1

Introduction

This chapter is a kind of overture. Simplified statements of three theorems are presented that set the tone for the whole text. Much broader versions of these theorems will appear later and several other themes around them will be introduced and developed as we move on. At this initial stage we choose to focus on the following special, yet significant, setting.

Let $A_1, \dots, A_m$ be invertible 2×2 real matrices and let $p_1, \dots, p_m$ be positive numbers with $p_1 + \dots + p_m = 1$. Consider

$$L^n = L_{n-1} \cdots L_1 L_0, \quad n \geq 1,$$

where the L_j are independent random variables with identical probability distributions, such that

$$\text{the probability of } \{L_j = A_i\} \text{ is equal to } p_i$$

for all $j \geq 0$ and $i = 1, \dots, m$. In brief, our goal is to describe the (almost certain) behavior of L^n as $n \to \infty$.

1.1 Existence of Lyapunov exponents

We begin with the following seminal result of Furstenberg and Kesten [56]:

Theorem 1.1 *There exist real numbers λ_+ and λ_- such that*

$$\lim_n \frac{1}{n} \log \|L^n\| = \lambda_+ \quad \text{and} \quad \lim_n \frac{1}{n} \log \|(L^n)^{-1}\|^{-1} = \lambda_-$$

with full probability.

The numbers λ_+ and λ_- are called *extremal Lyapunov exponents*. Clearly,

$$\lambda_+ \geq \lambda_- \tag{1.1}$$

because $\|B\| \geq \|B^{-1}\|^{-1}$ for any invertible matrix B. If B has determinant ± 1 then we even have $\|B\| \geq 1 \geq \|B^{-1}\|^{-1}$. Hence,

$$\lambda_+ \geq 0 \geq \lambda_- \tag{1.2}$$

when all matrices A_i, $1 \leq i \leq m$ have determinant ± 1.

1.2 Pinching and twisting

Next, we discuss conditions for the inequalities (1.1) and (1.2) to be strict.

Let $\mathscr{B}$ be the *monoid* generated by the matrices A_i, $i = 1, \dots, m$; that is, the set of all products $A_{k_1} \cdots A_{k_n}$ with $1 \leq k_j \leq m$ and $n \geq 0$ (for $n = 0$ interpret the product to be the identity matrix). We say that $\mathscr{B}$ is *pinching* if for any constant $\kappa > 1$ there exists some $B \in \mathscr{B}$ such that

$$\|B\| > \kappa \|B^{-1}\|^{-1}. \tag{1.3}$$

This means that the images of the unit circle under the elements of $\mathscr{B}$ are ellipses with arbitrarily large eccentricity. See Figure 1.1.

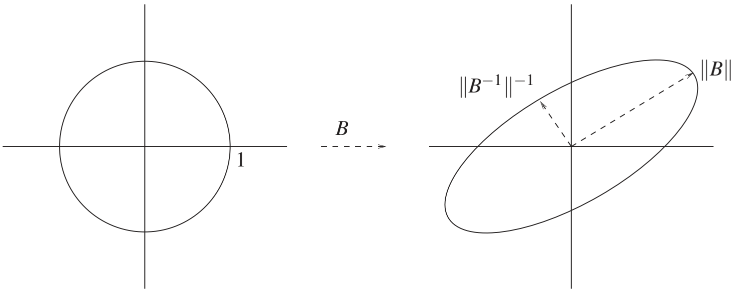

Figure 1.1 Eccentricity and pinching

We say that the monoid $\mathscr{B}$ is *twisting* if given any vector lines F, $G_1, \dots,$ $G_n \subset \mathbb{R}^2$ there exists $B \in \mathscr{B}$ such that

$$B(F) \notin \{G_1, \dots, G_n\}. \tag{1.4}$$

The following result is a variation of a theorem of Furstenberg [54]:

Theorem 1.2 *Assume $\mathscr{B}$ is pinching and twisting. Then $\lambda_- < \lambda_+$. In particular, if $|\det A_i| = 1$ for all $1 \leq i \leq m$ then both extremal Lyapunov exponents are different from zero.*

1.3 Continuity of Lyapunov exponents

The extremal Lyapunov exponents λ_+ and λ_- may be viewed as functions of the data

$$A_1, \ldots, A_m, p_1, \ldots, p_m.$$

Let the matrices A_j vary in the linear group $\mathrm{GL}(2)$ of invertible 2×2 matrices and the probability vectors $(p_1, \ldots, p_m)$ vary in the open simplex

$$\Delta^m = \{(p_1, \ldots, p_m) : p_1 > 0, \ldots, p_m > 0 \text{ and } p_1 + \cdots + p_m = 1\}.$$

The following result is part of a theorem of Bocker and Viana [35]:

Theorem 1.3 *The extremal Lyapunov exponents $\lambda_\pm$ depend continuously on $(A_1, \ldots, A_m, p_1, \ldots, p_m) \in \mathrm{GL}(2)^m \times \Delta^m$ at all points.*

Example 1.4 Let $m = 2$, with

$$A_1 = \begin{pmatrix} \sigma & 0 \\ 0 & \sigma^{-1} \end{pmatrix} \quad \text{and} \quad A_2 = R_\theta A_1 R_{-\theta}, \; R_\theta = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$$

for some $\sigma > 1$ and $\theta \in \mathbb{R}$. By Theorem 1.3, the Lyapunov exponents $\lambda_\pm$ depend continuously on the parameter σ and θ. Moreover, using Theorem 1.2, we have $\lambda_+ = 0$ if and only if $p_1 = p_2 = 1/2$ and $\theta = \pi/2 + n\pi$ for some $n \in \mathbb{Z}$.

1.4 Notes

Theorem 1.1 is a special case of the theorem of Furstenberg and Kesten [56], which is valid in any dimension $d \geq 2$. The full statement and the proof will appear in Chapter 3: we will deduce this theorem from an even more general statement, the subadditive ergodic theorem of Kingman [74]. Kingman's theorem will also be used in Chapter 4 to prove the fundamental result of the theory of Lyapunov exponents, the multiplicative ergodic theorem of Oseledets [92].

It is natural to ask whether the type of asymptotic behavior prescribed by Theorem 1.1 for the norm $\|L^n\|$ and conorm $\|(L^n)^{-1}\|^{-1}$ extends to the individual *matrix coefficients* $L^n_{i,j}$. Furstenberg, Kesten [56] proved that this is so if the coefficients of the matrices A_i, $1 \leq i \leq m$ are all strictly positive. The example in Exercise 1.3 shows that this assumption cannot be removed. On the other hand, the theorem of Oseledets theorem does contain such a description for the matrix *column vectors*.

Theorem 1.2 is also the tip of a series of fundamental results, which are to be discussed in Chapters 6 through 8. The full statement and proof of Furstenberg's theorem for 2-dimensional cocycles (Furstenberg [54]) will be given in

Chapter 6. The extension to any dimension will be stated and proved in Chapter 7: it will be deduced from the invariance principle (Ledrappier [81], Bonatti, Gomez-Mont and Viana [37], Avila and Viana [16], Avila, Santamaria and Viana [13]), a general tool that has several other applications, both for linear and nonlinear systems.

In dimension larger than 2, there is a more ambitious problem: rather than asking when $\lambda_- < \lambda_+$, one wants to know when *all* the Lyapunov exponents are distinct. That will be the subject of Chapter 8, which is based on Avila and Viana [14, 15].

Furstenberg and Kifer [57] proved continuity of the Lyapunov exponents of products of random matrices, restricted to the (almost) irreducible case. A variation of their argument will be given in Section 6.2.2. The reducible case requires a delicate analysis of the random walk defined by the cocycle in projective space. That was carried out by Bocker and Viana [35], in the 2-dimensional case, using certain discretizations of projective space. At the time of writing, Avila, Eskin and Viana [12] are extending the statement of the theorem to arbitrary dimension, using a very different strategy. The proof of Theorem 1.3 that we present in Chapter 10 is based on this more recent approach.

The problem of the dependence of Lyapunov exponents on the data can be formulated in the broader context of linear cocycles that we are going to introduce in Chapter 2. We will see in Chapter 9 that, in contrast, continuity often breaks down in that generality.

1.5 Exercises

The following elementary notions are used in some of the exercises that follow. We call a 2×2 matrix *hyperbolic* if it has two distinct real eigenvalues, *parabolic* if it has a unique real eigenvalue, with a one-dimensional eigenspace, and *elliptic* if it has two distinct complex eigenvalues. Multiples of the identity belong to neither of these three classes.

Exercise 1.1 Show that, in dimension $d = 2$, if $|\det A_i| = 1$ for all $1 \le i \le m$ then $\lambda_+ + \lambda_- = 0$.

Exercise 1.2 Calculate the extremal Lyapunov exponents for $m = 2$ and p_1, $p_2 > 0$ with $p_1 + p_2 = 1$ and

(1) $A_1 = \begin{pmatrix} \sigma & 0 \\ 0 & \sigma^{-1} \end{pmatrix}$ and $A_2 = \begin{pmatrix} \sigma^{-1} & 0 \\ 0 & \sigma \end{pmatrix}$, where $\sigma > 1$;

(2) $A_1 = \begin{pmatrix} \sigma & 0 \\ 0 & \sigma^{-1} \end{pmatrix}$ and $A_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, where $\sigma > 1$.

Exercise 1.3 (Furstenberg and Kesten [56]) Take $m = 2$ with $p_1 = p_2 = 1/2$ and

$$A_1 = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad A_2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Show that $\lim_n (1/n) \log |L^n_{i,j}|$ does not exist for any i, j, with full probability.

Exercise 1.4 Show that if some matrix A_i, $1 \le i \le m$ is either hyperbolic or parabolic then the monoid $\mathscr{B}$ is pinching.

Exercise 1.5 Show that the monoid $\mathscr{B}$ may be pinching even if all the matrices A_i, $1 \le i \le m$ are elliptic.

Exercise 1.6 Suppose that there exists $1 \le i \le m$ such that A_i is conjugate to an irrational rotation. Conclude that $\mathscr{B}$ is twisting.

Exercise 1.7 Suppose that there exist $1 \le i, j \le m$ such that A_i and A_j are either hyperbolic or parabolic and that they have no common eigenspace. Conclude that $\mathscr{B}$ is twisting (and pinching).

Exercise 1.8 Let A_i, $i = 1, 2$ be as in the second part of Exercise 1.2. Check that

$$\lambda_+(A_1, A_2, 1, 0) \ne \lim_{p_2 \to 0} \lambda_+(A_1, A_2, 1 - p_2, p_2).$$

Thus, the hypothesis $p_1 > 0, \ldots, p_m > 0$ cannot be removed in Theorem 1.3.

2

Linear cocycles

Linear cocycles are the basic object upon which this text is built. Here we define this concept and introduce a few examples. Special attention is given to (uniformly) hyperbolic cocycles, a class that is often used as a kind of paradigm for the behavior of more general systems.

Let $(M, \mathscr{B}, \mu)$ be a probability space and $f : M \to M$ be a measure-preserving map. Let $A : M \to \mathrm{GL}(d)$ be a measurable function with values in the *linear group* $\mathrm{GL}(d)$ of invertible $d \times d$ matrices with real coefficients. Sometimes we let A take values in the *special linear group* $\mathrm{SL}(d)$ of real $d \times d$ matrices with determinant ± 1. The *linear cocycle* defined by A over f is the transformation

$$F : M \times \mathbb{R}^d \to M \times \mathbb{R}^d, \quad (x, v) \mapsto (f(x), A(x)v). \tag{2.1}$$

Observe that $F^n(x, v) = (f^n(x), A^n(x))$ for every $n \geq 1$, where

$$A^n(x) = A(f^{n-1}(x)) \cdots A(f(x))A(x).$$

If f is invertible then so is F. Moreover, $F^{-n}(x, v) = (f^{-n}(x), A^{-n}(x))$ for all $n \geq 1$, where

$$A^{-n}(x) = A(f^{-n}(x))^{-1} \cdots A(f^{-1}(x))^{-1} = A^n(f^{-n}(x))^{-1}.$$

The Furstenberg–Kesten theorem (Theorem 1.1) extends to this setting, as follows: for any f-invariant probability measure μ such that $\log \|A^{\pm 1}\| \in L^1(\mu)$,

$$\lambda_+(x) = \lim_n \frac{1}{n} \log \|A^n(x)\| \quad \text{and} \quad \lambda_-(x) = \lim_n \frac{1}{n} \log \|A^n(x)^{-1}\|^{-1} \tag{2.2}$$

exist at μ-almost every point x. This fact will be proven in Chapter 3.

More generally, one may consider A to take values in the group $\mathrm{GL}(d, \mathbb{C})$ of invertible $d \times d$ matrices with complex coefficients, or the subgroup $\mathrm{SL}(d, \mathbb{C})$ of matrices with determinant in the unit circle. This gives rise to *complex linear cocycles* $M \times \mathbb{C}^d \to M \times \mathbb{C}^d$. Of course, every complex cocycle in dimension

d is also a real cocycle in dimension $2d$ and, conversely, every d-dimensional real linear cocycle defines a d-dimensional complex linear cocycle. The two theories, real and complex, are actually very similar. We focus on the real case, except where stated otherwise.

2.1 Examples

We illustrate this notion with three important classes of linear cocycles, arising from probability theory, dynamical systems, and spectral theory, respectively.

2.1.1 Products of random matrices

The situation we considered in Chapter 1 can be modeled by (a special case of) the following class of linear cocycles. Let $X = \mathrm{GL}(d)$ and $M = X^{\mathbb{Z}}$ (or $M = X^{\mathbb{N}}$) and

$$f : M \to M, \quad (\alpha_k)_k \mapsto (\alpha_{k+1})_k$$

be the shift map on X. Consider the function

$$A : M \to \mathrm{GL}(d), \quad (\alpha_k)_k \mapsto \alpha_0$$

and let $F : M \times \mathbb{R}^d \to M \times \mathbb{R}^d$ be the linear cocycle defined by A over f. Note that the kth iterate of F is given by

$$F^n\big((\alpha_k)_k, v\big) = \big((\alpha_{k+n})_k, \alpha_{n-1} \ldots \alpha_1 \alpha_0 v\big).$$

Given a probability measure p in the space $\mathrm{GL}(d)$, consider the product measure $\mu = p^{\mathbb{Z}}$ (or $\mu = p^{\mathbb{N}}$), which is characterized by

$$\mu\big(\{(\alpha_k)_k : \alpha_i \in E_i, \ldots, \alpha_j \in E_j\}\big) = p(E_i) \cdots p(E_j)$$

for every $i \le j$ and any measurable sets $E_1, \ldots, E_j \subset X$. It is clear that μ is invariant under the shift map.

We call a *locally constant linear cocycle* the following slightly more general construction. Let $(Y, \mathscr{Y}, q)$ be any probability space and then consider $N = Y^{\mathbb{Z}}$ endowed with the product σ-algebra $\mathscr{C} = \mathscr{Y}^{\mathbb{Z}}$ and the product measure $\nu = q^{\mathbb{Z}}$ (or $N = Y^{\mathbb{N}}$ endowed with $\mathscr{C} = \mathscr{Y}^{\mathbb{N}}$ and $\nu = q^{\mathbb{N}}$). Let $g : N \to N$ be the shift map. Moreover, let $B : N \to \mathrm{GL}(d)$ be any measurable function depending only on the zeroth coordinate; that is, of the form $B(y) = \beta(y_0)$ for some measurable function $\beta : Y \to \mathrm{GL}(d)$. Then consider the linear cocycle $G : N \times \mathbb{R}^d \to N \times \mathbb{R}^d$

defined by B over g. Note that G is semi-conjugate to a cocycle F as in the previous paragraph, with $p = \beta_* q$:

$$\begin{array}{ccc} N \times \mathbb{R}^d & \xrightarrow{G} & N \times \mathbb{R}^d \\ \Phi \downarrow & & \downarrow \Phi \\ M \times \mathbb{R}^d & \xrightarrow{F} & M \times \mathbb{R}^d \end{array}$$

with $\Phi\big((x_k)_k, v\big) = \big((\beta(x_k))_k, v\big)$. For this reason, the two cocycles are equivalent for most of our purposes.

2.1.2 Derivative cocycles

Consider a diffeomorphism $f : M \to M$ on the torus $M = \mathbb{T}^d$ of dimension $d \geq 1$. It is easy to construct smooth vector fields $X_1, \ldots, X_d$ on $\mathbb{T}^d$ such that $\{X_1(x), \ldots, X_d(x)\}$ is a basis of the tangent space $T_x M$, for every $x \in M$. One says that the torus is a *parallelizable* manifold. The *derivative cocycle* of f is

$$F : M \times \mathbb{R}^d \to M \times \mathbb{R}^d, \quad (x, v) \mapsto (f(x), A(x)v),$$

where $A(x) \in \mathrm{GL}(d)$ is the matrix, with respect to these bases, of the derivative $Df(x) : T_x M \to T_{f(x)} M$.

For more general diffeomorphisms, on non-parallelizable manifolds, the previous construction does not apply. However, one can still view the derivative map $Df : TM \to M$ as a linear cocycle, in the following more general sense. Let $\pi : \mathcal{V} \to M$ be a finite-dimensional vector bundle. This means that $\mathcal{V}$ is equipped with a family of homeomorphisms $h_\alpha : U_\alpha \times \mathbb{R}^d \to \pi^{-1}(U_\alpha)$ such that:

(i) $\{U_\alpha\}$ is an open cover of M;
(ii) $\pi \circ h_\alpha(x, v) = x$ for every $x \in U_\alpha$ and any α;
(iii) for every $x \in U_\alpha \cap U_\beta$ and any α, β, there exists a linear isomorphism $L_{\alpha,\beta}(x) : \mathbb{R}^d \to \mathbb{R}^d$ such that $h_\beta^{-1} \circ h_\alpha(x, v) = (x, L_{\alpha,\beta}(x)v)$ for every v.

The integer $d \geq 1$ is the *dimension* of the vector bundle. A *linear cocycle* on $\mathcal{V}$ over a transformation $f : M \to M$ is a measurable transformation $F : \mathcal{V} \to \mathcal{V}$ such that $\pi \circ F = f \circ \pi$ and the actions $F_x : \mathcal{V}_x \to \mathcal{V}_{f(x)}$ on the fibers are linear isomorphisms:

$$\begin{array}{ccc} \mathcal{V} & \xrightarrow{F} & \mathcal{V} \\ \pi \downarrow & & \downarrow \pi \\ M & \xrightarrow{f} & M \end{array}$$

For our purposes there is not much to gain from considering such generality. So, most of the time we will stick to the case of trivial fiber bundles; that is, to linear cocycles of the form (2.1).

2.1.3 Schrödinger cocycles

Consider $\ell^2 = \{(u_n)_{n\in\mathbb{Z}} : \sum_n |u_n|^2 < \infty\}$. The *Schrödinger operator* associated with a sequence $(V_n)_{n\in\mathbb{Z}}$ in $\mathbb{R}$, is defined by

$$H : \ell^2 \to \ell^2, \quad u = (u_n)_{n\in\mathbb{Z}} \mapsto H(u) = (u_{n+1} + u_{n-1} + V_n u_n)_{n\in\mathbb{Z}}. \tag{2.3}$$

In the most interesting models the sequence V_n is generated from a dynamical system $f : M \to M$ and a function $V : M \to \mathbb{R}$ (the so-called *potential*) through $V_n = V(f^n(x))$, for some $x \in M$. The most studied cases are:

(1) Random Schrödinger cocycles: Let $M = X^{\mathbb{Z}}$ be a shift space, $f : M \to M$ be the shift map, and $\mu = p^{\mathbb{Z}}$ be a Bernoulli measure on M. Fix $x \in M$ and then take $V_n = V(f^n(x))$, where the function $V : M \to \mathbb{R}$ is such that $V(x)$ depends only on the zeroth coordinate of $x \in M$.
(2) Quasi-periodic Schrödinger cocycles: Let μ the normalized Lebesgue measure on $M = \mathbb{T}^d$ and $f : \mathbb{T}^d \to \mathbb{T}^d$ be an irrational translation. Fix $x \in \mathbb{T}^d$ and take $V_n = V(f^n(x))$, where $V : \mathbb{T}^d \to \mathbb{R}$ is an analytic function.

A main objective is to understand the spectral theory of these operators (a theorem of Pastur [97] asserts that when the base system (f, μ) is ergodic the spectrum of H is the same for almost all choices of $x \in M$; the same is true for the absolutely continuous spectrum, the singular continuous spectrum and the pure point spectrum, by Kunz and Souillard [78]). Thus, one is led to studying the eigenvalue equation

$$H(u) = Eu, \quad \text{for } E \in \mathbb{R}. \tag{2.4}$$

While, by definition, the eigenvectors of H are the solutions of this equation in the space ℓ^2, it is useful to consider (2.4) for any real sequence $u = (u_n)_{n\in\mathbb{Z}}$. Note that the equation may be rewritten as

$$u_{n+1} + u_{n-1} + V(f^n(x))u_n = Eu_n$$

or, still equivalently,

$$\begin{pmatrix} u_{n+1} \\ u_n \end{pmatrix} = \begin{pmatrix} E - V(f^n(x)) & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} u_n \\ u_{n-1} \end{pmatrix}.$$

This suggests that we consider the linear cocycle $F_E : M \times \mathbb{R}^2 \to M \times \mathbb{R}^2$ defined over $f : M \to M$ by the function

$$A : M \to \mathrm{SL}(2), \quad A(y) = \begin{pmatrix} E - V(y) & -1 \\ 1 & 0 \end{pmatrix}.$$

We have just seen that $u = (u_n)_{n\in\mathbb{Z}}$ is a solution to $H(u) = Eu$ if and only if $U = (u_n, u_{n-1})_{n\in\mathbb{Z}}$ is a trajectory of the linear cocycle F_E. The behavior of these linear cocycles provides useful information about the spectral properties of the Schrödinger operator. For example, if the Lyapunov exponents of F_E are different from zero then E cannot be an eigenvalue of $H : \ell^2 \to \ell^2$. See Damanik [48] for much more information.

2.2 Hyperbolic cocycles

We are going to define an important class of cocycles whose behavior is particularly well understood. We focus on the two-dimensional setting, but we also comment briefly on the general case.

2.2.1 Definition and properties

Let M be a compact metric space and $f : M \to M$ be a homeomorphism. We call a continuous cocycle

$$F : M \times \mathbb{R}^2 \to M \times \mathbb{R}^2, \quad (x, v) \mapsto (f(x), A(x)v)$$

hyperbolic if there are $C > 0$ and $\lambda < 1$ and, for every $x \in M$, there exist transverse lines E^s_x and E^u_x in $\mathbb{R}^2$ such that

(1) $A(x)E^s_x = E^s_{f(x)}$ and $A(x)E^u_x = E^u_{f(x)}$
(2) $\|A^n(x)v^s\| \le C\lambda^n \|v^s\|$ and $\|A^{-n}(x)v^u\| \le C\lambda^n \|v^u\|$

for every $v^s \in E^s_x$, $v^u \in E^u_x$, $x \in M$, and $n \ge 1$.

Proposition 2.1 *Let $F : M \times \mathbb{R}^2 \to M \times \mathbb{R}^2$ be the linear cocycle defined by a continuous function $A : M \to \mathrm{SL}(2)$ over a homeomorphism $f : M \to M$. Then F is hyperbolic if and only if there exist constants $c > 0$ and $\sigma > 1$ such that $\|A^n(x)\| \ge c\sigma^n$ for all $x \in M$ and $n \ge 1$.*

Proof (We are going to use (2.2), whose proof will be given in Section 3.2. Similar arguments will appear in Section 3.4, for proving the multiplicative ergodic theorem in dimension 2.)