The Oseledets Multiplicative Ergodic Theorem in Lean 4

Marcel Morgenstern

June 29, 2026

Chapter 1

Introduction

This blueprint accompanies a *sorry-free* Lean 4 + Mathlib formalization of the **Oseledets multiplicative ergodic theorem** (MET) and a broad layer of companion results. It documents the mathematical content of the development and the dependency structure of its proof: every node carries a `\lean` annotation pointing at the corresponding Lean declaration, and a green checkmark records that the statement (and, where shown, its proof) is fully formalized. The API documentation generated by `doc-gen4` is linked from each node.

1.1 Setting

Throughout, (X, μ) is a probability space and $T : X \to X$ is an ergodic measure-preserving transformation. A *linear cocycle* of dimension d is generated by a measurable map

$$A : X \to \mathrm{Mat}_d(\mathbb{R}), \qquad \det(A\,x) \neq 0 \text{ for all } x,$$

whose n-step product along the orbit of x is

$$A^{(n)}(x) \;=\; A(T^{n-1}x) \cdots A(Tx)\, A(x), \qquad A^{(0)}(x) = I.$$

We impose the one-sided integrability condition

$$\log^{+} \|A\|,\ \log^{+} \|A^{-1}\| \in L^1(\mu),$$

where matrices act on EuclideanSpace $\mathbb{R}$ (Fin d) so that $\|\cdot\|$ is the L^2 operator norm (submultiplicative). For $v \neq 0$ the *Lyapunov exponent* in the direction v is the growth rate

$$\bar{\lambda}(x, v) \;=\; \limsup_{n \to \infty} \tfrac{1}{n} \log \left\| A^{(n)}(x)\, v \right\|.$$

1.2 The three headline theorems

The development proves three principal theorems, each formalized sorry-free and verified (by a guarded axiom audit) to depend only on the standard axioms {`propext`, `Classical.choice`, `Quot.sound`}.

- **One-sided MET (filtration form),** : there are finitely many distinct Lyapunov exponents $\lambda_1 > \cdots > \lambda_k$ and, for μ-a.e. x, a strictly decreasing, A-equivariant, measurable filtration
$$\text{EuclideanSpace } \mathbb{R} \text{ (Fin } d) = V_x^0 \supsetneq V_x^1 \supsetneq \cdots \supsetneq V_x^k = \{0\}$$

along which $\frac{1}{n}\log\left\|A^{(n)}(x)\,v\right\| \to \lambda_i$ for every $v \in V_x^i \setminus V_x^{i+1}$. This is the central result; it is proved in Chapter 5.

- **Two-sided splitting,** : when both the forward and backward filtrations are available, they are transverse and the space splits a.e. into an A-equivariant direct sum of *Oseledets subspaces* $\mathbb{R}^d = \bigoplus_i E_x^i$ on which the cocycle grows at the exact two-sided rate λ_i. Chapter 7.
- **Continuous-flow MET,** : the analogue for a continuous-time linear cocycle over a measurable flow, obtained by reduction to the time-one map together with a between-times sandwich estimate. Chapter 8.

1.3 Structure of the proof

The proof proceeds in layers, mirrored by the chapters of this blueprint. Chapter 2 sets up the cocycle, the operator norm and its measurability, and the Furstenberg–Kesten extremal exponents. Chapter 3 develops the ergodic-theoretic engine: the maximal ergodic inequality, the pointwise Birkhoff theorem, and—crucially—Kingman's subadditive ergodic theorem, which converts the subadditivity of the log-norms into almost-everywhere limits. Chapter 4 introduces the Lyapunov exponent as a measurable function, the Lyapunov spectrum, and the limsup filtration together with the measurability of its subspaces. Chapter 5 assembles the one-sided theorem through the Oseledets limit, the spectral upper bound and determinant squeeze, the spectral identification of the limsup filtration, and a top-gap envelope induction. The remaining chapters record the companion results (Chapter 6), the two-sided splitting (Chapter 7), and the continuous-flow theorem (Chapter 8).

Chapter 2

The linear cocycle and Furstenberg–Kesten

2.1 The measure-preserving system and the linear cocycle

The Oseledets theorem studies the long-term growth of products of matrices driven by a dynamical system. The data are a measure-preserving self-map $T : X \to X$ of a probability space (X, μ) and a matrix-valued generator $A : X \to \mathrm{Matrix}(\mathrm{Fin}\, d)(\mathrm{Fin}\, d)\, \mathbb{R}$. The fundamental object is the *iterated linear cocycle*: the product of the generator evaluated along the orbit of x. We use the convention that the newest factor sits on the left.

Definition 2.1 (The entrywise measurable structure on matrices)**.** A matrix is a function $m \to n \to \alpha$, and we equip $\mathrm{Matrix}\; m\, n\, \alpha$ with the Pi (product) σ-algebra induced from the entry type α. For finitely many entries over a Borel α this agrees with the Borel σ-algebra. Mathlib does not register this automatically because Matrix is a `def` rather than reducibly the underlying Pi type, so the ambient Pi instance does not transfer; we install it explicitly.

Definition 2.2 (The iterated linear cocycle)**.** Given $A : X \to \mathrm{Matrix}(\mathrm{Fin}\, d)(\mathrm{Fin}\, d)\, \mathbb{R}$ and $T : X \to X$, define $A^{(n)} =$ `cocycle` $A\, T\, n$ by recursion on n:

$$A^{(0)}(x) = 1, \qquad A^{(n+1)}(x) = A^{(n)}(Tx) \cdot A(x).$$

Unfolding, $A^{(n)}(x) = A(T^{n-1}x) \cdots A(Tx)\, A(x)$, with the newest factor on the left. The matrix norm throughout is the scoped L2 operator norm, which is submultiplicative; vectors live in $\mathrm{EuclideanSpace}\; \mathbb{R}\, (\mathrm{Fin}\, d)$ and the action is via `toEuclideanCLM`.

Theorem 2.3 (The cocycle identity)**.** *For all $m, n \in \mathbb{N}$ and $x \in X$,*

$$A^{(m+n)}(x) = A^{(m)}(T^n x) \cdot A^{(n)}(x).$$

Proof. Induction on n with x generalized. For $n = 0$ both sides equal $A^{(m)}(x)$. For the step, reassociate $m + (n + 1)$ and apply the defining recursion $A^{(k+1)}(x) = A^{(k)}(Tx) \cdot A(x)$ on both sides, then use the inductive hypothesis at Tx together with $T^{n+1}x = T^n(Tx)$ and associativity of matrix multiplication. □

Definition 2.4 (One-sided log-integrability of the generator)**.** The hypothesis `IntegrableLogNorm` $A\, \mu$ asserts that the positive part of the log-norm of the generator is integrable: $\log^+ \|A(\cdot)\| \in L^1(\mu)$, where $\log^+ t = \max(\log t, 0)$. This is the standard integrability assumption of the Furstenberg–Kesten and Oseledets theorems; combined with the same hypothesis for the inverse generator A^{-1} it pins both extremal Lyapunov exponents in $\mathbb{R}$.

Theorem 2.5 (Measurability of the cocycle iterates)**.** *If A and T are measurable then for each n the iterate $x \mapsto A^{(n)}(x)$ is measurable for the entrywise structure 2.1.*

Proof. Induction on n. The base case is a constant map. For the step, the recursion writes $A^{(n+1)}$ as the product $(A^{(n)} \circ T) \cdot A$; matrix multiplication is jointly measurable on the Pi structure because each entry of a product is the finite sum $\sum_k M_{ik} N_{kj}$ of products of measurable coordinate projections, so the result follows by composing the inductive hypothesis with T and multiplying by A. ☐

2.2 Measurability of the operator norm and the inverse

To feed the cocycle into the ergodic theory we must know that the operator norm and the matrix inverse are measurable on the entrywise structure. The subtlety is that Mathlib's `Measurable.norm` is stated for a `BorelSpace`, whereas our matrix σ-algebra is the Pi structure; the two coincide here because the L2 operator-norm topology is definitionally the Pi product topology.

Lemma 2.6 (The Pi structure is an opens-measurable space)**.** *The entrywise (Pi) measurable structure on* $\mathrm{Matrix}(\mathrm{Fin}\, d)(\mathrm{Fin}\, d)\, \mathbb{R}$ *is an `OpensMeasurableSpace` for the L2 operator-norm topology, since that topology is installed (via `replaceTopology`) to be definitionally the Pi product topology, of which the Pi σ-algebra is exactly the Borel structure.*

Theorem 2.7 (Measurability of the L2 operator norm)**.** *The map $M \mapsto \|M\|$ on* $\mathrm{Matrix}(\mathrm{Fin}\, d)(\mathrm{Fin}\, d)\, \mathbb{R}$ *is measurable.*

Proof. The norm is continuous for the operator-norm topology, and by 2.6 that topology's Borel structure is the entrywise Pi structure; continuous maps into a Borel codomain are measurable. ☐

Theorem 2.8 (Measurability of the determinant)**.** *The determinant $M \mapsto \det M$ is measurable.*

Proof. By the Leibniz formula $\det M = \sum_\sigma \mathrm{sgn}(\sigma) \prod_i M_{i,\sigma(i)}$, the determinant is a finite sum of finite products of measurable coordinate projections, hence measurable. ☐

Theorem 2.9 (Measurability of the matrix inverse)**.** *The inverse $M \mapsto M^{-1}$ is measurable on the entrywise structure.*

Proof. Writing $M^{-1} = (\det M)^{-1} \cdot \mathrm{adj}(M)$, each entry is a ratio of polynomials in the entries. The adjugate is measurable (each entry is a determinant of a row update, again a polynomial in the entries), and $(\det M)^{-1}$ is measurable by 2.8 and measurability of inversion on $\mathbb{R}$; the product is therefore measurable entrywise. ☐

2.3 Positivity and submultiplicativity of the log-norm

To take logarithms cleanly we need the iterate norms to be strictly positive, which forces the generator to be everywhere invertible ($\det A \neq 0$) and the dimension to be nonzero. These hypotheses also make the log of a product split as a genuine sum, which is what yields subadditivity rather than a mere inequality with junk values.

Lemma 2.10 (Invertibility of the iterates)**.** *If* $\det A(x) \neq 0$ *for every* x*, then* $\det A^{(n)}(x) \neq 0$ *for all* n, x*.*

Proof. Induction on n. The base case is $\det 1 = 1$. For the step, $\det(A^{(n)}(Tx) \cdot A(x)) = \det A^{(n)}(Tx) \cdot \det A(x)$, and both factors are nonzero by the inductive hypothesis and the assumption on A. ☐

Lemma 2.11 (The unit matrix has norm one)**.** *When* $d \neq 0$*,* $\|(1 : \mathrm{Matrix}(\mathrm{Fin}\, d)(\mathrm{Fin}\, d)\, \mathbb{R})\| = 1$ *for the L2 operator norm.*

Proof. For $d \neq 0$ the space EuclideanSpace $\mathbb{R}\,(\mathrm{Fin}\, d)$ is nontrivial. The star-algebra equivalence `toEuclideanCLM` sends the identity matrix to the identity operator and preserves the norm, and the operator norm of the identity on a nontrivial space is 1. There is no `NormOneClass` instance for the matrix operator norm, so this must be proved by hand; note that at $d = 0$ the statement is false ($\|1\| = 0$), which is why $d \neq 0$ is required. ☐

Lemma 2.12 (Positivity of the iterate norms)**.** *Assume* $\det A(x) \neq 0$ *for all* x *and* $d \neq 0$*. Then* $0 < \|A^{(n)}(x)\|$ *for every* n, x*. The analogous statement* `norm_inv_cocycle_pos` *holds for the inverse iterates* $\|(A^{(n)}(x))^{-1}\|$*.*

Proof. If $\|A^{(n)}(x)\| = 0$ then $A^{(n)}(x)$ is the zero matrix, whose determinant vanishes (as $d \neq 0$), contradicting 2.10. The inverse case is identical using $\det((A^{(n)}(x))^{-1}) = (\det A^{(n)}(x))^{-1} \neq 0$. ☐

Theorem 2.13 (Subadditivity of the log-norm cocycle)**.** *Assume* $\det A \neq 0$ *everywhere and* $d \neq 0$*. Then* $g_n(x) = \log \|A^{(n)}(x)\|$ *is a subadditive cocycle over* T*:*

$$g_{m+n}(x) \leq g_m(x) + g_n(T^m x).$$

Proof. Rewrite $m+n$ as $n+m$ and apply the cocycle identity 2.3 to get $A^{(m+n)}(x) = A^{(n)}(T^m x) \cdot A^{(m)}(x)$. Submultiplicativity of the operator norm and monotonicity of log give $\log \|A^{(m+n)}(x)\| \leq \log(\|A^{(n)}(T^m x)\| \cdot \|A^{(m)}(x)\|)$; both factors are strictly positive by 2.12, so log of the product splits as the sum $\log \|A^{(n)}(T^m x)\| + \log \|A^{(m)}(x)\|$, which is the required bound after commuting the summands. ☐

Theorem 2.14 (Subadditivity of the inverse log-norm cocycle)**.** *Under the same hypotheses,* $g_n(x) = \log \|(A^{(n)}(x))^{-1}\|$ *is a subadditive cocycle over* T*.*

Proof. As above, after 2.3 the inverse of the product reverses the order: $(A^{(m+n)}(x))^{-1} = (A^{(m)}(x))^{-1} \cdot (A^{(n)}(T^m x))^{-1}$. Submultiplicativity, monotone log, strict positivity of each inverse norm (2.12), and log of a product splitting as a sum give the subadditive inequality. ☐

2.4 Birkhoff-sum sandwich bounds

The two-sided integrability hypothesis controls the log-norm by Birkhoff sums of $\log^+ \|A\|$ and $\log^+ \|A^{-1}\|$ on both sides. These bounds drive both the integrability of each level and the bounded-below proviso of Kingman's theorem.

Lemma 2.15 (Upper Fekete bound)**.** *For* $\det A \neq 0$ *everywhere and* $d \neq 0$*,*

$$\log \|A^{(n)}(x)\| \leq \sum_{k<n} \log^+ \|A(T^k x)\|.$$

Proof. Induction on n. The base case is $0 \le 0$. For the step, the recursion gives $A^{(n+1)}(x) = A^{(n)}(Tx) \cdot A(x)$; submultiplicativity and log-splitting bound $\log\|A^{(n+1)}(x)\|$ by $\log\|A(x)\| + \log\|A^{(n)}(Tx)\|$, and then $\log\|A(x)\| \le \log^+\|A(x)\|$ together with the inductive hypothesis at Tx peels off one Birkhoff term. ☐

Lemma 2.16 (Lower bound via the inverse Birkhoff sum)**.** *Under the same hypotheses,*

$$-\sum_{k<n} \log^+\|(A(T^k x))^{-1}\| \le \log\|A^{(n)}(x)\|.$$

Proof. From $A^{(n)}(x) \cdot (A^{(n)}(x))^{-1} = 1$ and submultiplicativity, $1 = \|1\| \le \|A^{(n)}(x)\| \cdot \|(A^{(n)}(x))^{-1}\|$ (using 2.11 and 2.10), so taking logs gives $0 \le \log\|A^{(n)}(x)\| + \log\|(A^{(n)}(x))^{-1}\|$. Bounding the inverse log-norm above by its Birkhoff sum 2.17 and rearranging yields the claim. ☐

Lemma 2.17 (Upper Fekete bound for the inverse cocycle)**.** *Under the same hypotheses,*

$$\log\|(A^{(n)}(x))^{-1}\| \le \sum_{k<n} \log^+\|(A(T^k x))^{-1}\|.$$

Proof. Induction on n, exactly as for the forward bound but with the inverse generator: the recursion and `mul_inv_rev` give $(A^{(n+1)}(x))^{-1} = (A(x))^{-1} \cdot (A^{(n)}(Tx))^{-1}$, and submultiplicativity plus log-splitting and $\log\|(A(x))^{-1}\| \le \log^+\|(A(x))^{-1}\|$ peel off one Birkhoff term. ☐

2.5 Integrability of each level

Lemma 2.18 (Integral of a Birkhoff sum)**.** *For measure-preserving T and integrable f, $\int \sum_{k<n} f(T^k x)\, d\mu = n \int f\, d\mu$.*

Proof. Each composition $f \circ T^k$ is integrable and integral-preserving because T^k is measure-preserving, so $\int f \circ T^k\, d\mu = \int f\, d\mu$. Summing the n equal terms over the finite range gives $n \int f\, d\mu$. ☐

Theorem 2.19 (Integrability of the log-norm levels)**.** *Let T be measure-preserving for a finite measure μ, A measurable and everywhere invertible, $d \ne 0$, with both $\log^+\|A\|$ and $\log^+\|A^{-1}\|$ integrable. Then each $x \mapsto \log\|A^{(n)}(x)\|$ is integrable. The companion `integrable_logNorm_inv_cocycle` gives the same for the inverse iterates.*

Proof. The level g_n is sandwiched between $-B_n^-$ and B_n^+, where $B_n^\pm$ are the Birkhoff sums of $\log^+\|A\|$ and $\log^+\|A^{-1}\|$ (Lemmas 2.15, 2.16). Both $B_n^\pm$ are nonnegative and integrable (finite sums of integrable, m.p. precompositions), so $|g_n| \le B_n^+ + B_n^-$ pointwise. Since g_n is measurable (2.7 composed with the cocycle), domination by the integrable $B_n^+ + B_n^-$ gives integrability. ☐

2.6 The Furstenberg–Kesten theorems

Feeding the subadditive cocycles into Kingman's ergodic theorem produces the two extremal Lyapunov exponents as a.e.-constant limits. The bounded-below proviso of Kingman is supplied by the cross integrability hypothesis: for the top exponent the lower bound comes from $\log^+\|A^{-1}\|$, and vice versa. This is precisely where the second integrability hypothesis keeps the limit finite in $\mathbb{R}$ rather than $-\infty$.

Theorem 2.20 (Furstenberg–Kesten, top exponent)**.** *Let T be ergodic for a probability measure μ, let A be measurable and everywhere invertible with $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1(\mu)$. Then there is a constant $\lambda_1 \in \mathbb{R}$ (the top Lyapunov exponent) with*

$$\frac{1}{n} \log \left\|A^{(n)}(x)\right\| \xrightarrow{n\to\infty} \lambda_1 \qquad \textit{for } \mu\textit{-a.e. } x.$$

Proof. If $d = 0$ the matrix algebra is trivial, every norm is 0, and the limit is the constant 0. Otherwise set $d \neq 0$ and let $g_n = \log \left\|A^{(n)}\right\|$. It is a subadditive cocycle (2.13) with each level integrable (2.19). For the bounded-below proviso, the lower bound 2.16 and the Birkhoff integral identity 2.18 give $(\int g_{n+1})/(n+1) \geq -\int \log^+ \left\|A^{-1}\right\| \, d\mu$, a constant lower bound. The ergodic Kingman theorem 3.26 then returns the a.e. constant limit, which is the top exponent λ_1 (see 4.8). ☐

Theorem 2.21 (Furstenberg–Kesten, bottom exponent)**.** *Under the same hypotheses there is a constant $\lambda \in \mathbb{R}$ with*

$$\frac{1}{n} \log \left\|(A^{(n)}(x))^{-1}\right\| \xrightarrow{n\to\infty} \lambda \qquad \textit{for } \mu\textit{-a.e. } x,$$

so the bottom Lyapunov exponent $\lambda_k = -\lambda$ exists and is finite.

Proof. Identical to the top case with the inverse subadditive cocycle $g_n = \log \left\|(A^{(n)})^{-1}\right\|$ (2.14). The roles of the two integrability hypotheses swap: the bounded-below proviso now uses the lower bound $-\sum_{k<n} \log^+ \left\|A(T^k x)\right\| \leq \log \left\|(A^{(n)}(x))^{-1}\right\|$, obtained from $\|1\| \leq \left\|A^{(n)}\right\| \cdot \left\|(A^{(n)})^{-1}\right\|$ (2.11) and the forward Fekete bound 2.15, together with the integral identity 2.18. Ergodic Kingman 3.26 delivers the a.e. constant limit. ☐

Chapter 3

Ergodic theorems

This chapter develops the three classical ergodic theorems that drive the multiplicative ergodic theorem: the *maximal ergodic inequality* of Hopf and Garsia, the *pointwise (Birkhoff) ergodic theorem*, and *Kingman's subadditive ergodic theorem*. The first is the analytic gateway; the second turns the gateway into almost-everywhere convergence of Birkhoff averages; the third is the genuine engine of Oseledets, turning a subadditive sequence of log-norms into an almost-everywhere limit.

Throughout, $T : X \to X$ is a map on a measurable space carrying a measure μ, $\texttt{birkhoffSum}\, T\, g\, n\, x = \sum_{k<n} g(T^{[k]}x)$ is the n-th Birkhoff partial sum, and $\texttt{birkhoffAverage}\, \mathbb{R}\, T\, g\, n\, x = n^{-1}\, \texttt{birkhoffSum}\, T\, g\, n\, x$ is the Birkhoff average. The arguments are due to Garsia, Katznelson–Weiss, and Karlsson; the formalization keeps every quantity in $\mathbb{R}$ (or, where $-\infty$ cannot be excluded a priori, in `EReal`) to avoid junk values.

3.1 Maximal ergodic inequality

Definition 3.1 (Garsia's maximal function)**.** For a measurable T, a function $g : X \to \mathbb{R}$, $N : \mathbb{N}$ and x, the *maximal function* is

$$\texttt{maxBirkhoff}\, T\, g\, N\, x = \max_{0 \le k \le N} \texttt{birkhoffSum}\, T\, g\, k\, x,$$

the nonempty `Finset.sup'` of the Birkhoff partial sums over $k \in \texttt{range}\,(N+1)$. Since the $k = 0$ term equals $\texttt{birkhoffSum}\, T\, g\, 0\, x = 0$, the maximal function is its own positive part.

Lemma 3.2 (Nonnegativity of the maximal function)**.** *For all* g, N, x *one has* $0 \le \texttt{maxBirkhoff}\, T\, g\, N\, x$.

Proof. The index $k = 0$ lies in $\texttt{range}\,(N+1)$, and the corresponding term is $\texttt{birkhoffSum}\, T\, g\, 0\, x = 0$. The supremum over a nonempty finite set dominates each of its terms, so the maximal function is at least 0. ☐

Lemma 3.3 (Recursion for the maximal function)**.** *For all* g, N, x*,*

$$\texttt{maxBirkhoff}\, T\, g\, (N+1)\, x = \texttt{birkhoffSum}\, T\, g\, (N+1)\, x \sqcup \texttt{maxBirkhoff}\, T\, g\, N\, x.$$

Proof. This is the identity $\texttt{range}\,(N+2) = \texttt{insert}\,(N+1)\,(\texttt{range}\,(N+1))$ applied to the supremum: splitting off the top index gives the join of the new term with the previous maximum. Both inequalities follow from the universal/existential characterizations of `Finset.sup'`. The recursion is what makes integrability provable by induction on N. ☐

Lemma 3.4 (Garsia's pointwise inequality)**.** *On the set where the maximal function is positive, for every x with $0 < \mathtt{maxBirkhoff}\, T\, g\, N\, x$,*

$$\mathtt{maxBirkhoff}\, T\, g\, N\, x \leq g\, x + \mathtt{maxBirkhoff}\, T\, g\, N\, (T x).$$

Proof. Pull the constant $g\, x$ through the supremum: by the additive shift of a nonempty `sup'` and the recursion $\mathtt{birkhoffSum}\, T\, g\, (k+1)\, x = g\, x + \mathtt{birkhoffSum}\, T\, g\, k\, (T x)$, the right-hand side equals $\max_{0 \leq k \leq N} \mathtt{birkhoffSum}\, T\, g\, (k+1)\, x$, the maximum of the *shifted* partial sums. The maximum defining the left-hand side is attained at some index k_0; positivity forces $k_0 \geq 1$ (the index 0 gives the value 0), so $\mathtt{birkhoffSum}\, T\, g\, k_0\, x$ is one of the shifted sums and is therefore dominated by their maximum. ☐

Lemma 3.5 (Measurability and integrability of the maximal function)**.** *If T is measure-preserving and g is integrable, then each Birkhoff partial sum and the maximal function $\mathtt{maxBirkhoff}\, T\, g\, N$ are integrable; if moreover g and T are measurable, then $\mathtt{maxBirkhoff}\, T\, g\, N$ is measurable.*

Proof. Each summand $g \circ T^{[j]}$ is integrable because $T^{[j]}$ is measure-preserving, so the finite Birkhoff sum is integrable; measurability is likewise a finite sum of measurable compositions. Integrability of the maximal function follows by induction on N using the recursion theorem 3.3 and the fact that a join of two integrable functions is integrable; measurability is the measurability of a finite `sup'`. ☐

Proposition 3.6 (Garsia's inequality at a fixed level)**.** *For a measure-preserving T and a measurable integrable g, and every $N : \mathbb{N}$,*

$$0 \leq \int_{\{x \,:\, 0 < \mathtt{maxBirkhoff}\, T\, g\, N\, x\}} g \, d\mu.$$

Proof. Write E for the level set and $M = \mathtt{maxBirkhoff}\, T\, g\, N$. Integrating the pointwise inequality theorem 3.4 over E gives $\int_E (M - M \circ T) \leq \int_E g$. Now $\int_E M \circ T \leq \int_X M \circ T = \int_X M$ by nonnegativity and measure-preservation (via `integral_map` and $\mathtt{map}\, T\, \mu = \mu$, avoiding any embedding hypothesis), and $\int_X M = \int_E M$ since M vanishes off E. Chaining these cancels the two maximal-function integrals and leaves $0 \leq \int_E (M - M \circ T) \leq \int_E g$. ☐

Lemma 3.7 (Level sets exhaust the target set)**.** *The level sets $\{x : 0 < \mathtt{maxBirkhoff}\, T\, g\, N\, x\}$ are monotone in N, and their union over all N is exactly $\{x : \exists n,\ 0 < \mathtt{birkhoffSum}\, T\, g\, (n{+}1)\, x\}$.*

Proof. Monotonicity is `Finset.sup'` monotonicity over the nested ranges. For the union: $0 < \mathtt{maxBirkhoff}\, T\, g\, N\, x$ means some $\mathtt{birkhoffSum}\, T\, g\, k\, x > 0$ with $k \leq N$; the index 0 gives 0 and is excluded, so $k = n + 1$ for some n. Conversely a positive partial sum at index $n + 1$ makes the maximal function positive at level $N = n + 1$. ☐

Theorem 3.8 (Maximal ergodic inequality (Hopf–Garsia))**.** *For a measure-preserving T and an integrable $f : X \to \mathbb{R}$,*

$$0 \leq \int_{\{x \,:\, \exists n,\ 0 < \mathtt{birkhoffSum}\, T\, f\, (n+1)\, x\}} f \, d\mu.$$

Proof. First treat a measurable integrable g: by theorem 3.6 each level-set integral is nonnegative, and these integrals converge to the integral over the union theorem 3.7 by monotone convergence of set integrals, so the limit is nonnegative. For general integrable f, pass to a measurable representative $g =^{\text{a.e.}} f$; the Birkhoff sums of f and g agree a.e., so the two target sets agree a.e. (the one for g being measurable, the one for f null-measurable), and the integrands agree a.e., which transfers the inequality from g to f. ☐

3.2 Birkhoff

Theorem 3.9 (Conditional expectation commutes with the dynamics)**.** *For a finite measure, a measure-preserving measurable T and integrable g, the conditional expectation onto the σ-algebra $\mathtt{invariants}\,T$ of T-invariant sets is a.e. T-invariant:*

$$(\mu[g \mid \mathtt{invariants}\,T]) \circ T =^{a.e.} \mu[g \mid \mathtt{invariants}\,T].$$

Proof. Set-integral invariance of $h \circ T$ over a measurable invariant set reduces, via $\mathtt{map}\,T\,\mu = \mu$, to the integral of h. Hence $\mu[g \circ T \mid \mathcal{I}]$ and $\mu[g \mid \mathcal{I}]$ have equal integrals over every invariant set, so by uniqueness of the conditional expectation they agree a.e.; combining this with the commutation identity $\mu[g \circ T \mid \mathcal{I}] =^{\text{a.e.}} (\mu[g \mid \mathcal{I}]) \circ T$ (again by uniqueness, using that T is $(\mathcal{I}, \mathcal{I})$-measurable) yields the invariance. $\square$

Lemma 3.10 (Subexponential orbital tail)**.** *For a measure-preserving T and integrable g, the orbital tail $n^{-1}\,g(T^{[n]}x)$ tends to 0 for μ-a.e. x.*

Proof. A Borel–Cantelli argument. For each threshold $\delta = 1/(k+1)$ the series $\sum_n \mu\{x : (n+1)\delta \le |g\,x|\}$ is finite: the pointwise count of crossed thresholds is at most $|g\,x|/\delta$, and integrating (Tonelli) bounds the series by $\delta^{-1}\int |g|$. Measure-preservation transfers finiteness to the shifted orbit $g \circ T^{[n]}$, so a.e. only finitely many n cross any fixed threshold; choosing k large makes the tail eventually smaller than any ε. $\square$

Lemma 3.11 (A.e. boundedness of Birkhoff averages)**.** *For a finite measure, a measure-preserving T and integrable g, the Birkhoff averages $n \mapsto \mathtt{birkhoffAverage}\,\mathbb{R}\,T\,g\,(n+1)\,x$ are a.e. bounded above (and, applied to $-g$, a.e. bounded below).*

Proof. The maximal ergodic inequality applied to $g - c$ yields, for the maximal set $B_c = \{x : \exists n,\ c < \mathtt{birkhoffAverage}\,\mathbb{R}\,T\,g\,(n+1)\,x\}$, the estimate $c\,\mu B_c \le \int_{B_c} g \le \int |g|$. Taking $c = k \in \mathbb{N}$ gives $\mu B_k \le \int |g|/k \to 0$, so the intersection $\bigcap_k B_k$ is null. Off this null set the range of Birkhoff averages is bounded above. $\square$

Lemma 3.12 (Limsup invariance and vanishing perturbations)**.** *The pointwise limsup $x \mapsto \limsup_n \mathtt{birkhoffAverage}\,\mathbb{R}\,T\,g\,n\,x$ is a.e. T-invariant.*

Proof. The difference $A_n(g)(Tx) - A_n(g)(x) = n^{-1}(g(T^{[n]}x) - g\,x)$ tends to 0 a.e. by the tail estimate theorem 3.10. Two bounded sequences differing by a null sequence have equal limsup (proved via `limsup_add_const` and an ε-argument), and boundedness holds a.e. at both x and Tx by theorem 3.11; hence the limsup is a.e. T-invariant. $\square$

Proposition 3.13 (The core maximal-inequality step)**.** *For a finite measure, measure-preserving T, integrable g and $\varepsilon > 0$, the superlevel set where the limsup of the Birkhoff averages exceeds $\mu[g \mid \mathcal{I}] + \varepsilon$ is null.*

Proof. Write $L = \mu[g \mid \mathcal{I}]$ and $Ls = \limsup A_\bullet(g)$. The set $E = \{L + \varepsilon < Ls\}$ is a.e. T-invariant (theorem 3.9, theorem 3.12), hence a.e. equal to a genuinely invariant measurable set E'. Feed $\varphi = \mathbf{1}_{E'}(g - L - \varepsilon)$ to the maximal ergodic inequality theorem 3.8; on E' the partial sums of φ telescope (using orbit-constancy of L), and the maximal set equals E'. Thus $0 \le \int_{E'}(g - L - \varepsilon) = -\varepsilon\,\mu E'$ (using $\int_{E'} g = \int_{E'} L$), forcing $\mu E' = 0$. $\square$

Theorem 3.14 (Pointwise (Birkhoff) ergodic theorem)**.** *For a finite measure, a measure-preserving T and an integrable g, the Birkhoff averages converge μ-a.e. to the conditional expectation of g onto the invariant σ-algebra:*

$$\mathtt{birkhoffAverage}\,\mathbb{R}\,T\,g\,n\,x \to (\mu[g \mid \mathtt{invariants}\,T])(x).$$

Proof. Unioning the null superlevel sets of theorem 3.13 over $\varepsilon = 1/(k+1)$ gives $\limsup A_\bullet(g) \le \mu[g \mid \mathcal{I}]$ a.e.; applying the same to $-g$ (and using $\limsup(-a) = -\liminf a$) gives $\mu[g \mid \mathcal{I}] \le \liminf A_\bullet(g)$ a.e. With a.e. boundedness theorem 3.11, the sandwich $\limsup \le \mu[g \mid \mathcal{I}] \le \liminf$ forces convergence to $\mu[g \mid \mathcal{I}]$. ☐

Corollary 3.15 (Ergodic case of Birkhoff)**.** *For an ergodic T on a probability space and integrable g, the Birkhoff averages converge μ-a.e. to the space average $\int g\,d\mu$.*

Proof. By theorem 3.14 the limit is $\mu[g \mid \mathcal{I}]$, which is a.e. T-invariant theorem 3.9; ergodicity forces it a.e. constant. Its integral equals $\int g$ (conditional expectation preserves the integral) and equals the constant times $\mu(X) = 1$, so the constant is $\int g$. ☐

3.3 Kingman

Definition 3.16 (Subadditive cocycle)**.** A sequence $g : \mathbb{N} \to X \to \mathbb{R}$ is a *subadditive cocycle* over T when

$$g\,(m+n)\,x \le g\,m\,x + g\,n\,(T^{[m]}x) \qquad \text{for all } m, n, x.$$

For $g_n = \log\|A^{(n)}\|$ this is submultiplicativity of the operator norm composed with the cocycle identity.

Lemma 3.17 (Singleton and block subadditivity)**.** *For a subadditive cocycle and $n : \mathbb{N}$, one has $g\,(n+1)\,x \le \texttt{birkhoffSum}\,T\,(g\,1)\,(n+1)\,x$; more generally, for any decomposition of $[0, N)$ into $k+1$ consecutive blocks of lengths $\ell_0, \dots, \ell_k$, the cocycle is dominated by the sum of the block values along the orbit at the frontiers $T^{[\sum_{j<i} \ell_j]}x$.*

Proof. Both are inductions that peel off the last block and apply the defining subadditivity at the split point. The statement is restricted to nonempty decompositions (and to index $n+1$) because subadditivity at $(0,0)$ only forces $0 \le g\,0\,x$, the wrong sign for a one-sided bound at 0. ☐

Definition 3.18 (Normalized cocycle)**.** The *normalized cocycle* is $\texttt{cdiv}\,g\,n\,x = g\,(n+1)\,x/(n+1)$, the sequence whose a.e. limit is the content of Kingman's theorem; its `EReal` coercion is $\texttt{ecdiv}\,g\,n\,x$, used where $-\infty$ cannot be excluded a priori.

Lemma 3.19 (Fekete limit of the normalized integrals)**.** *For a measure-preserving T, an integrable subadditive cocycle g whose normalized integrals are bounded below, the sequence $(\int g\,(n+1)\,d\mu)/(n+1)$ converges to the Fekete constant γ.*

Proof. Integrating the cocycle inequality and using measure-preservation ($\int g\,n\,(T^{[m]}\cdot) = \int g\,n$) shows the integral sequence $a_n = \int g\,n$ is subadditive in Fekete's sense. The $(n+1)$-indexed lower bound is bridged by hand to a lower bound on a_n/n (the $n = 0$ term is 0), and Fekete's lemma delivers convergence of a_n/n, hence of the shifted sequence, to $\gamma = \inf_n a_n/n$. ☐

Lemma 3.20 (Invariance from a one-sided orbital bound)**.** *For a finite measure, measure-preserving T and an a.e. measurable F with $F\,x \le F\,(Tx)$ for a.e. x, one has $F \circ T =^{a.e.} F$.*

Proof. For each rational c the upper level set $\{c \le F\}$ is null-measurable (via a measurable representative) and a.e. contained in its preimage $T^{-1}\{c \le F\}$, which has equal finite measure; an a.e. subset of equal finite measure is a.e. equal. Ranging over all rational c and collecting the a.e. statements gives the invariance. ☐

Lemma 3.21 (Envelopes are a.e. measurable and T-invariant)**.** *For a finite measure, measure-preserving T, an integrable subadditive cocycle with normalized integrals bounded below, the limsup envelope $f_+(x) = \limsup_n \mathtt{cdiv}\, g\, n\, x$ (and likewise the liminf envelope $f\_$) is a.e. T-invariant.*

Proof. The subadditive bound $\mathtt{cdiv}\, g\, n\, x \le g\, 1\, x/(n+1) + g\, n\, (Tx)/(n+1)$ differs from $\mathtt{cdiv}\, g\, n\, (Tx)$ by a null sequence, so the vanishing-perturbation lemma for limsup gives the pointwise comparison $f_+(x) \le f_+(Tx)$ wherever the cocycle is a.e. bounded (above and below) at x and Tx. Feeding this into theorem 3.20 yields a.e. invariance. ☐

Lemma 3.22 (Integrability of the limsup envelope)**.** *Under the same hypotheses, the limsup envelope f_+ is integrable.*

Proof. The nonnegative Fatou defect $d_n(x) = \mathtt{birkhoffAverage}\, \mathbb{R}\, T\, (g\, 1)\, (n+1)\, x - \mathtt{cdiv}\, g\, n\, x \ge 0$ (singleton subadditivity) controls the envelope. The `ENNReal` Fatou inequality $\int^{-} \liminf u_n \le \liminf \int^{-} u_n$ applied to $u_n = \mathtt{ofReal}\, d_n$, together with $\int d_n = \int g\, 1 - a_{n+1}/(n+1) \to \int g\, 1 - \gamma < \infty$ (theorem 3.19) and Birkhoff convergence of $A_{n+1}(g\, 1)$ to a loose envelope B (theorem 3.14), shows $B - f_+ \ge 0$ has finite lower integral, hence is integrable; therefore $f_+ = B - (B - f_+)$ is integrable. ☐

Proposition 3.23 (Hard direction: $\limsup \le \liminf$ a.e.)**.** *Under the same hypotheses, for a.e. x the* ***`EReal`*** *limsup of the normalized cocycle is dominated by its liminf.*

Proof. This is the stopping-time / greedy block argument of Katznelson–Weiss and Karlsson. After the WLOG shift to the nonpositive process $\tilde{g}\,(n+1) \le 0$, fix $\varepsilon, M > 0$ and set $h = \mu[\max(f\_, -M) \mid \mathcal{I}]$. A greedy two-type partition of $[0, n)$ into "good" blocks (where the stopping time $\tau \le L$ realizing $\tilde{g}\, \tau \le \tau(h + \varepsilon)$ exists) and short singletons (bad/overrun) bounds, via block subadditivity theorem 3.17, $\tilde{g}\, n\, x/n \le (h\, x + \varepsilon)(1 - (L-1)/n - \mathtt{birkhoffAverage}\, \mathbf{1}_{B_L})$. Letting $n \to \infty$ (Birkhoff, theorem 3.14), then $L \to \infty$ (the bad sets B_L shrink to a null set), $M \to \infty$ and $\varepsilon \to 0$ yields $f_+ \le f\_$ a.e. The whole argument is carried in `EReal` to keep the bookkeeping clean near $-\infty$. ☐

Theorem 3.24 (Kingman core: a.e. existence of an integrable limit)**.** *For a finite measure, measure-preserving measurable T, an integrable subadditive cocycle with normalized integrals bounded below, there is an integrable G with $\mathtt{cdiv}\, g\, n\, x \to G\, x$ for μ-a.e. x.*

Proof. Take $G = f_+$, integrable by theorem 3.22. On the a.e. good set the `EReal` limsup e satisfies $\bot < e \le B < \top$ (finiteness from the loose envelope and the Fatou step), and $\liminf = \limsup = e$ theorem 3.23; in a complete linear order equal liminf and limsup force convergence to e. Transferring to $\mathbb{R}$ gives $\mathtt{cdiv}\, g\, n\, x \to e^{\mathrm{toReal}} = f_+\, x$. ☐

Theorem 3.25 (Kingman's subadditive ergodic theorem)**.** *For a finite measure, a measure-preserving T, an integrable subadditive cocycle g whose normalized integrals are bounded below, there is a T-invariant integrable G with*

$$n^{-1}\, g\, n\, x \to G\, x \qquad \textit{for } \mu\textit{-a.e. } x.$$

Proof. Take $G = f\_$. On the a.e. set where the normalized cocycle is bounded, $\liminf \le \limsup$ is trivial and $\limsup \le \liminf$ is the hard direction inside theorem 3.24, so $f\_ \overset{\text{a.e.}}{=} f_+$; the sandwich $\limsup \le f\_ \le \liminf$ yields pointwise convergence to $f\_$, and reindexing removes the $n = 0$ term. Invariance is theorem 3.21 (liminf variant), and integrability follows from $f\_ \overset{\text{a.e.}}{=} f_+$ with theorem 3.22. ☐

Corollary 3.26 (Kingman, ergodic case)**.** *For an ergodic T on a probability space, an integrable subadditive cocycle with normalized integrals bounded below, there is a constant c with $n^{-1}\, g\, n\, x \to c$ for μ-a.e. x.*

Proof. Kingman's theorem theorem 3.25 gives a T-invariant integrable limit G; ergodicity forces an a.e. T-invariant integrable function to be a.e. constant, so $G =^{\text{a.e.}} c$ and the limit is c. (That constant is the Fekete infimum γ; only a.e.-constancy is asserted, as this is what the multiplicative ergodic theorem consumes.) $\square$

Chapter 4

Lyapunov exponents and the limsup filtration

This chapter constructs, from the linear cocycle A over the measure-preserving system T, the geometric scaffolding of the Oseledets theorem: a per-vector growth rate, the finite Lyapunov spectrum it produces, and the decreasing flag of subspaces along which the cocycle grows at exactly the prescribed rates. Throughout, $A\colon X \to \mathrm{Mat}_{d\times d}(\mathbb{R})$ is invertible ($\det(Ax) \neq 0$) and acts on $v \in \mathbb{R}^d$ through the Euclidean representation $\texttt{toEuclideanCLM}(M)\colon v \mapsto M \cdot v$. Two abstract ingredients are imported: the Furstenberg–Kesten extremal exponents (the a.e. limits of $\frac{1}{n}\log\left\|A^{(n)}(x)\right\|$ and $\frac{1}{n}\log\left\|A^{(n)}(x)^{-1}\right\|$) and the subadditive ergodic theorem of Kingman (). The headline statement assembled downstream from this material is .

4.1 Ultrametric growth functions

The combinatorial heart of the construction is purely linear-algebraic: a scaling-invariant, non-Archimedean real function on the nonzero vectors of a real vector space has only finitely many values, and its sublevel sets are subspaces. We isolate this with no reference to dynamics.

Definition 4.1 (Ultrametric growth function)**.** Let E be a real vector space. A function $g\colon E \to \mathbb{R}$ is an *ultrametric growth function* when it is scaling-invariant,

$$g(c \cdot v) = g(v) \qquad (c \neq 0),$$

and non-Archimedean (strong triangle inequality),

$$g(v + w) \leq \max(g(v),\, g(w)) \qquad (v, w,\, v + w \neq 0).$$

The value $g(0)$ is never used; the side conditions $v \neq 0$ are carried explicitly to avoid extended-real arithmetic.

Lemma 4.2 (Strict ultrametric equality)**.** *If g is an ultrametric growth function and $g(v) \neq g(w)$ (with $v, w, v + w$ nonzero), then $g(v + w) = \max(g(v), g(w))$.*

Proof. By symmetry ($v+w = w+v$) assume $g(v) < g(w)$, so $\max = g(w)$. The non-Archimedean inequality gives $g(v + w) \leq g(w)$. Conversely $w = (v + w) + (-v)$ and $g(-v) = g(v)$ (scaling by -1), so $g(w) \leq \max(g(v + w), g(v))$. Were $g(v + w) < g(w)$, then both arguments of this max would be $< g(w)$, contradicting the bound. Hence $g(w) \leq g(v + w)$, giving equality. ☐

Lemma 4.3 (Sum of distinct-value vectors)**.** *Let g be an ultrametric growth function, s a nonempty finite index set, and $v\colon s \to E$ a family of nonzero vectors with $g \circ v$ injective on s. Then $\sum_{i\in s} v_i \neq 0$ and $g(\sum_{i\in s} v_i) = \sup_{i\in s} g(v_i)$.*

Proof. Strong induction on s, peeling off one element a. For the inductive step, the tail sum is nonzero with g-value $g(v_b)$ for some b in the tail; since $g(v_a) \neq g(v_b)$, 4.2 applies to $v_a + \sum_{\text{tail}} v_i$, yielding both nonvanishing and the equality of the value with the new maximum. The two conclusions are proved jointly because each step needs the tail subsum to be nonzero. ☐

Lemma 4.4 (Distinct values are independent)**.** *If g is an ultrametric growth function and $v\colon \iota \to E$ is a family of nonzero vectors with $g\circ v$ injective, then v is linearly independent over $\mathbb{R}$.*

Proof. Suppose $\sum_j c_j v_j = 0$ with some $c_i \neq 0$. Restrict to the support $t = \{j : c_j \neq 0\} \ni i$, nonempty. The scaled vectors $c_j v_j$ are nonzero, and $g(c_j v_j) = g(v_j)$ by scaling-invariance, so $g\circ(c\cdot v)$ is still injective on t. By 4.3 the support sum is nonzero, contradicting $\sum_{j\in t} c_j v_j = 0$. ☐

Lemma 4.5 (Finiteness of the value set)**.** *If E is finite-dimensional and g is an ultrametric growth function, then $\{g(v) : v \neq 0\}$ is finite, with at most $\dim_{\mathbb{R}} E$ elements.*

Proof. If the value set were infinite, pick $\dim_{\mathbb{R}} E + 1$ distinct values and witnessing nonzero vectors. By 4.4 these are linearly independent, exceeding $\dim_{\mathbb{R}} E$, a contradiction. ☐

Definition 4.6 (Sublevel submodule)**.** For an ultrametric growth function g and threshold $t \in \mathbb{R}$, the *sublevel set*

$$\mathrm{sublevel}(g,t) = \{\, v \mid v = 0 \ \vee\ g(v) \le t \,\}$$

is a submodule of E: it contains 0; closure under addition is the non-Archimedean inequality together with $\max(gv, gw) \le t$; closure under scaling is scaling-invariance. These submodules are monotone in t (`sublevel_mono`).

4.2 The upper Lyapunov growth function

We now instantiate the abstract machinery at the per-vector logarithmic growth rate of the cocycle.

Definition 4.7 (Defining sequence)**.** For $x \in X$ and $v \in \mathbb{R}^d$ the *growth sequence* is

$$\mathrm{growthSeq}(A,T,x,v)(n) = \frac{1}{n}\log\left\|A^{(n)}(x)\cdot v\right\|,$$

where $A^{(n)}(x) =$ `cocycle` $A\,T\,n\,x$ is the n-step cocycle iterate acting on v via `toEuclideanCLM`.

Definition 4.8 (Upper Lyapunov growth function)**.** The *upper Lyapunov growth function* is the lim sup of the growth sequence,

$$\bar\lambda(v) \;=\; \mathrm{lambdaBar}(A,T,x,v) \;=\; \limsup_{n\to\infty}\ \frac{1}{n}\log\left\|A^{(n)}(x)\cdot v\right\|.$$

The basic per-n sandwich is the submultiplicativity of the operator norm: $\left\|A^{(n)}(x)^{-1}\right\|^{-1}\|v\| \le \left\|A^{(n)}(x)\cdot v\right\| \le \left\|A^{(n)}(x)\right\|\|v\|$, which after taking $\frac{1}{n}\log$ pins growthSeq between two sequences differing from the Furstenberg–Kesten data by a term tending to 0. This yields both boundedness and finiteness.

Lemma 4.9 (Scaling invariance)**.** *For $c \neq 0$ and $v \neq 0$, $\bar{\lambda}(c \cdot v) = \bar{\lambda}(v)$.*

Proof. By linearity $\|A^{(n)}(x)(c \cdot v)\| = |c|\, \|A^{(n)}(x)v\|$, so the two growth sequences differ by $\frac{1}{n}\log|c|$, which $\to 0$. A lim sup is unchanged under a perturbation tending to zero (a robust helper proved directly on the defining sets $\{a : \forall^{\infty} n,\ u_n \leq a\}$, with no boundedness hypothesis). The bound is unconditional. □

Lemma 4.10 (Finiteness sandwich)**.** *Assume T ergodic on a probability space, A measurable and invertible, with $\log\|A\|$ and $\log\|A^{-1}\|$ integrable. Then there exist $\lambda_{\text{bot}} \leq \lambda_{\text{top}}$ such that for a.e. x and every $v \neq 0$, $\bar{\lambda}(v) \in [\lambda_{\text{bot}}, \lambda_{\text{top}}]$.*

Proof. Take λ_{top} and λ'_k from the Furstenberg–Kesten limits of $\frac{1}{n}\log\|A^{(n)}\|$ and $\frac{1}{n}\log\|(A^{(n)})^{-1}\|$, and set $\lambda_{\text{bot}} = -\lambda'_k$. The ordering follows from $\|A^{(n)}\|\|(A^{(n)})^{-1}\| \geq 1$. On the intersection of the two full-measure convergence sets, the upper sandwich bounds $\bar{\lambda}(v) = \limsup \text{growthSeq}$ above by λ_{top}, while the lower sandwich bounds the lim inf below by $-\lambda'_k$; since $\liminf \leq \limsup$ the value lies in the interval. □

Lemma 4.11 (Non-Archimedean inequality)**.** *For nonzero $v, w, v + w$, if the three growth sequences are bounded, then $\bar{\lambda}(v + w) \leq \max(\bar{\lambda}(v), \bar{\lambda}(w))$.*

Proof. From the triangle inequality $\|A^{(n)}(v + w)\| \leq \|A^{(n)}v\| + \|A^{(n)}w\| \leq 2\max(\|A^{(n)}v\|, \|A^{(n)}w\|)$, taking $\frac{1}{n}\log$ gives $\text{growthSeq}(v{+}w)(n) \leq \frac{1}{n}\log 2 + \max(\text{growthSeq}(v)(n), \text{growthSeq}(w)(n))$. The term $\frac{1}{n}\log 2 \to 0$, and $\limsup \max = \max \limsup$ for bounded sequences, giving the claim. □

Theorem 4.12 ($\bar{\lambda}$ is an ultrametric growth function, a.e.)**.** *Under the hypotheses of 4.10, for a.e. x the function $v \mapsto \bar{\lambda}(v)$ is an ultrametric growth function.*

Proof. Scaling-invariance is 4.9 (trivial on $v = 0$). The non-Archimedean axiom is 4.11, whose boundedness hypotheses are discharged on the full-measure Furstenberg–Kesten convergence set, where the growth sequence of every nonzero vector is bounded above and below by the two FK sandwich sequences. □

Theorem 4.13 (A-equivariance, a.e.)**.** *Under the same hypotheses, for a.e. x and every $v \neq 0$,*

$$\bar{\lambda}_x(v) = \bar{\lambda}_{Tx}(Ax \cdot v).$$

Proof. The cocycle identity $A^{(n+1)}(x) = A^{(n)}(Tx)\, A(x)$ gives $\text{growthSeq}_x(v)(n{+}1) = \frac{1}{n+1}\log\|A^{(n)}(Tx)(Ax \cdot v)\|$. Reindexing the lim sup by one, the two scalings differ by $(\frac{1}{n+1} - \frac{1}{n})\log\|\cdot\| = -\frac{1}{n+1} \cdot (\frac{1}{n}\log\|\cdot\|)$, which tends to 0 precisely because $\frac{1}{n}\log\|\cdot\|$ is bounded. The boundedness is needed at the image point Tx; it holds a.e. in x by pulling back the a.e. boundedness at a generic point through the measure-preserving T. □

4.3 The Lyapunov spectrum and the descending exponent list

Definition 4.14 (Lyapunov spectrum)**.** The *Lyapunov spectrum* at x is the finite set of realized values

$$\text{lyapunovSpectrum}(A, T, x) = \{\, \bar{\lambda}_x(v) : v \neq 0 \,\},$$

defined as a `Finset` via 4.5 on the good set where $\bar{\lambda}_x$ is an ultrametric growth function, and as $\emptyset$ off it. A value lies in it iff it is realized by some nonzero vector.

Definition 4.15 (Multiplicity count and descending list)**.** Write $k = \mathrm{specCard}(A, T, x)$ for the number of distinct exponents (the cardinality of the spectrum). The *exponent list* $\mathrm{specList}\colon \mathrm{Fin}\, k \to \mathbb{R}$ enumerates the spectrum in *strictly descending* order, $\lambda_0 > \lambda_1 > \cdots > \lambda_{k-1}$, obtained from the order embedding of the finset composed with index reversal. It is strictly antitone, every $\mathrm{specList}(i)$ lies in the spectrum, and every spectrum value is $\mathrm{specList}(i)$ for a unique i.

4.4 The limsup filtration

Definition 4.16 (Sublevel subspace)**.** The *sublevel subspace* at threshold t,

$$\mathrm{lambdaSublevel}(A, T, x, t) = \{\, v \mid v = 0 \ \vee \ \bar{\lambda}_x(v) \le t \,\},$$

is the submodule 4.6 of $\bar{\lambda}_x$ at t on the good set, and $\perp$ off it.

Definition 4.17 (Oseledets filtration / limsup flag)**.** The *limsup flag* at x is the family $\mathrm{vflag}(A, T, x)\colon \mathrm{Fin}\,(k+1) \to \mathrm{Submodule}$ with

$$\mathrm{vflag}(A, T, x)(j) = \begin{cases} \mathrm{lambdaSublevel}(A, T, x,\ \mathrm{specList}(j)) & j < k, \\ \perp & j = k. \end{cases}$$

With the descending enumeration, level j is the sublevel set at λ_j, so the flag decreases from the whole space down to $\perp$.

Lemma 4.18 (Extremal levels)**.** *On the good set,* $\mathrm{vflag}(A, T, x)(0) = \top$*; and unconditionally* $\mathrm{vflag}(A, T, x)(\mathrm{last}) = \perp$ *(`vflag_last`).*

Proof. For any $v \neq 0$, $\bar{\lambda}_x(v)$ lies in the spectrum, so $k > 0$ and $\mathrm{specList}(0)$ is the maximum of the spectrum; thus $\bar{\lambda}_x(v) \le \mathrm{specList}(0)$ and v lies in level 0. Level k is $\perp$ by definition. ☐

Theorem 4.19 (Strict decrease)**.** *On the good set,* $\mathrm{vflag}(A, T, x)(i{+}1) \subsetneq \mathrm{vflag}(A, T, x)(i)$ *for each interior index i.*

Proof. Inclusion: since specList is strictly antitone, $\mathrm{specList}(i{+}1) < \mathrm{specList}(i)$, so the sublevel at the smaller threshold is contained in that at the larger. Strictness: a witness w with $\bar{\lambda}_x(w) = \mathrm{specList}(i)$ exists (the value is realized); it lies in level i but its value exceeds $\mathrm{specList}(i{+}1)$, so it is not in level $i{+}1$. ☐

Lemma 4.20 (Stratum exactness)**.** *On the good set, if $v \in$* $\mathrm{vflag}(A, T, x)(i)$ *but $v \notin$* $\mathrm{vflag}(A, T, x)(i{+}1)$*, then* $\bar{\lambda}_x(v) = \mathrm{specList}(i) = \lambda_i$*.*

Proof. Membership in level i gives $\bar{\lambda}_x(v) \le \lambda_i$. Since $\bar{\lambda}_x(v)$ is a spectrum value, $\bar{\lambda}_x(v) = \lambda_j$ for some $j \ge i$. Non-membership in level $i{+}1$ rules out $\bar{\lambda}_x(v) \le \lambda_{i+1}$, forcing $j \le i$. Hence $j = i$ by injectivity of the strictly antitone list. ☐

Theorem 4.21 (A-equivariance of spectrum and flag, a.e.)**.** *Under the standing hypotheses, for a.e. x the spectrum is invariant,* $\mathrm{lyapunovSpectrum}(A, T, x) = \mathrm{lyapunovSpectrum}(A, T, Tx)$ *(`lyapunovSpectrum_equivariant_ae`), and the action of Ax maps each flag level (each sublevel subspace) at x onto the corresponding level at Tx:*

$$(Ax)_* \, \mathrm{lambdaSublevel}(A, T, x, t) = \mathrm{lambdaSublevel}(A, T, Tx, t).$$

Proof. The bijection $v \mapsto Ax \cdot v$ preserves $\bar{\lambda}$ by 4.13 (a.e.), hence carries witnesses at x to witnesses at Tx and conversely (using $(Ax)^{-1}$), giving the spectrum identity. The same value-preserving bijection sends $\{v = 0 \vee \bar{\lambda}_x(v) \le t\}$ onto $\{w = 0 \vee \bar{\lambda}_{Tx}(w) \le t\}$, which is the claimed image of sublevel sets. ☐

4.5 Measurability of the filtration

The Oseledets theorem requires the flag to vary measurably in x. Mathlib has no measurable structure on submodules, so we encode a subspace by its orthogonal-projection matrix.

Definition 4.22 (Projection matrix encoding)**.** For $K \leq \mathbb{R}^d$, $\mathrm{orthProjMatrix}(K)$ is the matrix of the orthogonal projection onto K, namely the preimage of K.`starProjection` under the star-algebra isomorphism `toEuclideanCLM`. A subspace is determined by this matrix, which lives in a space carrying the Borel/Pi measurable structure. Its (i, j) entry equals the i-th coordinate of the projection applied to the standard basis vector e_j (`orthProjMatrix_apply`).

Definition 4.23 (Measurable family of subspaces)**.** A subspace-valued map $V\colon X \to \mathrm{Submodule}\,\mathbb{R}\,\mathbb{R}^d$ is a *measurable family of subspaces* when $x \mapsto \mathrm{orthProjMatrix}(Vx)$ is measurable. Equivalently (`measurable_orthProjMatrix_iff`), for each standard basis index j the $\mathbb{R}^d$-valued map $x \mapsto (Vx)$.`starProjection`(e_j) is measurable.

Lemma 4.24 (Scalar growth is measurable)**.** *For fixed v, the map $x \mapsto \bar{\lambda}_x(v)$ is measurable.*

Proof. It is the $\limsup$ of the sequence $x \mapsto \frac{1}{n}\log\left\|A^{(n)}(x) \cdot v\right\|$. Each term is measurable: $x \mapsto A^{(n)}(x)$ is measurable (measurability of the cocycle), and $M \mapsto \|M \cdot v\|$ is continuous (a fixed-vector linear map of M, on a finite-dimensional space, post-composed with the norm), so the composite with log is measurable, and a $\limsup$ of measurable functions is measurable. ☐

Lemma 4.25 (Polynomial in a measurable matrix)**.** *For a fixed real polynomial q, the map $a \mapsto q(a)$ on $\mathrm{Mat}_{d\times d}(\mathbb{R})$ is measurable.*

Proof. Induction on q over the constant/sum/monomial generators, using that matrix addition and multiplication are measurable in each argument (`instMeasurableAdd Matrix` and the matrix `MeasurableMul` instance), whence $a \mapsto a^n$ is measurable (`measurable_matrix_pow`). ☐

Theorem 4.26 (CFC measurability via interpolating polynomial)**.** *Let $M\colon X \to \mathrm{Mat}_{d\times d}(\mathbb{R})$ be measurable with each Mx self-adjoint, and let $g\colon \mathbb{R} \to \mathbb{R}$. If a* fixed *polynomial q agrees with g on the spectrum of every Mx, then $x \mapsto \mathrm{cfc}\, g\,(Mx)$ is measurable.*

Proof. On the spectrum of Mx the continuous functional calculus of g coincides with that of q, and for a polynomial $\mathrm{cfc}\, q\,(Mx) = q(Mx)$. Thus pointwise $\mathrm{cfc}\, g\,(Mx) = q(Mx)$, which is measurable in x by 4.25. This uses only the bare Hermitian CFC instance, avoiding the isometric CFC (absent for real matrices) and any measurable selection. ☐

Theorem 4.27 (CFC measurability for continuous functions)**.** *Let M be measurable with each Mx self-adjoint, and $f\colon \mathbb{R} \to \mathbb{R}$ continuous. Then $x \mapsto \mathrm{cfc}\, f\,(Mx)$ is measurable.*

Proof. A single polynomial need not agree with f on the unbounded family of spectra, so approximate per point: by Weierstrass choose, for each k, a polynomial q_k with $|q_k - f| \leq 1/(k{+}1)$ on $[-k, k]$. Each spectrum is finite, hence in some $[-R, R]$, so $q_k \to f$ uniformly on $\mathrm{spectrum}(Mx)$ and $\mathrm{cfc}\, q_k\,(Mx) \to \mathrm{cfc}\, f\,(Mx)$. Each $x \mapsto \mathrm{cfc}\, q_k\,(Mx) = q_k(Mx)$ is measurable, and matrix-entrywise the metrizable limit upgrades to measurability of $x \mapsto \mathrm{cfc}\, f\,(Mx)$. ☐

These CFC tools deliver MeasurableSubspace for the concrete Oseledets flag once the Oseledets limit operator $\Lambda x = \lim_n \left((A^{(n)})^\top A^{(n)}\right)^{1/(2n)}$ is available: each flag projection is realized as a spectral band projector $P_i x = \mathrm{cfc}\, g_i\,(\Lambda x)$ of Λx for a continuous gap function g_i, so the projection matrix of its range equals $P_i x$ definitionally, and measurability of $x \mapsto P_i x$ follows from 4.27 (or, on the gapped good set where the spectrum is the fixed Lyapunov set, from 4.26). This is the measurability input feeding .

Chapter 5

The one-sided multiplicative ergodic theorem

This is the crux of the development. We work over a probability space (X, μ) with an ergodic measure-preserving transformation T and a measurable cocycle generator $A : X \to \mathrm{Mat}_{d\times d}(\mathbb{R})$ with $\det(Ax) \neq 0$, subject to the one-sided integrability $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1(\mu)$. The matrices act on $\mathbb{R}^d$ (the Euclidean space) through the operator $v \mapsto Av$, so that the relevant norm is the L^2 operator norm, which is submultiplicative; we write $A^{(n)}(x)$ for the cocycle $A(T^{n-1}x) \cdots A(x)$. The whole chapter is organized around a single positive semidefinite limiting matrix, the *Oseledets limit*

$$\Lambda(x) \;=\; \lim_{n\to\infty} \left((A^{(n)}(x))^\top A^{(n)}(x)\right)^{1/2n},$$

whose eigenspaces, once their growth rates are matched to the deterministic singular-value exponents, yield the filtration of the target theorem.

5.1 The Oseledets limit

The candidate approximants are the symmetric positive roots $q_n(x) := \left((A^{(n)}(x))^\top A^{(n)}(x)\right)^{1/2n}$, realized through the continuous functional calculus on the Gram matrix. The eigenvalues of $q_n(x)$ are the $1/n$-th roots of the singular values of $A^{(n)}(x)$; the heart of this section is that they converge a.e. to the exponentials of the deterministic Lyapunov exponents.

Theorem 5.1 (Deterministic singular-value exponents)**.** *There is an antitone sequence $\lambda^0 : \mathbb{N} \to \mathbb{R}$ (antitone on $[0, d)$) such that, for every $i < d$ and μ-a.e. x,*

$$\tfrac{1}{n} \log(\sigma_i(A^{(n)}(x))) \longrightarrow \lambda^0_i,$$

where σ_i is the i-th sorted singular value.

Proof. Package the ergodic limits $\Gamma_k = \lim \frac{1}{n} \log s_k(A^{(n)}(x))$ of the products of the top k singular values, $0 \leq k \leq d$, obtained from Kingman's subadditive ergodic theorem applied to the subadditive cocycle $\log \left\| \bigwedge^k A^{(n)} \right\|$ (the exterior-power functor turns submultiplicativity of the $\bigwedge^k$ operator norms into subadditivity). Then $\lambda^0_i := \Gamma_{i+1} - \Gamma_i$; the difference of the two a.e. Γ-limits gives the i-th singular-value exponent. Antitonicity of consecutive λ^0_i descends from the antitone ordering of the singular values inside each n, and chaining yields full antitonicity on $[0, d)$. $\square$

Lemma 5.2 (Block-value step function reproduces the spectrum)**.** *Let* $\operatorname{stepVal}\lambda^0\,D$ *be the step function* $e^{\lambda^0_{D-1}} + \sum_{k=1}^{D-1}(e^{\lambda^0_{k-1}} - e^{\lambda^0_k})\,\mathbf{1}_{(c_k,\infty)}$ *with thresholds* $c_k = e^{(\lambda^0_k+\lambda^0_{k-1})/2}$ *strictly inside the* k*-th gap. If* λ^0 *is antitone on* $[0, D)$ *and* $j < D$*, then* $\operatorname{stepVal}\lambda^0\,D\,(e^{\lambda^0_j}) = e^{\lambda^0_j}$*.*

Proof. At the argument $e^{\lambda^0_j}$ the threshold indicator $\mathbf{1}_{(c_k,\infty)}$ is 1 exactly when $k > j$ (since λ^0 is antitone and c_k lies strictly between λ^0_k and λ^0_{k-1}), so only the increments above index j survive. Those increments $e^{\lambda^0_{k-1}} - e^{\lambda^0_k}$ telescope to $e^{\lambda^0_j} - e^{\lambda^0_{D-1}}$, which added to the constant base $e^{\lambda^0_{D-1}}$ returns $e^{\lambda^0_j}$. □

Lemma 5.3 (Spectral deviation bound)**.** *For a self-adjoint matrix* M *and any function* g*,*

$$\|M - g(M)\| \;\le\; \sum_j \left|\mu_j - g(\mu_j)\right|,$$

where μ_j *ranges over the sorted eigenvalues of* M*.*

Proof. Writing $M = \operatorname{id}(M)$ gives $M - g(M) = (\operatorname{id} - g)(M)$ by linearity of the continuous functional calculus, and the operator norm of a self-adjoint matrix functional calculus is the largest absolute eigenvalue $\max_j \left|\mu_j - g(\mu_j)\right|$, which is bounded by the full nonnegative sum. □

Lemma 5.4 (Per-term band-projector convergence)**.** *For* μ*-a.e.* x *and every threshold index* $k \in [1, d)$*, the block term* $(e^{\lambda^0_{k-1}} - e^{\lambda^0_k}) \cdot P_n^{c_k}(x)$ *converges, where* P_n^c *is the band projector* $\mathbf{1}_{(c,\infty)}(q_n(x))$*.*

Proof. At a genuine gap $\lambda^0_k < \lambda^0_{k-1}$ the threshold c_k is strictly separated from the two limiting eigenvalue clusters, so once the sorted eigenvalues of $q_n(x)$ have converged the count of eigenvalues above c_k stabilizes and the corresponding spectral projector is Cauchy in operator norm. At a non-gap ($\lambda^0_{k-1} = \lambda^0_k$) the coefficient $e^{\lambda^0_{k-1}} - e^{\lambda^0_k}$ vanishes, so the term is constantly 0. □

Theorem 5.5 (Existence of the Oseledets limit)**.** *For* μ*-a.e.* x *the approximants* $q_n(x)$ *converge in the matrix metric to a single matrix* $\Lambda(x)$*.*

Proof. The eigenvalues $\mu_{j,n} = \sigma_j^{1/n}$ of $q_n(x)$ converge a.e. to the exponentials $e^{\lambda^0_j}$ (Theorem 5.1). Form the block approximant $\Lambda_n(x) := \operatorname{stepVal}\lambda^0\,d\,(q_n(x))$, a finite linear combination of band projectors. By Lemma 5.3, $\|q_n(x) - \Lambda_n(x)\| \le \sum_j \left|\mu_{j,n} - \operatorname{stepVal}(\mu_{j,n})\right|$; each summand is eventually $\left|\mu_{j,n} - e^{\lambda^0_j}\right| \to 0$ because the step function reproduces the exponentials on the spectrum (Lemma 5.2). Meanwhile $\Lambda_n(x)$ converges as a finite sum of convergent band-projector terms (Lemma 5.4). Adding the two convergences gives $q_n(x) \to \Lambda(x)$, and the limit is selected pointwise. □

Definition 5.6 (The named Oseledets limit)**.** $\Lambda(x) := (\,\lim_n (q_n(x))_{ij}\,)_{ij}$ is the entrywise real $\limsup$/limit of the matrix entries of $q_n(x)$; it is a total, measurable function of x.

Theorem 5.7 (The limit is the a.e. limit of the approximants)**.** *For* μ*-a.e.* x*,* $q_n(x) \to \Lambda(x)$ *in the matrix metric, and* Λ *is measurable.*

Proof. On the a.e. full convergence set of Theorem 5.5 the entrywise limit recovers the matrix limit (matrix convergence in finite dimensions is entrywise), so the entrywise $\liminf$ defining Λ equals the genuine limit. Measurability is entrywise: each entry is a limit of measurable functions of x, and a $\liminf$ of measurable $\mathbb{R}$-valued functions is measurable. □

Proposition 5.8 (Structure of the limit)**.** *For μ-a.e. x, $\Lambda(x)$ is self-adjoint and positive semidefinite.*

Proof. Self-adjointness $M^\top = M$ is an entrywise closed condition preserved under the matrix limit of the self-adjoint approximants $q_n(x)$. For positive semidefiniteness, the quadratic form $M \mapsto v^\top M v$ is continuous, so $v^\top \Lambda(x) v = \lim_n v^\top q_n(x) v \geq 0$ as a limit of nonnegatives. ☐

5.2 The per-vector lower bound

The lower half of the exact growth law isolates one band of the spectrum of $q_n(x)$ and shows that a vector with nonzero projection onto the band grows at least at the band rate.

Lemma 5.9 (Gram quadratic-form band bound)**.** *For self-adjoint Q, a band indicator $\chi = \mathbf{1}_{(c,\infty)}$, and a continuous $f \geq 0$ on $\operatorname{spec}(Q)$ with $a \leq f(t)$ whenever $c < t$,*

$$a\,\|\chi(Q)\,v\|^2 \;\leq\; \langle f(Q)\,v,\,v\rangle.$$

Proof. The band projector $\chi(Q)$ is a self-adjoint idempotent, so $\|\chi(Q)v\|^2 = \langle \chi(Q)v, v\rangle$. The gap operator $(f - a\chi)(Q)$ is positive semidefinite because $f - a\chi \geq 0$ on the spectrum (above c, $f \geq a = a\chi$; below, $\chi = 0$ and $f \geq 0$). Expanding $\langle (f - a\chi)(Q)v, v\rangle \geq 0$ gives the claim. ☐

Lemma 5.10 (Band lower bound for the cocycle)**.** *For $c \geq 0$ and $n \geq 1$,*

$$c^{2n}\,\|P_n^c(x)\,v\|^2 \;\leq\; \left\|A^{(n)}(x)\,v\right\|^2.$$

Proof. Raising $q_n(x) = (\mathrm{gram}_n)^{1/2n}$ to the $2n$-th power via the functional calculus recovers the Gram matrix $\mathrm{gram}_n = (A^{(n)})^\top A^{(n)}$ (the composed powers compose to the identity on the nonnegative spectrum). Apply Lemma 5.9 with $f(t) = t^{2n}$ and $a = c^{2n}$: above c one has $t^{2n} \geq c^{2n}$. The right-hand inner product is then $\langle \mathrm{gram}_n v, v\rangle = \left\|A^{(n)}(x)v\right\|^2$. ☐

Lemma 5.11 (The band correction vanishes)**.** *If $P_n^c(x) \to P$ with $Pv \neq 0$, then $\frac{1}{n}\log\|P_n^c(x)\,v\| \to 0$.*

Proof. The evaluation $M \mapsto Mv$ is continuous in finite dimensions, so $P_n^c(x)v \to Pv \neq 0$ and $\|P_n^c(x)v\| \to \|Pv\| > 0$. Hence the log converges to the finite number $\log\|Pv\|$, and dividing by $n \to \infty$ sends it to 0. ☐

Proposition 5.12 (Per-vector liminf lower bound)**.** *If $P_n^c(x) \to P$ with $c > 0$ and $Pv \neq 0$, and the cocycle growth sequence is cobounded, then*

$$\log c \;\leq\; \liminf_n \tfrac{1}{n}\log\left\|A^{(n)}(x)\,v\right\|.$$

Proof. Taking logs in Lemma 5.10 and dividing by $2n$ gives, eventually,

$$\log c + \tfrac{1}{n}\log\|P_n^c(x)v\| \;\leq\; \tfrac{1}{n}\log\left\|A^{(n)}(x)v\right\|.$$

The left side converges to $\log c$ since the band-correction term vanishes (Lemma 5.11). Passing to the $\liminf$ along the inequality, using its boundedness and the coboundedness of the right side (supplied by the Furstenberg–Kesten integrability of the top exponent), yields the bound. ☐

Lemma 5.13 (Band-projector nesting, kernel propagation)**.** *For thresholds $c \leq c'$ with limit band projectors P, P', if $Pv = 0$ then $P'v = 0$.*

Proof. The finite-n bands are nested: $\mathbf{1}_{(c,\infty)} \cdot \mathbf{1}_{(c',\infty)} = \mathbf{1}_{(c',\infty)}$ on the spectrum since $(c', \infty) \subseteq (c, \infty)$, so $P_n^c P_n^{c'} = P_n^{c'}$; passing to the limit gives $PP' = P'$. Both limit projectors are symmetric (limits of self-adjoint matrices), so transposing gives $P'P = P'$, whence $P'v = P'(Pv) = 0$. ☐

5.3 The spectral upper bound and the determinant squeeze

The upper half is the genuinely non-elementary step. For a vector v in the slow subspace of Λ at level e^t, one bounds $\|A^{(n)}(x)v\|$ by the restricted operator norm of $A^{(n)}$ on the slow subspace, whose growth exponent is pinned — non-circularly — by a determinant/volume squeeze.

Lemma 5.14 (Tempering)**.** *If T is measure-preserving and $g \geq 0$ with $g \in L^1(\mu)$, then for μ-a.e. x, $\frac{1}{n}\,g(T^n x) \to 0$.*

Proof. The series $\sum_n g(T^n x)/n^2$ is a.e. finite by integrability and invariance of μ, so its terms $g(T^n x)/n^2 \to 0$; the Borel–Cantelli/sublinear-growth argument then gives $g(T^n x) = o(n)$, i.e. $\frac{1}{n} g(T^n x) \to 0$. ☐

Lemma 5.15 (Slow-volume exponent squeeze)**.** *Given a volume cocycle whose top, slow and remaining log-exponent sequences satisfy a sum law* vol $=$ slow $+$ rest *with the appropriate limits, the slow restricted-operator-norm exponent obeys* $\limsup_n \frac{1}{n}\log(\mathrm{slow}_n) \leq \lambda_i$.

Proof. The total volume exponent is the Furstenberg–Kesten determinant limit $\sum_j \lambda_j^0$; the fast-block volume is the exterior-power Kingman limit; their difference forces the slow-block volume exponent. The squeeze converts a sum identity of limits into an upper bound for the slow factor once the fast and remaining factors are pinned, the angle/tilt between fast and slow blocks tempering to zero so no cross term inflates the slow volume. ☐

Theorem 5.16 (Spectral upper bound via the determinant squeeze)**.** *Let $v \neq 0$ lie in the slow subspace. Given the slow restricted-norm exponent bound $\limsup_n \frac{1}{n}\log r_n \leq \lambda_i$ and the restriction estimate $\|A^{(n)}(x)v\| \leq r_n \|v\|$ (valid for v slow), one has*

$$\limsup_n \ \tfrac{1}{n}\log\left\|A^{(n)}(x)\,v\right\| \ \leq\ \lambda_i.$$

Proof. Take logs in $\|A^{(n)}(x)v\| \leq r_n \|v\|$, divide by n, and pass to the lim sup: the $\|v\|$ factor contributes 0 and r_n contributes $\limsup_n \frac{1}{n}\log r_n \leq \lambda_i$ (Lemma 5.15). The restricted-norm exponent r_n depends only on the *global* volume cocycle and the tempered slow–fast angle, never on the growth of v itself, so the argument is non-circular. ☐

Theorem 5.17 (Upper bound on the slow space)**.** *On the ultrametric-growth good set, every vector v of the Λ-slow band* vslow(e^t) *with $\limsup_n \frac{1}{n}\log\|A^{(n)}v\| \leq t$ lies in the growth sublevel $\{v : \overline{\lambda}(x,v) \leq t\}$.*

Proof. The lim sup of $\frac{1}{n}\log\|A^{(n)}v\|$ is by definition the upper growth function $\overline{\lambda}(x,v)$; the hypothesis says it is $\leq t$, which is exactly membership in the sublevel lambdaSublevel t. The zero vector lies in every submodule. ☐

5.4 Spectral identification of the filtration

The two bounds match the spectral filtration of Λ with the analytic limsup filtration. The bridge is that the finite-n band projectors converge a.e. to the functional calculus indicator of Λ.

Theorem 5.18 (Band projectors converge to the CFC indicator)**.** *For μ-a.e. x, every $c > 0$ that is not one of the limiting eigenvalues $e^{\lambda_{\mathrm{sing}}(x,i)}$ of $\Lambda(x)$ satisfies*

$$P_n^c(x) \longrightarrow \mathbf{1}_{(c,\infty)}(\Lambda(x)).$$

Proof. Since c avoids the spectrum of $\Lambda(x)$, there is a gap $\delta > 0$ between c and every eigenvalue. Replace the discontinuous indicator by a continuous $\delta/2$-clamp surrogate χ that is Lipschitz and agrees with $\mathbf{1}_{(c,\infty)}$ at distance $\geq \delta/2$ from c. Once the sorted eigenvalues of $q_n(x)$ are within $\delta/2$ of their limits, χ and $\mathbf{1}_{(c,\infty)}$ agree on both spectra, so $P_n^c(x) = \chi(q_n(x))$ eventually; by Lipschitz continuity of the functional calculus and $q_n(x) \to \Lambda(x)$ (Theorem 5.7), $\chi(q_n(x)) \to \chi(\Lambda(x)) = \mathbf{1}_{(c,\infty)}(\Lambda(x))$. ☐

Theorem 5.19 (Reverse slow-flag inclusion)**.** *For μ-a.e. x and every t,* $\mathrm{lambdaSublevel}(x,t) \leq \mathrm{vslow}(x, e^t)$*.*

Proof. Contrapositively, a vector $v \notin \mathrm{vslow}(e^t)$ has nonzero component in the band of $\Lambda(x)$ above e^t; by Theorem 5.18 the finite band projectors converge to the corresponding CFC indicator with $Pv \neq 0$, so the per-vector lower bound (Proposition 5.12) forces $\liminf \frac{1}{n} \log \|A^{(n)} v\| > t$, hence $\overline{\lambda}(x,v) > t$ and $v \notin \mathrm{lambdaSublevel}\, t$. Kernel propagation across nested thresholds (Lemma 5.13) makes this consistent for all t. ☐

Theorem 5.20 (The slow flag equals the limsup sublevel)**.** *Under the spectral upper bound and the reverse inclusion, for μ-a.e. x and every t,*

$$\mathrm{vslow}(x, e^t) \;=\; \mathrm{lambdaSublevel}(x,t).$$

Proof. Two inclusions: the forward one (Theorem 5.17, from the upper bound) shows slow vectors grow slowly, hence lie in the sublevel; the reverse one (Theorem 5.19, from the lower bound) shows slowly-growing vectors lie in the slow space. Antisymmetry gives the identity, simultaneously for all t on the a.e. ultrametric-growth good set. ☐

5.5 Ruelle's reverse cofactor bound and the top-gap envelope

The upper bound on an individual vector reduces, after diagonalizing the limit, to controlling the overlap matrix between the sorted Gram eigenbasis at level n and the limiting eigenbasis. The graded overlap is controlled by a leakage induction; Ruelle's cofactor estimate then converts a one-sided forward decay into the full pairwise rate.

Lemma 5.21 (Ruelle reverse cofactor bound)**.** *Let S be orthogonal ($SS^\top = 1$) with the graded forward decay $|S_{ab}| \leq c \cdot e^{-\max(g_b - g_a, 0)}$. Then every entry obeys the reverse bound at the full pairwise rate:*

$$|S_{ij}| \;\leq\; (d-1)!\, c^{d-1}\, e^{-(g_i - g_j)}.$$

Proof. Since $S^{-1} = S^\top$, the entry S_{ij} is $(\det S)^{-1}$ times the cofactor $\mathrm{adj}(S)_{ji}$, and $|\det S| = 1$ because $SS^\top = 1$. Expanding the minor by the Leibniz formula, every surviving permutation term collects the level imbalance $g_i - g_j$ by telescoping the forward factors, each forward factor contributing at most $c\, e^{-\max(\cdot,0)}$; there are at most $(d-1)!$ such terms. ☐

Definition 5.22 (Top-gap fast-band-mass envelope)**.** $\mathrm{TopGapMassEnvelope}\, A\, T\, \lambda^0\, x$ asserts the uniform geometric leakage of fast-band mass across each genuine gap: for every cut tolerance δ there is a constant C controlling, eventually and uniformly, the band mass that crosses the top gap of each stratum.

Lemma 5.23 (Multi-source geometric envelope). *If a nonnegative sequence a obeys a one-step recursion fed by finitely many source sequences each decaying geometrically with ratio $\rho < 1$, then a_n is bounded by a fixed multiple of the summed source envelopes for all n.*

Proof. Each single source contributes a geometric partial sum bounded by $K/(1-\rho)$; summing the finitely many per-source envelopes and folding them through the linear one-step recursion gives a uniform bound on the chained quantity a_n. ☐

Lemma 5.24 (Per-stratum envelope step). *At a fixed cut strictly inside a gap of width $\geq G$, the one-step band-mass increment is the current mass damped by e^{-G} plus a tempered source term; iterating produces the per-stratum leakage envelope.*

Proof. In the sorted-Gram-eigenbasis block decomposition, the band mass above the cut at step $n+1$ is the mass at step n attenuated by the singular-value ratio across the gap (bounded by e^{-nG}-type damping after n iterates) plus the contribution injected by the one-step generator $A(T^n x)$, whose log-norm is tempered to $o(n)$. This is exactly the geometric one-step recursion fed by tempered sources of Lemma 5.23. ☐

Theorem 5.25 (The top-gap envelope, a.e.). *For μ-a.e. x, the top-gap fast-band-mass envelope* $\mathrm{TopGapMassEnvelope}\, A\, T\, \lambda^0\, x$ *holds.*

Proof. Fix the deterministic distinct gap $G > 0$ separating distinct exponents. On the a.e. set where every singular-value exponent converges (Theorem 5.1) and the one-step generator log-norm is tempered (Lemma 5.14), build the per-stratum leakage envelope for each gap pair (Lemma 5.24) and assemble them into the top-gap envelope: each stratum's fast-band mass crossing its top gap is geometrically damped, uniformly in the cut, by a single per-pair constant maximized over the finitely many pairs. ☐

5.6 Constancy of the spectrum

The deterministic exponent set is constant in x by construction, so once the per-point spectrum is identified with it, ergodic constancy is automatic.

Lemma 5.26 (Spectrum identity from two inclusions). *If at x every realized exponent is a deterministic one and every deterministic exponent is attained, then* $\mathrm{lyapunovSpectrum}(x) = \mathrm{distinctExp}\, \lambda^0\, d$.

Proof. Both directions are finite-set inclusions; antisymmetry of $\subseteq$ gives the equality of the two finite subsets of $\mathbb{R}$. ☐

Theorem 5.27 (Ergodic constancy of the spectrum). *Given a.e. that every realized value of the upper growth function is a deterministic exponent (upper inclusion) and every deterministic exponent is attained (lower inclusion), for μ-a.e. x,*

$$\mathrm{lyapunovSpectrum}(x) \;=\; \mathrm{distinctExp}\, \lambda^0\, d.$$

Proof. On the ultrametric-growth good set the two Finset inclusions are equivalent to native statements about the upper growth function $\overline{\lambda}(x, \cdot)$: the upper inclusion is the spectral-upper-bound output (each stratum value is a deterministic exponent), the lower inclusion is attainment from the lower bound. Lemma 5.26 then gives the identity with the deterministic constant set; since that set does not depend on x, the identification is T-invariant and the spectrum is a.e. constant — ergodicity adds nothing beyond the deterministic value already in hand. ☐

5.7 Assembling the target theorem

The forward graded-overlap bound (built from the top-gap envelope and Ruelle's reverse bound) gives the spectral upper bound hupper; combined with the reverse inclusion it yields the slow flag, and the lower bound supplies attainment. Together they discharge the three a.e. interfaces (spectrum, slow flag, exact growth) of the filtration assembly.

Theorem 5.28 (Filtration from the spectral upper bound)**.** *Assume the per-vector spectral upper bound on the slow flag, the reverse slow-flag inclusion, the two spectrum inclusions, and the band-projector convergence datum. Then there exist k, strictly decreasing $\lambda : \operatorname{Fin} k \to \mathbb{R}$, and a measurable family V forming a.e. a strictly decreasing A-equivariant flag along which $\frac{1}{n} \log \|A^{(n)}(x)v\| \to \lambda_i$ on each stratum.*

Proof. The deterministic exponents λ^0 come from Theorem 5.1. The spectrum interface is discharged by the two inclusions (constancy, Theorem 5.27); the slow-flag interface by vslow = lambdaSublevel (Theorem 5.20); the exact-growth interface by combining the unconditional upper half (the stratum value $\overline{\lambda} = \lambda_i$ on vflag), the lower half (Proposition 5.12 fed the band datum), and Furstenberg–Kesten boundedness into a two-sided limit. These three interfaces feed the generic slow-flag assembly. ☐

Theorem 5.29 (Filtration from the top-gap envelope)**.** *Under the standing ergodic, invertible, log-integrable hypotheses, and assuming the top-gap envelope* TopGapMassEnvelope *quantified over λ^0, the full Oseledets filtration conclusion holds.*

Proof. Diagonalize $\Lambda(x)$ by its limit eigenbasis with eigenvalues $e^{\lambda_{\text{sing}}}$ and slow-orthogonality. The forward graded-overlap bound, consuming the envelope (Definition 5.22), yields the one-sided forward decay of the overlap matrix between the sorted Gram eigenbasis and the limit eigenbasis; Ruelle's reverse cofactor estimate (Lemma 5.21) upgrades this to the full pairwise rate, which is exactly the slow-restriction bound feeding the spectral upper bound hupper on the limit slow space. The band-projector convergence (Theorem 5.18) supplies the reverse slow-flag inclusion and the lower-bound datum, and the spectrum inclusions follow from the slow-flag identity. These discharge all hypotheses of Theorem 5.28. ☐

Theorem 5.30 (One-sided Oseledets multiplicative ergodic theorem)**.** *Let μ be a probability measure, $T : X \to X$ ergodic measure-preserving, and $A : X \to \operatorname{Mat}_{d\times d}(\mathbb{R})$ measurable with* $\det(Ax) \neq 0$ *and* $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1(\mu)$. *Then there are k distinct Lyapunov exponents* $\lambda : \operatorname{Fin} k \to \mathbb{R}$*, strictly decreasing, and a measurable family of subspaces*

$$V : \operatorname{Fin}(k+1) \to X \to \operatorname{Submodule}_{\mathbb{R}}(\mathbb{R}^d)$$

with each $x \mapsto V_i\, x$ measurable, such that for μ-a.e. x: $V_0\, x = \top$, $V_k\, x = \bot$; the flag is strictly decreasing, $V_{i+1}\, x < V_i\, x$; it is A-equivariant, $A(x)\, V_i\, x = V_i\,(Tx)$; and along it the cocycle grows at the exact rate λ_i:

$$\tfrac{1}{n} \log \left\|A^{(n)}(x)\, v\right\| \;\longrightarrow\; \lambda_i \qquad \textit{for every } v \in V_i\, x \setminus V_{i+1}\, x.$$

Proof. If $d = 0$ the trivial flag $\top = \bot$ with no exponents discharges the statement. For $d > 0$, the top-gap envelope holds a.e. (Theorem 5.25), so the conditional assembly Theorem 5.29 applies directly and produces the exponents, the measurable equivariant flag, and the exact per-stratum growth limits. ☐

Chapter 6

Companion results and extensions

The Oseledets multiplicative ergodic theorem `oseledets_filtration` delivers, μ-a.e., a strictly decreasing A-equivariant flag $\mathbb{R}^d = V_0(x) \supsetneq \cdots \supsetneq V_k(x) = 0$ together with a strictly decreasing exponent list $\lambda_0 > \cdots > \lambda_{k-1}$ governing the growth $\frac{1}{n}\log\left\|A^{(n)}(x)v\right\| \to \lambda_i$ on each stratum $V_i \setminus V_{i+1}$. This chapter collects its companion results. Many are surprisingly cheap: they follow from the *statement* of the main theorem, quantified over an arbitrary witness of its conclusion, with no access to the construction. To make this precise we bundle the conclusion as a predicate.

6.1 The bundled predicate and uniqueness

Definition 6.1 (Bundled Oseledets filtration)**.** For a measure μ, map T, generator A, count k, exponent list $\lambda : \mathrm{Fin}\, k \to \mathbb{R}$ and flag $V : \mathrm{Fin}(k+1) \to X \to \mathrm{Submodule}$, the predicate $\mathrm{IsOseledetsFiltration}\ \mu\ T\ A\ k\ \lambda\ V$ asserts: λ is strictly antitone; each level V_i is a measurable subspace family; and μ-a.e. x carries the strictly decreasing A-equivariant flag $\mathbb{R}^d = V_0(x) \supsetneq \cdots \supsetneq V_k(x) = 0$ with exact growth rate λ_i on the stratum $V_i \setminus V_{i+1}$. This is byte-identical to the conclusion of the main theorem.

Theorem 6.2 (Repackaged existence)**.** *Under the standing hypotheses (ergodic T, invertible measurable A with $\log^+\|A\|, \log^+\left\|A^{-1}\right\| \in L^1$), there exist k, λ, V with* $\mathrm{IsOseledetsFiltration}\ \mu\ T\ A\ k\ \lambda\ V$*.*

Proof. Destructure the conclusion of `oseledets_filtration` and repackage its three conjuncts as the bundled predicate; the data match definitionally. ☐

Theorem 6.3 (Canonical sublevel characterization)**.** *If* $\mathrm{IsOseledetsFiltration}\ \mu\ T\ A\ k\ \lambda\ V$ *holds, then μ-a.e. each interior flag level is exactly a growth-sublevel set: for every i and vector v,*

$$v \in V_i(x) \iff v = 0 \ \vee\ \limsup_n \tfrac{1}{n}\log\left\|A^{(n)}(x)v\right\| \le \lambda_i.$$

Proof. At a good point pick the stratum index j of $v \neq 0$; the per-stratum convergence gives $\limsup = \lambda_j$. Membership in V_i forces $i \le j$, so by antitonicity $\lambda_j \le \lambda_i$; conversely if $v \notin V_i$ then $j < i$ and $\lambda_j > \lambda_i$ strictly, contradicting the bound. The disjunct $v = 0$ handles $\log 0$. No machinery beyond the a.e. block is used. ☐

Theorem 6.4 (Uniqueness of the spectrum and filtration)**.** *On a probability space, any two Oseledets filtration data (k, λ, V) and (k_2, λ_2, V_2) for the same cocycle agree: $k = k_2$, the exponents coincide under the index cast, and μ-a.e. the flags agree level by level.*

Proof. At a single good point the set of realized growth limits equals range λ and range λ_2; two strictly antitone enumerations of one finite set coincide, giving $k = k_2$ and $\lambda = \lambda_2$. Levelwise identity then follows from the sublevel characterization 6.3 applied to both data. ☐

6.2 The top exponent as operator-norm growth

Lemma 6.5 (Nontriviality)**.** *On a probability space with $0 < d$, any Oseledets filtration has $0 < k$.*

Proof. If $k = 0$ then at a good point $\mathbb{R}^d = V_0(x) = V_{\text{last}}(x) = 0$, forcing finrank $= 0$, contradicting $d > 0$. ☐

Theorem 6.6 (Top exponent = norm growth)**.** *On a probability space, with A invertible and $0 < k$, μ-a.e. the operator-norm growth rate of the cocycle converges to the top exponent:*

$$\tfrac{1}{n}\log\|A^{(n)}(x)\| \to \lambda_0.$$

Proof. Two-sided squeeze from the flag block. *Lower:* a vector v in the top stratum has $\frac{1}{n}\log\|A^{(n)}v\| \to \lambda_0$ and $\|A^{(n)}v\| \le \|A^{(n)}\|\,\|v\|$. *Upper:* the column-sum bound $\|M\| \le \sum_j \|Me_j\|$ on the L^2 operator norm, each basis vector being nonzero, gives eventually $\frac{1}{n}\log\|A^{(n)}\| \le \lambda_0 + \varepsilon$. Neither Furstenberg–Kesten nor singular values are needed. ☐

Corollary 6.7 (Identification of the Furstenberg–Kesten constant)**.** *Any constant c to which $\frac{1}{n}\log\|A^{(n)}(x)\|$ converges μ-a.e. (e.g. the Furstenberg–Kesten constant) equals λ_0.*

Proof. Both 6.6 and the hypothesis hold at a common good point; uniqueness of limits gives $\lambda_0 = c$. ☐

6.3 A.e.-constant multiplicities

Theorem 6.8 (Deterministic dimension profile)**.** *For ergodic T and invertible A, every Oseledets filtration has a deterministic dimension profile: there is a strictly decreasing $m : \mathrm{Fin}(k+1) \to \mathbb{N}$ with $m_0 = d$, $m_k = 0$ and, μ-a.e., $\operatorname{finrank} V_i(x) = m_i$.*

Proof. The dimension $x \mapsto \operatorname{finrank} V_i(x)$ is measurable via the trace of the orthogonal projector, and T-invariant a.e. because equivariance through the invertible $A(x)$ preserves dimension. Ergodicity makes each invariant $\mathbb{N}$-valued function a.e. constant. The profile structure (StrictAnti, endpoints) is read off one good point. ☐

Corollary 6.9 (Per-exponent multiplicities)**.** *For ergodic T and invertible A, each exponent λ_i carries a positive deterministic multiplicity $m_i = \dim V_i - \dim V_{i+1}$ with $\sum_i m_i = d$.*

Proof. Set m_i to the consecutive dimension drops of 6.8; positivity is strict antitonicity, and the telescoping sum equals $m_0 - m_k = d$. ☐

Theorem 6.10 (MET with multiplicities)**.** *Under the standing hypotheses there exist k, λ, V and a strictly decreasing m with $m_0 = d$, $m_k = 0$, such that* $\mathrm{IsOseledetsFiltration}\ \mu\ T\ A\ k\ \lambda\ V$ *holds and μ-a.e.* $\operatorname{finrank} V_i(x) = m_i$.

Proof. Obtain a witness from 6.2 and apply 6.8 to it. ☐

6.4 The Lyapunov spectrum

Definition 6.11 (Sorted spectrum)**.** The full Lyapunov spectrum with multiplicity is the total function exponents : $\mathrm{Fin}\, d \to \mathbb{R}$, whose i-th entry is the deterministic limit of $\frac{1}{n} \log \sigma_i(A^{(n)})$, sorted non-increasingly. The top entry is topExponent.

Theorem 6.12 (Defining σ-limit and order)**.** *For each sorted index i and μ-a.e. x, $\frac{1}{n} \log \sigma_i(A^{(n)}(x)) \to$ $\mathrm{exponents}_i$; moreover* exponents *is antitone.*

Proof. The deterministic exponent sequence is extracted by Classical.choose from the singular-value convergence theorem underlying the MET; antitonicity and the a.e. limit are its defining specification. ☐

Theorem 6.13 (Eigenvalue tie)**.** *μ-a.e., $\exp(\mathrm{exponents}_i)$ is the i-th sorted eigenvalue of the Oseledets limit matrix $\Lambda(x)$.*

Proof. The σ-limit identifies $\mathrm{lamSing}(x, i) = \mathrm{exponents}_i$ a.e.; combine with the eigenvalues of Λ being e^{lamSing}. ☐

6.5 Exponent sums

Theorem 6.14 (Sign characterizations of exponent sums)**.** *The sum* sumPosExp *of the strictly positive exponents is nonnegative, vanishes iff all exponents are ≤ 0, and is strictly positive iff some exponent is positive (and dually* $\mathrm{sumNegExp} \leq 0$ *with the mirror characterizations).*

Proof. Each summand of sumPosExp is strictly positive on the filter, so the sum is ≥ 0; a sum of nonnegatives vanishes iff the filter is empty, i.e. no exponent is positive. ☐

Theorem 6.15 (Telescoping identity for partial sums)**.** *For $k \leq d$, the ergodic growth rate Γ_k of the product of the top-k singular values equals the sum of the top-k exponents: $\Gamma_k = \sum_{i<k} \mathrm{exponents}_i$.*

Proof. Since $\mathrm{sprod}_k = \prod_{i<k} \sigma_i$, the normalized $\log \mathrm{sprod}_k$ is the finite sum of the per-index $\frac{1}{n} \log \sigma_i$, each converging to $\mathrm{exponents}_i$. The sum of convergents converges to the sum of limits; uniqueness against the defining limit of Γ_k closes it. ☐

6.6 Exterior (wedge) growth

Definition 6.16 (Exterior cocycle generator)**.** The k-th exterior generator $\mathrm{extGen}\, k\, A$ sends x to the k-th compound matrix $C_k(Ax)$ of $k \times k$ minors. Its iterated cocycle is the compound of the iterate: $\mathrm{cocycle}(\mathrm{extGen}\, k\, A)\, T\, n\, x = C_k(A^{(n)}(x))$.

Theorem 6.17 (k-volume growth rate)**.** *For $k \leq d$ and μ-a.e. x, the operator-norm growth of the compound cocycle converges to Γ_k: $\frac{1}{n} \log \lVert C_k(A^{(n)}(x)) \rVert \to \Gamma_k$, the k-dimensional volume growth rate.*

Proof. The operator norm of the compound matrix is sprod_k, the product of the top-k singular values; rewrite and apply the defining a.e. limit of Γ_k. ☐

Corollary 6.18 (Positive sum as a maximal partial sum)**.** *Writing $k_+ = \#\{i : 0 < \mathrm{exponents}_i\}$, the positive-exponent sum equals the partial sum $\Gamma_{k_+} = \sum_{i<k_+} \mathrm{exponents}_i$.*

Proof. By antitonicity the strictly positive entries are exactly the top k_+ indices, so the filtered positive sum coincides with the top-k_+ prefix sum, which is Γ_{k_+} via 6.15. ☐

6.7 The trace–determinant identity

Lemma 6.19 (Product of singular values is the absolute determinant)**.** *For every* n*,* x*:* $\operatorname{sprod} A\,T\,d\,n\,x = \left|\det A^{(n)}(x)\right|$.

Proof. Squaring, $\operatorname{sprod}_d^2 = \prod_i \sigma_i^2 = \det(M^\top M) = (\det M)^2$ for the symmetric Gram operator; take the nonnegative square root. No invertibility is needed. ☐

Theorem 6.20 (Determinant identity)**.** *The sum of all Lyapunov exponents equals the integral of* $\log|\det|$ *of the generator:*

$$\sum_i \operatorname{exponents}_i = \int_X \log|\det A(x)|\, d\mu.$$

Proof. Two a.e. limits of $\frac{1}{n}\log\left|\det A^{(n)}\right|$: it equals $\frac{1}{n}\log \operatorname{sprod}_d \to \Gamma_d = \sum_i \operatorname{exponents}_i$ by 6.19 and 6.15; and $\log\left|\det A^{(n)}\right|$ is the additive Birkhoff sum of $\log|\det A|$, whose ergodic average tends to $\int \log|\det A|$. Uniqueness of limits closes the identity. ☐

Corollary 6.21 (Volume contraction)**.** *If* $\sum_i \operatorname{exponents}_i < 0$ *then* μ*-a.e.* $\left|\det A^{(n)}(x)\right| \to 0$*.*

Proof. Since $\frac{1}{n}\log\left|\det A^{(n)}\right|$ tends to a negative constant, $\log\left|\det A^{(n)}\right| \to -\infty$, so its exponential tends to 0. ☐

6.8 The inverse / time-reversed spectrum

Theorem 6.22 (Inverse cocycle exponents)**.** *For each sorted index* i *and* μ*-a.e.* x*, the singular-value exponents of the inverse-matrix cocycle are the negated, reversed exponents of* A*:*

$$\tfrac{1}{n}\log\sigma_i((A^{(n)}(x))^{-1}) \to -\operatorname{exponents}_{\operatorname{rev} i}.$$

Proof. Singular-value reciprocity $\sigma_i(M^{-1}) = \sigma_{\operatorname{rev} i}(M)^{-1}$ for invertible M, applied to the iterate, turns the forward limit $\operatorname{exponents}_{\operatorname{rev} i}$ into its negative. ☐

Corollary 6.23 (Top of the reversed spectrum is minus the bottom)**.** μ*-a.e. the largest exponent of the inverse cocycle is* $-\operatorname{exponents}_{d-1}$*, the negative of the smallest forward exponent.*

Proof. Specialize 6.22 at $i = 0$, where $\operatorname{rev} 0 = d - 1$. ☐

6.9 Restriction to invariant subbundles

Definition 6.24 (Invariant subbundle)**.** An invariant subbundle is a measurable family of fibre subspaces $W(x) \leq \mathbb{R}^d$ that is A-equivariant a.e.: $A(x)\,W(x) = W(Tx)$.

Lemma 6.25 (Dimension interlacing)**.** *At each ambient flag level, the dimension captured by the subbundle is bounded by the ambient dimension:* $\dim\left(W(x) \cap V_i(x)\right) \leq \dim V_i(x)$*, so the restricted multiplicities are a sub-multiset of the ambient ones.*

Proof. Monotonicity of finrank under $W \cap V_i \leq V_i$. ☐

Theorem 6.26 (Restricted strict Oseledets filtration)**.** *For ergodic T, invertible A, and an invariant subbundle W, collapsing the constant-dimension levels of $i \mapsto W \cap V_i$ yields a genuine* strict *Oseledets filtration inside W: a strictly antitone λ' and a measurable family V' with, μ-a.e., $V'_0(x) = W(x)$, $V'_{k'}(x) = 0$, strictly descending and A-equivariant, with exact growth rate λ'_i per stratum and all levels $\leq W(x)$.*

Proof. Obtain a forward witness via 6.2, restrict the flag to W and read its a.e.-constant dimension profile (the restricted analogue of 6.8); enumerate the first-occurrence surviving indices via an order embedding to collapse repeats into a strict flag, inheriting equivariance and per-stratum growth. The top level is W, not $\top$, so the result is stated directly rather than through 6.1. ☐

6.10 The non-ergodic spectrum

Theorem 6.27 (Non-ergodic exponents)**.** *For merely measure-preserving T (no ergodicity) and invertible measurable A with the standing integrability, there is a family of T-invariant integrable functions $\lambda_i : X \to \mathbb{R}$ such that for each $i < d$ and μ-a.e. x, $\frac{1}{n} \log \sigma_i(A^{(n)}(x)) \to \lambda_i(x)$. The exponents become invariant functions rather than constants.*

Proof. The non-ergodic Kingman theorem applied to the subadditive cocycle $\log \operatorname{sprod}_k$ produces invariant integrable partial-sum limits G_k; set $\lambda_i = G_{i+1} - G_i$, with the per-$\sigma$ telescoping as in the ergodic case. ☐

Corollary 6.28 (Non-ergodic positive-exponent sum)**.** *Summing the positive parts* $\max(\lambda_i(x), 0)$ *over $i < d$ yields a single T-invariant integrable function $G_+ : X \to \mathbb{R}$, the non-ergodic analogue of the positive-exponent sum.*

Proof. A finite sum of positive parts of the invariant integrable functions of 6.27 is invariant and integrable. ☐

6.11 Regularity of the exponents

Theorem 6.29 (Fekete infimum representation)**.** *The partial-sum growth rate is the infimum over n of the normalized integrals:*

$$\Gamma_k = \inf_n \frac{1}{n+1} \int_X \log \operatorname{sprod}_k(n+1, x)\, d\mu.$$

Proof. The integral sequence is subadditive (Fekete), so it converges to its infimum; a Fatou estimate against the dominating Birkhoff averages of $\log^+ \|A^{\pm 1}\|$ identifies that limit with the a.e. growth rate Γ_k in both directions. ☐

Theorem 6.30 (Upper semicontinuity of partial sums and top exponent)**.** *Along a filter of generators $B_i \to A$ for which each fixed-n integral is continuous, the partial-sum rate is upper semicontinuous:* $\limsup_i \Gamma_k(B_i) \leq \Gamma_k(A)$*; specializing $k = 1$ gives the same for the top exponent. This is USC, not continuity — equality can fail when a spectral gap closes.*

Proof. Γ_k is an infimum of the per-n continuous normalized integrals (6.29); for each n, $\Gamma_k(B_i)$ is below the n-th integral, whose limit along the filter is the n-th integral of A, so $\limsup_i \Gamma_k(B_i)$ is at most every term of the infimum. ☐

Theorem 6.31 (Lower semicontinuity of the bottom exponent)**.** *The bottom exponent $\lambda_d = \Gamma_d - \Gamma_{d-1}$ is lower semicontinuous in the generator: $\lambda_d(A) \leq \liminf_i \lambda_d(B_i)$.*

Proof. Writing $\Gamma_d = \int \log|\det|$ (the determinant identity 6.20), which is continuous in the generator under the hypothesis, and Γ_{d-1} upper semicontinuous (6.30), the difference $\Gamma_d - \Gamma_{d-1}$ is lower semicontinuous. ☐

6.12 Singular one-sided bounds

Theorem 6.32 (Forward top value)**.** *For ergodic T and a possibly-singular measurable generator with only $\log^+\|A\| \in L^1$ (no invertibility, no inverse integrability), the normalized positive-part log-norms $\frac{1}{n}\log^+\|A^{(n)}(x)\|$ converge μ-a.e. to a constant λ_1^+.*

Proof. Apply the ergodic Kingman theorem to the subadditive, bounded-below, integrable cocycle $\log^+\|A^{(n)}\|$; only the forward integrability is used. ☐

Theorem 6.33 (Upper bound on the singular top exponent)**.** *Under the same singular hypotheses there is λ_1^+ with, μ-a.e.,*

$$\limsup_n \left(\tfrac{1}{n}\log\|A^{(n)}(x)\| : \mathrm{EReal}\right) \le \lambda_1^+.$$

The lim sup *is taken in* EReal *so the bound is unconditional even when the growth tends to $-\infty$; this is one-sided only.*

Proof. Termwise $\log \le \log^+$, then pass to the EReal lim sup; since the $\log^+$ sequence converges to λ_1^+ (6.32), its lim sup is λ_1^+. ☐

Theorem 6.34 (Sharp lim sup in the expanding case)**.** *There is a forward top value λ_1^+ (the a.e. limit of $\frac{1}{n}\log^+\|A^{(n)}\|$) such that, whenever $\lambda_1^+ > 0$, μ-a.e. the genuine log-norm* lim sup *is exactly λ_1^+:*

$$\limsup_n \left(\tfrac{1}{n}\log\|A^{(n)}(x)\| : \mathrm{EReal}\right) = \lambda_1^+.$$

The positivity hypothesis is essential; in the contracting case $\lambda_1^+ = 0$ the genuine growth may tend to $-\infty$ and equality fails.

Proof. The $\le$ half is 6.33. For $\ge$: where $\frac{1}{n}\log^+\|A^{(n)}\| \to \lambda_1^+ > 0$ the sequence is eventually positive, forcing $\log\|A^{(n)}\| > 0$, so $\log^+ = \log$ eventually; the two EReal sequences agree eventually, hence so do their lim sups, which equal λ_1^+. ☐

Theorem 6.35 (Singular top-k volume upper bound)**.** *For ergodic T and a possibly-singular generator with $\log^+\|A\| \in L^1$, there is a constant Γ_k^+ with, μ-a.e.,*

$$\limsup_n \left(\tfrac{1}{n}\log \mathrm{sprod}_k(x) : \mathrm{EReal}\right) \le \Gamma_k^+.$$

The top-k volume growth is bounded above unconditionally (the EReal lim sup *allows volume collapse), again one-sided only.*

Proof. Since $\mathrm{sprod}_k \ge 0$ is submultiplicative with no invertibility, the same $\log^+$-of-a-nonnegative-subadditive-quantity Kingman construction gives the a.e.-constant value Γ_k^+ and the termwise $\log \le \log^+$ domination transfers to the EReal lim sup. ☐

Chapter 7

The two-sided splitting

The one-sided multiplicative ergodic theorem produces, for a measurable matrix cocycle over an ergodic system, a decreasing measurable equivariant *filtration* $\mathbb{R}^d = V_0(x) \supsetneq V_1(x) \supsetneq \cdots \supsetneq V_k(x) = 0$ whose successive quotients carry the distinct Lyapunov exponents $\lambda_0 > \cdots > \lambda_{k-1}$. When the base dynamics is *invertible* and the generator is everywhere invertible with $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1$, the filtration upgrades to a genuine direct-sum *splitting* $\mathbb{R}^d = E_0(x) \oplus \cdots \oplus E_{k-1}(x)$, with each E_i equivariant and the cocycle growing at the exact rate $+\lambda_i$ forward and $-\lambda_i$ backward on $E_i \setminus \{0\}$.

This chapter develops that upgrade. The strategy applies the one-sided theorem twice — once forward to (T, A), once to a reflected *backward generator* over T^{-1} — and intersects the two filtrations. Two genuinely new analytic nodes carry the argument: an identification of the Kingman subadditive constant with the limit of integral means, and a backward-orbit growth envelope for restricted cocycles. Their consequence is the *transversality crux*: forward and backward exponents of any common nonzero vector add to a nonnegative number, so opposite-side sublevels meet in 0. A pure multiset reflection lemma then aligns the backward spectrum to the negated forward one, after which a telescoping-flag lattice argument assembles the splitting.

7.1 The invertible setup and the backward generator

The base dynamics is an invertible measure-preserving system $T : X \simeq_m X$ (a `MeasurableEquiv`) with $T^{-1} = T.\text{symm}$. The backward iterates of the cocycle $A^{(n)}(x) = \text{cocycle}\, A\, T\, n\, x$ are governed by the *backward generator*, a generator over T^{-1} whose cocycle reflects the forward one through inverse and orbit reversal.

Definition 7.1 (Backward generator)**.** For $A : X \to \text{Matrix}\,(\text{Fin}\, d)\,(\text{Fin}\, d)\,\mathbb{R}$ and $T : X \simeq_m X$, the backward generator is

$$\text{backwardGen}\, A\, T\, x \;=\; \left(A(T^{-1}x)\right)^{-1}.$$

Its cocycle is taken over T^{-1}.

Lemma 7.2 (Cocycle recursion, newest factor on the right)**.** *For any generator A and map T, $A^{(n+1)}(x) = A(T^{[n]}x) \cdot A^{(n)}(x)$.*

Proof. This is the companion of the standard recursion $A^{(n+1)}(x) = A^{(n)}(Tx) \cdot A(x)$: apply the additive law $A^{(m+n)} = A^{(m)}(T^{[n]}\cdot) \cdot A^{(n)}$ with $m = 1$ after commuting the summands. ☐

Lemma 7.3 (Backward cocycle identity)**.** *Writing* $B = \mathrm{backwardGen}\, A\, T$*, the cocycle of* B *over* T^{-1} *satisfies*

$$B^{(n)}(x) \;=\; \bigl(A^{(n)}(T^{-n}x)\bigr)^{-1}, \qquad B^{(n)}(T^{n}y) \;=\; \bigl(A^{(n)}(y)\bigr)^{-1}.$$

Proof. Induct on n. The successor step expands $B^{(n+1)}$ by the left recursion, applies the inductive hypothesis at $T^{-1}x$, rewrites $B = (A \circ T^{-1})^{-1}$ by theorem 7.2, and collapses the product of inverses by Matrix.mul_inv_rev; a left-inverse iterate identity $(T \circ T^{-1})^{[n]} = \mathrm{id}$ reindexes the orbit point. The dual form follows by substituting $x = T^n y$ and cancelling $(T^{-1} \circ T)^{[n]} = \mathrm{id}$. This identity is exactly what makes the second (backward) growth limit in the splitting theorem the cocycle $(A^{(n)}(T^{-n}x))^{-1}$, avoiding any moving-point singular-value bookkeeping. $\square$

Proposition 7.4 (Backward standing hypotheses)**.** *If* $\det A(x) \neq 0$ *for all* x*,* A *is measurable,* T *is measure-preserving, and* $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1(\mu)$*, then the backward system* (T^{-1}, B) *satisfies the same four hypotheses:* $\det B \neq 0$ *everywhere,* B *measurable, and* $\log^+ \|B\|, \log^+ \|B^{-1}\| \in L^1(\mu)$*.*

Proof. Invertibility of $B = (A \circ T^{-1})^{-1}$ is $\det(M^{-1}) = (\det M)^{-1} \neq 0$; measurability is composition of matrix inversion with the measurable T^{-1}. For integrability, T^{-1} is measure-preserving (Mathlib's MeasurePreserving.symm), so $\log^+ \|B\| = \log^+ \|A^{-1}\| \circ T^{-1}$ and $\log^+ \|B^{-1}\| = \log^+ \|A\| \circ T^{-1}$ (using $(M^{-1})^{-1} = M$) transfer by composition. Crucially T ergodic implies T^{-1} ergodic (Mathlib's Ergodic.symm), so the one-sided theorem applies verbatim to the backward system. $\square$

Lemma 7.5 (Biinvariant conull set)**.** *For* T *measure-preserving on a probability space and a conull measurable set* S*, there is a conull measurable* $S' \subseteq S$ *invariant under both* T *and* T^{-1}*.*

Proof. Take $S' = \bigl(\bigcap_n (T^{[n]})^{-1} S\bigr) \cap \bigl(\bigcap_n (T^{-[n]})^{-1} S\bigr)$. Each iterate is measure-preserving, so every preimage of the conull S is conull and the countable intersection is conull. Membership is "all forward and backward iterates land in S"; applying T or T^{-1} merely shifts the index, the cross term collapsing through $T^{-1} \circ T = \mathrm{id}$. $\square$

7.2 The strong one-sided export with dimensions

The crux argument needs the dimensions of the filtration spaces, which the headline one-sided statement quantifies away. We re-run the one-sided composition with concrete witnesses, exposing the spectrum λ and the filtration V together with a dimension formula.

Proposition 7.6 (Strong one-sided export)**.** *Under the one-sided hypotheses (with* $[\mathrm{NeZero}\, d]$*) there exist a decreasing exponent sequence* $\mathrm{lam0} : \mathbb{N} \to \mathbb{R}$ *realizing the a.e. per-index singular-value limits, and a measurable filtration* $V : \mathrm{Fin}(k+1) \to X \to \mathrm{Submodule}\, \mathbb{R}\, (\mathbb{R}^d)$ *where* $k = \mathrm{numExp}\, \mathrm{lam0}\, d$*, such that for a.e.* x*:* $V_0 x = \top$*,* $V_k x = \bot$*,* V *is strictly decreasing and* A*-equivariant, every* $v \in V_i x \setminus V_{i+1} x$ *grows at rate* $\mathrm{expEnum}\, \mathrm{lam0}\, d\, i$*, and*

$$\mathrm{finrank}\,(V_i x) \;=\; \#\{\, j < d : \mathrm{lam0}\, j \le \mathrm{expEnum}\, \mathrm{lam0}\, d\, i \,\}.$$

Proof. This is the committed one-sided assembly, re-run with the concrete witness $V = V'\, A\, T\, \mathrm{lam0}$: obtain lam0 from the singular-value exponents, discharge the top-gap mass envelope, build the spectral, slow-flag and growth interfaces exactly as the headline proof does, and read the structural block off the same conull set. The dimension formula is supplied by theorem 7.7 combined with the slow-flag identification $V'_i x = \mathrm{vslow}\,(\exp \lambda_i)\, x$. $\square$

Proposition 7.7 (Rank of the slow space)**.** *For a.e. x and all $t \in \mathbb{R}$,*

$$\operatorname{finrank}\left(\operatorname{vslow} A\, T\, (\exp t)\, x\right) \;=\; \#\{\, j < d : \operatorname{lam0} j \le t \,\}.$$

Proof. The sanitized limit operator $\widehat{\Lambda}$ has an orthonormal eigenbasis with eigenvalues $\exp(\operatorname{lamSing} x\, e)$, and a.e. $\operatorname{lamSing} x\, j = \operatorname{lam0} j$ (a countable conjunction over $j < d$). The slow space is the range of a spectral sublevel projector $\operatorname{cfc} f\, \widehat{\Lambda}$; theorem 7.8 shows it acts diagonally on the eigenbasis with eigenvalues in $\{0, 1\}$, and the rank of such a self-adjoint idempotent equals the number of 1-eigenvectors, which by exp-monotonicity of the threshold is $\#\{j : \operatorname{lam0} j \le t\}$. ☐

Lemma 7.8 (Functional calculus on an eigenvector)**.** *If M is self-adjoint and $Mv = c\, v$ for $v \neq 0$, then $(\operatorname{cfc} f\, M)\, v = f(c) \cdot v$.*

Proof. The eigenvalue c lies in the (finite) spectrum of M. Pick a Lagrange interpolating polynomial q with $q = f$ on the spectrum; then $\operatorname{cfc} f\, M = \operatorname{cfc} q\, M = q(M)$, and $q(M)v = q(c)v = f(c)v$ because v is an eigenvector. This is pointwise, so no measurability is incurred. ☐

7.3 The Kingman means identification

The repository's ergodic Kingman theorem exhibits a constant c with $(1/n)g_n \to c$ a.e., but defers identifying c with the Fekete infimum of the integral means. The two-sided theorem is the first consumer that needs the means form, in order to equate the Kingman constants of a subadditive cocycle over T and of its orbit-reversed version over T^{-1}.

Proposition 7.9 (Kingman constant as the limit of integral means)**.** *Under the hypotheses of the ergodic Kingman theorem, there is a constant c with*

$$\frac{1}{n+1}\int_X g_{n+1}\, d\mu \longrightarrow c \qquad \text{and} \qquad \tfrac{1}{n}\, g_n(x) \longrightarrow c \;\; \text{for a.e. } x.$$

Proof. Let L be the Fekete limit of the means (existing by subadditivity). For $c \le L$: iterate subadditivity $g_{mn}(x) \le \sum_{j<m} g_n(T^{[jn]}x)$, divide by mn and let $m \to \infty$; the left side converges to c along a subsequence, the right to the Birkhoff average of $(1/n)g_n$ for the measure-preserving $T^{[n]}$ (no ergodicity needed), whose integral is $(1/n)\int g_n$; integrate the inequality. For $c \ge L$: apply Fatou to the nonnegative sequence $A_n - \operatorname{cdiv}_n$, where A_n is the Birkhoff average of g_1 (nonnegative by single-step subadditivity) and cdiv_n tends to L; since $A_n - \operatorname{cdiv}_n \to B - c$ a.e. and $\int A_n = \int g_1$, Fatou yields $\int g_1 - L \ge \int g_1 - c$. ☐

7.4 Restricted cocycles and their exponent

To measure the growth of $A^{(n)}$ restricted to a filtration level V_i, we form a restricted cocycle by post-composing the orthogonal projector onto V_i. A *floor* term repairs the failure of everywhere-subadditivity off the good set, so that Kingman applies without an a.e. caveat.

Definition 7.10 (Floored restricted log-cocycle)**.** For a measurable family $V : X \to \operatorname{Submodule} \mathbb{R}\, (\mathbb{R}^d)$, set $\operatorname{restGen} A\, V\, x = A(x) \cdot P_{V(x)}$ (with P_K the orthogonal projector onto K), $\operatorname{sFloor} A\, T\, n\, x = \prod_{j<n} \left\|(A(T^{[j]}x))^{-1}\right\|^{-1}$, and

$$\operatorname{restLog} A\, V\, T\, n\, x \;=\; \log\!\Big(\left\|\operatorname{cocycle}(\operatorname{restGen} A\, V)\, T\, n\, x\right\| \,\sqcup\, \operatorname{sFloor} A\, T\, n\, x\Big).$$

Lemma 7.11 (Everywhere subadditivity)**.** $\mathrm{restLog}\, A\, V\, T$ *is an everywhere subadditive cocycle: for all* m, n, x*,* $\mathrm{restLog}\,(m+n)\,x \le \mathrm{restLog}\, m\, x + \mathrm{restLog}\, n\, (T^{[m]}x)$*.*

Proof. The floor is multiplicative, $\mathrm{sFloor}\,(m+n) = \mathrm{sFloor}\, m \cdot (\mathrm{sFloor}\, n) \circ T^{[m]}$, the norm of the restricted product is submultiplicative, and $\max(ab, cd) \le \max(a, c)\max(b, d)$ for nonnegatives; taking log (monotone) gives subadditivity with no exceptional set. The floor is the only mechanism that keeps the everywhere signature Kingman requires; on the good set it is dominated and disappears. ☐

Lemma 7.12 (Restricted Kingman, both directions of time)**.** *Kingman applied to* restLog *over* T *yields a constant* χ_V*; the orbit-reversed cocycle* $h_n(x) = \mathrm{restLog}_n(T^{-n}x)$ *is subadditive over* T^{-1} *with the same integral means, hence converges a.e. to the same* χ_V*.*

Proof. Over T, theorem 7.11 together with integrability of the endpoints gives a Kingman constant χ_V. The reversed family $h_n = \mathrm{restLog}_n \circ T^{-n}$ is subadditive over T^{-1} by an index juggle through the additive law, and $\int h_n = \int \mathrm{restLog}_n$ by measure-preservation of T^{-n}. Since T^{-1} is ergodic, theorem 7.9 forces the two Kingman constants — both being the common limit of the identical integral means — to coincide. ☐

Proposition 7.13 (Restricted exponent equals λ_i)**.** *For the forward level* V_i*, the restricted Kingman constant equals the* i*-th exponent:* $\chi_{V_i} = \mathrm{expEnum}\, \mathrm{lam0}\, d\, i$*.*

Proof. Lower bound $\ge$: on the good set, for $v \in V_i x$ the restricted cocycle equals $A^{(n)}(x)v$; a stratum witness from strictness of the filtration realizes the rate λ_i. Upper bound $\le$: pointwise, expand any $P_{V_i(y)}w$ in a classical orthonormal basis of $V_i(y)$; then $\|A^{(n)}Pw\| \le \sum_j \|A^{(n)}e_j\|$ with each e_j in a stratum of index $\ge i$, so its growth is $\le \lambda_i$; the log-of-a-finite-sum lemma passes the maximum through. Both bounds match the a.e.-constant Kingman limit. ☐

Proposition 7.14 (Backward-orbit growth envelope)**.** *For a.e.* x*,*

$$\limsup_n \tfrac{1}{n} \log\bigl\|A^{(n)}(T^{-n}x) \cdot P_{V_i(T^{-n}x)}\bigr\| \ \le\ \mathrm{expEnum}\, \mathrm{lam0}\, d\, i.$$

Proof. This is the analytic heart. By theorem 7.12 the reversed restricted log converges a.e. to the same constant χ_{V_i}, which theorem 7.13 identifies as λ_i. On the good set the floor is absorbed, so the lim sup of the genuine restricted norm along the backward orbit is bounded by λ_i. Only the $\le$ direction is consumed downstream. ☐

7.5 The transversality crux

The envelope feeds the central nonnegativity fact: a vector cannot decay backward strictly faster than $-\lambda_i$ while lying in V_i. The crux is pointwise — ergodicity is not used here — and directly kills nonzero intersection vectors of opposite-side sublevels.

Proposition 7.15 (Sublevels of opposite sign are transverse)**.** *Fix* x*. If the forward envelope for* V_x *holds at rate* a*,* U_x *decays backward at rate* $\le b$*, and* $a + b < 0$*, then* $V_x \sqcap U_x = \perp$*.*

Proof. Suppose $0 \ne v \in V_x \sqcap U_x$. Put $v_n = B^{(n)}(x)v$; by theorem 7.3 $v = A^{(n)}(T^{-n}x)\, v_n$ and, by forward equivariance along the backward orbit, $v_n \in V_i(T^{-n}x)$. Hence $\log\|v\| \le \text{restLog-envelope}_n(x) + \log\|v_n\|$, whose right side behaves like $n(a + b + 2\varepsilon)$, tending to $-\infty$. This contradicts $\|v\|$ being a positive constant. ☐

Proposition 7.16 (The a.e. crux)**.** *For a.e. x, for every forward level i and backward level s with* $\operatorname{expEnum} \mathrm{lam0}\, d\, i + \operatorname{expEnum} \mathrm{mu0}\, d\, s < 0$,

$$V_i x \ \sqcap \ W_s x \ = \ \bot.$$

Proof. Bundle all the a.e. facts — the forward envelope (one application of theorem 7.14 per forward level), the backward growth (flag descent in W) and equivariance along the orbit — onto a single biinvariant conull set via theorem 7.5. On it, every pair (i, s) with negative exponent sum satisfies the hypotheses of theorem 7.15, so the intersection is $\bot$. The quantifiers over the finite $\mathrm{Fin}\, k \times \mathrm{Fin}\, l$ are discharged simultaneously. ☐

Corollary 7.17 (Counting bound)**.** *For all $a, b \in \mathbb{R}$ with $a + b < 0$,*

$$\#\{\, j < d : \mathrm{lam0}\, j \le a \,\} \ + \ \#\{\, j < d : \mathrm{mu0}\, j \le b \,\} \ \le \ d.$$

Proof. Convert thresholds to filtration levels and apply theorem 7.16 at one good point: the two sublevel spaces meet in $\bot$, so by the Grassmann formula their dimensions sum to at most $\dim \mathbb{R}^d = d$. The dimension formula of theorem 7.6 (forward and backward) turns the dimensions into the stated counts. The bound is deterministic, so it holds outright. ☐

7.6 Spectral reflection

The backward spectrum must be aligned to the forward one. The determinant identity $\sum_j \mathrm{lam0}\, j = \int \log |\det A|$ — and its backward negation — together with the counting bound pin the backward exponents to the negated, reversed forward ones, by a purely combinatorial multiset argument that avoids any exterior-power calculus for $\|\bigwedge^q M^{-1}\|$.

Proposition 7.18 (Forward determinant sum)**.** *Any exponent sequence* lam0 *realizing the a.e. singular-value limits satisfies* $\sum_{j<d} \mathrm{lam0}\, j = \int_X \log |\det A(x)| \, d\mu$.

Proof. By uniqueness of a.e. limits at a common conull point, $\mathrm{lam0}\, j$ equals the chosen spectrum $\mathrm{exponents}\, j$ for each $j < d$; the sum of those is the integral of $\log |\det A| = \log \prod_j \sigma_j$ via the established singular-value/determinant identity. ☐

Proposition 7.19 (Backward sum is the negated forward sum)**.** *For a backward exponent sequence* mu0, $\sum_{j<d} \mathrm{mu0}\, j = -\sum_{j<d} \mathrm{lam0}\, j$.

Proof. Apply theorem 7.18 to the backward system (legitimate by theorem 7.4): $\sum \mathrm{mu0}\, j = \int \log |\det B|$. The pointwise identity $\log |\det B| = -\log \left|\det(A \circ T^{-1})\right|$ and the change of variables along the measure-preserving T^{-1} give $\int \log |\det B| = -\int \log |\det A| = -\sum \mathrm{lam0}\, j$. ☐

Proposition 7.20 (Reflection lemma)**.** *Let $p, q : \mathbb{N} \to \mathbb{R}$ be antitone on $[0, d)$. If $\#\{p \le a\} + \#\{q \le b\} \le d$ whenever $a + b < 0$, and $\sum_{j<d} q\, j = -\sum_{j<d} p\, j$, then $q\, j = -p\,(d - 1 - j)$ for all $j < d$.*

Proof. Apply the counting bound at b and $a = -q(j) - \varepsilon$: since antitone tuples are their own sorted enumerations, this gives the sorted domination $q\, j \ge -p\,(d - 1 - j)$ for each j. A pointwise $\ge$ whose total sum is an equality forces pointwise equality, so $q\, j = -p\,(d - 1 - j)$. This is pure finite combinatorics with no analytic content. ☐

Corollary 7.21 (Aligned backward index)**.** *Under the reflection* $\mathrm{mu0}\, j = -\mathrm{lam0}\,(d - 1 - j)$, *the backward count of distinct exponents equals the forward one, and the index* $\mathrm{sidx}\, i = \mathrm{cast}\,(\mathrm{Fin.rev}\, i)$ *satisfies* $\operatorname{expEnum} \mathrm{mu0}\, d\, (\mathrm{sidx}\, i) = -\operatorname{expEnum} \mathrm{lam0}\, d\, i$.

Proof. From theorem 7.20 the multisets of distinct values agree up to negation and reversal, so $\mathrm{numExp}\,\mathrm{mu0}\,d = \mathrm{numExp}\,\mathrm{lam0}\,d$ and the enumerations satisfy $\mathrm{expEnum}\,\mathrm{mu0}\,d\,a = -\mathrm{expEnum}\,\mathrm{lam0}\,d\,(\mathrm{Fin.rev}\,a)$. Substituting $a = \mathrm{sidx}\,i = \mathrm{Fin.rev}\,i$ and computing the index by omega gives the negated forward exponent. ☐

7.7 Measurability of the intersection bundle

The split bundle is $x \mapsto V_i x \sqcap W_{\mathrm{sidx}\,i}\, x$. Measurable subspaces do not close under $\sqcap$ for free, and Mathlib has no alternating-projection theorem; instead a single von Neumann-style power lemma supplies measurability.

Proposition 7.22 (Powers of the projector triple converge to the intersection projector)**.** *For subspaces K, L of $\mathbb{R}^d$, with P_K, P_L the orthogonal projectors,*

$$(P_K\, P_L\, P_K)^n \longrightarrow P_{K\sqcap L} \qquad (n \to \infty).$$

Proof. The matrix $S = P_K P_L P_K$ is self-adjoint, PSD, and a contraction, so its eigenvalues lie in $[0, 1]$; by theorem 7.23 its 1-eigenspace is exactly $K \sqcap L$. Diagonalizing by the spectral theorem, $c^n \to 1$ if $c = 1$ and $\to 0$ otherwise, so S^n converges to the orthogonal projection onto the 1-eigenspace, namely $P_{K\sqcap L}$. ☐

Lemma 7.23 (The intersection fixes exactly $K \sqcap L$)**.** *$(P_K\, P_L\, P_K)\, v = v$ if and only if $v \in K \sqcap L$.*

Proof. The "if" direction is immediate since both projectors fix vectors in $K \sqcap L$. For "only if", $\langle Sv, v\rangle = \|P_L P_K v\|^2$; the Cauchy–Schwarz equality chain $\|v\|^2 = \langle Sv, v\rangle \le \|v\|^2$ forces $v \in K$ and $P_K v \in L$, whence $v \in K \sqcap L$. ☐

Proposition 7.24 (Measurable intersection of measurable subspaces)**.** *If $x \mapsto V(x)$ and $x \mapsto W(x)$ are measurable subspace families, then so is $x \mapsto V(x) \sqcap W(x)$.*

Proof. Each $x \mapsto (P_{V(x)} P_{W(x)} P_{V(x)})^n$ is measurable (matrix powers of measurable matrices). By theorem 7.22 these converge pointwise to $P_{V(x)\sqcap W(x)}$, so the limiting projector is measurable entrywise as a limit of measurable functions; this exhibits the intersection family as measurable. ☐

7.8 Assembling the splitting

With the aligned reflection, the crux disjointness, and the dimension formulas in hand, the splitting is assembled by intersecting each forward level with the matching backward level and telescoping a flag into a direct sum.

Lemma 7.25 (Telescoping-flag lattice lemma)**.** *A descending flag $V_0 \supseteq \cdots \supseteq V_k = \bot$ in a modular lattice, with complements E_i satisfying $V_i = E_i \sqcup V_{i+1}$ and $E_i \sqcap V_{i+1} = \bot$, telescopes into an independent family (E_i) with $\bigsqcup_i E_i = V_0$.*

Proof. Induct on k. The cons step shows that prepending a head E_0 disjoint from the supremum of an already-independent tail preserves independence, using modularity to peel E_0 off the join $E_0 \sqcup C$ intersected with the tail bound. The supremum identity unfolds $\bigsqcup_i E_i = E_0 \sqcup \bigsqcup_i E_{i+1} = E_0 \sqcup V_1 = V_0$ via the first telescoping identity. ☐

Proposition 7.26 (Splitting at a point)**.** *Define* $E_i = V_i \sqcap W_{\operatorname{sidx} i}$. *Then* $\operatorname{finrank} E_i \geq 1$, *the telescoping identities* $V_i = E_i \sqcup V_{i+1}$ *and* $E_i \sqcap V_{i+1} = \bot$ *hold, and consequently* (E_i) *is independent with* $\bigsqcup_i E_i = V_0 = \top$.

Proof. The crux theorem 7.16 at the negated-exponent pair gives $V_{i+1} \sqcap W_{\operatorname{sidx} i} = \bot$, so by Grassmann $V_{i+1} \sqcup W_{\operatorname{sidx} i} = \top$, and the reflection theorem 7.21 fixes $\operatorname{finrank} W_{\operatorname{sidx} i} = d - \#\{\mathrm{lam0} < \lambda_i\}$. Combining with the dimension formula of theorem 7.6 yields $\operatorname{finrank} E_i = \#\{\mathrm{lam0} \leq \lambda_i\} - \#\{\mathrm{lam0} < \lambda_i\} \geq 1$, the telescoping totality $V_i = E_i \sqcup V_{i+1}$ by equal finrank, and disjointness from a second crux instance. theorem 7.25 then telescopes to independence and total supremum. ☐

Theorem 7.27 (The two-sided Oseledets splitting)**.** *Let* $T : X \simeq_m X$ *be an invertible ergodic measure-preserving transformation of a probability space, and let* $A : X \to \operatorname{Matrix}(\operatorname{Fin} d)(\operatorname{Fin} d)\,\mathbb{R}$ *be measurable with* $\det A(x) \neq 0$ *for all* x *and* $\log^+ \|A\|, \log^+ \|A^{-1}\| \in L^1(\mu)$. *Then there exist* $k \in \mathbb{N}$, *a strictly decreasing* $\lambda : \operatorname{Fin} k \to \mathbb{R}$, *and measurable subspace families* $E : \operatorname{Fin} k \to X \to \operatorname{Submodule} \mathbb{R}\,(\mathbb{R}^d)$ *such that for* μ*-a.e.* x*:*

- $\mathbb{R}^d = \bigoplus_i E_i(x)$ *is an internal direct sum with every* $E_i(x) \neq \bot$*;*
- *each* E_i *is* A*-equivariant:* $(A(x))(E_i(x)) = E_i(Tx)$*;*
- *for every nonzero* $v \in E_i(x)$*,*

$$\tfrac{1}{n} \log\left\|A^{(n)}(x)\, v\right\| \to \lambda_i \qquad \textit{and} \qquad \tfrac{1}{n} \log\left\|\left(A^{(n)}(T^{-n}x)\right)^{-1} v\right\| \to -\lambda_i.$$

Proof. For $d = 0$ take $k = 0$ (the empty internal sum of the zero module). For $d > 0$, apply theorem 7.6 forward to obtain $\mathrm{lam0}, V$ and, via theorem 7.4, to the backward system (T^{-1}, B) to obtain $\mathrm{mu0}, W$, each with its dimension formula. Run the crux theorem 7.16 and counting bound theorem 7.17; combine with theorem 7.19 through theorem 7.20 to get the reflection $\mathrm{mu0}\, j = -\mathrm{lam0}\,(d - 1 - j)$, hence $l = k$ and the index alignment theorem 7.21. Set $\lambda = \operatorname{expEnum} \mathrm{lam0}\, d$ and $E_i(x) = V_i x \sqcap W_{\operatorname{sidx} i}\, x$; measurability is theorem 7.24. On a conull set, theorem 7.26 provides the internal direct sum and nonzeroness. Equivariance is Submodule.map over $\sqcap$ (injective $A(x)$): the forward factor is the one-sided equivariance, and the backward factor is transported through $\operatorname{backwardGen} A\, T\,(Tx) = (A(x))^{-1}$ by theorem 7.3. For growth, a nonzero $v \in E_i(x)$ lies in the i-th forward stratum and the $\operatorname{rev} i$-th backward stratum (two crux instances exclude deeper levels), so the two one-sided growth limits apply; the backward limit is rewritten as $-\lambda_i$ through theorem 7.3 and theorem 7.21. ☐

Chapter 8

The continuous-flow multiplicative ergodic theorem

8.1 Overview

The discrete multiplicative ergodic theorem governs a single measure-preserving map $T : X \to X$ and the iterates of a matrix cocycle generated one step at a time. This chapter lifts that theorem to *continuous time*: the map T is replaced by a one-parameter measure-preserving flow $\varphi : \mathbb{R} \to X \to X$, and the iterated cocycle by a continuous-time linear cocycle $A : \mathbb{R} \to X \to \mathrm{Mat}_d(\mathbb{R})$. The conclusion is a finite Lyapunov spectrum $\lambda_1 > \cdots > \lambda_k$, a measurable filtration that is *flow-equivariant* at every real time, and the exact continuous-parameter growth $t^{-1} \log \|A(t,x)v\| \to \lambda_i$ as $t \to \infty$ over $\mathbb{R}$.

The strategy is a reduction, not a redevelopment of the ergodic machinery for $\mathbb{R}$. We set $T := \varphi(1)$ and read the discrete cocycle off the sampled flow cocycle; the proved discrete theorem 5.30 delivers the integer-time conclusion, and two analytic devices lift it to the continuous parameter: a *between-times sandwich* that controls growth on each interval $[n, n+1)$, and a *shift-invariance of the growth* lim sup that promotes the discrete (time-one) equivariance to equivariance at every real time. No continuous-time Kingman theorem is needed; the integer clock appears only as a technical reduction device. Throughout, X carries no topology: the flow is a measurable measure-preserving action of $(\mathbb{R}, +)$, and the cocycle is measurable in the state at each fixed time.

8.2 The continuous-time data

Definition 8.1 (Measure-preserving one-parameter flow)**.** A *measure-preserving one-parameter flow* on a measurable space X for a measure μ is a family $\varphi : \mathbb{R} \to X \to X$ together with the data $\varphi(0) = \mathrm{id}$, $\varphi(s+t) = \varphi(s) \circ \varphi(t)$ for all $s, t \in \mathbb{R}$, and a proof that each time-t map $\varphi(t)$ preserves μ. No topology on X is assumed; in particular no joint continuity in (t, x) is required.

Lemma 8.2 (Integer times are iterates of the time-one map)**.** *For a measure-preserving flow φ and $n \in \mathbb{N}$, the integer-time map of the flow is the n-fold iterate of its time-one map: $\varphi(n) = (\varphi(1))^{[n]}$.*

Proof. Induction on n. The base case $\varphi(0) = \mathrm{id} = (\varphi(1))^{[0]}$ is the time-zero law. For the step, write $n + 1 = n + 1$ as reals and apply additivity $\varphi(n+1) = \varphi(n) \circ \varphi(1)$, then the inductive hypothesis and $(\varphi(1))^{[n+1]} = (\varphi(1))^{[n]} \circ \varphi(1)$. ☐

Definition 8.3 (Continuous-time linear cocycle over a flow)**.** A *continuous-time linear cocycle* over a measure-preserving flow φ, valued in invertible $d \times d$ real matrices, is a family $A : \mathbb{R} \to X \to \mathrm{Mat}_d(\mathbb{R})$ with $A(0, x) = 1$, the cocycle identity (newest factor on the left) $A(t + s, x) = A(t, \varphi(s)x)\, A(s, x)$, a proof that $\det A(t, x) \neq 0$ for all t, x, and measurability of each time-t map $A(t, \cdot)$.

Proposition 8.4 (Reduction identity at integer times)**.** *For a flow cocycle A over φ, every $n \in \mathbb{N}$ and every x,*

$$A(n, x) = \mathrm{cocycle}\,(A(1, \cdot))\,(\varphi(1))\,n\,x,$$

i.e. at integer times the continuous-time cocycle equals the discrete iterated cocycle generated by its time-one map $A(1, \cdot)$ over the time-one dynamics $\varphi(1)$.

Proof. Induction on n. At $n = 0$ both sides are the identity. For the step, split $A((n + 1), x) = A(n, \varphi(1)x)\, A(1, x)$ by the cocycle identity at $t = n$, $s = 1$, match the recursion $\mathrm{cocycle}\,(n+1)\,x = \mathrm{cocycle}\,n\,(\varphi(1)x) \cdot A(1, x)$, and apply the inductive hypothesis at the point $\varphi(1)x$. ☐

8.3 Reduction to the discrete theorem

The discrete theorem requires the time-one generator to have integrable positive log-norm, both forward and inverse. These follow by evaluating the uniform dominating hypotheses at $s = 1$.

Lemma 8.5 (Integrability of the time-one log-norm)**.** *If $g \in L^1(\mu)$ dominates $\log^+ \|A(s, x)\|$ for all $s \in [0, 1]$ and all x, then the time-one map $A(1, \cdot)$ has integrable positive log-norm.*

Proof. The map $x \mapsto \log^+ \|A(1, x)\|$ is measurable (composition of the measurable $\log^+$, the operator norm, and the measurable time-one cocycle map) and is dominated pointwise by g, taking $s = 1 \in [0, 1]$ in the hypothesis. Dominated by an integrable function, it is integrable. ☐

Lemma 8.6 (Integrability of the inverse time-one log-norm)**.** *If $g' \in L^1(\mu)$ dominates $\log^+ \|(A(s, x))^{-1}\|$ for all $s \in [0, 1]$ and all x, then the inverse $(A(1, \cdot))^{-1}$ has integrable positive log-norm.*

Proof. Identical to 8.5, inserting the measurable matrix-inversion map and evaluating the dominating hypothesis at $s = 1$. ☐

Proposition 8.7 (Discrete filtration for the time-one data)**.** *Let μ be a probability measure, φ a measure-preserving flow with $\varphi(1)$ ergodic, and A a flow cocycle with the two uniform integrable dominators g, g'. Then there exist k, exponents $\lambda : \mathrm{Fin}\, k \to \mathbb{R}$, and a subspace family V forming an Oseledets filtration for the generator $A(1, \cdot)$ over the dynamics $\varphi(1)$.*

Proof. Apply the discrete theorem 5.30 to the ergodic map $\varphi(1)$ and generator $A(1, \cdot)$, feeding it invertibility ($\det A(1, \cdot) \neq 0$), measurability, and the two integrability inputs from 8.5 and 8.6. The output is the desired discrete Oseledets datum. ☐

8.4 Between integer times

The discrete theorem controls growth only along integer times; the next results bridge the gap to a continuous parameter. Both rest on the orbital sublinearity of an integrable function along the flow's integer orbit.

Lemma 8.8 (Error sublinearity along the integer orbit)**.** *For integrable g, g' and a measure-preserving flow φ, for almost every x the combined fluctuation along the integer orbit vanishes:*

$$n^{-1}(g(\varphi(n)x) + g'(\varphi(n)x)) \to 0 \quad (n \to \infty).$$

Proof. Apply the Birkhoff orbital-tail estimate (a.e. $n^{-1}h(T^{[n]}x) \to 0$ for integrable h) to $h = g + g'$ and the time-one map $T = \varphi(1)$, then rewrite the iterate orbit $(\varphi(1))^{[n]}x$ as $\varphi(n)x$ using 8.2. ☐

Theorem 8.9 (Between-times sandwich: continuous growth equals integer-time growth)**.** *Fix a flow φ, a flow cocycle A with uniform forward/inverse one-step controls g, g' on $[0,1]$, a point x, and a nonzero vector v. Suppose the integer fluctuation error vanishes and the integer-time average converges, $n^{-1}\log\|A(n,x)v\| \to L$. Then the continuous-time average converges to the same limit:*

$$t^{-1}\log\|A(t,x)v\| \to L \quad (t \to \infty).$$

Proof. Write $t = r + n$ with $n = \lfloor t \rfloor \geq 1$ and $r \in [0,1)$. The cocycle identity splits $A(t,x) = A(r,\varphi(n)x)\,A(n,x)$, so with $w = A(n,x)v$ one sandwiches $\log\|A(t,x)v\|$ between $\log\|w\| - \log\|(A(r,\varphi(n)x))^{-1}\|$ and $\log\|w\| + \log\|A(r,\varphi(n)x)\|$. Both correction terms are bounded in absolute value by $g(\varphi(n)x) + g'(\varphi(n)x)$ via the $[0,1]$ controls (using $\|M\|\,\|M^{-1}\| \geq 1$). Dividing by t, the error term vanishes by hypothesis, the discrete average converges to L along $\lfloor t \rfloor$, and the floor ratio $\lfloor t \rfloor / t \to 1$; a squeeze over the floor delivers the continuous-time limit. ☐

8.5 Equivariance at every real time

The discrete theorem gives equivariance one integer step at a time. To obtain equivariance at every real time t_0 we use the intrinsic growth characterization: a vector lies in level V_i iff it is zero or its discrete growth $\limsup$ is $\leq \lambda_i$. The fixed matrix $A(t_0, x)$ is a bounded bijection, so it perturbs the per-step log-norm by $o(n)$, hence leaves the growth $\limsup$ unchanged.

Lemma 8.10 (Fixed-time log-norm is sublinear)**.** *Fix a real time t_0. For almost every x, both $n^{-1}\log\|A(t_0,\varphi(n)x)\|$ and $n^{-1}\log\|(A(t_0,\varphi(n)x))^{-1}\|$ tend to 0 as $n \to \infty$.*

Proof. One builds an integrable function H dominating both $\log^+\|A(t_0,\cdot)\|$ and $\log^+\|(A(t_0,\cdot))^{-1}\|$: by induction on $\lfloor t_0 \rfloor$ one splits $A((\rho+n)+1,\cdot) = A(\rho+n,\varphi(1)\cdot)\,A(1,\cdot)$, bounds $\log^+$ of a product by the sum, and uses measure-preservation of $\varphi(1)$; negative times reduce to positive ones via $A(t_0,y) = (A(-t_0,\varphi(t_0)y))^{-1}$. The two-sided bound $|\log\|M\|| \leq \log^+\|M\| + \log^+\|M^{-1}\|$ turns H into an integrable dominator for the absolute log-norm; the Birkhoff orbital tail of H along the integer orbit and a squeeze finish the proof. ☐

Theorem 8.11 (Shift-invariance of the growth limsup)**.** *Fix a real time t_0. For almost every x and every test vector u, the discrete-time growth* $\limsup$ *of the cocycle applied to u at x equals that at $\varphi(t_0)x$ applied to the pushed-forward vector $A(t_0,x)\,u$.*

Proof. First, a.e. the discrete growth average $n^{-1}\log\|\mathrm{cocycle}\, n\, x\, u\|$ has bounded range, with upper and lower bounds from the Furstenberg–Kesten Fekete inequalities $\log\|\mathrm{cocycle}\| \leq \mathrm{birkhoffSum}(\log^+\|A(1,\cdot)\|)$ (and the inverse version), whose Birkhoff averages converge by the ergodic theorem. The cocycle identity and 8.4 give the shift relation $\mathrm{cocycle}\, n\, (\varphi(t_0)x) \cdot A(t_0,x) = A(t_0,\varphi(n)x) \cdot \mathrm{cocycle}\, n\, x$, so the two growth averages differ by $n^{-1}(\log\|A(t_0,\varphi(n)x)(\cdots)\| - \log\|\cdots\|)$, which is squeezed to 0 by 8.10. A difference tending to 0 between two range-bounded sequences leaves the $\limsup$ unchanged. ☐

Theorem 8.12 (Flow-equivariance of the filtration at every real time)**.** *Let V be the Oseledets filtration of the time-one data, and fix $t_0 \in \mathbb{R}$. Then for almost every x and every level i,*

$$\operatorname{map}\,(A(t_0,x))\,(V_i\,x) = V_i\,(\varphi(t_0)x).$$

Proof. Use the growth characterization $v \in V_i\,x \iff v = 0 \vee \limsup \le \lambda_i$ at x and, pulled back along the measure-preserving $\varphi(t_0)$, at $\varphi(t_0)x$. The map $P = A(t_0,x)$ is a bijection with inverse $A(t_0,x)^{-1}$. For a non-bottom level, prove both inclusions by le_antisymm: membership of v (resp. $P^{-1}v$) is transported through the equivalence of growth lim sups supplied by 8.11. The bottom level $V_k = \bot$ maps to $\bot$. ☐

8.6 The continuous-flow theorem

Theorem 8.13 (Continuous-flow multiplicative ergodic theorem)**.** *Let μ be a probability measure on X, let φ be a measure-preserving one-parameter flow whose time-one map $\varphi(1)$ is μ-ergodic, and let A be a continuous-time linear cocycle over φ valued in invertible $d \times d$ real matrices. Suppose $g, g' \in L^1(\mu)$ satisfy, for all $s \in [0,1]$ and all x,*

$$\log^+ \|A(s,x)\| \le g(x), \qquad \log^+ \left\|(A(s,x))^{-1}\right\| \le g'(x).$$

Then there exist $k \in \mathbb{N}$, a strictly decreasing sequence of exponents $\lambda : \operatorname{Fin} k \to \mathbb{R}$, and a measurable family of subspaces $V : \operatorname{Fin}(k+1) \to X \to \operatorname{Submodule} \mathbb{R}\,(\operatorname{EuclideanSpace} \mathbb{R}\,(\operatorname{Fin} d))$ such that:

- *λ is strictly decreasing and each V_i is a measurable subspace family;*
- ***(full flow equivariance)** for every $t \in \mathbb{R}$, almost every x has $A(t,x) \cdot V_i\,x = V_i\,(\varphi(t)x)$ for all i;*
- *almost every x carries the strict flag $\top = V_0\,x \supsetneq \cdots \supsetneq V_k\,x = \bot$, and on each stratum $v \in V_i\,x \setminus V_{i+1}\,x$ the continuous-time growth rate is exactly λ_i:*

$$t^{-1} \log \|A(t,x)v\| \to \lambda_i \quad (t \to \infty).$$

Proof. Take (k, λ, V) from the reduction 8.7; this supplies strict antitonicity, measurability, the strict flag, and the integer-time growth rates. Full flow equivariance at each t is 8.12. For the continuous-time growth, work a.e. on the discrete-conclusion set intersected with the error-sublinearity set of 8.8. Fix a stratum vector $v \ne 0$; rewriting via the reduction identity 8.4 turns the discrete stratum growth $n^{-1} \log \|\operatorname{cocycle} n\,x\,v\| \to \lambda_i$ into $n^{-1} \log \|A(n,x)v\| \to \lambda_i$. The between-times sandwich 8.9 then upgrades this integer-time limit to the continuous-parameter limit $t^{-1} \log \|A(t,x)v\| \to \lambda_i$. ☐