELCYN, J.M., LEDRAPPIER, F. & PRZYTYCKI, F. (1986). *Invariant manifolds, entropy and billiards; smooth maps with singularities*. 225

KEANE, M. (1975). Interval exchange transformations. *Mathematische Zeitschrift*, **141**, 25–31. 37

KRASNOPOLSKAYA, T.S., MELESHKO, V.V., PETERS, G.W.M. & MEIJER, H.E.H. (1999). Mixing in stokes flow in an annular wedge cavity. *European Journal of Mechanics-B/Fluids*, **18**, 793–822. 30

KROTTER, M.K., CHRISTOV, I.C., OTTINO, J.M. & LUEPTOW, R.M. (2012). Cutting and shuffling a line segment: Mixing by interval exchange transformations. *Int. J. Bifurc. Chaos*, **22**, 1230041. ix, x, 30, 38, 67, 68, 69, 71, 72, 73, 75, 92, 94, 95, 97, 99, 104, 117, 118, 216, 219, 222

LEONG, C.W. & OTTINO, J.M. (1989). Experiments on mixing due to chaotic advection in a cavity. *Journal of Fluid Mechanics*, **209**, 463–499. 13, 16

LESTER, D.R., METCALFE, G., TREFRY, M.G., ORD, A., HOBBS, B. & RUDMAN, M. (2009). Lagrangian topology of a periodically reoriented potential flow: Symmetry, optimization, and mixing. *Phys. Rev. E*, **80**, 036208. 39

LIANG, T.C. & WEST, M. (2009). Numerical evidence for cutoffs in chaotic microfluidic mixing. In *ASME 2008 Dynamic Systems and Control Conference*, 1405–1412, American Society of Mechanical Engineers Digital Collection. 42

LOUZGUINE-LUZGIN, D.V., LOUZGUINA-LUZGINA, L.V. & CHURYUMOV, A.Y. (2012). Mechanical properties and deformation behavior of bulk metallic glasses. *Metals*, **3**, 1–22. 40

MAJDA, A.J. & KRAMER, P.R. (1999). Simplified models for turbulent diffusion: theory, numerical modelling, and physical phenomena. *Physics reports*, **314**, 237–574. 31

MATHEW, G., MEZIĆ, I. & PETZOLD, L. (2005). A multiscale measure for mixing. *Physica D: Nonlinear Phenomena*, **211**, 23–46. 27, 28, 75

MELESHKO, V.V. & AREF, H. (1996). A blinking rotlet model for chaotic advection. *Physics of Fluids (1994-present)*, **8**, 3215–3217. 30

MILES, C.J. & DOERING, C.R. (2018). Diffusion-limited mixing by incompressible flows. *Nonlinearity*, **31**, 2346. 119

MUZZIO, F. & LIU, M. (1996). Chemical reactions in chaotic flows. *The Chemical Engineering Journal and the Biochemical Engineering Journal*, **64**, 117–127. 11

NOVAK, C.F. (2009). Discontinuity-growth of interval-exchange maps. **3**. 37

OTT, E. (2002). *Chaos in dynamical systems*. Cambridge university press. 12

OTTINO, J.M. (1989). *The kinematics of mixing: stretching, chaos, and transport*. Cambridge university press. 1, 10, 13

OTTINO, J.M. (1990). Mixing, chaotic advection, and turbulence. *Annual Review of Fluid Mechanics*, **22**, 207–254. 10

OTTINO, J.M., JANA, S.C. & CHAKRAVARTHY, V.S. (1994). From Reynolds's stretching and folding to mixing studies using horseshoe maps. *Physics of Fluids (1994-present)*, **6**, 685–699. 13

PARK, P.P., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2016). Mixing with piecewise isometries on a hemispherical shell. *Chaos*, **26**, 073115. 38, 119, 222

PARK, P.P., LYNN, T.F., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2017). Mixing and the fractal geometry of piecewise isometries. *Physical Review E*, **95**, 042208. 38, 119, 222

PARKER, T.S. & CHUA, L. (2012). *Practical numerical algorithms for chaotic systems*. Springer Science & Business Media. 53, 227

PIERREHUMBERT, R. (1991). Large-scale horizontal mixing in planetary atmospheres. *Phys. Fluids A*, **3**, 1250–1260. 36

PIERREHUMBERT, R. (1994). Tracer microstructure in the large-eddy dominated regime. *Chaos Solitons Fractals*, **4**, 1091–1110. 32, 41

POPOVYCH, O., PIKOVSKY, A. & ECKHARDT, B. (2007). Abnormal mixing of passive scalars in chaotic flows. *Phys. Rev. E*, **75**, 036308. 22, 36, 41, 66, 83, 88, 91, 118, 160, 189

PRIYE, A., HASSAN, Y.A. & UGAZ, V.M. (2013). Microscale chaotic advection enables robust convective dna replication. *Analytical chemistry*, **85**, 10536–10541. 2

REYNOLDS, O. (1894). Study of fluid motion by means of colored bands. *Scientific American*, **38**, 15509–15510. 11

ROTHSTEIN, D., HENRY, E. & GOLLUB, J.P. (1999). Persistent patterns in transient chaotic fluid mixing. *Nature*, **401**, 770–772. 33

SALMAN, H. & HAYNES, P. (2007). A numerical study of passive scalar evolution in peripheral regions. *Physics of Fluids*, **19**, 067101. 35

SCHLICK, C.P., CHRISTOV, I.C., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2013). A mapping method for distributive mixing with diffusion: Interplay between chaos and diffusion in time-periodic sine flow. *Physics of Fluids*, **25**, 052102. 41, 114, 115, 160

SCHLICK, C.P., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2014). Competitive autocatalytic reactions in chaotic flows with diffusion: Prediction using finite-time Lyapunov exponents. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **24**, 013109. 22

SCHLICK, C.P., ISNER, A.B., UMBANHOWAR, P.B., LUEPTOW, R.M. & OTTINO, J.M. (2015). On mixing and segregation: from fluids and maps to granular solids and advection–diffusion systems. *Industrial & Engineering Chemistry Research*, **54**, 10465–10471. 160, 224

SCHÖNFELD, F., HESSEL, V. & HOFMANN, C. (2004). An optimised split-and-recombine micro-mixer with uniform 'chaotic'mixing. *Lab on a Chip*, **4**, 65–69. 2, 39, 223

SHEARER, C. (1973). Mixing of highly viscous liquids: flow geometrics for stream-line subdivision and redistribution. *Chemical Engineering Science*, **28**, 1091–1098. 225

SINGH, M.K., ANDERSON, P.D., SPEETJENS, M.F. & MEIJER, H.E. (2008a). Optimizing the rotated arc mixer. *AIChE journal*, **54**, 2809–2822. 160

SINGH, M.K., KANG, T.G., MEIJER, H.E. & ANDERSON, P.D. (2008b). The mapping method as a toolbox to analyze, design, and optimize micromixers. *Microfluidics and Nanofluidics*, **5**, 313–325. 22, 224

SMITH, L.D. (2016). *Chaotic Advection in a Three-dimensional Volume-preserving Potential Flow*. Ph.D. thesis, Monash University. 28, 75, 92

SMITH, L.D., RUDMAN, M., LESTER, D.R. & METCALFE, G. (2016). Mixing of discontinuously deforming media. *Chaos*, **26**, 023113. 3, 37, 39, 40, 41, 121, 161, 163, 204, 216, 223, 225

SMITH, L.D., PARK, P.P., UMBANHOWAR, P.B., OTTINO, J.M. & LUEP-
TOW, R.M. (2017a). Predicting mixing via resonances: Application to spheri-
cal piecewise isometries. *Physical Review E*, **95**, 062210. 38, 119, 222

SMITH, L.D., RUDMAN, M., LESTER, D.R. & METCALFE, G. (2017b). Impact
of discontinuous deformation upon the rate of chaotic mixing. *Phys. Rev. E*,
95, 022213. 3, 41, 163, 205, 216, 225

SMITH, L.D., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M.
(2017c). Mixing and transport from combined stretching-and-folding and
cutting-and-shuffling. *Physical Review E*, **96**, 042213. 39, 41, 223

SMITH, L.D., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW,
R.M. (2018). Optimized mixing by cutting-and-shuffling. *arXiv preprint
arXiv:1803.08102*. 30, 38, 119

SMITH, L.D., UMBANHOWAR, P.B., LUEPTOW, R.M. & OTTINO, J.M.
(2019). The geometry of cutting and shuffling: An outline of possibilities for
piecewise isometries. *Physics Reports*. 38, 119, 196, 215, 225

SPENCER, R. & WILEY, R. (1951). The mixing of very viscous liquids. *Journal
of Colloid Science*, **6**, 133–145. 22

SPRINGHAM, J. & STURMAN, R. (2014). Polynomial decay of correlations in
linked-twist maps. *Ergodic Theory Dyn. Syst.*, **34**, 1724–1746. 1, 31

STROOCK, A.D., DERTINGER, S.K., AJDARI, A., MEZIĆ, I., STONE, H.A. &
WHITESIDES, G.M. (2002). Chaotic mixer for microchannels. *Science*, **295**,
647–651. 11

STURMAN, R. (2012). The role of discontinuities in mixing. *Adv. Appl. Mech.*,
45, 2012. 37, 40, 42

STURMAN, R. & SPRINGHAM, J. (2013). Rate of chaotic mixing and boundary
behavior. *Phys. Rev. E*, **87**, 012906. 1, 2, 31, 35, 49

STURMAN, R., OTTINO, J.M. & WIGGINS, S. (2006). *The mathematical foun-
dations of mixing: the linked twist map as a paradigm in applications: micro
to macro, fluids to solids*. Cambridge University Press. 10, 14, 48

STURMAN, R., MEIER, S., OTTINO, J.M. & WIGGINS, S. (2008). Linked twist map formalism in two and three dimensions applied to mixing in tumbled granular flows. *Journal of Fluid Mechanics*, **602**, 129–174. 3, 38, 39

SUKHATME, J. & PIERREHUMBERT, R.T. (2002). Decay of passive scalars under the action of single scale smooth velocity fields in bounded two-dimensional domains: From non-self-similar probability distribution functions to self-similar eigenmodes. *Phys. Rev. E*, **66**, 056302. 34, 41, 60

SWANSON, P.D. & OTTINO, J.M. (1990). A comparative computational and experimental study of chaotic mixing of viscous fluids. *Journal of Fluid Mechanics*, **213**, 227–249. 13

THIFFEAULT, J.L. (2008). *Scalar decay in chaotic mixing*, 3–36. Springer. 9, 32, 33, 144

THIFFEAULT, J.L. (2012). Using multiscale norms to quantify mixing and transport. *Nonlinearity*, **25**, R1. 28

THIFFEAULT, J.L. & CHILDRESS, S. (2003). Chaotic mixing in a torus map. *Chaos*, **13**, 502–507. 22, 32, 33, 35, 49, 59, 65, 82, 229

THIFFEAULT, J.L., FINN, M.D., GOUILLART, E. & HALL, T. (2008). Topology of chaotic mixing patterns. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **18**, 033123. 13

THIFFEAULT, J.L., GOUILLART, E. & DAUCHOT, O. (2011). Moving walls accelerate mixing. *Phys. Rev. E*, **84**, 036313. 2, 36, 223

TOUSSAINT, V., CARRIÈRE, P. & RAYNAL, F. (1995). A numerical Eulerian approach to mixing by chaotic advection. *Phys. Fluids*, **7**, 2587–2600. 35, 77

TSANG, Y.K., ANTONSEN JR, T.M. & OTT, E. (2005). Exponential decay of chaotically advected passive scalars in the zero diffusivity limit. *Physical Review E*, **71**, 066301. 34

ULAM, S.M. (1960). *A collection of mathematical problems*, vol. 8. Interscience Publishers. 22

VAIENTI, S. (1992). Ergodic properties of the discontinuous sawtooth map. *Journal of statistical physics*, **67**, 251–269. 40, 216

VIANA, M. (2006). Ergodic theory of interval exchange maps. *Revista Matemática Complutense*, **19**, 7–100. 41, 225

WANG, D.G. & CAMPBELL, C.S. (1992). Reynolds analogy for a shearing granular material. *Journal of Fluid Mechanics*, **244**, 527–527. 11

WANG, M. & CHRISTOV, I.C. (2018). Cutting and shuffling with diffusion: Evidence for cut-offs in interval exchange maps. *Physical Review E*, **98**, 022221. 4, 37, 41, 42, 68, 73, 76, 86, 87, 92, 107, 113, 114, 117, 222

WELANDER, P. (1955). Studies on the general development of motion in a two-dimensional, ideal fluid. *Tellus*, **7**, 141–156. 11

WIGGINS, S. (2003). *Introduction to applied nonlinear dynamical systems and chaos*, vol. 2. Springer Science & Business Media. 13

WIGGINS, S. (2013). *Chaotic transport in dynamical systems*, vol. 2. Springer Science & Business Media. 1

WONHAS, A. & VASSILICOS, J. (2002). Mixing in fully chaotic flows. *Physical Review E*, **66**, 051205. 24, 32, 34, 35, 41, 78, 124, 144, 159, 161, 163, 166, 180, 181, 182, 189, 221, 231, 232

YU, M., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2016). Cutting and shuffling of a line segment: Effect of variation in cut location. *Int. J. Bifurc. Chaos*, **26**, 1630038. ix, x, 30, 38, 67, 68, 69, 72, 73, 74, 75, 81, 92, 94, 95, 97, 99, 104, 117, 118, 196, 216, 222

YU, M., UMBANHOWAR, P.B., OTTINO, J.M. & LUEPTOW, R.M. (2019). Pattern formation in a fully three-dimensional segregating granular flow. *Physical Review E*, **99**, 062905. 39

ZAGGOUT, F.A. & GILBERT, A.D. (2012). Passive scalar decay in chaotic flows with boundaries. *Fluid Dynamics Research*, **44**, 025504. 35

ZHANG, Y. (2012). *Applications of transfer operator methods to the dynamics of low-dimensional piecewise smooth maps*. Ph.D. thesis. 161