Notes on ergodic theory

Michael Hochman[1]

January 27, 2013

[1]Please report any errors to mhochman@math.huji.ac.il

Contents

Preface

These are notes from an introductory course on ergodic theory given at the Hebrew University of Jerusalem in the fall semester of 2012.

The course covers the usual basic subjects, though relatively little about entropy (a subject that was covered in a course the previous year). On the less standard side, we have included a discussion of Furstenberg disjointness.

Chapter 1

Introduction

At its most basic level, dynamical systems theory is about understanding the long-term behavior of a map $T : X \to X$ under iteration. X is called the *phase space* and the points $x \in X$ may be imagined to represent the possible states of the “system”. The map T determines how the system evolves with time: time is discrete, and from state x it transitions to state Tx in one unit of time. Thus if at time 0 the system is in state x, then the state at all future times $t = 1, 2, 3, \ldots$ are determined: at time $t = 1$ it will be in state Tx, at time $t = 2$ in state $T(Tx) = T^2x$, and so on; in general we define

$$T^n x = \underbrace{T \circ T \circ \ldots \circ T}_{n}(x)$$

so $T^n x$ is the state of the system at time n, assuming that at time zero it is in state x. The “future” trajectory of an initial point x is called the (forward) orbit, denoted

$$O_T(x) = \{x, Tx, T^2x, \ldots\}$$

When T is invertible, $y = T^{-1}x$ satisfies $Ty = x$, so it represents the state of the world at time $t = -1$, and we write $T^{-n} = (T^{-1})^n = (T^n)^{-1}$. The one can also consider the full or two-sided orbit

$$O_T^{\pm}(x) = \{T^n x \,:\, n \in \mathbb{Z}\}$$

There are many questions one can ask. Does a point $x \in X$ necessarily return close to itself at some future time, and how often this happens? If we fix another set A, how often does x visit A? If we cannot answer this for all points, we would like to know the answer at least for typical points. What is the behavior of pairs of points $x, y \in X$: do they come close to each other? given another pair x', y', is there some future time when x is close to x' and y is close to y'? If $f : X \to \mathbb{R}$, how well does the value of f at time 0 predict its value at future times? How does randomness arise from deterministic evolution of time? And so on.

The set-theoretic framework developed so far there is relatively little that can be said besides trivialities, but things become more interesting when more structure is given to X and T. For example, X may be a topological space, and T continuous; or X may be a compact manifold and T a differentiable map (or k-times differentiable for some k); or there may be a measure on X and T may preserve it (we will come give a precise definition shortly). The first of these settings is called *topological dynamics,* the second *smooth dynamics*, and the last is *ergodic theory*. Our main focus in this course is ergodic theory, though we will also touch on some subjects in topological dynamics.

One might ask why these various assumptions are natural ones to make. First, in many cases, all these structures are present. In particular a theorem of Liouville from celestial mechanics states that for Hamiltonian systems, e.g. systems governed by Newton's laws, all these assumptions are satisfied. Another example comes from the algebraic setting of flows on homogeneous spaces. At the same time, in some situations only some of these structures is available; an example is can be found in the applications of ergodic theory to combinatorics, where there is no smooth structure in sight. Thus the study of these assumptions individually is motivated by more than mathematical curiosity.

In these notes we focus primarily on ergodic theory, which is in a sense the most general of these theories. It is also the one with the most analytical flavor, and a surprisingly rich theory emerges from fairly modest axioms. The purpose of this course is to develop some of these fundamental results. We will also touch upon some applications and connections with dynamics on compact metric spaces.

Chapter 2

Measure preserving transformations

2.1 Measure preserving transformations

Our main object of study is the following.

Definition 2.1.1. A measure preserving system is a quadruple $\mathcal{X} = (X, \mathcal{B}, \mu, T)$ where $(X, \mathcal{B}, \mu)$ is a probability space, and $T : X \to X$ is a measurable, measure-preserving map: that is

$$T^{-1}A \in \mathcal{B} \qquad \text{and} \qquad \mu(T^{-1}A) = \mu(A) \qquad \text{for all } A \in \mathcal{B}$$

If T is invertible and T^{-1} is measurable then it satisfies the same conditions, and the system is called invertible.

Example 2.1.2. Let X be a finite set with the σ-algebra of all subsets and normalized counting measure μ, and $T : X \to X$ a bijection. This is a measure preserving system, since measurability is trivial and

$$\mu(T^{-1}A) = \frac{1}{|X|}|T^{-1}A| = \frac{1}{|X|}|A| = \mu(A)$$

This example is very trivial but many of the phenomena we will encounter can already be observed (and usually are easy to prove) for finite systems. It is worth keeping this example in mind.

Example 2.1.3. The identity map on any measure space is measure preserving.

Example 2.1.4 (Circle rotation)**.** Let $X = S^1$ with the Borel sets $\mathcal{B}$ and normalized length measure μ. Let $\alpha \in \mathbb{R}$ and let $R_\alpha : S^1 \to S^1$ denote the rotation by angle α, that is, $z \mapsto e^{2\pi i\alpha}z$ (if $\alpha \notin 2\pi\mathbb{Z}$ then this map is not the identity). Then R_α preserves μ; indeed, it transforms intervals to intervals of equal length. If we consider the algebra of half-open intervals with endpoints in $\mathbb{Q}[\alpha]$, then T preserves this algebra and preserves the measure on it, hence it preserves the extension of the measure to $\mathcal{B}$, which is μ.

This example is sometimes described as $X = \mathbb{R}/\mathbb{Z}$, then the map is written additively, $x \mapsto x + \alpha$.

This example has the following generalization: let G be a compact group with normalized Haar measure μ, fix $g \in G$, and consider $R_g : G \to G$ given by $x \to gx$. To see that $\mu(T^{-1}A) = \mu(A)$, let $\nu(A) = \mu(g^{-1}A)$, and note that ν is a Borel probability measure that is right invariant: for any $h \in H$, $\nu(Bh) = \mu(g^{-1}Bh) = \mu(g^{-1}B) = \nu(B)$. This $\nu = \mu$.

Example 2.1.5 (Doubling map)**.** Let $X = [0,1]$ with the Borel sets and Lebesgue measure, and let $Tx = 2x \bmod 1$. This map is onto is ,not 1-1, in fact every point has two pre-images which differ by $\frac{1}{2}$, except for 1, which is not in the image. To see that T_2 preserves μ, note that for any interval $I = [a, a+r) \subseteq [0,1)$,

$$T_2^{-1}[a, a+r) = [\frac{a}{2}, \frac{a+r}{2}) \cup [\frac{a}{2} + \frac{1}{2}, \frac{a+r}{2} + \frac{1}{2})$$

which is the union of two intervals of length half the length; the total length is unchanged.

Note that TI is generally of larger length than I; the property of measure preservation is defined by $\mu(T^{-1}A) = \mu(A)$.

This example generalizes easily to $T_a x = ax \bmod 1$ for any $1 < a \in \mathbb{N}$. For non-integer $a > 1$ Lebesgue measure is not preserved.

If we identify $[0,1)$ with $\mathbb{R}/\mathbb{Z}$ then the example above coincides with the endomorphism $x \mapsto 2x$ of the compact group $\mathbb{R}/\mathbb{Z}$. More generally one can consider a compact group G with Haar measure μ and an endomorphism $T : G \to G$. Then from uniqueness of Haar measure one again can show that T preserves μ.

Example 2.1.6. (Symbolic spaces and product measures) Let A be a finite set, $|A| \geq 2$, which we think of as a discrete topological space. Let $X^+ = A^{\mathbb{N}}$ and $X = A^{\mathbb{Z}}$ with the product σ-algebras. In both cases there is a map which shifts "to the right",

$$(\sigma x)_n = x_{n+1}$$

In the case of X this is an invertible map (the inverse is $(\sigma x)_n = x_{n-1}$). In the one-sided case X^+, the shift is not 1-1 since for every sequence $x = x_1 x_2 \ldots \in A^{\mathbb{N}}$ we have $\sigma^{-1}(x) = \{x_0 x_1 x_2 \ldots : x_0 \in A\}$.

Let p be a probability measure on A and $\mu = p^{\mathbb{Z}}$, $\mu^+ = p^{\mathbb{N}}$ the product measures on X, X^+, respectively. By considering the algebra of cylinder sets $[a] = \{x : x_i = a_i\}$, where a is a finite sequence of symbols, one may verify that σ preserves the measure.

Example 2.1.7. (Stationary processes) In probability theory, a sequence $\{\xi_n\}_{n=1}^{\infty}$ of random variables is called *stationary* if the distribution of a consecutive n-tuple $(\xi_k, \ldots, \xi_{k+n-1})$ does not depend on where it behind; i.e. $(\xi_1, \ldots, \xi_n) =$

$(\xi_k, \ldots, \xi_{k+n-1})$ in distribution for every k and n. Intuitively this means that if we observe a finite sample from the process, the values that we see give no information about when the sample was taken.

From a probabilistic point of view it rarely matters what the sample space is and one may as well choose it to be $(X, \mathcal{B}) = (Y^{\mathbb{N}}, \mathcal{C}^{\mathbb{N}})$, where $(Y, \mathcal{C})$ is the range of the variables. On this space there is again defined the shift map $\sigma : X \to X$ given by $\sigma((y_n)_{n=1}^{\infty}) = (y_{n+1})_{n=1}^{\infty}$. For any $A_1, \ldots, A_n \in \mathcal{C}$ and k let

$$A^i = \underbrace{Y \times \ldots \times Y}_{k} \times A_1 \times \ldots \times A_n \times Y \times Y \times Y \times \ldots$$

Note that $\mathcal{B}$ is generated by the family of such sets. If P is the underlying probability measure, then stationarity means that for any $A_1, \ldots, A_n$ and k,

$$P(A^0) = P(A^k)$$

Since $A^k = \sigma^{-k} A^0$ this shows that the family of sets B such that $P(\sigma^{-1}B) = P(B)$ contains all the sets of the form above. Since this family is a σ-algebra and the sets above generate $\mathcal{B}$, we see that σ preserves P.

There is a converse to this: suppose that P is a σ-invariant measure on $X = Y^{\mathbb{N}}$. Define $\xi_n(y) = y_n$. Then (ξ_n) is a stationary process.

Example 2.1.8. (Hamiltonian systems) The notion of a measure-preserving system emerged from the following class of examples. Let $\Omega = \mathbb{R}^{2n}$; we denote $\omega \in \Omega$ by $\omega = (p, q)$ where $p, q \in \mathbb{R}^n$. Classically, p describes the positions of particles and q their momenta. Let $H : \Omega \to \mathbb{R}$ be a smooth map and consider the differential equation

$$\begin{aligned} \frac{d}{dt} p_i &= -\frac{\partial H}{\partial q_i} \\ \frac{d}{dt} \dot{q}_i &= \frac{\partial H}{\partial p_i} \end{aligned}$$

Under suitable assumptions, for every initial state $\omega = (p_0, q_0) \in \Omega$ and $t \in \mathbb{R}$ there is determines a unique solution $\gamma_\omega(t) = (p(t), q(t))$, and $\omega_t = \gamma_\omega(t)$ is the state of the world after evolving for a period of t started from ω.

Thinking of t as fixed, we have defined a map $T_t : \Omega \to \Omega$ by $T_t\omega = \gamma_\omega(t)$. Clearly

$$T_0(\omega) = \gamma_\omega(0) = \omega$$

We claim that this is an action of $\mathbb{R}$. Indeed, notice that $\sigma(s) = \gamma_\omega(t + s)$ satisfies $\sigma(0) = \gamma_\omega(t) = \omega_t$ and $\dot{\sigma}(s) = \dot{\gamma_{\omega_t}}(t + s)$, and so $A(\sigma, \dot{\sigma}) = A(\gamma_\omega(t + s), \dot{\gamma}_\omega(t + s)) = 0$. Thus by uniqueness of the solution, $\gamma_{\omega_t}(s) = \gamma_\omega(t + s)$. This translates to

$$T_{t+s}(\omega) = \gamma_\omega(t + s) = \gamma_{\omega_t}(s) = T_s\omega_t = T_s(T_t\omega)$$

and of course also $T_{t+s} = T_{s+t} = T_t T_s \omega$. Thus $(T_t)_{t \in \mathbb{R}}$ is action of $\mathbb{R}$ on Ω.

It often happens that Ω contains compact subsets which are invariant under the action. For example there may be a notion of energy $E : \Omega \to \mathbb{R}$ that is preserved, i.e. $E(T_t\omega) = E(\omega)$, and then the level sets $M = E^{-1}(e_0)$ are invariant under the action. E is nice enough, M will be a smooth and often compact manifold. Furthermore, by a remarkable theorem of Liouville, if the equation governing the evolution is a Hamiltonian equation (as is the case in classical mechanics) then the flow preserves volume, i.e. $\mathrm{vol}(T_tU) = \mathrm{vol}(U)$ for every t and open (or Borel) set U. The same is true for the volume form on M.

2.2 Recurrence

One of deep and basic properties of measure preserving systems is that they display "recurrence", meaning, roughly, that for typical x, anything that happens along its orbit happens infinitely often. This phenomenon was first discovered by Poincaré and bears his name.

Given a set A and $x \in A$ it will be convenient to say that x returns to A if $T^n x \in A$ for some $n > 0$; this is the same as $x \in A \cap T^{-n}A$. We say that x returns for A infinitely often if there are infinitely many such n.

The following proposition is, essentially, the pigeon-hole principle.

Proposition 2.2.1. *Let A be a measurable set, $\mu(A) > 0$. Then there is an n such that $\mu(A \cap T^{-n}A) > 0$.*

Proof. Consider the sets $A, T^{-1}A, T^{-2}A, \ldots, T^{-k}A$. Since T is measure preserving, all the sets $T^{-i}A$ have measure $\mu(A)$, so for $k > 1/\mu(A)$ they cannot be pairwise disjoint mod μ (if they were then $1 \geq \mu(X) \geq \sum_{i=1}^{k} \mu(T^{-i}A) > 1$, which is impossible). Therefore there are indices $0 \leq i < j \leq k$ such that $\mu(T^{-i}A \cap T^{-j}A) > 0$. Now,

$$T^{-i}A \cap T^{-j}A = T^{-i}(A \cap T^{-(j-i)}A)$$

so $\mu(A \cap T^{-(j-i)}A) > 0$, as desired. □

Theorem 2.2.2 (Poincare recurrence theorem)**.** *If $\mu(A) > 0$ then μ-a.e. $x \in A$ returns to A.*

Proof. Let

$$E = \{x \in A \, : \, T^n x \notin A \text{ for } n > 0\} = A \setminus \bigcup_{n=1}^{\infty} T^{-n}A$$

Thus $E \subseteq A$ and $T^{-n}E \cap E \subseteq T^{-n}E \cap A = \emptyset$ for $n \geq 1$ by definition. Therefore by the previous corollary, $\mu(E) = 0$. □

Corollary 2.2.3. *If $\mu(A) > 0$ then μ-a.e. $x \in A$ returns to A infinitely often.*

Proof. Let E be as in the previous proof. For any k-tuple $n_1 < n_2 < \ldots < n_k$, the set of points $x \in A$ which return to A only at times $n_1, \ldots, n_k$ satisfy $T^{n_k} x \in E$. Therefore,

$$\{x \in A \,:\, x \text{ returns to } A \text{ finitely often}\} = \bigcup_k \bigcup_{n_1 < \ldots < n_k} T^{-n_k} E$$

Hence the set on the left is the countable union of set of measure 0. □

In order to discuss of recurrence for individual points we suppose now assume that X is a metric space.

Definition 2.2.4. Let X be a metric space and $T : X \to X$. Then $x \in X$ is called *forward recurrent* if there is a sequence $n_k \to \infty$ such that $T^{n_k} x \to x$.

Proposition 2.2.5. *Let $(X, \mathcal{B}, \mu, T)$ by a measure-preserving system where X is a separable metric space and the open sets are measurable. Thenμ-a.e. x is forward recurrent.*

Proof. Let $A_i = B_{r_i}(x_i)$ be a countable sequence of balls that generate the topology. By Theorem 2.2.2, there are sets $A_i' \subseteq A_i$ of full measure such that every $x \in A_i'$ returns to A_i. Let $X_0 = X \setminus \bigcup (A_i \setminus A_i')$, which is of full μ-measure. For $x \in X_0$ if $x \in A_i$ then x returns to A_i, so it returns to within $|\operatorname{diam} A_n|$ of itself. Since x belongs to A_n of arbitrarily small diameter, x is recurrent. □

When the phenomenon of recurrence was discovered it created quite a stir. Indeed, by Liouville's theorem it applies to Hamiltonian systems, such as planetary systems and the motion of molecules in a gas. In these settings, Poincaré recurrence seems to imply that the system is stable in the strong sense that it nearly returns to the same configuration infinitely often. This question arose original in the context of stability of the solar system in a weaker sense, i.e., will it persist indefinitely or will the planets eventually collide with the sun, or fly off into deep space. Stability in the strong sense above contradicts our experience. One thing to note, however, is the time frame for this recurrence is enormous, and in the celestial-mechanical or thermodynamics context it does not say anything about the short-term stability of the systems.

Recurrence also implies that there are no quantities that only increase as time moves forwards; this is on the face of it in contradiction of the second law of thermodynamics, which asserts that the thermodynamic entropy of a mechanical system increases monotonely over time. A function $f : X \to \mathbb{R}$ is increasing (respectively, constant) along orbits if $f(Tx) \geq f(x)$ a.e. (respectively $f(Tx) = f(x)$ a.e.). This is the same as requiring that for a.e. x the sequence $f(x), f(Tx), f(T^2 x), \ldots$ is non-decreasing (respectively constant). Although superficially stronger, the latter condition follows because for fixed n,

$$\begin{aligned}
\mu(x \,:\, f(T^{n+1}(x)) < f(T^n x)) &= \mu(T^{-n}\{x \,:\, f(Tx) < f(x)\}) \\
&= \mu(x \,:\, f(Tx) < f(x)) \\
&= 0
\end{aligned}$$

and so the intersection of these events is still of measure zero. The same argument works for functions constant along orbits.

Corollary 2.2.6. *In a measure preserving system any measurable function that is increasing along orbits is a.s. constant along orbits.*

Proof. Let f be increasing along orbits. For $\delta > 0$ let

$$J(\delta) = \{x \in X \,:\, f(Tx) \geq f(x) + \delta\}$$

We must show that $\mu(J(\delta)) = 0$ for all $\delta > 0$, since then $\mu(\bigcup_{n=1}^{\infty} J(1/n)) = 0$, which implies that $f(Tx) = f(x)$ for a.e. x.

Suppose there were some $\delta > 0$ such that $\mu(J(\delta)) > 0$. For $k \in \mathbb{Z}$ let

$$J(\delta, k) = \{x \in J(\delta) \,:\, k\frac{\delta}{2} \leq f(x) < (k+1)\frac{\delta}{2}\}$$

Notice that $J(\delta) = \bigcup_{k\in\mathbb{Z}} J(\delta, k)$, so there is some k with $\mu(J(\delta, k)) > 0$. On the other hand, if $x \in J(\delta, k)$ then

$$f(Tx) \geq f(x) + \delta \geq k\frac{\delta}{2} + \delta > (k+1)\frac{\delta}{2}$$

so $Tx \notin J(\delta, k)$. Similarly for any $n \geq 1$, since f is increasing on orbits, $f(T^n x) \geq f(Tx) > (k+1)\frac{\delta}{2}$, so $T^n x \notin J(\delta, k)$. We have shown that no point of $J(\delta, k)$ returns to $J(\delta, k)$, contradicting Poincaré recurrence. □

The last result highlights the importance of measurability. Using the axiom of choice one can easily choose a representative x from each orbit, and using it define $f(T^n x) = n$ for $n \geq 0$ (and also $n < 0$ if T is invertible). Then we have a function which is strictly increasing along orbits; but by the corollary, it cannot be measurable.

2.3 Induced action on functions and measures

Given a map $T : X \to Y$ there is an induced map $\widehat{T}$ on functions with domain Y, given by

$$\widehat{T}f(x) = f(Tx)$$

On the space $f : Y \to \mathbb{R}$ or $f : Y \to \mathbb{C}$ the operator $\widehat{T}$ has some obvious properties: it is linear, positive ($f \geq 0$ implies $\widehat{T}f \geq 0$), multiplicative ($\widehat{T}(fg) = \widehat{T}f \cdot \widehat{T}g$). Also $|\widehat{T}f| = \widehat{T}|f|$ and $\widehat{T}(f^c) = (\widehat{T}f)^c$.

When $(X, \mathcal{B})$ and $(Y, \mathcal{C})$ are measurable spaces and T is measurable, the induced map $\widehat{T}$ acts on the space of measurable functions on Y.

Similarly, in the measurable setting T induces a map on measures. Write $\mathcal{M}(X)$ and $\mathcal{P}(X)$ for the spaces signed measures and probability measures on $(X, \mathcal{B})$, respectively, and similarly for Y. Then $\widehat{T} : \mathcal{M}(X) \to \mathcal{M}(Y)$ is given by

$$(\widehat{T}\mu)(A) = \mu(T^{-1}A) \qquad \text{for measurable } A \subseteq Y$$

This is called the push-forward of μ and is sometimes denoted $T_*\mu$ or $T_\#\mu$. It is easy to check that $\widehat{T}\mu \in \mathcal{M}(Y)$ and that $\mu \mapsto \widehat{T}\mu$ this is a measure on Y. It is easy to see that this operator is also linear, i.e. $\widehat{T}(a\mu + b\nu) = a\widehat{T}\mu + b\widehat{T}\nu$ for scalars a, b.

Lemma 2.3.1. *$\nu = \widehat{T}\mu$ is the unique measure satisfying $\int f\,d\nu = \int \widehat{T}f\,d\mu$ for every bounded measurable function $f : Y \to \mathbb{R}$ (or for every $f \in C(X)$ if X is a compact).*

Proof. For $A \in \mathcal{C}$ note that $\widehat{T}1_A(x) = 1_A(Tx) = 1_{T^{-1}A}(x)$, hence

$$\int \widehat{T}1_A\,d\mu = \mu(T^{-1}A) = (\widehat{T}\mu)(A) = \int 1_A\,d\widehat{T}\mu$$

This shows that $\nu = \widehat{T}\mu$ has the stated property when f is an indicator function. Every bounded measurable function (and in particular every continuous function if X is compact) is the pointwise limit of uniformly bounded sequence of linear combinations of indicator functions, so the same holds by dominated convergence (note that $f_n \to f$ implies $\widehat{T}f_n \to \widehat{T}f$).

Uniqueness follows from the fact that ν is determined by the values of $\int f\,d\nu$ as f ranges over bounded measurable functions,or, when X is compact, over continuous functions. □

Corollary 2.3.2. *Let $(X, \mathcal{B}, \mu)$ be a measure space and $T : X \to X$ a measurable map. Then T preserves μ if and only if $\int f\,d\mu = \int \widehat{T}f\,d\mu$ for every bounded measurable $f : X \to \mathbb{R}$ (or $f \in C(X)$ if X is compact)*

Proof. T preserves μ if and only if $\mu = \widehat{T}\mu$, so this is a special case of the previous lemma. □

Proposition 2.3.3. *Let $f : X \to Y$ be a map between measurable spaces, $\mu \in \mathcal{P}(X)$ and $\nu \in \widehat{T}\mu \in \mathcal{P}(Y)$. Then $\widehat{T}$ maps $L^p(\nu)$ isometrically into $L^p(\mu)$ for every $1 \leq p \leq \infty$.*

Proof. First note that if f is an a.e. defined function then $\widehat{T}f$ is also, because if E is the nullset where f is not defined then $T^{-1}E$ is the set where $\widehat{T}f$ is not defined, and $\mu(T^{-1}E) = \nu(E) = 0$. Thus $\widehat{T}$ acts on equivalence classes of measurable functions mod μ. Now, for $1 \leq p < \infty$ we have

$$\left\|\widehat{T}f\right\|_p^p = \int |\widehat{T}f|^p\,d\mu = \int \widehat{T}(|f|^p)\,d\mu = \int |f^p|\,d\nu = \|f\|_p^p$$

For $p = \infty$ the claim follows from the identity $\|f\|_\infty = \lim_{p\to\infty} \|f\|_p$. □

Corollary 2.3.4. *In a measure preserving system $\widehat{T}$ is a norm-preserving self-map of L^p, and if T is invertible then $\widehat{T}$ is an isometry of L^p.*

The operator $\widehat{T}$ on L^2 is sometimes called the *Koopman operator*. When T is invertible it is a unitary operator and opens up the door for using spectral techniques to study the underlying system. We will return to this idea later.

We will follow the usual convention and write T instead of $\widehat{T}$. This introduces slight ambiguity but the meaning should usually be clear from he context.

2.4 Dynamics on metric spaces

Many (perhaps most) spaces and maps studied in ergodic theory have additional topological structure, and there is a developed dynamical theory for system of this kind. Here we will discuss only a few aspects of it, especially those which are related to ergodic theory.

Definition 2.4.1. A topological dynamical system is a pair (X,T) where X is a compact metric space and $T: X \to X$ is continuous.

It is sometimes useful to allow compact non-metrizable spaces but in this course we shall not encounter them.

Before we begin discussing such systems we review some properties of the space of measures. Let $\mathcal{M}(X)$ denote the linear space of signed (finite) Borel measures on X and $\mathcal{P}(X) \subseteq \mathcal{M}(X)$ the convex space of Borel probability measures. Two measures $\mu, \nu \in \mathcal{M}(X)$ are equal if and only if $\int f d\mu = \int f d\nu$ for all $f \in C(X)$, so the maps $\mu \mapsto \int f d\mu$, $f \in C(X)$, separate points.

Definition 2.4.2. The weak-* topology on $\mathcal{M}(X)$ (or $\mathcal{P}(X)$) is the weakest topology that make the maps $\mu \mapsto \int f\, d\mu$ continuous for all $f \in C(X)$. In particular,

$$\mu_n \to \mu \text{ if and only if } \int f d\mu_n \to \int f d\mu \text{ for all } f \in C(X)$$

Proposition 2.4.3. *The weak-* topology is metrizable and compact.*

For the proof see Appendix 9.

Let (X,T) be a topological dynamical system. It is clear that the induced map T on functions preserves the space $C(X)$ of continuous functions.

Lemma 2.4.4. *$T: C(X) \to C(X)$ is contracting in $\|\cdot\|_\infty$, and if the original map $T: X \to X$ is onto, the induced $T: C(X) \to C(X)$ is an isometry. $T: \mathcal{P}(X) \to \mathcal{P}(X)$ is continuous.*

Proof. The first part follows from

$$\|Tf\|_\infty = \max_{x \in X} |f(T(x))| = \max_{y \in T(X)} |fyx)| \le \|f\|_\infty$$

and the fact that there is equality if $TX = X$.

For the second part, if $\mu_n \to \mu$ then for $f \in C(X)$,

$$\int f\, dT\mu_n = \int f \circ T\, d\mu_n \to \int f \circ T\, d\mu = \int f\, dT\mu$$

This shows that $T\mu_n \to T\mu$, so T is continuous. $\square$

The following result is why ergodic theory is useful in studying topological systems.

Proposition 2.4.5. *Every topological dynamical system (X, T) has invariant measures.*

Proof. Let $x \in X$ be an arbitrary initial point and define

$$\mu_N = \frac{1}{N}\sum_{n=0}^{N-1} \delta_{T^n x}$$

Note that

$$\int f\, d\mu_N = \frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)$$

Passing to a subsequence $N(k) \to \infty$ we can assume by compactness that $\mu_{N(k)} \to \mu \in \mathcal{P}(X)$. We must show that $\int f\, d\mu = \int f \circ T\, d\mu$ for all $f \in C(X)$. Now,

$$\begin{aligned}
\int f\, d\mu - \int f \circ T\, d\mu &= \lim_{k\to\infty} \int (f - f \circ T)\, d\mu_{N(l)} \\
&= \lim_{k\to\infty} \frac{1}{N(k)} \sum_{n=0}^{N(k)-1} \int (f - f\circ T)(T^n x) \\
&= \lim_{k\to\infty} \frac{1}{N(k)} \left(f(T^{N(k)-1}x) - f(x) \right) \\
&= 0
\end{aligned}$$

because f is bounded. □

There are a number of common variations of this proof. We could have defined $\mu_N = \frac{1}{N}\sum_{n=0}^{N-1} \delta_{T^n x_N}$ (with the initial point x_N varying with N), of begun with an arbitrary measure μ and $\mu_N = \frac{1}{N}\sum_{n=0}^{N-1} T^n \mu$. The proof would then show that any accumulation point of μ_N is T-invariant.

We denote the space of T-invariant measures by $\mathcal{P}_T(X)$.

Corollary 2.4.6. *In a topological dynamical system (X, T), $\mathcal{P}_T(X)$ is non-empty, compact and convex.*

Proof. We already showed that it is non-empty, and convexity is trivial. For compactness we need only show it is closed. We know that

$$\mathcal{P}_T(X) = \bigcap_{f \in C(X)} \{\mu \in \mathcal{P}(X) \,:\, \int (f - f\circ T) d\mu = 0\}$$

Each of the sets in the intersection is the pre-image of 0 under the map $\mu \mapsto \int (f - f \circ T) d\mu$; since $f - f \circ T$ is continuous this map is continuous and so $\mathcal{P}_T(X)$ is the intersection of closed sets, hence closed. □

Corollary 2.4.7. *Every topological dynamical system (X, T) contains recurrent points.*

Proof. Choose any invariant measure $\mu \in \mathcal{P}_T(X)$ and apply Proposition 2.2.5 to the measure preserving system $(X, \mathcal{B}, \mu, T)$. □

2.5 Some technicalities

Much of ergodic theory holds in general probability spaces, but some of the results require assumptions on the measure space in order to avoid pathologies. There are two possible and theories available which for our purposes are essentially equivalent: the theory of Borel spaces and of Lebesgue spaces. We will work in the Borel category. In this section we review without proof the main facts we will use. These belong to descriptive set theory and we will not prove them.

Definition 2.5.1. A Polish space is an uncountable topological space whose topology is induced by a complete, separable metric.

Note that a metric space can be Polish even if the metric isn't complete. For example $[0,1)$ is not complete in the usual metric but it is homeomorphic to $[0,\infty)$, which is complete (and separable), and so $[0,1)$ is Polish.

Definition 2.5.2. A standard Borel space is a pair $(X,\mathcal{B})$ where X is an uncountable Polish space and $\mathcal{B}$ is the Borel σ-algebra. A standard measure space is a σ-finite measure on a Borel space.

Definition 2.5.3. Two measurable spaces $(X,\mathcal{B})$ and $(Y,\mathcal{C})$ are isomorphic if there is a bijection $f:X\to Y$ such that both f and f^{-1} are measurable.

Theorem 2.5.4. *Standard Borel spaces satisfy the following properties.*

1. *All Borel spaces are isomorphic.*
2. *Countable products of Borel spaces are Borel.*
3. *If A is an uncountable measurable subset of a Borel space, then the restriction of the σ-algebra to A again is a Borel space.*
4. *If f is a measurable injection (1-1 map) between Borel spaces then the image of a Borel set is Borel. In particular it is an isomorphism from the domain to its image.*

Another important operation on measure spaces is the factoring relation.

Definition 2.5.5. A factor between measurable spaces $(X,\mathcal{B})$ and $(Y,\mathcal{C})$ is a measurable onto map $f:X\to Y$. If there are measures μ,ν on X,Y, respectively, then the factor is required also to satisfy $f\mu=\nu$.

Given a factor $f:(X,\mathcal{B})\to(Y,\mathcal{C})$ between Borel spaces, we can pull back the σ-algebra $\mathcal{C}$ and obtain a sub-σ-algebra $\mathcal{E}\subseteq\mathcal{B}$ by

$$\mathcal{E}=f^{-1}\mathcal{C}=\{\pi^{-1}C\,:\,C\in\mathcal{C}\}$$

Note that $\mathcal{C}$ is countably generated (since it is isomorphic to the Borel σ-algebra of a separable metric space), so $\mathcal{E}$ is countably generated as well.

This procedure can to some extent be reversed. Let (X, B) be a Borel space and $\mathcal{E} \subseteq \mathcal{B}$ a countably generated sub-σ algebra. Partition X into the atoms of $\mathcal{E}$, that is, according to the the equivalence relation

$$x \sim y \qquad \Longleftrightarrow \qquad 1_E(x) = 1_E(y) \text{ for all } E \in \mathcal{E}$$

Let $\pi : X \to X/\sim$ denote the factor map and

$$\mathcal{E}/\sim = \{E \subseteq X/\sim : \pi^{-1}E \in \mathcal{E}\}$$

Then the quotient space $X/\mathcal{E} = (X/\sim, \mathcal{E}/\sim)$ is a measurable space and π is a factor map. Notice also that $\pi^{-1} : (\mathcal{E}/\sim) \to \mathcal{E}$ is 1-1, so $\mathcal{E}/\sim$ is countably generated. Also, the atoms of $\sim$ are measurable, since if f $\mathcal{E}$ is generated by $\{E_n\}$ then the atom of x is just $\bigcap F_n$ where $F_n = E_n$ if $x \in E_n$ and $F_n = X \setminus E_n$ otherwise. Hence $\mathcal{E}/\sim$ separates points in $X/\sim$.

These two operations are not true inverses of each other: it is not in general true that if $\mathcal{E} \subseteq \mathcal{B}$ is countably generated then $(X/\sim, \mathcal{E}/\sim)$ is a Borel space. But if one introduces a measure then it is true up to measure 0.

Theorem 2.5.6. *Let μ be a probability measure on a Borel space $(X, \mathcal{B}, \mu)$, and $\mathcal{E} \subseteq \mathcal{B}$ a countably generated infinite sub-σ-algebra. Then there is a measurable subset $X_0 \subseteq X$ of full measure such that the quotient space of $X_0/\mathcal{E}$ is a Borel space.*

As we mentioned above, there is an alternative theory available with many of the same properties, namely the theory of Lebesgue spaces. These are measure spaces arising as the completions of σ-finite measures on Borel spaces. In this theory all definitions are modulo sets of measure zero, and all of the properties above hold. In particular the last theorem can be stated more cleanly, since the removal of a set of measure 0 is implicit in the definitions. Many of the standard texts in ergodic theory work in this category. The disadvantage of Lebesgue spaces is that it makes it cumbersome to consider different measures on the same space, since the σ-algebra depends non-trivially on the measure. This is the primary reason we work in the Borel category.

Chapter 3

Ergodicity

3.1 Ergodicity

In this section and the following ones we will study how it may be decomposed into simpler systems.

Definition 3.1.1. Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system. A measurable set $A \subseteq X$ is invariant if $T^{-1}A = A$. The system is *ergodic* if there are no non-trivial invariant sets; i.e. every invariant set has measure 0 or 1.

If A is invariant then so is $X \setminus A$. Indeed,

$$T^{-1}(X \setminus A) = T^{-1}X \setminus T^{-1}A = X \setminus A$$

Thus, ergodicity is an irreducibility condition: a non-ergodic system the dynamics splits into two (nontrivial) parts which do not "interact", in the sense that an orbit in one of them never enters the other.

Example 3.1.2. Let X be a finite set with normalized counting measure, and $T : X \to X$ a 1-1 map. If X consists of a single orbit then the system is ergodic, since any invariant set that is not empty contains the whole orbit. In general, X splits into the disjoint (finite) union of orbits, and each of these orbits is invariant and of positive measure. Thus the system is ergodic if and only if it consists of a single orbit.

Note that every (invertible) system splits into the disjoint union of orbits. However, these typically have measure zero, so do not in themselves prevent ergodicity.

Example 3.1.3. By taking disjoint unions of measure preserving systems with the normalized sum of the measures, one gets many examples of non-ergodic systems.

Definition 3.1.4. A function $f : X \to Y$ for some set Y is invariant if $f(Tx) = f(x)$ for all $x \in X$.

The primary example is 1_A when A is invariant.

Lemma 3.1.5. *The following are equivalent:*

1. $(X, \mathcal{B}, \mu, T)$ *is ergodic.*
2. *If* $T^{-1}A = A \bmod \mu$ *then* $\mu(A) = 0$ *or* 1.
3. *Every measurable invariant function is constant a.e.*
4. *If* $f \in L^1$ *and* $Tf = f$ *a.e. then* f *is a.e. constant.*

Proof. (1) and (3) are equivalent since an invariant set A produces the invariant function 1_A, while if f is invariant and not a.e. constant then there is a measurable set U in the range of f such that $0 < \mu(f^{-1}U) < 1$. But this set is clearly invariant.

Exactly the same argument shows that (2) and (4) are equivalent.

We complete the proof by showing the equivalence of (3) and (4). Clearly (4) implies (3). Conversely, suppose that $f \in L^1$ and $Tf = f$ a.e.. Let $g = \limsup f(T^n x)$. Clearly g is T-invariant (since $g(Tx)$ is the limsup of the shifted sequence $f(T^{n+1}x)$, and is the same as the limsup of $f(T^n x)$, which is $g(x)$). The proof will be done by showing that $g = f$ a.e. This is true at a point x if $f(T^n x) = f(x)$ for all $n \geq 0$, and for this it is enough that $f(T^{n+1}x) = f(T^n x)$ for all $n \geq 0$; equivalently, that $T^n x \in \{Tf = f\}$ for all n, i.e. that $x \in \bigcap T^{-n}\{Tf = f\}$. But this is an intersection of sets of measure 1 and hence holds for a.e. x, as desired. □

Example 3.1.6 (Irrational circle rotation)**.** Let $R_\alpha(x) = e^{2\pi i\alpha}x$ be an irrational circle rotation ($\alpha \notin \mathbb{Q}$) on S^1 with Lebesgue measure. We claim that this system is ergodic. Indeed, let $\chi_n(z) = z^n$ (the characters of the compact group S^1) and consider an invariant function $f \in L^\infty(\mu)$. Since $f \in L^2$, it can be represented in L^2 as a Fourier series $f = \sum a_n \chi_n$. Now,

$$T\chi_n(z) = (e^{2\pi i\alpha}z)^n = e^{2\pi i n\alpha}z^n = 2^{2\pi i n\alpha}\chi_n$$

so from

$$f = Tf = \sum a_n T\chi_n = \sum e^{2\pi i n\alpha} a_n \chi_n$$

Comparing this to the original expansion we have $a_n = e^{2\pi i n\alpha} a_n$. Thus if $a_n \neq 0$ then $e^{2\pi i n\alpha} = 1$, which, since $\alpha \notin \mathbb{Q}$, can occur only if $n = 0$. Thus $f = a_0\chi_0$, which is constant .

Non-ergodicity means that one can split the system into two parts that don't "interact". The next proposition reformulates this in a positive way: ergodicity means that every non-trivial sets do "interact".

Proposition 3.1.7. *The following are equivalent:*

1. $(X, \mathcal{B}, \mu, T)$ *is ergodic.*
2. *For any* $B \in \mathcal{B}$*, if* $\mu(B) > 0$ *then* $\bigcup_{n=N}^{\infty} T^{-n}B = X \bmod \mu$ *for every* N.

3. *If $A, B \in \mathcal{B}$ and $\mu(A), \mu(B) > 0$ then $\mu(A \cap T^{-n}B) > 0$ for infinitely many n.*

Proof. (1) implies (2): Given B let $B' = \bigcup_{n=N}^{\infty} T^{-n}B$ and note that

$$T^{-1}(B') = \bigcup_{n=N}^{\infty} T^{-1}T^{-n}B = \bigcup_{n=N+1}^{\infty} T^{-n}B \subseteq B'$$

Since $\mu(T^{-1}B') = \mu(B')$ we have $B' = T^{-1}B' \bmod \mu$, hence by ergodicity $B' = X \bmod \mu$.

(2) implies (3): Given A, B as in (3) we conclude from (2) that, for every N, $\mu(A \cap \bigcup_{n=N}^{\infty} T^{-n}B) = \mu(A)$, hence there some $n > N$ with $\mu(T^{-n}A) > 0$. This implies that there are infinitely many such n.

Finally if (3) holds and if A is invariant and $\mu(A) > 0$, then taking $B = X \setminus A$ clearly $A \cap \bigcup T^{-n}B = \emptyset$ for all n so $\mu(B) = 0$ by (3). Thus every invariant set is trivial. $\square$

3.2 Mixing

Although a wide variety of ergodic systems can be constructed or shown abstractly to exist, it is surprisingly difficult to verify ergodicity of naturally arising systems. In fact, in most cases where ergodicity can be proved because the system satisfies a stronger "mixing" property.

Definition 3.2.1. $(X, \mathcal{B}, \mu, T)$ is called *mixing* if for every pair A, B of measurable sets,

$$\mu(A \cap T^{-n}B) \to \mu(A)\mu(B) \qquad \text{as } n \to \infty$$

It is immediate from the definition that mixing systems are ergodic. The advantage of mixing over ergodicity is that it is enough to verify it for a "dense" family of sets A, B. It is better to formulate this in a functional way.

Lemma 3.2.2. *For fixed $f \in L^2$ and n, the map $(f, g) \mapsto \int f \cdot T^n g \, d\mu$ is multilinear and $\left\| \int f \cdot T^n g \, d\mu \right\|_2 \leq \|f\|_2 \|g\|_2$.*

Proof. Using Cauchy-Schwartz and the previous lemma,

$$\int f \cdot T^n g \, d\mu \leq \|f\|_2 \|T^n g\|_2 = \|f\|_2 \|g\|_2$$

$\square$

Proposition 3.2.3. *$(X, \mathcal{B}, \mu, T)$ is mixing if and only if for every $f, g \in L^2$,*

$$\int f \cdot T^n g \, d\mu \to \int f \, d\mu \cdot \int g \, d\mu \qquad \text{as } n \to \infty$$

Furthermore this limit holds for ail $f, g \in L^2$ if and only if it holds for f, g in a dense subset of L^2.

Proof. We prove the second statement first. Suppose the limit holds for $f, g \in V$ with $V \subseteq L^2$ dense. Now let $f, g \in L^2$ and for $\varepsilon > 0$ let $f', g' \in V$ with $\|f - f'\| < \varepsilon$ and $\|g - g'\| < \varepsilon$. Then

$$\begin{aligned}
\left\| \int f \cdot T^n g \, d\mu \right\| &\leq \left\| \int (f - f' + f') \cdot T^n (g - g' + g') \, d\mu \right\| \\
&\leq \left\| \int (f - f') \cdot T^n g \, d\mu \right\| + \left\| \int f \cdot T^n (g - g') \, d\mu \right\| + \\
&\quad + \left\| \int (f - f') \cdot T^n (g - g') \, d\mu \right\| + \left\| \int f' \cdot T^n g' \, d\mu \right\| \\
&\leq \varepsilon \|g\| + \|f\| \varepsilon + \varepsilon^2 + \left\| \int f' \cdot T^n g' \, d\mu \right\|
\end{aligned}$$

Since $\left\| \int f' \cdot T^n g' \, d\mu \right\| \to 0$ and ε was arbitrary this shows that $\left\| \int f \cdot T^n g \, d\mu \right\| \to 0$, as desired.

For the first part, using the identities $\int 1_A \, d\mu = \mu(A)$, $T^n 1_A = 1_{T^{-n}A}$ and $1_A 1_B = 1_{A \cap B}$, we see that mixing is equivalent to the limit above for indicator functions, and since the integral is multilinear in f, g it holds for linear combinations of indicator functions and these combinations are dense in L^2, we are done by what we proved above. $\square$

Example 3.2.4. Let $X = A^{\mathbb{Z}}$ for a finite set A, take the product σ-algebra, and μ a product measure with marginal given by a probability vector $p = (p_a)_{a \in A}$. Let $\sigma : X \to X$ be the shift map $(\sigma x)_n = x_{n+1}$. We claim that this map is mixing and hence ergodic.

To prove this note that if $f(x) = \widetilde{f}(x_1, \ldots, x_k)$ depends on the first k coordinates of the input, then $\sigma^n f(x) = \widetilde{f}(x_{k+1}, \ldots, x_{k+n})$. If f, g are two such functions then for n large enough, $\sigma^n g$ and f depend on different coordinates, and hence, because μ is a product measure, they are independent in the sense of probability theory:

$$\int f \cdot \sigma^n g \, d\mu = \int f \, d\mu \cdot \int \sigma^n g \, d\mu = \int f \, d\mu \cdot \int g \, d\mu$$

so the same is true when taking $n \to \infty$. Mixing follows from the previous proposition.

3.3 Kac's return time formula

We pause to give a nice application of ergodicity to estimation of the "recurrence rate" of points to a set.

Let $(X, \mathcal{B}, \mu, T)$ be ergodic and let $\mu(A) > 0$. Since $X_0 = \bigcup_{n=1}^{\infty} T^{-n} A$, and its measure is at least $\mu(A)$ which is positive, by ergodicity $\mu(X_0) = 1$. Thus for a.e. there is a minimal $n \geq 1$ with $T^n x \in A$; we denote this number by $r_A(x)$

and note that r_A is measurable, since

$$\{r_A < k\} = A \cap (\bigcup_{1\leq i<k} T^{-i}A)$$

Theorem 3.3.1 (Kac's formula)**.** *Assume that T is invertible. Then $\int_A r_A\, d\mu = 1$; in particular, $\mathbb{E}(r_A|A) = 1/\mu(A)$, so the expected time to return to A starting from A is $1/\mu(A)$.*

Proof. Let $A_n = A \cap \{r_A = n\}$. Then

$$\int_A r_A d\mu = \sum_{n=1}^{\infty} n\mu(A_n) = \sum_{n=1}^{\infty}\sum_{k=1}^{n} \mu(T^k A_n)$$

The proof will be completed by showing that the sets $\{T^k A_n : n \in \mathbb{N}, \ldots 1 \leq k \leq n\}$ are pairwise disjoint and that their union has full measure. Indeed, for a.e. $x \in X$ there is a least $m \geq 1$ such that $y = T^{-m}x \in A$. Let $n = r_A(y)$. Clearly $m \leq n$, since if $n < m$ and $T^n y \in A$ then $T^n y = T^n T^{-m} x = T^{-(m-n)}x \in A$ and $m - n \geq 1$ is smaller than m. Thus $x \in T^m A_n$. This shows that the union of the given family is X up to a null set.

To show that the family is disjoint, suppose $x \in T^{m'}A_{n'}$ for some $(m', n') \neq (m, n)$. We cannot have $m' < m$ because then $T^{-m'}x \in A_{n'} \subseteq A$ would contradict minimality of m. We cannot have $m' > m$ because this would imply that $r_A(T^{-m'}x) \geq m' > m$, and at the same time $T^{-m}x = T^{m'-m}(T^{-m'}x) \in A_n \subseteq A$, implying $r_A(T^{-m'}x) \leq m' - m < m'$, a contradiction. Finally, $m = m'$ and $n \neq n'$ is impossible because then then $T^{-m}x \in A_n \cap A_{n'}$, despite $A_n \cap A_{n'} \neq \emptyset$. □

Even under the stated ergodicity assumption this result strengthens Poincare recurrence. First, it shows now only that a.e. $x \in A$ returns to A, if shows that it does so in finite expected time, and identifies this expectation. Simple examples show that the formula is incorrect in the non-ergodic case.

The invertability assumption is not necessary. We shall later see how to remove it.

3.4 Ergodic measures as extreme points

It is clear that $\mathcal{P}_T(X)$ is convex; in this section we will prove a nice algebraic characterization of the ergodic measures as precisely the extreme points of $\mathcal{P}_T(X)$. Recall that a point in a convex set is an extreme point if it cannot be written as a convex combination of other points in the set.

Proof. Let $f = d\nu/d\mu$. Given t let $E = \{f < t\}$; it suffices to show that this set is invariant μ-a.e. We first claim that the sets $E \setminus T^{-1}E$ and $T^{-1}E \setminus E$ are of the same μ-measure. Indeed,

$$\begin{aligned}
\mu(E \setminus T^{-1}E) &= \mu(E) - \mu(E \cap T^{-1}E) \\
\mu(T^{-1}E \setminus E) &= \mu(T^{-1}E) - \mu(E \cap T^{-1}E)
\end{aligned}$$

and since $\mu(E) = \mu(T^{-1}E)$, the right hand sides are equal, and hence also the left hand sides.

Now

$$\nu(E) = \int_E f\,d\mu = \int_{E\cap T^{-1}E} f\,d\mu + \int_{E\setminus T^{-1}E} f\,d\mu$$

On the other hand

$$\nu(E) = \nu(T^{-1}E) = \int_{T^{-1}E\cap E} f\,d\mu + \int_{(T^{-1}E)\setminus E} f\,d\mu$$

Subtracting we find that

$$\int_{E\setminus T^{-1}E} f\,d\mu = \int_{T^{-1}E\setminus E} f\,d\mu$$

On the left hand side the integral is over a subset of E, where $f < t$, so the integral is $< t\mu(E \setminus T^{-1}E)$; on the right it is over a subset of $X \setminus E$, where $f > t$, so the integral is $\geq t\mu(T^{-1}E \setminus E)$. Equality is possible only if the measure of these sets is 0, and since $\mu(E) = \mu(T^{-1}E)$, the set difference can be a μ-nullset if and only if $E = T^{-1}E \bmod \mu$, which is the desired invariance of E. □

Remark 3.4.1. If T is invertible, there is an easier argument: since $T\mu = \mu$ and $T\nu = \nu$ we have $dT\nu/dT\mu = d\nu/d\mu = f$. Now, for any measurable set A,

$$\int_A dT\nu = \int 1_A dT\nu = \int 1_A \circ T\,d\nu = \int 1_A \circ T\,f d\mu = \int 1_A f \circ T^{-1}\,dT\mu = \int_A f \circ T^{-1}\,dT\mu$$

This shows that $f \circ T^{-1} = dT\nu/dT\mu = f$. But of course we have used invertability.

Proposition 3.4.2. *The ergodic invariant measures are precisely the extreme points of $\mathcal{P}_T(X)$.*

Proof. If $\mu \in \mathcal{P}_T(X)$ is non-ergodic then there is an invariant set A with $0 < \mu(A) < 1$. Then $B = X \setminus A$ is also invariant. Let $\mu_A = \frac{1}{\mu(A)}\mu|_A$ and $\mu_B = \frac{1}{\mu(B)}\mu|_B$ denote the normalized restriction of μ to A, B. Clearly $\mu = \mu(A)\mu_A + \mu(B)\mu_B$, so μ is a convex combination of μ_A, μ_B, and these measures are invariant:

$$\begin{aligned}
\mu_A(T^{-1}E) &= \frac{1}{\mu(A)}\mu(A \cap T^{-1}E) \\
&= \frac{1}{\mu(A)}\mu(T^{-1}A \cap T^{-1}E) \\
&= \frac{1}{\mu(A)}\mu(T^{-1}(A \cap E)) \\
&= \frac{1}{\mu(A)}\mu(A \cap E) \\
&= \mu_A(E)
\end{aligned}$$

Thus μ is not an extreme point of $\mathcal{P}_T(X)$.

Conversely, suppose that $\mu = \alpha\nu + (1-\alpha)\theta$ for $\nu, \theta \in \mathcal{P}_T(X)$ and $\nu \neq \mu$. Clearly $\mu(E) = 0$ implies $\nu(E) = 0$, so $\nu \ll \mu$, and by the previous lemma $f = d\nu/d\mu \in L^1(\mu)$ is invariant. Since $1 = \nu(X) = \int f d\mu$, we know that $f \neq 0$, and since $\nu \neq \mu$ we know that f is not constant. Hence μ is not ergodic by Lemma 3.1.5. □

As an application we find that distinct ergodic measures are also separated at the spacial level:

Corollary 3.4.3. *Let μ, ν be ergodic measures for a measurable map T of a measurable space $(X, \mathcal{B})$. Then either $\mu = \nu$ or $\mu \perp \nu$.*

Proof. Suppose $\mu \neq \nu$ and let $\theta = \frac{1}{2}\mu + \frac{1}{2}\nu$. Since this is a nontrivial representation of θ as a convex combination, it is not ergodic, so there is a nontrivial invariant set A. By ergodicity, A must have μ-measure 0 or 1 and similarly for ν. They cannot be both 0 since this would imply $\theta(A) = 0$, and they cannot both have measure 1, since this would imply $\theta(A) = 1$. Therefore one is 0 and one is 1. This implies that A supports one of the measures and $X \setminus A$ the other, so $\mu \perp \nu$. □

3.5 Ergodic decomposition I

Having described those systems that are "indecomposable", we now turn to study how a non-ergodic system may decompose into ergodic ones. One can begin such a decomposition immediately from the definitions: if $\mu \in \mathcal{P}_T(X)$ is not ergodic then there is an invariant set A and μ is a convex combination of the invariant measures $\mu_A, \mu_{X\setminus A}$, which are supported on disjoint invariant sets. If μ_A, μ_B are not ergodic we can split each of them further as a convex combination of mutually singular invariant measures. Iterating this procedure we obtain representations of μ as convex combinations of increasingly "fine" mutually singular measures. While at each finite stage the component measures need not be ergodic, in a sense they are getting closer to being ergodic, since at each stage we eliminate a potential invariant set. One would like to pass to a limit, in some sense, and represent the original measure is a convex combination of ergodic ones.

Example 3.5.1. If T is a bijection of a finite set with normalized counting measure then the measure splits as a convex combination of uniform measures on orbits, each of which is ergodic.

In general, it is too much to ask that a measure be a convex combination of ergodic ones.

Example 3.5.2. Let $X = [0, 1]$ with Borel sets and Lebesgue measure μ and T the identity map. In this case the only ergodic measures are atomic, so we cannot write μ as a finite convex combination of ergodic measures.

The idea of decomposing μ is also motivated by the characterization of ergodic measures as the extreme points of $\mathcal{P}_T$ and the fact that in finite-dimensional vector spaces a point in a convex set is a convex combination of extreme points. There are also infinite-dimensional versions of this, and if $\mathcal{P}_T(X)$ can be made into a compact convex set satisfying some other mild conditions one can apply Choquet's theorem. However, we will take a more measure-theoretic approach through the measure integration and disintegration.

3.6 Measure integration

Given a measurable space $(X, \mathcal{B})$, a family $\{\nu_x\}_{x\in X}$ of probability measures on $(Y, \mathcal{C})$ is measurable if for every $E \in \mathcal{C}$ the map $x \mapsto \nu_x(E)$ is measurable (with respect to $\mathcal{B}$). Equivalently, for every bounded measurable function $f : Y \to \mathbb{R}$, the map $x \mapsto \int f(y)\, d\nu_x(y)$ is measurable.

Given a measure $\mu \in \mathcal{P}(X)$ we can define the probability measure $\nu = \int \nu_x d\mu(x)$ on Y by

$$\nu(E) = \int \nu_x(E)\, d\mu(x)$$

For bounded measurable $f : Y \to \mathbb{R}$ this gives

$$\int f\, d\nu = \int (\int f\, d\nu_x)\, d\mu(x)$$

and the same holds for $f \in L^1(\nu)$ by approximation (although f is defined only on a set E of full ν-measure, we have $\nu_x(E) = 1$ for μ-a.e. x, so the inner integral is well defined μ-a.e.).

Example 3.6.1. Let X be finite and $\mathcal{B} = 2^X$. Then

$$\int \nu_x\, d\mu(x) = \sum_{x\in X} \mu(x) \cdot \nu_x$$

Any convex combination of measures on Y can be represented this way, so the definition above generalizes convex combinations.

Example 3.6.2. Any measure μ on $(X, \mathcal{B})$ the family $\{\delta_x\}_{x\in X}$ is measurable since $\delta_x(E) = 1_E(x)$, and $\mu = \int \delta_x\, d\mu(x)$ because

$$\mu(X) = \int 1_E(x) d\mu(x) = \int \nu_x(E)\, d\mu(x)$$

In this case the parameter space was the same as the target space.

In particular, this representation shows that Lebesgue measure on $[0, 1]$ is an integral of ergodic measures for the identity map.

Example 3.6.3. $X = [0,1]$ and $Y = [0,1]^2$. For $x \in [0,1]$ let ν_x be Lebesgue measure on the fiber $\{x\} \times [0,1]$. Measurability is verified using the definition of the product σ-algebra, and by Fubini's theorem

$$\nu(E) = \int \nu_x(E)d\mu(x) = \int_0^1 \int_0^1 1_E(x,y)dy\,dx = \int\int_E 1dxdy$$

so ν is just Lebesgue measure on $[0,1]^2$.

One could also represent ν as $\int \nu_{x,y}\,d\nu(x,y)$ where $\nu_{x,y} = \nu_x$. Written this way each fiber measure appears many times.

3.7 Measure disintegration

We now reverse the procedure above and study how a measure may be decomposed as an integral of other measures. Specifically, we will study the decomposition of a measure with respect to a partition.

Example 3.7.1. Let $(X, \mathcal{B}, \mu)$ be a probability space and let $P = \{P_1, \ldots, P_n\}$ be finite partition of it, i.e. P_i are measurable, $P_i \cap P_j = \emptyset$ for $i \neq j$, and $X = \bigcup P_i$. For simplicity assume also that $\mu(P_i) > 0$. Let $\mathcal{P}(x)$ denote the unique P_i containing x and let μ_x denote the conditional measure on it, $\mu_x = \frac{1}{\mu(\mathcal{P}(x))}\mu|_{\mathcal{P}(x)}$. Then it is easy to check that $\mu = \int \mu_x\,d\mu(x)$.

Alternatively we can define $Y = \{1, \ldots, n\}$ with a probability measure given by $P(\{i\}) = \mu(P_i)$. Let $\mu_i = \frac{1}{\mu(P_i)}\mu|_{P_i}$. Then $\mu = \sum \mu(P_i)\mu_i = \int \mu_i\,dP(i)$.

Our goal is to give a similar decomposition of a measure with respect to an infinite (usually uncountable) partition $\mathcal{E}$ of X. Then the partition elements $E \in \mathcal{E}$ typically have measure 0, and the formula $\frac{1}{\mu(E)}\mu|_E$ no longer makes sense. As in probability theory one can define the conditional probability of an event E given that $x \in E$ as the conditional expectation $\mathbb{E}(1_E|\mathcal{P})$ evaluated at x (conditional expectation is reviewed in the Appendix). This would appear to give the desired decomposition: define $\mu_x(E) = \mathbb{E}(1_E|\mathcal{E})(x)$. For any countable algebra this does give a countably additive measure defined for μ-a.e. x. The problem is that $\mu_x(E)$ is defined only for a.e. x but we want to define $\mu_x(E)$ for all measurable sets. Overcoming this problem is a technical but nontrivial chore which will occupy us for the rest of the section.

For a measurable space $(X, \mathcal{B})$ and a sub-σ-algebra $\mathcal{E} \subseteq \mathcal{B}$ generated by a countable sequence $\{E_n\}$. Write $x \sim_{\mathcal{E}} y$ if $1_E(x) = 1_E(y)$ for every $E \in \mathcal{E}$, or equivalently, $1_{E_n}(x) = 1_{E_n}(y)$ for all n. This is an equivalence relation. The atoms of $\mathcal{E}$ are by definition the equivalence classes of $\sim_{\mathcal{E}}$, which are measurable, being intersections of sequences F_n of the form $F_n \in \{E_n, X \setminus E_n\}$. We denote $\mathcal{E}(x)$ the atom containing x.

In the next theorem we assume that the space is compact, which makes the Riesz representation theorem available as a means for of defining measures. We shall discuss this restriction afterwards.

Theorem 3.7.2. *Let X be compact metric space, $\mathcal{B}$ the Borel algebra, and $\mathcal{E} \subseteq \mathcal{B}$ a countably generated sub-σ-algebra. Then there is an $\mathcal{E}-$measurable family $\{\mu_y\}_{y \in X} \subseteq \mathcal{P}(X)$ such that μ_y is supported on $\mathcal{E}(y)$ and*

$$\mu = \int \mu_y \, d\mu(y)$$

Furthermore if $\{\mu'_y\}_{y \in X}$ is another such system then $\mu_y = \mu'_y$ a.e.

Note that $\mathcal{E}$-measurability has the following consequence: For μ-a.e. y, for every $y' \in \mathcal{E}(y)$ we have $\mu_{y'} = \mu_y$ (and, since since $\mu_y(\mathcal{E}(y)) = 1$, it follows that $\mu_{y'} = \mu_y$ for μ_y-a.e y').

Definition 3.7.3. The representation $\mu = \int \mu_y \, d\mu(y)$ in the proof is often called the *disintegration* of μ over $\mathcal{E}$.

We adopt the convention that y denotes the variable of $\mathcal{E}$-measurable functions.

Let $V \subseteq C(X)$ be a countable dense $\mathbb{Q}$-linear subspace with $1 \in V$. For $f \in V$ let

$$\overline{f} = \mathbb{E}(f|\mathcal{E})$$

(see the Appendix for a discussion of conditional expectation). Since V is countable there is a subset $X_0 \subseteq X$ of full measure such that $\overline{f}$ is defined everywhere on X_0 for $f \in V$ and $f \mapsto \overline{f}$ is $\mathbb{Q}$-linear and positive on X_0, and $\overline{1} = 1$ on X_0. Thus, for $y \in X_0$ the functions $\Lambda_y : V \to \mathbb{R}$ given by

$$\Lambda_y(f) = \overline{f}(y)$$

are positive $\mathbb{Q}$-linear functionals on the normed space $(V, \|\cdot\|_\infty)$, and they are continuous, since by positivity of conditional expectation $\left\|\overline{f}\right\|_\infty \leq \|f\|_\infty$. Thus Λ_y extends to a positive $\mathbb{R}$-linear functional$\Lambda_y : C(X) \to \mathbb{R}$. Note that $\Lambda_y 1 = \overline{1}(y) = 1$. Hence, by the Riesz representation theorem, there exists $\mu_y \in \mathcal{P}(X)$ such that

$$\Lambda_y f = \int f(x) \, d\mu_y(x)$$

For $y \in X \setminus X_0$ define μ_y to be some fixed measure to ensure measurability.

Proposition 3.7.4. *$y \to \mu_y$ is $\mathcal{E}$-measurable and $\mathbb{E}(1_A|\mathcal{E})(y) = \mu_y(A)$ μ-a.e., for every $A \in \mathcal{B}$.*

Proof. Let $\mathcal{A} \subseteq \mathcal{B}$ denote the family of sets $A \in \mathcal{B}$ such that $y \mapsto \mu_y(A)$ measurable from $(X, \mathcal{E})$ to $(X, \mathcal{B})$ and $\mathbb{E}(1_A|\mathcal{E})(y) = \mu_y(A)$ μ-a.e. We want to show that $\mathcal{A} = \mathcal{B}$.

Let $\mathcal{A}_0 \subseteq \mathcal{B}$ denote the family of sets $A \subseteq X$ such that 1_A is a pointwise limit of a uniformly bounded sequence of continuous functions. First, $\mathcal{A}_0$ is an algebra: clearly $X, \emptyset \in \mathcal{A}$, if $f_n \to 1_A$ then $1 - f_n \to 1_{X \setminus A}$, and if also $g_n \to 1_B$ then $f_n g_n \to 1_A 1_B = 1_{A \cap B}$.

We claim that $\mathcal{A}_0 \subseteq \mathcal{A}$. Indeed, if $f_n \to 1_A$ and $\|f_n\|_\infty \leq C$ then

$$\int f_n \, d\mu_y \to \int 1_A \, d\mu_y = \mu_y(A)$$

by dominated convergence, so $y \mapsto \mu_y(A)$ is the pointwise limit of the functions $y \mapsto \int f_n \, d\mu_y$, which are the same a.e. as the measurable functions $\overline{f}_n = \mathbb{E}(f_n|\mathcal{E}) : (X, \mathcal{E}) \to (X, \mathcal{B})$. This establishes measurability of the limit function $y \mapsto \mu_y(A)$ and also proves that this function is $\mathbb{E}(1_A|\mathcal{E})$ a.e., since $\mathbb{E}(\cdot|\mathcal{E})$ is continuous in L^1 and $f_n \to 1_A$ boundedly. This proves $\mathcal{A}_0 \subseteq \mathcal{A}$.

Now, $\mathcal{A}_0$ contains the closed sets, since if $A \subseteq X$ then $1_A = \lim f_n$ for $f_n(x) = \exp(-n \cdot d(x, A))$. Thus $\mathcal{A}_0$ generates the Borel σ-algebra $\mathcal{B}$.

Finally, we claim that $\mathcal{A}$ is a monotone class. Indeed, if $A_1 \subseteq A_2 \subseteq \ldots$ belong to $\mathcal{B}'$ and $A = \bigcup A_n$, then $\mu_y(A) = \lim \mu_y(A_n)$, and so $y \mapsto \mu_y(A)$ is the pointwise limit of the measurable functions $y \mapsto \mu_y(A_n)$. The latter functions are just $\mathbb{E}(1_{A_n}|\mathcal{E})$ and, since $1_{A_n} \to 1_A$ in L^1, by continuity of conditional expectation, $\mathbb{E}(1_{A_n}|\mathcal{E}) \to \mathbb{E}(1_A|\mathcal{E})$ in L^1. Hence $\mu_y(A) = \mathbb{E}(1_A|\mathcal{E})$ a.e. as desired.

Since $\mathcal{A}$ is a monotone class containing the sub-algebra of $\mathcal{A}_0$ and $\mathcal{A}_0$ generates $\mathcal{B}$, by the monotone class theorem we have $\mathcal{B} \subseteq \mathcal{A}$. Thus $\mathcal{A} = \mathcal{B}$, as desired. $\square$

Proposition 3.7.5. $\mathbb{E}(f|\mathcal{E})(y) = \int f \, d\mu_y$ *μ-a.e. for every* $f \in L^1(\mu)$.

Proof. We know that this holds for $f = 1_A$ by the previous proposition. Both sides of the claimed equality are linear and continuous under monotone increasing sequences. Approximating by simple functions this gives the claim for positive $f \in L^1$ and, taking differences, for all $f \in L^1$. $\square$

Proposition 3.7.6. μ_y *is μ-a.s. supported* $\mathcal{E}(y)$*, that is,* $\mu_y(\mathcal{E}(y)) = 1$ *ν-a.e.*

Proof. For $E \in \mathcal{E}$ we have

$$1_E(y) = \mathbb{E}(1_E|\mathcal{E})(y) = \int 1_E \, d\mu_y = \mu_y(E)$$

and it follows that $\mu_y(E) = 1_E(y)$ a.e. Let $\{E_n\}_{n=1}^\infty$ generate $\mathcal{E}$, and choose a set of full measure on which the above holds for all $E = E_n$. For y in this set let $F_n \in \{E_n, X \setminus E_n\}$ be such that $\mathcal{E}(y) = \bigcap F_n$. By the above $\mu_y(F_n) = 1$, and so $\mu_y(\mathcal{E}(y)) = 1$, as claimed. $\square$

Proposition 3.7.7. *If* $\{\mu'_y\}_{y \in Y}$ *is another family with the same properties then* $\mu'_y = \mu_y$ *for μ-a.e.* y.

Proof. For $f \in L^1(\mu)$ define $f'(y) = \int f \, d\mu'_y$. This is clearly a linear operator defined on $L^1(X, \mathcal{B}, \mu)$, and its range is $L^1(X, \mathcal{E}, \mu)$ because

$$\int |f'| \, d\mu \leq \int (\int |f| \, d\mu_y) \, d\mu(y) = \int |f| \, d\mu = \|f\|_1$$

The same calculation shows that $\int f'\,d\mu = \int f\,d\mu$. Finally, for $E \in \mathcal{E}$ we know that μ_y is supported on E for μ-a.e. $y \in E$ and on $X \setminus E$ for μ-a.e. $y \in X \setminus E$. Thus μ-a.s. we have

$$(1_E f)'(y) = \int 1_E f\, d\mu'_y = 1_E(y) \int f\, d\mu'_y = 1_E \cdot f'$$

By a well-known characterization of conditional expectation, $f' = \mathbb{E}(f|\mathcal{E}) = \overline{f}$ (see the Appendix). $\square$

It remains to address the compactness assumption on X. Examples show that one the disintegration theorem does require some assumption; it does not hold for arbitrary measure spaces and sub-σ-algebras. We will not eliminate the compactness assumption so much as explain why it is not a large restriction.

We can now formulate the disintegration theorem as follows.

Theorem 3.7.8. *Let μ be a probability measure on a standard Borel space $(X, \mathcal{B}, \mu)$ and $\mathcal{E} \subseteq \mathcal{B}$ a countably generated sub-σ-algebra. Then there is an $\mathcal{E}$-measurable family $\{\mu_y\}_{y \in Y} \subseteq \mathcal{P}(X, \mathcal{B})$ such that μ_y is supported on $\mathcal{E}(y)$ and*

$$\mu = \int \mu_y \, d\mu(y)$$

Furthermore if $\{\mu'_y\}_{y \in X}$ is another such system then $\mu_y = \mu'_y$.

3.8 Ergodic decomposition II

Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system on a Borel space. Let $\mathcal{I} \subseteq \mathcal{B}$ denote the family of T-invariant measurable sets. It is easy to check that $\mathcal{I}$ is a σ-algebra.

The σ-algebra $\mathcal{I}$ in general is not countably generated. Consider for example the case of an invertible ergodic transformation on a Borel space, such as an irrational circle rotation or two-sided Bernoulli shift. Then $\mathcal{I}$ consists only of sets of measure 0 and 1. If $\mathcal{I}$ were countably generated by $\{I_n\}_{n=1}^{\infty}$, say, then for each n either $\mu(I_n) = 1$ or $\mu(X \setminus I_n) = 1$. Set $F_n = I_n$ or $F_n = X \setminus I_n$ according to these possibilities. Then $F = \bigcap F_n$ is an invariant set of measure 1 and is an atom of $\mathcal{I}$. But the atoms of $\mathcal{I}$ are the orbits, since each point in X is measurable and hence every countable set is. But this would imply that μ is supported on a single countable orbit, contradicting the assumption that it is non-atomic.

We shall work instead with a fixed countably generated μ-dense sub-σ-algebra $\mathcal{I}_0$ of $\mathcal{I}$. Let $L^1(X, \mathcal{I}, \mu)$ is a closed subspace of $L^1(X, \mathcal{B}, \mu)$, and since the latter is separable, so is the former. Choose a dense countable sequence $f_n \in L^1(X, \mathcal{I}, \mu)$, choosing representatives of the functions that are genuinely $\mathcal{I}$ measurable, not just modulo a $\mathcal{B}$-measurable nullset. Now consider the countable family of sets $A_{n,p,q} = \{p < f_n < q\}$, where $p, q \in \mathbb{Q}$, and let $\mathcal{I}_0$ be the σ-algebra that they generate. Clearly $\mathcal{I}_0 \subseteq \mathcal{I}$ and all of the f_n are $\mathcal{I}_0$-measurable, so $L^1(X, \mathcal{I}_0, \mu) = L^1(X, \mathcal{I}, \mu)$. In particular, $\mathcal{I}$ is contained in the μ-completion of $\mathcal{I}_0$.

Theorem 3.8.1 (Ergodic decomposition theorem)**.** *Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system on a Borel space and let $\mathcal{I}, \mathcal{I}_0$ be as above. Then there is an $\mathcal{I}_0$-measurable (and in particular $\mathcal{I}$-measurable) disintegration $\mu = \int \mu_x \, d\mu(x)$ of μ such that a.e. μ_y is T-invariant, ergodic, and supported on $\mathcal{I}_0(y)$. Furthermore the representation is unique in the sense that if $\{\mu'_y\}$ is any other family with the same properties then $\mu_y = \mu'_y$ for μ-a.e. y.*

Let $\{\mu_y\}_{y \in X}$ be the disintegration of μ relative to $\mathcal{I}_0$, we need only show that for μ-a.e. y the measure μ_y is T-invariant and ergodic.

Claim 3.8.2. For μ-a.e. y, μ_y is T-invariant.

Proof. Define $\mu'_y = T\mu_y$. This is an $\mathcal{E}$ measurable family since for any $E \in \mathcal{B}$, $\mu'_y(E) = \mu_y(T^{-1}E)$ so measurability of $y \mapsto \mu'_y(E)$ follows from that of $y \mapsto \mu_y(E)$. We claim that $\{\mu'_y\}_{y \in X}$ is a disintegration of μ over $\mathcal{I}_0$. Indeed, for any $E \in \mathcal{B}$,

$$\begin{aligned}
\int (\int \mu'_y(E)) \, d\mu(y) &= \int (\int \mu_y(T^{-1}E)) \, d\mu(y) \\
&= \mu(T^{-1}E) \\
&= \mu(E)
\end{aligned}$$

Also $T^{-1}\mathcal{I}_0(y) = \mathcal{I}_0(y)$ (since $\mathcal{I}_0(y) \in \mathcal{I}$) so

$$\mu'_y(\mathcal{I}_0(y)) = \mu_y(T^{-1}\mathcal{I}_0(y)) = \mu_y(\mathcal{I}_0(y)) = 1$$

so μ'_y is supported on $\mathcal{E}(y)$. Thus, $\{\mu'_y\}_{y \in X}$ is an $\mathcal{E}$-measurable disintegration of μ, so $\mu'_y = \mu_y$ a.e. This is exactly the same as a.e. invariance of μ_y. $\square$

Claim 3.8.3. For μ-a.e. y, μ_y is ergodic.

Proof. This can be proved by purely measure-theoretic means, but we will give a proof that uses the mean ergodic theorem, Theorem 4.2.3 below. Let $\mathcal{F} \subseteq C(X)$ be a dense countable family. Then

$$\frac{1}{N} \sum_{n=1}^{N} T^n f \to \mathbb{E}(f|\mathcal{I}) = \mathbb{E}(f|\mathcal{I}_0)$$

in $L^2(\mathcal{B}, \mu)$. For each $f \in \mathcal{F}$, we can ensure that this holds a.e. along an appropriate subsequence, and by a diagonal argument we can construct a subsequence $N_k \to \infty$ such that $\frac{1}{N_k} \sum_{n=1}^{N_k} T^n f \to \mathbb{E}(f|\mathcal{I}_0)$ for all $f \in \mathcal{F}$, a.e. Since $\mu = \int \mu_y d\mu(y)$ this holds μ_y-a.e. for μ-a.e. y. Now, for such a y, in the measure preserving system $(X, \mathcal{B}, \mu_y, T)$, for $f \in \mathcal{F}$ we have $\frac{1}{N_k} \sum_{n=1}^{N_k} T^n f \to \mathbb{E}_{\mu_y}(f|\mathcal{I})$ in L^2; since $f \in \mathcal{F}$ is bounded and the limit is a.s. equal to $\mathbb{E}(f|\mathcal{I}_0)$, we have

$\mathbb{E}_{\mu_y}(f|\mathcal{I}) = \mathbb{E}(f|\mathcal{I}_0)$ μ-a.e. But the right hand side is $\mathcal{I}_0$-measurable, hence μ_y-a.e. constant. We have found that for $f \in \mathcal{F}$ the conditional expectation $\mathbb{E}_{\mu_y}(f|\mathcal{I})$ is μ_y-a.e. constant. $\mathcal{F}$ is dense in $C(X)$ and therefore in $L^1(\mathcal{B}, \mu_y)$, and $\mathbb{E}_{\mu_y}(\cdot|\mathcal{I})$ is continuous, we the image of $\mathbb{E}_{\mu_y}(\cdot|\mathcal{I})$ is contained in the constant functions. But if $g \in L^1(\mathcal{B}, \mu_y)$ is invariant it is $\mathcal{I}$-measurable and $\mathbb{E}_{\mu_y}(g|\mathcal{I}) = g$ is constant. Thus all invariant functions in $L^1(\mathcal{B}, \mu_y)$ are constant, which implies that $(X, \mathcal{B}, \mu_y, T)$ is ergodic. □

Our formulation of the ergodic decomposition theorem represents μ as an integral of ergodic measures parametrized by $y \in X$ (in an $\mathcal{I}$-measurable way). Sometimes the following formulation is given, in which $\mathcal{P}_T(X)$ is given the σ-algebra generated by the maps $\mu \mapsto \mu(E)$, $E \in \mathcal{B}$; this coincides with the Borel structure induced by the weak-* topology when X is given the structure of a compact metric space. One can show that the set of ergodic measures is measurable, for example because in the topological representation they are the extreme points of a weak-* compact convex set.

Theorem 3.8.4 (Ergodic decomposition, second version)**.** *Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system on a Borel space. Then there is a unique probability measure θ on $\mathcal{P}_T(X)$ supported on the ergodic measure and such that $\mu = \int \nu \, d\theta(\nu)$.*

Chapter 4

The ergodic theorem

4.1 Preliminaries

We have seen that in a measure preserving system, a.e. $x \in A$ returns to A infinitely often. Now we will see that more is true: these returns occur with a definite frequency which, in the ergodic case, is just $\mu(A)$; in the non-ergodic case the limit is $\mu_x(A)$, where μ_x is the ergodic component to which x belongs.

This phenomenon is better formulated at an analytic level in terms of averages of functions along an orbit. To this end let us introduce some notation. Let $T : V \to V$ be a linear operator of a normed space V, and suppose T is a contraction, i.e. $\|Tf\| \leq \|f\|$. This is the case when T is induced from a measure-preserving transformation (in fact we have equality). For $v \in V$ define

$$S_N v = \frac{1}{N} \sum_{n=0}^{N-1} T^n v$$

Note that in the dynamical setting, the frequency of visits x to A up to time N is $S_N 1_A(x) = \frac{1}{N}\sum_{n=0}^{N-1} 1_A(T^n x)$. Clearly S_N is linear, and since T is a contraction $\|T^n v\| \leq \|v\|$ for $n \geq 1$, so by the triangle inequality, $\|S_N v\| \leq \frac{1}{N}\sum_{n=0}^{N-1} \|T^n v\| \leq \|v\|$. Thus S_N are also contractions. This has the following useful consequence.

Lemma 4.1.1. *Let $T : V \to V$ as above and let $S : V \to V$ be another bounded linear operator. Suppose that $V_0 \subseteq V$ is a dense subset and that $S_N v \to Sv$ as $N \to \infty$ for all $v \in V_0$. Then the same is true for all $v \in V$.*

Proof. Let $v \in V$ and $w \in V_0$. Then

$$\limsup_{N\to\infty} \|S_N v - Sv\| \leq \limsup_{N\to\infty} \|S_N v - S_N w\| + \limsup_{N\to\infty} \|S_N w - Sv\|$$

Since $\|S_N v - S_N w\| = \|S_N(v - w)\| \leq \|v - w\|$ and $S_N w \to Sw$ (because $w \in V_0$), we have

$$\limsup_{N\to\infty} \|S_N v - Sv\| \leq \|v - w\| + \|Sw - Sv\| \leq (1 + \|S\|) \cdot \|v - w\|$$

Since $\|v - w\|$ can be made arbitrarily small, the lemma follows. $\square$

4.2 Mean ergodic theorem

Historically, the first ergodic theorem is von-Neuman's "mean" ergodic theorem, which can be formulated in a purely Hilbert-space setting (and it is not hard to adapt it to L^P). Recall that if $T : V \to V$ is a bounded linear operator of a Hilbert space then $T^* : V \to V$ is the adjoint operator, characterized by $\langle v, Tw\rangle = \langle T^*v, w\rangle$ for $v, w \in V$, and satisfies $\|T^*\| = \|T\|$.

Lemma 4.2.1. *Let $T : V \to V$ be a contracting linear operator of a Hilbert space. Then $v \in V$ is T-invariant if and only if it is T^*-invariant.*

Remark 4.2.2. When T is unitary (which is one of the main cases of interest to us) this lemma is trivial. Note however that without the contraction assumption this is false even in $\mathbb{R}^d$.

Proof. Since $(T^*)^* = T$ it suffices to prove that $T^*v = v$ implies $Tv = v$.

$$\begin{aligned}
\|v - Tv\|^2 &= \langle v - Tv, v - Tv\rangle \\
&= \|v\|^2 + \|Tv\|^2 - \langle Tv, v\rangle - \langle v, Tv\rangle \\
&= \|v\|^2 + \|Tv\|^2 - \langle v, T^*v\rangle - \langle T^*v, v\rangle \\
&= \|v\|^2 + \|Tv\|^2 - \langle v, v\rangle - \langle v, v\rangle \\
&= \|Tv\|^2 - \|v\|^2 \\
&\leq 0
\end{aligned}$$

where the last inequality is because T is a contraction. $\square$

Theorem 4.2.3 (Hilbert-space mean ergodic theorem)**.** *Let T be a linear contraction of a Hilbert space V, i.e. $\|Tv\| \leq \|v\|$. Let $V_0 \leq V$ denote the closed subspace of T-invariant vectors (i.e. $V_0 = \ker(T - I)$) and π the orthogonal projection to V_0. Then*

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n v \to \pi v \qquad \textit{for all } v \in V$$

Proof. If $v \in V_0$ then $S_N v = v$ and so $S_N v \to v = \pi v$ trivially. Since $V = V_0 \oplus V_0^\perp$ and S_N is linear, it suffices for us to show that $S_N v \to 0$ for $v \in V_0^\perp$. The key insight is that $V_0^\perp$ can be identified as the space of *co-boundaries*,

$$V_0^\perp = \overline{\{v - Tv \,:\, v \in V\}} \tag{4.1}$$

assuming this, by Lemma 4.1.1 we must only show that $S_N(v - Tv) \to 0$ for $v \in V$, and this follows from

$$\begin{aligned} S_N(v - Tv) &= \frac{1}{N}\sum_{n=0}^{N-1} T^n(v - Tv) \\ &= \frac{1}{N}(w - T^{N+1}w) \\ &\to 0 \end{aligned}$$

where in the last step we used $\left\|w - T^{N+1}w\right\| \leq \|w\| + \left\|T^{N+1}w\right\| \leq 2\|w\|$.

To prove (4.1) it suffices to show that $w \perp \{v - Uv \,:\, v \in V\}$ implies $w \in V_0$. Suppose that $w \perp (v - Uv)$ for all $v \in V$. Since

$$\begin{aligned} \langle w, v - Uv\rangle &= \langle w, v\rangle - \langle w, Uv\rangle \\ &= \langle w, v\rangle - \langle U^* w, v\rangle \\ &= \langle w - U^* w, v\rangle \end{aligned}$$

we conclude that $\langle w - U^* w, v\rangle = 0$ for all $v \in V$, hence $w - U^* w = 0$. Hence $Uw = w$ and by the lemma $Uw = w$, as desired. □

Now let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system and let T denote also the Koopman operator induced on L^2 by T. Then the space V_0 of T-invariant vectors is just $L^2(X, \mathcal{I}, \mu)$, where $\mathcal{I} \subseteq \mathcal{B}$ is the σ-algebra of invariant sets, and the orthogonal projection π to V_0 is just the conditional expectation operator, $\pi f = \mathbb{E}(f|\mathcal{I})$ (see the Appendix). We derive the following:

Corollary 4.2.4 (Dynamical mean ergodic theorem)**.** *Let $(X, \mathcal{B}, \mu, T)$ be a measure-preserving system, let $\mathcal{I}$ denote the σ-algebra of invariant sets, and let π denote the orthogonal projection from $L(X, \mathcal{B}, \mu)$ to the closed subspace $L^2(X, \mathcal{I}, \mu)$. Then for every $f \in L^2$,*

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n f \to \mathbb{E}(f|\mathcal{I}) \qquad \text{in } L^2$$

In particular, if the system is ergodic then the limit is constant:

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n f \to \int f\, d\mu \qquad \text{in } L^2$$

Specializing to $f = 1_A$, and noting that L^2-convergence implies, for example, convergence in probability, the last result says that on an arbitrarily large part of the space, the frequency of visits of an orbit to A up to time N is arbitrarily close to $\mu(A)$, if N is large enough.

4.3 The pointwise ergodic theorem

Very shortly after von Neumann's mean ergodic theorem (and appearing in print before it), Birkhoff proved a stronger version in which convergence takes place a.e. and in L^1.

Theorem 4.3.1 (Pointwise ergodic theorem)**.** *Let $(X, \mathcal{B}, \mu, T)$ be a measure-preserving system, let $\mathcal{I}$ denote the σ-algebra of invariant sets. Then for any $f \in L^1(\mu)$,*

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n f \to \mathbb{E}(f|\mathcal{I}) \qquad \textit{a.e. and in } L^1$$

In particular, if the system is ergodic then the limit is constant:

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n f \to \int f\, d\mu \qquad \textit{a.e. and in } L^1$$

We shall see several proofs of this result. The first and most "standard" proof follows the same scheme as the mean ergodic theorem: one first establishes the statement for a dense subspace $V \subseteq L^1$, and then uses some continuity property to extend to all of L^1. The first step is nearly identical to the proof of the mean ergodic theorem.

Proposition 4.3.2. *There is a dense subspace $V \subseteq L^1$ such that the conclusion of the theorem holds for every $f \in V$.*

Proof. We temporarily work in L^2. Let V_1 denote the set of invariant $f \in L^2$, for which the theorem holds trivially because $S_N f = f$ for all N. Let $V_2 \subseteq L^2$ denote the linear span of functions of the form $f = g - Tg$ for $g \in L^\infty$. The theorem also holds for these, since

$$\left\|g + T^{N+1}g\right\|_\infty \leq \|g\|_\infty + \left\|T^{N+1}g\right\|_\infty = 2\,\|g\|_\infty$$

and therefore

$$\frac{1}{N}\sum_{n=0}^{N-1} T^n(g - Tg) = \frac{1}{N}(g - T^{N+1}g) \to 0 \qquad \text{a.e. and in } L^1$$

Set $V = V_1 + V_2$. By linearity of S_N, the theorem holds for $f \in V_1 + V_2$. Now, L^∞ is dense in L^2 and T is continuous on L^2, so $\overline{V}_2 = \overline{\{g - Tg \,:\, g \in L^2\}}$. In the proof of the mean ergodic theorem we saw that $L^2 = V_1 \oplus \overline{V}_2$, so $V = V_1 \oplus V_2$ is dense in L^2, and hence in L^1, as required. $\square$

By Lemma 4.1.1, this proves the ergodic theorem in the sense of L^1-convergence for all $f \in L^1$. In order to similarly extend the pointwise version to all of L^1 we need a little bit of "continuity", which is provided by the following.

Theorem 4.3.3 (Maximal inequality). *Let $f \in L^1$ with $f \geq 0$ and $S_N f = \frac{1}{N}\sum_{n=0}^{N-1} T^n f$. Then for every t,*

$$\mu\left(x : \sup_N S_N f(x) > t\right) \leq \frac{1}{t}\int f\, d\mu$$

Before giving the proof let us show how this finishes the proof of the ergodic theorem. Write $S = \mathbb{E}(\cdot|\mathcal{I})$, which is a bounded linear operator on L^1, let $f \in L^1$ and $g \in V$. Then

$$\begin{aligned} |S_N f - Sf| &\leq |S_N f - S_N g| + |S_N g - Sg| \\ &\leq S_N|f-g| + |S_N g - Sf| \end{aligned}$$

Now, $S_N g \to Sg$ a.e., hence $|S_N g - Sf| \to |S(g-f)| \leq S|f-g|$ a.e. Thus,

$$\limsup_{N\to\infty} |S_N f - Sf| \leq \limsup_{N\to\infty} S_N|f-g| + S|g-f|$$

If the left hand side is $> \varepsilon$ then at least one of the terms on the right is $> \varepsilon/2$. Therefore,

$$\mu\left(\limsup_{N\to\infty} |S_N f - Sf| > \varepsilon\right) \leq \mu\left(\limsup_{N\to\infty} S_N|f-g| > \varepsilon/2\right) + \mu\left(S|g-f| > \varepsilon/2\right)$$

Now, by the maximal inequality, the first term on the right side is bounded by $\frac{1}{\varepsilon/2}\|f-g\|$, and by Markov's inequality and the identity $\int Sh\, d\mu = \int h\, d\mu$, the second term is bounded by $\frac{1}{\varepsilon/2}\|g-f\|$ as well. Thus for any $\varepsilon > 0$ and $g \in V$ we have found that

$$\mu\left(\limsup_{N\to\infty} |S_N f - Sf| > \varepsilon\right) \leq \frac{4}{\varepsilon}\|f-g\|$$

For each fixed $\varepsilon > 0$, the right hand side can be made arbitrarily close to 0, hence $\limsup |S_N f - Sf| = 0$ a.e. which is just $S_N f \to Sf = \mathbb{E}(f|\mathcal{I})$, as claimed.

We now return to the maximal inequality which will be proved by reducing it to a purely combinatorial statement about functions on the integers. Given a function $\widehat{f} : \mathbb{N} \to [0,\infty)$ and a set $\emptyset \neq I \subseteq \mathbb{N}$, the average of $\widehat{f}$ over I is denoted

$$S_I \widehat{f} = \frac{1}{|I|}\sum_{i\in I} \widehat{f}(i)$$

In the following discussion we write $[i,j]$ also for integer segments, i.e. $[i,j] \cap \mathbb{Z}$.

Proposition 4.3.4 (Discrete maximal inequality). *Let $\widehat{f} : \mathbb{N} \to [0,\infty)$. Let $J \subseteq I \subseteq \mathbb{N}$ be finite intervals, and for each $j \in J$ let $I_j \subseteq I$ be a sub-interval of I whose left endpoint is j. Suppose that $S_{I_j}\widehat{f} > t$ for all $j \in J$. Then*

$$S_I \widehat{f} > t \cdot \frac{|J|}{|I|}$$

Proof. Suppose first that the intervals $\{I_j\}$ are disjoint. Then together with $U = I \setminus \bigcup I_j$ they form a partition of I, and by splitting the average $S_I\widehat{f}$ according to this partition, we have the identity

$$S_I\widehat{f} = \frac{|U|}{|I|}S_U\widehat{f} + \sum \frac{|I_j|}{|I|}S_{I_j}\widehat{f}$$

Since $\widehat{f} \geq 0$ also $S_U\widehat{f} \geq 0$, and so

$$S_I\widehat{f} \geq \sum \frac{|I_j|}{|I|}S_{I_j}\widehat{f} \geq \frac{1}{|I|}\sum t|I_j| \geq t\frac{|\bigcup I_j|}{|I|}$$

Now, $\{I_j\}_{j\in J}$ is not a disjoint family, but the above applies to every disjoint sub-collection of it. Therefor we will be done if we can extract from $\{I_j\}_{j\in J}$ a disjoint sub-collection whose union is of size at least $|J|$. This is the content of the next lemma. □

Lemma 4.3.5 (Covering lemma)**.** *Let $I, J, \{I_j\}_{j\in J}$ be intervals as above. Then there is a subset $J_0 \subseteq J$ such that (a) $J \subseteq \bigcup_{i\in J_0} I_j$ and (b) the collection of intervals $\{J_i\}_{i\in J_0}$ is pairwise disjoint.*

Proof. Let $I_j = [j, j+N(j)-1]$. We define $J_0 = \{j_k\}$ by induction using a greedy procedure. Let $j_1 = \min J$ be the leftmost point. Assuming we have defined $j_1 < \ldots < j_k$ such that $I_{j_1}, \ldots, I_{j_k}$ are pairwise disjoint and cover $J\cap[0, j_k+N(j_k)-1]$. As long as this is not all of J, define

$$j_{k+1} = \min\{I \setminus [0, j_k + N(j_k) - 1]\}$$

It is clear that the extended collection satisfies the same conditions, so we can continue until we have covered all of J. □

We return now to the dynamical setting. Each $x \in X$ defines a function $\widehat{f} = \widehat{f}_x : \mathbb{N} \to [0, \infty)$ by evaluating f along the orbit:

$$\widehat{f}(i) = f(T^i x)$$

Let

$$A = \{\sup_N S_N f > t\}$$

and note that if $T^j x \in A$ then there is an $N = N(j)$ such that $S_N f(T^j x) > t$. Writing

$$I_j = [j, j + N(j) - 1]$$

this is the same as

$$S_{I_j}\widehat{f} > t$$

Fixing a large M (we eventually take $M \to \infty$), consider the interval $I = [0, M-1]$ and the collection $\{I_j\}_{j\in J}$, where

$$J = J_x = \{0 \leq j \leq M-1 \,:\, T^j x \in A \text{ and } I_j \subseteq [0, M-1]\}$$

The proposition then gives

$$S_{[0,M-1]}\widehat{f} > t \cdot \frac{|J|}{M}$$

In order to estimate the size of J we will restrict to intervals of some bounded length $R > 0$ (which we eventually will send to infinity). Let

$$A_R = \{\sup_{0 \le N \le R} S_N f > t\}$$

Then

$$J \supseteq \{0 \le j \le M - R - 1 \,:\, T^j x \in A_R\}$$

and if we write $h = 1_{A_R}$, then we have

$$\begin{aligned} |J| &\ge \sum_{j=0}^{M-R-1} \widehat{h}(j) \\ &= (M - R - 1) S_{[0,M-R-1]} \widehat{h} \end{aligned}$$

With this notation now in place,the above becomes

$$S_{[0,M-1]}\widehat{f}_x > t \cdot \frac{M-R-1}{M} \cdot S_{[0,M-R-1]}\widehat{h}_x \tag{4.2}$$

and notice that the average on the right-hand side is just frequency of visits to A_R up to time M.

We now apply a general principle called the *transference principle*, which relates the integral $\int g\,d\mu$ of a function $g : X \to \mathbb{R}$ its discrete averages $S_I \widehat{g}$ along orbits: using $\int g = \int T^n g$, we have

$$\begin{aligned} \int g\,d\mu &= \frac{1}{M} \sum_{m=0}^{M-1} \int T^m g\,d\mu \\ &= \int \left(\frac{1}{M} \sum_{m=0}^{M-1} T^m g \right) d\mu \\ &= \int S_{[0,M-1]}\widehat{g}_x\,d\mu(x) \end{aligned}$$

Applying this to f and using 4.2, we obtain

$$\begin{aligned} \int f\,d\mu &= S_{[0,M-1]}\widehat{f}_x \\ &> t \cdot \frac{M-R-1}{M} \cdot \int h\,d\mu \\ &= t \cdot (1 - \frac{R-1}{M}) \cdot \int 1_{A_R}\,d\mu \\ &= t \cdot (1 - \frac{R-1}{M}) \cdot \mu(A_R) \end{aligned}$$

Letting $M \to \infty$, this is

$$\int f \, d\mu > t \cdot \mu(A_R)$$

Finally, letting $R \to \infty$ and noting that $\mu(A_R) \to \mu(A)$, we conclude that $\int f \, d\mu > t \cdot \mu(A)$, which is what was claimed.

Example 4.3.6. Let $(\xi_n)_{n=1}^{\infty}$ be an independent identically distributed sequence of random variables represented by a product measure on $(X, \mathcal{B}, \mu) = (\Omega, \mathcal{F}, P)^{\mathbb{N}}$, with $\xi_n(\omega) = \xi(\omega_n)$ for some $\xi \in L^1(\Omega, \mathcal{F}, P)$. Let $\sigma : X \to X$ be the shift, which preserves μ and is ergodic, and $\xi_n = \xi_0(\sigma^n)$. Since the shift acts ergodically on product measures, the ergodic theorem implies

$$\frac{1}{N} \sum_{n=0}^{N-1} \xi_n = \frac{1}{N} \sum_{n=0}^{N-1} \sigma^n \xi_0 \to \mathbb{E}(\xi_0 | \mathcal{I}) = \mathbb{E}\xi_0 \qquad \text{a.e.}$$

Thus the ergodic theorem generalizes the law of large numbers. However it is a very broad generalization: it holds for any stationary process $(\xi_n)_{n=1}^{\infty}$ without any independence assumption, as long as the process is ergodic.

When T is invertible it is also natural to consider the two-sided averages $\overline{S}_N = \frac{1}{2N+1} \sum_{n=-N}^{N} T^n f$. Up to an extra term $\frac{1}{2N+1} f$, this is just $\frac{1}{2} S_N(T, f) + \frac{1}{2} S_N T^{-1}, f)$, where we write $S_N(T, f)$ to emphasize which map is being used. Since both of these converge in L^1 and a.e. to the same function $\mathbb{E}(f|\mathcal{I})$, the same is true for $\overline{S}_N f$.

4.4 Generic points

The ergodic theorem is an a.e. statement relative to a given L^1 function, and, anyway, L^1 functions are only a.e. Therefore it is not clear how to interpret the statement that the orbit of an individual point distributes well in the space. There is an exception: When the space is a compact metric space, one can use the continuous functions as test functions to define a more robust notion.

Definition 4.4.1. Let (X, T) be a topological dynamical system. A point $x \in X$ is *generic* for a Borel measure $\mu \in \mathcal{P}(X)$ if it satisfies the conclusion of the ergodic theorem for every continuous function, i.e.

$$\frac{1}{N} \sum_{n=0}^{N-1} T^n f(x) \to \int f \, d\mu \qquad \text{for all } f \in C(X) \tag{4.3}$$

We have already seen that any measure μ that satisfies the above is T-invariant.

Lemma 4.4.2. *Let $\mathcal{F} \subseteq C(X)$ be a countable $\|\cdot\|_\infty$-dense set. If 4.3 holds for every $f \in \mathcal{F}$ then x is generic for μ.*

Proof. By a familiar calculation, given $f \in C(X)$ and $g \in \mathcal{F}$,

$$\begin{aligned}
\limsup_{N\to\infty} |S_N f(x) - \int f\,d\mu| &\leq \limsup_{N\to\infty} |S_N f(x) - S_N g(X)| + \limsup_{N\to\infty} |S_N g(x) - \int f\,d\mu| \\
&\leq \limsup_{N\to\infty} S_N |f-g|(x) + \limsup_{N\to\infty} |\int g\,d\mu - \int f\,d\mu| \\
&\leq 2\,\|f-g\|_\infty
\end{aligned}$$

since g can be made arbitrarily close to f we are done. □

Proposition 4.4.3. *If μ is T-invariant with ergodic decomposition $\mu = \int \mu_x\,d\mu(x)$. Then μ-a.e. x is generic for μ_x.*

Proof. Since $\mu = \int \mu_x\,d\mu(x)$, it suffices to show that for μ-a.e. x, for μ_x-a.e. y, y is generic for μ_x. Thus we may assume that μ is ergodic and show that a.e. point is generic for it. To do this, fix a $\|\cdot\|_\infty$-dense, countable set $\mathcal{F} \subseteq C(X)$. By the ergodic theorem, $S_N f(x) \to \int f$ a.e., for every $f \in \mathcal{F}$, so since $\mathcal{F}$ is countable there is a set of measure one on which this holds simultaneously for all $f \in \mathcal{F}$. The previous lemma implies that each of these points is generic for μ. □

This allows us to give a new interpretation of the ergodic decomposition when $T : X \to X$ is a continuous map of a compact metric space. For a given ergodic measure μ, let $\mathcal{G}_\mu$ denote the set of generic points for μ. Since a measure is characterized by its integral against continuous functions, if $\mu \neq \nu$ then $\mathcal{G}_\mu \cap \mathcal{G}_\nu = \emptyset$. Finally, it is not hard to see that $\mathcal{G}_\mu$ is measurable and $\mu(\mathcal{G}_\mu) = 1$ by the proposition above. Thus we may regard $\mathcal{G}_\mu$ as the ergodic component of μ. One can also show that $\mathcal{G} = \bigcup \mathcal{G}_\mu$, the set of points that are generic for ergodic measures, is measurable, because these are just the points such that ergodic averages exist against every continuous function, or equivalently every function in a dense countable subset of $C(X)$. Now, for any invariant measure ν with ergodic decomposition $\nu = \int \nu_x\,d\nu(x)$,

$$\nu(\mathcal{G}) = \int \nu_x(\mathcal{G})\,d\nu(x) = 1$$

because ν_x are a.s. ergodic and $\mathcal{G}_{\nu_x} \subseteq \mathcal{G}$. Thus on a set of full ν-measure sets $\mathcal{G}_\mu$ give a partition that coincides with the ergodic decomposition. Note, however, that this partition *does not depend on* ν (in the ergodic decomposition theorem it is not a-priori clear that such a decomposition can be achieved).

Example 4.4.4. Let $X = \{0,1\}^{\mathbb{N}}$ and let $\mu_0 = \delta_{000\ldots}$ and $\mu_1 = \delta_{111\ldots}$. These are ergodic measures for the shift σ. Now let $x \in X$ be the point such that $x_n = 0$ for $k^2 \leq n < (k+1)^2$ if k is even, and $x_n = 1$ for $k^2 \leq n < (k+1)^2$ if k is odd. Thus

$$x = 1110000011111110000000000111\ldots$$

We claim that x is generic for the non-ergodic measure $\mu = \frac{1}{2}\mu_0 + \frac{1}{2}\mu_1$. It suffices to prove that for any ℓ,

$$\begin{aligned} \frac{1}{N}\sum_{n=0}^{N-1} 1_{0^\ell}(T^n x) &\to \frac{1}{2} \\ \frac{1}{N}\sum_{n=0}^{N-1} 1_{1^\ell}(T^n x) &\to \frac{1}{2} \end{aligned}$$

where $0^\ell, 1^\ell$ are the sets of points beginning with ℓ consecutive 0s and ℓ consecutive 1s, respectively. The proofs are similar so we show this for 0^ℓ. Notice that $1_{0^\ell}(T^n x) = 1$ if $k^2 \leq n < (k+1)^2 - \ell$ and k is even, and $1_{0^\ell}(T^n x) = 0$ otherwise. Now, each N satisfies $k^2 \leq N < (k+2)^2$ for some even k. Then

$$\sum_{n=0}^{N-1} 1_{0^\ell}(T^n x) = \sum_{j=1}^{k/2}((2j+1)^2 - \ell) - (2j)^2) = \sum_{j=1}^{k/2}(4j+1-\ell) = (\frac{1}{2}k^2 + O(k))$$

Also $N - k^2 \leq (k+1)^2 - k^2 = O(k)$. Therefore $S_N 1_{0^\ell}(x) \to \frac{1}{2}$ as claimed.

Example 4.4.5. With (X, σ) as in the previous example, let $y_n = 0$ if $2^k \leq n < 2^{k+1}$ for k even and $y_n = 1$ otherwise. Then one can show that x is not generic for any measure, ergodic or not.

Our original motivation for considering ergodic averages was to study the frequency of visits of an orbit to a set. Usually 1_A is not continuous even when A is topologically a nice set (e.g. open or closed), so generic points do not have to behave well with respect to visit frequencies. The following shows that this can be overcome with slightly stronger assumption on A and x.

Lemma 4.4.6. *If x is generic for μ, and if U is open and C is closed, then*

$$\begin{aligned} \liminf \frac{1}{N}\sum_{n=0}^{N-1} 1_U(T^n x) &\geq \mu(U) \\ \limsup \frac{1}{N}\sum_{n=0}^{N-1} 1_C(T^n x) &\leq \mu(C) \end{aligned}$$

Proof. Let $f_k \in C(X)$ with $f_k \nearrow 1_U$ (e.g. $f_n(y) = 1 - e^{-kd(y,U^c)}$). Then $1_U \geq f_n$ and so

$$\liminf \frac{1}{N}\sum_{n=0}^{N-1} 1_U(T^n x) \geq \lim \frac{1}{N}\sum_{n=0}^{N-1} f_k(T^n x) = \int f_k d\mu \to \mu(U)$$

The other inequality is proves similarly using $g_n \searrow 1_C$. $\square$

Proposition 4.4.7. *If x is generic for μ, $A \subseteq X$ and $\mu(\partial A) = 0$ then $\frac{1}{N}\sum_{n=0}^{N-1} 1_A(T^n x) \to \mu(A)$.*

Proof. Let $U = interior(A)$ and $C = \overline{A}$, so $1_U \leq 1_A \leq 1_C$. By the lemma,

$$\liminf S_N 1_A \geq \liminf S_N 1_U \geq \mu(U)$$

and

$$\limsup S_N 1_A \leq \limsup S_N 1_C \leq \mu(C)$$

But by our assumption, $\mu(U) = \mu(C) = \mu(A)$, and we find that

$$\mu(A) = \liminf S_N 1_A \leq \limsup S_N 1_A \leq \mu(A)$$

So all are equalities, and $S_N 1_A \to \mu(A)$. ☐

4.5 Unique ergodicity and circle rotations

When can the ergodic theorem be strengthened from a.e. point to every point? Once again the question does not make sense for L^1 functions, since these are only defined a.e., but it makes sense for continuous functions.

Definition 4.5.1. A topological system (X, T) is uniquely ergodic if there is only one invariant probability measure, which in this case is denoted μ_X.

Proposition 4.5.2. *Let (X, T) be a topological system and $\mu \in \mathcal{P}_T(X)$. The following are equivalent.*

1. *Every point is generic for μ.*
2. *$S_N f \to \int f \, d\mu$ uniformly, for every $f \in C(X)$.*
3. *(X, T) is uniquely ergodic and μ is its invariant measure.*

Proof. (1) implies (3): If $\nu \neq \mu$ were another invariant measure there would be points that are generic for it, contrary to (1).

(3) implies (2): Suppose (2) fails, so there is an $f \in C(X)$ such that $\left\|S_N f - \int f d\mu\right\|_\infty \not\to 0$. Then there is some sequence $x_k \in X$ and integers $N_k \to \infty$ such that $S_{N_k} f(x_k) \to c \neq \int f d\mu$. Let ν be an accumulation point of $\frac{1}{N_k} \sum_{n=1}^{N_k} \delta_{T^n x_k}$. This is a T-invariant measure and $\int f d\nu = c$ so $\nu \neq \mu$, contradicting (3).

(2) implies (1) is immediate. ☐

Proposition 4.5.3. *Let $X = \mathbb{R}/\mathbb{Z}$ and $\alpha \notin \mathbb{Q}$. The map $T_\alpha x = x + \alpha$ on X is uniquely ergodic with invariant measure $\mu = Lebesgue$.*

We give two proofs.

Proof number 1. We know that μ is ergodic for T_α so a.e. x is generic. Fix one such x. Let $y \in X$ be any other point. then there is a $\beta \in \mathbb{R}$ such that $y = T_\beta x$.

For any function $f \in C(X)$,

$$\begin{aligned}
\frac{1}{N}\sum_{n=0}^{N-1} T_\alpha^n f(y) &= \frac{1}{N}\sum_{n=0}^{N-1} f(y+\alpha n) \\
&= \frac{1}{N}\sum_{n=0}^{N-1} f(x+\alpha_n+\beta) \\
&= \frac{1}{N}\sum_{n=0}^{N-1} (T_\beta f)(T_\alpha^n x) \\
&\to \int T_\beta f\, d\mu = \int f\, d\mu
\end{aligned}$$

Therefore every point is generic for μ and T_α is uniquely ergodic. □

Our second proof is based on a more direct calculation that does not rely on the ergodic theorem.

Definition 4.5.4. A sequence (x_k) in a compact metric space X equidistributes for a measure μ if $\frac{1}{N}\sum_{n=1}^{N} \delta_{x_n} \to \mu$ weak-*.

Lemma 4.5.5 (Weyl's equidistribution criterion)**.** *A sequence $(x_k) \subseteq \mathbb{R}/\mathbb{Z}$ equidistributes for Lebesgue measure μ if and only if for every m,*

$$\frac{1}{N}\sum_{n=0}^{N-1} e^{2\pi i m x_n} \to \begin{cases} 0 & m=0 \\ 1 & m \neq 0 \end{cases}$$

Proof. Let $\chi_m(t) = e^{s\pi i m t}$. The linear span of $\{\chi_m\}_{m\in\mathbb{Z}}$ is dense in $C(\mathbb{R}/\mathbb{Z})$ by Fourier analysis so equidistribution of (x_k) is equivalent to $S_N\chi_m(x) \to \int \chi_m d\mu$ for every m. This is what the lemma says. □

Proof number 2. Fix $t \in \mathbb{R}/\mathbb{Z}$ and $x_k = t + \alpha k$. For $m = 0$ the limit in Weyl's criterion is automatic so we only need to check $m \neq 0$. Then

$$\frac{1}{N}\sum_{n=0}^{N-1} e^{2\pi i m x_n} = \frac{1}{N} e^{2\pi i m t} \cdot \sum_{n=0}^{N-1} (e^{2\pi i m \alpha})^n = \frac{1}{N} e^{2\pi i t} \cdot \frac{e^{2\pi i m \alpha N} - 1}{e^{2\pi i m \alpha} - 1} = 0$$

(note that $\alpha \notin \mathbb{Q}$ ensures that the denominator is not 0, otherwise the summation formula is invalid). □

Corollary 4.5.6. *For any open or closed set $A \subseteq \mathbb{R}/\mathbb{Z}$, for every $x \in \mathbb{R}/\mathbb{Z}$, $S_N 1_A(x) \to Leb(A)$.*

Proof. The boundary of an open or closed is countable and hence of Lebesgue measure 0. □

Example 4.5.7 (Benford's law)**.** Many samples of numbers collected in the real world exhibit the interesting feature that the most significant digit is not uniformly distributed. Rather, 1 is the most common digit, with frequency approximately 0.30; the frequency of 2 is about 0.18; the frequency of 3 is about 0.13; etc. More precisely, the frequency of the digit k is approximately $\log_{10}(1+\frac{1}{d})$.

We will show that a similar distribution of most significant digits holds for powers of b whenever b is not a rational power of 10. The main observation is that the most significant base-10 digit of $x \in [1,\infty)$ is determined by $y = \log_{10} x \bmod 1$, and is equal to k if $y \in I_k = [\log_{10} k, \log_{10}(k+1))$. Therefore, the asymptotic frequency of k being the most significant digits of b^n is

$$\begin{aligned}
\lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} 1_{I_k}(\log_{10} b^n) &= \lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} 1_{I_k}(n\frac{\ln b}{\ln 10}) \\
&= Leb(I_k) \\
&= Leb[\log_{10} k, \log_{10}(k+1)] \\
&= \log_{10}(1+\frac{1}{k})
\end{aligned}$$

since this is just the frequency of visits of the orbit of 0 to $[\log_{10} k, \log_{10}(k+1)]$ under the map $t \mapsto t + \ln b/\ln 10 \bmod 1$, and $\ln b/\ln 10 \notin \mathbb{Q}$ by assumption (it would be rational if and only if b is a rational power of 10).

4.6 Sub-additive ergodic theorem

Theorem 4.6.1 (Subadditive ergodic theorem)**.** *Let $(X,\mathcal{B},\mu,T)$ be an ergodic measure-preserving system. Suppose that $f_n \in L^1(\mu)$ satisfy the subadditivity relation*

$$f_{m+n}(x) \le f_m(x) + f_n(T^m x)$$

and are uniformly bounded above, i.e. $f_n \le L$ for some L. Then $\lim_{n\to\infty} \frac{1}{n} f_n(x)$ exists a.e. and is equal to the constant $\lim_{n\to\infty} \frac{1}{n}\int f_n$.

Before giving the proof we point out two examples. First, if $f_n = \sum_{k=0}^{n-1} T^k g$ then f_n satisfies the hypothesis, so this is a generalization of the usual ergodic theorem (for ergodic T).

For a more interesting example, let $A_n = A(T^n x)$ be a stationary sequence of $d\times d$ matrices (for example, if the entries are i.i.d.). Let $f_n = \log \|A_1 \cdot \ldots \cdot A_n\|$ satisfies the hypothesis. Thus, the subadditive ergodic theorem implies that random matrix products have a Lyapunov exponent – their norm growth is asymptotically exponential.

Proof. Let us first make a simple observation. Suppose that $\{1,\ldots,N\}$ is partitioned into intervals $\{[a_i,b_i)\}_{i\in I}$. Then subadditivity implies

$$f_N(x) \le \sum_{i\in I} f_{b_i-a_i}(T^{a_i}x)$$

Let
$$a = \liminf \frac{1}{n} f_n$$
We claim that a is invariant. Indeed,
$$\frac{1}{n} f_n(Tx) \geq \frac{1}{n}(f_{n+1}(x) - f_1(x))$$
From this it follows that $a(Tx) \geq a(x)$ so by ergodicity a is constant.

Fix $\varepsilon > 0$. Since $\liminf \frac{1}{n} f_n = a$ there is an N such that the set
$$A = \{x \,:\, \frac{1}{n} f_n(x) < a + \varepsilon \text{ for some } 0 \leq n \leq N\}$$
satisfies $\mu(A) > 1 - \varepsilon$.

Now fix a typical point x. By the ergodic theorem, for every large enough M,
$$\frac{1}{M} \sum_{n=0}^{M-1} 1_A(T^n x) > 1 - \varepsilon$$
Fix such an M and let
$$I_0 = \{0 \leq n \leq M - N \,:\, T^n x \in A\}$$
For $i \in I_0$ there is a $0 \leq k_i \leq N$ such that $\frac{1}{k} f_{k_i}(T^i x) < a + \varepsilon$. Let $U_i = [i, n + k_n)$. Applying the covering lemma, Lemma 4.3.5, there is a subset $I_1 \subseteq I_0$ such that $\{U_i\}_{i \in I_1}$ are pairwise disjoint and $|\bigcup_{i \in I_1} U_i| \geq |I_0| > (1-\varepsilon)M$. By construction also $\bigcup_{i \in I_1} U_i \subseteq [0, M)$.

Choose an enumeration $\{U_i\}_{i \in I_2}$ of the complementary intervals in $[0, M) \setminus \bigcup_{i \in I_i} U_i$, so that $\{U_i\}_{i \in I_1 \cup I_2}$ is a partition of $[0, M)$. Writing $U_i = [a_i, b_i)$ and using the comment above, we find that
$$\begin{aligned}
\frac{1}{M} f_M(x) &\leq \frac{1}{M}\left(\sum_{i \in I_1} f_{b_i - a_i}(T^{a_i} x) + \sum_{i \in I_2} f_{b_i - a_i}(T^{a_i} x)\right) \\
&\leq \frac{\sum_{n \in I_1} |U_i|}{M}(a + \varepsilon) + \frac{\sum_{n \in I_2} |U_i|}{M} \|f\|_\infty \\
&\leq (a + \varepsilon) + \varepsilon \|f\|_\infty
\end{aligned}$$
Since this holds for all large enough M we conclude that $\limsup \frac{1}{M} f_M \leq a = \liminf \frac{1}{n} f_n$ so the limit exists and is equal to a.

It remains to identify $a = \lim \frac{1}{n} \int f_n$. First note that
$$\int f_{m+n} \leq \int f_m d\mu + \int f_n \circ T^m d\mu = \int f_m d\mu + \int f_n d\mu$$
so $a_n = \int f_n d\mu$ is subadditive, hence the limit $a' = \lim \frac{1}{n} a_n$ exists. By Fatou's lemma (since $f_n \leq L$ we can apply it to $-f_n$) we get
$$a = \int \limsup \frac{1}{n} f_n d\mu \geq \limsup \int \frac{1}{n} f_n d\mu = \lim \frac{1}{n} a_n = a'$$

Suppose the inequality were strict, $a' < a - \varepsilon$ for some $\varepsilon > 0$ and let n be such that $a_n < a - \varepsilon$. Note that for every $0 \leq p \leq n-1$ we have the identity

$$f_N(x) \leq f_p(x) + \sum_{k=0}^{[N/n]-1} f_n(T^{kn+p}x) + f_{N-p-n([N/n]-1))}(T^{p+n([N/n]-1))}x)$$

Averaging this over $0 \leq p < n$, we have

$$\frac{1}{N} f_N \leq S_N(\frac{1}{n} f_n) + O(\frac{n}{N})$$

This by the ergodic theorem,

$$\lim_{N\to\infty} \frac{1}{N} f_N \leq \lim_{N\to\infty} S_N(\frac{1}{n} f_n) = \int \frac{1}{n} f_n < a - \varepsilon$$

which is a contradiction to the definition of a. □

4.6.1 Group actions

Let G be a countable group. A measure preserving action of G on a measure space $(X, \mathcal{B}, \mu)$ is, first of all, an action, that is a map $G \times X \to X$, $(g, x) \mapsto gx$, such that $g(hx) = (gh)(x)$ for all $g, h \in G$ and $x \in X$. In addition, for each $g \in G$ the map $T_g : x \mapsto gx$ must be measurable and measure-preserving. It is convenient to denote the action by $\{T_g\}_{g\in G}$.

An invariant set for the action is a set $A \in \mathcal{B}$ such tat $T_g A = A$ for all $g \in G$. If every such set satisfies $\mu(A) = 0$ or $\mu(X \setminus A) = 0$, then the action is ergodic. There is an ergodic decomposition theorem for such actions, but for simplicity (and without loss of generality) we will assume that the action is ergodic.

For a function $f : X \to \mathbb{R}$ the function $T_g f = f \circ T_{g^{-1}} : X \to \mathbb{R}$ has the same regularity, and $\{T_g\}_{g\in G}$ gives an isometric action on L^p for all $1 \leq p \leq \infty$. Given a finite set $E \subseteq G$ let $S_E f$ be the functions defined by

$$S_E f(x) = \sum_{g\in E} f(T_g x)$$

As before, this is a contraction in L^p. We say that a sequence $E_n \subseteq G$ of finite sets *satisfies the ergodic theorem along* $\{E_n\}$ if $S_{E_n} f \to \int f$, in a suitable sense (e.g. in L^2 or a.e.) for every ergodic action and every suitable f.

Definition 4.6.2. A group G is *amenable* if there is a sequence of sets $E_n \subseteq G$ such that for every $g \in G$,

$$\frac{|E_n g \Delta E_n|}{|E_n|} \to 0$$

Such a sequence $\{E_n\}$ is called a *Følner sequence*.

For example, $\mathbb{Z}^d$ is a amenable because $E_n = [-n, n]^d \cap \mathbb{Z}^d$ satisfies

$$|(E_n + u) \cap E_n| = |E_{n-\|u\|_\infty}| = |E_n| + o(1)$$

The class of amenable groups is closed under taking subgroups and countable increasing unions, and if G and $N \lhd G$ are amenable so is G/N. Groups of sub-exponential growth are amenable; the free group is not amenable, but there are amenable groups of exponential growth.

Theorem 4.6.3. *If $\{E_n\}$ is a Følner sequence in an amenable group G then the ergodic theorem holds along $\{E_n\}$ in the L^2 sense (the mean ergodic theorem).*

Proof. Let
$$V_0 = \text{span}\{f - T_g f \,:\, f \in L^2\,,\, g \in G\}$$
One can show exactly as before that $V_0^\perp$ consists of the invariant functions (in this case, the constant functions, because we are assuming the action is ergodic). Then one must only show that $S_{E_n}(f - T_g f) \to 0$ for $f \in L^2$. But this is immediate from the Følner property, since

$$S_{E_n} f - S_{E_n} T_g f = S_{E_n \setminus E_n r^{-1}} f$$

and therefore

$$\left\| \frac{1}{|E_n|} S_{E_n}(f - T_g f) \right\|_2 \leq \frac{1}{|E_n|} |E_n \setminus E_n g^{-1}| \cdot \|f\|_2 \leq \frac{|E_n \Delta E_n g^{-1}|}{|E_n|} \|f\|_2 \to 0$$

This proves the mean ergodic theorem. □

The proof of the pointwise ergodic theorem for amenable groups is more delicate and does not hold for every Følner sequence. However, one can reduce it as before to a maximal inequality. What one then needs is an analog of the discrete maximal inequality, which now concerns functions $\widehat{f} : G \to [0, \infty)$, and requires an analog of the covering Lemma 4.3.5. Such a result is known under a stronger assumption on $\{E_n\}$, namely assuming that $|\bigcup_{k<n} E_k^{-1} E_n| \leq C|E_n|$ for some constant C and all n. Every Følner sequence has a subsequence that satisfies this, and so every amenable group has a sequence along which the pointwise ergodic theorem holds a.e. and in L^1.

Outside of amenable groups one can also find ergodic theorems. The simplest to state is for the free group $\mathbb{F}_s$ on s generates $g_1^{\pm 1}, \ldots, g_s^{\pm 1}$. This is a non-amenable group which can be identified with the set of words in the generators that don't contain any occurrence of uu^{-1}. The group operation is concatenation follows by reduction, that is, repeatedly deleting any pair ss^1. For example the product of words $aba^{-1}c$ and $c^{-1}abb$ is

$$aba^{-1}cc^{-1}abb = aba^{-1}abb = abbb$$

the right hand side is reduced.

Let $E_n \subseteq \mathbb{F}_s$ denote the set of reduced words of length $\leq n$.

Theorem 4.6.4 (Nevo-Stein, Bufetov)**.** *If $\mathbb{F}_s$ acts ergodically by measure preserving transformations on $(X, \mathcal{B}, \mu)$ then for every $S_{E_n} f \to \int f$ for every $f \in L^1(\mu)$.*

There is a major difference between the proof of this result and in the amenable case. Because $|E_n \Delta E_n g^{-1}|/|E_n| \not\to 0$, the there is no trivial reason for the averages of co-boundaries to tend to 0. Consequently there is no natural dense set of functions in L^1 for which convergence holds. In any case, the maximal inequality is not valid either. The proof in non-amenable cases takes completely different approaches (but we will not discuss them here).

4.6.2 Hopf's ergodic theorem

Another generalization is to the case of a measure-preserving transformation T of a measure space $(X, \mathcal{B}, \mu)$ with $\mu(X) = \infty$ (but σ-finite). Ergodicity is defined as before – all invariant sets are of measure 0 or their complement is of measure 0. It is also still true that $T : L^2(\mu) \to L^2(\mu)$ is norm-preserving, and so the mean ergodic theorem holds: $S_N f \to \pi f$ for $f \in L^2$, where π is the projection to the subspace of invariant L^2 functions. Now, however, the only constant function that is integrable is 0, and we find that $S_N f \to 0$ in L^2. In fact this is true in L^1 and a.e. The meaning is, however, the same: if we take a set of finite measure A, this says that the fraction of time an orbit spends in A is the same as the relative size of A compared to Ω; in this case $\mu(A)/\mu(\Omega) = 0$.

Instead of asking about the absolute time spent in A, it is better to consider two sets A, B of positive finite measure. Then an orbit visits both with frequency 0, but one may expect that the frequency of visits to A is $\mu(A)/\mu(B)$-times the frequency of visits to B. This is actually he case:

Theorem 4.6.5 (Hopf)**.** *If T is an ergodic measure-preserving transformation of $(X, \mathcal{B}, \mu)$ with $\mu(X) = \infty$, and if $f, g \in L^1(\mu)$ and $\int g d\mu \neq 0$, then*

$$\frac{\sum_{n=0}^{N-1} T^n f}{\sum_{n=0}^{N-1} T^n g} \xrightarrow[N\to\infty]{} \frac{\int f d\mu}{\int g d\mu} \qquad a.e.$$

Since the right hand side is usually not 0, one cannot expect this to hold in L^1.

Hopf's theorem can also be generalized to group actions, but the situation there is more subtle, and it is known that not all amenable groups have sequences E_n such that $\sum_{E_n} T^g f / \sum_{E_n} T^g h \to \int f / \int h$. See **??**.

Chapter 5

Some categorical constructions

5.1 Isomorphism and factors

Definition 5.1.1. Two measure preserving systems $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ are isomorphic if there are invariant subsets $X_0 \subseteq X$ and $Y_0 \subseteq Y$ of full measure and a bijection $\pi : X_0 \to Y_0$ such that π, π^{-1} are measurable, $\pi\mu = \nu$, and $\pi \circ T = S \circ \pi$. The last condition means that the following diagram commutes:

$$\begin{array}{ccc} X_0 & \xrightarrow{T} & X_0 \\ \downarrow \pi & & \downarrow \pi \\ Y_0 & \xrightarrow{S} & Y_0 \end{array}$$

It is immediate that ergodicity and mixing are isomorphism invariants. Also, π induces an isometry $L^2(\mu) \to L^2(\nu)$ in the usual manner and the induced maps of T, S on these spaces commute with π, so the induced maps T, S are unitarily equivalent in the Hilbert-space sense. The same is true for the associated L^p spaces.

Example 5.1.2. Let $\alpha \in \mathbb{R}$ and $X = \mathbb{R}/\mathbb{Z}$ with Lebesgue measure μ, and $T_\alpha x = x + \alpha \bmod 1$. Then $T_\alpha, T_{-\alpha}$ are isomorphic via the isomorphism $x \to -x$.

Example 5.1.3. Let $X = \{0,1\}^{\mathbb{N}}$ with μ the product measure $\frac{1}{2}, \frac{1}{2}$ and the shift T, and $Y = [0,1]$ with $\nu =$ Lebesgue measure and $Sx = 2x \bmod 1$. Let $\pi : X \to Y$ be the map $\pi(x) = \sum_{n=1}^{\infty} x_n 2^{-n}$, or $\pi(x) = 0.x_1x_2x_3\ldots$ in binary notation. Then it is well known that $\pi\mu = \nu$, and we have

$$S(\pi x) = S(0.x_1x_2\ldots) = 0.x_2x_3\ldots = \pi(Tx)$$

Thus π is a factor map between the corresponding systems. Furthermore it is an isomoprhism, since if we take $X_0 \subseteq X$ to be all eventually-periodic sequences and $Y_0 = Y \setminus \mathbb{Q}$. These are invariant sets; $\mu(X_0) = 1$, since there are countably many eventually periodic sequences and each has measure 0; and $\nu(Y_0) = 1$. Finally $\pi : X_0 \to Y_0$ is 1-1 and onto, since there is an inverse given by the binary expansion, which is measurable. This proves that the two systems are isomorphic.

Example 5.1.4. An irrational rotation is not isomorphic to a shift space with a product measure. This can be seen in many ways, one of which is the following. Note that there is a sequence $n_k \to \infty$ such that $T_\alpha^{n_k} 0 \to 0$; this follows from the fact that 0 equidistributes for Lebesgue measure, so its orbit must return arbitrarily close to x. Since $x = T_x 0$, we find that

$$T^{n_k} x = T^{n_k} T_x 0 = T_x T^{n_k} 0 \to T_x 0 = x$$

so $T^{n_k} x \to x$ for all $x \in \mathbb{R}/\mathbb{Z}$. It follows from dominated convergence that for every $f \in L^2(\mu) \cap L^\infty(\mu)$ we have $T^{n_k} f \to f$ in L^2, hence

$$\int T^{n_k} f \cdot f \to \int (f^2) d\mu$$

On the other hand if $(A^{\mathbb{Z}}, \mathcal{C}, \nu, S)$ is a shifts space with a product measure then we have already seen that it is mixing, hence for every $f \in L^2(\nu)$ we have

$$\int S^{n_k} f \cdot f \to (\int f)^2 d\nu$$

By Cauchy-Schwartz, we generally have $(\int f)^2 \neq \int (f^2)$ so the operators S, T cannot be unitarily equivalent.

Definition 5.1.5. A measure preserving system $(Y, \mathcal{C}, \nu, S)$ is a factor of a measure preserving system $(X, \mathcal{B}, \nu, T)$ if there are invariant subsets $X_0 \subseteq X$, $Y_0 \subseteq Y$ of full measure and a measurable map $\pi : X_0 \to Y_0$ such that $\pi\mu = \nu$ and $\pi \circ T = S \circ \pi$.

This is the same as an isomorphism, but without requiring π to be 1-1 or onto. Note that $\pi\mu = \nu$ means that π is automatically "onto" in the measure sense: if $A \subseteq Y_0$ and $\pi^{-1}(A) = \emptyset$ then $\nu(A) = 0$.

Remark 5.1.6. When X, Y are standard Borel spaces (i.e. as measurable spaces they are isomorphic to complete separable metric spaces with the Borel σ-algebra), one can always assume that $\pi : X_0 \to Y_0$ is onto (even though the image of a measurable set is not in general measurable, in the standard setting one can find a measurable subset of the image that has full measure, and restrict to it).

Example 5.1.7. Let T_α be rotation by α. Let $k \in \mathbb{Z} \setminus 0$. Then $\pi : x \mapsto kx \bmod 1$ maps Lebesgue measure to Lebesgue measure on $\mathbb{R}/\mathbb{Z}$ and

$$\pi T_\alpha x = k(x + \alpha) = T_{k\alpha} \pi x$$

Thus $T_{k\alpha}$ is a factor of T_α. Note that unless $k = \pm 1$ this is not an isomorphism, since $|\pi^{-1} y| = k$ for all y.

Example 5.1.8. Let $A = \{1, \ldots, n\}$, $p = (p_1, \ldots, p_n)$ be a non-degenerate probability vector, $X = A^{\mathbb{Z}}$ with the product σ-algebra and product measure $\mu = p^{\mathbb{Z}}$. Let B be a set and $\pi : A \to B$ any map, and let $q_b = p_{\pi^{-1}b}$ the push-forward probability vector. Let $Y = B^{\mathbb{Z}}$ and $\nu = q^{\mathbb{Z}}$. Finally let S be the shift (on both spaces) and extend π to a map $X \to Y$ pointwise:

$$\pi(\ldots, x_{-1}, x_0, x_1, \ldots) = (\ldots \pi(x_{-1}), \pi(X_0), \pi(x_1), \ldots)$$

Then by considering cylinder sets it is easy to show that this $\pi\mu = \nu$, and clearly $S\pi - \pi S$. This Y is a factor of X.

Proposition 5.1.9. *Let $(X, \mathcal{B})$ and $(Y, \mathcal{C})$ be measurable spaces and $T : X \to X$ and $S : Y \to Y$ measurable maps, and $\pi : X \to Y$ $\pi : X \to Y$ be a measurable map such that $\pi T = S\pi$. .*

1. *If μ is an invariant measure for T then $\nu = \pi\mu$ is an invariant measure for S and π is a factor map between $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$.*

2. *If the spaces are standard Borel spaces and if ν is an invariant measure for S, then there exist invariant measures μ for T such that $\nu = \pi\mu$ and π is a factor map between $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ (but no uniqueness is claimed).*

Proof. The first part is an exercise.

The second part is less trivial. We give a proof for the case that X, Y are compact metric spaces, $\mathcal{B}, \mathcal{C}$ the Borel σ-algebras, and T, S, π are continuous. In this situation, we first need a non-dynamical version:

Lemma 5.1.10. *There is a measure μ_0 on X such that $\pi\mu = \nu$.*

Proof No. 1 (almost elementary). Start by constructing a sequence ν_n of atomic measures on Y with $\nu_n \to \nu$ weakly, i.e. $\int g\,d\nu_n \to \int g\,d\nu$ for all $g \in C(Y)$. To get such a sequence, given n choose a finite partition $\mathcal{E}_n$ of Y into measurable sets of diameter $< 1/n$ (for instance cover Y by balls B_i of radius $< 1/n$ and set $E_i = B_i \setminus \bigcup_{j<i} B_j$). For each $E \in \mathcal{E}_n$ choose x_E and set $\nu_n = \sum_{E \in \mathcal{E}_n} \nu(E) \cdot \delta_{x_E}$. One may verify that $\nu_n \to \nu$.

Now, each ν_n can be lifted to a probability measure μ_n on X such that $\pi\mu_n = \nu_n$: to see this, if $\nu_n = \sum w_i \cdot \delta_{y_i}$ choose $x_i \in \pi^{-1}(y_i)$ (there may be many choices, choose one), and set $\mu_n = \sum w_i \cdot \delta_{x_i}$.

Since the space of Borel probability measures on X is compact in the weak-* topology, by passing to a subsequence we can assume $\mu_n \to \mu$. Clearly μ is a probability measures; we claim $\pi\mu = \nu$. It is enough to show that $\int g\,d(\pi\mu) = \int g\,d\nu$ for every $g \in C(Y)$. Using the identity $\int g\,d\nu_n = \int g \circ \pi\,d\mu_n$ (which is equivalent to $\nu_n = \pi\mu_n$) we have

$$\int g\,d\nu = \lim \int g\,d\nu_n = \int g \circ \pi\,d\mu_n = \int g \circ \pi\,d\mu = \int g\,d(\pi\mu)$$

as claimed. $\square$

Proof No. 2 (function-analytic). . First a few general remarks. A linear functional μ^* on $C(X)$ is positive if it takes non-negative values on non-negative functions. This property implies boundedness: to see this note that for any $f \in C(X)$ we have $\|f\|_\infty - f \geq 0$, hence by linearity and positivity $\mu^*(\|f\|_\infty) - \mu^*(f) \geq 0$, giving

$$\mu^*(f) \leq \mu^*(\|f\|_\infty) = \|f\|_\infty \cdot \mu^*(1)$$

Similarly, using $f + \|f\|_\infty \geq 0$ we get $\mu^*(f) \geq -\|f\|_\infty$. Combining the two we have $|\mu^*(f)| \leq C\|f\|_\infty$, where $C = \mu^*(1)$.

Since a positive functional μ^* is bounded it corresponds to integration against a regular signed Borel measure μ, and since $\int f\,d\mu = \mu^*(f) \geq 0$ for continuous $f \geq 0$, regularity implies that μ is a positive measure. Hence a linear functional $\mu^* \in C(X)^*$ corresponds to a probability measure if and only if it be positive and $\mu^*(1) = 1$ (this is the normalization condition $\int 1\,d\mu = 1$).

We now begin the proof. Let $\nu^* : C(Y) \to \mathbb{R}$ be bounded positive the linear functional $g \mapsto \int g\,d\nu$. The map $\pi^* : C(Y) \to C(X)$, $g \mapsto g \circ \pi$, embeds $C(Y)$ isometrically as a subspace $V = \pi(C(Y)) < C(X)$, and lifts ν^* to a bounded linear functional $\mu_0^* : V \to \mathbb{R}$ (given by $\mu_0^*(g \circ \pi) = \nu^*(g)$).

Consider the positive cone $P = \{f \in C(X) \,:\, f \geq 0\}$, and let $s \in C(X)^*$ be the functional

$$s(f) = \sup\{0, -f(x) \,:\, x \in X\}$$

It is easy to check that s is a seminorm, that $s|_P \equiv 0$ and that $-\mu_0^*(f) \leq s(f)$ on V. Hence by Hahn-Banach we can extend $-\mu_0^*$ to a functional $-\mu^*$ on $C(X)$ satisfying $-\mu^* \leq s$, which for $f \in P$ implies $\mu^*(f) \geq -s(f) = 0$, so μ^* is positive. By the previous discussion there is a Borel probability measure μ such that $\int f\,d\mu = \mu^*(f)$; for $f = g \circ \pi$ this means that

$$\int g\,d\pi\mu = \int g \circ \pi\,d\mu = \mu^*(g \circ \pi) = \mu_0^*(g \circ \pi) = \nu^*(g) = \int g\,d\nu$$

so μ is the desired measure. □

Now let $\mu_n = \frac{1}{n}\sum_{k=0}^{n-1} T^k\mu_0$ and let μ be a subsequential limit of μ_n. It is easy to check int he usual way that μ is T-invariant. Also,

$$\pi\mu_n = \frac{1}{n}\sum_{k=0}^{n-1}\pi(T^k\mu) = \frac{1}{n}\sum_{k=0}^{n-1}T^k\pi\mu = \frac{1}{n}\sum_{k=0}^{n-1}\nu = \nu$$

Since the space $\mathcal{M}$ of measure on X projecting to ν is weak-* closed, and $\mu_n \in \mathcal{M}$, also their limit point $\mu \in \mathcal{M}$, as claimed. □

Example 5.1.11. Let A, B be finite sets, $A^{\mathbb{Z}}, B^{\mathbb{Z}}$ the product spaces and S the shift. Let $\pi : A^{2n+1} \to B$ be any map and extend π to $A^{\mathbb{Z}} \to B^{\mathbb{Z}}$ by

$$(\pi x)_i = \pi(x_{i-n}, \ldots, x_i, \ldots x_{i+n})$$

This commutes with the shift and so any invariant measure on $A^{\mathbb{Z}}$ projects to an invariant measure on $B^{\mathbb{Z}}$ and vice versa.

For example if $A = B = \{0,1\}$, $n = 1$ and $\pi(a,b,c) = b + c \bmod 1$ then we get a factor map as above and one may verify that the $(\frac{1}{2}, \frac{1}{2})$ product measure is mapped to itself; but the factor map is non-trivial since each sequence has two pre-images.

If $\pi : X \to Y$ is a factor map between measure preserving systems $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ (already restricted to the invariant subsets of full measure). Then

$$\pi^{-1}\mathcal{C} = \{\pi^{-1}C \,:\, C \in \mathcal{C}\}$$

is a sub-σ-algebra of $\mathcal{B}$ and it is invariant since $\pi T = S\pi$. Noe also that π is an isometry between $L^2(\pi^{-1}\mathcal{C}, \mu)$ and $L^2(\mathcal{C}, \nu)$. Thus

There is also a converse:

Proposition 5.1.12. *Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system on a standard Borel space and let $\mathcal{C}' \subseteq \mathcal{B}$ be an invariant, countably generated σ-algebra. Then there is a factor $(Y, \mathcal{C}, \nu, S)$, with $(Y, \mathcal{C})$ a standard Borel space, and factor map π such that $\mathcal{C}' = \pi^{-1}\mathcal{C}$.*

The proof relies on the analogous non-dynamical fact that a countably generated sub-σ-algebra in standard Borel space always arises as the pullback via a measurable map of some other standard Borel space. We shall not go into details.

5.2 Product systems

Another basic construction is to take products:

Definition 5.2.1. Let $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ be measure preserving systems. Let $T \times S : X \times Y \to X \times Y$ denote the map $T \times S(x,y) = (Tx, Sy)$. Then $(X \times Y, \mathcal{B} \times \mathcal{C}, \mu \times \nu, T \times S)$ is called the **product** of X and Y.

Claim 5.2.2. The product of measure preserving systems is measure preserving.

Proof. Let $\theta = \mu \times \nu$ and $R = T \times S$. It is enough to check that $\theta(R^{-1}(A \times B)) = \theta(A \times B)$ for $A \subseteq X$, $B \subseteq Y$, since these sets generate the product algebra and the family invariant sets whose measure is preserved is a σ-algebra. For such A, B,

$$\theta(R^{-1}(A \times B)) = \theta(T^{-1}A \times S^{-1}B) = \mu(T^{-1}A)\nu(T^{-1}B) = \mu(A)\nu(B) = \theta(a \times B)$$

□

Remark 5.2.3. Observe that the coordinate projections π_1 and π_2 are factor maps from the product system to the original systems.

The definition generalizes to products of finitely many or countably many systems.

We turn tot he construction of inverse limits. Let $(X_n, \mathcal{B}_n, \mu_n, T_n)$ be measure preserving systems, and let $\pi_n : X_{n+1} \to X_n$ be factor maps, and suppose that π_n are defined everywhere and onto their image. Let $X_\infty \subseteq \times_{n=1}^{\infty} X_n$ denote the set of sequences $(\ldots, x_n, x_{n-1}, \ldots, x_1)$, with $x_n \in X_n$, such that $\pi_n(x_{n+1}) = x_n$.

5.3 Natural extension

When $(X, \mathcal{B}, \mu, T)$ is an ergodic m.p.s. on a standard Borel space, there is canonical invertible m.p.s. $(\widetilde{X}, \widetilde{\mathcal{B}}, \widetilde{\mu}, \widetilde{T})$ and factor map $\pi : \widetilde{X} \to X$ such that if $(Y, \mathcal{C}, \nu, S)$ is another invertible system and $\tau : Y \to X$ a factor map, τ factors through π, that is there is a factor map $\sigma : Y \to \widetilde{X}$ with $\tau = \pi\sigma$:

$$\begin{tikzcd} Y \arrow[r, "\sigma"] \arrow[dr, "\tau"'] & \widetilde{X} \arrow[d, "\pi"] \\ & X \end{tikzcd}$$

We now construct $\widetilde{X}$. Let $\pi_n : X^{\mathbb{Z}} \to X$ denote the coordinate projections, which are measurable with respect tot he product algebra, and let

$$\widetilde{X} = \{x \in X^{\mathbb{Z}} \,:\, Tx_n = x_{n+1}\}$$

This is the intersection of the measurable sets $\{x \in X^{\mathbb{Z}} : \pi_{n+1}x = T \circ \pi_n x\}$ so $\widetilde{X}$ is measurable. The shift map $\widetilde{T}$ is measurable and $\widetilde{X}$ is clearly invariant. It is also easy to check that $\pi_n : \widetilde{X} \to X$ is a factor map: $T\pi_n = \pi_n \widetilde{T}$.

As for uniqueness, if $\sigma : (Y, \mathcal{C}, \nu, S) \to (X, \mathcal{B}, T, \mu)$ we can define $\tau : Y \to X^{\mathbb{Z}}$ by $\tau(y)_n = \sigma(T^n y)$. Then the image is $\widetilde{X}$, $\sigma = \pi\tau$ is automatic, and one can show that $\tau\nu = \widetilde{\mu}$; again, we omit the details.

Lemma 5.3.1. *If a m.p.s. is ergodic then so is its natural extension.*

Proof. Let X be the original system and $\widetilde{X}$ its natural extension, $\pi : \widetilde{X} \to X$ the factor map. Then from the construction above it is clear that $f = \lim \mathbb{E}(f|\mathcal{B}_n)$ for $f \in L^2(\widetilde{\mu})$, where $\mathcal{B}_n = \widetilde{T}^n \pi^{-1}\mathcal{B}$. It is clear that $(\widetilde{X}, \mathcal{B}_n, \widetilde{\mu}, \widetilde{T}) \cong (X, \mathcal{B}, \mu, T)$ and is therefore ergodic. It is also clear that $\mathbb{E}(\widetilde{T}f|\mathcal{B}_n) = \widetilde{T}\mathbb{E}(f|\mathcal{B}_n)$, since $\mathcal{B}_n$ is invariant. Therefore if $f \in L^2(\widetilde{\mu})$ is $\widetilde{T}$-invariant then $\mathbb{E}(f|\mathcal{B}_n)$ is invariant. Since it corresponds to an invariant function on the ergodic system $(X, \mathcal{B}, \mu, T)$ it is constant. Since $f = \lim \mathbb{E}(f|\mathcal{B}_n)$, f is constant. Therefore $\widetilde{X}$ has no non-constant invariant functions, so it is ergodic. $\square$

Example 5.3.2. Let $(X, \mathcal{B}, \mu, T)$ be an ergodic measure preserving system and $A \in \mathcal{B}$. Define the return time $r_A(x)$ as usual. If T is invertible, Kac's formula (Theorem 3.3.1) ensures that $\int_A r_A d\mu = 1$. We can now prove this in the non-invertible case: let $(\widetilde{X}, \widetilde{T}.\widetilde{\mu})$ and $\pi : \widetilde{X} \to X$ be the natural extension. Let $\widetilde{A} = \pi^{-1}A$ and $r_{\widetilde{A}}$ the return time function in $\widetilde{X}$. Since $\widetilde{X}$ is ergodic, by Kac's formula $\int r_{\widetilde{A}} d\widetilde{\mu} = 1$. But clearly $r_{\widetilde{A}} = r_A \circ \pi$ so $\int r_{\widetilde{A}} d\widetilde{\mu} = \int r_A d\pi\widetilde{\mu} = \int r_A d\mu$, and the claim follows.

5.4 Inverse limits

A very similar construction gives the following theorem.

Theorem 5.4.1. *Let $(X_n, \mathcal{B}_n, T_n, \mu_n)$ be m.p.s. and $\pi_n : X_n \to X_{n-1}$ factor maps. Then there is a measure-preserving system $(\widetilde{X}, \widetilde{\mathcal{B}}, \widetilde{\mu}, \widetilde{T})$ and factor maps $\tau_n : \widetilde{X} \to X_n$ such that*

1. *$\tau_{n-1} = \pi_n \tau_n$.*
2. *$\widetilde{X}$ is unique in the sense that if Y is another m.p.s. and $\sigma_n : Y \to X_n$ satisfies $\sigma_{n-1} = \pi_n \sigma_n$, then there is a factor map $\tau : Y \to \widetilde{X}$ such that $\sigma_n = \tau_n \tau$.*
3. *If all the X_ns are ergodic, so is $\widetilde{X}$.*

5.5 Skew products

Let $(X, \mathcal{B}, \mu, T)$ be a m.p.s. and G a compact group with normalized Haar measure m (for example $\mathbb{R}/\mathbb{Z}$ with Lebesgue measure). Fix a probability space$(Y, \mathcal{C}, \nu)$ and let $G = \mathrm{Aut}(\nu)$ be the group of invertible ν-preserving maps $Y \to Y$. This group can be identified with a subgroup of the bounded linear operators on $L^2(\nu)$ with the operator topology (strong or weak) and given the induced Borel structure.

Now suppose that $\Phi : x \to \varphi_x$ is a measurable map $X \to G$. Now form $X_\Phi = \times Y$ and $T_\Phi : X \times Y \to X \times Y$ by $T_\Phi(x, y) = (Tx, \Phi(x)y)$.

Lemma 5.5.1. *Let $\theta = \mu \times \nu$. Then θ is T_Φ-invariant.*

Proof. Using Fubini, for $A \subseteq X \times Y$ we have:

$$\begin{aligned}
\theta(T_\Phi^{-1} A) &= \int T_\Phi^{-1}(A) d\theta \\
&= \int 1_A(T_\gamma x) d\theta(x) \\
&= \int \left(\int 1_A(Tx, \varphi_x y) d\nu(y) \right) d\mu(x) \\
&= \int \left(\int 1_A(Tx, y) d\nu(y) \right) d\mu(x) \\
&= \int \left(\int 1_A(Tx, y) d\mu(x) \right) d\nu(y) \\
&= \int \left(\int 1_A(x, y) d\mu(x) \right) d\nu(y) \\
&= \theta(A)
\end{aligned}$$

as claimed. □

The m.p.s. (X_Φ, θ, T_Φ) is called a *skew-product* over X (the base) with cocycle Φ.

Note that $\pi(x, y) = x$ is a factor map from $(X \times Y, \theta, T_\Phi)$ to (X, μ, T). In a sense there is a converse:

Theorem 5.5.2 (Rohlin)**.** *Let $\pi : Y \to X$ be a factor map between standard Borel spaces. Then Y is isomorphic to a skew-product over X.*

Note that we already know that the measure on Y disintegrates into conditional measures. What is left to do is identify all the fibers with a single space. Then Φ is defined as the transfer map between fibers. We omit the technical details.

Chapter 6

Weak mixing

6.1 Weak mixing

Ergodicity is the most basic "mixing" property; it ensures that any two sets A, B intersect at some time, i.e. there is an n with $\mu(A \cap T^{-n}B) > 0$. It also ensures that every orbit "samples" the entire space correctly. The notion of weak mixing is a natural generalization to pairs (and later k-tuples) of points. Thus we are interested in understanding the implications of independent pairs x, y of points sampling the product space. This is ensured by the following definition.

Definition 6.1.1. A m.p.s. $(X, \mathcal{B}, \mu, T)$ is **weak mixing** if $(X \times X, \mathcal{B} \times \mathcal{B}, \mu \times \mu, T \times T)$ is ergodic.

This property has many equivalent and surprising forms which we will now discuss. Let us first give some example.

Since X is a factor of $X \times X$ and the factor of an ergodic system is ergodic, weak mixing implies ergodicity. The converse is false.

Example 6.1.2. The translation of $\mathbb{R}/\mathbb{Z}$ by $\alpha \notin \mathbb{Q}$ is ergodic but not weak mixing. Indeed, let d be a translation-invariant metric on $\mathbb{R}/\mathbb{Z}$, so that $d(T_\alpha x, T_\alpha y) = d(x, y)$. Then the function $d : \mathbb{R}/\mathbb{Z} \times \mathbb{R}/\mathbb{Z} \to \mathbb{R}$ is $T \times T$-invariant and is not a.e. constant with respect to Lebesgue measure on the product space (each level set is a sub-manifold of measure 0). Therefore the product is not ergodic.

Example 6.1.3. More generally, if (X, d) is a compact metric space, $T : X \to X$ an isometry, and μ an invariant measure supported on more than one point, then the same argument shows that (X, μ, T) is not weak mixing.

Example 6.1.4. Let $X = \{0, 1\}^{\mathbb{Z}}$ with the product measure $\mu = \{\frac{1}{2}, \frac{1}{2}\}^{\mathbb{Z}}$ and the shift S. Then $X \times X \cong \{00, 01, 10, 11\}^{\mathbb{Z}}$ with the product measure $\mu \times \mu = \{\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\}^{\mathbb{Z}}$, which is ergodic. Therefore (X, μ, T) is weak mixing.

In the last example, X was strong mixing: $\mu(A \cap T^{-n}B) \to \mu(A), (B)$ for all measurable A, B, or: $\int f \cdot T^n g \to \int f \int g$ in L^2 for $f, g \in L^2$. Weak mixing has a similar characterization. Before stating it we need a definition. For a subset $I \subseteq \mathbb{N}$ we define the upper density to be

$$d(I) = \limsup_{N\to\infty} \frac{|I \cap \{1,\ldots,N\}|}{N}$$

Definition 6.1.5. A sequence $a_n \in \mathbb{R}$ **converges in density** to $a \in \mathbb{R}$, denoted $a_n \xrightarrow{D} a$ or $\text{D-lim}\, a_n = a$, if

$$d(\{n \,:\, |a_n - a| > \varepsilon\}) = 0 \qquad \text{for all } \varepsilon > 0$$

Compare this to the usual notion of convergence, where we require the set above to be finite rather than 0-density. Since the union of finitely many sets of zero density has zero density, this notion of limit has the usual properties (with the exception that a subsequence may not have the same limit). One can also show the following:

Lemma 6.1.6. *For a bounded sequence a_n, the following are equivalent:*

1. $a_n \xrightarrow{D} a$.
2. $\frac{1}{N}\sum_{n=0}^{N} |a_n - a| = 0$.
3. $\frac{1}{N}\sum_{n=0}^{N} (a_n - a)^2 = 0$.
4. *There is a subset $I = \{n_1 < n_2 < \ldots\} \subseteq \mathbb{N}$ with $d(I) = 1$ and $\lim_{k\to\infty} a_{n_k} = a$.*

We leave the proof as an exercise.

Theorem 6.1.7. *For a m.p.s. $(X, \mathcal{B}, \mu, T)$ the following are equivalent:*

1. *X is weak mixing.*
2. $\frac{1}{N}\sum_{n=0}^{N-1} \left|\int f \cdot T^n g d\mu - \int f d\mu \int g d\mu\right| \to 0$ *for all $f, g \in L^2(\mu)$.*
3. $\frac{1}{N}\sum_{n=0}^{N-1} |\mu(A \cap T^{-n}B) - \mu(A)\mu(B)| \to 0$ *for every $A, B \in \mathcal{B}$.*
4. $\mu(A \cap T^{-n}B) \xrightarrow{D} \mu(A)\mu(B)$ *for every $A, B \in \mathcal{B}$.*

This is based on the following:

Lemma 6.1.8. *For a m.p.s. $(Y, \mathcal{C}, \nu, s)$ the following are equivalent:*

1. *Y is ergodic.*
2. $\frac{1}{N}\sum_{n=0}^{N} (\mu(A \cap T^{-n}B) - \mu(A)\mu(B)) \to 0$ *for every $A, B \in \mathcal{C}$.*
3. $\frac{1}{N}\sum_{n=0}^{N} \int \left(f \cdot T^n g - \int f d\mu \cdot \int g d\mu\right) \to 0$ *for every $f, g \in L^2(\nu)$.*

Proof. The equivalence of the last two conditions is standard, we prove equivalence of the first two.

If the limit holds then, it implies that $\mu(A \cap T^{-n}B) = \int 1_A \cdot T^n 1_B d\mu > 0$ infinitely often if $\mu(A)\mu(B) > 0$. This gives ergodicity.

Conversely, if the system is ergodic then by the mean ergodic theorem, $S_N g = \frac{1}{N}\sum_{n=0}^{N-1} T^n g \to \int g$ in L^2 for any $g \in L^2$. So by continuity of the inner product,

$$\begin{aligned}
\lim_{N\to\infty} \sum_{n=0}^{N-1} \mu(A \cap T^{-n}B) &= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N} \left(\int 1_A \cdot T^n 1_B d\mu \right) \\
&= \lim_{N\to\infty} \langle 1_A, S_N 1_B \rangle \\
&= \left\langle 1_A, \int 1_B d\mu \right\rangle \\
&= \langle 1_A, \mu(B) \rangle \\
& \quad \mu(A)\mu(B)
\end{aligned}$$

which is what we wanted. □

Proof of the Proposition. Since $|\mu(A\cap T^{-n}B) - \mu(A)\mu(B)| \leq 1$, the equivalence of (3) and (4) is Lemma 6.1.6. The equivalence of (2) and (3) is standard be approximating L^2 functions by simple functions and using Cauchy-Schwartz. So we have to prove that (1) $\iff$ (2).

We may suppose that X is ergodic, since otherwise (1) fails trivially and (2) fails already without absolute values by the lemma. Then for $f, g \in L^\infty(\mu)$ we know from the lemma that

$$\frac{1}{N} \sum_{n=0}^{N-1} \int f \cdot T^n g \to \int f \int g \tag{6.1}$$

Suppose that X is weak mixing. Let $f' = f(x)f(y) \in L^2(\mu \times \mu)$ and define $g' \in L^2(\mu \times \mu)$ similarly. By ergodicity of $X \times X$

$$\frac{1}{N} \sum_{n=0}^{N-1} \int f' \cdot (T \times T)^n g' d\mu \times \mu \to \int f' \int g'$$

but

$$\begin{aligned}
\int f'(T\times T)^n g' d\mu \times \mu &= \int \int f(x)f(y)g(T^n x)g(T^n y)\, d\mu(x)d\mu(y) \\
&= (\int f \cdot T^n g d\mu)^2
\end{aligned}$$

and

$$\int f' d\mu \times \mu \cdot \int g' d\mu \times \mu = (\int f d\mu)^2 (\int g d\mu)^2$$

Thus we have proved:

$$\frac{1}{N}\sum_{n=0}^{N-1}(\int f\cdot T^n g d\mu)^2 \to (\int f\int g)^2 \tag{6.2}$$

Combining this with $\frac{1}{N}\sum_{n=0}^{N-1}\int f\cdot T^n g \to \int f\int g$, we find that

$$\frac{1}{N}\sum_{n=0}^{N-1}(\int f\cdot T^n g d\mu - (\int f\int g))^2 \to 0 \tag{6.3}$$

and since the terms are bounded, this implies (2).

In the opposite direction, assume (4), which is equivalent to (2). We must prove that $X\times X$ is ergodic, or equivalently, that for every $F,G\in L^2(\mu\times\nu)$,

$$\frac{1}{N}\sum_{n=0}^{N-1}\int F\cdot T^n G\,d\mu\times\nu \to \int Fd\mu\times\nu\cdot\int Gd\mu\times\nu$$

By approximation it is enough to prove this when $F(x_1,x_2)=f_1(x_1)f_2(x_2)$ and $G(x_1,x_2)=g_1(x_1)g_2(x_2)$. Furthermore we may assume that f_1,f_2,g_1,g_2 are simple, and even indicator functions $1_A,1_{A'},1_B,1_{B'}$. Thus we want to prove that for A,A',B,B',

$$\frac{1}{N}\sum_{n=0}^{N-1}\mu(A\cap T^{-n}B)\mu(A\cap T^{-n}B')\to\mu(A)\mu(B)\mu(A')\mu(B')$$

But $\mu(A\cap T^{-n}B)\to\mu(A)\mu(B)$ in density and $\mu(A'\cap T^{-n}B')\to\mu(A')\mu(B')$ in density, so the same is true for their product; and hence the averages converge as desired. □

Corollary 6.1.9. *If $(X,\mathcal{B},\mu,T)$ is mixing then it is weak mixing.*

Proof. $\mu(A\cap T^{-n}B)\to\mu(A)\mu(B)$ implies it in density. Now apply the proposition. □

6.2 Weak mixing as a multiplier property

Proposition 6.2.1. *$(X,\mathcal{B},\mu,T)$ is weak mixing if and only if $X\times Y$ is ergodic for every ergodic system $(Y,\mathcal{C},\nu,S)$.*

Proof. One direction is trivial: if $X\times Y$ is ergodic whenever Y is ergodic then this is true in particular for the 1-point system. Then $X\times Y\cong X$ so X is ergodic. It then follows taking $Y=X$ that $X\times X$ is ergodic, so X is weak mixing.

In the opposite direction we must prove that for every $F,G\in L^2(\mu\times\nu)$,

$$\frac{1}{N}\sum_{n=0}^{N-1}\int F\cdot(T\times S)^n G\,d\mu\times\nu\to\int Fd\mu\times\nu\int Gd\mu\times\nu$$

As before it is enough to prove this when $F(x_1,x_2) = f_1(x_1)f_2(x_2)$ and $G(x_1,x_2) = g_1(x_1)g_2(x_2)$, and it reduces to

$$\frac{1}{N}\sum_{n=0}^{N-1}\int f_1(x)T^n f_2(x)d\mu(x)\cdot\int g_2(x)S^n g_2(x)d\nu(x) \to \int f_1 d\mu\int f_2 d\mu\int g_1 d\nu\int g_2 d\nu$$

Splitting $L^2(\mu)$ into constant functions and their orthogonal complement (functions of integral 0), it is enough to prove this for f_1 in each of these spaces. If f_1 is constant then $\int f_1(x)T^n f_2(x)d\mu(x) = \int f_1 d\mu \int f_2 d\nu$ and the claim follows from ergodicity of S. On the other hand if $\int f_1 d\mu = 0$ we have

$$\left(\frac{1}{N}\sum_{n=0}^{N-1}\int f_1(x)T^n f_2(x)d\mu(x)\cdot\int g_2(x)S^n g_2(x)d\nu(x)\right)^2$$

$$\leq \frac{1}{N}\sum_{n=0}^{N-1}\left(\int f_1(x)T^n f_2(x)d\mu(x)\right)^2\cdot\frac{1}{N}\sum_{n=0}^{N-1}\left(\int g_2(x)S^n g_2(x)d\nu(x)\right)^2$$

but

$$\frac{1}{N}\sum_{n=0}^{N-1}\left(\int f_1(x)T^n f_2(x)d\mu(x)\right)^2 =\leq \frac{1}{N}\sum_{n=0}^{N-1}\left(\int f_1(x)T^n f_2(x)d\mu(x) - \int f_1\int f_2\right)^2 \to 0$$

by weak mixing of X and we are done. □

Corollary 6.2.2. *If X is weak mixing so is $X\times X$ and $X\times X\times\ldots\times X$.*

Proof. For any ergodic Y, $(X\times X)\times Y = X\times(X\times Y)$. Since $X\times Y$ is ergodic so is $X\times(X\times Y)$. The general claim follows in the same way. □

More generally,

Corollary 6.2.3. *If $X_1, X_2,\ldots$ are weak mixing so are $X_1\times X_2\times\ldots$.*

Also

Corollary 6.2.4. *If $(X,\mathcal{B},\mu,T)$ is weak mixing then so is T^n for all $n\in\mathbb{N}$ (if T is invertible, also negative n).*

Proof. Since $T\times T$ is ergodic if and only if $T^{-1}\times T^{-1}$ is ergodic, weak mixing of T and T^{-1} are equivalent, so we only need to consider $n>0$.

First we show that t weak mixing implies T^m is ergodic. Otherwise let $f\in L^2$ be a T^m invariant and non-constant function. Consider the system $Y=\{0,\ldots,m-1\}$ and $S(y)=y+1 \bmod m$ with uniform measure. Since X,T is weak mixing, $X\times Y, T\times S$ is ergodic. Let $F(x,i)=f(T^{m-i}x)$. Then

$$F(Tx,Si) = f(T^{m+1-(i+1)}x) = f(T^{m-i}x) = F(x,i)$$

so F is $T\times S$ invariant and non-constant, a contradiction. Thus T^m is ergodic.

Applying this to $T\times T$, which is weak mixing, we find that $(T\times T)^m$ is ergodic, equivalently $T^m\times T^m$ is ergodic, so T^m is weak mixing. □

6.3 Isometric factors

An isometric m.p.s. is a m.p.s. $(Y, \mathcal{C}, \nu, S)$ where Y is a compact metric space, $\mathcal{C}$ the Borel algebra, and S preserves the metric. We say that system is nontrivial if ν is not a single atom.

We have already seen that nontrivial isometric systems are not weak mixing, since if d is the metric then $d : Y \times Y \to \mathbb{R}$ is a non-trivial invariant function. In this section we prove the following.

Theorem 6.3.1. *Let $(X, \mathcal{B}, \mu, T)$ be an invertible m.p.s. on a standard Borel space. Then it is weak mixing if and only if it does not have nontrivial isometric factors (up to isomorphism).*

One direction is easy: If Y is an isometric factor then $X \times X$ factors onto $Y \times Y$ and so, since the latter is not ergodic, neither is the former, For the converse direction we must show that if X is not weak mixing then it has non-trivial isometric factors. We can assume that X is ergodic, since otherwise the ergodic decomposition gives a factor to a compact metric space with the identity map, which is an isometry.

Recall that if $(X, \mathcal{B}, \mu)$ is a probability space then there is a pseudo-metric on $\mathcal{B}$ defined by

$$d(A, B) = \mu(A \Delta B) = \|1_A - 1_B\|_1$$

Identifying sets that differ in measure 0 dives a metric on equivalence classes, and the resulting space may be identified with the space of 0-1 valued functions in L^1, which is the same as the set of indicator functions. This is an isometry and since the latter space is closed in L^1, it is complete.

Now suppose that X is not weak mixing and let $A \subseteq X \times X$ be a non-trivial invariant set. Consider the map $X \to L^1$ given by $x \to 1_{A_x}$ where

$$A_x = \{y \in X \,:\, (x, y) \in A\}$$

The map is measurable, we use the fact that the Borel structure of the unit ball in L^1 in the norm and weak topologies coincide (this fact is left as an exercise). Then we only need to check that

$$x \mapsto \int 1_{A_x}(y) g(y) \, d\mu(y)$$

is measurable for every $g \in L^\infty$. For fixed g, this clearly holds when A is a product set or a union of product sets, and the general case follows from the monotone class theorem.

Now, notice that

$$\begin{aligned} TA_x &= \{Ty \,:\, (x, y) \in A\} \\ &= \{y \,:\, (x, T^{-1}y) \in A\} \\ &= \{y \,:\, (x, T^{-1}y) \in T^{-1}A\} \\ &= \{y \,:\, (Tx, y) \in A\} \\ &= A_{Tx} \end{aligned}$$

so $\pi : x \to A_x$ commutes with the action of T on X and L^1. Finally, the action of T on L^1 is an isometry. Therefore we have proved:

Claim 6.3.2. If T is not weak mixing then there is a complete metric space (Y, d), and isometry $T : Y \to Y$ and a Borel map $\pi : X \to Y$ such that $T\pi = \pi T$.

Let $\nu = \pi\mu$, the image measure; it is preserved. Thus (Y, ν, T) is almost the desired factor, except that the space Y is not compact (and there is another technicality we will mention later). To fix these problems we need a few general facts.

Definition 6.3.3. Let (Y, d) be a complete metric space. A subset $Z \subseteq Y$ is called ***totally bounded*** *if for every ε there is a finite set $Z_\varepsilon \subseteq Y$ such that $Z \subseteq \bigcup_{z \in Z_\varepsilon} B_\varepsilon(z)$.*

Lemma 6.3.4. *Let $Z \subseteq Y$ as above. Then $\overline{Z}$ is compact if and only if Z is totally bounded.*

Proof. This is left as an exercise. □

Lemma 6.3.5. *Let (Y, d) be a complete separable metric space, $S : Y \to Y$ an isometry and μ an invariant and ergodic Borel probability measure. Then* $\operatorname{supp} \mu$ *is compact.*

Proof. Let $C = \operatorname{supp} \mu$. This is a closed set and is clearly invariant so we only need to show that it is compact. For this it is enough to show that it is totally bounded.

Choose a μ-typical point y. By the ergodic theorem, its orbit is dense in C. Furthermore since S is an isometry, $B_r(S^n y) = S^n B_r(y)$. Now let $z \in C$. There is an n with $d(z, S^n y) < r$ so $B_r(T^n y) \subseteq B_{2r}(z)$, hence

$$\mu(B_r(y)) \leq \mu(B_{2r}(z))$$

This is true for every r.

Now, let $\{z_i\}$ be a maximal set of r-separated points in C. the set must be finite, because $B_{r/2}(z_i)$ are disjoint balls of mass uniformly bounded below. Therefore $B_{2r}(z_i)$ is a finite cover of C, and since r was arbitrary, C is totally bounded. □

Let ν be, as before, the image of μ under the map $\pi : X \to L^1$. Then $\operatorname{supp} \nu$ is compact and we can replace X by $\pi^{-1}(\operatorname{supp} \nu)$, which has full measure. We are done.

6.4 Eigenfunctions

Definition 6.4.1. Let T be an operator on a Hilbert space H. Then $\lambda \in \mathbb{C}$ is an eigenvalue and u a corresponding eigenfunction if $Uu = \lambda u$. We denote the set of eigenvalues by $\Sigma(U)$.

As in the finite-dimensional case, unitary operators have only eigenvalues of modulus 1, because if $Uv = \lambda v$ and $\|v\| = 1$ then

$$\lambda\overline{\lambda} = \langle \lambda v, \lambda v\rangle = \langle Uv, Uv\rangle = \langle U^*Uv, v\rangle = \langle v, v\rangle = 1$$

Lemma 6.4.2. *If $Uv = \lambda v$ and $Uv' = \lambda' v'$ and $\lambda \neq \lambda'$ then $v \perp v'$. In particular, if H is separable then $|\Sigma(U)| \leq \aleph_0$.*

Proof. We can assume $\|v\| = \|v'\| = 1$. Then the same calculation as above shows that $\lambda\overline{\lambda'}\,\langle v, v'\rangle = \langle v, v'\rangle$, which is possible if and only if $\langle v, v'\rangle = 0$. $\square$

From now on let $(X, \mathcal{B}, \mu, T)$ be an invertible m.p.s. and denote also by T the induced unitary operator on $L^2(\mu)$.

Notice that an invariant function is an eigenvector of eigenvalue 1, and the space of such functions is always at least 1-dimensional, since it contains the constant functions. Therefore,

Lemma 6.4.3. *T is ergodic if and only if 1 is a simple eigenvalue (the corresponding eigenspace has complex dimension 1).*

If T is ergodic and $e^{2\pi i\alpha}$ is an eigenvalue with eigenfunction $f \in L^2(\mu)$, then $|f|$ satisfies

$$T|f| = |Tf| = |e^{2\pi i\alpha} f| = |f|$$

so $|f|$ is invariant. Therefore,

Corollary 6.4.4. *If T is ergodic then all eigenfunctions are constant.*

By convention, we always assume that eigenfunctions have modulus 1.

Therefore if $\alpha, \beta \in \Sigma(T)$ with corresponding eigenfunctions f, g, then $f \cdot g$ has modulus 1, hence is in L^2, and

$$T(fg) = Tf \cdot Tg = \alpha f \cdot \beta g = \alpha\beta \cdot fg$$

so $\alpha\beta \in \Sigma(T)$ and fg is an eigenfunction for it. Similarly considering $\overline{f}$ we find that $\overline{\alpha} = \alpha^{-1} \in \Sigma(T)$. This shows that

Corollary 6.4.5. *$\Sigma(T) \subseteq \mathbb{C}$ is a (multiplicative) group and if T is ergodic then all eigenvalues are simple.*

Proof. The first statement is immediate. For the second, note that if f, g are distinct eigenfunctions for the same α then $f\overline{g}$ is an eigenfunction for $\alpha\overline{\alpha} = 1$, hence $f\overline{g} = 1$, or (using $\overline{g} = g^{-1}$) we have $f = g$. $\square$

Note that $\Sigma(U)$ is not generally a group if U is unitary.

The following observation is trivial:

Lemma 6.4.6. *If $\pi : (X, T) \to (Y, S)$ is a factor map between measure preserving systems then $\Sigma(S) \subseteq \Sigma(T)$.*

Proof. If $\lambda \in \Sigma(S)$ let f be a corresponding eigenfunction, $Sf = \lambda f$ then

$$T(f\pi) = f\pi T = fS\pi = \lambda f\pi$$

This shows that $\lambda \in \Sigma(T)$. □

Theorem 6.4.7. *$(X, \mathcal{B}, \mu, T)$ is ergodic then it is weak mixing if and only if 1 is the only eigenvalue, i.e. $\Sigma(T) = \{1\}$.*

One direction of the proof is easy. Note that from a dynamical point of view, an eigenfunction f for $\alpha \in \Sigma$ of $(X, \mathcal{B}, \mu, T)$ is a factor map to rotation on S^1, since if f is an eigenfunction with eigenvalue α and $R_\alpha z = \alpha z$ is the rotation on S^1 then the eigenfunction equation is just

$$fT(x) = \alpha f(x) = R_\alpha(f(x))$$

so $fT = R_\alpha f$, and the image measure $f\mu$ is a probability measure on S^1 invariant under R_α. Of course, R_α is an isometry of S^1, and if $\alpha \neq 1$ then $f\mu$ is not a single atom; thus we have found a non-trivial isometric factor of X and so it is not weak mixing.

For the other direction we need the following. For a group G and $g \in G$ let $L_g : G \to G$ denote the map $h \mapsto gh$. A dynamical system in which G is a compact group and the map is L_g is called a group translation. Note that Haar measure is automatically invariant, since by definition it is the unique Borel probability measure that is invariant under all translations. If Haar measure is also ergodic the we say the translation is ergodic.

Proposition 6.4.8. *Let (Y, d) be a compact metric space and $S : Y \to Y$ an isometry with a dense orbit. Then there is a compact metric group G and $g \in G$ and a homeomorphism $\pi : Y \to G$ such that $L_g\pi = \pi S$. Furthermore if ν is an invariant measure on Y then it is ergodic and $\pi\nu$ is Haar measure on G.*

Proof. Consider the group Γ of isometries of Y with the sup metric,

$$d(\gamma, \gamma') = \sup_{y \in Y} d(\gamma(y), \gamma'(y))$$

Then (Γ, d) is a complete metric space, and note that it is right invariant: $d(\gamma \circ \delta, \gamma \circ \delta) = d(\gamma, \gamma')$.

Let $y_0 \in Y$ have dense orbit and set $Y_0 = \{S^n y_0\}_{n \in \mathbb{Z}}$. If the orbit is finite, $Y = Y_0$ is a finite set permuted cyclically by S, so the statement is trivial. Otherwise $y \in Y_0$ uniquely determines n such that $S^n y_0 = y$ and we can define $\pi : Y_0 \to \Gamma$ by $y \mapsto S^n \in \Gamma$ for this n.

We claim that π is an isometry. Fix $y, y' \in Y_0$, so $y = S^n y_0$ and $y' = S^{n'} y_0$, so

$$d(\pi y, \pi y') = \sup_{z \in Y} d(S^n z, S^{n'} z)$$

Given $z \in Y$ there is a sequence $n_k \to \infty$ such that $S^{n_k} y_0 \to z$. But then

$$\begin{aligned} d(S^n z, S^{n'} z) &= d(S^n(\lim S^{n_k} y_0), S^{n'}(\lim S^{n_k} y_0)) \\ &= \lim d(S^n S^{n_k} y_0, S^{n'} S^{n_k} y_0) \\ &= \lim d(S^{n_k}(S^n y_0), S^{n_k}(S^{n'} y_0)) \\ &= \lim d(S^n y_0, S^{n'} y_0) \\ &= d(S^n y_0, S^{n'} y_0) \\ &= d(y, y') \end{aligned}$$

Thus $d(\pi y, \pi y') = d(y, y')$ and π is an isometry $Y_0 \hookrightarrow \Gamma$. Furthermore, for $y = S^n y_0 \in Y_0$,

$$\pi(Sy) == \pi(SS^n y) = S^{n+1} = L_S S^n = L_S \pi(y)$$

It follows that π extends uniquely to an isometry with $Y \hookrightarrow \Gamma$ also satisfying $\pi(Sy) = S(\pi y)$. The image $\pi(Y_0)$ is compact, being the continuous image of the compact set Y. Since $\pi(Y_0) = \{S^n\}_{n \in \mathbb{Z}}$ and this is a group its closure is also a group G.

Finally, suppose ν is an invariant measure on Y. Then $m = \pi\nu$ is L_S invariant on G. Since it is invariant under L_S it is invariant under $\{L_S^n\}_{n \in \mathbb{Z}}$, and this is a dense set of elements in G. Thus m it is invariant under every translation in G, and there is only one such measure up to normalization: Haar measure. The same argument applies to every ergodic components of m (w.r.t. L_S) and shows that the ergodic components are also Haar measure. Thus m is L_S-ergodic and since π is an isomorphism, (Y, ν, S) is ergodic. □

Corollary 6.4.9. *If $(X, \mathcal{B}, \mu, T)$ is ergodic but not weak mixing, then a non-trivial ergodic group translation as a factor.*

Finally, for the existence of eigenfunctions we will rely on a well-known result from group theory.

Theorem 6.4.10 (Peter-Weyl)**.** *If G is a compact topological group with Haar measure m then the characters form an orthonormal basis for $L^2(m)$ (a character is a continuous homomorphisms $\xi : G \to S^1 \subseteq \mathbb{C}$).*

Corollary 6.4.11. *Let G be a compact group with Haar measure m and $g \in G$. Then $L^2(m)$ is spanned by eigenfunctions of L_g.*

Proof. If ξ is a character then

$$L_g \xi(x) = \xi(gx) = \xi(g)\xi(x)$$

so ξ is an eigenfunction with eigenvalue $\xi(g)$. Since the characters span $L^2(m)$ we are done. □

In summary, we have seen that if $(X, \mathcal{B}, \mu, T)$ is ergodic but not weak mixing then it has a nontrivial isometric factor. This is isomorphic to a non-trivial group rotation, which has a non-trivial eigenfunction. This eigenfunction lifts via the composition of all these factor maps to a non-trivial eigenfunction on X. This completes the characterization of weak mixing in terms of existence of eigenfunctions.

6.5 Spectral isomorphism and the Kronecker factor

We will now take a closer look at systems that are isomorphic to group rotations.

Let $\mathbb{T}^\infty = (S^1)^\mathbb{N}$, which is a compact metrizable group with the product topology. Let m be the infinite product of Lebesgue measure, which is invariant under translations in $\mathbb{T}^\infty$. Given $\alpha \in \mathbb{T}^\infty$ let $L_\alpha : \mathbb{T}^\infty \to \mathbb{T}^\infty$ as usual be the translation map. Note that

$$L_\alpha^n x = \alpha^n x$$

Lemma 6.5.1. *The orbit closure G_α of $0 \in \mathbb{T}^\infty$ under L_α (that is, the closure of $\{\alpha^n : n \in \mathbb{N}\}$) is the closed subgroup generated by α.*

Proof. It is clear that it is contained in the group in question, is closed, contains α, and is a semigroup. The latter is because it is the closure of the semigroup $\{\alpha^n\}_{n\in\mathbb{N}}$; explicitly, if $x, y \in G_\alpha$ then $x = \lim \alpha^{n_i}$ and $y = \lim 0^{m_j}$, so

$$xy = (\lim \alpha^{m_j})(\lim \alpha^{n_i}) \lim \alpha^{m_i+n_i} \in G_\alpha$$

Thus, we only need to check that $x \in G_\alpha$ implies $x^{-1} \in G_\alpha$. Let $x = \lim \alpha^{n_i}$. We can assume $n_{i+1} > 2n_i$. Passing to a subsequence, we can also assume that $\alpha^{n_{i+1}-2n_i} \to y$. But then

$$xy = (\lim \alpha^{n_i})(\lim \alpha^{n_{i+1}-2n_i}) = \lim \alpha^{n_{i+1}-n_i} = xx^{-1} = 1$$

so $y = x^{-1}$. □

Lemma 6.5.2. *Let $\alpha \in \mathbb{T}^\infty = (S^1)^\mathbb{N}$ and G_α be as above. Let m_α be the Haar measure on G_α, equivalently, the unique L_α-invariant measure. Then $\Sigma(L_\alpha) = \langle \alpha_1, \alpha_2, \ldots \rangle \subseteq S^1$, the discrete group generated by the coordinates α_i of α.*

Proof. Let $\pi_n : \mathbb{T}^\infty \to S^1$ denote the n-th coordinate projection. Clearly $\pi_n(L_\alpha x) = \alpha_n x_n = \alpha_n \pi_n(x)$, so the functions π_n are eigenfunctions of $(G_\alpha, m_\alpha, L_\alpha)$. Let $\mathcal{A}$ denote the $\mathbb{C}$-algebra generated by $\{\pi_n\}$. This is an algebra of continuous functions that separate points in $\mathbb{T}^\infty$ and certainly in G_α, so they are dense in $L^2(m_\alpha)$. Since this algebra consists precisely of the eigenfunctions with eigenvalues in $\langle \alpha_1, \alpha_2, \ldots \rangle$ we are done. □

Proposition 6.5.3. *Let $(X, \mathcal{B}, \mu, T)$ be an ergodic measure preserving system on a standard Borel space. Let $\Sigma(T) = \{\alpha_1, \alpha_2, \ldots\}$ and $\alpha = (\alpha_1, \alpha_2, \ldots) \in \mathbb{T}^\infty$. Then*

1. *$(G_\alpha, m_\alpha, L_\alpha)$ is a factor of X*
2. *If $L^2(\mu)$ is spanned by eigenfunctions then the factor map is an isomorphism: $X \cong G_\alpha$.*
3. *If $\tau : X \to Y$ is a factor map to an isometric system (Y, S, ν) then π factors through G_α.*

Proof. We may restrict to a subset of X of full measure where all the eigenfunctions f_i of α_i are defined, and satisfy $Tf_i = \alpha_i f_i$. Let $F : X \to \mathbb{T}^\infty$ denote the map $F(x) = (f_1(x), f_2(x), \ldots)$. Then it is immediate that $F(Tx) = L_\alpha F(x)$. Let $\nu = F\mu$, which is an L_α-invariant measure on $\mathbb{T}^\infty$, and $y \in \operatorname{supp} \nu$ be a point with dense orbit (which exists by ergodicity). Consider the map $L_{y^{-1}} x \mapsto y^{-1}x$, which commutes with L_α (because $\mathbb{T}^\infty$ is abelian) and note that $L_{y^{-1}}$ maps the L_α-orbit of y to the L_α-orbit of 1. Writing $m = L_{y^{-1}}\nu$ it follows that m is L_α-invariant and $\operatorname{supp} m = G_\alpha$. Hence $m = m_\alpha$. Also $\pi : x \mapsto L_{y^{-1}} F(x)$ is a factor map from X to G_α.

Suppose that the eigenfunctions span $L^2(\mu)$. Since each eigenfunction is lifted by π from an eigenfunction of G_α, we find that $L^2(\pi^{-1}\mathcal{B}_\alpha) = L^2(\mu)$, where $\mathcal{B}_\alpha$ is the Borel algebra of G_α. It follows that π is 1-1 a.e. and by standardness it is an isomorphism.

Finally suppose $\tau : X \to Y$ as in the statement. Then (Y, X) is a group rotation, its eigenvectors are dense in $L^2(\nu)$, and we have an isomorphism $\sigma : Y \to G_\beta$ where $\beta = (\beta_1, \beta_2, \ldots)$ enumerates $\Sigma(S)$. Now, each eigenvector f of (G_β, S) lifts to one of X and since the multiplicity is 1, this is *the* eigenvector for its eigenvalue. Thus the coordinates of α include those of β. Let $u : G_\alpha \to \mathbb{T}^\infty$ denote projection to the coordinates corresponding to eigenvalues of S; then the map $u \circ F = \sigma \circ \tau$ defined above. The claim follows. □

Let us point out two consequences of this theorem. First,

Definition 6.5.4. An ergodic measure preserving system has discrete spectrum if L^2 is spanned by eigenfunctions.

Corollary 6.5.5. *Discrete spectrum systems are isomorphic if and only if the induced unitary operators are unitarily equivalent.*

Proof. This follows from the theorem above and the fact that two diagonalizable unitary operators are unitarily equivalent if and only if they have the same eigenvalues (counted with multiplicities), and ergodicity implies that all eigenvalues are simple. □

Second, part (3) of the theorem above shows that every measure preserving system has a maximal isometric factor. This factor is called the **Kronecker**

factor. The factor is canonical, although the factor map is not – one can always post-compose it with a translation of the group.

We emphasize that in general it is false that unitary equivalence implies ergodic-theoretic isomorphism. The easiest example to state is that the product measures $(1/2, 1/2)^{\mathbb{Z}}$ and $(1/3, 1/3, 1/3)^{\mathbb{Z}}$ with the shift map have unitarily isomorphic induced actions on L^2, but they are not isomorphic.

6.6 Spectral methods

Our characterization of weak mixing is, in the end, purely a Hilbert-space statement. Thus one should be able to prove the existence of eigenfunctions without use of the underlying dynamical system. This can be done with the help of the spectral theorem. Let us first give a brief review of the version we will use.

Let us begin with an example of a unitary operator. Let μ be a probability measure on the circle S^1 and let $M : L^2(\mu) \to L^2(\mu)$ be given by $(Mf)(z) = zf(z)$. Note that U preserves norm, since $|zf(z)| = |f(z)|$ for $z \in S^1$ and hence μ-a.e. z; it is invertible since the inverse is given by multiplication by $\overline{z}$.

The spectral theorem says that any unitary operator can be represented in this way on any invariant subspace for which it has an cyclic vector.

Theorem 6.6.1 (Spectral theorem for unitary operators)**.** *Let $U : H \to H$ be a unitary operator and $v \in H$ a unit vector such that $\overline{\{U^n v\}_{n=-\infty}^{\infty}} = H$. Then there is a probability measure $\mu_v \in \mathcal{P}(S^1)$ and a unitary operator $V : L^2(\mu) \to H$ such that $U = VMV^{-1}$, where $M : L^2(\mu) \to L^2(\mu)$ is as above. Furthermore $V(1) = v$.*

We give the main idea of the proof. The measure μ_v is characterized by the statement because its Fourier transform $\widehat{\mu}_v : \mathbb{Z} \to \mathbb{R}$ is given by

$$\widehat{\mu}_v(n) = \int z^n \, d\mu_v = \langle M^n 1, 1 \rangle_{L^2(\mu_v)} = \langle U^n v, v \rangle_H$$

Reversing this argument, in order to construct μ_v one starts with the sequence $a_n = \langle U^n 1, 1 \rangle_H$. This sequence is positive definite in the sense that for any sequence $\lambda_i \in \mathbb{C}$ and any n, $\sum_{i,j=1}^{n} \lambda_i \overline{\lambda_j} a_{i-j} \geq 0$:

$$\begin{aligned}
\sum_{i,j=1}^{\infty} \lambda_i \overline{\lambda_j} a_{i-j} &= \sum_{i,j=-n}^{n} \lambda_i \overline{\lambda_j} \left\langle U^{i-j} v, v \right\rangle_H \\
&= \sum_{i,j=-n}^{n} \left\langle U^i \lambda_i v, U^j \lambda_j v \right\rangle_H \\
&= \left\langle \sum_{i=-n}^{n} U^i \lambda_i v, \sum_{j=-n}^{n} U^j \lambda_j v \right\rangle_H \\
&= \left\| \sum_{i=-n}^{n} U^i \lambda_i v \right\|_2^2 \\
&\geq 0
\end{aligned}$$

Therefore, by a theorem of Hergolz (also Bochner) a_n is the Fourier transform of a probability measure on S^1 (note that $a_0 == \|v\|^2 = 1$).

One first defines V on complex polynomials $p(z) = \sum_{n=0}^{d} b_n z^n$ by $Vp = \sum_{n=0}^{d} b_n U^n v$. One can check that this preserves inner products; it suffices to check for monomials, and indeed

$$\langle V(bz^m), V(cz^n) \rangle = b\overline{c} \langle U^m v, U^n v \rangle = b\overline{c} \cdot a_{m-n} = b\overline{c} \int z^{m-n} d\mu_v = \int (bz^m)(\overline{cz^n}) d\mu_v = \langle bz^m, bz^n \rangle_{L^2(\mu_v)}$$

Since polynomials are dense in $L^2(\mu)$ it remains to extend V to measurable functions. The technical details of carrying this out can be found in many textbooks.

Lemma 6.6.2. *Let $U : H \to H$ be unitary and v a cyclic unit vector for U with spectral measure μ. Then $\alpha \in \Sigma(U)$ if and only if $\mu_v(\alpha) > 0$.*

Proof. If α is an atom of μ_v let $f = 1_{\{\alpha\}}$. This is a non-zero vector in $L^2(\mu_v)$, and $Mf(z) = zf(z) = \alpha f(z)$. Hence $\alpha \in \Sigma(M)$ and by the spectral theorem $\alpha \in \Sigma(U)$.

Conversely, suppose that $\mu_v(\{\alpha\}) = 0$. Consider the operator $U_\alpha(w) = \overline{\alpha} Uw$, which can easily be seen to be unitary. Clearly w is an eigenfunction with eigenvalue α if and only if $U_\alpha w = w$. Thus it suffices for us to prove that $\frac{1}{N} \sum_{n=0}^{N-1} U_\alpha^n w \to 0$ for all w, and, since v is cyclic and the averaging operator is linear and continuous, it is enough to check this for v. Transferring the problem to (S^1, μ_v, M), we must show that $\frac{1}{N} \sum_{n=0}^{N-1} \overline{\alpha}^n z^n \to 0$ in $L^2(\mu_v)$. We have

$$\frac{1}{N} \sum_{n=0}^{N-1} \overline{\alpha}^n z^n = \frac{1}{N} \frac{(\overline{\alpha} z)^N - 1}{\overline{\alpha} z - 1}$$

This converges to 0 at every point $z \neq \alpha$, hence μ_v-a.e., and it is bounded. Hence by bounded convergence, it tends to 0 in $L^2(\mu_v)$, as required. $\square$

Proposition 6.6.3. *Let U, H, v, μ_v be as above. If μ_v is continuous (has no atoms), then*

$$\frac{1}{N}\sum_{n=0}^{N-1} |(w, U^n w')| \to 0$$

for every $w, w' \in H_v$.

Proof. Using the fact that w, w' can be approximated in L^2 by linear combinations of $\{U^n v\}$, it is enough to prove this for $w, w' \in \{U^n v\}$. Since the statement we are trying to prove is formally unchanged if we replace w by $U^{\pm 1}w$ or w' by $U^{\pm 1}w'$, we may assume that $w = w' = v$. Also, we may square the summand, as we have seen this does not affect the convergence to 0 of the averages. Passing to the spectral setting, we have

$$\begin{aligned}
\frac{1}{N}\sum_{n=0}^{N-1} |(v, U^n v)|^2 &= \frac{1}{N}\sum_{n=0}^{N-1} \left|\int z^n d\mu_v\right|^2 \\
&= \frac{1}{N}\sum_{n=0}^{N-1} \left(\int z^n d\mu_v \cdot \int \overline{z}^n d\mu_v\right) \\
&= \frac{1}{N}\sum_{n=0}^{N-1} \left(\int\int z^n \overline{y}^n d\mu_v(y) d\mu_v(z)\right) \\
&= \int\int \left(\frac{1}{N} \cdot \frac{(z\overline{y})^N - 1}{z\overline{y} - 1}\right) d\mu_v \times \mu_v(z, y)
\end{aligned}$$

The integrand is bounded by 1 and tends pointwise to 0 off the diagonal $\{y = z\}$, which has $\mu_v \times \mu_v$-measure 0, since μ_v is non-atomic. Therefore by bounded convergence, the expression tends to 0. □

Corollary 6.6.4. *If $(X, \mathcal{B}, \mu, T)$ is ergodic then it is weak mixing if and only if μ_f is continuous (has no atoms) for every $f \perp 1$ (equivalently, the maximal spectral type is non-atomic except for an atom at 1), if and only if $\Sigma(T) = \{1\}$.*

Proof. Suppose for $f \perp 1$ the spectral measure μ_f is continuous. By the last proposition, $\frac{1}{N}\sum_{n=0}^{N-1} |\int f \cdot T^n f \, d\mu| \to 0$. For general f, g we can write $f = f' + \int f$, $g = g' + \int g$, where $f', g' \perp 1$. Substituting $f = f' + \int f$ into

$$\frac{1}{N}\sum_{n=0}^{N-1} |\int f \cdot T^n f \, d\mu - (\int f \, d\mu)(\int g d\mu)|$$

we obtain the expression

$$\frac{1}{N}\sum_{n=0}^{N-1} |\int f' \cdot T^n g' \, d\mu|$$

which by assumption $\to 0$. This was one of our characterizations of weak mixing.

Conversely suppose T is weak mixing. Then it has no eigenfunctions except 1 (this was the trivial direction of the eigenfunction characterization), so if $f \perp 1$ also $\overline{\mathrm{span}\{U^n f\}} \perp 1$ and so, since on this subspace T has no eigenfunctions, μ_f is continuous.

We already know that weak mixing implies $\Sigma(T) = \{1\}$. In the other direction, if T is not weak mixing, we just saw that there is some $f \perp 1$ with $\mu_f(\alpha) > 0$ for some α, and by the previous lemma, $\alpha \in \Sigma(T)$. $\square$

In a certain sense, we can now "split" the dynamics of a non-weak-mixing system into an isometric part, and a weak mixing part:

Corollary 6.6.5. *Let (X, B, μ, T) be ergodic. Then $L^2(\mu) = U \oplus V$, where $U = L^2(\mu, \mathcal{E})$ for $\mathcal{E} \subseteq \mathcal{B}$ the Kronecker factor, and V is an invariant subspace such that $T|_V$ is a weak-mixing in the sense that $\frac{1}{N}\sum_{n=0}^{N-1} |\int f \cdot T^n g \, d\mu| \to 0$ for $g \in V$.*

One should note that, in general, the subspace V in the corollary does not correspond to a factor in the dynamical sense.

An important consequence is the following:

Theorem 6.6.6. *Let $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ be ergodic measure preserving systems. Then $X \times Y$ is ergodic if and only if $\Sigma(T) \cap \Sigma(S) = \{1\}$.*

Proof. Let $Z = X \times Y$, $R = T \times S$, $\theta = \mu \times \nu$. If $\alpha \neq 1$ is a common eigenvalue of T, S with eigenfunctions f, g, then $h(x, y) = \overline{g}(y) \cdot f(x)$ is a non-trivial invariant function, since

$$h(R(z, y)) = \overline{g}(Sy) \cdot f(Tx) = \overline{\alpha g}(y) \alpha f(x) = h(x, y)$$

and so Z is not ergodic.

Conversely, write $L^2(\mu) = V_{wm} \oplus V_{pp}$, where $T|_{V_{pp}}$ as in the previous corollary, where $T|_{V_{pp}}$ has no eigenvalues, and decompose $L^2(\nu) = W_{wm} \oplus W_{pp}$ similarly. We must show that

$$\frac{1}{N}\sum_{n=0}^{N-1} \int h \cdot R^n h d\theta \to (\int h)^2$$

for every $h \in L^2(\theta)$ and it suffices to check this for $h = fg$, $f \in L^2(\mu)$, $g \in L^2(\nu)$, since the span of these is dense in L^2. Then $\int hR^n h d\theta = \int fT^n f d\mu \cdot \int gS^n g d\nu$. Also, since we can write $f = f_{wm} + f_{pp}$ and $g = g_{wm} + g_{pp}$ for $f_{wm} \in V_{wm}$ etc. we can expand the expression above, and obtain a sum of terms of the form

$$\frac{1}{N}\sum_{n=0}^{N-1} \int h \cdot R^n h d\theta \to (\int h)^2 = \sum_{i,j,s,t \in \{wm, pp\}} \left(\frac{1}{N} \sum_{n=0}^{N-1} (\int f_i T^n f_j d\mu)(\int g_s S^n g_t d\nu) \right) \tag{6.4}$$

Consider the terms in parentheses; they are all bounded independently of n. So if $i, j = wm$ we can bound

$$\left| \frac{1}{N} \sum_{n=0}^{N-1} (\int f_i T^n f_j d\mu)(\int g_s S^n g_t d\nu) \right| \leq C \cdot \frac{1}{N} \sum_{n=0}^{N-1} |\int f_i T^n f_j d\mu| \to 0$$

and similarly if $s, t = wm$. Also if $i = wm$, $s = pp$, then $T^n f_j = \alpha^n f_j$ for some α, and $\overline{f}_j \perp f_i$. Hence

$$\int f_i T^n f_j d\mu = \alpha^n \int f_i f_j d\mu = \alpha \left\langle f_i, \overline{f}_j \right\rangle = 0$$

Thus in 6.4 we only need to consider the case $i, j, s, t = pp$. In this case we can expand each of the functions as a series in normalized, distinct eigenfunctions: $f_{pp} = \sum \varphi_k$ and $g_{pp} = \sum \psi_k$ where $T\varphi_k = \alpha_k \varphi_k$ and $S\psi_k = \beta_k \psi_k$. We assume $\varphi_0 = const$ and $\psi_0 = const$. Expanding again using linearity, we must consider terms of the form

$$\frac{1}{N} \sum_{n=0}^{N-1} (\int \varphi_i T^n \varphi_j d\mu)(\int \psi_s S^n \psi_t d\nu) = \frac{1}{N} \sum_{n=0}^{N-1} (\alpha_j^n \int \varphi_i \varphi_j d\mu)(\beta_t^n \int \psi_s \psi_t d\nu)$$

Now, the first integral is 0 unless $\varphi_j = \overline{\varphi}_i$ and the second is 0 unless $\psi_t = \overline{\psi}_s$. If this is the case we have, writing $c_{i,s} = \|\varphi_j\|^2 \|\psi_j\|^2$

$$= \frac{1}{N} \sum_{n=0}^{N-1} \alpha_j^n \beta_t^n c_{i,s} = \begin{cases} c_{i,s} & \alpha_j = \overline{\beta}_t \\ c_{i,s} \frac{1}{N} \frac{(\alpha\beta)^N - 1}{\beta - 1} & \text{otherwise} \end{cases} \xrightarrow[N \to \infty]{} \begin{cases} c_{i,s} & \alpha_j = \overline{\beta}_t \\ 0 & \text{otherwise} \end{cases}$$

Since $\Sigma(T) \cap \Sigma(S) = \{1\}$ the limit is thus 0 except for $i = j = s = t = 0$. In the latter case, $c_{0,0} = \int \varphi_0^2 d\mu \int \psi_0^2 d\nu = (\int f)^2 (\int g)^2$, so this was the limit we wanted. $\square$

Chapter 7

Disjointness and a taste of entropy theory

7.1 Joinings and disjointness

Definition 7.1.1. A joining of measure preserving systems $(X, \mathcal{B}, \mu, T)$ and $(Y, \mathcal{C}, \nu, S)$ is a measure θ on $X \times Y$ that is invariant under $T \times S$ and projects to μ, ν ,respectively, under the coordinate projections.

Remark 7.1.2. There is a more general notion of a joining of X, Y, namely, a preserving system $(Z, \mathcal{E}, \theta, R)$ together with factor maps $\pi_X : Z \to X$ and $\pi_Y : Z \to Y$. This gives a joining in the previous sense by taking the image of the measure θ under $z \mapsto (\pi_X(z), \pi_Y(z))$.

Joinings always exist since we could take $\theta = \mu \times \nu$ with the coordinate projection. This is called the trivial joining.

Another example arises when $\varphi : X \to Y$ is an isomorphism. Then the graph map $g : x \mapsto (x, \varphi x)$ pushes μ forwards to a measure on $X \times Y$ that is invariant under $T \times S$ and projects in a 1-1 manner under the coordinate maps to μ, ν respectively. In particular this joining is different from the product joining unless X consists of a single point.

We saw that when X, Y are isomorphic there are non-trivial joinings. The following can be viewed, then, as an extreme form of non-isomorphism.

Definition 7.1.3. X, Y are disjoint if the only joining is the trivial one. Then we write $X \perp Y$.

If X is not a single point then it cannot be disjoint from itself. Indeed, the graph of T is a non-trivial joining. Also if X, Y are ergodic but $X \times Y$ is not the ergodic components of θ_z of $\mu \times \nu$ must project to μ, ν under the coordinate maps, $\pi_X \theta_z$ is T-invariant and $\int \theta_z = \theta$ implies $\int \pi_X \theta_z = \mu$, and by ergodicity of μ this implies $\pi \theta_z = \mu$ for a.e. z; similarly for $\pi_Y \theta$. Thus if the ergodic decomposition is non-trivial then the ergodic components are non-trivial joinings.

Example 7.1.4. Let $X = \mathbb{Z}/m\mathbb{Z}, Y = \mathbb{Z}/n\mathbb{Z}$ with the maps $i \mapsto i+1$. If $\pi : X \to Y$ is a factor between these systems (i.e. preserves measure) then, since it is measure preserving, the fibers $\pi^{-1}(i)$ all have the same cardinality, hence $n|m$. Conversely if $m|n$ there exists a factor map, given by $x \mapsto x \bmod m$.

It is now a simple algebraic fact that if θ joining of X, Y, then it consists of a coset in $(\mathbb{Z}/m\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$ generated by $(i,j) \mapsto (i,j) + (1,1)$. Thus the systems are disjoint if and only if $(\mathbb{Z}/m\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$ is a cyclic group of order mn generated by $(1,1)$, which is the same as saying that mn is the least common multiple of m and n. This is equivalent to $\gcd(m,n) = 1$.

Since any common divisor of m, n gives rise to a common factor of the systems, in this setting disjointness is equivalent to the absence of common nontrivial factors. It also shows that if X, Y are disjoint, then any factor of X is disjoint from any factor of Y (since if $\gcd(m,n) = 1$ then $\gcd(m', n') = 1$ for any $m'|m$ and$n'|n$.

These phenomena hold to some extent in the general ergodic setting.

Proposition 7.1.5. *Let X, Y be invertible systems on standard Borel spaces (neither assumption is not really necessary). Suppose that $\pi_X X \to W$ and $\pi_Y Y \to W$ are factors and $(W, \mathcal{E}, \theta, R)$ is non-trivial. Then $X \not\perp Y$.*

Proof. Let $\mu = \int \mu_z d\theta(x)$ and $\nu = \int \nu_z d\theta(z)$ denote the decomposition of μ over the factor W. Define a measure τ on $X \times Y$ by

$$\tau = \int \mu_z \times \nu_z d\theta(z)$$

We claim that τ is a non-trivial joining of X and Y.

First, it is invariant:

$$(T \times S)\tau = \int T \times S(\mu_z \times \nu_z) d\theta(z) = \int T\mu_z \times S\nu_z d\theta(z) = \int \mu_{Rz} \times \nu_{Rz} d\theta(z) = \int \mu_z \times \nu_z d\theta(z) = \tau$$

Second, it is distinct from $\mu \times \nu$. Indeed, let $E \subseteq W$ be a non-trivial set and $E' = \pi_X^{-1} E \times \pi_Y^{-1} E$. Then

$$\tau(E') = \int \mu_z(\pi_X^{-1} E)\nu_z(\pi_Y^{-1} E) d\theta(z) = \int 1_E(z) 1_E(z) d\theta(z) = \theta(E)$$

On the other hand

$$(\mu \times \nu)(E') = \mu(\pi_X^{-1} E) \cdot \nu(\pi_Y^{-1}(E)) = \theta(E)^2$$

since $\theta(E) \equiv 0, 1$ we have $\theta(E) \neq \theta(E)^2$ so $\tau \neq \mu \times \nu$. $\square$

The converse is false in the ergodic setting: there exist systems with no common factor but non-trivial joinings.

Proposition 7.1.6. *If $X_1 \perp X_2$ and $X_1 \to Y_1$ and $X_2 \to Y_2$, then $Y_1 \perp Y_2$.*

Proof. Let us suppose $X_1 \to Y_1$ and $X_2 \to Y_2$, and $\theta \in \mathcal{P}(Y_1 \times Y_2)$ is a non-trivial joining. Let $\mu_1 = \int \mu_{1,y} d\nu_1(y)$ and $\mu_2 = \int \mu_2, y \mathrm{d}\nu_2(y)$ be the disintegrations. Define a measure τ on $X_1 \times X_2$ by

$$\tau = \int \mu_{1,y_!} \times \mu_{2,y_2} d\theta(y_1, y_2)$$

One can check in a similar way to the previous proposition that τ is invariant and projects to μ_1, μ_2 under the coordinate projections. It is not the product measure because under the maps $X_1 \times X_2 \to Y_1 \times Y_2$ given by applying the original factor maps to each coordinate, the image measure is θ which the image of $\mu_1 \times \mu_2$ is $\nu_1 \times \nu_2$. □

Proposition 7.1.7. *Let (X, Y) and (Y, S) be uniquely ergodic topological systems. and μ, ν invariant measures respectively. Suppose $\mu \perp \nu$.*

1. *For any generic point x for μ and y for ν, the point $(x, y) \in X \times Y$ is generic for $\mu \times \nu$.*
2. *If X, Y are uniquely ergodic then so is $X \times Y$.*

Proof. Clearly any accumulation point of $\frac{1}{N}\sum_{n=0}^{N-1} 1_{(T^n x, S^n y)}$ projects under coordinate maps to μ, ν, hence is a joining. Since the only joining is $\mu \times \nu$ there is only one accumulation point, sot he averages converge to $\mu \times \nu$, which is the same as saying that (x, y) is generic for $\mu \times \nu$.

If X, Y are uniquely ergodic then every invariant measure on $X \times Y$ is a joining (since it projects to invariant measures under the coordinate maps, and these are necessarily μ, ν). Therefore by disjointness the only invariant measure on $X \times Y$ is $\mu \times \nu$. □

Remark 7.1.8. In the ergodic setting, unlike the arithmetic setting, there do exists examples of systems X, Y that are not disjoint, but have no common factor. The examples are highly nontrivial to construct, however.

7.2 Spectrum, disjointness and the Wiener-Wintner theorem

Theorem 7.2.1. *If X is weak mixing and Y has discrete spectrum then $X \perp Y$.*

In fact this is a result of the following more general fact.

Theorem 7.2.2. *Let X, Y be ergodic m.p.s. and Y discrete spectrum. If $\Sigma(X) \cap \Sigma(Y) = \{1\}$ then they are disjoint.*

Proof. We may assume $Y = G$ is a compact abelian group with Haar measure m and $S = L_a$ a translation.

Suppose that θ is a joining. Define the maps $\widetilde{L}_g : X \times G \to X \times G$ by

$$\widetilde{L}_g(x, h) = (x, L_g h) = (x, gh)$$

and let
$$\theta_g = \widetilde{L}_g \theta$$
These are invariant since L_g and L_a commute, hence $\widetilde{L}_g$ commutes with $T \times L_a$. Also, $\pi_1 \theta_g = \pi_1 \theta = \mu$ and $\pi_2 \theta_g = L_g \pi_2 \theta = L_g m_G = m_G$, so θ_g is a joining.

Let
$$\theta' = \int \theta_g dm_G(g)$$
This is again a joining. Now, for any $F \in L^2(\theta)$,
$$\begin{aligned}
\int F(x,h) d\theta'(x,h) &= \int F(x,h) d\theta_g(x,h) dm_G(g) \\
&= \int \widetilde{L}_g F(x,h) d\theta(x,h) dm_G(g) \\
&= \int F(x,gh) d\theta(x,h) dm_G(g) \\
&= \int F(x,gh) dm_G(g) d\theta(x) \\
&= \int F(x,g) dm_G(g) d\theta(x)
\end{aligned}$$
which means that $\theta' = \mu \times m_G$ is a product measure.

Now, from Theorem **??** we know that $X \times Y$ is ergodic hence θ' is ergodic. But $\theta' = \int \theta_g dm_G(g)$ and the θ_g are invariant. Thus by uniqueness of the ergodic decomposition θ_g are m_G-a.s. equal to θ'. Since they are all images of each other under the $\widetilde{L}_g$ maps, they are all ergodic, in particular $\theta = \theta_e = \theta'$, which is what we wanted to prove. □

Corollary 7.2.3. *Among systems with discrete spectrum, disjointness is equivalent to absence of nontrivial common factors.*

Theorem 7.2.4 (Wiener-Wintner theorem)**.** *Let $(X, \mathcal{B}, \mu, T)$ be an ergodic measure preserving system, with X compact metric and T continuous (this is no restriction assuming the space is standard), and $f \in L^1$. Then for a.e. x the limit $\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \alpha^n f(T^n x)$ exists for every $\alpha \in S^1$.*

Proof. Fix f. If we fix a countable sequence $\alpha_i \in S^1$, then the conclusion for these values is obtained as follows. Given α consider the product system $X \times S^1$ and $T \times R_\alpha$, and the function $g(x,z) = zf(x)$. Then by the ergodic theorem $\lim \frac{1}{N} \sum_{n=0}^{N-1} (T \times R_\alpha)^n g$ exists a.s., hence for μ-a.e. x there is a $w \in S^1$ such that $\lim \frac{1}{N} \sum_{n=0}^{N-1} (T \times R_\alpha)^n g(x,w)$ exists, but this is just $\lim w \frac{1}{N} \sum_{n=0}^{N-1} \alpha^n f(T^n x)$ so the claim follows.

First assume f is continuous. Applying the above to the collection of α such that α or some power α^n are in $\Sigma(T)$, we obtain a set of full measure where the associated averages converge so it is enough to prove there is a set of full measure where the averages converge for the other values of α. For these α consider the system $X \times S^1$ as above. Now (X,T) and (S^1, R_α) are disjoint (since $\Sigma(R_a) = \{\alpha^n\}$ so by choice of α, $\Sigma(T) \cap \Sigma(R_\alpha) = \{1\}$) so the only invariant measure is $\mu \times m$. It follows that if x is a generic point for T then (x,z) is generic for $\mu \times m$ (this is proved similarly to the statement about products of disjoint uniquely ergodic systems). Therefore the ergodic averages of g converge, and hence $\frac{1}{N}\sum_{n=0}^{N-1} \alpha^n f(x)$.

We have proved the following: for continuous f there is a set of full measure of x such that $\frac{1}{N}\sum_{n=0}^{N-1} \alpha^n f(x)$ converges for every α. Now if $f \in L^1$ we can approximate f by continuous functions, and note that the limsup of the difference of the averages in question is bounded by $\|f - f'\|_1$ for every α and a.s. x. This gives, for fixed f, a set of measure 1 where the desired limit converges for every α. $\square$

7.3 Shannon entropy: a quick introduction

Entropy originated as a measure of randomness. It was introduced in information theory by Shannon in 1948. Later the ideas turned out relevant in ergodic theory and were adapted by Kolmogorov, with an important modification due to Sinai. We will quickly cover both of these topics now.

We begin with the non-dynamical setting. Suppose we are given a discrete random variable. How random is it? Clearly a variable that is uniformly distributed on 3 points is more random than one that is uniformly distributed on 2 point, and the latter is more random than a non-uniform measure on 2 points. It turns out that the way to quantify this randomness is through entropy. Given a random variable X on a probability space $(\Omega, \mathcal{F}, P)$, and values in a countable set A, then its entropy $H(X)$ is defined by

$$H(X) = -\sum_{a\in A} P(X=a)\log P(X=a)$$

with the convention that log is in base 2 and $0\log 0 = 0$. This quantity may in general be infinite.

Evidently this definition is not affected by the actual values in A, and depends only on the probabilities of each value, $P(X=a)$. We shall write$dist(X)$ for the distribution of X which is the probability vector

$$dist(X) = (P(X=a))_{a\in range(X)}$$

The entropy is then a function of $dist(X)$; fixing the size k of the range, the function is

$$H(t_1,\ldots,t_k) = -\sum t_i \log t_i$$

Again we define $0\log 0 = 0$, which makes $t\log t$ continuous on $[0,1]$.

Lemma 7.3.1. *$H(\cdot)$ is strictly concave.*

Remark if p,q are distributions on Ω then $tp+(1-t)q$ is also a distribution on Ω for all $0 \le t \le 1$.

Proof. Let $f(t) = -t \log t$. Then

$$\begin{aligned} f'(t) &= -\log t - 1 \\ f''(t) &= -\frac{1}{t} \end{aligned}$$

so f is strictly concave on $(0, \infty)$. Now

$$H(tp+(1-t)q) = \sum_i f(tp_i+(1-t)q_i) \geq \sum_i (tf(p_i) + (1-t)f(q_i)) = tH(p)+(1-t)H(q)$$

with equality if and only if $p_i = q_i$ for all i. □

Lemma 7.3.2. *If X takes on k values with positive probability then $0 \leq H(X) \leq \log k$. the left is equality if and only if $k = 1$ and the right is equality if and only if $dist(X)$ is uniform, i.e. $p(X = x) = 1/k$ for each of the values.*

Proof. The inequality $H(X) \geq 0$ is trivial. For the second, note that since H is strictly concave on the convex set of probability vectors it has a unique maximum. By symmetry this must be $(1/k, \ldots, 1/k)$. □

Definition 7.3.3. The **joint entropy** of random variables X, Y is $H(X, Y) = H(Z)$ where $Z = (X, Y)$.

The **conditional entropy** of X given that another random variable Y (defined on the same probability space as X) takes on the value y is is an entropy associated to the conditional distribution of X given $Y = y$, i.e.

$$H(X|Y = y) = H(dist(X|Y = y)) = -\sum_x p(X = x|Y = y) \log p(X = x|Y = y)$$

The **conditional entropy** of X given Y is the average of these over y,

$$H(X|Y) = \sum_y p(Y = y) \cdot H(dist(X|Y = y))$$

Example 7.3.4. If X, Y are independent $\frac{1}{2}, \frac{1}{2}$ coin tosses then (X, Y) takes 4 values with probability $1/4$ so $H(X, Y) = \log 4$.

If X, Y are correlated fair coin tosses then (X, Y) is distributed non-uniformly on its 4 values and the entropy $H(X, Y) < \log 4$.

Lemma 7.3.5. .

1. *$H(X, Y) = H(X) + H(Y|X)$.*
2. *$H(X, Y) \geq H(X)$ with equality if and only if Y is a function of X.*
3. *$H(X|Y) \leq H(X)$ with equality if and only if X, Y are independent.*
4. *More generally, $H(X|YZ) \leq H(X|Y)$.*

Proof. Write

$$\begin{aligned}
H(X,Y) &= -\sum_{x,y} p((X,Y)=(x,y))\log p((X,Y)=(x,y)) \\
&= \sum_{y} p(Y=y)\sum_{x} p(X=x|Y=y)\left(-\log\frac{p((X,Y)=(x,y))}{p(Y=y)} - \log p(Y=y)\right) \\
&= -\sum_{y} p(Y=y)\log p(Y=y)\sum_{x} p(X=x|Y=y) \\
&\quad -\sum_{y} p(Y=y)\sum_{x} p(X=x|Y=y)\log p(X=x|Y=y) \\
&= H(Y)+H(X|Y)
\end{aligned}$$

Since $H(Y|X)\geq 0$, the second inequality follows, and it is an equality if and only if $H(X|Y=y)=0$ for all y which Y attains with positive probability. But this occurs if and only if on each event $\{Y=y\}$, the variable X takes one value. This means that X is determined by Y.

The third inequality follows from concavity:

$$\begin{aligned}
H(X|Y) &= \sum_{y} p(Y=y)H(dist(X|Y=y)) \\
&\leq H(\sum_{y} p(Y=y)dist(X|Y=y)) \\
&= H(X)
\end{aligned}$$

and equality if and only if and only if $dist(X|Y=y)$ are all equal to each other and to $dist(X)$, which is the same as independence.

The last inequality follows similarly from concavity. □

It is often convenient to re-formulate entropy for partitions rather then random variables. Given a countable partition $\beta=\{B_i\}$ of Ω let

$$dist(\beta) = (P(B))_{B\in\beta}$$

and define the entropy by

$$H(\beta) = H(dist(\beta)) = -\sum_{B\in\beta} P(B)\log P(B)$$

This is the entropy of the random variable that assigns to $\omega\in\Omega$ the unique set B_i containing ω. This set is denoted $\beta(\omega)$. Conversely if X is a random variable then it induces a partition ofΩ and the entropy of X and of this partition coincide.

Given partitions α,β their join is the partition

$$\alpha\vee\beta = \{A\cap B \,:\, A\in\alpha \,:\, B\in\beta\}$$

This is the partition induced by the pair (X, Y) of random variables $X : \omega \mapsto \alpha(\omega)$ and $Y : \omega \mapsto \beta(\omega)$. If we define the conditional entropy of α on a set B by

$$H(\alpha|B) = -\sum_{A\in\alpha} P(A|B)\log P(A|B)$$

and

$$H(\alpha|\beta) = \sum_{B\in\beta} P(B)\cdot H(\alpha|B)$$

then we have the identities corresponding to Lemma **??**. Specifically, we say that α **refines** β, or $\alpha \prec \beta$, if every $A \in \alpha$ is a subset of some $B \in \beta$. Then we have

Lemma 7.3.6. .

1. $H(\alpha \vee \beta) = H(\alpha) + H(\beta|\alpha)$.
2. $H(\alpha \vee \beta) \geq H(\alpha)$ *with equality if and only if* $\alpha \prec \beta$.
3. $H(\alpha|\beta) \leq H(\alpha)$ *with equality if and only if* α, β *are independent.*

Definition 7.3.7. Given a discrete random variable X and a σ-algebra $\mathcal{B} \subseteq \mathcal{F}$, the conditional entropy $H(X|\mathcal{B})$ is

$$\int H(dist(X|\mathcal{B}(\omega))\,dP(\omega) = \int H_{P_\omega}(X)\,dP(\omega)$$

where $dist(X|\mathcal{B}(\omega))$ is the distribution of X given the atom $\mathcal{B}(\omega)$, and P_ω is the disintegrated measure at ω given $\mathcal{B}$. For a partition α we define $H(\alpha|\mathcal{B})$ similarly.

Proposition 7.3.8. *Suppose α is a finite partition and $\beta_1 \succ \beta_2 \succ \ldots$ a sequence of refining partitions and $\mathcal{B} = \sigma(\beta_1, \beta_2, \ldots) = \bigvee \beta_n$ the generated σ-algebra. Then*

$$H(\alpha|\mathcal{B}) = \lim_{n\to\infty} H(X|\beta_n)$$

Equivalently, if X is a finite-valued random variable and Y_n are discrete random variables then

$$H(X|Y_1, Y_2, \ldots) = \lim_{n\to\infty} H(X|Y_1 \ldots Y_n)$$

Remark 7.3.9. The same is true when α is a countable partition but the proof is slightly longer and we won't need this more general version.

Proof. By the martingale convergence theorem, for each $A \in \alpha$ we have

$$P(A|\beta_n(\omega)) = \mathbb{E}(1_A|\beta_n)(\omega) \to \mathbb{E}(1_A|\mathcal{B})(\omega) = P(A|\mathcal{B}(\omega)) \qquad \text{a.e.}$$

This just says that $dist(\alpha|\beta_n) \to dist(\alpha|\mathcal{B})$ pointwise as probability vectors. Since $H(t_1, \ldots, t_{|\alpha|})$ is continuous this implies that $H(dist(\alpha|\beta_n)) \to H(dist(\alpha|\mathcal{B})$, which is the desired conclusion. □

Remark 7.3.10. There is a beautiful axiomatic description of entropy as the only continuous functional of random variables satisfying the conditional entropy formula. Suppose that $H_m(t_1,\ldots,t_m)$ are functions on the space of m-dimensional probability vectors, satisfying

1. $H_2(\cdot,\cdot)$ is continuous,
2. $H_2(\frac{1}{2},\frac{1}{2}) = 1$,
3. $H_{k+m}(p_1,p_2,\ldots p_k,q_1,\ldots,q_m) = (\sum p_i)H_k(p'_1,\ldots,p'_k)+(\sum q_i)H_m(q'_1,\ldots,q'_k)$, where $p'_i = p_i/\sum p_i$ and $q'_i = q_i/\sum q_i$.

Then $H_m(t) = -\sum t_i \log t_i$. We leave the proof as an exercise.

7.4 Digression: applications of entropy

We describe here two applications of entropy. The first demonstrates how entropy can serve as a useful analog of cardinality, but with better analytical properties. The basic connection is that if X is uniform on its range Σ, then $H(X) = \log|\Sigma|$.

Proposition 7.4.1 (Loomis-Whitney Theorem)**.** *Let $A \subseteq \mathbb{R}^3$ be a finite set and $\pi_{xy},\pi_{xz},\pi_{yz}$ the projections to the coordinate planes. Then the image of A under one of the projections is of size at least $|A|^{2/3}$.*

If we make the same statement in 2 dimensions, it is trivial, since if $|\pi_x(A)| \leq \sqrt{|A|}$ and $|\pi_y(A)| < \sqrt{|A|}$ then $A \subseteq \pi_x(A)\times\pi_y(A)$ has cardinality $< |A|$, which is impossible. This argument does not work in three dimensions.

Lemma 7.4.2 (Shearer's inequality)**.** *If X,Y,Z are random variables then*

$$H(X,Y,Z) \leq \frac{1}{2}\left(H(X,Y) + H(X,Z) + H(Y,Z)\right)$$

Proof. Write

$$\begin{array}{ccccccc}
H(X,Y,Z) & = & H(X) & + & H(Y|X) & + & H(Z|X,Y) \\
H(X,Y) & = & H(X) & + & H(Y|X) & & \\
H(Y,Z) & = & & & H(Y) & + & H(Z|Y) \\
H(X,Z) & = & H(X) & & & + & H(Z,X)
\end{array}$$

Using $H(Y) \geq H(Y|X)$ and $H(Z|X) \geq H(Z|X,Y)$ and $H(Z|Y) \geq H(Z|X,Y)$, the sum of last 3 lines is at least equal to the first line. This was the claim. □

Proof of Loomis-Whitney. Let $P = (X,Y,Z)$ be a random variable uniformly distributed over A. Thus

$$H(P) = H(X,Y,Z) = \log|A|$$

On the other hand by Shearer's inequality,

$$H(P) \leq \frac{1}{2}\left(H(X,Y) + H(X,Z) + H(Y,Z)\right) = \frac{1}{2}\left(H(\pi_{xy}(P)) + H(\pi_{xz}(P)) + H(\pi_{zy}(P))\right)$$

so at least one $H(\pi_{xy}(P)), H(\pi_{xz}(P)), H(\pi_{yz}(P))$ is $\geq \frac{2}{3}\log|A|$. Since $\pi(P)$ is supported on $\pi(A)$ we also have $H(\pi_{x,z}(P)) \leq \log|\operatorname{supp}\pi_{x,y}A|$, etc. In conclusion for one of the projections, $\pi(|A|) \geq \frac{2}{3}\log|A|$, which is what we claimed. □

We now turn to the original use of Shannon entropy in information theory, where $H(X)$ should be thought of as a quantitative measure of the amount of randomness tin the variable X, Suppose we want to record the value of X using a string of 0s and 1s. Such an association $c : \Sigma \rightarrow \{0,1\}^*$ is called a code. We shall require that the code be 1-1, and for simplicity we require it to be a **prefix code**, which means that if $a, b \in \Sigma$ then neither of a, b is a prefix of the other. Let $|c(a)|$ denote the length of the codeword $c(a)$.

Lemma 7.4.3. *If $c : \Sigma \rightarrow \{0,1\}^*$ is a prefix code then $\sum_{a\in\Sigma} 2^{-|c(a)|} \leq 1$. Conversely, if $\ell : \Sigma \to \mathbb{N}$ are given and $\sum_{a\in\Sigma} 2^{-\ell(a)}$ then there is a prefix code $c : \Sigma \to \{0,1\}^*$ with $|c(a)| = \ell(a)$.*

Theorem 7.4.4. *Let $\{\ell_i\}_{i\in\Sigma}$ be integers. Then the following are equivalent:*

1. $\sum 2^{-\ell_i} \leq 1$.
2. There is a prefix code with lengths ℓ_i.

Proof. We can assume $\Sigma = \{1, \ldots, n\}$. Let $L = \max \ell_i$ and order $\ell_1 \leq \ell_2 \leq \ldots \leq \ell_n = L$. Identify $\bigcup_{i\leq L}\{0,1\}^i$ with the full binary tree of height L, so each vertex has two children, one connected to the vertex by an edge marked 0 and the other by an edge marked 1. Each vertex is identified with the labels from the root to the vertex; the root corresponds to the empty word and the leaves (at distance L from the root) correspond to words of length L.

$2 \implies 1$: The prefix condition means that $c(i), c(j)$ are not prefixes of each other if $i \neq j$; consequently no leaf of the tree has both as an ancestor. Writing A_i for the leaves descended from i, we have $|A_i| = 2^{L-\ell_i}$ and the sets are disjoint, therefore

$$\sum 2^{L-\ell_i} = \sum |A_i| = |\bigcup A_i| \leq |\{0,1\}^L| = 2^L$$

dividing by 2^L gives (1) .

Conversely a greedy procedure allows us to construct a prefix code for ℓ_i as above. The point is that if we have defined a prefix code on $\{1, \ldots, k-1\}$ then the set of leaves below the codewords must have size

$$|\bigcup_{i<k} A_i| \leq \sum_{i<k} |A_i| = \sum_{i<k} 2^{L-\ell_i} < 2^L$$

The strict inequality is because $\sum 2^{-\ell_i} \leq 1$, and the sum above includes at least one term less than the full sum. Therefore we can choose a codeword for k that is the ancestor of a leaf not in $\bigcup_{i<k} A_i$. This extends the prefix code to k and we continue until all codewords are defined. $\square$

Now, given a code $c(\cdot)$, the average coding length (w.r.t. X) is $\mathbb{E}(|c(X)|)$. We wish to find the prefix code that minimizes this. quantity. This is equivalent to optimizing $\mathbb{E}(\ell(X))$ over all functions $\ell : \Sigma \to \mathbb{N}$ satisfying $\sum 2^{-\ell_i} \leq 1$.

Theorem 7.4.5. *If c is a prefix code, then the expected coding length is $\geq H(X)$ and equality is achieved if and only if $|c(i)| = -\log P(X = i)$.*

Proof. Let ℓ_i be the coding length of i and $p_i = P(X = i)$. We know that $\sum 2^{-\ell_i} \leq 1$. Consider

$$\begin{aligned} \Delta &= H(p) - \sum p_i \ell_i \\ &= -\sum p_i (\log p_i + \ell_i) \end{aligned}$$

Let $r_i = 2^{-\ell_i} / \sum 2^{-\ell_i}$, so $\sum r_i = 1$ and $\ell_i \geq -\log r_i$ (because $\sum 2^{-\ell_i} \leq 1$).

$$\begin{aligned} \Delta &\leq -\sum p_i (\log p_i - \log r_i) \\ &= -\sum p_i \left(\log \frac{p_i}{r_i}\right) \\ &= \sum p_i \left(\log \frac{r_i}{p_i}\right) \end{aligned}$$

Using the concavity of the logarithm,

$$\leq \log\left(\sum p_i (\frac{r_i}{p_i})\right) = \log 1 = 0$$

Equality occurs unless $r_i / p_i = 1$. $\square$

Theorem 7.4.6 (Achieving optimal coding length (almost))**.** *There is a prefix code whose average coding length is $H(X) + 1$*

Proof. Set $\ell_i = \lceil -\log p_i \rceil$. Then

$$\sum 2^{-\ell_i} \leq \sum 2^{-\log p_i} = \sum p_i = 1$$

and since $\ell_i \leq -\log p_i + 1$, the expected coding length is

$$\sum p_i \ell_i \leq H(X) + 1$$

Thus **up to one extra bit we achieved the optimal coding length**. $\square$

Corollary 7.4.7. *If (X_n) is a stationary finite-valued process with entropy $h(X) = \lim \frac{1}{n} H(X_1 \ldots X_n)$, then for every $\varepsilon > 0$ if n is large enough we can code $X_1 \ldots X_n$ using $h + \varepsilon$ bits per symbol (average coding length $\leq (h + \varepsilon) n$).*

Proof. We can code $X_1, \ldots, X_n$ using $H(X_1, \ldots, X_n) + 1$ bits on average. Dividing by n and assuming $n > 2/\varepsilon$, and also that $\frac{1}{n}H(X_1 \ldots X_n) < h + \varepsilon/2$, this gives a per-bit coding rate of

$$\frac{1}{n}(H(X_1, \ldots, X_n) + 1) \leq h + \varepsilon/2 + \varepsilon/2 = h + \varepsilon$$

□

7.5 Entropy of a stationary process

$X = (X_n)_{n=-\infty}^{\infty}$ be a stationary sequence of random variables with values in a finite set Σ; as we know, such a sequence may be identified with a shift-invariant measure μ on $\Sigma^{\mathbb{Z}}$, with $X_n(\omega) = \omega_n$ for $\omega = (\omega_n) \in \Sigma^{\mathbb{Z}}$, and $X_n(\omega) = X_0(S^n\omega)$, where S is the shift. The partition induced by X_0 on $\Sigma^{\mathbb{Z}}$ is the partition according tot he 0-th coordinate, and we denote α. Then X_n induces the partition $T^n\alpha = \{T^{-n}A \,:\, A \in \alpha\}$. Since $(X_n(\omega))_{n=-\infty}^{\infty}$ determines ω, we see that the partitions $T^n\alpha$, $n \in \mathbb{Z}$, separate points, so $\mathcal{B} = \bigvee_{n=-\infty}^{\infty} T^n\alpha$.

Definition 7.5.1. The entropy $h(X)$ of the process $X = (X_n)$ is

$$\lim_{n\to\infty} \frac{1}{n} H(X_1, \ldots, X_n)$$

Let us first show that the limit exists: Let $H_n(X) = H(X_1, \ldots \ldots X_n)$. For each n, m,

$$\begin{aligned}
H_{m+n}(X) &= H(X_1 \ldots X_m, X_{m+1}, \ldots, X_{n+m}) \\
&= H(X_1, \ldots, X_m) + H(X_{m+1}, \ldots, X_{m+n} | X_1, \ldots, X_m) \\
&\leq H(X_1, \ldots, X_m) + H(X_{m+1}, \ldots, X_{m+n}) \\
&= H(X_1, \ldots, X_m) + H(X_1, \ldots, X_n) \\
&= H_m(X) + H_n(X)
\end{aligned}$$

Here the inequality is because conditioning cannot increase entropy, and then we used stationarity, which implies $dist(X_{m+1}, \ldots, X_{m+n}) = dist(X_1, \ldots, X_n)$. e have shown that the sequence $H_n(X)$ is subadditive so the limit $h(X) = \lim \frac{1}{n} H_n(X)$ exists (and equals $\inf \frac{1}{n} H_n(X)$).

Example 7.5.2. If X_n are i.i.d. then $H(X_1, \ldots, X_n) = \sum H(X_i) = nH(X_1)$ and so $h(X) = H(X_1)$.

Example 7.5.3. Let μ be the S-invariant measure on a periodic point $\omega = S^N\omega$. Then (X_n) can be obtained by choosing a random shift of the sequence ω. Since $X_1, \ldots, X_n$ takes on at most N different values, $H(X_1, \ldots, X_N) \leq \log N$ and so $\frac{1}{n} H(X_1 \ldots X_N) \to 0 = h(X)$.

Example 7.5.4. Let $\theta \in \mathbb{R} \neq \mathbb{Q}$ and $R_\theta : \mathbb{R}/\mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ translation, for which Lebesgue measure μ is the unique invariant measure. Let $X_1(x) = 1_{[0,1/2]}(x)$ and $X_n = X_1(R_\theta^n x)$. We claim that the process $X = (X_n)$ has entropy 0.

Indeed, the partition determined by X_n is an interval of length $1/2$ in $\mathbb{R}/\mathbb{Z}$ and so $(X_1, \ldots, X_n)$ determines a partition that is the join of n intervals. This partition is the partition into intervals that is determined by $2n$ endpoints and so consists of $2n$ intervals. Hence $H(X_1, \ldots, X_n) \leq \log 2n$ and so $\frac{1}{n} h(X_1 \ldots X_n) \to 0 = h(X)$.

Theorem 7.5.5. $h(X) = H(X_0 | X_{-1}, X_{-2}, \ldots)$.

Proof. It is convenient to write $X_m^n = X_m, \ldots, X_n$. Now,

$$\begin{aligned} H(X_1^n) &= H(X_1 \ldots X_{n-1}) + H(X_n | X_1, \ldots, X_{n-1}) \\ &= H(X_1 \ldots X_{n-1}) + H(X_0 | X_{-n}, \ldots, X_{-1}) \end{aligned}$$

by stationarity. By induction,

$$H(X_1^n) = \sum_{i=0}^{n-1} H(X_0 | X_{-i}^{-1})$$

so

$$\begin{aligned} h(X) &= \lim_{n\to\infty} \frac{1}{n} H(X_1, \ldots, X_n) \\ &= \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} H(X_0 | X_{-i}^{-1}) \end{aligned}$$

Since $H(X_0 | X_{-\infty}^{-1}) = \lim_{n\to\infty} H(X_0 | X_{-n}^{-1})$, the summands in the averages above converge, so the averages converge to the same limit, which is what we claimed. □

Definition 7.5.6. A process $(X_n)_{n=-\infty}^{\infty}$ is **deterministic** if $X_{-\infty}^{-1}$ a.s. determines X_0 (and by induction also $X_1, X_2, \ldots$).

Corollary 7.5.7. *(X_n) is deterministic if and only if $h(X) = 0$.*

Proof. $H(X_0 | X_{-\infty}^{-1}) = 0$ if and only if a.e. the conditional distribution of X_0 given $X_{-\infty}^{-1}$ is atomic, which is the same as saying that X_0 is measurable with respect to $\sigma(X_{-\infty}^{-1})$, which is the same as determinism. □

Remark 7.5.8. Note that the entropy of the time-reversal $(X_{-n})_{n=-\infty}^{\infty}$ is the same as of (X_n), because the entropy of the initial n variables is the same in both cases. It follows that a process and its time-reversal are either both deterministic (if the entropy is 0) or neither is. This conclusion is nonobvious!

Definition 7.5.9. For a stationary process $(X_n)_{n=-\infty}^{\infty}$, the **tail σ-algebra** is the σ-algebra $\mathcal{T} = \bigcap_{n=1}^{\infty} \sigma(X_{-\infty}^{-n})$. The process has **trivial tail**, or is a **Kolmogorov process**, if $\mathcal{T}$ is the trivial algebra (modulo nullsets).

Intuitively, a process with trivial tail is a process in which the remote past has no effect on the present (and future).

Example 7.5.10. If (X_n) are i.i.d. then $\sigma(X_{-\infty}^{-n})$ is independent of X_{-n+1}^{∞}. so $\mathcal{T}$ is independent of all the variables X_n and so, since it is measurable with respect tot hem , is independent of itself; so it is trivial.

Proposition 7.5.11. *If $X = (X_n)$ has trivial tail and $H(X_0) > 0$ then $h(X) >$ 0.*

Proof. Since $H(X_0|X_{-\infty}^{-n}) \to H(X_0|\mathcal{T}) = H(X_0)$. Therefore there is some n_0 such that such that $H(X_0|X_{-\infty}^{-n_0}) > 0$. Now,

$$\begin{aligned} H(X_0, \ldots, X_n) &\geq H(X_0, X_{n_0}, X_{2n_0}, \ldots, X_{[n/n_0]n_0}) \\ &= \sum_{i=0}^{[n/n_0]-1} H(X_{in_0}|X_0, X_{n_0}, \ldots, X_{(i-1)n_0}) \\ &\geq \left[\frac{n}{n_0}\right] H(X_0|X_{-\infty}^{-n_0}) \end{aligned}$$

so

$$h(X) = \lim_{n\to\infty} \frac{1}{n} H(X_0, \ldots, X_n) \geq \frac{1}{n_0} H(X_0|X_{-\infty}^{-n_0}) > 0$$

as claimed. □

Corollary 7.5.12. *The only process that is both deterministic and has trivial tail is the trivial process $X_n = a$ for all n.*

7.6 Couplings, joinings and disjointness of processes

Let $(X_n), (Y_n)$ be stationary processes taking values in A, B respectively. A **coupling** is a stationary process (Z_n) with $Z_n = (X'_n, Y'_n)$ and the marginals (X'_n), (Y'_n) have the same distribution as $(X_n), (Y_n)$, respectively. This is the probabilistic analog of a joining. Indeed, identify $(X_n), (Y_n)$ with invariant measures μ, ν on $A^{\mathbb{Z}}, B^{\mathbb{Z}}$,respectively. The coupling (Z_n) can be identified with an invariant measure θ on $(A \times B)^{\mathbb{Z}}$, and the assumption on the distribution of (X'_n), (Y'_n) is the same as saying that the projections $(A \times B)^{\mathbb{Z}} \to A^{\mathbb{Z}}$ maps $\theta \mapsto \mu$ and the other projection maps $\theta \mapsto \nu$; thus θ is a joining of the systems $(A^{\mathbb{Z}}, \mu, S)$ and $(B^{\mathbb{Z}}, \nu, S)$. On the other hand any joining gives an associated coupling, $Z_n = \pi_n(z)$ where $\pi_n : (A \times B)^{\mathbb{Z}} \to A \times B$ is the projection to the n-th coordinate.

Definition 7.6.1. We say that two processes are disjoint if the only coupling is the independent one; equivalently, the associated shift-invariant measures are disjoint.

Proposition 7.6.2. *If $(X_n), (Y_n)$ have positive entropy, they are not disjoint.*

Proof. We prove this for $0,1$-valued processes, the general proof is the same.

If U, V are random variables taking on values $0, 1$ with positive probability, then there is a non-trivial coupling that can be constructed as follows: let $p = P(U = 1)$ and $q = P(V = 1)$. Let $Z \sim U[0,1]$ and set $U' = 1_{\{Z<p\}}$ and $V' = 1_{\{Z<q\}}$. Clearly $P(U' = 1) = p$ and $P(V' = 1) = q$, but U', V' are not independent because, assuming w.l.o.g. that $p < q$, if $U' = 1$ then $Z' < p < q$ so $V' = 1$. This shows that

$$H(V'|U') = 0$$

Similarly if $q < p$ then

$$H(V'|U') = 0$$

Since $H(U', V') = H(U') + H(V'|U') = H(V') + H(U'|V')$, we see that in any case,

$$H(U', V') \le \max\{H(U), H(V)\}$$

We construct a stochastic process $(U_n, V_n)_{-\infty}^{\infty}$ with $U_{-\infty}^{\infty} \sim X_{-\infty}^{\infty}$ and $V_{-\infty}^{\infty} \sim Y_{-\infty}^{\infty}$ as follows. Choose $U_{-\infty}^{-1} \sim X_{-\infty}^{-1}$ and $V_{-\infty}^{-1} \sim Y_{-\infty}^{-1}$ independently. Now assume we have defined a distribution on sequences $U_{-\infty}^{-n}, V_{-\infty}^{-n}$ for some $n \ge 0$ such that $U_{-\infty}^{n} \sim X_{-\infty}^{n}$ and $V_{-\infty}^{n} \sim Y_{-\infty}^{n}$, we almost surely have a conditional probability $p = P(X_{n+1} = 1 | X_{-\infty}^{n} = U_{-\infty}^{n})$ and $q = P(Y_{n+1} = 1 | Y_{-\infty}^{n} = V_{-\infty}^{n})$. For each realization of $U_{-\infty}^{n}, V_{-\infty}^{n}$ in a set of probability 1, this allows us to define U_{n+1}, V_{n+1} as above, using an independent auxiliary random variable $Z_{n+1} \sim U[0,1]$.

Let θ denote the measure associated to the sequence $(W_n) = (U_n, V_n)$. By construction θ is not a product measure and projects to μ, ν (the measures associated to $(X_n), (Y_n)$) on the marginals, but it is not invariant. To fix this we can pass to a limit

$$\theta' = \lim \frac{1}{N_k} \sum_{n=0}^{N_k} S^n \theta$$

This measure is invariant, and still projects to μ, ν on the marginals (because the set of measures with this property is closed). It remains to show that θ' it not the product measure. Suppose, by way of contradiction, that it is a product measure. Let W'_n be the associated process; we obtain a contradiction by showing that $h(W) < h(X) + h(Y)$. To see this, note that

$$h(W') = h(W'_0 | (W')_{-\infty}^{-1}) = \lim_{T \to \infty} H(W'_0 | (W')_{-T}^{-1})$$

so it suffices to prove that

$$H(W'_0 | (W')_{-T}^{-1}) < h(X) + h(Y)$$

for some T. Now, for a given T we have

$$H(W'_0 | (W')_{-T}^{-1}) = \lim_{k \to \infty} \frac{1}{n} \sum_{i=N_k}^{n} H(W_i | W_{i-T}^{i-1})$$

Now, for $i \geq 1$ note that $H(W_i|W_{i-T}^{i-1})$ is determined in a manner independent of i by the conditional distribution of (U_i, V_i) given $(U_j, V_j)_{j=i-T}^{i-1}$ and the latter is converging to an independent coupling of the marginal distributions, so this limit is

$$H(W_0'|(W')_{-T}^{-1}) = H(W_1|W_{-T+1}^0)$$

and this $\to H(W_0|W_{-\infty}^{-1})$ as $T \to \infty$. Therefore it is enough to show that

$$H(W_1|W_{-\infty}^0) < h(X) + h(Y)$$

But, by construction, for each realization $w_{-\infty}^0 = (x, y)_{-\infty}^0$ of $W_{-\infty}^0$ we have

$$\begin{aligned} H(W_1|W_{-\infty}^0 = w_{-\infty}^0) &\leq \max\{H(X_1|X_{-\infty}^0 = x_{-\infty}^0), H(Y_1|Y_{-\infty}^0 = y_{-\infty}^0)\} \\ &\leq H(X_1|X_{-\infty}^0 = x_{-\infty}^0) + H(Y_1|Y_{-\infty}^0 = y_{-\infty}^0) \end{aligned}$$

and we have a strict inequality with positive probability, because each term is positive with positive probability, and the pasts are independent. Therefore

$$\begin{aligned} H(W_1|W_{-\infty}^0) &= \int H(W_1|W_{-\infty}^0 = w_{-\infty}^0)\, dP(w_{-\infty}^0) \\ &< \int H(X_1|X_{-\infty}^0 = x_{-\infty}^0) + H(Y_1|Y_{-\infty}^0 = y_{-\infty}^0)\, d\mu \times \nu((x, y)_{-\infty}^0) \\ &= h(X) + h(Y) \end{aligned}$$

as desired. □

Theorem 7.6.3. *Let (X_n) be a process with trivial tail and (Y_n) deterministic process, taking values in finite sets A, B respectively. Then $(X_n) \perp (Y_n)$.*

Lemma 7.6.4. *Let $Z = (X_n, Y_n)_{n=-\infty}^{\infty}$ be a stationary process with values in $A \times B$. Then*

$$h(Z) = h(Y) + H(X_0|X_{-\infty}^{-1}, Y_{-\infty}^{\infty})$$

Proof. Expand $H(Z_1, \dots, Z_n) = H(X_1, \dots, X_n Y_1, \dots, Y_n)$ as

$$\begin{aligned} H(Y_1, \dots, Y_n, X_1, \dots, X_n) &= H(Y_n) + H(Y_1, \dots, Y_{n-1}, X_1, \dots X_n|Y_n) \\ &= H(Y_n) + H(Y_{n-1}|Y_n) + H(Y_1, \dots, Y_{n-2}, X_1, \dots X_n|Y_{n-1}Y_n) \\ &\vdots \\ &= \sum_{i=0}^{n-1} H(Y_{n-i}|Y_1^{n-i-1}) + H(X_1 \dots X_n|Y_1^n) \\ &\vdots \\ &= \sum_{i=0}^{n-1} H(Y_{n-i}|Y_1^{n-i-1}) + \sum_{i=0}^{n-1} H(X_{n-i}|X_1^{n-i-1}, Y_1^n) \end{aligned}$$

Dividing by n, the first term converges to $h(Y)$. The second term, after shifting indices, is an average of terms of the form $H(X_1|X_{-k}^{-1}Y_{-k}^m)$ where $k \to -\infty$ and $m \to \infty$, and these tend uniformly to $H(X_0|X_{-\infty}^{-1}, Y_{-\infty}^{\infty})$, so the average does too. □

Lemma 7.6.5. *Let* $Y_n^M = (Y_{n-M}, \ldots\ldots Y_{n+M})$*. Then* $h(Y^M) = h(Y)$*.*

Lemma 7.6.6. $H(X_0|X_{-M}, X_{-2M}, \ldots) \to H(X_0)$ *as* $M \to \infty$*.*

Of the theorem. Suppose $Z_n = (X_n, Y_n)$ is a coupling. Let us first show that X_0 is independent of $\sigma(Y_{-\infty}^{\infty})$. If not then there is some M and $\varepsilon > 0$ such that $H(X_0|Y_{-M}^M) < H(X_0) - \varepsilon$. Consider $Y_n^M = (Y_{n-M}, \ldots\ldots Y_{n+M})$ and $Z_n^M = X_n, (Y_{n-M}, \ldots, Y_{n+M})$. Then Z_n is stationary. Consider the stationary process (Z_{Mn}^M). Then

$$h(Z^M) \geq h((X_{Mn})_{n=1}^{\infty}) \geq h(X_0) - \varepsilon$$

assuming M is large enough. But applying the previous lemma,

$$h(Z^M) = h(Y_n^M) + H(X_0|(X_{kn})_{k=-\infty}^{-1}, (\overline{Y}_n^M)_{-\infty}^{-1}) \leq 0 + H(X_0|Y_0^M) = H(X_0|Y_{-M}^M) < H(X_0) - \varepsilon$$

a contradiction.

To show that the coupling is independent, repeat this argument with (X^M, Y) in place of (X, Y), noting that X^M also has trivial tail. □

7.7 Kolmogorov-Sinai entropy

Let $(X, \mathcal{B}, \mu, T)$ be an invertible measure-preserving system. For a partition $\alpha = \{A_i\}$ of X we define

$$h_\mu(T, \alpha) = \lim_{n\to\infty} \frac{1}{n} H_\mu(\bigvee_{i=1}^{n} T^{-i}\alpha)$$

This is the same as the entropy of the process $(X_n^\alpha)_{n=-\infty}^{\infty}$, where

$$X_n^\alpha(x) = i \qquad \Longleftrightarrow \qquad T^n x \in A_i$$

Definition 7.7.1. The Kolmogov-Sinai entropy $h_\mu(T)$ of the system is

$$h_\mu(T) = \sup_\alpha h_\mu(T, \alpha)$$

This number is non-negative and may be infinite.

Proposition 7.7.2. $h_\mu(T)$ *is an isomorphism invariant of the system (that is, isomorphic systems have equal entropies).*

Proof. Since $H_\mu(\bigvee_{i=1}^n T^{-i}\alpha)$ depends only on the masses of the atoms of $\bigvee_{i=1}^n T^{-i}\alpha$, and the partition $\beta = \pi\alpha$ of Y has the property that $\pi(\bigvee_{i=1}^n T^{-i}\alpha) = \bigvee_{i=1}^n S^{-i}\beta$, the two have equal entropy. It follows that $h_\mu(T, \alpha) = h_\nu(S, \pi\alpha)$. This shows that the supremum of entropies of partitions of Y is at least as large as the supremum of the entropies of partitions of X. By symmetry, we have equality. □

Definition 7.7.3. A partition α of X is **generating** if $\bigvee_{i-=\infty}^{\infty} T^{-i}\alpha = \mathcal{B} \bmod \mu$.

Example 7.7.4. the partition of $A^{\mathbb{Z}}$ according tot he first coordinate is generating.

In order for entropy to be useful, we need to know how to compute it. This will be possibly using

Theorem 7.7.5. *If α is a generating partition, then $h_\mu(T) = h_\mu(T, \alpha)$.*

Proof. Let α be a generating partition and β another partition. We only need to show that $h_\mu(T, \alpha) \geq h_\mu(T, \beta)$. Consider the process (X_n, Y_n) where X_n is the process determined by α and Y_n the one determined by β (so (X_n, Y_n) is determined by $\alpha \vee \beta$). By definition $h_\mu(T, \alpha) = h((X_n))$ and $h_\mu(T, \beta) = h((Y_n))$. Now, since $H(X_1^n, Y_1^n) \geq H(Y_1^n)$ we have $h((X_n, Y_n)) \geq h((Y_n))$. On the other hand by a previous lemma,

$$h((X_n, Y_n)) = h((X_n) + H(Y_0 | Y_{-\infty}^0, X_{-\infty}^\infty)$$

Since $\sigma(X_{-\infty}^\infty)$ is the full σ-algebra up to nullsets, the conditional expectation on the right is 0. Thus, $h_\mu(T, \alpha) = h((X_n)) \geq h((Y_n)) = h_\mu(T, \beta)$, as desired. $\square$

Example 7.7.6. Let $p = (p_1, \ldots, p_n)$ be a probability vector and $\mu_p = p^{\mathbb{Z}} \in \mathcal{P}(\{1, \ldots, n\}^{\mathbb{Z}})$, $S=$ the shift. Then the partition α according to the 0-coordinate is generating, so $h_{\mu_p}(S) = h((X_n))$ where (X_n) is the process associated to α. This is an i.i.d. process with marginal p so its entropy if $H(p)$. Thus, $\mu_p \cong \mu_q$ implies that $H(p) = H(q)$. In particular, this shows that $(1/2, 1/2)^{\mathbb{Z}} \neq (1/3, 1/3, 1/3)^{\mathbb{Z}}$.

7.8 Application to filterling

See Part I, Section 9 of H. Furstenberg, *Disjointness in ergodic theory, minimal sets, and a problem in diophantine approximation*, 1967, available online at
http://www.kent.edu/math/events/conferences/cbms2011/upload/furst-disjointness.pdf

Chapter 8

Rohlin's lemma

8.1 Rohlin lemma

Theorem 8.1.1 (Rohlin's Lemma)**.** *Let $(X, \mathcal{B}, \mu, T)$ be an invertible measure preserving system, and suppose that for every $\delta > 0$ there is a set A with $\mu(A) < \delta$ and $\mu(X \setminus \bigcup_{n=0}^{\infty} TA) = 0$. Then for every $\varepsilon > 0$ and integer $N \geq 1$ there is a set B such that $B, TB, \ldots, T^{N-1}B$ are pairwise disjoint, and their union has mass $> 1 - \varepsilon$.*

Remark 8.1.2. We will discuss the hypothesis soon but note for now that if μ is non-atomic and T is ergodic then it is automatically satisfied, since in fact for any set B of positive measure, $C = \bigcup_{n=0}^{\infty} T^n B$ satisfies $C \subseteq T^{-1}C$, hence by ergodicity $\mu(C \Delta X) = 0$.

Thus the theorem has the following heuristic implication: any two measure preserving maps behave in an identical manner on an arbitrarily large fraction of the space, on which it acts simply as a "shift"; the differences are "hidden" in the exceptional ε of mass.

Proof. Let ε, B be given and choose A as in the hypothesis for $\delta = \varepsilon/N$. Let $r_A(x) = \min\{n > 0 : T^n x \in A\}$ and A into $A_n = r_A^{-1}(x) \cap A$, that is,

$$A_n = \{x \in A : T^i x \in A \text{ for } 0 \leq i < n \text{ but } T^n x \in A\}$$

Note that

$$\bigcup_{n=0}^{\infty} T^n A = \bigcup_{n} \bigcup_{i=0}^{n-1} T^i A_n$$

Both are X up to measure 0, but the union on the right hand side is disjoint.

Now, fix n and let

$$A'_n = \bigcup_{i=0}^{[n/N]-1} T^{iN} A_n$$

Also set

$$E_n = \bigcup_{i=N[n/N]}^{n-1} T^i A_n$$

Notice that

$$\bigcup_{i=0}^{n-1} T^i A = E_n \cup \bigcup_{i=0}^{N-1} T^i A'_n$$

and the union is disjoint.

Since the above holds for all N, the set

$$B = \bigcup_n A'_n$$

has the property that $B, TB, \ldots, T^{N-1}B$ are pairwise disjoint, and $\mu(X \setminus \bigcup_{i=0}^{N-1} T^i B) = \mu(E)$, where $E = \bigcup E_n$. Finally, to estimate $\mu(E)$, we have

$$\mu(E) = \sum \mu(E_n) \leq \sum N\mu(A_n) = N\mu(A) < N\frac{\varepsilon}{N} = \varepsilon$$

as desired. □

Definition 8.1.3. For a measure-preserving system $(X, \mathcal{B}, \mu, T)$ let

$$Per(T) = \bigcup_{n=1}^{\infty} \{x \in X \,:\, x = T^n x\}$$

The system is *aperiodic* if $\mu(Per(T)) = 0$.

Proposition 8.1.4. *If $(X, \mathcal{B}, \mu, T)$ is aperiodic and $(X, \mathcal{B})$ is standard, then the hypothesis of the previous proposition is satisfied.*

Proof. First, we may assume that $Per(T) = \emptyset$ by replacing X with $X \setminus Per(T)$. Let $\varepsilon > 0$ and consider the class $\mathcal{A}$ of measurable sets A with the property that $\mu(A) \leq \varepsilon\mu(\bigcup_{n=0}^{\infty} T^n A)$. Note that $\mathcal{A}$ is non-empty since it contains $\emptyset$, and it is closed under monotone increasing unions of its members.

Consider the partial order $\prec$ on $\mathcal{A}$ given by $A \prec A'$ if $A \subseteq A'$ and $\mu(A) < \mu(A')$. Any maximal chain is countable, since there are no uncountable bounded increasing sequences of reals. Therefore the chain has a maximal element, namely the union of its members. We shall show that every such maximal element A must satisfy $\mu(\bigcup_{n=0}^{\infty} T^n A) = 1$.

To see this first suppose that $A \in \mathcal{A}$ and $\mu(\bigcup_{n=0}^{\infty} T^n A) < 1$. Set $X' = X \setminus \bigcup_{n=0}^{\infty} T^n A$ which is an invariant set (up to a nullset, which we also remove as necessary). All we must show is that there exists $E \subseteq X'$ with $\mu(E) \leq \varepsilon\mu(\bigcup_{n=0}^{\infty} T^n E)$, then $A \cup E \succ A$. To see this let d be a compatible metric on X, with $\mathcal{B}$ the Borel algebra. Let $N = [1/\varepsilon]$. By Lusin's theorem we can find $X'' \subseteq X'$ with $\mu(X'') > 0$ and all the maps $T, \ldots, T^N$ are continuous on X''. For $x \in X''$, since all the points $x, Tx, \ldots, T^N x$ are distinct, there is a relative ball $B = B_{r(x)}(x) \cap X''$ such that $B, TB, \ldots, T^N B$ are pairwise disjoint, which implies $\mu(E) \leq \varepsilon\mu(\bigcup_{n=0}^{\infty} T^n B)$. Moreover the same holds for any $B' \subseteq B$. Choose such a B' containing X from some fixed countable basis for the topology of X'' (here we use standardness, though actually only separability is needed). Since these B' form a countable cover of X'' one of them must have positive mass. This is the desired set E. □

8.2 The group of automorphisms and residuality

Fix $\Omega = ([0,1], \mathcal{B}, m)$ Lebesgue measure on the unit interval with Borel sets. Let Aut denote the group of measure-preserving maps of Ω, with two maps identified if they agree on a Borel set of full measure.

We introduce a topology on Aut whose sub-basis is given by sets of the form

$$U(T, A, \varepsilon) = \{S \in \text{Aut} : \mu(S^{-1}A \Delta T^{-1}A) < \varepsilon\}$$

where $T \in \text{Aut}$, $A \in \mathcal{B}$ and $\varepsilon > 0$. Note that $T \in U(T, A, \varepsilon)$.

We may also identify Aut with a subgroup of the group of unitary (or bounded linear) operators of $L^2(m)$. As such it inherits the strong and weak operator topologies from the group of bounded linear operators. These are the topologies whose bases are given by sets

$$\{S : \|Tf - Sf\|_2 < \varepsilon\} \qquad \text{for given operator } T,\ f \in L^2\ \varepsilon > 0$$

and

$$\{S : |(Tf, Sf) - (f, f)| < \varepsilon\} \qquad \text{for given operator } T,\ f \in L^2\ \varepsilon > 0$$

respectively. When restricted to the group of unitary operators these are equivalent bases, as can be seen from the identity

$$\|Tf - Sf\|_2^2 = \|Tf\|_2^2 - 2(Tf, Sf) + \|Sf\|_2^2 = 2(f, f) - 2(Tf, Sf)$$

Now. the topology we have defined is clearly weaker than the strong operator topology, since

$$\mu(S^{-1}A \Delta T^{-1}A) = \|1_{S^{-1}A} - 1_{TA}\|_2^2 = \|S1_A - T1_A\|_2^2$$

so

$$U(T, 1_A, \varepsilon^2) \subseteq U(T^{-1}, A, \varepsilon)$$

On the other hand for step functions $f = \sum_{i=1}^k a_i 1_{A_i}$ we can show that it is also stronger: setting $a = \sum |a_i|$,

$$\bigcap_{i=1}^{k} U(T, A_i, \varepsilon/ak) \subseteq U(T, f, \varepsilon)$$

For general f one can argue by approximation. Hence, all three topologies on Aut agree. It also shows that the topology makes composition continuous and makes tha map $T \mapsto T^{-1}$ continuous (since this is true in the unitary group).

Definition 8.2.1. Denote by $\mathcal{D}_n$ the partition of $[0,1)$ into the dyadic intervals $[k/2^n,(k+1)/2^n)$.

Lemma 8.2.2. Aut *is closed in the group of unitary operators with the strong operator topology.*

Proof. Let $T_n \in$ Aut. Since T_n arise from maps of $[0,1]$, each T_n is not only a unitary map of L^2, it also preserves pointwise multiplication of functions in L^∞: $T_n(fg) = T_n f \cdot T_n g$. It is easy to see that if $T_n \to T \in U(L^2(m))$ in the strong operator topology. Then is it a simple matter to verify that T also preserves pointwise multiplication. But it is then a classical fact that T arises from a measure preserving map of the underlying space. By using the fact that $T_n^{-1} \to T^{-1}$ we similarly see that T^{-1} arises from a measure preserving map. Now the relation $T_n T_n^{-1} A = T_n^{-1} T_n A = A$ imples that T, T^{-1} are inverses, so $T, T^{-1} \in$ Aut. $\square$

Corollary 8.2.3. Aut *is Polish in our topology.*

Definition 8.2.4. A set is called **dyadic** (of generation n) if it is the union of elements of $\mathcal{D}_n$. Let $\mathcal{D}_n^*$ denote the algebra of these sets.

$T \in$ Aut is **dyadic** (of generation n) if it permutes the elements of $\mathcal{D}_n$ (with the map of intervals realized by isometries).

A dyadic automorphism is **cyclic** if the permutation is a cycle.

Note that if π is a permutation of $\mathcal{D}_n$ then there is an automorphism $S_\pi \in \mathrm{Aut}_D$ such that $S_\pi : I \to \pi I$ in a linear and orientation preserving manner.

Proposition 8.2.5. *The cyclic permutations are dense in* Aut. *Furthermore for every n_0, the cyclic permutations of order $\geq 2^{n_0}$ are dense.*

Proof. Fix n_0. Let $U_i = U(A_i, T_i, \varepsilon_i)$, $i = 1, \ldots, k$, be given. We must show that $\mathrm{Aut}_C \cap \bigcap_{i=1}^k U_i \neq \emptyset$.

First let us show that $\mathrm{Aut}_D \cap (\bigcap U_i) \neq \emptyset$. Let $\mathcal{A}$ denote the coarsest partition of X that refines the partitions $\{A_i, X \setminus A_i\}$. Let $T\mathcal{A}$ be its image, which is similarly defined by the sets TA_i. Fix an auxiliary parameter $\delta > 0$. Given n and $A \in \mathcal{A}$ let

$$\begin{aligned}\mathcal{E}_n(A) &= \{I \in \mathcal{D}_n : m(I \cap A) > (1-\delta)m(I)\}\\ \mathcal{F}_n(A) &= \{I \in \mathcal{D}_n : m(I \cap TA) > (1-\delta)m(I)\}\end{aligned}$$

and write

$$\begin{aligned}E_n(A) &= \cup \mathcal{E}_n(A)\\ F_n &= \cup \mathcal{F}_n(A)\end{aligned}$$

By Lebesgue's differentiation theorem, we can choose an $n > n_0$ so large that for every $A \in \mathcal{A}$,

$$\begin{aligned}m(A \Delta E_n(A)) &< \delta\\ m(TA \Delta F_n(A)) &< \delta\end{aligned}$$

Of course, $m(A) = m(TA)$, so the above means that $|m(E_n(A)) - m(F_n(A))| < 2\delta$. If $m(E_n(A)) < m(F_n(A))$, remove dyadic intervals from $\mathcal{F}_n(A)$ of total mass $< 2\delta$, to make the mass of $E_n(A), F_n(A)$ equal. Then all the inequalities above will hold with 3δ instead of δ. Similarly if $m(F_n(A)) < m(E_n(A))$.

Thus we assume $m(E_n(A)) = m(F_n(A))$, which is equivalent to $|\mathcal{E}_n(A)| = |\mathcal{F}_n(A)|$. Let $\pi : \mathcal{D}_n \to \mathcal{D}_n$ be a permutation defined by choosing an arbitrary bijection $\pi : \mathcal{E}_n(A) \to \mathcal{F}_n(A)$, and then extending to a permutation where it is not defined. Let S_π be the associated permutation. For $A \in \mathcal{A}$ clearly $S_\pi(E_n(A)) = F_n(A)$ so

$$\begin{aligned} m(S_\pi A \Delta TA) &\leq m(S_\pi A \Delta S_\pi E_n(A)) + m(S_\pi E_n(A) \Delta F_n(A)) \leq m(F_n(A) \Delta TA) \\ &\leq m(A \Delta E_n(A)) + 0 + m(F_n(A) + TA) \\ &< 6\delta \end{aligned}$$

Assuming δ was small this shows that $S_\pi \in \bigcap U_i$.

The permutation π above need may not be cyclic, so we have only shown so far that Aut_D is dense. To improve this to Aut_C proceed as follows. Fix a large N and consider the partition $\mathcal{D}_{n+N}$ into dyadic intervals of level $n+N$. Let π' be a permutation of $\mathcal{D}_{n+N}$ that has the property that if $I' \in \mathcal{D}_{n+N}$ and $I' \subseteq I$ for $I \in \mathcal{D}_n$, then $\pi' I' \subseteq \pi I$. There are many such π' and we may choose one so that every cycle of π' covers a complete cycle of π. For example, if $I_1 \to I_2 \to I_1$ is a cycle or order 2 of π, then enumerate the $(n+N)$-adic subintervals of I_1 as $I'_{1,1}, \ldots, I'_{1,k}$ and similarly $I'_{2,1}, \ldots, I'_{2,k}$ the sub-intervals of I_2, and define

$$\pi' : I'_{1,1} \to I'_{2,1} \to I'_{1,2} \to I'_{2,2} I'_{1,3} \to I'_{2,3} \to \ldots \to I'_{1,k} \to I'_{2,k} \to I'_{1,1}$$

Now, if $S' = S_{\pi'}$ is the automorphism associated to π' and S to π, then $S'I = SI$ for every $I \in \mathcal{D}_n$, so the argument above shows that $m(S'A, TA) < 6\delta$ for $A \in \mathcal{A}$. Also, it is clear that π' has the same number of cycles as π, which is at most 2^n. By changing the definition of π' on at most 2^n of the $n+N$-adic intervals, we obtain cyclic permutation π'' and associated $S'' = S_{\pi''} \in \mathrm{Aut}_C$ with

$$m(S'' I \Delta S' I) \leq 2^n 2^{-n+N} = 2^{-N} \qquad \text{for } I \in \mathcal{D}_n$$

Thus, assuming $2^{-N} < \delta$, we will have $m(S'A, TA) < 8\delta$ for $A \in \mathcal{A}$. Thus $S'' \in \mathrm{Aut}_C \cap \bigcap U_i$, as desired. □

Theorem 8.2.6 (Rohlin density phenomenon). *Let $(X, \mathcal{B}, \mu, T)$ be an aperiodic measure preserving transformation on a standard measure space. Then the isomorphism class*

$$[T] = \{S \in \mathrm{Aut} : S \cong T\}$$

is dense in Aut.

Proof. Take $S \in \mathrm{Aut}$, A and ε, we must show $[T] \cap U(S, A, \varepsilon) \neq \emptyset$. By the previous proposition, it suffices to show this for $S \in \mathrm{Aut}_C$ of arbitrarily high order N. Now, by Rohlin's lemma we can find $B \in \mathcal{B}$ such that $B, TB, \ldots, T^{N-1}B$ are disjoint and fill $> 1 - \varepsilon$ of X. Let $I_1, \ldots, I_N$ denote an ordering of dyadic intervals as permuted by S. Let I'_N denote $I_n \setminus I''_n$ where I''_n is the interval at the right end of I_n of length ε/N. Thus $m(I'_n) = (1-\varepsilon)/N = \mu(B)$ and S maps $I'_n \to I'_{n+1}$.

Using standardness we can find an isomorphism $\pi_0 : B \to I_1$. Now for $x \in T^n B$, $0 \leq n < N$, define $\pi(x) = S^n(\pi_0(T^{-n}x))$. Finally, extend π to an isomorphism $X \to [0,1]$ using standardness again.

Let $S' : [0,1] \to [0,1]$ be given by $S'(y) = \pi T \pi^{-1} y$. By construction, $S' = S$ on $\bigcup I'_n$. From this it follows easily that for any dyadic interval J of level $< \log_2 N$, we have $m(S'J\Delta SJ) < 2\varepsilon$. The claim therefore follows then A is a finite union of dyadic intervals (by taking N large relative to their lengths); and for general A by approximation. $\square$

Take a moment to appreciate this theorem. It immediately implies that the group rotations are dense in Aut (in fact, any particular group rotation is); the isomorphism class of every nontrivial product measure is dense in Aut; etc. Stated another way, for any two non-atomic invertible ergodic transformations of a standard measure space, one can find realizations of them in Aut that are arbitrarily close. Hence, essentially nothing can be learned about an automorphism by observing it at a "finite resolution".

Theorem 8.2.7. *Let*

$$WM = \{T \in \mathrm{Aut} : T \textit{ is weak mixing}\}$$

Then WM is a dense G_δ in Aut. *In particular the set of ergodic automorphisms contains a dense G_δ.*

Proof. Let $\{f_i\}$ be a countable dense subset of $L^2(m)$. Let

$$U = \bigcap_i \bigcap_j \bigcap_n \bigcup_k \{T \in \mathrm{Aut} : |(f_i, T^k f_j) - \int f_i \int f_j| < 1/n\}$$

Since $T \mapsto (f_i, T^k f_j)$ is continuous, the innermost set is open, and so this is a G_δ set. We claim $U = WM$. To see that $WM \subseteq U$, note that for every weak mixing transformation and every $i, j,$, we know that

$$|(f_i, T^k f_j) - \int f_i \int f_j| \xrightarrow{density} 0$$

so for every n there is a k with the difference $< 1/n$, hence $T \in U$.

In the other direction, if $T \in \mathrm{Aut} \setminus WM$, then there is a non-constant eigenfunction $T\varphi = \lambda\varphi$, $\|\varphi\| = 1$, which we may assume has zero integral. Then Taking f_i, f_j close to φ we will have $|(f_i, T^k f_j)| > |(\varphi, T^k \varphi)| - 1/2 = 1/2$ for all k, thus for $n = 3$ and these i, j we do not have $|(f_i, T^k f_j) - \int f_i \int f_j| < 1/n\}$ and therefore $T \notin U$.

Finally, WM is dense by the previous theorem, since it contains aperiodic automorphisms. $\square$

Theorem 8.2.8. *Let*

$$SM = \{T \in \text{Aut} : T \textit{ is strongly mixing}\}$$

Then SM is meager in Aut*.*

Proof. Let $I = [0, 1/2]$. Let $o(S)$ denote the order of a cyclic automorphism $S \in \text{Aut}_C$ and let

$$V = \bigcap_n \bigcup_{\{S \in \text{Aut}_C \,:\, o(S) \geq n\}} \bigcap_{i=0}^{o(S)} U(S, S^i I, \frac{1}{n \cdot o(S)})$$

where $o(S)$ is the order of S. This is a G_δ set and it is dense since Aut_C is dense. So we only need to prove that $V \cap SM = \emptyset$. Indeed, if $T \in V$ then for arbitrarily large n there is a cyclic automorphism S of order $k = o(S) \geq n$ with $T \in U(S, I, 1/nk)$. Now note that $m(I \cap S^k I) = m(I)$. It is an easy induction to show that $m(S^i I \Delta T^i I) \leq i/nk$ for $1 \leq i \leq k$, since for $i = 1$ it follows from $T \in U(S, I, 1/nk$, and assuming we know it for $i - 1$, we can write

$$\begin{aligned}
m(S^i I \Delta T^i I) &= m(S(S^{i-1}I)\Delta T(T^{i-1}I)) \\
&\leq m(S(S^{i-1}I)\Delta T(S^{i-1}I)) + m(T(S^{i-1}I)\Delta T(T^{i-1}I)) \\
&\leq \frac{1}{nk} + m(S^{i-1}I\Delta T^{i-1}I) \\
&\leq \frac{1}{nk} + \frac{i-1}{nk}
\end{aligned}$$

where in the second inequality we used $T \in U(S, S^{i-1}I, 1/nk)$ and the fact that T is measure-preserving. We conclude that

$$m(S^k I \Delta T^k I) \leq \frac{k}{nk} = \frac{1}{n}$$

so

$$m(I \Delta T^k I) \leq m(I \Delta S^k I) + m(S^k I \Delta T^k I) = 0 + \frac{1}{n}$$

Since $k \to \infty$ as $n \to \infty$, we do not have $m(I \Delta T^k I) \to m(I)^2 = \frac{1}{4}$. This shows that T is not strong mixing. $\square$

Chapter 9

Appendix

9.1 The weak-* topology

Proposition 9.1.1. *Let X be a compact metric space. Then $\mathcal{P}(X)$ is metrizable and compact in the weak-* topology.*

Proof. Let $\{f_i\}_{i=1}^{\infty}$ be a countable dense subset of the unit ball in $C(X)$. Define a metric on $\mathcal{P}(X)$ by

$$d(\mu,\nu) = \sum_{i=1}^{\infty} 2^{-i} |\int f_i d\mu - \int f_i d\nu|$$

It is easy to check that this is a metric. We must show that the topology induced by this metric is the weak-* topology.

If $\mu_n \to \mu$ weak-* then $\int f_i d\mu_n - \int f_i d\mu \to 0$ as $n \to \infty$, hence $d(\mu_n, \mu) \to 0$.

Conversely, if $d(\mu_n, \mu) \to 0$, then $\int f_i d\mu_n \to \int f_i d\mu$ for every i and therefore for every linear combination of the f_is. Given $f \in C(X)$ and $\varepsilon > 0$ there is a linear combination g of the f_i such that $\|f - g\|_\infty < \varepsilon$. Then

$$\begin{aligned} |\int f d\mu_n - \int f d\mu| &< |\int f d\mu_n - \int g d\mu_n| + |\int g d\mu_n - \int g d\mu| + |\int g d\mu - \int f d\mu| \\ &< \varepsilon + |\int g d\mu_n - \int g d\mu| + \varepsilon \end{aligned}$$

and the right hand side is $< 3\varepsilon$ when n is large enough. Hence $\mu_n \to \mu$ weak-*.

Since the space is metrizable, to prove compactness it is enough to prove sequential compactness, i.e. that every sequence $\mu_n \in \mathcal{P}(X)$ has a convergent subsequence. Let $V = \mathrm{span}_{\mathbb{Q}}\{f_i\}$, which is a countable dense $\mathbb{Q}$-linear subspace of $C(X)$. The range of each $g \in V$ is a compact subset of $\mathbb{R}$ (since X is compact and g continuous) so for each $g \in V$ we can choose a convergent subsequence of $\int g d\mu_n$. Using a diagonal argument we may select a single subsequence $\mu_{n(j)}$ such that $\int g \mu_{n(j)} \to \Lambda(g)$ as $j \to \infty$ for every $g \in V$. Now, Λ is a $\mathbb{Q}$-linear functional because

$$\begin{aligned}
\Lambda(af_i + bf_j) &= k \lim \int (af_i + bf_j)\, d\mu_{n(k)} \\
&= \lim_{k\to\infty} a \int f_i d\mu_{n(k)} + b \int f_j d\mu_{n(k)} \\
&= a\Lambda(f_i) + b\Lambda(f_j)
\end{aligned}$$

Λ is also uniformly continuous because, if $\|f_i - f_j\|_\infty < \varepsilon$ then

$$\begin{aligned}
|\Lambda(f_i - f_j)| &= \left| \lim_{k\to\infty} \int (f_i - f_j)\, d\mu_{n(k)} \right| \\
&\leq \lim_{k\to\infty} \int |f_i - f_j| d\mu_{n(k)} \\
&\leq \varepsilon
\end{aligned}$$

Thus Λ extends to a continuous linear functional on $C(X)$. Since Λ is positive (i.e. non-negative on non-negative functions), sos is its extension, so by the Riesz representation theorem there exists $\mu \in \mathcal{P}(X)$ with $\Lambda(f) = \int f d\mu$. By definition $\int g d\mu - \int g d\mu_{n(k)} \to 0$ as $k \to \infty$ for $g \in V$, hence this is true for the f_i, so $d(\mu_{n(k)}, \mu) \to 0$ Hence $\mu_{n(k)} \to \mu$ weak-* . □

9.2 Conditional expectation

When (X, B, μ) is a probability space, $f \in L^1$, and A a set of positive measure, then the conditional expectation of f on A is usually defined as $\frac{1}{\mu(A)} \int_A f\, d\mu$. When A has measure 0 this formula is meaningless, and it is not clear how to give an alternative definition. But if $\mathcal{A} = \{A_i\}_{i\in I}$ is a partition of X into measurable sets (possibly of measure 0), one can sometimes give a meaningful definition of the conditional expectation of f on $\mathcal{A}(x)$ for a.e. x, where $\mathcal{A}(x)$ is the element A_i containing x. Thus the conditional expectation of$f$ on $\mathcal{A}$ is a function that assigns to a.e. x the conditional expectation of f on the set $\mathcal{A}(x)$. Rather than partitions, we will work with σ-algebras; the connection is made by observing that if $\mathcal{E}$ is a countably-generated σ-algebra then the partition of X into the atoms of $\mathcal{E}$ is a measurable partition.

Theorem 9.2.1. *Let $(X, \mathcal{B}, \mu)$ be a probability space and $\mathcal{E} \subseteq \mathcal{B}$ a sub-σ algebra. Then there is a linear operator $L^1(X, \mathcal{B}, \mu) \to L^1(X, \mathcal{E}, \mu)$ satisfying*

1. *Chain rule: $\int \mathbb{E}(f|\mathcal{E})\, d\mu = \int f\, d\mu$.*
2. *Product rule: $\mathbb{E}(gf|\mathcal{E}) = g \cdot \mathbb{E}(f|\mathcal{E})$ for all $g \in L^\infty(X, \mathcal{E}, \mu)$.*

Proof. We begin with existence. Let $f \in L^1(X, \mathcal{B}, \mu)$ and let μ_f be the finite signed measure $d\mu_f = f d\mu$. Then $\mu_f \ll \mu$ in the measure space $(X, \mathcal{B}, \mu)$ and this remains true in $(X, \mathcal{E}, \mu)$. Let $\mathbb{E}(f|\mathcal{E}) = d\mu_f / d\mu \in L^1(X, \mathcal{E}, \mu)$, the Radon-Nykodim derivative of μ_f with respect to μ in $(X, \mathcal{E}, \mu)$.

The domain of this map is $L^1(X,\mathcal{B},\mu)$ and its range is in $L^1(X,\mathcal{E},\mu)$ by the properties of $d\mu_f/d\mu$.

Linearity follows from uniqueness of the Radon-Nykodim derivative and the definitions. The chain rule is also immediate:

$$\int \mathbb{E}(f|\mathcal{E})\,d\mu = \int \frac{d\mu_f}{d\mu}\,d\mu = \int f\,d\mu$$

For the product rule, let $g \in L^\infty(X,\mathcal{E},\mu)$. We must show that $g \cdot \frac{d\mu_f}{d\mu} = \frac{d\mu_{gf}}{d\mu}$ in $(X,\mathcal{E},\mu)$. Equivalently we must show that

$$\int_E g\frac{d\mu_f}{d\mu}\,d\mu = \int_E \frac{d\mu_{gf}}{d\mu}\,d\mu \qquad \text{for all } E \in \mathcal{E}$$

Now, for $A \in \mathcal{E}$ and $g = 1_A$ we have

$$\begin{aligned}
\int_E 1_A\frac{d\mu_f}{d\mu}\,d\mu &= \int_{A\cap E}\frac{d\mu_f}{d\mu}\,d\mu \\
&= \mu_f(A\cap E) \\
&= \int_{A\cap E} f\,d\mu \\
&= \int_E 1_A f\,d\mu \\
&= \int_E \frac{d\mu_{1_A f}}{d\mu}\,d\mu
\end{aligned}$$

so the identity holds. By linearity of these integrals in the g argument it holds linear combinations of indicator functions. For arbitrary $g \in L^\infty$ we can take a uniformly bounded sequence of such functions converging pointwise to g, and pass to the limit using dominated convergence. This proves the product rule.

To prove uniqueness, let $T : L^1(X,\mathcal{B},\mu) \to L^1(X,\mathcal{E},\mu)$ be an operator with these properties. Then for $f \in L^1(X,\mathcal{B},\mu)$ and $E \in \mathcal{E}$,

$$\begin{aligned}
\int_E Tf\,d\mu &= \int 1_E Tf\,d\mu \\
&= \int T(1_E f)\,d\mu \\
&= \int 1_E f\,d\mu \\
&= \int_E f\,d\mu
\end{aligned}$$

where the second equality uses the product rule and the third uses the chain rule. Since this holds for all $E \in \mathcal{E}$ we must have $Tf = d\mu_f/d\mu$. □

Proposition 9.2.2. *The conditional expectation operator satisfies the following properties:*

1. *Positivity:* $f \geq 0$ *a.e. implies* $\mathbb{E}(f|\mathcal{E}) \geq 0$ *a.e.*
2. *Triangle inequality:* $|\mathbb{E}(f|\mathcal{I})| \leq \mathbb{E}(|f|\,|\mathcal{I})$.
3. *Contraction:* $\|\mathbb{E}(f|\mathcal{E})\|_1 \leq \|f\|_1$*; in particular,* $\mathbb{E}(\cdot|\mathcal{E})$ *is* L^1*-continuous.*
4. *Sup/inf property:* $\mathbb{E}(\sup f_i|\mathcal{E}) \geq \sup \mathbb{E}(f_i|\mathcal{E})$ *and* $\mathbb{E}(\inf f_i|\mathcal{E}) \leq \inf \mathbb{E}(f_i|\mathcal{E})$ *for any countable family* $\{f_i\}$.
5. *Jensen's inequality: if* g *is convex then* $g(\mathbb{E}(f|\mathcal{E})) \leq \mathbb{E}(g \circ f|\mathbb{E})$.
6. *Fatou's lemma:* $\mathbb{E}(\liminf f_n|\mathcal{E}) \leq \liminf \mathbb{E}(f_n|\mathcal{E})$.

Remark 9.2.3. Properties (2)–(6) are consequences of positivity only.

Proof. (1) Suppose $f \geq 0$ and $\mathbb{E}(f|\mathcal{E}) \not\geq 0$, so $\mathbb{E}(f|\mathcal{E}) < 0$ on a set $A \in \mathcal{E}$ of positive measure. Applying the product rule with $g = 1_A$, we have

$$\mathbb{E}(1_A f|\mathcal{E}) = 1_A \mathbb{E}(f|\mathcal{E})$$

hence, replacing f by 1_A, we can assume that $f \geq 0$ and $\mathbb{E}(f|\mathcal{E}) < 0$. But this contradicts the chain rule since $\int f\,d\mu \geq 0$ and $\int \mathbb{E}(f|\mathcal{E})\,d\mu < 0$.

(2) Decompose f into positive and negative parts, $f = f^+ - f^-$, so that $|f| = f^+ + f^-$. By positivity,

$$\begin{aligned}
|\mathbb{E}(f|\mathcal{E})| &= |\mathbb{E}(f^+|\mathcal{E}) - \mathbb{E}(f^-|\mathcal{E})| \\
&\leq |\mathbb{E}(f^+|\mathcal{E})| + |\mathbb{E}(f^-|\mathcal{E})| \\
&= \mathbb{E}(f^+|\mathcal{E}) + \mathbb{E}(f^-|\mathcal{E}) \\
&= \mathbb{E}(f^+ + f^-|\mathcal{E}) \\
&= \mathbb{E}(|f|\,|\mathcal{E})
\end{aligned}$$

(3) We compute:

$$\begin{aligned}
\|\mathbb{E}(f|\mathcal{E})\|_1 &= \int |\mathbb{E}(f|\mathcal{E})|\,d\mu \\
&\leq \int \mathbb{E}(|f|\,|\mathcal{E})|\,d\mu \\
&= \int |f|\,d\mu \\
&= \|f\|_1
\end{aligned}$$

where we have used the triangle inequality and the chain rule.

(4) We prove the sup version. By monotonicity and continuity it suffices to prove this for finite families and hence for two functions. The claim now follows from the identity $\max\{f_1, f_2\} = \frac{1}{2}(f_1 + f_2 + |f_1 - f_2|)$, linearity, and the triangle inequality.

(5) For an affine function $g(t) = at + b$,

$$\mathbb{E}(g \circ f|\mathcal{E}) = \mathbb{E}(af + b|\mathcal{E}) = a\mathbb{E}(f|\mathcal{E}) + b = g \circ \mathbb{E}(f|\mathcal{E})$$

If g is convex then $g = \sup g_i$ where $\{g_i\}_{i\in I}$ is a countable family of affine functions. Thus

$$\begin{aligned}
\mathbb{E}(g \circ f|\mathcal{E}) &= \mathbb{E}(\sup_i g_i \circ f|\mathcal{E}) \\
&\geq \sup_i \mathbb{E}(g_i \circ f|\mathcal{E}) \\
&= \sup_i g_i \circ \mathbb{E}(f|\mathcal{E}) \\
&= g \circ \mathbb{E}(f|\mathcal{E})
\end{aligned}$$

(6) Since $\inf_{k>n} f_k \nearrow \liminf f_k$ as $n \to \infty$ the convergence is also in L^1, so by continuity and positivity the same holds after taking the conditional expectation. Thus, using the inf property,

$$\begin{aligned}
\liminf_{n\to\infty} \mathbb{E}(f_n|\mathcal{E}) &= \lim_{n\to\infty} \inf_{k>n} \mathbb{E}(f_k|\mathcal{E}) \\
&\geq \lim_{n\to\infty} \mathbb{E}(\inf_{k>n} f_k|\mathcal{E}) \\
&= \mathbb{E}(\liminf_{n\to\infty} f_n|\mathcal{E}) \square
\end{aligned}$$

Corollary 9.2.4. *The restriction of the conditional expectation operator to* $L^2(X,\mathcal{B},\mu)$ *coincides with the orthogonal projection* $\pi : L^2(X,\mathcal{B},\mu) \to L^2(X,\mathcal{E},\mu)$.

Proof. Write $\pi = \mathbb{E}(\cdot|\mathcal{E})$. If $f \in L^2$ then by by convexity of $t \to t^2$ and Jensen's inequality (which is immediate for simple functions and hence holds for $f \in L^1$ by approximation),

$$\begin{aligned}
\|\pi f\|_2 &= \int |\mathbb{E}(f|\mathcal{E})|^2 \, d\mu \\
&\leq \int \mathbb{E}(|f|^2|\mathcal{E}) \, d\mu \\
&= \int |f|^2 \, d\mu \qquad \text{by the chain rule} \\
&= \|f\|_2
\end{aligned}$$

Thus π maps L^2 into the subspace of $\mathcal{E}$-measurable L^2 functions, hence $\pi : L^2(X,\mathcal{B},m) \to L^2(X,\mathcal{E},\mu)$. We will now show that π is the identity on $L^2(X,\mathcal{E},\mu)$ and is π. Indeed, if $g \in L^2(X,\mathbb{E},\mu)$ then for every $A \in \mathcal{E}$

$$\begin{aligned}
\pi g &= \mathbb{E}(g \cdot 1|\mathcal{E}) \\
&= g \cdot \mathbb{E}(1|\mathcal{E})
\end{aligned}$$

Since $\int \mathbb{E}(1|\mathcal{E}) = \int 1 = 1$, this shows that π is the identity on $L^2(X,\mathcal{E},)$. Next if $f, g \in L^2$ then $fg \in L^1$, and

$$\begin{aligned}
\langle f, \pi g\rangle &= \int f \cdot \mathbb{E}(g|\mathcal{E})\, d\mu \\
&= \int \mathbb{E}\left(f \cdot \mathbb{E}(g|\mathcal{E})\right)\, d\mu && \text{by the chain rule} \\
&= \int \mathbb{E}\left(f|\mathcal{E}\right) \mathbb{E}(g|\mathcal{E})\, d\mu && \text{by the product rule} \\
&= \int \mathbb{E}\left(\mathbb{E}(f|\mathcal{E}) \cdot g\right)\, d\mu && \text{by the product rule} \\
&= \int \mathbb{E}(f|\mathcal{E}) \cdot g\, d\mu && \text{by the chain rule} \\
&= \langle \pi f, g\rangle
\end{aligned}$$

so π is self-adjoint. □

9.3 Regularity

I'm not sure we use this anywhere, but for the record:

Lemma 9.3.1. *A Borel probability measure on a complete (separable) metric space is regular.*

Proof. It is easy to see that the family of sets A with the property that

$$\begin{aligned}
\mu(A) &= \inf\{\mu(U) \,:\, U \supseteq A \text{ is open}\} \\
&= \sup\{\mu(C) \,:\, C \subseteq A \text{ is closed}\}
\end{aligned}$$

contains all open and closed sets, and is a σ-algebra. Therefore every Borel set A has this property. We need to verify that in the second condition we can replace closed by compact. Clearly it is enough to show that for every closed set C and every $\varepsilon > 0$ there is a compact $K \subseteq C$ with $\mu(K >> \mu(C) - \varepsilon$.

Fix C and $\varepsilon > 0$. For every n we can find a finite family $B_{n,1}, \ldots, B_{n,k(n)}$ of δ-balls whose union $B_n = \bigcup B_{n,i}$ intersects A in a set of measure $> \mu(A) - \varepsilon/2^n$. Let $K_0 = C \cap \bigcap B_n$, so that $\mu(K_0) > \mu(C) - \varepsilon$. By construction K_0 is pre-compact, and $K = \overline{K_0} \subseteq C$, so K has the desired property. □