

Omri Sarig

Lecture Notes on Ergodic Theory

Penn State, Fall 2008

December 12, 2008

(Prepared using the Springer svmono author package)

Contents

1	Basic definitions and constructions	1
1.1	What is ergodic theory and how it came about.....	1
1.2	The abstract setup of ergodic theory	3
1.3	The probabilistic point of view.....	4
1.4	Ergodicity and mixing	5
1.5	Examples	7
1.5.1	Circle rotations	7
1.5.2	Angle Doubling	8
1.5.3	Bernoulli Schemes	9
1.5.4	Finite Markov Chains.....	11
1.5.5	The geodesic flow on a hyperbolic surface	16
1.6	Basic constructions	19
1.6.1	Skew-products	20
1.6.2	Factors	22
1.6.3	The natural extension.....	23
1.6.4	Induced transformations	24
1.6.5	Suspensions and Kakutani skyscrapers	26
	Problems	27
	References	29
2	Ergodic Theorems	31
2.1	The Mean Ergodic Theorem	31
2.2	The Pointwise Ergodic Theorem	32
2.3	The non-ergodic case	34
2.3.1	Conditional expectations and the limit in the ergodic theorem	35
2.3.2	Conditional probabilities	36
2.3.3	The ergodic decomposition	38
2.4	The Subadditive Ergodic Theorem	39
2.5	The Multiplicative Ergodic Theorem	43
2.5.1	Preparations from Multilinear Algebra	43
2.5.2	Proof of the Multiplicative Ergodic Theorem	48

2.5.3	The Multiplicative Ergodic Theorem for Invertible Cocycles	57
	Problems	60
	References	62
3	Spectral Theory	63
3.1	The spectral approach to ergodic theory	63
3.2	Weak mixing	65
3.2.1	Definition and characterization	65
3.2.2	Spectral measures and weak mixing	66
3.3	The Koopman operator of a Bernoulli scheme	69
	Problems	71
	References	75
4	Entropy	77
4.1	Information content and entropy	77
4.2	Properties of the entropy of a partition	79
4.2.1	The entropy of $\alpha \vee \beta$	79
4.2.2	Convexity properties	80
4.2.3	Information and independence	81
4.3	The Metric Entropy	81
4.3.1	Definition and meaning	81
4.3.2	The Shannon–McMillan–Breiman Theorem	84
4.3.3	Sinai’s Generator theorem	85
4.4	Examples	87
4.4.1	Bernoulli schemes	87
4.4.2	Irrational rotations	87
4.4.3	Markov chains	87
4.4.4	Expanding Markov Maps of the Interval	88
4.5	Abramov’s Formula	89
4.6	Topological Entropy	91
4.6.1	The Adler–Konheim–McAndrew definition	91
4.6.2	Bowen’s definition	94
4.6.3	The variational principle	95
	Problems	97
	References	98
A	The isomorphism theorem for standard measure spaces	99
A.1	Polish spaces	99
A.2	Standard probability spaces	100
A.3	Atoms	101
A.4	The isomorphism theorem	102
A	The Monotone Class Theorem	107
Index		109

Chapter 1

Basic definitions and constructions

1.1 What is ergodic theory and how it came about

Dynamical systems and ergodic theory. Ergodic theory is a part of the theory of dynamical systems. At its simplest form, a *dynamical system* is a function T defined on a set X . The *iterates* of the map are defined by induction $T^0 := id$, $T^n := T \circ T^{n-1}$, and the aim of the theory is to describe the behavior of $T^n(x)$ as $n \rightarrow \infty$.

More generally one may consider the *action of a semi-group* of transformations, namely a family of maps $T_g : X \rightarrow X$ ($g \in G$) satisfying $T_{g_1} \circ T_{g_2} = T_{g_1 g_2}$. In the particular case $G = \mathbb{R}^+$ or $G = \mathbb{R}$ we have a family of maps T_t such that $T_t \circ T_s = T_{t+s}$, and we speak of a *semi-flow* or a *flow*.

The original motivation was classical mechanics. There X is the set of all possible states of given dynamical system (sometimes called *configuration space* or *phase space*), and $T : X \rightarrow X$ is the *law of motion* which prescribes that if the system is at state x now, then it will evolve to state $T(x)$ after one unit of time. The *orbit* $\{T^n(x)\}_{n \in \mathbb{Z}}$ is simply a record of the time evolution of the system, and the understanding the behavior of $T^n(x)$ as $n \rightarrow \infty$ is the same as understanding the behavior of the system at the far future. Flows T_t arise when one insists on studying continuous, rather than discrete time. More complicated group actions, e.g. \mathbb{Z}^d -actions, arise in material science. There $x \in X$ codes the configuration of a d -dimensional lattice (e.g. a crystal), and $\{T_{\underline{v}} : \underline{v} \in \mathbb{Z}^d\}$ are the symmetries of the lattice.

The theory of dynamical systems splits into subfields which differ by the structure which one imposes on X and T :

1. *Differentiable dynamics* deals with actions by differentiable maps on smooth manifolds;
2. *Topological dynamics* deals with actions of continuous maps on topological spaces, usually compact metric spaces;
3. *Ergodic theory* deals with measure preserving actions of measurable maps on a measure space, usually assumed to be finite.

It may seem strange to assume so little on X and T . The discovery that such meagre assumptions yield non trivial information is due to Poincaré, who should be considered the progenitor of the field.

Poincaré's Recurrence Theorem and the birth of ergodic theory. Imagine a box filled with gas, made of N identical molecules. Classical mechanics says that if we know the positions $\underline{q}_i = (q_i^1, q_i^2, q_i^3)$ and momenta $\underline{p}_i = (p_i^1, p_i^2, p_i^3)$ of the i -th molecule for all $i = 1, \dots, N$, then we can determine the positions and momenta of each molecule at time t by solving Hamilton's equations

$$\begin{aligned} \dot{p}_i^j(t) &= -\partial H / \partial q_i^j \\ \dot{q}_i^j(t) &= \partial H / \partial p_i^j. \end{aligned} \quad (1.1)$$

$H = H(q_1, \dots, q_N; p_1, \dots, p_N)$, the *Hamiltonian*, is the total energy of the system.

It is natural to call $(\underline{q}, \underline{p}) := (q_1, \dots, q_N; p_1, \dots, p_N)$ the *state* of the system. Let X denote the collection of all possible states. If we assume (as we may) that the total energy is bounded above, then this is a open bounded subset of \mathbb{R}^{6N} . Let

$$T_t : (\underline{q}, \underline{p}) \mapsto (\underline{q}(t), \underline{p}(t))$$

denote the map which gives solution of (1.1) with initial condition $(\underline{q}(0), \underline{p}(0))$. If H is sufficiently regular, then (1.1) had a unique solution for all t and every initial condition. The uniqueness of the solution implies that T_t is a flow. The law of conservation of energy implies that $\underline{x} \in X \Rightarrow T_t(\underline{x}) \in X$ for all t .

Question: Suppose the system starts at a certain state $(\underline{q}(0), \underline{p}(0))$, will it eventually return to a state close to $(\underline{q}(0), \underline{p}(0))$?

For general H , the question seems intractable because (1.1) is strongly coupled system of an enormous number of equations ($N \sim 10^{24}$). Poincaré's startling discovery is that the question is trivial, if viewed from the right perspective. To understand his solution, we need to recall a classical fact, known as *Liouville's theorem*: The Lebesgue measure m on X satisfies $m(T_t E) = m(E)$ for all t and all measurable $E \subset X$ (problem 1.1).

Here is Poincaré's solution. Define $T := T_1$, and observe that $T^n = T_n$. Fix $\varepsilon > 0$ and consider the set W of all states $\underline{x} = (\underline{q}, \underline{p})$ such that $d(\underline{x}, T^n(\underline{x})) > \varepsilon$ for all $n \geq 1$. Divide W into finitely many disjoint pieces W_i of diameter less than ε .

For each fixed i , the sets $T^{-n}(W_i)$ ($n \geq 1$) are pairwise disjoint: If $T^{-n}(W_i) \cap T^{-(n+k)}(W_i) \neq \emptyset$, then there is some $\underline{x} \in W_i \cap T^{-k}(W_i) \neq \emptyset$ and

1. $\underline{x} \in T^{-k}(W_i)$ implies that $T^k(\underline{x}) \in W_i$ whence $d(\underline{x}, T^k \underline{x}) \leq \text{diam}(W_i) < \varepsilon$, whereas
2. $\underline{x} \in W_i \subset W$ implies that $d(\underline{x}, T^k \underline{x}) > \varepsilon$ by the definition of W .

Since $\{T^{-n}W_i\}_{n \geq 1}$ are pairwise disjoint, $m(X) \geq \sum_{k \geq 1} m(T^{-k}W_i)$. But $T^{-k}(W_i)$ all have the same measure (Liouville theorem), and $m(X) < \infty$, so we must have $m(W_i) = 0$. Summing over i we get that $m(W) = 0$. In summary, a.e. \underline{x} has the property that $d(T^n(\underline{x}), \underline{x}) < \varepsilon$ for some $n \geq 1$. Considering the countable collection $\varepsilon = 1/n$, we obtain the following:

Poincaré's Recurrence Theorem: For almost every $\underline{x} = (q(0), p(0))$, if the system is at state \underline{x} at time zero, then it will return arbitrarily close to this state infinitely many times in the arbitrarily far future.

Poincaré's Recurrence Theorem is a tour de force, because it turns a problem which looks intractable to a triviality by simply looking at it from a different angle. The only thing the solution requires is the existence of a finite measure on X such that $m(T^{-1}E) = m(E)$ for all measurable sets E . This startling realization raises the following mathematical question: What other dynamical information can one extract from the existence of a measure m such that $m = m \circ T^{-1}$? Of particular interest was the justification of the following "assumption" made by Boltzmann in his work on statistical mechanics:

The Ergodic Hypothesis: For certain invariant measures μ , many functions $f : X \rightarrow \mathbb{R}$, and many states $\underline{x} = (q, p)$, the time average of f $\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t(\underline{x})) dt$ exists, and equals the space average of f , $\frac{1}{\mu(X)} \int f d\mu$.

(This is not Boltzmann's original formulation.) The ergodic hypothesis is a quantitative version of Poincaré's recurrence theorem: If f is the indicator of the ε -ball around a state \underline{x} , then the time average of f is the frequency of times when $T_t(\underline{x})$ is ε -away from \underline{x} , and the ergodic hypothesis is a statement on its value.

1.2 The abstract setup of ergodic theory

The proof of Poincaré's Recurrence Theorem suggests the study of the following setup.

Definition 1.1. A *measure space* is a triplet (X, \mathcal{B}, μ) where

1. X is a set, sometime called the *space*.
2. \mathcal{B} is a σ -*algebra*, namely a collection of subsets of X which contains the empty set, and which is closed under complements and countable unions. The elements of \mathcal{B} are called *measurable sets*.
3. $\mu : \mathcal{B} \rightarrow [0, \infty]$, called the *measure*, is a σ -additive function, namely a function s.t if $E_1, E_2, \dots \in \mathcal{B}$ are pairwise disjoint, then $\mu(\bigsqcup_i E_i) = \sum_i \mu(E_i)$.

If $\mu(X) = 1$ then we say that μ is a *probability measure* and (X, \mathcal{B}, μ) is a *probability space*.

In order to avoid measure theoretic pathologies, we will always assume that (X, \mathcal{B}, μ) is the completion (see problem 1.2) of a *standard measure space*, namely a measure space (X, \mathcal{B}', μ') , where X is a complete, metric, separable space and \mathcal{B}' is its Borel σ -algebra. It can be shown that such spaces are *Lebesgue spaces*, namely measure spaces which are isomorphic to a compact interval equipped with the Lebesgue σ -algebra and measure, and a finite or countable collection of points with positive measure. See the appendix for details.

Definition 1.2. A *measure preserving transformation* (mpt) is a quartet (X, \mathcal{B}, μ, T) where (X, \mathcal{B}, μ) is a measure space, and

1. T is measurable: $E \in \mathcal{B} \Rightarrow T^{-1}E \in \mathcal{B}$;
2. m is T -invariant: $m(T^{-1}E) = m(E)$ for all $E \in \mathcal{B}$.

A *probability preserving transformation* (ppt) is a mpt on a probability space.

This is the minimal setup needed to prove (problem 1.3):

Theorem 1.1 (Poincaré's Recurrence Theorem). Suppose (X, \mathcal{B}, μ, T) is a p.p.t. If E is a measurable set, then for almost every $x \in E$ there is a sequence $n_k \rightarrow \infty$ such that $T^{n_k}(x) \in E$.

Poincaré's theorem is not true for general infinite measure preserving transformations, as the example $T(x) = x + 1$ on \mathbb{Z} demonstrates.

Having defined the objects of the theory, we proceed to declare when do we consider two objects to be isomorphic:

Definition 1.3. Two m.p.t. $(X_i, \mathcal{B}_i, \mu_i, T_i)$ are called *measure theoretically isomorphic*, if there exists a measurable map $\pi : X_1 \rightarrow X_2$ such that

1. there are $X'_i \in \mathcal{B}_i$ such that $m_i(X_i \setminus X'_i) = 0$ and such that $\pi : X'_1 \rightarrow X'_2$ is invertible with measurable inverse;
2. for every $E \in \mathcal{B}_2$, $\pi^{-1}(E) \in \mathcal{B}_1$ and $m_1(\pi^{-1}E) = m_2(E)$;
3. $T_2 \circ \pi = \pi \circ T_1$ on X_1 .

One of the main aims of ergodic theorists is to devise method for deciding whether two mpt's are isomorphic.

1.3 The probabilistic point of view.

Much of the power and usefulness of ergodic theory is due to the following probabilistic interpretation of the abstract set up discussed above. Suppose (X, \mathcal{B}, μ, T) is a ppt. We think of

1. X as of a *sample space*, namely the collection of all possible states ω of a random system;
2. \mathcal{B} as the collection of all *measurable events*, namely all sets $E \subset X$ such that we have enough information to answer the question "is $\omega \in E$?";
3. μ is the *probability law*: $\Pr[\omega \in E] := \mu(E)$;
4. measurable functions $f : X \rightarrow \mathbb{R}$ are *random variables* $f(\omega)$;
5. the sequence $X_n := f \circ T^n$ ($n \geq 1$) is a *stochastic process*, whose distribution is given by the formula

$$\Pr[X_{i_1} \in E_{i_1}, \dots, X_{i_k} \in E_{i_k}] := \mu \left(\bigcap_{j=1}^k \{\omega \in X : f(T^{i_j} \omega) \in E_{i_j}\} \right).$$

The invariance of μ guarantees that such stochastic processes are always *stationary*: $\Pr[X_{i_1+m} \in E_{i_1+m}, \dots, X_{i_k} \in E_{i_k+m}] = \Pr[X_{i_1} \in E_{i_1}, \dots, X_{i_k} \in E_{i_k}]$ for all m .

The point is that we can ask what are the properties of the stochastic processes $\{f \circ T^n\}_{n \geq 1}$ arising out of the ppt (X, \mathcal{B}, μ, T) , and thus bring in tools and intuition from probability theory to the study of dynamical systems. Note that we have found a way of studying stochastic phenomena in a context which is, a priori, completely deterministic (if we know the state of the system at time zero is x , then we know with full certainty that its state at time n is $T^n(x)$). The modern treatment of the question “how come a deterministic system can behave randomly” is based on this idea.

1.4 Ergodicity and mixing

Suppose (X, \mathcal{B}, μ, T) is a mpt. A measurable set $E \in \mathcal{B}$ is called *invariant*, if $T^{-1}(E) = E$. Evidently, in this case T can be split into two measure preserving transformations $T|_E : E \rightarrow E$ and $T|_{E^c} : E^c \rightarrow E^c$, which do not interact.

Definition 1.4. A mpt (X, \mathcal{B}, μ, T) is called *ergodic*, if every invariant set E satisfies $\mu(E) = 0$ or $\mu(X \setminus E) = 0$. We say μ is an ergodic measure.

Proposition 1.1. Suppose (X, \mathcal{B}, μ, T) is a mpt on a complete measure space, then the following are equivalent:

1. μ is ergodic;
2. if $E \in \mathcal{B}$ and $\mu(T^{-1}E \Delta E) = 0$, then $\mu(E) = 0$ or $\mu(X \setminus E) = 0$;
3. if $f : X \rightarrow \mathbb{R}$ is measurable and $f \circ T = f$ a.e., then there is $c \in \mathbb{R}$ s.t. $f = c$ a.e.

Proof. Suppose μ is ergodic, and E is measurable s.t. $\mu(E \Delta T^{-1}E) = 0$. We construct a measurable set E_0 such that $T^{-1}E_0 = E_0$ and $\mu(E_0 \Delta E) = 0$. By ergodicity $\mu(E_0) = 0$ or $\mu(X \setminus E_0) = 0$. Since $\mu(E \Delta E_0) = 0$ implies that $\mu(E) = \mu(E_0)$ and $\mu(X \setminus E) = \mu(X \setminus E_0)$ we get that either $\mu(E) = 0$ or $\mu(X \setminus E) = 0$.

The set E_0 we use is $E_0 := \{x \in X : T^k(x) \in E \text{ infinitely often.}\}$. It is obvious that this set is measurable and invariant. To estimate $\mu(E_0 \Delta E)$ note that

- (a) if $x \in E_0 \setminus E$, then there exists some k s.t. $x \in T^{-k}(E) \setminus E$;
- (b) if $x \in E \setminus E_0$, then there exists some k s.t. $x \notin T^{-k}(E)$, whence $x \in E \setminus T^{-k}(E)$.

Thus $E_0 \Delta E \subset \bigcup_{k \geq 1} E \Delta T^{-k}(E)$. We now use the following “triangle inequality” : $\mu(A_1 \Delta A_3) \leq \mu(A_1 \Delta A_2) + \mu(A_2 \Delta A_3)$ ($A_i \in \mathcal{B}$) (prove!):

$$\begin{aligned} \mu(E_0 \Delta E) &\leq \sum_{k=1}^{\infty} \mu(E \Delta T^{-k}E) \leq \sum_{k=1}^{\infty} \sum_{i=0}^k \mu(T^{-i}E \Delta T^{-(i+1)}E) \\ &= \sum_{k=1}^{\infty} k \mu(E \Delta T^{-1}E) \quad (\because \mu \circ T^{-1} = \mu). \end{aligned}$$

Since $\mu(E\Delta T^{-1}E) = 0$, $\mu(E_0\Delta E) = 0$ and we have shown that (1) implies (2).

Next assume (2). and let f be a measurable function s.t. $f \circ T = f$ almost everywhere. For every t , $[f > t]\Delta T^{-1}[f > t] \subset [f \neq f \circ T]$, so

$$\mu([f > t]\Delta T^{-1}[f > t]) = 0.$$

By assumption, this implies that either $\mu[f > t] = 0$ or $\mu[f \leq t] = 0$. In other words, either $f > t$ a.e., or $f \leq t$ a.e. Define $c := \sup\{t : f > t \text{ a.e.}\}$. then $f = c$ almost everywhere, proving (3). The implication (3) \Rightarrow (2) is obvious: take $f = 1_E$. \square

An immediate corollary is that ergodicity is an invariant of measure theoretic isomorphism: If two mpt are isomorphic, then the ergodicity of one implies the ergodicity of the other.

The next definition is motivated by the probabilistic notion of *independence*. Suppose (X, \mathcal{B}, μ) is a *probability space*. We think of elements of \mathcal{B} as of “events”, we interpret measurable functions $f : X \rightarrow \mathbb{R}$ as “random variables”, and we view μ as a “probability law” $\mu(E) = \mathbb{P}[x \in E]$. Two events $E, F \in \mathcal{B}$ are called *independent*, if $\mu(E \cap F) = \mu(E)\mu(F)$ (because in the case $\mu(E), \mu(F) \neq 0$ this is equivalent to saying that $\mu(E|F) = \mu(E), \mu(F|E) = \mu(F)$).

Definition 1.5. A probability preserving transformation (X, \mathcal{B}, μ, T) is called *mixing* (or *strongly mixing*), if for all $E, F \in \mathcal{B}$, $\mu(E \cap T^{-k}F) \xrightarrow[k \rightarrow \infty]{} \mu(E)\mu(F)$. (There is no notion of strong mixing for infinite measure spaces.)

In other words, $T^{-k}(F)$ is “asymptotically independent” of E . It is easy to see that strong mixing is an invariant of measure theoretic isomorphism.

It can be shown that the sets E, F in the definition of mixing can be taken to be equal (problem 1.12).

Proposition 1.2. *Strong mixing implies ergodicity.*

Proof. Suppose E is invariant, then $\mu(E) = \mu(E \cap T^{-n}E) \xrightarrow[n \rightarrow \infty]{} \mu(E)^2$, whence $\mu(E)^2 = \mu(E)$. It follows that $\mu(E) = 0$ or $\mu(E) = 1 = \mu(X)$. \square

Just like ergodicity, strong mixing can be defined in terms of functions. Before we state the condition, we recall a relevant notion from statistics. The *correlation coefficient* of $f, g \in L^2(\mu)$ is defined to be

$$\rho(f, g) := \frac{\int fg d\mu - \int f d\mu \cdot \int g d\mu}{\|f - \int f d\mu\|_2 \|g - \int g d\mu\|_2}.$$

The numerator is equal to

$$\text{Cov}(f, g) := \int [(f - \int f)(g - \int g)] d\mu,$$

called the *covariance* of f, g . The idea behind this quantity is that if f, g are weakly correlated then they will not always deviate from their means in the same direction, leading to many cancelations in the integral, and a small net result. If f, g are

strongly correlated, there will be less cancelations, and a larger net result. (The denominator in the definition of ρ is not important - it is there to force $\rho(f, g)$ to take values in $[-1, 1]$.)

Proposition 1.3. *A ppt (X, \mathcal{B}, μ, T) is strongly mixing iff for every $f, g \in L^2$, $\int f g \circ T^n d\mu \xrightarrow{n \rightarrow \infty} \int f d\mu \int g d\mu$, equivalently $\text{Cov}(f, g \circ T^n) \xrightarrow{n \rightarrow \infty} 0$.*

Proof. The condition implies mixing (take $f = 1_E, g = 1_F$). We show the other direction. We need the following trivial observations:

1. Since $\mu \circ T^{-1} = \mu$, $\|f \circ T\|_p = \|f\|_p$ for all $f \in L^p$ and $1 \leq p \leq \infty$;
2. $\text{Cov}(f, g)$ is bilinear in f, g ;
3. $|\text{Cov}(f, g)| \leq 4\|f\|_2\|g\|_2$.

The first two statements are left as an exercise. For the third we use the Cauchy-Schwarz inequality (twice) to observe that

$$\begin{aligned} |\text{Cov}(f, g)| &\leq \|f - \int f\|_2 \|g - \int g\|_2 \\ &\leq (\|f\|_2 + \|f\|_1)(\|g\|_2 + \|g\|_1) \leq (2\|f\|_2)(2\|g\|_2). \end{aligned}$$

Assume that μ is mixing, and let f, g be two elements of L^2 . If f, g are indicators of measurable sets, then $\text{Cov}(f, g \circ T^n) \rightarrow 0$ by mixing. If f, g are finite linear combinations of indicators, $\text{Cov}(f, g \circ T^n) \rightarrow 0$ because of the bilinearity of the covariance. For general $f, g \in L^2$, we can find for every $\varepsilon > 0$ finite linear combinations of indicators $f_\varepsilon, g_\varepsilon$ s.t. $\|f - f_\varepsilon\|_2, \|g - g_\varepsilon\|_2 < \varepsilon$. By the observations above,

$$\begin{aligned} |\text{Cov}(f, g \circ T^n)| &\leq |\text{Cov}(f - f_\varepsilon, g \circ T^n)| + |\text{Cov}(f_\varepsilon, g_\varepsilon \circ T^n)| + |\text{Cov}(f_\varepsilon, (g - g_\varepsilon) \circ T^n)| \\ &\leq 4\varepsilon\|g\|_2 + o(1) + 4(\|f\|_2 + \varepsilon)\varepsilon, \text{ as } n \rightarrow \infty. \end{aligned}$$

It follows that $\limsup |\text{Cov}(f, g \circ T^n)| \leq 4\varepsilon(\|f\|_2 + \|g\|_2 + \varepsilon)$. Since ε is arbitrary, the limsup, whence the limit itself, is equal to zero. \square

1.5 Examples

We illustrate these definitions by examples.

1.5.1 Circle rotations

Let $\mathbb{T} := [0, 1)$ equipped with the Lebesgue measure m , and define for $\alpha \in [0, 1)$ $R_\alpha : \mathbb{T} \rightarrow \mathbb{T}$ by $R_\alpha(x) = x + \alpha \pmod{1}$. R_α is called a *circle rotation*, because the map $\pi(x) = \exp[2\pi ix]$ is an isomorphism between R_α and the rotation by the angle $2\pi\alpha$ on S^1 .

Proposition 1.4.

1. R_α is measure preserving for every α ;
2. R_α is ergodic iff $\alpha \notin \mathbb{Q}$;
3. R_α is never strongly mixing.

Proof. A direct calculation shows that the Lebesgue measure m satisfies $m(R_\alpha^{-1}I) = m(I)$ for all intervals $I \subset [0, 1)$. Thus the collection $\mathcal{M} := \{E \in \mathcal{B} : m(R_\alpha^{-1}E) = m(E)\}$ contains the algebra of finite disjoint unions of intervals. It is easy to check \mathcal{M} is a monotone class, so by the monotone class theorem (see appendix) \mathcal{M} contains all Borel sets. It clearly contains all null sets. Therefore it contains all Lebesgue measurable sets. Thus $\mathcal{M} = \mathcal{B}$ and (1) is proved.

We prove (2). Suppose first that $\alpha = p/q$ for $p, q \in \mathbb{N}$. Then $R_\alpha^q = id$. Fix some $x \in [0, 1)$, and pick ε so small that the ε -neighborhoods of $x + k\alpha$ for $k = 0, \dots, q-1$ are disjoint. The union of these neighborhoods is an invariant set of positive measure, and if ε is sufficiently small then it is not equal to \mathbb{T} . Thus R_α is not ergodic.

Next assume that $\alpha \notin \mathbb{Q}$. Suppose E is an invariant set, and set $f = 1_E$. Expand f to a Fourier series:

$$f = \sum_{n \in \mathbb{Z}} \widehat{f}(n) e^{2\pi i n t} \quad (\text{convergence in } L^2).$$

The invariance of E dictates $f = f \circ R_\alpha$. The Fourier expansion of $f \circ R_\alpha$ is

$$f \circ R_\alpha = \sum_{n \in \mathbb{Z}} e^{2\pi i n \alpha} \widehat{f}(n) e^{2\pi i n t}.$$

Equating coefficients, we see that $\widehat{f}(n) = \widehat{f}(n) \exp[2\pi i n \alpha]$. Thus either $\widehat{f}(n) = 0$ or $\exp[2\pi i n \alpha] = 1$. Since $\alpha \notin \mathbb{Q}$, $\widehat{f}(n) = 0$ for all $n \neq 0$. We obtain that $f = \widehat{f}(0)$ a.e., whence $1_E = m(E)$ almost everywhere. This can only happen if $m(E) = 0$ or $m(E) = 1$, proving the ergodicity of m .

To show that m is not mixing, we consider the function $f(x) = \exp[2\pi i x]$. This function satisfies $f \circ R_\alpha = \lambda f$ with $\lambda = \exp[2\pi i \alpha]$ (such a function is called an *eigenfunction*). For every α there is a sequence $n_k \rightarrow \infty$ s.t. $n_k \alpha \pmod{1} \rightarrow 0$ (Dirichlet theorem), thus

$$\|f \circ R_\alpha^{n_k} - f\|_2 = |\lambda^{n_k} - 1| \xrightarrow[k \rightarrow \infty]{} 0.$$

It follows that $F := \operatorname{Re}(f) = \cos(2\pi x)$ satisfies $\|F \circ R_\alpha^{n_k} - F\|_2 \xrightarrow[k \rightarrow \infty]{} 0$, whence $\int F \circ R_\alpha^{n_k} F dm \xrightarrow[k \rightarrow \infty]{} \int F^2 dm \neq (\int F)^2$, and m is not mixing. \square

1.5.2 Angle Doubling

Again, we work with $\mathbb{T} := [0, 1]$ equipped with the Lebesgue measure m , and define $T : \mathbb{T} \rightarrow \mathbb{T}$ by $T(x) = 2x \pmod{1}$. T is called the *angle doubling map*, because the map $\pi(x) := \exp[2\pi i x]$ is an isomorphism between T and the map $e^{i\theta} \mapsto e^{2i\theta}$ on S^1 .

Proposition 1.5. *The angle doubling map is probability preserving, and strong mixing, whence ergodic.*

Proof. It is convenient to work with binary expansions $x = 0.d_1d_2d_3\dots$, ($d_i = 0, 1$), because with this representation $T(0.d_1d_2\dots) = 0.d_2d_3\dots$. For every finite n -word of zeroes and ones (d_1, \dots, d_n) , define the sets (called “cylinders”)

$$[d_1, \dots, d_n] := \{x \in [0, 1) : x = 0.d_1 \dots d_n \varepsilon_1 \varepsilon_2 \dots, \text{ for some } \varepsilon_i \in \{0, 1\}\}.$$

This is a (dyadic) interval, of length $1/2^n$.

It is clear that $T^{-1}[d_1, \dots, d_n] = [* , d_1, \dots, d_n]$ where $*$ stands for “0 or 1”. Thus, $m(T^{-1}[d]) = m[0, d] + m[1, d] = 2 \cdot 2^{-(n+1)} = 2^{-n} = m[d]$. We see that $\mathcal{M} := \{E \in \mathcal{B} : m(T^{-1}E) = m(E)\}$ contains the algebra of finite disjoint unions of cylinders. Since \mathcal{M} is obviously a monotone class, and since the cylinders generate the Borel σ -algebra (prove!), we get that $\mathcal{M} = \mathcal{B}$, whence T is measure preserving.

We prove that T is mixing. Suppose f, g are indicators of cylinders: $f = 1_{[a_1, \dots, a_n]}$, $g = 1_{[b_1, \dots, b_m]}$. Then for all $k > n$,

$$\int f \cdot g \circ T^k dm = m[a, \underbrace{*\dots*}_{k-1}, b] = m[a]m[b].$$

Thus $\text{Cov}(f, g \circ T^k) \xrightarrow[k \rightarrow \infty]{} 0$ for all indicators of cylinders. Every L^2 -function can be approximated in L^2 by a finite linear combination of indicators of cylinders (prove!), one can proceed as in the proof of proposition 1.3 to show that $\text{Cov}(f, g \circ T^k) \xrightarrow[k \rightarrow \infty]{} 0$ for all L^2 functions. \square

1.5.3 Bernoulli Schemes

Let S be a finite set, called the *alphabet*, and let $X := S^{\mathbb{N}}$ be the set of all one-sided infinite sequences of elements of S . Impose the following metric on X :

$$d((x_n)_{n \geq 0}, (y_n)_{n \geq 0}) := 2^{-\min\{k : x_k \neq y_k\}}. \quad (1.2)$$

The resulting topology is generated by the collection of *cylinders*:

$$[a_0, \dots, a_{n-1}] := \{\underline{x} \in X : x_i = a_i \text{ (} 0 \leq i \leq n-1)\}.$$

It can also be characterized as being the product topology on $S^{\mathbb{N}}$, when S is given the discrete topology. In particular this topology is compact.

The *left shift* is the transformation $T : (x_0, x_1, x_2, \dots) \mapsto (x_1, x_2, \dots)$. The left shift is continuous.

Next fix a probability vector $\underline{p} = (p_a)_{a \in S}$, namely a vector of positive numbers whose sum is equal to one.

Definition 1.6. The *Bernoulli measure* corresponding to \underline{p} is the unique measure on the Borel σ -algebra of X such that $\mu[a_0, \dots, a_{n-1}] = p_{a_0} \cdots p_{a_{n-1}}$ for all cylinders.

It is useful to recall why such a measure exists. The collection of cylinders is a *semi-algebra*, namely a collection \mathcal{S} such that

1. \mathcal{S} is closed under intersections, and
2. for every $A \in \mathcal{S}$, $X \setminus A$ is a finite disjoint union of elements from \mathcal{S} .

Carathéodory's Extension Theorem says that every σ -additive function on a semi-algebra has a unique extension to a σ -additive function on the σ -algebra generated by \mathcal{S} . Thus it is enough to check that

$$m : [a_0, \dots, a_{n-1}] \mapsto p_{a_0} \cdots p_{a_{n-1}}$$

is σ -additive on \mathcal{S} , namely that if $[\underline{a}]$ is a countable disjoint union of cylinders $[\underline{b}^j]$, then $\mu[\underline{a}] = \sum \mu[\underline{b}^j]$.

Each cylinder is open and compact (prove!), so such unions are necessarily finite. Let N be the maximal length of the cylinder $[\underline{b}^j]$. Since $[\underline{b}^j] \subseteq [\underline{a}]$, we can write $[\underline{b}^j] = [\underline{a}, \underline{\beta}^j] = \bigsqcup_{|\underline{c}|=N-|\underline{b}^j|} [\underline{a}, \underline{\beta}^j, \underline{c}]$, and a direct calculation shows that

$$\sum_{|\underline{c}|=N-|\underline{b}^j|} \mu[\underline{a}, \underline{\beta}^j, \underline{c}] = \mu[\underline{a}, \underline{\beta}^j] \left(\sum_c p_c \right)^{N-|\underline{b}^j|} = \mu[\underline{a}, \underline{\beta}^j] \equiv \mu[\underline{b}^j].$$

Summing over j , we get that

$$\sum_j \mu[\underline{b}^j] = \sum_j \sum_{|\underline{c}|=N-|\underline{b}^j|} \mu[\underline{a}, \underline{\beta}^j, \underline{c}].$$

Now $[\underline{a}] = \bigsqcup_j [\underline{b}^j] = \bigsqcup_j \bigsqcup_{|\underline{c}|=N-|\underline{b}^j|} [\underline{a}, \underline{\beta}^j, \underline{c}]$, so the collection of $(\underline{\beta}^j, \underline{c})$ is equal to the collection of all possible words \underline{w} of length $N - |\underline{a}|$ (otherwise the right hand side misses some sequences). Thus

$$\sum_j \mu[\underline{b}^j] = \sum_{|\underline{w}|=N-|\underline{a}|} \mu[\underline{a}, \underline{w}] = \mu[\underline{a}] \left(\sum_w p_w \right)^{N-|\underline{a}|} = \mu[\underline{a}],$$

proving the σ -additivity of μ on \mathcal{S} . The existence and uniqueness of μ follows from Carathéodory's theorem.

Proposition 1.6. Suppose $X = \{0, 1\}^{\mathbb{N}}$, μ is the $(\frac{1}{2}, \frac{1}{2})$ -Bernoulli measure, and σ is the left shift, then $(X, \mathcal{B}(X), \mu, T)$ is measure theoretically isomorphic to the angle doubling map.

Proof. The isomorphism is $\pi(x_0, x_1, \dots) := \sum 2^{-n} x_n$. This is a bijection between

$$X' := \{\underline{x} \in \{0, 1\}^{\mathbb{N}} : \exists n \text{ s.t. } x_m = 1 \text{ for all } m \geq n\}$$

and $[0, 1)$ (prove that $\mu(X') = 1$), and it is clear that $\pi \circ \sigma = T \circ \pi$. Since the image of a cylinder of length n is a dyadic interval of length 2^{-n} , π preserves the measures of cylinders. The collection of measurable sets which are mapped by π to sets of the same measure is a σ -algebra. Since this σ -algebra contains all the cylinders and all the null sets, it contains all measurable sets. \square

Proposition 1.7. *Every Bernoulli scheme is mixing, whence ergodic.*

The proof is the same as in the case of the angle doubling map, and is therefore omitted.

1.5.4 Finite Markov Chains

We saw that the angle doubling map is isomorphic to a dynamical system acting as the left shift on a space of sequences (a Bernoulli scheme). Such representations appear frequently in the theory of dynamical systems, but more often than not, the space of sequences is slightly more complicated than the set of all sequences.

1.5.4.1 Subshifts of finite type

Let S be a finite set, and $A = (t_{ij})_{S \times S}$ a matrix of zeroes and ones without columns or rows made entirely of zeroes.

Definition 1.7. The *subshift of finite type* with *alphabet* S and *transition matrix* A is the set $\Sigma_A^+ = \{\underline{x} = (x_0, x_1, \dots) \in S^{\mathbb{N}} : t_{x_i x_{i+1}} = 1 \text{ for all } i\}$, together with the metric $d(\underline{x}, \underline{y}) := 2^{-\min\{k: x_k \neq y_k\}}$ and the action $\sigma(x_0, x_1, x_2, \dots) = (x_1, x_2, \dots)$.

This is a compact metric space, and the left shift map $\sigma : \Sigma_A^+ \rightarrow \Sigma_A^+$ is continuous. We think of Σ_A^+ as of the space of all infinite paths on a directed graph with vertices S and edge $a \rightarrow b$ connecting $a, b \in S$ such that $t_{ab} = 1$.

Let Σ_A^+ be a topologically mixing SFT with set of states S , $|S| < \infty$, and transition matrix $A = (A_{ab})_{S \times S}$.

- A *stochastic matrix* is a matrix $P = (p_{ab})_{a, b \in S}$ with non-negative entries, such that $\sum_b p_{ab} = 1$ for all a , i.e. $P1 = 1$. The matrix is called *compatible* with A , if $A_{ab} = 0 \Rightarrow p_{ab} = 0$.
- A *probability vector* is a vector $\underline{p} = \langle p_a : a \in S \rangle$ of non-negative entries, s.t. $\sum p_a = 1$
- A *stationary probability vector* is a probability vector $\underline{p} = \langle p_a : a \in S \rangle$ s.t. $\underline{p}P = \underline{p}$: $\sum_a p_a p_{ab} = p_b$.

Given a probability vector \underline{p} and a stochastic matrix P compatible with A , one can define a *Markov measure* μ on Σ_A^+ (or Σ_A) by

$$\mu[a_0, \dots, a_{n-1}] := p_{a_0} p_{a_0 a_1} \cdots p_{a_{n-2} a_{n-1}}$$

The stochasticity of P guarantees that this measure is finitely additive on the algebra of cylinders, and σ -subadditivity is trivial because of compactness. Thus this gives a Borel probability measure on Σ_A^+ .

Proposition 1.8. μ is shift invariant iff \underline{p} is stationary w.r.t. P . Any stochastic matrix has a stationary probability vector.

Proof. To see the first half of the statement, we note that μ is stationary iff $\mu[* , \underline{b}] = \mu[\underline{b}]$ for all $[\underline{b}]$, which is equivalent to

$$\sum_a p_a p_{ab_0} p_{b_0 b_1} \cdots p_{b_{n-2} b_{n-1}} = \sum_a p_{b_0} p_{b_0 b_1} \cdots p_{b_{n-2} b_{n-1}}.$$

Canceling the identical terms on both sides gives $\sum_a p_a p_{ab_0} = p_{b_0}$. Thus μ is shift invariant iff \underline{p} is P -stationary.

We now show that every stochastic matrix has a stationary probability vector. Consider the right action of P on the simplex Δ of probability vectors in \mathbb{R}^N :

$$\Delta := \{(x_1, \dots, x_N) : x_i \geq 0, \sum x_i = 1\}, \quad T(\underline{x}) = \underline{x}P.$$

We have $T(\Delta) \subseteq \Delta$, since $\sum_a (T\underline{x})_a = \sum_a \sum_b x_b p_{ba} = \sum_b x_b \sum_a p_{ba} = \sum_b x_b = 1$. By Brouwer's fixed point theorem (continuous mapping of a closed convex set) T has a fixed point. This is the stationary probability vector. \square

Thus every stochastic matrix determines (at least one) shift invariant measure. We ask when is this measure ergodic, and when is it mixing.

1.5.4.2 Ergodicity and mixing

There are obvious obstructions to ergodicity and mixing. To state them concisely, we introduce some terminology. Suppose A is a transition matrix. We say that a connects to b in n steps, and write $a \xrightarrow{n} b$, if there is a path of length $n+1$ on the directed graph representing Σ_A^+ which starts at a and ends at b . In terms of A , this means that there are states b_1, \dots, b_{n-1} s.t. $t_{ab_1} t_{b_1 b_2} \cdots t_{b_{n-1} b} > 0$ (see problem 1.5).

Definition 1.8. A transition matrix $A = (t_{ab})_{a,b \in S}$ is called *irreducible*, if for every $a, b \in S$ there exists an n s.t. $a \xrightarrow{n} b$. The *period* of an irreducible transition matrix A is the number $p := \gcd\{n : a \xrightarrow{n} a\}$ (this is independent of a).¹ An irreducible transition matrix is called *aperiodic* if its period is equal to one.

For example, the SFT with transition matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ is irreducible with period two.

¹ Let p_a, p_b denote the periods of a, b , and let $\Lambda_b := \{n : b \xrightarrow{n} b\}$. By irreducibility, there are α, β s.t. $a \xrightarrow{\alpha} b \xrightarrow{\beta} a$. In this case for all $n \in \Lambda_b$, $a \xrightarrow{\alpha} b \xrightarrow{n} b \xrightarrow{\beta} a$, whence $p_a | \gcd(\alpha + \beta + \Lambda_b) | \gcd(\Lambda_b - \Lambda_b)$. Now $\gcd(\Lambda_b - \Lambda_b) | \gcd(\Lambda_b)$, because $\gcd(\Lambda_b) \in \Lambda_b - \Lambda_b$. Thus $p_a | p_b$. By symmetry, $p_b | p_a$ and we obtain $p_a = p_b$.

If A is not irreducible, then any (globally supported) Markov chain measure on Σ_A^+ is non-ergodic. To see why, pick $a, b \in S$ s.t. a does not connect to b in any number of steps. The set

$$E := \{\underline{x} \in \Sigma_A^+ : x_i \neq b \text{ for all } i \text{ sufficiently large}\}$$

is a shift invariant set which contains $[a]$, but which is disjoint from $[b]$. Thus it has positive, non-full measure w.r.t. every (globally supported) Markov chain – an obstruction for ergodicity.

If A is irreducible, but not aperiodic, then any Markov chain measure on Σ_A^+ is non-mixing, because the function

$$f := 1_{[a]}$$

satisfies $ff \circ T^n \equiv 0$ for all n not divisible by the period. This means that $\int ff \circ T^n d\mu$ is equal to zero on a subsequence, and therefore cannot converge to $\mu[a]^2$.

We claim that these are the only possible obstructions to ergodicity and mixing. The proof is based on the following fundamental fact, whose proof will be given at the next section.

Theorem 1.2 (Ergodic Theorem for Markov Chains). *Suppose P is a stochastic matrix, and write $P^n = (p_{ab}^{(n)})$, then P has a stationary probability vector \underline{p} , and*

1. *if P is irreducible, then $\frac{1}{n} \sum_{k=1}^n p_{ab}^{(k)} \xrightarrow{n \rightarrow \infty} p_b$ ($a, b \in S$);*
2. *if P is irreducible and aperiodic, then $p_{ab}^{(k)} \xrightarrow{n \rightarrow \infty} p_b$ ($a, b \in S$).*

Corollary 1.1. *A shift invariant Markov chain measure on a SFT Σ_A^+ is ergodic iff A is irreducible, and mixing iff A is irreducible and aperiodic.*

Proof. Let μ be a Markov chain measure with stochastic matrix P and stationary probability vector \underline{p} . For all cylinders $[a] = [a_0, \dots, a_{n-1}]$ and $[b] = [b_0, \dots, b_{m-1}]$,

$$\begin{aligned} \mu([a] \cap \sigma^{-k}[b]) &= \mu \left(\bigoplus_{\underline{\xi} \in \mathcal{W}_{k-n}} [a, \underline{\xi}, b] \right), \quad \mathcal{W}_\ell := \{\underline{\xi} = (\xi_1, \dots, \xi_\ell) : [a, \underline{\xi}, b] \neq \emptyset\} \\ &= \mu[a] \cdot \sum_{\underline{\xi} \in \mathcal{W}_{k-n}} p_{a_{n-1}\xi_1} \cdots p_{\xi_{k-n-1}b_0} \cdot \frac{\mu[b]}{p_{b_0}} = \mu[a]\mu[b] \cdot \frac{p_{a_{n-1}b_0}^{(k-n)}}{p_{b_0}}. \end{aligned}$$

If A is irreducible, then by theorem 1.2,

$$\frac{1}{n} \sum_{k=0}^{n-1} \mu([a] \cap \sigma^{-k}[b]) \xrightarrow{n \rightarrow \infty} \mu[a]\mu[b].$$

We claim that this implies ergodicity. Suppose E is an invariant set, and fix $\varepsilon > 0$, arbitrarily small. There are cylinders $A_1, \dots, A_N \in \mathcal{S}$ s.t. $\mu(E \Delta \biguplus_{i=1}^N A_i) < \varepsilon$.² Thus

$$\mu(E) = \mu(E \cap \sigma^{-k}E) = \sum_{i=1}^N \mu(A_i \cap \sigma^{-k}E) \pm \varepsilon = \sum_{i,j=1}^N \mu(A_i \cap \sigma^{-k}A_j) \pm 2\varepsilon.$$

Averaging over k , and passing to the limit, we get

$$\begin{aligned} \mu(E) &= \sum_{i,j=1}^N \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mu(A_i \cap \sigma^{-k}A_j) \pm 2\varepsilon = \sum_{i,j=1}^N \mu(A_i)\mu(A_j) \pm 2\varepsilon \\ &= \left(\sum_{i=1}^N \mu(A_i) \right)^2 \pm 2\varepsilon = [\mu(E) \pm \varepsilon]^2 \pm 2\varepsilon. \end{aligned}$$

Passing to the limit $\varepsilon \rightarrow 0^+$, we obtain $\mu(E) = \mu(E)^2$, whence $\mu(E) = 0$ or 1 .

Now assume that A is irreducible and aperiodic. The ergodic theorem for Markov chains says that $\mu(\underline{a} \cap \sigma^{-k}\underline{b}) \xrightarrow[k \rightarrow \infty]{} \mu(\underline{a})\mu(\underline{b})$. Since any measurable sets E, F can be approximated by finite disjoint unions of cylinders, an argument similar to the previous one shows that $\mu(E \cap \sigma^{-k}F) \xrightarrow[k \rightarrow \infty]{} \mu(E)\mu(F)$ for all $E, F \in \mathcal{B}$ and so μ is mixing. \square

1.5.4.3 Proof of the Ergodic Theorem for Markov chains

Suppose first that P is an irreducible and aperiodic stochastic matrix. This implies that there is some power m such that all the entries of P^m are strictly positive.³

Let $N := |S|$ denote the number of states, and consider the N -simplex $\Delta := \{(x_1, \dots, x_N) : x_i \geq 0, \sum x_i = 1\}$ (the set of all probability vectors). Since P is stochastic, the map $T(\underline{x}) = \underline{x}P$ maps Δ continuously into itself. By the Brouwer Fixed Point Theorem, there is a probability vector \underline{p} s.t. $\underline{p}P = \underline{p}$ (this is the stationary probability vector).

The set $C := \Delta - \underline{p}$ is a compact convex neighborhood of the origin, and

$$T(C) \subset C, \quad T^m(C) \subset \text{int}[C] \quad (\text{we mean the relative interior}).$$

The image is in the interior because all entries of P^m are positive, so all coordinates $T^m(\underline{x}) = \underline{x}P^m$ are positive.

² Proof: The collection of sets E satisfying this approximation property is a σ -algebra which contains all cylinders, therefore it is equal to \mathcal{B} .

³ Begin by proving that if A is irreducible and aperiodic, then for every a there is an N_a s.t. $a \xrightarrow{n} a$ for all $n > N_a$. Use this to show that for all a, b there is an N_{ab} s.t. $a \xrightarrow{n} b$ for all $n > N_{ab}$. Take $m = \max\{N_{ab}\}$.

Now consider $L := \text{span}(C)$ (an $N - 1$ -dimensional space). This is an invariant space for T , whence for P^t . We claim that all the eigenvalues of $P^t|_L$ have absolute value less than 1:

1. *Eigenvalues of modulus larger than one are impossible*, because P is stochastic, so $\|\underline{v}P\|_1 \leq \|\underline{v}\|_1$, so the spectral radius of P^t cannot be more than 1.
2. *Roots of unity are impossible*, because in this case P^{km} has eigenvalue one for some k . The eigenvector \underline{v} is real valued. Normalizing it we can ensure $\underline{v} \in \partial C$. But P^{km} cannot have fixed points on ∂C , because $P^{km}(C) \subset \text{int}[C]$.
3. *Eigenvalues $e^{i\theta}$ with $\theta \notin 2\pi\mathbb{Q}$ are impossible*, because if $e^{i\theta}$ is an eigenvalue then there are two real eigenvectors $\underline{u}, \underline{v} \in \partial C$ such that the action of P on $\text{span}\{\underline{u}, \underline{v}\}$ is conjugate to $\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}$, namely an irrational rotation. This means $\exists k_n \rightarrow \infty$ s.t. $\underline{u}P^{mk_n} \rightarrow \underline{u} \in \partial C$. But this cannot be the case because $P^{mk_n}[C] \subset \text{int}[C]$, and by compactness, this cannot intersect ∂C .

In summary the spectral radius of $P^t|_L$ is less than one.

But $\mathbb{R}^N = L \oplus \text{span}\{\underline{p}\}$. If we decompose a general vector $\underline{v} = \underline{q} + t\underline{p}$ with $\underline{q} \in L$, then the above implies that $\underline{v}P^n = t\underline{p} + O(\|P^n|_L\|)\|\underline{q}\| \xrightarrow[n \rightarrow \infty]{} t\underline{p}$. It follows that $p_{ab}^{(n)} \xrightarrow[n \rightarrow \infty]{} p_b$ for all a, b .

This is almost the ergodic theorem for irreducible aperiodic Markov chains, the only thing which remains is to show that P has a unique stationary vector. Suppose \underline{q} is another probability vector s.t. $\underline{q}P = \underline{q}$. We can write $p_{ab}^{(n)} \rightarrow p_b$ in matrix form as follows:

$$P^n \xrightarrow[n \rightarrow \infty]{} Q, \text{ where } Q = (q_{ab})_{S \times S} \text{ and } q_{ab} = p_b.$$

This means that $\underline{q} = \underline{q}P^n \rightarrow \underline{q}Q$, whence $\underline{q}Q = \underline{q}$, so $q_a = \sum_b q_b q_{ba} = \sum_b q_b p_a = p_a$.

Consider now the periodic irreducible case. Let A be the transition matrix associated to P (with entries $t_{ab} = 1$ when $p_{ab} > 0$ and $t_{ab} = 0$ otherwise). Fix once and for all a state v . Working with the SFT Σ_A^+ , we let

$$S_k := \{b \in S : v \xrightarrow{n} b \text{ for some } n = k \pmod{p}\} \quad (k = 0, \dots, p-1).$$

These sets are pairwise disjoint, because if $b \in S_{k_1} \cap S_{k_2}$, then $\exists \alpha_i = k_i \pmod{p}$ and $\exists \beta$ s.t. $v \xrightarrow{\alpha_i} b \xrightarrow{\beta} v$ for $i = 1, 2$. By the definition of the period, $p | \alpha_i + \beta$ for $i = 1, 2$, whence $k_1 = \alpha_1 = -\beta = \alpha_2 = k_2 \pmod{p}$.

It is also clear that every path of length ℓ which starts at S_k , ends at $S_{k'}$ where $k' = k + \ell \pmod{p}$. In particular, every path of length p which starts at S_k ends at S_k . This means that if $p_{ab}^{(p)} > 0$, then a, b belong to the same S_k .

It follows that P^p is conjugate, via a coordinate permutation, to a block matrix with blocks $(p_{ab}^{(p)})_{S_k \times S_k}$. Each of the blocks is stochastic, irreducible, and aperiodic. Let $\underline{\pi}^{(k)}$ denote the stationary probability vectors of the blocks.

By the first part of the proof,

$$p_{ab}^{(\ell p)} \xrightarrow{\ell \rightarrow \infty} \pi_b^{(k)} \text{ for all } a, b \text{ in the same } S_k, \text{ and } p_{ab}^{(n)} = 0 \text{ for } n \not\equiv 0 \pmod{p}.$$

More generally, if $a \in S_{k_1}$ and $b \in S_{k_2}$, then

$$\begin{aligned} \lim_{\ell \rightarrow \infty} p_{ab}^{(\ell p + k_2 - k_1)} &= \lim_{\ell \rightarrow \infty} \sum_{\xi \in S_{k_2}} p_{a\xi}^{(k_2 - k_1)} p_{\xi b}^{(\ell p)} = \sum_{\xi \in S_{k_2}} p_{a\xi}^{(k_2 - k_1)} \pi_b^{(k_2)} \text{ (by the above)} \\ &= \pi_b^{(k_2)} \sum_{\xi \in S} p_{a\xi}^{(k_2 - k_1)} = \pi_b^{(k_2)}. \quad (\because p_{a\xi}^{(k_2 - k_1)} = 0 \text{ when } \xi \notin S_{k_2}) \end{aligned}$$

A similar calculation shows that $\lim_{k \rightarrow \infty} p_{ab}^{(kp + \alpha)} = 0$ when $\alpha \not\equiv k_2 - k_1 \pmod{p}$. We conclude that the following limit exists

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n p_{ab}^{(k)} = \pi_b := \frac{1}{p} \sum_{k=0}^{p-1} \pi_b^{(k)}.$$

The limit $\underline{\pi} = (\pi_b)_{b \in S}$ is a probability vector.

We claim that $\underline{\pi}$ is the unique stationary probability vector of P . The limit theorem for $p_{ab}^{(n)}$ can be written in the form $\frac{1}{n} \sum_{k=0}^{n-1} P^k \rightarrow Q$ where $Q = (q_{ab})_{S \times S}$ and $q_{ab} = \pi_b$. As before this implies that $\underline{\pi}P = \underline{\pi}$ and that any probability vector \underline{q} such that $\underline{q}P = \underline{q}$, we also have $\underline{q}Q = \underline{q}$, whence $\underline{q} = \underline{\pi}$. \square

Remark 1. The theorem has the following nice probabilistic interpretation. Imagine that we distribute mass on the states of S according to a probability distribution $\underline{q} = (q_a)_{a \in S}$. Now shift mass from one state to another using the rule that a p_{ab} -fraction of the mass at a is moved to b . The new mass distribution is $\underline{q}P$ (check). After n steps, the mass distribution is $\underline{q}P^n$. The previous theorem says that, in the aperiodic case, the mass distribution converges to the stationary distribution — the equilibrium state. It can be shown that the rate of convergence is exponential (problem 1.7).

Remark 2: The ergodic theorem for Markov chains has an important generalization to all matrices with non-negative entries, see problem 1.6.

1.5.5 The geodesic flow on a hyperbolic surface

The *hyperbolic plane* is the surface $\mathbb{H} := \{z \in \mathbb{C} : \text{Im}(z) > 0\}$ equipped with the Riemannian metric $ds = 2|dz|/\text{Im}(z)$, which gives it constant curvature (-1) .

It is known that the orientation preserving *isometries* (i.e. distance preserving maps) are the Möbius transformations which preserve \mathbb{H} . They form the group

$$\begin{aligned} \text{Möb}(\mathbb{H}) &= \left\{ \frac{az+b}{cz+d} : a, b, c, d \in \mathbb{R}, ad-bc=1 \right\} \\ &\simeq \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, c, d \in \mathbb{R}, ad-bc=1 \right\} =: \text{PSL}(2, \mathbb{R}). \end{aligned}$$

The *geodesics* (i.e. length minimizing curves) on the hyperbolic plane are vertical half-lines, or circle arcs which meet $\partial\mathbb{H}$ at right angles. Here is why: Suppose $\omega \in TM$ is a unit tangent vector which points directly up, then it is not difficult to see that the geodesic at direction ω is a vertical line. For general unit tangent vectors ω , find an element $\varphi \in \text{Möb}(\mathbb{H})$ which rotates them so that $d\varphi(\omega)$ points up. The geodesic in direction ω is the φ -preimage of the geodesic in direction $d\varphi(\omega)$ (a vertical half-line). Since Möbius transformations map lines to lines or circles in a conformal way, the geodesic of ω is a circle meeting $\partial\mathbb{H}$ at right angles.

The *geodesic flow* of \mathbb{H} is the flow g^t on the *unit tangent bundle* of \mathbb{H} ,

$$T^1\mathbb{H} := \{\omega \in T\mathbb{H} : \|\omega\| = 1\}$$

which moves ω along the geodesic it determines at unit speed.

To describe this flow it useful to find a convenient parametrization for $T^1\mathbb{H}$. Fix $\omega_0 \in T^1\mathbb{H}$ (e.g. the unit vector based at i and pointing up). For every ω , there is a unique $\varphi_\omega \in \text{Möb}(\mathbb{H})$ such that $\omega = d\varphi_\omega(\omega_0)$, thus we can identify

$$T^1\mathbb{H} \simeq \text{Möb}(\mathbb{H}) \simeq \text{PSL}(2, \mathbb{R}).$$

It can be shown, that in this coordinate system, the geodesic flow takes the form

$$g^t \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix}.$$

To verify this, it is enough to calculate the geodesic flow on $\omega_0 \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$.

Next we describe the Riemannian volume measure on $T^1\mathbb{H}$ (up to normalization). Such a measure must be invariant under the action of all isometries. In our coordinate system, the isometry $\varphi(z) = (az+b)/(cz+d)$ acts by

$$\varphi \begin{pmatrix} x & y \\ z & w \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x & y \\ z & w \end{pmatrix}.$$

Since $\text{PSL}(2, \mathbb{R})$ is a locally compact topological group, there is only one Borel measure on $\text{PSL}(2, \mathbb{R})$ (up to normalization), which is left invariant by all left translations on the group: the Haar measure of $\text{PSL}(2, \mathbb{R})$. Thus the Riemannian volume measure is a left Haar measure of $\text{PSL}(2, \mathbb{R})$, and this determines it up to normalization.

It is a particular feature of $\text{PSL}(2, \mathbb{R})$ that its left Haar measure is also invariant under right translations. It follows that the geodesic flow preserves the volume measure on $T^1\mathbb{H}$.

But this measure is infinite, and it is not ergodic (prove!). To obtain ergodic flows, we need to pass to compact quotients of \mathbb{H} . These are called *hyperbolic surfaces*.

A hyperbolic surface is a Riemannian surface M such that every point in M has a neighborhood V which is isometric to an open subset of \mathbb{H} . Recall that a Riemannian surface is called *complete*, if every geodesic can be extended indefinitely in both directions.

Theorem 1.3 (Killing–Hopf Theorem). *Every complete connected hyperbolic surface is isometric to $\Gamma \backslash \mathbb{H}$, where*

1. Γ is a subgroup of $\text{Möb}(\mathbb{H})$;
2. every point $z \in \mathbb{H}$ is in the interior of some open set $U \subset \text{PSL}(2, \mathbb{R})$ s.t. $\{g(U) : g \in \Gamma\}$ are pairwise disjoint;
3. the Riemannian structure on $\{\Gamma z' : z' \in U\}$ is the one induced by the Riemannian structure on F .

If we identify Γ with a subgroup of $\text{PSL}(2, \mathbb{R})$, then we get the identification $T^1(\Gamma \backslash \mathbb{H}) \simeq \Gamma \backslash \text{PSL}(2, \mathbb{R})$. It is clear that Haar measure on $\text{PSL}(2, \mathbb{R})$ induces a unique locally finite measure on $\Gamma \backslash \text{PSL}(2, \mathbb{R})$, and that the geodesic flow on $T^1(\Gamma \backslash \mathbb{H})$ takes the form

$$g^t(\Gamma \omega) = \Gamma \omega \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix},$$

and preserves this measure.

Definition 1.9. A measure preserving flow $g^t : X \rightarrow X$ is called *ergodic*, if every measurable set E such that $g^{-t}(E) = E$ for all t satisfies $m(E) = 0$ or $m(E^c) = 0$.

This is equivalent to asking that all L^2 -functions f such that $f \circ g^t = f$ a.e. for all $t \in \mathbb{R}$ are constant (prove!).

Theorem 1.4. *If $\Gamma \backslash \mathbb{H}$ is compact, then the geodesic flow on $T^1(\Gamma \backslash \mathbb{H})$ is ergodic.*

Proof. Consider the following flows:

$$h_{st}^t(\Gamma \omega) = \Gamma \omega \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$$

$$h_{un}^t(\Gamma \omega) = \Gamma \omega \begin{pmatrix} 1 & 0 \\ t & 1 \end{pmatrix}$$

If we can show that any geodesic invariant function f is also invariant under these flows then we are done, because it is known that

$$\left\langle \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ t & 1 \end{pmatrix}, \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix} \right\rangle = \text{PSL}(2, \mathbb{R})$$

(prove!), and any $\text{PSL}(2, \mathbb{R})$ -invariant function on $\Gamma \backslash \text{PSL}(2, \mathbb{R})$ is constant.

Since our measure is induced by the the Haar measure, the flows h_{un}^t, h_{st}^t are measure preserving. A matrix calculation shows:

$$g^s h_{st}^t g^{-s} = h_{st}^{te^{-s}} \xrightarrow{s \rightarrow \infty} id$$

$$g^{-s} h_{un}^t g^s = h_{un}^{te^s} \xrightarrow{s \rightarrow -\infty} id$$

Step 1. If $f \in L^2$, then $f \circ h^{te^{-s}} \xrightarrow{s \rightarrow \infty} f$.

Proof. Approximate by continuous functions of compact support, and observe that h^t is an isometry of L^2 .

Step 2. If $f \in L^2$ and $f \circ g^s = f$, then $f \circ h_{un}^t = f$ and $f \circ h_{st}^t = f$.

Proof. $\|f \circ h_{st}^t - f\| = \|f \circ g^s \circ h_{st}^t - f\|_2 = \|f \circ g^s \circ h_{st}^t \circ g^{-s} - f\|_2 \rightarrow 0$.

Thus f is h_{st}^t -invariant. A similar calculation shows that it is h_{un}^t -invariant, and we are done. \square

1.6 Basic constructions

In this section we discuss several standard methods for creating new measure preserving transformations from old ones. These constructions appear quite frequently in applications.

Products

Recall that the product of two measure spaces $(X_i, \mathcal{B}_i, \mu_i)$ ($i = 1, 2$) is the measure space $(X_1 \times X_2, \mathcal{B}_1 \otimes \mathcal{B}_2, \mu_1 \times \mu_2)$ where $\mathcal{B}_1 \times \mathcal{B}_2$ is the smallest σ -algebra which contains all set of the form $B_1 \times B_2$ where $B_i \in \mathcal{B}_i$, and $\mu_1 \times \mu_2$ is the unique measure such that $(\mu_1 \times \mu_2)(B_1 \times B_2) = \mu_1(B_1)\mu_2(B_2)$.

This construction captures the idea of *independence* from probability theory: if $(X_i, \mathcal{B}_i, \mu_i)$ are the probability models of two random experiments, and these experiments are “independent”, then $(X_1 \times X_2, \mathcal{B}_1 \otimes \mathcal{B}_2, \mu_1 \times \mu_2)$ is the probability model of the pair of the experiments: If $E_i \in \mathcal{B}_i$, then

$F_1 := E_1 \times X_2$ is the event “in experiment 1, E_1 happened”

$F_2 := X_1 \times E_2$ is the event “in experiment 2, E_2 happened”,

and $F_1 \cap F_2 = E_1 \times E_2$; so $(\mu_1 \times \mu_2)(F_1 \cap F_2) = (\mu_1 \times \mu_2)(F_1)(\mu_1 \times \mu_2)(F_2)$, showing that the events F_1, F_2 are independent.

Definition 1.10. The *product* of two measure preserving systems $(X_i, \mathcal{B}_i, \mu_i, T_i)$ ($i = 1, 2$) is the measure preserving system $(X_1 \times X_2, \mathcal{B}_1 \otimes \mathcal{B}_2, \mu_1 \times \mu_2, T_1 \times T_2)$, where $(T_1 \times T_2)(x_1, x_2) = (T_1 x_1, T_2 x_2)$.

(Check that S is measure preserving.)

Proposition 1.9. *The product of two ergodic mpt is not always ergodic. The product of two mixing mpt is always mixing.*

Proof. The product of two (ergodic) irrational rotations $S := R_\alpha \times R_\alpha : \mathbb{T}^2 \rightarrow \mathbb{T}^2$, $S(x, y) = (x + \alpha, y + \alpha) \pmod{1}$ is not ergodic: $F(x, y) = y - x \pmod{1}$ is a non-constant invariant function. (See problem 1.8.)

The product of two mixing mpt is however mixing. To see this set $\mu = \mu_1 \times \mu_2$, $S = T_1 \times T_2$, and $\mathcal{S} := \{A \times B : A \in \mathcal{B}_1, B \in \mathcal{B}_2\}$. For any $E_1 := A_1 \times B_1, E_2 := A_2 \times B_2 \in \mathcal{S}$,

$$\begin{aligned} \mu(E_1 \cap S^{-n} E_2) &= \mu((A_1 \times B_1) \cap (T_1 \times T_2)^{-n}(A_2 \times B_2)) \\ &= \mu((A_1 \cap T^{-n} A_2) \cap (B_1 \cap T^{-n} B_2)) \\ &= \mu_1(A_1 \cap T^{-n} A_2) \mu_2(B_1 \cap T^{-n} B_2) \\ &\xrightarrow{n \rightarrow \infty} \mu_1(A_1) \mu_2(B_1) \mu_1(A_2) \mu_2(B_2) = \mu(A_1 \times B_1) \mu(A_2 \times B_2). \end{aligned}$$

Since \mathcal{S} is a semi-algebra which generates $\mathcal{B}_1 \otimes \mathcal{B}_2$, any element of $\mathcal{B}_1 \otimes \mathcal{B}_2$ can be approximated by a finite disjoint elements of \mathcal{S} , and therefore $\mu(E \cap S^{-n} F) \rightarrow \mu(E) \mu(F)$ for all $E, F \in \mathcal{B}$. \square

1.6.1 Skew-products

We start with an example. Let μ be the $(\frac{1}{2}, \frac{1}{2})$ -Bernoulli measure on the two shift $\Sigma_2^+ := \{0, 1\}^{\mathbb{N}}$. Let $f : \Sigma_2^+ \rightarrow \mathbb{Z}$ be the function $f(x_0, x_1, \dots) = (-1)^{x_0}$. Consider the transformation

$$T_f : \Sigma_2^+ \times \mathbb{Z} \rightarrow \Sigma_2^+ \times \mathbb{Z}, T_f(\underline{x}, k) = (\sigma(\underline{x}), k + f(x)),$$

where $\sigma : \Sigma_2^+ \rightarrow \Sigma_2^+$ is the left shift. This system preserves the (infinite) measure $\mu \times m_{\mathbb{Z}}$ where $m_{\mathbb{Z}}$ is the counting measure on \mathbb{Z} . The n -th iterate is

$$T_f^n(\underline{x}, k) = (\sigma^n \underline{x}, k + X_0 + \dots + X_{n-1}), \text{ where } X_i := (-1)^{x_i}.$$

What we see in the second coordinate is the symmetric random walk on \mathbb{Z} , started at k , because (1) the steps X_i take the values ± 1 , and (2) $\{X_i\}$ are independent because of the choice of μ . We say that the second coordinate is a “random walk on \mathbb{Z} driven by the noise process $(\Sigma_2^+, \mathcal{B}, \mu, \sigma)$ ”.

Here is a variation on this example. Suppose T_0, T_1 are two measure preserving transformations of the same measure space (Y, \mathcal{C}, ν) . Consider the transformation

$(X \times Y, \mathcal{B} \otimes \mathcal{C}, \mu \times \nu, T_f)$, where

$$T_f(x, y) = (Tx, T_{f(x)}y).$$

The n -th iterate takes the form $T_f^n(x, y) = (\sigma^n x, T_{x_{n-1}} \cdots T_{x_0} y)$. The second coordinate looks like the random concatenation of elements of $\{T_0, T_1\}$. We say that T_f is a “random dynamical system driven by the noise process (X, \mathcal{B}, μ, T) ”.

These examples suggest the following abstract constructions.

Suppose (X, \mathcal{B}, μ, T) is a mpt, and G is a locally compact Polish⁴ topological group, equipped with a left invariant Haar measure m_G . Suppose $f : X \rightarrow G$ is measurable.

Definition 1.11. The *skew-product* with cocycle f over the basis (X, \mathcal{B}, μ, T) is the mpt $(X \times G, \mathcal{B} \otimes \mathcal{B}(G), \mu \times m_G, T_f)$, where $T_f : X \times G \rightarrow X \times G$ is the transformation $T_f(x, g) = (Tx, gf(x))$.

(Check, using Fubini’s theorem, that this is a mpt.) The n -th iterate $T_f^n(x, g) = (T^{n-1}x, f(x)f(Tx) \cdots f(T^{n-1}x)g)$, a “random walk on G ”.

Now imagine that the group G is acting in a measure preserving way on some space (Y, \mathcal{C}, ν) . This means that there are measurable maps $T_g : Y \rightarrow Y$ such that $\nu \circ T_g^{-1} = \nu$, $T_{g_1 g_2} = T_{g_1} T_{g_2}$, and $(g, y) \mapsto T_g(y)$ is a measurable from $X \times G$ to Y .

Definition 1.12. The *random dynamical system* on (Y, \mathcal{C}, ν) with action $\{T_g : g \in G\}$, cocycle $f : X \rightarrow G$, and noise process (X, \mathcal{B}, μ, T) , is the system $(X \times Y, \mathcal{B} \otimes \mathcal{C}, \mu \times \nu, T_f)$ given by $T_f(x, y) = (Tx, T_{f(x)}y)$.

(Check using Fubini’s theorem that this is measure preserving.) Here the n -th iterate is $T_f^n(x, y) = (T^n x, T_{f(T^n x)} \cdots T_{f(Tx)} T_{f(x)} y)$.

It is obvious that if a skew-product (or a random dynamical system) is ergodic or mixing, then its base is ergodic or mixing. The converse is not always true. The ergodic properties of a skew-product depend in a subtle way on the interaction between the base and the cocycle.

Here are two important obstructions to ergodicity and mixing for skew-products. In what follows G is a polish group and \widehat{G} is its group of characters,

$$\widehat{G} := \{\gamma : G \rightarrow S^1 : \gamma \text{ is a continuous homomorphism}\}.$$

Definition 1.13. Suppose (X, \mathcal{B}, μ, T) is a ppt and $f : X \rightarrow G$ is Borel.

1. f is called *arithmetic* w.r.t. μ , if $\exists h : X \rightarrow S^1$ measurable, and $\gamma \in \widehat{G}$ non-constant, such that $\gamma \circ f = h/h \circ T$ a.e.
2. f is called *periodic* w.r.t. μ , if $\exists h : X \rightarrow S^1$ measurable, $|\lambda| = 1$, and $\gamma \in \widehat{G}$ non-constant, such that $\gamma \circ f = \lambda h/h \circ T$ a.e.

Proposition 1.10. Let (X, \mathcal{B}, μ, T) be a ppt, $f : X \rightarrow G$ Borel, and $(X \times G, \mathcal{B} \otimes \mathcal{B}(G), \mu \times m_G, T_f)$ the corresponding skew-product. If f is arithmetic, then T_f is not ergodic, and if f is periodic, then T_f is not mixing.

⁴ “Polish”=has a topology which makes it a complete separable metric space.

Proof. Suppose f is arithmetic. The function $F(x, y) := h(x)\gamma(y)$ satisfies

$$F(Tx, yf(x)) = h(Tx)\gamma(y)\gamma(f(x)) = h(Tx)\gamma(y)h(x)/h(Tx) = F(x, y),$$

and we have a non-constant invariant function, meaning that T_f is not ergodic. Now suppose f is periodic, then $F(x, y) = h(x)\gamma(y)$ satisfies

$$F(Tx, yf(x)) = h(Tx)\gamma(y)\gamma(f(x)) = h(Tx)\gamma(y)\lambda h(x)/h(Tx) = \lambda F(x, y),$$

whence $F \circ T_f = \lambda F$. Pick $n_k \rightarrow \infty$ s.t. $\lambda^{n_k} \rightarrow 1$, then $\text{Cov}(F, F \circ T_f^{n_k}) \xrightarrow[k \rightarrow \infty]{} \int F^2 - (\int F)^2$. Since $F \neq \int F$ a.e., the limit is non-zero and we get a contradiction to mixing. \square

1.6.2 Factors

When we construct skew-products over a base, we enrich the space. There are some-time constructions which deplete the space.

Definition 1.14. A mpt transformation (X, \mathcal{B}, μ, T) is called a (*measure theoretic*) *factor* of a mpt transformation (Y, \mathcal{C}, ν, S) , if there are sets of full measure $X' \subset X, Y' \subset Y$ such that $T(X') \subset X', S(Y') \subset Y'$, and a measurable *onto* map $\pi : Y' \rightarrow X'$ such that $\mu \circ \pi^{-1} = \nu$ and $\pi \circ S = T \circ \pi$ on Y' . We call π the *factor map*.

If (X, \mathcal{B}, μ, T) is a factor of (Y, \mathcal{C}, ν, S) , then it is customary to call (Y, \mathcal{C}, ν, S) an *extension* of (X, \mathcal{B}, μ, T) .

There are three principle examples:

1. Any measure theoretic isomorphism between two mpt is a factor map between them.
2. A skew product $T_f : X \times Y \rightarrow X \times Y$ is an extension of its base $T : X \rightarrow X$. The factor map is $\pi : X \times Y \rightarrow X, \pi(x, y) = x$.
3. Suppose (X, \mathcal{B}, μ, T) is an mpt and T is measurable w.r.t. a smaller σ -algebra $\mathcal{C} \subset \mathcal{B}$, then (X, \mathcal{C}, μ, T) is a factor of (X, \mathcal{B}, μ, T) . The factor map is the identity.

We dwell a bit more on the third example. In probability theory, σ -algebras model *information*: a set E is “measurable”, if we can answer the question “is ω in E ?” using the information available to use. For example, if a real number $x \in \mathbb{R}$ is unknown, but we can “measure” $|x|$, then the information we have on x is modeled by the σ -algebra $\{E \subset \mathbb{R} : E = -E\}$, because we can determined whether $x \in E$ only for symmetric sets E . By decreasing the σ -algebra, we are forgetting some information. For example if instead of knowing $|x|$, we only know whether $0 \leq |x| \leq 1$ or not, then our σ -algebra is the finite σ -algebra generated by $[-1, 1]$.

Here is a typical example. Suppose we have a dynamical system (X, \mathcal{B}, μ, T) , and we cannot “measure” x , but we can “measure” $f(x)$ for some measurable

$f : X \rightarrow \mathbb{R}$. Then the information we have by observing the dynamical system is encoded in the sub sigma algebra

$$\mathcal{C} := \sigma(f \circ T^n : n \geq 0)$$

defined to be the smallest σ -algebra with respect to which $f \circ T^n$ are all measurable.⁵ The dynamical properties we feel in this case are those of the factor (X, \mathcal{C}, μ, T) and not of the (X, \mathcal{B}, μ, T) . For example, it could be the case that $\mu(E \cap T^{-n}F) \rightarrow \mu(E)\mu(F)$ for all $E, F \in \mathcal{C}$ but not for all $E, F \in \mathcal{B}$ — and then we will observe “mixing” simply because our information is not sufficient to observe the non-mixing in the system.

1.6.3 The natural extension

Every ppt is a factor of an *invertible* ppt. Moreover, there is a extension of this type which is “minimal” in the sense that it is a factor of any other invertible extension. This extension is unique up to isomorphism, and is called the *natural extension*. We describe the construction.

Definition 1.15. Suppose (X, \mathcal{B}, μ, T) is a ppt, and that $T(X) = X$.⁶ The *natural extension* of (X, \mathcal{B}, μ, T) is the system $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{T})$, where

1. $\tilde{X} := \{\underline{x} = (\dots, x_{-1}, x_0, x_1, x_2, \dots) : x_i \in X, T(x_i) = x_{i+1} \text{ for all } i\}$;
2. $\tilde{\mathcal{B}} := \{\pi^{-1}(E) : E \in \mathcal{B}\}$, where $\pi(\underline{x}) = x_0$;
3. $\tilde{\mu}$ is the measure $\tilde{\mu}(\pi^{-1}E) := \mu(E)$;
4. \tilde{T} is the left shift.

Theorem 1.5. *The natural extension of (X, \mathcal{B}, μ, T) is an invertible extension of (X, \mathcal{B}, μ, T) , and is the factor of any other invertible extension of (X, \mathcal{B}, μ, T) . Ergodicity and mixing are preserved under natural extensions.*

Proof. It is clear that the natural extension is an invertible extension, and that $\pi \circ \tilde{T} = T \circ \pi$. Since $TX = X$, every point has a preimage, and so π is onto. Thus $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{T})$ is an invertible extension of (X, \mathcal{B}, μ, T) .

Suppose (Y, \mathcal{C}, ν, S) is another invertible extension, and let $\pi_Y : Y \rightarrow X$ be the factor map (defined a.e. on Y). We show that (Y, \mathcal{C}, ν, S) extends $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{T})$.

Let $(\tilde{Y}, \tilde{\mathcal{C}}, \tilde{\nu}, \tilde{S})$ be the natural extension of (Y, \mathcal{C}, ν, S) . It is isomorphic to (Y, \mathcal{C}, ν, S) , with the isomorphism given by $\vartheta(y) = (y_k)_{k \in \mathbb{Z}}$, $y_k := S^k(y)$. Thus it is enough to show that $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{T})$ is a factor of $(\tilde{Y}, \tilde{\mathcal{C}}, \tilde{\nu}, \tilde{S})$. Here is the factor map: $\theta : (y_k)_{k \in \mathbb{Z}} \mapsto (\pi_Y(y_k))_{k \in \mathbb{Z}}$.

⁵ Such a σ -algebra exists: take the intersection of all sub- σ -algebras which make $f \circ T^n$ all measurable, and note that this intersection is not empty because it contains \mathcal{B} .

⁶ Actually all that we need is that $T^n X$ is measurable for all n . In general, the measurable (or even continuous) forward image of a measurable set is not measurable.

We show that ergodicity is preserved under natural extensions: Suppose \tilde{E} is a \tilde{T} -invariant $\tilde{\mathcal{B}}$ -measurable set. By the definition of $\tilde{\mathcal{B}}$, we can write $\tilde{E} = \pi^{-1}(E)$, $E \in \mathcal{B}$. We saw above that π is onto, therefore $E = \pi\tilde{E}$. We have

$$\begin{aligned} T^{-1}E &= \{x : T(x) \in E\} = \pi\{(\dots, x_{-1}, x_0, x_1, \dots) \in \tilde{X} : x_1 \in E\} \\ &= \pi\{(\dots, x_{-1}, x_0, x_1, \dots) \in \tilde{X} : \tilde{T}(\dots, x_{-1}, x_0, x_1, \dots) \in \pi^{-1}E\} \\ &= \pi\tilde{T}^{-1}(\pi^{-1}E) = \pi\tilde{T}^{-1}\tilde{E} = \pi\tilde{E} = E. \end{aligned}$$

Thus E has full measure, whence $\tilde{E} = \pi^{-1}E$ has full measure. The proof that mixing is preserved under natural extensions is left to the reader. \square

1.6.4 Induced transformations

Suppose (X, \mathcal{B}, μ, T) is a probability preserving transformation, and let $A \in \mathcal{B}$ be a set of positive measure. By Poincaré's Recurrence Theorem, for a.e. $x \in A$ there is some $n \geq 1$ such that $T^n(x) \in A$. Define

$$\varphi_A(x) := \min\{n \geq 1 : T^n x \in A\},$$

with the minimum of the empty set being understood as infinity. Note that $\varphi_A < \infty$ a.e. on A .

Definition 1.16. The *induced transformation* on A is $(A_0, \mathcal{B}(A), \mu_A, T_A)$, where $A := \{x \in A : \varphi_A(x) < \infty\}$, $\mathcal{B}(A) := \{E \cap A : E \in \mathcal{B}\}$, μ_A is the measure $\mu_A(E) := \mu(E|A) = \mu(E \cap A)/\mu(A)$, and $T_A : A_0 \rightarrow A_0$ is $T_A(x) = T^{\varphi_A(x)}(x)$.

Theorem 1.6. Suppose (X, \mathcal{B}, μ, T) is a ppt, and $A \in \mathcal{B}$ has positive finite measure.

1. $\mu_A \circ T_A^{-1} = \mu_A$;
2. if T is ergodic, then T_A is ergodic (but the mixing of T does not imply the mixing of T_A);
3. **Kac Formula:** If μ is ergodic, then $\int f d\mu = \int_A \sum_{k=0}^{\varphi_A-1} f \circ T^k d\mu$ for every $f \in L^1(X)$. In particular $\int_A \varphi_A d\mu_A = 1/\mu(A)$.

Proof. Given $E \subset A$ measurable,

$$\begin{aligned} \mu(E) &= \underbrace{\mu(T^{-1}E \cap A)}_{\mu(T_A^{-1}E \cap [\varphi_A=1])} + \mu(T^{-1}E \cap A^c) \\ &= \underbrace{\mu(T^{-1}E \cap A)}_{\mu(T_A^{-1}E \cap [\varphi_A=1])} + \underbrace{\mu(T^{-2}E \cap T^{-1}A^c \cap A)}_{\mu(T_A^{-1}E \cap [\varphi_A=2])} + \mu(T^{-2}E \cap T^{-1}A^c \cap A^c) \\ &= \dots = \sum_{j=1}^N \mu(T_A^{-1}E \cap [\varphi_A = j]) + \mu(T^{-N}E \cap \bigcap_{j=0}^{N-1} T^{-j}A^c). \end{aligned}$$

Passing to the limit as $N \rightarrow \infty$, we see that $\mu(E) \geq \mu_A(T_A^{-1}E)$. Working with $A \setminus E$, and using the assumption that $\mu(X) < \infty$, we get that $\mu(A) - \mu(E) \leq \mu(A) - \mu(T_A^{-1}E)$ whence $\mu(E) = \mu(T_A^{-1}E)$. Since μ_A is proportional to μ on $\mathcal{B}(A)$, we get $\mu_A = \mu_A \circ T_A^{-1}$.

We now assume that T is ergodic, and prove that T_A is ergodic. Since T is ergodic, and the set

$$\{x : T^n(x) \in A \text{ for infinitely many } n \geq 0\}$$

is a T -invariant set of non-zero measure (bounded below by $\mu(A)$), a.e. $x \in X$ has some $n \geq 0$ s.t. $T^n(x) \in A$. Thus $r_A(x) := \min\{n \geq 0 : T^n x \in A\} < \infty$ a.e. in X (not just a.e. in A).

Suppose $f : A_0 \rightarrow \mathbb{R}$ is a T_A -invariant L^2 -function. Define

$$F(x) := f(T^{r_A(x)}x).$$

This makes sense a.e. in X , because $r_A < \infty$ almost everywhere. This function is T -invariant, because $T^{r_A(Tx)}(Tx) \in \{T^{r_A(x)}(x), T_A(T^{r_A(x)}x), T_A^2(T^{r_A(x)}x)\}$, and $f \circ T_A = f$ almost everywhere. Since T is ergodic, F is constant. Since $F|_A = f$, f is constant. Thus the ergodicity of T implies the ergodicity of T_A .

Here is an example showing that the mixing of T does not imply the mixing of T_A . Let Σ^+ be a SFT with states $\{a, 1, 2, b\}$ and allowed transitions

$$a \rightarrow 1; 1 \rightarrow 1, b; b \rightarrow 2; 2 \rightarrow a.$$

Let $A = \{x : x_0 = a, b\}$. Any shift invariant Markov chain measure μ on Σ^+ is mixing, because Σ^+ is irreducible and aperiodic ($1 \rightarrow 1$). But T_A is not mixing, because $T_A[a] = [b]$ and $T_A[b] = [a]$, so $[a] \cap T_A^{-n}[a] = \emptyset$ for all n odd.

Next we prove the Kac formula. Suppose first that $f \in L^\infty(X, \mathcal{B}, \mu)$ and $f \geq 0$.

$$\begin{aligned} \int f d\mu &= \int_A f d\mu + \int f \cdot 1_{A^c} d\mu = \int_A f d\mu + \int f \circ T \cdot 1_{T^{-1}A^c} d\mu \\ &= \int_A f d\mu + \int f \circ T \cdot 1_{T^{-1}A^c \cap A} d\mu + \int f \circ T \cdot 1_{T^{-1}A^c \cap A^c} d\mu \\ &= \int_A f d\mu + \int_A f \circ T \cdot 1_{[\varphi_A > 1]} d\mu + \int f \circ T^2 \cdot 1_{T^{-2}A^c \cap T^{-1}A^c} d\mu \\ &= \dots = \int_A \sum_{j=0}^{N-1} f \circ T^j \cdot 1_{[\varphi_A > j]} d\mu + \int f \circ T^N \cdot 1_{\bigcap_{j=1}^N T^{-j}A^c} d\mu. \end{aligned}$$

The first term tends, as $N \rightarrow \infty$, to

$$\int_A \sum_{j=0}^{\infty} f \circ T^j \sum_{i=j+1}^{\infty} 1_{[\varphi_A = i]} d\mu \equiv \int_A \sum_{j=0}^{\varphi_A - 1} f \circ T^j d\mu.$$

The second term is bounded by $\|f\|_\infty \mu\{x : T^j(x) \notin A \text{ for all } k \leq N\}$. This bound tends to zero, because $\mu\{x : T^j(x) \notin A \text{ for all } k\} = 0$ because T is ergodic and recurrent (fill in the details). This proves the Kac formula for all L^∞ functions.

Every non-negative L^1 -function is the increasing limit of L^∞ functions. By the monotone convergence theorem, the Kac formula must hold for all non-negative L^1 -function.

Every L^1 -function is the difference of two non-negative L^1 -functions ($f = f \cdot 1_{[f>0]} - |f| \cdot 1_{[f<0]}$). It follows that the Kac formula holds for all absolutely integrable functions. \square

1.6.5 Suspensions and Kakutani skyscrapers

The operation of inducing can be “inverted”, as follows. Let (X, \mathcal{B}, μ, T) be a ppt, and $r : X \rightarrow \mathbb{N}$ an integrable measurable function.

Definition 1.17. The *Kakutani skyscraper* with base (X, \mathcal{B}, μ, T) and height function r is the system $(X_r, \mathcal{B}(X_r), \nu, S)$, where

1. $X_r := \{(x, n) : x \in X, 0 \leq n \leq r(x) - 1\}$;
2. $\mathcal{B}(X_r) = \{E \in \mathcal{B}(X) \otimes \mathcal{B}(\mathbb{N}) : E \subseteq X_r\}$, where $\mathcal{B}(\mathbb{N}) = 2^{\mathbb{N}}$;
3. ν is the unique measure such that $\nu(B \times \{k\}) = \mu(B) / \int r d\mu$;
4. S is defined by $S(x, n) = (x, n+1)$, when $n < r(x) - 1$, and $S(x, n) = (Tx, 0)$, when $n = r(x) - 1$.

(Check that this is a ppt.)

We think of X_r as a skyscraper made of stories $\{(x, k) : r(x) > k\}$; the orbits of S climb up the skyscraper until they reach the top floor possible, and then move to the ground floor according to T .

If we induce a Kakutani skyscraper on $\{(x, 0) : x \in X\}$, we get a system which is isomorphic to (X, \mathcal{B}, μ, T) .

Proposition 1.11. *A Kakutani skyscraper over an ergodic base is ergodic, but there are non-mixing skyscrapers over mixing bases.*

The proof is left as an exercise.

There is a straightforward important continuous-time version of this construction: Suppose (X, \mathcal{B}, μ, T) is a ppt, and $r : X \rightarrow \mathbb{R}^+$ is a measurable function such that $\inf r > 0$.

Definition 1.18. The *suspension semi-flow* with base (X, \mathcal{B}, μ, T) and height function r is the semi-flow $(X_r, \mathcal{B}(X_r), \nu, T_s)$, where

1. $X_r := \{(x, t) \in X \times \mathbb{R} : 0 \leq t < r(x)\}$;
2. $\mathcal{B}(X_r) = \{E \in \mathcal{B}(X) \otimes \mathcal{B}(\mathbb{R}) : E \subseteq X_r\}$;
3. ν is the measure such that $\int_{X_r} f d\nu = \int_X \int_0^{r(x)} f(x, t) dt d\mu(x) / \int_X r d\mu$;
4. $T_s(x, t) = (T^n x, t + s - \sum_{k=0}^{n-1} r(T^k x))$, where n is s.t. $0 \leq t + s - \sum_{k=0}^{n-1} r(T^k x) < r(T^n x)$.

(Check that this is a measure preserving semi-flow.)

Suspension flows appear in applications in the following way. Imagine a flow T_t on a manifold X . It is often possible to find a curve $S \subset X$ such that (almost) every orbit of the flow intersects S transversally infinitely many times. Such a curve is called a *Poincaré section*. If it exists, then one can define a map $T_S : S \rightarrow S$ which maps $x \in S$ into $T_t x$ with $t := \min\{s > 0 : T_s(x) \in S\}$. This map is called the *Section map*. The flow itself is isomorphic to a suspension flow over its Poincaré section.

Problems

1.1. Proof of Liouville's theorem in section 1.1

- (a) Write $\underline{x} := (q, p)$ and $\underline{y} := T_t(q, p)$. Use Hamilton's equations to show that the Jacobian matrix of $\underline{y} = \underline{y}(\underline{x})$ satisfies $\frac{\partial \underline{y}}{\partial \underline{x}} = I + tA + O(t^2)$ as $t \rightarrow 0$, where $\text{tr}(A) = 0$.
 (b) Show that for every matrix A , $\det(I + tA + O(t^2)) = 1 + t\text{tr}(A) + O(t^2)$ as $t \rightarrow 0$.
 (c) Prove that the Jacobian of T_t is equal to one for all t . Deduce Liouville's theorem.

1.2. The completion of a measure space. Suppose (X, \mathcal{B}, μ) is a measure space. A set $N \subset X$ is called a *null set*, if there is a measurable set $E \supseteq N$ such that $\mu(E) = 0$. A measure space is called *complete*, if every null set is measurable. Every measure space can be completed, and this exercise shows how to do this.

- (a) Let \mathcal{B}_0 denote the collection of all sets of the form $E \cup N$ where $E \in \mathcal{B}$ and N is a null set. Show that \mathcal{B}_0 is a σ -algebra.
 (b) Show that μ has a unique extension to a σ -additive measure on \mathcal{B}_0 .

1.3. Prove Poincaré's Recurrence Theorem for a general probability preserving transformation (theorem 1.1).

1.4. Fill in the details in the proof above that the Markov chain measure corresponding to a stationary probability vector and a stochastic matrix exists, and is shift-invariant measure.

1.5. Suppose Σ_A^+ is a SFT with transition matrix A . Write $A^n = (t_{ab}^{(n)})$. Prove that $t_{ab}^{(n)}$ is the number of paths of length n starting at a and ending at b . In particular: $a \xrightarrow{n} b \Leftrightarrow t_{ab}^{(n)} > 0$.

1.6. The Perron–Frobenius Theorem⁷: Suppose $A = (a_{ij})$ is a matrix all of whose entries are non-negative, and let $B := (b_{ij})$ be the matrix $b_{ij} = 1$ if $a_{ij} > 0$ and $b_{ij} = 0$ if $a_{ij} = 0$. Assume that B is irreducible, then A has a positive eigenvalue λ with the following properties:

- (i) There are positive vectors \underline{r} and $\underline{\ell}$ s.t. $\underline{\ell}A = \lambda\underline{\ell}$, $A\underline{r} = \lambda\underline{r}$.
 (ii) The eigenvalue λ is simple.

⁷ Perron first proved this in the aperiodic case. Frobenius later treated the periodic irreducible case.

- (iii) The spectrum of $\lambda^{-1}A$ consists of 1, several (possibly zero) roots of unity, and a finite subset of the open unit disc. In this case the limit $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \lambda^{-k} A^k$ exists.
- (iv) If B is irreducible and aperiodic, then the spectrum of $\lambda^{-1}A$ consists of 1 and a finite subset of the open unit disc. In this case the limit $\lim_{n \rightarrow \infty} \lambda^{-n} A^n$ exists.

1. Prove the Perron–Frobenius theorem in case A is stochastic, first in the aperiodic case, then in the general case.
2. Now consider the case of a non-negative matrix:
 - a. Use a fixed point theorem to show that $\lambda, \underline{\ell}, \underline{r}$ exist;
 - b. Set $\underline{1} := (1, \dots, 1)$ and let V be the diagonal matrix such that $V\underline{1} = \underline{r}$. Prove that $\lambda^{-1}V^{-1}AV$ is stochastic.
 - c. Prove the Perron-Frobenius theorem.

1.7. Suppose $P = (p_{ab}^{(n)})_{S \times S}$ is an irreducible aperiodic stochastic matrix. Use the spectral description of P obtained in problem 1.6 to show that $p_{ab}^{(n)} \rightarrow p_b$ exponentially fast.

1.8. Show that the product of n irrational rotations $R_{\alpha_1}, \dots, R_{\alpha_n}$ is ergodic iff $(\alpha_1, \dots, \alpha_n)$ are independent over the irrationals.

1.9. Suppose $g^t : X \rightarrow X$ is a measure preserving flow. The *time one map* of the flow is the measure preserving map $g^1 : X \rightarrow X$. Give an example of an ergodic flow whose time one map is not ergodic.

1.10. The adding machine

Let $X = \{0, 1\}^{\mathbb{N}}$ equipped with the σ -algebra \mathcal{B} generated by the cylinders, and the Bernoulli $(\frac{1}{2}, \frac{1}{2})$ -measure μ . The *adding machine* is the ppt (X, \mathcal{B}, μ, T) defined by the rule $T(1^n 0^*) = (0^n 1^*)$, $T(1^\infty) = 0^\infty$. Prove that the adding machine is invertible and probability preserving. Show that $T(\underline{x}) = \underline{x} \oplus (10^\infty)$ where \oplus is “addition with carry to the right”.

1.11. Prove proposition 1.11.

1.12. Show that a ppt (X, \mathcal{B}, μ, T) is mixing whenever $\mu(A \cap T^{-n}A) \xrightarrow{n \rightarrow \infty} \mu(A)^2$ for all $A \in \mathcal{B}$. *Guidance:*

1. $\int 1_A f \circ T^n d\mu \xrightarrow{n \rightarrow \infty} \mu(A) \int f d\mu$ for all $f \in \overline{\text{span}}_{L^2} \{1, 1_A \circ T, 1_A \circ T^2, \dots\}$.
2. $\int 1_A f \circ T^n d\mu \xrightarrow{n \rightarrow \infty} \mu(A) \int f d\mu$ for all $f \in L^2$.
3. $\int g f \circ T^n d\mu \xrightarrow{n \rightarrow \infty} \int g d\mu \int f d\mu$ for all $f, g \in L^2$.

1.13. Show that a Kakutani skyscraper over an invertible transformation is invertible, and find a formula for its inverse.

1.14. Conservativity

Let (X, \mathcal{B}, μ, T) be a measure preserving transformation on an *infinite* σ -finite,

measure space.⁸ A set $W \in \mathcal{B}$ is called *wandering*, if $\{T^{-n}W : n \geq 0\}$ are pairwise disjoint. A mpt is called *conservative*, if every wandering set has measure zero.

1. Show that any ppt is conservative. Give an example of a non-conservative mpt on a σ -finite infinite measure space.
2. Show that Poincaré's recurrence theorem extends to conservative mpt.
3. Suppose (X, \mathcal{B}, μ, T) is a conservative ergodic mpt, and let A be a set of finite positive measure. Show that the induced transformation $T_A : A \rightarrow A$ is well-defined a.e. on A , and is an ergodic ppt.
4. Prove Kac formula for conservative ergodic transformations under the previous set of assumptions.

Notes for chapter 1

The standard references for measure theory are [7] and [4]. The standard references for ergodic theory of probability preserving transformations are [1] and [3]. The standard reference on ergodic theory of mpt on infinite measure spaces is [1]. Our proof of the Perron-Frobenius theorem is taken from [3]. Kac's formula has very simple proof when T is invertible. The proof we use has the merit that it works for non-invertible transformations, and that it extends to the (conservative) infinite measure setting. It is taken from [2]. The ergodicity of the geodesic flow was first proved by E. Hopf. The short proof we use was found much later by Gelfand, and is reproduced in [2].

References

1. Aaronson, J.: *Introduction to infinite ergodic theory*, Mathematical Surveys and Monographs 50, AMS, 1997. xii+284pp.
2. Bekka, M.B. and Mayer, M.: *Ergodic theory and topological dynamics of group actions on homogeneous spaces*. London Mathematical Society LNM 269, 2000. x+200pp.
3. Brin, M. and Stuck, G.: *Introduction to dynamical systems*. Cambridge University Press, Cambridge, 2002. xii+240 pp.
4. Halmos, P. R.: *Measure Theory*. D. Van Nostrand Company, Inc., New York, N. Y., 1950. xi+304 pp.
5. Krengel, U.: *Ergodic theorems*. de Gruyter Studies in Mathematics 6 1985. viii+357pp.
6. Petersen, K.: *Ergodic theory*. Corrected reprint of the 1983 original. Cambridge Studies in Advanced Mathematics, 2. Cambridge University Press, Cambridge, 1989. xii+329 pp.
7. Royden, H. L.: *Real analysis*. Third edition. Macmillan Publishing Company, New York, 1988. xx+444 pp.
8. Walters, P.: *An introduction to ergodic theory*. Graduate Texts in Mathematics, 79. Springer-Verlag, New York-Berlin, 1982. ix+250 pp.

⁸ A measure space is called σ -finite, if its sample space is the countable union of finite measure sets.

Chapter 2

Ergodic Theorems

2.1 The Mean Ergodic Theorem

Theorem 2.1 (von Neumann's Mean Ergodic Theorem). *Suppose (X, \mathcal{B}, μ, T) is a ppt. If $f \in L^2$, then $\frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \xrightarrow[n \rightarrow \infty]{L^2} \bar{f}$ where $\bar{f} \in L^2$ is invariant. If T is ergodic, then $\bar{f} = \int f d\mu$.*

Proof. Observe that since T is measure preserving, $\|f \circ T\|_2 = \|f\|_2$ for all $f \in L^2$ (prove this, first for indicator functions, then for all L^2 -functions).

Suppose $f = g - g \circ T$ where $g \in L^2$ (in this case we say that f is a *coboundary* with *transfer function* $g \in L^2$), then it is obvious that

$$\left\| \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \right\|_2 = \frac{1}{N} \|g \circ T^N - g\|_2 \leq 2\|g\|_2/N \xrightarrow{N \rightarrow \infty} 0.$$

Thus the theorem holds for all elements of $\mathcal{C} := \{g - g \circ T : g \in L^2\}$.

We claim that the theorem holds for all elements of $\overline{\mathcal{C}}$ (L^2 -closure). Suppose $f \in \overline{\mathcal{C}}$, then for every $\varepsilon > 0$, there is an $F \in \mathcal{C}$ s.t. $\|f - F\|_2 < \varepsilon$. Choose N_0 such that for every $N > N_0$, $\left\| \frac{1}{N} \sum_{k=0}^{N-1} F \circ T^k \right\|_2 < \varepsilon$, then for all $N > N_0$

$$\begin{aligned} \left\| \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \right\|_2 &\leq \left\| \frac{1}{N} \sum_{k=0}^{N-1} (f - F) \circ T^k \right\|_2 + \left\| \frac{1}{N} \sum_{k=0}^{N-1} F \circ T^k \right\|_2 \\ &\leq \frac{1}{N} \sum_{k=0}^{N-1} \|(f - F) \circ T^k\|_2 + \varepsilon < 2\varepsilon. \end{aligned}$$

This shows that $\left\| \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \right\|_2 \xrightarrow{N \rightarrow \infty} 0$.

Next we claim that $\overline{\mathcal{C}}^\perp = \{\text{invariant functions}\}$. Suppose $f \perp \overline{\mathcal{C}}$, then

$$\begin{aligned}\|f - f \circ T\|_2^2 &= \langle f - f \circ T, f - f \circ T \rangle = \|f\|_2^2 - 2\langle f, f \circ T \rangle + \|f \circ T\|_2^2 \\ &= 2\|f\|_2^2 - 2\langle f, f - (f - f \circ T) \rangle = 2\|f\|_2^2 - 2\|f\|_2^2 = 0.\end{aligned}$$

Thus $f = f \circ T$ a.e., and we have proved that $L^2 = \overline{\mathcal{C}} \oplus \{\text{invariant functions}\}$.

We saw above that the MET holds for all elements of $\overline{\mathcal{C}}$, and it is obvious for all invariant functions. Therefore the MET holds for all L^2 -functions.

The proof shows that the limit \bar{f} is the projection of f on the space of invariant functions, and is thus invariant. If T is ergodic, then it is constant. This constant is equal to $\int f d\mu$, because if $f_n \rightarrow f$ in L^2 on a finite measure space, then $f_n \rightarrow f$ in L^1 (Cauchy–Schwarz), and so $\int f_n \rightarrow \int f$.

Remark 1. The proof shows that the limit \bar{f} is the projection of f on the space of invariant functions.

Remark 2. The proof only uses the fact that $Uf = f \circ T$ is an isometry of L^2 . In fact it works for all contractions, see problem 2.1.

Remark 3. If $f_n \xrightarrow[n \rightarrow \infty]{L^2} f$, then $\langle f_n, g \rangle \xrightarrow[n \rightarrow \infty]{} \langle f, g \rangle$ for all $g \in L^2$. Specializing to the case $f_n = \frac{1}{n} \sum_{k=0}^{n-1} 1_B \circ T^k$, $g = 1_A$ we obtain the following corollary of the MET:

Corollary 2.1. A ppt (X, \mathcal{B}, μ, T) is ergodic iff for all $A, B \in \mathcal{B}$,

$$\frac{1}{n} \sum_{k=0}^{n-1} \mu(A \cap T^{-k}B) \xrightarrow[n \rightarrow \infty]{} \mu(A)\mu(B).$$

So ergodicity is mixing “on the average”.

2.2 The Pointwise Ergodic Theorem

Theorem 2.2 (Birkhoff’s Pointwise Ergodic Theorem). Let (X, \mathcal{B}, μ, T) be a ppt.

If $f \in L^1$, then the following limit exists μ -almost everywhere: $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k$.

The limit is a T -invariant function. If T is ergodic, it is equal a.e. to $\int f d\mu$.

Proof. We begin by proving convergence when $f = 1_B$, $B \in \mathcal{B}$. Define

$$A_n(x) := \frac{1}{N} \sum_{k=0}^{N-1} 1_B(T^k x)$$

$$\bar{A}(x) := \limsup_{n \rightarrow \infty} A_n(x)$$

$$\underline{A}(x) := \liminf_{n \rightarrow \infty} A_n(x)$$

Fix ε and $M > 0$ (to be determined later), and set

$$\tau(x) := \min\{n > 0 : A_n(x) > \bar{A}(x) - \varepsilon\}, \quad B' := B \cup [\tau > M].$$

For a given N , we ‘color’ the time interval $0, 1, 2, \dots, N-1$ as follows:

- If $\tau(T^0x) > M$, color 0 red; If $\tau(T^0x) \leq M$ color the next $\tau(x)$ times blue, and move to the first uncolored k
- If $\tau(T^kx) > M$, color k red; Otherwise color the next $\tau(T^kx)$ times blue, and move to the first uncolored k

Continue in this way until all times colored, or until $\tau(T^kx) > N - k$.

This partitions $\{0, 1, \dots, N-1\}$ into red segments, and (possibly consecutive) blue segments of length $\leq M$, and one last segment of length $\leq M$. Note:

- If k is red, then $T^kx \in B'$, so $\sum_{\text{red } k's} 1_{B'}(T^kx) \geq \text{number of red } k's$
- The average of 1_B on each blue segment of length $\tau(T^kx)$ is larger than $\bar{A}(T^kx) - \varepsilon = \bar{A}(x) - \varepsilon$. It follows that

$$\sum_{\text{blue segment}} 1_{B'}(T^kx) \geq \text{length of blue segment} \times (\bar{A} - \varepsilon).$$

$$\text{Summing over all segments: } \sum_{\text{blue } k's} 1_{B'}(T^kx) \geq (\bar{A} - \varepsilon) \times \text{number of blues.}$$

Since $\bar{A} - \varepsilon < 1$, we get

$$\begin{aligned} \sum_{k=0}^{N-1} 1_{B'}(T^kx) &\geq \text{number of colored times} \times (\bar{A} - \varepsilon) \\ &\geq (N - M)(\bar{A} - \varepsilon). \end{aligned}$$

Integrating, we get $N\mu(B') \geq (N - M)(\int \bar{A}d\mu - \varepsilon)$ whence

$$\mu(B') \geq \bar{A} - \varepsilon$$

or $\mu(B) \geq \int \bar{A}d\mu - (\varepsilon + \mu[\tau > M])$. Now choose M to be so small that $\mu[\tau > M] < \varepsilon$. We get $\mu(B) \geq \bar{A} - 2\varepsilon$, or $\int \bar{A}d\mu \leq \mu(B) + 2\varepsilon$.

Applying this argument to B^c , we get $\int \underline{A}d\mu \leq \mu(B) + 2\varepsilon$. It follows that

$$0 \leq \int (\bar{A} - \underline{A})d\mu \leq 4\varepsilon \xrightarrow{\varepsilon \rightarrow 0^+} 0.$$

Thus $\int (\bar{A} - \underline{A})d\mu = 0$, whence $\bar{A} = \underline{A}$ μ -almost everywhere. This proves almost everywhere convergence for $f = 1_B$

It automatically follows that the pointwise ergodic theorem holds for all simple functions $\sum_{i=1}^n \alpha_i 1_{B_i}$, $B_i \in \mathcal{B}$.

Any L^∞ function is the uniform limit of simple functions. Therefore the pointwise ergodic theorem holds for all bounded measurable functions.

Now suppose $f \in L^1$. Find some $\varphi \in L^\infty$ such that $\|f - \varphi\|_1 < \varepsilon$. Let $A_n(g)$, $\bar{A}(g)$, $\underline{A}(g)$ be the corresponding objects for a function g . We have

$$\bar{A}(f) \leq \bar{A}(\varphi) + \bar{A}(f - \varphi) \equiv \bar{A}(\varphi) + \underline{A}(\varphi - f)$$

Integrating, we have

$$\begin{aligned} \int \bar{A} d\mu &\leq \int \bar{A}(\varphi) d\mu + \int \liminf_{n \rightarrow \infty} A_n(\varphi - f) d\mu \\ &\leq \int \bar{A}(\varphi) d\mu + \liminf_{n \rightarrow \infty} \int A_n(\varphi - f) d\mu \leq \int \bar{A}(\varphi) d\mu + \|\varphi - f\|_1. \end{aligned}$$

Applying this argument for $-f, -\varphi$ gives

$$\int \bar{A} d\mu \geq \int \underline{A}(\varphi) d\mu - \|\varphi - f\|_1.$$

Since $\bar{A}(\varphi) = \underline{A}(\varphi)$ (because it is bounded), we get

$$\int (\bar{A}(f) - \underline{A}(f)) d\mu < 2\|\varphi - f\|_1 < 2\varepsilon \xrightarrow{\varepsilon \rightarrow 0^+} 0.$$

Again, $\bar{A}(f) = \underline{A}(f)$ almost everywhere and we obtain pointwise convergence. \square

Remark: If $f \in L^1$ then the Birkhoff averages of f converge in L^1 . (For bounded functions this is the bounded convergence theorem; for functions with small L^1 norm, the L^1 -norm of the averages remains small; any L^1 function is the sum of a bounded functions and an L^1 function with small norm.)

2.3 The non-ergodic case

The almost sure limit in the pointwise ergodic theorem is clear when the map is ergodic: $\frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \xrightarrow{N \rightarrow \infty} \int f d\mu$. In this section we ask what is the limit in the non-ergodic case.

If f belongs to L^2 , the limit is the projection of f on the space of invariant functions, because of the Mean Ergodic Theorem and the fact that every sequence of functions which converges in L^2 has a subsequence which converges almost everywhere to the same limit.¹ But if $f \in L^1$ we cannot speak of projections. The right notion in this case is that of the *conditional expectation*.

¹ Proof: Suppose $f_n \xrightarrow{L^2} f$. Pick a subsequence n_k s.t. $\|f_{n_k} - f\|_2 < 2^{-k}$. Then $\sum_{k \geq 1} \|f_{n_k} - f\|_2 < \infty$. This means that $\|\sum |f_{n_k} - f|\|_2 < \infty$, whence $\sum (f_{n_k} - f)$ converges absolutely almost surely. It follows that $f_{n_k} - f \rightarrow 0$ a.e.

2.3.1 Conditional expectations and the limit in the ergodic theorem

Let (X, \mathcal{B}, μ) be a probability space. Let $\mathcal{F} \subset \mathcal{B}$ be a σ -algebra. We think of $F \in \mathcal{F}$ as of the collection of all sets F for which we have sufficient information to answer the question “is $x \in F$?”. The functions we have sufficient information to calculate are exactly the \mathcal{F} -measurable functions, as can be seen from the formula

$$f(x) := \inf\{t : x \in [f < t]\}.$$

Suppose g is *not* \mathcal{F} -measurable. What is the ‘best guess’ for $g(x)$ given the information \mathcal{F} ?

Had g been in L^2 , then the “closest” \mathcal{F} -measurable function (in the L^2 -sense) is the projection of g on $L^2(X, \mathcal{F}, \mu)$. The defining property of the projection Pg of g is $\langle Pg, h \rangle = \langle g, h \rangle$ for all $h \in L^2(X, \mathcal{F}, \mu)$. The following definition mimics this case when g is not necessarily in L^2 :

Definition 2.1. The *conditional expectation* of $f \in L^1(X, \mathcal{B}, \mu)$ given \mathcal{F} is the unique $L^1(X, \mathcal{F}, \mu)$ -element $\mathbb{E}(f|\mathcal{F})$ which is

1. $\mathbb{E}(f|\mathcal{F})$ is \mathcal{F} -measurable;
2. $\forall \varphi \in L^\infty \mathcal{F}$ -measurable, $\int \varphi \mathbb{E}(f|\mathcal{F}) d\mu = \int \varphi f d\mu$.

Note: $\mathbb{E}(f|\mathcal{F})$ is only determined almost everywhere.

Proposition 2.1. *The conditional expectation exists for every L^1 element, and is unique up sets of measure zero.*

Proof. Consider the measures $\nu_f := f d\mu|_{\mathcal{F}}$ and $\mu|_{\mathcal{F}}$ on (X, \mathcal{F}) . Then $\nu_f \ll \mu$. The function $\mathbb{E}(f|\mathcal{F}) := \frac{d\nu_f}{d\mu}$ (Radon-Nikodym derivative) is \mathcal{F} -measurable, and it is easy to check that it satisfies the conditions of the definition of the conditional expectation. The uniqueness of the conditional expectation is left as an exercise. \square

Proposition 2.2.

1. $f \mapsto \mathbb{E}(f|\mathcal{F})$ is linear, and a contraction in the L^1 -metric;
2. $f \geq 0 \Rightarrow \mathbb{E}(f|\mathcal{F}) \geq 0$ a.e.;
3. if φ is convex, then $\mathbb{E}(\varphi \circ f|\mathcal{F}) \geq \varphi(\mathbb{E}(f|\mathcal{F}))$;
4. if h is \mathcal{F} -measurable, then $\mathbb{E}(hf|\mathcal{F}) = h\mathbb{E}(f|\mathcal{F})$;
5. If $\mathcal{F}_1 \subset \mathcal{F}_2$, then $\mathbb{E}[\mathbb{E}(f|\mathcal{F}_1)|\mathcal{F}_2] = \mathbb{E}(f|\mathcal{F}_2)$.

We leave the proof as an exercise.

Theorem 2.3. *Let (X, \mathcal{B}, μ, T) be a p.p.t, and $f \in L^1(X, \mathcal{B}, \mu)$. Then*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k = \mathbb{E}(f|\mathcal{I}nv(T)),$$

where $\mathcal{I}nv(T) := \{E \in \mathcal{B} : E = T^{-1}E\}$. Alternatively, $\mathcal{I}nv(T)$ is the σ -algebra generated by all T -invariant functions.

Proof. Set $\bar{f} := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k$ on the set where the limit exists, and zero otherwise. Then \bar{f} is T -invariant. For every $\varphi \in L^\infty$ which is T -invariant,

$$\begin{aligned} \int \varphi \bar{f} d\mu &= \int \varphi \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k d\mu + O\left(\|\varphi\|_\infty \left\| \frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k - \bar{f} \right\|_1\right) = \\ &= \frac{1}{N} \sum_{k=0}^{N-1} \int \varphi \circ T^k f \circ T^k d\mu + o(1) \xrightarrow{N \rightarrow \infty} \int \varphi f d\mu, \end{aligned}$$

because the convergence in the ergodic theorem is also in L^1 . \square

2.3.2 Conditional probabilities

Recall that a *standard probability space* is a probability space (X, \mathcal{B}, μ) where X is a complete, metric, separable space, and \mathcal{B} is its Borel σ -algebra.

Theorem 2.4 (Existence of Conditional Probabilities). *Let μ be a Borel probability measure on a standard probability space (X, \mathcal{B}, μ) , and let $\mathcal{F} \subset \mathcal{B}$ be a σ -algebra. There exist Borel probability measures $\{\mu_x\}_{x \in X}$ s.t.:*

1. $x \mapsto \mu_x(E)$ is \mathcal{F} -measurable for every $E \in \mathcal{B}$;
2. if f is μ -integrable, then $x \mapsto \int f d\mu_x$ is integrable, and $\int f d\mu = \int \left(\int f d\mu_x \right) d\mu$;
3. if f is μ -integrable, then $\int f d\mu_x = \mathbb{E}(f|\mathcal{F})(x)$ for μ -a.e. x .

Definition 2.2. The measures μ_x are called the *conditional probabilities* of \mathcal{F} . Note that they are only determined almost everywhere.

Proof. By the isomorphism theorem for standard spaces (see appendix), there is no loss of generality in assuming that X is compact (indeed, we may take X to be a compact interval).

Fix a countable dense set $\{f_n\}_{n=0}^\infty$ in $C(X)$ s.t. $f_0 \equiv 1$. Let $\mathcal{A}_\mathbb{Q}$ be the algebra generated by these functions over \mathbb{Q} . It is still countable.

Choose for every $g \in \mathcal{A}_\mathbb{Q}$ an \mathcal{F} -measurable version $\bar{\mathbb{E}}(g|\mathcal{F})$ of $\mathbb{E}(g|\mathcal{F})$ (recall that $\mathbb{E}(g|\mathcal{F})$ is an L^1 -function, namely not a function at all but an equivalence class of functions). Consider the following collection of conditions:

1. $\forall \alpha, \beta \in \mathbb{Q}, g_1, g_2 \in \mathcal{A}_\mathbb{Q}, \bar{\mathbb{E}}(\alpha g_1 + \beta g_2|\mathcal{F})(x) = \alpha \bar{\mathbb{E}}(g_1|\mathcal{F})(x) + \beta \bar{\mathbb{E}}(g_2|\mathcal{F})(x)$
2. $\forall g \in \mathcal{A}_\mathbb{Q}, \min g \leq \bar{\mathbb{E}}(g|\mathcal{F})(x) \leq \max g$

This is countable collection of \mathcal{F} -measurable conditions, each of which holds with full μ -probability. Let X_0 be the set of x 's which satisfies all of them. This is an \mathcal{F} -measurable set of full measure.

We see that for each $x \in X_0$, $\varphi_x[g] := \bar{\mathbb{E}}(g|\mathcal{F})(x)$ is linear functional on $\mathcal{A}_\mathbb{Q}$, and $\|\varphi_x\| \leq 1$. It follows that φ_x extends uniquely to a positive bounded linear functional on $C(X)$. This is a measure μ_x .

Step 1. $\int (\int f d\mu_x) d\mu(x) = \int f d\mu$ for all $f \in C(X)$.

Proof. This is true for all $f \in \mathcal{A}_Q$ by definition, and extends to all $C(X)$ because \mathcal{A}_Q is dense in $C(X)$. (But for $f \in L^1$ it is not even clear that the statement makes sense, because μ_x could live on a set with zero μ -measure!)

Step 2. $x \mapsto \mu_x(E)$ is \mathcal{F} -measurable for all $E \in \mathcal{B}$.

Exercise: Prove this using the following steps

1. The indicator function of any open set is the pointwise limit of a sequence of continuous functions $0 \leq h_n \leq 1$, thus the step holds for open sets.
2. The collection of sets whose indicators are pointwise limits of a bounded sequence of continuous functions forms an algebra. The step holds for every set in this algebra.
3. The collection of sets for which step 1 holds is a monotone class which contains a generating algebra.

Step 3. If $f = g$ μ -a.e., then $f = g$ μ_x -a.e. for μ -a.e. x .

Proof. Suppose $\mu(E) = 0$. Choose open sets $U_n \supseteq E$ such that $\mu(U_n) \rightarrow 0$. Choose continuous functions $0 \leq h_n^\varepsilon \leq 1$ s.t. h_n^ε vanish outside U_n , h_n^ε are non-zero inside U_n , and $h_n^\varepsilon \xrightarrow{\varepsilon \rightarrow 0^+} 1_{U_n}$ (e.g. $h_n^\varepsilon(\cdot) := [\text{dist}(x, U_n^c) / \text{diam}(X)]^{1/\varepsilon}$).

By construction $1_E \leq 1_{U_n} \equiv \lim_{\varepsilon \rightarrow 0^+} h_n^\varepsilon$, whence

$$\begin{aligned} \int \mu_x(E) d\mu(x) &\leq \int \int \lim_{\varepsilon \rightarrow 0^+} h_n^\varepsilon d\mu_x d\mu \leq \\ &\leq \lim_{\varepsilon \rightarrow 0^+} \int \int h_n^\varepsilon d\mu_x d\mu = \lim_{\varepsilon \rightarrow 0^+} \int h_n^\varepsilon d\mu \leq \lim_{\varepsilon \rightarrow 0^+} \int h_n^\varepsilon d\mu \leq \mu(U_n) \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

It follows that $\mu_x(E) = 0$ a.e.

Step 4. For all f μ -absolutely integrable, $\mathbb{E}(f|\mathcal{F})(x) = \int f d\mu_x$ μ -a.e.

Proof. Find $g_n \in C(X)$ such that

$$f = \sum_{n=1}^{\infty} g_n \mu\text{-a.e.}, \text{ and } \sum \|g_n\|_{L^1(\mu)} < \infty.$$

Then

$$\begin{aligned} \mathbb{E}(f|\mathcal{F}) &= \sum_{n=1}^{\infty} \mathbb{E}(g_n|\mathcal{F}), \text{ because } \mathbb{E}(\cdot|\mathcal{F}) \text{ is a bounded operator on } L^1 \\ &= \sum_{n=1}^{\infty} \int_X g_n d\mu_x \text{ a.e.}, \text{ because } g_n \in C(X) \\ &= \int_X \sum_{n=1}^{\infty} g_n d\mu_x \text{ a.e.}, \text{ (justification below)} \\ &= \int_X f d\mu_x \text{ a.e.} \end{aligned}$$

Here is the justification: $\int \sum |g_n| d\mu_x < \infty$, because the integral of this expression, by the monotone convergence theorem is less than $\sum \|g_n\|_1 < \infty$. \square

2.3.3 The ergodic decomposition

Theorem 2.5 (The Ergodic Decomposition). *Let μ be an invariant Borel probability measure of a Borel map T on a standard probability space X . Let $\{\mu_x\}_{x \in X}$ be the conditional probabilities w.r.t. $\mathfrak{Inv}(T)$. Then*

1. $\mu = \int_X \mu_x d\mu(x)$ (i.e. this holds when applies to L^1 -functions or Borel sets);
2. μ_x is invariant for μ -a.e. $x \in X$;
3. μ_x is ergodic for μ -a.e. $x \in X$.

Proof. By the isomorphism theorem for standard probability spaces, there is no loss of generality in assuming that X is a compact metric space, and that \mathcal{B} is its σ -algebra of Borel sets.

For every $f \in L^1$, $\int f d\mu = \int \mathbb{E}(f | \mathfrak{Inv}(T)) d\mu(x) = \int_X \int_X f d\mu_x d\mu(x)$. This shows (1). We have to show that μ_x is invariant and ergodic for μ -a.e. x .

Fix a countable set $\{f_n\}$ which is dense in $C(X)$, and choose Borel versions $\overline{\mathbb{E}}_\mu(f_n | \mathfrak{Inv}(T))(x)$. By the ergodic theorem, there is a set of full measure Ω such that for all $x \in \Omega$,

$$\int f_n d\mu_x = \overline{\mathbb{E}}_\mu(f_n | \mathfrak{Inv}(T))(x) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f_n(T^k x) \text{ for all } n.$$

Step 1. μ_x is T -invariant for a.e. $x \in \Omega$.

Proof. For every n ,

$$\begin{aligned} \int f_n \circ T d\mu_x &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f_n(T^{k+1} x) \text{ a.e. (by the PET)} \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f_n(T^k x) = \overline{\mathbb{E}}_\mu(f_n | \mathfrak{Inv}(T))(x) = \int f_n d\mu_x. \end{aligned}$$

Let Ω' be the set of full measure for which the above holds for all n . Since $\{f_n\}$ is $\|\cdot\|_\infty$ -dense in $C(X)$, we have $\int f \circ T d\mu_x = \int f d\mu_x$ for all $x \in \Omega'$. Fix $x \in \Omega'$. $C(X)$ is $\|\cdot\|_{L^1(\mu_x)}$ -dense in $L^1(\mu_x)$, so $\int f \circ T d\mu_x = \int f d\mu_x$ for all μ_x -integrable functions. This means that $\mu_x \circ T^{-1} = \mu_x$ for all $x \in \Omega'$.

Step 2. μ_x is ergodic for all $x \in \Omega$.

Proof. With $\{f_n\}_{n=1}^\infty$ as above, let $\Omega := \{x : \forall k, \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f_k(T^k x) = \int f_k d\mu_x\}$. This is a set of full measure because of the ergodic theorem. Now

$$\begin{aligned}
0 &= \lim_{N \rightarrow \infty} \left\| \frac{1}{N} \sum_{k=0}^{N-1} f_n \circ T^k - \int_X f_n d\mu_x \right\|_{L^1(\mu)} \quad (\because \text{the convergence in the PET is in } L^1) \\
&= \lim_{N \rightarrow \infty} \int_X \left\| \frac{1}{N} \sum_{k=0}^{N-1} f_n \circ T^k - \int_X f_n d\mu_x \right\|_{L^1(\mu_x)} d\mu(x) \quad (\because \mu = \int_X \mu_x d\mu) \\
&= \int_X \lim_{N \rightarrow \infty} \left\| \frac{1}{N} \sum_{k=0}^{N-1} f_n \circ T^k - \int_X f_n d\mu_x \right\|_{L^1(\mu_x)} d\mu(x)
\end{aligned}$$

because of the bounded convergence theorem. But this means that

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{k=0}^{n-1} f_n \circ T^k - \int_X f_n d\mu_x \right\|_{L^1(\mu_x)} = 0 \quad \mu\text{-a.e.}$$

Since n ranges over a countable set, we get that for μ -a.e. x ,

$$\frac{1}{n} \sum_{k=0}^{n-1} f_n \circ T^k \xrightarrow[n \rightarrow \infty]{L^1(\mu_x)} \int f_n d\mu_x \quad \text{for all } n.$$

But $\{f_n\}_{n \geq 1}$ is dense in $L^1(\mu_x)$, because it is dense in $C(X)$. Therefore we have L^1 convergence for all $L^1(\mu_x)$ -functions. We just showed that for a.e. x , the ergodic averages converge in $L^1(\mu_x)$ to constants. This means that all T -invariant functions are μ_x -a.e. constant. \square

2.4 The Subadditive Ergodic Theorem

We begin with two examples.

Example 1 (Random walks on groups) Let (X, \mathcal{B}, μ, T) be the Bernoulli scheme with probability vector $\underline{p} = (p_1, \dots, p_d)$. Suppose G is a group, and $f : X \rightarrow G$ is the function $f(x_0, x_2, \dots) = g_{x_0}$, where $g_1, \dots, g_n \in G$. The expression

$$f_n(x) := f(x)f(Tx) \cdots f(T^{n-1}x)$$

describes the position of a random walk on G , which starts at the identity, and whose steps have the distribution $\Pr[\text{step} = g_i] = p_i$. What can be said on the behavior of this random walk?

In the special case $G = \mathbb{Z}^d$ or $G = \mathbb{R}^d$, $f_n(x) = f(x) + f(Tx) + \cdots + f(T^{n-1}x)$, and the ergodic theorem² says that $\frac{1}{n} f_n(x)$ has an almost sure limit, equal to $\int f d\mu = \sum p_i g_i$. So: the random walk has speed $\|\sum p_i g_i\|$, and direction $\sum p_i g_i / \|\sum p_i g_i\|$. (Note that if $G = \mathbb{Z}^d$, the direction need not lie in G .)

² applied to the each coordinate of the vector valued function $f = (f^1, \dots, f^d)$.

Example 2 (The derivative cocycle) Suppose $T : V \rightarrow V$ is a diffeomorphism acting on an open set $V \subset \mathbb{R}^d$. The derivative of T at $x \in V$ is a linear transformation $(dT)(x)$ on \mathbb{R}^d , $v \mapsto [(dT)(x)]v$. By the chain rule,

$$(dT^n)(x) = (dT)(T^{n-1}x) \circ (dT)(T^{n-2}x) \circ \cdots \circ (dT)(x).$$

If we write $f(x) := (dT)(x) \in \text{GL}(d, \mathbb{R})$, then we see that

$$(dT^n)(x) = f(T^{n-1}x)f(T^{n-2}x) \cdots f(Tx)f(x)$$

is a “random walk” on $\text{GL}(d, \mathbb{R})$. (But notice the order of multiplication!)

What is its “speed”? Is there an asymptotic “direction”?

The problem of describing the “direction” of random walk on a group is deep, and at the front line of research to this day (even in the case of matrix groups!) We postpone it for the moment, and focus on the conceptually easier task of defining the “speed”. To do this, we assume that G possesses a right invariant metric (e.g. $\text{dist}(A, B) = |\log \|AB^{-1}\||$ on $\text{GL}(d, \mathbb{R})$), and we ask for the asymptotic behavior of $g^{(n)}(x) := \text{dist}(id, f_n(x))$ as $n \rightarrow \infty$.

Key observation: $g^{(n+m)} \leq g^{(n)} + g^{(m)} \circ T^n$, because

$$\begin{aligned} g^{(n+m)} &= \text{dist}(id, f_{n+m}) = \text{dist}(id, f_m \circ T^n \cdot f_n) \\ &\leq \text{dist}(id, f_n) + \text{dist}(f_n, f_m \circ T^n \cdot f_n) \\ &= \text{dist}(id, f_n) + \text{dist}(id, f_m \circ T^n) \quad (\text{right invariance}) \\ &= g^{(n)} + g^{(m)} \circ T^n. \end{aligned}$$

We say that $\{g^{(n)}\}_n$ is a *subadditive cocycle*.

Theorem 2.6 (Kingman’s Subadditive Ergodic Theorem). *Let (X, \mathcal{B}, m, T) be a probability preserving transformation, and suppose $g^{(n)} : X \rightarrow \mathbb{R}$ is a sequence of absolutely integrable functions such that $g^{(n+m)} \leq g^{(n)} + g^{(m)} \circ T^n$ for all n, m . Then the limit $g := \lim_{n \rightarrow \infty} g^{(n)}/n \geq -\infty$ exists almost surely, and is an invariant function.*

Proof. We begin by observing that it is enough to treat the case when $g^{(n)}$ are all non-positive. Indeed, the functions

$$h^{(n)} := g^{(n)} - (g^{(1)} + g^{(1)} \circ T + \cdots + g^{(1)} \circ T^{n-1})$$

are non-positive, satisfy $h^{(n+m)} \leq h^{(n)} + h^{(m)} \circ T^n$, and $h^{(n)}/n$ converges a.e. to an invariant function iff $g^{(n)}/n$ converges a.e. to an invariant function, because $(g^{(n)} - h^{(n)})/n = (g^{(1)} + g^{(1)} \circ T + \cdots + g^{(1)} \circ T^{n-1})/n \rightarrow \mathbb{E}(g^{(1)} | \mathcal{I}_{\text{inv}})$ a.e. by Birkhoff’s ergodic theorem.

Assume then that $g^{(n)}$ are all non-negative. Define $G(x) := \liminf_{n \rightarrow \infty} g^{(n)}(x)/n$ (the limit may be equal to $-\infty$). We claim that $G \circ T = G$ almost surely. Starting from the subadditivity inequality $g^{(n+1)} \leq g^{(n)} \circ T + g^{(1)}$, we see that $G \leq G \circ T$. Suppose

there were a set of positive measure E where $G \circ T > G + \varepsilon$. Then for every $x \in E$, $G(T^n x) > G(x) + \varepsilon$, in contradiction to the Poincaré Recurrence Theorem. Thus $G = G \circ T$ almost surely.

Henceforth we work the set of full measure $X_0 := \bigcap_{n \geq 1} [G \circ T^n = G]$.

Fix $M > 0$, and define $G_M := G \vee (-M)$. This is an invariant function on X_0 . We aim at showing $\limsup_{n \rightarrow \infty} \frac{g^{(n)}}{n} \leq G_M$ a.s.. Since M is arbitrary, this implies that $\limsup_{n \rightarrow \infty} g^{(n)}/n \leq G = \liminf_{n \rightarrow \infty} g^{(n)}/n$, whence the theorem.

Fix $x \in X_0$, $N \in \mathbb{N}$, and $\varepsilon > 0$. Call $k \in \mathbb{N} \cup \{0\}$

- “good”, if $\exists \ell \in \{1, \dots, N\}$ s.t. $g^{(\ell)}(T^k x)/\ell \leq G_M(T^k x) + \varepsilon = G_M(x) + \varepsilon$;
- “bad”, if $g^{(\ell)}(T^k x)/\ell > G_M(x) + \varepsilon$ for all $\ell = 1, \dots, N$.

Color the integers $1, \dots, n$ inductively as follows, starting from $k = 1$. Let k be the smallest non-colored integer,

- If $k \leq n - N$ and k is “bad”, color it red;
- If $k \leq n - N$ and k is “good”, find the smallest $1 \leq \ell \leq N$ s.t. $g^{(\ell)}(T^k x)/\ell \leq G_M(T^k x) + \varepsilon$ and color the segment $[k, k + \ell)$ blue;
- If $k > n - N$, color k white.

Repeat this procedure until all integers $1, \dots, n$ are colored.

The “blue” part can be decomposed into segments $[\tau_i, \tau_i + \ell_i)$, with ℓ_i s.t. $g^{(\ell_i)}(T^{\tau_i} x)/\ell_i \leq G_M(x) + \varepsilon$. Let b denote the number of these segments.

The “red” part has size $\sum_{k=1}^n 1_{B(N, M, \varepsilon)}(T^k x)$, where

$$B(N, M, \varepsilon) := \{x \in X_0 : g^{(\ell)}(x)/\ell > G_M(x) + \varepsilon \text{ for all } 1 \leq \ell \leq N\}.$$

Let r denote the size of the red part.

The “white” part has size at most N . Let w be this size.

By the sub-additivity condition

$$\begin{aligned} \frac{g^{(n)}(x)}{n} &\leq \frac{1}{n} \sum_{i=1}^b g^{(\ell_i)}(T^{\tau_i} x) + \underbrace{\frac{1}{n} \sum_{k \text{ red}} g^{(1)}(T^k x) + \frac{1}{n} \sum_{k \text{ white}} g^{(1)}(T^k x)}_{\text{non-positive}} \\ &\leq \frac{1}{n} \sum_{i=1}^b g^{(\ell_i)}(T^{\tau_i} x) \leq \frac{1}{n} \sum_{i=1}^b (G_M(x) + \varepsilon) \ell_i = \frac{b}{n} (G_M(x) + \varepsilon). \end{aligned}$$

Now $b = n - (r + w) = n - \sum_{k=1}^n 1_{B(N, M, \varepsilon)}(T^k x) + O(1)$, so by the Birkhoff ergodic theorem, for almost every x , $b/n \xrightarrow[n \rightarrow \infty]{} 1 - \mathbb{E}(1_{B(N, M, \varepsilon)} | \mathfrak{I}nv)$. Thus

$$\limsup_{n \rightarrow \infty} \frac{g^{(n)}(x)}{n} \leq (G_M(x) + \varepsilon)(1 - \mathbb{E}(1_{B(N, M, \varepsilon)} | \mathfrak{I}nv)) \text{ almost surely.}$$

Now N was arbitrary, and for fixed M and ε , $B(N, M, \varepsilon) \downarrow \emptyset$ as $N \uparrow \infty$, because $G_M \geq G = \liminf_{\ell \rightarrow \infty} g^{(\ell)}/\ell$. It is not difficult to deduce from this that $\mathbb{E}(1_{B(N, M, \varepsilon)} | \mathfrak{I}nv) \xrightarrow[N \rightarrow \infty]{} 0$

almost surely.³ Thus

$$\limsup_{n \rightarrow \infty} \frac{g^{(n)}(x)}{n} \leq G_M(x) + \varepsilon \text{ almost surely.}$$

Since ε was arbitrary, $\limsup_{n \rightarrow \infty} g^{(n)}/n \leq G_M$ almost surely, which proves the theorem by the discussion above. \square

Proposition 2.3. *Suppose m is ergodic, then the limit in Kingman's ergodic theorem is the constant $\inf[(1/n) \int g^{(n)} dm]$ (possibly equal to $-\infty$).*

Proof. Let $G := \lim g^{(n)}/n$. Subadditivity implies that $G \circ T \leq G$. Recurrence implies that $G \circ T = G$. Ergodicity implies that $G = c$ a.e., for some constant $c = c(g) \geq -\infty$. We claim that $c \leq \inf[(1/n) \int g^{(n)} dm]$. This is because

$$\begin{aligned} c &= \lim_{k \rightarrow \infty} \frac{1}{kn} g^{(kn)} \leq \lim_{k \rightarrow \infty} \frac{1}{k} \left(\frac{g^{(n)}}{n} + \frac{g^{(n)}}{n} \circ T^n + \dots + \frac{g^{(n)}}{n} \circ T^{n(k-1)} \right) \\ &= \frac{1}{n} \int g^{(n)} dm \quad (\text{Birkhoff's ergodic theorem}), \end{aligned}$$

proving that $c \leq (1/n) \int g^{(n)} dm$ for all n .

To prove the other inequality we first note (as in the proof of Kingman's subadditive theorem) that it is enough to treat the case when $g^{(n)}$ are all non-positive. Otherwise work with $h^{(n)} := g^{(n)} - (g^{(1)} + \dots + g^{(1)} \circ T^{n-1})$. Since $g^{(1)} \in L^1$,

$$\frac{1}{n} (g^{(1)} + \dots + g^{(1)} \circ T^{n-1}) \xrightarrow{n \rightarrow \infty} \int g^{(1)} dm \text{ pointwise and in } L^1.$$

Thus $c(g) = \lim \frac{g^{(n)}}{n} = c(h) + \int g^{(1)} = \inf[(1/n) \int h^{(n)} + \int S_n g^{(1)}] = \inf[(1/n) \int g^{(n)}] dm$.

Suppose then that $g^{(n)}$ are all non-positive. Fix N , and set $g_N^{(n)} := \max\{g^{(n)}, -nN\}$. This is, again, subadditive because

$$\begin{aligned} g_N^{(n+m)} &= \max\{g^{(n+m)}, -(n+m)N\} \leq \max\{g^{(n)} + g^{(m)} \circ T^n, -(n+m)N\} \\ &\leq \max\{g_N^{(n)} + g_N^{(m)} \circ T^n, -(n+m)N\} \equiv g_N^{(n)} + g_N^{(m)} \circ T^n. \end{aligned}$$

By Kingman's theorem, $g_N^{(n)}/n$ converges pointwise to a constant $c(g_N)$. By definition, $-N \leq g_N^{(n)}/n \leq 0$, so by the bounded convergence theorem,

$$c(g_N) = \lim_{n \rightarrow \infty} \frac{1}{n} \int g_N^{(n)} dm \geq \inf \frac{1}{n} \int g_N^{(n)} dm \geq \inf \frac{1}{n} \int g^{(n)} dm. \quad (2.1)$$

³ Suppose $0 \leq f_n \leq 1$ and $f_n \downarrow 0$. The conditional expectation is monotone, so $\mathbb{E}(f_n | \mathcal{F})$ is decreasing at almost every point. Let φ be its almost sure limit, then $0 \leq \varphi \leq 1$ a.s., and by the BCT, $\mathbb{E}(\varphi) = \mathbb{E}(\lim \mathbb{E}(f_n | \mathcal{F})) = \lim \mathbb{E}(\mathbb{E}(f_n | \mathcal{F})) = \lim \mathbb{E}(f_n) = \mathbb{E}(\lim f_n) = 0$, whence $\varphi = 0$ almost everywhere.

Case 1: $c(g) = -\infty$. In this case $g^{(n)}/n \rightarrow -\infty$, and for every N there exists $N(x)$ s.t. $n > N(x) \Rightarrow g_N^{(n)}(x) = -N$. Thus $c(g_N) = -N$, and (2.1) implies $\inf[(1/n) \int g^{(n)} dm] = -\infty = c(g)$.

Case 2: $c(g)$ is finite. Take $N > |c(g)| + 1$, then for a.e. x , if n is large enough, then $g^{(n)}/n > c(g) - \varepsilon > -N$, whence $g_N^{(n)} = g^{(n)}$. Thus $c(g) = c(g_N) \geq \inf \frac{1}{n} \int g^{(n)} dm$ and we get the other inequality. \square

Here is a direct consequence of the subadditive ergodic theorem (historically, it predates the subadditive ergodic theorem):

Theorem 2.7 (Furstenberg–Kesten). *Let (X, \mathcal{B}, μ, T) be a ppt, and suppose $A : X \rightarrow \text{GL}(d, \mathbb{R})$ is a measurable function s.t. $\|A\| \in L^1$. If $A_n(x) := A(T^{n-1}x) \cdots A(x)$, then the following limit exists a.e. and is invariant: $\lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_n(x)\|$.*

The following immediate consequence will be used in the proof of the Oseledets theorem for invertible cocycles:

Remark: *Suppose (X, \mathcal{B}, m, T) is invertible, and let $g^{(n)}$ be a subadditive cocycle s.t. $g^{(1)} \in L^1$. Then for a.e. x , $\lim_{n \rightarrow \infty} g^{(n)} \circ T^{-n} / n$ exists and equals $\lim_{n \rightarrow \infty} g^{(n)} / n$.*

Proof. Since $g^{(n)}$ is subadditive, $g^{(n)} \circ T^{-n}$ is subadditive:

$$g^{(n+m)} \circ T^{-(n+m)} \leq [g^{(n)} \circ T^m + g^{(m)}] \circ T^{-(n+m)} = g^{(n)} \circ T^{-n} + [g^{(m)} \circ T^{-m}] \circ T^{-n}.$$

Let $m = \int m_y d\pi(y)$ be the ergodic decomposition of m . Kingman's ergodic theorem and the previous remark say that for π -a.e. y ,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{g^{(n)} \circ T^{-n}}{n} &= \inf \frac{1}{n} \int g^{(n)} \circ T^{-n} dm_y = \inf \frac{1}{n} \int g^{(n)} dm_y \quad m_y \text{ a.e.} \\ &= \lim_{n \rightarrow \infty} \frac{g^{(n)}}{n} \quad m_y \text{ a.e.} \end{aligned}$$

Thus the set where the statement of the remark fails has zero measure with respect to all the ergodic components of m , and this means that the statement is satisfied on a set of full m -measure. \square

2.5 The Multiplicative Ergodic Theorem

2.5.1 Preparations from Multilinear Algebra

Multilinear forms. Let $V = \mathbb{R}^n$ equipped with the usual euclidean metric. A *linear functional* on V is a linear map $\omega : V \rightarrow \mathbb{R}$. The set of linear functional is denoted by V^* . Any $v \in V$ determines $v^* \in V^*$ via $v^* = \langle v, \cdot \rangle$. Any linear function is of this form.

A k -multilinear function is a function $T : V^k \rightarrow \mathbb{R}$ such that for all i and $v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_k \in V$, $T(v_1, \dots, v_{i-1}, \cdot, v_{i+1}, \dots, v_k)$ is a linear functional.

The set of all k -multilinear functions on V is denoted by $T^k(V)$. The *tensor product* of $\omega \in T^k(V)$ and $\eta \in T^\ell(V)$ is $\omega \otimes \eta \in T^{k+\ell}(V)$ given by

$$(\omega \otimes \eta)(v_1, \dots, v_{k+\ell}) := \omega(v_1, \dots, v_k) \eta(v_{k+1}, \dots, v_{k+\ell}).$$

The tensor product is bilinear and associative, but it is not commutative.

The dimension of $T^k(V)$ is n^k . Here is a basis: $\{e_{i_1}^* \otimes \dots \otimes e_{i_k}^* : 1 \leq i_1, \dots, i_k \leq n\}$. To see this note that every element in $T^k(\Omega)$ is completely determined by its action on $\{(e_{i_1}, \dots, e_{i_k}) : 1 \leq i_1, \dots, i_k \leq n\}$.

Define an *inner product* on $T^k(V)$ by declaring the above basis to be orthonormal.

Alternating multilinear forms. A multilinear form ω is called *alternating*, if it satisfies $\exists i \neq j (v_i = v_j) \Rightarrow \omega(v_1, \dots, v_n) = 0$. Equivalently,

$$\omega(v_1, \dots, v_i, \dots, v_j, \dots, v_n) = -\omega(v_1, \dots, v_j, \dots, v_i, \dots, v_n).$$

(to see the equivalence, expand $\omega(v_1, \dots, v_i + v_j, \dots, v_j + v_i, \dots, v_n)$). The set of all k -alternating forms is denoted by $\Omega^k(V)$.

Any multilinear form ω gives rise to an alternating form $\text{Alt}(\omega)$ via

$$\text{Alt}(\omega) := \frac{1}{k!} \sum_{\sigma \in S_k} \text{sgn}(\sigma) \sigma \cdot \omega,$$

where S_k is the group of k -permutations, and the action of a permutation σ on $\omega \in T^k(V)$ is given by $(\sigma \cdot \omega)(v_1, \dots, v_k) = \omega(v_{\sigma(1)}, \dots, v_{\sigma(k)})$. The normalization $k!$ is to guarantee $\text{Alt}|_{\Omega^k(V)} = \text{id}$, and $\text{Alt}^2 = \text{Alt}$. Note that Alt is linear.

Lemma 2.1. $\text{Alt}[\text{Alt}(\omega_1 \otimes \omega_2) \otimes \omega_3] = \text{Alt}(\omega_1 \otimes \omega_2 \otimes \omega_3)$.

Proof. We show that if $\text{Alt}(\omega) = 0$, then $\text{Alt}(\omega \otimes \eta) = 0$ for all η . Specializing to the case $\omega = \text{Alt}(\omega_1 \otimes \omega_2) - \omega_1 \otimes \omega_2$ and $\eta = \omega_3$, we get (since $\text{Alt}^2 = \text{Alt}$)

$$\text{Alt}[(\text{Alt}(\omega_1 \otimes \omega_2) - \omega_1 \otimes \omega_2) \otimes \omega_3] = 0,$$

which is equivalent to the statement of the lemma.

Suppose $\omega \in T^k(V)$, $\eta \in T^\ell(V)$, and $\text{Alt}(\omega) = 0$. Let $G := \{\sigma \in S_{k+\ell} : \sigma(i) = i \text{ for all } i = k+1, \dots, k+\ell\}$. This is a normal subgroup of $S_{k+\ell}$, and there is natural isomorphism $\sigma \mapsto \sigma' := \sigma|_{\{1, \dots, k\}}$ from G to S_k . Let $S_{k+\ell} = \bigsqcup_j G\sigma_j$ be the corresponding right coset decomposition, then

$$(k+\ell)! \text{Alt}(\omega \otimes \eta)(v_1, \dots, v_{k+\ell}) =$$

$$\begin{aligned}
&= \sum_j \sum_{\sigma \in G} \operatorname{sgn}(\sigma \sigma_j) (\sigma \sigma_j) \cdot (\omega \otimes \eta)(v_1, \dots, v_{k+\ell}) \\
&= \sum_j \operatorname{sgn}(\sigma_j) \eta(v_{\sigma_j(k+1)}, \dots, v_{\sigma_j(k+\ell)}) \sum_{\sigma \in G} \operatorname{sgn}(\sigma) (\sigma \cdot \omega)(v_{\sigma_j(1)}, \dots, v_{\sigma_j(k)}) \\
&= \sum_j \operatorname{sgn}(\sigma_j) \eta(v_{\sigma_j(k+1)}, \dots, v_{\sigma_j(k+\ell)}) \sum_{\sigma' \in S_k} \operatorname{sgn}(\sigma') (\sigma' \cdot \omega)(v_{\sigma_j(1)}, \dots, v_{\sigma_j(k)}) \\
&= \sum_j \operatorname{sgn}(\sigma_j) \eta(v_{\sigma_j(k+1)}, \dots, v_{\sigma_j(k+\ell)}) k! \operatorname{Alt}(\omega)(v_{\sigma_j(1)}, \dots, v_{\sigma_j(k)}) = 0.
\end{aligned}$$

□

Using this ‘antisymmetrization operator’, we define the following product, called *exterior product* or *wedge product*: If $\omega \in \Omega^k(V)$, $\eta \in \Omega^\ell(V)$, then

$$\omega \wedge \eta := \frac{(k+\ell)!}{k!\ell!} \operatorname{Alt}(\omega \otimes \eta).$$

The wedge product is bilinear, and the previous lemma shows that it is associative. It is almost anti commutative: If $\omega \in \Omega^k(V)$, $\eta \in \Omega^\ell(V)$, then

$$\omega \wedge \eta = (-1)^{k\ell} \eta \wedge \omega.$$

We’ll see the reason for the peculiar normalization later.

Proposition 2.4. $\{e_{i_1}^* \wedge \dots \wedge e_{i_k}^* : 1 \leq i_1 < \dots < i_k \leq n\}$ is a basis for $\Omega^k(V)$, whence $\dim \Omega^k(V) = \binom{n}{k}$.

Proof. Suppose $\omega \in \Omega^k(V)$, then $\omega \in T^k(V)$ and so $\omega = \sum a_{i_1, \dots, i_k} e_{i_1}^* \otimes \dots \otimes e_{i_k}^*$, where the sum ranges over all k -tuples of numbers between 1 and n . If $\omega \in \Omega^k(V)$, then $\operatorname{Alt}(\omega) = \omega$ and so

$$\omega = \sum a_{i_1, \dots, i_k} \operatorname{Alt}(e_{i_1}^* \otimes \dots \otimes e_{i_k}^*).$$

Fix $\xi := e_{i_1}^* \otimes \dots \otimes e_{i_k}^*$. If $i_\alpha = i_\beta$ for some $\alpha \neq \beta$, then the permutation σ_0 which switches $\alpha \leftrightarrow \beta$ preserves ξ . Thus for all $\sigma \in S_k$,

$$\operatorname{sgn}(\sigma \sigma_0) (\sigma \sigma_0) \cdot \xi = -\operatorname{sgn}(\sigma) \sigma \cdot \xi$$

and we conclude that $\operatorname{Alt}(\xi) = 0$. If, on the other hand, i_1, \dots, i_k are all different, then it is easy to see using lemma 2.1 that

$$\operatorname{Alt}(e_{i_1}^* \otimes e_{i_2}^* \otimes \dots \otimes e_{i_k}^*) = \frac{1}{k!} e_{i_1}^* \wedge \dots \wedge e_{i_k}^*.$$

Thus $\omega = \frac{1}{k!} \sum a_{i_1, \dots, i_k} e_{i_1}^* \wedge \dots \wedge e_{i_k}^*$, and we have proved that the set of forms in the statement spans $\Omega^k(V)$. To see that this set is independent, note that we can determine the coefficient of $e_{i_1}^* \wedge \dots \wedge e_{i_k}^*$ by evaluating the form on $(e_{i_1}, \dots, e_{i_k})$. □

Corollary 2.2. $e_1^* \wedge \cdots \wedge e_n^*$ is the determinant. This is the reason for the peculiar normalization in the definition of \wedge .

Proof. The determinant is an alternating n -form, and $\dim \Omega^n(V) = 1$, so the determinant is proportional to $e_1^* \wedge \cdots \wedge e_n^*$. Since the values of both forms on the standard basis is one (because $e_1^* \wedge \cdots \wedge e_n^* = n! \text{Alt}(e_1^* \otimes \cdots \otimes e_n^*)$), they are equal. \square

We define an inner product on $\Omega^k(V)$ by declaring the basis in the proposition to be orthonormal. Let $\|\cdot\|$ be the resulting norm.

Lemma 2.2. For $v \in V$, let $v^* := \langle v, \cdot \rangle$, then

- (a) $\|\omega \wedge \eta\| \leq \|\omega\| \|\eta\|$.
- (b) $\langle v_1^* \wedge \cdots \wedge v_k^*, w_1^* \wedge \cdots \wedge w_k^* \rangle = \det(\langle v_i, w_j \rangle)$.
- (c) If $\{u_1, \dots, u_n\}$ is an orthonormal basis for V , then $\{u_{i_1}^* \wedge \cdots \wedge u_{i_k}^* : 1 \leq i_1 < \cdots < i_k \leq n\}$ is an orthonormal basis for $\Omega^k(V)$.
- (d) If $\text{span}\{v_1, \dots, v_k\} = \text{span}\{u_1, \dots, u_k\}$, then $v_1^* \wedge \cdots \wedge v_k^*$ and $u_1^* \wedge \cdots \wedge u_k^*$ are proportional.

Proof. Write for $I = (i_1, \dots, i_k)$ such that $1 \leq i_1 < \cdots < i_k \leq n$, $e_I^* := e_{i_1}^* \wedge \cdots \wedge e_{i_k}^*$. Represent $\omega := \sum \alpha_I e_I^*$, $\eta := \sum \beta_J e_J^*$, then

$$\|\omega \wedge \eta\|^2 = \left\| \sum_{I,J} \alpha_I \beta_J e_I^* \wedge e_J^* \right\|^2 = \left\| \sum_{I \cap J = \emptyset} \pm \alpha_I \beta_J e_{I \cup J}^* \right\|^2 = \sum_{I \cap J = \emptyset} \alpha_I^2 \beta_J^2 \leq \|\omega\|^2 \|\eta\|^2.$$

Take two multi indices I, J . If $I = J$, then the inner product matrix is the identity matrix. If $I \neq J$, then $\exists \alpha \in I \setminus J$ and then the α -row and column of the inner product matrix will be zero. Thus the formula holds for any pair e_I^*, e_J^* . Since part (b) of the lemma holds for all basis vectors, it holds for all vectors. Part (c) immediately follows.

Next we prove part (d). Represent $v_i = \sum \alpha_{ij} w_j$, then

$$\begin{aligned} v_1^* \wedge \cdots \wedge v_k^* &= \text{const. Alt}(v_1^* \otimes \cdots \otimes v_k^*) = \text{const. Alt} \left(\sum_j \alpha_{1j} u_j^* \otimes \cdots \otimes \sum_j \alpha_{kj} u_j^* \right) \\ &= \text{const.} \sum \alpha_{1j_1} \cdots \alpha_{kj_k} \text{Alt}(u_{j_1}^* \otimes \cdots \otimes u_{j_k}^*). \end{aligned}$$

The terms where j_1, \dots, j_k are not all different are annihilated by Alt. The terms where j_1, \dots, j_k are all different are mapped by Alt to a form which proportional to $u_1^* \wedge \cdots \wedge u_k^*$. Thus the result of the sum is proportional to $u_1^* \wedge \cdots \wedge u_k^*$. \square

Exterior product of linear operators Let $A : V \rightarrow V$ be a linear operator. The k -th exterior product of A is $A^{\wedge k} : \Omega^k(V) \rightarrow \Omega^k(V)$ given by

$$(A^{\wedge k} \omega)(v_1, \dots, v_k) := (A^t v_1, \dots, A^t v_k).$$

The transpose is used to get $A^{\wedge k}(v_1^* \wedge \cdots \wedge v_k^*) = (Av_1)^* \wedge \cdots \wedge (Av_k)^*$.

Theorem 2.8. $\|A^{\wedge k}\| = \lambda_1 \cdots \lambda_k$, where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ are the eigenvalues of $(A^t A)^{1/2}$, listed in decreasing order with multiplicities.

Proof. The matrix AA^t is symmetric, so it can be orthogonally diagonalized. Let $\{v_1, \dots, v_n\}$ be an orthonormal basis of eigenvectors, listed so that $(AA^t)v_i = \lambda_i^2 v_i$. Then $\{v_I^* : I \subseteq \{1, \dots, d\}, |I| = k\}$ is an orthonormal basis for $\Omega^k(\mathbb{R}^d)$, where we are using the multi index notation

$$v_I^* = v_{i_1}^* \wedge \cdots \wedge v_{i_k}^*,$$

where $i_1 < \cdots < i_k$ is an ordering of I .

Given $\omega \in \Omega^k(\mathbb{R}^d)$, write $\omega = \sum \omega_I v_I^*$, then

$$\begin{aligned} \|A^{\wedge k} \omega\|^2 &= \langle A^{\wedge k} \omega, A^{\wedge k} \omega \rangle = \left\langle \sum_I \omega_I A^{\wedge k} v_I^*, \sum_J \omega_J A^{\wedge k} v_J^* \right\rangle \\ &= \sum_{I, J} \omega_I \omega_J \langle A^{\wedge k} v_I^*, A^{\wedge k} v_J^* \rangle. \end{aligned}$$

$$\begin{aligned} \text{Now, } \langle A^{\wedge k} v_I^*, A^{\wedge k} v_J^* \rangle &= \langle (A^t v_{i_1})^* \wedge \cdots \wedge (A^t v_{i_k})^*, (A^t v_{j_1})^* \wedge \cdots \wedge (A^t v_{j_k})^* \rangle \\ &= \det \left(\langle A^t v_{i_\alpha}, A^t v_{j_\beta} \rangle \right) \quad (\text{Lemma 2.2(b)}) \\ &= \det \left(\langle v_{i_\alpha}, AA^t v_{j_\beta} \rangle \right) = \det \left(\langle v_{i_\alpha}, \lambda_{i_\beta}^2 v_{j_\beta} \rangle \right) \\ &= \prod_{j \in J} \lambda_j^2 \det \left(\langle v_{i_\alpha}, v_{i_\beta} \rangle \right) \\ &= \prod_{j \in J} \lambda_j^2 \langle v_I^*, v_J^* \rangle = \begin{cases} \prod_{i \in I} \lambda_i^2 & I = J \\ 0 & I \neq J \end{cases} \end{aligned}$$

Thus $\|A^{\wedge k} \omega\|^2 = \sum_I \omega_I^2 \prod_{i \in I} \lambda_i^2 \leq \|\omega\|^2 \prod_{i=1}^k \lambda_i^2$. It follows that $\|A^{\wedge k}\| \leq \lambda_1 \cdots \lambda_k$.

To see that the inequality is in fact an equality, consider the case $\omega = v_I^*$ where $I = \{1, \dots, k\}$: $\|A^{\wedge k} \omega\| = \langle v_I^*, v_I^* \rangle = (\lambda_1 \cdots \lambda_k)^2 = (\lambda_1 \cdots \lambda_k)^2 \|\omega\|^2$. \square

Exterior products and angles between vector spaces The angle between vector spaces $V, W \subset \mathbb{R}^d$ is

$$\angle(V, W) := \min \{ \arccos \langle v, w \rangle : v \in V, w \in W, \|v\| = \|w\| = 1 \}.$$

So if $V \cap W \neq \{0\}$ iff $\angle(V, W) = 0$, and $V \perp W$ iff $\angle(V, W) = \pi/2$.

Proposition 2.5. If (w_1^*, \dots, w_k^*) is a basis of W , and (v_1^*, \dots, v_ℓ^*) is a basis of V , then $\|(v_1^* \wedge \cdots \wedge v_\ell^*) \wedge (w_1^* \wedge \cdots \wedge w_k^*)\| \leq \|v_1^* \wedge \cdots \wedge v_\ell^*\| \cdot \|w_1^* \wedge \cdots \wedge w_k^*\| \cdot |\sin \angle(V, W)|$.

Proof. If $V \cap W \neq \{0\}$ then both sides are zero, so suppose $V \cap W = \{0\}$, and pick an orthonormal basis e_1, \dots, e_{n+k} for $V \oplus W$. Let $w \in W, v \in V$ be unit vectors s.t. $\angle(V, W) = \angle(v, w)$, and write $v = \sum v_i e_i, w = \sum w_j e_j$, then

$$\begin{aligned}
\|v^* \wedge w^*\|^2 &= \left\| \sum_{i,j} v_i w_j e_i^* \wedge e_j^* \right\|^2 = \left\| \sum_{i < j} (v_i w_j - v_j w_i) e_i^* \wedge e_j^* \right\|^2 = \sum_{i < j} (v_i w_j - v_j w_i)^2 \\
&= \frac{1}{2} \sum_{i,j} (v_i w_j - v_j w_i)^2 \quad (\text{the terms where } i = j \text{ vanish}) \\
&= \frac{1}{2} \sum_{i,j} (v_i^2 w_j^2 + v_j^2 w_i^2 - 2v_i w_i \cdot v_j w_j) = \frac{1}{2} \left[2 \sum_i v_i^2 \sum_j w_j^2 - 2 \left(\sum_i v_i w_i \right)^2 \right] \\
&= \|v\|^2 \|w\|^2 - \langle v, w \rangle^2 = 1 - \cos^2 \angle(v, w) = \sin^2 \angle(V, W).
\end{aligned}$$

Complete v to an orthonormal basis $(v, v'_2, \dots, v'_\ell)$ of V , and complete w to an orthonormal basis (w, w'_2, \dots, w'_k) of W . Then

$$\begin{aligned}
&\|(v^* \wedge v_2'^* \wedge \dots \wedge v_\ell'^*) \wedge (w^* \wedge w_2'^* \wedge \dots \wedge w_k'^*)\| \\
&\leq \|v^* \wedge w^*\| \cdot \|v_2'^* \wedge \dots \wedge v_\ell'^*\| \cdot \|w_2'^* \wedge \dots \wedge w_k'^*\| = |\sin \angle(V, W)| \cdot 1 \cdot 1,
\end{aligned}$$

because of orthonormality. By lemma 2.2

$$\begin{aligned}
v_1^* \wedge \dots \wedge v_\ell'^* &= \pm \|v_1^* \wedge \dots \wedge v_\ell'^*\| \cdot v^* \wedge v_2'^* \wedge \dots \wedge v_\ell'^* \\
w_1^* \wedge \dots \wedge w_k'^* &= \pm \|w_1^* \wedge \dots \wedge w_k'^*\| \cdot w^* \wedge w_2'^* \wedge \dots \wedge w_k'^*
\end{aligned}$$

and the proposition follows. \square

2.5.2 Proof of the Multiplicative Ergodic Theorem

Let (X, \mathcal{B}, m, f) be a ppt, and $A : X \rightarrow \text{GL}(d, \mathbb{R})$ some Borel map. We define $A_n := A \circ f^{n-1} \cdots A$, then the *cocycle identity* holds: $A_{n+m}(x) = A_n(f^m x) A_m(x)$.

Theorem 2.9 (Multiplicative Ergodic Theorem). *Let (X, \mathcal{B}, T, m) be a ppt, and $A : X \rightarrow \text{GL}(d, \mathbb{R})$ a Borel function s.t. $\ln \|A(x)^{\pm 1}\| \in L^1(m)$, then*

$$\Lambda(x) := \lim_{n \rightarrow \infty} [A_n(x)^t A_n(x)]^{1/2n}$$

exists a.e., and $\lim_{n \rightarrow \infty} \frac{1}{n} \ln \|A_n(x) \Lambda(x)^{-n}\| = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \|(A_n(x) \Lambda(x)^{-n})^{-1}\| = 0$ a.s.

Proof. The matrix $\sqrt{A_n(x)^t A_n(x)}$ is symmetric, therefore it can be orthogonally diagonalized. Let $\exists \lambda_n^1(x) < \dots < \lambda_n^{s_n(x)}(x)$ be its different eigenvalues, and $\mathbb{R}^d = W_n^{\lambda_n^1(x)}(x) \oplus \dots \oplus W_n^{\lambda_n^{s_n(x)}(x)}(x)$ the orthogonal decomposition of \mathbb{R}^d into the corresponding eigenspaces. The proof has the following structure:

Part 1: Let $t_n^1(x) \leq \dots \leq t_n^d(x)$ be a list of the eigenvalues of $\sqrt{A_n(x)^t A_n(x)}$ with multiplicities, then for a.e. x , there is a limit $t_i(x) = \lim_{n \rightarrow \infty} [t_n^i(x)]^{1/n}$, $i = 1, \dots, d$.

Part 2: Let $\lambda_1(x) < \dots < \lambda_{s(x)}(x)$ be a list of the different values of $\{t_i(x)\}_{i=1}^d$. Divide $\{t_n^i(x)\}_{i=1}^d$ into $s(x)$ subsets of values $\{t_n^i(x) : i \in I_n^j\}$, ($1 \leq j \leq s(x)$) in such a way that $t_i(x)^{1/n} \rightarrow \lambda_j(x)$ for all $i \in I_n^j$. Let

$$\begin{aligned} U_n^j(x) &:= \text{sum of the eigenspaces of } t_n^i(x), i \in I_n^j. \\ &= \text{the part of the space where } B_n(x) \text{ dilates by approximately } \lambda_j(x)^n. \end{aligned}$$

We show that the spaces $U_n^j(x)$ converge as $n \rightarrow \infty$ to some limiting spaces $U^j(x)$ (in the sense that the orthogonal projections on $U_n^j(x)$ converge to the orthogonal projection on $U^j(x)$).

Part 3: The theorem holds with $\Lambda(x) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ given by $v \mapsto \lambda_i(x)$ on $U^i(x)$.

Part 1 is proved by applying the sub additive ergodic theorem for a cleverly chosen sub-additive cocycle (“Ragunathan’s trick”). Parts 2 and 3 are (non-trivial) linear algebra.

Part 1: Set $g_i^{(n)}(x) := \sum_{j=d-i+1}^d \ln t_n^j(x)$. This quantity is finite, because $A_n^t A_n$ is invertible, so none of its eigenvalues vanish.

The sequence $g_i^{(n)}$ is subadditive! This is because the theory of exterior products says that $\exp g_i^{(n)} =$ product of the i largest e.v.’s of $\sqrt{A_n(x)^t A_n(x)} = \|A_n(x)^{\wedge i}\|$, so

$$\begin{aligned} \exp g_i^{(n+m)}(x) &= \|A_{n+m}^{\wedge i}(x)\| = \|A_m(T^n x)^{\wedge i} A_n(x)^{\wedge i}\| \leq \|A_m(T^n x)^{\wedge i}\| \|A_n(x)^{\wedge i}\| \\ &= \exp[g_i^{(m)}(T^n x) + g_i^{(n)}(x)], \end{aligned}$$

whence $g_i^{(n+m)} \leq g_i^{(n)} + g_i^{(m)} \circ T^n$.

We want to apply Kingman’s subadditive ergodic theorem. First we need to check that $g_i^{(1)} \in L^1$. We use the following fact from linear algebra: if λ is an eigenvalue of a matrix B , then $\|B^{-1}\|^{-1} \leq |\lambda| \leq \|B\|$.⁴ Therefore

$$\begin{aligned} |\ln t_n^i(x)| &\leq \frac{1}{2} \max\{|\ln \|A_n^t A_n\||, |\ln \|(A_n^t A_n)^{-1}\||\} \\ &\leq \max\{|\ln \|A_n(x)\||, |\ln \|A_n(x)^{-1}\||\} \\ &\leq \sum_{k=0}^{n-1} \left(|\ln \|A(T^k x)\|| + |\ln \|A(T^k x)^{-1}\|| \right) \\ \therefore |g_i^{(n)}(x)| &\leq ni \sum_{k=0}^{n-1} \left(|\ln \|A(T^k x)\|| + |\ln \|A(T^k x)^{-1}\|| \right). \end{aligned} \quad (2.2)$$

In the particular case $n = 1$, we get that $\|g_i^{(1)}\|_1 \leq \int |\ln \|A\|| + \int |\ln \|A^{-1}\|| d\mu < \infty$.

Thus Kingman’s ergodic theorem says that $\lim \frac{1}{n} g_i^{(n)}(x)$ exists almost surely, and belongs to $[-\infty, \infty)$. In fact the limit is finite almost everywhere, because (2.2) and

⁴ Proof: Let v be an eigenvector of λ with norm one, then $|\lambda| = \|Bv\| \leq \|B\|$ and $1 = \|B^{-1}Bv\| \leq \|B^{-1}\| \|Bv\| = \|B^{-1}\| |\lambda|$.

the Pointwise Ergodic Theorem imply that

$$\lim_{n \rightarrow \infty} \frac{1}{n} |g_i^{(n)}| \leq \mathbb{E} (|\ln \|A\|| + |\ln \|A^{-1}\||) \mathfrak{J} \nu < \infty \text{ a.e.}$$

Taking differences, we see that the following limit exists a.e.:

$$\ln t_i(x) := \lim_{n \rightarrow \infty} \frac{1}{n} [g_{d-i+1}^{(n)}(x) - g_{d-i}^{(n)}(x)] = \lim_{n \rightarrow \infty} \frac{1}{n} \ln t_i^n(x).$$

Thus $[t_n^i(x)]^{1/n} \xrightarrow[n \rightarrow \infty]{} t_i(x)$ almost surely, for some $t_i(x) \in \mathbb{R}$.

Part 2: Fix x s.t. $[t_n^i(x)]^{1/n} \xrightarrow[n \rightarrow \infty]{} t_i(x)$ for all $1 \leq i \leq d$. Henceforth we work with this x only, and write for simplicity $A_n = A_n(x)$, $t_i = t_i(x)$ etc.

Let $s = s(x)$ be the number of the different t_i . List the *different* values of these quantities in an increasing order: $\lambda_1 < \lambda_2 < \dots < \lambda_s$. Set $\chi_j := \log \lambda_j$. Fix $0 < \delta < \min\{\chi_{j+1} - \chi_j\}$. Since $(t_n^i)^{1/n} \rightarrow \lambda_i$, the following sets eventually stabilize and are independent of n :

$$I_j := \{i : |(t_n^i)^{1/n} - \lambda_i| < \delta\} \quad (j = 1, \dots, s).$$

Define, relative to $\sqrt{A_n^t A_n}$,

- $U_n^j := \bigoplus_{i \in I_j} [\text{eigenspace of } t_n^i(x)]$;
- $V_n^r := \bigoplus_{j \leq r} U_n^j$
- $\tilde{V}_n^r := \bigoplus_{j \geq r} U_n^j$

The linear spaces U_n^1, \dots, U_n^s are orthogonal, since they are eigenspaces of different eigenvalues for a symmetric matrix $(\sqrt{A_n^t A_n})$. We show that they converge as $n \rightarrow \infty$, in the sense that their orthogonal projections converge.

The proof is based on the following technical lemma. Denote the projection of a vector v on a subspace W by $v|W$, and write $\chi_i := \log \lambda_i$.

Technical lemma: For every $\delta > 0$ there exists constants $K_1, \dots, K_s > 1$ and N s.t. for all $n > N$, $t = 1, \dots, s$, $k \in \mathbb{N}$, and $u \in V_n^r$,

$$\|u|\tilde{V}_{n+k}^{r+t}\| \leq K_t \|u\| \exp(-n(\chi_{r+t} - \chi_r - \delta t))$$

We give the proof later. First we show how it can be used to finish parts 2 and 3.

We show that V_n^r converge as $n \rightarrow \infty$. Since the projection on U_n^i is the projection on V_n^i minus the projection on V_n^{i-1} , it will then follow that the projections of U_n^i converge.

Fix N large. We need it to be so large that

1. I_j are independent of n for all $n > N$;
2. The technical lemma works for $n > N$ with δ as above.

There will be other requirements below.

Fix an orthonormal basis $(v_n^1, \dots, v_n^{d_r})$ for V_n^r ($d_r = \dim(V_n^r) = \sum_{j \leq r} |I_j|$). Write

$$v_n^i = \alpha_n^i w_{n+1}^i + u_{n+1}^i, \text{ where } w_{n+1}^i \in V_{n+1}^r, \|w_{n+1}^i\| = 1, u_{n+1}^i \in \tilde{V}_{n+1}^{r+1}.$$

Note that $\|u_{n+1}^i\| = \|v_n^i\| \|\tilde{V}_{n+1}^{r+1}\| \leq K_1 \exp(-n(\chi_{r+1} - \chi_r - \delta))$. Using the identity $\alpha_n^i = \sqrt{1 - \|u_{n+1}^i\|^2}$, it is easy to see that for some constants C_1 and $0 < \theta < 1$ independent of n and (v_n^i) ,

$$\|v_n^i - w_{n+1}^i\| \leq C_1 \theta^n.$$

($\theta := \max_r \exp[-(\lambda_{r+1} - \lambda_r - \delta)]$ and $C_1 := 2K_1$ should work.)

The system $\{w_{n+1}^i\}$ is very close to being orthonormal:

$$\langle w_{n+1}^i, w_{n+1}^j \rangle = \langle w_{n+1}^i - v_n^i, w_{n+1}^j \rangle + \langle v_n^i, v_n^j \rangle + \langle v_n^i, w_{n+1}^j - v_n^j \rangle = \delta_{ij} + O(\theta^n),$$

because $\{v_n^i\}$ is an orthonormal system. It follows that for all n large enough, w_{n+1}^i are linearly independent. A quick way to see this is to note that

$$\begin{aligned} \|(w_{n+1}^1)^* \wedge \dots \wedge (w_{n+1}^{d_r})^*\|^2 &= \det(\langle w_{n+1}^i, w_{n+1}^j \rangle) \quad (\text{lemma 2.2}) \\ &= \det(I + O(\theta^n)) \neq 0, \text{ provided } n \text{ is large enough,} \end{aligned}$$

and to observe that wedge produce of a linearly dependent system vanishes.

It follows that $\{w_{n+1}^1, \dots, w_{n+1}^{d_r}\}$ is a linearly independent subset of V_{n+1}^r . Since $\dim(V_{n+1}^r) = \sum_{j \leq r} |I_j| = d_r$, this is a basis for V_{n+1}^r .

Let (v_{n+1}^i) be the orthonormal basis obtained by applying the Gram–Schmidt procedure to (w_{n+1}^i) . We claim that there is a global constant C_2 such that

$$\|v_n^i - v_{n+1}^i\| \leq C_2 \theta^n. \quad (2.3)$$

Write $v_i = v_{n+1}^i, w_i = w_{n+1}^i$, then the Gram–Schmidt process is to set $\bar{v}_i = u_i / \|u_i\|$, where u_i are defined by induction by $u_1 := w_1, u_i := w_i - \sum_{j < i} \langle w_i, \bar{v}_j \rangle \bar{v}_j$. We construct by induction global constants C_2^i s.t. $\|\bar{v}_i - w_i\| \leq C_2^i \theta^n$, and then take $C_2 := \max\{C_2^i\}$. When $i = 1$, we can take $C_2^1 := C_1$, because $\bar{v}_1 = w_1$, and $\|w_1 - v_n^1\| \leq C_1 \theta^n$. Suppose we have constructed C_2^1, \dots, C_2^{i-1} . Then

$$\|u_i - w_i\| \leq \sum_{j < i} |\langle w_i, \bar{v}_j \rangle| \leq \sum_{j < i} |\langle w_i, w_j \rangle| + \|w_j - \bar{v}_j\| \leq \left(2C_1(i-1) + \sum_{j < i} C_2^j \right) \theta^n,$$

because $|\langle w_i, w_j \rangle| = |\langle w_i - v_n^i, w_j \rangle + \langle v_n^i, w_j - v_n^j \rangle + \langle v_n^i, v_n^j \rangle| \leq 2C_1 \theta^n$. Call the term in the brackets K , and assume n is so large that $K\theta^n < 1/2$, then $\|u_i\| - 1 \leq \|u_i - w_i\| \leq K\theta^n$, whence

$$\|\bar{v}_i - w_i\| = \left\| \frac{u_i - \|u_i\|w_i}{\|u_i\|} \right\| \leq \frac{\|u_i - w_i\| + |1 - \|u_i\||}{\|u_i\|} \leq 4K\theta^n$$

and we can take $C_2^i := 4K$. This proves (2.3).

Starting from the orthonormal basis (v_n^i) for V_n^r , we have constructed an orthonormal basis (v_{n+1}^i) for V_{n+1}^r such that $\|v_n^i - v_{n+1}^i\| \leq C_2\theta^n$. Continue this procedure by induction, and construct the orthonormal bases (v_{n+k}^i) for V_{n+k}^r . By (2.3), these bases form Cauchy sequences: $v_{n+k}^i \xrightarrow[k \rightarrow \infty]{} v^i$.

The limit vectors must also be orthonormal. Denote their span by V^r . The projection on V^r takes the form

$$\sum_{i=1}^{d_r} \langle v^i, \cdot \rangle v^i = \lim_{k \rightarrow \infty} \sum_{i=1}^{d_r} \langle v_{n+k}^i, \cdot \rangle v_{n+k}^i = \lim_{k \rightarrow \infty} \text{proj}_{V_{n+k}^r}.$$

Thus $V_{n+k}^r \rightarrow V^r$.

Part 3: We saw that $\text{proj}_{U_n^i(x)} \xrightarrow[n \rightarrow \infty]{} \text{proj}_{U^i(x)}$ for some linear spaces $U_i(x)$. Set $\Lambda(x) \in \text{GL}(\mathbb{R}^d)$ to be the matrix representing

$$\Lambda(x) = \sum_{j=1}^{s(x)} e^{\chi_j(x)} \text{proj}_{U_j(x)}.$$

Since $U_i(x)$ are limits of U_n^i , they are orthogonal, and they sum up to \mathbb{R}^d . It follows that Λ is invertible, symmetric, and positive.

Let W_n^i be the eigenspace of $t_n^i(x)$ for $\sqrt{A_n^t A_n}$, then for all $v \in \mathbb{R}^d$,

$$\begin{aligned} (A_n^t A_n)^{1/2n} v &= (\sqrt{A_n^t A_n})^{1/n} v = \sum_{i=1}^d t_n^i(x)^{1/n} \text{proj}_{W_n^i}(v) \\ &= \sum_{j=1}^s \sum_{i \in I_j} t_n^i(x)^{1/n} \text{proj}_{W_n^i}(v) \\ &= \sum_{j=1}^s e^{\chi_j(x)} \sum_{i \in I_j} \text{proj}_{W_n^i}(v) + o(\|v\|), \end{aligned}$$

where $o(\|v\|)$ denotes a vector with norm $o(\|v\|)$

$$= \sum_{j=1}^s e^{\chi_j(x)} \text{proj}_{U_n^j}(v) + o(\|v\|) \xrightarrow[n \rightarrow \infty]{} \Lambda(x)v.$$

Thus $(A_n^t A_n)^{1/2n} \rightarrow \Lambda$.

We show that $\frac{1}{n} \log \|(A_n \Lambda^{-n})^{\pm 1}\| \xrightarrow[n \rightarrow \infty]{} 0$. It's enough to show that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_n v\| = \chi_r := \log \lambda_r \text{ uniformly on the unit ball in } U_r. \quad (2.4)$$

To see that this is enough, note that $\Lambda v = \sum_{r=1}^s e^{\chi_r} v|U_r$; for all $\delta > 0$, if n is large enough, then for every v ,

$$\begin{aligned} \|A_n \Lambda^{-n} v\| &\leq \sum_{r=1}^s e^{-n\chi_r} \|A_n(v|U_r)\| = \sum_{r=1}^s e^{-n\chi_r} e^{n(\chi_r+\delta)} \|v\| \leq s e^{n\delta} \|v\| \quad (v \in \mathbb{R}^d) \\ \|A_n \Lambda^{-n} v\| &= e^{-n\chi_r} \|A_n v\| = e^{\pm n\delta} \|v\| \quad (v \in U_r) \end{aligned}$$

Thus $\|A_n \Lambda^{-n}\| \asymp e^{-n\delta}$ for all δ , whence $\frac{1}{n} \log \|A_n \Lambda^{-n}\| \rightarrow 0$ a.e.

To see that $\frac{1}{n} \log \|(A_n \Lambda^{-n})^{-1}\| \rightarrow 0$, we use a duality trick.

Define for a matrix C , $C^\# := (C^{-1})^t$, then $(C_1 C_2)^\# = C_1^\# C_2^\#$. Thus $(A^\#)_n = (A_n)^\#$, and $B_n^\# := \sqrt{(A^\#)_n^t (A^\#)_n} = (\sqrt{A_n^t A_n})^\# = (\sqrt{A_n^t A_n})^{-1}$. Thus we have the following relation between the objects associated to $A^\#$ and A :

1. the eigenvalues of $B_n^\#$ are $1/t_n^d \leq \dots \leq 1/t_n^1$ (the order is flipped)
2. the eigenspace of $1/t_n^i$ for $B_n^\#$ is the eigenspace of t_n^i for B_n
3. $\chi_j^\# = -\chi_{s-j+1}$
4. $(U_n^j)^\# = U_n^{s-j+1}$, $(V_n^r)^\# = \tilde{V}_n^{s-r+1}$, $(\tilde{V}_n^r)^\# = V_n^{s-r+1}$
5. $\Lambda^\# = \Lambda^{-1}$.

Thus $\|(\Lambda^n A_n^{-1})\| = \|(\Lambda^n A_n^{-1})^t\| = \|A_n^\# (\Lambda^\#)^{-n}\|$, so the claim $\frac{1}{n} \log \|\Lambda^n A_n^{-1}\| \xrightarrow{n \rightarrow \infty} 0$ a.e. follows from what we did above, applied to $A^\#$.

Here is another consequence of this duality: There exist $K_1^\#, \dots, K_t^\#$ s.t. for all δ , there is an N s.t. for all $n > N$, if $u \in U_n^r$, then for all k

$$\|u|V_{n+k}^{r-t}\| \leq K_t^\# \exp[-n(\chi_r - \chi_{r-t} - \delta)]. \quad (2.5)$$

To see this note that $V_{n+k}^{r-t} = (\tilde{V}_{n+k}^{s-r+t+1})^\#$ and $U_n^r \subset (V_n^{s-r+1})^\#$, and apply the technical lemma to the cocycle generated by $A^\#$.

We prove (2.4). Fix $\delta > 0$ and N large (we see how large later), and assume $n > N$. Suppose $v \in U_r$ and $\|v\| = 1$. Write $v = \lim v_{n+k}$ with $v_{n+k} := v|U_{n+k}^k \in U_{n+k}^r$. Note that $\|v_{n+k}\| \leq 1$. We decompose v_{n+k} as follows

$$v_{n+k} = (v_{n+k}|V_n^{r-1}) + (v_{n+k}|U_n^r) + \sum_{t=1}^{s-r} (v_{n+k}|U_n^{r+t}),$$

and estimate the size of the image of each of the summands under A_n .

First summand:

$$\begin{aligned} \|A_n(v_{n+k}|V_n^{r-1})\|^2 &\equiv \langle B_n^2(v_{n+k}|V_n^{r-1}), (v_{n+k}|V_n^{r-1}) \rangle \\ &= e^{2n(\chi_{r-1}+o(1))} \|v_{n+k}|V_n^{r-1}\|^2 \leq e^{2n(\chi_{r-1}+o(1))}. \end{aligned}$$

Thus the first summand is less than $\exp[n(\chi_{r-1} + o(1))]$.

Second Summand:

$$\begin{aligned}
\|A_n(v_{n+k}|U_n^r)\|^2 &= \langle B_n^t(v_{n+k}|U_n^r), (v_{n+k}|U_n^r) \rangle \\
&= e^{2n(\chi_r + o(1))} \|v_{n+k}|U_n^r\|^2 = e^{2n(\chi_r + o(1))} (\|v|U_n^r\| \pm (\|(v_{n+k} - v)|U_n^r\|))^2 \\
&= e^{2n(\chi_r \pm \delta)} (\|v|U_n^r\| \pm \|v_{n+k} - v\|)^2 = e^{2n(\chi_r \pm \delta)} [1 + o(1)] \text{ uniformly in } v.
\end{aligned}$$

Thus the second summand is $[1 + o(1)] \exp[n(\chi_r \pm \delta)]$ uniformly in $v \in U_r$, $\|v\| = 1$.

Third Summand: For every t ,

$$\begin{aligned}
\|A_n(v_{n+k}|U_n^t)\|^2 &= \langle B_n(v_{n+k}|U_n^t), (v_{n+k}|U_n^t) \rangle \\
&\leq e^{2n(\chi_{r+t} + o(1))} \|v_{n+k}|U_n^{r+t}\|^2 \\
&\equiv e^{2n(\chi_{r+t} + o(1))} \left(\sup_{u \in U_n^{r+t}, \|u\|=1} \langle v_{n+k}, u \rangle \right)^2, \text{ because } \|x|W\| = \sup_{w \in W, \|w\|=1} \langle x, w \rangle \\
&\leq e^{2n(\chi_{r+t} + o(1))} \left(\sup_{u \in U_n^{r+t}, \|u\|=1} \sup_{v \in V_{n+k}^r, \|v\| \leq 1} \langle v, u \rangle \right)^2 \\
&= e^{2n(\chi_{r+t} + o(1))} \sup_{u \in U_n^{r+t}, \|u\|=1} \|u|V_{n+k}^r\|^2 \\
&\leq (K_r^\#)^2 e^{2n(\chi_{r+t} + o(1))} \exp[-2n(\chi_r - \chi_{r+t} - o(1))], \text{ by (2.5)} \\
&= (K_r^\#)^2 e^{2n(\chi_r + o(1))}.
\end{aligned}$$

Note that the cancellation of χ_{r+t} — this is the essence of the technical lemma. We get: $\|A_n(v_{n+k}|U_n^t)\| = O(\exp[n(\chi_r + o(1))])$. Summing over $t = 1, \dots, s-r$, we get that third summand is $O(\exp[n(\chi_r + o(1))])$.

Putting these estimates together, we get that

$$\|A_n v_{n+k}\| \leq \text{const.} \exp[n(\chi_r + o(1))] \text{ uniformly in } k, \text{ and on the unit ball in } U_r.$$

“Uniformity” means that the $o(1)$ can be made independent of v and k . It allows us to pass to the limit as $k \rightarrow \infty$ and obtain

$$\|A_n v\| \leq \text{const.} \exp[n(\chi_r + o(1))] \text{ uniformly on the unit ball in } U_r.$$

On the other hand, an orthogonality argument shows that

$$\begin{aligned}
\|A_n v_{n+k}\|^2 &= \langle B_n^2 v_{n+k}, v_{n+k} \rangle \\
&= \|\text{1st summand}\|^2 + \|\text{2nd summand}\|^2 + \|\text{3rd summand}\|^2 \\
&\geq \|\text{2nd summand}\|^2 = [1 + o(1)] \exp[2n(\chi_r + o(1))].
\end{aligned}$$

Thus $\|A_n v_{n+k}\| \geq [1 + o(1)] \exp[n(\chi_r + o(1))]$ uniformly in v, k . Passing to the limit as $k \rightarrow \infty$, we get $\|A_n v\| \geq \text{const.} \exp[n(\chi_r + o(1))]$ uniformly on the unit ball in U_r . These estimates imply (2.4).

Proof of the technical lemma: We are asked to estimate the norm of the projection of a vector in V_n^r on V_{n+k}^{r+t} . We do this in three steps:

1. $V_n^r \rightarrow V_{n+1}^{r+t}$, all $t > 0$;
2. $V_n^r \rightarrow V_{n+k}^{r+1}$, all $k > 0$;
3. $V_n^r \rightarrow V_{n+k}^{r+t}$, all $t, k > 0$.

Step 1. The technical lemma for $k = 1$: Fix $\delta > 0$, then for all n large enough and for all $r' > r$, if $u \in V_n^{r'}$, then $\|u|V_{n+1}^{r'}\| \leq \|u\| \exp(-n(\chi_{r'} - \chi_r - \delta))$.

Proof. Fix ε , and choose $N = N(\varepsilon)$ so large that $t_n^i = e^{\pm n\varepsilon} t_i$ for all $n > N$, $i = 1, \dots, d$. For every $t = 1, \dots, s$, if $u \in V_n^r$, then

$$\begin{aligned} \|A_{n+1}u\| &= \sqrt{\langle A_{n+1}^t A_{n+1}u, u \rangle} \\ &= \sqrt{\langle A_{n+1}^t A_{n+1}(u|\tilde{V}_{n+1}^{r+t}), (u|\tilde{V}_{n+1}^{r+t}) \rangle + \langle A_{n+1}^t A_{n+1}(u|V_{n+1}^{r+t-1}), (u|V_{n+1}^{r+t-1}) \rangle} \\ &\quad \text{(because } V_{n+1}^{r+t-1}, \tilde{V}_{n+1}^{r+t} \text{ are orthogonal, } A_{n+1}^t A_{n+1} \text{-invariant,} \\ &\quad \text{and } \mathbb{R}^d = V_{n+1}^{r+t-1} \oplus \tilde{V}_{n+1}^{r+t}) \\ &= \sqrt{\|A_{n+1}(u|\tilde{V}_{n+1}^{r+t})\|^2 + \|A_{n+1}(u|V_{n+1}^{r+t-1})\|^2} \\ &\geq \|A_{n+1}(u|\tilde{V}_{n+1}^{r+t})\| = e^{(\chi_{r+t} \pm \varepsilon)(n+1)} \|u|\tilde{V}_{n+1}^{r+t}\|. \end{aligned}$$

On the other hand

$$\begin{aligned} \|A_{n+1}u\| &= \|A(T^n x)A_n(x)u\| \leq \|A(T^n x)\| \sqrt{\langle A_n^t A_n u, u \rangle} \\ &\leq \|A(T^n x)\| e^{n(\chi_r \pm \varepsilon)} \|u\| \\ &= e^{n(\chi_r \pm \varepsilon) + o(n)} \|u\|, \end{aligned}$$

because by the ergodic theorem

$$\frac{1}{n} \log \|A(T^n x)\| = \frac{1}{n} \sum_{k=0}^n \log \|A(T^k x)\| - \frac{1}{n} \sum_{k=0}^{n-1} \log \|A(T^k x)\| \xrightarrow[n \rightarrow \infty]{} 0 \text{ a.e.}$$

By further increasing N , we can arrange $|o(n)| < n\varepsilon$, which gives

$$e^{(\chi_{r+t} - \varepsilon)(n+1)} \|u|\tilde{V}_{n+1}^{r+t}\| \leq e^{n(\chi_r + 2\varepsilon)},$$

whence $\|u|\tilde{V}_{n+1}^{r+t}\| \leq e^{-n(\chi_{r+t} - \chi_r - 3\varepsilon)}$. Now take $\varepsilon := \delta/3$.

Step 2. Fix $\delta > 0$. Then for all n large enough and for all k , if $u \in V_n^r$, then $\|u|\tilde{V}_{n+k}^{r+1}\| \leq \|u\| \sum_{j=0}^{k-1} \exp(-(n+j)(\chi_{r+1} - \chi_r - \delta))$. Thus $\exists K_1$ s.t.

$$\|u|\tilde{V}_{n+k}^{r+1}\| \leq K_1 \|u\| \exp[-n(\chi_{r+1} - \chi_r - \delta)].$$

Proof. We use induction on k . The case $k = 1$ is dealt with in step 1. We assume by induction that the statement holds for $k - 1$, and prove it for k . Decompose

$$u|\widetilde{V}_{n+k}^{r+1} = [(u|V_{n+k-1}^r)|\widetilde{V}_{n+k}^{r+1}] + [(u|\widetilde{V}_{n+k-1}^{r+1})|\widetilde{V}_{n+k}^{r+1}].$$

- *First summand:* $u|V_{n+k-1}^r \in V_{n+k-1}^r$, so by step 1 the norm of the first summand is less than $\|u|V_{n+k-1}^r\| \exp[-(n+k-1)(\chi_{r+1} - \chi_r - \delta)]$, whence less than $\|u\| \exp[-(n+k-1)(\chi_{r+1} - \chi_r - \delta)]$.
- *Second summand:* The norm is at most $\|u|\widetilde{V}_{n+k-1}^{r+1}\|$. By the induction hypothesis, this is less than $\|u\| \sum_{j=0}^{k-2} \exp(-(n+j)(\chi_{r+1} - \chi_r - \delta))$.

We get the statement for k , and step 2 follows by induction.

As a result, we obtain the existence of a constant $K_1 > 1$ for which $u \in V_n^r$ implies $\|u|\widetilde{V}_{n+k}^{r+1}\| \leq K_1 \|u\| \exp(-n(\chi_{r+1} - \chi_r - \delta))$.

Step 3. $\exists K_1, \dots, K_{s-1} > 1$ s.t. for all n large enough and for all k , $u \in V_n^r$ implies $\|u|\widetilde{V}_{n+k}^{r+\ell}\| \leq K_\ell \|u\| \exp(-n(\chi_{r+\ell} - \chi_r - \ell\delta))$ ($\ell = 1, \dots, s-r$).

Proof. We saw that K_1 exists. We assume by induction that K_1, \dots, K_{t-1} exist, and construct K_t . Fix $0 < \delta_0 < \min_j \{\chi_{j+1} - \chi_j\} - \delta$; the idea is to first prove that if $u \in V_n^r$, then

$$\|u|\widetilde{V}_{n+k}^{r+t}\| \leq \|u\| \left(\sum_{\tau=1}^{t-1} K_\tau \right) \left(\sum_{j=0}^{k-1} e^{-\delta_0 j} \right) \left(\sum_{j=0}^{k-1} \exp[-(n+j)(\chi_{r+t} - \chi_r - t\delta)] \right) \quad (2.6)$$

Once this is done, step 3 follows with $K_t := (\sum_{\tau=1}^{t-1} K_\tau) (\sum_{j \geq 0} e^{-\delta_0 j})^2$.

We prove (2.6) using induction on k . When $k = 1$ this is because of step 1. Suppose, by induction, that (2.6) holds for $k - 1$. Decompose:

$$u|\widetilde{V}_{n+k}^{r+t} = \underbrace{u|V_{n+k-1}^r}_{A} |\widetilde{V}_{n+k}^{r+t} + \underbrace{\sum_{r < r' < r+t} u|U_{n+k-1}^{r'}}_B |\widetilde{V}_{n+k}^{r+t} + \underbrace{u|\widetilde{V}_{n+k-1}^{r+t}}_C |\widetilde{V}_{n+k}^{r+t}$$

- *Estimate of $\|A\|$:* By step 1, $\|A\| \leq \|u\| \exp(-(n+k-1)(\chi_{r+t} - \chi_r - \delta))$.
- *Estimate of $\|B\|$:* By step 1, and the induction hypothesis (on t):

$$\begin{aligned} \|B\| &\leq \sum_{r < r' < r+t} \|u|U_{n+k-1}^{r'}\| \exp(-(n+k-1)(\chi_{r+t} - \chi_{r'} - \delta)) \\ &\leq \sum_{r < r' < r+t} \|u|\widetilde{V}_{n+k-1}^{r'}\| \exp(-(n+k-1)(\chi_{r+t} - \chi_{r'} - \delta)) \\ &\leq \sum_{r < r' < r+t} K_{r'-r} \|u\| \exp(-n(\chi_{r'} - \chi_r - (r' - r)\delta)) \times \\ &\quad \times \exp(-(n+k-1)(\chi_{r+t} - \chi_{r'} - \delta)) \end{aligned}$$

$$\begin{aligned}
&\leq \|u\| \left(\sum_{t'=1}^{t-1} K_{t'} \right) e^{-(k-1)(\chi_{r+t} - \chi_{r'} - \delta)} \exp(-n(\chi_{r+t} - \chi_r - t\delta)) \\
&\leq \|u\| \left(\sum_{t'=1}^{t-1} K_{t'} \right) e^{-\delta_0(k-1)} \exp(-n(\chi_{r+t} - \chi_r - t\delta)).
\end{aligned}$$

- Estimate of $\|C\|$: $\|C\| \leq \|u\| \tilde{V}_{n+k-1}^{r+t}$. By the induction hypothesis on k ,

$$\|C\| \leq \|u\| \left(\sum_{t'=1}^{t-1} K_{t'} \right) \left(\sum_{j=0}^{k-2} e^{-\delta_0 j} \right) \left(\sum_{j=0}^{k-2} \exp[-(n+j)(\chi_{r+t} - \chi_r - t\delta)] \right).$$

It is not difficult to see that when we add these bounds for $\|C\|$, $\|B\|$ and $\|A\|$, the result is smaller than the RHS of (2.6) for k . This completes the proof by induction of (2.6). As explained above, step 3 follows by induction. \square

Corollary 2.3. Let $\chi_1(x) < \dots < \chi_{s(x)}(x)$ denote the logarithms of the (different) eigenvalues of $\Lambda(x)$. Let U_{χ_i} be the eigenspace of $\Lambda(x)$ corresponding to $\exp \chi_i$. Set $V_\chi := \bigoplus_{\chi' \leq \chi} U_{\chi'}$.

1. $\chi(x, v) := \lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_n(x)v\|$ exists a.s. and is invariant.
2. $\chi(x, v) = \chi_i$ on $V_{\chi_i} \setminus V_{\chi_{i-1}}$
3. If $\|A^{-1}\|, \|A\| \in L^\infty$, then $\frac{1}{n} \log |\det A_n(x)| = \sum k_i \chi_i$, where $k_i = \dim U_{\chi_i}$.

$\{\chi_i(x)\}$ are called the Lyapunov exponents of x . $\{V_{\chi_i}\}$ is called the Lyapunov filtration of x . Property (2) implies that $\{V_\chi\}$ is A -invariant: $A(x)V_\chi(x) = V_\chi(Tx)$. Property (3) is sometimes called regularity.

Remark: $V_{\chi_i} \setminus V_{\chi_{i-1}}$ is A -invariant, but if $A(x)$ is not orthogonal, then U_{χ_i} doesn't need to be A -invariant. When T is invertible, there is a way of writing $V_{\chi_i} = \bigoplus_{j \leq i} H_j$ so that $A(x)H_j(x) = H_j(Tx)$ and $\chi(x, \cdot) = \chi_j$ on $H_j(x)$, see the next section.

2.5.3 The Multiplicative Ergodic Theorem for Invertible Cocycles

Suppose $A : X \rightarrow \text{GL}(n, \mathbb{R})$. There is a unique extension of the definition of $A_n(x)$ to non-positive n 's, which preserves the cocycle identity: $A_0 := id$, $A_{-n} := (A_n \circ T^{-n})^{-1}$. (Start from $A_{n-n} = A_0 = id$ and use the cocycle identity.)

The following theorem establishes a compatibility between the Lyapunov spectra and filtrations of A_n and A_{-n} .

Theorem 2.10. Let (X, \mathcal{B}, m, T) be an invertible probability preserving transformation, and $A : X \rightarrow \text{GL}(d, \mathbb{R})$ a Borel function s.t. $\ln \|A(x)^{\pm 1}\| \in L^1$. There are invariant Borel functions $p(x)$, $\chi_1(x) < \dots < \chi_{p(x)}(x)$, and a splitting $\mathbb{R}^d = \bigoplus_{i=1}^{p(x)} H^i(x)$ s.t.

1. $A_n(x)H^i(x) = H^i(T^n x)$ for all $n \in \mathbb{Z}$

2. $\lim_{n \rightarrow \pm\infty} \frac{1}{|n|} \log \|A_n(x)v\| = \pm\chi_i(x)$ on the unit sphere in $H^i(x)$.
3. $\frac{1}{n} \log \sin \angle(H^i(T^n x), H^j(T^n x)) \xrightarrow{n \rightarrow \infty} 0$.

Proof. Fix x , and let $t_n^1 \leq \dots \leq t_n^d$ and $\bar{t}_n^1 \leq \dots \leq \bar{t}_n^d$ be the eigenvalues of $(A_n^t A_n)^{1/2}$ and $(A_{-n}^t A_{-n})^{1/2}$. Let $t_i := \lim(t_n^i)^{1/n}$, $\bar{t}_i := \lim(\bar{t}_n^i)^{1/n}$. These limits exist almost surely, and $\{\log t_i\}, \{\log \bar{t}_i\}$ are lists of the Lyapunov exponents of A_n and A_{-n} , repeated with multiplicity. The proof of the Oseledets theorem shows that

$$\begin{aligned}
\sum_{k=d-i+1}^d \log \bar{t}_i &\equiv \lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_{-n}^{\wedge i}\| \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(|\det A_{-n}| \cdot \|((A_{-n})^{-1})^{\wedge(d-i)}\| \right) \quad (\text{write using e.v.'s}) \\
&\equiv \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(|(\det A_n \circ T^{-n})|^{-1} \cdot \|A_n^{\wedge(d-i)} \circ T^{-n}\| \right) \\
&\equiv \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(|\det A_n|^{-1} \cdot \|A_n^{\wedge(d-i)}\| \right) \quad (\text{remark after Kingman's Theorem}) \\
&= \sum_{k=d-i+1}^d \log t_i - \sum_{k=1}^d \log t_i = - \sum_{k=1}^i \log t_i.
\end{aligned}$$

Since this is true for all i , $\log t_i = -\log \bar{t}_{d-i+1}$.

It follows that if the Lyapunov exponents of A_n are $\chi_1 < \dots < \chi_s$, then the Lyapunov exponents of A_{-n} are $-\chi_s < \dots < -\chi_1$.

Let $V^1(x) \subset V^2(x) \subset \dots \subset V^s(x)$ be the Lyapunov filtration of A_n :

$$V^i(x) := \{v : \lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_n(x)v\| \leq \chi_i(x)\}.$$

Let $\bar{V}^1(x) \supset \bar{V}^2(x) \supset \dots \supset \bar{V}^s(x)$ be the following *decreasing* filtration, given by

$$\bar{V}^i(x) := \{v : \lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_{-n}(x)v\| \leq -\chi_i(x)\}.$$

These filtrations are invariant: $A(x)V^i(x) = V^i(Tx)$, $A(x)\bar{V}^i(x) = \bar{V}^i(Tx)$.

Set $H^i(x) := V^i(x) \cap \bar{V}^i(x)$. We must have $A(x)H^i(x) = H^i(Tx)$.

We claim that $\mathbb{R}^d = \bigoplus H^i(x)$ almost surely. It is enough to show that for a.e. x , $\mathbb{R}^d = V^i(x) \oplus \bar{V}^{i+1}(x)$, because

$$\begin{aligned}
\mathbb{R}^d &\equiv \bar{V}^1 = \bar{V}^1 \cap [V^1 \oplus \bar{V}^2] && (V^1 \oplus \bar{V}^2 = \mathbb{R}^d) \\
&= H^1 \oplus [\bar{V}^1 \cap \bar{V}^2] = H^1 \oplus \bar{V}^2 && (\bar{V}^1 \supseteq \bar{V}^2) \\
&= H^1 \oplus [\bar{V}^2 \cap (V^2 \oplus \bar{V}^3)] && (V^2 \oplus \bar{V}^3 = \mathbb{R}^d) \\
&= H^1 \oplus H^2 \oplus \bar{V}^3 = \dots = H^1 \oplus \dots \oplus H^s.
\end{aligned}$$

Since the spectra of $\Lambda, \bar{\Lambda}$ agree with matching multiplicities, $\dim V^i + \dim \bar{V}^{i+1} = d$. Thus it is enough to show that $E := \{x : V^i(x) \cap \bar{V}^{i+1}(x) \neq \{0\}\}$ has zero measure for all i .

Assume otherwise, then by the Poincaré recurrence theorem, for almost every $x \in E$ there is a sequence $n_k \rightarrow \infty$ for which $T^{n_k}(x) \in E$. By the Oseledets theorem, for every $\delta > 0$, there is $N_\delta(x)$ such that for all $n > N_\delta(x)$,

$$\|A_n(x)u\| \leq \|u\| \exp[n(\chi_i + \delta)] \quad \text{for all } u \in V^i \cap \bar{V}^{i+1}, \quad (2.7)$$

$$\|A_{-n}(x)u\| \leq \|u\| \exp[-n(\chi_{i+1} - \delta)] \quad \text{for all } u \in V^i \cap \bar{V}^{i+1}. \quad (2.8)$$

If $n_k > N_\delta(x)$, then $A_{n_k}(x)u \in V^i(T^{n_k}x) \cap \bar{V}^{i+1}(T^{n_k}x)$ and $T^{n_k}(x) \in E$, so

$$\|u\| = \|A_{-n_k}(T^{n_k}x)A_{n_k}(x)u\| \leq \|A_{n_k}(x)u\| \exp[-n_k(\chi_{i+1} - \delta)],$$

whence $\|A_{n_k}(x)u\| \geq \|u\| \exp[n_k(\chi_{i+1} - \delta)]$. By (2.7),

$$\exp[n_k(\chi_{i+1} + \delta)] \leq \exp[n_k(\chi_i + \delta)],$$

whence $|\chi_{i+1} - \chi_i| < 2\delta$. But δ was arbitrary, and could be chosen to be much smaller than the gaps between the Lyapunov exponents. With this choice, we get a contradiction which shows that $m(E) = 0$.

Thus $\mathbb{R}^d = \bigoplus H^i(x)$. Evidently, $V^i = V^i \supseteq \bigoplus_{j \leq i} H^j$ and $V^i \cap \bigoplus_{j > i} H^j \subseteq V^i \cap \bar{V}^{i+1} = \{0\}$, so $V^i = \bigoplus_{j \leq i} H^j$. In the same way $\bar{V}^i = \bigoplus_{j \geq i} H^j$. It follows that $H^i \subset (V^i \setminus V^{i-1}) \cap (\bar{V}^i \setminus \bar{V}^{i+1})$. Thus $\lim_{n \rightarrow \pm\infty} \frac{1}{|n|} \log \|A_n v\| = \pm \chi_i$ on the unit sphere in H^i .

Next we study the angle between $H^i(x)$ and $\tilde{H}^i(x) := \bigoplus_{j \neq i} H^j(x)$. Pick a basis $(v_1^i, \dots, v_{m_i}^i)$ for $H^i(x)$. Pick a basis $(w_1^i, \dots, w_{m_i}^i)$ for $\tilde{H}^i(x)$. Since $A_n(x)$ is invertible, $A_k(x)$ maps $(v_1^i, \dots, v_{m_i}^i)$ onto a basis of $H^i(T^k x)$, and $(w_1^i, \dots, w_{m_i}^i)$ onto a basis of $\tilde{H}^i(T^k x)$. Thus if $v := \bigwedge v_j^i$, $w := \bigwedge w_j^i$, then

$$|\sin \angle(H^i(T^k x), \tilde{H}^i(T^k x))| \geq \frac{\|A_n(x)^{\wedge d}(v \wedge w)\|}{\|A_n(x)^{\wedge m_i} v\| \cdot \|A_n(x)^{\wedge (d-m_i)} w\|}.$$

We view $A_n^{\wedge p}$ as an invertible matrix acting on $\text{span}\{e_{i_1}^* \wedge \dots \wedge e_{i_p}^* : i_1 < \dots < i_p\}$ via $(A_n(x)e_{i_1})^* \wedge \dots \wedge (A_n(x)e_{i_p})^*$. It is clear

$$\Lambda_p(x) := \lim_{n \rightarrow \infty} ((A_n^{\wedge p})^* (A_n^{\wedge p}))^{1/2n} = \left(\lim_{n \rightarrow \infty} (A_n^* A_n)^{1/2n} \right)^{\wedge p} = \Lambda(x)^{\wedge p},$$

thus the eigenspaces of $\Lambda_p(x)$ are the tensor products of the eigenspaces of $\Lambda(x)$. This determines the Lyapunov filtration of $A_n(x)^{\wedge p}$, and implies – by Oseledets theorem – that if $v_j \in V_{\chi_{k(j)}} \setminus V_{\chi_{k(j)-1}}$, and v_1, \dots, v_k are linearly independent, then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \|A_n(x)^{\wedge p} \omega\| = \sum_{j=1}^p \chi_{k(j)}, \text{ for } \omega := v_1 \wedge \cdots \wedge v_p.$$

It follows that $\lim_{n \rightarrow \infty} \frac{1}{n} \log |\sin \angle(H^i(T^n x), \tilde{H}^i(T^n x))| \geq 0$. \square

Problems

2.1. The Mean Ergodic Theorem for Contractions

Suppose H is a Hilbert space, and $U : H \rightarrow H$ is a bounded linear operator such that $\|U\| \leq 1$. Prove that $\frac{1}{N} \sum_{k=0}^{N-1} U^k f$ converges in norm for all $f \in H$, and the limit is the projection of f on the space $\{f : Uf = f\}$.

2.2. Ergodicity as a mixing property

Prove that a ppt (X, \mathcal{B}, μ, T) is ergodic, iff for every $A, B \in \mathcal{B}$, $\frac{1}{N} \sum_{k=0}^{N-1} \mu(A \cap T^{-k}B) \xrightarrow{N \rightarrow \infty} \mu(A)\mu(B)$.

2.3. Use the pointwise ergodic theorem to show that any two different ergodic invariant probability measures for the same transformation are mutually singular.

2.4. Ergodicity and extremality

An invariant probability measure μ is called *extremal*, if it cannot be written in the form $\mu = t\mu_1 + (1-t)\mu_2$ where μ_1, μ_2 are different invariant probability measures, and $0 < t < 1$. Prove that an invariant measure is extremal iff it is ergodic, using the following steps.

1. Show that if E is a T -invariant set of non-zero measure, then $\mu(\cdot|E)$ is a T -invariant measure. Deduce that if μ is not ergodic, then it is not extremal.
2. Show that if μ is ergodic, and $\mu = t\mu_1 + (1-t)\mu_2$ where μ_i are invariant, and $0 < t < 1$, then
 - a. For every $E \in \mathcal{B}$, $\frac{1}{N} \sum_{k=0}^{N-1} 1_E \circ T^k \xrightarrow{N \rightarrow \infty} \mu(E)$ μ_i -a.e. ($i = 1, 2$).
 - b. Conclude that $\mu_i(E) = \mu(E)$ for all $E \in \mathcal{B}$ ($i = 1, 2$). This shows that ergodicity implies extremality.

2.5. Prove that the Bernoulli $(\frac{1}{2}, \frac{1}{2})$ -measure is the invariant probability measure for the adding machine (Problem 1.10), by showing that all cylinders of length n must have the same measure as $[0^n]$. Deduce from the previous problem that the adic machine is ergodic.

2.6. Explain why when $f \in L^2$, $\mathbb{E}(f|\mathcal{F})$ is the projection of f on $L^2(\mathcal{F})$. Prove:

1. If $\mathcal{F} = \{\emptyset, A, X \setminus A\}$, then $\mathbb{E}(1_B|\mathcal{F}) = \mu(B|A)1_A + \mu(B|A^c)1_{A^c}$
2. If $\mathcal{F} = \{\emptyset, X\}$, then $\mathbb{E}(f|\mathcal{F}) = \int f d\mu$
3. If $X = [-1, 1]$ with Lebesgue measure, and $\mathcal{F} = \{A : A \text{ is Borel and } -A = A\}$, then $\mathbb{E}(f|\mathcal{F}) = \frac{1}{2}[f(x) + f(-x)]$

2.7. Prove:

1. $f \mapsto \mathbb{E}(f|\mathcal{F})$ is linear, and a contraction in the L^1 -metric
2. $f \geq 0 \Rightarrow \mathbb{E}(f|\mathcal{F}) \geq 0$ a.e.
3. if φ is convex, then $\mathbb{E}(\varphi \circ f|\mathcal{F}) \leq \varphi(\mathbb{E}(f|\mathcal{F}))$
4. if h is \mathcal{F} -measurable, then $\mathbb{E}(hf|\mathcal{F}) = h\mathbb{E}(f|\mathcal{F})$
5. If $\mathcal{F}_1 \subset \mathcal{F}_2$, then $\mathbb{E}[\mathbb{E}(f|\mathcal{F}_2)|\mathcal{F}_1] = \mathbb{E}(f|\mathcal{F}_1)$

2.8. The Martingale Convergence Theorem (Doob)

Suppose (X, \mathcal{B}, μ) is a probability space, and $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots$ are σ -algebras all of which are contained in \mathcal{B} . Let $\mathcal{F} := \sigma(\bigcup_{n \geq 1} \mathcal{F}_n)$ (the smallest σ -algebra containing the union). If $f \in L^1$, then $\mathbb{E}(f|\mathcal{F}_n) \xrightarrow[n \rightarrow \infty]{} \mathbb{E}(f|\mathcal{F})$ a.e. and in L^1

Prove this theorem, using the following steps (W. Parry). It is enough to consider non-negative $f \in L^1$.

1. Prove that $\mathbb{E}(f|\mathcal{F}_n) \xrightarrow[n \rightarrow \infty]{} \mathbb{E}(f|\mathcal{F})$ using the following observations:
 - a. The convergence holds for all elements of $\bigcup_{n \geq 1} L^1(X, \mathcal{F}_n, \mu)$;
 - b. $\bigcup_{n \geq 1} L^1(X, \mathcal{F}_n, \mu)$ is dense in $L^1(X, \mathcal{F}, \mu)$.
2. Set $E_a := \{x : \max_{1 \leq n \leq N} \mathbb{E}(f|\mathcal{F}_n)(x) > a\}$. Show that $\mu(E_a) \leq \frac{1}{a} \int f d\mu$. (Hint: $E = \bigcup_{n \geq 1} \{x : \mathbb{E}(f|\mathcal{F}_n)(x) > \lambda, \text{ and } \mathbb{E}(f|\mathcal{F}_k)(x) \leq \lambda \text{ for } k = 1, \dots, n-1\}$.)
3. Prove that $\mathbb{E}(f|\mathcal{F}_n) \xrightarrow[n \rightarrow \infty]{} \mathbb{E}(f|\mathcal{F})$ a.e. for every non-negative $f \in L^1$, using the following steps. Fix $f \in L^1$. For every $\varepsilon > 0$, choose n_0 and $g \in L^1(X, \mathcal{F}_{n_0}, \mu)$ such that $\|\mathbb{E}(f|\mathcal{F}) - g\|_1 < \varepsilon$.
 - a. Show that $|\mathbb{E}(f|\mathcal{F}_n) - \mathbb{E}(f|\mathcal{F})| \leq \mathbb{E}(|f - g||\mathcal{F}_n)| + |\mathbb{E}(f|\mathcal{F}) - g|$ for all $n \geq n_0$. Deduce that

$$\begin{aligned} \mu \left[\limsup_{n \rightarrow \infty} |\mathbb{E}(f|\mathcal{F}_n) - \mathbb{E}(f|\mathcal{F})| > \sqrt{\varepsilon} \right] &\leq \mu \left[\sup_n \mathbb{E}(|f - g||\mathcal{F}_n)| > \frac{1}{2} \sqrt{\varepsilon} \right] \\ &\quad + \mu \left[|\mathbb{E}(f|\mathcal{F}) - g| > \frac{1}{2} \sqrt{\varepsilon} \right] \end{aligned}$$

- b. Show that $\mu \left[\limsup_{n \rightarrow \infty} |\mathbb{E}(f|\mathcal{F}_n) - \mathbb{E}(f|\mathcal{F})| > \sqrt{\varepsilon} \right] \xrightarrow[\varepsilon \rightarrow 0^+]{} 0$. (Hint: Prove first that for every L^1 function F , $\mu[|F| > a] \leq \frac{1}{a} \|F\|_1$.)
- c. Finish the proof.

2.9. Hopf's ratio ergodic theorem

Let (X, \mathcal{B}, μ, T) be a conservative ergodic mpt on a σ -finite measure space. If $f, g \in L^1$ and $\int g \neq 0$, then $\frac{\sum_{k=0}^{n-1} f \circ T^k}{\sum_{k=0}^{n-1} g \circ T^k} \xrightarrow[n \rightarrow \infty]{} \frac{\int f d\mu}{\int g d\mu}$ almost everywhere.

Prove this theorem using the following steps (R. Zweimüller). Fix a set $A \in \mathcal{B}$ s.t. $0 < \mu(A) < \infty$, and let $(A, \mathcal{B}_A, T_A, \mu_A)$ denote the induced system on A (problem 1.14). For every function F , set

$$S_n F := F + F \circ T + \cdots + F \circ T^{n-1}$$

$$S_n^A F := F + F \circ T_A + \cdots + F \circ T_A^{n-1}$$

1. Read problem 1.14, and show that a.e. x has an orbit which enters A infinitely many times. Let $0 < \tau_1(x) < \tau_2(x) < \cdots$ be the times when $T^{\tau}(x) \in A$.
2. Suppose $f \geq 0$. Prove that for every $n \in (\tau_{k-1}(x), \tau_k(x)]$ and a.e. $x \in A$,

$$\frac{(S_{k-1}^A f)(x)}{(S_{k-1}^A 1_A)(x)} \leq \frac{(S_n f)(x)}{(S_n 1_A)(x)} \leq \frac{(S_k^A f)(x)}{(S_k^A 1_A)(x)}.$$

3. Verify that $S_j^A 1_A = j$ a.e. on A , and show that $(S_n f)(x)/(S_n 1_A)(x) \xrightarrow[n \rightarrow \infty]{} \frac{1}{\mu(A)} \int f d\mu$ a.e. on A .
4. Finish the proof.

Notes for chapter 2

For a comprehensive reference to ergodic theorems, see [2]. The mean ergodic theorem was proved by von Neumann. The pointwise ergodic theorem was proved by Birkhoff. By now there are many proofs of this theorem. The one we use is taken from [1], where it is attributed to Kamae — who apparently found it using ideas from nonstandard analysis. The subadditive ergodic theorem was first proved by Kingman. The proof we give is due to Steele [5]. The multiplicative ergodic theorem is due to Oseledets. The proof we use is due to Raghunathan and Ruelle, and is taken from [4]. The Martingale convergence theorem (problem 2.8) is due to Doob. The proof sketched in problem 2.8 is taken from [2]. The proof of Hopf's ratio ergodic theorem sketched in problem 2.9 is due to R. Zweimüller and is taken from [6].

References

1. Keane, M.: *Ergodic theory and subshifts of finite type*. In *Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989)*, 35–70, Oxford Sci. Publ., Oxford Univ. Press, New York, 1991.
2. Krengel, U.: *Ergodic theorems*. de Gruyter Studies in Mathematics **6** 1985. viii+357pp.
3. Parry, W.: *Topics in ergodic theory*. Cambridge Tracts in Mathematics, **75** Cambridge University Press, Cambridge-New York, 1981. x+110 pp.
4. Ruelle, D.: *Ergodic theory of differentiable dynamical systems*. Inst. Hautes Études Sci. Publ. Math. No. **50** (1979), 27–58.
5. Steele, J. M.: *Kingman's subadditive ergodic theorem*. Ann. Inst. H. Poincaré Probab. Statist. **25** (1989), no. 1, 93–98.
6. Zweimüller, R.: *Hopf's ratio ergodic theorem by inducing*. Colloq. Math. **101** (2004), no. 2, 289–292.

Chapter 3

Spectral Theory

3.1 The spectral approach to ergodic theory

A basic problem in ergodic theory is to determine whether two ppt are measure theoretically isomorphic. This is done by studying *invariants*: properties, quantities, or objects which are equal for any two isomorphic systems. The idea is that if two ppt have different invariants, then they cannot be isomorphic. Ergodicity and mixing are examples of invariants for measure theoretic isomorphism.

An effective method for inventing invariants is to look for a *weaker* equivalence relation, which is better understood. Any invariant for the weaker equivalence relation is automatically an invariant for measure theoretic isomorphism. The spectral point of view is based on this approach.

The idea is to associate to the ppt (X, \mathcal{B}, μ, T) the operator $U_T : L^2(X, \mathcal{B}, \mu) \rightarrow L^2(X, \mathcal{B}, \mu)$, $U_T f = f \circ T$. This is an isometry of L^2 (i.e. $\|U_T f\|_2 = \|f\|_2$ and $\langle U_T f, U_T g \rangle = \langle f, g \rangle$). It is useful here to think of L^2 as a Hilbert space over \mathbb{C} .

Definition 3.1. Two ppt (X, \mathcal{B}, μ, T) , (Y, \mathcal{C}, ν, S) are called *spectrally isomorphic*, if their associated L^2 -isometries U_T and U_S are *unitarily equivalent*, namely if there exists a linear operator $W : L^2(X, \mathcal{B}, \mu) \rightarrow L^2(Y, \mathcal{C}, \nu)$ s.t.

1. W is invertible;
2. $\langle Wf, Wg \rangle = \langle f, g \rangle$ for all $f, g \in L^2(X, \mathcal{B}, \mu)$;
3. $WU_T = U_S W$.

It is easy to see that any two measure theoretically isomorphic ppt are spectrally isomorphic, but we will see later that there are Bernoulli schemes which are spectrally isomorphic but not measure theoretically isomorphic.

Definition 3.2. A property of ppt is called a *spectral invariant*, if whenever it holds for (X, \mathcal{B}, μ, T) , it holds for all ppt which are spectrally isomorphic to (X, \mathcal{B}, μ, T) .

Proposition 3.1. *Ergodicity and mixing are spectral invariants.*

Proof. Suppose (X, \mathcal{B}, μ, T) is a ppt, and let U_T be as above. The trick is to phrase ergodicity and mixing in terms of U_T .

Ergodicity is equivalent to the statement “all invariant L^2 -functions are constant”, which is the same as saying that $\dim\{f : U_T f = f\} = 1$. Obviously, this is a spectral invariant.

Mixing is equivalent to the following statement: $\dim\{f : U_T f = f\} = 1$, and

$$\langle f, U_T^n g \rangle \xrightarrow{n \rightarrow \infty} \langle f, 1 \rangle \overline{\langle g, 1 \rangle} \text{ for all } f, g \in L^2.$$

To see that this property is preserved by spectral isomorphisms, note that if $\dim\{f : U_T f = f\} = 1$, then any unitary equivalence W satisfies $W1 = c$ with $|c| = 1$. \square

The spectral point of view immediately suggests the following invariant.

Definition 3.3. Suppose (X, \mathcal{B}, μ, T) is a ppt. If $f : X \rightarrow \mathbb{C}$, $f \in L^2$ satisfies $f \circ T = \lambda f$, then we say that f is an *eigenfunction* and that λ is an *eigenvalue*. The *point spectrum* T is the set $H(T) := \{\lambda \in \mathbb{C} : f \circ T = \lambda f\}$.

$H(T)$ is a countable subgroup of the unit circle (problem 3.1). Evidently $H(T)$ is a spectral invariant of T .

It is easy to see using Fourier expansions that for the irrational rotation R_α , $H(R_\alpha) = \{\alpha^k : k \in \mathbb{Z}\}$ (problem 3.2), thus irrational rotations by different angles are non-isomorphic.

Here are other related invariants:

Definition 3.4. Given a ppt (X, \mathcal{B}, μ, T) , let $V_d := \overline{\text{span}\{\text{eigenfunctions}\}}$. We say that (X, \mathcal{B}, μ, T) has

1. *discrete spectrum* (sometime called *pure point spectrum*), if $V_d = L^2$,
2. *continuous spectrum*, if $V_d = \{\text{constants}\}$ (i.e. is smallest possible),
3. *mixed spectrum*, if $V_d \neq L^2, \{\text{constants}\}$.

Any irrational rotation has discrete spectrum (problem 3.2). Any mixing transformation has continuous spectrum, because a non-constant eigenfunction $f \circ T = \lambda f$ satisfies

$$\langle f, f \circ T^{n_k} \rangle \xrightarrow{n \rightarrow \infty} \|f\|_2^2 \neq \left(\int f\right)^2$$

along any $n_k \rightarrow \infty$ s.t. $\lambda^{n_k} \rightarrow 1$. (To see that $\|f\|_2^2 \neq \left(\int f d\mu\right)^2$ for all non-constant functions, apply Cauchy-Schwarz to $f - \int f$, or note that non-constant L^2 functions have positive variance.)

The invariant $H(T)$ is tremendously successful for transformations with discrete spectrum:

Theorem 3.1 (Discrete Spectrum Theorem). *Two ppt with discrete spectrum are measure theoretically isomorphic iff they have the same group of eigenvalues.*

But this invariant cannot distinguish transformations with continuous spectrum. In particular - it is unsuitable for the study of mixing transformations.

3.2 Weak mixing

3.2.1 Definition and characterization

We saw that if a transformation is mixing, then it does not have non-constant eigenfunctions. But the absence of non-constant eigenfunctions is not equivalent to mixing (see problems 3.8–3.10 for an example). Here we study the dynamical significance of this property. First we give it a name.

Definition 3.5. A ppt is called *weak mixing*, if every $f \in L^2$ s.t. $f \circ T = \lambda f$ a.e. is constant almost everywhere.

Theorem 3.2. *The following are equivalent for a ppt (X, \mathcal{B}, μ, T) on a Lebesgue space:*

1. *weak mixing;*
2. *for all $E, F \in \mathcal{B}$, $\frac{1}{N} \sum_{k=0}^{N-1} |\mu(E \cap T^{-k}F) - \mu(E)\mu(F)| \xrightarrow{N \rightarrow \infty} 0$;*
3. *for every $E, F \in \mathcal{B}$, $\exists \mathcal{N} \subset \mathbb{N}$ of density zero (i.e. $|\mathcal{N} \cap [1, N]|/N \xrightarrow{N \rightarrow \infty} 0$) s.t.*

$$\mu(E \cap T^{-n}F) \xrightarrow{\mathcal{N} \not\ni n \rightarrow \infty} \mu(E)\mu(F);$$
4. *$T \times T$ is ergodic.*

Proof. We prove (2) \Rightarrow (3) \Rightarrow (4) \Rightarrow (1). The remaining implication (1) \Rightarrow (2) requires additional preparation, and will be shown later.

The implication (2) \Rightarrow (3) is a general fact from calculus (Koopman–von Neumann Lemma): If a_n is a bounded sequence of non-negative numbers, then $\frac{1}{N} \sum_{n=1}^N a_n \rightarrow 0$ iff there is a set of zero density $\mathcal{N} \subset \mathbb{N}$ s.t. $a_n \xrightarrow{\mathcal{N} \not\ni n \rightarrow \infty} 0$ (Problem 3.3).

We show that (3) \Rightarrow (4). Let \mathcal{S} be the semi-algebra $\{E \times F : E, F \in \mathcal{B}\}$ which generates $\mathcal{B} \otimes \mathcal{B}$, and fix $E_i \times F_i \in \mathcal{S}$. By (3), $\exists \mathcal{N}_i \subset \mathbb{N}$ of density zero s.t.

$$\mu(E_i \cap T^{-n}F_i) \xrightarrow{\mathcal{N}_i \not\ni n \rightarrow \infty} \mu(E_i)\mu(F_i) \quad (i = 1, 2).$$

The set $\mathcal{N} = \mathcal{N}_1 \cup \mathcal{N}_2$ also has zero density, and

$$\mu(E_i \cap T^{-n}F_i) \xrightarrow{\mathcal{N} \not\ni n \rightarrow \infty} \mu(E_i)\mu(F_i) \quad (i = 1, 2).$$

Writing $m = \mu \times \mu$ and $S = T \times T$, we see that this implies that

$$m[(E_1 \times E_2) \cap S^{-n}(F_1 \times F_2)] \xrightarrow{\mathcal{N} \not\ni n \rightarrow \infty} m(E_1 \times F_1)m(E_2 \times F_2),$$

whence $\frac{1}{N} \sum_{k=0}^{N-1} m[(E_1 \times F_1) \cap S^{-k}(E_2 \times F_2)] \xrightarrow{N \rightarrow \infty} m(E_1 \times F_1)m(E_2 \times F_2)$. In summary, $\frac{1}{N} \sum_{k=0}^{N-1} m[A \cap S^{-k}B] \xrightarrow{N \rightarrow \infty} m(A)m(B)$ for all $A, B \in \mathcal{S}$.

Since \mathcal{S} generates $\mathcal{B} \otimes \mathcal{B}$ the above holds for all $A, B \in \mathcal{B} \otimes \mathcal{B}$, and this implies that $T \times T$ is ergodic.

Proof that (4) \Rightarrow (1): Suppose T were not weak mixing, then T has a non-constant eigenfunction f with eigenvalue λ . The eigenvalue λ has absolute value equal to one, because $|\lambda| \|f\|_2 = \|f \circ T\|_2 = \|f\|_2$. Thus

$$F(x, y) = f(x)\overline{f(y)}$$

is $T \times T$ -invariant. Since f is non-constant, F is non-constant, and we get a contradiction to the ergodicity of $T \times T$.

The proof that (1) \Rightarrow (2) is presented in the next section. \square

3.2.2 Spectral measures and weak mixing

It is convenient to introduce the following notation $U_T^n := (U_T^*)^{|n|}$ where $n < 0$, where U_T^* is the unique operator s.t. $\langle U_T^* f, g \rangle = \langle f, U_T g \rangle$ for all $g \in L^2$. This makes sense even if U_T is not invertible. The reader can check that when U_T is invertible, $U_T^{-1} = (U_T)^{-1}$, so that there is no risk of confusion.

We are interested in the behavior of $U_T^n f$ as $n \rightarrow \pm\infty$. To study it, it is enough to study $U_T : H_f \rightarrow H_f$, where $H_f := \overline{\text{span}}\{U_T^n f : n \in \mathbb{Z}\}$.

It turns out that $U_T : H_f \rightarrow H_f$ is unitarily equivalent to the operator $M : g(z) \mapsto zg(z)$ on $L^2(S^1, \mathcal{B}(S^1), \nu_f)$ where ν_f is some finite measure on S^1 , called the *spectral measure* of f , which contains all the information on $U_T : H_f \rightarrow H_f$.

To construct it, we need the following important tool from harmonic analysis. Recall that the n -th *Fourier coefficient* of μ is the number $\widehat{\mu}(n) = \int_{S^1} z^n d\mu$.

Theorem 3.3 (Herglotz). *A sequence $\{r_n\}_{n \in \mathbb{Z}}$ is the sequence of Fourier coefficients of a positive Borel measure on S^1 iff $r_{-n} = \overline{r_n}$ and $\{r_n\}$ is positive definite:*

$$\sum_{n,m=-N}^N r_{n-m} a_n \overline{a_m} \geq 0 \text{ for all sequences } \{a_n\} \text{ and } N. \text{ This measure is unique.}$$

It is easy to check that $r_n = \langle U_T^n f, f \rangle$ is positive definite (to see this expand $\langle \sum_{n=-N}^N a_n U_T^n f, \sum_{m=-N}^N a_m U_T^m f \rangle$ noting that $\langle U_T^n f, U_T^m f \rangle = \langle U_T^{n-m} f, f \rangle$).

Definition 3.6. Suppose (X, \mathcal{B}, μ, T) is a ppt, and $f \in L^2 \setminus \{0\}$. The *spectral measure* of f is the unique measure ν_f on S^1 s.t. $\langle f \circ T^n, f \rangle = \int_{S^1} z^n d\nu_f$ for $n \in \mathbb{Z}$.

Proposition 3.2. *Let $H_f := \overline{\text{span}}\{U_T^n f : n \in \mathbb{Z}\}$, then $U_T : H_f \rightarrow H_f$ is unitarily equivalent to the operator $g(z) \mapsto zg(z)$ on $L^2(S^1, \mathcal{B}(S^1), \nu_f)$.*

Proof. By the definition of the spectral measure,

$$\begin{aligned} \left\| \sum_{n=-N}^N a_n z^n \right\|_{L^2(\nu_f)}^2 &= \left\langle \sum_{n=-N}^N a_n z^n, \sum_{m=-N}^N a_m z^m \right\rangle = \sum_{n,m=-N}^N a_n \overline{a_m} \int_{S^1} z^{n-m} d\nu_f(z) \\ &= \sum_{n,m=-N}^N a_n \overline{a_m} \langle U_T^{n-m} f, f \rangle = \sum_{n,m=-N}^N a_n \overline{a_m} \langle U_T^n f, U_T^m f \rangle = \left\| \sum_{n=-N}^N a_n U_T^n f \right\|_{L^2(\mu)}^2 \end{aligned}$$

In particular, if $\sum_{n=-N}^N a_n U_T^n f = 0$ in $L^2(\mu)$, then $\sum_{n=-N}^N a_n z^n = 0$ in $L^2(\nu_f)$. It follows that $W : U_T^n f \mapsto z^n$ extends to a linear map from $\text{span}\{U_T^n f : n \in \mathbb{Z}\}$ to $L^2(\nu_f)$.

This map is an isometry, and it is bounded. It follows that W extends to an linear isometry $W : H_f \rightarrow L^2(\nu_f)$. The image of W contains all the trigonometric polynomials, therefore $W(H_f)$ is dense in $L^2(\nu_f)$. Since W is an isometry, its image is closed (exercise). It follows that W is an isometric bijection from H_f onto $L^2(\nu_f)$.

Since $(WU_T)[g(z)] = z[Wg(z)]$ on $\text{span}\{U_T^n f : n \in \mathbb{Z}\}$, $WU_T g(z) = zg(z)$ on H_f , and so W is the required unitary equivalence. \square

Proposition 3.3. *If T is weak mixing ppt on a Lebesgue space, then all the spectral measures of $f \in L^2$ s.t. $\int f = 0$ are non-atomic (this explains the terminology “continuous spectrum”).*

Proof. Suppose $f \in L^2$ has integral zero and that ν_f has an atom $\lambda \in S^1$. We construct an eigenfunction (with eigenvalue λ). Consider the sequence $\frac{1}{N} \sum_{n=0}^{N-1} \lambda^{-n} U_T^n f$. This sequence is bounded in norm, therefore has a weakly convergent subsequence (here we use the fact that L^2 is separable — a consequence of the fact that (X, \mathcal{B}, μ) is a Lebesgue space):

$$\frac{1}{N_k} \sum_{n=0}^{N_k-1} \lambda^{-n} U_T^n f \xrightarrow{N \rightarrow \infty} g.$$

The limit g must satisfy $\langle U_T g, h \rangle = \langle \lambda g, h \rangle$ (check!), therefore it must be an eigenfunction with eigenvalue λ .

But it could be that $g = 0$. We rule this out using the assumption that $\nu_f\{\lambda\} \neq 0$:

$$\begin{aligned} \langle g, f \rangle &= \lim_{k \rightarrow \infty} \frac{1}{N_k} \sum_{n=0}^{N_k-1} \lambda^{-n} \langle U_T^n f, f \rangle = \lim_{k \rightarrow \infty} \frac{1}{N_k} \sum_{n=0}^{N_k-1} \int \lambda^{-n} z^n d\nu_f(z) \\ &= \nu_f\{\lambda\} + \lim_{k \rightarrow \infty} \frac{1}{N_k} \sum_{n=0}^{N_k-1} \int_{S^1 \setminus \{\lambda\}} \lambda^{-n} z^n d\nu_f(z) \\ &= \nu_f\{\lambda\} + \lim_{k \rightarrow \infty} \int_{S^1 \setminus \{\lambda\}} \frac{1}{N_k} \frac{1 - \lambda^{-N_k} z^{N_k}}{1 - \lambda^{-1} z} d\nu_f(z). \end{aligned}$$

The limit is equal to zero, because the integrand tends to zero and is uniformly bounded (by one). Thus $\langle g, f \rangle = \nu_f\{\lambda\} \neq 0$, whence $g \neq 0$. \square

Lemma 3.1. *Suppose T is a ppt on a Lebesgue space. If T is weak mixing, then for*

$$\text{every } f \in L^2, \frac{1}{N} \sum_{k=0}^{N-1} |\int f \cdot f \circ T^k d\mu - (\int f d\mu)|^2 \xrightarrow{N \rightarrow \infty} 0.$$

Proof. It is enough to treat the case when $\int f d\mu = 0$. Let ν_f denote the spectral measure of f , then

$$\begin{aligned}
\frac{1}{N} \sum_{k=0}^{N-1} \left| \int f \cdot f \circ T^n d\mu \right|^2 &= \frac{1}{N} \sum_{k=0}^{N-1} |\langle U_T^n f, f \rangle|^2 = \frac{1}{N} \sum_{k=0}^{N-1} \left| \int_{S^1} z^n d\nu_f(z) \right|^2 \\
&= \frac{1}{N} \sum_{k=0}^{N-1} \left(\int_{S^1} z^n d\nu_f(z) \right) \overline{\left(\int_{S^1} z^n d\nu_f(z) \right)} \\
&= \frac{1}{N} \sum_{k=0}^{N-1} \int_{S^1} \int_{S^1} z^n \bar{w}^n d\nu_f(z) d\nu_f(w) \\
&= \int_{S^1} \int_{S^1} \frac{1}{N} \left(\sum_{k=0}^{N-1} z^n \bar{w}^n \right) d\nu_f(z) d\nu_f(w)
\end{aligned}$$

The integrand tends to zero and is bounded outside $\Delta := \{(z, w) : z = w\}$. If we can show that $(\nu_f \times \nu_f)(\Delta) = 0$, then it will follow that $\frac{1}{N} \sum_{k=0}^{N-1} \left| \int f \cdot f \circ T^n d\mu \right|^2 \xrightarrow{N \rightarrow \infty} 0$. This is indeed the case: T is weak mixing, so by the previous proposition ν_f is non-atomic, whence $(\nu_f \times \nu_f)(\Delta) = \int_{S^1} \nu_f\{w\} d\nu_f(w) = 0$ by Fubini-Tonelli.

It remains to note that by the Koopman - von Neumann theorem, for every bounded non-negative sequence a_n , $\frac{1}{N} \sum_{k=1}^N a_n^2 \rightarrow 0$ iff $\frac{1}{N} \sum_{k=1}^N a_n \rightarrow 0$, because both conditions are equivalent to saying that a_n converges to zero outside a set of indices of density zero. \square

We can now complete the proof of the theorem in the previous section:

Proposition 3.4. *If T is weak mixing, then for all $f, g \in L^2$,*

$$\frac{1}{N} \sum_{k=0}^{N-1} \left| \int g \cdot f \circ T^n d\mu - \left(\int f d\mu \right) \left(\int g d\mu \right) \right| \xrightarrow{N \rightarrow \infty} 0. \quad (3.1)$$

Proof. Assume first T is invertible, then U_T is invertible, with a bounded inverse (equal to $U_{T^{-1}}$). Fix $f \in L^2$, and set

$$S(f) := \overline{\text{span}}\{U_T^k f : k \in \mathbb{Z}\}.$$

Write $L^2 = S(f) + \{\text{constants}\} + [S(f) + \{\text{constants}\}]^\perp$.

1. Every $g \in S(f)$ satisfies (3.1), because $S(f)$ is generated by functions of the form $g := U_T^k f$, and these functions satisfy (3.1) by Lemma 3.1.
2. Every constant g satisfies (3.1) trivially.
3. Every $g \perp S(f) \oplus \{\text{constants}\}$ satisfies (3.1) because $\langle g, f \circ T^n \rangle$ is eventually zero.

It follows that every $g \in L^2$ satisfies (3.1).

Now consider the case of a non-invertible ppt. Let $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{T})$ be the natural extension. A close look at the definition of $\tilde{\mathcal{B}}$ shows that if $\tilde{f} : \tilde{X} \rightarrow \mathbb{R}$ is $\tilde{\mathcal{B}}$ -measurable, then the value of $\tilde{f}(\dots, x_{-1}, x_0, x_1, \dots)$ is completely determined by x_0 . Moreover, $\tilde{f} : \tilde{X} \rightarrow \mathbb{C}$ is of the form $f \circ \tilde{\pi}$ where f is \mathcal{B} -measurable. Thus every eigenfunction for \tilde{T} is a lift of an eigenfunction for T . It follows that if T is weak mixing, then \tilde{T} is weak mixing.

By the first part of the proof, \tilde{T} satisfies (3.1). Since T is a factor of \tilde{T} , it also satisfies (3.1). \square

3.3 The Koopman operator of a Bernoulli scheme

In this section we analyze the Koopman operator of an invertible Bernoulli scheme. The idea is to produce an orthonormal basis for L^2 which makes the action of U_T transparent.

We cannot expect to diagonalize U_T : Bernoulli schemes are mixing, so they have no non-constant eigenfunctions. But we shall see that we can get the following nice structure:

Definition 3.7. An invertible ppt is said to have *countable Lebesgue spectrum* if L^2 has an orthonormal basis of the form $\{1\} \cup \{f_{\lambda,j} : \lambda \in \Lambda, j \in \mathbb{Z}\}$ where Λ is countable, and $U_T f_{\lambda,j} = f_{\lambda,j+1}$ for all i, j .

The reason for the terminology is that the spectral measure of each $f_{\lambda,j}$ is proportional to the Lebesgue measure on S^1 (problem 3.6).

Example. The invertible Bernoulli scheme with probability vector $(\frac{1}{2}, \frac{1}{2})$ has countable Lebesgue spectrum.

Proof. The phase space is $X = \{0, 1\}^{\mathbb{Z}}$. Define for every finite non-empty $A \subset \mathbb{Z}$ the function $\varphi_A(x) := \prod_{j \in A} (-1)^{x_j}$. Define $\varphi_{\emptyset} := 1$. Then,

1. if $A \neq B$, then $\varphi_A \perp \varphi_B$;
2. $\text{span}\{\varphi_A : A \subset \mathbb{Z} \text{ finite}\}$ is algebra of functions which separates points, and contains the constants.

By the Stone-Weierstrass theorem, $\overline{\text{span}}\{\varphi_A : A \subset \mathbb{Z} \text{ finite}\} = L^2$, so $\{\varphi_A\}$ is an orthonormal basis of L^2 . This is called the *Fourier-Walsh system*.

Note that $U_T \varphi_A = \varphi_{A+1}$, where $A+1 := \{a+1 : a \in A\}$. Take Λ the set of equivalence classes of the relation $A \sim B \Leftrightarrow \exists c \text{ s.t. } A = c+B$. Let A_λ be a representative of $\lambda \in \Lambda$. The basis is $\{1\} \cup \{\varphi_{A_\lambda+n} : \lambda \in \Lambda, n \in \mathbb{Z}\} = \{\text{Fourier Walsh functions}\}$. \square

It is not easy to produce such bases for other Bernoulli schemes. But they exist. To prove this we introduce the following sufficient condition for countable Lebesgue spectrum, which turns out to be satisfied by many smooth dynamical systems:

Definition 3.8. An invertible ppt (X, \mathcal{B}, μ, T) is called a *K automorphism* if there is a σ -algebra $\mathcal{A} \subset \mathcal{B}$ s.t.

1. $T^{-1} \mathcal{A} \subset \mathcal{A}$;
2. \mathcal{A} generates \mathcal{B} : $\sigma(\bigcup_{n \in \mathbb{Z}} T^{-n} \mathcal{A}) = \mathcal{B} \pmod{\mu}$;¹
3. the tail of \mathcal{A} is trivial: $\bigcap_{n=0}^{\infty} T^{-n} \mathcal{A} = \{\emptyset, X\} \pmod{\mu}$.

¹ $\mathcal{F}_1 \subset \mathcal{F}_2 \pmod{\mu}$ is for all $F_1 \in \mathcal{F}_2$ there is a set $F_2 \in \mathcal{F}_2$ s.t. $\mu(F_1 \Delta F_2) = 0$, and $\mathcal{F}_1 = \mathcal{F}_2 \pmod{\mu}$ iff $\mathcal{F}_1 \subset \mathcal{F}_2 \pmod{\mu}$ and $\mathcal{F}_2 \subset \mathcal{F}_1 \pmod{\mu}$.

Proposition 3.5. *Every invertible Bernoulli scheme has the K property.*

Proof. Let $(S^{\mathbb{Z}}, \mathcal{B}(S^{\mathbb{Z}}), \mu, T)$ be a Bernoulli scheme, i.e. $\mathcal{B}(S^{\mathbb{Z}})$ is the sigma algebra generated by cylinders ${}_{-k}[a_{-k}, \dots, a_{\ell}] := \{x \in S^{\mathbb{Z}} : x_i = a_i \text{ } (-k \leq i \leq \ell)\}$, T is the left shift map, and $\mu({}_k[a_{-k}, \dots, a_{\ell}]) = p_{a_{-k}} \cdots p_{a_{\ell}}$.

Call a cylinder *non-negative*, if it is of the form ${}_0[a_0, \dots, a_n]$. Let \mathcal{A} be the sigma algebra generated by all non-negative cylinders. It is clear that $T^{-1}\mathcal{A} \subset \mathcal{A}$ and that $\bigcup_{n \in \mathbb{Z}} T^{-n}\mathcal{A}$ generates $\mathcal{B}(S^{\mathbb{Z}})$. We show that the measure of every element of $\bigcap_{n=0}^{\infty} T^{-n}\mathcal{A}$ is either zero or one. Probabilists call the elements of this intersection *tail events*. The fact that every tail event for a sequence of independent identically distributed random variables has probability zero or one is called ‘‘Kolmogorov’s zero–one law’’.

Two measurable sets A, B are called *independent*, if $\mu(A \cap B) = \mu(A)\mu(B)$. For Bernoulli schemes, any two cylinders with non-overlapping set of indices is independent (check). Thus for every cylinder B of length $|B|$,

$$B \text{ is independent of } T^{-|B|}A \text{ for all non-negative cylinders } A.$$

It follows that B is independent of every element of $T^{-|B|}\mathcal{A}$ (a monotone class theorem argument). Thus every cylinder B is independent of every element of $\bigcap_{n \geq 1} T^{-n}\mathcal{A}$. Thus every element of \mathcal{B} is independent of every element of $\bigcap_{n \geq 1} T^{-n}\mathcal{A}$ (another monotone class theorem argument).

This means that every $E \in \bigcap_{n \geq 1} T^{-n}\mathcal{A}$ is independent of itself. Thus $\mu(E) = \mu(E \cap E) = \mu(E)^2$, whence $\mu(E) = 0$ or 1 . \square

Proposition 3.6. *Every K automorphism on a non-atomic standard probability space has countable Lebesgue spectrum.*

Proof. Let (X, \mathcal{B}, μ, T) be a K automorphism of a non-atomic standard probability space. Since (X, \mathcal{B}, μ) is a non-atomic standard space, $L^2(X, \mathcal{B}, \mu)$ is (i) infinite dimensional, and (ii) separable.

Let \mathcal{A} be a sigma algebra in the definition of the K property. Set $V := L^2(X, \mathcal{A}, \mu)$. This is a closed subspace of $L^2(X, \mathcal{B}, \mu)$, and

1. $U_T(V) \subseteq V$, because $T^{-1}\mathcal{A} \subset \mathcal{A}$;
2. $\bigcup_{n \in \mathbb{Z}} U_T^n(V)$ is dense in $L^2(X, \mathcal{B}, \mu)$, because $\bigcup_{n \in \mathbb{Z}} T^{-n}\mathcal{A}$ generates \mathcal{B} , so every $B \in \mathcal{B}$ can be approximated by a finite disjoint union of elements of $\bigcup_{n \in \mathbb{Z}} T^{-n}\mathcal{A}$;
3. $\bigcap_{n=1}^{\infty} U_T^n(V) = \{\text{constant functions}\}$, because $\bigcap_{n \geq 1} T^{-n}\mathcal{A} = \{\emptyset, X\} \text{ mod } \mu$.

Now let $W := V \ominus U_T(V)$ (the orthogonal complement of $U_T(V)$ in V). For all $n > 0$, $U_T^n(W) \subset U_T^n(V) \subset U_T(V) \perp W$. Thus $W \perp U_T^n(W)$ for all $n > 0$. Since U_T^{-1} is an isometry, $W \perp U_T^n(W)$ for all $n < 0$. It follows that

$$L^2(X, \mathcal{B}, \mu) = \{\text{constants}\} \oplus \bigoplus_{n \in \mathbb{Z}} U_T^n(W) \quad (\text{orthogonal sum}).$$

If $\{f_{\lambda} : \lambda \in \Lambda\}$ is an orthonormal basis for W , then the above implies that

$$\{1\} \cup \{U_T^n f_\lambda : \lambda \in \Lambda\}$$

is an orthonormal basis of $L^2(X, \mathcal{B}, \mu)$ (check!).

This is almost the full countable Lebesgue spectrum property. It remains to show that $|\Lambda| = \aleph_0$. $|\Lambda| \leq \aleph_0$ because $L^2(X, \mathcal{B}, \mu)$ is separable. We show that Λ is infinite by proving $\dim(W) = \infty$. We use the following fact (to be proved later):

$$\forall N \exists A_1, \dots, A_N \in \mathcal{A} \text{ pairwise disjoint sets, with positive measure.} \quad (3.2)$$

Suppose we know this. Pick $f \in W \setminus \{0\}$ ($W \neq \{0\}$, otherwise $L^2 = \{\text{constants}\}$ and (X, \mathcal{B}, μ) is atomic). Set $w_i := f1_{A_i} \circ T$ with A_1, \dots, A_N as above, then (i) w_i are linearly independent (because they have disjoint supports); (ii) $w_i \in V$ (because $T^{-1}A_i \in T^{-1}\mathcal{A} \subset \mathcal{A}$, so w_i is \mathcal{A} -measurable); and (iii) $w_i \perp U_T(V)$ (check, using $f \in W$). It follows that $\dim(W) \geq N$. Since N was arbitrary, $\dim(W) = \infty$.

Here is the proof of (3.2). Since (X, \mathcal{B}, μ) is non-atomic, $\exists B_1, \dots, B_N \in \mathcal{B}$ pairwise disjoint with positive measure. By assumption, $\bigcup_{n \in \mathbb{Z}} T^n \mathcal{A}$ generates \mathcal{B} , thus we can approximate B_i arbitrarily well by elements of $\bigcup_{n \in \mathbb{Z}} T^n \mathcal{A}$. By assumption, $\mathcal{A} \subseteq T \mathcal{A}$. This means that we can approximate B_i arbitrarily well by sets from $T^n \mathcal{A}$ by choosing n sufficiently large. It follows that $L^2(X, T^n \mathcal{A}, \mu)$ has dimension at least N . This forces $T^n \mathcal{A}$ to contain at least N pairwise disjoint sets of positive measure. It follows that \mathcal{A} contains at least N pairwise disjoint sets of positive measure. \square

Corollary 3.1. *All systems with countable Lebesgue spectrum, whence all invertible Bernoulli schemes, are spectrally isomorphic.*

Proof. Problem 3.7. \square

But it is not true that all Bernoulli schemes are measure theoretically isomorphic. To prove this one needs new (non-spectral) invariants. Enter the *measure theoretic entropy*, which we discuss in the next chapter.

Problems

3.1. Suppose (X, \mathcal{B}, μ, T) is an ergodic ppt on a Lebesgue space, and let $H(T)$ be its group of eigenvalues.

1. show that if f is an eigenfunction, then $|f| = \text{const. a.e.}$, and that if $\lambda, \mu \in H(T)$, then so do $1, \lambda\mu, \lambda/\mu$.
2. Show that eigenfunctions of different eigenvalue are orthogonal. Deduce that $H(T)$ is a countable subgroup of the unit circle.

3.2. Prove that the irrational rotation R_α has discrete spectrum, and calculate $H(R_\alpha)$.

3.3. Koopman - von Neumann Lemma

Suppose a_n is a bounded sequence of non-negative numbers. Prove that $\frac{1}{N} \sum_{n=1}^N a_n \rightarrow$

0 iff there is a set of zero density $\mathcal{N} \subset \mathbb{N}$ s.t. $a_n \xrightarrow[\mathcal{N} \not\rightarrow \infty]{} 0$. Guidance: Fill in the details in the following argument.

1. Suppose $\mathcal{N} \subset \mathbb{N}$ has density zero and $a_n \xrightarrow[\mathcal{N} \not\rightarrow \infty]{} 0$, then $\frac{1}{N} \sum_{n=1}^N a_n \rightarrow 0$.
2. Now assume that $\frac{1}{N} \sum_{n=1}^N a_n \rightarrow 0$.
 - a. Show that $\mathcal{N}_m := \{k : a_k > 1/m\}$ form an increasing sequence of sets of density zero.
 - b. Fix $\varepsilon_i \downarrow 0$, and choose $k_i \uparrow \infty$ such that if $n > k_i$, then $(1/n)|\mathcal{N}_i \cap [1, n]| < \varepsilon_i$. Show that $\mathcal{N} := \bigcup_i \mathcal{N}_i \cap (k_i, k_{i+1}]$ has density zero.
 - c. Show that $a_n \xrightarrow[\mathcal{N} \not\rightarrow \infty]{} 0$.

3.4. Here is a sketch of an alternative proof of proposition 3.4, which avoids natural extensions (B. Parry). Fill in the details.

1. Set $H := L^2$, $V := \bigcap_{n \geq 0} U_T^n(H)$, and $W := H \ominus U_T H := \{g \in H, g \perp U_T H\}$.
 - a. $H = V \oplus [(U_T H)^\perp + (U_T^2 H)^\perp + \dots]$
 - b. $\{U_T^k H\}$ is decreasing, $\{(U_T^k H)^\perp\}$ is increasing.
 - c. $H = V \oplus \bigoplus_{k=1}^\infty U_T^k W$ (orthogonal space decomposition).
2. $U_T : V \rightarrow V$ has a bounded inverse (hint: use the fact from Banach space theory that any bounded linear operator between mapping one Banach space *onto* another Banach space which is one-to-one, has a bounded inverse).
3. (3.1) holds for any $f, g \in V$.
4. if $g \in U_T^k W$ for some k , then (3.1) holds for all $f \in L^2$.
5. if $g \in V$, but $f \in U_T^k W$ for some k , then (3.1) holds for f, g .
6. (3.1) holds for all $f, g \in L^2$.

3.5. Show that every invertible ppt with countable Lebesgue spectrum is mixing, whence ergodic.

3.6. Suppose (X, \mathcal{B}, μ, T) has countable Lebesgue spectrum. Show that $\{f \in L^2 : \int f = 0\}$ is spanned by functions f whose spectral measures ν_f are equal to the Lebesgue measure on S^1 .

3.7. Show that any two ppt with countable Lebesgue spectrum are spectrally isomorphic.

3.8. Cutting and Stacking and Chacon's Example

This is an example of a ppt which is weak mixing but not mixing. The example is a certain map of the unit interval, which preserves Lebesgue's measure. It is constructed using the method of "cutting and stacking" which we now explain.

Let $A_0 = [1, \frac{2}{3})$ and $R_0 := [\frac{2}{3}, 1]$ (thought of as reservoir).

Step 1: Divide A_0 into three equal subintervals of length $\frac{2}{9}$. Cut a subinterval B_0 of length $\frac{2}{9}$ from the left end of the reservoir.

- Stack the three thirds of A_0 one on top of the other, starting from the left and moving to the right.
- Stick B_0 between the second and third interval.
- Define a partial map f_1 by moving points vertically in the stack. The map is defined everywhere except on $R \setminus B_0$ and the top floor of the stack. It can be viewed as a partially defined map of the unit interval.

Update the reservoir: $R_1 := R \setminus B_0$. Let A_1 be the base of the new stack (equal to the rightmost third of A_0).

Step 2: Cut the stack vertically into three equal stacks. The base of each of these thirds has length $\frac{1}{3} \times \frac{2}{9}$. Cut an interval B_1 of length $\frac{1}{3} \times \frac{2}{9}$ from the left side of the reservoir R_1 .

- Stack the three stacks one on top of the other, starting from the left and moving to the right.
- Stick B_1 between the second stack and the third stack.
- Define a partial map f_2 by moving points vertically in the stack. This map is defined everywhere except the union of the top floor floor and $R_1 \setminus B_1$.

Update the reservoir: $R_2 := R_1 \setminus B_1$. Let A_2 be the base of the new stack (equal to the rightmost third of A_1).

Step 3: Cut the stack vertically into three equal stacks. The base of each of these thirds has length $\frac{1}{3^2} \times \frac{2}{9}$. Cut an interval B_2 of length $\frac{1}{3^2} \times \frac{2}{9}$ from the left side of the reservoir R_2 .

- Stack the three stacks one on top of the other, starting from the left and moving to the right.
- Stick B_2 between the second stack and the third stack.
- Define a partial map f_3 by moving points vertically in the stack. This map is defined everywhere except the union of the top floor floor and $R_2 \setminus B_2$.

Update the reservoir: $R_3 := R_2 \setminus B_2$. Let A_3 be the base of the new stack (equal to the rightmost third of A_2).

Continue in this manner, to obtain a sequence of partially defined maps f_n . There is a canonical way of viewing the intervals composing the stacks as of subintervals of the unit interval. Using this identification, we may view f_n as partially defined maps of the unit interval.

1. Show that f_n is measure preserving where it is defined (the measure is Lebesgue's measure). Calculate the Lebesgue measure of the domain of f_n .
2. Show that f_{n+1} extends f_n (i.e. the maps agree on the intersection of their domains). Deduce that the common extension of f_n defines an invertible probability preserving map of the open unit interval. This is Chacon's example. Denote it by (I, \mathcal{B}, m, T) .
3. Let ℓ_n denote the height of the stack at step n . Show that the sets $\{T^i(A_n) : i = 0, \dots, \ell_n, n \geq 1\}$ generate the Borel σ -algebra of the unit interval.

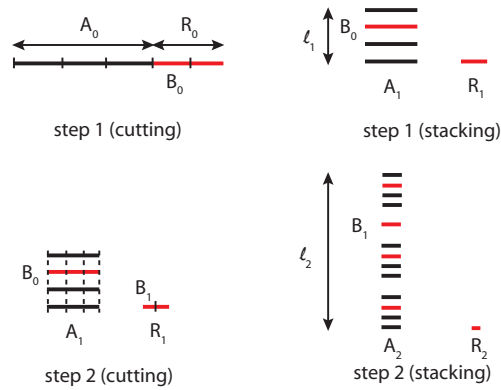


Fig. 3.1 The construction of Chacon's example

3.9. (Continuation) Prove that Chacon's example is weak mixing using the following steps. Suppose f is an eigenfunction with eigenvalue λ .

1. We first show that if f is constant on A_n for some n , then f is constant everywhere. (A_n is the base of the stack at step n .)
 - a. Let ℓ_k denote the height of the stack at step k . Show that $A_{n+1} \subset A_n$, and $T^{\ell_n}(A_{n+1}) \subset A_n$. Deduce that $\lambda^{\ell_n} = 1$.
 - b. Prove that $\lambda^{\ell_{n+1}} = 1$. Find a recursive formula for ℓ_n . Deduce that $\lambda = 1$.
 - c. The previous steps show that f is an invariant function. Show that any invariant function which constant on A_n is constant almost everywhere.
2. We now consider the case of a general L^2 -eigenfunction.
 - a. Show, using Lusin's theorem, that there exists an n such that f is nearly constant on most of A_n . (Hint: part 3 of the previous question).
 - b. Modify the argument done above to show that any L^2 -eigenfunction is constant almost everywhere.

3.10. (Continuation) Prove that Chacon's example is not mixing, using the following steps.

1. Inspect the image of the top floor of the stack at step n , and show that for every n and $0 \leq k \leq \ell_{n-1}$, $m(T^k A_n \cap T^{k+\ell_n} A_n) \geq \frac{1}{3} m(T^k A_n)$.

2. Use problem 3.8 part 3 and an approximation argument to show that for every Borel set E and $\varepsilon > 0$, $m(E \cap T^{\ell_n} E) \geq \frac{1}{3}m(E) - \varepsilon$ for all n . Deduce that T cannot be mixing.

Notes to chapter 3

The spectral approach to ergodic theory is due to von Neumann. For a thorough modern introduction to the theory, see Nadkarni's book [1]. Our exposition follows in parts the books by Parry [2] and Petersen [1]. A proof of the discrete spectrum theorem mentioned in the text can be found in Walters' book [3]. A proof of Herglotz's theorem is given in [2].

References

1. Nadkarni, M. G.: *Spectral theory of dynamical systems*. Birkhäuser Advanced Texts: Birkhäuser Verlag, Basel, 1998. x+182 pp.
2. Parry, W.: *Topics in ergodic theory*. Cambridge Tracts in Mathematics, 75. Cambridge University Press, Cambridge-New York, 1981. x+110 pp.
3. Petersen, K.: *Ergodic theory*. Corrected reprint of the 1983 original. Cambridge Studies in Advanced Mathematics 2 Cambridge University Press, Cambridge, 1989. xii+329 pp.
4. Walters, P.: *An introduction to ergodic theory*. Graduate Texts in Mathematics, 79 Springer-Verlag, New York-Berlin, 1982. ix+250 pp.

Chapter 4

Entropy

In the end of the last chapter we saw that every two Bernoulli schemes are spectrally isomorphic (because they have countable Lebesgue spectrum). The question whether any two Bernoulli schemes are measure theoretically isomorphic was a major open question in the field. It was solved by Kolmogorov and Sinai, through the invention of a new invariant: *entropy*. Later, Ornstein proved that this invariant is complete within the class of Bernoulli schemes.

4.1 Information content and entropy

Let $\alpha = \{A_1, \dots, A_N\}$ be a measurable partition of (X, \mathcal{B}, μ) , and suppose $T : X \rightarrow X$ is measurable. Let

$\alpha(x) :=$ The element of α which contains x .

The *itinerary* of x is $(\alpha(x), \alpha(Tx), \alpha(T^2x), \dots)$, a sequence taking values A_1, \dots, A_N .

Suppose x is not known, but $(\alpha(x), \dots, \alpha(T^{n-1}x))$ is known; How much uncertainty do we have regarding $\alpha(T^n x)$?

Example 1: Irrational Rotations. $R_\theta(z) = e^{i\theta}z$ with $0 < \theta < \frac{1}{100}$ irrational, and $\alpha := \{A_0, A_1\}$ where $A_0 := \{e^{i\theta} : 0 \leq \theta < \pi\}$, and $A_1 := \{e^{i\theta} : \pi \leq \theta < 2\pi\}$. If the five first symbols are

$(1, 1, 1, 0, 0, \dots)$

then we are *certain* that the next one is 0. This is the case whenever there is a 1 there. In the case $(0, 0, 0, 0, 0, \dots)$ we can guess that the next one is 0 with certainty of 99%. Thus knowing $\alpha(T^5x)$ isn't 'worth much', if we already know $(\alpha(x), \alpha(Tx), \dots, \alpha(T^4x))$. We can guess it with high certainty anyway.

Example 2: Angle Doubling. $T(z) = z^2$, same partition. Knowing the first five symbols tells us nothing on the sixth one: It is zero or one with probability 50%. So the 'information content' of $\alpha(T^5x)$ given $(\alpha(x), \alpha(Tx), \dots, \alpha(T^4x))$ is constant.

The following question arises: Let (X, \mathcal{B}, μ) be a probability space. Suppose $x \in X$ is unknown. How to quantify the ‘information content’ $I(A)$ of the the statement ‘ x belongs to A ’

Our guiding principle is to think of the information content of an event E as of the uncertainty lost when learning that $x \in A$. Thus the information content of an event of small probability is large. Here are some intuitively clear requirements that a good definition of $I(A)$ should satisfy:

1. $I(A)$ should be a continuous function of the probability of A ;
2. $I(A)$ should be non-negative, decreasing in $\mu(A)$, and if $\mu(A) = 1$ then $I(A) = 0$;
3. If A, B are independent, i.e. $\mu(A \cap B) = \mu(A)\mu(B)$, then $I(A \cap B) = I(A) + I(B)$.

Proposition 4.1. *The only functions $\varphi : [0, 1] \rightarrow \mathbb{R}^+$ such that $I(A) = \varphi[\mu(A)]$ satisfies the above axioms for all probability spaces (X, \mathcal{B}, μ) are $c \ln t$ with $c < 0$.*

We leave the proof as an exercise. This leads to the following definition.

Definition 4.1 (Shannon). Let (X, \mathcal{B}, μ) be a probability space.

1. The *Information Content* of a set $A \in \mathcal{B}$ is $I_\mu(A) := -\log \mu[A]$
2. The *Information Function* of a finite measurable partition is

$$I_\mu(\alpha)(x) := \sum_{A \in \alpha} I_\mu(A) 1_A(x) = - \sum_{A \in \alpha} \log \mu(A) 1_A(x)$$

3. The *Entropy* of a finite measurable partition is the average of the information content of its elements:

$$H_\mu(\alpha) := \int_X I_\mu(\alpha) d\mu = - \sum_{A \in \alpha} \mu(A) \log \mu(A).$$

Conventions: The base of the log is 2; $0 \log 0 = 0$.

There are important conditional versions of these notions:

Definition 4.2. Let (X, \mathcal{B}, μ) be a probability space, and suppose \mathcal{F} is a sub- σ -algebra of \mathcal{B} . We use the notation $\mu(A|\mathcal{F})(x) := \mathbb{E}(1_A|\mathcal{F})(x)$ (as L^1 -elements).

1. The *information content* of A given \mathcal{F} is $I_\mu(A|\mathcal{F})(x) := -\log \mu(A|\mathcal{F})(x)$
2. The *information function* of a finite measurable partition α given \mathcal{F} is $I_\mu(\alpha|\mathcal{F}) := \sum_{A \in \alpha} I_\mu(A|\mathcal{F}) 1_A$
3. The *conditional entropy* of α given \mathcal{F} is $H_\mu(\alpha|\mathcal{F}) := \int I_\mu(\alpha|\mathcal{F}) d\mu$.

Convention: Let α, β be partitions; We write $H_\mu(\alpha|\beta)$ for $H_\mu(\alpha|\sigma(\beta))$, where $\sigma(\beta) :=$ smallest σ -algebra which contains β .

The following formulæ are immediate:

$$H_\mu(\alpha|\mathcal{F}) = - \int_X \sum_{A \in \alpha} \mu(A|\mathcal{F})(x) \log \mu(A|\mathcal{F})(x) d\mu(x)$$

$$H_\mu(\alpha|\beta) = - \sum_{B \in \beta} \mu(B) \sum_{A \in \alpha} \mu(A|B) \log \mu(A|B), \text{ where } \mu(A|B) = \frac{\mu(A \cap B)}{\mu(B)}.$$

4.2 Properties of the entropy of a partition

We need some notation and terminology. Let α, β be two countable partitions.

1. $\sigma(\alpha)$ is the smallest σ -algebra which contains α ;
2. $\alpha \leq \beta$ means that $\alpha \subseteq \sigma(\beta) \pmod{\mu}$, i.e. every element of α is equal up to a set of measure zero to an element of $\sigma(\beta)$. Equivalently, $\alpha \leq \beta$ if every element of α is equal up to a set of measure zero to a union of elements of β . We say that β is *finer* than α , and that α is *coarser* than β .
3. $\alpha = \beta \pmod{\mu}$ iff $\alpha \subseteq \beta \pmod{\mu}$ and $\beta \subseteq \alpha \pmod{\mu}$.
4. $\alpha \vee \beta$ is the smallest partition which is finer than both α and β . Equivalently, $\alpha \vee \beta := \{A \cap B : A \in \alpha, B \in \beta\}$.

If $\mathcal{F}_1, \mathcal{F}_2$ are two σ -algebras, then $\mathcal{F}_1 \vee \mathcal{F}_2$ is the smallest σ -algebra which contains $\mathcal{F}_1, \mathcal{F}_2$.

4.2.1 The entropy of $\alpha \vee \beta$

It is useful to think of a partition $\alpha = \{A_1, \dots, A_n\}$ as of the “information” which element of α contains an unknown x .

We state and prove a formula which says that the information content of α and β is the information content of α plus the information content of β given the knowledge α .

Theorem 4.1 (The Basic Identity). *Suppose α, β are measurable countable partitions, and assume $H_\mu(\alpha), H_\mu(\beta) < \infty$, then*

1. $I_\mu(\alpha \vee \beta | \mathcal{F}) = I_\mu(\alpha | \mathcal{F}) + I_\mu(\beta | \mathcal{F} \vee \sigma(\alpha))$;
2. $H_\mu(\alpha \vee \beta) = H_\mu(\alpha) + H_\mu(\beta | \alpha)$.

Proof. We calculate $I_\mu(\beta | \mathcal{F} \vee \sigma(\alpha))$:

$$I_\mu(\beta | \mathcal{F} \vee \sigma(\alpha)) = - \sum_{B \in \beta} 1_B \log \mu(B | \mathcal{F} \vee \sigma(\alpha))$$

Claim: $\mu(B | \mathcal{F} \vee \sigma(\alpha)) = \sum_{A \in \alpha} 1_A \frac{\mu(B \cap A | \mathcal{F})}{\mu(A | \mathcal{F})}$.

1. This expression is $\mathcal{F} \vee \sigma(\alpha)$ -measurable
2. Observe that $\mathcal{F} \vee \sigma(\alpha) = \{\biguplus_{A \in \alpha} A \cap F_A : F_A \in \mathcal{F}\}$ (this is a σ -algebra which contains α and \mathcal{F}). Thus every $\mathcal{F} \vee \sigma(\alpha)$ -measurable function is of the form $\sum_{A \in \alpha} 1_A \varphi_A$ with φ_A \mathcal{F} -measurable. It is therefore enough to check test functions of the form $1_A \varphi$ with $\varphi \in L^\infty(\mathcal{F})$. For such functions

$$\begin{aligned}
\int 1_A \varphi \sum_{A' \in \alpha} 1_{A'} \frac{\mu(B \cap A' | \mathcal{F})}{\mu(A' | \mathcal{F})} d\mu &= \int 1_A \varphi \frac{\mu(B \cap A | \mathcal{F})}{\mu(A | \mathcal{F})} d\mu = \\
&= \int \mathbb{E}(1_A | \mathcal{F}) \varphi \frac{\mu(B \cap A | \mathcal{F})}{\mu(A | \mathcal{F})} d\mu \\
&= \int \varphi \mu(B | \mathcal{F}) d\mu = \int \varphi 1_B d\mu.
\end{aligned}$$

Using the claim, we see that

$$\begin{aligned}
I_\mu(\beta | \mathcal{F} \vee \alpha) &= - \sum_{B \in \beta} 1_B \log \sum_{A \in \alpha} 1_A \frac{\mu(B \cap A | \mathcal{F})}{\mu(A | \mathcal{F})} \\
&= - \sum_{B \in \beta} \sum_{A \in \alpha} 1_{A \cap B} \log \frac{\mu(B \cap A | \mathcal{F})}{\mu(A | \mathcal{F})} \\
&= - \sum_{B \in \beta} \sum_{A \in \alpha} 1_{A \cap B} \log \mu(B \cap A | \mathcal{F}) + \sum_{A \in \alpha} \sum_{B \in \beta} 1_{A \cap B} \log \mu(A | \mathcal{F}) \\
&= I_\mu(\alpha \vee \beta | \mathcal{F}) - I_\mu(\alpha | \mathcal{F}).
\end{aligned}$$

This proves the first part of the theorem.

Integrating, we get $H_\mu(\alpha \vee \beta | \mathcal{F}) = H_\mu(\alpha | \mathcal{F}) + H_\mu(\beta | \mathcal{F} \vee \alpha)$. If $\mathcal{F} = \{\emptyset, X\}$, then $H_\mu(\alpha \vee \beta) = H_\mu(\alpha) + H_\mu(\beta | \alpha)$. \square

4.2.2 Convexity properties

Lemma 4.1. *Let $\varphi(t) := -t \log t$, then for every probability vector (p_1, \dots, p_n) and $x_1, \dots, x_n \in [0, 1]$ $\varphi(p_1 x_1 + \dots + p_n x_n) \geq p_1 \varphi(x_1) + \dots + p_n \varphi(x_n)$, with equality iff all the x_i with i s.t. $p_i \neq 0$ are equal.*

Proof. This is because $\varphi(\cdot)$ is strictly concave. Let $m := \sum p_i x_i$. If $m = 0$ then the lemma is obvious, so suppose $m > 0$. It is an exercise in calculus to see that $\varphi(t) \leq \varphi(m) + \varphi'(m)(t - m)$ for $t \in [0, 1]$, with equality iff $t = m$. In the particular case $m = \sum p_i x_i$ and $t = x_i$ we get

$$p_i \varphi(x_i) \leq p_i \varphi(m) + \varphi'(m)(p_i x_i - p_i m) \text{ with equality iff } p_i = 0 \text{ or } x_i = m.$$

Summing over i , we get $\sum p_i \varphi(x_i) \leq \varphi(m) + \varphi'(m)(\sum p_i x_i - m) = \varphi(m)$. There is an equality iff for every i $p_i = 0$ or $x_i = m$. \square

Proposition 4.2 (Convexity properties). *Let α, β, γ be countable measurable partitions with finite entropies, then*

1. $\alpha \leq \beta \Rightarrow H_\mu(\alpha | \gamma) \leq H_\mu(\beta | \gamma)$
2. $\alpha \leq \beta \Rightarrow H_\mu(\gamma | \alpha) \geq H_\mu(\gamma | \beta)$

Proof. The basic identity shows that $\beta \vee \gamma$ has finite entropy, and so $H_\mu(\beta|\gamma) = H_\mu(\alpha \vee \beta|\gamma) = H_\mu(\alpha|\gamma) + H_\mu(\beta|\gamma \vee \alpha) \geq H_\mu(\alpha|\gamma)$.

For the second inequality, note that $\varphi(t) = -t \log t$ is strictly concave (i.e. its negative is convex), therefore by Jensen's inequality

$$\begin{aligned} H_\mu(\gamma|\alpha) &= \int \sum_{C \in \gamma} \varphi[\mathbb{E}(C|\sigma(\alpha))] d\mu = \int \sum_{C \in \gamma} \varphi[\mathbb{E}(\mathbb{E}(C|\sigma(\beta))|\sigma(\alpha))] d\mu \geq \\ &\geq \int \sum_{C \in \gamma} \mathbb{E}(\varphi[\mathbb{E}(1_C|\sigma(\beta))]) d\mu = \sum_{C \in \gamma} \int \varphi[\mathbb{E}(1_C|\sigma(\beta))] d\mu \equiv H_\mu(\gamma|\beta), \end{aligned}$$

proving the inequality. \square

4.2.3 Information and independence

We say that two partitions are *independent*, if $\forall A \in \alpha, B \in \beta, \mu(A \cap B) = \mu(A)\mu(B)$. This the same as saying that the random variables $\alpha(x), \beta(x)$ are independent.

Proposition 4.3 (Information and Independence). $H_\mu(\alpha \vee \beta) \leq H_\mu(\alpha) + H_\mu(\beta)$ with equality iff α, β are independent.

Proof. $H_\mu(\alpha \vee \beta) = H_\mu(\alpha) + H_\mu(\beta)$ iff $H_\mu(\alpha|\beta) = H_\mu(\alpha)$. But

$$H_\mu(\alpha|\beta) = - \sum_{B \in \beta} \mu(B) \sum_{A \in \alpha} \mu(A|B) \log \mu(A|B).$$

Let $\varphi(t) = -t \log t$. We have:

$$\sum_{A \in \alpha} \sum_{B \in \beta} \mu(B) \varphi[\mu(A|B)] = \sum_{A \in \alpha} \varphi[\mu(A)].$$

But φ is strictly concave, so $\sum_{B \in \beta} \mu(B) \varphi[\mu(A|B)] \leq \varphi[\mu(A)]$, with equality iff $\mu(A|B)$ are equal for all $B \in \beta$ s.t. $\mu(B) \neq 0$.

We conclude that $\mu(A|B) = c(A)$ for all $B \in \beta$ s.t. $\mu(B) \neq 0$. For such B , $\mu(A \cap B) = c(A)\mu(B)$. Summing over B , gives $c(A) = \mu(A)$ and we obtain the independence condition. \square

4.3 The Metric Entropy

4.3.1 Definition and meaning

Definition 4.3 (Kolmogorov, Sinai). The *metric entropy* of a ppt (X, \mathcal{B}, μ, T) is defined to be

$h_\mu(T) := \sup\{h_\mu(T, \alpha) : \alpha \text{ is a countable measurable partition s.t. } H_\mu(\alpha) < \infty\},$

where $h_\mu(T, \alpha) := \lim_{n \rightarrow \infty} \frac{1}{n} H_\mu\left(\bigvee_{i=0}^{n-1} T^{-i} \alpha\right).$

Proposition 4.4. *The limit which defines $h_\mu(T, \alpha)$ exists.*

It can be shown that the supremum is attained by *finite* measurable partitions (problem 4.9).

Proof. Write $\alpha_n := \bigvee_{i=0}^{n-1} T^{-i} \alpha$. Then $a_n := H_\mu(\alpha_n)$ is subadditive, because $a_{n+m} := H_\mu(\alpha_{n+m}) \leq H_\mu(\alpha_n) + H_\mu(T^{-n} \alpha_m) = a_n + a_m$.

We claim that any sequence of numbers $\{a_n\}_{n \geq 1}$ which satisfies $a_{n+m} \leq a_n + a_m$ converges to a limit (possibly equal to minus infinity), and that this limit is $\inf\{a_n/n\}$. Fix n . Then for every m , $m = kn + r$, $0 \leq r \leq n - 1$, so

$$a_m \leq ka_n + a_r.$$

Dividing by m , we get that for all $m > n$

$$\frac{a_m}{m} \leq \frac{ka_n + a_r}{kn + r} \leq \frac{a_n}{n} + \frac{a_r}{m},$$

whence $\limsup(a_m/m) \leq a_n/n$. Since this is true for all n , $\limsup a_m/m \leq \inf a_n/n$. But it is obvious that $\liminf a_m/m \geq \inf a_n/n$, so the limsup and liminf are equal, and their common value is $\inf a_n/n$.

We remark that in our case the limit is not minus infinity, because $H_\mu(\bigvee_{i=0}^{n-1} T^{-i} \alpha)$ are all non-negative. \square

$H_\mu(\alpha_n)$ is the average information content in the first n -digits of the α -itinerary. Dividing by n gives the average ‘information per unit time’. Thus the entropy measure the maximal rate of information production the system is capable of generating.

It is also possible to think of entropy as a measure of unpredictability. Let’s think of T as moving *backward* in time. Then $\alpha_1^\infty := \sigma(\bigcup_{n=1}^\infty T^{-n} \alpha)$ contains the information on the past of the itinerary. Given the future, how unpredictable is the present, on average? This is measured by $H_\mu(\alpha|\alpha_1^\infty)$.

Theorem 4.2. *If $H_\mu(\alpha) < \infty$, then $h_\mu(T, \alpha) = H_\mu(\alpha|\alpha_1^\infty)$, where $\alpha_1^\infty = \sigma(\bigcup_{n=1}^\infty T^{-n} \alpha)$.*

Proof. We show that $h_\mu(T, \alpha) = H_\mu(\alpha|\alpha_1^\infty)$. Observe that

$$H_\mu(\alpha|\alpha_0^n) = H_\mu(\alpha_0^{n-1}) - H_\mu(T^{-1} \alpha_0^{n-1}) = H_\mu(\alpha_0^n) - H_\mu(\alpha_0^{n-1}).$$

Summing over n , we obtain

$$H_\mu(\alpha_n) - H_\mu(\alpha) = \sum_{k=1}^n H_\mu(\alpha|\alpha_1^k)$$

Dividing by n and passing to the limit we get

$$h_\mu(T, \alpha) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n H_\mu(\alpha | \alpha_1^k)$$

It is therefore enough to show that $H_\mu(\alpha | \alpha_1^k) \xrightarrow[k \rightarrow \infty]{} H_\mu(\alpha | \alpha_1^\infty)$.

This is dangerous!! It is true that $H_\mu(\alpha | \alpha_1^k) = \int I_\mu(\alpha | \alpha_1^k) d\mu$ and that by the martingale convergence theorem

$$I_\mu(\alpha | \alpha_1^k) \xrightarrow[k \rightarrow \infty]{} I_\mu(\alpha | \alpha_1^\infty) \text{ a.e. .}$$

But the claim that the integral of the limit is equal to the limit of the integrals requires justification.

If $|\alpha| < \infty$, then we can bypass the problem by writing

$$H_\mu(\alpha | \alpha_1^k) = \int \sum_{A \in \alpha} \varphi[\mu(A | \alpha_1^k)] d\bar{\mu}, \text{ with } \varphi(t) = -t \log t,$$

and noting that *this* function is bounded (by $|\alpha| \max \varphi$). Thus the BCT applies and gives $H_\mu(\alpha | \alpha_1^k) \xrightarrow[k \rightarrow \infty]{} H_\mu(\alpha | \alpha_1^\infty)$.

If $|\alpha| = \infty$ (but $H_\mu(\alpha) < \infty$) then we need to be more clever, and appeal to the following lemma (proved below):

Lemma 4.2 (Chung–Neveu). *Suppose α is a countable measurable partition with finite entropy, then the function $f^* := \sup_{n \geq 1} I_\mu(\alpha | \alpha_1^n)$ is absolutely integrable.*

The result now follows from the dominated convergence theorem. \square

Here is the proof of the Chung Neveu Lemma. Fix $A \in \alpha$, then we may decompose $A \cap [f^* > t] = \bigsqcup_{m \geq 1} A \cap B_m(t; A)$, where

$$B_m(t; A) := \{x \in X : m \text{ is the minimal natural number s.t. } -\log_2 \mu(A | \alpha_1^m) > t\}.$$

We have

$$\begin{aligned} \mu[A \cap B_m(t; A)] &= \mathbb{E}_\mu(1_A 1_{B_m(t; A)}) = \mathbb{E}_\mu((1_A 1_{B_m(t; A)} | \sigma(\alpha_1^m))) \\ &= \mathbb{E}_\mu(1_{B_m(t; A)} \mathbb{E}_\mu(1_A | \sigma(\alpha_1^m))), \text{ because } B_m(t; A) \in \sigma(\alpha_1^m) \\ &\equiv \mathbb{E}_\mu(1_{B_m(t; A)} 2^{-\log_2 \mu(A | \sigma(\alpha_1^m))}) \\ &\leq \mathbb{E}_\mu(1_{B_m(t; A)} 2^{-t}) = 2^{-t} \mu[B_m(t; A)]. \end{aligned}$$

Summing over m we see that $\mu(A \cap [f^* > t]) \leq 2^{-t}$. Of course we also have $\mu(A \cap [f^* > t]) \leq \mu(A)$. Thus $\mu(A \cap [f^* > t]) \leq \min\{\mu(A), 2^{-t}\}$.

We now use the following fact from measure theory: If $g \geq 0$, then $\int g d\mu = \int_0^\infty \mu[g > t] dt$.¹

¹ Proof: $\int_X g d\mu = \int_X \int_0^\infty 1_{[0 \leq t < g(x)]}(x, t) dt d\mu(x) = \int_0^\infty \int_X 1_{[g > t]}(x, t) d\mu(x) dt = \int_0^\infty \mu[g > t] dt$.

$$\begin{aligned}
\int_A f^* d\mu &= \int_0^\infty \mu(A \cap [f^* > t]) dt = \int_0^\infty \min\{\mu(A), 2^{-t}\} dt \\
&\leq \int_0^{-\log_2 \mu(A)} \mu(A) dt + \int_{-\log_2 \mu(A)}^\infty 2^{-t} dt = -\mu(A) \log_2 \mu(A) - \frac{2^{-t}}{\ln 2} \Big|_{-\log_2 \mu(A)}^\infty \\
&= -\mu(A) \log_2 \mu(A) + \mu(A) / \ln 2.
\end{aligned}$$

Summing over $A \in \alpha$ we get that $\int f^* d\mu \leq H_\mu(\alpha) + (\ln 2)^{-1} < \infty$. \square

4.3.2 The Shannon–McMillan–Breiman Theorem

Theorem 4.3 (Shannon–McMillan–Breiman). *Let (X, \mathcal{B}, μ, T) be an ergodic ppt, and α a countable measurable partition of finite entropy, then*

$$\frac{1}{n} I_\mu(\alpha_0^{n-1}) \xrightarrow{n \rightarrow \infty} h_\mu(T, \alpha) \text{ a.e.}$$

In particular, if $\alpha_n(x) :=$ element of α_n which contains x , then $-\frac{1}{n} \log \alpha_n(x) \xrightarrow{n \rightarrow \infty} h_\mu(T, \alpha)$ a.e.

Proof. We start with the basic identity $I_\mu(\alpha_0^{n-1}) \equiv I_\mu(\alpha_1^{n-1} \vee \alpha) = I_\mu(\alpha_1^{n-1}) + I_\mu(\alpha | \alpha_1^{n-1})$. This gives

$$\begin{aligned}
I_\mu(\alpha_0^n) &= I_\mu(\alpha | \alpha_1^n) + I_\mu(\alpha_0^{n-1}) \circ T \\
&= I_\mu(\alpha | \alpha_1^n) + [I_\mu(\alpha | \alpha_1^{n-1}) + I_\mu(\alpha_0^{n-2}) \circ T] \circ T \\
&= \dots = \sum_{k=0}^{n-1} I_\mu(\alpha | \alpha_1^{n-k}) \circ T^k \\
&= \sum_{k=1}^n I_\mu(\alpha | \alpha_1^k) \circ T^{n-k}
\end{aligned}$$

By the Martingale Convergence Theorem, $I_\mu(\alpha | \alpha_1^k) \xrightarrow{k \rightarrow \infty} I_\mu(\alpha | \alpha_1^\infty)$. The idea of the proof is to use this to say

$$\begin{aligned}
\lim_{n \rightarrow \infty} \frac{1}{n} I_\mu(\alpha_n) &\equiv \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n I_\mu(\alpha | \alpha_1^k) \circ T^{n-k} \\
&\stackrel{?}{=} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n I_\mu(\alpha | \alpha_1^\infty) \circ T^{n-k} \equiv \frac{1}{n} \sum_{k=0}^{n-1} I_\mu(\alpha | \alpha_1^\infty) \circ T^k \\
&= \int I_\mu(\alpha | \alpha_1^\infty) d\mu \quad (\text{Ergodic Theorem}) \\
&= H_\mu(\alpha | \alpha_1^\infty) = h_\mu(T, \alpha).
\end{aligned}$$

The point is to justify the question mark. Write $f_n := I_\mu(\alpha|\alpha_1^n)$ and $f_\infty = I_\mu(\alpha|\alpha_1^\infty)$. It is enough to show

$$\int \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |f_k - f_\infty| \circ T^{n-k} d\mu = 0.$$

(This implies that the limsup is zero almost everywhere.) Set $F_n := \sup_{k>n} |f_k - f_\infty|$. Then $F_n \rightarrow 0$ almost everywhere. We claim that $F_n \rightarrow 0$ in L^1 . This is because of the dominated convergence theorem and the fact that $F_n \leq 2f^* := 2 \sup_m f_m \in L^1$ (Chung–Neveu Lemma). Fix some large N , then

$$\begin{aligned} & \int \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} |f_{n-k} - f_\infty| \circ T^k d\mu = \\ &= \int \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-N-1} |f_{n-k} - f_\infty| \circ T^k d\mu + \int \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=n-N}^{n-1} |f_{n-k} - f_\infty| \circ T^k d\mu \\ &\leq \int \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-N-1} F_N \circ T^k d\mu + \int \lim_{n \rightarrow \infty} \frac{1}{n} \left(\sum_{k=0}^{N-1} 2f^* \circ T^k \right) \circ T^{N-n} d\mu \\ &= \int F_N d\mu + \lim_{n \rightarrow \infty} \frac{1}{n} \int \sum_{k=0}^{N-1} 2f^* \circ T^k d\mu = \int F_N d\mu. \end{aligned}$$

Since $F_N \rightarrow 0$ in L^1 , $\int F_N d\mu \rightarrow 0$, and this proves that the integral of the limsup is zero. \square

4.3.3 Sinai's Generator theorem

Let $\mathcal{F}_1, \mathcal{F}_2$ be two sub- σ -algebras of a probability space (X, \mathcal{B}, μ) . We write $\mathcal{F}_1 \subseteq \mathcal{F}_2 \pmod{\mu}$, if $\forall F_1 \in \mathcal{F}_1, \exists F_2 \in \mathcal{F}_2$ s.t. $\mu(F_1 \triangle F_2) = 0$. We write $\mathcal{F}_1 = \mathcal{F}_2 \pmod{\mu}$, if both inclusions hold $\pmod{\mu}$. For example, $\mathcal{B}(\mathbb{R}) = \mathcal{B}_0(\mathbb{R}) \pmod{\text{Lebesgue}}$'s measure. For every partition α , let

$$\alpha_{-\infty}^\infty = \bigvee_{i=-\infty}^{\infty} T^{-i} \alpha, \quad \alpha_0^\infty := \bigvee_{i=0}^{\infty} T^{-i} \alpha$$

denote the smallest σ -algebras generated by, respectively, $\bigcup_{i=-\infty}^{\infty} T^{-i} \alpha$ and $\bigcup_{i=0}^{\infty} T^{-i} \alpha$.

Definition 4.4. A countable measurable partition α is called a *generator* for an invertible (X, \mathcal{B}, μ, T) if $\bigvee_{i=-\infty}^{\infty} T^{-i} \alpha = \mathcal{B} \pmod{\mu}$, and a *strong generator*, if $\bigvee_{i=0}^{\infty} T^{-i} \alpha = \mathcal{B} \pmod{\mu}$.

(This latter definition makes sense in the non-invertible case as well)

Example: $\alpha = \{[0, \frac{1}{2}), [\frac{1}{2}, 1]\}$ is a strong generator for $Tx = 2x \pmod 1$, because $\bigvee_{i=0}^{\infty} T^{-i}\alpha = \sigma(\bigcup_n \alpha_0^{n-1})$ is the Borel σ -algebra (it contains all dyadic intervals, whence all open sets).

Theorem 4.4 (Sinai's Generator Theorem). *Let (X, \mathcal{B}, μ, T) be a ppt. If α is a generator of finite entropy, then $h_\mu(T) = h_\mu(T, \alpha)$.*

Proof. Fix a finite measurable partition β ; Must show that $h_\mu(T, \beta) \leq h_\mu(T, \alpha)$.

Step 1. $h_\mu(T, \beta) \leq h_\mu(T, \alpha) + H_\mu(\beta|\alpha)$

$$\begin{aligned} \frac{1}{n} H_\mu(\beta_0^{n-1}) &= \frac{1}{n} [H_\mu(\alpha_0^{n-1}) + H_\mu(\beta_0^{n-1}|\alpha_0^{n-1})] \\ &\leq \frac{1}{n} \left[H_\mu(\alpha_0^{n-1}) + \sum_{k=0}^{n-1} H_\mu(T^{-k}\beta|\alpha_0^{n-1}) \right] \\ &\leq \frac{1}{n} \left[H_\mu(\alpha_0^{n-1}) + \sum_{k=0}^{n-1} H_\mu(T^{-k}\beta|T^{-k}\alpha) \right] \\ &= \frac{1}{n} H_\mu(\alpha_0^{n-1}) + H_\mu(\beta|\alpha). \end{aligned}$$

Now pass to the limit.

Step 2. For every N , $h_\mu(T, \beta) \leq h_\mu(T, \alpha) + H_\mu(\beta|\alpha_{-N}^N)$

Repeat the previous step with α_{-N}^N instead of α , and check that $h_\mu(T, \alpha_{-N}^N) = h_\mu(T, \alpha)$.

Step 3. $H_\mu(\beta|\alpha_{-N}^N) \xrightarrow{N \rightarrow \infty} H_\mu(\beta|\mathcal{B}) = 0$.

$$\begin{aligned} H_\mu(\beta|\alpha_{-N}^N) &= \int I_\mu(\beta|\alpha_{-N}^N) d\mu = - \sum_{B \in \beta} \int 1_B \log \mu(B|\alpha_{-N}^N) d\mu \\ &= - \sum_{B \in \beta} \int \mu(B|\alpha_{-N}^N) \log \mu(B|\alpha_{-N}^N) d\mu \\ &= \sum_{B \in \beta} \int \varphi[\log \mu(B|\alpha_{-N}^N)] d\mu \xrightarrow{N \rightarrow \infty} \sum_{B \in \beta} \int \varphi[\log \mu(B|\mathcal{B})] d\mu = 0, \end{aligned}$$

because $\mu(B|\mathcal{B}) = 1_B$, $\varphi[1_B] = 0$, and $|\beta| < \infty$

This proves that $h_\mu(T, \alpha) \geq \sup\{h_\mu(T, \beta) : |\beta| < \infty\}$. Problem 4.9 says that this supremum is equal to $h_\mu(T)$, so we are done. \square

4.4 Examples

4.4.1 Bernoulli schemes

Proposition 4.5. *The entropy of the Bernoulli shift with probability vector \underline{p} is $-\sum p_i \log p_i$. Thus the $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ -Bernoulli scheme and the $(\frac{1}{2}, \frac{1}{2})$ -Bernoulli scheme are not isomorphic.*

Proof. $\alpha = \{[1], \dots, [n]\}$ is a strong generator, and

$$H_\mu(\alpha_0^{n-1}) = - \sum_{x_0, \dots, x_{n-1}} p_{x_0} \cdots p_{x_{n-1}} (\log p_{x_0} + \cdots + \log p_{x_{n-1}}) = -n \sum p_i \log p_i.$$

□

4.4.2 Irrational rotations

Proposition 4.6. *The irrational rotation has entropy zero w.r.t. the Haar measure.*

Proof. The reason is that it is an invertible transformation with a strong generator. We first explain why any invertible map with a strong generator must have zero entropy. Suppose α is such a strong generator. Then

$$\begin{aligned} h_\mu(T, \alpha) &= H_\mu(\alpha | \alpha_1^\infty) = H_\mu(T\alpha | T(\alpha_1^\infty)) = \\ &= H_\mu(T\alpha | \alpha_0^\infty) = H_\mu(T\alpha | \mathcal{B}) = 0, \text{ because } T\alpha \subset \mathcal{B}. \end{aligned}$$

We now claim that $\alpha := \{A_0, A_1\}$ (the two halves of the circle) is a strong generator. It is enough to show that for every ε , $\bigcup_{n \geq 1} \alpha_0^{n-1}$ contains open covers of the circle by open arcs of diameter $< \varepsilon$ (because this forces α_0^∞ to contain all open sets).

It is enough to manufacture one arc of diameter less than ε , because the translations of this arc by $k\alpha$ will eventually cover the circle.

But such an arc necessarily exists: Choose some n s.t. $n\alpha \bmod 1 \in (0, \varepsilon)$. Then $A_1 \setminus T^{-n}A_1 = (A_1 \setminus [A_1 - n\alpha])$ is an arc of diameter less than ε .

4.4.3 Markov chains

Proposition 4.7. *Suppose μ is a translation invariant Markov measure with transition matrix $P = (p_{ij})$ and probability vector (p_i) . Then $h_\mu(\sigma) = -\sum p_i p_{ij} \log p_{ij}$.*

Proof. The natural partition $\alpha = \{[a] : a \in S\}$ is a strong generator.

$$\begin{aligned}
H_\mu(\alpha_0^n) &= - \sum_{\xi_0, \dots, \xi_n \in S} \mu[\underline{\xi}] \log \mu[\underline{\xi}] \\
&= - \sum_{\xi_0, \dots, \xi_n \in S} p_{\xi_0} p_{\xi_0 \xi_1} \cdots p_{\xi_{n-1} \xi_n} (\log p_{\xi_0} + \log p_{\xi_0 \xi_1} + \cdots + \log p_{\xi_{n-1} \xi_n}) \\
&= - \sum_{j=0}^{n-1} \sum_{\xi_0, \dots, \xi_n \in S} p_{\xi_0} p_{\xi_0 \xi_1} \cdots p_{\xi_{n-1} \xi_n} \log p_{\xi_j \xi_{j+1}} \\
&\quad - \sum_{\xi_0, \dots, \xi_n \in S} p_{\xi_0} p_{\xi_0 \xi_1} \cdots p_{\xi_{n-1} \xi_n} \log p_{\xi_0} \\
&= - \sum_{j=0}^{n-1} \sum_{\xi_0, \dots, \xi_n \in S} p_{\xi_0} p_{\xi_0 \xi_1} \cdots p_{\xi_{j-1} \xi_j} \cdot p_{\xi_{j+1} \xi_{j+2}} \cdots p_{\xi_{n-1} \xi_n} p_{\xi_j \xi_{j+1}} \log p_{\xi_j \xi_{j+1}} \\
&\quad - \sum_{\xi_0 \in S} p_{\xi_0} \log p_{\xi_0} \\
&= - \sum_{j=0}^{n-1} \sum_{\xi_j, \dots, \xi_n \in S} \mu(\sigma^{-j}[\underline{\xi}_j]) \cdot p_{\xi_{j+1} \xi_{j+2}} \cdots p_{\xi_{n-1} \xi_n} p_{\xi_j \xi_{j+1}} \log p_{\xi_j \xi_{j+1}} \\
&\quad - \sum_{\xi_0 \in S} p_{\xi_0} \log p_{\xi_0} \\
&= - \sum_{j=0}^{n-1} \sum_{\xi_j, \xi_{j+1} \in S} p_{\xi_j} p_{\xi_j \xi_{j+1}} \log p_{\xi_j \xi_{j+1}} \sum_{\xi_{j+2}, \dots, \xi_{n-1} \in S} p_{\xi_{j+1} \xi_{j+2}} \cdots p_{\xi_{n-1} \xi_n} \\
&\quad - \sum_{\xi_0 \in S} p_{\xi_0} \log p_{\xi_0} \\
&= - \sum_{j=0}^{n-1} \sum_{\xi_j, \xi_{j+1} \in S} p_{\xi_j} p_{\xi_j \xi_{j+1}} \log p_{\xi_j \xi_{j+1}} - \sum_{\xi_0 \in S} p_{\xi_0} \log p_{\xi_0} \\
&= n \left(- \sum_{i,j} p_i p_{ij} \log p_{ij} \right) - \sum_i p_i \log p_i
\end{aligned}$$

Now divide by $n + 1$ and pass to the limit. □

4.4.4 Expanding Markov Maps of the Interval

Theorem 4.5. Suppose $T : [0, 1] \rightarrow [0, 1]$ and $\alpha = \{I_1, \dots, I_N\}$ is a partition into intervals s.t.

1. α is a Markov partition
2. The restriction of T to α is C^1 , monotonic, and $|T'| > \lambda > 1$
3. T has an invariant measure μ .

Then $h_\mu(T) = - \int \log \frac{d\mu}{d\mu \circ T} d\mu$, where $(\mu \circ T)(E) = \sum_{A \in \alpha} \mu[T(A \cap E)]$.

Proof. One checks that the elements of α_0^{n-1} are all intervals of length $O(\lambda^{-n})$. Therefore α is a strong generator, whence

$$h_\mu(T) = h_\mu(T, \alpha) = H_\mu(\alpha|\alpha_1^\infty) = \int I_\mu(\alpha|\alpha_1^\infty) d\mu.$$

We calculate $I_\mu(\alpha|\alpha_1^\infty)$.

First note that $\alpha_1^\infty = T^{-1}(\alpha_0^\infty) = T^{-1}\mathcal{B}$, thus $I_\mu(\alpha|\alpha_1^\infty) = - \sum_{A \in \alpha} 1_A \log \mu(A|T^{-1}\mathcal{B})$.

We need to calculate $\mathbb{E}(\cdot|T^{-1}\mathcal{B})$. For this purpose, introduce the operator $\widehat{T} : L^1 \rightarrow L^1$ given by

$$(\widehat{T}f)(x) = \sum_{Ty=x} \frac{d\mu}{d\mu \circ T}(y) f(y).$$

Exercise: Verify: $\forall \varphi \in L^\infty$ and $f \in L^1$, $\int \varphi \widehat{T}f d\mu = \int \varphi \circ T \cdot f d\mu$.

We claim that $\mathbb{E}(f|T^{-1}\mathcal{B}) = (\widehat{T}f) \circ T$. Indeed, the $T^{-1}\mathcal{B}$ -measurable functions are exactly the functions of the form $\varphi \circ T$ with φ \mathcal{B} -measurable; Therefore $(\widehat{T}f) \circ T$ is $T^{-1}\mathcal{B}$ -measurable, and

$$\int \varphi \circ T \cdot \widehat{T}f \circ T d\mu = \int \varphi \cdot \widehat{T}f d\mu = \int \varphi \circ T \cdot f d\mu,$$

proving the identity.

We can now calculate and see that

$$\begin{aligned} I_\mu(\alpha|\alpha_1^\infty) &= - \sum_{A \in \alpha} 1_A(x) \log \mathbb{E}(1_A|T^{-1}\mathcal{B})(x) \\ &= - \sum_{A \in \alpha} 1_A(x) \log \sum_{Ty=x} \frac{d\mu}{d\mu \circ T}(y) 1_A(y) \equiv - \sum_{A \in \alpha} 1_A(x) \log \frac{d\mu}{d\mu \circ T}(x) \\ &= - \log \frac{d\mu}{d\mu \circ T}(x). \end{aligned}$$

We conclude that $h_\mu(T) = - \int \log \frac{d\mu}{d\mu \circ T}(x) d\mu(x)$. □

4.5 Abramov's Formula

Suppose (X, \mathcal{B}, μ, T) is a ppt. A set A is called *spanning*, if $X = \bigcup_{n=0}^{\infty} T^{-n}A \pmod{\mu}$. If T is ergodic, then every set of positive measure is spanning.

Theorem 4.6 (Abramov). *Suppose (X, \mathcal{B}, μ, T) is a ppt on a Lebesgue space, let A be a spanning measurable set, and let $(A, \mathcal{B}_A, \mu_A, T_A)$ be the induced system, then $h_{\mu_A}(T_A) = \frac{1}{\mu(A)} h_\mu(T)$.*

Proof. (Scheller) We prove the theorem in the case when T is invertible. The non-invertible case is handled by passing to the natural extension.

The idea is to show, for as many partitions α as possible, that $h_\mu(T, \alpha) = \mu(A)h_{\mu_A}(T_A, \alpha \cap A)$, where $\alpha \cap A := \{E \cap A : E \in \alpha\}$. As it turns out, this is the case for all partitions s.t. (a) $H_\mu(\alpha) < \infty$; (b) $A^c \in \alpha$; and (c) $\forall n, T_A[\varphi_A = n] \in \sigma(\alpha)$. Here, as always, $\varphi_A(x) := 1_A(x) \inf\{n \geq 1 : T^n x \in A\}$ (the first return time).

To see that there are such partitions, we let

$$\xi_A := \{A^c\} \cup T_A \eta_A, \text{ where } \eta_A := \{[\varphi_A = n] : n \in \mathbb{N}\}$$

(the coarsest possible) and show that $H_\mu(\xi_A) < \infty$. A routine calculation shows that $H_\mu(\xi_A) = H_\mu(\{A, A^c\}) + \mu(A)H_{\mu_A}(T_A \eta_A) \leq 1 + H_{\mu_A}(\eta_A)$. It is thus enough to show that $-\sum p_n \log_2 p_n < \infty$, where $p_n := \mu_A[\varphi_A = n]$. This is because $\sum n p_n = 1/\mu(A)$ (Kac formula) and the following fact from calculus: probability vectors with finite expectations have finite entropy.²

Assume now that α is a partition which satisfies (a)–(c) above. We will use throughout the following fact:

$$A, A^c, [\varphi_A = n] \in \alpha_1^\infty. \quad (4.1)$$

Here is why: $[\varphi_A = n] = T^{-n} T_A[\varphi_A = n] \in T^{-n} \alpha \subset \alpha_1^\infty$. Since $AA = \bigcup_{n \geq 1} [\varphi_A = n]$, we automatically have $A, A^c \in \alpha_1^\infty$.

Let α be a finite entropy countable measurable partition of X such that A^c is an atom of α and such that $\alpha \geq \xi_A$. In what follows we use the notation $A \cap \alpha := \{B \cap A : B \in \alpha\}$, $\alpha_1^\infty \cap A := \{B \cap A : B \in \alpha_1^\infty\}$. Since $H_\mu(\alpha) < \infty$,

$$\begin{aligned} h_\mu(T, \alpha) &= H_\mu(\alpha | \alpha_1^\infty) \\ &= \int \sum_{B \in A \cap \alpha} 1_B \log \mu(B | \alpha_1^\infty) d\mu + \int 1_{A^c} \log \mu(A^c | \alpha_1^\infty) d\mu \\ &= \int \sum_{B \in A \cap \alpha} 1_B \log \mu(B | \alpha_1^\infty) d\mu, \text{ because } A^c \in (\xi_A)_1^\infty \subseteq \alpha_1^\infty \\ &= \mu(A) \int_A \sum_{B \in A \cap \alpha} 1_B \log \mu_A(B | \alpha_1^\infty) d\mu_A, \end{aligned}$$

because $A \in \alpha_1^\infty$ and $B \subset A$ imply $\mathbb{E}_\mu(1_B | \mathcal{F}) = 1_A \mathbb{E}_{\mu_A}(1_B | A \cap \mathcal{F})$.

It follows that $h_\mu(T, \alpha) = \mu(A)H_{\mu_A}(A \cap \alpha | A \cap \alpha_1^\infty)$. We will show later that

$$A \cap \alpha_1^\infty = \bigvee_{i=1}^{\infty} T_A^{-i}(A \cap \alpha) \quad (4.2)$$

This implies that $h_\mu(T, \alpha) = \mu(A)h_{\mu_A}(T_A, A \cap \alpha)$. Passing to the supremum over all α which contain A^c as an atom, we obtain

² Proof: Enumerate (p_n) in a decreasing order: $p_{n_1} \geq p_{n_2} \geq \dots$. If $C = \sum n p_n$, then $C \geq \sum_{i=1}^k n_i p_{n_i} \geq p_{n_k}(1 + \dots + k)$, whence $p_{n_k} = O(k^{-2})$. Since $-x \log x = O(x^{1-\varepsilon})$ as $x \rightarrow 0^+$, this means that $-p_{n_k} \log p_{n_k} = O(k^{-(2-\varepsilon)})$, and so $-\sum p_n \log p_n = -\sum p_{n_k} \log p_{n_k} < \infty$.

$$\begin{aligned}\mu(A)h_{\mu_A}(T_A) &= \sup\{h_{\mu}(T, \alpha) : \alpha \geq \xi_A, A^c \in \alpha, H_{\mu}(\alpha) < \infty\} \\ &= \text{Entropy of } (X, \mathcal{B}', \mu, T), \mathcal{B}' := \sigma\left(\bigcup\{\alpha_{-\infty}^{\infty} : A^c \in \alpha, H_{\mu}(\alpha) < \infty\}\right).\end{aligned}$$

(See problem 4.11).

Now $\mathcal{B}' = \mathcal{B} \bmod \mu$, because A is spanning, so $\forall E \in \mathcal{B}, E = \bigcup_{n=0}^{\infty} T^{-n}(T^n E \cap A) \bmod \mu$, whence $E \in \mathcal{B}' \bmod \mu$. This shows Abramov's formula, given (4.2).

The proof of (4.2):

Proof of \subseteq : Suppose B is an atom of $A \cap \bigvee_{j=1}^n T^{-j}\alpha$, then B has the form $A \cap \bigcap_{j=1}^n T^{-j}A_j$ where $A_j \in \alpha$. Let $j_1 < j_2 < \dots < j_N$ be an enumeration of the j 's s.t. $A_j \subset A$ (possibly an empty list). Since A^c is an atom of α , $A_j = A^c$ for j not in this list, and so $B = \bigcap_{k=1}^{N-1} T_A^{-k}(A_{j_k} \cap [\varphi_A = j_{k+1} - j_k]) \cap T_A^{-N}[\varphi_A > n - j_N]$. Since $\eta_A \leq \alpha \cap A$, $B \in \bigvee_{i=1}^{\infty} T_A^{-1}(\alpha \cap A)$.

Proof of \supseteq : $T_A^{-1}(\alpha \cap A) \leq A \cap \bigvee_{i=1}^{\infty} T^{-i}\alpha$, because if $B \in \alpha \cap A$, then

$$T_A^{-1}B = \bigcup_{n=1}^{\infty} T^{-n}(B \cap T_A[\varphi_A = n]) \in \bigvee_{n=1}^{\infty} T^{-n}\alpha \quad (\because T_A \eta_A \leq \xi_A \leq A \cap \alpha).$$

The same proof shows that $T_A^{-1}(T^{-n}\alpha \cap A) \leq A \cap \bigvee_{i=1}^{\infty} T^{-i}\alpha$. It follows that

$$T_A^{-2}(\alpha \cap A) \leq T_A^{-1}\left(A \cap \bigvee_{i=1}^{\infty} T^{-i}\alpha\right) \subseteq A \cap \bigvee_{i=1}^{\infty} T_A^{-1}(A \cap T^{-i}\alpha) \subseteq A \cap \bigvee_{i=1}^{\infty} T^{-i}\alpha.$$

Iterating this procedure we see that $T_A^{-n}(\alpha \cap A) \leq A \cap \bigvee_{i=1}^{\infty} T^{-i}\alpha$ for all n , and \supseteq follows. \square

4.6 Topological Entropy

Suppose $T : X \rightarrow X$ is a continuous mapping of a compact topological space (X, d) . Such a map can have many different invariant Borel probability measures. For example, the left shift on $\{0, 1\}^{\mathbb{N}}$ has an abundance of Bernoulli measures, Markov measures, and there are many others.

Different measures may have different entropies. What is the largest value possible? We study this question in the context of continuous maps on topological spaces which are *compact* and *metric*.

4.6.1 The Adler–Konheim–McAndrew definition

Let (X, d) be a compact metric space, and $T : X \rightarrow X$ a continuous map. Some terminology and notation:

1. an *open cover* of X is a collection of open sets $\mathcal{U} = \{U_\alpha : \alpha \in \Lambda\}$ s.t. $X = \bigcup_{\alpha \in \Lambda} U_\alpha$;
2. if $\mathcal{U} = \{U_\alpha : \alpha \in \Lambda\}$ is an open cover, then $T^{-k}\mathcal{U} := \{T^{-k}U_\alpha : \alpha \in \Lambda\}$. Since T is continuous, this is another open cover.
3. if \mathcal{U}, \mathcal{V} are open covers, then $\mathcal{U} \vee \mathcal{V} := \{U \cap V : U \in \mathcal{U}, V \in \mathcal{V}\}$.

Since X is compact, every open cover of X has a finite subcover. Define

$$N(\mathcal{U}) := \min\{\#\mathcal{V} : \mathcal{V} \subseteq \mathcal{U} \text{ is finite, and } X = \bigcup \mathcal{V}\}.$$

It is easy to check that $N(\cdot)$ is subadditive in the following sense:

$$N(\mathcal{U} \vee \mathcal{V}) \leq N(\mathcal{U}) + N(\mathcal{V}).$$

Definition 4.5. Suppose $T : X \rightarrow X$ is a continuous mapping of a compact metric space (X, d) , and let \mathcal{U} be an open cover of X . The *topological entropy* of \mathcal{U} is

$$h_{\text{top}}(T, \mathcal{U}) := \lim_{n \rightarrow \infty} \frac{1}{n} \log_2 N(\mathcal{U}_0^{n-1}), \text{ where } \mathcal{U}_0^{n-1} := \bigvee_{i=0}^{n-1} T^{-i}\mathcal{U}.$$

The limit exists because of the subadditivity of $N(\cdot)$: $a_n := \log N(\mathcal{U}_0^{n-1})$ satisfies $a_{m+n} \leq a_m + a_n$, so $\lim a_n/n$ exists.

Definition 4.6. Suppose $T : X \rightarrow X$ is a continuous mapping of a compact metric space (X, d) , then the *topological entropy* of T is the (possibly infinite)

$$h_{\text{top}}(T) := \sup\{h_{\text{top}}(T, \mathcal{U}) : \mathcal{U} \text{ is an open cover of } X\}.$$

The following theorem was first proved by Goodwyn.

Theorem 4.7. Suppose T is a continuous mapping of a compact metric space, then every invariant Borel probability measure μ satisfies $h_\mu(T) \leq h_{\text{top}}(T)$.

Proof. Eventually everything boils down to the following inequality, which can be checked using Lagrange multipliers: For every probability vector (p_1, \dots, p_k) ,

$$-\sum_{i=1}^k p_i \log_2 p_i \leq \log k, \quad (4.3)$$

with equality iff $p_1 = \dots = p_k = 1/k$.

Suppose μ is an invariant probability measure, and let $\alpha := \{A_1, \dots, A_k\}$ be a measurable partition.

We approximate α by a partition into sets with better topological properties. Fix $\varepsilon > 0$ (to be determined later), and construct *compact* sets

$$B_j \subset A_j \text{ s.t. } \mu(A_j \setminus B_j) < \varepsilon \quad (j = 1, \dots, k).$$

Let $B_0 := X \setminus \bigcup_{j=1}^k B_j$ be the remainder (of measure less than $k\varepsilon$), and define $\beta = \{B_0, B_1, \dots, B_k\}$.

Step 1 in the proof of Sinai's theorem says that $h_\mu(T, \alpha) \leq h_\mu(T, \beta) + H_\mu(\alpha|\beta)$. We claim that $H_\mu(\alpha|\beta)$ can be made uniformly bounded by a suitable choice of $\varepsilon = \varepsilon(\alpha)$:

$$\begin{aligned} H_\mu(\alpha|\beta) &= - \sum_{A \in \alpha} \sum_{B \in \beta} \mu(A \cap B) \log_2 \mu(A|B) \\ &= - \sum_{B \in \beta \setminus \{B_0\}} \sum_{A \in \alpha} \mu(A \cap B) \log_2 \mu(A|B) - \sum_{A \in \alpha} \mu(A \cap B_0) \log_2 \mu(A|B_0) \\ &= \sum_{i=1}^k \mu(B_i) \log_2 1 - \sum_{A \in \alpha} \mu(A \cap B_0) \log_2 \mu(A|B_0) \\ &= -\mu(B_0) \sum_{A \in \alpha} \mu(A|B_0) \log_2 \mu(A|B_0) \leq -\mu(B_0) \log(\#\alpha) \leq k\varepsilon \cdot \log_2 k. \end{aligned}$$

If we choose $\varepsilon < 1/(k \log_2 k)$, then we get $H_\mu(\alpha|\beta) \leq 1$, and

$$h_\mu(T, \alpha) \leq h_\mu(T, \beta) + 1. \quad (4.4)$$

We now create an open cover from β by setting $\mathcal{U} := \{B_0 \cup B_1, \dots, B_0 \cup B_k\}$. This is a cover. To see that it is open note that

$$\begin{aligned} B_0 \cup B_j &= B_0 \cup (A_j \setminus B_0) \quad (\because A_j \cap B_0 = A_j \setminus B_j) \\ &= B_0 \cup A_j = B_0 \cup \left(X \setminus \bigcup_{i \neq j} A_i \right) = B_0 \cup \left(X \setminus \bigcup_{i \neq j} B_i \right). \end{aligned}$$

We compare the number of elements in \mathcal{U}_0^{n-1} to the number of elements in β_0^{n-1} . Every element of \mathcal{U}_0^{n-1} is of the form

$$(B_0 \uplus B_{i_0}) \cap T^{-1}(B_0 \uplus B_{i_1}) \cap \dots \cap T^{-(n-1)}(B_0 \uplus B_{i_{n-1}}).$$

This can be written as a pairwise disjoint union of 2^n elements of β_0^{n-1} (some of which may be empty sets). Thus every element of \mathcal{U}_0^{n-1} contains at most 2^n elements of β_0^{n-1} . Forming the union over a sub cover of \mathcal{U}_0^{n-1} with cardinality $N(\mathcal{U}_0^{n-1})$, we get that $\#\beta_0^{n-1} \leq 2^n N(\mathcal{U}_0^{n-1})$.

We now appeal to (4.3): $H_\mu(\beta_0^{n-1}) \leq \log_2(\#\beta_0^{n-1}) \leq H(\mathcal{U}_0^{n-1}) + n$. Dividing by n and passing to the limit as $n \rightarrow \infty$, we see that $h_\mu(T, \beta) \leq h_{\text{top}}(\mathcal{U}) + 1$. By (4.4), $h_\mu(T, \alpha) \leq h_{\text{top}}(\mathcal{U}) + 2 \leq h_{\text{top}}(T) + 2$.

Passing to the supremum over all α , we get that $h_\mu(T) \leq h_{\text{top}}(T) + 2$, and this holds for all continuous mappings T and invariant Borel measures μ . In particular, this holds for T^n (note that $\mu \circ (T^n)^{-1} = \mu$): $h_\mu(T^n) \leq h_{\text{top}}(T^n) + 2$. But $h_\mu(T^n) = nh_\mu(T)$ and $h_{\text{top}}(T^n) = nh_{\text{top}}(T)$ (problems 4.4 and 4.13). Thus we get upon division by n that $h_\mu(T) \leq h_{\text{top}}(T) + (2/n) \xrightarrow{n \rightarrow \infty} 0$, which proves the theorem. \square

In fact, $h_{\text{top}}(T) = \sup\{h_{\mu}(T) : \mu \text{ is an invariant Borel probability measure}\}$. But to prove this we need some more preparations. These are done in the next section.

4.6.2 Bowen's definition

We assume as usual that (X, d) is a compact metric space, and that $T : X \rightarrow X$ is continuous. For every n we define a new metric d_n on X as follows:

$$d_n(x, y) := \max_{0 \leq i \leq n-1} d(T^i x, T^i y).$$

This is called *Bowen's metric*. It depends on T . A set $F \subset X$ is called (n, ε) -separated, if for every $x, y \in F$ s.t. $x \neq y$, $d_n(x, y) > \varepsilon$.

Definition 4.7.

1. $s_n(\varepsilon, T) := \max\{\#(F) : F \text{ is } (n, \varepsilon)\text{-separated}\}$.
2. $s(\varepsilon, T) := \limsup_{n \rightarrow \infty} \frac{1}{n} \log s_n(\varepsilon, T)$
3. $\bar{h}_{\text{top}}(T) := \lim_{\varepsilon \rightarrow 0^+} s(\varepsilon, T)$.

Theorem 4.8 (Bowen). *Suppose T is a continuous mapping of a compact metric space X , then $h_{\text{top}}(T) = \bar{h}_{\text{top}}(T)$.*

Proof. Suppose \mathcal{U} is an open cover all of whose elements have diameters less than ε . We claim that $N(\mathcal{U}_0^{n-1}) \geq s_n(\varepsilon, T)$ for all n . To see this suppose F is an (n, ε) -separated set of maximal cardinality. Each $x \in F$ is contained in some $U_x \in \mathcal{U}_0^{n-1}$. Since the d -diameter of every element of \mathcal{U} is less than δ , the d_n -diameter of every element of \mathcal{U}_0^{n-1} is less than δ . Thus the assignment $x \mapsto U_x$ is one-to-one, whence

$$N(\mathcal{U}_0^{n-1}) \geq s_n(\varepsilon, T).$$

It follows that $s(\varepsilon, T) \leq h_{\text{top}}(T, \mathcal{U}) \leq h_{\text{top}}(T)$, whence $\bar{h}_{\text{top}}(T) \leq h_{\text{top}}(T)$.

To see the other inequality we use *Lebesgue numbers*: a number δ is called a Lebesgue number for an open cover \mathcal{U} , if for every $x \in X$, the ball with radius δ and center x is contained in some element of \mathcal{U} . (Lebesgue numbers exist because of compactness.)

Fix ε and let \mathcal{U} be an open cover with Lebesgue number bigger than or equal to ε . It is easy to check that for every n , ε is a Lebesgue number for \mathcal{U}_0^{n-1} w.r.t. d_n .

Let F be an $(n, \varepsilon/2)$ -separated set of *maximal* cardinality, i.e. $\#F = s_n(\varepsilon)$. Then any point y we add to F will break its (n, ε) -separation property, and so

$$\forall y \in X \exists x \in F \text{ s.t. } d_n(x, y) \leq \varepsilon/2.$$

It follows that the sets $\overline{B_n(x; \varepsilon/2)} := \{y : d_n(x, y) \leq \varepsilon/2\}$ ($x \in F$) cover X .

Every $\overline{B_n(x, \varepsilon/2)}$ ($x \in F$) is contained in some element of \mathcal{U}_0^{n-1} , because \mathcal{U}_0^{n-1} has Lebesgue number ε w.r.t d_n . The union of these elements covers X . We found a sub cover of \mathcal{U}_0^{n-1} of cardinality at most $\#F = s_n(\varepsilon)$. This shows that

$$N(\mathcal{U}_0^{n-1}) \leq s_n(\varepsilon).$$

We just proved that for every open cover \mathcal{U} with Lebesgue number at least ε , $h_{\text{top}}(T, \mathcal{U}) \leq s(\varepsilon)$. It follows that

$$\sup\{h_{\text{top}}(T, \mathcal{U}) : \mathcal{U} \text{ has Lebesgue number at least } \varepsilon\} \leq s(\varepsilon).$$

We now pass to the limit $\varepsilon \rightarrow 0^+$. The left hand side tends to the supremum over all open covers, which is equal to $h_{\text{top}}(T)$. We obtain $h_{\text{top}}(T) \leq \lim_{\varepsilon \rightarrow 0^+} s(\varepsilon)$. \square

Corollary 4.1. *Suppose T is an isometry, then all its invariant probability measures have entropy zero.*

Proof. If T is an isometry, then $d_n = d$ for all n , therefore $s(\varepsilon, T) = 0$ for all $\varepsilon > 0$, so $\bar{h}_{\text{top}}(T) = 0$. The theorem says that $h_{\text{top}}(T) = 0$. The corollary follows from Goodwyn's theorem. \square

4.6.3 The variational principle

The following theorem was first proved under additional assumptions by Dinaburg, and then in the general case by Goodman. The proof below is due to Misiurewicz.

Theorem 4.9 (Variational principle). *Suppose $T : X \rightarrow X$ is a continuous map of a compact metric space, then $h_{\text{top}}(T) = \sup\{h_{\mu}(T) : \mu \text{ is an invariant Borel measure}\}$.*

Proof. We have already seen that the topological entropy is an upper bound for the metric entropies. We just need to show that this is the least upper bound.

Fix ε , and let F_n be a sequence of (n, ε) -separated sets of maximal cardinality (so $\#F_n = s_n(\varepsilon, T)$). Let

$$\nu_n := \frac{1}{\#F_n} \sum_{x \in F_n} \delta_x,$$

where δ_x denotes the Dirac measure at x (i.e. $\delta_x(E) = 1_E(x)$). These measure are not invariant, so we set

$$\mu_n := \frac{1}{n} \sum_{k=0}^{n-1} \nu_n \circ T^{-k}.$$

Any weak star limit of μ_n will be T -invariant (check).

Fix some sequence $n_k \rightarrow \infty$ s.t. $\mu_{n_k} \xrightarrow{w^*} \mu$ and s.t. $\frac{1}{n_k} \log s_{n_k}(\varepsilon, T) \xrightarrow{n \rightarrow \infty} s(\varepsilon, T)$. We show that the entropy of μ is at least $s(\varepsilon, T)$. Since $s(\varepsilon, T) \xrightarrow{\varepsilon \rightarrow 0^+} h_{\text{top}}(T)$, this will prove the theorem.

Let $\alpha = \{A_1, \dots, A_N\}$ be a measurable partition of X s.t. (1) $\text{diam}(A_i) < \varepsilon$; and (2) $\mu(\partial A_i) = 0$. (Such a partition can be generated from a cover of X by balls of radius less than $\varepsilon/2$ and boundary of zero measure.) It is easy to see that the d_n -diameter of α_0^{n-1} is also less than ε . It is an exercise to see that every element of α_0^{n-1} has boundary with measure μ equal to zero.

We calculate $H_{V_n}(\alpha_0^{n-1})$. Since F_n is (n, ε) -separated and every atom of α has d_n -diameter less than ε , α_0^{n-1} has $\#F_n$ elements whose v_n measure is $1/\#F_n$, and the other elements of α_0^{n-1} have measure zero. Thus

$$H_{V_n}(\alpha_0^{n-1}) = \log_2 \#F_n = \log_2 s_n(\varepsilon, T).$$

We now “play” with $H_{V_n}(\alpha_0^{n-1})$ with the aim of bounding it by something involving a sum of the form $\sum_{i=0}^{n-1} H_{V_n \circ T^{-i}}(\alpha_0^{q-1})$. Fix q , and $j \in \{0, \dots, q-1\}$, then

$$\begin{aligned} \log_2 s_n(\varepsilon, T) &= H_{V_n}(\alpha_0^{n-1}) \leq H_{V_n}(\alpha_0^{j-1} \vee \bigvee_{i=0}^{\lfloor n/q \rfloor - 1} T^{-(qi+j)} \alpha_0^{q-1} \vee \alpha_{q(\lfloor n/q \rfloor - 1) + j + 1}^{n-1}) \\ &\leq \sum_{i=0}^{\lfloor n/q \rfloor - 1} H_{V_n \circ T^{-(qi+j)}}(\alpha_0^{q-1}) + 2q \log_2 \# \alpha. \end{aligned}$$

Summing over $j = 0, \dots, q-1$, we get

$$\begin{aligned} q \log_2 s_n(\varepsilon, T) &\leq n \cdot \frac{1}{n} \sum_{k=0}^{n-1} H_{V_n \circ T^{-k}}(\alpha_0^{q-1}) + 2q \log_2 \# \alpha \\ &\leq n H_{\mu_n}(\alpha_0^{q-1}) + 2q \log_2 \# \alpha, \end{aligned}$$

because $\mu_n = \frac{1}{n} \sum_{i=0}^{n-1} V_n \circ T^{-i}$ and $\phi(t) = -t \log_2 t$ is concave. Thus

$$\frac{1}{n_k} \log_2 s_{n_k}(\varepsilon, T) \leq \frac{1}{q} H_{\mu_{n_k}}(\alpha_0^{q-1}) + \frac{2}{n_k} \log_2 \# \alpha, \quad (4.5)$$

where $n_k \rightarrow \infty$ is the subsequence chosen above.

Since every $A \in \alpha_0^{n-1}$ satisfies $\mu(\partial A) = 0$, $\mu_{n_k}(A) \xrightarrow[k \rightarrow \infty]{w^*} \mu(A)$ for all $A \in \alpha_0^{n-1}$ (sandwich $u \leq 1_A \leq v$ with u, v continuous s.t. $\int |u - v| d\mu \ll 1$). It follows that $H_{\mu_{n_k}}(\alpha_0^{q-1}) \xrightarrow[k \rightarrow \infty]{} H_\mu(\alpha_0^{q-1})$. Passing to the limit $k \rightarrow \infty$ in (4.5), we have $s(\varepsilon, T) \leq \frac{1}{q} H_\mu(\alpha) \xrightarrow[q \rightarrow \infty]{} h_\mu(T, \alpha) \leq h_\mu(T)$. Thus $h_\mu(T) \geq s(\varepsilon, T)$. Since $s(\varepsilon, T) \xrightarrow[\varepsilon \rightarrow 0^+]{} h_{\text{top}}(T)$ the theorem is proved. \square

Problems

4.1. Prove: $H_\mu(\alpha|\beta) = - \sum_{B \in \beta} \mu(B) \sum_{A \in \alpha} \mu(A|B) \log \mu(A|B)$, where $\mu(A|B) = \frac{\mu(A \cap B)}{\mu(B)}$.

4.2. Prove: if $H_\mu(\alpha|\beta) = 0$, then $\alpha \subseteq \beta \pmod{\mu}$.

4.3. Prove that $h_\mu(T)$ is an invariant of measure theoretic isomorphism.

4.4. Prove that $h_\mu(T^n) = nh_\mu(T)$.

4.5. Prove that if T is invertible, then $h_\mu(T^{-1}) = h_\mu(T)$.

4.6. Entropy is affine

Let T be a measurable map on X , and μ_1, μ_2 be two T -invariant probability measures. Set $\mu = t\mu_1 + (1-t)\mu_2$ ($0 \leq t \leq 1$). Show: $h_\mu(T) = th_{\mu_1}(T) + (1-t)h_{\mu_2}(T)$.
Guidance: Start by showing that for all $0 \leq x, y, t \leq 1$,

$$0 \leq \varphi(tx + (1-t)y) - [t\varphi(x) + (1-t)\varphi(y)] \leq -t \log t - (1-t) \log(1-t)$$

4.7. Let (X, \mathcal{B}, μ) be a probability space. If α, β are two measurable partitions of X , then we write $\alpha = \beta \pmod{\mu}$ if $\alpha = \{A_1, \dots, A_n\}$ and $\beta = \{B_1, \dots, B_n\}$ where $\mu(A_i \triangle B_i) = 0$ for all i . Let \mathfrak{P} denote the set of all countable measurable partitions of X , modulo the equivalence relation $\alpha = \beta \pmod{\mu}$. Show that

$$\rho(\alpha, \beta) := H_\mu(\alpha|\beta) + H_\mu(\beta|\alpha)$$

induces a metric on \mathfrak{P} .

4.8. Let (X, \mathcal{B}, μ, T) be a ppt. Show that $|h_\mu(T, \alpha) - h_\mu(T, \beta)| \leq H_\mu(\alpha|\beta) + H_\mu(\beta|\alpha)$.

4.9. Use the previous problem to show that the supremum in the definition of metric entropy is attained by *finite* measurable partitions.

4.10. Suppose $\alpha = \{A_1, \dots, A_n\}$ is a finite measurable partition. Show that for every ε , there exists $\delta = \delta(\varepsilon, n)$ such that if $\beta = \{B_1, \dots, B_n\}$ is measurable partition s.t. $\mu(A_i \triangle B_i) < \delta$, then $\rho(\alpha, \beta) < \varepsilon$.

4.11. Entropy via generating sequences of partitions

Suppose (X, \mathcal{B}, μ) is a probability space, and \mathcal{A} is an algebra of \mathcal{F} -measurable subsets (namely a collection of sets which contains \emptyset and which is closed under finite unions, finite intersection, and forming complements). Suppose \mathcal{A} generates \mathcal{B} (i.e. \mathcal{B} is the smallest σ -algebra which contains \mathcal{A}).

1. For every $F \in \mathcal{F}$ and $\varepsilon > 0$, there exists $A \in \mathcal{A}$ s.t. $\mu(A \triangle F) < \varepsilon$.
2. For every \mathcal{F} -measurable finite partition β and $\varepsilon > 0$, there exists an \mathcal{A} -measurable finite partition α s.t. $\rho(\alpha, \beta) < \varepsilon$.

3. If $T : X \rightarrow X$ is probability preserving, then

$$h_\mu(T) = \sup\{h_\mu(T, \alpha) : \alpha \text{ is an } \mathcal{A}\text{-measurable finite partition}\}.$$

4. Suppose $\alpha_1 \leq \alpha_2 \leq \dots$ is an increasing sequence of finite measurable partitions such that $\sigma(\bigcup_{n \geq 1} \alpha_n) = \mathcal{B} \pmod{\mu}$, then $h_\mu(T) = \lim_{n \rightarrow \infty} h_\mu(T, \alpha_n)$.

4.12. Show that the entropy of the product of two ppt is the sum of their two entropies.

4.13. Show that $h_{\text{top}}(T^n) = nh_{\text{top}}(T)$.

Notes to chapter 4

The notion of entropy as a measure of information is due to Shannon, the father of the modern theory of information. Kolmogorov had the idea to adapt this notion to the ergodic theoretic context for the purposes of inventing an invariant which is able to distinguish Bernoulli schemes. This became possible once Sinai has proved his generator theorem — which enables the calculation of this invariant for Bernoulli schemes. Later, in the 70's, Ornstein has proved that entropy is a complete invariant for Bernoulli schemes: they are isomorphic iff they have the same entropy. The maximum of the possible entropies for a topological Markov shift was first calculated by Parry, who also found the maximizing measure. The material in this chapter is all classical, [3] and [1] are both excellent references. For an introduction to Ornstein's isomorphism theorem, see [2].

References

1. Petersen, K.: *Ergodic theory*. Corrected reprint of the 1983 original. Cambridge Studies in Advanced Mathematics, 2. Cambridge University Press, Cambridge, 1989. xii+329 pp.
2. Rudolph, D.: *Fundamentals of measurable dynamics: Ergodic theory on Lebesgue spaces*, Oxford Science Publications, 1990. x+168 pp.
3. Walters, P.: *An introduction to ergodic theory*. Graduate Texts in Mathematics, 79. Springer-Verlag, New York-Berlin, 1982. ix+250 pp.

Appendix A

The isomorphism theorem for standard measure spaces

A.1 Polish spaces

Definition A.1. A *polish space* is a metric space (X, d) which is

1. complete (every Cauchy sequence has a limit);
2. and separable (there is a countable dense subset).

Every compact metric space is polish. But a polish space need not be compact, or even locally compact. For example,

$$\mathbb{N}^{\mathbb{N}} := \{x = (x_1, x_2, x_3, \dots) : x_k \in \mathbb{N}\}$$

equipped with the metric $d(x, y) := \sum_{k \geq 1} 2^{-k} |x_k^{-1} - y_k^{-1}|$ is a non locally compact polish metric space.

Notation: $B(x, r) := \{y \in X : d(x, y) < r\}$ (the open ball with center x and radius r).

Proposition A.1. Suppose (X, d) is a polish space, then

1. Second axiom of countability: *There exists a countable family of open sets \mathcal{U} such that every open set in X is a union of a subfamily of \mathcal{U} .*
2. Lindelöf property: *Every cover of X by open sets has a countable sub-cover.*
3. *The intersection of any decreasing sequence of closed balls whose radii tend to zero is a single point.*

Proof. Since X is separable, it contains a countable dense set $\{x_n\}_{n \geq 1}$. Define

$$\mathcal{U} := \{B(x_n, r) : n \in \mathbb{N}, r \in \mathbb{Q}\}.$$

This is a countable collection, and we claim that it satisfies (1). Take some open set U . For every $x \in U$ there are

1. $R > 0$ such that $B(x, R) \subset U$ (because U is open);
2. $x_n \in B(x, R/2)$ (because $\{x_n\}$ is dense); and

3. $r \in \mathbb{Q}$ such that $d(x, x_n) < r < R/2$.

It is easy to check that $x \in B(x_n, r) \subset B(x, 2r) \subset B(x, R) \subset U$. Thus for every $x \in U$ there is $U_x \in \mathcal{U}$ s.t. $x \in U_x \subset U$. It follows that U is a union of elements from \mathcal{U} and (1) is proved. (2) is an immediate consequence.

To see (3), suppose $B_n := \overline{B(z_n, r_n)}$ is a sequence of closed balls such that $B_n \supset B_{n+1}$ and $r_n \rightarrow 0$. It is easy to verify that $\{z_n\}$ is a Cauchy sequence. Since X is complete, it converges to a limit z . This limit belongs to B_n because $z_k \in B_k \subset B_n$ for all $k > n$, and B_n is closed. Thus the intersection of B_n contains at least one point. It cannot contain more than one point, because its diameter is zero (because it is bounded by $\text{diam}[B_n] \leq r_n \rightarrow 0$). \square

A.2 Standard probability spaces

Definition A.2. A *standard probability space* is a probability space (X, \mathcal{B}, μ) where X is polish, \mathcal{B} is the σ -algebra of Borel sets of X , and μ is a Borel probability measure.

Theorem A.1. Suppose (X, \mathcal{B}, μ) is a standard probability space, then

1. **Regularity:** Suppose $E \in \mathcal{B}$. For every $\varepsilon > 0$ there exists an open set U and a closed set F such that $F \subset E \subset U$ and $\mu(U \setminus F) < \varepsilon$.
2. **Separability:** There exists a countable collection of measurable sets $\{E_n\}_{n \geq 1}$ such that for every $E \in \mathcal{B}$ and $\varepsilon > 0$ there exists some n s.t. $\mu(E \Delta E_n) < \varepsilon$. Equivalently, $L^p(X, \mathcal{B}, \mu)$ is separable for all $1 \leq p < \infty$.

Proof. Say that a set E satisfies the *approximation property* if for every ε there are a closed set F and an open set U s.t. $F \subset E \subset U$ and $\mu(U \setminus E) < \varepsilon$.

Open balls $B(x, r)$ have the approximation property: Take $U = B(x, r)$ and $F = \overline{B(x, r - \frac{1}{n})}$ for n sufficiently large (these sets increase to $B(x, r)$ so their measure tends to that of $B(x, r)$).

Open sets U have the approximation property. The approximating open set is the set itself. To find the approximating closed set use the second axiom of countability to write the open set as the countable union of balls B_n , and approximate each B_n from within by a closed set F_n such that $\mu(B_n \setminus F_n) < \varepsilon/2^{n+1}$. Then $\mu(U \setminus \bigcup_{n \geq 1} F_n) < \varepsilon/2$. Now take $F := \bigcup_{i=1}^N F_i$ for N large enough.

Thus the collection $\mathcal{C} := \{E \in \mathcal{B} : E \text{ has the approximation property}\}$ contains the open sets. Since it is a σ -algebra (check!), it must be equal to \mathcal{B} , proving (1).

We prove separability. Polish spaces satisfy the second axiom of countability, so there is a countable family of open balls $\mathcal{U} = \{B_n : n \in \mathbb{N}\}$ such that every open set is the union of a countable subfamily of \mathcal{U} . This means that every open set can be approximated by a *finite* union of elements of \mathcal{U} to arbitrary precision. By the regularity property shown above, every measurable set can be approximated by a finite union of elements of \mathcal{U} to arbitrary precision. It remains to observe that $\mathcal{E} = \{\text{finite unions of elements of } \mathcal{U}\}$ is countable.

The separability of L^p for $1 \leq p < \infty$ follows from the above and the obvious fact that the collection $\{\sum_{i=1}^N \alpha_i 1_{E_i} : N \in \mathbb{N}, \alpha_i \in \mathbb{Q}, E_i \in \mathcal{B}\}$ is dense in L^p (prove!). The other direction is left to the reader. \square

The following statement will be used in the proof of the isomorphism theorem.

Lemma A.1. *Suppose (X, \mathcal{B}, μ) is a standard probability space and E is a measurable set of positive measure, then there is a point $x \in X$ such that $\mu[E \cap B(x, r)] \neq 0$ for all $r > 0$.*

Proof. Fix $\varepsilon_n \downarrow 0$. Write X as a countable union of open balls of radius ε_1 (second axiom of countability). At least one of these, B_1 , satisfies $\mu(E \cap B_1) \neq 0$. Write B_1 as a countable union of open balls of radius ε_2 . At least one of these, B_2 , satisfies $\mu[E \cap B_1 \cap B_2] \neq 0$. Continue in this manner. The result is a decreasing sequence of open balls with shrinking diameters $B_1 \supset B_2 \supset \dots$ which intersect E at a set of positive measure.

The sequence of centers of these balls is a Cauchy sequence. Since X is polish, it converges to a limit $x \in X$. This x belongs to the closure of each B_n .

For every r find n so large that $\varepsilon_n < r/2$. Since $x \in \overline{B_n}$, $d(x, x_n) \leq \varepsilon_n$, and this implies that $B(x, r) \supseteq B(x_n, \varepsilon_n) = B_n$. Since B_n intersects E with positive measure, $B(x, r)$ intersects E with positive measure. \square

A.3 Atoms

Definition A.3. An *atom* of a measure space (X, \mathcal{B}, μ) is a measurable set A of non-zero measure with the property that for all other measurable sets B contained in A , either $\mu(B) = \mu(A)$ or $\mu(B) = 0$. A measure space is called *non-atomic*, if it has no atoms.

Proposition A.2. *For standard spaces (X, \mathcal{B}, μ) , every atom is of the form $\{x\} \cup \text{null set}$ for some x s.t. $\mu\{x\} \neq 0$.*

Proof. Suppose A is an atom. Since X can be covered by a countable collection of open balls of radius $r_1 := 1$, $A = \bigcup_{i \geq 1} A_i$ where A_i are measurable subsets of A of diameter at most r_1 . One of those sets, A_{i_1} , has non-zero measure. Since A is an atom, $\mu(A_{i_1}) = \mu(A)$. Setting $A^{(1)} := A_{i_1}$, we see that

$$A^{(1)} \subset A, \text{diam}(A^{(1)}) \leq r_1, \mu(A^{(1)}) = \mu(A).$$

Of course $A^{(1)}$ is an atom.

Now repeat this argument with $A^{(1)}$ replacing A and $r_2 := 1/2$ replacing r_1 . We obtain an atom $A^{(2)}$ s.t.

$$A^{(2)} \subset A, \text{diam}(A^{(2)}) \leq r_2, \mu(A^{(2)}) = \mu(A).$$

We continue in this manner, to obtain a sequence of atoms $A \supset A^{(1)} \supset A^{(2)} \supset \dots$ of the same measure, with diameters $r_k = 1/k \rightarrow 0$. The intersection $\bigcap A^{(k)}$ is non-empty, because its measure is $\lim \mu(A^{(k)}) = \mu(A) \neq 0$. But its diameter is zero. Therefore it is a single point x , and by construction $x \in A$ and $\mu\{x\} = \mu(A)$. \square

Lemma A.2. *Suppose (X, \mathcal{B}, μ) is a non-atomic standard probability space, and $r > 0$. Every \mathcal{B} -measurable set E can be written in the form $E = \biguplus_{i=1}^{\infty} F_i \uplus N$ where $\mu(N) = 0$, F_i are closed, $\text{diam}(F_i) < r$, and $\mu(F_i) \neq 0$.*

Proof. Since every measurable set is a finite or countable disjoint union of sets of diameter less than r (prove!), it is enough to treat sets E such that $\text{diam}(E) < r$.

Standard spaces are regular, so we can find a closed set $F_1 \subset E$ such that $\mu(E \setminus F_1) < \frac{1}{2}$. If $\mu(E \setminus F_1) = 0$, then stop. Otherwise apply the argument to $E \setminus F_1$ to find a closed set $F_2 \subset E \setminus F_1$ of positive measure such that $\mu[E \setminus (F_1 \cup F_2)] < \frac{1}{2^2}$. Continuing in this manner we obtain pairwise disjoint closed sets $F_i \subset E$ such that $\mu(E \setminus \bigcup_{i=1}^n F_i) < 2^{-i}$ for all n , or until we get to an n such that $\mu(E \setminus \biguplus_{i=1}^n F_i) = 0$.

If the procedure did not stop at any stage, then the lemma follows with $N := E \setminus \bigcup_{i \geq 1} F_i$.

We show what to do in case the procedure stops after n steps. Set $F = F_n$, the last closed set. The idea is to split F into countably many disjoint closed sets, plus a set of measure zero.

Find an x such that $\mu[F \cap B(x, r)] \neq 0$ for all $r > 0$ (previous lemma). Since X is non-atomic, $\mu\{x\} = 0$. Since $B(x, r) \downarrow \{x\}$, $\mu[F \cap B(x, r)] \xrightarrow[n \rightarrow \infty]{} 0$. Choose $r_n \downarrow 0$ for which $\mu[F \cap B(x, r_n)]$ is strictly decreasing. Define

$$C_1 := F \cap \overline{B(x, r_1)} \setminus B(x, r_2), \quad C_2 := F \cap \overline{B(x, r_3)} \setminus B(x, r_4) \text{ and so on.}$$

This is an *infinite* sequence of closed pairwise disjoint sets of positive measure inside F . By the construction of F they are disjoint from F_1, \dots, F_{n-1} and they are contained in E .

Now consider $E' := E \setminus (\bigcup_{i=0}^{n-1} F_i \cup \bigcup_i C_i)$. Applying the argument in the first paragraph to E' , we write it as a finite or countable disjoint union of closed sets plus a null set. Adding these sets to the collection $\{F_1, \dots, F_{n-1}\} \cup \{C_i : i \geq 1\}$ gives us the required decomposition of E . \square

A.4 The isomorphism theorem

Definition A.4. Two measure spaces $(X_i, \mathcal{B}_i, \mu_i)$ ($i = 1, 2$) are called *isomorphic* if there are measurable subsets of full measure $X'_i \subset X_i$ and a measurable bijection $\pi : X'_1 \rightarrow X'_2$ with measurable inverse such that $\mu_2 = \mu_1 \circ \pi^{-1}$.

Theorem A.2 (Isomorphism theorem). *Every non-atomic standard probability spaces is isomorphic to the unit interval equipped with the Lebesgue measure.*

Proof. Fix a decreasing sequence of positive numbers ε_n which tends to zero. Using lemma A.2, decompose $X = \bigsqcup_{j=1}^{\infty} F(j) \sqcup N$ where $F(j)$ are pairwise disjoint closed sets of positive measure and diameter less than ε_1 , and N is a null set.

Applying lemma A.2 to each $F(j)$, decompose $F(j) = \bigsqcup_{k=1}^{\infty} F(j,k) \sqcup N(j)$ where $F(j,k)$ are pairwise disjoint closed sets of positive measure and diameter less than ε_2 , and $N(j)$ is a null set.

Continuing in this way we obtain a family of sets $F(x_1, \dots, x_n), N(x_1, \dots, x_n)$, ($n, x_1, \dots, x_n \in \mathbb{N}$) such that

1. $F(x_1, \dots, x_n)$ are closed, have positive measure, and $\text{diam}[F(x_1, \dots, x_n)] < \varepsilon_n$;
2. $F(x_1, \dots, x_{n-1}) = \bigsqcup_{y \in \mathbb{N}} F(x_1, \dots, x_{n-1}, y) \sqcup N(x_1, \dots, x_{n-1})$;
3. $\mu[N(x_1, \dots, x_n)] = 0$.

Set $X' := \bigcap_{n \geq 1} \bigsqcup_{x_1, \dots, x_n \in \mathbb{N}} F(x_1, \dots, x_n)$. It is a calculation to see that $\mu(X \setminus X') = 0$. The set X' has ‘tree like structure’: every $x \in X'$ determines a unique sequence $(x_1, x_2, \dots) \in \mathbb{N}^{\mathbb{N}}$ such that $x \in F(x_1, \dots, x_n)$ for all n . Define $\pi : X' \rightarrow [0, 1]$ by

$$\pi(x) = \frac{1}{x_1 + \frac{1}{x_2 + \dots}}$$

This map is one-to-one on X' , because if $\pi(x) = \pi(y)$, then $1/(x_1 + 1/(x_2 + \dots)) = 1/(y_1 + 1/(y_2 + \dots))$ whence $x_k = y_k$ for all k ;¹ this means that $x, y \in F(x_1, \dots, x_n)$ for all n , whence $d(x, y) \leq \varepsilon_n \xrightarrow{n \rightarrow \infty} 0$.

This map is onto $[0, 1] \setminus \mathbb{Q}$, because every irrational $t \in [0, 1]$ has an infinite continued fraction expansion $1/(a_1 + 1/(a_2 + \dots))$, so $t = \pi(x)$ for the unique $x \in \bigcap_{n \geq 1} F(a_1, \dots, a_n)$. (This intersection is non-empty because it is the decreasing intersection of closed sets of shrinking diameters in a complete metric space.)

We claim that $\pi : X' \rightarrow [0, 1] \setminus \mathbb{Q}$ is Borel measurable. Let $[a_1, \dots, a_n]$ denote the collection of all irrationals in $[0, 1]$ whose continued fraction expansion starts with (a_1, \dots, a_n) . We call such sets ‘‘cylinders’’. We have $\pi^{-1}[a_1, \dots, a_n] = F(a_1, \dots, a_n)$ a closed set, so the π -preimage of every cylinder is Borel measurable. Thus

$$\mathcal{C} := \{E \in \mathcal{B}([0, 1] \setminus \mathbb{Q}) : \pi^{-1}(E) \in \mathcal{B}\}$$

contains the cylinders. It is easy to check that \mathcal{C} is a σ -algebra. The cylinders generate $\mathcal{B}([0, 1] \setminus \mathbb{Q})$ (these are intervals whose length tends to zero as $n \rightarrow \infty$). It follows that $\mathcal{C} = \mathcal{B}([0, 1] \setminus \mathbb{Q})$ and the measurability of π is proved.

Next we claim that $\pi^{-1} : [0, 1] \setminus \mathbb{Q} \rightarrow X'$ is Borel measurable. This is because $\pi[F(a_1, \dots, a_n)] = [a_1, \dots, a_n]$ and an argument similar to the one in the previous paragraph.

It follows that $\pi : (X, \mathcal{B}, \mu) \rightarrow ([0, 1] \setminus \mathbb{Q}, \mathcal{B}([0, 1] \setminus \mathbb{Q}), \mu \circ \pi^{-1})$ is an isomorphism of measure spaces. There is an obvious extension of $\mu \circ \pi^{-1}$ to $\mathcal{B}([0, 1])$ obtained by declaring $\mu(\mathbb{Q}) := 0$. Let m denote this extension. Then we get an

¹ Hint: apply the transformation $x \mapsto [1/x]$ to both sides.

isomorphism between (X, \mathcal{B}, μ) to $([0, 1], \mathcal{B}([0, 1]), m)$ where m is some Borel probability measure on $[0, 1]$. Since μ is non-atomic, m is non-atomic.

We now claim that $([0, 1], \mathcal{B}([0, 1]), m)$ is isomorphic to $([0, 1], \mathcal{B}([0, 1]), \lambda)$, where λ is the Lebesgue measure.

Consider first the distribution function of m , $s \mapsto m[0, s]$. This is a monotone increasing function (in the weak sense). We claim that it is continuous. Otherwise it has a jump J at some point x_0 :

$$m[0, x_0 + \varepsilon] - m[0, x_0] > J \text{ for all } \varepsilon > 0.$$

This means that $m\{x_0\} \geq J$, which cannot be the case since m is non-atomic.

It follows that the following definition makes sense for all t :

$$\vartheta(t) := \min\{s \geq 0 : m([0, s]) = t\}.$$

This is a monotone increasing map $\vartheta : [0, 1] \rightarrow [0, 1]$. We aim at showing that ϑ is an isomorphism with the unit interval equipped with Lebesgue's measure λ .

Step 1. $m([0, 1] \setminus \vartheta[0, 1]) = 0$.

The reader should note that in general $\vartheta[0, 1] \neq [0, 1]$ as sets: If $m[[s', s]] = 0$ for some $s' < s$, then $s \notin \vartheta[0, 1]$. The opposite is also true: If $s \notin \vartheta([0, 1])$, then $\exists s' < s$ such that $m[0, s] = m[0, s']$, whence $m[s', s] = 0$.

Define for every $s \notin \vartheta[0, 1]$, $I(s) := [s_1, s_2]$, where

$$s_1 := \min\{s' < s : m[0, s'] = m[0, s]\} \text{ and } s_2 := \max\{s' > s : m[0, s'] = m[0, s]\}$$

Then $s \in I(s)$, $|I(s)| \geq |s - s'| > 0$, and since m is not atomic, $m[I(s)] = m[s_1, s_2] = 0$. Moreover, for any two $s', s'' \in [0, 1] \setminus \vartheta[0, 1]$, either $I(s') = I(s'')$, or $I(s') \cap I(s'') = \emptyset$. Since there can be at most countably many disjoint intervals of positive length, the collection $\{I(s) : s \in [0, 1] \setminus \vartheta[0, 1]\}$ is countable. It follows that $[0, 1] \setminus \vartheta[0, 1]$ is covered by a countable collection of sets of m -measure zero. The step follows.

Step 2. ϑ is one-to-one on $[0, 1]$.

Observe first that since m has no atoms, $m([0, s]) = m([0, s])$ for all s . This implies the equation $m[0, \vartheta(t)] = t$. It immediately follows that ϑ is one-to-one.

Step 3. ϑ is measurable, with measurable inverse.

ϑ is monotonic, so it maps intervals to intervals. This implies that ϑ and ϑ^{-1} are measurable (prove!).

Step 4. $m \circ \vartheta = \lambda$.

Recall the equation $m[0, \vartheta(t)] = t$ found above. It implies that $m[(\vartheta(s), \vartheta(t))] = t - s$ for all $0 < s < t < 1$, and this is enough to deduce that $m \circ \vartheta = \lambda$ (prove!).

Steps 1–4 show that $\vartheta : ([0, 1], \mathcal{B}([0, 1]), m) \rightarrow ([0, 1], \mathcal{B}([0, 1]), \lambda)$ is an isomorphism. Composing this with π , we get an isomorphism between (X, \mathcal{B}, μ) and the unit interval equipped with Lebesgue's measure. \square

We comment on the atomic case. A standard probability space (X, \mathcal{B}, μ) can have at most countably many atoms (otherwise it will contain an uncountable collection of pairwise disjoint sets of positive measure, which cannot be the case). Let $\{x_i : i \in \Lambda\}$ be a list of the atoms, where $\Lambda \subset \mathbb{N}$. Then

$$\mu = \mu' + \sum_{i \in \Lambda} \mu\{x_i\} \delta_{x_i} \quad (\delta_x = \text{Dirac measure})$$

where μ' is non-atomic.

Suppose w.l.o.g that $X \cap \mathbb{N} = \emptyset$. The map

$$\pi : X \rightarrow X \cup \Lambda, \quad \pi(x) = \begin{cases} x & x \notin \{x_i : i \in \Lambda\} \\ i & x = x_i \end{cases}$$

is an isomorphism between X and the measure space obtained by adding to (X, \mathcal{B}, μ') atoms with right mass at points of Λ . The space (X, \mathcal{B}, μ') is non-atomic, so it is isomorphic to $[0, \mu'(X)]$ equipped with the Lebesgue measure. We obtain the following generalization of the isomorphism theorem: *Every standard probability space is isomorphic to the measure space consisting of a finite interval equipped with Lebesgue's measure, and a finite or countable collection of atoms.*

Definition A.5. A measure space is called a *Lebesgue space*, if it is isomorphic to the measure space consisting of a finite interval equipped with the Lebesgue measurable sets and Lebesgue's measure, and a finite or countable collection of atoms.

Note that the σ -algebra in the definition is the Lebesgue σ -algebra, not the Borel σ -algebra. (The Lebesgue σ -algebra is the completion of the Borel σ -algebra with respect to the Lebesgue measure, see problem 1.2.) The isomorphism theorem and the discussion above say that the *completion* of a standard space is a Lebesgue space. So the class of Lebesgue probability spaces is enormous!

Appendix A

The Monotone Class Theorem

Definition A.1. A sequence of sets $\{A_n\}$ is called *increasing* (resp. *decreasing*) if $A_n \subseteq A_{n+1}$ for all n (resp. $A_n \supseteq A_{n+1}$ for all n).

Notation: $A_n \uparrow A$ means that $\{A_n\}$ is an increasing sequence of sets, and $A = \bigcup A_n$. $A_n \downarrow A$ means that $\{A_n\}$ is a decreasing sequence of sets, and $A = \bigcap A_n$.

Proposition A.1. Suppose (X, \mathcal{B}, μ) is a measure space, and $A_n \in \mathcal{B}$.

1. if $A_n \uparrow A$, then $\mu(A_n) \xrightarrow[n \rightarrow \infty]{} \mu(A)$;
2. if $A_n \downarrow A$ and $\mu(A_n) < \infty$ for some n , then $\mu(A_n) \xrightarrow[n \rightarrow \infty]{} \mu(A)$.

Proof. For (1), observe that $A = \bigcup_{n \geq 1} A_{n+1} \setminus A_n$ and use σ -additivity. For (2), fix n_0 s.t. $\mu(A_{n_0}) < \infty$, and observe that $A_n \downarrow A$ implies that $(A_{n_0} \setminus A_n) \uparrow (A_{n_0} \setminus A)$. \square

The example $A_n = (n, \infty)$, $\mu =$ Lebesgue measure on \mathbb{R} , shows that the condition in (2) cannot be removed.

Definition A.2. Let X be a set. A *monotone class* of subsets of X is a collection \mathcal{M} of subsets of X which contains the empty set, and such that if $A_n \in \mathcal{M}$ and $A_n \uparrow A$ or $A_n \downarrow A$, then $A \in \mathcal{M}$.

Recall that an *algebra* of subsets of a set X is a collection of subsets of X which contains the empty set, and which is closed under finite unions, finite intersections, and forming the complement.

Theorem A.1 (Monotone Class Theorem). A monotone class which contains an algebra, also contains the sigma-algebra generated by this algebra.

Proof. Let \mathcal{M} be a monotone class which contains an algebra \mathcal{A} . Let $\mathcal{M}(\mathcal{A})$ denote the intersection of all the collections $\mathcal{M}' \subset \mathcal{M}$ such that (a) \mathcal{M}' is a monotone class, and (b) $\mathcal{M}' \supseteq \mathcal{A}$. This is a monotone class (check!). In fact it is the *minimal* monotone class which contains \mathcal{A} . We prove that it is a σ -algebra. Since $\mathcal{M}(\mathcal{A}) \subset \mathcal{M}$, this completes the proof.

We begin by claiming that $\mathcal{M}(\mathcal{A})$ is closed under forming complements. Suppose $E \in \mathcal{M}(\mathcal{A})$. The set

$$\mathcal{M}' := \{E' \in \mathcal{M}(\mathcal{A}) : (E')^c \in \mathcal{M}(\mathcal{A})\}$$

contains \mathcal{A} (because \mathcal{A} is an algebra), and it is a monotone class (check!). But $\mathcal{M}(\mathcal{A})$ is the minimal monotone class which contains \mathcal{A} , so $\mathcal{M}' \supset \mathcal{M}(\mathcal{A})$. It follows that $E \in \mathcal{M}'$, whence $E^c \in \mathcal{M}(\mathcal{A})$.

Next we claim that $\mathcal{M}(\mathcal{A})$ has the following property:

$$E \in \mathcal{M}(\mathcal{A}), A \in \mathcal{A} \implies E \cup A \in \mathcal{M}(\mathcal{A}).$$

Again, the reason is that the collection \mathcal{M}' of sets with this property contains \mathcal{A} , and is a monotone class.

Now fix $E \in \mathcal{M}(\mathcal{A})$, and consider the collection

$$\mathcal{M}'' := \{F \in \mathcal{M}(\mathcal{A}) : E \cup F \in \mathcal{M}(\mathcal{A})\}.$$

By the previous paragraph, \mathcal{M}'' contains \mathcal{A} . It is clear that \mathcal{M}'' is a monotone class. Thus $\mathcal{M}(\mathcal{A}) \subseteq \mathcal{M}''$, and as a result $E \cup F \in \mathcal{M}(\mathcal{A})$ for all $F \in \mathcal{M}(\mathcal{A})$. But $E \in \mathcal{M}(\mathcal{A})$ was arbitrary, so this means that $\mathcal{M}(\mathcal{A})$ is closed under finite unions.

Since $\mathcal{M}(\mathcal{A})$ is closed under finite unions, and countable increasing unions, it is closed under general countable unions.

Since $\mathcal{M}(\mathcal{A})$ is closed under forming complements and taking countable unions, it is a sigma algebra. By definition this sigma algebra contains \mathcal{A} and is contained in \mathcal{M} . \square

Index

- Abramov formula, 89
- Action, 1
- adding machine, 28, 60
- algebra, 98
- Alt, 44
- arithmeticity, 21
- atom, 103

- Bernoulli scheme, 9, 69
 - entropy of, 87
- Bowen's metric, 94

- Carathéodory Extension Theorem, 10
- Chacon's example, 72
- Chung-Neveu Lemma, 83
- coboundary, 31
- complete
 - measure space, 27
 - Riemannian surface, 18
- conditional
 - entropy, 78
 - expectation, 35
 - probabilities, 36
- configuration, 1
- conservative mpt, 29
- covariance, 6
- cutting and stacking, 72
- cylinders, 9

- Dynamical system, 1

- eigenfunction, 64
- eigenvalue, 64
- entropy
 - conditional, 78
 - of a measure, 81
 - of a partition, 78
 - topological, 91
- ergodic, 5
 - decomposition, 38
 - hypothesis, 3
 - ergodicity and countable Lebesgue spectrum, 72
 - ergodicity and extremality, 60
 - ergodicity and mixing, 60
 - flows, 18, 28
 - theory, 1
- Ergodic Theorem
 - Mean, 31, 60
 - Multiplicative, 48
 - Pointwise, 32
 - Ratio, 61
 - Subadditive, 40
- ergodicity and mixing, 32
- extension, 22
- exterior product, 45
 - and angles, 47
 - of linear operators, 46
- extremal measure, 60

- factor, 22
- flow, 1
- Fourier Walsh system, 69
- Furstenberg-Kesten theorem, 43

- generator, 85
- geodesic flow, 17

- Herglotz theorem, 66
- hyperbolic
 - plane, 16
 - surface, 18

- independence, 6

- for partitions, 81
- induced transformation, 24
 - Abramov formula, 89
 - entropy of, 89
 - for infinite mpt, 29
 - Kac formula, 24
 - Kakutani skyscraper, 26
- information
 - conditional, 78
 - content, 78
 - function, 78
- invariant set, 5
- isomorphism, 4
 - measure theoretic, 4
 - of measure spaces, 104
 - spectral, 63
- isomorphism theorem for measure spaces, 104
- iterates, 1
- itinerary, 77

- K automorphism, 69
- Kac formula, 24, 29
- Kakutani skyscraper, 26

- Lebesgue number, 94
- Liouville's theorem, 2

- Martingale convergence theorem, 61
- measure, 3
- measure preserving transformation, 4
- measure space, 3
 - Lebesgue, 3, 107
 - non-atomic, 103
 - sigma finite, 29
 - standard, 3
- mixing
 - and countable Lebesgue spectrum, 72
 - weak, 65
- mixing and ergodicity, 32
- monotone class, 109
- mpt, 4
- multilinear function, 44
 - alternating, 44

- natural extension, 23

- orbit, 1

- partition
 - finer or coarser, 79
 - wedge product of, 79
- periodicity, 21
- Perron-Frobenius Theorem, 27
- Phase space, 1

- Poincaré Recurrence Theorem, 3
- Poincaré section, 27
- Polish, 21
- polish space, 21, 101
- positive definite, 66
- probability
 - preserving transformation, 4
 - stationary probability vector, 11
 - measure, 3
 - space, 3
 - vector, 11
- product
 - of measure spaces, 19
 - of mpt, 20
- PSL(2, \mathbb{R}), 16

- regularity (of a measure space), 102
- rotation, 7
- Rotations, 28

- section map, 27
- semi algebra, 10
- semi-flow, 1
- sigma algebra, 3
- sigma finiteness, 29
- skew product, 21
- spectral
 - invariant, 63
 - isomorphism, 63
 - measure, 66
- spectrum
 - and the K property, 70
 - continuous, 64, 67
 - countable Lebesgue, 69
 - discrete, 64
 - Lebesgue, 72
 - mixed, 64
 - point, 64
 - pure point, 64
- standard
 - probability space, 102
- stationary
 - probability vector, 11
- stationary stochastic process, 5
- stochastic matrix, 11
- stochastic process, 4
- subadditive
 - cocycle, 40
 - ergodic theorem, 40
- subshift of finite type, 11
- suspension, 26

- tail events, 70
- tensor product, 44

time one map, 28
Topological entropy, 91
transition matrix
 aperiodic, 12
 irreducible, 12
 period, 12
unitary equivalence, 63

wandering set, 29
wedge product
 of σ -algebras, 79
 of multilinear forms, 45
 of partitions, 85
Zero one law, 70