Statistical properties of chaotic systems: from 1D maps to high dimensions

Caroline Wormell

Supervised by Georg Gottwald

A thesis submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

Applied Mathematics
University of Sydney

April 2020

Statement of originality

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

Caroline Wormell, December 24, 2019

Contents

Statement of originality **3**

List of Figures **6**

List of Tables **8**

Acknowledgements **9**

List of Related Manuscripts **10**

Abstract **12**

Chapter 1. Introduction **13**

Chapter 2. Spectral Galerkin methods for transfer operators in uniformly expanding dynamics **18**

2.1. Introduction 18
2.2. Set-up and main theorems 21
2.3. Proofs of results 27
2.4. Algorithms 40
2.5. Numerical results 43
2.6. Discussion 47

Chapter 3. Accurate and efficient numerical methods for intermittent dynamics **51**

3.1. Introduction 51
3.2. Main theorems 54
3.3. Return maps and Abel function 57
3.4. Calculating statistical properties via inducing 61
3.5. Numerical results 65
3.6. Discussion 67

Chapter 4. Linear response theory for macroscopic observables in high-dimensional deterministic systems **68**

4.1. Introduction 68

4.2. Linear response theory 71
4.3. Model 73
4.4. Summary of results 76
4.5. Macroscopic observables of uncoupled microscopic subsystems 78
4.6. Linear response of macroscopic observables of distinguished variables driven by microscopic subsystems 88
4.7. Linear response of macroscopic observables of microscopic subsystems with mean field coupling 96
4.8. A homoclinic tangency in the macroscopic dynamics of a mean-field coupled system 112
4.9. Discussion 120

Chapter 5. Conclusion 123

Appendix A. Some results relevant for Chapter 2 126
A1. The relationship between uniform expansion and uniform C-expansion . 126
A2. Results on conditions (DD_1) and (P) 129
A3. Proof of Lemma 2.12 131
A4. Explicit bounds on the norm of the solution operator in BV 135

Appendix B. Proof of Lemmas 3.5-3.8 138

Appendix C. Statistical test for linear response given time series ... 141

Appendix D. Numerical methods used in Chapter 4 143
D1. Model reduction for chaotic microscopic sub-systems 143
D2. Numerical method to compute the thermodynamic limit $M \to \infty$ for uniformly expanding maps 146

Bibliography 149

List of Figures

2.1 Heatmaps of transfer operator coefficients . . . 26
2.2 Density of acim of Lanford map . . . 44
2.3 Convergence of spectral Galerkin acim estimates for Lanford map . . . 46
2.4 Acim for circle map using adaptive algorithm . . . 47
2.5 Convergence of acim estimates of C^4 circle map . . . 48

3.1 Dynamics of LSV map . . . 52
3.2 Induced map of the LSV map . . . 52
3.3 Abel function of an LSV map . . . 55
3.4 Acims of LSV maps . . . 66

4.1 General set-up of high-dimensional systems in Chapter 4 . . . 72
4.2 Plot of logistic map, perturbations and coupling function . . . 74
4.3 Raised cosine distribution . . . 75
4.4 Mean response for logistic maps with smooth parameter distribution . . . 82
4.5 Acim of $q_{n+1}^{(j)}$ averaged over different parameter distributions . . . 83
4.6 Mean response for inhomogeneous collection of 2D maps . . . 86
4.7 Individual responses for collection of 2D maps . . . 87
4.8 Mean response for logistic maps with discrete parameter distribution . . . 88
4.9 Empirical probability distribution of distinguished variable Q . . . 90
4.10 Moments of distinguished variable Q . . . 91
4.11 Response for distinguished variable with diffusive scaling . . . 92
4.12 Empirical probability density of macroscopic variable for $\gamma = 1$ . . . 94
4.13 Response for distinguished variable with deterministic scaling . . . 95
4.14 Response for distinguished variable with short time series . . . 96
4.15 Sketch of macroscopic dynamics mediated by the microscopic reservoir. . 99
4.16 Response for coupled uniformly expanding maps with trivial dynamics . . 102

4.17 Response for coupled logistic maps with smoothly distributied parameters 103
4.18 Response discrepancy for coupled logistic maps at finite size 104
4.19 Response for stochastically driven logistic map for various noise sizes 105
4.20 Response for coupled logistic maps with discrete parameter distribution . 107
4.21 Plot of (4.43) and its invariant measure for coupling variable fixed 108
4.22 Attractor and bifurcation diagram of (4.43) 109
4.23 Response for coupled uniformly expanding maps 110
4.24 Dynamics and attractor for coupled logistic maps 111
4.25 Spectrum of macroscopic dynamics' fixed point Jacobian 116
4.26 Homoclinic tangency 118
4.27 Norms associated with homoclinic tangency orbit 119
4.28 Histogram of L^2 distances of time series from fixed point 120

A.1 A uniformly expanding map which is not uniformly C-expanding 126

List of Tables

4.1 Summary of main linear response theory results 77

Acknowledgements

I thank Georg Gottwald for his guidance, enthusiasm and firm support over the years. He has been a wonderful supervisor.

Thank you also to the many people of over the last four years I've had discussions with and asked things of: in particular, Alexey Korepanov for his request for numerics on intermittent systems and patience with me since then, Ian Melbourne for a discussion on intermittent maps, and Sheehan Olver for conversations about ApproxFun.

Thank you to my parents, who as well as feeding and housing me for a lot of this time have given me a lot of valuable advice.

Thank you to my fellow postgraduates for the help and encouragement in maths and administration, hang-out sessions, outings and often delightful friendship: to Maddie , Ishraq, Adarsh, Carol, Wenqi, Chunlei, Eric and everyone else, including at the overseas institutions I visited. Thanks also to my other friends and housemates for good times.

During this thesis I was supported by a Research Training Program stipend, and obtained travel funding through the Postgraduate Research Support Scheme, the K.E. Bullen Prize, and the SFB 1294 (Data Assimilation) Research Fellowship.

List of Related Manuscripts

This thesis is partly based on the following published articles:

- Chapter 2 of this thesis is published as Wormell (2019).
- Chapter 4 was published as Wormell and Gottwald (2019), with the following exceptions:
  - Section 4.6, the last two paragraphs of Section 4.3, and small parts of Section 4.4 were published in Wormell and Gottwald (2018);
  - Section 4.8 and the second and last paragraphs of Section 4.9 are not based on published material.

  Both aforementioned papers were joint work with Georg Gottwald. In both cases the research direction and mathematical analysis were developed jointly, and I contributed to the writing of the manuscript, performed the numerical experiments and created the figures.

For all of the above papers, I was corresponding author.

I also developed the numerical software package Poltergeist.jl, discussed in Chapter 2, and made public at `https://github.com/wormell/Poltergeist.jl` as well as in the official Julia package repository.

Caroline Wormell, April 17, 2020

As supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

Georg Gottwald, April 17, 2020

Abstract

Many important physical systems, such as the Earth's climate, are chaotic: as such we would like (probabilistic) predictions of these systems into the far future. These systems' long-term statistical behaviour is mathematically encoded by various objects, which can be studied functional-analytically using the so-called transfer operator. While rigorous study of many simple, usually one-dimensional, chaotic systems is theoretically tractable, for more complex, high-dimensional systems it is not: as a result the statistical properties of simple chaotic systems are often used as models for those of more complex systems, an idea that has been partially formalised as the Gallavotti-Cohen chaotic hypothesis.

To study these simple dynamical systems better, we will in the first part of this thesis rigorously develop efficient, powerful numerics for two classes of one-dimensional maps: uniformly expanding Markov maps, and intermittent maps. To do this, we harness the smooth structure of these problems, in particular by discretising transfer operators using spectral basis functions. We obtain highly accurate numerical estimates of statistical properties of these maps: later, we apply the methods profitably to a numerical continuation problem associated with more complex systems.

In the second part of the thesis we investigate high-dimensional systems, focusing on the differentiability of the response of statistical properties to dynamical perturbations (linear response theory). Although there are rigorous examples of one-dimensional maps that do not have differentiable responses, it is commonly believed that complex, high-dimensional chaotic systems generally do. We examine this belief through a study of model classes of "high-dimensional systems". We provide a comprehensive picture of the response properties of these systems and give broadly-applicable criteria governing the response's differentiability. In particular, we find classes of maps that, despite being composed of microscopic subsystems with non-differentiable responses, obey linear response theory, and vice versa.

Chapter 1

Introduction

Of the physical processes in the world that are most vital to us, a great many are chaotic dynamical systems such as the Earth's climate system. We are naturally interested in predicting the future behaviour of chaotic systems. Unfortunately, it is almost by definition impossible to derive information about the future state of a chaotic system far into the future, given uncertain initial conditions (Lorenz, 1963).

On the other hand, "lack of information" in fact corresponds to a probabilistic structure: in the far future one will almost certainly see the same statistical behaviour regardless of initial conditions. For the Earth's atmospheric and oceanic system, this means that—modulo slower changes in the Earth's climate over time, for example caused by changes in solar radiation due to oscillatory changes in the earth's axis and orbit around the sun (so-called Milankovitch cycles) or caused by human greenhouse gas emissions—we have some expectations of temperatures, rainfall, likelihoods of extreme events, temporal autocorrelations of climate events, sensitivity to forcing (for example by the aforementioned processes) and so on at times beyond the predictive time length of any single forceast.

These statistical properties may be codified through various mathematical objects, which typically give rise to convergence in a weak sense (see for example Lasota and Mackey (1994)). Suppose our dynamical system is a map $f : X \to X$, and we have a smooth observable function $A : X \to \mathbb{R}$. These maps are typically endowed with *invariant measures* on the phase space X: they are invariant in the sense that they are equal to their push-forwards under the map f. These measures encode geometric structure of the system and may be studied using ergodic theory, but typically a chaotic map possesses an uncountable number of invariant measures, even ergodic invariant measures, many of which are supported on dynamically unstable sets of zero Lebesgue measure. When it comes to studying the long-time statistics of a (topologically mixing) chaotic system, the measure of interested is a *physical measure*. These are invariant measures μ such that Birkhoff averages converge to averages over μ for Lebesgue-almost all initial conditions (Young, 2002):

$$N^{-1}\sum_{n=1}^{N} A(f^n(x)) \xrightarrow{N\to\infty} \int_X A \,\mathrm{d}\mu, \ x \text{ a.e..} \tag{1.1}$$

The physical measure represents the asymptotic likelihood of a system being in a certain state over long times. One may then consider various ways to quantitatively measure the convergence to this "average behaviour": for example, diffusion coefficients $\sigma_f^2(A)$ can be given to quantify corrections in a central limit theorem to the limit in (1.1). Furthermore, supposing one has a smooth family of maps f^ε parameterised by ε, one may also study the so-called response of the respective physical measures μ^ε to small changes in ε: typically this is done in a weak sense by studying the regularity and derivatives of functions of the perturbation:

$$\varepsilon \mapsto \int_X A \, \mathrm{d}\mu^\varepsilon, \tag{1.2}$$

for A a smooth observable. In particular, the derivatives of this response, where they exist, may be calculated using linear response theory (LRT) (Ruelle, 2009a; Baladi, 2014).

The tool that has, over the last twenty years, become standard for theoretical and practical study of these statistical properties is the transfer (or Perron-Frobenius) operator (Young, 1998; Baladi, 2000, 2017). This is a linear endomorphism $\mathcal{L}$ on $L^1(X)$ (or some other function space on X) associated to f which essentially tracks the evolution of measure (i.e. likelihoods of events) on the phase space of the dynamical systems. In hyperbolic systems (which globally enforce a splitting between stable and unstable directions in phase space) and for certain other, low-dimensional systems, transfer operator techniques have been used to prove existence and give formulae for many statistical properties: for hyperbolic diffeomorphisms and flows these include the existence of a physical measure (Lasota and Yorke, 1973; Gora and Boyarsky, 1989) and bounds on mixing rates (Baladi, 2000; Field et al., 2007; Baladi and Liverani, 2012; Korepanov et al., 2016), and for hyperbolic diffeomorphisms, large deviations laws for Birkhoff sums (Melbourne and Nicol, 2008; Crimmins and Froyland, 2019b) and the existence of an almost sure invariance principle, i.e. a functional central limit theorem for Birkhoff sums (Melbourne and Nicol, 2005; Melbourne and Török, 2004; Gouëzel, 2010) and (Dolgopyat, 2004; De Simoi and Liverani, 2015, 2016).

While hyperbolic systems have been generally well-studied (Pollicott, 1985; Ruelle, 1997; Baladi, 2017), as have one-dimensional non-hyperbolic systems (Rugh, 1999; Gouëzel, 2004a; Ruelle, 2009b; Baladi and Smania, 2012), research into most deterministic non-hyperbolic systems in dimensions greater than one has been rather limited. This is a result of the extraordinarily complicated structure of the geometry that emerges (Benedicks and Carleson, 1991), which makes even some two-dimensional model systems intractable to altogether rigorous study (Blumenthal et al., 2017). Because of this perhaps, simple one-dimensional systems have often come to stand for the more complex non-hyperbolic ones: in practice many of

the same phenomena known for low-dimensional simple chaotic systems manifest in the more complex chaotic systems as well, and as such the hope is that a "good" understanding of the theoretically tractable systems informs study of systems that occur in applications[1].

As well as being integral to exploring systems of direct interest, numerics are a useful tool for the theorist studying model systems. Outsourcing the work of performing calculations to a computer enables discovery or visualisation of examples and phenomena that a mathematician working alone may not themselves easily grasp or come to. With the use of rigorously validated interval arithmetic they may even be used to prove mathematical results (Tucker, 2011). Rigorously justified numerical methods have been developed for many simple chaotic systems. Often this is done through discretisations of the transfer operator, in particular via Ulam's method (Dellnitz and Junge, 1999; Froyland, 1999; Dellnitz et al., 2001; Galatolo and Nisoli, 2014; Bahsoun et al., 2016; Crimmins and Froyland, 2019b), or by exploiting the theory of dynamical zeta functions which is an alternate, and in certain cases highly effective, method of extracting statistical properties (Pollicott and Jenkinson, 2000; Jenkinson and Pollicott, 2005; Pollicott and Vytnova, 2016; Jenkinson et al., 2018). These existing methods, however, have certain limitations. Ulam's method operates only in function spaces of low regularity: consequently, it exhibits slow convergence (Bose and Murray, 2001). Moreover, the insufficient degree of regularity restricts it from deducing anything on the linear response of deterministic systems, i.e. the derivatives of functions (1.2) (Bahsoun et al., 2018). On the other hand, zeta function methods have better convergence but are only practicable for analytic maps with a small number of branches, as they require computation of exponentially growing numbers of periodic orbits (Jenkinson et al., 2018).

In the first part of my thesis I will develop transfer operator-based numerical methods for exemplary classes of one-dimensional maps, which harness good regularity properties, in particularly of Fourier and Chebyshev approximations (Trefethen, 2013), for fast convergence: in Chapter 2 for uniformly expanding maps, and in Chapter 3 for non-uniformly expanding intermittent maps. Through efficient and extremely accurate estimation of quantities such as physical measures and mixing rates, these methods enable exploration of these simple systems, which can be useful in making broader statements about chaotic systems.

[1]Hyperbolic systems are themselves hardly seen in the world, the most well-known and indeed common example being chain linkages (Hunt and MacKay, 2003; Magalhães and Pollicott, 2013).

Coming back to the premise invoked above that the statistical behaviour of general chaotic dynamical systems can be deduced from simple one-dimensional maps, we focus in the second part of this thesis on linear response theory, a salient area where this premise may not hold. Linear response theory has long been believed to apply in higher dimensions, which has been corroborated by many successes in applying linear response theory to a large variety of climate models (Bell, 1980; North et al., 1993; Majda et al., 2010; Lucarini and Sarno, 2011; Cooper et al., 2013; Gritsun and Dymnikov, 1999; Dymnikov and Gritsoun, 2001; Ring and Plumb, 2008; Ragone et al., 2016; Lembo et al., 2019), although this success is not universal (Kirk-Davidoff, 2009; Cooper and Haynes, 2013; Chekroun et al., 2014). To explain this, physicists have invoked the chaotic hypothesis of Gallavotti and Cohen (1995a,b), under which the dynamics of complex chaotic systems on their attractors are hypothesised to be "for all intents and purposes" hyperbolic systems. This invocation, however, is not valid: while hyperbolic systems do have linear response (Ruelle, 1997), the chaotic hypothesis only speaks to the dynamics when restricted to the attractor (for one fixed parameter): linear response depends on the manifold structure of the systems in a neighbourhood of the attractor as well (Ruelle, 2009a). Moreover, from the work of Baladi and others logistic maps, which are a family of one-dimensional non-hyperbolic maps, are known to have responses which are not differentiable, even in the sense of Whitney (Baladi and Smania, 2008, 2010; Baladi, 2014). This motivates the question: how might linear response theory be valid for macroscopic observabless in high-dimensional systems when more simple systems, which may in fact be their constituents, do not have linear response?

As a first model for the broad class of "complex chaotic systems" we will study large ensembles of simple one-dimensional chaotic systems which may either evolve independently or be coupled via a mean-field, and we will consider the response of the mean fields of these ensembles. In Chapter 4 we will combine a (non-rigorous) stochastic reduction of the macroscopic dynamics of these systems with a numerical study to obtain a comprehensive survey of the linear response properties of these systems. In Section 4.5 we will make a detailed study of the uncoupled case. Through this study, we find a system where the microscopic variables are logistic maps whose parameters are randomly selected from a distribution: we find that despite the microscopic variables not having LRT a linear response emerges in the macroscopic mean field observable. Following on from this in Section 4.6 we will consider the linear response theory of a distinguished variable driven by the uncoupled system, a standard scenario treated in many statistical physics works (Ford et al., 1965; Zwanzig, 1973; Kupferman et al., 2002). In Section 4.7 we will consider the linear response for a system that is coupled to itself via a mean-field. We find that for systems in finite size generate emergent stochastic behaviour in the mean-field, which induces

LRT in the system, regardless of the composition of the microscopic variables. Furthermore, we will describe a self-coupled system, comprised of microscopic variables that in isolation satisfy LRT, which nonetheless in the thermodynamic limit fails to satisfy linear response; in Section 4.8 we will, using the numerics developed in Chapter 2, find a homoclinic tangency in this system. This implies that the macroscopic dynamics are actually non-hyperbolic, which is a violation of the Galavotti-Cohen hypothesis itself (Gallavotti and Cohen, 1995a,b; Gallavotti, 2019).

Furthermore, in Chapters 2 and 4 we will provide more detailed background on the statistical properties of one-dimensional maps and linear response theory, respectively.

Chapter 2

Spectral Galerkin methods for transfer operators in uniformly expanding dynamics

2.1. Introduction

One-dimensional full-branch Markov uniformly expanding maps are an important class of chaotic dynamical systems: as well as being common toy models, more complex chaotic systems may be reduced to this class, for example by constructing carefully-chosen return maps Young (1998). Mathematically, these maps are endomorphisms f on a compact and connected one-dimensional manifold Λ for which there are a set of disjoint open intervals $(\mathcal{O}_\iota)_{\iota\in I}$ of full measure such that, for all $\iota \in I$, $f|_{\mathcal{O}_\iota}$ is injective with $\overline{f(\mathcal{O}_\iota)} = \Lambda$, and on $\mathcal{O}_\iota$, f is differentiable with $|f'| \geq \lambda > 1$. A standard example of such a map is the Lanford map, defined on $[0, 1]$ with $f(x) = 2x + \frac{1}{2}x(1 - x) \mod 1$ Lanford (1998).

Many significant properties of these maps can be determined from the transfer operator $\mathcal{L} : BV(\Lambda) \to BV(\Lambda)$, which describes the action of the map f pushing forward signed measure densities on the phase space. The transfer operator has the property that for all $\phi \in BV(\Lambda)$ and $\psi \in L^\infty(\Lambda)$,

$$\int_\Lambda \mathcal{L}\phi\, \psi \,\mathrm{d}x = \int_\Lambda \phi\, \psi \circ f \,\mathrm{d}x,$$

and for the maps we consider has explicit action

$$(\mathcal{L}\phi)(x) = \sum_{f(y)=x} \frac{1}{|f'(y)|}\phi(y), \tag{2.1}$$

where $BV(\Lambda)$ denotes the space of functions of bounded variation on Λ. The prime example of a mathematical object that can be derived from the transfer operator is that the density ρ of the system's physical measure, which is an absolutely continuous invariant measure (acim), is the unique 1-eigenfunction of $\mathcal{L}$ with $\int_\Lambda \rho\, dx = 1$. Another example is that the central limit theorem for the long-time average of an observable $\phi \in BV(\Lambda)$ for initial condition x_0 sampled from a BV density can be written

$$\frac{\sum_{i=0}^n \phi(f^i(x_0)) - n\mathbb{E}\phi}{\sigma_f(\phi)\sqrt{n}} \xrightarrow{n\to\infty}_d N(0, 1),$$

where the expectation value $\rho(\phi)$ is given

$$\rho(\phi) = \int_\Lambda \phi\, \rho\, dx \tag{2.2}$$

and the diffusion coefficient $\sigma_f(\phi)$ is defined via the Green-Kubo formula

$$\sigma_f^2(\phi) = \int_\Lambda \phi \sum_{n=-\infty}^{\infty} \mathcal{L}^{|n|}\big((\phi - \rho(\phi))\ \rho\big)\, dx. \tag{2.3}$$

Transfer operator problems cannot in general be solved analytically, and numerical approaches are therefore of prime importance. One scheme that has been widely studied in the literature is Ulam's method. In Ulam's method the transfer operator is projected onto the subspace generated by the characteristic functions of this partition of the phase space and statistical properties are estimated using the projected transfer operator (Dellnitz and Junge, 1999; Dellnitz et al., 2001). Ulam's method is effective for a range of families of chaotic systems (Froyland et al., 2011; Froyland, 2007; Murray, 2010), and in particular Ulam estimates for a variety of statistical quantities have been proven to converge for uniformly expanding maps (Froyland, 2007; Bahsoun et al., 2016). Higher-order extensions of Ulam's method to compactly supported piecewise continuous functions have also been used (Tian et al., 2016), in particular to compute quantities such as linear response that require a higher degree of regularity (Bahsoun et al., 2018); theory for a wavelet-based method has also been developed (Holschneider, 1996). On the other hand Pollicott, Jenkinson and others have presented a completely different approach, wherein one computes statistical properties using the theory of dynamical zeta functions: this involves computing sums over periodic orbits of the system (Jenkinson and Pollicott, 2005; Jenkinson et al., 2018). Zeta function-based methods have furnished the most accurate estimates in the literature, but due to the rate of exponential growth of the number periodic points of a given length are functionally limited to analytic maps possessing only a few branches.

In this chapter we will pursue a different approach: we will construct a so-called spectral Galerkin approximation, whereby one studies the transfer operator of the dynamical system of interest and functions it acts on in a basis of orthogonal polynomials and orthogonally project to finite-dimensional spaces $E_N \subseteq BV$ spanned by low-index elements in the basis. Spectral methods are the gold standard for solving high regularity operator problems in other areas, including ODE and PDE theory Boyd (2001); Trefethen (2013). In the realm of chaotic systems, spectral methods have previously been applied to transfer operators of stochastic flows (Froyland et al., 2013; Denner et al., 2015): the theory is much more straightforward here than in the deterministic case, as elliptic operator theory can be applied.

Since the publication of the work in this chapter, spectral methods have also been rigorously applied in other deterministic settings beyond the one we consider, in particular to the transfer operator eigendata of Anosov maps in higher dimensions (Crimmins and Froyland, 2019a) and of analytic expanding maps of the interval (Bandtlow and Slipantschuk, 2020).

We consider the Fourier exponential basis

$$e_k(x) = e^{ikx},\ k \in \mathbb{Z}, \tag{2.4}$$

which is orthogonal in $L^2([0, 2\pi))$, and the Chebyshev polynomial basis

$$T_k(x) = \cos k \cos^{-1} x,\ k \in \mathbb{N}, \tag{2.5}$$

which is orthogonal on $[-1, 1]$ with respect to the weight $(1 - x^2)^{-1/2}$.

In our theoretical results we find that, providing the maps under consideration exhibit sufficient regularity, these kinds of spectral methods provide up to exponentially fast convergence with a small numerical outlay. Our main theoretical results are that spectral Galerkin estimates of acims and the resolvent of the transfer operator $\mathcal{L}$ at the eigenvalue 1 converge exponentially fast in the approximation order N for analytic maps, and as $O(N^{2.5-r})$ for C^r maps (we will be more specific about the kinds of map we consider in Section 2.2.1). The algorithmic outlay of our method is $O(N^3)$. These results are summarised respectively in Corollaries 2.4 and 2.5 in Section 2.2.2.

We obtain these results by defining a so-called solution operator that allows one to access transfer operator resolvent data at eigenvalue 1 (Theorem 2.1 in Section 2.2.2), and then showing that spectral Galerkin approximations of this solution operator converge at the aforementioned rates (Theorem 2.3). These rates of convergence are determined via bounds on entries of the spectral basis matrix representations of transfer operators, proved in Theorems 2.6-2.7.

This theoretical work carries out some of the directions for further research suggested in Baladi and Holschneider (1999), which proved convergence of eigenvalue and eigenvector estimates of transfer operators of circle maps in a wavelet basis (a transformation of a Fourier basis). In particular, we extend from periodic intervals to non-periodic intervals, and establish quantitative convergence rates for the invariant density and resolvent data at the eigenvalue 1.

The chapter is structured as follows. In Section 2.2 we define the classes of maps pertinent to our results, and introduce the main theorems, which we then prove in Section 2.3. In Section 2.4 we describe the algorithms we use that demonstrate the possiblities of transfer operator spectral methods, and in Section 2.5 we present numerical results illustrating these algorithms. Finally in Section 2.6 we discuss the methods presented and consider possible extensions.

2.2. Set-up and main theorems

We first set up the problem, introducing the maps under consideration. We then present the major results: Theorem 2.1 characterising an operator that explicitly solves many typical transfer operator problems; the main theorem, Theorem 2.3, which gives convergence of spectral operator estimates and Corollaries 2.4 and 2.5, which give convergence of acims and other statistical properties as a result; and finally Theorems 2.6 and 2.7 bounding the magnitude of transfer operator spectral coefficient matrix entries L_{jk}, which are central to the proof of Theorem 2.3. We finally present two rigorously validated bounds on the Lyapunov exponent and a diffusion coefficient of the Lanford map, obtained via a rigorous implementation of our spectral method.

2.2.1. Dynamical systems under consideration. We first introduce the two generic classes of maps we will consider; we will then introduce a set of so-called distortion conditions that maps from these classes may optionally hold, and which determine the spectral method's rates of convergence.

2.2.1.1. *Classes of maps.* We define two main classes of maps: circle maps U_P and interval maps U_{NP}. Maps in U_P are defined on the one-dimensional torus and must be continuous and differentiable on the whole domain, whereas maps in U_{NP} are defined on a (non-periodic) interval, and there is no requirement for any continuity or differentiability between branches of the map. For example, a Markovian tent map may lie in U_{NP}, whereas maps in U_P must have a derivative defined everywhere.

A map $f : \Lambda \to \Lambda$ is in U_P if it satisfies the following axioms:

- Its domain Λ is a circle, which we suppose to be canonically $\mathbb{R}/2\pi\mathbb{Z}$.
- It is piecewise C^2 with Lipschitz-bounded distortion, that is,

$$\sup_{x\in\Lambda} \frac{|f''(x)|}{|f'(x)|^3} < \infty.$$

- It is uniformly expanding, that is,

$$\lambda := \inf_{x\in\Lambda} |f'(x)| > 1. \tag{E}$$

Maps in U_P are circle maps, and can be extended to bijective lifts $\hat{f} : [0, 2\pi] \to [0, 2\beta\pi]$ for some $\beta \in \{2, 3, 4, \ldots\}$. We denote the inverse of $\hat{f}$ by $v = \hat{f}^{-1}$, and for consistency with the notation for U_{NP} define $v_\iota(x) := v(x + 2\iota\pi)$ for $x \in [0, 2\pi]$ and $\iota \in I := \{0, 1, \ldots, \beta - 1\}$.

A map $f : \Lambda \to \Lambda$ is in U_{NP} if it satisfies the following axioms:

- Its domain Λ is an interval, which we suppose to be the canonical interval for Chebyshev expansions $[-1, 1]$.
- It is full-branch Markov, i.e. there are open disjoint intervals $\mathcal{O}_\iota, \iota \in I$ whose union is of full measure in Λ such that $f|_{\mathcal{O}_\iota}$ extends continuously to a bijective function $\hat{f}_\iota : \overline{\mathcal{O}_\iota} \to \Lambda$.
- The functions $\hat{f}_\iota$ are all C^2 and furthermore, the map f has Lipschitz-bounded distortion. We will find it useful to formulate the Lipschitz distortion condition in terms of $v_\iota := \hat{f}_\iota^{-1}$, as

$$\sup_{x\in\Lambda,\iota\in I} \left| \frac{v_\iota''(x)}{v_\iota'(x)} \right| < \infty. \tag{DD$_1$}$$

- It satisfies a *uniform C-expansion condition*[1]

$$\check{\lambda} = \inf_{x\in\cup_{\iota\in I}\mathcal{O}_\iota} \frac{\sqrt{1-x^2}}{\sqrt{1-f(x)^2}} |f'(x)| > 1. \tag{CE}$$

- It satisfies a *partition spacing condition*

$$\sup \left\{ \frac{|\mathcal{O}_\iota|}{d(\mathcal{O}_\iota, \partial\Lambda)} : d(\mathcal{O}_\iota, \partial\Lambda) > 0 \right\} = \boldsymbol{\xi} < \infty. \tag{P}$$

The latter two conditions we introduce to control the high oscillatory behaviour of the spectral basis functions' images under the action of the transfer operator near the endpoints of the interval. They are not especially onerous conditions: uniformly expanding maps typically satisfy (CE), and a uniformly expanding map satisfying all conditions of U_{NP} except (CE) will have an iterate in U_{NP} (see Appendix A1 for a discussion of C-expansion); (P) is always satisfied for maps with finitely many branches. The bounded distortion condition (DD$_1$) is a standard regularity condition necessary for, among other things, a spectral gap in BV.

We also consider maps that satisfy the conditions of U_P (resp. U_{NP}) except that the associated expansion parameter in (E) (resp. (CE)) need only be positive, rather than strictly greater than 1. We denote the class of such maps $\bar{U}_P$ (resp. $\bar{U}_{NP}$).

2.2.1.2. *Distortion conditions.* To obtain good convergence results we will optionally impose the following generalised distortion conditions on our maps.

The first set of distortion conditions are equivalent to uniform bounds on derivatives of the distortion $\log|v_\iota'|$. A map satisfies distortion condition (DD$_r$) for some $r \in \mathbb{N}^+$ if

$$\sup_{\iota\in I, x\in\Lambda} \left| \frac{v_\iota^{(n+1)}(x)}{v_\iota'(x)} \right| = C_n < \infty,\ n = 1, \ldots, r. \tag{DD$_r$}$$

[1]This condition can be reformulated as requiring $|(\cos^{-1} \circ f \circ \cos)'| \geq \check{\lambda} > 1$.

The second set of distortion conditions are equivalent to uniform bounds on the first derivative of the distortion on a complex neighbourhood of the map's domain Λ. For circle maps, the neighbourhood is the closed complex strip $\Lambda_\delta^\beta = \{x+iy \mid x \in \mathbb{R}/2\beta\pi\mathbb{Z}, |y| \leq \delta\}$, for a given $\delta > 0$. For intervals, the neighbourhood is $\check{\Lambda}_\delta$, defined to be a Bernstein ellipse[2] of parameter e^δ. We assume that v' and v'_ι respectively extend holomorphically to these sets. A map satisfies (AD_δ) for some $\delta > 0$ if

$$(\mathrm{AD}_\delta) \qquad \begin{cases} \sup_{z\in\Lambda_\delta^\beta} \left|\frac{v''(z)}{v'(z)}\right| = C_{1,\delta} < \infty, & \Lambda = \mathbb{R}/2\pi\mathbb{Z}, \\ \sup_{\iota\in I, z\in\check{\Lambda}_\delta} \left|\frac{v''_\iota(z)}{v'_\iota(z)}\right| = C_{1,\delta} < \infty, & \Lambda = [-1,1]. \end{cases}$$

We associate with each distortion condition a *spectral rate of convergence*. We formulate these rates of convergence $\kappa(\cdot)$ as function classes:

$$\kappa\left(\mathrm{DD}_r\right) = \left\{x \mapsto C(1+x)^{-r} : C > 0\right\}, \tag{2.6}$$

$$\kappa\left(\mathrm{AD}_\delta\right) = \left\{x \mapsto Ce^{-\zeta x} : C > 0,\ \zeta \in (0,\delta]\right\}. \tag{2.7}$$

2.2.2. Main results. Before formulating the main theoretical results of this paper, we introduce an operator derived from the transfer operator that explicitly generates acims and other statistical properties. Many statistical properties can be computed using resolvent data of $\mathcal{L}$ at its eigenvalue 1: the so-called *solution operator* is the inverse of a bounded perturbation of $\mathrm{id} - \mathcal{L}$ which allows the resolvent data to be recovered.

We define the *solution operator inverse*

$$\mathcal{K} = \mathrm{id} - \mathcal{L} + u\mathscr{S} \tag{2.8}$$

and the solution operator

$$\mathcal{S} = \mathcal{K}^{-1} = (\mathrm{id} - \mathcal{L} + u\mathscr{S})^{-1}, \tag{2.9}$$

where the functional $\mathscr{S}$ is the total Lebesgue integral on Λ and u is a function in the domain of $\mathcal{L}$ such that $\mathscr{S}u = 1$.

For any transfer operator $\mathcal{L}$ with a spectral gap (i.e. with a simple eigenvalue at 1 and the remaining spectrum bounded inside a disk of radius less than 1), the solution operator therefore solves for two important quantities, according to the following theorem:

Theorem 2.1. *Let $\mathcal{L} : E \to E$ be a transfer operator with a spectral gap. Choose $u \in E$ with $\mathscr{S}u = 1$.*

Then $\mathcal{S} = (\mathrm{id} - \mathcal{L} + u\mathscr{S})^{-1}$ is well-defined and bounded as an operator on E, and

[2] A Bernstein ellipse of parameter $\rho > 1$ is an ellipse in the complex plane centred at 0 with semi-major axis of length $\frac{1}{2}(\rho + \rho^{-1})$ along the real axis and semi-minor axis $\frac{1}{2}(\rho - \rho^{-1})$ along the imaginary axis.

(a) If ρ is the unique acim with $\mathscr{S}\rho = 1$,

$$\rho = \mathcal{S}u. \tag{2.10}$$

(b) For any $\phi \in \ker \mathscr{S}$,

$$\sum_{n=0}^{\infty} \mathcal{L}^n \phi = \mathcal{S}\phi. \tag{2.11}$$

Remark 2.2. *As a result of Theorem 2.1, many important statistical quantities can be simply expressed using the solution operator and the acim. For example, the Green-Kubo formula for the diffusion coefficient given in (2.3) can be rewritten using Theorem 2.1(b) as*

$$\sigma_f^2(A) = \int_\Lambda A\ (2\mathcal{S} - \mathrm{id})(\mathrm{id} - \rho\mathscr{S})(\rho A)\, dx. \tag{2.12}$$

This closed formula enables effective rigorous calculation of diffusion coefficients, as we will see in Section 2.5.

We now provide some notation to enable us to state the main theorem, which proves the convergence of the spectral methods. We define the finite-dimensional subspaces $(E_N)_{N\in\mathbb{N}^+}$

$$E_N = \begin{cases} \mathrm{span}\{e_{-N}, \ldots, e_N\}, \Lambda = [0, 2\pi) \\ \mathrm{span}\{T_0, \ldots, T_N\}, \Lambda = [-1, 1], \end{cases}$$

where the Fourier exponentials e_k and Chebyshev polynomials T_k are defined respectively in (2.4-2.5). We also define the corresponding orthogonal projections $\mathcal{P}_N$ onto the E_N in the L^2 space in which the bases are orthogonal ($L^2([0, 2\pi])$ for the Fourier basis $(e_k)_{k\in\mathbb{Z}}$, and $L^2([-1, 1], (1 - x^2)^{-1/2})$ for the Chebyshev basis $(T_k)_{k\in\mathbb{N}}$. We also define the spectral Galerkin operator discretisations, which are as endomorphisms on the finite-dimensional space E_N:

$$\mathcal{L}_N = \mathcal{P}_N \mathcal{L}|_{E_N} \tag{2.13}$$

and

$$\mathcal{S}_N := \mathcal{K}_N^{-1} := (\mathrm{id} - \mathcal{L}_N + u\mathscr{S}|_{E_N})^{-1}, \tag{2.14}$$

where the function u is taken to be in E_N. (A typical choice of u is $u = 1/|\Lambda|$.)

Our main theorem can then be formulated as follows, recalling the definitions of $\kappa(\cdot)$ in (2.6-2.7):

Theorem 2.3. *Suppose $f \in U_P$ or U_{NP}, and satisfies a distortion bound* (DD_r) *(resp.* (AD_δ)*). Then there exist functions $K, \bar{K} \in \kappa(\mathrm{DD}_r)$ (resp. $\kappa(\mathrm{AD}_\delta)$) such that for sufficiently large N and all $\phi \in E_N$,*

$$\|\mathcal{L}_N\phi - \mathcal{L}\phi\|_{BV} < N\sqrt{N}K(N)\|\phi\|_{BV},$$

and

$$\|\mathcal{S}_N\phi - \mathcal{S}\phi\|_{BV} < N\sqrt{N}\bar{K}(N)\|\phi\|_{BV}.$$

For ease of expression, in the rest of this section we use the notation (D) to denote either of (DD_r) or (AD_δ).

Theorem 2.3 together with Theorem 2.1 directly implies the convergence of estimates of statistical quantities. In particular, the following corollary gives spectral convergence of the acim.

Corollary 2.4. *Suppose $f \in U_P$ or U_{NP}, and satisfies a distortion bound* (D).

Let $\rho_N = \mathcal{S}_N u$. Then there exists $K \in \kappa(\mathrm{D})$ such that for all N sufficiently large

$$\|\rho_N - \rho\|_{BV} < N\sqrt{N}K(N).$$

The next corollary gives strong convergence of $\sum_{n=1}^{\infty} \mathcal{L}^n$, and consequently many important statistical estimates (see (2.12) for an example).

Corollary 2.5. *Suppose $f \in U_P$ or U_{NP}, and satisfies a distortion bound* (D). *Then there exists $K \in \kappa(\mathrm{D})$ such that for N large enough and all $\phi \in E_N \cap \ker \mathscr{S}$,*

$$\left\| \mathcal{S}_N\phi - \sum_{n=0}^{\infty} \mathcal{L}^n\phi \right\|_{BV} < N\sqrt{N}K(N)\|\phi\|_{BV}.$$

Since the operators $\mathcal{L}_N$ and $\mathcal{S}_N$ are endomorphisms on E_N, Theorem 2.3 and Corollary 2.5 show that the spectral method converges in operator norm within E_N. When attempting to estimate, for example, $\mathcal{S}\phi$ for some $\phi \notin E_N$, one can simply substitute ϕ for its spectral discretisation $\mathcal{P}_N\phi$, and propagate through the calculation the error arising from this substitution.

Critical to proving Theorem 2.3 are bounds on the entries L_{jk} of the transfer operator considered as a matrix with respect to the Fourier or Chebyshev orthogonal basis, which are stated in the following two analogous theorems for transfer operators on periodic and non-periodic domains: the situation is illustrated in Figure 2.1. Abstractly, these results reformulate the characterisation of the transfer operator of a uniformly-expanding map as the sum of a strictly upper-triangular operator and a compact operator developed by Holschneider (1996) and Baladi and Holschneider (1999) in the context of C^∞ circle maps in wavelet bases. The important development of our approach is the large amount of quantitative information generated, which allows us to prove convergence rates and provide rigorous concrete bounds for specific maps.

Theorem 2.6. *Suppose f is in the class $\bar{U}_P$ satisfying some distortion bound* (D)*, with $\lambda_1 \leq f' \leq \lambda_2$. Suppose L is the matrix representing the transfer operator of f in a Fourier exponential basis.*

Then for every $p_1 > \lambda_1^{-1}$ and $p_2 < \lambda_2^{-1}$ there exists $K \in \kappa(\mathrm{D})$ such that for $j/k > p_1$ or $k = 0$,

$$|L_{jk}| \leq K(|j - p_1 k|),$$

and for $j/k < p_2$ or $k = 0$,

$$|L_{jk}| \leq K(|j - p_2 k|).$$

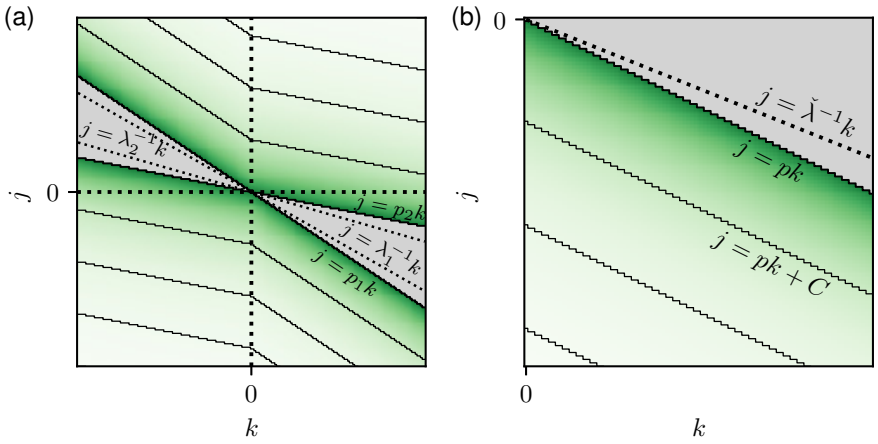

FIGURE 2.1. Heatmaps of maximum possible magnitudes of coefficients L_{jk} of transfer operator matrix of a system described in: (a) Theorem 2.6; (b) Theorem 2.7. Shown are contours of constant magnitude of coefficients (thin black lines) and, in grey block colour, coefficients not characterised by the theorems. Note that because the full Fourier and Chebyshev bases are indexed by $\mathbb{Z}$ and $\mathbb{N}$ respectively, the matrix indices range over these values.

Theorem 2.7. *Suppose f is in the class $\bar{U}_{NP}$, satisfying some distortion bound* (D)*. Suppose L is the matrix representing the transfer operator of f in a Fourier exponential basis.*

Then for every $p > \check{\lambda}$ there exists $K \in \kappa(\mathrm{D})$ such that for $j/k > p$ or $k = 0$,

$$|L_{jk}| \leq K(|j - pk|).$$

Remark 2.8. *One can prove similar results for transfer operators with general weights (c.f. (2.1)):*

$$(\mathcal{L}_g\phi)(x) = \sum_{f(y)=x} g(y)\phi(y). \tag{2.15}$$

This class of operator includes transfer operators and composition operators $C_v : \phi \mapsto \phi \circ v$.

From the previous results, we are able to prove extremely accurate rigorous bounds on maps in U_P and U_{NP} satisfying sufficiently strong distortion conditions. In particular, we prove the following bound on the Lanford map:

Theorem 2.9. *Consider the Lanford map* $f : [0,1] \to [0,1]$, $f(x) = 2x + \frac{1}{2}x(1-x) \mod 1$.

(a) The Lanford map's Lyapunov exponent $\lambda := \int_\Lambda \log|f'|\,\rho\,dx$ *lies in the range*

$$\begin{aligned}\lambda = 0.657\ 661\ 780\ 006\ 597\ 677\ 541\ 582\ 413\ 823\ 832\ 065\ 743\ 241\ 069\\ 580\ 012\ 201\ 953\ 952\ 802\ 691\ 632\ 666\ 111\ 554\ 023\ 759\ 556\ 459\\ 752\ 915\ 174\ 829\ 642\ 156\ 331\ 798\ 026\ 301\ 488\ 594\ 89 \pm 2\times 10^{-128}.\end{aligned}$$

(b) The diffusion coefficient for the Lanford map with observable $\phi(x) = x^2$ *lies in the range*

$$\begin{aligned}\sigma_f^2(\phi) = 0.360\ 109\ 486\ 199\ 160\ 672\ 898\ 824\ 186\ 828\ 576\ 749\ 241\ 669\ 997\\ 797\ 228\ 864\ 358\ 977\ 865\ 838\ 174\ 403\ 103\ 617\ 477\ 981\ 402\ 783\\ 211\ 083\ 646\ 769\ 039\ 410\ 848\ 031\ 999\ 960\ 664\ 7 \pm 6\times 10^{-124}.\end{aligned}$$

These bounds are derived in Section 2.5.1.

2.3. Proofs of results

Our attack on the theorems in Section 2.2.2 is structured as follows.

We begin by proving Theorem 2.1 characterising the solution operator. This proof uses standard quasi-compactness properties of the transfer operator.

We then turn to proving the entry bound results (Theorems 2.6 and 2.7). These results stem from more general properties of Fourier series representations of composition operators (Lemma 2.11), which we prove using oscillatory integral techniques. Because it is necessary to make a non-diffeomorphic cosine transformation to obtain Fourier basis functions from Chebyshev polynomials, some work is required to prove appropriate bounds on derivatives after the transformation.

We then go on to prove Theorem 2.3. We consider a perturbation of the transfer operator $\mathcal{L}$ that is block-upper-triangular in the relevant spectral basis (in the Fourier case, under the basis order $e_0, e_1, e_{-1}, e_2, e_{-2}, \ldots$), with the finite matrix L_N forming

the first block on the diagonal. Since the solution operator of such a perturbation is a composition of upper block-diagonal operators, the first diagonal block can thus be approximated only from knowledge of L_N (Lemma 2.15). Using that the BV-norm of our perturbation can be bounded using spectral matrix coefficients (Lemma 2.16), we obtain the main result.

We begin with the proof of Theorem 2.1, which gives the properties of the solution operator $\mathcal{S} = (I - \mathcal{L} + u\mathscr{S})^{-1}$ (see (2.9)).

Proof of Theorem 2.1. Split E as $V_\perp \oplus V$ where $V_\perp = \operatorname{span}\{u\}$ and $V = \ker \mathscr{S}$. Since $V_\perp$ and V are closed subspaces of E there exists a bounded operator $\mathcal{N} : E \to V_\perp$ such that one may also define $\mathrm{id} - \mathcal{N} : E \to V$. We now consider the action with respect to this splitting of the putative solution operator inverse, $\mathcal{K} = \mathrm{id} - \mathcal{L} + u\mathscr{S}$.

Since $\mathscr{S}(\mathcal{L} - \mathrm{id}) = 0$, we have for any element $\phi \in V$ that

$$\mathcal{K}\phi = (\mathrm{id} - \mathcal{L})\phi + u\mathscr{S}\phi = (\mathrm{id} - \mathcal{L})|_V \phi \in V. \tag{2.16}$$

Similarly, for any scalar α we have

$$\mathcal{K}\alpha u = (\mathrm{id} - \mathcal{L})(\alpha u) + u\mathscr{S}\alpha u = \alpha u + (\mathrm{id} - \mathcal{L})(\alpha u), \tag{2.17}$$

where the sum follows the splitting of $E = V_\perp \oplus V$.

Since the transfer operator $\mathcal{L}$ has a spectral gap, the spectral radius of $\mathcal{L}|_V$ is strictly less than 1, and the operator

$$\mathcal{Q} := (\mathrm{id} - \mathcal{L}|_V)^{-1} = \sum_{n=0}^{\infty} \mathcal{L}^n|_V$$

is bounded as an endomorphism on V.

Back-solving (2.16-2.17) thus gives that for any $\psi \in V$,

$$\mathcal{S}\phi = \mathcal{K}^{-1}\psi = \mathcal{Q}\psi,$$

and for any scalar α that

$$\mathcal{S}\alpha u = \alpha u - \mathcal{Q}(\mathrm{id} - \mathcal{L})\alpha u = \lim_{n\to\infty} \mathcal{L}^n \alpha u = \alpha\rho.$$

Since $\alpha\rho = \rho\mathscr{S}(\alpha u)$, we can use these results to write the solution operator

$$\mathcal{S} = \mathcal{Q}(\mathrm{id} - \mathcal{N}) + \rho\mathscr{S}\mathcal{N}, \tag{2.18}$$

which is clearly bounded.

It clearly follows from (2.18) that $\mathcal{S}u = \rho$ and $\mathcal{S}\phi = \mathcal{Q}\phi$ for $\phi \in V$. □

Remark 2.10. *The solution operator can be written as the following expression*

$$\mathcal{S} = u\mathscr{S} + \sum_{n=0}^{\infty} \mathcal{L}^n(\mathrm{id} + (\mathcal{L}u - 2u)\mathscr{S}). \tag{2.19}$$

We now set about proving Theorems 2.6 and 2.7, which place bounds on the magnitudes of the entries of transfer operator matrices in Fourier and Chebyshev bases.

We begin by proving similar kinds of bounds on the coefficients of a matrix associated with a more general operator $\mathcal{M}$ on the circle $\mathbb{R}/2\pi\mathbb{Z}$. We introduce $\mathcal{M}$ as a generalised transfer operator (2.15) where instead of using the inverse of the map, one uses a general function v which may be non-injective. Bounds on elements of the Fourier basis transfer operator matrix for $\mathcal{M}$ imply similar bounds on transfer operators in Fourier and Chebyshev bases.

Lemma 2.11. *Let v be a differentiable function from $\mathbb{R}/2\beta\pi\mathbb{Z}$, $\beta \in \mathbb{Z}^+$ to $\mathbb{R}/2\pi\mathbb{Z}$ such that $v'(\mathbb{R}/2\beta\pi\mathbb{Z}) = \tilde{\mu} = [\mu_2, \mu_1]$, and let h be a continuous function on the circle $\mathbb{R}/2\beta\pi\mathbb{Z}$.*

Let $\mathcal{M}$ be the endomorphism on $L_2([0, 2\pi])$ defined by

$$\mathcal{M}: \ \phi \mapsto \sum_{b=1}^{\beta} h\left(x + 2\pi b\right) \phi\left(v\left(x + 2\pi b\right)\right). \tag{2.20}$$

Let M be the corresponding bi-infinite matrix in the Fourier complex exponential basis.

Then:

(a) The entries of M are bounded uniformly by $\|h\|_1/2\pi$.

(b) Suppose that for $n = 1, \ldots, r$, $\sup|v^{(n+1)}| \leq \Upsilon_n < \infty$ and $\sup|h^{(n)}/h| \leq H_n < \infty$. Then there exist constants $W_{r,n}$ such that for $j \notin k\tilde{\mu}$,

$$|M_{jk}| \leq \frac{\|h\|_1}{2\pi} \sum_{n=0}^{r} \frac{W_{r,n}|k|^n}{d(j, k\tilde{\mu})^{n+r}}. \tag{2.21}$$

Each $W_{r,n}$ is bounded by a linear combination of H_l, $l \leq r - n$, whose coefficients are polynomials in Υ_l, $l \leq r - n$.

(c) Suppose v and h extend analytically to the complex strip $\Lambda_\delta^\beta = [0, 2\beta\pi) + i[-\delta, \delta]$, and on this strip $\sup|h'/h| \leq H_{1,\delta} < \infty$ and $\sup|v''| \leq \Upsilon_{1,\delta} < \infty$.

Choose any $\tilde{p} = [p_1, p_2]$ such that $\tilde{\mu} \subset \int \tilde{p}$. Define $\zeta = \min\left\{2\Upsilon_{1,\delta}^{-1} d(\tilde{\mu}, \mathbb{R}\backslash\tilde{p}), \delta\right\}$. Then $\zeta > 0$ and

$$|M_{jk}| \leq \frac{\|h\|_1}{2\pi} e^{\zeta\left(H_{1,\delta} - d(j,\tilde{p})\right)}. \tag{2.22}$$

Proof of Lemma 2.11. The matrix element M_{jk} is the jth Fourier coefficient of the function $\mathcal{M}e_k$, so using the orthogonality of Fourier bases in L^2 and (2.20), we have that

$$M_{jk} = \frac{1}{2\pi} \int_0^{2\pi} \sum_{b=1}^{\beta} h\left(x + 2\pi b\right) e^{ikv(x+2\pi b)} e^{-ijx} dx,$$

which using the 2π-periodicity of e^{ijx} we can rewrite as a single integral

$$M_{jk} = \frac{1}{2\pi}\int_0^{2\beta\pi} h(x)e^{i(kv(x)-jx)}dx. \tag{2.23}$$

We obtain (a) from this equation simply by taking absolute values.

For (b), we use that the integrand in (2.23) is oscillatory when the derivative of $kv(x) - jx$ is bounded away from zero, that is, when $j/k \notin [\mu_2, \mu_1]$. As a result, we can improve the bound we got in the first part by repeatedly integrating by parts.

Starting from (2.23), we separate the integrand into two terms

$$\left(\frac{h(x)}{i(kv'(x)-j)}\right)\left(i(kv'(x)-j)e^{i(kv(x)-jx)}\right),$$

so as to integrate by parts, differentiating the left term and integrating the right. Because the right term integrates to zero, the boundary terms in the integration by parts formula cancel, and we are left with an integral of the same form as (2.23) on which we can repeat the process. Thus we obtain a family of expressions

$$M_{jk} = \frac{(-1)^n}{2\pi}\int_0^{2\beta\pi} h_n(x)e^{i(kv(x)-jx)}dx,\ n \le r,$$

with each h_n being $(r-n)$-times differentiable and defined by the recurrence relation

$$h_0 = h,\ h_{n+1} = -i\left[\frac{h_n}{j-kv'}\right]'.$$

We find by induction that

$$h_n = i^n \sum_{l=0}^{n} \frac{k^l w_{n,l}(x)}{(j-kv'(x))^{n+l}},$$

with $w_{n,l}$ having the recurrence relation

$$\begin{aligned}
w_{n,l} &= w'_{n-1,l} + (n+l-1)v''w_{n-1,l-1}, && 0 < l < n,\\
w_{n,0} &= w'_{n-1,0}, && n > 0,\\
w_{n,n} &= 2nv''w_{n-1,n-1}, && n > 0,\\
w_{0,0} &= h.
\end{aligned}$$

By induction, we see that each $w_{n,l}$ has the form

$$w_{n,l} = \sum_{l'=0}^{n-l} \omega_{n,l,l'}(v'', \dots, v^{(n-l+2)})h^{(l')},$$

where $\omega_{n,l,l'}$ are degree l homogeneous polynomials with positive coefficients. (The $\omega_{n,0,l'}$ are constants as a result, and thus issues of existence of derivatives do not arise.)

Setting

$$W_{n,l} = \sup_{x\in[0,2\pi]} \frac{|w_{n,l}(x)|}{|h(x)|} \leq \sum_{l'=0}^{n-l} \omega_{n,l,l'}(\Upsilon_1, \ldots, \Upsilon_{n-l+1}) H_{l'},$$

we have

$$|M_{jk}| \leq \frac{1}{2\pi} \int_0^{2\beta\pi} \sum_{l=0}^{n} \frac{W_{n,l}|h(x)||k|^l}{|j - kv'(x)|^{n+l}} dx,$$

from which (2.21) follows by Hölder's inequality.

For (c), we use the $2\beta\pi$-periodicity of the integrand of (2.23) to move the contour of integration. When $j/k < p_2$, we shift the contour of integration by $-i\zeta \operatorname{sgn} k$ in the complex plane so

$$M_{jk} = \frac{1}{2\pi} \int_0^{2\beta\pi} h(x - i\zeta \operatorname{sgn} k) e^{ikv(x-i\zeta \operatorname{sgn} k) - ij(x - i\zeta \operatorname{sgn} k)} dx. \tag{2.24}$$

We now use our bounds on derivatives of h and v to bound elements of this expression, beginning with the argument of the exponential.

Applying Taylor's theorem to $\Im v(x + i\xi)$, we have

$$\Im v(x - i\zeta \operatorname{sgn} k) = -\zeta \operatorname{sgn} k v'(x) - \frac{1}{2}\zeta^2 \Im v''(\xi)$$

for some $\xi \in \Lambda_\delta^\beta$. This gives us that

$$\begin{aligned} \Re(ikv(x - i\zeta \operatorname{sgn} k)) &\leq \zeta k \operatorname{sgn} k |v'(x)| + |k| \frac{1}{2}\zeta^2 \Upsilon_{1,\delta} \\ &\leq \zeta |k| \left(\mu_1 + \frac{\zeta \Upsilon_{1,\delta}}{2} \right) \\ &\leq \zeta |k| p_1, \end{aligned}$$

where the last inequality results from the definition of ζ in the statement of the lemma.

We can bound $h(x - i\zeta \operatorname{sgn} k)$ by using that the Lipschitz constant of $\log h$ on Λ_δ^β is $\sup |h'/h| \leq H_{1,\delta}$. As a result,

$$|h(x - i\zeta \operatorname{sgn} k)| \leq |h(x)| e^{|i\zeta \operatorname{sgn} k| H_1} = |h(x)| e^{\zeta H_{1,\delta}}.$$

Thus when we take absolute values on (2.24) we obtain that

$$|M_{jk}| \leq \frac{1}{2\pi} \int_0^{2\beta\pi} |h(x)| e^{\zeta H_{1,\delta}} e^{\zeta \operatorname{sgn} k (kp_1 - j)} dx.$$

Using that $\operatorname{sgn} k = \operatorname{sgn}(j - p_1 k)$ for $j > p_1 k$ and Hölder's inequality yields (2.22).

The proof of (c) for $j/k < p_2$ is analogous, with the contour shifted in the opposite direction.

□

Given Lemma 2.11, Theorem 2.6 is an elementary result. It is necessary only to check that the conditions for the theorem imply the conditions for the lemma, and vice versa for the results.

Proof of Theorem 2.6. From (2.1), the transfer operator $\mathcal{L}$ of a map $f \in \bar{U}_P$ has action

$$\mathcal{L}\phi(x) = \sum_{n=1}^{b} \sigma v'(x + 2b\pi)\phi(v(x + 2b\pi)),$$

where $\sigma = \operatorname{sgn} v'(0)$. (Note that v is monotonic and so $\sigma v' = |v'|$).

Since $\lambda_2^{-1} < |v'| < \lambda_1^{-1}$, we can apply Lemma 2.11 with $h = \sigma v'$.

Suppose that f satisfies (DD_r). Then we can set $\Upsilon_n = C_n$ for all $n \le r$, as the definition of C_n in (DD_r) and of Υ_n in Lemma 2.11 are the same. We can also set

$$|v^{n+1}| \le \left|\frac{v^{(n+1)}}{v'}\right| |v'| \le \frac{C_n}{\min\{|\lambda_1|, |\lambda_2|\}} = H_n < \infty.$$

This gives us what we need for Lemma 2.11(b), and so there exist $W_{r,n}$ such that

$$L_{jk} \le \frac{\|v'\|_1}{2\pi} \sum_{n=0}^{r} \frac{W_{r,n}|k|^n}{|j - \lambda_m^{-1} k|^{n+r}}$$

where λ_m is λ_1 for $j/k > \lambda_1^{-1}$ and λ_2 for $j/k < \lambda_2^{-1}$.

We can eliminate the sum by using that for $j/k > p_1$,

$$\frac{|k|^r}{|j - \lambda_m^{-1} k|^r} = \frac{1}{|j/k - \lambda_1^{-1}|^r} \le \frac{1}{(p_1 - \lambda_1^{-1})^r},$$

and similarly for p_2. Furthermore, since v' does not change sign, $\|v'\|_1 = |v(2\pi\beta) - v(0)| = 2\pi$.

Thus there exists a constant C depending on the distortion constants C_r, expansion bounds $\lambda_{1,2}$ and constants p_1, p_2 such that for $j/k \notin [p_2, p_1]$ or $k = 0$,

$$L_{jk} \le \frac{C}{|j - \lambda_m^{-1} k|} \le \frac{C}{|j - p_m^{-1} k|},$$

which implies the bound for maps in (DD_r) from Theorem 2.6.

Similarly, suppose that f satisfies (DD_r). Then $\Upsilon_{1,\delta} = C_{1,\delta} < \infty$, and

$$\sup_{v \in \Lambda_\delta^\beta} |v'| \le e^{\delta C_{1,\delta}} \cdot \sup_{x \in [0, 2\beta\pi)} |v'(x)| < \infty,$$

and hence by Lemma 2.11(c) there exists $C > 0$ and $\zeta \in (0, \delta]$ such that for $j/k > p_1$, $|L_{jk}| < Ce^{-\zeta|j-p_2k|}$, and similarly for $j/k < p_2$. □

Theorem 2.7 also follows from Lemma 2.11, using the following relation between Chebyshev polynomials and Fourier modes:

$$T_k(\cos\theta) = \frac{1}{2}e^{ik\theta} + \frac{1}{2}e^{-ik\theta}.$$

However, because the cosine function on $[0, 2\pi)$ is two-to-one with critical points at 0 and π, the proof is less straightforward than for Theorem 2.6. In particular, we will have to address how to turn the transfer operator of a map in $\bar{U}_{NP}$ into the sum of operators of the form (2.20), with regard to the two-to-one nature of the transformation. We will then need to examine how distortion bounds translate quantitatively under this transformation. Once we have done these, the bounds follow easily.

Proof of Theorem 2.7. From the definition of transfer operators (2.1) and the orthogonality relation for the Chebyshev basis, we obtain the following formula for Chebyshev basis matrix elements of transfer operators of maps in $\bar{U}_{NP}$:

$$L_{jk} = \frac{t_j}{\pi} \sum_{\iota \in I} \int_{-1}^{1} \frac{\sigma_\iota}{\sqrt{1-x^2}} v_\iota'(x) T_k(v_\iota(x))) T_j(x)\ dx,$$

where $\sigma_\iota = \operatorname{sgn} v_\iota'$, $t_j = 2 - \delta_{j0}$, and the sum is taken over the branches of the map. Under the transformation $x = \cos\theta$ and using that $T_k(x) = \cos(k \cos^{-1} x)$, we find that L_{jk} is related to a Fourier basis matrix entry for a weighted transfer operator:

$$L_{jk} = \frac{t_j}{\pi} \sum_{\iota \in I} \sigma_\iota \int_0^\pi v_\iota'(\cos\theta) \cos(k \cos^{-1} v_\iota(\cos\theta))) \cos j\theta\ d\theta. \tag{2.25}$$

Based on this, we set $h_\iota = v_\iota' \circ \cos$ for each $\iota \in I$. These functions h_ι are 2π-periodic.

Defining $\nu_{\iota+} := \cos^{-1} \circ v_\iota \circ \cos$ and $\nu_{\iota-} = 2\pi - \nu_{\iota+}$, we find

$$\begin{aligned} L_{jk} &= \frac{t_j}{\pi} \sum_{\iota \in I} \sigma_\iota \int_0^\pi h_\iota(\theta) \cos(k\nu_{\iota+}(\theta)) \cos j\theta\ d\theta \\ &= \frac{t_j}{4\pi} \sum_{\iota \in I} \sigma_\iota \int_0^\pi h_\iota(\theta) \sum_{\pm} \left(e^{i(k\nu_{\iota\pm}(\theta) - j\theta)} + e^{i(k\nu_{\iota\pm}(\theta) + j\theta)}\right) d\theta. \end{aligned}$$

Continuing $\nu_{\iota\pm}$ differentiably to the interval $[0, 2\pi]$ and using that the integrands are symmetric about π, we can finally rewrite the transfer operator in the form

$$L_{jk} = \frac{t_j}{8\pi} \sum_{\iota \in I, \pm} \sigma_\iota \int_0^{2\pi} h_\iota(\theta) \left(e^{i(k\nu_{\iota\pm}(\theta) - j\theta)} + e^{i(k\nu_{\iota\pm}(\theta) + j\theta)}\right) d\theta. \tag{2.26}$$

If neither or both of $v_\iota(-1)$ and $v_\iota(1)$ are a singular point of the $\cos^{-1}$ transformation (i.e. -1 or 1), then the $\nu_{\iota\pm}$ are differentiably defined on the circle $\mathbb{R}/2\pi\mathbb{Z}$. If one of these values is, then $\nu_{\iota+}$ will continue across the critical points on either side to $\nu_{\iota-}$ and so their concatenation ν_ι is a differentiable map on $\mathbb{R}/4\pi\mathbb{Z}$. Thus, if we define the sets $I_c = \{\iota \in I : |v_\iota(\{\pm 1\}) \cap \{\pm 1\}| = 1\}$ and $I' = (I \backslash I_c \times \{+,-\}) \cup I_c$, and set $\beta_{\iota'} = 1 + \mathbf{1}_{I_c}(\iota')$, we have

$$L_{jk} = \frac{t_j}{8\pi} \sum_{\iota \in I'} \sigma_\iota \int_0^{2\pi} \sum_{b=0}^{\beta_{\iota'}-1} h_{\iota'}(\theta + 2\pi b) e^{ik\nu_{\iota'}(\theta+2\pi b)} \left(e^{-ij\theta} + e^{ij\theta}\right) d\theta. \tag{2.27}$$

Clearly, the summands are two-element sums of Fourier coefficient matrix elements of operators of the form 2.20. The following lemma, whose proof is for ease of exposition in Appendix A3, shows that that the relevant bounds on $\nu_{\iota'}$ and h_ι hold uniformly for all ι:

Lemma 2.12. *Suppose $f \in \bar{U}_{NP}$ with partition spacing constant $\boldsymbol{\xi}$ and I' is defined as above. Then*

(a) If the v_ι satisfy (DD_r) with the same distortion constants C_n, $n \le r$, then for these n there exist $\Upsilon_n, H_n < \infty$ depending only on $C_m, m \le n$ and $\boldsymbol{\xi}$ such that

$$\sup_{\theta \in [0,2\pi], \iota' \in I'} \nu_{\iota'}^{(n+1)}(\theta) \le \Upsilon_n$$

and

$$\sup_{\theta \in [0,2\pi], \iota' \in I'} \left| \frac{h_{\iota'}^{(n)}(\theta)}{h_{\iota'}(\theta)} \right| \le H_n < \infty.$$

(b) If the v_ι obey (AD_δ) with the same distortion constant $C_{1,\delta}$, then there exists $\zeta \in (0, \delta]$, $\Upsilon_{1,\zeta}, H_{1,\zeta} < \infty$ depending only on ζ, $C_{1,\delta}$ and partition spacing constant $\boldsymbol{\xi}$ such that

$$\sup_{\theta \in \Lambda_\zeta^{\beta_{\iota'}}, \iota' \in I'} \nu_{\iota'}''(\theta) \le \Upsilon_{1,\zeta}$$

and

$$\sup_{\theta \in \Lambda_\zeta^{\beta_{\iota'}}, \iota' \in I', \pm} \left| \frac{h_{\iota'}'(\theta)}{h_{\iota'}(\theta)} \right| \le H_{1,\zeta} < \infty.$$

Setting $\tilde{\mu} = [-\check{\lambda}^{-1}, \check{\lambda}^{-1}]$ and $\tilde{p} = [-p, p]$, Lemma 2.12 means we can apply Lemma 2.11 to each summand in (2.27). Up to a constant factor G to be discussed later we have that if f satisfies (DD_r) then there exist $W_{r,n}$ such that for $j > pk \ge 0$,

$$|L_{jk}| \le 2\frac{t_j}{8\pi} G \sum_{n=0}^{r} \frac{W_{r,n} k^n}{(j - \check{\lambda}^{-1} k)^{n+r}}.$$

Similarly, if f satisfies (AD_δ) there exists $\zeta' \in (0, \zeta]$ such that for $j > pk \geq 0$,

$$|L_{jk}| \leq 2\frac{t_j}{8\pi} G e^{\zeta'(H_{1,\zeta}-(j-pk))}$$

which gives the decay rates stated in Theorem 2.7 by the same means as in the proof of Theorem 2.6.

However, we need to check that the constant factor

$$G = \sum_{\iota' \in I'} \int_0^{2\beta_{\iota'}\pi} |v'_\iota(\cos\theta)| d\theta$$

is in fact finite.

We convert back to a sum over I by collapsing the sum over $\pm$ for $\iota \in I \backslash I_c$, obtaining

$$G = 2\sum_{\iota \in I} \int_0^{2\pi} |v'_\iota(\cos\theta)| d\theta.$$

We then make the two-to-one change of variable $x = \cos\theta$ to find that

$$\begin{aligned} G &= 4\sum_{\iota \in I} \int_{-1}^{1} |v'_\iota(x)| \frac{1}{\sqrt{1-x^2}} dx \\ &\leq 4\sum_{\iota \in I} \int_{-1}^{1} \frac{(1+2C_1)|\mathcal{O}_\iota|}{2\sqrt{1-x^2}} dx = 4\pi(1+2C_1) < \infty, \end{aligned}$$

where the first inequality is a result of Lemma A.3(b).

This concludes the proof of Theorem 2.7. □

Remark 2.13. *Elements of Fourier and Chebyshev transfer operator matrices are uniformly bounded, with*

$$|L_{jk}| \leq 1$$

for maps in $\bar{U}_P$, and

$$|L_{jk}| \leq (2 - \delta_{j0})(2 + 4C_1)$$

for maps in $\bar{U}_{NP}$.

This follows by applying Lemma 2.11(a) in the proofs of Theorems 2.6-2.7.

Remark 2.14. *The uniform C-expansion condition (CE) is the natural expansion condition for any choice of spectral basis on an interval. Our reasoning is as follows. If one wishes to use oscillatory integral techniques on these basis functions as in Lemma 2.11, it is best for the wavelength of the basis functions to be approximately spatially constant. However, wavelengths of sufficiently high-order spectral basis functions on intervals will always be much smaller towards the endpoints. Potential theory (Trefethen, 2013) tells us that the optimal transformation to even out high-order basis functions across the interval is always the cosine transformation.*

We now turn to proving the main theorem, Theorem 2.3, and its corollaries. Our idea is to perturb $\mathcal{L}$ such that the associated coefficient matrix is block-upper-triangular (in the Fourier case, with the ordering of basis elements $e_0, e_1, e_{-1}, e_2, \ldots$). This isolates the top block $E_N \to E_N$, which then approximates the corresponding $E_N \times E_N$ block of the full, unperturbed transfer operator, yielding convergence on domain E_N.

We summarise this using the following lemma, where we do not require our Banach space E to be BV.

Lemma 2.15. *Let E be a Banach space such that $\mathbb{E}_N := (\mathrm{id} - \mathcal{P}_N)\mathcal{L}\mathcal{P}_N$ is an endomorphism on E.*

Suppose $\mathcal{L}$ has a spectral gap on E. Then

$$\|\mathcal{L}_N - \mathcal{L}|_{E_N}\|_E = \|\mathbb{E}_N\|_E \tag{2.28}$$

and

$$\|\mathcal{S}_N - \mathcal{S}|_{E_N}\|_E \leq \frac{\|\mathcal{S}\|_E \|\mathbb{E}_N\|_E}{1 - \|\mathcal{S}\|_E \|\mathbb{E}_N\|_E}. \tag{2.29}$$

Proof. The first equality (2.28) arises simply because $\mathcal{L}_N - \mathcal{L}|_{E_N} = (\mathrm{id} - \mathcal{P}_N)\mathcal{L}_N|_{E_N} = \mathbb{E}_N$.

Let

$$\tilde{\mathcal{L}}_N := \mathcal{L} - \mathbb{E}_N = \mathcal{L}_N + \mathcal{L}(\mathrm{id} - \mathcal{P}_N).$$

Recalling that we defined $\mathscr{S}$ to be the Lebesgue integral functional and u an element of E_N with $\mathscr{S}u = 1$, let us also define $\tilde{\mathcal{S}}_N = (\mathrm{id} - \tilde{\mathcal{L}}_N + u\mathscr{S})^{-1}$. If $\|\mathbb{E}_N\|$ is small enough, this is well-defined, since $\tilde{\mathcal{S}}_N = (\mathrm{id} + \mathcal{S}\mathbb{E}_N)^{-1}\mathcal{S}$ and thus

$$\|\tilde{\mathcal{S}}_N - \mathcal{S}\| \leq \frac{\|\mathcal{S}\|\|\mathbb{E}_N\|}{1 - \|\mathcal{S}\|\|\mathbb{E}_N\|}. \tag{2.30}$$

For $\phi \in E_N$, we have that

$$\tilde{\mathcal{S}}_N^{-1}\phi = \phi - \tilde{\mathcal{L}}_N\phi + u\mathscr{S}\phi = \phi - \mathcal{L}_N\phi + u\mathscr{S}\phi \in E_N,$$

and thus $\tilde{\mathcal{S}}_N^{-1}|_{E_N}$ is an endomorphism on E_N, is equal to $\mathcal{S}_N^{-1}$. Consequently, $\mathcal{S}_N|_{E_N} = \tilde{\mathcal{S}}_N|_{E_N}$, which combined with (2.30) yields as desired (2.29). □

The following lemma is then required to connect $\|\mathbb{E}_N\|_{BV}$ to spectral matrix coefficients.

Lemma 2.16. *Suppose $\mathcal{F} : BV(\Lambda) \to BV(\Lambda)$ is an operator for Λ either $[0, 2\pi)$ or $[-1, 1]$. Let the matrix $D = (k\delta_{jk})_{j,k\in\mathbb{Z}}$ and $\check{D} = (k\delta_{(j-1)k})_{j,k\in\mathbb{N}}$.*

If $\mathcal{F}$ has Fourier coefficient matrix F, then

$$\|\mathcal{F}\|_{BV} \leq 2\pi(\|DF\|_{\ell^2} + \|F\|_{\ell^2}). \tag{2.31}$$

Similarly, if $\mathcal{F}$ has Chebyshev coefficient matrix F, then

$$\|\mathcal{F}\|_{BV} \le 2\pi(\left\|\check{C}\check{D}F\check{C}^{-1}\right\|_{\ell^2} + \|\check{C}F\check{C}^{-1}\|_{\ell^2}), \tag{2.32}$$

where $\check{C} = (t_k^{-1/2}\delta_{jk})_{j,k\in\mathbb{N}}$.

Proof. Consider first the Fourier case. Then since $\frac{1}{\sqrt{2\pi}}\|\cdot\|_{L_2} \le \|\cdot\|_{BV} \le \sqrt{2\pi}\|\cdot\|_{H^1}$,

$$\|\mathcal{F}\|_{BV} \le 2\pi\|\mathcal{F}\|_{L_2\to H^1} = 2\pi\left(\|\mathcal{D}\mathcal{F}\|_{L_2} + \|\mathcal{F}\|_{L_2}\right). \tag{2.33}$$

By the Plancherel equality, $\|\mathcal{D}\mathcal{F}\|_{L_2} = \|DF\|_{\ell^2}$ and $\|\mathcal{F}\|_{L_2} = \|F\|_{\ell^2}$. This gives the required bound in (2.31).

Consider instead the Chebyshev case. Define the Jacobi weight function $\mathfrak{j}(x) = \sqrt{1-x^2}$, and the Sobolev spaces $\check{H}^k \subset L_2([-1,1],1/\mathfrak{j})$, $k \ge 0$ with norm

$$\|\phi\|_{\check{H}^k} = \sum_{n=0}^{k}\int_{-1}^{1}\mathfrak{j}^{2n-1}|\phi^{(n)}|^2dx. \tag{2.34}$$

Note that $\check{H}^0 = L^2([-1,1],1/\mathfrak{j})$.

If G is the set of even functions on $\mathbb{R}/2\pi\mathbb{Z}$, simple trigonometric manipulations show that the operator $\mathcal{C}: \phi \mapsto \frac{1}{2}\phi\circ\cos$ is an isometry from $\check{H}^k$ to $G\cap H^k([0,2\pi))$ and similarly from $BV([-1,1])$ to $G\cap BV([0,2\pi))$. Thus,

$$\begin{aligned}\|\mathcal{F}\|_{BV([-1,1])} &= \|\mathcal{C}\mathcal{F}\mathcal{C}^{-1}\|_{G\cap BV([0,2\pi))}\\ &\le 2\pi\left(\|\mathcal{D}\mathcal{C}\mathcal{F}\mathcal{C}^{-1}\|_{G\cap L_2} + \|\mathcal{C}\mathcal{F}\mathcal{C}^{-1}\|_{G\cap L_2}\right),\end{aligned}$$

where the inequality comes from (2.33). We can then convert back to $\check{H}^0$ to get the inequality

$$\|\mathcal{F}\|_{BV([-1,1])} \le 2\pi\left(\|\mathcal{C}^{-1}\mathcal{D}\mathcal{C}\mathcal{F}\|_{\check{H}^0} + \|\mathcal{F}\|_{\check{H}^0}\right).$$

We can convert these operator norms into matrix norms as follows. The Chebyshev polynomial basis is an orthogonal basis for $\check{H}^0$ with $\|T_k\|_{\check{H}^0} = \sqrt{\pi/t_k}$ and furthermore the functions

$$\mathcal{C}^{-1}\mathcal{D}\mathcal{C}T_k = k\sin(k\cos^{-1}x)$$

are orthogonal in $\check{H}^0$ with norms $k\sqrt{\pi/t_k}$ respectively. The resulting Plancherel equality results in (2.32). □

We now have the requisite results to tie together to prove Theorem 2.3.

Proof of Theorem 2.3. Maps in U_P have a spectral gap in BV as they are uniformly expanding with bounded distortion. Since maps in U_{NP} have a forward iterate that is uniformly expanding with bounded distortion by Theorem A.1 in Appendix A1, they also have a spectral gap in BV.

Suppose E_N is the Fourier coefficient matrix of $\mathbb{E}_N$ and the expansion coefficient of the associated map f is $\lambda > 1$. Then given $p \in (\lambda, 1)$, there exists an appropriate spectral decay function K such that when $|j| \geq |k|$,

$$|L_{jk}| \leq K(|j| - p|k|).$$

Now suppose f satisfies (DD_r) for some $r \geq 2$. Then $K(M) = CM^{-r}$ for some $C > 0$, and so

$$\begin{aligned}
\|E_N\|_{\ell^2}^2 &\leq \sum_{k=-N}^{N} \sum_{j=N+1}^{\infty} \left(|L_{jk}|^2 + |L_{-jk}|^2\right) \\
&\leq \sum_{k=-N}^{N} \sum_{j=N+1}^{\infty} 2C^2(j - p|k|)^{-2r} \\
&\leq \frac{2C^2}{2r-1} \sum_{k=-N}^{N} (N - p|k|)^{1-2r},
\end{aligned}$$

by converting to an integral. We can then take the supremum of the summands to obtain

$$\sum_{k=-N}^{N} (N - p|k|)^{1-2r} \leq (2N+1)(N - pN)^{1-2r} \leq \frac{3}{(1-p)^{2r-1}} N^{2-2r}$$

and thus

$$\|E_N\|_{\ell^2}^2 \leq \frac{6}{(2r-1)(1-p)^{2r-1}} N^2 K(N)^2.$$

Similarly,

$$\|DE_N\|_{\ell^2}^2 \leq \frac{6}{(2r-2)(1-p)^{2r-2}} N^3 K(N)^2,$$

where D is as in Lemma 2.16.

Hence as a result of Lemma 2.16, there exists a function $K' \in \kappa(\mathrm{DD}_r)$ such that $\|\mathbb{E}_N\| \leq N\sqrt{N}K'(N)$.

Suppose f instead satisfies (AD_δ). Then for some $\zeta \in (0, \delta]$ there exists $p > 1$ such that for all $|j| \geq |k|$, $L_{jk} \leq Ce^{-\zeta(|j|-p|k|)}$. Consequently,

$$\|E_N\|_{\ell^2}^2 \leq \sum_{k=-N}^{N} \sum_{j=N+1}^{\infty} 2e^{-2\zeta(j-p|k|)} \leq \frac{4N}{\zeta^2} e^{-2\zeta(1-p)N}$$

with a comparable result for DE_N. Thus, there exists a function $K' \in \kappa(\mathrm{AD}_\delta)$ such that

$$\|\mathbb{E}_N\| \leq NK'(N) \leq N\sqrt{N}K'(N).$$

Similarly, we get the same results up to constants in the Chebyshev case: the C matrices are unproblematic as $\|C\|_{\ell^2} = 1$ and $\|C^{-1}\|_{\ell^2} = 2$.

We therefore have by Lemma 2.15 that if f satisfies some distortion condition (D) then there exists $K' \in \kappa(\mathrm{D})$ such that for any N and ϕ in $\mathcal{L}_N$,

$$\|\mathcal{L}_N\phi - \mathcal{L}\phi\|_{BV} \leq N\sqrt{N}K'(N)\|\phi\|_{BV}$$

and if N is sufficiently large,

$$\|\mathcal{S}_N\phi - \mathcal{S}\phi\|_{BV} \leq \frac{\|\mathcal{S}\|_{BV}N\sqrt{N}K'(N)}{1 - \|\mathcal{S}\|_{BV}N\sqrt{N}K'(N)} \leq 2N\sqrt{N}\|\mathcal{S}\|_{BV}K'(N)\|\phi\|_{BV},$$

which is what was required for Theorem 2.3. □

Corollary 2.4 is a direct result of this convergence and Theorem 2.1:

Proof of Corollary 2.4. We know from Theorem 2.1 that $\rho = \mathcal{S}u$. We have also defined $\rho_N = \mathcal{S}_N u$, recalling that u lies in E_N. As a result, by Theorem 2.3,

$$\|\rho_N - \rho\|_{BV} = \|\mathcal{S}_N u - \mathcal{S}u\|_{BV} \leq N\sqrt{N}\bar{K}(N)\|u\|_{BV},$$

as required. □

Note that here, unlike in Theorem 2.3, we actually have that estimates converge in norm to the true values.

Corollary 2.5 also follows directly from Theorems 2.1 and 2.3.

Proof of Corollary 2.5. We know from Theorem 2.1 that on V, the space of zero integral functions, the solution operator $\mathcal{S}$ is identical to $\sum_{n=0}^{\infty}\mathcal{L}^n$. We then need only apply the second part of Theorem 2.3 to get the required inequality. □

Remark 2.17. *When a map satisfies* (AD_δ)*, one might be interested in the best rate of decay one can get for* $\|\mathbb{E}_N\|_{BV}$*, which controls the convergence of estimates. In the non-periodic case one can show that*

$$\lim_{N\to\infty}\frac{1}{N}\log\|\mathbb{E}_N\|_{BV} = \sup_{\iota\in I, x\in[0,2\pi]}|\Im\nu_\iota(x+i\zeta)| - \zeta. \tag{2.35}$$

The value of z *where the supremum in (2.35) is maximised will have* $\Im\nu_\iota'(z) = 0$*; if this value of* z *varies continuously with* ζ*, then it will have a maximum when* $|\Re\nu_\iota'(z)|$ *is* 1 *or* -1*. Thus, one expects the right-hand side of (2.35) to be maximised for*

$$\zeta = \min\left\{\inf\left\{|\Im z| \mid z \in (\nu_\iota')^{-1}(\{\pm 1\}), \iota \in I\right\}, \delta\right\}. \tag{2.36}$$

The result is the same in the periodic case but with v_ι *substituted for* ν_ι*.)*

2.4. Algorithms

Our results suggest a variety of possible algorithms to capture, given a map f, statistical properties that can be expressed as $\mathcal{S}\phi$ for some ϕ, such as acims (2.10) and diffusion coefficients (2.12). We present two possible algorithms a practitioner might wish to use to calculate invariant measures: one that gives rigorous bounds on statistical properties but is somewhat cumbersome for exploratory use, and one that gives accurate but non-validated estimates that is much more convenient to use. In this section we describe the two algorithms, and then explain how in both algorithms we calculate elements of the transfer operator matrix. We will demonstrate the algorithms in Section 2.5.

Algorithm 1 is a traditional fixed-order spectral method, implemented in interval arithmetic. It requires as input the map inverses v_ι and their derivatives, a spectral order N, various bounds on elements of the transfer operator matrix in the relevant spactral basis, and a bound on the norm of the solution operator[3] $\mathcal{S}$; it then outputs an estimate for the acim ρ with a rigorously validated BV error.

By contrast, Algorithm 2 is an adaptive-order spectral method that is not rigorously validated: it uses an adaptive QR factorisation of the solution operator inverse $\mathcal{K}$ to solve the linear problem and test for convergence (Olver and Townsend, 2013; Hansen, 2010). It requires as input only an algorithm to calculate the map f and outputs an estimate for ρ whose error is not rigorously bounded but is of the order of $\|\mathcal{S}\|_{BV}\epsilon^{1-\theta}$, where ϵ is the floating-point precision and θ is a small number depending on the order of differentiability of f.

Algorithm 2 is extremely well-suited for numerical exploration. Because the only required input is the map itself, Algorithm 2 requires a minimum of drudge work on the part of the user. It is typically also extremely fast: just with a personal computer, Algorithm 2 gives estimates of statistical quantities of a simple analytic map accurate to 14 decimal places in less than one-tenth of a second (see Section 2.5). Because our spectral methods are very accurate in an easily verifiable way, an adaptive, non-validated method is also highly reliable. We have consequently made an implementation of Algorithm 2 available in the open-source Julia package Poltergeist (Wormell, 2017).

In the presentation of the algorithms and the following discussion we assume that the Fourier and Chebyshev spectral bases have been relabeled as $(b_k)_{k\in\mathbb{N}^+}$. We also

[3]Available theoretical bounds typically scale exponentially with the distortion bound C_1 (see Korepanov et al. (2016) and Appendix A4). However, at least in the analytic case, the spectrally fast convergence dominates the large theoretical bounds. It is only necessary to control floating-point error using an appropriately high numerical precision. Alternatively and more generally, one may apply the approach of Galatolo and Nisoli (2014).

Input: Map inverses and derivatives $v_\iota, v'_\iota,\ \iota \in I$; spectral order N; aliasing bounds $A^{(N)}_{jk}$ for $j,k = 1,\dots,N$; bound $b^{\mathcal{S}} \geq \|\mathcal{S}\|_{BV}$; bound $b^{\mathbb{E}_N} \geq \|\mathbb{E}_N\|_{BV}$ (see Lemma 2.15); bounds $b^L_{jk} \leq |L_{jk}|$.

Output: High-precision floating-point vector $\tilde{\rho}$ containing spectral coefficients of acim estimate; rigorous BV error bound $\bar{\epsilon}_{\text{obs}}$

Check that $b^{\mathbb{E}_N} b^{\mathcal{S}} < 1$: if this is not the case increase N;
Set the number of floating-point bits to be greater than $-\log_2(N^4 * b^{\mathbb{E}_N})$;
Initialise $N \times N$ matrix of intervals $L^{(N)}$;
for $k \leftarrow 1$ **to** N **do**
    Calculate interpolant values $q^{(k,N)} = \{\mathcal{L}(b_k)(x_{l,N})\}_{l=1}^N$ in interval arithmetic, using (2.1);
    Calculate spectral coefficients of the interpolant $p^{(k,N)} = FFT(q^{(k,N)})$ ($DCT(q^{(k,N)})$ in the Chebyshev case);
    for j $\leftarrow 1$ **to** N **do**
        Calculate spectral coefficient matrix entry $L^{(N)}_{jk}$ as $q^{(k,N)}_j$ plus aliasing error $[-A^{(N)}_{jk}, A^{(N)}_{jk}]$;
        Refine interval estimate $L^{(N)}_{jk}$ by intersecting it with $[-b^L_{jk}, b^L_{jk}]$;
    end
end
Calculate $u^{(N)} = \{[\delta_{j0}/|\Lambda|, \delta_{j0}/|\Lambda|]\}_{j=1}^N$;
Calculate row vector of intervals $\mathscr{S}^{(N)} = (\mathscr{S} b_j)_{j=1}^N$ using standard formulae (Trefethen, 2013);
Calculate the spectral coefficient matrix of $\mathcal{S}_N^{-1}$, $K^{(N)} = I - L^{(N)} + \mathscr{S}^{(N)} u^{(N)}$, where I is an $N \times N$ identity matrix;
Calculate $\rho^{(N)} = K^{(N)} \backslash u^{(N)}$;
Calculate $\tilde{\rho} = \{\text{midpoint}(\rho_j^{(N)})\}_{j=1}^N$;
Calculate a bound $\bar{\epsilon}_{\text{interval}} > \|\rho^{(N)} - \tilde{\rho}\|_{BV}$;
Calculate $\bar{\epsilon}_{\text{finite}} = 1/(1/(b^{\mathbb{E}_N} b^{\mathcal{S}}) - 1)$;
Calculate $\bar{\epsilon}_{\text{obs}} = \bar{\epsilon}_{\text{interval}} + \bar{\epsilon}_{\text{finite}}$;

Algorithm 1: Rigorous algorithm to compute acims.

implicitly assume that the Fourier exponential basis has been transformed to sines and cosines so that real functions have real spectral coefficients.

In both algorithms, one calculates L_N by columns, using that the kth column of L_N consists of the first N spectral coefficients of $\mathcal{L}b_k$. The most effective way to estimate these coefficients is by calculating an interpolant. The idea of this is as follows.

Input: Map f; map derivative f' (optional; may be calculated automatically using dual number routines (Revels et al., 2016)); tolerance ϵ

Output: Adaptive order k_{opt}; floating-point vector $\tilde{\boldsymbol{\rho}}$ containing spectral coefficients of acim estimate $\tilde{\rho}_{N_{\text{opt}}}$

`# Extendable vectors encode an infinite vector with finitely many non-zero entries, ragged matrices' columns are extendable vectors. These will encode infinite-dimensional objects approximating` u, $\mathcal{K}$.

Initialise empty ragged matrix H, which will hold Householder vectors for row-reduction;

Initialise empty ragged matrix $\hat{K}$, which will hold row-reduced coefficients of solution operator inverse $\mathcal{K}$;

Calculate extendable vector $\hat{u} = (\delta_{j1}/|\Lambda|)_{j\geq 1}$ containing coefficients of u that will be progressively row-reduced;

Set k, the number of columns of matrix $\hat{K}$, to be 0;

repeat `# Loop between calculating columns of` $\hat{K}$ `and row-reducing`

  Increment k by 1;

  Set the interpolation order $M = 4$;

  repeat `# Calculating optimum order interpolant of` $\mathcal{L}b_k$

    Set $M \leftarrow 2M$;

    Calculate values of the interpolant $q^{(k)} = \{\mathcal{L}(b_k)(x_{l,M})\}_{l=1}^{M}$ using (2.1) with Newton iteration for the transfer operator;

    Calculate spectral coefficients of the interpolant $p^{(k)} = FFT(q^{(k)})$ ($DCT(q^{(k)})$ in the Chebyshev case);

  until *the interpolant has converged according to the reasoning in Aurentz and Trefethen (2017)*;

  Calculate $\kappa^{(k)}$, which will become the kth column of $\hat{K}$, as an extendable vector $\{\delta_{jk} + \mathscr{S}_k\delta_{j1}\}_{j\geq 1} - p^{(k)}$, where $\mathscr{S}_k = (\mathscr{S}b_k)$ is calculated from Chebyshev and Fourier integral formulae (Trefethen, 2013);

  Apply previous Householder transformations encoded as column vectors of H to $\kappa^{(k)}$;

  Calculate Householder vector h that will row-reduce $\kappa^{(k)}$ considered as the kth column of $\hat{K}$;

  Apply h to $\kappa^{(k)}$;

  Right-concatenate $\kappa^{(k)}$ onto $\hat{K}$;

  `# Note` $\hat{K}$ `is row-reduced and so upper-triangular`

  Apply h to $\hat{u}$;

  Right-concatenate h onto H;

until $\max\{u_j\}_{j\geq k+1} \leq \epsilon/|\Lambda|^{-1}$ `# i.e. negligible benefit from larger k`;

Set $N_{\text{opt}} = k$;

Calculate $\tilde{\boldsymbol{\rho}} = \{\hat{K}_{jk}\}_{j,k=1}^{N_{\text{opt}}} \backslash \{\hat{u}_j\}_{j=1}^{N_{\text{opt}}}$ via backsolving;

Algorithm 2: Algorithm to compute acims using adaptive interpolation and infinite-dimensional adaptive QR solver.

Using (2.1) one evaluates the function $\mathcal{L}b_k$ at N special interpolation nodes $x_{l,N}$: in the Fourier case these interpolation nodes are evenly-spaced on the periodic interval, in the Chebyshev case, the interpolation nodes are given by Chebyshev nodes of the first kind (Trefethen, 2013; Boyd, 2001). One then applies the Fast Fourier Transform (resp. Discrete Cosine Transform) to the vector $((\mathcal{L}b_k)(x_{l,N}))_{l=1,\dots,N}$. The resulting length-$N$ vector contains the spectral coefficients of the unique function $p^{(k,N)} \in E_N$ which matches $\mathcal{L}b_k$ at the interpolation nodes. The so-called interpolant $p^{(k,N)}$ is a close approximation of $\mathcal{L}b_k$: the difference between the jth spectral coefficient of $p^{(k,N)}$ and that of $\mathcal{L}b_k$ (the so-called aliasing error) is guaranteed to be smaller than some bound $A_{jk;N}$. This bound can be determined from aliasing formulae standard in approximation theory (Trefethen, 2013) combined with bounds on higher-order spectral elements of $\mathcal{L}b_k$ (e.g. from Theorems 2.6-2.7).

These algorithms generalise very easily to other transfer operator problems of the form $\psi = \mathcal{S}\phi$. This can be done by formulating the problem as $\mathcal{K}\psi = \phi$ and thus substituting ρ and u (when it is not constituting the solution operator) for ψ and ϕ respectively in the algorithms. As an example, the formula for diffusion coefficients $\sigma_f^2(A)$ (2.12) requires computing $\psi = \mathcal{S}[(\mathrm{id} - \rho\mathcal{K})(\rho A)]$ for a given observable A. The input $(\mathrm{id} - \rho\mathcal{K})(\rho A)$ may be approximated by projection onto Chebyshev nodes, and estimating the final expression reduces to simply numerically solving the equation

$$\mathcal{K}\psi = (\mathrm{id} - \rho\mathcal{K})(\rho A)$$

using the same approximation methods as above.

2.5. Numerical results

In Section 2.5.1 we will prove some rigorous bounds on basic statistical properties of the Lanford map using the rigorous Algorithm 1; we will then demonstrate the adaptive Algorithm 2 using the Lanford map and a non-smooth circle map, assessing the adaptive algorithm's accuracy and the spectral method's rate of convergence.

2.5.1. Rigorous bounds on statistical quantities: the Lanford map. The Lanford map, $f : [0,1] \to [0,1]$

$$f(x) = 2x + \frac{1}{2}x(1-x) \mod 1$$

is a common test case for rigorous estimation of statistical quantities of maps (Galatolo and Nisoli, 2014; Jenkinson et al., 2018; Bahsoun et al., 2016). By linearly rescaling of $[0,1]$ onto $[-1,1]$ we can apply our spectral method to it.

The Lanford map's uniform expansion parameter is $\lambda = \frac{3}{2}$ and its distortion bound (on $[0,1]$) is $C_1 = \frac{4}{9}$. Applying (A.12), we find that $\|\mathcal{S}\|_{BV} \leq 9235$.

By considering explicit bounds that will be given in Lemma 2.11, we choose $\zeta = \cosh^{-1}\frac{7}{4}$, as it is close to the optimal value for ζ given in Remark 2.17. Using

Mathematica (Wolfram Research, Inc., 2013) we show that as a result of Remark 2.17,

$$|L_{jk}| \leq t_j \sqrt{7 + \frac{\sqrt{33}}{2}} e^{\cosh^{-1}(4-\sqrt{6})k - \cosh^{-1}\frac{7}{4}j}. \tag{2.37}$$

To calculate an estimate of the acim of this map, we implement Algorithm 1 with $N = 2048$. We find the truncation error to be $\|\mathbb{E}_N\|_{BV} \leq 6.75 \times 10^{-133}$, and choose the floating-point precision to be 512 significand bits.

Consequently, we obtain an acim estimate $\tilde{\rho}$ with the rigorously validated error bound

$$\|\tilde{\rho} - \rho\|_{BV} \leq 6.3 \times 10^{-129}.$$

This estimate is plotted in Figure 2.2. The Chebyshev coefficients of $\tilde{\rho}$ are available in `Lanford-acim.zip`.

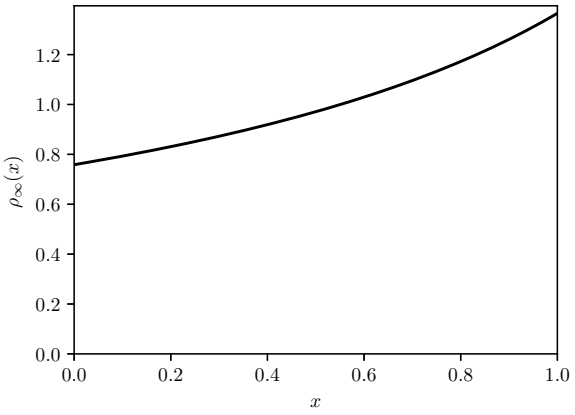

Figure 2.2. The density of the absolutely continuous invariant measure for the Lanford map, obtained by Algorithm 1.

We then use this estimate of the acim ρ to calculate the Lyapunov exponent of the Lanford map

$$\lambda = \int_\Lambda \log|f'(x)|\rho(x)dx.$$

using Clenshaw-Curtis quadrature (Trefethen, 2013). on $\tilde{\rho}\log|f'| = \tilde{\rho}\log(2-3x)$ to obtain the rigorous estimate given in Theorem 2.9(a).

We then calculated the diffusion coefficient of the observable $\phi(x) = x^2$ by evaluating the natural finite-order approximation of formula (2.12), using Clenshaw-Curtis quadrature. obtaining the rigorous bound given in Theorem 2.9(b).

The results together were obtained in 9 hours over 15 hyper-threaded cores of a research server running 2 E5-2667v3 CPUs with 128GB of memory. The most time-consuming operation was inverting the Galerkin projection of the solution operator inverse matrix $\mathcal{K}_{2048} = \mathcal{P}_{2048}\mathcal{K}|_{E_{2048}}$ (see (2.14)): this process took up 94% of the runtime, which may stem partly from using an unoptimised routine. Once $\mathcal{K}_{2048}^{-1}$, i.e. the solution operator matrix, was supplied, all the statistical quantities were calculated on a personal computer in seconds.

2.5.2. Adaptive algorithms. We now present results from the adaptive Algorithm 2, and illustrate the algorithm's convergence by comparison with a fixed-order version of Algorithm 2.

We have implemented Algorithm 2 in Julia, an open-source dynamic scientific computing language. This implementation is publically available in the package Poltergeist (Wormell, 2017). Poltergeist is integrated with ApproxFun, a comprehensive function approximation package written in Julia (Olver, 2019); thus, standard manipulations of functions and operators may readily be applied to invariant measures, transfer operators and so on.

Using Poltergeist, we present empirical convergence results for the Lanford map (for comparison with rigorous methods), and a circle map which is C^4 but not analytic.

2.5.2.1. *The Lanford map.* The Lanford map experiment in Section 2.5.1 can be repeated in Poltergeist in a few lines of Julia code:

```
using Poltergeist, ApproxFun
f_lift(x) = 5x/2 - x^2/2; d = 0..1
f = modulomap(f_lift,d);
K = SolutionInv(f);
rho = acim(K);
L_exp = lyapunov(f,rho)
sigmasq_A = birkhoffvar(K,Fun(x->x^2,d))
```

This code instantiates a `MarkovMap` object `f` and creates a `QROperator` object `K`, which stands in for the corresponding solution operator inverse $\mathcal{K}$ (recalling the definition of the solution operator inverse (2.8)). The `acim` function carries out Algorithm 2 by calling ApproxFun's adaptive QR solver (Hansen, 2010) on the equation $\mathcal{K}\rho = u$. The output is an ApproxFun `Fun` object containing $\tilde{\boldsymbol{\rho}}_N$, the Chebyshev coefficients of the adaptive acim estimate. The Lyapunov exponent and diffusion coefficient are calculated using special commands defined in the package that call appropriate ApproxFun integration and QR solving routines, in the latter case via (2.12). Once the relevant functions have compiled using Julia's just-in-time compiler, the last five lines of the code will run in less than 0.12 seconds on a personal computer.

By applying Algorithm 2 with fixed orders N, the exponential convergence of ρ_N with N predicted in Theorem 2.3 was seen to hold in practice. Indeed, only $N_{\text{opt}} = 24$ columns of the transfer operator were required for convergence using Algorithm 2 (see Figure 2.3).

The estimate of Algorithm 2 for ρ is in fact remarkably accurate: the ℓ^∞ error on Chebyshev coefficients is less than 8×10^{-15} (40 times the floating point precision) and the BV error on the acim estimate is 3×10^{-13} (around 1300 times the floating point precision). The Lyapunov exponent estimate was correct almost to within the floating point precision, with the error compared to the rigorous estimate being 2.2×10^{-16}: this level of accuracy appears fortuitous rather than representative. More realistically, the estimate for $\sigma_f^2(A)$ was accurate to about 25 times floating point precision (1.4×10^{-15}).

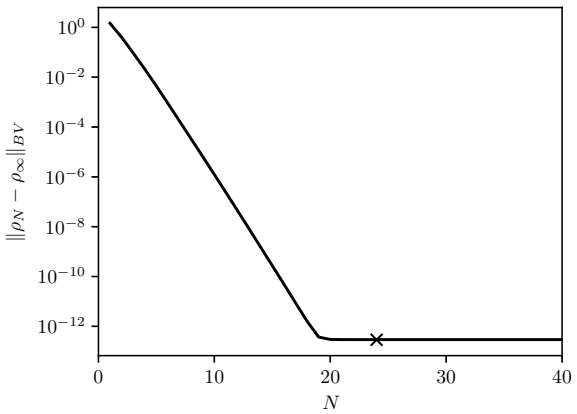

FIGURE 2.3. Exponential convergence with N of floating-point estimates of ρ_N for the Lanford map. The error of the adaptive estimate for $\rho_{N_{\text{opt}}}$ from Algorithm 2 is shown as a cross for comparison.

2.5.2.2. *A non-analytic circle map.* We now consider a circle map which does not satisfy an analytic distortion condition (AD_δ) but rather a differentiable distortion condition (DD_r).

Define the uniformly expanding, triple-covering circle map $g : [0, 2\pi) \to [0, 2\pi)$ via the inverse of its lift:

$$v_g(x) = \frac{x}{3} + \sum_{m=0}^{\infty} 2^{-\frac{33}{8}m} \cos\left(2^m \left(1 - \cos\frac{x}{3}\right)\right).$$

The map g is $C^{4.125-\epsilon}$ and thus satisfies distortion condition (DD_3) but not (DD_4).

We implement the acim-finding process in a similar fashion to the Lanford map, although to optimise for speed we also supply `CircleMap` with the derivative for the lift:

```
g = CircleMap(v_g,0..2pi,diff=v_g_dash,dir=Reverse)
Lg = Transfer(g);
rho_g = acim(Lg);
```

This routine took approximately 9 minutes to run on a personal computer and required the evaluation of $N_{\text{opt}} = 2747$ columns of the transfer operator. It produced an acim estimate (plotted in Figure 2.4) whose BV error we estimate to be approximately 4.8×10^{-10}, by comparison with an estimate obtained using high-precision floating-point arithmetic and $N = 6144$ columns.

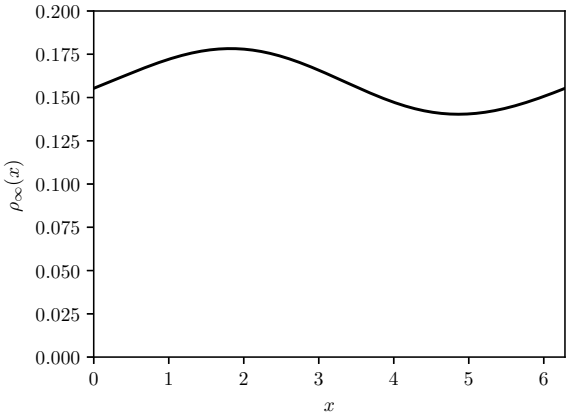

FIGURE 2.4. Invariant measure estimate for g using Algorithm 2.

The convergence of ρ_N is illustrated in Figure 2.5. The BV error on ρ_N is estimated to be $O(N^{\epsilon-2.125})$, which is better than the estimate in Theorem 2.3 of $O(N^{-1.5})$. We conjecture that acim estimates of $C^{r+\alpha}$ circle maps (i.e. those satisfying "$(\mathrm{DD}_{r-1+\alpha})$") converge in BV as $O(N^{2-r-\alpha+\varepsilon})$ for all $\varepsilon > 0$, supposing that the transfer operator matrix coefficients have decay rates that interpolate those in Theorems 2.6-2.7.

2.6. Discussion

In this chapter we have demonstrated that adaptive spectral methods allow for very fast and user-friendly computation of statistical properties, via an implementation of transfer operator spectral methods in the Julia package Poltergeist. Using Poltergeist, the quantities in Theorem 2.9 may be estimated in under 0.1 seconds on

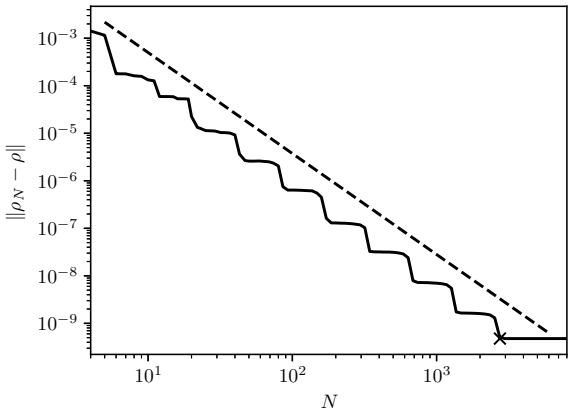

FIGURE 2.5. The convergence with N of floating-point estimates of ρ_N for g. The error of the adaptive estimate for $\rho_{N_{\text{opt}}}$ using Algorithm 2 is plotted with a cross. The slope of a function $K(N) = CN^{-2.125}$ is plotted with a dashed line. Error estimates are by comparison with an $N = 6144$ high-precision floating point acim estimate.

a personal computer to 13 decimal places of accuracy; the package also allows for computation of a great many statistical quantities not included here, such as transfer operator spectra, correlation functions and linear response.

To give a rigorous illustration of the power of spectral Galerkin methods for transfer operators, we have also applied a rigorously justified spectral method to the Lanford map to get validated bounds for the Lyapunov exponent and diffusion coefficient to 123 decimal places (see Theorem 2.9 in Section 2.2.2).

The rates of convergence we obtain compare very favourably with other approaches. While set-based approaches (Ulam's method) cover a much larger class of maps than we consider, they have an optimal convergence rate of only $O(\log(N)/N)$ irrespective of regularity, where N is the size of the Ulam matrix. Spectral methods are also significantly more efficient than algorithms that use periodic orbits to calculate statistical quantities: in the case of analytic maps, these algorithms converge superexponentially in the order (i.e. maximum length of periodic orbits), but the number of periodic orbits that must be computed, and hence the computational cost, grows exponentially with the order (Jenkinson et al., 2018; Jenkinson and Pollicott, 2005). In terms of computational power P, which is the relevant quantity in practice, periodic orbit algorithms have error $O(e^{-k(\log P)^2})$ as opposed to the spectral method's convergence rate of $O(e^{-kP^{1/3}})$.

Consequently, the illustrative numerical bounds we obtain in Theorem 2.9 are well beyond the practicable capabilities of other numerical results. The previous best rigorous bound on the Lyapunov exponent found in the literature was $\lambda = 0.6575 \pm 0.0015$ obtained by Ulam's method in Galatolo and Nisoli (2014). The diffusion coefficient of $A(x) = x^2$ was calculated rigorously in Bahsoun et al. (2016) to less than one significant figure of accuracy using Ulam's method; in Jenkinson et al. (2018) an estimate was given correct to 55 significant figures, of which 17 were rigorously validated. For comparison, we use a comparable amount of computational power to obtain 123 validated significant figures for both the Lyapunov exponent and diffusion coefficient (see Section 2.5.1 for further details).

Similarly, our adaptive algorithms provide a much higher degree of accuracy for practical uses than previous algorithms. This is of great advantage in broader study of chaotic dynamics, as having rigorous (or at least very reliable) algorithms at one's disposal allows one to easily explore mathematical phenomena for which analytical results may not yet exist. An example of such an endeavour is provided in Gottwald et al. (2016), where the authors make use of a Fourier spectral method to explore a particular rate of convergence in linear response theory, directing a subsequent proof; another is in Chapter 4, wherein we use a Chebyshev spectral method to study the macroscopic dynamics of a mean-field coupled chaotic system in the thermodynamic limit and show that it violates various hypotheses about the behaviour of complex chaotic systems. We hope that our spectral methods become a useful tool in theoretical and numerical study of chaotic systems.

There are several further directions for research. Numerical results indicate that the actual rates of convergence are slightly better than what we have proven (see Section 2.5.2), and a different theoretical approach may yield the optimal convergence rates. While we have used function spaces with bounded variation norms, most of our results are largely agnostic to the function space used: that said, the exposition would likely be substantially simplified by using weighted Sobolev Hilbert spaces suited to the orthogonal polynomial bases. Furthermore, with a proof of convergence of the spectral method on a scale of these spaces one could obtain estimates of linear response to dynamical perturbations. Numerical experiments demonstrate that eigenvalues and eigenfunctions of L_N using the spectral method converge in norm to those of L, as proved in the periodic case by Baladi and Holschneider (1999). The observed rates of convergence are the standard spectral rates. Interestingly, numerical experiments have also shown that convergence of estimates of the Fredholm determinant of $(\mathrm{id} - z\mathcal{L})^{-1}$ using adaptive QR factorisations give accurate estimates of transfer operator eigenvalues: this may imply the possibility of computing dynamical zeta functions using spectral methods (notwithstanding that the most important property of the zeta functions are their zeros, which are given by the transfer operator's eigenvalues).

Our results may be extended to higher dimensions, possibly including maps with contracting directions. A significant problem with extending to higher dimensions, however, is that the number of basis functions necessary to compute estimates of a given accuracy will increase exponentially with the dimension: this may be remediable to a limited extent by using bases of smooth, compactly-supported wavelets (Holschneider, 1996; Baladi and Holschneider, 1999), which could lower the complexity as a result of their sparse structure.

Finally, by constructing efficient numerical inducing schemes, it seems likely that our methods can provide fast, accurate and more or less rigorous estimates of statistical properties for almost all major classes of one-dimensional chaotic maps, such as non-Markov expanding maps, quadratic maps, and, as we will discuss in the next chapter, intermittent maps.

Chapter 3

Accurate and efficient numerical methods for intermittent dynamics

3.1. Introduction

Intermittent dynamics, wherein long periods of regular dynamics alternate with bursts of chaotic dynamics, is a feature of many physical systems around a bifurcation between chaotic and regular dynamics, such as in turbulence (Pomeau and Manneville, 1980). The ergodic and statistical behaviour of intermittent dynamics is commonly studied using a prototypical class of so-called Pomeau-Manneville-type maps, which we denote by PM (an example of which is plotted in Figure 3.1). These maps are defined on an interval, which we take to be $[0,1]$, and their phase spaces can be divided into a "good" set $[a,1]$, on which the map is uniformly expanding, and a "bad" set close to an unstable but linearly neutral fixed point at 0:

$$f(x) = \begin{cases} f_b(x) := xh(x^\alpha)), & x \in [0,a) \\ f_g(x), & x \in [a,1], \end{cases} \tag{3.1}$$

where $\alpha > 0$, $f_b' \geq 1$, $h(0) = 1$, $h'(0) > 0$ and $f_g : [a,1] \to [0,1]$ is a full-branch expanding Markov map in class U_{NP}, as defined in Section 2.2.1. For simplicity we will assume f_g satisfies analytic distortion condition (AD_δ), and h extends analytically into the complex plane, but our proofs may be appropriately modified to cover the differentiable case.

A standard example of such a map is the Liverani-Saussol-Vaienti map (Liverani et al., 1999) on the interval $[0,1]$ with $a = \frac{1}{2}$, given by

$$f(x) = \begin{cases} x(1+(2x)^\alpha), & x \in [0, \frac{1}{2}) \\ 2x - 1, & x \in [\frac{1}{2}, 1]. \end{cases} \tag{3.2}$$

This map and its typical dynamics for $\alpha = 0.8$ are shown in Figure 3.1.

Maps of class PM are endowed with absolutely continuous invariant measures (acims), which are finite for $\alpha \in (0,1)$, and have summable correlations for $\alpha \in (0, \frac{1}{2})$ (Gouëzel, 2004b).

The standard framework for theoretical study of intermittent maps is via the so-called induced map $f^\tau : [a,1] \circlearrowleft$ (i.e. $f(x)$ iterated $\tau(x)$ times), where $\tau : [a,1] \to \mathbb{N}^+$

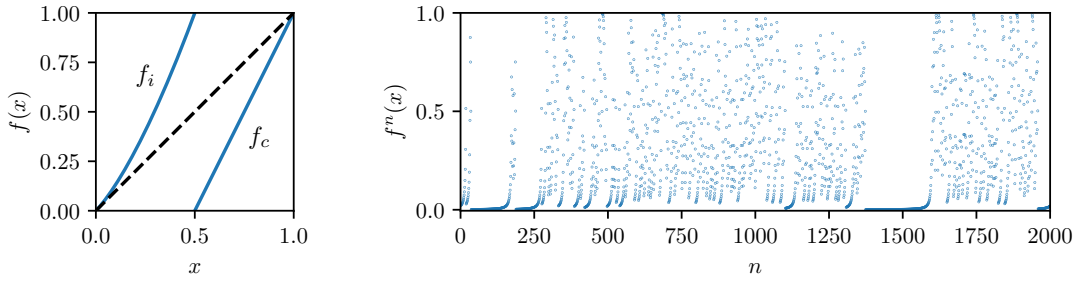

FIGURE 3.1. Left: Graph of the LSV map (3.2) for $\alpha = 0.8$. Right: Time series of its dynamics.

is the return time to the inducing set, the "good" set $[a, 1]$:

$$\tau(x) := \inf\{n \in \mathbb{N}^+ : f^n(x) \in [a, 1]\}. \tag{3.3}$$

The induced map is uniformly expanding (see Figure 3.2), and it is therefore possible to apply results on uniformly expanding dynamics to it, as well as numerical methods such as those developed in Chapter 2.

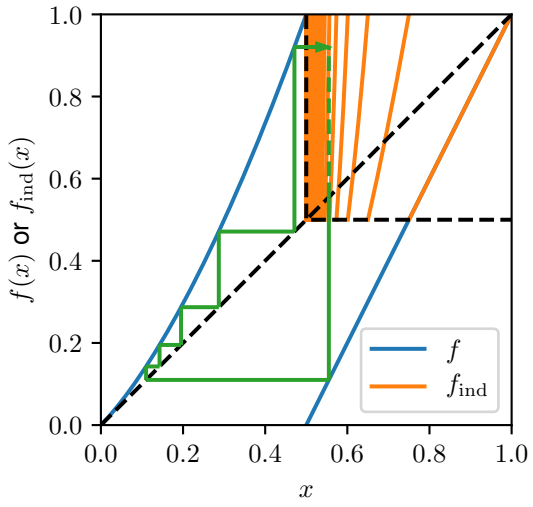

FIGURE 3.2. The LSV map (3.2) (blue) for $\alpha = 0.8$ with its induced map on $[\frac{1}{2}, 1]$ (orange).

The non-mixing dynamics near the fixed point poses a problem for obtaining accurate numerical estimates for these maps. This is true even when one naively

attempts to estimate expectations of observables by taking Birkhoff sums: expectations converge as a stable law for $\alpha \in [\frac{1}{2}, 1]$ (Gouëzel, 2004b), and do not converge for $\alpha \geq 1$ since for these values of α the acim has infinite measure (Aaronson, 1997).

Nonetheless, a variety of transfer-operator based numerical methods have been considered. Various authors have proposed applying Ulam's method or modifications thereof to the full intermittent system: that is, partitioning the phase space into intervals and calculating statistical properties from an associated discretisation of the transfer operator (Murray, 2010; Froyland et al., 2011; Galatolo and Nisoli, 2014). With a suitable choice of partitions estimates of the invariant measure were found to converge for $\alpha < 1$ as $\mathcal{O}(N^{-(1-\alpha)})$, where N is the cardinality of the partition: the source of this slow convergence is the map's weak expansion near the neutral fixed point.

By discretising the transfer operator of the induced map, the slow convergence with respect to the partition size was notionally avoided by Bahsoun et al. (2015) and the standard convergence rate of Ulam's method, $N^{-1} \log N$ in the partition cardinality N, was obtained. However, to calculate an Ulam-like discretisation of the transfer operator, one must repeatedly evaluate the induced map: this requires iterating the full map past the neutral fixed point, a procedure whose computational expense grows as N^{α}. Furthermore, the number of evaluations of the induced map required to estimate the transfer operator to a given accuracy is proportional to the norm of the derivative of the induced map, which increases polynomially with the length of the orbit of the full dynamics corresponding to the step of the induced map. The computational time required to estimate an order N Ulam matrix is consequently $\mathcal{O}(N^{\alpha+2})$ as opposed to the $\mathcal{O}(N)$ for uniformly expanding maps, and perhaps due to the inefficiency of this method, this numerical approach has not yet been implemented. Consequently, it is clear that for good numerics it is not enough just compute with the induced map: it is necessary also to avoid iterating through the full dynamics.

In this chapter, we will present a numerical method that both solves this problem and harnesses the smooth structure to attain fast convergence rates with low numerical overhead. There are three main ingredients at play. The first ingredient solves the problem of efficiently computing the induced map: this is achieved by employing the Abel function, a concept developed in the area of functional iterative equations (Abel, 1826; Kuczma et al., 1990). The Abel function conjugates the map f close to the neutral fixed point to a unit shift, thus allowing the induced map to be calculated in closed form. Furthermore, we find that the Abel function possesses an asymptotic expansion that enables efficient computation. The second ingredient is the Chebyshev Galerkin transfer operator method presented in the previous chapter, which enables statistical properties of the induced map to be computed exponentially accurately. The Chebyshev method requires pointwise evaluation of the

action of the transfer operator of the induced map, which has an infinite number of branches, each of which contribute to the transfer operator. To treat this efficiently, the final ingredient is the Euler-Maclaurin formula, which allows us to very efficiently evaluate infinite sums over the branches. The end result is that statistical properties of the induced map can practically be computed to very high accuracy: because many statistical properties of the full map can be obtained by summing appropriate statistics of the induced map over backward orbits (see Proposition 3.3 and Remark 3.4 below), the full map's statistical properties can then be accurately computed by the same methods.

This chapter is organised as follows. In Section 3.2 we state the main theorems that form the core ideas in our numerical methods; in Section 3.3 we prove Theorems 3.1-3.2; in Section 3.4 we give explicit bounds the convergence of Euler-Maclaurin summation over backward orbits. In Section 3.5 we give some rigorously validated numerical results and in Section 3.6 we consider possible extensions of the work.

3.2. Main theorems

We first state two theorems which allow us to efficiently compute the induced map by means of Abel functions, which allows us to express the induced map

$$f^\tau(x) = (\underbrace{f \circ \cdots \circ f}_{\tau(x) \text{ times}})(x)$$

in closed form. A cursory background on Abel functions is given in Section 3.3

In Theorem 3.1 we show that the induced map and return times can both be expressed in terms of an Abel function, if it exists; in Theorem 3.2 we show that such a function exists and is the principal Abel function: we give an asymptotic expansion for it around the neutral fixed point.

Theorem 3.1. *For maps in class PM, the return map $f^\tau : [a,1] \to [a,1]$ has the explicit expression*

$$f^\tau = A^{-1}(\{A(f_g(x))\}) \tag{3.4}$$

where $\{y\}$ denotes the fractional part of y.

The return time $\tau : [a,1] \to \mathbb{N}$ given by (3.3) is also explicitly given in terms of the Abel function by

$$\tau = \begin{cases} \lfloor A(f_g(x)) \rfloor + 1, & x \neq a \\ 1, & x = a, \end{cases} \tag{3.5}$$

where the bijection $A : [0,1] \to [0,\infty]$ is an Abel function with

$$A(f_b(x)) = A(x) - 1 \tag{3.6}$$

for $x \in [0, a]$ and $A(1) = 0$.

Theorem 3.2. *For maps in class PM there exists a (principal) Abel function $A : [0, 1] \to [0, \infty]$ such that*

(a) A satisfies

$$A(x) = \lim_{k\to\infty} \frac{1}{\alpha h'(0)} \left(f_b^{-k}(x)^{-\alpha} - f_b^{-k}(1)^{-\alpha} \right). \tag{3.7}$$

(b) There is an analytic extension of A into the complex plane having asymptotic expansion uniformly as $z \to 0, \Re z^{-1} \geq R_1 < \infty$

$$A(z) \sim a_{-1} z^{-\alpha} + a_\ell \log x + a_0 + \sum_{n=1}^{\infty} a_n z^{n\alpha}. \tag{3.8}$$

Furthermore, this expansion, truncated after the $z^{N\alpha}$ term for $N \geq 1$, has error

$$E_N(z) = \mathcal{O}(N^{N+2}|z|^{-(N+1)\alpha})$$

with explicit constants given in the proof.

An example of a principal Abel function is plotted on the map's real domain in Figure 3.3.

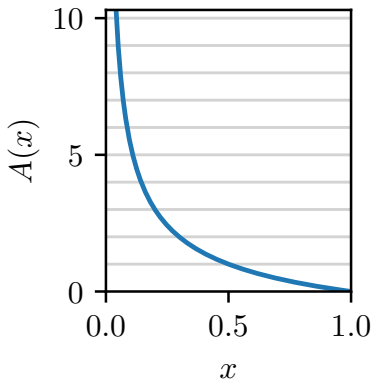

Figure 3.3. Plot of the principal Abel function for the LSV map for $\alpha = 0.8$.

Because the induced map has many desirable properties for the computation of statistical properties, in particular being uniformly expanding, we will compute statistical properties of the full map using those of the induced map. This will at various

points require the computation of sums over backward orbits of the intermittent dynamics: for example, using the chain rule the transfer operator of the induced map has the form

(3.9)
$$\mathcal{L}_{ind}\phi(z) = \sum_{f^\tau(y)=z} |(f^\tau)'(y)|^{-1}\phi(y) = \sum_{n\in\mathbb{N}, f_b^n(f_g(y))=z} (f_b^n)'(f_g(y))^{-1}|f_g'(y)|^{-1}\phi(y).$$

To deal with these sums in a unified way, we will suppose that the summands of these systems can be written as functions Q of the backward orbit $x = f_g(y)$, the derivative $d = (f_b^n(x))^{-1}$, and the orbit index x: for example, from (3.9) we can see that the induced map's transfer operator has summand function $Q_{\mathcal{L}_{ind},\phi}(x, d, n)$

$$Q_{\mathcal{L}_{ind},\phi}(x, d, n) = d(\mathcal{L}_g\phi)(x),$$

where $\mathcal{L}_g$ is the transfer operator of $f_g : [a, 1] \to [0, 1]$:

$$(\mathcal{L}_g\phi)(x) = \sum_{f_g(y)=x} |f_g'(y)|^{-1}\phi(y).$$

We introduce the operator $\mathbb{T}[n; \cdot]$ which acts on the summand functions $Q : (0, 1] \times \mathbb{R} \times \mathbb{N} \to \mathbb{R}$ so as to output the nth summand:

$$\mathbb{T}[Q](n; x) = Q(f_b^{-n}(x), (f_b^{-n})'(x), n) \tag{3.10}$$
$$= Q(A^{-1}(A(x) + n), (A'(x)(A^{-1})'(A(x) + n), n). \tag{3.11}$$

We further introduce the operator $\mathbb{S}$ which acts on the summand functions Q to output the sum over all n:

$$\mathbb{S}[Q](x) = \sum_{n=0}^{\infty} \mathbb{T}[Q](n; x).. \tag{3.12}$$

In Section 3.4 we will show that, when Q extends to a complex analytic function, these sums $\mathbb{S}[Q](x)$ may be very efficiently estimated using the Euler-Maclaurin formula.

The proposition provides the recipe to compute the transfer operator $\mathcal{L}_{ind}$ and the acim of the induced map, via Abel functions. As in Chapter 2, we consider the solution operator, but now for the induced map: $\mathcal{S}_{ind} = (\mathrm{id} - \mathcal{L}_{ind} + u \int)^{-1}$, where $\int$ is the Lebesgue integral functional on $[a, 1]$ and $u(x) = \frac{1}{1-a}$.

Proposition 3.3. *Let $f \in PM$ as in (3.1). The induced map's invariant probability measure ρ_{ind} is given by $\mathcal{S}_{ind}u$.*

Furthermore, the induced map's transfer operator $\mathcal{L}_{ind} : BV([a, 1]) \to BV([a, 1])$ can be written as $\mathcal{L}_{ind}[\phi](x) = \mathbb{S}[Q_{\mathcal{L}_{ind},\phi}](x)$, where $Q_{\mathcal{L}_{ind},\phi}(x, d, n) = d(\mathcal{L}_g\phi)(x)$.

Remark 3.4. *Many statistics of the full dynamics may be efficiently computed through similar formulations. In particular, we have the following formulae for some statistical quantities associated the full dynamics:*

a) Expectations of functions ψ of the return time to the inducing set are given by

$$\rho_{ind}(\psi(\tau)) := \int_a^1 \psi(\tau(x))\rho_{ind}(x)\mathrm{d}x = \int_a^1 \mathbb{S}[Q_{\tau,\psi}](x)\mathrm{d}x, \tag{3.13}$$

where

$$Q_{\tau,\psi}(x,d,n) = \psi(n)d(\mathcal{L}_g\phi)(x).$$

b) The full invariant measure $\rho\mathrm{d}x$ evaluated pointwise is given by

$$\rho(x) = \mathbb{S}[Q_{\mathcal{L}_{ind},\rho_{ind}}](x). \tag{3.14}$$

This is normalised so its restriction to $[a,1]$ is a probability measure; for $\alpha < 1$ it may be renormalised to a probability measure on the full set by the constant factor $\frac{1}{\rho_{ind}(\tau)}$.

c) The average of an observable $A : [0,1] \to \infty$ over ρ is given by

$$\rho(A) := \int_0^1 A(x)\rho(x)\,\mathrm{d}x.$$

Analyticity-preserving properties of $\mathbb{S}$ ensure that $\rho(z)$ extends into the complex plane sufficiently as to allow for accurate quadrature.

d) For $\alpha < \frac{1}{2}$, the diffusion coefficient of an observable $A : [0,1] \to \infty$ is given by

$$\sigma_f^2(A) = \rho(\phi_A + \mathbb{S}[Q_{\mathcal{L}_{ind},\mathcal{S}_{ind}(\phi_A|_{[a,1]})}]),$$

where $\phi_A(x) = \mathbb{S}[Q_{\phi,A}(x)]$ and $Q_{\phi,A}(x,d,n) = d\,\phi(x)(A(x) - \rho(A))$.

Theorems 3.1-3.2 and Proposition 3.3 can be exploited to design algorithms to capture statistical properties of intermittent maps with high accuracy. We give some results from such an algorithm in Section 3.5.

3.3. Return maps and Abel function

Given an iterated function of one dimension $x_{n+1} = g(x_n)$, a function A is considered an Abel function of g if it satisfies the Abel functional equation $A(g(x)) = A(x) + 1$ and is invertible (at least locally). The existence and behaviour of Abel functions around fixed points which are linearly neutral and stable (as opposed to unstable, the case we consider) have been studied in statistics by Szekeres (1958) and others (Kuczma et al., 1990). This corresponds to studying the local inverse f_b^{-1} of our map near the fixed point 0, because the fixed point is linearly neutral and *unstable*. Consequently, our definition assumes iteration of f decrements the Abel

function A (as in Theorem 3.1) rather than incrementing it, as is standard in the literature.

Around a fixed point there are an infinite number of continuous or even smooth solutions to the Abel functional equation. It is possible to define a so-called principal Abel function via a certain iterative equation (Szekeres, 1958), which may be seen to be equivalent to (3.7): principal Abel functions have the best regularity properties of all possible solutions to the Abel equation.

We begin by proving Theorem 3.1, which states that the induced map and return time can be appropriately computed using a monotonic function A satisfying the Abel equation (3.6).

Proof of Theorem 3.1. Suppose A is a bijection $[0,1] \to [0,\infty]$ satisfying $A(f_b(x)) = A(x) - 1$ and $A(1) = 0$. For $x \in (a,1]$, the return time, which measures the number of iterates required to return to the inducing set, is given by

$$\tau(x) = \inf\{n \in \mathbb{N}^+ : f^n(x) \in [a,1]\}.$$

This definition implies that $f^j(x) \in (0,a)$ for $0 < j < \tau(x)$, and consequently $f^j(x) = f_b(f^{j-1}(x)) = A^{-1}(A(f^{j-1}(x) - 1))$ for $0 < j < \tau(x)$. As a result,

$$f^\tau(x) = f^{\tau(x)}(x) = f_b^{\tau(x)-1}(f(x)) = A^{-1}(A(f(x)) - \tau(x) + 1),$$

and since $x \in (a,1]$, $f(x) = f_g(x)$.

Since $f^\tau(x) \in (a,1]$ and since $A : [0,1] \to [0,\infty]$ is a bijection with $A(1) = 0$, we find that A is decreasing and so $0 \le A(f^\tau(x)) < A(a) = A(f_b(a)) + 1 = A(1) + 1 = 1$, and consequently $A(f(x)) - \tau(x) + 1 \in [0,1)$. Using that $\tau(x)$ is an integer we obtain (3.4) and (3.5). □

We will now prove the existence of a principal Abel function with nice asymptotic properties (Theorem 3.2). We will do this by showing that an analytic function satisfying part (b) of the theorem must have asymptotic properties as given in part (a). Using results in Szekeres (1958), we then prove the existence of such a function.

Proof of Theorem 3.2. In this setting we find it convenient to transform to coordinates $z = x^\alpha$, considering the conjugated inverse map which we define

$$\hat{f}(z) := z\hat{h}(z) := f_b^{-1}(x)^\alpha.$$

By the implicit function theorem, $\hat{f}$ is uniquely defined for z in a complex neighbourhood of $[0, a^\alpha]$, and in particular for $|z| \le R$ for some $R > 0$. We consider the principal Abel function for this map, having $\hat{A}(\hat{f}(z)) = \hat{A}(z) + 1$, and set $A(x) = \hat{A}(z)$.

Let the power series at 0 of $\hat{f}(z) = z + \sum_{n=1}^\infty \hat{h}_n(-z)^{n+1}$: in particular, $\hat{h}_1 = \alpha h'(0)$. We have that $\hat{f}^{-1}$ is analytic in a neighbourhood of zero with $\hat{f}^{-1}(z) \sim$

$z + \hat{h}_1 z^2 + \ldots$, and consequently that $\hat{f}$ is similarly analytic near zero with $\hat{f}(z) \sim z - \hat{h}_1 z^2 + \ldots$.

For $n \geq 0$ define the following functions, which are holomorphic except at zero,

$$\hat{A}_n(z) = a_{-1} z^{-1} + a_\ell \alpha^{-1} \log z + C + \sum_{i=1}^{n} a_i z^i, \tag{3.15}$$

with constant C to be determined later, such that as $z \to 0$,

$$\hat{A}_n(\hat{f}(z)) - \hat{A}_n(z) + 1 =: D_n(z) = \mathcal{O}(z^{n+2}). \tag{3.16}$$

Let us define the function g by

$$\frac{1}{\hat{f}(z)} = \frac{1}{z} + \hat{h}_1 + g(z) z. \tag{3.17}$$

From the Taylor expansion of $\hat{f}$ at $z = 0$ we can see that the magnitude $|g(z)| \leq G$ for all $|z| \leq R$ and some constant $G < \infty$.

The error D_n is given by the following lemma, whose proof is in Appendix B:

Lemma 3.5. *Let*

$$r_n = \min\{R, 0.4(\hat{h}_1 + \sqrt{0.4G})^{-1}\} n^{-1}, \tag{3.18}$$

$$d_2 = 1 + 2.5 e^{3/5} (1 + 0.4 G \hat{h}_1^{-2}), \tag{3.19}$$

$$d_1 = \frac{1 + G\hat{h}_1^{-2}}{d_2^2}. \tag{3.20}$$

For $|z| \leq r_n$,

$$|D_n(z)| \leq d_1 \left(\frac{d_2 |z|}{r_n} \right)^{n+2}.$$

The following lemmas, whose proofs are in Appendix B give bounds on iterates of $\hat{f}$:

Lemma 3.6. *Let* $\aleph \in (0,1)$ *and* $R_1 = \min\{R, \aleph G^{-1} \hat{h}_1\}$.

Then for all z *with* $\Re z^{-1} \geq R_1$, *and all* $k \in \mathbb{N}$,

$$\Re z^{-1} + (1 + \aleph)\hat{h}_1 k \geq \Re \hat{f}^k(z)^{-1} \geq \Re z^{-1} + (1 - \aleph)\hat{h}_1 k, \tag{3.21}$$

and

$$|z^{-1} + k\hat{h}_1| + \aleph \hat{h}_1 k \geq |\hat{f}^k(z)|^{-1} \geq |z^{-1} + k\hat{h}_1| - \aleph \hat{h}_1 k \geq R_1. \tag{3.22}$$

Lemma 3.7. *Let* $\aleph, R_1$ *be as before, and let* $\bar{\beta} - 1 \geq \delta \geq 0$. *Then for all* z *with* $\Re z^{-1} \geq R_1$,

$$\sum_{k=0}^{\infty} |\hat{f}^k(z)|^{\bar{\beta}} k^\delta \leq \beth_{\bar{\beta},\delta} |z|^{\bar{\beta}-\delta-1},$$

where

$$\beth_{\bar{\beta},\delta} := (1-\aleph)^{-\bar{\beta}}\hat{h}_1^{-\delta-1}\left(\frac{1}{\delta+1}+\frac{1}{\bar{\beta}-\delta-1}+\hat{h}_1 R_1^{-1}\right).$$

Define the sets

$$S_r = \{z \in \mathbb{C} \mid 0 < |z| < r, \Re z^{-1} \geq R_1\}$$

for $r > 0$, where R_1 is given in Lemma 3.6, and consider a function $\hat{A} : S_r \to \mathbb{C}$ such that for all $z \in S_r$ and $n > 0$,

$$\lim_{k\to\infty} \hat{A}(\hat{f}^k(z)) - \hat{A}_n(\hat{f}^k(z)) = 0. \tag{3.23}$$

Later, we will show that $A(z^\alpha)$ where A is given in (3.7) is such a function. We will prove bounds on $|\hat{f}^n(z)|$ for $z \in S_{R'}$ for some $R' < R$, which will allow us to bound the error between $\hat{A}$ and the $\hat{A}_n$ on this set.

Now defining

$$E_n(z) = \hat{A}(z) - \hat{A}_n(z),$$

we have that

$$E_n(z) = E_n(\hat{f}(z)) - (\hat{A}(\hat{f}(z)) - \hat{A}(z)) + (\hat{A}_n(\hat{f}(z)) - \hat{A}_n(z)) = E_n(z) + D_n(z)$$

and thus by (3.23) and the fact that $f^{-k}(z) \to 0$ as $k \to \infty$, we have that

$$E_n(z) = \sum_{k=1}^{\infty} D_n(\hat{f}^k(z))$$

and thus by Lemma 3.5

$$|\hat{A}(z) - \hat{A}_n(z)| \leq d_1 d_2^{n+2} r_n^{-(n+2)} \sum_{k=1}^{\infty} |\hat{f}^k(z)|^{n+2}.$$

When $\Re z^{-1} \geq R_1$ we can apply Lemma 3.7 to obtain that for $n \geq 1$ and $|z| \leq \min\{R_1, r_n\}$,

$$\begin{aligned} |\hat{A}(z) - \hat{A}_n(z)| &\leq d_1 d_2^{n+2} r_n^{-(n+2)} \beth_{n+2,0} |z|^{n+1} \\ &\leq d_3 d_2^{n+2} r_n^{-(n+2)} |z|^{-n+1}, \end{aligned} \tag{3.24}$$

where

$$d_3 = \beth_{n+2,0} d_2 \tag{3.25}$$

Consequently $\hat{A}$ satisfies the appropriate asymptotic expansion generated by the $\hat{A}_n$ with the error bound given by (3.24).

By change of coordinates $x^\alpha = z$ we have that

$$E_n(z) = d_3 (d_2/r_n)^{n+2} \left(|z|^{-\alpha} - dR\right)^{-(n+1)}.$$

We now show that $\hat{A}$ is in fact given by the principal Abel function that we desired. Let

$$\hat{A}(x) = \lim_{k\to\infty} \frac{\hat{f}^k(1) - \hat{f}^k(x)}{\hat{f}^k(1) - \hat{f}^{k+1}(1)}.$$

Note that $\hat{A}(x^\alpha)$ is just our principal Abel function (3.7), and $\hat{A}(\hat{f}(z)) = \hat{A}(z) - 1$ where $\hat{A}$ is defined. As a result of Lemma 7 in Szekeres (1958), $\hat{A}$ extends into the complex plane, and for a correct choice of C in (3.15) for every $\mathfrak{C}$ there exists a $\mathfrak{D}$ so that this map satisfies

$$\lim_{z\to 0} \left(\hat{A}(z) - \hat{A}_0(z)\right) = 0$$

for $|\Re z^{-1}| > \mathfrak{D}$, $|\Im z^{-1}| \le \mathfrak{C}$. By Lemma 3.6 we have that $\Re \hat{f}^n(z)^{-1}$ is increasing in n, and the following lemma, proved in Appendix B, gives that $\Im \hat{f}^n(z)^{-1}$ is bounded:

Lemma 3.8. *For all z with $\Re z^{-1} \ge R_1$,*

$$\sup_{k\in\mathbb{N}} \left|\Im \hat{f}^k(z)^{-1}\right| < \infty.$$

As a result, $\hat{f}^n(z)$ goes to 0 as $n \to \infty$.

Furthermore, as a result of the monotonicity of f_b and (3.7), A is clearly monotonically increasing on $[0,1] \to [0,\infty]$; because it is analytic and unbounded it must be a bijection. □

Remark 3.9. *The Thaler map is an interval map with an explicitly known invariant measure that has a neutral fixed point of order $1+\alpha$ at zero:*

$$f(z) = z\left[1 + z^{\alpha-1}\left((1+z)^{1-\alpha} - 1\right)\right]^{1/(\alpha-1)}.$$

This map is not in class PM as the series expansion of $f(z)/z$ at zero contains integer powers of z as well as of z^α. However, one could extend the methods in this paper accordingly.

3.4. Calculating statistical properties via inducing

Because the infinite sums required to evaluate statistical properties, such as in Proposition 3.3 and Remark 3.4, are summing over smooth functions evaluated on a lattice, we can use the Euler-Maclaurin formula to approximate these sums with exponentially decreasing errors. We state a general theorem that in particular allows us to obtain rigorous bounds on the error of these approximations.

Define the small, bounded sets $\mathcal{R}_s := \{x \in \mathbb{C} : \Re(x^{-\alpha}) \ge s^{-1}\}$, and its transform to z coordinates, $\hat{\mathcal{R}}_s := \{z \in \mathbb{C} : \Re(z^{-1}) \ge s^{-1}\}$.

We will first find it useful to define a constant encoding the regularity of our of our map,

$$G' = \sup_{z \in \hat{\mathcal{R}}_{R_1}} \left| \frac{d}{dz} z g(z) \right|$$

and a radius

$$Z := \min\{R_1, (2G')^{-1/2}, (2G' \beth_{2,0})^{-1}\}, \tag{3.26}$$

which will be used to specify the region inside which the Euler-Maclaurin formula may be used.

Theorem 3.10. *Suppose $Q(x, d, n)$ is analytic such that for some $\bar{Q}$, some non-negative β, γ, δ with $\bar{\beta} := (\beta + (1+\alpha)\gamma)/\alpha > 1 + \delta$, and for all $z \in \mathcal{R}_R$, all $d \leq 1$ and all n with $\Re n > 0$,*

$$|Q(z, d, n)| \leq \bar{Q} |z|^\beta |d|^\gamma |n|^\delta.$$

Let $\rho > 0$.
Then for all z such that

$$n^*_\rho := \inf\{n : f_b^{-n} z \in \mathcal{R}_{(Z^{-\alpha} + 2\hat{h}_1 + \rho)^{-1/\alpha}}\}$$

is defined, then

$$\mathbb{S}[Q](z) = \sum_0^{n^*_\rho - 1} \mathbb{T}[Q](n; z) + \frac{1}{2}\mathbb{T}[Q](n^*_\rho; z) - \int_0^{f_b^{-n^*_\rho}(z)} A'(\zeta) Q\left(\zeta, \frac{A'(z)}{A'(\zeta)}, A(\zeta) - A(z)\right) \mathrm{d}\zeta \\ - \sum_{k=1}^{K} \frac{B_{2k}}{(2k)!} \frac{\partial^{2k-1}}{\partial n^{2k-1}} \mathbb{T}[Q](n^*_\rho; z) + \mathcal{E}_K, \tag{3.27}$$

where B_p are the Bernoulli numbers, and $\mathcal{E}_K = O(\rho^{-K+\delta-(\beta+(1+\alpha)\gamma)/\alpha})$, with an explicit bound given in (3.30-3.31).

We will find the following proposition useful in proving this theorem:

Proposition 3.11. *Let $m \in \mathbb{C}$. If we restrict $\hat{A}$ to act on $\hat{\mathcal{R}}_Z$, then for any $z_0 \in \hat{\mathcal{R}}_{(Z^{-1} + 2\hat{h}_1 |m|)^{-1}}$,*

$$|A^{-1}(A(z_0) + m)^{-\alpha} - z_0^{-\alpha}| \leq 2\hat{h}_1 |m|.$$

Proof of Theorem 3.10. A simple application of the Euler-Maclaurin formula (Abramowitz and Stegun, 1973) gives most of the terms in (3.27); we convert the integral expression

$$\begin{aligned}\int_{n_\rho^*}^{\infty} \mathbb{T}[Q](n,z)dn &= \int_{n_\rho^*}^{\infty} Q\left(A^{-1}(n+A(z)), \frac{A'(z)}{A'(A^{-1}(n+A(z)))}, n\right) dn \\ &= \int_{n_\rho^*+A(z)}^{\infty} Q\left(A^{-1}(n), \frac{A'(z)}{A'(A^{-1}(n))}, n-A(z)\right) dn \\ &= \int_0^{T^{n_\rho^*}(z)} -A'(\zeta) Q\left(\xi, \frac{A'(z)}{A(\xi)}, n-A(z)\right) d\zeta.\end{aligned}$$

The remainder term $\mathcal{E}_K$ can be bounded (Lehmer, 1940; Abramowitz and Stegun, 1973) by

$$|\mathcal{E}_K| \le \frac{2}{(2\pi)^{2K+1}} \int_{n_\rho^*}^{\infty} \left| \frac{\partial^{2K+1}}{\partial n^{2K+1}} \mathbb{T}[Q](n;z) \right| dn. \tag{3.28}$$

From Lemma 3.6 we have that if $s^{-1} = \Re z^{-\alpha}$ with $s^{-1} > Z^{-1} + 2\hat{h}_1$, then $z \in \mathcal{R}_s$ with

$$A^{-1}(n+A(z)) = f_b^{-n}(z) \in \mathcal{R}_{(s^{-1}+n(1-\aleph)\hat{h}_1)^{-1}}$$

for integer n, and from Proposition 3.11 that for $n \in (0,1)$ that

$$A^{-1}(A(z)-n) \in \mathcal{R}_{(s^{-1}-2\hat{h}_1)^{-1}};$$

consequently for all $n > 0$,

$$A^{-1}(A(z)+n) \in \mathcal{R}_{(s^{-1}+\hat{h}_1((1-\aleph)n-2))^{-1}}. \tag{3.29}$$

Thus, for any z, $n \ge n_\rho^*$, and $m \in B(0, \rho + (1-\aleph)(n-n_\rho^*)/2)$, $\zeta := A^{-1}(A(z) + n+m) \in \mathcal{R}_Z$ and so

$$\begin{aligned}|\mathbb{T}[Q](n+m;z)| &= \left| Q\left(\zeta, \frac{A'(z)}{A'(\zeta)}, n+m\right)\right| \\ &\le \bar{Q}|\zeta|^\beta | \frac{|A'(z)|^\gamma}{|A'(\zeta)|^\gamma} |n+m|^\delta \\ &= \bar{Q}|A'(z)|^\gamma |\zeta|^\beta \left| \frac{\zeta^{1+\alpha}}{\zeta^{1+\alpha}A'(\zeta)} \right|^\gamma |n+m|^\delta \\ &\le \bar{Q}|A'(z)|^\gamma |\zeta|^{\beta+(1+\alpha)\gamma} 2^{|\gamma|} h_1^\gamma |n+m|^\delta,\end{aligned}$$

by Lemma 3.12.

We know from (3.29) and Lemma 3.12 that $\zeta \in \mathcal{R}_Z$, and thus if $\bar{\beta} = (\beta + (1+\alpha)\gamma)/\alpha \geq 0$,

$$|\mathbb{T}[Q](n+m;z)| \leq \bar{Q}|A'(z)|^\gamma Z^{-\bar{\beta}} 2^{|\gamma|} h_1^\gamma (n+\rho+(1-\aleph)(n-n_\rho^*)/2)^\delta.$$

By applying Proposition 3.11 we have that

$$|\zeta|^{-\alpha} \leq |T^{n_\rho^*}(z)|^{-\alpha} + 2h_1(n+|m|-n_\rho^*) = |T^{n_\rho^*}(z)|^{-\alpha} + 2h_1(3-\aleph)(n-n_\rho^*)/2$$

and thus if $\bar{\beta} \leq 0$,

$$|\mathbb{T}[Q](n+m;z)| \leq \bar{Q}|A'(z)|^\gamma 2^{|\gamma|} h_1^\gamma \left(|T^{n_\rho^*}(z)|^{-\alpha} + 2h_1 \frac{3-\aleph}{2}(n-n_\rho^*)\right)^{-\bar{\beta}} \left(n+\rho+\frac{1-\aleph}{2}(n-n_\rho^*)\right)^\delta.$$

We then have by Cauchy's formula that

$$\left|\frac{\partial^{2K+1}}{\partial n^{2K+1}}\mathbb{T}[Q](n;z)\right| \leq (2K+1)!(\rho+(1-\aleph)(n-n_\rho^*)/2)^{-(2K+2)} \cdot \sup\left\{|\mathbb{T}[Q](n+m;z)| \mid m \in B(0,\rho+(1-\aleph)(n-n_\rho^*)/2)\right\};$$

thus, from (3.28) and with some simplification,

$$|\mathcal{E}_K| \leq \frac{(2K+1)!}{(2\pi\rho)^{2K+1}} \frac{4\bar{Q}|A'(z)|^\gamma}{1-\aleph} \left(1+\rho^{-1}n_\rho^*\right)^\delta W, \tag{3.30}$$

where

$$W = \begin{cases} \frac{\rho^\delta}{2K+1-\delta} 2^{|\gamma|} h_1^\gamma Z^{-\bar{\beta}}, & \bar{\beta} \geq 0, \\ \frac{\rho^{\delta-\bar{\beta}}}{2K+1+\bar{\beta}-\delta} 2^{|\gamma|} h_1^\gamma \max\left\{\rho^{-1}|T^{n_\rho^*}(z)|, 2h_1(1+2(1-\aleph)^{-1})\right\}^{-\bar{\beta}}, & \bar{\beta} < 0, \end{cases} \tag{3.31}$$

recalling that we defined $\bar{\beta} := (\beta + (1+\alpha)\gamma)/\alpha$.

Finally, since the integrand in the integral in (3.27) is $\mathcal{O}(z^{-(\alpha+1)(1-\gamma)+\beta+\alpha\delta})$, we know that it will converge if $\bar{\beta} + \delta < 1$. □

The following lemma, used in the proof of Theorem 3.10 bounds the derivative of the Abel function:

Lemma 3.12. *For any $z \in \mathcal{R}_Z$,*

$$(2h_1)^{-1} \leq |z^{\alpha+1}A'(z)| \leq 2h_1^{-1}.$$

Proof. It is possible to show from the definition of A in (3.7) that

$$z^{-2}\hat{A}'(z) = \lim_{k\to\infty} \frac{z^2(\hat{f}^k)'(z)}{\hat{h}_1(\hat{f}^k(z))^2} = \hat{h}_1^{-1} \prod_{k=0}^{\infty} \xi(\hat{f}^k(z)),$$

where

$$\xi(z) := \frac{z^2}{\hat{f}(z)^2}\hat{f}'(z).$$

From the definition of $g(z)$ we have that

$$\xi(z) = 1 + z^2(zg)'$$

and so for $z \in \hat{\mathcal{R}}_R$ such that $|z| \leq (2G')^{-1/2}$,

$$|\log|\xi(z)|| \leq \log 2\, 2G'|z|^2.$$

By Lemma 3.6 then, for $z \in \hat{\mathcal{R}}_{\min\{R,(2G')^{-1/2}\}}$,

$$\begin{aligned}\left|\log \hat{h}_1|z^{-2}\hat{A}'(z)|\right| l &\leq \log 2\, 2G' \sum_{k=0}^{\infty} |\hat{f}^k(z)|^2 \\ &\leq \log 2\, 2G' \beth_{2,0}|z| \\ &\leq \log 2,\end{aligned}$$

for $z \in \hat{\mathcal{R}}_Z$ as required. □

Remark 3.13. *The choice of K which minimises the bound on $\mathcal{E}_K$ for given ρ is asymptotically $K \approx \pi\rho - \frac{1}{2}$, which gives an error $\log|\mathcal{E}_K| \lesssim -2\pi\rho$.*

3.5. Numerical results

We have implemented numerics suggested by the preceding work, specifically rigorously validated algorithms (c.f. Section 2.5.1) to compute acims and return times. Our algorithm first estimates the coefficients of the Abel function's asymptotic expansion (3.8) by matching Taylor coefficients at zero of the Abel equation (3.6). This immediately enables accurate evaluation of the Abel function for x near the fixed point: away from the fixed point accurate estimates may be calculated by numerically iterating backwards until for some $k \in \mathbb{N}$ an iterate $f_b^{-k}(x)$ was sufficiently close to 0 is reached, and using that $A(f_b^{-k}(x)) = A(x) + k$. The algorithm computes a Chebyshev Galerkin matrix as in Algorithm 1: the action of the transfer operator on Chebyshev basis functions is evaluated pointwise by using Proposition 3.3 and (3.27). By the Chebyshev method the acim of the induced map can be rigorously estimated as in Algorithm 1. Estimates of the return time are obtained by a rigorous computation of the return time formula in Remark 3.4(a) using the Euler-Maclaurin formula (3.27); pointwise estimates of the full map's acim are obtained similarly using Remark 3.4(b).

We applied these methods to LSV maps (3.2) for various α. Plots of the acims obtained using our method are given in Figure 3.4 with comparisons to estimates obtained using long time series. An example of the rigorous estimate is given in the following theorem:

Theorem 3.14. *For the LSV map with parameter $\alpha = 0.95$, the expected return time to the set $[\frac{1}{2}, 1]$ is*

$$\rho_{ind}(\tau) = 14.073\ 323\ 220\ 001\ 939\ 529\ 241\ 549\ 699\ 610\ 756\ 609\ 803\ 3171 \pm 10^{-43}.$$

It is illustrative of the power of the method, particularly of the Abel function numerics, to contrast this with an estimate of the expected return time obtained via iterating the LSV map: the sample of 10^8 iterates used in Figure 3.4 furnished an estimate $\rho_{ind}(\psi(\tau)) \approx 8.63$, a 40% error. This large error arises because the distribution of the return time $\tau|_{[\frac{1}{2},1]}$, which the iterates sample, is for maps with $\alpha = 0.95$ very heavy-tailed: in fact it becomes non-integrable at the nearby value $\alpha = 1$. The Euler-Maclaurin summation however allows these tails to be summed over very easily, regardless of their decay rates.

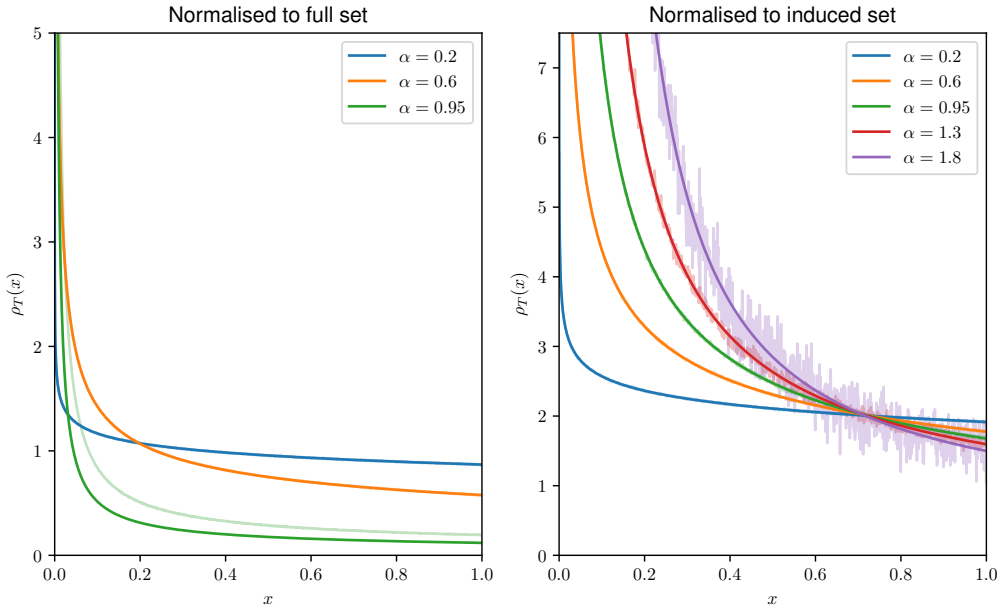

FIGURE 3.4. Absolutely continuous invariant measures of the LSV map for varying α, normalised to a probability measure: (left) on the full domain when $\alpha < 1$, i.e. for finite absolutely continuous invariant measures, (right) for on the inducing domain. In semitransparency, histograms of long time series of 10^8 iterates of LSV maps for various α (see discussion of long time series estimates in Section 2.5).

The results for different values of α were each obtained in 6 hours over 15 hyper-threaded cores of a research server running 2 E5-2667v3 CPUs with 128GB of memory. The number of basis elements used in the Chebyshev Galerkin method was $N = 512$, and 256-bit extended floating point arithmetic was used using the ValidatedNumerics library in Julia (Benet and Sanders, 2019).

Let us briefly note that at a preliminary stage in the development of these numerical methods, we implemented adaptive algorithms to compute acims of the induced and full systems using floating-point arithmetic. These algorithms were similar in spirit to, and made use of, the Poltergeist package discussed in Section 2.5.2. In place of the Euler-Maclaurin formula algorithms used, a poorly optimised version of the already less numerically efficient Abel-Plana formula (Olver, 1997), the adaptive method could obtain acim estimates accurate to 13 decimal places in around 20 seconds. With good numerical optimisation and using the Euler-Maclaurin formula we believe that it would be possible to obtain these estimates in around 2 seconds.

3.6. Discussion

We have presented, and partially demonstrated, a set of numerical methods to compute statistical properties of intermittent maps. The methods are made up of essentially three components: Abel functions to "solve" for the induced map and orbits near the neutral fixed point; Chebyshev Galerkin methods to compute statistics of the induced map; and the Euler-Maclaurin formula to resolve sums over the neutral fixed point. All of these, in particular the Abel function, may be used separately in the computation of statistical properties of these systems, but each component harnesses and preserves smoothness, in fact analyticity, which is what facilitates the exponential convergence of the methods. (To this end, it would be useful for the Chebyshev Galerkin methods to be adapted to spaces of analytic functions.)

The methods, as we have seen, are extremely accurate and work more or less equally well for large $\alpha > 1$ or near critical thresholds for statistics as for small α. This is in great contrast to Birkhoff averaging or Ulam-style methods. We have chiefly presented here a rigorously validated numerical algorithm which is necessarily rather slow, as interval arithmetic and extended precision floating point arithmetic are not integrated into computer architecture in the same way as double precision. A user-friendly, adaptive implementation of these methods would open up intermittent maps (and thus infinite ergodic dynamics) for numerical exploration.

Chapter 4

Linear response theory for macroscopic observables in high-dimensional deterministic systems

4.1. Introduction

Since its introduction in the 1960s, linear response theory (LRT) has been widely used across numerous disciplines to quantify the change of the mean behaviour of observables in a perturbed environment. LRT is valid, in essence, provided the invariant measure varies differentiably with respect to the perturbation; consequently LRT allows for a Taylor expansion of the perturbed invariant measure around the unperturbed invariant measure. Hence, when valid, LRT provides an expression of the average of some observable when subjected to small perturbations from an unperturbed state – the system's so called *response* – entirely in terms of statistical information from the unperturbed system.

The development of the theory occured in statistical mechanics in the context of thermostatted Hamiltonian systems (Kubo, 1966; Balescu, 1975; Zwanzig, 2001; Marconi et al., 2008) but found applications far beyond this realm; recent years have seen an increased interest in LRT and its applications. Climate scientists in particular have successfully applied LRT to eke out valuable information about the change of certain atmospheric and oceanic observables under changed climatic conditions. Applications include atmospheric toy models (Majda et al., 2010; Lucarini and Sarno, 2011; Abramov and Majda, 2007, 2008; Cooper and Haynes, 2011; Cooper et al., 2013), barotropic models (Bell, 1980; Gritsun and Dymnikov, 1999; Abramov and Majda, 2009), quasi-geostrophic models (Dymnikov and Gritsoun, 2001), atmospheric models (North et al., 1993; Cionni et al., 2004; Gritsun et al., 2002; Gritsun and Branstator, 2007; Gritsun et al., 2008; Ring and Plumb, 2008; Gritsun, 2010) and coupled climate models (Langen and Alexeev, 2005; Kirk-Davidoff, 2009; Fuchs et al., 2014; Ragone et al., 2016).

Linear response theory is well-known to be theoretically justified in stochastic systems (Hänggi, 1978; Hairer and Majda, 2010), however in deterministic systems the picture is more complicated. The seminal work by Ruelle (Ruelle, 1997, 1998, 2009a,b) rigorously established that LRT is valid in uniformly hyperbolic Axiom A

systems. Success in reliably estimating the response of a physical system, as exemplified by the above applications in the climate sciences, prompted scientists to believe that general chaotic dynamical systems obeyed LRT. This belief was proven wrong by Baladi and co-workers (Baladi and Smania, 2008, 2010; Baladi, 2014; Baladi et al., 2015; De Lima and Smania, 2018) who showed that simple dynamical systems such as the logistic map violate LRT and support an invariant measure that changes non-smoothly with respect to the perturbation. This raises the question of how a high-dimensional dynamical system, despite its constituent subsystems typically individually violating LRT, may exhibit linear response.

The majority of the scientific community believes that the interaction between the microscopic constituents in typical high-dimensional systems leads to an emergence of LRT at the macroscopic level. How exactly this is achieved and what the conditions are for the dynamical systems for which LRT is guaranteed, however, remains an open question. In the literature the validity of LRT in high-dimensional deterministic systems is often justified by appealing to the *chaotic hypothesis* of Gallavotti and Cohen (Gallavotti and Cohen, 1995a,b; Gallavotti, 2019), which is as stated in Gallavotti (2008):

Conjecture 4.1 (Chaotic hypothesis)**.** *The macroscopic dynamics of a chaotic system on its attractor can be regarded for all practical purposes as an Anosov evolution.*

However, even under this hypothesis one cannot relate the equivalent Anosov systems for different perturbations, which is the focus of LRT. In particular, for dissipative systems the response of the attracting dynamics to perturbations depends on the properties of the flow outside the attractor as well as on it: off the attractor the flow may be non-hyperbolic, and hence leading to a breakdown of linear response.

In a recent paper (Gottwald et al., 2016) we showed that breakdown of LRT might not be detectable using uncertainty quantification when analyzing time series unless the time series is very long (exceeding 1 million data points even for simple one-dimensional systems such as the logistic map, for example) and/or the observables are sensitive to the non-smooth change of the invariant measure. Consequently, the apparent observed validity of LRT in climate science might be a finite size effect of the time series used.

Here we follow a different avenue, drawing on the existence of linear response theory for stochastic dynamical systems. We argue that certain deterministic chaotic systems have stochastic limits for macroscopic observables and in various circumstances these limits imply they are amenable to LRT. Statistical limit laws of deterministic dynamical systems have recently been proven for slow variables in multiscale systems (Melbourne and Stuart, 2011; Gottwald and Melbourne, 2013; Kelly

and Melbourne, 2014) and (Dolgopyat, 2004; De Simoi and Liverani, 2015, 2016). In both cases the diffusive limit of the macroscopic observables relies on the central limit theorem via a summation of infinitely many weakly dependent variables. We consider here high-dimensional dynamical systems whose microscopic constituents may violate LRT, and "macroscopic" observables which are either mean-fields, or distinguished degrees of freedom driven by mean-fields. Distinguished degrees of freedom in high-dimensional weakly coupled systems have been studied extensively in the statistical mechanics literature (Ford et al., 1965; Zwanzig, 1973; Ford and Kac, 1987; Stuart and Warren, 1999; Kupferman et al., 2002; Givon et al., 2004). We consider a variety of microscopic subsystems with different linear response properties, and as well as treating cases where the microscopic variables evolve independently, study the more realistic case when the microscopic dynamics are globally coupled via a mean field. The dynamics of the latter case have been given a significant amount of study, particularly in the case where the coupling is strictly attractive, and complex emergent dynamics at the level of the mean-field have been observed (Kaneko, 1990; Pikovsky and Kurths, 1994; Ershov and Potapov, 1995, 1997; Shibata et al., 1999; Sélley and Bálint, 2016). We shall provide a systematic macroscopic reduction using statistical limit theorems for the mean-field coupled dynamics, and use this reduction to study a range of interesting dynamical scenarios in the context of linear response. To validate our findings, we will use a statistical test that we recently developed, which allows one to probe for the validity of LRT in a given time series (Gottwald et al., 2016). We provide a comprehensive picture of the linear response behaviour of macroscopic observables, for uncoupled, for distinguished-variable driven and for mean field coupled systems, which we summarise in Table 4.1. We find that the existence of LRT depends in an intricate way on the combination of effective stochastic behaviour of the macroscopic observable, the macroscopic dynamics of the thermodynamic limit, the scaling of the coupling, and on the smoothing property of heterogeneously distributed dynamical parameters of the microscopic subsystems. Indeed, we will present a case when all individual microscopic subsystems violate LRT, but the collective macroscopic dynamics obeys LRT, and cases where all the all individual microscopic subsystems obey LRT when uncoupled, but the collective coupled macroscopic dynamics violates LRT.

In Section 4.8 we will make a deeper investigation of the case where we find a violation of LRT from hyperbolic (in fact uniformly expanding) components. We will use the spectral methods from Chapter 2 and numerical continuation techniques to uncover a so-called homoclinic tangency in the reduced macroscopic dynamics. This is a tangency between stable and unstable directions in the dynamics (i.e. a violation of hyperbolicity): as a result of this tangency the system is in violation of the Gallavotti-Cohen chaotic hypothesis.

The chapter is organized as follows. Section 4.2 briefly reviews LRT. We introduce the high-dimensional systems under consideration in Section 4.3 and summarise our results in Section 4.4. Sections 4.5-4.7 provide numerical evidence and an analytical treatment corroborating the results summarized in Table 4.1 respectively for uncoupled mean field, distinguished particle and mean field coupled systems. In Section 4.8, we present a numerical example of a homoclinic tangency in the macroscopic dynamics of a mean field coupled system. We conclude with a discussion and an outlook in Section 2.6.

4.2. Linear response theory

We briefly review some basic notation of linear response theory. Consider a family of dynamical systems $f_\varepsilon : D \to D$ on some space D where the map f_ε depends smoothly on the parameter ε and where for each ε the dynamical system admits a unique invariant physical measure μ_ε. We recall from Chapter 1 that an invariant measure is called physical if for a set of initial conditions of nonzero Lebesgue measure the temporal average of a typical observable converges to the spatial average over this measure. LRT is concerned with the change of the average of an observable $\phi : D \to \mathbb{R}$,

$$\mathbb{E}^\varepsilon[\phi] = \int_D \phi \, \mathrm{d}\mu_\varepsilon$$

upon varying ε. A system exhibits *linear response* at $\varepsilon = \varepsilon_0$, if the derivative

$$\mathbb{E}^{\varepsilon_0}[\phi]' := \frac{\partial}{\partial \varepsilon} \mathbb{E}^\varepsilon[\phi]_{|\varepsilon_0}$$

exists. A sufficient condition for this is that the invariant measure μ_ε is differentiable with respect to ε. This derivative can be expressed entirely in terms of the invariant measure μ_{ε_0} of the unperturbed system using so-called linear response formulae (Ruelle, 2009a, 1998; Baladi, 2014). The average of an observable of the perturbed state is then expressed to first order as

$$\mathbb{E}^\varepsilon[\phi] \approx \mathbb{E}^{\varepsilon_0}[\phi] + (\varepsilon - \varepsilon_0)\, \mathbb{E}^{\varepsilon_0}[\phi]'.$$

If the derivative exists, then this expansion expresses the remarkable result that the average of the perturbed state is determined up to $o(\varepsilon - \varepsilon_0)$ by the properties of the unperturbed system. If however it does not exist, we say there is a breakdown of linear response, which manifests itself in a rough dependency of averages of the observable on the perturbation ε.

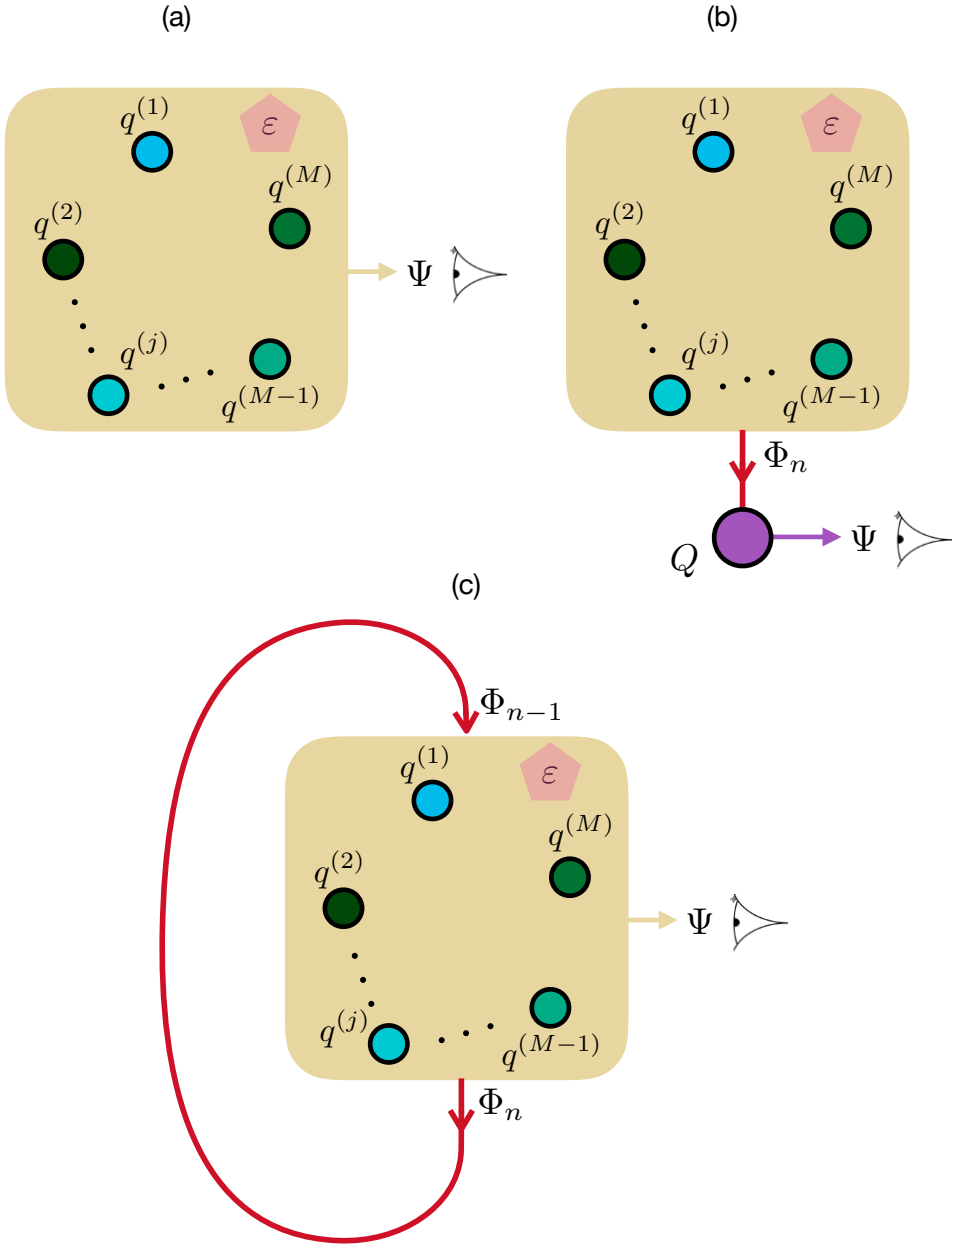

Figure 4.1. General set-up of high-dimensional systems in this chapter. We consider the behaviour of macroscopic observables Ψ which are constructed either from: (a) a large system of M uncoupled microscopic units $q^{(j)}$; (b) a large system in which the M microscopic units drive a macroscopic distinguished variable Q via a mean field variable Φ; or (c) a large system in which the M microscopic units are coupled via a mean field variable Φ. Perturbations ε are applied globally to the dynamics of all heat bath variables $q^{(j)}$.

4.3. Model

We consider here three model scenarios, which are illustrated in Figure 4.1. These are high-dimensional systems composed of $M \gg 1$ chaotic microscopic degrees of freedom $q^{(j)}$, $j = 1, \ldots, M$, which evolve in discrete time n according to their individual parameters $a^{(j)}$.

In the first scenario (Figure 4.1(a)) these degrees of freedom are uncoupled, with

$$q_{n+1}^{(j)} = f(q_n^{(j)}; a^{(j)}, \varepsilon). \tag{4.1}$$

We observe their dynamics via a mean-field Ψ:

$$\Psi_n = \Psi(q_n^{(1)}, q_n^{(2)}, \cdots, q_n^{(M)}) = \frac{1}{M^\gamma} \sum_{j=1}^{M} \psi(q_n^{(j)}), \tag{4.2}$$

where ψ is some observable of microscopic variables. In this work we will always consider smooth observables (at least Hölder continuous), ensuring that the statistical limit laws we later invoke exist.

In the second scenario (Figure 4.1(b)), we consider additionally a macroscopic distinguished variable Q. The microscopic dynamics evolve as in (4.1); the distinguished variable is driven by a mean field Φ_n given by

$$\Phi_n = \frac{1}{M^\gamma} \sum_{j=1}^{M} \phi(q_n^{(j)}; a, \varepsilon), \tag{4.3}$$

for some microscopic observable ϕ, where either $\gamma = \frac{1}{2}$, a diffusive scaling limit, or $\gamma = 1$, a deterministic scaling limit. Except in Section 4.6.1, we choose $\phi(q; a, \varepsilon) \equiv \phi(q)$ to be independent of a and ε. The dynamics of the distinguished variable are given by

$$Q_{n+1} = \tilde{f}(Q_n, \Phi_n), \tag{4.4}$$

and we observe the distinguished variable

$$\Psi_n = \Psi(Q_n). \tag{4.5}$$

In the third scenario (Figure 4.1(c)), the dynamics of the $q_n^{(j)}$ are coupled by a mean field Φ with

$$q_{n+1}^{(j)} = f(q_n^{(j)}, \Phi_n; a^{(j)}, \varepsilon), \tag{4.6}$$

and the mean field is defined as in (4.3) with $\gamma = 1$. We observe a mean-field Ψ given by (4.2). (Note that in this scenario Ψ and Φ are structurally identical.)

To study the effect of the structure of the microscopic variables on the macroscopic dynamics we will consider three different types of microscopic dynamics f.

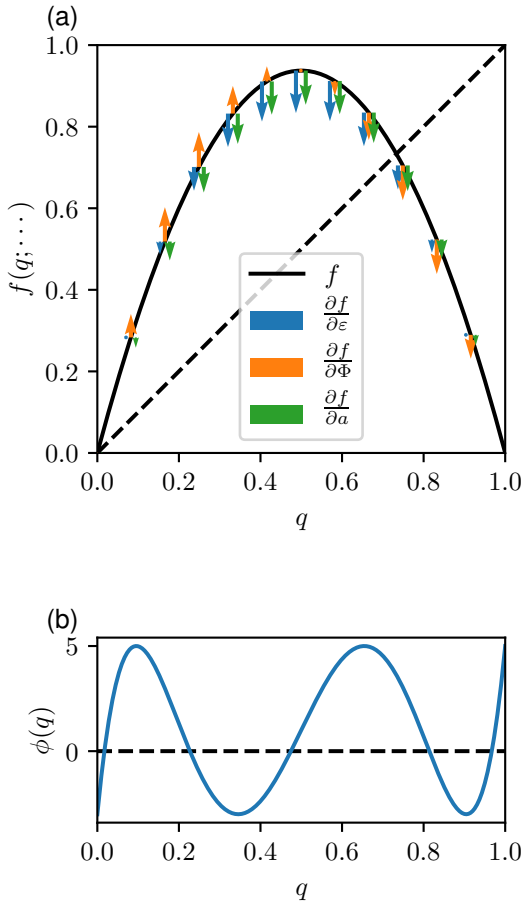

FIGURE 4.2. (a) Plot of $q_{n+1}^{(j)}$ under logistic dynamics (4.7) as a function of $q_n^{(j)}$ for $r_n^{(j)} > 1/2$, $a^{(j)} = 3.75$, $\Phi_n = 0$ and $\varepsilon = 0$. Effects of perturbations in ε, Φ and a on the map are indicated by arrows. (b) Plot of coupling function $\phi(q)$.

We will consider the case of microscopic dynamics which when viewed in isolation obeys LRT, such as uniformly expanding maps (the specific maps we will study are described in Section 4.7.2.1). We will also consider the case when the microscopic dynamics when viewed in isolation does not obey LRT. The simplest such system is the logistic map as established by Baladi and co-workers (Baladi and Smania, 2008, 2010; Baladi, 2014; Baladi et al., 2015; De Lima and Smania, 2018). We shall distinguish

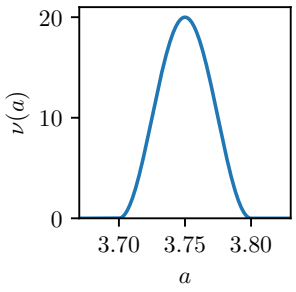

FIGURE 4.3. Probability density function $\nu(a)$ of the raised cosine distribution (4.8) with compact support on the interval $[3.7, 3.8]$.

two subcases here: one where the parameters of the logistic map are drawn from a smooth heterogeneous distribution and one where they are drawn from a non-smooth distribution. For concreteness, we consider perturbations of the following modified logistic map,

$$\left(q_{n+1}^{(j)}, r_{n+1}^{(j)}\right) = \begin{cases} \left(q_n^{(j)}, 2r_n^{(j)}\right) & r_n^{(j)} < \frac{1}{2} \\ \left(a^{(j)}\, q_n^{(j)}(1-q_n^{(j)}) + h(q_n^{(j)}, \Phi_n) + \varepsilon g(q_n^{(j)}), 2r_n^{(j)} - 1\right) & r_n^{(j)} \geq \frac{1}{2} \end{cases}, \tag{4.7}$$

where the logistic map parameters $a^{(j)}$ are sampled from a distribution $\nu(a)da$. The action of this map on q is plotted in Figure 4.2(a). We will introduce the external and coupling perturbation functions g and h in the respective following subsections, as well as the mean field observable ϕ.

The inclusion of the mixing doubling map dynamics r_n ensures that the overall dynamics is mixing even when the logistic parameters $a^{(j)}$ correspond to regular dynamics. The inclusion of the cocycle r_n, however, does not alter the invariant measure of the logistic map for constant Φ_n and the marginal invariant measure of $q^{(j)}$ the invariant measure of a logistic map at parameter $a^{(j)}$. Hence, notwithstanding any dynamics of Φ_n, the microscopic dynamics (4.7) violates LRT while being mixing.

In Section 4.5.2.2 we will see that the regularity of $\nu(a)$ is crucial in establishing LRT. We therefore consider here two cases: the case when $\nu(a)$ is smooth, in particular at least once-differentiable with respect to a, and the case when $\nu(a)$ is non-smooth, for example when $\nu(a)\mathrm{d}a$ is a linear combination of delta functions. We choose as a smooth distribution the raised cosine distribution supported on the interval $[3.7, 3.8]$, which is given by

$$\nu(a) = \mathbf{1}_{[3.7,3.8]} \frac{1}{0.1} \left(1 + \cos\left(\frac{a - 3.75}{0.05}\pi\right)\right). \tag{4.8}$$

We have chosen this distribution as it is both compactly supported and resembles a Gaussian distribution (see Figure 4.3). Furthermore, a large proportion of logistic parameters between $[3.7, 3.8]$ give chaotic dynamics (Galias, 2017). For a non-smooth distribution we choose the discrete distribution

$$\nu \mathrm{d}a = \frac{1}{3}(\delta_{3.72} + \delta_{3.75} + \delta_{3.78}), \tag{4.9}$$

which has a similar distribution of moments.

4.4. Summary of results

Our main results for these different dynamical scenarios and cases are summarised in Table 4.1. We differentiate between the thermodynamic limit $M = \infty$ and the case of a large, finite heat bath size M (which may not necessarily approach a smooth limit as $M \to \infty$). We summarise the dynamical mechanisms leading to the comprehensive picture provided in Table 4.1, which are to the best of our knowledge hitherto unknown. The following sections will establish these findings in detail.

- Macroscopic mean field observables generated by an appropriate heterogeneous set of microscopic chaotic systems may exhibit linear response, even if the individual members of those systems may not individually have LRT (Section 4.5.2.2).
- In the thermodynamic limit, macroscopic observables obey a law of large numbers. If the microscopic dynamics is mixing, this leads in the case of no back-coupling to trivial macroscopic dynamics (Section 4.5.1); if the microscopic dynamics is coupled via its mean field and provided the microscopic dynamics collectively obeys LRT, one can derive a smooth non-Markovian closure for macroscopic variables (Section 4.7.1). In the latter case, if the macroscopic dynamics converges to a fixed point or to a limit cycle, the macroscopic mean field observables satisfy LRT in the thermodynamic limit (Section 4.7.2).
- However, the reduced macroscopic dynamics may also converge to a chaotic dynamical system which is demonstrably non-hyperbolic (Section 4.8) and which violates LRT (Section 4.7.3). This is possible even if the individual microscopic dynamics is uniformly hyperbolic.
- In finite ensembles with $M < \infty$ the mean field involves an $\mathcal{O}(1/\sqrt{M})$ correction to the thermodynamic mean field dynamics, which may not obey LRT (Section 4.5.1). The possible violation of LRT of macroscopic observables, however, is not detectable for practical purposes, and the observed linear response is determined by the linear response property of the thermodynamic limit. We call this behaviour *approximate LRT*.

microscopic subsystem		macroscopic observable: uncoupled mean field	distinguished var. $\gamma = \frac{1}{2}$	distinguished var. $\gamma = 1$	self-coupled mean field
f satisfies LRT	finite M	✓	✓	✓	✓
	$M \to \infty$	✓	✓	✗	✯
f violates LRT with smooth $\nu(a)da$	finite M	(✓)	(✓)	(✓)	(✓)
	$M \to \infty$	✓	✓	✗	✯
f violates LRT with non-smooth $\nu(a)da$	finite M	✗	✗	✗	(✓)
	$M \to \infty$	✗	✗	✗	✗

Table 4.1. Summary of our main LRT results. The checkmarks ✓ denote cases when the macroscopic observable Ψ enjoys LRT. The bracketed checkmarks (✓) denote cases of approximate LRT, when LRT is satisfied for practical purposes. The cross-marks ✗ denote cases when LRT is violated for the macroscopic observable. The star ✯ denotes cases when LRT may or may not be satisfied depending on the linear response of the limiting dynamics of the macroscopic observable (see Section 4.7).

- In finite mean-field coupled systems such as (4.6), macroscopic mean fields typically satisfy a central limit theorem. As a result, the back-coupling of the mean field introduces a small "noise" into the microscopic systems, which can induce linear response. The statistical properties of this dynamic self-generated noise and its linear response properties are determined by the linear response property of the thermodynamic limit. In the case of failure of linear response in the thermodynamic limit, the convergence to the thermodynamic limit is approached for finite large M through the creation of saddle-node bifurcations (with associated multistability) which become increasingly dense in ε (Section 4.7.2.3).
- When the mean-field is used to drive a distinguished variable, it induces the same LRT properties as the noise/deterministic signal it approximates in a central limit theorem. Thus, for finite M the mean-field will induce LRT in the distinguished variable, provided its parameters vary smoothly with the perturbation (Section 4.6.2); furthermore if the mean-field has a diffusive coupling (i.e. normalised by $M^{-1/2}$ instead of M^{-1}) then the central limit theorem will continue to hold in the thermodynamic limit under the same conditions (Section 4.6.1).

In the following we provide numerical evidence and theoretical arguments corroborating these results. We first consider the case of macroscopic observables of an

uncoupled heat bath; we then extend our study first to observables of distinguished variables driven by such a heat bath, and then to the case of macroscopic observables of a mean field coupled heat bath.

4.5. Macroscopic observables of uncoupled microscopic subsystems

We are concerned with the behaviour of averages of the macroscopic observable Ψ. We distinguish here two averages; the average with respect to initial conditions of $q^{(j)}$, which we denote by $\mathbb{E}$, and the average over the independently chosen logistic map parameters distributed according to $\nu(a)$ which we denote by angular brackets $\langle\cdot\rangle$. In real systems (for which the parameters $a^{(j)}$ are selected once only), the average relevant for linear response is $\mathbb{E}$, the expectation with respect to initial conditions.

We describe a stochastic reduction of the mean field dynamics in Section 4.5.1 and then in Section 4.5.2 discuss the linear response properties for each of the three kinds of microscopic subsystems that we outlined in Section 4.3: in Table 4.1 these are covered in the rows corresponding to the uncoupled macroscopic observables.

4.5.1. Stochastic reduction of mean field dynamics. The average with respect to initial conditions is written as

$$\mathbb{E}\psi(q^{(j)}) = \int \psi(q)\,\mathrm{d}\mu^{a^{(j)}}(q),$$

where $\mu^{a^{(j)}}(q^{(j)})$ is the invariant measure of $q^{(j)}$. The Law of Large Numbers then reads as

$$\langle \mathbb{E}\Psi \rangle = \iint \psi(q)\,\mathrm{d}\mu^{a}(q) d\nu(a). \tag{4.10}$$

(In view of Section 4.7 where the mean field coupling is considered and the $q^{(j)}$ depend on a time-varying driver, we remark that in that case averages are computed with a time dependent measure $\mu_n^{a^{(j)}}(q^{(j)})$.)

We first establish the case of LRT for a finite heat bath. For large but finite system size M, both averages are equipped with their own finite size correction, described by the central limit theorem. In equilibrium each ensemble member $q_n^{(j)}$, at a given time n, is an independent sample from the invariant measure $\mu^{a^{(j)}}$. Macroscopic observables Ψ, as defined in (4.2), can be approximated using the central limit theorem and the independence of the $q^{(j)}$ by

$$\Psi_n = \mathbb{E}\Psi + \frac{1}{\sqrt{M}}\zeta_n + o(1/\sqrt{M}), \tag{4.11}$$

where the expectation value

$$\mathbb{E}\Psi = \frac{1}{M}\sum_{j=1}^{M}\int \psi(q)\,\mathrm{d}\mu^{a^{(j)}}(q)$$

is over initial conditions $q^{(j)}$ at fixed $a^{(j)}$. The random mean-zero Gaussian process ζ_n has autocovariance function $C^\zeta(m)$ with

$$\begin{aligned} C^\zeta(m) = \mathrm{cov}(\zeta_n,\zeta_{n+m}) &= \lim_{M\to\infty}\frac{1}{M}\sum_{j=1}^{M}\mathbb{E}[\psi_0^{(j)}\psi_m^{(j)}] \\ &= \langle\mathbb{E}[\psi_0\psi_m]\rangle. \end{aligned} \tag{4.12}$$

The stochastic process ζ_n can be thought of as a moving average process of infinite order, which is uniquely defined by its mean and its covariance $R(m)$.

The existence of a central limit theorem is guaranteed for unimodal maps using results of Lyubich (2002) who proved that almost every non-regular logistic parameter satisfies the so-called Collet-Eckmann condition (Collet and Eckmann, 1983), which then implies the existence of good statistical properties including the central limit theorem (Alves et al., 2004; Melbourne and Nicol, 2008). We remark that the parameters determining the process ζ_n in the finite-M case, such as the covariance (4.12), have the same LRT properties as the associated thermodynamic limit $\langle\mathbb{E}^\varepsilon\Psi\rangle$.

The independent sampling of the $a^{(j)}$ allows for a further application of the central limit theorem, and we can write

$$\mathbb{E}\Psi_n = \langle\mathbb{E}\Psi_n\rangle + \frac{1}{\sqrt{M}}\eta + o(1/\sqrt{M}), \tag{4.13}$$

where the random variable η is, for fixed ε, a mean-zero Gaussian variable. As a function of ε, η is a Gaussian process with covariance

$$\langle\eta^\varepsilon\eta^{\varepsilon'}\rangle = \langle\mathbb{E}^\varepsilon[\psi]\mathbb{E}^{\varepsilon'}[\psi]\rangle - \langle\mathbb{E}^\varepsilon[\psi]\rangle\langle\mathbb{E}^{\varepsilon'}[\psi]\rangle, \tag{4.14}$$

and typically is no more differentiable with respect to ε than $\mathbb{E}^\varepsilon[\psi]$, which implies that LRT is violated for finite M if the microscopic subsystems do not individually satisfy LRT. However, for finite $M \gg 1$ the response term $\langle\mathbb{E}\Psi_n\rangle$ dominates over the contribution of η and the violation of LRT can only be detected for vanishingly small values of ε. We call this instance of LRT for all practical purposes *approximate LRT*.

We remark that, notwithstanding the rough parameter selection error discussed above, and recalling that the linear response of the process ζ_n is determined by the linear response property of the associated thermodynamic limit, the overall linear

response of Ψ_n depends entirely on whether the thermodynamic limit $\langle \mathbb{E}\Psi_n \rangle$ satisfies LRT or not. We discuss this question in the next section.

4.5.2. LRT of thermodynamic limit mean field observables of uncoupled microscopic subsystems. We now investigate the response of $\langle \mathbb{E}\Psi_n \rangle$, i.e. the thermodynamic limit. We distinguish between three cases: when the microscopic dynamics satisfies LRT, and when the the microscopic dynamics does not satisfy LRT and has a distribution $\nu(a)$ of the parameters which is either smooth or non-smooth.

4.5.2.1. *The microscopic subsystems satisfy LRT.* If the microscopic dynamics obeys LRT, as is the case for uniformly expanding maps such as (4.43), which will be considered in Section 4.7.2.1, it is clear that LRT also holds for macroscopic observables defined in (4.2). For finite heat bath sizes M, we have

$$\frac{d}{d\varepsilon}\mathbb{E}^{\varepsilon}\Psi_n = \frac{1}{M}\sum_{j=1}^{M}\frac{d}{d\varepsilon}\mathbb{E}\psi(q_n^{(j)})$$

and the macroscopic observable Ψ obeys LRT since the M subsystems individually satisfy LRT with uniformly bounded $\frac{d}{d\varepsilon}\,\mathrm{d}\mu^{(a_j,\varepsilon)}$. The validity of LRT carries over to the thermodynamic limit with

$$\frac{d}{d\varepsilon}\langle \mathbb{E}^{\varepsilon}\Psi_n \rangle = \iint \psi^{\varepsilon}(q)\frac{d}{d\varepsilon}\,\mathrm{d}\mu^{(a,\varepsilon)}(q)\nu(a)da.$$

Note that we may allow for a ν-measure zero subset of subsystems at any given ε to individually violate LRT, and still obtain LRT for the macroscopic observable Ψ in the $M \to \infty$ limit. In this case, however, the η correction may not be differentiable, and we observe approximate LRT.

4.5.2.2. *The microscopic subsystems do not satisfy LRT but are appropriately heterogeneous.* To model microscopic dynamics that violate LRT we consider the modified logistic map (4.7). We will perturb the map in ε by the function (see (4.7))

$$g(q) = 4(q(1-q))^2, \tag{4.15}$$

and set $h \equiv 0$. The effect of g is plotted in Figure 4.2(a). We will draw parameters from the smooth raised cosine distribution on $[3.8, 3.9]$ (4.8), which is three-times continuously differentiable (i.e. $\nu(a)$ is C^3). We use $\psi(q, r) = q$ for our mean-field observable Ψ.

Figure 4.4 provides numerical evidence that, for these maps, the macroscopic observable Ψ with $\psi(x) = x$ has linear response for a wide range in ε. To determine the smoothness of $\langle \mathbb{E}^{\varepsilon}\Psi \rangle$, we determine its Chebyshev coefficients on a Chebyshev roots grid of 1000 points, using the Chebyshev transform routines in the Julia package ApproxFun (Olver, 2019). It is well known that any smooth function $f(x)$ can be expressed as an infinite series of Chebyshev polynomials $T_k(x)$ as $f(x) = \sum_{k=0}^{\infty} f_k T_k(x)$

and the degree of differentiability of the function is given by the decay of its Chebyshev coefficients f_k (Trefethen, 2013). We find that the Chebyshev coefficients of $\langle \mathbb{E}^\varepsilon \Psi \rangle$ decay as $\mathcal{O}(k^{-4})$, which is indicative of $\langle \mathbb{E}^\varepsilon \Psi \rangle$ being between C^{3^-} and C^4 differentiable over a large interval: this level of differentiability, as we will discuss below, is connected to the smoothness of the raised-cosine distribution (4.8), which is C^3. We have also employed the test statistics for higher-order linear response developed in Gottwald et al. (2016) to test the null-hypothesis of $\langle \mathbb{E}^\varepsilon \Psi \rangle$ being well-approximated by a linear combination of $T_k(0.1^{-1}(\varepsilon + 0.1)), k = 0, \ldots, 60$ for $\varepsilon \in [-0.2, 0]$, i.e. that the response is in fact a smooth function. This test is summarised in Appendix C. We used the aforementioned Chebyshev grid simulating $1,000,000$ different randomly selected parameters with 50 runs of 3000 timesteps each, and quantified Birkhoff variance both within parameters and between parameters (i.e. arising from the random parameter selection), using standard ANOVA methods (Rice, 2006). We obtain a p-value of 0.26, consistent with the null-hypothesis of a smooth response.

We note that in Figure 4.4 the response of $\mathbb{E}^\varepsilon \Psi$ for systems with finite M have (barely) noticeable rough deviations from the $M \to \infty$ limit: these non-smooth deviations arise from the finite sampling of parameters $a^{(j)}$ from $\nu(a)$, approximated by the random variable η defined in (4.13), as discussed above.

We now provide a heuristic argument how averaging over a smooth distribution such as the raised cosine distribution (4.8) can lead to LRT for the macroscopic observable Ψ. Let us first recall the dynamic reason of why LRT is violated in the logistic map. We follow here Ruelle (2009b) in our exposition. The critical point $q = c$ with $f'(c) = 0$ leads to a non-uniform compression of the phase space around $q = c$: an initially smooth initial density which contains the critical point in its support is pushed forward under the dynamics to a non-smooth density with a spike with an inverse square-root singularity at $q = f(c)$. This compression is repeated to produce further inverse square-root singularities at locations $q_n = f^n(c)$ of amplitudes asymptotically proportional to $\alpha^{-n/2}$ (and thus contain a probability mass of order $\alpha^{n/2}$), where $1 < \alpha$ denotes the Lyapunov multiplier of the logistic map. The result is that the invariant density contains an infinite number of spikes of decreasing amplitude. The effect of the perturbation, by modifying the forward orbit of the critical point $(f^n(c))_{n \in \mathbb{N}}$, is to displace these spikes. Because the map f is chaotic and thus exponentially sensitive to perturbations, spikes move with an instantaneous speed of the order of α^n per unit change of the perturbation. This scenario is illustrated in Figure 4.5(a) where we show the absolutely continuous invariant measure (acim), averaged over the heat bath, of an perturbed and of a slightly perturbed dynamics. The high speed of the small spikes (i.e. those with large n) in conjunction with their relatively large probability mass implies that the

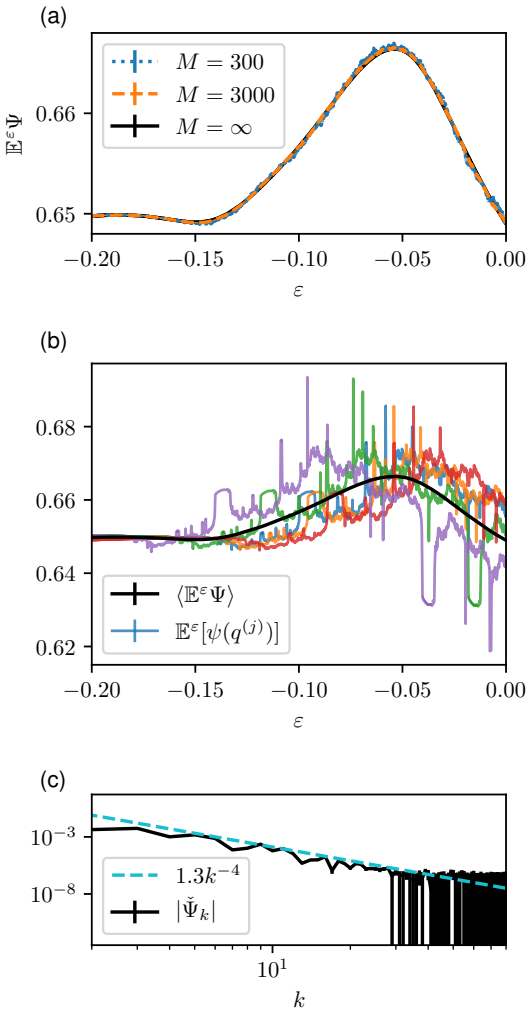

FIGURE 4.4. Response term $\mathbb{E}^\varepsilon\Psi$ for a perturbation of the form (4.15) for an uncoupled heat bath scenario in the case when the microscopic dynamics is given by the logistic map (4.7), which does not satisfy LRT, and the logistic map parameters are sampled from the raised cosine distribution (4.8). For different values of ε we employ a total of 10^5 iterates to estimate $\mathbb{E}^\varepsilon\Psi$ as a temporal average. (a) Plots for finite M: 95% confidence intervals were estimated from 10 realisations differing in the initial conditions of the heat bath, and are not visible. Thermodynamic limit curve (black), confidence intervals also not visible, was estimated from 50 realisations of 3000 iterates for 10^6 parameters $a^{(j)}$ independently selected for each ε. (b) Thermodynamic limit of $\mathbb{E}^\varepsilon\Psi$ (black), with LRT-violating response of microscopic variables $\mathbb{E}^\varepsilon\psi(q^{(j)})$ (coloured lines), estimated from 10 realisations of 10^6 iterates each. (c) Estimate of Chebyshev coefficients $\sum_{k=0}^{\infty}\check{\Psi}_k T_k(0.1^{-1}(x+0.1)) := \langle\mathbb{E}^\varepsilon\Psi\rangle$.

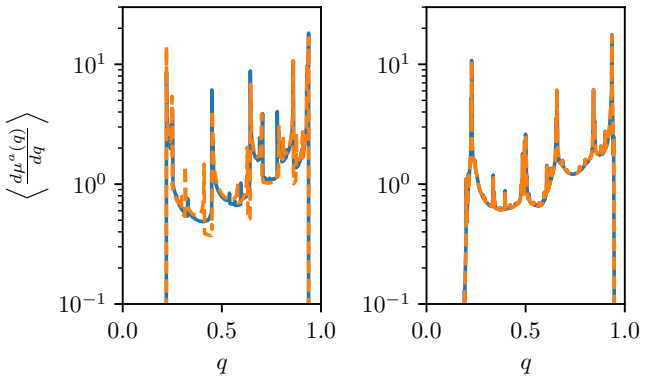

FIGURE 4.5. Histogram of the averaged acim of the $q^{(j)}$ for the logistic map system (4.7) with (a) $\nu = \delta_{3.75}$ and (b) ν the raised-cosine distribution (4.8), for $\varepsilon = -5 \times 10^{-4}$ (orange dashes) and $\varepsilon = 0$ (blue line).

sum of their (distributional) derivatives diverges, leading to breakdown of linear response. The reader is referred to Ruelle (2009b) for details and to Gottwald et al. (2016) for a numerical illustration. For comparison we also depict in Figure 4.5(b) the averaged invariant measure for the case when the parameters are heterogeneously drawn from a raised-cosine distribution and for which we showed above that LRT is valid. The averaging over the heterogenously drawn parameters clearly implies a smoothed invariant measure of the logistic map.

This may be seen as analogous to a recent heuristic argument for linear response in general non-hyperbolic systems (Ruelle, 2018): rough contributions to the response caused by singularities in the physical measures that arise from stable manifold-unstable manifold tangencies average out if these singularities distribute themselves suitably evenly.

We may understand more concretely the differentiability of the response, and in particular its relationship with the differentiability of the parameter distribution ν, by considering the simplified, more specific case where ε is an additive perturbation of the inhomogeneous microscopic map parameters $a^{(j)}$, i.e. $f(q; a, \varepsilon)$ can be written as $f(q; a + \varepsilon)$, for example as in the system we will consider in Section 4.6 in the distiguished variable scenario. When ε is an additive perturbation of the parameters

$a^{(j)}$, it is possible to apply a change of variables and write

$$\begin{aligned}\langle \mathbb{E}^{\varepsilon}\Psi\rangle &= \iint \psi(q)\,\mathrm{d}\mu^{a+\varepsilon}(q)\nu(a)\,\mathrm{d}a \\ &= \iint \psi(q)\,\mathrm{d}\mu^{a}(q)\nu(a-\varepsilon)\,\mathrm{d}a. \end{aligned} \tag{4.16}$$

The linear response term is then readily evaluated as

$$\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\langle \mathbb{E}^{\varepsilon}\Psi\rangle = -\iint \psi(q)\,\mathrm{d}\mu_n^a(q)\frac{\mathrm{d}}{\mathrm{d}a}\nu(a)\,\mathrm{d}a,$$

which implies that LRT is valid provided that the system is appropriately heterogeneous with integrable density $\mathrm{d}\nu(a)/\mathrm{d}a$, i.e. $\nu(a)$ is of bounded variation. Furthermore, from the integral formulation of the expectation in (4.16) it is clear that the regularity of the response depends on the regularity of ν: indeed, to achieve higher-order response, say of order ℓ, ν must have (weak) derivatives of order ℓ. This can be achieved if a is drawn from a distribution ν in Sobolev space $W^{\ell,1}$. Thus, we would expect to see three-times differentiable response of Ψ with a distribution $\nu \in W^{3,1}$ such as the raised-cosine distribution.

This particular argument is completely independent of the actual microscopic dynamics, and it generalises to systems beyond logistic maps—provided the maps can be written $f(\cdot; a, \varepsilon) = f(\cdot; a + \varepsilon)$.

In the specific case where the microscopic dynamics evolves under unimodal maps such as those studied here, we can make a more concrete argument for the effect of the smearing out of the fast displacement of the small spikes upon perturbation. It is conjectured by Avila et al. (2003) that there exists an ε-dependent analytic function $\alpha(a, \varepsilon)$ of the invariant measures, such that the map with parameters (a, ε) is topologically conjugate to the map with parameters $(\alpha(a, \varepsilon), 0)$. Unimodal maps, at least those of Benedicks-Carleson type, have linear response within topological conjugacy classes (Baladi and Smania, 2012), and as a result we can say

$$\begin{aligned}\langle \mathbb{E}^{\varepsilon}\Psi\rangle &= \iint \psi(q)\,\mathrm{d}\mu_n^{a,\varepsilon}\nu(a)\,\mathrm{d}a \\ &= \iint \psi(q)\,\mathrm{d}\mu_n^{\alpha(a,\varepsilon),0}\nu(a)\,\mathrm{d}a + \text{h.o.t.} \\ &= \iint \psi(q)\,\mathrm{d}\mu_n^{\alpha,0}\nu(a(\alpha,\varepsilon))\frac{\mathrm{d}a}{\mathrm{d}\alpha}\,\mathrm{d}\alpha + \text{h.o.t.}, \end{aligned} \tag{4.17}$$

where the higher order terms capture the response of the μ_n under the topological conjugacy-preserving parameter change from $(\alpha(a, \varepsilon), 0)$ to (a, ε), and are thus similarly smooth. Hence the existence of linear response $d\langle \mathbb{E}^{\varepsilon}\Psi\rangle/d\varepsilon$ (and by a similar argument higher-order response) is guaranteed, provided that the distribution of the

parameters of the logistic map $\nu(a)$ is at least once continuously differentiable (and provided $\alpha(a,\varepsilon)$ is analytic in a sufficiently uniform way (Avila et al., 2003)). The linear response of $\langle \mathbb{E}^\varepsilon \Psi \rangle$ was numerically confirmed in Figure 4.4. Since in at least one-dimensional systems topological conjugacy classes form manifolds of finite codimension, we believe this argument will generalise to more general maps, provided the space of parameters $a^{(j)}$ is sufficiently high-dimensional.

However, we caution that our smearing argument may not generalise to other systems, at least when the support of the parameters a is one-dimensional. In Figure 4.6 we present numerical evidence demonstrating that mean-field averaging fails to improve the linear response of a unimodal map of the torus $(\mathbb{R}/\mathbb{Z})^2$ for heterogeneously distributed parameters, given by

$$x_{n+1}^{(j)} = x_n^{(j)} + a^{(j)} y \sin \pi x_n^{(j)} \mod 1 \tag{4.18}$$

$$y_{n+1}^{(j)} = y_n^{(j)} + a^{(j)} \sin \pi (x_n^{(j)} + y_n^{(j)}) + \varepsilon \mod 1, \tag{4.19}$$

for $j = 1, \ldots, M$. The parameters $a^{(j)}$ are again distributed according to a raised cosine distribution with support on $[3.7, 4.3]$

$$\nu(a) = \mathbf{1}_{[3.7,4.3]} \frac{1}{0.1} \left(1 + \cos \left(\frac{a-3}{0.3} \pi \right) \right), \tag{4.20}$$

and the mean-field observable is given by $\Psi = \frac{1}{M} \sum_{j=1}^{M} \psi(x^{(j)}, y^{(j)})$ where $\psi(x,y) = x$, as before. We tested the null-hypothesis of $\langle \mathbb{E}^\varepsilon \Psi \rangle$ being well-approximated by a linear combination of $T_k(0.05^{-1}(\varepsilon - 0.05)), k = 0, \ldots, 100$ for $\varepsilon \in [0, 0.1]$, and obtained a p-value $p = 0.49$, consistent with the null hypothesis. However in Figure 4.6 we see that the estimated Chebyshev coefficients decay approximately as $\mathcal{O}(k^{-1.5})$ which is slower than $\mathcal{O}(k^{-4})$ seen in the one-dimensional unimodal example, indicating a rather low-order differentiability. This level corresponds quite closely to that obtained for the expectation value $\mathbb{E}^\varepsilon \psi(x^{(j)}, y^{(j)})$ of a single microscopic systems, as illustrated in Figure 4.7; hence the averaging over parameters appears only to be smoothing out the large jumps arising from periodic windows but does not improve the degree of differentiability.

4.5.2.3. *The microscopic subsystems do not satisfy LRT and are not appropriately heterogeneous.* If the microscopic dynamics does not obey LRT and the logistic map parameters are non-smoothly distributed, then LRT fails for macroscopic observables (4.2), independent of whether the heat bath is finite or infinite. In this case the averaging over the heat bath variables does not provide the necessary smearing of the non-smoothness of the perturbed invariant measures $\mu^{a,\varepsilon}$. To illustrate this,

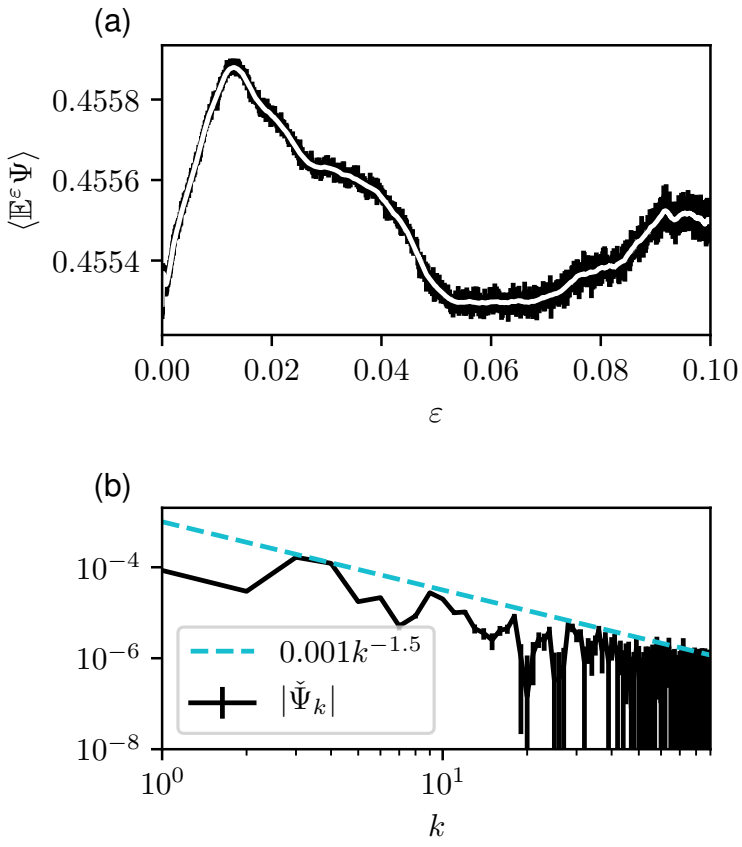

FIGURE 4.6. Response $\mathbb{E}^\varepsilon \Psi$ for an uncoupled heat bath scenario for the map (4.19) where the parameters are sampled from a raised cosine distribution (4.20). (a) Infinite M limit with confidence intervals (black) and 21-point moving average with confidence intervals (white) from 15 realisations of 6000 iterates for 10^6 parameters $a^{(j)}$ independently selected for each ε. (b) Estimate of Chebyshev coefficients $\sum_{k=0}^{\infty} \check{\Psi}_k T_k(0.05^{-1}(x - 0.05)) := \langle \mathbb{E}^\varepsilon \Psi \rangle$.

consider the following non-smooth parameter distribution

$$\nu(a) = \sum_{k=1}^{p} w_k \delta(a - a_k), \tag{4.21}$$

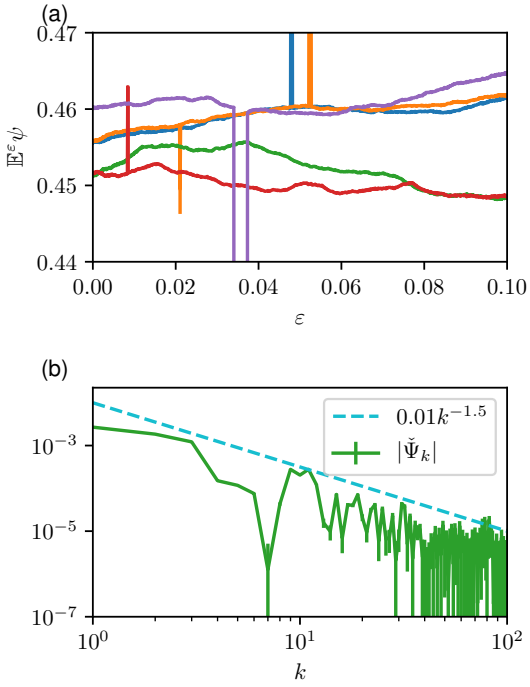

FIGURE 4.7. Individual response terms $\mathbb{E}^\varepsilon\psi(x^{(j)}, y^{(j)})$ with confidence intervals for the map (4.19) where the parameters selected from the raised cosine distribution (4.20). (a) Response for five randomly selected microscopic variables. The large jumps of the response outside the figure correspond to regions of regular dynamics. (b) Estimate of Chebyshev coefficients $\sum_{k=0}^{\infty}\check{\Psi}_k T_k(0.05^{-1}(x+-0.05)) := \langle\mathbb{E}^\varepsilon\Psi\rangle$ for one of the variables in (a).

where at least one of the logistic map parameters a_k corresponds to chaotic dynamics. The invariant measures $\mu^{a_k,\varepsilon}$ for fixed parameters a_k are not differentiable with respect to the perturbation size ε per assumption. The averaging over the independent heat bath variables only involves finitely many logistic parameter values, and hence in this situation averaging is not able to smear the effect of the non-smoothness of the finite number of associated invariant measures μ^{a_j}. In Figure 4.8 we show the response $\mathbb{E}^\varepsilon\Psi$ for microscopic logistic variables as in Section 4.5.2.2 with ν a discrete distribution. As expected, the response exhibits non-smooth behaviour upon varying the strength of the perturbation ε. The response term $\mathbb{E}^\varepsilon\Psi$ quickly converges as $M \to \infty$, and for $M = 300$ is almost indistinguishable by eye from the response $\langle\mathbb{E}^\varepsilon\Psi\rangle$ in the thermodynamic limit.

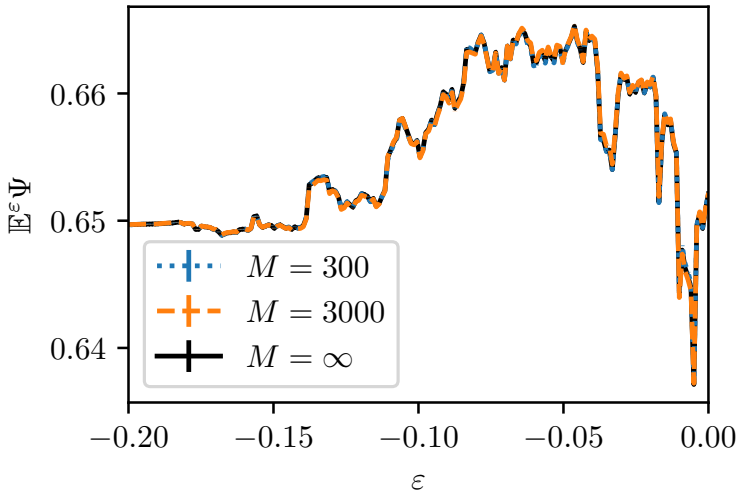

FIGURE 4.8. Response $\mathbb{E}^\varepsilon\Psi$ for a perturbation of the form (4.15) for an uncoupled heat bath scenario in the case when the microscopic dynamics is given by the logistic map (4.7), which does not satisfy LRT, and the logistic map parameters are distributed as in (4.21) with $\nu = \frac{1}{3}(\delta_{3.72}+\delta_{3.75}+\delta_{3.78})$. Error bars were estimated from 10 realisations differing in the initial conditions of the heat bath, and are not visible.

4.6. Linear response of macroscopic observables of distinguished variables driven by microscopic subsystems

We now extend the uncoupled case by considering a distinguished particle Q driven by a mean-field Φ as in (4.3). This mean-field may either have a diffusive scaling (where $\gamma = \frac{1}{2}$ in (4.3)), or a deterministic scaling (where $\gamma = 1$). In this section we in turn derive reductions to macroscopic dynamics for these two different scalings. As we have made a comprehensive study of the LRT of the driving mean-field in the previous section, we will here explicitly study LRT properties only for most unfavourable situation where both microscopic and macroscopic variables have logistic dynamics: for other microscopic variables the situation is given in Table 4.1. We find that the different scalings lead to qualitatively different dynamics of the distinguished variables, and consequently the outlook for LRT is different.

We recall that the dynamics of the microscopic heat bath are given from (4.1,4.3,4.4) by

$$q_{n+1}^{(j)} = f(q_n^{(j)}; a^{(j)}, \varepsilon), \tag{4.22}$$

$$\Phi_n = \frac{1}{M^\gamma} \sum_{j=1}^{M} \phi(q_n^{(j)}; a, \varepsilon), \tag{4.23}$$

$$Q_{n+1} = \tilde{f}(Q_n, \Phi_n). \tag{4.24}$$

We choose f to be the modified logistic maps (4.7), this time with $g(q) = q(1-q)$. Note that this means that we can write

$$f(q, r; a, \varepsilon) = f(q, r; a + \varepsilon).$$

We will select the parameter distribution $\nu(a)$ to be the raised cosine distribution as in Figure 4.3, this time supported on $a \in [3.8, 3.9]$.

Furthermore, so as to consider the LRT "worst case" for the macroscopic dynamics, we will choose the macroscopic variable also to be a logistic map:

$$\tilde{f}(Q; \Phi) = (A_0 + A_1\Phi)Q(1-Q). \tag{4.25}$$

where A_0, A_1 are constants that will be specified in the presentation of each numerical experiment. We will in our numerical simulations use the macroscopic observable $\Psi(Q) = Q$.

4.6.1. $\gamma = \frac{1}{2}$: Weak coupling with diffusive limit. We begin by showing that for each ε the macroscopic variable Q asymptotically satisfies a stochastic limit system in the thermodynamic limit $M \to \infty$ when $\gamma = \frac{1}{2}$. We consider driving terms Φ_n with mean-zero functions $\phi(\cdot; a)$, $\mathbb{E}^\varepsilon[\phi(\cdot; a)] = 0$ for all a, where the $\mathbb{E}$ average, as in the previous section, is with respect to the invariant measure of the unresolved microscopic variable for a single value of a. The driving term Φ_n (4.23) again contains a sum over independent identically distributed random variables for each time n, which are mean-zero by stipulation. Hence, for $\gamma = \frac{1}{2}$, the central limit theorem assures that the driving term Φ_n converges to a random Gaussian variable $\zeta_n \sim \mathcal{N}(0, \sigma^2)$ with $\sigma^2 = \langle \mathbb{E}^\varepsilon[\phi(q^{(j)})^2] \rangle$, where the angular brackets again denote the average over the measure of the logistic map parameters $\nu(a)\mathrm{d}a$. As in Section 4.5.1, the ζ_n define a stationary mean-zero Gaussian stochastic process, whose covariance is readily determined, similarly to (4.12), as

$$R(m) = \mathrm{cov}(\zeta_n, \zeta_{n+m}) = \langle \mathbb{E}[\phi_0 \phi_m] \rangle. \tag{4.26}$$

The process Q_n hence converges weakly to the stochastic process defined by

$$Q_{n+1} = \tilde{f}(Q_n; \zeta_n) \tag{4.27}$$

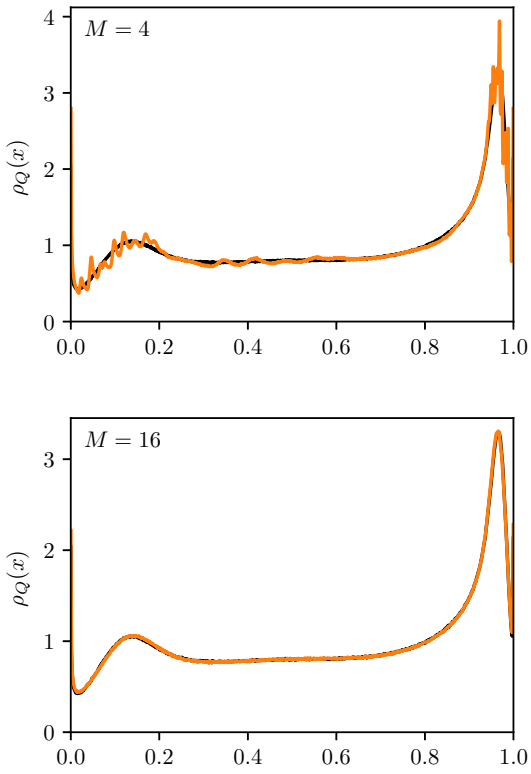

FIGURE 4.9. Empirical probability density $\rho_Q(x)$ (orange) of the distinguished variable Q for $\gamma = \frac{1}{2}$ as estimated from simulations of the original deterministic system (4.22-4.24) for different values of the size M of the microscopic sub-system. Top: $M = 4$. Bottom: $M = 16$. The continuous black line depicts the invariant density of the stochastic limit system (4.27). We used $A_0 = 3.91$, $A_1 = 0.05$ and $\varepsilon = 0$.

which for the logistic dynamics we consider (4.25) are

$$Q_{n+1} = (A_0 + A_1\zeta_n)\, Q_n(1 - Q_n). \tag{4.28}$$

Figure 4.9 illustrates the convergence of the full dynamics (4.22)-(4.24) to the stochastic limit system (4.27) in distribution by comparing the respective empirical measures for several values M of the size of the microscopic sub-system. The microscopic dynamics is run unperturbed with $\varepsilon = 0$. Here we chose the mean-zero (conditional on the parameter a) functions $\phi(x; a, \varepsilon) = x^2 - f(x; a, \varepsilon)^2$ to generate the driving sum Φ_n. We used a time series of $N = 4 \times 10^7$ and determined the empirical measure of the full system (4.22)-(4.24) by binning using 1000 bins. Details

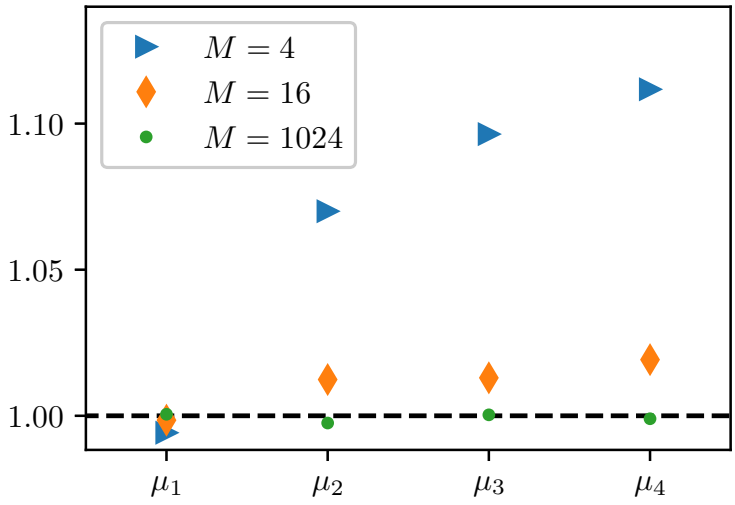

Figure 4.10. First four centred moments μ_i, $i = 1, \cdots, 4$ (not to be confused with measures μ discussed elsewhere), of the distinguished variable Q for $\gamma = \frac{1}{2}$ as estimated from simulations of the original deterministic system (4.22-4.24) for fixed time $n = 6$ for several values of the size of the microscopic sub-system: $M = 4$ (blue triangles), $M = 16$ (orange diamonds) and $M = 1024$ (green dots). We depict the moments scaled by the respective moments of the stochastic limit system (4.27) so that the asymptotic limit is 1 for all moments. Parameters as in Figure 4.9.

on how to determine the statistics of the limiting diffusive system (4.28) are given in Appendix D1. It is remarkable that with only $M = 16$ microscopic variables the eye can barely distinguish the empirical density from the density of the diffusive limit equation (4.27). We further show convergence of the first four moments of Q when increasing M in Figure 4.10. It is seen that for accurate convergence of higher order moments to the values of their stochastic limiting equation (4.27) larger system sizes M are required.

We present in Figure 4.11 results of the linear response for an observable $\Psi(Q) = Q$. The microscopic sub-system is perturbed homogeneously with $a_1^{(j)} = 1$ for all j. It is clearly seen that the perturbation ε induces a smooth change in the observable for large M, indicative of the validity of LRT. We employ here the test for linear response developed in Gottwald et al. (2016), which we describe in Appendix C and report the p-values testing the null hypothesis of linear response. We compute averages for

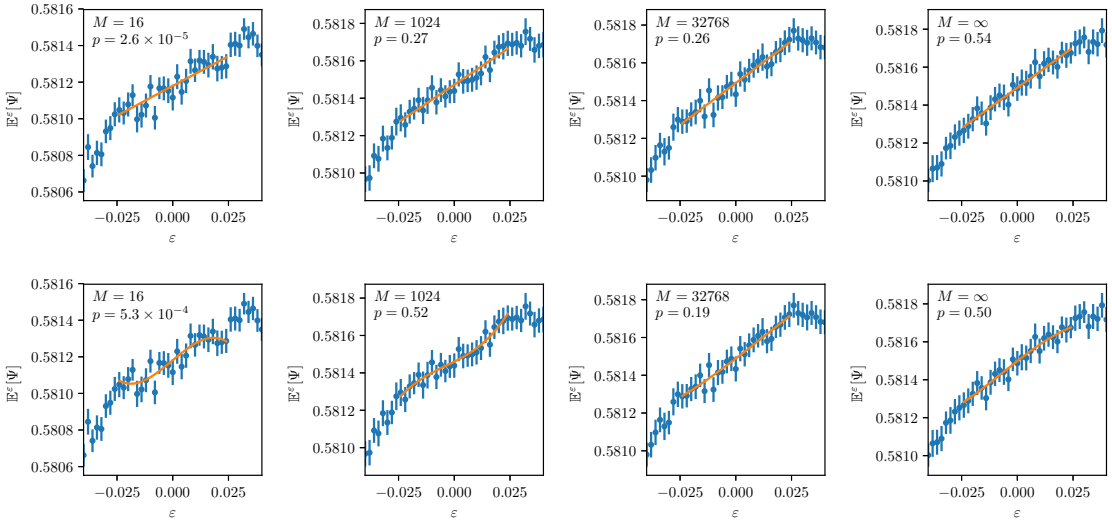

Figure 4.11. Top row: Linear fit of response of an observable $\Psi(Q) = Q$ for the deterministic system (4.22-4.24) for $\gamma = \frac{1}{2}$ for different values of the size M of the microscopic sub-system. (a): $M = 16$. (b): $M = 1024$. (c) $M = 32768$. (d): Stochastic limit system (4.27). Bottom row: Cubic fit of the same. All experiments used a time series of length $N = 2 \times 10^5$. The error bars were estimated from $K = 200$ realizations differing in the initial conditions. We used $A_0 = 3.91$, $A_1 = 0.05$.

several values of ε from long simulations of length $N = 5 \times 10^6$. The error bars shown in Figure 4.11 are estimated from $K = 200$ realizations differing in the initial conditions of the microscopic variables. For small values of $M = 16$ the p-value is $\mathcal{O}(10^{-5})$, rejecting the null hypothesis of linear response, whereas for $M = 2^{10}$ the p-value is 0.27, consistent with linear response. We also show results of the linear response for the stochastic limit system (4.27), illustrating that the thermodynamic limit implies linear response with a p-value of $p = 0.54$. Note that although the invariant density of the resolved degree of freedom Q has sufficiently converged to the invariant density of the stochastic limit system (4.27) for $M = 16$ (cf. Figure 4.9), this size is not sufficiently large to assure linear response. However, as the sample size increases, we see convergence to a curve that from statistical tests is consistent with a C^3 response (see Figure 4.11): from the discussion in Section 4.5.2.2, this is because the raised cosine distribution ν lies in $W^{3,1}$. The induced linear response here can be understood as arising from the convergence to the stochastic system

(4.28): dynamical systems driven by Gaussian noise are known to have LRT (Hairer and Majda, 2010).

4.6.2. $\gamma = 1$: Weak coupling with deterministic limit. We begin again for the $\gamma = 1$ deterministic scaling case by considering the stochastic reduction of the distinguished variable dynamics in the large M limit. We will let the driving term Φ_n be generated by a function ϕ with non-vanishing mean: for concreteness, we will consider $\phi(q) = q^2$. Since each unresolved degree of freedom generates an invariant measure, for $\gamma = 1$ the driving variable Φ_n converges to a constant according to the law of large numbers with $\Phi_n \to \bar{\Phi} = \langle \mathbb{E}[\phi] \rangle$. In the thermodynamic limit therefore the limiting equation of (4.24) is a deterministic map

$$Q_{n+1} = F(Q; \bar{\Phi}) \tag{4.29}$$

which, for the macroscopic dynamics we consider, is

$$Q_{n+1} = (A_0 + A_1 \bar{\Phi}) Q_n (1 - Q_n). \tag{4.30}$$

Figure 4.12 illustrates the convergence of the invariant measure of the deterministic map (4.22)-(4.24) to the averaged deterministic limit system (4.29) in distribution upon increasing the size M of the microscopic sub-system. We used again a time series of $N = 4 \times 10^7$ and determined the empirical measure by binning using 1000 bins. We see that for $M = 1024$ convergence to the rough limiting invariant measure of the deterministic logistic map (4.30) with its narrow peaks has not been fully achieved. This is due to finite sample size M. From Section 4.5 we have up to $\mathcal{O}(1/\sqrt{M})$ that

$$\Phi_n = \langle \mathbb{E}[\phi] \rangle + \frac{1}{\sqrt{M}} \eta^\varepsilon + \frac{1}{\sqrt{M}} \zeta_n, \tag{4.31}$$

where ζ_n is the mean-zero Gaussian process with covariance matrix (4.26), and η^ε is a mean-zero Gaussian random variable which, as in Section 4.5.1, is a Gaussian *process* in ε with covariance function

$$\langle \eta^\varepsilon \eta^{\varepsilon'} \rangle = \langle \mathbb{E}^\varepsilon[\phi] \mathbb{E}^{\varepsilon'}[\phi] \rangle - \langle \mathbb{E}^\varepsilon[\phi] \rangle \langle \mathbb{E}^{\varepsilon'}[\phi] \rangle.$$

In general, η^ε is non-differentiable in ε, which implies that LRT is violated for macroscopic observables $\Psi(Q)$, even for the random finite-size driver Φ_n given by (4.31). However, if the variation in $\mathbb{E}[\phi]$ over the parameter values sampled by ν is small by comparison with the typical variance $R(0) = \mathbb{E}[(\phi - \mathbb{E}[\phi])^2]$ for these parameters (e.g. if the support of ν is sufficiently small), then the small, rough contribution of $\frac{1}{\sqrt{M}} \eta^\varepsilon$ to the response of $\Psi(Q)$ is dominated by the (linear) response generated by $\langle \mathbb{E}[\phi] \rangle + \frac{1}{\sqrt{M}} \zeta_n$. We remark, however, that if the support of ν is too small and the parameters are therefore less heterogeneous, LRT is only valid for a small range of perturbation sizes ε.

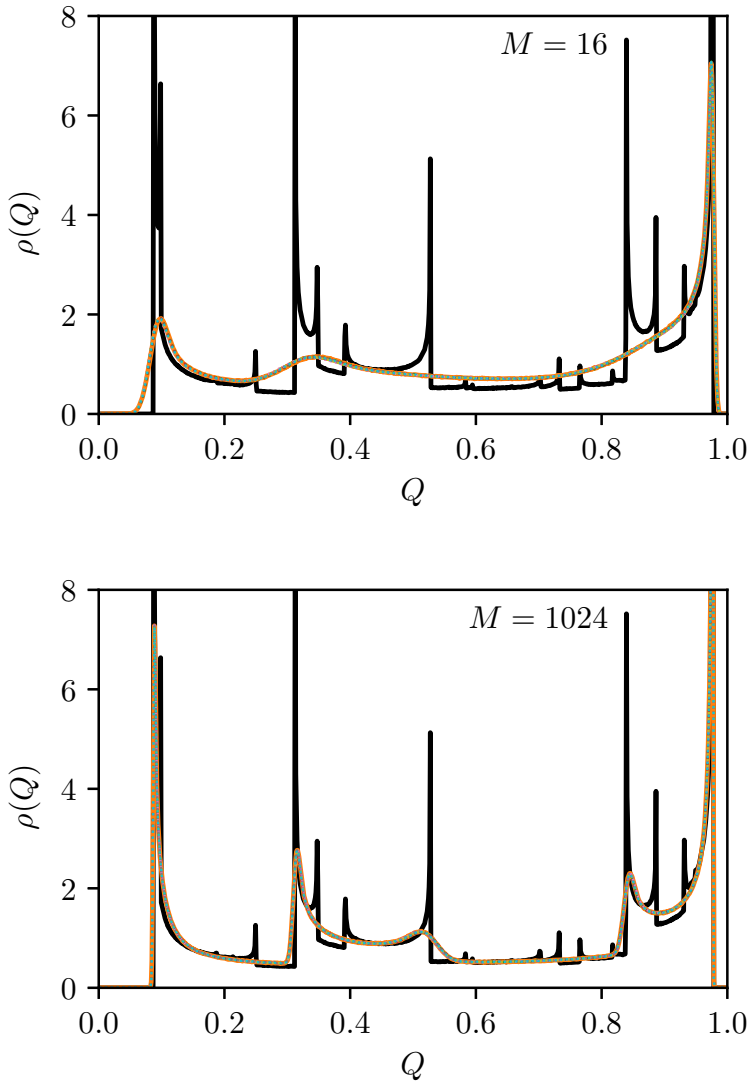

Figure 4.12. Empirical probability density $\rho_Q(x)$ (orange line) of the macroscopic variable Q for $\gamma = 1$ as estimated from simulations of the original deterministic system (4.22-4.24) for different values of the size M of the microscopic sub-system. Top: $M = 16$. Bottom: $M = 1024$. The continuous black line depicts the invariant density of the deterministic logistic map limit system (4.27); the thin dotted lines, which are indistinguishable from $\rho_Q(x)$, represent invariant densities of the logistic map (4.24) with the stochastic driving Φ_n given by (4.31) for various realisations of η^ε. We used $A_0 = 3.847$, $A_1 = 0.147$ and $\varepsilon = 0$.

To illustrate the role of finite size effects, we present in Figure 4.12 also results of simulations of the logistic map (4.25) with Φ_n stochastically generated by (4.31), mimicking random finite size effects in approximating the deterministic limit $\Phi_n = \langle \mathbb{E}[\phi] \rangle$. It is seen that for finite M the peaks are smoothed by sampling noise, and the random logistic map reproduces the invariant density of the macroscopic variable Q of the full deterministic model driven by the microscopic dynamics.

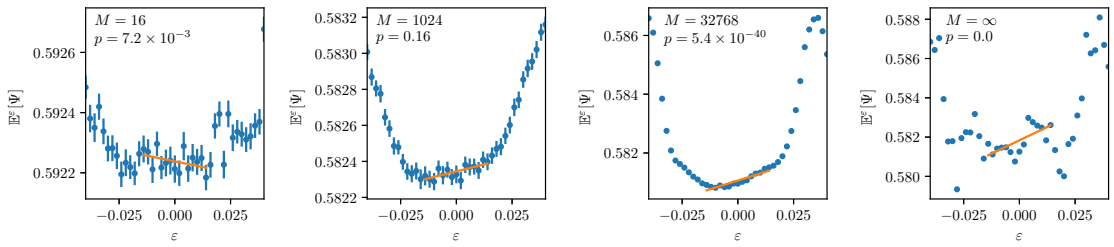

FIGURE 4.13. Linear fit of response of an observable $\Psi(Q) = Q$ for the deterministic system (4.22-4.24) for $\gamma = 1$ for different values of the size M of the microscopic sub-system. (a): $M = 16$. (b): $M = 1024$. (c) $M = 32768$. (d): Deterministic limit system (4.29). All experiments used a time series of length $N = 2 \times 10^5$. The error bars were estimated from $K = 200$ realizations differing in the initial conditions. We used $A_0 = 3.847$ and $A_1 = 0.147$.

Given that the thermodynamic limit system is deterministic, one might be tempted to conclude that linear response is not valid. Figure 4.13 shows the linear response as a function of perturbation ε for several values of the microscopic sub-system size M. For small values of M LRT is clearly violated with a p-value of $\mathcal{O}(10^{-3})$, as expected. For very large values of $M = 2^{15}$ LRT is violated with a p-value of $\mathcal{O}(10^{-40})$, consistent with the LRT-violating deterministic limit system (4.29). Remarkably and maybe surprisingly, decreasing the size M from these very large values to intermediate values of $M = 1024$ we observe that the numerical results are consistent with LRT and the p-value increases dramatically to around 0.16. This can be explained by the finite size corrections (4.31) to the deterministic limit $\Phi_n = \langle \mathbb{E}[\phi] \rangle$ provided by the central limit theorem. We note that the p-value for $M = 1024$ indicates marginal evidence in favour of breakdown of LRT associated with the (small) contribution of the non-differentiable η^ε term. Just as in the $\gamma = \frac{1}{2}$ case it is necessary for LRT to hold in the case of finite sample size, that the parameters $a^{(j)}$ are inhomogeneously distributed with a sufficiently smooth distribution $\nu(a)$.

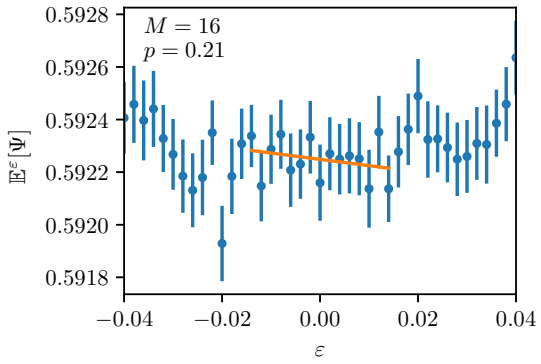

FIGURE 4.14. Linear response of an observable $\Psi(Q) = Q$ for the deterministic system (4.22-4.24) for $\gamma = 1$ with $M = 16$ estimated from a time series of length $N = 2\times10^4$. The error bars were estimated from $K = 200$ realizations differing in the initial conditions. We used $A_0 = 3.847$ and $A_1 = 0.147$.

In Gottwald et al. (2016) it was found that even if a system does not obey linear response one might not be able to reject the null hypothesis of linear response with sufficient statistical significance when the data length N of the time series is not sufficiently long. In Figure 4.14 we show the linear response as a function of ε for a microscopic sub-system of size $M = 16$ for $N = 2 \times 10^4$. While for $N = 2 \times 10^5$ linear response was rejected with $p = 7.2 \times 10^{-3}$, linear response is now consistent with the given data with a p-value of $p = 0.21$. It is found that decreasing the length of the time series allows for a larger range in the perturbation size ε for which linear response is consistent with the data.

4.7. Linear response of macroscopic observables of microscopic subsystems with mean field coupling

In the two previous sections we considered microscopic subsystems that were uncoupled. We now consider linear response when the heat bath variables $q^{(j)}$ couple via the mean field Φ_n. In Section 4.7.1 we derive a non-Markovian closure of the mean-field dynamics, along the lines of the reduction derived in Section 4.5, that is deterministic in the thermodynamic limit and stochastic for finite M; in Sections 4.7.2 and 4.7.3 we study the mean-field dynamics and its linear response using this macroscopic closure.

4.7.1. Surrogate approximation of the mean field dynamics. Before we can study the response of the mean field coupled system to external perturbations εg,

we need to understand the implied macroscopic dynamics Φ_n generated by the system for the externally unperturbed system with $\varepsilon = 0$. To do so we view the system as driven by a prescribed time-dependent external driver d_n rather than the mean field Φ_n, as illustrated in Figure 4.15 (which should be compared with Figure 4.1(b)). Hence we replace the mean field coupled dynamics (4.6) by

$$q_{n+1}^{(j)} = f(q_n^{(j)}, d_n; a^{(j)}, \varepsilon) \tag{4.32}$$

for a prescribed external driver d_n. In the thermodynamic limit of the mean field coupled system (4.6) we will see that the macroscopic mean field dynamics is deterministic (see (4.40) further down), and the driver $d_n = \Phi_n$ is indeed prescribed by the initial conditions, which is simply the initial distribution of the $q^{(j)}$. For large but finite M, the $q_n^{(j)}$ conditioned on the history of the mean field $(\Phi_{n-1}, \ldots)$ can be considered as independently distributed and one can again view the mean field Φ_n as an externally prescribed noisy driver d_n with specified statistical properties. The surrogate system (4.32) with the external driver d_n chosen as a random draw of the Gaussian process (defined below in (4.33)) provides an accurate representation of the statistical behaviour of the original mean field coupled system (4.6). We have checked that both yield the same linear response $\mathbb{E}^\varepsilon\Psi$, and now set out to study the linear response of the original mean field coupled system via the surrogate system.

Let us now determine the statistical properties of a macroscopic mean field observable for the driven surrogate system (4.32). The mean fields Φ_n and Ψ_n are again Gaussian process with (now time-dependent) statistical properties given again by statistical limit laws, and we write in particular

$$\Phi_n = \langle \mathbb{E}^d[\Phi_n] \rangle + \frac{1}{\sqrt{M}}\zeta_n + \frac{1}{\sqrt{M}}\eta_n + o\left(\frac{1}{\sqrt{M}}\right), \tag{4.33}$$

where

$$\mathbb{E}^d[\Phi_n] = \mathbb{E}[\Phi_n | d_k;\, k < n] \tag{4.34}$$

denotes the conditional expectation over the past history of the driver and averages now involve time-dependent measures $\mu_n^{a^{(j)}}$. The autocovariances of the mean-zero Gaussian process ζ_n are given by a central limit theorem approximation of $\Phi_n - \mathbb{E}^d[\Phi_n]$ with

$$\mathrm{cov}[\zeta_n, \zeta_{n-k}] = \langle \mathrm{cov}[\phi(q_n^{(j)}), \phi(q_{n-k}^{(j)})] \rangle, \tag{4.35}$$

where the covariance is defined using the conditional average over the history of the driver (cf. (4.12)). Note that the autocovariance is not a function of $n - m$ due to the non-Markovian nature of the dynamics. Similarly, a central limit theorem approximation of $\mathbb{E}^d[\Phi_n] - \langle \mathbb{E}^d[\Phi_n] \rangle$, defines the mean-zero Gaussian process η_n with

autocovariance

$$\langle \eta_n^{\varepsilon}, \eta_m^{\varepsilon'} \rangle = \langle \mathbb{E}^{d,\varepsilon}[\phi(q_n^{(j)})] \mathbb{E}^{d,\varepsilon'}[\phi(q_m^{(j)})] \rangle - \langle \mathbb{E}^{d,\varepsilon}[\phi(q_n^{(j)})] \rangle \langle \mathbb{E}^{d,\varepsilon'}[\phi(q_m^{(j)})] \rangle, \tag{4.36}$$

where again the conditional expectation values $\mathbb{E}$ are used (cf. (4.14)). Note that the Gaussian processes ζ_n and η_n are independent.

The impulsive response of Φ_n at a given time to a perturbation of the driving process $d_n \mapsto d_n + \theta_n$, where $|\theta_n| \ll 1$, can be, at least formally, captured by the susceptibility function

$$R_n(z) = \sum_{k=1}^{\infty} \chi_{n,k}\, z^k, \tag{4.37}$$

defined for complex z with $|z| \leq 1$ (Ruelle, 2004). The fluctuation coefficients $\chi_{n,k}$ describe the change of the mean field induced by the drivers θ_n as

$$\langle \mathbb{E}^{d+\theta}[\Phi_n] \rangle - \langle \mathbb{E}^{d}[\Phi_n] \rangle = \sum_{k=1}^{\infty} \chi_{n,k}\, \theta_{n-k}. \tag{4.38}$$

The fluctuation coefficients $\chi_{n,k}$ of Φ_n are given as an average over the microscopic fluctuation coefficients $\chi_{n,k}^a$ of $\phi(q_n^{(j)})$ as

$$\chi_{n,k} = \int \chi_{n,k}^a\, \nu(a) da. \tag{4.39}$$

A necessary condition for LRT with respect to a bounded driver θ_n is the summability of the coefficients χ_k. Once LRT with respect to the driver d_n can be shown, we can proceed to study the linear response with respect to the external perturbation with $\varepsilon \neq 0$ (recall that $d_n = \Phi_n$ (cf. Figure 4.1)). Note that if $\langle \mathbb{E}^d[\Phi_n] \rangle$ does not satisfy LRT with respect to a perturbation of the driver d_n, then it cannot be expected to satisfy LRT with respect to external perturbation.

In the thermodynamic limit, provided the microscopic dynamics is mixing, we can use that the measures μ_n^a are the physical invariant measures generated by the cocycle $f(\cdot, \Phi_n; a, \varepsilon)$ to create a closure of the dynamics of Φ_n as a deterministic recurrence relation:

$$\Phi_n = \langle \mathbb{E}^{d=\Phi}[\Phi_n] \rangle =: F(\Phi_{n-1}, \Phi_{n-2}, \ldots; \varepsilon). \tag{4.40}$$

If the mixing times of the $q^{(j)}$ are much shorter than a delay k_*, then the effect of the driving Φ_{n-k} for $k > k_*$ is minimal and the mean field dynamics is effectively Markovian in a space of dimension k_* or less. The linear response with respect to perturbations, $\mathbb{E}^{\varepsilon}\Phi_n$, is now determined by the properties of the deterministic macroscopic dynamics (4.40). In the following sections we shall consider cases and

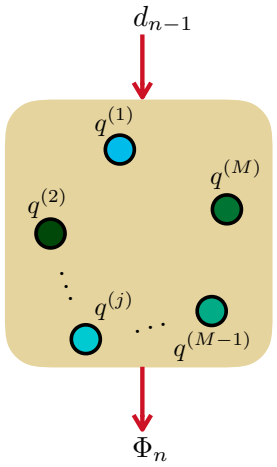

FIGURE 4.15. Sketch of the macroscopic dynamics Φ_n mediated by the microscopic reservoir.

conditions on the deterministic macroscopic dynamics (4.40) when $\mathbb{E}^\varepsilon \Phi_n$ enjoys linear response and when it does not.

For finite size M the response approximates that of the thermodynamic limit, as in the case of an uncoupled heat bath discussed in Section 4.5. However, the central limit theorem approximation (4.33) assures that the microscopic dynamics (4.32), which is driven by $d_n = \Phi_n$, is essentially stochastic with a noise process ζ_n that has decay of temporal correlations (since the $q^{(j)}$ exhibit decay of correlations). This self-generated dynamic noise induces linear response for finite size mean field coupled heat baths: this is similar to the results for finite M and in the diffusive-scaled coupling seen in Section 4.6, and can again be understood as a consequence of results by Hairer and Majda (Hairer and Majda, 2010).

In Section 4.7.2-4.7.3 we will consider the dynamics of the system in the thermodynamic limit. The first case is when $\Phi_n = \langle \mathbb{E}\Phi_n \rangle$ approaches a fixed point $\bar{\Phi}$ for $M \to \infty$, the second case is when the mean field Φ_n itself exhibits nontrivial dynamics. Whereas in the first case the linear response of the macroscopic observable Ψ is determined by the properties of the microscopic dynamics, in the latter case it is entirely determined by the response of the macroscopic dynamics.

4.7.2. Trivial dynamics of the mean field observable. Let us first look at the case of the mean field at a stable fixed point $\bar{\Phi}$, in the sense that the mean field remains bounded when perturbed from $\bar{\Phi}$ and, when the collective dynamics has LRT, $\bar{\Phi}$ is a stable fixed point of Φ_n in the thermodynamic limit.

To understand the stability, we can apply the external driving framework expounded in the previous section to the dynamics of our system about the equilibrium $d_n \equiv \bar{\Phi}$. Stability is in fact assured provided that the complex susceptibility function $R(z)$ does not have any roots inside the unit disk. This follows by considering $\theta_n = \langle \mathbb{E}\Phi_n \rangle - \bar{\Phi} = \sum_{k=0}^{\infty} \chi_{n,k}\theta_{n-k}$ with $\theta_n \sim \lambda^n$ in (4.38) which leads to

$$R(\lambda^{-1}) - 1 = 0,$$

and hence, for unstable $|\lambda| > 1$, to the above condition for instability for the susceptibility function $R(z)$.

If stability is ensured, the linear response of the fixed point $\bar{\Phi}$ with respect to external perturbations εg is established by the implicit function theorem from the deterministic macroscopic dynamics (4.40). Once an external perturbation εg is applied, the fixed point depends on ε, and we write

$$0 = F(\bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \cdots; \varepsilon) - \bar{\Phi}^\varepsilon. \tag{4.41}$$

In the following numerical experiments the $M \to \infty$ limit is computed by estimating a solution to this algebraic equation.

Differentiation with respect to the external perturbation yields

$$\begin{aligned} 0 &= \frac{\mathrm{d}\bar{\Phi}^\varepsilon}{\mathrm{d}\varepsilon}\left(\frac{\partial}{\partial\bar{\Phi}^\varepsilon}F(\bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \cdots; \varepsilon) - 1\right) \\ &\quad + \frac{\partial}{\partial\varepsilon}F(\bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \cdots; \varepsilon) \\ &= \frac{\mathrm{d}\bar{\Phi}^\varepsilon}{\mathrm{d}\varepsilon}\left(\sum_{k=1}^{\infty}\chi_k - 1\right) + \frac{\partial}{\partial\varepsilon}F(\bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \cdots; \varepsilon). \end{aligned}$$

This immediately yields that

$$\frac{\mathrm{d}\bar{\Phi}^\varepsilon}{\mathrm{d}\varepsilon} = \frac{\frac{\partial}{\partial\varepsilon}F}{1 - R(1)}, \tag{4.42}$$

and hence the existence of linear response, provided $R(1) - 1 \neq 0$.

As for the uncoupled scenario, we shall now discuss the linear response behaviour for the three different cases of the microscopic dynamics, which are covered by the rows in Table 4.1 corresponding to the coupled macroscopic observables.

4.7.2.1. *The microscopic subsystems satisfy LRT.* We consider here the case of uniformly expanding dynamics of the microscopic systems, such that each subsystem individually satisfies LRT. In particular, we choose the following uniformly expanding

map

$$q_{n+1} = \frac{\mathfrak{d}(q_n) + K_n\left(1 - \sqrt{0.03(1-0.97K_n^2) + 0.97(\mathfrak{d}(q_n)+K_n)^2}\right)}{1-0.97K_n^2}, \tag{4.43}$$

where $K_n = \tanh(\varepsilon\Phi_n - 2)$, $q_n \in [-1,1]$ and $\mathfrak{d}(q) = 2q - \operatorname{sign} q$ is the doubling map on $[-1,1]$. All microscopic degrees of freedom $q^{(j)}$ evolve according to the same map but with randomly distributed initial conditions. (Note that having identical microscopic subsystems implies that $\langle \mathbb{E}^\varepsilon \Phi\rangle = \mathbb{E}^\varepsilon \Phi$.) This map is full-branch uniformly expanding for fixed K_n. It is carefully constructed to allow for nontrivial mean field dynamics for larger values of ε which will be discussed in Section 4.7.3. We choose the coupling function $\phi(q) = -\frac{23}{30} + \frac{7}{2}q^2 - 2q^4$ to generate the mean field Φ_n. For simplicity, we choose the mean field observable $\Psi = \Phi$. The dynamics in the thermodynamic limit $M = \infty$ was computed in Poltergeist (Wormell, 2017) using the Chebyshev methods in Chapter 4 (see Appendix D2 for more details). For small values of ε the mean field Φ_n converges to a stable fixed point, and the macroscopic observable $\Psi = \Phi$ satisfies LRT as shown in Figure 4.16. The variation about $\Phi = \bar{\Phi}$ can be shown to converge to the limiting distribution of the mean zero stochastic process ζ_n with autocovariance (4.36). We will see later in Section 4.7.3 that for larger values of ε, the mean field exhibits nontrivial chaotic dynamics, violating LRT.

4.7.2.2. *The microscopic subsystems do not satisfy LRT but are appropriately heterogeneous.* We consider a mean field coupled system of LRT-violating modified logistic maps (4.7). We choose the external perturbation $g(q) = 4(q(1-q))^2$ as in (4.15), and mean-field coupling

$$h(q,\Phi) = (1-2q)q(1-q)\tanh\Phi. \tag{4.44}$$

The effect of g and h on the map are plotted in Figure 4.2. We remark that a naive choice of mean field coupling with $h(q,\Phi) = \Phi_n$ would just lead back to the standard logistic map for some $p_n^{(j)} = \alpha q_n^{(j)} + \beta$ with a modified logistic map parameter $a^{(j)} = a^{(j)}(\Phi_n)$. The mean field Φ_n is given by (4.3) and is constructed using

$$\phi(q) = 4T_5(2q-1) + 1,$$

where $T_5(x) = 16x^5 - 20x^3 + 5x$ is the 5th Chebyshev polynomial, which oscillates between ± 1 in the domain (see Figure 4.2(b)).

We draw the parameters $a^{(j)}$ of the logistic map from the smooth raised-cosine distribution on $[3.7, 3.8]$ (4.8). In this case, the macroscopic dynamics Φ_n converges to a stable fixed point $\Phi_n \to \bar{\Phi}$ for $\varepsilon < -0.075$. The associated linear response is clearly visible in Figure 4.17. In fact, as discussed in Section 4.5.2.2, nonlinear third order response holds for the three times continuously differentiable raised cosine distribution (4.8). We remark that for $\varepsilon > -0.075$ the mean field exhibits nontrivial

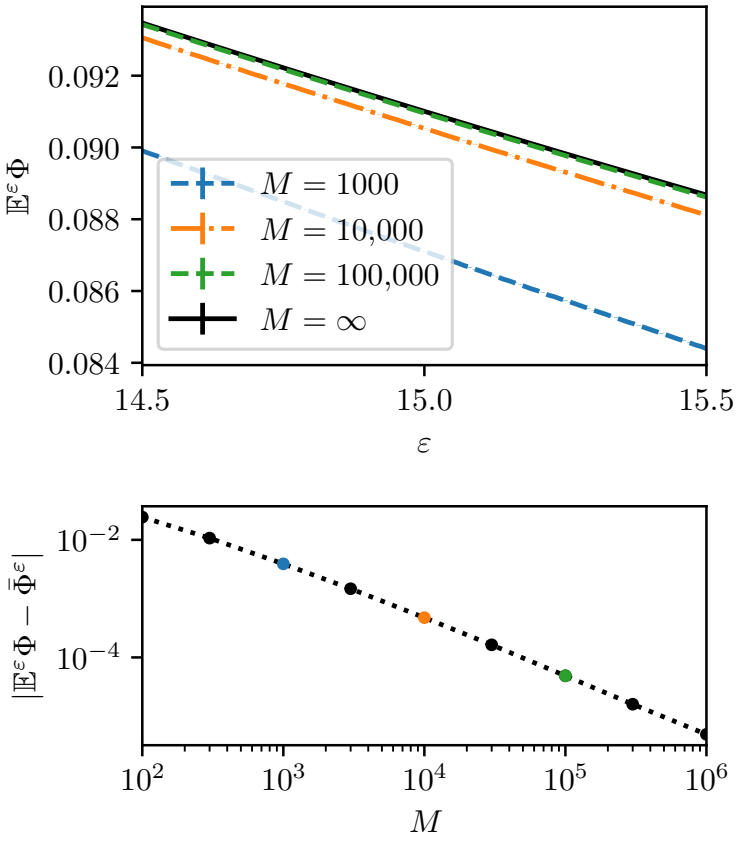

FIGURE 4.16. (a) Response term $\mathbb{E}^\varepsilon\Phi$ for the uniformly expanding map (4.43) with mean field coupling under trivial dynamics; (b) Difference $|\mathbb{E}^\varepsilon\Phi - \bar{\Phi}^\varepsilon|$ for $\varepsilon = 15$, exhibiting $\mathcal{O}(1/M)$ convergence.

dynamics in the thermodynamic limit, and we observe a breakdown of LRT to be discussed in Section 4.7.3.

In Figure 4.18 we see very slow convergence of the mean $\mathbb{E}^\varepsilon\Psi$ to its limiting value: in particular, it is slower than the $\mathcal{O}(1/M)$ rate for uniformly expanding dynamics leading to trivial dynamics (cf Figure 4.16), and seemingly slower still than the $\mathcal{O}(1/\sqrt{M})$ rate that we might expect from sampling errors of η. Although smooth families of microscopic logistic maps allow for linear response of macroscopic observables to *constant-in-time perturbations* as discussed in Section 4.5.2.2, and in fact numerical experiments (not shown) suggest that the susceptibility function χ has

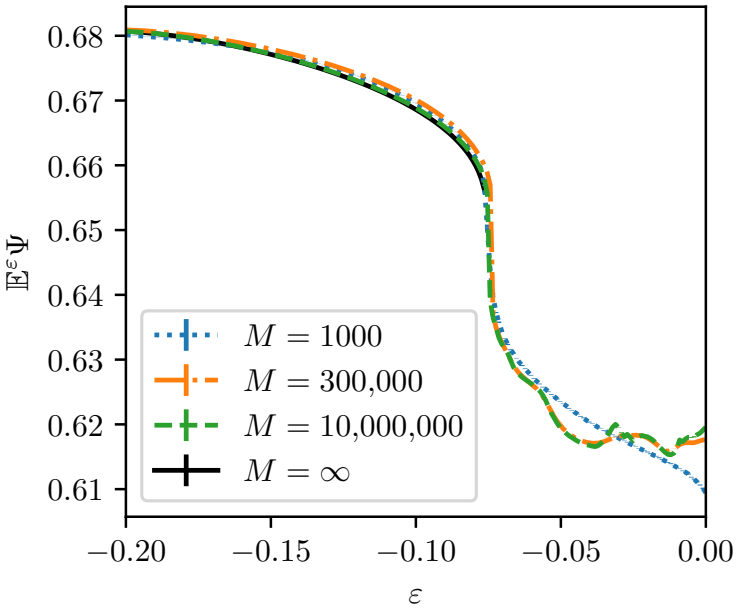

FIGURE 4.17. Response $\mathbb{E}^\varepsilon\Psi$ for the modified logistic map (4.7) with mean field coupling (4.44). The parameters $a^{(j)}$ are drawn from the raised-cosine distribution (4.8). Error bars were estimated from 10 realisations of 10^5 iterates, differing in the initial conditions of the heat bath, and are not visible.

summable decay, they do not appear to have linear response with respect to stochastic perturbations. We argue that this counter-intuitive lack of linear response with respect to the (self-generated) stochastic perturbations arises from the noise-induced destruction of narrow periodic windows that have "extreme" values of $\mathbb{E}^{a,\varepsilon}\psi(q)$ compared with the neighbouring, more stochastically stable chaotic parameters. Thus at these periodic parameter values the macroscopic dynamics exhibits a disproportionately large response to the introduction of noise. We illustrate this in Figure 4.19, where we plot the response of a single logistic map with additive noise of variance σ^2. Here the noise models the finite size effects of the heat bath with $\sigma \sim 1/\sqrt{M}$. One sees clearly that periodic windows can be destroyed by very small amounts of noise ($\sigma = 10^{-6}$). One also sees that associated with the destruction of these periodic windows is a very large response in the average $\mathbb{E}^a\psi$.

The statistical properties of the macroscopic observable for dynamics of a finite size heat bath can be modelled again by a surrogate system. Writing

$$\Phi_n = \bar{\Phi}^{\varepsilon,M} + \frac{1}{\sqrt{M}}\zeta_n,$$

where ζ_n is the Gaussian central limit theorem correction term to $\Phi_n \equiv \bar{\Phi}^{\varepsilon,\infty}$ with covariance given by (4.35). The macroscopic dynamics (4.41) then becomes

$$0 = \mathbb{E}F(\bar{\Phi}^{\varepsilon,M} + \frac{1}{\sqrt{M}}\zeta_{n-1}, \bar{\Phi}^{\varepsilon,M} + \frac{1}{\sqrt{M}}\zeta_{n-2}, \ldots) - \bar{\Phi}^{\varepsilon,M}.$$

The response $\mathbb{E}^\varepsilon\Psi$ for this surrogate macroscopic dynamics is shown in Figure 4.18, labelled CLT approximation, and is barely distinguishable from the response of the original macroscopic dynamics.

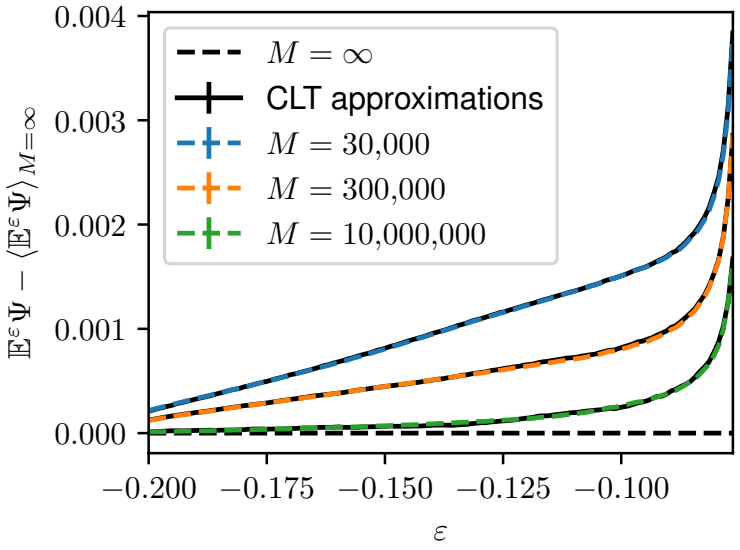

Figure 4.18. Difference between the response $\mathbb{E}^\varepsilon\Psi$ for finite M and for the thermodynamic limit for the modified logistic map (4.7) with mean field coupling (4.44). The parameters $a^{(j)}$ of the logistic map are drawn from the raised-cosine distribution (4.8). For each value of M the response of the corresponding central limit theorem (CLT) approximation using noise estimated from $M = 10^6$ was used. Error bars were estimated from 10 realisations of 10^5 iterates, differing in the initial conditions of the heat bath, and are not visible.

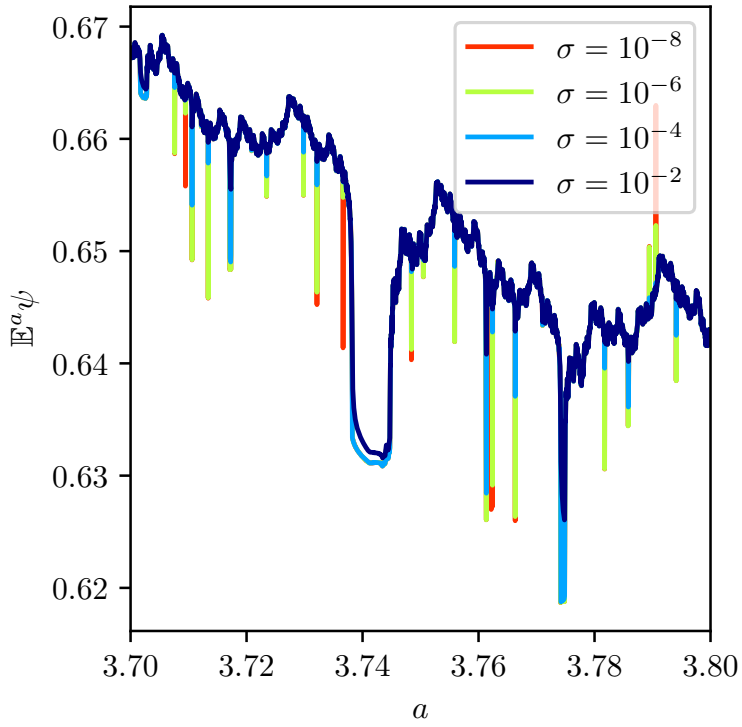

FIGURE 4.19. Response $\mathbb{E}^a\psi$ for a single stochastically driven logistic map $q_{n+1} = aq_n(1-q_n) + \sigma\xi_n$, where ξ_n is *i.i.d.* Gaussian noise and the observable $\psi(q) = q$. The response is recorded at increments of $da = 10^{-8}$, thus for small σ only a subset of narrow periodic windows are captured.

4.7.2.3. *The microscopic subsystems do not satisfy LRT and are not appropriately heterogeneous.* For a non-smooth distribution of the logistic map parameters $a^{(j)}$, namely the discrete distribution (4.9), the dynamics of Section 4.7.2.2 also converges to a stable fixed point $\bar{\Phi}^{\varepsilon,M}$ for $\varepsilon < -0.075$. The mean field $\bar{\Phi}^{\varepsilon,M}$ varies smoothly with respect to ε for almost all ε, but $\mathbb{E}^\varepsilon\Psi$ experiences saddle node bifurcations on increasingly dense sets as M approaches the thermodynamic limit $M \to \infty$. This is illustrated in Figure 4.20. Looking at (4.42), we see that linear response is violated where the fixed point loses stability and $\partial F(\bar{\Phi}^\varepsilon, \bar{\Phi}^\varepsilon, \cdots)/\partial\bar{\Phi}^\varepsilon = R(1) = 1$.

For finite M, the macroscopic equation can be modelled as $\Phi_n^{\varepsilon,M} = \bar{\Phi}^{\varepsilon,M} + \zeta_n/\sqrt{M}$ where

$$\bar{\Phi}^{\varepsilon,M} = \mathbb{E}^\varepsilon F(\Phi_n^{\varepsilon,M}, \Phi_n^{\varepsilon,M}, \cdots).$$

The derivatives $\partial \mathbb{E}^\varepsilon F(\Phi_n^\varepsilon, \Phi_n^\varepsilon, \cdots)/\partial \bar{\Phi}^\varepsilon$ are well defined and thus for appropriate bath sizes M we observe again approximate LRT for all practical purposes. However, as $M \to \infty$, saddle-node bifurcations become visible as a result of the diminishing effect of stability-providing noise, and in the thermodynamic limit we observe failure of LRT (see inset in Figure 4.20). The failure of linear response through saddle-node bifurcations is accompanied by $\mathbb{E}^\varepsilon \Psi$ experiencing multistability, with multiple very close stable equilibria, demonstrated in Figure 4.20[1] . This can be understood as coming from the fact that $\mathbb{E}^\varepsilon F(\Phi_n^{\varepsilon,M}, \Phi_n^{\varepsilon,M}, \cdots)$ is essentially a smoothed out version of the rough logistic map response $\mathbb{E}^\varepsilon F(\bar{\Phi}, \bar{\Phi}, \cdots)$: as $M \to \infty$ the smoothing decreases, leading to increasing numbers of roots of the equation.

4.7.3. Nontrivial dynamics of the mean field observable. The mean field Φ or any macroscopic observable Ψ may itself exhibit non-trivial dynamics of varying complexity in the thermodynamic limit $M = \infty$. The overall response behaviour is then determined by the macroscopic dynamics rather than by the properties of the microscopic subsystems. We show the emergence of non-trivial chaotic macroscopic dynamics which violates LRT. The first one, surprisingly, involves a heat bath which evolves under uniformly expanding dynamics when uncoupled, and the second one involves microscopic dynamics that individually violate LRT.

To generate emergent nontrivial macroscopic dynamics of the mean field, we again use the uniformly expanding map (4.43) with the even Lebesgue-measure zero coupling function $\phi(q) = -\frac{23}{30} + \frac{7}{2}q^2 - 2q^4$ to generate the mean field Φ_n. We show in Figure 4.21 the map and its invariant measure, where the dynamics in the thermodynamic limit $M = \infty$ was computed using the spectral method presented in Chapter 2, and specifically with the numerical package Poltergeist (see Appendix D2 for more details). The map and coupling function K_n are judiciously chosen to yield nontrivial dynamics for the mean field Φ_n, mediating dynamics akin to a unimodal map for Φ_n.

The map is constructed such that when the $q^{(j)}$ are approximately evenly distributed, $\Phi_n \approx 0$, causing an extreme value $K_n \approx \tanh(-2) \approx -0.96$: this pushes the $q^{(j)}$ strongly towards $q = -1$ which leads to a larger value Φ_{n+1}, concentrating around $\Phi = 2/\varepsilon$. For these values of Φ_{n+1}, $K_{n+1} \approx 0$, and thus in the next step the $q^{(j)}$ are spread more evenly over the interval $[-1, 1]$, mapping Φ_{n+1} back around zero. The concentration in the first step provides the folding and the sensitivity of K_n to small changes in Φ_n for large ε provides the stretching necessary for chaotic dynamics.

[1]When randomly searching for equilibria it is important to make sure that, as well as randomly initialising the $q_0^{(j)}$, the *distribution* from which the $q_0^{(j)}$ are sampled is also randomly initialised, as up to an error term of $\mathcal{O}(M^{-1/2})$ the macroscopic dynamics are deterministic functions of the initial measures of the microscopic variables μ_0^a.

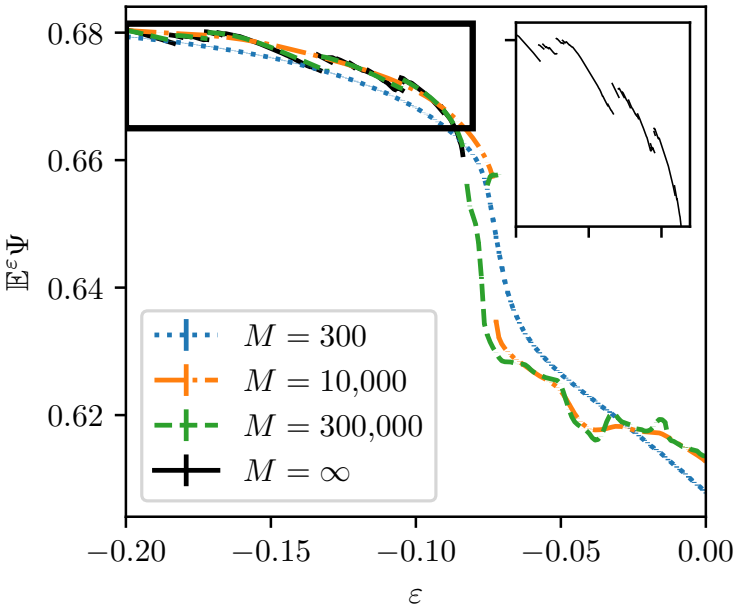

Figure 4.20. Response $\mathbb{E}^\varepsilon\Psi$, including multistability, for the modified logistic map (4.7) with mean field coupling (4.44). The parameters $a^{(j)}$ are drawn from the discrete distribution (4.9). Error bars were estimated from 200 realisations differing in the initial conditions (including initialising distributions) of the heat bath, and are not visible. The inset illustrates the occurrence of saddle-node bifurcations in the infinite-dimensional limit.

In Figure 4.22 we show the map $\Phi_{n+1} = F(\Phi_n, \Phi_{n-1}, \ldots)$ generated by the dynamical system (4.43) in the thermodynamic limit $M = \infty$ for $\varepsilon = 30$. The dynamics is clearly chaotic with the leading Lyapunov exponent $\lambda_1 = 0.18 > 0$ ($\lambda_2 = -0.43$ and $\lambda_3 = -0.81$). The dynamics of the macroscopic observable $\Psi_n = \Phi_n$ exhibits a complex bifurcation cascade upon varying ε, depicted in Figure 4.22. For $\varepsilon \leq 18.4159$, the macroscopic dynamics has a stable fixed point; upon increasing the perturbation ε a period-doubling cascade leads to chaotic, apparently unimodal-like, dynamics intermingled with periodic windows for values of $\varepsilon > 26.1649$. One can clearly see dark scars in the bifurcation diagram in the chaotic region of $\varepsilon > 26.1649$. This is reminiscent of the logistic map (Collet and Eckmann, 2007) where the scars denote narrow intervals of $\langle\Phi_n\rangle$ with increased probability, corresponding to large spikes in the invariant measure, which (unlike small spikes) vary smoothly with respect to

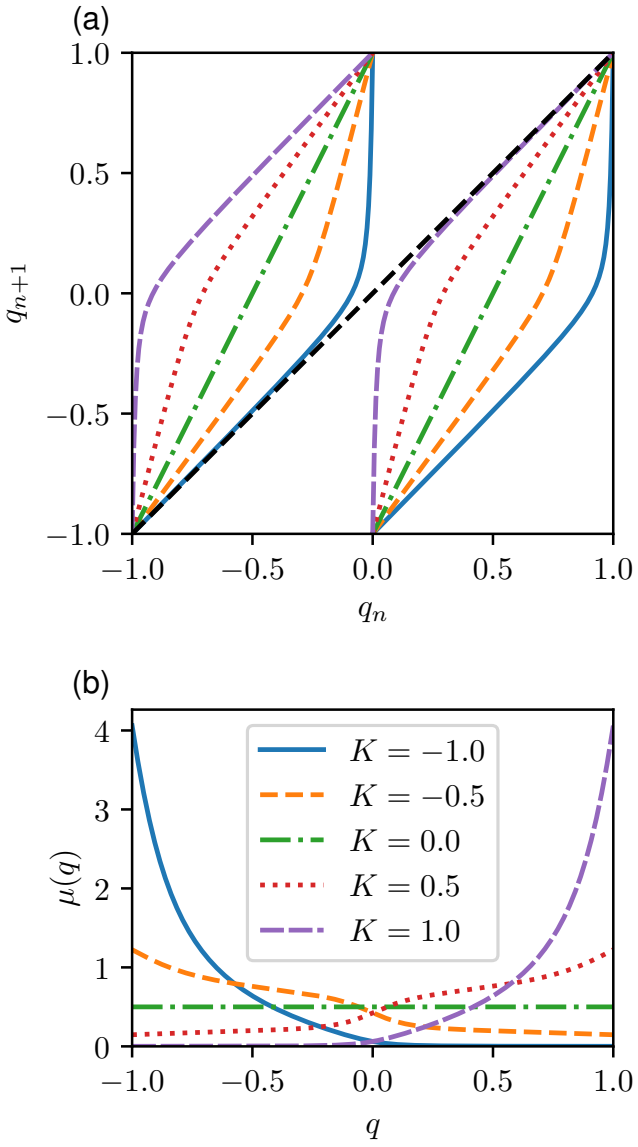

FIGURE 4.21. (a) Plot of the uniformly expanding map (4.43) and (b) its invariant measure for various values of $K_n \equiv K$.

perturbations ε.

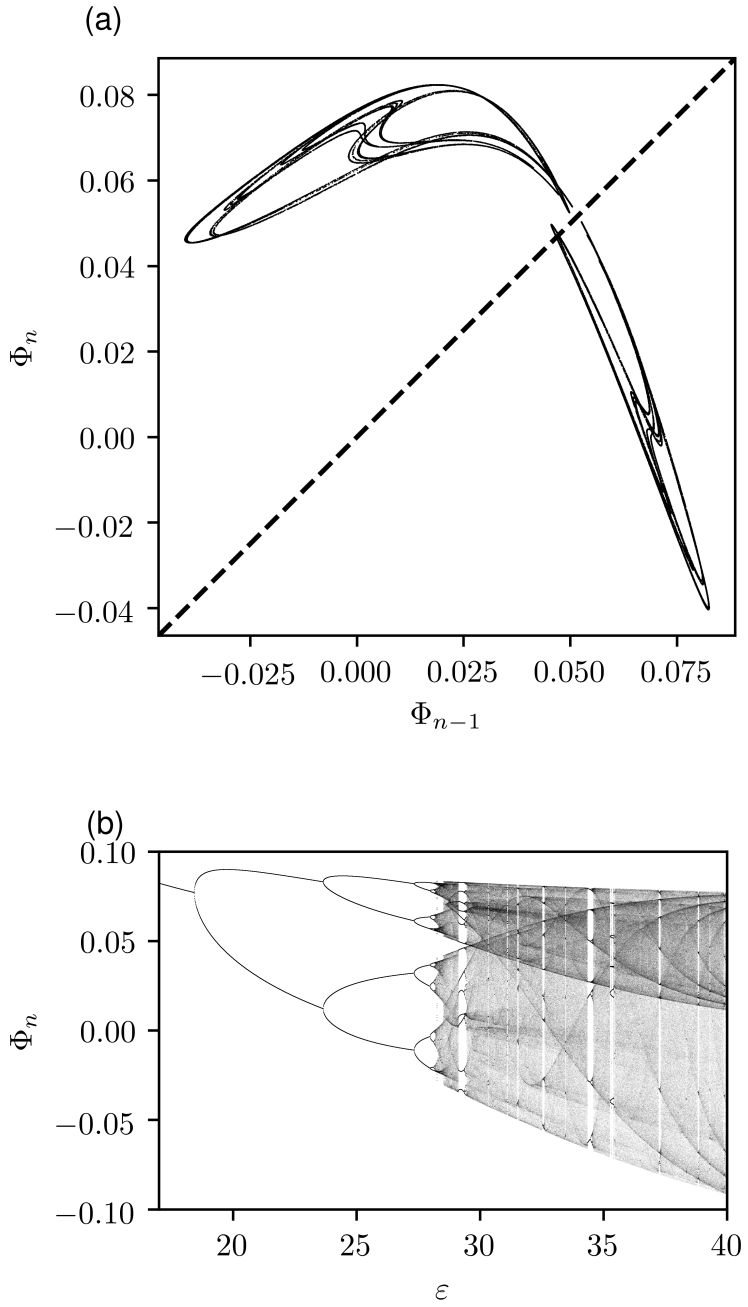

FIGURE 4.22. (a): 2D projection of the attractor onto delay coordinates of the macroscopic map $\Phi_{n+1} = F(\Phi_n, \Phi_{n-1}, \ldots)$ generated by the uniformly expanding map (4.43) for $\varepsilon = 30$. The system has two periodic components separated by a gap around the unstable fixed point $\Phi_{n-k} \equiv 0.51258$. (b): Bifurcation diagram of the map (4.43) showing period doubling bifurcations and chaotic dynamics.

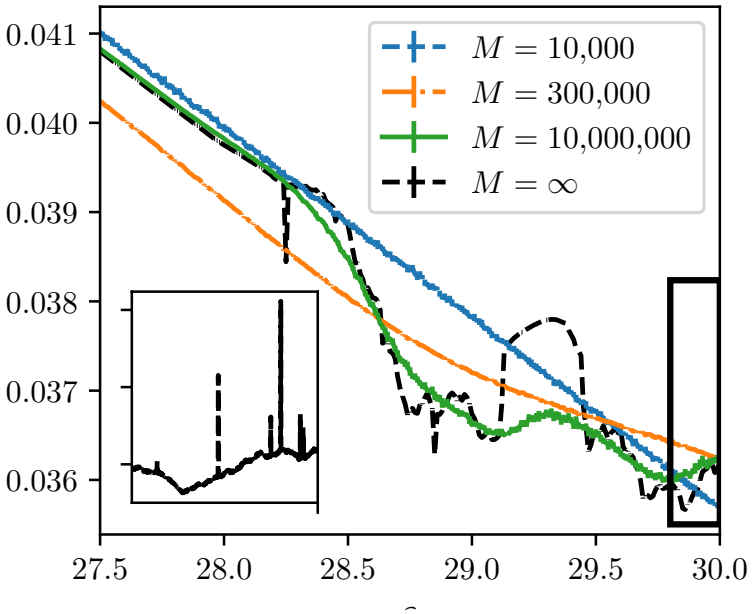

Figure 4.23. Response $\mathbb{E}^{\varepsilon}\Psi$ of the uniformly expanding map (4.43) for finite M response, showing convergence to the thermodynamic limit $M \to \infty$. The black box shows the region which is magnified on the right. Error bars were estimated from 10 realisations differing in the initial conditions of the heat bath, and are not visible.

In Figure 4.23 we show the linear response term $\mathbb{E}^{\varepsilon}\Psi$ of the uniformly expanding map (4.43) for several finite M heat bath sizes and for the thermodynamic limit $M = \infty$ for $\varepsilon \in [27.5, 30]$, clearly illustrating the breakdown of LRT. We recall that the same map exhibits LRT for small values of ε, where the macroscopic mean field converges to a stable fixed point, for the same parameters (cf. Figure 4.16).

For curiosity and to further study the effect of the self-generated noise on the LRT behaviour of macroscopic observables in the mean field coupled case, we provide another example of nontrivial chaotic mean field dynamics which violates LRT. We revisit the mean field coupled dynamics of microscopic subsystems which do not satisfy LRT discussed in Section 4.7.2.2, and consider the modified logistic map (4.7) with mean field coupling (4.44) where the parameters $a^{(j)}$ of the logistic map are drawn from the smooth raised-cosine distribution (4.8). We recall that

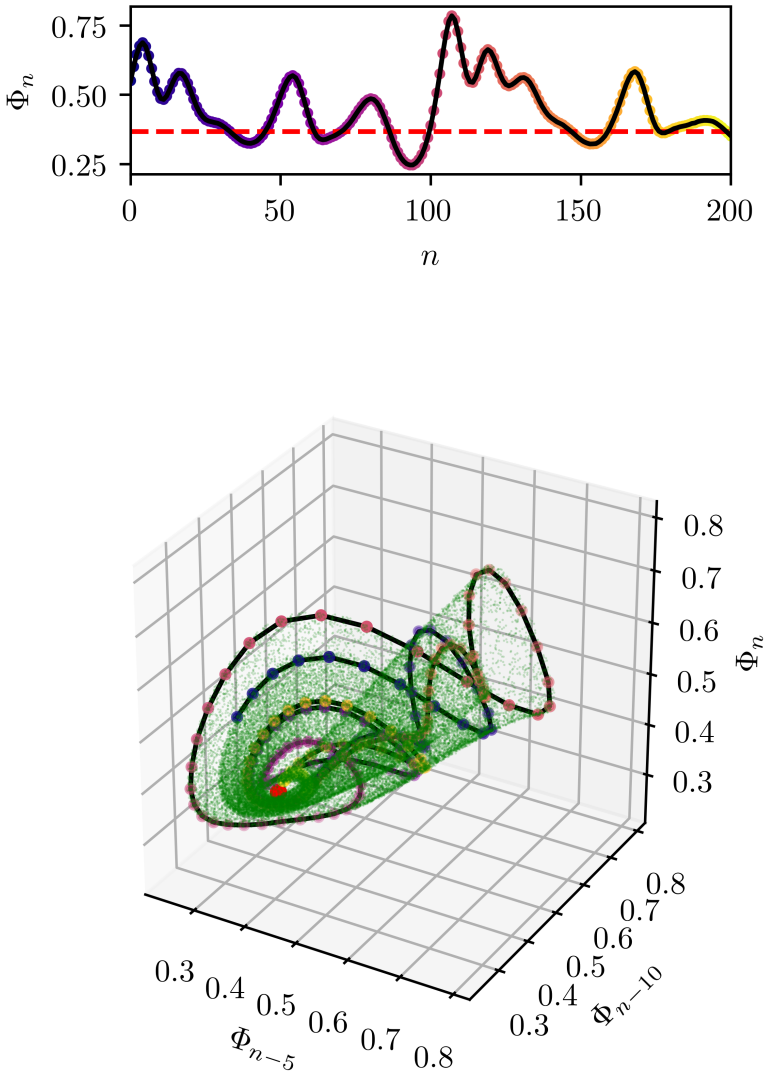

FIGURE 4.24. Top: Time series of the macroscopic map $\Phi_{n+1} = F(\Phi_n, \Phi_{n-1}, \ldots)$ generated by the modified logistic map (4.7) with mean field coupling (4.44) for $\varepsilon = 0$, approximated by a finite ensemble of size $M = 10^7$. The red dotted line shows an unstable fixed point of the system. The parameters $a^{(j)}$ are drawn from the raised-cosine distribution (4.8). Bottom: Projection onto delay coordinates of the attractor and dynamics of the same map. The red dot near the centre of the attractor denotes an unstable fixed point of the system.

for $\varepsilon \approx [-0.2, -0.075]$ the macroscopic dynamics (4.40) was trivial and $\mathbb{E}^\varepsilon\psi$ satisfies LRT (cf. Figure 4.17). The stable fixed point loses stability at $\varepsilon \approx -0.075$ through a saddle-node bifurcation (not shown), from which emanates a stable limit cycle centred around an unstable fixed point with $\mathbb{E}\Psi \approx 0.615$: as ε increases, this bifurcates to chaos and for a wide range of values $\varepsilon > -0.075$ nontrivial chaotic macroscopic dynamics of (4.40) is observed. Figure 4.24 illustrates the macroscopic dynamics for $\varepsilon = 0$, which exhibit Shilnikov-type chaos. The associated response term $\mathbb{E}^\varepsilon\Psi$ was shown in Figure 4.17 for several finite heat bath sizes M and for the thermodynamic limit $M = \infty$ for $\varepsilon \in [-0.2, 0]$, clearly illustrating the transition to LRT violating macroscopic dynamics around $\varepsilon = -0.075$. Note that the finite size response is smoothed due to the self-generated noise process ζ_n.

The above examples of high-dimensional systems, which appear to exhibit non-uniformly hyperbolic chaotic collective behaviour, are in disagreement with the often invoked assumption that macroscopic observables of high-dimensional systems obey linear response. This is the more surprising as the non-uniformly hyperbolic chaotic behaviour is robust (modulo periodic windows) with respect to the external perturbation, different choices of the coupling function, different weightings in the coupling, etc. In the next section, we establish directly that the system is indeed non-hyperbolic, and that the non-hyperbolicity persists under generic perturbations, through finding a homoclinic tangency.

4.8. A homoclinic tangency in the macroscopic dynamics of a mean-field coupled system

In this section we show that, unlike Anosov systems, there are families of mean-field coupled systems that do not obey linear response theory. While the chaotic hypothesis cannot be applied to linear response (as discussed in Section 4.1), we will demonstrate in this section an actual violation of the chaotic hypothesis itself in one of these systems, in the sense that the limiting macroscopic dynamics are non-Anosov (i.e. non-uniformly hyperbolic). In particular we show, numerically, that the dynamics involve a homoclinic tangency (see Section 4.8.1). The aim of this is to show that the system defined in Section 4.7.2.1 is a high-dimensional example of a system with non-hyperbolic macroscopic dynamics, thus violating the chaotic hypothesis of Gallavotti and Cohen (1995a,b) (see Section 4.1).

This section is structured as follows. In Section 4.8.1 we will properly define relevant geometrical objects, including homoclinic tangencies; in Section 4.8.2 we describe the thermodynamic limit of the mean-field coupled system; in Section 4.8.3 we outline the numerical methods used to find the homoclinic tangency, which finally we present in Section 4.8.4.

4.8.1. Anosov maps and homoclinic tangencies. Suppose that for M a subset of a Banach space, one has a map $H : M \to M$ with $x \in M$. The stable manifold of x is the set of points whose forward orbits converge to the forward orbit of x:

$$\mathcal{V}_x^s = \{y \in M : \lim_{n\to\infty} d_M(H^n x, H^n y) = 0\},$$

where d_M is the metric on M.

Similarly, the unstable manifold of x is the set of points with *backward* orbits converging to that of x:

$$\mathcal{V}_x^u = \{y \in M : \lim_{n\to\infty} d_M(H^{-n} x, H^{-n} y) = 0\}.$$

We can extend these notions of stable and unstable manifolds onto the tangent bundles. The *stable subspace* at x, E_x^s, is the set of tangent vectors at x which converge to zero under the action of T:

$$E_x^s = \{v \in T_x M : \lim_{n\to\infty} D_x H^n \, v = 0\},$$

where $D_x H$ is the Jacobian of H. Similarly the *unstable subspace* at x is given as

$$E_x^u = \{v \in T_x M : \lim_{n\to\infty} D_x H^{-n} \, v = 0\}.$$

These are respectively tangent to stable and unstable manifolds (Bowen, 2008).

If, as in our situation, H is not a diffeomorphism, then unstable manifolds and subspaces are ill-defined. However, it is possible to define the unstable manifold (resp. subspace) of a backward orbit $(x_{-n})_{n\in\mathbb{N}}$. If x_* is a fixed point then for convenience we will define $\mathcal{V}_{x_*}^u$ (resp. $\mathcal{E}_{x_*}^u$) to be the unstable manifold (resp. subspace) of the orbit $x_{-n} \equiv x_*$.

If x is a fixed point, then $\mathcal{V}_x^s$ are the set of points with orbits converging to x, and $\mathcal{V}_x^u$ are the set of points with orbits emanating from x; furthermore, provided that the Jacobian $D_x H$ is hyperbolic, E_x^s is the span of the stable eigenspaces and E_x^u the span of the unstable eigenspaces.

A separation between unstable and stable subspaces are key properties of most well-behaved chaotic systems. A system is *uniformly hyperbolic* if at every point the tangent space allows for a splitting $T_x M = E_x^s \oplus E_x^u$, and there are constants $c > 0$, $\lambda < 1$ such that for all $x \in M$,

$$\|D_x H^n{}_{|E_x^s}\| \leq c\lambda^n,$$
$$\|D_x H^{-n}{}_{|E_x^u}\| \leq c\lambda^n.$$

Uniformly hyperbolic diffeomorphisms with compact, transitive attractors are known as *Anosov diffeomorphisms*; systems that are Anosov on the attractor are known as *Axiom A diffeomorphisms* (Bowen, 2008).

While Axiom A diffeomorphisms form a fairly restricted subclass of chaotic dynamical systems,they are widely cited as being representative of the large-scale dynamics of many physical chaotic systems Gritsun (2010); Lucarini and Sarno (2011); Nijsse et al. (2019) under the so-called chaotic hypothesis of Gallavotti and Cohen (see Section 4.1).

One generic mechanism to generate non-hyperbolic dynamics is via *homoclinic tangencies*. A *homoclinic tangency* in a map $H : M \circlearrowleft$ is a fixed point p of H together with a different point $q \in \mathcal{V}_p^s \cap \mathcal{V}_p^u$ such that $\mathcal{V}_p^s = \mathcal{V}_q^s$ and $\mathcal{V}_p^u = \mathcal{V}_q^u$ are tangent at q. This property means that the tangent spaces E_q^s and E_q^u have non-trivial intersection, implying non-hyperbolicity of the given map H.

4.8.2. Mean-field system. We will consider the mean-field limit system whose microscopic variables are uniformly expanding maps, which was introduced in Section 4.7.2.1. Recall that the system is defined on $q \in [-1,1]^M$ as

$$q_{n+1} = \frac{\mathfrak{d}(q_n) + K_n\left(1 - \sqrt{0.03(1-0.97K_n^2) + 0.97(\mathfrak{d}(q_n)+K_n)^2}\right)}{1-0.97K_n^2},$$
$$\Phi_n = \frac{1}{M}\sum_{j=1}^{M} \phi(q_n^{(j)}),$$

where $K_n = \tanh(\varepsilon\Phi_n - 2)$, $\mathfrak{d}(q) = 2q - \operatorname{sign} q$ and

$$\phi(q) := -\tfrac{23}{30} + \tfrac{7}{2}q^2 - 2q^4.$$

The maps f and the coupling function ϕ are plotted in Figure 4.21.

In the thermodynamic limit we may reduce to macroscopic dynamics (c.f. Section 4.7.1): here instead of writing this as a delay system in the coupling variable Φ_n as in (4.40)), we write explicitly the dynamics of the microscopic variable distributions μ_n

$$\mu_{n+1} = \mathcal{L}_{\varepsilon;\Phi_n}\mu_n,$$
$$\Phi_n = \int_{-1}^{1} \phi\,\mu_n\,\mathrm{d}q,$$

where $\mathcal{L}_{\varepsilon;\Phi}$ is the transfer operator of $f_{\varepsilon,\Phi}$ and we interpret the μ_n as densities with respect to Lebesgue. In particular, we can write the macroscopic dynamics as a map on probability distributions

$$\mu_{n+1} = G_\varepsilon[\mu_n] := \mathcal{L}_{\varepsilon;\varphi[\mu_n]}\mu_n, \tag{4.45}$$

where the functional $\varphi[\mu] = \int \phi\mu\,\mathrm{d}q$. We will restrict the domain of G to be (appropriately smooth) functions of total integral equal to one, i.e. the μ must be densities of probablility measures.

The transfer operators $\mathcal{L}_{\varepsilon;\Phi_n}$ are well-defined as endomorphisms on a range of Banach spaces, including $BV([-1,1])$ and Hardy spaces (Lasota and Yorke, 1973; Bandtlow and Jenkinson, 2008). Because these systems have uniformly bounded uniform expansion and distortion bounds, many relevant quantities of these maps and their cocycles, including the norm, spectral gap and the gap between Lyapunov exponents of the cocycle are uniformly bounded (Baladi et al., 1996; Bandtlow and Jenkinson, 2008; Korepanov, 2015).

4.8.3. Finding the homoclinic tangency in the mean-field coupled system. To find the homoclinic tangency we will need to compute stable and unstable directions in the system, which entails computing Jacobians. Our macroscopic system (4.45) has Jacobian

$$D_\mu G_\varepsilon = \mathcal{L}_{\varepsilon;\varphi[\mu]} - \frac{\mathrm{d}}{\mathrm{d}q}(X_{\varepsilon;\varphi[\mu]} G_\varepsilon(\mu))\varphi, \tag{4.46}$$

where

$$X_{\varepsilon;\Phi} \circ f_{\varepsilon;\Phi} := \frac{\partial}{\partial \Phi} f_{\varepsilon;\Phi}.$$

Thus, the Jacobian is the sum of the transfer operator $\mathcal{L}_{\varepsilon;\varphi[\mu]}$ and a rank-one operator.

Since the $f_{\varepsilon,\Phi}$ are in class U_{NP} with uniform distortion bound (AD_δ) as defined in Section 2.2.1, we know from Chapter 2 that we can use an exponentially convergent Chebyshev Galerkin method to approximate the transfer operator. If μ is analytic, then the rank-one operator is bounded for any reasonable choices of function space, and its image lies in a space of analytic functions. Hence, we will be able to approximate the rank-one operator exponentially closely too, and so we can obtain not only a close approximation of the dynamics of G_ε but also those of its Jacobian. The same arguments imply that we can closely approximate most other relevant smooth properties such as response to perturbations and so on, although we do not give any formal proof.

To find a homoclinic tangency, we first need to find a fixed point μ^*_ε of the macroscopic dynamics G_ε for given ε. If we define $\mu^{\mathrm{static}}_{\varepsilon,\Phi}$ to be the acim of $f_{\varepsilon,\Phi}$, then $\mu^*_\varepsilon = \mu^{\mathrm{static}}_{\varepsilon,\Phi^*_\varepsilon}$ where Φ^*_ε solves

$$\varphi[\mu^{\mathrm{static}}_{\varepsilon,\Phi^*_\varepsilon}] - \Phi^*_\varepsilon = 0.$$

For $\varepsilon \approx 30$ there is a robust fixed point around $\Phi^*_\varepsilon \approx 0.051$, which can be accurately estimated using Chebyshev transfer operator methods, for example in Poltergeist (Wormell, 2017).

We can then determine the stable and unstable directions of the fixed point by numerically estimating the spectrum of the fixed point's Jacobian $D_{\mu^*_\varepsilon} G_\varepsilon$ using the

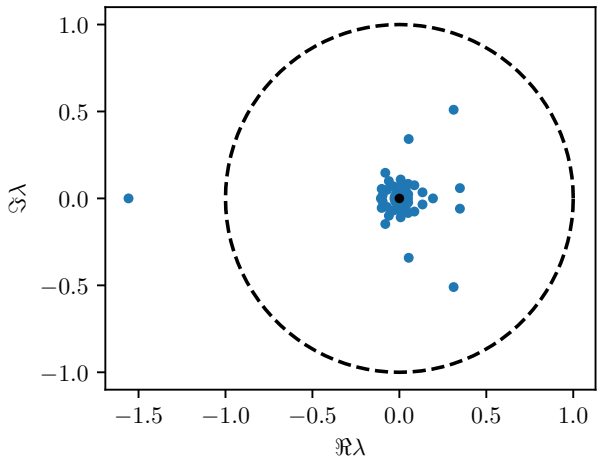

Figure 4.25. Spectrum of the Jacobian of G_ε at the fixed point for $\varepsilon = 30$, obtained using Poltergeist.jl (Wormell, 2017).

Chebyshev Galerkin methods. For ε around 30 we find that when we restrict to perturbations of zero total integral (as the) the Jacobian is hyperbolic, with one negative unstable eigenvalue λ^u_ε and the leading stable eigenvalues being a complex pair (see Figure 4.25).

The unstable subspace is therefore the span of the leading right eigenvector v^u_ε of $D_{\mu^*_\varepsilon} G_\varepsilon$, and the stable subspace is the kernel of the left eigendistribution w^s_ε of $D_{\mu^*_\varepsilon} G_\varepsilon$. Both of these objects can be accurately estimated by Chebyshev Galerkin transfer operator methods, with the convergence of w^s_ε occuring in a dual space.

We can use this to compute a local parametrisation of the unstable manifold $\mathcal{V}^u_\varepsilon := \mathcal{V}^u_{\mu^*_\varepsilon}$ via a parameter $a \in \mathbb{R}$:

$$\mathcal{V}^u_\varepsilon(a) = \mu^*_\varepsilon + v^u_\varepsilon a + \frac{1}{2} h_\varepsilon a^2 + \mathcal{O}(a^3), \tag{4.47}$$

where

$$h_\varepsilon = ((\lambda^u_\varepsilon)^2 \,\mathrm{id} - D_{\mu^*_\varepsilon} G_\varepsilon)^{-1} \mathrm{Hess}_{\mu^* \varepsilon} G_\varepsilon[v^u_\varepsilon, v^u_\varepsilon],$$

such that

$$\mathcal{V}^u_\varepsilon(\lambda^u_\varepsilon a) = G_\varepsilon(\mathcal{V}^u_\varepsilon(a)).$$

Thus, for any $a \in \mathbb{R}$, $\{\mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a) : m \in \mathbb{Z}\}$ is an orbit originating from μ^*_ε.

We now describe how to find a homoclinic tangency. If our orbit is to be a homoclinic orbit, then we need that

$$\lim_{m\to\infty} \mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a) - \mu^*_\varepsilon = 0. \tag{4.48}$$

If this homoclinic orbit is also to have a tangency between stable and unstable directions, we need additionally (since $\mathcal{V}^u_\varepsilon$ is also the unstable manifold of $\mathcal{V}^u_\varepsilon(a)$'s backward orbit out of the fixed point) that

$$\lim_{m\to\infty} \frac{\mathrm{d}}{\mathrm{d}a}\mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a) = 0. \tag{4.49}$$

which means for the homoclinic orbit that $\frac{\mathrm{d}}{\mathrm{d}a}\mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a)$ must become parallel with the stable subspace at the fixed point, and so

$$\lim_{m\to\infty} w^s_\varepsilon \mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a) = 0.$$

There are two parameters to vary to find a homoclinic tangency: the element of the orbit a and the parameter value ε. While it may seem natural to find points $\mathcal{V}^u_\varepsilon(a_{\text{homoclinic}}(\varepsilon))$ on homoclinic orbits for each ε and vary ε to obtain a stable-unstable manifold tangency, this is problematic as $a_{\text{homoclinic}}(\varepsilon)$ has a fold singularity at the homoclinic tangency. Instead, we first found values $a(\varepsilon)$ such that $\frac{\mathrm{d}}{\mathrm{d}a}\mathcal{V}^u_\varepsilon((\lambda^u_\varepsilon)^m a)$ lies in the kernel of w^s_ε for some large m (i.e. find a "tangency"), and then varied ε to obtain a homoclinic orbit.

4.8.4. Results. We estimated an external parameter $\varepsilon = \varepsilon_{\text{ht}}$ that manifested a homoclinic tangency and a point on the homoclinic orbit $\mathcal{V}^u_\varepsilon(a_{\text{ht}})$ to be

$$\begin{aligned}\varepsilon_{\text{ht}} &= 30.06183139229653\\ a_{\text{ht}} &= 7.027598895378105 \times 10^{-5},\end{aligned}$$

where the fixed point's Jacobian's unstable eigenfunction v^u_ε is normalised so that its first Chebyshev coefficient is positive and the ℓ^2 norm of its Chebyshev coefficients is 1. We obtained a good guess of a homoclinic orbit via random search, and then applied a bisection algorithm to obtain the values $(a_{\text{ht}}, \varepsilon_{\text{ht}})$. Again, we used our implementation of the Chebyshev Galerkin transfer operator methods in Poltergeist.jl for computation (Wormell, 2017).

The homoclinic orbit is plotted in a projection to mean-field delay coordinates in Figure 4.26. In Figure 4.27 are plotted norms of the distance from the homoclinic orbit and of unstable tangent vector under the action of the Jacobian, confirming the limits (4.48-4.49) constitutive of a homoclinic tangency. In our numerical results these norms decay over the orbit by around seven decimal digits from their maximum value: this is around half of the number of digits of accuracy (≈ 13) available using the Chebyshev methods at double floating point accuracy, a result of the rather

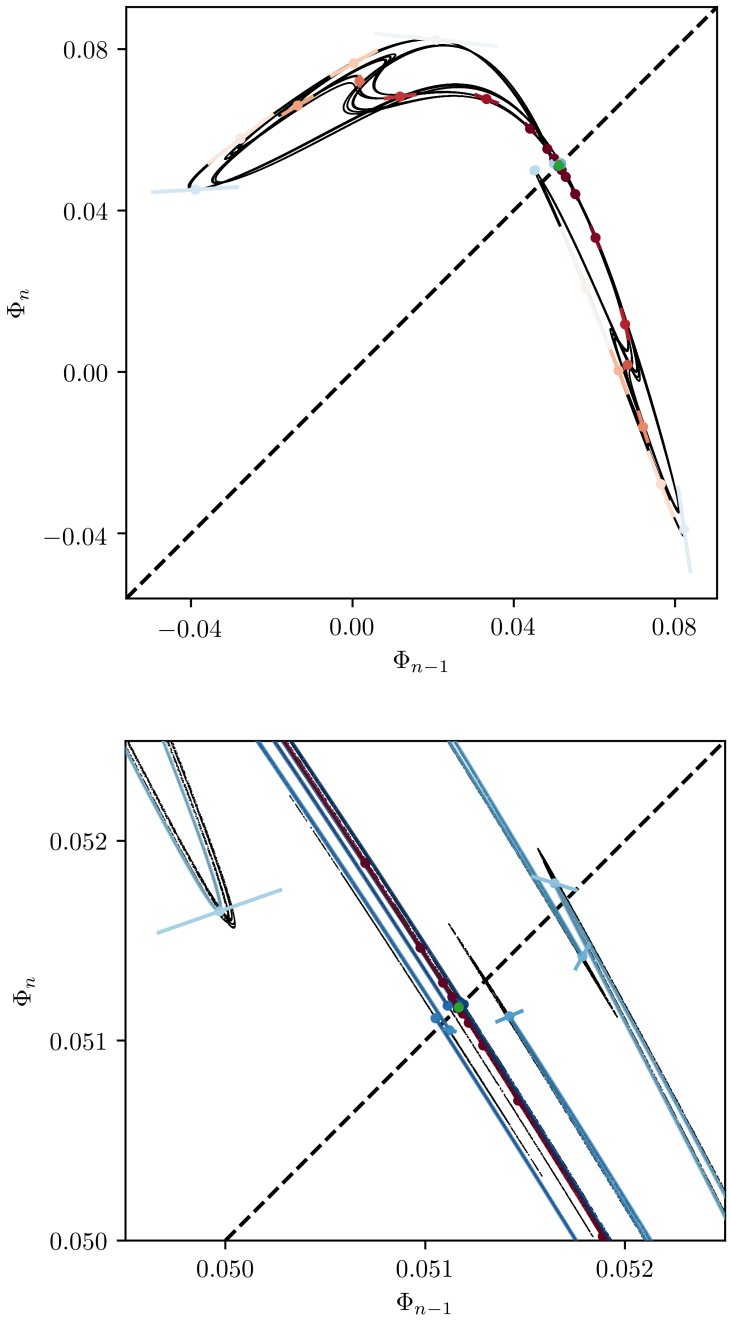

FIGURE 4.26. Plot of the homoclinic tangency (red to blue) for $\varepsilon = \varepsilon_{\text{ht}}$ at two different scales. The attractor is plotted in black and the fixed point in green. At bottom, the unstable manifold $\mathcal{V}^u_\varepsilon$ is locally plotted around each point.

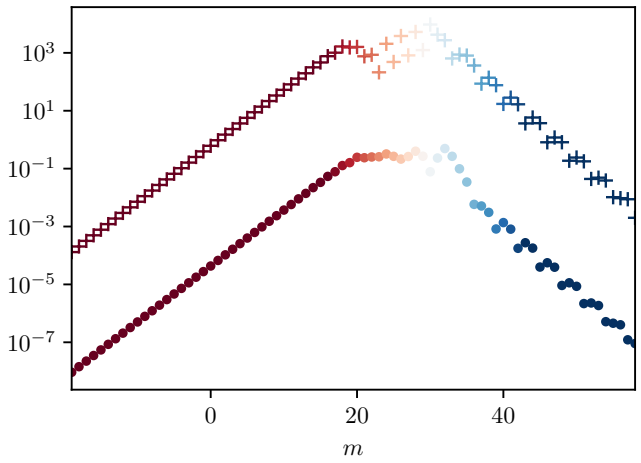

FIGURE 4.27. In dots, the H^1 norm of the distance from the homoclinic orbit to the fixed point, as in (4.48). In crosses, the H^1 norm of the unstable tangent vector iterated under G, as in (4.49).

primitive shooting continuation method used. This method of continuation attempts to choose the parameter a so as to "hit" the saddle fixed point μ^*_ε; numerical round-off errors are then magnified by expansion in the unstable direction, leading to reduced accuracy. However, because the relevant parameters and geometry in the system are $\mathcal{O}(1)$ (c.f. Figures 4.22(b) and 4.26), our results are clear-cut evidence for a homoclinic tangency at this parameter value. This accuracy could be improved through the use of extended precision floating point arithmetic and/or through a continuation method based on boundary-value solutions.

Finally, we give numerical evidence that the fixed point actually lies on the attractor in Figure 4.28, and we note that the dynamics at ε_{ht} are chaotic with leading Lyapunov exponent $\lambda_1 \approx 0.199 > 0$. Because the attractor contains the fixed point and the system is chaotic (so cannot be confined to the stable manifold of the fixed point), the attractor must also contain the unstable manifold of the fixed point and thus the homoclinic tangency. As a result, we can conclude that the system is non-hyperbolic and thus not consistent with the chaotic hypothesis as introduced in Section 4.1.

Furthermore, because non-hyperbolic maps are not guaranteed to have linear response, and furthermore because homoclinic tangencies cause poor behaviour of the map under perturbation—for example, structural instability and the emergence of an infinite number of sinks under generic perturbations (Palis et al., 1995)—this failure of non-hyperbolicity is likely behind the failure of linear response observed in

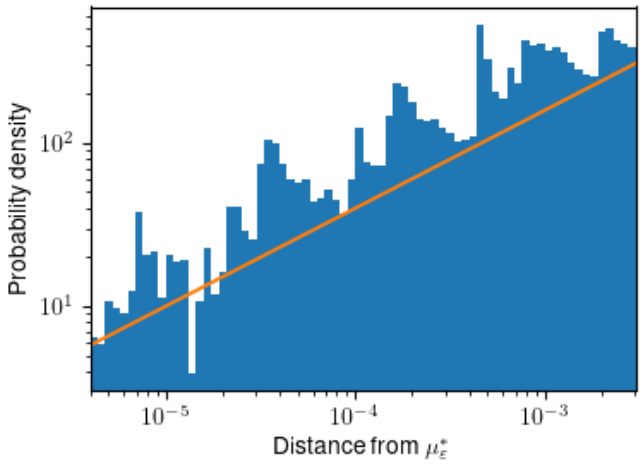

FIGURE 4.28. In blue, a histogram of L^2 distances from the fixed point μ^*_ε of a time series of iterates of $G_{\varepsilon_0 \mathrm{ht}}$. A time series was used containing $400,000$ iterates that had L^2 distance from $\mu^*_{\varepsilon_{\mathrm{ht}}}$ of less than 0.003. The observed density is consistent with a local correlation dimension of around 1.6 at the fixed point (orange line).

Section 4.7.3 (c.f. Figure 4.17).

To make the computation of the homoclinic tangency strictly rigorous, two clear theoretical challenges arise. The first is to rigorously validate the numerical findings: while there is good reason to believe that our use of Chebyshev methods here is justified, much of the theoretical work has yet to be done, in particular in bounding convergence of the spectral data of the Jacobian DG (see 4.45-4.46) and obtaining rigorous bounds on local approximations of the stable and unstable manifolds of the fixed points. The second is to demonstrate that comparable behaviour holds around a homoclinic tangency in this transfer operator setting as in finite dimensions. Although transfer operators are non-invertible and highly non-normal, in appropriate Hardy spaces the transfer operator is compact with good control of approximation numbers (Bandtlow and Jenkinson, 2008): this may provide an opening for analysis.

4.9. Discussion

In this chapter, we have established conditions under which macroscopic mean field observables enjoy linear response. We considered three scenarios: macroscopic observables of an uncoupled collection of microscopic subunits, observables of macroscopic distinguished variables driven by an uncoupled microscopic collection,as well

as macroscopic observables of microscopic subunits which are coupled via their mean field. We found that linear response is possible even in the case when the microscopic systems individually violate LRT, provided the microscopic dynamics is heterogeneous with parameters drawn from a sufficiently smooth distribution. We also found that for back-coupled systems of finite size, LRT (for small enough perturbations) is expected for any kind of microscopic dynamics, independent of the smoothness of $\nu(a)$: this can be understood as the result of emergent, self-generated stochastic effects. We further established that in the thermodynamic limit of infinite M, the mean field dynamics of a self-coupled system can exhibit attracting dynamics that appears non-uniformly hyperbolic and certainly fails to have LRT, even when the microscopic subunits are individually uniformly expanding; this presents a counter example to the widely believed hypothesis that macroscopic observables of high-dimensional systems typically obey linear response.

Furthermore, we have found a numerical example of a chaotic system whose macroscopic dynamics in the thermodynamic limit has homoclinic tangencies and are thus non-hyperbolic. This violates the chaotic hypothesis of Gallavotti-Cohen as formulated in Gallavotti (2019), which posits that macroscopic dynamics should evolve as if they were hyperbolic. That said, what is more commonly claimed is that "typical" dynamics are hyperbolic, and we do not claim that these mean-field coupled systems (or any mean-field coupled systems) are necessarily typical complex systems. Certainly many chaotic systems that naturally occur tend to be locally rather than globally coupled, which would mean the kind of central limit theorem-type noise as discussed in the previous section would persist in the thermodynamic limit, and thus these systems should obey LRT and have many of the desirable statistical properties associated with stochastic dynamics.

Our results rely on the existence of statistical limit laws such as the central limit theorem. These are proved for strongly chaotic systems, and in particular for uniformly expanding maps as well as for smooth unimodal maps. We follow here Gottwald and Melbourne (2014) and assume that typical dynamical systems are strongly chaotic and hence enjoy good statistical properties, so that our results carry over to typical dynamical microscopic systems. To ensure the existence of the central limit theorem and the convergence of the deviations ζ_n to a Gaussian process with decay of correlations, we require the observables $\psi(q)$ to be at least Hölder continuous and the external forcing ε to be such that the perturbed system is mixing.

We presented here results for mean field observables Ψ of the form (4.2). We remark that our results carry over for more general (e.g. weighted) mean field variables provided those weights are sufficiently smoothly distributed, and indeed we expect broadly similar results for more general "macroscopic" observables.

To reduce the complexity of expression we have enforced mixing dynamics, with no chaotic synchronisation, for example by including the hidden r-dynamics in (4.7). It would be interesting to study the case when the microscopic dynamics is not restricted in this way, for example if periodic dynamics were allowed. We have only discussed the existence of LRT and have not considered fluctuation-dissipation formulae to provide a compact analytical formula for the response term. This may require treatment of the non-Markovian dynamics of the macroscopic variable as well as the interplay of the perturbed microscopic dynamics and the macroscopic dynamics, the latter having been studied in the context of slow-fast systems (Abramov, 2010).

We have corroborated our findings with detailed numerical simulations and have provided several heuristic arguments based on statistical limit laws: we hope these arguments can be made rigorous. We remark that in the case of uncoupled systems, rather than averaging over the heat bath, one may express the invariant measures $\mu^{(a_j}$ as an infinite sum over the systems' unstable periodic orbits (UPOs) (Ruelle, 2004; Pollicott, 1986; Cvitanovic and Eckhardt, 1991; Eckhardt and Grossmann, 1994). This approach may be more effective when trying to rigorously justify the convergence of the deviations ζ_n to a Gaussian process. It seems likely also possible to apply ideas from a recent argument by Ruelle for linear response in non-hyperbolic systems, based on the statistical smearing-out of singularities in the physical measures (Ruelle, 2018).

Chapter 5

Conclusion

This thesis is concerned with the statistical behaviour of simple model systems. It speaks to two fundamental questions: how we can study chaotic model systems numerically, and to what extent we can use results from simple model systems to extend our knowledge of general chaotic dynamics.

In the first part of the thesis, we applied ourselves to the former question through the development of highly accurate numerics for Markovian uniformly expanding and intermittent one-dimensional systems. The numerical methods were highly accurate as well as being applicable to a variety of problems, while having low computational overhead. We accomplished this through discretisation of appropriate transfer operators. Key to the fast convergence of the transfer operator method was the discretisations used: we chose projections onto orthogonal spectral basis functions (Fourier bases and Chebyshev polynomials). Using these discretisations we were able to make use of the smooth structure of the underlying dynamical system, which facilitated fast convergence of the numerical estimates. The spectral discretisation also facilitates the estimation of statistical properties such as linear response where higher regularity function spaces are required. At the point of these methods' development, transfer operator numerics using these bases were new in the literature, although they have since been studied in two-dimensional systems Crimmins and Froyland (2019a); Slipantschuk et al. (2019).

The efficiency of the spectral transfer operator methods are manifested in the publically available software package Poltergeist: we have seen that it can compute acims of analytic uniformly-expanding systems to 13 decimal places of accuracy, in only a small fraction of a second. While implementations of Ulam's method have existed for a long time (Dellnitz et al., 2001) they are subject to the usual limitations of Ulam's method, in particular slow convergence. To obtain accuracy of up to 13 decimal places is impossible using Ulam's method given current computational resources. This level of accuracy allows one to effectively perform complex or difficult numerical work: for example, numerical continuation without boundary value solvers necessitates a loss of a large fraction of the number of digits of accuracy: to obtain a

credible level of accuracy in the output it is therefore necessary to be able to numerically evaluate the system to a high degree of accuracy. Poltergeist also incorporates the user-friendliness of modern numerical function approximation packages (Driscoll et al., 2014; Olver, 2019): optimised estimates of relevant statistical properties it is necessary only to input the map and the domains of its branches.

Furthermore, because the conceptual basis of the spectral methods developed in this thesis is simply the discretisation of linear-algebraic problems, the same ideas that were treated rigorously in Chapter 2 may be applied to other questions that can be posed in terms of transfer operators. A productive illustration of this was given in Section 4.7, where we probed the dynamics and response theory of a high-dimensional system that could be expressed in terms of the transfer operators of a 1D map, and then in Section 4.8 where we calculated a homoclinic tangency via numerical continuation. Continuation as applied in Section 4.8 is a difficult numerical problem (Krauskopf et al., 2007), and it is due to the Chebyshev methods' high accuracy and good regularity properties that it was possible to obtain convincing evidence of a homoclinic tangency in this system. So while the choice of a globally coupled system allowed for the mean-field reduction, the finding of a provable violation of the chaotic hypothesis was facilitated by our numerical techniques.

For intermittent maps, the other component, apart from transfer operator methods, that enabled fast convergence of estimates of statistical properties, was the efficient numerical evaluation of induced maps and their associated transfer operators. Key to this was the Abel function, which allowed the induced maps to be expressed in close form. Although Abel functions have been studied in the statistical branching process literature, they have not hitherto been applied to the study of intermittent maps. Previously existing numerics were only effective for small values $\alpha \lesssim 0.4$ (Murray, 2010; Galatolo and Nisoli, 2014), and failed completely in the infinite ergodic setting $\alpha \geq 1$. Hence, in particular our numerics furnish the ability to numerically explore infinite ergodic theory.

Thus, as illustrated by the homoclinic tangency, the methods we presented in the first part of the thesis enable theoretical advances that would have been difficult to come to or visualise by hand and to which previously existing numerics were not appliable.

In Chapter 4 we then sought to understand the question of the relationship between the statistical properties of more complex, high-dimensional chaotic systems and low-dimensional systems, largely in the field of linear response theory. Because higher dimensional systems are complicated and very few things may be proved about them, we took a less rigorous approach: our work was based on statistical reductions that were numerically justified rather then strictly rigorous. We started with large collections of simple, well-understood microscopic systems and coupled them

together in some structurally very basic ways: in all-to-one (Section 4.6) and all-to-all (Section 4.7) networks. We then investigated the dynamics and LRT properties of these systems: to explore the effect of the systems' microscopic LRT properties on the macroscopic response, we considered variations of the models with different microscopic constituents.

In this investigation we uncovered several new phenomena. One of these was the collective structural stability to perturbations: out of an appropriately inhomogeneous collection of microscopic systems that were structurally unstable and violated LRT, linear response emerged for the mean field. We also observed that moderately sized mean fields of deterministic chaotic variables approximated stochastic processes, to the extent that driving another system with such a mean field induced linear response in that system. This effect was particularly striking in the case where a finite sized system was self-coupled via a mean-field: regardless of their microscopic constituents these, the macroscopic observables obeyed LRT. This aligns with recent results on LRT for stochastic systems (Hairer and Majda, 2010).

These phenomena are largely due to the network structure of the systems, relying on the mean-field reduction (which was possible due to the all-to-all coupling) rather than emerging organically from the dynamics. While this may be to some extent an artefact of the method of enquiry, in some sense these phenomena are therefore intrinsically high-dimensional, and we must conclude that certain different factors are at play in high-dimensional chaotic dynamics than in low-dimensional ones.

However, while distinguished variables driven by heat baths and globally coupled systems are both natural first steps in understanding high-dimensional dynamics, they cannot be seen as representative of many natural dissipative systems such as the climate system, which has multiple spatial and time scales and where a lot of interactions are local rather than global (Hasselmann, 1976; Palmer and Williams, 2010; Dijkstra, 2005). While our mean-field coupled thermodynamic limit system behaved like a low-dimensional deterministic chaotic system, it may be an atypical example because of its global structure. Because of the propensity in real systems for local (i.e. low network degree) interactions, most high-dimensional chaotic systems are likely more like stochastic systems than simple deterministic chaotic systems. Consequently, LRT in macroscopic systems such as turbulent fluids and the Earth's climate may largely arise as a result of emergent stochastic effects.

It would therefore be interesting to study high-dimensional coupled chaotic systems with different network structures that better model real world systems. We imagine that the phenomena we have studied in this work could be influential in determining whether these systems satisfy LRT. We remark that the fundamental assumption of such a study is that coupled networks are an effective model for complex high dimensional chaotic systems (Tsonis et al., 2006; Donges et al., 2009).

Appendix A

Some results relevant for Chapter 2

A1. The relationship between uniform expansion and uniform C-expansion

In Chapter 2 we stipulated that maps on non-periodic domains satisfy a so-called uniform C-expansion condition rather than the usual uniform expansion condition. Neither of these conditions imply the other: in fact it is not hard to construct non-pathological examples of uniformly-expanding maps which are not uniformly C-expanding (see Figure A.1).

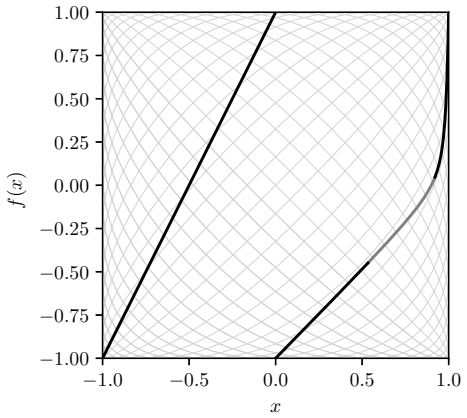

Figure A.1. In black, an example of a map $f \in U^u_{NP}$ ($\lambda = 0.98^{-1}$) which is not uniformly C-expanding ($\check{\lambda} \approx 0.763$). Non-C-expanding parts of f marked in mid grey. In light grey, lines of unit C-expansion (i.e. curves $\psi(x)$ for which $(\cos^{-1} \circ \psi \circ \cos)' = \pm 1$).

However in practice, maps in U^u_{NP} are generally also in U_{NP}. For example, all piecewise linear maps in U^u_{NP} lie in U_{NP}. (In particular, if f is the k-tupling map, the uniform C-expansion parameter for f is $\check{\lambda} = \sqrt{k}$.) A map in U^u_{NP} typically fails to be in U_{NP} if the tangent to the graph points towards the sides of the graph's plot. This necessitates a kink in the map, so as to preserve the Markov structure. An example of the situation is illustrated in Figure A.1.

However, if we consider only maps that are Markov with bounded distortion, we find close connections between C-expansion and classical expansion. In fact, a positive lower bound on one implies a positive lower bound on the other, which may be seen by an adaptation of the proof of Theorem A.1 below. Importantly, uniformly C-expanding maps eventually become uniformly expanding under iteration and vice versa, according to the following theorem.

Theorem A.1. *Suppose $f \in U^u_{NP}$ (resp. $f \in U_{NP}$). Then there exists $n_* \in \mathbb{N}$ such that $f^n \in U_{NP}$ (resp. U^u_{NP}) for all $n \geq n_*$. Each f^n satisfies the same distortion conditions as f, with possibly different constants.*

Remark A.2. *Since iterates of a map have an exponentially growing number of branches, for computational purposes it may be more effective to simply compute a conjugacy of a map which is C-expanding.*

It is in fact possible to construct, for a map $f \in U^u_{NP}$ (resp. U_{NP}), an analytic diffeomorphism η_f such that $f_c = \eta_f \circ f \circ \eta_f^{-1} \in U_{NP}$ (resp. U^u_{NP}).

Furthermore, if f satisfies (DD_r) *then so will f_c, and if f satisfies* (AD_δ) *for some $\delta > 0$ then there exists $\delta' > 0$ for which f_c satisfies* ($\mathrm{AD}_{\delta'}$).

Thus, maps in U^u_{NP} and in U_{NP} have the same dynamical properties and can additionally be converted from one class to the other. We emphasise that the crucial assumption here is bounded distortion.

We now prove the results stated above, beginning with Theorem A.1.

Proof of Theorem A.1. Suppose $f \in U^u_{NP}$ with $|f'| > \lambda$ and distortion constant C_1. Then $f^n \in U^u_{NP}$ with $|(f^n)'| > \lambda^n$ and distortion constant bounded by $C_1 \frac{1-\lambda^{-n-1}}{1-\lambda^{-1}}$ (Gouëzel, 2004a). Let us use the notation $f^n = g$ with branches $\mathcal{P}_\iota, \iota \in I^n$.

Suppose $x \in \mathcal{P}_\iota$ for some $\iota \in I^n$. Then

$$\frac{1-|x|}{|\operatorname{sgn}(xg'(x)) - g(x)|} \geq \frac{v_{\iota_x}(\operatorname{sgn}(xg'(x))) - x}{|g(v_{\iota_x}(\operatorname{sgn}(xg'(x)))) - g(x)|},$$

and by the intermediate value theorem there exists $w \in \mathcal{P}_\iota$ such that

$$\frac{v_{\iota_x}(\operatorname{sgn}(xg'(x))) - x}{|g(v_{\iota_x}(\operatorname{sgn}(xg'(x)))) - g(x)|} = \frac{1}{|g'(w)|}.$$

Using Lemma A.3(a) we find that

$$\frac{1}{|g'(w)|} > \frac{e^{-2C_1}}{|g'(x)|},$$

and so for all $x \in \cup_{\iota\in I}\mathcal{O}_\iota$,

$$\sqrt{\frac{1-x^2}{1-(g(x))^2}}|g'(x)| \geq \sqrt{\frac{1+|x|}{|\operatorname{sgn}(xg'(x))+g(x)|}e^{-2C_1}|g'(x)|} \geq \sqrt{\frac{1}{2}e^{-2C_1}\lambda^n} > 1$$

for n sufficiently large.

The map f^n is full-branch Markov with bounded distortion, by the above satisfies (CE), and by Lemma A.5 satisfies (P). Consequently, $f^n \in U_{NP}$.
Now, suppose $f \in U_{NP}$ and let $n \in \mathbb{N}^+$. For the remainder of this proof, we will use subscript notation for forward iterates: $x_n = f^n(x)$. We additionally call two points $x, y \in \Lambda$ *n-companions* if there exists a sequence $(\iota_j)_{j=1,\dots,n}$ such that $x_{j-1}, y_{j-1} \in \mathcal{O}_{\iota_j}$ for $j \leq n$.

Given $x \in \Lambda$, choose y, z such that x, y and z are n-companions, $y_n = -1$ and $z_n = 1$. Then by the mean value theorem, there exists w between y and z such that

$$|(f^n)'(w)| = \frac{|z_n - y_n|}{|z-y|} \geq \frac{2}{\pi}\check{\lambda}^n. \tag{A.1}$$

Now, since w lies between y and z, it is an n-companion of x, y and z. We will therefore relate $|(f^n)'(w)|$ to $|(f^n)'(x)|$ using bounded distortion. We expand their quotient out using the chain rule and rewrite:

$$\begin{aligned}\frac{|(f^n)'(x)|}{|(f^n)'(w)|} &= \prod_{j=1}^{n}\frac{|f'(x_{j-1})|}{|f'(w_{j-1})|}\\ &= \prod_{j=1}^{n}\frac{|v'_{\iota_j}(x_j)^{-1}|}{|v'_{\iota_j}(w_j)^{-1}|}\\ &= e^{\sum_{j=1}^{n}\left(\log|v'_{\iota_j}(w_j)|-\log|v'_{\iota_j}(x_j)|\right)}\\ &\geq e^{-\sum_{j=1}^{n}\left|\log|v'_{\iota_j}(w_j)|-\log|v'_{\iota_j}(x_j)|\right|}. \end{aligned}\tag{A.2}$$

We then bound the summands using (DD_1) and the fact that $v''_\iota/v'_\iota = (\log|v'_\iota|)'$:

$$\left|\log|v'_{\iota_j}(w_j)| - \log|v'_{\iota_j}(x_j)|\right| \leq C_1|w_j - x_j| \leq C_1\check{\lambda}^{j-n}\pi.$$

The sum in (A.2) can thus be collapsed to give

$$\frac{|(f^n)'(x)|}{|(f^n)'(w)|} \geq e^{-\sum_{j=1}^{n}C_1\check{\lambda}^{j-n}\pi} > e^{-C_1\pi(1-\check{\lambda})^{-1}}.$$

Combining this with (A.1) gives us that

$$|(f^n)'(x)| \geq e^{-C_1\pi(1-\check{\lambda})^{-1}}\frac{2}{\pi}\check{\lambda}^n,$$

which implies that f^n is uniformly expanding for sufficiently large n. Since as before f^n satisfies all the non-expansion conditions to be in U^u_{NP}, we have that $f^n \in U^u_{NP}$. □

A2. Results on conditions (DD_1) and (P)

In this appendix, we prove some properties possessed by maps in U_{NP} used in Chapter 2. We first give some "non-local" properties of the bounded distortion condition (DD_1), and then prove that (P) is preserved under iteration.

We first prove a lemma relating bounded distortion constants to bounds on derivatives of the map. The properties summarised in Lemma A.3 are mostly standard, but we improve the upper bound in part (b) from the exponentially large e^{2C_1} to a computationally more useful $1 + 2C_1$.

Lemma A.3. *Suppose $f : [-1, 1] \to [-1, 1]$ is full-branch Markov with bounded distortion. Suppose the distortion constant of f is C_1. Then for all $\iota \in I$:*
(a) For all $x, w \in [-1, 1]$,

$$e^{-2C_1} \leq \frac{|v'_\iota(x)|}{|v'_\iota(w)|} \leq e^{2C_1}.$$

(b) For all $x \in [-1, 1]$,

$$e^{-2C_1}\frac{|\mathcal{O}_\iota|}{2} \leq |v'_\iota(x)| \leq (1 + 2C_1)\frac{|\mathcal{O}_\iota|}{2}.$$

Proof of Lemma A.3. Part (a) is a standard result (Gouëzel, 2004a; Korepanov et al., 2016).

To prove (b), we have that as a result of the intermediate value theorem there exists some $w \in [-1, 1]$ such that

$$v'_\iota(w) = \frac{v_\iota(1) - v_\iota(-1)}{2} = \frac{|\mathcal{O}_\iota|}{2}.$$

By part (a), $e^{-2C_1} \leq \frac{|v'_\iota(x)|}{|v'_\iota(w)|}$. Additionally, the fundamental theorem of calculus gives that

$$v'_\iota(x) = v'_\iota(w) + \int_w^x v''_\iota(\xi)d\xi,$$

and consequently

$$\begin{aligned} |v'_\iota(x)| &\leq |v'_\iota(w)| + \int_x^w C_1|v'_\iota(\xi)||d\xi| \\ &\leq \frac{|\mathcal{O}_\iota|}{2} + C_1\int_{-1}^1 |v'_\iota(\xi)|d\xi = \left(\frac{1}{2} + C_1\right)|\mathcal{O}_\iota|, \end{aligned}$$

as required.

□

Remark A.4. *Similarly, suppose that a map $f \in \bar{U}_{NP}$ obeys analytic distortion condition (AD_δ) with constant $C_{1,\delta}$. Then for all $x, w \in \check{\Lambda}_\delta$,*

$$e^{-2C_{1,\delta}\cosh\delta} \leq \frac{|v_\iota'(x)|}{|v_\iota'(w)|} \leq e^{2C_{1,\delta}\cosh\delta}.$$

We now prove that the partition spacing condition (P) is preserved under composition. Consequently, U_{NP} and $\bar{U}_{NP}$ are closed under composition.

Lemma A.5. *Suppose f and g are Markov maps on $[-1,1]$ satisfying (P), and that in addition f has bounded distortion with parameter $C_1^{(f)}$ and g has uniform expansion parameter $\lambda^{(g)} > 0$.*

Then $g \circ f$ satisfies (P).

Proof. Let $\mathcal{O}_{\phi\gamma} = v_\phi^{(f)}\left(v_\gamma^{(g)}(\Lambda)\right)$ be a branch set of $g \circ f$. Let $p \in \partial\Lambda$, i.e. $p = \pm 1$.

Since $\mathcal{O}_{\phi\gamma} = v_\phi^{(f)}\left(\mathcal{O}_\gamma^{(g)}\right)$, by Lemma A.3(c) we have

$$\frac{|\mathcal{O}_{\phi\gamma}|}{\left|\mathcal{O}_\gamma^{(g)}\right|} \leq \left(1 + 2C_1^{(f)}\right)\frac{\left|\mathcal{O}_\phi^{(f)}\right|}{2},$$

and thus a preliminary bound on our ratio of interest:

$$\frac{|\mathcal{O}_{\phi\gamma}|}{d(\mathcal{O}_{\phi\gamma}, p)} \leq \frac{\left(1 + 2C_1^{(f)}\right)\frac{1}{2}\left|\mathcal{O}_\gamma^{(g)}\right|\left|\mathcal{O}_\phi^{(f)}\right|}{d(\mathcal{O}_{\phi\gamma}, p)}. \tag{A.3}$$

We are interested in intervals for which $p \notin \mathcal{O}_{\phi\gamma}$. If $p \in \mathcal{O}_{\phi\gamma}$, then we need $p \in \mathcal{O}_\phi^{(f)}$ and $\hat{f}_\phi(p) =: \tau \in \mathcal{O}_\gamma^{(g)}$. Note that since $p \in \partial\mathcal{O}_\phi^{(f)}$, then $\tau \in \partial\Lambda$. Therefore, intervals $\mathcal{O}_{\phi\gamma}$ which do not contain p either have $p \notin \mathcal{O}_\phi^{(f)}$ or $\tau \notin \mathcal{O}_\gamma^{(f)}$.

We split into cases accordingly. In the first case where $p \notin \mathcal{O}_\phi^{(f)}$, we have that since $\mathcal{O}_\gamma^{(g)} = v_\gamma^{(g)}([-1,1])$, its length must be less than $2/\lambda^{(g)}$. Since $\mathcal{O}_{\phi\gamma} \subseteq \mathcal{O}_\phi^{(f)}$, we must have $d\left(\mathcal{O}_{\phi\gamma}, p\right) \geq d\left(\mathcal{O}_\phi^{(f)}, p\right)$. Therefore from (A.3),

$$\frac{|\mathcal{O}_{\phi\gamma}|}{d(\mathcal{O}_{\phi\gamma}, p)} \leq (1 + 2C_1^{(f)})\frac{1}{\lambda^{(g)}}\boldsymbol{\xi}^{(f)},$$

where we used that $|\mathcal{O}_\gamma^{(g)}| < |\Lambda|/\lambda^{(g)}$ from the expansion assumption.

For the second case, let $q \in \partial\Lambda$ be such that $v_\phi^{(f)}(q)$ lies in between $\mathcal{O}_{\phi\gamma}$ and p and let $r \in \partial\mathcal{O}_\gamma^{(f)}$ such that $v_\phi^{(f)}(r)$ is the nearest point in $\mathcal{O}_{\phi\gamma}$ to p (and thus, q). Then

$d(\mathcal{O}_{\phi\gamma}, p)$ is the length of the interval $[v_\phi^{(f)}(r), p]$, which is bigger than the length of the interval $[v_\phi^{(f)}(r), v_\phi^{(f)}(q)] = v_\phi^{(f)}([r,q])$. Using Lemma A.3 the length of this last interval can be bounded:

$$\frac{\left|v_\phi^{(f)}([r,q])\right|}{|[r,q]|} = \frac{\int_r^q \left|(v_\phi^{(f)})'(x)\right| |dx|}{|[r,q]|} \geq e^{-2C_1^{(f)}} \frac{1}{2} \left|\mathcal{O}_\phi^{(f)}\right|.$$

Furthermore, the distance between r and q is precisely $d(\mathcal{O}_\gamma^{(g)}, q)$.

Combining these results with (A.3), we find that

$$\frac{|\mathcal{O}_{\phi\gamma}|}{d(\mathcal{O}_{\phi\gamma}, p)} \leq \left(1 + 2C_1^{(f)}\right) e^{2C_1^{(f)}} \frac{\mathcal{O}_\gamma^{(g)}}{d(\mathcal{O}^{(g)}, q)} \leq \left(1 + 2C_1^{(f)}\right) e^{2C_1^{(f)}} \boldsymbol{\xi}^{(g)}.$$

Combining the two cases, then, we find that

$$\boldsymbol{\xi}^{(g\circ f)} \leq \left(1 + 2C_1^{(f)}\right) \max\left\{e^{2C_1^{(f)}} \boldsymbol{\xi}^{(g)}, \frac{1}{\lambda^{(g)}} \boldsymbol{\xi}^{(f)}\right\},$$

as required. □

A3. Proof of Lemma 2.12

In this appendix we will prove Lemma 2.12, which states that standard properties of f (e.g. differentiability of the distortion) imply the properties of $\cos^{-1} \circ f \circ \cos$ required to apply Lemma 2.11 in the proof of Theorem 2.3.

We remark that while the partition position condition (P) is crucial for the proof in general, it is not necessary if one restricts to maps that only satisfy (DD_1).

We also emphasise that this lemma gives very loose bounds for the Υ_n and $\Upsilon_{1,\delta}$, and that in practice one is best served by calculating these constants directly from the ν_ι.

We will first state and prove two lemmas upon which Lemma 2.12 relies, and then prove the latter.

Lemma A.6. *Suppose the map $f \in \bar{U}_{NP}$ is piecewise C^{n+1}, choose $\sigma \in \partial\Lambda = \{\pm 1\}$ and let $\tau = v_\iota(\sigma)$. Define the gradient of the chord*

$$\hat{S}_{\iota,\sigma}(x) = \frac{\tau - v_\iota(x)}{\sigma - x}.$$

Then

$$\hat{S}_{\iota,\sigma}^{(n)}(w) = \frac{1}{n+1} v^{(n+1)}(w)$$

for some w directly between σ and x.

Proof. We can show by induction that for all $n \geq 0$

$$\hat{S}^{(n)}_{\iota,\sigma}(x) = n! \frac{\tau - \sum_{m=0}^{n} \frac{1}{m!} v_\iota^{(m)}(x)(\sigma - x)^m}{(\sigma - x)^{n+1}}.$$

Since $v_\iota(\sigma) = \tau$, the lemma follows by Taylor's theorem. □

Lemma A.7. *Suppose $f \in \bar{U}_{NP}$, $\iota \in I$ and $\sigma \in \partial\Lambda = \{\pm 1\}$ such that $v_\iota(\sigma) \notin \partial\Lambda$. Let $\tau_{\iota,\sigma} = \sigma \operatorname{sgn} v_\iota'(0)$ and*

$$T_{\iota,\sigma}(x) = 1 - v_\iota(x)/\tau_{\iota,\sigma}.$$

Then for $x \in [-1,1]$, $T_{\iota,\sigma}(x) \geq \boldsymbol{\xi}^{-1}|\mathcal{O}_\iota|$. If f satisfies analytic distortion condition (AD_δ) then there exists $\zeta \in (0,\delta]$ and $\mathfrak{K}_\zeta > 0$ such that for $z \in \check{\Lambda}_\zeta$ $|T_{\iota,\sigma}(z)| \geq \mathfrak{K}_\zeta|\mathcal{O}_\iota|$.

Proof. Recalling that $\tau_{\iota,\sigma}^2 = 1$ we can write

$$T_{\iota,\sigma}(x) = \tau_{\iota,\sigma}(\tau_{\iota,\sigma} - v_\iota(\sigma)) + \tau_{\iota,\sigma}(v_\iota(\sigma) - v_\iota(x)).$$

Since $\tau_{\iota,\sigma} - v_\iota(\sigma)$ has the same sign as $\tau_{\iota,\sigma}$, the first term can be written as a positive quantity $|\tau_{\iota,\sigma} - v_\iota(\sigma)|$ which is equal to $d(\tau_{\iota,\sigma}, \mathcal{O}_\iota)$. By the partition spacing condition (P), we have $d(\tau_{\iota,\sigma}, \mathcal{O}_\iota) \geq \boldsymbol{\xi}^{-1}|\mathcal{O}_\iota|$.

Furthermore, one can apply Taylor's theorem to the second term to get that $\tau_{\iota,\sigma}(v_\iota(\sigma) - v_\iota(x)) = \tau_{\iota,\sigma}(\sigma - x)v_\iota'(w)$ for some w between x and σ. Since v_ι' keeps its sign on $[-1,1]$ and the sign of $\sigma - x$ is simply the sign of σ, the definition of $\tau_{\iota,\sigma}$ means that the second term is positive on $[-1,1]$. Thus, for $x \in [-1,1]$, $T_{\iota,\sigma}(x) \geq \boldsymbol{\xi}^{-1}|\mathcal{O}_\iota|$.

On the analytic domain $\check{\Lambda}_\zeta$ the situation is more complicated. We write that

$$\Re T_{\iota,\sigma}(x) = d(\tau_{\iota,\sigma}, \mathcal{O}_\iota) - \tau_{\iota,\sigma}\Re\left(v_\iota(x) - v_\iota(\Re x)\right) - \tau_{\iota,\sigma}\left(v_\iota(\Re x) - v_\iota(\sigma)\right), \tag{A.4}$$

and bound terms from below.

Set $\mathfrak{C}_\zeta = e^{2C_{1,\delta}\cosh\zeta}$. We have that for any point x in $\check{\Lambda}_\zeta$, $|v'(x)| \leq \mathfrak{C}_\zeta \frac{1}{2}|\mathcal{O}_\iota|$ as a result of Remark A.4. We will use this fact in the following discussion.

The Bernstein ellipse $\check{\Lambda}_\zeta$ has major axis $\cosh\zeta \cdot [-1,1]$ and minor axis $i\sinh\zeta \cdot [-1,1]$. As a consequence every point w in $\check{\Lambda}_\zeta$ has $\Re w \leq \cosh\zeta$ and $|\Im w| \leq \sinh\zeta$.

We have by Taylor's theorem that

$$\tau_{\iota,\sigma}\Re\left(v_\iota(x) - v_\iota(\Re x)\right) = \tau_{\iota,\sigma}\Re\left(v_\iota'(\Re x)i\Im x - \frac{v_\iota''(w)}{2}\Im x^2\right)$$

for w between x and $\Re x$ (i.e. in $\check{\Lambda}_\zeta$). Thus,

$$|\tau_{\iota,\sigma}\Re\left(v_\iota(x) - v_\iota(\Re x)\right)| \leq \frac{C_{1,\delta}\mathfrak{C}_\zeta|\mathcal{O}_\iota|}{4}\sinh^2\zeta.$$

Furthermore,

$$\tau_{\iota,\sigma}\left(v_\iota(\Re x) - v_\iota(\sigma)\right) = \sigma^{-1}\operatorname{sgn} v_\iota'(0)(\Re x - \sigma)v_\iota'(w)$$

for w between $\Re x$ and σ, i.e. in $\check{\Lambda}_\zeta \cap \mathbb{R}$. Since $v_\iota' \neq 0$ on $\check{\Lambda}_\zeta$ because of the bounded distortion condition (AD_δ), and v_ι' must be real on $\check{\Lambda}_\zeta \cap \mathbb{R}$ as it is real on $[-1,1]$ and analytic on the whole interval, we have $\operatorname{sgn} v_\iota'(w) = \operatorname{sgn} v_\iota'(0)$ and so

$$\begin{aligned}\tau_{\iota,\sigma}\left(v_\iota(\Re x) - v_\iota(\sigma)\right) &= (\Re x/\sigma - 1)|v_\iota'(w)| \\ &\leq (\cosh\zeta - 1)|v_\iota'(w)| \\ &\leq (\cosh\zeta - 1)\mathfrak{C}_\zeta \frac{|\mathcal{O}_\iota|}{2}.\end{aligned}$$

As a result we have from (A.4)

$$|T_{\iota,\sigma}(x)| \geq \Re T_{\iota,\sigma}(x) \geq \left(\boldsymbol{\xi}^{-1} - \frac{\mathfrak{C}_\zeta}{4}\left(C_{1,\delta}\sinh^2\zeta + 2\cosh\zeta - 2\right)\right)|\mathcal{O}_\iota|.$$

When ζ is small enough, the term multiplying $|\mathcal{O}_\iota|$ is positive. □

With these lemmas in hand, we can now prove Lemma 2.12.

Proof of Lemma 2.12. We begin with the first part of part (a), bounding derivatives of the ν_ι. We will do this by first proving a formula for the derivatives of ν_ι and then bounding terms in this formula to get overall bounds.

Let $\pi_n := n \mod 2$. We claim that

$$\nu_\iota^{(n+1)}(\cos^{-1}x) = \sum_{q+r+s\leq n} a_{q,r,s,n}(x) Y_{\iota,1}^{q,n}(x) Y_{\iota,-1}^{r,n}(x) v^{(s+1)}(x), \tag{A.5}$$

where $a_{q,r,s,n}$ are polynomials in x with coefficients independent of f, and

$$Y_{\iota,\sigma}^{m,n}(x) = \begin{cases} (1-x\sigma^{-1})^{\frac{\pi_n}{2}}\left(S_{\iota,\sigma}^{-1/2}\right)^{(m)}, & v_\iota(\sigma) \in \{-1,1\}, \\ (1-x\sigma^{-1})^{\frac{\pi_n}{2}}\left(T_{\iota,\sigma}^{-1/2}\right)^{(m)}, & v_\iota(\sigma) \notin \{-1,1\}. \end{cases} \tag{A.6}$$

We prove this claim by induction. Suppose without loss of generality that $\operatorname{sgn} v_\iota' = 1$.

When $n = 0$, we have that

$$\nu_\iota'(\cos^{-1}x) = \sqrt{\frac{1-x}{1-v_\iota(x)}}\sqrt{\frac{1+x}{1+v_\iota(x)}}v_\iota'(x).$$

From (A.6), we find that

$$Y_{\iota,\sigma}^{0,0}(x) = \sqrt{\frac{1-x\sigma^{-1}}{1-v_\iota(x)\sigma^{-1}}},$$

and thus (A.5) follows for $n = 0$.

Suppose, then, that (A.5) is true for some n. Then

$$\nu_\iota^{(n+2)}(\cos^{-1}(x)) = \sqrt{1-x^2}(\nu_\iota^{(n+1)} \circ \cos^{-1})'(x).$$

All we need to show is that $\sqrt{1-x\sigma^{-1}}Y_{\iota,\sigma}^{m,n}(x)$ and $\sqrt{1-x\sigma^{-1}}(Y_{\iota,\sigma}^{m,n})'(x)$ can be written as a product of $Y_{\iota,\sigma}^{m,n+1}(x)$ (and for the derivative possibly also $Y_{\iota,\sigma}^{m+1,n+1}(x)$), and polynomials in x. In the case where $v_\iota(\sigma) \in \{-1,1\}$, we have

$$\begin{aligned}\sqrt{1-x\sigma^{-1}}Y_{\iota,\sigma}^{m,n}(x) &= (1-x\sigma^{-1})^{\frac{\pi_n+1}{2}}\left(S_{\iota,\sigma}^{-1/2}\right)^{(m)} \\ &= (1-x\sigma^{-1})^{\pi_n}Y_{\iota,\sigma}^{m,n+1}(x)\end{aligned}$$

and

$$\begin{aligned}\sqrt{1-x\sigma^{-1}}(Y_{\iota,\sigma}^{m,n})'(x) &= (1-x\sigma^{-1})^{\frac{\pi_n+1}{2}}\left(S_{\iota,\sigma}^{-1/2}\right)^{(m+1)} \\ &\quad - \pi_n\sigma^{-1}(1-x\sigma^{-1})^{\frac{\pi_n-1}{2}}\left(S_{\iota,\sigma}^{-1/2}\right)^{(m+1)} \\ &= (1-x\sigma^{-1})^{\pi_n}Y_{\iota,\sigma}^{m+1,n+1}(x) \\ &\quad - \pi_n\sigma^{-1}(1-x\sigma^{-1})^{\pi_n-1}Y_{\iota,\sigma}^{m,n+1}(x) \\ &= (1-x\sigma^{-1})^{\pi_n}Y_{\iota,\sigma}^{m+1,n+1}(x) - \pi_n\sigma^{-1}Y_{\iota,\sigma}^{m,n+1}(x),\end{aligned}$$

where in the last line we removed the $(1-x\sigma^{-1})^{\pi_n-1}$ element from the last term by using that the last term is zero unless $\pi_n = 1$. The relation when $v_\iota(\sigma) \notin \{-1,1\}$ is clearly analogous, from which the claim falls.

We now attempt to bound the expression in (A.5). To bound the $Y_{\iota,\sigma}^{m,n}$, we need to bound derivatives of $S_{\iota,\sigma}^{-1/2}$ and $T_{\iota,\sigma}^{-1/2}$. One may show by induction that for $n \geq 1$ there exist multivariate polynomials q_n such that for any function U,

$$(U^{-1/2})^{(n)} = U^{-1/2}q_n\left(\frac{U'}{U}, \ldots, \frac{U^{(n)}}{U}\right). \tag{A.7}$$

By Lemma A.6, we have that when $v_\iota(\sigma) \in \{-1,1\}$

$$|S_{\iota,\sigma}^{(n)}(x)| = \frac{1}{n+1}|v^{(n+1)}(w)|$$

for some $w \in [-1,1]$. Using distortion bound (DD_n) and Lemma A.3 we can bound this again to get that

$$|S_{\iota,\sigma}^{(n)}(x)| \leq \frac{C_n e^{2C_1}}{n+1}|v'(x)|.$$

We also have that $|S_{\iota,\sigma}(x)| = |v'(w)| > e^{-2C_1}|v'(x)|$ for some $w \in [-1,1]$.

Substituting these bounds into (A.7) we find that

$$\left|\left(S_{\iota,\sigma}^{-1/2}\right)^{(n)}\right| \leq e^{C_1}|v'|^{-1/2}|q_n|\left(\frac{C_1e^{4C_1}}{2}, \ldots, \frac{C_ne^{4C_1}}{n+1}\right)$$

when $v_\iota(\sigma) \in \{-1, 1\}$.

Similarly, we have that $|T^{(n)}_{\iota,\sigma}(x)| = |v^{(n)}(x)| \leq C_n|v'(x)|$ and, by Lemma A.7, when $v_\iota(\sigma) \notin \{-1, 1\}$ that $|T_{\iota,\sigma}(x)| \geq \boldsymbol{\xi}^{-1}|\mathcal{O}_\iota| \geq 2\boldsymbol{\xi}^{-1}e^{-2C_1}|v'(x)|$. These bounds can be substituted into (A.7) similarly to give

$$\left|\left(T^{-1/2}_{\iota,\sigma}\right)^{(n)}\right| \leq \sqrt{\frac{\boldsymbol{\xi}}{2}}e^{C_1}|v'|^{-1/2}|q_n|\left(\frac{C_1\boldsymbol{\xi}e^{2C_1}}{2}, \ldots, \frac{C_n\boldsymbol{\xi}e^{2C_1}}{2}\right)$$

when $v_\iota(\sigma) \notin \{-1, 1\}$.

Thus, there exist constants $\mathfrak{k}^{m,n}$ depending on the distortion constants and partition spacing constant such that for all $\iota \in I$ and $\sigma \in \{-1, 1\}$, we have $|Y^{m,n}_{\iota,\sigma}(x)| \leq \mathfrak{k}^{m,n}|v'(x)|^{-1/2}$.

Returning to (A.5), we have that since $|v^{(s+1)}(x)| \leq C_s|v'(x)|$,

$$|\nu^{(n+1)}_\iota(\cos^{-1} x)| \leq \sum_{q+r+s\leq n} |a_{q,r,s,n}|(1)\ \mathfrak{k}^{q,n}\mathfrak{k}^{r,n}C_s,$$

for $x \in [-1, 1]$, and thus $|\nu^{(n+1)}_\iota(\theta)|$ is bounded by the same constant for $\theta \in [0, 2\pi\beta_{\iota'})$.

The proof of the first part of part (b) is essentially the same as the above with $n = 1$. The major difference is that we apply Remark A.4 and the second bound in Lemma A.7 instead of Lemma A.3 and the first bound, respectively. We also use that $\cos^{-1}\check{\Lambda}_\zeta = \Lambda^{\beta_{\iota'}}_\zeta$ so bounds on $\nu^{(n)}(\theta)$ transfer directly to bounds on $\nu^{(n)}(\cos^{-1}(x))$.

The second parts of (a) and (b) are much more straightforward. In both cases we seek to bound

$$\left|\frac{h^{(n)}_\iota}{h_\iota}\right| = \frac{\left|(v'_\iota \circ \cos)^{(n)}\right|}{|v'_\iota \circ \cos|} \tag{A.8}$$

on appropriate domains. The nth derivative of $v'_\iota \circ \cos$ can be written as a linear combination of $v^{(m+1)}_\iota \circ \cos, m \leq n$ with coefficients of trignometric polynomials. Trigonometric polynomials are bounded on $[0, 2\pi\beta_{\iota'}]$ and $\Lambda^{\beta_{\iota'}}_\zeta$; on these respective domains, the $|v^{(m+1)}_\iota \circ \cos|$ are bounded by $C_m|v'_\iota \circ \cos|$ and by $C_{1,\delta}|v'_\iota \circ \cos|$ for $m = 1$ respectively. Thus, we find that (A.8) are bounded by constants depending on C_m, $m \leq n - 1$, and in the analytic case on ζ (which parameterised $\check{\Lambda}_\zeta$) and $C_{1,\zeta}$. □

A4. Explicit bounds on the norm of the solution operator in BV

In Korepanov et al. (2016), explicit a priori bounds on decay of correlations were stated in the Lipschitz norm. Specifically, if a map on $[0, 1]$ has expansion coefficient λ and (DD_1) distortion constant C_1, then with V the space of zero-integral functions

on $[0,1]$, the following bound holds:

$$\begin{aligned} R &= \frac{2C_1}{1-\lambda^{-1}} \\ D &= 4e^R(1+R), \\ \xi &= \frac{1}{2}e^{-R}(1-\lambda^{-1}), \\ \|\mathcal{L}^n|_V\|_{\mathrm{Lip}} &\leq De^{-\xi n}. \end{aligned} \tag{A.9}$$

In this appendix we sketch how these explicit bounds work through to bound $\|\mathcal{S}\|_{BV}$.

Let $\mathrm{Lip}([0,1])$ be the space of Lipschitz functions on the interval $[0,1]$ with the usual norm.

Suppose that $\|\mathcal{L}^n|_V\|_{\mathrm{Lip}} \leq K_n < 1/2$. Suppose that $g \in BV([0,1]) \cap V$ with $\|g\|_{BV} = 1$. Let $\hat{g}_{\mathfrak{n}}$ be the piecewise linear interpolant to g at the points $0, \frac{1}{\mathfrak{n}}, \frac{2}{\mathfrak{n}}, \ldots, 1$. It can be seen that $\mathrm{Lip}\,\hat{g}_{\mathfrak{n}} \leq \mathfrak{n}$ and $\|\hat{g}_{\mathfrak{n}} - g\|_1 \leq \frac{1}{2\mathfrak{n}}$.

Consequently,

$$\begin{aligned} \|\mathcal{L}^n g\|_1 &\leq \|\mathcal{L}^n \hat{g}_{\mathfrak{n}}\|_1 + \|\mathcal{L}^n(g - \hat{g}_{\mathfrak{n}})\|_1 \\ &\leq \frac{1}{5}\|\mathcal{L}^n \hat{g}_{\mathfrak{n}}\|_{\mathrm{Lip}} + \|g - \hat{g}_{\mathfrak{n}}\|_1 \\ &\leq \frac{K_n}{5}\left(\mathrm{Lip}\,\hat{g}_{\mathfrak{n}} + \|\hat{g}_{\mathfrak{n}}\|_\infty\right) + \|g - \hat{g}_{\mathfrak{n}}\|_1 \\ &\leq \frac{K_n}{5}\left(\mathfrak{n}+1\right) + \frac{1}{2\mathfrak{n}}, \end{aligned}$$

where we used that $\|h\|_{\mathrm{Lip}} \geq 5\|h\|_1$ and $\|h\|_{BV} \geq \|h\|_\infty$ for $h \in V$.

Setting $\mathfrak{n} = \lceil K_n^{-1/2} \rceil$, we have

$$\|\mathcal{L}^n g\|_1 \leq \frac{\sqrt{K_n}(7 + 4\sqrt{K_n})}{10} \leq \sqrt{K_n}.$$

Hence, as a result of the standard BV Lasota-Yorke inequality (Galatolo and Nisoli, 2014) we find that

$$\|\mathcal{L}^{m+n} g\|_{BV} \leq \frac{5}{4}|\mathcal{L}^{m+n} g|_{BV} \leq \frac{5}{4}(\lambda^{-m} C_1 \sqrt{K_n}). \tag{A.10}$$

Using that $\|\mathcal{L}^n|_V\|_{\mathrm{Lip}} \le De^{-\xi n}$ from (A.9), and choosing

$$n = \left\lceil \frac{4 + 2\log(\max\{C_1, 1\}\sqrt{D}))}{\xi} \right\rceil$$
$$m = \left\lceil \frac{2}{\log \lambda} \right\rceil,$$

we have

$$\|\mathcal{L}^n|_V\| \le \frac{e^{-4}}{\min\{1, C_1^{-2}\}D^{-1}} =: K_n < 1/2.$$

Consequently from (A.10) we have that $\|\mathcal{L}^{m+n}\|_{BV} \le \frac{5}{2}e^{-2} \le \frac{2}{5}$.

As a result,

$$\left\|\sum_{k=0}^{\infty} \mathcal{L}^k\right\|_{BV} \le \left\|\sum_{k=0}^{\infty} \mathcal{L}^{(m+n)k}\right\|_{BV} \left\|\sum_{k=0}^{m+n-1} \mathcal{L}^k\right\|_{BV} \le \frac{5}{3}(m+n)C', \tag{A.11}$$

where $C' := 1 + \frac{1}{3}\frac{C_1}{1-\lambda^{-1}} \ge \sup_{n\in\mathbb{N}} \|\mathcal{L}^n\|_{BV} \le$ This bounding property of C' is the result of the Lasota-Yorke inequality and the fact that $\|g\|_{BV} \ge 3\|g\|_1$ for $g \in BV \cap V$.

As a result of (2.19), we finally obtain the a priori bound on the solution operator

$$\|\mathcal{S}\|_{BV} \le 1 + \frac{5}{3}(m+n)C'(3+C'). \tag{A.12}$$

Appendix B

Proof of Lemmas 3.5-3.8

Proof of Lemma 3.5. Matching power series coefficients at $z = 0$, we have that

$$a_{-1} = \hat{h}_1^{-1} \tag{B.1}$$

$$a_\ell = \hat{h}_2 \hat{h}_1^{-2} - 1 \tag{B.2}$$

$$a_n = \frac{1}{n\hat{h}_1} \frac{D_{n-1}^{(n+1)}(0)}{(n+1)!}. \tag{B.3}$$

Suppose

$$r_n = \min\{R, cn^{-1}(\hat{h}_1 + \sqrt{Gc})^{-1}\} \tag{B.4}$$

for some $c \in (0,1)$ and let

$$M_{n,r} = \sup_{|z|\le r} |D_n(z)| .$$

We have as a result of (B.3) that for any $r \le r_n$

$$|a_n| \le \frac{1}{n\hat{h}_1} r^{-n-1} M_{n-1,r}.$$

Consequently, for $n \ge 1$ and $r \le r_n$ we have that

$$\begin{aligned} M_{n,r} &\le M_{n-1,r} + |a_n| \sup_{|z|\le r} |\hat{f}(z)^n - z^n| \\ &\le M_{n-1,r} + \frac{r^{-n-1} M_{n-1,r}}{n\hat{h}_1} \sup_{|z|\le r} r^n |(z^{-1}\hat{f}(z))^n - 1| \\ &\le M_{n-1,r} \left(1 + \frac{e^{n(z^{-1}\hat{f}(z)-1)} - 1}{n\hat{h}_1 r} \right). \end{aligned}$$

Now, by our stipulation on r_n we have that

$$|\hat{h}_1 z + g(z) z^2| \le \hat{h}_1 r_n + G r_n^2 \le cn^{-1}, \tag{B.5}$$

and so

$$|z^{-1}\hat{f}(z) - 1| = \left| -\frac{\hat{h}_1 + g(z)z}{z^{-1} + \hat{h}_1 + g(z)} \right| \le \frac{cn^{-1}}{1-c}.$$

Consequently,

$$\begin{aligned} M_{n,r_n} &\leq \left(1 + \frac{e^{c/(1-c)}}{n\hat{h}_1 r_n}\right) M_{n-1,r_n} \\ &\leq \left(1 + c^{-1}e^{c/(1-c)}\left(1 + \sqrt{G\hat{h}_1^{-2}c}\right)\right) M_{n-1,r_n} \\ &=: d_2 M_{n-1,r_n} \leq d_2 M_{n-1,r_{n-1}} \end{aligned}$$

and thus

$$M_{n,r} \leq d_2^n M_{0,r_0}. \tag{B.6}$$

We now aim to bound

$$M_{0,r} = \sup_{|z|\leq r} \left| a_{-1}((\hat{f}(z))^{-1} - z^{-1}) + a_\ell \log(z^{-1}\hat{f}(z)) \right|.$$

From (B.2) it can be shown that $a_\ell = \hat{h}_1^{-2} g(0)$, and thus $|a_\ell| \leq \hat{h}_1^{-1} G$. Furthermore, (B.5) gives that

$$|z^{-1}\hat{f}(z)| \geq 1 - n^{-1}c \geq 1 - c,$$

giving that

$$\begin{aligned} M_{0,r_0} &\leq \hat{h}_1^{-1}(\hat{h}_1 + Gr_0) + \hat{h}_1^{-2} G \log((1-c)^{-1}) \\ &\leq 1 + \hat{h}_1^{-2} G(c - \log(1-c)) := d_1 d_2^2 \end{aligned}$$

where in the last line we used (B.4).

Thus, since $D_n(z) = \mathcal{O}(z^{n+2})$ as $z \to 0$, for all $|z|$ smaller than r, where r is as in (B.4),

$$|D_n(z)| \leq (|z|/r_n)^{n+2} \sup_{|w|=r_n} |D_n(w)| \leq d_1 d_2^2 \, d_2^n |z|^{n+2} r_n^{-(n+2)}.$$

Choosing $c = 0.4$ we finally obtain the required bounds. □

Proof of Lemma 3.6. We proceed by induction on (3.21) and (3.22). The base case clearly holds as $\hat{f}^0(z) = z$. Suppose that (3.21) and (3.22)hold for some $k \in \mathbb{N}$. Then $|\hat{f}^k(z)| \leq (\mathbb{R}\hat{f}^k(z)^{-1})^{-1} \leq R_1$, where $R_1 := \min\{R, \aleph G^{-1}\hat{h}_1\}$. Because $R_1 \leq R$ we can apply (3.17), giving

$$|\hat{f}^{k+1}(z)^{-1} - \hat{f}^k(z)^{-1} - \hat{h}_1| \leq G|\hat{f}^k(z)| \leq \aleph G^{-1}\hat{h}_1. \tag{B.7}$$

Since

$$\left| \Re\hat{f}^{k+1}(z)^{-1} - \Re\hat{f}^k(z)^{-1} - \hat{h}_1 \right| \leq \left| \hat{f}^{k+1}(z)^{-1} - \hat{f}^k(z)^{-1} - \hat{h}_1 \right|,$$

we obtain from (B.7) that (3.21) must also hold for $k+1$. Furthermore, since

$$|\hat{f}^{k+1}(z)^{-1} - \hat{f}^k(z)^{-1} - \hat{h}_1| = (\hat{f}^{k+1}(z)^{-1} - (k+1)\hat{h}_1) - (\hat{f}^k(z)^{-1} - k\hat{h}_1),$$

the inequality (B.7) implies (3.22) for $k+1$. □

Proof of Lemma 3.7. From Lemma 3.6 we have

$$\begin{aligned}\sum_{k=0}^{\infty} |\hat{f}^k(z)|^{\bar{\beta}} k^{\delta} &\le \sum_{k=0}^{\infty} (|z^{-1} + \hat{h}_1 k|^{-1} - \aleph \hat{h}_1 k)^{-\bar{\beta}} k^{\delta} \\ &\le \sum_{k=0}^{\infty} \left(\max\{z^{-1}, \hat{h}_1 k\} - \aleph \hat{h}_1 k \right)^{-\bar{\beta}} k^{\delta}.\end{aligned}$$

The summand is increasing for $k \le \hat{h}_1^{-1} z^{-1}$ and decreasing for larger k. Thus we can use an integral bound:

$$\begin{aligned}\sum_{k=0}^{\infty} |\hat{f}^k(z)|^{\bar{\beta}} k^{\delta} &\le \int_0^{\infty} \left(\max\{z^{-1}, \hat{h}_1 k\} - \aleph \hat{h}_1 k \right)^{-\bar{\beta}} k^{\delta} \, \mathrm{d}k + (1-\aleph)^{-\bar{\beta}} |z|^{\bar{\beta}-\delta} \hat{h}_1^{-\delta} \\ &= |z|^{-\bar{\beta}-\delta-1} \hat{h}_1^{-\delta-1} (1-\aleph)^{-\bar{\beta}} \left(\frac{{}_2F_1(\bar{\beta}, \delta+1, \delta+2, \aleph)}{(\delta-1)(1-\aleph)^{-\bar{\beta}}} \right. \\ &\quad \left. + \frac{1}{\bar{\beta}-\delta-1} + \hat{h}_1 |z| \right),\end{aligned}$$

which using that ${}_2F_1(\bar{\beta}, \delta+1, \delta+2, \aleph) \le (1-\aleph)^{-\bar{\beta}}$ and $|z| \le R_1^{-1}$ gives the desired bound. □

Proof of Lemma 3.8. We know that $g(z)$ is analytic for complex $|z| \le R$; as a result, if we define $g_1 := g'(0)$ and $g_2(z) = z^{-1}(g(z) - g_1)$ we have that g_2 is bounded for $|z| \le R$ by some constant $G_2 < \infty$. Since g maps real inputs to real inputs, we also know that g_1 is real. Combining this with (3.17), we have for $|z| \le R_1 \le R$ that

$$\hat{f}(z)^{-1} - z^{-1} = \hat{h}_1 + g_1 z + g_2(z) z^2,$$

and so taking imaginary parts,

$$\Im \hat{f}(z)^{-1} - \Im z^{-1} = g_1 \Im z + \Im(g_2(z) z^2) = -g_1 |z|^2 \Im z^{-1} + \Im(g_2(z) z^2).$$

We can then bound the growth in the imaginary part of z^{-1} under iteration by $\hat{f}$:

$$|\Im \hat{f}(z)^{-1}| \le (1 + g_1 |z|^2) |\Im z^{-1}| + G_2 |z|^2.$$

Since for $\Re z \le R_1$ we have from Lemma 3.6 that $|\hat{f}^k(z)| \le R$ for all $k \in \mathbb{N}$, we obtain the linear recurrence relation

$$|\Im \hat{f}^{k+1}(z)^{-1}| \le (1 + g_1 |z|^2) |\hat{f}^k(z)^{-1}| + G_2 |z|^2.$$

Since by Lemma 3.6, $|z|^2 = \mathcal{O}(k^{-2})$ for all $\Re z \le R_1$, iterates of this equation are bounded, as required. □

Appendix C

Statistical test for linear response given time series

A statistical test, developed by this author, to probe the linear or higher-order response of a chaotic system from time series data at various parameter values was proposed in Gottwald et al. (2016). In this appendix we summarise the principles of this test, which we use in Chapter 4.

Suppose that a system has response $\mathbb{E}^\varepsilon\Psi$ which has certain regularity properties for $\varepsilon \in [\varepsilon_1, \varepsilon_2]$ (or around some ε_0), and suppose this regularity property means that there exist functions $\varphi_i(\varepsilon), i = 1, \dots, I$ and (unknown) coefficients $\beta_i \in \mathbb{R}$ such that

$$\mathbb{E}^\varepsilon\Psi \approx \sum_{i=1}^{I} \beta_i \varphi_i. \tag{C.1}$$

is a good approximation in $L^2([\varepsilon_1, \varepsilon_2])$.

For example, if $\mathbb{E}^\varepsilon\Psi$ is C^1 and $\varepsilon_2 - \varepsilon_1$ is sufficiently small, then $\mathbb{E}^\varepsilon\Psi$ can be well-approximated with a Taylor expansion about ε_1: thus, $\varphi_0 \equiv 1$ and $\varphi_1(\varepsilon) = \varepsilon - \varepsilon_1$ form a good basis for approximation, and we would expect an L^2 error of size $o(\varepsilon_1 - \varepsilon_2)$. If instead $\mathbb{E}^\varepsilon\Psi$ is smooth (e.g. C^r) on a larger interval, then we could choose Chebyshev polynomials as a basis for approximation $\varphi_i(\varepsilon) = T_{i-1}((2\varepsilon - \varepsilon_1 - \varepsilon_2)/(\varepsilon_2 - \varepsilon_1))$ for $i = 1, \dots, I$, with an L^2 error of $O(I^r)$.

Suppose that for perturbation values ε_j, $j = 1, \dots, J$ we have time series of the observable's dynamics $(\Psi_{j,n})_{n=1,\dots,N}$ where the time series length N is sufficiently large. Supposing that (as is typical for many systems (Gottwald and Melbourne, 2013)) Ψ_n obeys a central limit theorem for each selected parameter, then for large enough N the Birkhoff averages for each ε_j have Gaussian approximations

$$\bar{\Psi}_j^N := \frac{1}{N}\sum_{j=1}^{N} \Psi_{j,n} = \mathbb{E}^{\varepsilon_j}\Psi + \sigma(\varepsilon_j)\xi_j/\sqrt{N}, \tag{C.2}$$

where ξ_j are *i.i.d.* standard normal variables and the Birkhoff variance $\sigma^2(\varepsilon_j)$ can be estimated by various means, including taking multiple time series for each ε_j, or subsampling.

If we define the vector with coefficients $y_j = \bar{\Psi}_j^N$ and the matrix with coefficients $X_{ji} = \sqrt{N}\sigma(\varepsilon_j)^{-1}\varphi_i(\varepsilon_j)$, then we can write (C.1-C.2) as the linear equation

$$y = X\beta + \xi,$$

where $\xi \sim \mathcal{N}(0, I_{J\times J})$. This is of course just a standard linear statistical model, and we can use the theory of these models (Rice, 2006) to test the null hypothesis that the approximation of the response by the ϕ_i (C.1) is an equality, i.e. that $\mathbb{E}^\varepsilon\Psi$ has linear (or smooth) response.

Defining the least-squares projection matrix

$$H = X(X^TX)^{-1}X^T$$

and the Pearson chi-square test statistic

$$\chi^2 = y^T(I - H)y,$$

we have that if the approximate equality in (C.1) is exact, then χ^2 has chi-squared distribution χ^2_{J-I} where I is the number of basis functions φ_i.

If χ^2_{obs} is the observed value of the test statistic, the p-value for the test for linear (or higher-order) response is then given by

$$p = P(\chi^2_{J-I} \geq \chi^2_{\text{obs}}),$$

provided the error associated with the non-exact nature of the approximation (C.1) is appropriately small. This error is small if

$$\mathbb{E}\chi^2 - \mathbb{E}\chi^2_{J-I} = N\left\|(I - H)(\mathbb{E}^{\varepsilon_j}\Psi/\sigma(\varepsilon_j))_{j=1,\dots,J}\right\|_{\ell^2},$$

which, supposing σ is a reasonably smooth function of ε and the ε_j are uniformly spaced, estimates the minimum possible $L^2(\sigma^2)$ error in approximations of the response of the form in (C.1), multiplied by the sample sizes N.

Appendix D

Numerical methods used in Chapter 4

D1. Model reduction for chaotic microscopic sub-systems

This appendix describes how to compute the statistics of the several stochastic limiting systems of the system comprised of a distinguished degree of freedom driven by a heat bath (4.1-4.4). We consider the limiting system for large ensemble size M (4.27) for $\gamma = \frac{1}{2}$, which we recall here as

$$Q_{n+1} = (A_0 + A_1\zeta_n)\, Q_n(1 - Q_n). \tag{D.1}$$

We also consider the (deterministic) thermodynamic limit system (4.29) for $\gamma = 1$, which we also recall here

$$Q_{n+1} = (A_0 + \langle \mathbb{E}[\phi] \rangle A_1)\, Q_n(1 - Q_n). \tag{D.2}$$

We finally also consider the (stochastic) finite-size system for $\gamma = 1$

$$Q_{n+1} = (A = A_0 + \Phi_n A_1)\, Q_n(1 - Q_n), \tag{D.3}$$

where Φ_n is given by (4.31), which is recalled here as

$$\Phi_n = \langle \mathbb{E}[\phi] \rangle + \frac{1}{\sqrt{M}}\eta + \frac{1}{\sqrt{M}}\zeta_n. \tag{D.4}$$

The random variable η accounts for the random variation in the selection of the parameters $a^{(j)}$ and the random process ζ_n accounts for the dynamics of the microscopic variables. (However, as can be seen from Figure 4.12, setting $\eta \equiv 0$, i.e. replacing it with its expectation, gives a remarkably good approximation of the invariant measure, at least in the system we consider.)

In order to simulate these systems we need to estimate $\langle \mathbb{E}[\phi] \rangle$ and, for the stochastic systems, also $R(m) = \langle \mathbb{E}[\phi_0 \phi_m] \rangle - (\mathbb{E}[\phi])^2, m \in \mathbb{N}$. We describe first how we estimate these parameters from Monte Carlo simulations of the logistic map, and then describe how we sample the stochastic process ζ_n with the covariance parameters given by $R(m)$.

D1.1. Estimating parameters. We need to estimate the expectation values for K perturbation sizes ε_i with $i = 1, \cdots, K$. Since we set here $a^{(j)} = 1$ for all

microscopic variables, at each ε_i we write the averages over the microscopic dynamics as

$$\langle \mathbb{E}^{\varepsilon_i}[\phi] \rangle = \int_{\mathbb{R}} \int \phi(x,a)\, \mathrm{d}\mu^{\alpha}(x)\, \nu(\alpha - \varepsilon_i)\, \mathrm{d}\alpha, \tag{D.5}$$

and

$$\langle \mathbb{E}[\phi_0 \phi_m] \rangle_{\varepsilon_i} = \int_{\mathbb{R}} \int \phi(x_0,\alpha)\phi(f^m(x;\alpha)\, \mathrm{d}\mu^{\alpha}(x)\, \nu(\alpha - \varepsilon_i)\, \mathrm{d}\alpha \tag{D.6}$$

for $i = 1, \ldots, K$ and $m = 1, \ldots, \infty$. Here, ν is the density function of the logistic map parameters and is chosen here as the raised cosine distribution on $[3.8, 3.9]$ (c.f. (4.8)), and we used that in Section 4.6 we can write $\phi(q; a, \varepsilon) = \phi(q; a + \varepsilon)$ and $f(q; a, \varepsilon) = f(q; a + \varepsilon)$.

From now on it is understood that all observables, expectations with respect to measures and so on are for a fixed parameter α: we therefore drop the α and (j) superscripts for ease of exposition, and write for expectations over the microscopic invariant measures $\mathbb{E}[\chi] = \int \chi(x)\, \mathrm{d}\mu(x)$.

We use a trapezoidal rule to estimate the integrals in (D.5) and (D.6), using a grid of $30,001$ values of the logistic map parameters α evenly spaced on $[3.7, 4.0]$ (to allow for the support of ν as well as the range of the perturbation). This is used for each ε_i.

The expectations (D.5) and (D.6) can be entirely determined by simulations of a standard logistic map without coupling to the expanding r-dynamics (4.7). We recall that the marginals in q of the physical measures of (4.7) and of a logistic map with the same parameters are equal. Thus, if we denote by $\varphi_n = \phi(x_n, a + \varepsilon_i)$ such that $x_{n+1} = (a + \varepsilon_i) x_n (1 - x_n)$ with $x_0 = q_0$, we have that $\mathbb{E}[\phi] = \mathbb{E}[\varphi]$.

To estimate the averages of the auto-correlations (D.6) we define $N(m)$ as the number of evolution steps of the q-dynamics up to physical time m which were done according to the logistic map (i.e. discarding all those instances when the r-dynamics forces q not to vary). Note that $N(m)$ has a binomial distribution $N(m) \sim \mathrm{B}(m, \frac{1}{2})$. Hence by definition we have

$$\phi(q_m) = \varphi_{N(m)},$$

and we can write

$$\begin{aligned} \mathbb{E}[\phi_0 \phi_m] &= \mathbb{E}[\varphi_0 \varphi_{N(m)}] \\ &= \sum_{i=0}^{m} 2^{-i} \binom{m}{i} \mathbb{E}[\varphi_0 \varphi_i]. \end{aligned}$$

To estimate the autocorrelations of the logistic variables, we will distinguish between regular and chaotic parameter values. In almost all cases there is no way to

determine that a given parameter value is chaotic: we therefore use the list of regular parameter values given in Galias (2017) and make the heuristically validated assumption that the vast majority of the remaining parameter values are chaotic.

For regular values of α, when the logistic map x_n with parameter α has a stable periodic orbit, calculating the stable periodic orbit allows for an accurate evaluation of the expectation. We use the database of periodic windows given in Galias (2017) to identify both the values of α and the stable periodic orbits.

For chaotic values of α we estimate expectations and lag-correlations of the logistic map with parameter α via Monte-Carlo simulation of the logistic map x_n, using 10 separate initialisations with 399,168 time steps each. This number of time steps was chosen as it has a large number of prime factors, and therefore will give more accurate estimates for short periodic windows outside the database, or for chaotic values where the acim has multiple connected components (i.e., f is not mixing but f^p is for some $p > 1$).

D1.2. Sampling the stochastic process $\boldsymbol{\zeta_n}$. The limiting process ζ_n is a stationary mean-zero Gaussian process given by lag-covariance function $R(m)$. Assuming sufficiently fast decay of the lag-covariance function, we can write this process as a moving-average process of infinite order

$$\zeta_n = \sum_{m=0}^{\infty} \beta_m X_{n-m}$$

with a deterministic sequence $(\beta_m)_{m\in\mathbb{N}} \in \ell_2$ and *i.i.d.* standard normal random variables X_n.

The moving average coefficients β_m and the covariance function R_m are related by

$$R(m) = \sum_{k=0}^{\infty} \beta_k \beta_{m+k}.$$

The coefficients can be extracted from the covariance function via the generating functions

$$\mathcal{B}(z) := \sum_{m=0}^{\infty} \beta_m z^m$$

and

$$\mathcal{R}(z) := \sum_{m=-\infty}^{\infty} R(|m|) z^m,$$

for which the relation $\mathcal{R}(z) = \mathcal{B}(z)\mathcal{B}(z^{-1})$ holds. If we restrict to the complex unit circle, setting $z = e^{i\theta}$, we find that $\mathcal{R}(e^{i\theta}) = \mathcal{B}(e^{i\theta})\mathcal{B}(e^{-i\theta}) = |\mathcal{B}(e^{i\theta})|^2$ since $\beta_m \in \mathbb{R}$.

Assuming that $\mathcal{R}(e^{i\theta}) \neq 0$, we have that

$$\frac{1}{2}\log \mathcal{R}(e^{i\theta}) = \Re \log \mathcal{B}(e^{i\theta}),$$

and hence, we can write, using that the β_m are real,

$$\log \mathcal{B}(z) = \sum_{m=0}^{\infty} b_m z^m$$

with $b_m \in \mathbb{R}$. The b_m may be calculated via Fourier cosine transform using that

$$\frac{1}{2}\log \mathcal{R}(e^{i\theta}) = \sum_{m=0}^{\infty} b_m \cos m\theta.$$

The b_m coefficients allow one to evaluate $\mathcal{B}(e^{i\theta})$, from which the moving average coefficients β_m are obtained via an additional Fourier transform.

D2. Numerical method to compute the thermodynamic limit $M \to \infty$ for uniformly expanding maps

In the thermodynamic limit of infinite M the strong law of large numbers holds and

$$\Phi_n = \langle \mathbb{E}\Phi_n \rangle = \int \phi(q)\, \mathrm{d}\mu_n(q),$$

where μ_n is the (time-varying) physical measure of the system, which evolves as

$$\mu_{n+1} = \mathcal{L}_{K_n}\mu_n,$$

where $K_n = \tanh(\varepsilon\Phi_n - 2)$ and $\mathcal{L}_{K_n}$ is the transfer operator of the system (4.43): we recall from Chapter 2 that the action of the transfer operator is given explicitly by 2.1.

Because for all fixed $K_n = K$ the map (4.43) is uniformly expanding, the physical measures μ_n are absolutely continuous with respect to Lebesgue, and we can write them as $\mu_n(q)dq$. Furthermore, because the map (4.43) is analytic and hence infinitely many times differentiable, it is possible to approximate the measure density and transfer operator dynamics very accurately using the Chebyshev spectral Galerkin methods proposed in Chapter 2. We have implemented an adaptive-order spectral approximation of the measure density in the Julia package Poltergeist.jl (Wormell, 2017), which allows us to simulate the dynamics of μ_n. The core routine, which outputs μ_{n+1} and Φ_{n+1} given inputs μ_n, ε and the driving d_n (by default Φ_n), is defined as follows (note that Julia recognises Unicode characters):

```
function F($\mu$_n, $\eps$, d_n = sum($\phi$*$\mu$_n))
    K_n = tanh(d_n*$\eps$-2)
    f_n = f_map(K_n)
        # create MarkovMap object
    $\mu$_n1 = transfer(f_n,$\mu$_n)
        # compute $\mu_{n+1}$
    return $\mu$_n1, sum($\phi$*$\mu$_n1)
end
```

More details of the algorithm and some examples of its use may be found at `https://github.com/wormell/PoltergeistExamples/blob/master/WeakSelfCoupling-LimitingSystem.ipynb`.

After the first time the routine is called (during which Julia compiles the code), the algorithm takes around 8×10^{-4} seconds on a standard laptop to compute each μ_{n+1} from μ_n, and has an approximation error of only around 10^{-13}: by comparison, if one aims to estimate μ_n as a Monte-Carlo approximation with a large ensemble of $M = \mathcal{O}(10^8)$, a relatively large approximation error of 10^{-3} is incurred.

Bibliography

Aaronson, J., 1997. An introduction to infinite ergodic theory. American Mathematical Soc.

Abel, N., 1826. Untersuchung der functionen zweier unabhängig veränderlichen größen x und y, wie $f(x, y)$, welche die eigenschaft haben, daß $f(z, f(x, y))$ eine symmetrische function von z, x und y ist. Journal für die reine und angewandte Mathematik 1826 (1), 11–15.

Abramov, R. V., 2010. Approximate linear response for slow variables of dynamics with explicit time scale separation. Journal of Computational Physics 229 (20), 7739–7746.

Abramov, R. V., Majda, A. J., 2007. Blended response algorithms for linear fluctuation-dissipation for complex nonlinear dynamical systems. Nonlinearity 20 (12), 2793.

Abramov, R. V., Majda, A. J., 2008. New approximations and tests of linear fluctuation-response for chaotic nonlinear forced-dissipative dynamical systems. J. Nonlinear Sci. 18 (3), 303–341.

Abramov, R. V., Majda, A. J., 2009. A new algorithm for low-frequency climate response. Journal of the Atmospheric Sciences 66 (2), 286–309.

Abramowitz, M., Stegun, I. A., 1973. Handbook of mathematical functions with formulas, graphs, and mathematical tables. Dover Publications, New York.

Alves, J. F., Luzzatto, S., Pinheiro, V., 2004. Lyapunov exponents and rates of mixing for one-dimensional maps. Ergodic Theory Dynam. Systems 24, 637–657.

Aurentz, J. L., Trefethen, L. N., 2017. Chopping a Chebyshev series. ACM Transactions on Mathematical Software 43 (4), 33.

Avila, A., Lyubich, M., de Melo, W., 2003. Regular or stochastic dynamics in real analytic families of unimodal maps. Inventiones mathematicae 154 (3), 451–550.

Bahsoun, W., Bose, C., Duan, Y., 2015. Rigorous pointwise approximations for invariant densities of non-uniformly expanding maps. Ergodic Theory and Dynamical Systems 35 (4), 1028–1044.

Bahsoun, W., Galatolo, S., Nisoli, I., Niu, X., 2016. Rigorous approximation of diffusion coefficients for expanding maps. Journal of Statistical Physics 163 (6), 1486–1503.

Bahsoun, W., Galatolo, S., Nisoli, I., Niu, X., 2018. A rigorous computational approach to linear response. Nonlinearity 31 (3), 1073.

Baladi, V., 2000. Positive transfer operators and decay of correlations. Vol. 16. World Scientific, Shanghai.

Baladi, V., Aug 2014. Linear response, or else. In: ICM Seoul 2014, Proceedings, Volume III. pp. 525–545.

Baladi, V., 2017. The quest for the ultimate anisotropic Banach space. Journal of Statistical Physics 166 (3-4), 525–557.

Baladi, V., Benedicks, M., Schnellmann, D., 2015. Whitney-Hölder continuity of the SRB measure for transversal families of smooth unimodal maps. Invent. Math. 201 (3), 773–844.

Baladi, V., Holschneider, M., 1999. Approximation of nonessential spectrum of transfer operators. Nonlinearity 12 (3), 525.

Baladi, V., Kondah, A., Schmitt, B., 1996. Random correlations for small perturbations of expanding maps. Random and Computational Dynamics 4 (2/3), 179–204.

Baladi, V., Liverani, C., 2012. Exponential decay of correlations for piecewise cone hyperbolic contact flows. Communications in Mathematical Physics 314 (3), 689–773.

Baladi, V., Smania, D., 2008. Linear response formula for piecewise expanding unimodal maps. Nonlinearity 21 (4), 677–711.

Baladi, V., Smania, D., 2010. Alternative proofs of linear response for piecewise expanding unimodal maps. Ergodic Theory and Dynamical Systems 30 (01), 1–20.

Baladi, V., Smania, D., 2012. Linear response for smooth deformations of generic nonuniformly hyperbolic unimodal maps. In: Annales scientifiques de l'École Normale Supérieure. Vol. 45. pp. 861–926.

Balescu, R., 1975. Equilibrium and Non-equilibrium Statistical Mechanics. John Wiley & Sons, New York.

Bandtlow, O. F., Jenkinson, O., 2008. Explicit eigenvalue estimates for transfer operators acting on spaces of holomorphic functions. Advances in Mathematics 218 (3), 902–925.

Bandtlow, O. F., Slipantschuk, J., 2020. Lagrange approximation of transfer operators associated with holomorphic data. arXiv preprint arXiv:2004.03534.

Bell, T. L., 1980. Climate sensitivity from fluctuation dissipation: Some simple model tests. Journal of the Atmospheric Sciences 37 (8), 1700–1707.

Benedicks, M., Carleson, L., 1991. The dynamics of the hénon map. Annals of Mathematics 133 (1), 73–169.

Benet, L., Sanders, D. P., 2019. ValidatedNumerics. Available at `https://github.com/JuliaIntervals/ValidatedNumerics.jl` and in the Julia package repository.

Blumenthal, A., Xue, J., Young, L.-S., 2017. Lyapunov exponents for random perturbations of some area-preserving maps including the standard map. Annals of Mathematics, 285–310.

Bose, C., Murray, R., 2001. The exact rate of approximation in Ulam's method. Discrete and continuous dynamical systems 7 (1), 219–235.

Bowen, R., 2008. Equilibrium states and the ergodic theory of Anosov diffeomorphisms, 2nd Edition. Lecture notes in mathematics, 470. Springer, Berlin.

Boyd, J. P., 2001. Chebyshev and Fourier spectral methods. Courier Corporation, Mineola, NY.

Chekroun, M. D., Neelin, J. D., Kondrashov, D., McWilliams, J. C., Ghil, M., 2014. Rough parameter dependence in climate models and the role of Ruelle-Pollicott resonances. Proceedings of the National Academy of Sciences 111 (5), 1684–90.

Cionni, I., Visconti, G., Sassi, F., May 2004. Fluctuation dissipation theorem in a general circulation model. Geophysical Research Letters 31 (9), L09206.

Collet, P., Eckmann, J.-P., 1983. Positive Liapunov exponents and absolute continuity for maps of the interval. Ergodic Theory Dynam. Systems 3 (1), 13–46.

Collet, P., Eckmann, J.-P., 2007. Concepts and results in chaotic dynamics: a short course. Springer Science & Business Media, Berlin.

Cooper, F., Haynes, P., 2013. Assessment of the fluctuation-dissipation theorem as an estimator of the tropospheric response to forcing. Quart. J. Roy. Met. Soc.Submitted.

Cooper, F. C., Esler, J. G., Haynes, P. H., 2013. Estimation of the local response to a forcing in a high dimensional system using the fluctuation-dissipation theorem. Nonlin. Processes Geophys. 20 (2), 239–248.

Cooper, F. C., Haynes, P. H., 05 2011. Climate Sensitivity via a Nonparametric Fluctuation-Dissipation Theorem. Journal of the Atmospheric Sciences 68 (5), 937–953.

Crimmins, H., Froyland, G., 2019a. Fourier approximation of the statistical properties of Anosov maps on tori. arXiv preprint arXiv:1906.04905.

Crimmins, H., Froyland, G., 2019b. Stability and approximation of statistical limit laws for multidimensional piecewise expanding maps. Annales Henri Poincaré 20 (9), 3113–3161.

Cvitanovic, P., Eckhardt, B., 1991. Periodic orbit expansions for classical smooth flows. Journal of Physics A: Mathematical and General 24 (5), L237–L241.

De Lima, A., Smania, D., 2018. Central limit theorem for the modulus of continuity of averages of observables on transversal families of piecewise expanding unimodal maps. Journal of the Institute of Mathematics of Jussieu 17 (3), 673–733.

De Simoi, J., Liverani, C., 2015. The Martingale Approach after Varadhan and Dolgopyat. In: Hyperbolic dynamics, fluctuations and large deviations. Vol. 89 of Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, RI, pp. 311–339.

De Simoi, J., Liverani, C., 2016. Statistical properties of mostly contracting fast-slow partially hyperbolic systems. Invent. Math. 206 (1), 147–227.

Dellnitz, M., Froyland, G., Junge, O., 2001. The algorithms behind GAIO — set oriented numerical methods for dynamical systems. In: Ergodic theory, analysis, and efficient simulation of dynamical systems. Springer, pp. 145–174.

Dellnitz, M., Junge, O., 1999. On the approximation of complicated dynamical behaviour. SIAM J. Numer. Anal. 36, 491–515.

Denner, A., Junge, O., Matthes, D., 2015. Computing coherent sets using the Fokker-Planck equation. Journal of Computational Dynamics 3 (2), 163.

Dijkstra, H., 2005. Nonlinear Physical Oceanography: A Dynamical Systems Approach to the Large Scale Ocean Circulation and El Niño, 2nd Edition. Springer.

Dolgopyat, D., 2004. On differentiability of SRB states for partially hyperbolic systems. Invent. Math. 155 (2), 389–449.

Donges, J. F., Zou, Y., Marwan, N., Kurths, J., 2009. Complex networks in climate dynamics. The European Physical Journal Special Topics 174 (1), 157–179.

Driscoll, T. A., Hale, N., Trefethen, L. N., 2014. Chebfun guide.

Dymnikov, V. P., Gritsoun, A. S., 2001. Climate model attractors: chaos, quasi-regularity and sensitivity to small perturbations of external forcing. Nonlinear Processes in Geophysics 8 (4/5), 201–209.

Eckhardt, B., Grossmann, S., Dec 1994. Correlation functions in chaotic systems from periodic orbits. Phys. Rev. E 50, 4571–4576.

Ershov, S. V., Potapov, A. B., 1995. On mean field fluctuations in globally coupled maps. Physica D: Nonlinear Phenomena 86 (4), 523–558.

Ershov, S. V., Potapov, A. B., 1997. On mean field fluctuations in globally coupled logistic-type maps. Physica D: Nonlinear Phenomena 106 (1-2), 9–38.

Field, M., Melbourne, I., Török, A., 2007. Stability of mixing and rapid mixing for hyperbolic flows. Annals of mathematics, 269–291.

Ford, G. W., Kac, M., 1987. On the quantum Langevin equation. J. Statist. Phys. 46 (5-6), 803–810.

Ford, G. W., Kac, M., Mazur, P., 1965. Statistical mechanics of assemblies of coupled oscillators. J. Mathematical Phys. 6, 504–515.

Froyland, G., 1999. Using Ulam's method to calculate entropy and other dynamical invariants. Nonlinearity 12 (1), 79.

Froyland, G., 2007. On Ulam approximation of the isolated spectrum and eigenfunctions of hyperbolic maps. Dynamical Systems 17 (3), 671–689.

Froyland, G., Junge, O., Koltai, P., 2013. Estimating long-term behavior of flows without trajectory integration: The infinitesimal generator approach. SIAM Journal on Numerical Analysis 51 (1), 223–247.

Froyland, G., Murray, R., Stancevic, O., 2011. Spectral degeneracy and escape dynamics for intermittent maps with a hole. Nonlinearity 24 (9), 2435.

Fuchs, D., Sherwood, S., Hernandez, D., 2014. An exploration of multivariate fluctuation dissipation operators and their response to sea surface temperature perturbations. Journal of the Atmospheric Sciences 72, 472–486.

Galatolo, S., Nisoli, I., 2014. An elementary approach to rigorous approximation of invariant measures. SIAM Journal on Applied Dynamical Systems 13 (2), 958–985.

Galias, Z., 2017. Systematic search for wide periodic windows and bounds for the set of regular parameters for the quadratic map. Chaos: An Interdisciplinary Journal of Nonlinear Science 27 (5), 053106.

Gallavotti, G., 2008. Chaotic hypothesis. Scholarpedia 3 (1), 5906, revision #126984.

Gallavotti, G., 2019. Nonequilibrium and fluctuation relation. arXiv preprint arXiv:1906.10069.

Gallavotti, G., Cohen, E. G. D., Apr 1995a. Dynamical ensembles in Nonequilibrium Statistical Mechanics. Phys. Rev. Lett. 74, 2694–2697.

Gallavotti, G., Cohen, E. G. D., Sep 1995b. Dynamical ensembles in stationary states. Journal of Statistical Physics 80 (5-6), 931–970.

Givon, D., Kupferman, R., Stuart, A., 2004. Extracting macroscopic dynamics: Model problems and algorithms. Nonlinearity 17 (6), R55–127.

Gora, P., Boyarsky, A., 1989. Absolutely continuous invariant measures for piecewise expanding C^2 transformations in $\mathbb{R}^n$. Israel Journal of Mathematics 67 (3), 272–286.

Gottwald, G. A., Melbourne, I., 2013. Homogenization for deterministic maps and multiplicative noise. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science 469 (2156).

Gottwald, G. A., Melbourne, I., 2014. A test for a conjecture on the nature of attractors for smooth dynamical systems. Chaos: An Interdisciplinary Journal of Nonlinear Science 24 (2), 024403.

Gottwald, G. A., Wormell, J. P., Wouters, J., 2016. On spurious detection of linear response and misuse of the fluctuation–dissipation theorem in finite time series. Physica D: Nonlinear Phenomena 331, 89–101.

Gouëzel, S., 2004a. Central limit theorem and stable laws for intermittent maps. Probability Theory and Related Fields 128 (1), 82–122.

Gouëzel, S., 2004b. Sharp polynomial estimates for the decay of correlations. Israel Journal of Mathematics 139 (1), 29–65.

Gouëzel, S., 2010. Almost sure invariance principle for dynamical systems by spectral methods. Ann. Probability 38, 1639–1671.

Gritsun, A., 2010. Unstable periodic orbits and sensitivity of the barotropic model of the atmosphere. Russian Journal of Numerical Analysis and Mathematical Modelling 25 (4), 303–321.

Gritsun, A., Branstator, G., 07 2007. Climate response using a three-dimensional operator based on the fluctuation-dissipation theorem. Journal of the Atmospheric

Sciences 64 (7), 2558–2575.

Gritsun, A., Branstator, G., Dymnikov, V., 2002. Construction of the linear response operator of an atmospheric general circulation model to small external forcing. Russ. J. Numer. Anal. Math. Modelling 17, 399–416.

Gritsun, A., Branstator, G., Majda, A., 09 2008. Climate response of linear and quadratic functionals using the fluctuation-dissipation theorem. Journal of the Atmospheric Sciences 65 (9), 2824–2829,2831–2841.

Gritsun, A., Dymnikov, V., 1999. Barotropic atmosphere response to small external actions: Theory and numerical experiments. Izv. Akad. Nauk. Fiz. Atmos. Okeana. Biol. 35, 565–581.

Gritsun, A. S., 2010. Construction of response operators to small external forcings for atmospheric general circulation models with time periodic right-hand sides. Izvestiya, Atmospheric and Oceanic Physics 46 (6), 748–756.

Hairer, M., Majda, A. J., 2010. A simple framework to justify linear response theory. Nonlinearity 23 (4), 909.

Hänggi, P., 1978. Stochastic processes 2: response theory and fluctuation theorems. Helvetica Physica Acta 51 (2), 202–219.

Hansen, A. C., 2010. Infinite-dimensional numerical linear algebra: theory and applications. In: Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences. The Royal Society, p. rspa20090617.

Hasselmann, K., 1976. Stochastic climate models. Part 1: Theory. Tellus 28 (6), 473–485.

Holschneider, M., 1996. Wavelet analysis of transfer operators acting on n-dimensional Hölder, Zygmund Triebel spaces. Mathematical Physics Preprint Archive MP-ARC-96-169.

Hunt, T., MacKay, R., 2003. Anosov parameter values for the triple linkage and a physical system with a uniformly chaotic attractor. Nonlinearity 16 (4), 1499.

Jenkinson, O., Pollicott, M., 2005. Orthonormal expansions of invariant densities for expanding maps. Advances in Mathematics 192 (1), 1–34.

Jenkinson, O., Pollicott, M., Vytnova, P., 2018. Rigorous computation of diffusion coefficients for expanding maps. Journal of Statistical Physics 170 (2), 221–253.

Kaneko, K., 1990. Globally coupled chaos violates the law of large numbers but not the central-limit theorem. Physical review letters 65 (12), 1391.

Kelly, D., Melbourne, I., 2014. Deterministic homogenization for fast-slow systems with chaotic noises. arXiv:1409.5748 [math.PR].

Kirk-Davidoff, D. B., Feb. 2009. On the diagnosis of climate sensitivity using observations of fluctuations. Atmos. Chem. Phys. 9 (3), 813–822.

Korepanov, A., 2015. Linear response for intermittent maps with summable and nonsummable decay of correlations. arXiv:1508.06571 [math.DS].

Korepanov, A., Kosloff, Z., Melbourne, I., 2016. Explicit coupling argument for uniformly expanding maps. arXiv preprint arXiv:1602.03795.

Krauskopf, B., Osinga, H. M., Galán-Vioque, J., 2007. Numerical continuation methods for dynamical systems. Springer.

Kubo, R., 1966. The fluctuation-dissipation theorem. Reports on Progress in Physics 29 (1), 255.

Kuczma, M., Choczewski, B., Ger, R., 1990. Iterative functional equations. No. 32 in Encyclopedia of Mathematics and its Applications. Cambridge University Press.

Kupferman, R., Stuart, A. M., Terry, J. R., Tupper, P. F., 2002. Long-term behaviour of large mechanical systems with random initial data. Stoch. Dyn. 2 (4), 533–562.

Lanford, I. O. E., 1998. Informal remarks on the orbit structure of discrete approximations to chaotic maps. Experiment. Math. 7 (4), 317–324.

Langen, P. L., Alexeev, V. A., 2005. Estimating $2 \times CO_2$ warming in an aquaplanet GCM using the fluctuation-dissipation theorem. Geophysical Research Letters 32 (23), l23708.

Lasota, A., Mackey, M. C., 1994. Chaos, fractals, and noise: stochastic aspects of dynamics. Vol. 97 of Berlin. Springer Science & Business Media.

Lasota, A., Yorke, J. A., 1973. On the existence of invariant measures for piecewise monotonic transformations. Transactions of the American Mathematical Society, 481–488.

Lehmer, D. H., 1940. On the maxima and minima of Bernoulli polynomials. The American Mathematical Monthly 47 (8), 533–538.

Lembo, V., Lucarini, V., Ragone, F., 2019. Beyond Forcing Scenarios: Predicting Climate Change through Response Operators in a Coupled General Circulation Model. arXiv preprint arXiv:1912.03996.

Liverani, C., Saussol, B., Vaienti, S., 1999. A probabilistic approach to intermittency. Ergodic Theory Dynam. Systems 19, 671–685.

Lorenz, E. N., 1963. Deterministic nonperiodic flow. Journal of the Atmospheric Sciences 20 (2), 130–141.

Lucarini, V., Sarno, S., 2011. A statistical mechanical approach for the computation of the climatic response to general forcings. Nonlinear Processes in Geophysics 18 (1), 7–28.

Lyubich, M., 2002. Almost every real quadratic map is either regular or stochastic. Ann. of Math. (2) 156 (1), 1–78.

Magalhães, M., Pollicott, M., 2013. Geometry and dynamics of planar linkages. Communications in Mathematical Physics 317 (3), 615–634.

Majda, A. J., Abramov, R., Gershgorin, B., 2010. High skill in low-frequency climate response through fluctuation dissipation theorems despite structural instability. Proceedings of the National Academy of Sciences 107 (2), 581–586.

Marconi, U. M. B., Puglisi, A., Rondoni, L., Vulpiani, A., 2008. Fluctuation–dissipation: Response theory in statistical physics. Physics Reports 461, 111–195.

Melbourne, I., Nicol, M., 2005. Almost sure invariance principle for nonuniformly hyperbolic systems. Commun. Math. Phys. 260, 131–146.

Melbourne, I., Nicol, M., 2008. Large deviations for nonuniformly hyperbolic systems. Trans. Amer. Math. Soc. 360 (12), 6661–6676.

Melbourne, I., Stuart, A., 2011. A note on diffusion limits of chaotic skew-product flows. Nonlinearity 24, 1361–1367.

Melbourne, I., Török, A., 2004. Statistical limit theorems for suspension flows. Israel Journal of Math. 144, 191–209.

Murray, R., 2010. Ulam's method for some non-uniformly expanding maps. Discrete. Contin. Dyn. Syst 26 (3), 1007–1018.

Nijsse, F. J., Cox, P. M., Huntingford, C., Williamson, M. S., 2019. Decadal global temperature variability increases strongly with climate sensitivity. Nature Climate Change 9 (8), 598–601.

North, G. R., Bell, R. E., Hardin, J. W., 1993. Fluctuation dissipation in a general circulation model. Climate Dynamics 8 (6), 259–264.

Olver, F., 1997. Asymptotics and special functions. AK Peters/CRC Press.

Olver, S., 2019. ApproxFun. Available at `https://github.com/JuliaApproximation/ApproxFun.jl` and in the Julia package repository.

Olver, S., Townsend, A., 2013. A fast and well-conditioned spectral method. SIAM Review 55 (3), 462–489.

Palis, J., Palis, Júnior, J., Takens, F., 1995. Hyperbolicity and sensitive chaotic dynamics at homoclinic bifurcations: Fractal dimensions and infinitely many attractors in dynamics. Cambridge University Press.

Palmer, T., Williams, P. (Eds.), 2010. Stochastic Physics and Climate Modelling. Cambridge University Press, Cambridge.

Pikovsky, A. S., Kurths, J., Mar 1994. Do globally coupled maps really violate the law of large numbers? Phys. Rev. Lett. 72, 1644–1646.

Pollicott, M., 1985. On the rate of mixing of Axiom A flows. Invent. Math. 81, 413–426.

Pollicott, M., 1986. Meromorphic extensions of generalised zeta functions. Invent. Math. 85 (1), 147–164.

Pollicott, M., Jenkinson, O., 2000. Computing invariant densities and metric entropy. Communications in Mathematical Physics 211 (3), 687–703.

Pollicott, M., Vytnova, P., 2016. Linear response and periodic points. Nonlinearity 29 (10), 3047.

Pomeau, Y., Manneville, P., 1980. Intermittent transition to turbulence in dissipative dynamical systems. Communications in Mathematical Physics 74 (2), 189–197.

Ragone, F., Lucarini, V., Lunkeit, F., 2016. A new framework for climate sensitivity and prediction: a modelling perspective. Climate Dynamics 46, 1459–1471.

Revels, J., Lubin, M., Papamarkou, T., 2016. Forward-mode automatic differentiation in Julia. arXiv:1607.07892 [cs.MS].

Rice, J., 2006. Mathematical statistics and data analysis. Thomson Learning, Belmont, CA.

Ring, M. J., Plumb, R. A., 2008. The response of a simplified GCM to axisymmetric forcings: Applicability of the fluctuation–dissipation theorem. Journal of the Atmospheric Sciences 65 (12), 3880–3898.

Ruelle, D., 1997. Differentiation of SRB states. Communications in Mathematical Physics 187 (1), 227–241.

Ruelle, D., 1998. General linear response formula in statistical mechanics, and the fluctuation-dissipation theorem far from equilibrium. Phys. Lett. A 245 (3-4), 220–224.

Ruelle, D., 2004. Thermodynamic Formalism: The Mathematical Structure of Equilibrium Statistical Mechanics, 2nd Edition. Cambridge Mathematical Library. Cambridge University Press.

Ruelle, D., 2009a. A review of linear response theory for general differentiable dynamical systems. Nonlinearity 22 (4), 855–870.

Ruelle, D., 2009b. Structure and f-dependence of the a.c.i.m. for a unimodal map f of Misiurewicz type. Communications in Mathematical Physics 287 (3), 1039–1070.

Ruelle, D., 2018. Linear response theory for diffeomorphisms with tangencies of stable and unstable manifolds—a contribution to the Gallavotti-Cohen chaotic hypothesis. Nonlinearity 31 (12), 5683.

Rugh, H. H., 1999. Intermittency and regularized Fredholm determinants. Inventiones mathematicae 135 (1), 1–24.

Sélley, F., Bálint, P., Aug 2016. Mean-field coupling of identical expanding circle maps. Journal of Statistical Physics 164 (4), 858–889.

Shibata, T., Chawanya, T., Kaneko, K., 1999. Noiseless collective motion out of noisy chaos. Physical review letters 82 (22), 4424.

Slipantschuk, J., Richter, M., Chappell, D., Tanner, G., Just, W., Bandtlow, O., 2019. Transfer operator approach to ray-tracing in circular domains. arXiv preprint arXiv:1910.02938.

Stuart, A. M., Warren, J. O., Nov 1999. Analysis and Experiments for a Computational Model of a Heat Bath. Journal of Statistical Physics 97 (3), 687–723.

Szekeres, G., 1958. Regular iteration of real and complex functions. Acta Mathematica 100 (3-4), 203–258.

Tian, H., Ding, J., Rhee, N. H., 2016. Approximations of Frobenius-Perron operators via piecewise quadratic functions. Dynamic Systems and Applications 25 (3), 557–575.

Trefethen, L. N., 2013. Approximation theory and approximation practice. Siam, Philadelphia, PA.

Tsonis, A. A., Swanson, K. L., Roebber, P. J., 2006. What do networks have to do with climate? Bulletin of the American Meteorological Society 87 (5), 585–596.

Tucker, W., 2011. Validated numerics: a short introduction to rigorous computations. Princeton University Press.

Wolfram Research, Inc., 2013. Mathematica 10.0.

Wormell, C. L., 2017. Poltergeist. Available at `https://github.com/wormell/Poltergeist.jl` and in the Julia package repository.

Wormell, C. L., 2019. Spectral Galerkin methods for transfer operators in uniformly expanding dynamics. Numerische Mathematik 142 (2), 421–463.

Wormell, C. L., Gottwald, G. A., Sep 2018. On the validity of linear response theory in high-dimensional deterministic dynamical systems. Journal of Statistical Physics 172 (6), 1479–1498.

Wormell, C. L., Gottwald, G. A., 2019. Linear response for macroscopic observables in high-dimensional systems. Chaos: An Interdisciplinary Journal of Nonlinear Science 29 (11), 113127.

Young, L.-S., 1998. Statistical properties of dynamical systems with some hyperbolicity. Annals of Mathematics 147 (3), 585–650.

Young, L.-S., 2002. What are SRB measures, and which dynamical systems have them? Journal of Statistical Physics 108 (5-6), 733–754.

Zwanzig, R., 1973. Nonlinear generalized Langevin equations. J. Stat. Phys. 9, 215–220.

Zwanzig, R., 2001. Nonequilibrium Statistical Mechanics. Oxford University Press, Oxford.