Ergodic Theory

Lecture Notes

Contents

Chapter 0

Preliminaries from Measure Theory

0.1. Algebras and σ-Algebras

Throughout this section, we use X to denote an arbitrary set. If A is a subset of X, we write $A^c = X \setminus A$ for the set-complement of A relative to X. We also define $\mathscr{P}(X)$ to be the power set of X, that is, the set formed by all subsets of X.

Definition 1 (Algebras and σ-algebras)**.** Let $\mathscr{A} \subseteq \mathscr{P}(X)$ be a collection of subsets of X. Then $\mathscr{A}$ is called an *algebra over* X if it satisfies:

(i) (Contains the empty set as an element) We have $\emptyset \in \mathscr{A}$;
(ii) (Closed under complements) If $A \in \mathscr{A}$, then $A^c \in \mathscr{A}$;
(iii) (Closed under finite unions) If $A_1, \ldots, A_k \in \mathscr{A}$, then $A_1 \cup \ldots \cup A_k \in \mathscr{A}$.

Moreover, $\mathscr{A}$ is called a *σ-algebra over* X if in addition to being an algebra, it also satisfies:

(iv) (Closed under countable unions) If $\{A_n\}_{n \in I} \subseteq \mathscr{A}$, $I \subseteq \mathbb{N}$, is a countable family of sets in $\mathscr{A}$, then $\bigcup_{n \in I} A_n \in \mathscr{A}$.

If $\mathscr{A}$ is an algebra (or a σ-algebra) of subsets of X, then a subset of X is said to be *$\mathscr{A}$-measurable* if it belongs to $\mathscr{A}$.

Here are some first examples of algebras over a set X.

Example 2 (Algebras)**.**

- The collection of subsets of X which are either finite or co-finite (meaning that their complement is fininte) is an algebra.
- The collection of all finite unions of intervals of the form $(-\infty, b], (a, b], (a, \infty)$, for $a, b \in \mathbb{R}$, is an algebra on the real numbers $\mathbb{R}$.

Note that any σ-algebra is an algebra but the converse is not true. Indeed, the second algebra provided in Example 2 above is not a σ-algebra.

A σ-algebra is also closed under countable intersections, that is, given a σ-algebra $\mathscr{A}$ and a countable family of sets $\{A_n\}_{n\in I} \subseteq \mathscr{A}$, $I \subseteq \mathbb{N}$, we have that $\bigcap_{n\in I} A_n \in \mathscr{A}$. This follows from De Morgan's law $(\bigcap_{n\in I} A_n)^c = \bigcup_{n\in I} A_n^c \in \mathscr{A}$, and property (iv). Here are some basic examples of σ-algebras.

Example 3 (σ-Algebras)**.**
- $\mathscr{P}(X)$ and $\{\emptyset, X\}$ are σ-algebras.
- For any subset $A \subseteq X$, $\mathscr{A} = \{\emptyset, A, X\setminus A, X\}$ is a σ-algebra.
- The collection of subsets of X which are either countable or co-countable (meaning that their complement is countable) is a σ-algebra.
- Given two σ-algebras $\mathscr{A}_1, \mathscr{A}_2 \subseteq \mathscr{P}(X)$, we have that $\mathscr{A}_1 \cap \mathscr{A}_2$ is also a σ-algebra. More generally, for any (possibly uncountable) family of σ-algebras $\mathscr{A}_i \subseteq \mathscr{P}(X)$, $i \in I$, the intersection $\bigcap_{n\in I} \mathscr{A}_n$ is a σ-algebra.

There is a very natural way of generating σ-algebras from a collection of subsets:

Definition 4 (σ-algebra generated by a collection of subsets)**.** Let $\mathscr{F} \subseteq \mathscr{P}(X)$ be a collection of subsets of X. The smallest σ-algebra containing $\mathscr{F}$, that is, the intersection of all σ-algebras containing $\mathscr{F}$ is called the *σ-algebra generated by $\mathscr{F}$*, and is usually denoted by $\sigma(\mathscr{F})$.

There are many families of subsets that generate useful σ-algebra, we will cover in this section some of them. Here are two simple examples of generated σ-algebras.

Example 5 (Generated σ-algebras)**.**
- The collection of subsets of X which are countable or whose complements are countable is the σ-algebra generated by the singletons of X.
- Let X_1, X_2 be two sets, and $\mathscr{A}_1, \mathscr{A}_2$ be σ-algebras on X_1 and X_2 respectively. We define $\mathscr{A}_1 \otimes \mathscr{A}_2$ to be the σ-algebra on the Cartesian product $X = X_1 \times X_2$ generated by all the subsets of the form $A_1 \times A_2 \subseteq X$, where $A_1 \in \mathscr{A}_1$ and $A_2 \in \mathscr{A}_2$. Note that $\mathscr{A}_1 \otimes \mathscr{A}_2$ is called the *product σ-algebra* generated by $\mathscr{A}_1$ and $\mathscr{A}_2$

A special case of σ-algebras generated by a collection of subsets are the σ-algebras generated by the open subsets with respect to some topology. This type of σ-algebra will be one of our main focus in Ergodic Theory as we will study the dynamical properties of dynamical systems on topological spaces such as the torus.

Definition 6 (Borel σ-algebra over any topological space)**.** Let (X, τ) be a topological space. The σ-algebra generated by the open subsets of X is called the *Borel σ-algebra on X* and we usually denote it by $\mathscr{B}_X$, or simply $\mathscr{B}$. Its elements are called the *Borel measurable* subset of X.

We give here two examples of such σ-algebras that will be used latter in this section.

Example 7 (Borel σ-algebra)**.**
- Consider $\mathbb{R}^d$ endowed with its usual topology. Then, the Borel-σ-algebra on

$\mathbb{R}^d$ is the σ-algebra generated by the open balls $B_r(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^d : |\mathbf{x} - \mathbf{y}| < r\}$. It contains all closed subsets of $\mathbb{R}^d$, but not all subsets of $\mathbb{R}^d$.

- Consider a finite set Σ, usually called the *alphabet*, containing n elements, usually refereed to as the *letters* of Σ. The collection of all infinite strings in these letters is defined as the product space $\Sigma^{\mathbb{N}}$. Observe that the natural topology on Σ is the discrete topology, whose basis consists of singletons, i.e. individual letters. The Borel σ-algebra on $\Sigma^{\mathbb{N}}$ is the σ-algebra generated by the algebra of cylinder sets, where cylinder sets consist of the open sets of $x \in \Sigma^{\mathbb{N}}$ (with respect to the product topology of $\Sigma^{\mathbb{N}}$) that have finitely many coordinates fixed.

We now define the notion of monotone class, which gives rise to another characterization of σ-algebras.

Definition 8 (Monotone class)**.** A monotone class $\mathscr{M} \subseteq \mathscr{P}(X)$ is a collection of subsets of X having the following properties:

(i) if $A_1, A_2, \dots \in \mathscr{M}$ and $A_1 \subseteq A_2 \subseteq \dots$, then $\bigcup_{n=1}^{\infty} A_n \in \mathscr{M}$
(ii) if $B_1, B_2, \dots \in \mathscr{M}$ and $B_1 \supseteq B_2 \supseteq \dots$, then $\bigcap_{n=1}^{\infty} B_n \in \mathscr{M}$

Note that both $\mathscr{P}(X)$ and $\{\emptyset, X\}$ are monotone classes. Thus any collection of subsets is contained in a monotone class. The following theorem gives an alternative characterization of the σ-algebra generated by an algebra.

Theorem 9 (Monotone Class Theorem)**.** *Let $\mathscr{A} \subseteq \mathscr{P}(X)$ be an algebra and let $\mathscr{S}$ be the smallest monotone class containing $\mathscr{A}$. Then we have $\sigma(\mathscr{A}) = \mathscr{S}$.*

0.2. Measures and Measure Spaces

A *measure* is a function that assigns a non-negative number to certain subsets of a set X in a manner consistent with the algebra of Boolean set operations, including unions, intersections, and complements. Measures provide the mathematical foundation for modeling quantities such as mass, length, area, volume, and, most importantly, probability. The subsets to which a measure can be assigned are called the *measurable sets*.

Definition 10 (Measurable space and measurable set)**.** An ordered pair $(X, \mathscr{A})$, where X is a set and $\mathscr{A} \subseteq \mathscr{P}(X)$ is a σ-algebra, is called a *measurable space*, and any set $A \in \mathscr{A}$ is called a *measurable set*.

Definition 11 (Measure and measure space)**.** A *measure* μ on a measurable space $(X, \mathscr{A})$ is a set function $\mu \colon \mathscr{A} \to [0, \infty]$ such that:

(i) $\mu(\emptyset) = 0$;

(ii) For any countable (or finite) sequence of pairwise disjoint sets $(A_n)_{n\in\mathbb{N}} \in \mathscr{A}$, we have

$$\mu\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mu(A_n). \qquad (\sigma\text{-additivity})$$

If $(X, \mathscr{A})$ measurable space and μ is a measure on it then the triple $(X, \mathscr{A}, \mu)$ is called a *measure space*.

The main structure of interest in classical ergodic theory is that of a probability space.

Definition 12 (finite and σ-finite measure space, probability space)**.** A measure space $(X, \mathscr{A}, \mu)$ is said to be a finite measure space if μ satisfies $\mu(X) < \infty$, and if in addition $\mu(X) = 1$, $(X, \mathscr{A}, \mu)$ is called a probability space.
$(X, \mathscr{A}, \mu)$ is called a σ-finite measure space if X is a countable union of elements of $\mathscr{A}$ of finite measure.

We now state different useful properties about measures.

Proposition 13. *Given a measure space $(X, \mathscr{A}, \mu)$, we have the following properties:*

(i) (Finite unions) For any positive integer n and disjoint sets $A_1, A_2, \ldots, A_n \in \mathscr{A}$, using the fact that $\mu(\emptyset) = 0$, we have

$$\mu\left(\bigcup_{k=1}^{n} A_k\right) = \sum_{k=1}^{n} \mu(A_k).$$

(ii) (Monotonicity) If $A, B \in \mathscr{A}$ and $A \subseteq B$, then $\mu(A) \leqslant \mu(B)$.

(iii) (Countable subadditivity) For any countable family of sets $\{A_n\}_{n\in I} \subseteq \mathscr{A}$, $I \subseteq \mathbb{N}$, not necessarily disjoint, we have

$$\mu\left(\bigcup_{n=1}^{\infty} A_n\right) \leqslant \sum_{n=1}^{\infty} \mu(A_n).$$

(iv) (Continuity) If $A_1 \subseteq A_2 \subseteq \cdots \in \mathscr{A}$, then

$$\lim_{k\to\infty} \mu(A_k) = \mu\left(\bigcup_{n=1}^{\infty} A_n\right),$$

and if $A_1 \supseteq A_2 \supseteq \cdots \in \mathscr{A}$, and $\mu(A_1) < \infty$, then

$$\lim_{k\to\infty} \mu(A_k) = \mu\left(\bigcap_{n=1}^{\infty} A_n\right).$$

Throughout the course we will extensively use the notion of "almost everywhere" (or "for almost every"). In short, a property holds almost everywhere on a set X if the subsets of elements for which it doesn't hold has zero measure. During the course, as we deal with probability measures, one way of seeing this notion is as follows: If we pick at random an element $x \in X$, then the probability that x satisfies the given property is 1. Here is the formal definition.

Definition 14 (Almost everywhere)**.** Let $(X, \mathscr{A}, \mu)$ be a measure space. We say that a property holds μ-almost everywhere on X (sometimes abbreviated as μ-a.e.) if it the set of elements for which the property does not hold has zero measure with respect to μ.

Examples

Below, we provide several examples of important measure spaces, many of which will appear again as we delve deeper into ergodic theory throughout the course.

Null measure. Let X be a non-empty set, let $\mathscr{A}$ be a σ-algebra on X and define $\mu(A) = 0$, $\forall A \in \mathscr{A}$. Then $(X, \mathscr{A}, \mu)$ is a measure space and μ is refereed to as the *null measure* on $(X, \mathscr{A})$.

Counting measure. Let X be a set, and for any $A \in \mathscr{P}(X)$ define $\mu(A) = |A|$, where $|A|$ denotes the cardinality of A. Then $(X, \mathscr{P}(X), \mu)$ is a measure space and μ is called the *counting measure* on X. This measure is finite when X is finite, it is σ-finite when X is countable, and it is not σ-finite when X is uncountable.

Dirac δ-measure. Let $(X, \mathscr{A})$ be a measurable space, and $x \in X$. Then we define the *Dirac measure* δ_x by

$$\delta_x(A) = \begin{cases} 1 & \text{if} \quad x \in A, \\ 0 & \text{if} n \end{cases}$$

The Dirac measure is a probability measure, and it represents the almost sure outcome x in the measurable space.

Restriction of a measure. Let $(X, \mathscr{A}, \mu)$ be a measure space and $A \in \mathscr{A}$. We define the measure ν by $\nu(B) = \mu(B \cap A), \forall B \in \mathscr{A}$, to be the restriction of μ to A. Then $(X, \mathscr{A}, \nu)$ is measure space and $\nu(B) = 0$, $\forall B \in \mathscr{A}$ with $B \subseteq X \setminus A$.

Conditional measure. Let $(X, \mathscr{A}, \mu)$ be a measure space and $A \in \mathscr{A}$ with $\mu(A) > 0$. We define for every $B \in \mathscr{A}$,

$$\mu|_A(B) = \mu(B|A) = \frac{\mu(A \cap B)}{\mu(A)}.$$

The set function $B \mapsto \mu|_A(B)$ is a measure on $\mathscr{A}$ called the *conditional measure with respect to* A. If μ is a finite measure (resp. probability measure) then the conditional measure with respect to A is also a finite measure (resp. probability measure).

Product measure. Let $(X_1, \mathscr{A}_1, \mu_1)$ and $(X_2, \mathscr{A}_2, \mu_2)$ be two measure spaces. Let $\mathscr{A} = \mathscr{A}_1 \otimes \mathscr{A}_2$ be the product σ-algebra on the product space $X = X_1 \times X_2$. We define the product measure $\mu = \mu_1 \times \mu_2$ (also sometimes denoted $\mu_1 \otimes \mu_2$) to be the unique measure on the measurable space $(X, \mathscr{A})$ which satisfies $\mu(A_1 \times A_2) = \mu_1(A_1)\mu_2(A_2)$ for every $A_1 \in \mathscr{A}_1$ and $A_2 \in \mathscr{A}_2$.

Probability measure on $\Sigma^{\mathbb{N}}$. Let $X = \Sigma^{\mathbb{N}}$ be the set of all infinite strings whose letters are in the finite alphabet Σ as previously defined. Denote by $\mathscr{A}$ the Borel σ-algebra on $\Sigma^{\mathbb{N}}$ generated by the cylinder sets. Let μ_0 be any probability measure on Σ. We define $\mu = \mu_0^{\mathbb{N}}$ to be the product measure on $\Sigma^{\mathbb{N}}$, which is the unique measure satisfying for every cylinder set I,

$$\mu(I) = \prod_{i \in F} \mu_0(\{x_i\})$$

where F is the finite set of the indices of the fixed coordinated of I.

Borel measure. In order to define Borel measures, we recall two definitions from topology.

Definition 15. A topological space X is Hausdorff if for any distinct points $x, y \in X$, there exists open neighborhoods U, V of x and y respectively such that U and V are disjoint.

Definition 16. A topological space X is locally compact if every $x \in X$ has a compact neighborhood.

Now, let X be a locally compact Hausdorff topological space and $\mathscr{B}$ the Borel σ-algebra defined on X. Then, any measure μ defined on $\mathscr{B}$ is called a Borel measure. If $\mu(X) = 1$, we say that μ is a Borel probability measure.

Radon measure. Let X be a locally compact Hausdorff topological space, $\mathscr{B}$ the Borel σ-algebra defined on X, and μ a finite Borel measure on $\mathscr{B}$. If in addition μ is tight, in the sense that for all $\varepsilon > 0$ there exists a compact set $K \subseteq X$ such that $\mu(X \setminus K) < \varepsilon$ (or equivalently $\mu(K) \geqslant \mu(X) - \varepsilon$), μ is called a Radon measure. These conditions guarantee that the measure is compatible in some sense with the topology of the space. We state the following theorem, without providing a proof, which gives a sufficient condition for a Borel measure to be Radon under a commonly encountered topological assumption on X:

Theorem 17. *Let X be a locally compact Hausdorff topological space in which every open set is σ-compact (that is, a countable union of compact sets). Then every Borel measure on X that is finite on compact sets is a Radon measure.*

It follows that every finite Borel measure on a compact metric Hausdorff topological space X is automatically Radon, as X verifies the conditions of the theorem.

Notice that if we only assume X to be a locally compact, second-countable Hausdorff topological space, then because it follows that X is σ-compact, X still verifies the conditions of the theorem.
An useful property of the Radon measure is that it makes the map $f \mapsto \int f \,\mathrm{d}\mu$, where $f \in L^1(X)$, continuous (recalls about integration theory are given in the next section).
The following measures are examples of Radon measures: the Lebesgue measure on an Euclidean space, the Haar measure on any locally compact topological group, the Dirac measure on any topological space.

Lebesgue measure. The Lebesgue measure is the unique measure μ on the Borel-σ-algebra $\mathscr{B}_{\mathbb{R}}$ such that for every interval $I \subseteq \mathbb{R}$, the measure $\mu(I)$ is the length of I.
We observe that the restriction of μ to the Borel-σ-algebra $\mathscr{B}_{[0,1]}$ of subsets of $[0,1]$ is the so called uniform distribution from probability theory.
We can generalize this idea to higher dimensions. Indeed, for the lower dimensions $n = 1,2$, the Lebesgue measure coincides with the notions of area and volume. For higher dimensions, it is also called n-dimensional volume.
More generally, if we consider the measurable space $(\mathbb{R}^n, \mathscr{B}_{\mathbb{R}^n})$, the Lebesgue measure μ on $\mathscr{B}_{\mathbb{R}^n}$ is the unique measure such that if A is a cartesian product of intervals $I_1 \times \cdots \times I_n$, then A is Lebesgue measurable (in the sense that we can attribute a Lebesgue measure to A) and $\mu(A) = \prod_{i=1}^n l(I_i)$, where l denotes the length of the interval I_i, i.e, l is the Lebesgue measure in one dimension.
We list, without proof, some of the properties of the Lebesgue measure on $\mathscr{B}_{\mathbb{R}^n}$:

(i) (translation invariance) If $A \subseteq \mathbb{R}^n$ is Lebesgue measurable, and $x \in \mathbb{R}^n$, then $A + x = \{y \in \mathbb{R}^n : y + x \in A\}$ is Lebesgue measurable and $\mu(A + x) = \mu(A)$. In particular, $A \subseteq \mathbb{R}^n$ is Lebesgue measurable if, and only if, all translates of A is Lebesgue measurable.

(ii) (dilation and scaling) Let $c > 0$, $A \subseteq \mathbb{R}^n$ be Lebesgue measurable, and let $cA = \{cy \in \mathbb{R}^n : y \in A\}$, then cA is Lebesgue measurable and $\mu(cA) = c^n \mu(A)$.

(iii) More generally, if T is a linear transformation and A is a Lebesgue measurable subset of $\mathbb{R}^n$, then $T(A)$ is a Lebesgue measurable set of measure $|\det(T)|\mu(A)$.

(iv) Finite or countable sets are Lebesgue measurable and have Lebesgue measure 0, and there exist uncountable Lebesgue measurable sets of measure 0. As an example, we can consider the Cantor set (when $n = 1$). Moreover, there exists sets which are not Lebesgue measurable.

Finally, note that the Haar measure (to be seen in the section about topological groups) on a locally compact Hausdorff topological group can be thought of as the natural generalization of the Lebesgue measure to a general locally compact Hausdorff topological group.

Atomic, non-atomic, and continuous measures. In order to define discrete and continuous measures we will need the following definition.

Definition 18 (Atom). Given a measure space $(X,\mathscr{A},\mu)$, a set $A \in \mathscr{A}$ is called an atom if:

(i) $\mu(A) > 0$, and
(ii) For any measurable set $B \subseteq A$ with $\mu(B) < \mu(A)$ we have $\mu(B) = 0$.

A σ-finite measure μ on a measurable space $(X,\mathscr{A})$ is called *purely atomic* if every measurable set of positive measure contains an atom. On the contrary, a σ-finite measure which has no atoms is called *non-atomic*. Equivalently, μ is *non-atomic* if for every measurable set A such that $\mu(A) > 0$ there exists a measurable subset B of A such that $0 < \mu(B) < \mu(A)$.
Finally, a σ-finite measure μ is called *continuous* if for any $A \in \mathscr{A}$ and any $c \in \mathbb{R}$ such that $0 < c < \mu(A)$, there exists a measurable subset B of A such that $\mu(B) = c$. Note that any continuous measure is non-atomic.

There are two important existence theorems for measures, the Carathéodory Extension Theorem and the Riesz Representation Theorem.

Definition 19 (Pre-measure). Let $\mathscr{A}$ be an algebra on a set X. A set function $\mu_0 : \mathscr{A} \to [0,\infty]$ is called a *pre-measure* on $(X,\mathscr{A})$ if $\mu_0(\varnothing) = 0$ and, for every countable (or finite) sequence $A_1, A_2, \ldots \in \mathscr{A}$ of pairwise disjoint sets whose union lies in $\mathscr{A}$, we have

$$\mu_0\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mu_0(A_n). \qquad (\sigma\text{-additivity})$$

Theorem 20 (Carathéodory). *Let $\mathscr{A}$ be an algebra on a set X. Any pre-measure μ_0 on $\mathscr{A}$ extends to a measure μ on the σ-algebra $\sigma(\mathscr{A})$ generated by $\mathscr{A}$. Moreover, if μ_0 is σ-finite then this extension is unique and σ-finite.*

Let X be a locally compact Hausdorff space. We write $C(X)$ for the space of all continuous complex-valued functions on X. Within this space, we distinguish the subspace $C_0(X)$ of functions *vanishing at infinity*, meaning

$$f \in C_0(X) \quad \Longleftrightarrow \quad \forall \varepsilon > 0,\ \{x \in X : |f(x)| \geqslant \varepsilon\} \text{ is compact.}$$

Equipped with the uniform norm $\|f\|_\infty = \sup_{x\in X} |f(x)|$, $C_0(X)$ is a Banach space.

When X is compact, every continuous function automatically vanishes at infinity, so $C_0(X) = C(X)$. In the non-compact case, $C_0(X)$ forms a proper subspace of $C(X)$, consisting precisely of those continuous functions that tend to 0 at infinity.

Theorem 21 (Riesz-Markov-Kakutani representation theorem). *Let X be a locally compact Hausdorff topological space in which every open set is σ-compact (cf. Theorem 17). For any continuous linear functional $l : C_c(X) \to \mathbb{C}$ there exists a unique* *complex-valued Radon measure μ on X such that*

$$l(f) = \int_X f(x)\,\mathrm{d}\mu(x),$$

for all $f \in C_c(X)$.

0.3. Measurable Functions and Integrals

Throughout this section we let $(X, \mathscr{A}, \mu)$ be a measure space. Natural classes of measurable functions on X are built up from simpler functions, just as the σ-algebra $\mathscr{A}$ may be built up from simpler collections of sets. Given a set $A \subseteq X$, we denote by $\mathbf{1}_A : X \to \{0,1\}$ the indicator function of A, that is,

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A, \end{cases} \qquad \forall x \in X.$$

Definition 22 (Simple function)**.** A function $f : X \to \mathbb{R}$ is called *simple* if

$$f(x) = \sum_{j=1}^{m} c_j \mathbf{1}_{A_j}(x), \qquad \forall x \in X,$$

where $c_j \in \mathbb{R}$ and the $A_j \in \mathscr{A}$ are disjoint sets $\forall j = 1, \ldots, m$. The *integral* of f is then defined to be

$$\int f \,\mathrm{d}\mu = \sum_{j=1}^{m} c_j \mu(A_j). \tag{0.3.1}$$

Definition 23 (Measurable function)**.** A function $g : X \to \mathbb{R}$ is called measurable if $g^{-1}(A) \in \mathscr{A}$ for any (Borel) measurable set $A \in \mathscr{B}_{\mathbb{R}}$.

Note that simple functions are always measurable functions. Below, we outline several methods for generating new measurable functions from existing measurable ones.

Proposition 24. *Let $f, g : X \to \mathbb{R}$ be measurable, and $c \in \mathbb{R}$. Then, the following functions are measurable:*

(i) cf
(ii) $f + g$
(iii) fg
(iv) $|f|$
(v) $\min\{f, g\}$ *and* $\max\{f, g\}$
(vi) *$Re(f)$ and $Im(f)$, where we understand $Re(f)$ (resp. $Im(f)$) as the unique function such that $Re(f(x)) = Re(f)(x)$ (resp. $Im(f(x)) = Im(f)(x)$) $\forall x \in X$.*

The integral of simple functions has already been defined in (0.3.1). Our next goal is to extend this definition to all measurable functions. To achieve this, we rely on the following key approximation result.

Proposition 25. *Let $g: X \to [0,\infty)$ be a measurable function taking non-negative values. There exists a pointwise increasing sequence of simple functions $(f_n)_{n\in\mathbb{N}}$ (in the sense that $f_n(x) \leqslant f_{n+1}(x)$ for all $x \in X$ and $n \in \mathbb{N}$) such that $\lim_{n\to\infty} f_n(x) = g(x)$ for each $x \in X$.*

Definition 26 (Integral of non-negative measurable function)**.** Let $g: X \to \mathbb{R}_{\geqslant 0}$ be a measurable function taking non-negative values, and let $(f_n)_{n\in\mathbb{N}}$ be a pointwise increasing sequence of simple functions converging to g as guaranteed by Proposition 25. Then the *integral* of g is defined to be

$$\int g \, \mathrm{d}\mu = \lim_{n\to\infty} \int f_n \, \mathrm{d}\mu.$$

Moreover, g is called *integrable* if $\int g \, \mathrm{d}\mu < \infty$.

Observe that the expression $\int g \, \mathrm{d}\mu$ defined above is guaranteed to exist since $f_n(x) \leqslant f_{n+1}(x)$ for all $n \in \mathbb{N}$ and $x \in X$. One can show that this is well-defined, i.e., that it is independent of the choice of the sequence of simple functions.

We now extend the notion of integral for any measurable functions.

Definition 27 (Integral of measurable function)**.** Given a measurable function $g: X \to \mathbb{R}$, g has in general a unique decomposition $g = g_+ - g_-$, where $g_+(x) = \max\{g(x), 0\}$ and $g_- = \max\{-g(x), 0\}$ for every $x \in X$. Note that both g_+ and g_- are measurable. The function g is said to be *integrable* if both g_+ and g_- are integrable, and the *integral* of g is defined as

$$\int g \, \mathrm{d}\mu = \int g_+ \, \mathrm{d}\mu - \int g_- \, \mathrm{d}\mu.$$

Consider now a measurable complex-valued function $g: X \to \mathbb{C}$, which we can decompose as $g = Re(g) + iIm(g)$, where both $Re(g)$ and $Im(g)$ are measurable. Then g is said to be integrable if $|g| = \sqrt{Re(g)^2 + Im(g)^2} \geqslant 0$ satisfies :

$$\int |g| \, \mathrm{d}\mu < +\infty$$

and the integral of g is then defined as :

$$\int g \, \mathrm{d}\mu = \int Re(g) \, \mathrm{d}\mu + i \int Im(g) \, \mathrm{d}\mu \in \mathbb{C}$$

so that we have :

$$Re\left(\int g \, \mathrm{d}\mu\right) = \int Re(g) \, \mathrm{d}\mu$$
$$Im\left(\int g \, \mathrm{d}\mu\right) = \int Im(g) \, \mathrm{d}\mu$$

Note that since we have $|Re(g)| \leqslant |g|$ and $|Im(g)| \leqslant |g|$, the condition $\int |g| \, \mathrm{d}\mu < +\infty$ implies that both $Re(g)$ and $Im(g)$ are integrable.

Here is a way of determining if a given function is integrable or not.

Proposition 28. *Let $f, g : X \to \mathbb{R}$. If f is integrable and g is measurable with $|g| \leqslant f$, then g is integrable.*

Being integrable is preserved under restriction to a measurable set, and we give the definition of the integrable restricted to a measurable set:

Definition 29. Let $f : X \to \mathbb{R}$ be an integrable function, and A be a measurable set. The integral of f over A is defined as

$$\int_A f \, \mathrm{d}\mu = \int \mathbf{1}_A f \, \mathrm{d}\mu.$$

0.4. L^p Spaces

We now recall some definitions and facts about L^p spaces, which are function spaces defined using a natural generalization of the p-norm for finite-dimensional vector spaces. L^p spaces form an important class of Banach spaces in functional analysis, and of topological vector spaces. In the course of Ergodic Theory we will use various results about functional analysis and in particular about L^p spaces. Further recalls about functional analysis are given in the next section.

Definition 30 ($\mathscr{L}^p$ spaces)**.** Let $(X, \mathscr{A}, \mu)$ be a measure space. For $1 \leqslant p < \infty$, we define the set $\mathscr{L}^p(X, \mathscr{A}, \mu)$ (sometimes also denoted $\mathscr{L}^p(\mu)$) to be the set of all measurable functions $f : X \to \mathbb{C}$ such that $\int |f|^p \, \mathrm{d}\mu < \infty$.

Definition 31 (L^p spaces)**.** We define an equivalence relation on $\mathscr{L}^p(\mu)$ by $f \sim g$ if $\int |f - g|^p \, \mathrm{d}\mu = 0$ and we write $L^p(\mu) = \mathscr{L}^p(\mu)/\sim$ for the space of equivalence classes. Elements of $L^p(\mu)$ will be described as functions rather than equivalence classes, but it is important to remember that this is an abuse of notation.
Furthermore, we define the norm $\|\cdot\|_{L^p}$ by:

$$\|f\|_{L^p} = \left(\int |f|^p \, \mathrm{d}\mu \right)^{1/p}$$

We now give the definition of the $L^p(\mu)$ in the case $p = \infty$.

Definition 32 (Essential supremum)**.** The *essential supremum* is the generalization to measurable functions of the supremum of a continuous function, and is defined by

$$\operatorname{ess\,sup} f = \inf\{\alpha \in \mathbb{R} : \mu(\{x \in X : f(x) > \alpha\} = 0)\}.$$

Let the norm $\|\cdot\|_{L^\infty}$ be given by

$$\|f\|_{L^\infty} = \operatorname{ess\,sup} |f|.$$

The space $\mathscr{L}^\infty(\mu)$ is then defined to be the set of measurable functions f such that $\|f\|_{L^\infty} < \infty$. Once again, $L^\infty(\mu)$ is defined to be $\mathscr{L}^\infty(\mu)/\sim$.

Proposition 33. *For every $1 \leqslant p \leqslant \infty$, the space $L^p(\mu)$ is complete with respect to the norm $\|\cdot\|_{L^p}$, and hence is a Banach space.*

Proposition 34. *For $1 \leqslant p < q \leqslant \infty$ we have $L^p(\mu) \supseteq L^q(\mu)$ for any finite measure space.*

Finally we turn to integration of functions of several variables.

Theorem 35 (Fubini–Tonelli)**.** *Let $(X, \mathscr{A}, \mu)$ and $(Y, \mathscr{B}, \nu)$ be two measure spaces and let f be a non-negative integrable function on the product space $(X \times Y, \mathscr{A} \otimes \mathscr{B}, \mu \otimes \nu)$. Then, for μ-almost every $x \in X$ the function $y \mapsto f(x,y)$ is integrable, and for ν-almost every $y \in X$ the function $x \mapsto f(x,y)$ is integrable, and we have*

$$\int_{X \times Y} f(x,y) d(\mu \times \nu)(x,y) = \int_X \left(\int_Y f(x,y)\, \mathrm{d}\nu(y) \right) \mathrm{d}\mu(x) = \int_Y \left(\int_X f(x,y)\, \mathrm{d}\mu(x) \right) \mathrm{d}\nu(y).$$

0.5. Convergence Theorems

The most important distinction between integration on L^p spaces as defined above and Riemann integration on bounded Riemann-integrable functions is that the L^p functions are closed under several natural limiting operations, allowing for the following important convergence theorems. We start with the Monotone Convergence Theorem.

Theorem 36 (Monotone Convergence Theorem)**.** *Suppose $f_1 \leqslant f_2 \leqslant \ldots$ is a pointwise increasing sequence of non-negative real-valued measurable functions on the measure space $(X, \mathscr{A}, \mu)$ which converges almost everywhere to a function f on X. Then f is measurable and*

$$\int f \, \mathrm{d}\mu = \lim_{n \to \infty} \int f_n \, \mathrm{d}\mu.$$

In particular, if $\lim_{n\to\infty} \int f_n \, \mathrm{d}\mu < \infty$, then f is integrable.

When the $f_n, n \in \mathbb{N}$, are integrable, the assumption that f_n is non-negative for every $n \in \mathbb{N}$ can be dropped by considering instead the non-negative sequence of measurable function $g_n = f_n - f_1$, which is also a pointwise increasing sequence.
Next, we state Fatou's lemma, which is not only needed to prove the dominated

convergence theorem below but it includes also a statement of the behavior of the integral under pointwise (or almost everywhere) convergence: The integral is lower semi-continuous under almost everywhere convergence.

Theorem 37 (Fatou's lemma – lim inf version). *Let $(f_n)_{n\in\mathbb{N}}$ be a sequence of non-negative real-valued measurable functions on the measure space $(X,\mathscr{A},\mu)$. Then, $f = \liminf_{n\to\infty} f_n$ is measurable and*

$$\liminf_{n\to\infty}\int f_n \, \mathrm{d}\mu \geqslant \int f \, \mathrm{d}\mu = \int \liminf_{n\to\infty} f_n \, \mathrm{d}\mu$$

In particular, if f_n is integrable for every $n \in \mathbb{N}$, then f is also integrable.

For completeness, we also add the reverse version of Fatou's lemma.

Corollary 38 (Fatou's lemma – lim sup version). *Let $(f_n)_{n\in\mathbb{N}}$ be a sequence of real-valued measurable functions on the measure space $(X,\mathscr{A},\mu)$. If there exists an integrable function g on X such that $f_n \leqslant g$, $\forall n \in \mathbb{N}$, then :*

$$\limsup_{n\to\infty}\int f_n \, \mathrm{d}\mu \leqslant \int \limsup_{n\to\infty} f_n \, \mathrm{d}\mu$$

Contrary to the Monotone Convergence Theorem, the hypothesis that f_n is non-negative for every $n \in \mathbb{N}$ cannot be dropped.
Finally, we state the Dominated Convergence Theorem, which formulates sufficient conditions under which almost everywhere convergence yields an integrable function and such that limit and integral are interchangeable. Note that this is an important difference with Riemann integral.

Theorem 39 (Dominated Convergence Theorem). *Let $(X,\mathscr{A},\mu)$ be a measure space. If $h\colon X \to \mathbb{R}$ is a non-negative integrable function, and $(f_n)_{n\in\mathbb{N}}$ is a sequence of measurable complex-valued functions on $(X,\mathscr{A},\mu)$ which are dominated by h in the sense that $|f_n| \leqslant h, \forall n \in \mathbb{N}$, and $\lim_{n\to\infty} f_n = f$ exists almost everywhere, then f is integrable and*

$$\int f \, \mathrm{d}\mu = \lim_{n\to\infty}\int f_n \, \mathrm{d}\mu.$$

Chapter 1

Measure Preserving Systems

1.1. Definition and Examples

Most of the material in this lecture notes is also contained, for instance, in [Wal82] and in [EW11].

Definition 40 (Measure preserving transformation)**.** Given a probability space $(X, \mathscr{A}, \mu)$, we say that a measurable map $T\colon X \to X$ *preserves the measure* or is a *measure preserving transformation* if for every $A \in \mathscr{A}$ we have $\mu(T^{-1}A) = \mu(A)$.

Recall that for any probability space $(X, \mathscr{A}, \mu)$ and any measurable map $T\colon X \to X$, the measure $T\mu$ defined via

$$T\mu(A) = \mu(T^{-1}A), \qquad \forall A \in \mathscr{A},$$

is a probability measure on $\mathscr{A}$ called the *push-forward of μ under T*. If $T\mu = \mu$, we say that the measure μ is *invariant* under the map T. This invariance implies that the map T does not change the measure of any measurable set, or in other words, for any $A \in \mathscr{A}$, we have $\mu(T^{-1}(A)) = \mu(A)$. Thus, saying that T preserving the measure μ (as defined in Definition 40) is equivalent to stating that μ is invariant under T; the two terms express the same property and and we will use them interchangeably throughout these lecture notes.

Example 41. Imagine a computer program with the capability to generate uniformly at random and without bias a real number x in the interval $[0,1) = \{x \in \mathbb{R} : 0 \leqslant x < 1\}$. Then there is a 50% chance that a number generated with this program lies in the interval $[0,1/2)$, and a 20% chance that the generated number lies in the interval $[3/5, 4/5)$, just as an example. Now consider a second, considerably simpler, program that receives as an input a real number $x \in [0,1)$ and produces as an output the number $y = 2x \bmod 1$. If you first run program number one to produce x and then apply program number two to "transform" x to y, then has this procedure still generated a “random” real number between 0 and 1? In particular, is there still a 50% chance for y to belong to the interval $[0,1/2)$, and a 20% chance for it to belong to $[3/5,4/5)$? The answer is yes! The first program produces a number chosen at random with respect to the Lebesgue measure on $[0,1)$ and the second program corresponds to the transformation $T\colon [0,1)\to[0,1)$ given by $T(x)=2x \bmod 1$. Since T preserves the Lebesgue measure on $[0,1)$, the second program does not introduce any subsidiary bias, meaning that the second number can also be thought of as chosen at random with respect to the Lebesgue measure on $[0,1)$.

Definition 42 (Measure preserving system)**.** A *measure preserving system* is a quadruple $(X,\mathscr{A},\mu,T)$ where $(X,\mathscr{A},\mu)$ is a probability space and $T\colon X\to X$ is a measure preserving transformation.

Examples

The following examples illustrate the above definitions and serve as a guide for the concepts and results presented throughout this course.

One point system. If $X=\{x\}$ is a singleton then there exist only one σ-algebra $\mathscr{A}$ and only one probability measure μ on X, namely $\mathscr{A}=\{\emptyset,\{x\}\}$ and $\mu(\emptyset)=0$ and $\mu(\{x\})$=1. Let $T:X\to X$ be the identity map. Then $(X,\mathscr{A},\mu,T)$ is a (rather trivial) measure preserving system, called the *one point system*.

Identity systems. Let $(X,\mathscr{A},\mu)$ be an arbitrary probability space and let $T=\mathrm{id}_X$ be the identity map on X. Since the push-forward of μ under id_X is always equal to μ, $(X,\mathscr{A},\mu,\mathrm{id}_X)$ is a measure preserving system. Systems of this kind are often referred to as *identity systems*.

Rotation on m points. Given an integer $m\geqslant 2$, let $X=\{0,1,\dots,m-1\}$, which we can identify with the finite cyclic group of order m. Let $\mathscr{A}$ be the power set of $\{0,1,\dots,m-1\}$ and let $T\colon\{0,1,\dots,m-1\}\to\{0,1,\dots,m-1\}$ be the map

$$T(x)=x+1 \bmod m.$$

Finally, let μ be the probability measure uniquely determined by $\mu(\{i\})=1/m$ for all $i=0,1,\dots,m-1$. The resulting measure-preserving system $(X,\mathscr{A},\mu,T)$ is called *rotation on m points*.

Circle rotations. Let $X=[0,1)$, endowed with the Borel σ-algebra $\mathscr{B}_{[0,1)}$ (see Definition 6 to recall the definition of Borel σ-algebra) and the Lebesgue measure μ. Given $\alpha\in\mathbb{R}$ we consider the map $T=T_\alpha\colon X\to X$ given by $Tx=x+\alpha \bmod 1$. The fact that T preserves the measure μ follows from the basic properties of Lebesgue measure.

Alternatively, we can identify the interval $[0,1)$ with the compact group $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ in the obvious way. The Lebesgue measure on $[0,1)$ gets identified with the Haar measure on $\mathbb{T}$, and T becomes the map $Tx = x + \tilde{\alpha}$ (where $\tilde{\alpha} = \alpha + \mathbb{Z} \in \mathbb{T}$). This map clearly preserves the Haar measure.

The reason to call this system a circle rotation is that the 1-dimensional torus $\mathbb{T}$ is isometrically isomorphic to the circle $S^1 \subseteq \mathbb{C}$, viewed as a group under multiplication. The map T under this identification becomes the rotation $T : z \mapsto \theta z$, where $\theta = e^{2\pi i \alpha} \in S^1$.

Compact group rotations. The previous two examples are special cases of so-called *group rotations*: If $(G,+)$ is a compact abelian group, endowed with the Borel σ-algebra $\mathscr{B}_G$ and the (normalized) Haar measure m_G, then for any fixed $\alpha \in X$, the map $R : x \mapsto x + \alpha$ preserves m_G and hence $(G, \mathscr{B}_G, m_G, R)$ is a measure preserving system.

The doubling map. The next example of a measure-preserving system is one that we have already encountered in Example 41 above. Take $(X, \mathscr{B}_X, \mu)$ to be the unit interval $[0,1)$ equipped with its Borel σ-algebra and Lebesgue measure. Let $T : X \to X$ be the *doubling map* $T(x) = 2x \bmod 1$. Let us show that this transformation preserves the measure: Given an interval $[a,b) \subseteq [0,1)$, the pre-image $T^{-1}([a,b))$ is the union of two intervals, each half the length of the original interval:

$$T^{-1}\big([a,b)\big) = \left[\frac{a}{2}, \frac{b}{2}\right) \cup \left[\frac{a+1}{2}, \frac{b+1}{2}\right).$$

This shows that the Lebesgue measure of $[a,b)$ and $T^{-1}([a,b))$ are identical. Since T^{-1} preserves the measure of all intervals and since intervals generate the Borel σ-algebra on $[0,1)$, it follows that T is a measure-preserving transformation.

More generally, for any positive integer p the map $T(x) = px \bmod 1$ preserves the Lebesgue measure, giving rise to a class of measure-preserving systems whose dynamical behavior can offer new insights on base-p digit expansions of the real numbers.

Toral endomorphisms and toral automorphisms. The transformations $T(x) = px \bmod 1$ for $p \in \mathbb{N}$ introduced in the previous example are 1-dimensional instances of so-called toral endomorphisms. For higher-dimensions, these are defined as follows. Given a matrix $A \in GL(n,\mathbb{Z})$, one can construct the measure preserving system $(X, \mathscr{A}, \mu, T)$, where $X = [0,1)^n$, $\mathscr{B}_X$ the Borel σ-algebra on $[0,1)^n$, μ the n-dimensional Lebesgue measure restricted to $[0,1)^n$, and T is defined by $Tx = Ax \bmod \mathbb{Z}^n$. Whenever $\det(A) \neq 0$, we call T a linear toral endomorphism.

Note that in general, A is not invertible in $GL(n,\mathbb{Z})$. However, if $\det(A) = \pm 1$ then A^{-1} exists, and belongs to $GL(n,\mathbb{Z})$. Such a matrix is called unimodular. In this case, T is said to be a toral automorphism, and its inverse transformation T^{-1} is given by $T^{-1}x = A^{-1}x \mod \mathbb{Z}^n$.

Arnold's cat map. In the case $n=2$, we define Arnold's cat map to be the toral automorphism where $A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \in GL(2,\mathbb{Z})$. The induced map is therefore given by $T(x,y) = (2x+y \mod 1, x+y \mod 1)$. It was named after Vladimir Arnold, who demonstrated its effects in the 1960s using an image of a cat, hence the name. Note that $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, that is, the square is sheared one unit up, then two units to the right, and all regions outside the unit square are reduced modulo $\mathbb{Z}^2$ to lie in the unit square. The following picture is showing how the linear map stretches the unit square and how its pieces are rearranged when the modulo operation is performed.

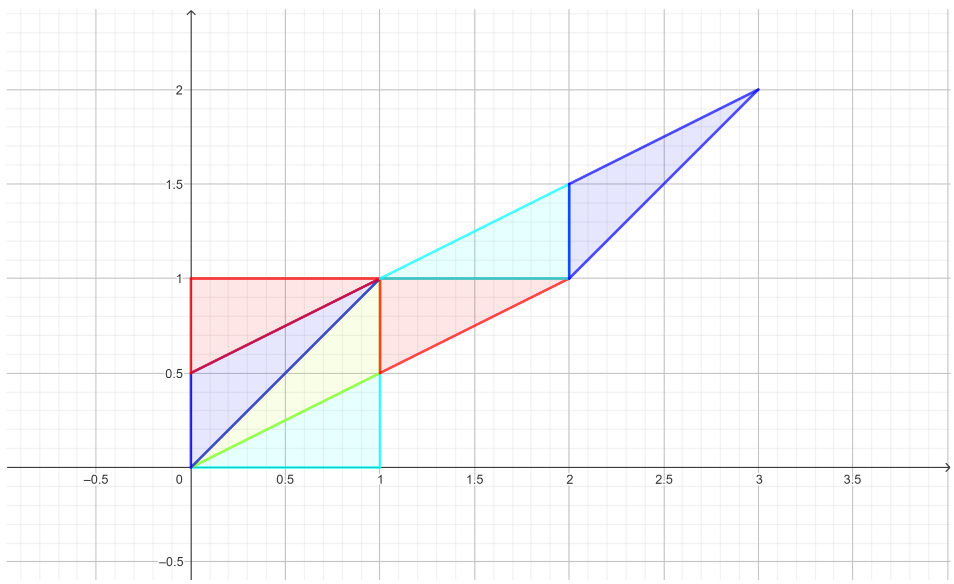

Figure 1.1: Visualization of the effect of Arnold's cat map on the unit square

A central concern of ergodic theory is the dynamical behavior of a measure preserving system when it is allowed to run for a long time, and one of the main object of study is the notion of periodicity, i.e the question of how and when orbits in dynamical systems return to their initial position. In this sense, Arnold's cat map is an interesting example as it exhibits various interesting properties based on periodicity. Indeed, a noticeable property is that for any $n \in \mathbb{N}$, the number of points with period n (returning to their initial position after n iterations) is exactly $|\lambda_1^n + \lambda_2^n - 2|$, where λ_1 and λ_2 are the eigenvalues of the matrix A. In fact, yhe set of points with a periodic orbit is dense on the torus. Actually, it can be shown that a point is periodic if and only if its coordinates are rational.

An interesting application of Arnold's cat map, and more generally, chaotic maps, is in the domain of image encryption. Indeed, instead of a torus, we consider an $N \times N$ pixels picture and the following sequence :

$$\begin{pmatrix} x_{n+1} \\ y_{n+1} \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x_n \\ y_n \end{pmatrix} \mod N$$

which describes the position of a given pixel after n iteration, where initially we pick $x_0, y_0 \in \{0, 1, \ldots, N-1\}$. One of this map's features is that when iteratively applied to an image, the result apparently looks randomized in a first place, but it always returns to its initial state after a number of steps depending on the size of the image. As it can be seen in the picture below, the original image of the cat is sheared and then wrapped around in the first iterations of the transformation. After some iterations, the various pixels of the original picture appear rather mixed together in a random manner, yet at various iterations, we can somewhat distinguish multiple smaller appearances of the cat arranged in a repeating structure, and it ultimately returns to the original image.

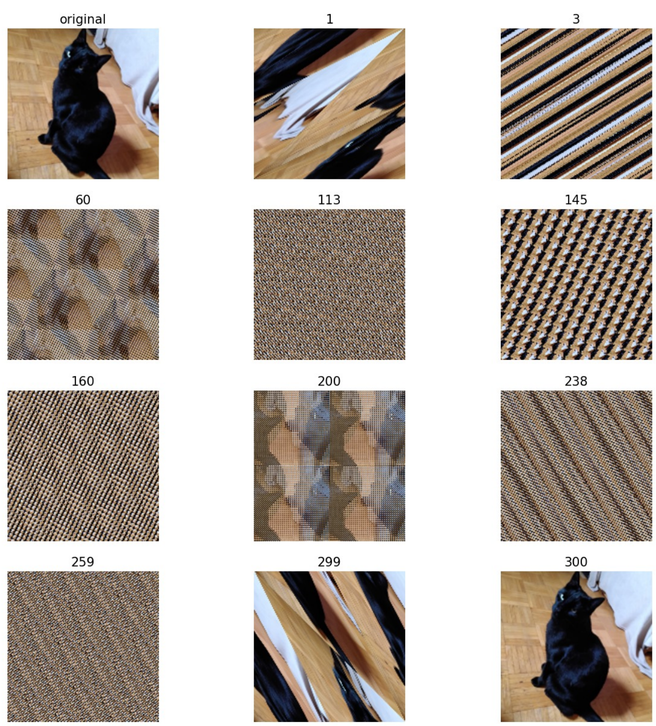

Figure 1.2: Visualization of the effect of Arnold's cat map on the unit square

Bernoulli schemes. Let $X = \{0,1\}^{\mathbb{N}}$ be the space of all (one-sided) infinite strings of 0's and 1's. Giving $\{0,1\}$ the discrete topology, we can endow X with the product topology[1]. In view of Tychonoff's theorem, X is compact. Let $\mathscr{B}_X$ be the Borel σ-algebra on X generated by the cylinder sets (see example 7 for the definition). Given $p \in (0,1)$, let μ_0 be the measure on $\{0,1\}$ given by $\mu_0(\{1\}) = p$ and $\mu_0(\{0\}) = 1-p$, and let $\mu = \mu_0^{\mathbb{N}}$ be the product probability measure on X already defined in the first chapter. There is a natural map $T\colon X \to X$ that preserves this measure μ, called the *left-shift*: For $(x_n)_{n=1}^\infty \in X$ define $T((x_n)_{n=1}^\infty) = (y_n)_{n=1}^\infty$ where $y_n = x_{n+1}$ for all $n \in \mathbb{N}$. The resulting measure preserving system $(X, \mathscr{B}_X, \mu, T)$ appears naturally in symbolic dynamics and is related to so-called *Bernoulli processes* in probability and statistics.

Instead of sequences consisting of 0's and 1's, one can also consider sequences using elements from any other alphabet Σ. In general, a measure preserving system is called a *Bernoulli scheme* if it is of the form $(X, \mathscr{B}_X, \mu, T)$ where $X = \Sigma^{\mathbb{N}}$, $\mathscr{B}_X$ is the σ-algebra of Borel sets on X generated by cylinder sets, T is the left shift and $\mu = \mu_0^{\mathbb{N}}$ is the product measure of some arbitrary probability measure μ_0 on Σ.

Baker's transformation. This example offers another way of generalizing the doubling map to two dimensions. Consider the probability space $(X, \mathscr{B}_X, \mu)$, where $X = [0,1)^2$ is the unit square, $\mathscr{B}_X$ is the Borel σ-algebra on X and μ is the two-dimensional Lebesgue measure. We define the Baker's map $T\colon [0,1)^2 \to [0,1)^2$ by

$$T(x,y) = \begin{cases} (2x, \frac{y}{2}), & \text{for } 0 \leqslant x < \frac{1}{2},\ 0 \leqslant y < 1, \\ (1-2x, 1-\frac{y}{2}), & \text{for } \frac{1}{2} \leqslant x < 1,\ 0 \leqslant y < 1. \end{cases}$$

Then, T is an invertible, measurable and measure preserving transformation.
We can define an analogous map $S : X \to X$ by :

$$S(x,y) = \begin{cases} (3x, \frac{y}{3}), & \text{for } 0 \leqslant x < \frac{1}{3},\ 0 \leqslant y < 1, \\ (2-3x, \frac{2-y}{3}), & \text{for } \frac{1}{3} \leqslant x < \frac{2}{3},\ 0 \leqslant y < 1, \\ (3x-2, \frac{y+2}{3}), & \text{for } \frac{2}{3} \leqslant x < 1,\ 0 \leqslant y < 1, \end{cases}$$

which is also invertible, measurable and a measure preserving transformation. To visualize what the map S does on the unit square, one can see that it represents the process of making the well-known French delicacy puff pastry, used in croissants and various other pastries. The idea is as follows: you have a piece of dough (represented by the unit square) with the lower half being the dough and the upper half being the butter. You then stretch the dough by 3 times its original length and consider the dough as composed of 3 parts, each of length one. We then fold it just as a baker would do it (hence the name Baker's transformation), namely, we put the second part on top of the first part, and the third part on top of everything, without cutting the dough, and finally, we compress the result in order to get back the unit square.

[1]This means that a set $U \subseteq X$ is open iff for every $x \in U$ there exists $n \in \mathbb{N}$ such that if $y \in X$ satisfies $y_i = x_i$ for all $i \leqslant n$, then $y \in U$.

This process is a chaotic map from the unit square into itself and it has this very niece property that it will efficiently mix the dough and the butter in order to form a very homogeneous buttered dough. In ergodic theory, we call this phenomenon strong mixing, which will be covered in Chapter 6.

Product systems. One way to construct new measure preserving systems out of given ones is by taking their product. Given two measure preserving systems $(X,\mathscr{A},\mu,T)$ and $(Y,\mathscr{B},\nu,S)$, we define their *product* to be the measure preserving system $(Z,\mathscr{C},\lambda,R)$, where $(Z,\mathscr{C},\lambda)=(X\times Y,\mathscr{A}\otimes\mathscr{B},\mu\times\nu)$ is the product of the probability spaces $(X,\mathscr{A},\mu)$ and $(Y,\mathscr{B},\nu)$ (see Example 5 for the definition of product algebra, and page 10 for product measure), and $R\colon Z\to Z$ is defined as $R(x,y)=(Tx,Sy)$.

Skew-products. Let $X=[0,1)^2$, let $\mathscr{B}_X$ be the Borel σ-algebra and let μ be the Lebesgue measure. Fix $\alpha\in\mathbb{R}$ and let $T\colon X\to X$ be the map $T(x,y)=(x+\alpha \bmod 1, y+x \bmod 1)$. Then $(X,\mathscr{B}_X,\mu,T)$ is a measure preserving system called a *skew-product*.

To see why T preserves the measure, observe that it suffices to check that for any $f,g\in C([0,1))$

$$\int_0^1\int_0^1 f(x+\alpha \bmod 1)g(x+y \bmod 1)\,\mathrm{d}y\,\mathrm{d}x=\int_0^1\int_0^1 f(x)g(y)\,\mathrm{d}y\,\mathrm{d}x,$$

which can be verified directly.

1.2. Recurrence

At the end of the XIX'th century, the french mathematician Henry Poincaré put an end to a myth acquired since Newton, that the universe is deterministic in the sense that knowing the past uniquely determines the future. Newton perfectly described the action of gravitational forces between two celestial bodies, but these laws don't apply as well to systems with more than two bodies.

It is in this context that Poincaré, in his work, considered systems with 3 celestial bodies. Newton's equations applied at these 3 bodies lead to a very complex differential equation that cannot be solved. He showed that in the special case where one body has zero mass, and the other two have a circular movement, then, the three bodies will eventually return infinitely many times to their original position. This initial observation led to the statement of Poincaré's Recurrence Theorem, which was proved 30 years later by Carathéodory using measure theory.

Here is the first theorem of ergodic theory.

Poincaré's Recurrence Theorem. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and let $A\in\mathscr{A}$ with $\mu(A)>0$. Then for some $n\in\mathbb{N}$ we have*

$$\mu(A\cap T^{-n}A)>0. \tag{1.2.1}$$

Proof. Since T is measure preserving, for any $n\in\mathbb{N}$ the set $T^{-n}A$ has the same measure as the set A. Since the ambient space X has measure 1 and $A, T^{-1}A, T^{-2}A,\ldots$ is an infinite sequence of sets with the same (positive) measure, by the pigeonhole principle there must exist $i>j$ with $\mu(T^{-i}A\cap T^{-j}A)>0$. Letting $n=i-j$, we obtain

$$\mu(A\cap T^{-n}A)=\mu\big(T^{-j}(A\cap T^{-n}A)\big)=\mu(T^{-i}A\cap T^{-j}A)>0.$$

□

Corollary 43. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and let $A\in\mathscr{A}$. Then for μ-a.e. $x\in A$ there exists $n\in\mathbb{N}$ such that $T^n x\in A$, i.e. x returns to A at time n.*

Proof. Let $B:=\{x\in A: T^n x\notin A \text{ for all } n\in\mathbb{N}\}$; we need to show that $\mu(B)=0$. If $\mu(B)>0$, then by Poincaré's Recurrence Theorem one can find $m\in\mathbb{N}$ such that $B\cap T^{-m}B$ has positive measure and, in particular, is non-empty. But if $y\in B\cap T^{-m}B$ then $T^m y\in B\subseteq A$, contradicting the fact that $y\in B$. This contradiction implies $\mu(B)=0$. □

Poincaré's Recurrence Theorem and its many generalizations, variations, and applications, form a sub-field of ergodic theory called the *theory of recurrence*. Broadly speaking, it focuses on the question of when and how close orbits in dynamical systems return to their initial position. The recurrence properties of measure preserving systems can provide important information about their dynamical behavior. Also, as we will discover in this course, there exist remarkable synergies between the theory of recurrence and problems in number theory and additive combinatorics.

1.3. Ergodicity

Poincaré's Recurrence Theorem asserts that the orbit $x, Tx, T^2x,\ldots$ of a typical point $x\in X$ returns to its initial location. But it doesn't provide any further information about the distribution of the orbit within the space. This is where the notion of ergodicity comes into play.

The word *ergodic* is derived from Ludwig Boltzmann's 'ergodic hypothesis' in thermodynamics, which describes a Hamiltonian system[2] with the property that the time spent in a certain region of the space is proportional to the spacial volume of that region. In the language of measure preserving systems, this means that the

[2]As an example of a Hamiltonian system, the reader can consider the *Lorentz gas model* commonly used to describe the kinetic movements of gas molecules in a chamber.

amount of time that an orbit $x, Tx, T^2x, T^3x, \ldots$ of a typical point $x \in X$ spends in a certain measurable set is proportional to the measure of that set. For example, if A has measure 1/2 then, asymptotically, half of all $n \in \mathbb{N}$ satisfy $T^n x \in A$. What we have just described is in fact the conclusion of *Birkhoff's Pointwise Ergodic Theorem*, one of the fundamental results in ergodic theory (discussed in Chapter 4) and equivalent to ergodicity.

Although Boltzmann initially conjectured that all naturally occurring systems satisfy the ergodic hypothesis, it was shown by John von Neumann that this is not the case, which is why today we distinguish between ergodic and non-ergodic systems.

Definition 44 (Ergodicity)**.** A measure preserving system $(X, \mathscr{A}, \mu, T)$ is *ergodic* if for every set $A \in \mathscr{A}$,

$$T^{-1}A = A \implies \mu(A) = 0 \text{ or } \mu(A) = 1.$$

Henceforth, let us call a set $A \in \mathscr{A}$ *strictly invariant* if $T^{-1}A = A$ and *almost everywhere invariant* if $\mu(A \triangle T^{-1}A) = 0$. Similarly, we call a measurable function $f\colon X \to \mathbb{C}$ *strictly invariant* if $f(Tx) = f(x)$ for all $x \in X$ and *almost everywhere invariant* if $f(Tx) = f(x)$ for μ-a.e. $x \in X$.

The next proposition provides four equivalent characterizations of the notion of ergodicity.

Proposition 45. *Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system. The following are equivalent:*

(i) *$(X, \mathscr{A}, \mu, T)$ is ergodic;*
(ii) *If $A \in \mathscr{A}$ is almost everywhere invariant then either $\mu(A) = 0$ or $\mu(A) = 1$;*
(iii) *If $f\colon X \to \mathbb{C}$ is measurable and strictly invariant then f is equal to a constant almost everywhere.*
(iv) *If $f\colon X \to \mathbb{C}$ is measurable and almost everywhere invariant then f is equal to a constant almost everywhere.*

Proof. The implication (ii) $\implies$ (i) is trivial. The reverse implication (i) $\implies$ (ii) follows readily from the observation that if $A \in \mathscr{A}$ is almost everywhere invariant then the set $A' = \bigcup_{m=0}^{\infty} \bigcap_{j=m}^{\infty} T^{-j}A$ is strictly invariant and satisfies $\mu(A) = \mu(A')$.

The implications (iv) $\implies$ (iii) $\implies$ (i) also do not require a proof, since they are immediate. To prove (iii) $\implies$ (iv), let $f\colon X \to \mathbb{C}$ be a measurable and almost everywhere invariant function. Let $A_f = \{x \in X : f(Tx) = f(x)\}$ and note that A_f has full measure and is almost everywhere invariant. Therefore the set $A_f' = \bigcup_{m=0}^{\infty} \bigcap_{j=m}^{\infty} T^{-j}A_f$ also has full measure and is strictly invariant. Now the function

$$f'(x) = \begin{cases} f(x), & \text{if } x \in A_f' \\ 0, & \text{otherwise} \end{cases}$$

is strictly invariant and almost everywhere equal to f. By (iii) it follows that f' is almost everywhere equal to a constant, which implies that f is almost everywhere equal to a constant.

Finally, let us prove (i) $\Longrightarrow$ (iii). Suppose $f : X \to \mathbb{R}$ is a measurable and strictly invariant function. Recall (see Definition 32) that the essential supremum of a measurable function is defined as

$$\operatorname{ess\,sup} f = \inf\{\alpha \in \mathbb{R} : \mu(\{x \in X : f(x) > \alpha\} = 0)\}.$$

For any $\alpha < \operatorname{ess\,sup} f$ consider the set $A_\alpha = \{x \in X : f(x) < \alpha\}$ and observe that if f is strictly invariant then A_α is strictly invariant. Note that A_α cannot have full measure, because α is smaller than the essential supremum of f. Therefore, in light of (i), A_α must have zero measure. But if A_α has zero measure for all $\alpha < \operatorname{ess\,sup} f$ then this implies that f is almost everywhere equal to $\operatorname{ess\,sup} f$, finishing the proof. Whenever $f : X \to \mathbb{C}$ is complex-valued, we can decompose $f = Re(f) + iIm(f)$, where both $Re(f)$ and $Im(f)$ are measurable and strictly invariant real-valued functions. We can therefore apply the above argument to both of them to deduce that they are equal to a constant almost everywhere, and deduce that f is equal to a (complex) constant almost everywhere. □

Examples

Finite systems. Let $X := \{1, \ldots, n\}$ be a finite of cardinality n, let $\mathscr{A} = \mathscr{P}(X)$, and let μ be the normalized counting measure on X, that is,

$$\mu(A) = \frac{|A|}{|X|}, \qquad \forall A \subseteq X.$$

Then $(X, \mathscr{A}, \mu)$ is a finite probability space. A map $T : X \to X$ preserves the measure μ if and only if it is a bijection from X to X. In other words, T is a permutation. Moreover, T is ergodic if, and only if it has only one orbit, that is, for every $x, y \in X$, there exists $k \in \mathbb{N}$ such that $y = T^k x$. For instance, the cycle $(1\ 2\ \ldots\ n)$ constitutes an ergodic transformation on $\{1, \ldots, n\}$, since the invariant subsets are $\emptyset$ and X. On the other hand, the permutation $(1\ 2)(3\ \ldots\ n)$ is not ergodic since the sets $\{1, 2\}$ and $\{3, \ldots, n\}$ are invariant subsets which have measure $\frac{2}{n}$ and $\frac{n-2}{n}$ respectively.

Circle rotations. Consider the probability space $(X, \mathscr{A}, \mu)$ where $X = [0, 1)$, $\mathscr{A}$ is the Borel σ-algebra on X, and μ is the Lebesgue measure. Given $\alpha \in \mathbb{R}$, consider the rotation by alpha $T : X \to X$ defined by $Tx = x + \alpha \mod 1$. We already argued that T is a measure preserving transformation. Now we can ask ourselves the following question: Is T ergodic?

As motivation, we can first consider the case when $\alpha = 1/4$. Observe that the set $A = [0, 1/8) \cup [1/4, 3/8) \cup [1/2, 5/8) \cup [3/4, 7/8)$ is T-invariant and satisfies $\mu(A) = 1/2$; this implies that the transformation is not ergodic.

More generally, one can show that T is ergodic if, and only if α is irrational.

Doubling map. Consider the probability space $(X, \mathscr{A}, \mu)$ where $X = [0,1)$, $\mathscr{A}$ is the Borel σ-algebra on X, and μ is the Lebesgue measure. This time, consider the doubling map $Tx = 2x \bmod 1$. It is left as an exercise to show that T is ergodic. More generally, this results also holds for non-integer values > 1. Even more generally, one can show that this results still holds for the product probability space $[0,1)^2, \mathscr{A} \otimes \mathscr{A}, \mu \otimes \mu)$, and the map $T \times T(x,y) = (px \mod 1, py \mod 1)$.
Finally, using multi-dimensional Fourier analysis, we can find an analogous result for toral endomorphisms over n-toruses $\mathbb{T}^n = (\mathbb{R}/\mathbb{Z})^n$. We have already seen that any $A \in GL_n(\mathbb{Z})$ induces a map $T_A : \mathbb{T}^n \to \mathbb{T}^n$ preserving the Lebesgue measure induced on $\mathbb{T}^n$. A well-known result is that T_A is ergodic if, and only if, no eigenvalue of A is a root of unity.

Induced transformation. Let $(X, \mathscr{A}, \mu)$ be a probability space and $T : X \to X$ a measure preserving transformation on it. Fix some $A \in \mathscr{A}$ with $\mu(A) > 0$. In light of Poincaré's Recurrence Theorem, it follows that almost every $x \in A$ returns infinitely often to A under the action of T. For every $x \in A$ we define $n(x) := \inf\{n \in \mathbb{N} : T^n x \in A\}$ to be the first return time of x to A.
By Poincaré's Recurrence Theorem, $n(x)$ is finite for almost every $x \in A$, hence, without loss of generality, we can assume that we remove the set of measure zero on which $n(x) = \infty$ and call the new set A. Consider the σ-algebra on $\mathscr{A}|_A$, which consists of the restriction of $\mathscr{A}$ on A, i.e $\mathscr{A}|_A := \{B \cap A : B \in \mathscr{A}\}$. We now define $\mu|_A$ to be the probability measure on A defined by :

$$\mu|_A(B) = \frac{\mu(B)}{\mu(A)}, \qquad \forall B \in \mathscr{A}|_A.$$

Hence, $(A, \mathscr{A}|_A, \mu|_A)$ is a probability space.
Finally, define the map $T_A : A \to A$ by $T_A x = T^{n(x)} x$, for $x \in A$. Then, this map is measurable with respect to $\mathscr{A}|_A$ and is a measure preserving transformation.
Moreover, if T is ergodic on $(X, \mathscr{A}, \mu)$, then T_A is ergodic on $(A, \mathscr{A}|_A, \mu|_A)$. If we additionally add the assumption that $\mu\left(\bigcup_{k\geqslant 1} T^{-k} A\right) = 1$, then the converse is also true (i.e, T_A ergodic implies T ergodic).

Chapter 2

Von Neumann's Mean Ergodic Theorem

2.1. Hilbert Spaces

In order to work through this section, we will need some notions about Hilbert spaces, and bounded operators on Hilbert spaces. We start by recalling the definition of an inner product on a complex linear space X :

Definition 46. An inner product on a complex linear space X is a map :

$$\langle \cdot, \cdot \rangle : X \times X \to \mathbb{C}$$

such that, for all $x, y, z \in X$ and $\mu, \lambda \in \mathbb{C}$:

(i) $\langle \lambda x + \mu y, z \rangle = \lambda \langle x, z \rangle + \mu \langle y, z \rangle$ (linear in the first argument)
(ii) $\langle x, y \rangle = \overline{\langle y, x \rangle}$ (Hermitian symmetric)
(iii) $\langle x, x \rangle \geqslant 0$, with equality if, and only if $x = 0$ (positive definite)

We call a linear space together with an inner product an inner product space, or a pre-Hilbert space.

Notice that according to the first two properties of the definition, we must have $\langle x, \lambda y + \mu z \rangle = \overline{\lambda} \langle x, y \rangle + \overline{\mu} \langle x, z \rangle$, for all $x, y, z \in X$ and $\mu, \lambda \in \mathbb{C}$. This inner product induces a norm on X defined for all $x \in X$ as $\|x\| = \sqrt{\langle x, x \rangle}$, so that any inner product space is a normed linear space.

Definition 47. An inner product space which is complete with respect to the induced norm is called a Hilbert space.

We then give the definition of a bounded linear operator on a Hilbert space:

Definition 48. Let $\mathscr{H}$ be a Hilbert space. A map $U: \mathscr{H} \to \mathscr{H}$ is a bounded linear operator if it is linear (that is, $U(\lambda x + y) = \lambda U(x) + U(y)$, for all $x, y \in \mathscr{H}$, $\lambda \in \mathbb{C}$), and there exists a constant $c > 0$ such that $\|Ux\| \leqslant c\|x\|$, $\forall x \in \mathscr{H}$, where $\|\cdot\|$ is the induced norm on $\mathscr{H}$.

As an inner product can be viewed as the abstract information of an angle between vectors, it is natural to define the notion of orthogonality in Hilbert spaces :

Definition 49. Let $\mathscr{H}$ be a Hilbert space, and $A, B \subseteq \mathscr{H}$. Then :

(i) $x, y \in \mathscr{H}$ are orthogonal, written $x \perp y$, if $\langle x, y \rangle = 0$

(ii) We write $A \perp B$ if $x \perp y$ for all $x \in A$ and $y \in B$

(iii) The orthogonal complement of a subset A is $A^{\perp} := \{x \in \mathscr{H} : x \perp y, \forall y \in A\}$

Finally, we give the definition of direct sum :

Definition 50. Let $\mathscr{H}$ be a Hilbert space, and $A, B \subseteq \mathscr{H}$ two closed subspaces such that $A \perp B$. Then, the direct sum of A with B, noted $A \oplus B \subseteq \mathscr{H}$, is the subspace given by all the elements of the form $x + y$ where $x \in A$ and $y \in B$.

2.2. Koopman Operator

Definition 51 (Koopman operator)**.** Given a measure preserving transformation $T: X \to X$ on a probability space $(X, \mathscr{A}, \mu)$, we call the linear operator $U_T : L^2(X, \mathscr{A}, \mu) \to L^2(X, \mathscr{A}, \mu)$ given by

$$U_T f = f \circ T$$

the associated *Koopman operator*.

The Koopman operator is well defined because T preserves the measure μ and therefore composition with T preserves measure-zero equivalency classes and square-integrability.

Lemma 52. *The operator U_T is an isometry, which means $\langle U_T f, U_T g \rangle = \langle f, g \rangle$ for all $f, g \in L^2(X, \mathscr{A}, \mu)$. In particular, $\|U_T f\|_{L^2} = \|f\|_{L^2}$ for all $f \in L^2(X, \mathscr{A}, \mu)$.*

Proof. Let $f, g \in L^2(X, \mathscr{A}, \mu)$. Since T preserves the measure, we have

$$\langle U_T f, U_T g \rangle = \int_X f(Tx)\overline{g(Tx)}\, \mathrm{d}\mu(x) = \int_X f(x)\overline{g(x)}\, \mathrm{d}\mu(x) = \langle f, g \rangle,$$

which proves that U_T is isometric. □

More generally, for any bounded linear operator $U: \mathscr{H}_1 \to \mathscr{H}_2$ from one Hilbert space to another, we say that U is an *isometry* if $\langle Uf, Ug \rangle_{\mathscr{H}_2} = \langle f, g \rangle_{\mathscr{H}_1}$ for any $f, g \in \mathscr{H}_1$, where $\langle \cdot, \cdot \rangle_{\mathscr{H}_1}$ and $\langle \cdot, \cdot \rangle_{\mathscr{H}_2}$ are the inner products of $\mathscr{H}_1$ and $\mathscr{H}_2$ respectively.

If additionally U is a Hilbert space isomorphism (i.e invertible), we say that U is *unitary*. Thus, for any measure preserving transformation T, the associated Koopman operator U_T is unitary whenever T is invertible, and is always an isometry by the previous lemma.

2.3. The Splitting $\mathscr{H}_{\text{inv}} \oplus \mathscr{H}_{\text{erg}}$

Henceforth, we denote by $\mathscr{H}_{\text{inv}}$ the space of almost everywhere invariant functions in $L^2(X, \mathscr{A}, \mu)$,

$$\mathscr{H}_{\text{inv}} = \{f \in L^2(X, \mathscr{A}, \mu) : U_T f = f\}.$$

In view of Proposition 45, the system $(X, \mathscr{A}, \mu, T)$ is ergodic if and only if $\mathscr{H}_{\text{inv}}$ consists only of almost everywhere constant functions.

A function $f \in L^2(X, \mathscr{A}, \mu)$ is called a *coboundary* if it satisfies the coboundary equation

$$f = g - g \circ T \tag{2.3.1}$$

for some $g \in L^2(X, \mathscr{A}, \mu)$. Note that the set of all coboundaries forms a subspace of $L^2(X, \mathscr{A}, \mu)$, but not a closed subspace. Let $\mathscr{H}_{\text{erg}}$ denote its closure,

$$\mathscr{H}_{\text{erg}} = \overline{\{f \in L^2(X, \mathscr{A}, \mu) : f \text{ is a coboundary}\}}. \tag{2.3.2}$$

Note that $\mathscr{H}_{\text{inv}}$ and $\mathscr{H}_{\text{inv}}$ are both invariant subspaces of $L^2(X, \mathscr{A}, \mu)$ under U_T, by which we mean that $U_T \mathscr{H}_{\text{inv}} \subseteq \mathscr{H}_{\text{inv}}$ and $U_T \mathscr{H}_{\text{erg}} \subseteq \mathscr{H}_{\text{erg}}$. The first claim follows from the observation that if f is almost everywhere invariant, then so is $U_T f$, and the second claim follows because if f is a coboundary then so is $U_T f$.

The following result says that $\mathscr{H}_{\text{erg}}$ is the orthocomplement of $\mathscr{H}_{\text{inv}}$.

Theorem 53. *We have $\mathscr{H}_{\text{inv}} \perp \mathscr{H}_{\text{erg}}$ and $\mathscr{H}_{\text{inv}} \oplus \mathscr{H}_{\text{erg}} = L^2(X, \mathscr{A}, \mu)$.*

Proof. For notational convenience, let us write $\mathscr{C}$ for the set $\{f \in L^2(X, \mathscr{A}, \mu) : f \text{ is a coboundary}\}$. It suffices to show $\mathscr{C}^\perp = \mathscr{H}_{\text{inv}}$, because this implies that the closure of $\mathscr{C}$ coincides with the orthocomplement of $\mathscr{H}_{\text{inv}}$, which by definition equals $\mathscr{H}_{\text{erg}}$. Let us first show $\mathscr{C}^\perp \subseteq \mathscr{H}_{\text{inv}}$. Suppose $f \in \mathscr{C}^\perp$, which simply means $\langle f, g \rangle = 0$ for all $g \in \mathscr{C}$. Then we have

$$\begin{aligned}
\|f - U_T f\|_{L^2}^2 &= \|f\|_{L^2}^2 + \|U_T f\|_{L^2}^2 - 2\text{Re}\langle f, U_T f\rangle \\
&= 2\|f\|_{L^2}^2 - 2Re\langle f, U_T f\rangle \\
&= 2\langle f, f\rangle - 2Re\langle f, U_T f\rangle \\
&= 2Re\langle f, f - U_T f\rangle = 0.
\end{aligned}$$

Hence $f \in \mathscr{H}_{\text{inv}}$ as was to be shown.

To prove the reverse inclusion $\mathscr{H}_{\mathrm{inv}} \subseteq \mathscr{C}^{\perp}$, we need to show $\langle f, h\rangle = 0$ for all $f \in \mathscr{C}$ and $h \in \mathscr{H}_{\mathrm{inv}}$. If $f \in \mathscr{C}$ then, by the definition of a coboundary, there exists some $g \in L^2(X, \mathscr{A}, \mu)$ for which $f = g - U_T g$ holds. Hence for any $h \in \mathscr{H}_{\mathrm{inv}}$ we have

$$\langle f, h\rangle = \langle g, h\rangle - \langle U_T g, h\rangle = \langle g, h\rangle - \langle U_T g, U_T h\rangle = \langle g, h\rangle - \langle g, h\rangle = 0,$$

showing that $h \in \mathscr{C}^{\perp}$ and finishing the proof. □

The following is an immediate corollary of Theorem 53.

Corollary 54. *For every $f \in L^2(X, \mathscr{A}, \mu)$ there exist unique $f_{\mathrm{inv}} \in \mathscr{H}_{\mathrm{inv}}$ and $f_{\mathrm{erg}} \in \mathscr{H}_{\mathrm{erg}}$ such that*

$$f = f_{\mathrm{inv}} + f_{\mathrm{erg}}. \tag{2.3.3}$$

Note that f_{inv} in (2.3.3) is exactly the orthogonal projection of f onto the space $\mathscr{H}_{\mathrm{inv}}$ and, likewise, f_{erg} is the orthogonal projection of f onto the space $\mathscr{H}_{\mathrm{erg}}$.

2.4. The Mean Ergodic Theorem

Here is Von Neumann's Mean Ergodic Theorem.

Mean Ergodic Theorem (General Case)**.** *Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system. For every $f \in L^2(X, \mathscr{A}, \mu)$ we have*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} U_T^n f = f_{\mathrm{inv}} \qquad \textit{in } L^2\textit{-norm}, \tag{2.4.1}$$

where f_{inv} is the orthogonal projection of f onto $\mathscr{H}_{\mathrm{inv}}$ as guaranteed by (2.3.3).

Proof. According to (2.3.3) we can write $f = f_{\mathrm{inv}} + f_{\mathrm{erg}}$. Hence

$$\frac{1}{N} \sum_{n=0}^{N-1} U_T^n f \;=\; \left(\frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\mathrm{inv}}\right) + \left(\frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\mathrm{erg}}\right).$$

Clearly, we have $\frac{1}{N}\sum_{n=0}^{N-1} U_T^n f_{\mathrm{inv}} = f_{\mathrm{inv}}$, because f_{inv} is invariant under U_T. Thus, to finish the proof of (2.4.1), it suffices to show

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\mathrm{erg}} = 0 \qquad \text{in } L^2\text{-norm} \tag{2.4.2}$$

for all $f_{\mathrm{erg}} \in \mathscr{H}_{\mathrm{erg}}$. In view of (2.3.2), we can assume that f_{erg} is a coboundary, i.e., there exists $g \in L^2(X, \mathscr{A}, \mu)$ such that $f_{\mathrm{erg}} = g - U_T g$. But if $f_{\mathrm{erg}} = g - U_T g$ then the sum in (2.4.2) is telescoping, yielding

$$\frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\mathrm{erg}} = \frac{U_T g - U_T^N g}{N}.$$

Since $U_T g - U_T^N g$ has norm at most $2\|g\|_{L^2}$, we obtain

$$\left\|\frac{1}{N}\sum_{n=0}^{N-1} U_T^n f_{\mathrm{erg}}\right\|_{L^2} \leqslant \frac{2\|g\|_{L^2}}{N}$$

and (2.4.2) follows. □

Mean Ergodic Theorem (Ergodic Case)**.** *Let $(X,\mathscr{A},\mu,T)$ be an ergodic measure preserving system. Then for every $f \in L^2(X,\mathscr{A},\mu)$ we have*

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1} U_T^n f = \int f \,\mathrm{d}\mu \tag{2.4.3}$$

in L^2-norm.

Proof. In light of (2.4.1), it suffices to show that if the system $(X,\mathscr{A},\mu,T)$ is ergodic then $f_{\mathrm{inv}} = \int f \,\mathrm{d}\mu$. So assume $(X,\mathscr{A},\mu,T)$ is ergodic and let $f \in L^2(X,\mathscr{A},\mu)$ be aritrary. According to Corollary 54, there exist unique $f_{\mathrm{inv}} \in \mathscr{H}_{\mathrm{inv}}$ and $f_{\mathrm{erg}} \in \mathscr{H}_{\mathrm{erg}}$ such that $f = f_{\mathrm{inv}} + f_{\mathrm{erg}}$. By definition, f_{inv} is an almost everywhere invariant function. Therefore, by part (iv) of Proposition 45, f_{inv} is almost everywhere equal to a constant, which we denote by c. To finish the proof of (2.4.3), it only remains to show that $\int f \,\mathrm{d}\mu = c$. Let $\mathbf{1}$ denote the function that is constant equal to 1 everywhere. Then

$$\int f \,\mathrm{d}\mu = \langle f,\mathbf{1}\rangle = \langle f_{\mathrm{inv}},\mathbf{1}\rangle + \langle f_{\mathrm{erg}},\mathbf{1}\rangle = c + \langle f_{\mathrm{erg}},\mathbf{1}\rangle.$$

Since $\mathbf{1}$ is invariant under the transformation T and f_{erg} is orthogonal to all invariant functions, we have $\langle f_{\mathrm{erg}},\mathbf{1}\rangle = 0$, showing that $\int f \,\mathrm{d}\mu = c$ as desired. □

2.5. Uniform Mean Ergodic Theorem

The mean ergodic theorem possesses a "uniform" version where the Cesàro averages $\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}$ are replaced by the more general uniform Cesàro averages $\lim_{N-M\to\infty}\frac{1}{N-M}\sum_{n=M}^{N-1}$. More precisely, we say that the uniform Cesàro average of a sequence $(u_n)_{n\in\mathbb{N}}$ (in a Hilbert space $\mathscr{H}$ with norm $\|.\|$) exists and equals u, and write

$$\lim_{N-M\to\infty}\frac{1}{N-M}\sum_{n=M}^{N-1} u_n = u,$$

if for all $\varepsilon > 0$ there exists $K \in \mathbb{N}$ such that for all $N,M \in \mathbb{N}$ with $N-M \geqslant K$ we have

$$\left\|\left(\frac{1}{N-M}\sum_{n=M}^{N-1} u_n\right) - u\right\| \leqslant \varepsilon.$$

Uniform Mean Ergodic Theorem. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system. For every $f\in L^2(X,\mathscr{A},\mu)$ we have*

$$\lim_{N-M\to\infty}\frac{1}{N-M}\sum_{n=M}^{N-1}U_T^n f=f_{\mathrm{inv}}\qquad \textit{in } L^2\textit{-norm,} \tag{2.5.1}$$

where f_{inv} is the orthogonal projection of f onto $\mathscr{H}_{\mathrm{inv}}$ as guaranteed by (2.3.3).

Proof. The proof of the Uniform Mean Ergodic Theorem is essentially identical to the proof of Mean Ergodic Theorem. One needs to replace all occurrences of Cesàro averages with uniform Cesàro averages, but otherwise the argument stays the same. □

2.6. Consequences of the Mean Ergodic Theorem

Corollary 55. *A measure preserving system $(X,\mathscr{A},\mu,T)$ is ergodic if and only if for every $A,B\in\mathscr{A}$,*

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\mu(T^{-n}A\cap B)=\mu(A)\mu(B). \tag{2.6.1}$$

Proof. If the system is not ergodic, then by definition there exists a strictly invariant set $A\in\mathscr{B}$ with $0<\mu(A)<1$. Taking $B=X\backslash A$, we see that $\mu(A)\mu(B)>0$ but $T^{-n}A\cap B=\emptyset$ for every n, contradicting (2.6.1).

If the system is ergodic then we proceed as follows. First observe that $\mathbf{1}_{T^{-n}A}=\mathbf{1}_A\circ T^n=U_T^n\mathbf{1}_A$. This implies $\mu(T^{-n}A\cap B)=\int U_T^n\mathbf{1}_A\cdot\mathbf{1}_B\,\mathrm{d}\mu$ and hence

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\mu(T^{-n}A\cap B)=\lim_{N\to\infty}\int\Big(\frac{1}{N}\sum_{n=0}^{N-1}U_T^n\mathbf{1}_A\Big)\cdot\mathbf{1}_B\,\mathrm{d}\mu.$$

By ergodicity, it follows from the Mean Ergodic Theorem that $\frac{1}{N}\sum_{n=0}^{N-1}U_T^n\mathbf{1}_A\to\mu(A)$ as $N\to\infty$ in L^2-norm. Since norm convergence in L^2 implies weak convergence in L^2, we get

$$\lim_{N\to\infty}\int\Big(\frac{1}{N}\sum_{n=0}^{N-1}U_T^n\mathbf{1}_A\Big)\cdot\mathbf{1}_B\,\mathrm{d}\mu=\int\mu(A)\cdot\mathbf{1}_B\,\mathrm{d}\mu=\mu(A)\mu(B),$$

completing the proof. □

Setting $A=B$ in Corollary 55 we see that, in ergodic systems, one can improve Poincaré's Recurrence Theorem by finding $n\in\mathbb{N}$ such that $\mu(T^{-n}A\cap A)$ is arbitrarily close to $\mu^2(A)$. One can in fact obtain a stronger version of this fact, which also applies to non-ergodic systems.

Definition 56. A set $S \subseteq \mathbb{N}$ is called *syndetic* if it has bounded gaps. More precisely, S is syndetic if there exists $L \in \mathbb{N}$ such that every interval $\{n, n+1, \dots, n+L-1\}$ of length L contains some element of S.

Khintchine's recurrence theorem. *Let $(X, \mathscr{B}, \mu, T)$ be a measure preserving system, let $A \in \mathscr{B}$ and let $\varepsilon > 0$. Then there exists $n \in \mathbb{N}$ such that $\mu(A \cap T^{-n}A) > \mu^2(A) - \varepsilon$, and moreover the set*

$$\big\{n \in \mathbb{N} : \mu(A \cap T^{-n}A) > \mu^2(A) - \varepsilon\big\}$$

is syndetic.

Proof. Suppose $\{n \in \mathbb{N} : \mu(A \cap T^{-n}A) > \mu^2(A) - \varepsilon\}$ is not syndetic. Then its complement contains arbitrarily long intervals, i.e., there exists a sequence of intervals $[M_k, N_k)$ with $N_k - M_k \to \infty$ as $k \to \infty$ and such that $\mu(A \cap T^{-n}A) \leqslant \mu^2(A) - \varepsilon$ for all $n \in [M_k, N_k)$. Applying Uniform Mean Ergodic Theorem to the indicator function $\mathbf{1}_A$ of A we have

$$\lim_{k \to \infty} \frac{1}{N_k - M_k} \sum_{n=M_k}^{N_k - 1} \mu(T^{-n}A \cap A) = \langle f_{\mathrm{inv}}, \mathbf{1}_A \rangle,$$

where f_{inv} is the orthogonal projection of $\mathbf{1}_A$ onto $\mathscr{H}_{\mathrm{inv}}$. Since it is an orthogonal projection, it follows that $\langle f_{\mathrm{inv}}, \mathbf{1}_A \rangle = \|f_{\mathrm{inv}}\|_{L^2}^2$. We now use the Cauchy-Schwarz inequality to get

$$\|f_{\mathrm{inv}}\|_{L^2}^2 \geqslant \left(\int f_{\mathrm{inv}} \, \mathrm{d}\mu \right)^2 = \langle f_{\mathrm{inv}}, \mathbf{1} \rangle^2 = \langle \mathbf{1}_A, \mathbf{1} \rangle^2 = \mu(A)^2.$$

It follows that

$$\lim_{k \to \infty} \frac{1}{N_k - M_k} \sum_{n=M_k}^{N_k - 1} \mu(T^{-n}A \cap A) \geqslant \mu(A)^2,$$

contradicting the assumption that $\mu(A \cap T^{-n}A) \leqslant \mu^2(A) - \varepsilon$ for all $n \in [M_k, N_k)$. $\square$

Chapter 3

Uniform Distribution of Sequences

3.1. Uniform Distribution Modulo 1

Definition 57. The *density* (sometimes also called the *natural density* or the *asymptotic density*) of a set $A \subseteq \mathbb{N}$ is defined as

$$d(A) = \lim_{N\to\infty} \frac{|A \cap \{1,\dots,N\}|}{N}$$

whenever this limit exists. If this limit does not exist then we say that the density of A does not exist.

Here are some examples of subsets of the natural numbers and their respective densities:

- $d(\mathbb{N}) = 1$;
- $d(2\mathbb{N}) = 0.5$;
- $d(\square\text{-free}) = \frac{6}{\pi^2}$, where $\square$-free denotes the set of squarefree numbers;
- $d(\mathbb{P}) = 0$, where $\mathbb{P}$ is the set of prime numbers.

Given a real number x we call $\lfloor x \rfloor = \max\{n \in \mathbb{Z} : n \leqslant x\}$ the *integer part* of x and $\{x\} = x - \lfloor x \rfloor$ the *fractional part* of x. Just as the interval $[0,1)$ is often identified with the (1-dimensional) torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, the map $x \mapsto \{x\}$, which sends a number to its fractional part, is often identified with the natural projection of $\mathbb{R}$ onto $\mathbb{T}$ given by $x \mapsto x \bmod 1$ (sometimes also written as $x \mapsto x \bmod \mathbb{Z}$).

Definition 58. We say a sequence of real numbers $(x_n)_{n\in\mathbb{N}}$ is *uniformly distributed mod* 1 if for every $0 \leqslant a \leqslant b \leqslant 1$ we have

$$\lim_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} = (b-a). \tag{3.1.1}$$

Remark 59. A sequence $(x_n)_{n\in\mathbb{N}}$ is uniformly distributed mod 1 if and only if for all $0 \leqslant a \leqslant b \leqslant 1$ the set $\{n \in \mathbb{Z} : \{x_n\} \in [a,b)\}$ has density $(b-a)$.

3.2. Weyl's Criterion

The following result gives necessary and sufficient conditions for a sequence to be uniformly distributed mod 1. We use $e(x)$ to abbreviate $e^{2\pi i x}$.

Weyl's Equidistribution Criterion. *Let $(x_n)_{n\in\mathbb{N}}$ be a sequence of real numbers. The following are equivalent:*

(i) $(x_n)_{n\in\mathbb{N}}$ is uniformly distributed mod 1;

(ii) For any continuous function $f\colon[0,1]\to\mathbb{C}$,

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} f(\{x_n\}) = \int_0^1 f(x)\,\mathrm{d}x;$$

(iii) For every $k\in\mathbb{Z}\setminus\{0\}$,

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} e(kx_n) = 0.$$

Proof of (i) $\Longrightarrow$ (ii). Suppose $(x_n)_{n\in\mathbb{N}}$ is uniformly distributed mod 1. Letting $\mathbf{1}_{[a,b)}$ denote the indicator function of the interval $[a.b)$, we can rewrite (3.1.1) as

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\mathbf{1}_{[a,b)}(\{x_n\}) = (b-a). \tag{3.2.1}$$

Let $f\colon[0,1]\to\mathbb{C}$ be continuous. Since continuous functions on compact sets are uniformly continuous, for every $\varepsilon>0$ there exists $M\in\mathbb{N}$ such that for all $x,y\in[0,1]$ we have

$$|x-y|\leqslant\tfrac{1}{M}\implies|f(x)-f(y)|\leqslant\varepsilon. \tag{3.2.2}$$

Let $y_j=\frac{j}{M}$, $j=0,1,\dots,M$, and define

$$f_M(x)=\sum_{j=0}^{M-1} f(y_j)\mathbf{1}_{[y_j,y_{j+1})}(x).$$

It follows from (3.2.2) that for any $x\in[0,1]$ we have $|f(x)-f_M(x)|\leqslant\varepsilon$. In particular, $|f(\{x_n\})-f_M(\{x_n\})|\leqslant\varepsilon$ for all $n\in\mathbb{N}$. Therefore

$$\limsup_{N\to\infty}\left|\frac{1}{N}\sum_{n=1}^{N} f(\{x_n\})-\frac{1}{N}\sum_{n=1}^{N} f_M(\{x_n\})\right|\leqslant\varepsilon. \tag{3.2.3}$$

Using (3.2.1) we see that

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} f_M(\{x_n\})=\sum_{j=0}^{M} f(y_j)(y_{j+1}-y_j).$$

Since the right hand side of the above equation is a (left) Riemann sum of f over the interval $[0,1]$ with respect to the partition induced by $y_0, y_1, \ldots, y_M$, we conclude that

$$\lim_{M\to\infty}\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} f_M(\{x_n\}) = \lim_{M\to\infty}\sum_{j=0}^{M} f(y_j)(y_{j+1}-y_j) = \int_0^1 f(x)\,\mathrm{d}x.$$

Therefore, taking the limit as $M\to\infty$ in (3.2.3) yields

$$\limsup_{N\to\infty}\left|\frac{1}{N}\sum_{n=1}^{N} f(\{x_n\}) - \int_0^1 f(x)\,\mathrm{d}x\right| \leqslant \varepsilon.$$

Since $\varepsilon > 0$ was chosen arbitrarily, this shows that the limit of $\frac{1}{N}\sum_{n=1}^{N} f(\{x_n\})$ as $N\to\infty$ exists and equals $\int_0^1 f(x)\,\mathrm{d}x$. □

Proof of (ii) $\Longrightarrow$ (iii). Observe that the function $x\mapsto e(kx)$ is continuous and for $k\neq 0$ we have $\int_0^1 e(kx)\,\mathrm{d}x = 0$. Since $e(k\{x_n\}) = e(kx_n)$ for all n, we see that (iii) follows from (ii) by choosing $f(x) = e(kx)$. □

For the proof of the implication (iii) $\Longrightarrow$ (i) we rely on a classical result from analysis. Given a topological space X, let $C(X)$ denote the space of all continuous functions from X to $\mathbb{C}$ and let $\|f\|_\infty = \sup_{x\in X}|f(x)|$ be the supremum norm.

Stone-Weierstrass Theorem. *Suppose X is a compact Hausdorff space and $\mathscr{A}$ is a subalgebra of $C(X)$ closed under complex conjugation and containing a non-zero constant function. Then $\mathscr{A}$ is dense in $C(X)$ (with respect to the supremum norm) if and only if it separates points.*

By a *trigonometric polynomial* on $[0,1]$ we mean any function of the form

$$x\mapsto c_1 e(k_1 x) + \ldots + c_\ell e(k_\ell x)$$

for $\ell\in\mathbb{N}$, $c_1,\ldots,c_\ell\in\mathbb{C}$, and $k_1,\ldots,k_\ell\in\mathbb{Z}$. The following is a well-known corollary of the Stone-Weierstrass Theorem.

Corollary 60. *Any continuous function $f\colon[0,1]\to\mathbb{C}$ satisfying $f(0)=f(1)$ can be approximated in supremum norm by trigonometric polynomials.*

Proof. By identifying the unit interval $[0,1)$ with the torus $\mathbb{T}=\mathbb{R}/\mathbb{Z}$, we can identify any continuous function $f\colon[0,1]\to\mathbb{C}$ satisfying $f(0)=f(1)$ with a continuous function on $\mathbb{T}$. In particular, we can view trigonometric polynomials as a functions on $\mathbb{T}$.

Note that the set of all trigonometric polynomials is closed under pointwise addition, pointwise multiplication, complex conjugation, and scalar multiplication. Therefore, it forms a subalgebra of $C(\mathbb{T})$ closed under complex conjugation. This subalgebra also contains all non-zero constant functions and separates points. Indeed, the former is obvious and the latter follows from the observation that the function $x \mapsto e(x)$ itself already separates points in $\mathbb{T}$, because for any $x, y \in [0,1)$ with $x \neq y$ one has $e(x) \neq e(y)$. It thus follows from the Stone-Weierstrass Theorem that any continuous function on $\mathbb{T}$ can be apprximated in supremum norm by trigonometric polynomials. Consequently, any continuous function $f\colon [0,1] \to \mathbb{C}$ satisfying $f(0) = f(1)$ can be approximated in supremum norm by trigonometric polynomials. □

Proof of (iii) $\implies$ (i). It suffices to show that for any $0 \leqslant a \leqslant b \leqslant 1$ one has

$$\liminf_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} \geqslant (b-a). \tag{3.2.4}$$

Indeed, assuming that (3.2.4) holds, we have

$$\frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} = 1 - \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [0,a)\}|}{N} - \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [b,1)\}|}{N}$$

and hence

$$\limsup_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} \leqslant 1-(a-0)-(1-b) = (b-a). \tag{3.2.5}$$

Then (3.2.4) and (3.2.5) together prove that $(x_n)_{n\in\mathbb{N}}$ is uniformly distributed mod 1.

For the proof of (3.2.4), let $\varepsilon > 0$ be arbitrary. By approximating $\mathbf{1}_{[a,b)}(x)$ from below, we can find a continuous function $f\colon [0,1] \to [0,1]$ supported on $[a,b)$ and with $\int_0^1 f(x)\,\mathrm{d}x \geqslant (b-a)-\varepsilon$. Without loss of generality, we can assume that $f(0) = f(1) = 0$. Using Corollary 60, we can now find a trigonometric polynomial $P(x) = c_1 e(k_1 x) + \ldots + c_\ell e(k_\ell x)$ such that $\|f-P\|_\infty \leqslant \varepsilon$. It follows that

$$\left|\int_0^1 f(x)\,\mathrm{d}x - \int_0^1 P(x)\,\mathrm{d}x\right| \leqslant \varepsilon \tag{3.2.6}$$

as well as

$$\left|\frac{1}{N}\sum_{n=1}^{N} f(\{x_n\}) - \frac{1}{N}\sum_{n=1}^{N} P(\{x_n\})\right| \leqslant \varepsilon, \qquad \forall N \in \mathbb{N}. \tag{3.2.7}$$

Using $\mathbf{1}_{[a,b)}(x) \geqslant f(x)$ for all $x \in [0,1]$, we have

$$\liminf_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} \geqslant \liminf_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} f(\{x_n\}). \tag{3.2.8}$$

Next, it follows from (iii) that for all $k \in \mathbb{Z}$ and $c \in \mathbb{C}$,

$$\lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} ce(k\{x_n\}) = \begin{cases} c, & \text{if } k = 0, \\ 0, & \text{otherwise.} \end{cases}$$

On the other hand, a straightforward calculation reveals

$$\int ce(kx)\,\mathrm{d}x = \begin{cases} c, & \text{if } k = 0, \\ 0, & \text{otherwise.} \end{cases}$$

This shows that for all $k \in \mathbb{Z}$ and $c \in \mathbb{C}$,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} ce(k\{x_n\}) = \int ce(kx)\, \mathrm{d}x.$$

Since $P(x) = c_1 e(k_1 x) + \ldots + c_\ell e(k_\ell x)$, it also follows that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} P(\{x_n\}) = \int_0^1 P(x)\, \mathrm{d}x. \tag{3.2.9}$$

Putting together (3.2.7) and (3.2.9), we get

$$\left| \liminf_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(\{x_n\}) - \int_0^1 P(x)\, \mathrm{d}x \right| \leqslant \varepsilon.$$

Combining this with (3.2.7) gives

$$\liminf_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(\{x_n\}) \geqslant \int_0^1 f(x)\, \mathrm{d}x - 2\varepsilon. \tag{3.2.10}$$

Finally, using $\int_0^1 f(x)\, \mathrm{d}x \geqslant (b-a) - \varepsilon$, it follows from (3.2.8) and (3.2.10) that

$$\liminf_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{x_n\} \in [a,b)\}|}{N} \geqslant (b-a) - 3\varepsilon.$$

Given that $\varepsilon > 0$ can be made arbitrarily small, (3.2.4) follows. □

Remark 61. In part (ii) of Weyl's Equidistribution Criterion, the assumption that the test function $f : [0,1] \to \mathbb{C}$ is continuous can be weakened to f being Riemann integrable. The equivalences remain valid, since continuity was only used to ensure convergence of the (left) Riemann sum, a property that holds by definition for all Riemann integrable functions. Thus the proof carries over verbatim.

The following theorem was proved in 1909 and 1910 separately by Hermann Weyl, Wacław Sierpiński and Piers Bohl, and variants of it continue to be studied to this day.

Weyl's Equidistribution Theorem. *For any irrational number α the sequence $(n\alpha)_{n\in\mathbb{N}}$ is uniformly distributed mod* 1.

Proof. In view of Weyl's Equidistribution Criterion, it suffices to show that for every $k \in \mathbb{Z}\setminus\{0\}$ we have

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} e(kn\alpha) = 0.$$

Taking $e(k\alpha) = \lambda$, we see that $e(kn\alpha) = \lambda^n$ and hence $\frac{1}{N}\sum_{n=1}^{N} e(kn\alpha) = \frac{1}{N}\sum_{n=1}^{N} \lambda^n$. Note also that $k\alpha$ is not an integer, because α is irrational, and hence $\lambda \neq 1$. Since

$\sum_{n=1}^{N} \lambda^n$ is a geometric sum, it can be calculated explicitly as

$$\sum_{n=1}^{N} \lambda^n = \lambda\Big(\frac{1-\lambda^N}{1-\lambda}\Big).$$

Therefore

$$\Big|\frac{1}{N}\sum_{n=1}^{N} e(kn\alpha)\Big| = \Big|\frac{1}{N}\sum_{n=1}^{N} \lambda^n\Big| = \Big|\frac{\lambda}{N}\Big(\frac{1-\lambda^N}{1-\lambda}\Big)\Big| \leqslant \frac{2}{N|1-\lambda|}.$$

Since the rightmost expression in the above equation converges to zero as $N \to \infty$, we are done. □

3.3. Benford's Law

A surprising phenomenon in data science is that the leading digits of many data sets are not uniformly distributed from 1 through 9, but rather exhibit a profound bias. For example, the first few elements in the sequence 2^n, $n = 0, 1, 2, \ldots$, are

$$\mathbf{1}, \mathbf{2}, \mathbf{4}, \mathbf{8}, \mathbf{1}6, \mathbf{3}2, \mathbf{6}4, \mathbf{1}28, \mathbf{2}56, \mathbf{5}12, \mathbf{1}024, \mathbf{2}048, \mathbf{4}096, \mathbf{8}192, \ldots$$

where the leading digits have been highlighted. As it turns out, the number 1 appears as the leading digit in this sequence about 30% of the time, while 9 appears as the leading digit less than 5% of the time. This phenomenon is called *Benford's law*.

Theorem 62. *For $k = 1, \ldots, 9$ we have*

$$\lim_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : 1^{\text{st}} \textit{ digit of } 2^n \textit{ equals } k\}|}{N} = \log_{10}(k+1) - \log_{10}(k).$$

Proof. Notice that the leading digit of 2^n equals k if and only if there exists $m \in \mathbb{N}$ such that $k10^m \leqslant 2^n < (k+1)10^m$, or equivalently, there exists $m \in \mathbb{N}$ such that

$$\log_{10}(k) \leqslant n\log_{10}(2) - m < \log_{10}(k+1).$$

This can only happen when $m = \lfloor n\log_{10}(2)\rfloor$, because $\log_{10}(k)$ and $\log_{10}(k+1)$ are numbers between 0 and 1. Hence the leading digit of 2^n equals k if and only if $\{n\log_{10}(2)\} \in [\log_{10}(k), \log_{10}(k+1))$. Since $\log_{10}(2)$ is irrational, the sequence $(n\log_{10}(2))_{N\in\mathbb{N}}$ is uniformly distributed mod 1, due to the Weyl's Equidistribution Theorem. We obtain

$$\lim_{N\to\infty} \frac{|\{1 \leqslant n \leqslant N : \{n\log_{10}(2)\} \in [\log_{10}(k), \log_{10}(k+1))\}|}{N} = \log_{10}(k+1) - \log_{10}(k)$$

as desired. □

3.4. Uniform Distribution in Metric Spaces

Recall, a *metric space* is a pair (X, d_X) where X is a set and $d_X : X \times X \to [0,\infty)$ is a function satisfying the following axioms of a *metric*:

- *(Positivity).* $x \neq y \iff d_X(x,y) > 0$.
- *(Symmetry).* $d_X(x,y) = d_X(y,x)$.
- *(Triangle inequality).* $d_X(x,z) \leqslant d_X(x,y) + d_X(y,z)$.

The *Borel σ-algebra*, denoted by $\mathscr{B}_X$, is the smallest σ-algebra on X containing all open balls in X. If X is a compact metric space then any Borel probability measure μ on X (i.e., any probability measure defined on the Borel σ-algebra $\mathscr{B}_X$) is a *Radon measure*, which means for all $A \in \mathscr{B}_X$ we have

$$\begin{aligned} &\textit{(inner regularity)} && \mu(A) = \sup\{\mu(K) : K \subseteq A \text{ compact}\}, \\ &\textit{(outer regularity)} && \mu(A) = \inf\{\mu(U) : U \supseteq A \text{ open}\}. \end{aligned}$$

The same statement is true if instead of a compact metric space one has a locally compact and σ-compact Hausdorff space, but for the purposes of this course it is enough to restrict our attention to compact metric spaces.

Definition 63. Let μ be a Borel probability measure on a compact metric space. A sequence $(x_n)_{n\in\mathbb{N}}$ of points in X are said to be *uniformly distributed according to μ* if for every continuous function $f : X \to \mathbb{C}$ one has

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(x_n) = \int_X f \, \mathrm{d}\mu.$$

A (Borel measurable) function $f : X \to \mathbb{C}$ is called *Riemann integrable with respect to μ* if the set of discontinuities of f has zero measure with respect to μ. A (Borel) set $A \subseteq X$ is called *Jordan measurable with respect to μ* if its boundary $\partial A = \overline{A} \setminus A^\circ$ has zero measure with respect to μ. It follows right away from the definition that a set is Jordan measurable if and only if its indicator function is Riemann integrable.

The following proposition can be viewed as a variant of Weyl's Equidistribution Criterion for arbitrary compact metric spaces. The idea behind the proof is also similar and omitted from these notes.

Proposition 64. *Let μ be a Borel probability measure on a compact metric space (X, d_X) and $(x_n)_{n\in\mathbb{N}}$ a sequence of points in X. The following are equivalent:*

(i) $(x_n)_{n\in\mathbb{N}}$ *is uniformly distributed according to μ;*

(ii) *For any Riemann integrable function $f : X \to \mathbb{C}$,*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(x_n) = \int_X f \, \mathrm{d}\mu;$$

(iii) For every Jordan measurable set $A \subseteq X$

$$d(\{n \in \mathbb{N} : x_n \in A\}) = \mu(A).$$

Examples

Prime Numbers. The Prime Number Theorem states that

$$\{p \leqslant N : p \text{ prime}\}| \sim \frac{N}{\log(N)}.$$

The prime number theorem in arithmetic progressions, also known as Dirichelt's prime number theorem, asserts that for any coprime positive integers $q, r \in \mathbb{N}$ one has

$$\{p \leqslant N : p \equiv r \bmod q,\ p \text{ prime}\}| \sim \frac{1}{\varphi(q)} \frac{N}{\log(N)},$$

where φ is Euler's totient function. It follows that the sequence of prime numbers appears with equal frequency in all coprime residue classes modulo q. In other words, if $p_1 < p_2 < p_3 < \ldots$ is an increasing enumeration of the primes then the sequence $(p_n \bmod q)_{n \in \mathbb{N}}$ is uniformly distributed according to the normalized counting measure on $(\mathbb{Z}/q\mathbb{Z})^* = \{0 \leqslant r < q : \gcd(q, r) = 1\}$.

Chapter 4

Birkhoff's Pointwise Ergodic Theorem

The Ergodic Theorems, both mean and pointwise, embody one the main principles of ergodic theory, specifically that time-averages are equal to space-averages:

$$\underbrace{\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)}_{\text{time-averages}} = \underbrace{\int f \, \mathrm{d}\mu.}_{\text{space-averages}}$$

4.1. The Maximal Inequality and the Maximal Ergodic Theorem

In measure theory, *Markov's inequality* states that if $(X,\mathscr{A},\mu)$ is a measure space, $f: X \to \mathbb{R}$ a measurable function, and $\varepsilon > 0$ then

$$\mu\big(\{x \in X : |f(x)| \geqslant \varepsilon\}\big) \leqslant \frac{1}{\varepsilon}\int |f| \, \mathrm{d}\mu.$$

Applying Markov's inequality to the ergodic average $\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)$ and using the triangle inequality yields

$$\mu\left(\left\{x \in X : \left|\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)\right| \geqslant \varepsilon\right\}\right) \leqslant \frac{1}{\varepsilon}\int |f| \, \mathrm{d}\mu. \tag{4.1.1}$$

The following results, called the Maximal Ergodic Theorem, provides a significant strengthening of (4.1.1) and can be viewed as a uniform version of Markov's inequality for ergodic averages.

Maximal Ergodic Theorem. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system. For any real-valued $f \in L^1(X,\mathscr{A},\mu)$ and $\varepsilon > 0$ we have*

$$\mu\left(\left\{x \in X : \sup_{N\geqslant 1}\left|\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)\right| \geqslant \varepsilon\right\}\right) \leqslant \frac{1}{\varepsilon}\int |f| \,\mathrm{d}\mu. \tag{4.1.2}$$

The proof of the Maximal Ergodic Theorem hinges on a technical result called the maximal inequality.

Maximal Inequality. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system. For $f \in L^1(X,\mathscr{A},\mu)$ a real-valued function define $S_0 = 0$ and*

$$S_m(x) = \sum_{n=0}^{m-1} f(T^n x), \qquad m \geqslant 1,$$

and let $F_N(x) = \max_{0\leqslant m\leqslant N} S_m(x)$ for all $x \in X$. Then

$$\int_{\{x\in X : F_N(x)>0\}} f \,\mathrm{d}\mu \geqslant 0$$

for all $N \geqslant 1$.

Proof. First, observe that $F_N(x) \geqslant S_m(x)$ for all $m = 0,1,\dots,N$, and therefore

$$F_N(Tx) + f(x) \geqslant S_m(Tx) + f(x) = S_{m+1}(x).$$

Hence

$$F_N(Tx) + f(x) \geqslant \max_{1\leqslant m\leqslant N} S_m(x), \qquad \forall x \in X. \tag{4.1.3}$$

Since $S_0 = 0$ we have

$$F_N(x) = \begin{cases} \max_{1\leqslant m\leqslant N} S_m(x), & \text{if } F_N(x) > 0, \\ 0, & \text{otherwise.} \end{cases}$$

So if $P = \{x \in X : F_N(x) > 0\}$ then (4.1.3) implies

$$F_N(Tx) + f(x) \geqslant F_N(x), \qquad \forall x \in P.$$

Thus,

$$\begin{aligned} \int_P f \,\mathrm{d}\mu &\geqslant \int_P F_N(x)\,\mathrm{d}\mu - \int_P F_N(Tx)\,\mathrm{d}\mu \\ &= \int_X F_N(x)\,\mathrm{d}\mu - \int_P F_N(Tx)\,\mathrm{d}\mu && \text{(since } F_N(x) = 0 \text{ for } x \notin P) \\ &\geqslant \int_X F_N(x)\,\mathrm{d}\mu - \int_X F_N(Tx)\,\mathrm{d}\mu && \text{(since } F_N(x) \geqslant 0 \text{ for all } x \in X) \\ &= 0. && \text{(since } T \text{ is measure-preserving)} \end{aligned}$$

□

Proof of the Maximal Ergodic Theorem. By decomposing f into $f = f_+ - f_-$, where $f_+ = f \cdot \mathbf{1}_{\{x:f(x)>0\}}$ and $f_- = -f \cdot \mathbf{1}_{\{x:f(x)<0\}}$, and treating each component separately, we may assume without loss of generality that f is non-negative.

By applying the Maximal Inequality to the function $f(x)-\varepsilon$ we obtain

$$\int_{P_M} f(x) - \varepsilon \, \mathrm{d}\mu \geqslant 0 \tag{4.1.4}$$

where $P_M = \{x \in X : \sup_{1 \leqslant N \leqslant M} \frac{1}{N} \sum_{n=0}^{N-1} f(T^n x) > \varepsilon\}$. Let

$$P = \left\{ x \in X : \sup_{N \geqslant 1} \frac{1}{N} \sum_{n=0}^{N-1} f(T^n x) \geqslant \varepsilon \right\}$$

and note that $P = \bigcup_{M \in \mathbb{N}} P_M$. Thus (4.1.4) and the dominated convergence theorem imply

$$\int_{P} f(x) - \varepsilon \, \mathrm{d}\mu \geqslant 0. \tag{4.1.5}$$

From (4.1.5) we deduce that $\int_P f \, \mathrm{d}\mu \geqslant \varepsilon\mu(P)$. Since $\int_P f \, \mathrm{d}\mu \leqslant \int |f| \, \mathrm{d}\mu$, the claim follows. □

4.2. The Pointwise Ergodic Theorem

Pointwise Ergodic Theorem (General Case). *Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system. For every $f \in L^2(X, \mathscr{A}, \mu)$ we have*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} f(T^n x) = f_{\mathrm{inv}}(x) \qquad \text{for } \mu\text{-a.e. } x \in X,$$

where f_{inv} is as guaranteed by (2.3.3).

Proof. Let $\mathscr{L}$ denote the space of all $f \in L^2(X, \mathscr{A}, \mu)$ for which the limit

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} f(T^n x)$$

exists for μ-almost every $x \in X$. Our goal is to show that $\mathscr{L} = L^2(X, \mathscr{A}, \mu)$.

Clearly, $\mathscr{L}$ is closed under finite linear combinations and contains $\mathscr{H}_{\mathrm{inv}}$. Thus, to conclude $\mathscr{L} = L^2(X, \mathscr{A}, \mu)$ it suffices to show $\mathscr{H}_{\mathrm{erg}} \subseteq \mathscr{L}$, because $\mathscr{H}_{\mathrm{inv}} \oplus \mathscr{H}_{\mathrm{erg}} = L^2(X, \mathscr{A}, \mu)$ by Theorem 53. Let f be an arbitrary element in $\mathscr{H}_{\mathrm{erg}}$. Fix $\varepsilon > 0$, and let $h = g - g \circ T$ be a coboundary with $g \in L^\infty(X, \mathscr{A}, \mu)$ and $\int |f - h| \, \mathrm{d}\mu \leqslant \varepsilon^2$, which is possible because coboundaries are dense in $\mathscr{H}_{\mathrm{erg}}$. Applying the Maximal Ergodic

Theorem to the functions $\mathrm{Re}(f) - \mathrm{Re}(h)$ and $\mathrm{Im}(f) - \mathrm{Im}(h)$ yields

$$\mu\left(\left\{x \in X : \sup_{N \geqslant 1}\left|\mathrm{Re}\left(\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) - h(T^n x)\right)\right| \geqslant \varepsilon\right\}\right) \leqslant \frac{1}{\varepsilon}\int |\mathrm{Re}(f) - \mathrm{Re}(h)|\, \mathrm{d}\mu,$$

$$\mu\left(\left\{x \in X : \sup_{N \geqslant 1}\left|\mathrm{Im}\left(\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) - h(T^n x)\right)\right| \geqslant \varepsilon\right\}\right) \leqslant \frac{1}{\varepsilon}\int |\mathrm{Im}(f) - \mathrm{Im}(h)|\, \mathrm{d}\mu.$$

Using $\int |f - h|\, \mathrm{d}\mu \leqslant \varepsilon^2$ and replacing $\sup_{N \geqslant 1}$ with $\limsup_{N \to \infty}$ yields

$$\mu\left(\left\{x \in X : \sup_{N \geqslant 1}\left|\mathrm{Re}\left(\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) - h(T^n x)\right)\right| \geqslant \varepsilon\right\}\right) \leqslant \varepsilon,$$

$$\mu\left(\left\{x \in X : \sup_{N \geqslant 1}\left|\mathrm{Im}\left(\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) - h(T^n x)\right)\right| \geqslant \varepsilon\right\}\right) \leqslant \varepsilon.$$

Combines, we thus have

$$\mu\left(\left\{x \in X : \limsup_{N \to \infty}\left|\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) - h(T^n x)\right| \geqslant \varepsilon\right\}\right) \leqslant 2\varepsilon. \tag{4.2.1}$$

Since $h = g - g \circ T$ is a coboundary with $g \in L^\infty(X, \mathscr{A}, \mu)$, its ergodic average is telescoping almost everywhere, giving

$$\lim_{N \to \infty} \frac{1}{N}\sum_{n=0}^{N-1} h(T^n x) = 0, \qquad \text{for } \mu - a.e.\ x \in X.$$

So (4.2.1) is equivalent to

$$\mu\left(\left\{x \in X : \limsup_{N \to \infty}\left|\frac{1}{N}\sum_{n=0}^{N-1} f(T^n x)\right| \geqslant \varepsilon\right\}\right) \leqslant 2\varepsilon. \tag{4.2.2}$$

Since ε was arbitrary, this implies that

$$\lim_{N \to \infty} \frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) = 0, \qquad \text{for } \mu - a.e.\ x \in X,$$

proving that $f \in \mathscr{L}$ as desired. □

Pointwise Ergodic Theorem (Ergodic Case). *Let $(X, \mathscr{A}, \mu, T)$ be an ergodic measure preserving system. Then for every $f \in L^2(X, \mathscr{A}, \mu)$,*

$$\lim_{N \to \infty} \frac{1}{N}\sum_{n=0}^{N-1} f(T^n x) = \int f\, \mathrm{d}\mu, \qquad \text{for } \mu\text{-a.e. } x \in X.$$

4.3. Consequences of the Pointwise Ergodic Theorem

Given a measure preserving system $(X,\mathscr{A},\mu,T)$, a set $A\in\mathscr{A}$, and a point $x\in X$, we call

$$R(x,A)=\{n\in\mathbb{N}:T^nx\in A\}$$

the *set of visits* of x to A. It describes the times at which the orbit of the point x under the transformation T "visits" the set A.

The following result is a consequence of the Pointwise Ergodic Theorem. It tells us that in ergodic systems generic points visit sets with the right frequency.

Corollary 65. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system. The following are equivalent.*

(i) $(X,\mathscr{A},\mu,T)$ *is ergodic.*

(ii) For every $A\subseteq\mathscr{A}$ with $\mu(A)>0$ and almost every $x\in X$ the set of visits $R(x,A)$ is non-empty.

(iii) For every $A\subseteq\mathscr{A}$ and almost every $x\in X$ the set of visits $R(x,A)$ has density $\mu(A)$, i.e.,

$$\lim_{N\to\infty}\frac{|\{1\leqslant n\leqslant N:T^nx\in A\}|}{N}=\mu(A).$$

Proof. The implication (i) $\Longrightarrow$ (iii) follows directly from the Pointwise Ergodic Theorem. The implication (iii) $\Longrightarrow$ (ii) is immediate because sets with positive density are always non-empty. Finally, we prove (ii) $\Longrightarrow$ (i) by contradiction. Assume $(X,\mathscr{A},\mu,T)$ is not ergodic, which means there exists $A\in\mathscr{A}$ that is invariant under T and satisfies $0<\mu(A)<1$. Since the complement $X\setminus A$ has positive measure, it follows from (ii) that there exists a set $X'\subseteq X$ of full measure such that $R(x,X\setminus A)\neq\emptyset$ for all $x\in X'$. Since A has positive measure, the intersection $X'\cap A$ is non-empty. In particular, there exists some $x_0\in X'\cap A$. Since $x_0\in X'$ we have $R(x,X\setminus A)\neq\emptyset$, but since $x_0\in A$ and A is invariant, we have $T^nx_0\in A$ for all $n\in\mathbb{N}$ and hence $R(x,X\setminus A)=\emptyset$. We have arrived at a contradiction. $\square$

Corollary 66. *Let (X,d_X) be a compact metric space, μ a Borel probability measure on X, and $T\colon X\to X$ an ergodic measure preserving transformation. Then for μ-almost every $x\in X$ the orbit $(T^nx)_{n\in\mathbb{N}}$ is uniformly distributed according to μ (see Definition 63).*

Proof. Let $f_1,f_2,f_3,\ldots\in C(X)$ be a sequence of continuous functions on X such that $\{f_i:i\in\mathbb{N}\}$ is a dense subset of $C(X)$ with respect to the supremum norm $\|.\|_\infty$. By thePointwise Ergodic Theorem, for every $i \in \mathbb{N}$ there exists a set of full measure $X_i \subseteq X$ such that for all $x \in X_i$,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} f_i(T^n x) = \int f_i \,\mathrm{d}\mu. \tag{4.3.1}$$

Let $X' = \bigcap_{i\in\mathbb{N}} X_i$ and note that X' has full measure. Since (4.3.1) holds for all $x \in X'$ and since any continuous function $f \in C(X)$ can be uniformly approximated by a subsequence of $(f_i)_{i\in\mathbb{N}}$, we conclude that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} f(T^n x) = \int f \,\mathrm{d}\mu$$

holds for all $f \in C(X)$ and all $x \in X'$. This proves that the orbit of any point in X' is uniformly distributed according to μ. □

4.4. Borel's Theorem on Normal Numbers

Let p be an integer greater than or equal to 2. Recall that any real number $x \in [0,1)$ possesses a *base-p digit expansion*,

$$x = \sum_{i=1}^{\infty} d_i p^{-i}, \qquad d_1, d_2, \ldots \in \{0,1,\ldots,p-1\}.$$

In this setting, the nubmers $d_1, d_2, \ldots$ are called the *base-p digits* of x. It is natural to ask about the frequency with which each digits appears in the expansion of x. If all possible finite combinations of digits appear with the expected frequency in the expansion of x then x is called a 'normal number'.

Definition 67. A number $x = \sum_{i=1}^{\infty} d_i p^{-i}$ is called *normal in base-p* if for all $k \geqslant 1$ and all $c_1, \ldots, c_k \in \{0,1,\ldots,p-1\}$ the set

$$\{n \in \mathbb{N} \cup 0 : d_{n+1} = c_1, \ldots, d_{n+k} = c_k\}$$

has density p^{-k}.

Borel's Theorem on Normal Numbers. *For any $p \geqslant 2$, Lebesgue-almost every $x \in [0,1)$ is normal in base-p.*

Proof. Given $x \in [0,1)$ let $d_i(x)$ denote the i-th digit of x in base-p, such that

$$x = \sum_{i=1}^{\infty} d_i(x) p^{-i}.$$

Consider the set $C = \{x \in X : d_1(x) = c_1, \ldots, d_k(x) = c_k\}$. A straightforward calculation reveals that the Lebesgue measure of C equals p^{-k}. Also, note that the map

$T(x) = px \bmod 1$ is an ergodic Lebesgue-measure-preserving transformation on $[0,1)$. Hence, by the Pointwise Ergodic Theorem we have for almost-every $x \in [0,1)$ that

$$\lim_{N\to\infty} \frac{1}{N}\sum_{n=0}^{N-1} \mathbf{1}_C(T^n x) \;=\; p^{-k}.$$

We leave it to the reader to check that

$$\mathbf{1}_C(T^n x) = 1 \iff d_{n+1} = c_1, \ldots, d_{n+k} = c_k,$$

which implies that $\lim_{N\to\infty} \frac{1}{N}\sum_{n=0}^{N-1} \mathbf{1}_C(T^n x)$ is precisely the natural density of the set $\{n \in \mathbb{N} \cup 0 : d_{n+1} = c_1, \ldots, d_{n+k} = c_k\}$, finishing the proof. □

4.5. Continued Fractions and the Gauss-Map

4.5.1. Continued Fractions

Definition 68. A finite continued fraction is an expression of the form

$$[a_0; a_1, a_2, \ldots, a_n] = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \ldots \cfrac{1}{a_{n-1} + \cfrac{1}{a_n}}}}$$

with $a_0 \in \mathbb{Z}$ and $a_1, \ldots a_n \in \mathbb{N}$.

Every rational number can be represented as a continued fraction. However, this representation is not unique, because for $a_n \geqslant 2$ we have

$$a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \ldots \cfrac{1}{a_{n-1} + \cfrac{1}{a_n}}}} = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \ldots \cfrac{1}{a_{n-1} + \cfrac{1}{(a_n - 1) + \cfrac{1}{1}}}}}$$

and hence $[a_0; a_1, a_2, \ldots, a_{n-1}, a_n]$ and $[a_0; a_1, a_2, \ldots, a_{n-1}, a_n - 1, 1]$ are two distinct representations of the same rational number. Aside from this modification at the tail of the continued fraction, every rational number has a unique representation. Usually the first, shorter expansion is chosen as the canonical representation of a rational number.

The continued fraction expansion of a rational number can be computed using Euclid's algorithm, as illustrated by the following example.

Example 69. Consider the rational number $\frac{97}{17}$. Applying Euclid's algorithm to the numbers 97 and 17 we obtain

$$97 = \mathbf{5} \cdot 17 + 12$$

$$\begin{aligned} 17 &= \mathbf{1} \cdot 12 + 5 \\ 12 &= \mathbf{2} \cdot 5 + 2 \\ 5 &= \mathbf{2} \cdot 2 + 1 \\ 2 &= \mathbf{2} \cdot 1 + 0, \end{aligned}$$

which proves that 97 and 17 are coprime. Using the sequence of quotients obtained from the algorithm, we can write down the (canonical) finite continued fraction expansion of $\frac{97}{17}$ as

$$\frac{97}{17} = [5;1,2,2,2] = 5 + \frac{1}{1 + \frac{1}{2 + \frac{1}{2 + \frac{1}{2}}}}.$$

Definition 70. A *(simple) continued fraction* is an expression of the form

$$[a_0; a_1, a_2, a_3, \ldots] = a_0 + \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{a_3 + \frac{1}{\ddots}}}} \tag{4.5.1}$$

where $a_0 \in \mathbb{Z}$ and $a_1, a_2, a_3, \ldots \in \mathbb{N}$.

As we will learn in this section, any continued fraction corresponds to a unique irrational number, and any irrational number possesses a unique continued fraction expansion. For example:

- $\sqrt{2} = [1;2,2,2,2,\ldots]$. The expansion is periodic with period 1. A well-known theorem (which we won't prove in this course) asserts that a number has a periodic continued fraction expansion if and only if it is a quadratic surd[1].
- $e = [2;1,2,1,1,4,1,1,6,1,1,8,1,\ldots,1,2n,1,\ldots]$, where e is Euler's constant. It follows an almost periodic structure, the pattern repeats with a period of 3 except that 2 is added to the intermediate value in each cycle.
- $\pi = [3;7,15,1,292,1,1,1,2,1,3,1,\ldots]$. No pattern has been observed in the continued fraction expansion of π.
- $\varphi = [1;1,1,1,1,1,1,1,1,1,1,1,\ldots]$, where $\varphi = \frac{1+\sqrt{5}}{2}$ denotes the golden ration. It is often regarded as the "most" irrational number, since its continued fraction expansion implies that it is the most difficult to approximate by rational numbers.

Definition 71. A finite truncation of a continued fraction

$$[a_0; a_1, a_2, \ldots, a_n] = a_0 + \frac{1}{a_1 + \frac{1}{a_2 + \ldots \frac{1}{a_{n-1} + \frac{1}{a_n}}}} = \frac{p_n}{q_n}$$

is called the *n-th convergent* to $[a_0; a_1, a_2, a_3, \ldots]$.

[1]A *quadratic surd* (often also called a *quadratic irrational*) is any number that is a root of a quadratic polynomial with rational coefficients that is irreducible over the rational numbers.

Proposition 72. *Let $[a_0;a_1,a_2,a_3,\ldots]$ be a continued fraction and let $\alpha_n = \frac{p_n}{q_n}$ denote its n-th convergent (with p_n and q_n coprime).*

(i) For all $n \geqslant 1$,

$$\begin{aligned} p_{n+1} &= a_{n+1}p_n + p_{n-1}, \\ q_{n+1} &= a_{n+1}q_n + q_{n-1}, \\ (-1)^n &= p_{n+1}q_n - p_n q_{n+1}. \end{aligned}$$

(ii) The limit $\alpha = \lim_{n\to\infty} \alpha_n = \lim_{n\to\infty} \frac{p_n}{q_n}$ exists, is irrational, and satisfies

$$\left|\alpha - \frac{p_n}{q_n}\right| < \frac{1}{q_n q_{n+1}}. \tag{4.5.2}$$

(iii) One has $\alpha_0 < \alpha_2 < \alpha_4 < \ldots < \alpha < \ldots < \alpha_5 < \alpha_3 < \alpha_1$.

Proof of (i). We proceed by induction on n. Note that $p_0 = a_0$, $q_0 = 1$, $p_1 = a_0a_1 + 1$ and $q_1 = a_1$. So if we set $p_{-1} = 1$ and $q_{-1} = 0$ then the formulas hold for $n = 0$.

Now suppose the formulas have already been verified some $n \geqslant 0$. Define $\tilde{p} = a_{n+2}p_{n+1} + p_n$ and $\tilde{q} = a_{n+2}q_{n+1} + q_n$; our goal is to show $(\tilde{p},\tilde{q}) = (p_{n+2}, q_{n+2})$ and $\tilde{p}q_{n+1} - p_{n+1}\tilde{q} = (-1)^{n+1}$. We have

$$\begin{aligned} \frac{p_{n+2}}{q_{n+2}} = [a_0;a_1,\ldots,a_{n+1},a_{n+2}] = [a_0;a_1,\ldots,a_{n+1} + 1/a_{n+2}] \\ &= \frac{\left(a_{n+1} + 1/a_{n+2}\right)p_n + p_{n-1}}{\left(a_{n+1} + 1/a_{n+2}\right)q_n + q_{n-1}} \\ &= \frac{a_{n+2}\left(a_{n+1}p_n + p_{n-1}\right) + p_n}{a_{n+2}\left(a_{n+1}q_n + q_{n-1}\right) + q_n} \\ &= \frac{\tilde{p}}{\tilde{q}}. \end{aligned}$$

Also,

$$\tilde{p}q_{n+1} - p_{n+1}\tilde{q} = p_n q_{n+1} - p_{n+1}q_n = (-1)^{n+1}$$

which proves that $\tilde{p}$ and $\tilde{q}$ are coprime. Thus $(\tilde{p},\tilde{q}) = (p_{n+2}, q_{n+2})$ as desired. □

Proof of (ii). Note that $p_{n+1}q_n - p_n q_{n+1} = (-1)^n$ implies

$$\frac{p_n}{q_n} = \frac{p_{n-1}}{q_{n-1}} + \frac{(-1)^{n+1}}{q_n q_{n-1}}$$

and hence

$$\alpha = \lim_{n\to\infty} \frac{p_n}{q_n} = a_0 + \sum_{j=1}^{\infty} \frac{(-1)^{j+1}}{q_j q_{j-1}},$$

which is an absolutely convergent series because $q_n \geqslant 2^{(n-2)/2}$. Moreover,

$$\left|\alpha - \frac{p_n}{q_n}\right| = \left|\sum_{j=n+1}^{\infty} \frac{(-1)^{j+1}}{q_j q_{j-1}}\right| < \frac{1}{q_n q_{n+1}}$$

as desired.

It remains to show that α is irrational. Suppose α is rational, i.e., $\alpha = \frac{a}{b}$ for some $a \in \mathbb{Z}$ and $b \in \mathbb{N}$. Then (4.5.2) implies $|q_n a - b p_n| \to 0$ as $n \to \infty$. Since $q_n a - b p_n$ is an integer, we must have $q_n a = b p_n$ for all but finitely many $n \in \mathbb{N}$, and hence $a/b = p_n/q_n$ for all but finitely many $n \in \mathbb{N}$. This contradicts the facts that p_n and q_n are coprime, and $q_n \to \infty$ as $n \to \infty$. □

Proof of (iii). The fact that $\alpha_0 < \alpha_2 < \alpha_4 < \ldots < \alpha < \ldots < \alpha_5 < \alpha_3 < \alpha_1$ follows from

$$\frac{p_n}{q_n} = a_0 + \sum_{j=1}^{n} \frac{(-1)^{j+1}}{q_j q_{j-1}}$$

and the fact that the terms in this sum are decreasing and have alternating signs. □

Proposition 73. *Every irrational number has an unique continued fraction expension.*

Proof. We first show existence by constructing a continued fraction expansion for α inductively: $x_0 = \alpha$, $a_n = \lfloor x_n \rfloor$, $x_{n+1} = \frac{1}{x_n - a_n}$ $\forall n \in \mathbb{N}$. This is well defined because $x_n - a_n > 0$ since if x_n is an integer for some n, then α must be rational. The inequality $0 < x_n - a_n < 1$ implies that $x_{n+1} > 1$, so $a_{n+1} \geqslant 1$ for $n \in \mathbb{N} \cup \{0\}$. By construction, $\alpha = [a_0; a_1, a_2, a_3, \ldots]$ This shows existence.

We now show uniqueness. Suppose that there exist a continued fraction expansion $[a_0'; a_1', a_2', a_3', \ldots]$ corresponding to α different than the one constructed above. Let m be the first index such that they differ, that is $a_m \neq a_m'$. Then $\beta = x_m - a_m$ is either negative or greater than 1. But $\beta = [0; a_{m+1}, a_{m+2}, a_{m+3}, \ldots] \in (0,1)$. This is a contradiction and this proves the uniqueness of the continued fraction expansion for α. □

Definition 74. A fraction p/q is called a *best approximate* to a real number α if for any other fraction a/b with denominator less than or equal to q one has

$$|q\alpha - p| < |b\alpha - a|.$$

Note that if p/q is a best approximate to α then, in particular,

$$\left|\alpha - \frac{p}{q}\right| < \left|\alpha - \frac{a}{b}\right|$$

for any other fraction a/b with denominator less than or equal to q.

Example 75. Let $\alpha > 0$ be an irrational number and consider the line $y = \alpha x$. By definition, the points on the integer lattice that are closest to the graph of this line are the best approximates to α. It turns out that these points are in one-to-one correspondence with the convergents to the continued fraction expansion of α.

For example, if $\alpha = 1/\sqrt{2}$ then its continued fraction expansion is given by $[0;1,2,2,2,2,\ldots]$ and the corresponding convergents are $\frac{1}{1},\frac{2}{3},\frac{5}{7},\frac{12}{17},\ldots$. In Fig. 4.1, we see the graph of the function $y = x/\sqrt{2}$ and the lattice points closest to it, which are $(1,1)$, $(2,3)$, $(5,7)$, $(12,17)$, and so forth. Notice that these points alternate between approximating the graph from above and from below, which is the geometric counterpart to property (iii) in Proposition 72.

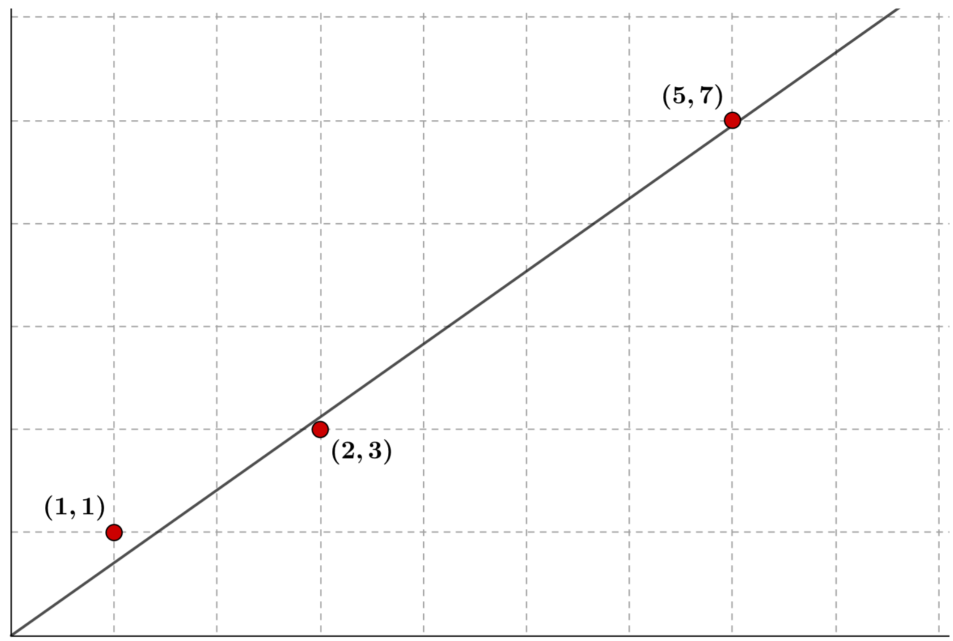

Figure 4.1: Graph of $y = \alpha x$ for $\alpha = 1/\sqrt{2}$.

Theorem 76. *Let α be irrational. The best approximates to α are exactly the convergents of the continued fraction expansion of α.*

The proof of Theorem 76 is omitted.

Now that we have defined continued fractions, we will focus on the properties of their expansions. In the next subsection we will prove the following theorem.

Theorem 77. *Each of the following properties holds for Lebesgue almost every real number $x = [a_0;a_1,a_2,\ldots]$:*

(i) The digit n appears in the expansion $[a_0;a_1,a_2,\ldots]$ of x with frequency

$$\frac{2\log(n+1) - \log(n) - \log(n+2)}{\log 2}.$$

(ii) $\lim_{n\to\infty} \frac{a_1+a_2+\ldots+a_n}{n} = \infty$.

(iii) $\lim_{n\to\infty}(a_1 a_2 \cdot \ldots \cdot a_n)^{1/n} = C$ *where* $C = \prod_{k=1}^{\infty}\left(\frac{(k+1)^2}{k(k+2)}\right)^{\log(k)/\log(2)}$.

(iv) If p_n/q_n *are the convergents to* $x = [a_0; a_1, a_2, \ldots]$ *then*

$$\lim_{n\to\infty} \frac{1}{n} \log\left|x - \frac{p_n}{q_n}\right| = -\frac{\pi^2}{6\log 2}.$$

In particular, this means $|x - p_n/q_n| = O(e^{-\lambda n})$ *for all* $0 < \lambda < \pi^2/6\log 2$.

4.5.2. Gauss map and Gauss measure

Our next goal is to introduce a dynamical approach to the theory of continued fractions. The *Gauß map* is a map $T\colon [0,1) \to [0,1)$ defined via

$$T(x) = \begin{cases} \frac{1}{x} \bmod 1, & \text{if } x \neq 0, \\ 0, & \text{if } x = 0. \end{cases}$$

There is an explicit Borel probability measure on $[0,1)$ for which T is measure preserving, called the *Gauß measure*, defined as

$$\mu(B) = \frac{1}{\log 2}\int_B \frac{1}{1+x}\,\mathrm{d}x, \qquad \text{for all Borel-measurable } B \subseteq [0,1).$$

Proposition 78. *The Gauß map preserves the Gauß measure.*

Proof. It suffices to show $\mu(T^{-1}[0,s]) = \mu([0,s])$ for all $s > 0$ because intervals of this form generate the Borel σ-algebra on $[0,1)$. Note that

$$T^{-1}[0,s] = \{0\} \cup \{x \in (0,1) : \tfrac{1}{x} - \lfloor \tfrac{1}{x} \rfloor \leqslant s\} = \bigcup_{k\in\mathbb{N}} \left[\tfrac{1}{k+s}, \tfrac{1}{k}\right]$$

is a disjoint union. It follows that

$$\begin{aligned}
\mu\big(T^{-1}[0,s]\big) &= \frac{1}{\log 2} \sum_{k\in\mathbb{N}} \int_{1/(k+s)}^{1/k} \frac{1}{1+x}\,\mathrm{d}x \\
&= \frac{1}{\log 2} \sum_{k\in\mathbb{N}} \Big(\log\big(1+\tfrac{1}{k}\big) - \log\big(1+\tfrac{1}{k+s}\big)\Big) \\
&= \frac{1}{\log 2} \sum_{k\in\mathbb{N}} \Big(\log\big(1+\tfrac{s}{k}\big) - \log\big(1+\tfrac{s}{k+1}\big)\Big) \\
&= \frac{1}{\log 2} \sum_{k\in\mathbb{N}} \int_{s/(k+1)}^{s/k} \frac{1}{1+x}\,\mathrm{d}x \\
&= \mu([0,s])
\end{aligned}$$

completing the proof. □

Proposition 79. *The Gauß map is ergodic with respect to the Gauß measure.*

The proof of Proposition 79 is omitted.

The Gauß map and the Gauß measure are tightly connected to the theory of continued fractions. Recall that any irrational number $x \in [0,1)$ has a unique continued fraction expansion

$$x = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \ddots}}}, \qquad a_1, a_2, \ldots \in \mathbb{N},$$

which we write as $[0; a_1, a_2, \ldots]$.

Note that if $x = [0; a_1, a_2, \ldots]$, then $T(x) = [0; a_2, a_3, \ldots]$. Thus T acts as the left shift on the continued fraction representation of a number.

Next, observe that the first digit a_1 in the continued fraction expansion of a real number $x = [0; a_1, a_2, \ldots] \in [0,1)$ satisfies $a_1 = k$ if and only if

$$x \in \left(\tfrac{1}{k+1}, \tfrac{1}{k}\right]. \tag{4.5.3}$$

In other words, the continued fraction expansion of all numbers in $(1/2, 1]$ starts with the digit 1, all numbers in $(1/3, 1/2]$ starts with the digit 2, all numbers in $(1/4, 1/3]$ starts with the digit 3, and so on. Let

$$a(x) = \sum_{k \in \mathbb{N}} k\, \mathbf{1}_{\left(\frac{1}{k+1}, \frac{1}{k}\right]}(x). \tag{4.5.4}$$

In view of the above observation, $a(x)$ is a function that maps every $x \in [0,1)$ to the first digit in its continued fraction expansion. Since the Gauß map acts as the left shift on the continued fraction representation of a number, it follows that we can recover the entire continued fraction expansion of x through its orbit $x, Tx, T^2x, \ldots$ under the Gauß map. More precisely, for any irrational $x = [0; a_1, a_2, \ldots]$ we have

$$a_{n+1} = a(T^n x), \qquad \forall n \in \mathbb{N} \cup \{0\}. \tag{4.5.5}$$

This establishes a direct link between the dynamical behaviour of the Gauß map T and continued fraction representations of irrationals in the $[0,1)$ interval.

Let us now give a proof of Theorem 77.

Proof of Theorem 77, part (i). The frequency of the digit k in $[a_0; a_1, a_2, \ldots]$ is

$$\lim_{N \to \infty} \frac{1}{N} \left| \{1 \leqslant n \leqslant N : a_n = k\} \right|$$

Recall that $a_n = k$ if and only if the $(n-1)$-th iterate of x under the Gauß map T lands in the interval $(1/(k+1), 1/k]$ (cf. (4.5.3) above). Therefore

$$\begin{aligned} \lim_{N \to \infty} \frac{1}{N} \left| \{1 \leqslant n \leqslant N : a_n = k\} \right| &= \lim_{N \to \infty} \frac{1}{N} \left| \{0 \leqslant n \leqslant N-1 : a(T^n x) = k\} \right| \\ &= \lim_{N \to \infty} \frac{1}{N} \sum_{n=0}^{N-1} \mathbf{1}_{\left(\frac{1}{k+1}, \frac{1}{k}\right]}(T^n x). \end{aligned}$$

By Birkhoff's Pointwise Ergodic Theorem, we have

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\mathbf{1}_{\left(\frac{1}{k+1},\frac{1}{k}\right]}(T^n x) = \mu\left(\left(\frac{1}{k+1},\frac{1}{k}\right]\right) \tag{4.5.6}$$

for μ-almost every $x \in [0,1)$. Since a subset of $[0,1)$ is a conull set with respect to the Gauß measure if and only if it is a conull set with respect to the Lebesgue measure, it follows that (4.5.6) also holds for Lebesgue-almost every $x \in [0,1)$. It is now a straightforward calculation that gives

$$\mu\left(\left(\frac{1}{k+1},\frac{1}{k}\right]\right) = \frac{1}{\log 2}\int_{1/(k+1)}^{1/k}\frac{1}{1+x}\,\mathrm{d}x = \frac{2\log(n+1)-\log(n)-\log(n+2)}{\log 2},$$

finishing the proof. □

Proof of Theorem 77, part (ii). We have

$$\lim_{N\to\infty}\frac{a_1+a_2+\ldots+a_N}{N} = \lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}a(T^n x)$$

where $a(x)$ is as defined in (4.5.4). Note that $a(x)$ is not an integrable function, which means we need to truncate it at a finite level to be able to apply the Pointwise Ergodic Theorem. Let

$$a_M(x) = \sum_{k=1}^{M} k\,\mathbf{1}_{\left(\frac{1}{k+1},\frac{1}{k}\right]}(x)$$

and note that $a_M(x)$ converges to $a(x)$ pointwise as $M\to\infty$. Since $a_M(x)$ is bounded, we can apply the Pointwise Ergodic Theorem and obtain

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}a_M(T^n x) = \sum_{k=1}^{M} k\mu\left(\left(\frac{1}{k+1},\frac{1}{k}\right]\right) \geqslant \frac{1}{2\log 2}\sum_{k=1}^{M}\frac{1}{k+1}$$

for μ-almost every $x\in[0,1)$, which also gives it for Lebesgue-almost every $x\in[0,1)$. The claim follows by taking $M\to\infty$ and noting that $\sum_{k=1}^{\infty}\frac{1}{k+1}=\infty$. □

Proof of Theorem 77, part (iii). Define $f(x)=\log(a(x))$, where $a(x)$ is as in (4.5.4). Then

$$\lim_{N\to\infty}\log\left((a_1a_2\cdot\ldots\cdot a_N)^{1/N}\right) = \lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}f(T^n x)$$

The claim follows from the Pointwise Ergodic Theorem and a straightforward calculation that shows

$$\int f\,\mathrm{d}\mu = \sum_{k=1}^{\infty}\frac{\log k}{\log 2}\int_{1/(k+1)}^{1/k}\frac{1}{1+x}\,\mathrm{d}x = \log(C),$$

where $C=\prod_{k=1}^{\infty}\left(\frac{(k+1)^2}{k(k+2)}\right)^{\log(k)/\log(2)}$. □

The proof of Theorem 77, part (iv) is omitted.

Chapter 5

Classifying Measure Preserving Systems

In every area of contemporary mathematics, one of the core objectives is to develop a formal approach for comparing and categorizing the main objects of interest. For example, in group theory, it is important to understand when two groups are isomorphic, or when one group embeds into another. In addition, it is useful to sort groups into different classes based on important attributes, such as free groups, nilpotent groups, cyclic groups, torsion-free groups, etc. Analyzing these special classes offers a more concrete sense of the different behavior that exists within the category of groups.

In ergodic theory, the main objects of interest are measure preserving systems. Naturally, it would be useful to understand when the dynamical behavior of two measure preserving systems is independent, or correlated, or identical. At the same time, it would be useful to develop basic notions that allow us to distinguish measure preserving systems with different behavior.

5.1. Factors, Extensions, and Isomorphisms

The morphisms in the category of measure preserving systems are factor maps.

Definition 80 (Factor map). Let $(X,\mathscr{A},\mu,T)$ and $(Y,\mathscr{B},\nu,S)$ be measure preserving systems. A measurable map $\pi\colon X\to Y$ is a *factor map* if

- the push-forward of μ under π equals ν;
- π intertwines T and S, by which we mean that $(S\circ\pi)(x)=(\pi\circ T)(x)$ for μ-almost every $x\in X$.

The last condition is equivalent to the commutativity of the following diagram μ-almost everywhere:

$$\begin{array}{ccc} X & \xrightarrow{T} & X \\ \downarrow{\scriptstyle\pi} & & \downarrow{\scriptstyle\pi} \\ Y & \xrightarrow{S} & Y \end{array}$$

Definition 81 (Factors and Extensions)**.** If there exists a factor map $\pi\colon X \to Y$ from a measure preserving system $(X,\mathscr{A},\mu,T)$ to another $(Y,\mathscr{B},\nu,S)$, then $(Y,\mathscr{B},\nu,S)$ is called a *factor* of $(X,\mathscr{A},\mu,T)$. Equivalently, $(X,\mathscr{A},\mu,T)$ is said to be an *extension* of $(Y,\mathscr{B},\nu,S)$.

Since in measure theory sets with 0 measure are considered negligible, the definition of factor maps and isomorphisms need to be flexible enough so that, in particular, two systems that differ only on a 0 measure set are isomorphic.

Definition 82 (Isomorphism)**.** A measurable map $\phi\colon X \to Y$ between two measure preserving systems $(X,\mathscr{A},\mu,T)$ and $(Y,\mathscr{B},\nu,S)$ is called an *isomorphism* if it is a factor map $\phi\colon X \to Y$ and a factor map $\psi\colon Y \to X$ such that $\phi\circ\psi$ and $\psi\circ\phi$ are almost everywhere the identity, i.e., $\psi\big(\phi(x)\big) = x$ for μ-a.e. $x \in X$ and $\phi\big(\psi(y)\big) = y$ for ν-a.e. $y \in Y$. Two measure preserving systems are *isomorphic* if there exists an isomorphism between them.

5.2. Introduction to topological groups

Definition 83. Let G be a group.

- A topology τ on G is said to be a *group topology* if the maps $\kappa\colon G\times G \to G$ and $\iota\colon G \to G$ defined by defined by $\kappa(x,y) = xy$ and $\iota(x) = x^{-1}$ are continuous with respect to the product topology on $G\times G$.
- A *topological group* is a pair (G,τ) where G is a group and τ is a group topology on G.

If the group topology τ is Hausdorff (resp., compact, locally compact, connected, metrizable, etc.), then the topological group (G,τ) is called Hausdorff (resp., compact, locally compact, connected, metrizable, etc.). Likewise, if G is cyclic (resp., abelian, nilpotent, torsionfree, etc.) then the topological group (G,τ) is called cyclic (resp. abelian, nilpotent, torsionfree, etc.).

Examples

Indiscrete groups. Let G be any group and let $\tau_{\mathrm{I}} = \{\emptyset, G\}$ denote the trivial topology on G (sometimes also refereed to as the indiscrete topology). Then (G,τ_{I}) is a topological group since any map into an indiscrete space is continuous. Since the trivial topology is finite, (G, τ_{I}) is a compact topological group. It is Hausdorff if and only if G is the trivial group.

Discrete groups. Let G be a group and denote by τ_{D} the discrete topology on G. Then (G, τ_{D}) is a topological group because any map whose domain is a discrete space is continuous. Any discrete space is Hausdorff and locally compact, so (G, τ_{D}) is a locally compact Hausdorff group. It is compact if and only if G is finite; and it is σ-compact if and only if G is countable.

The real line. The additive group $(\mathbb{R}, +)$ endowed with its usual topology is a topological group which we call the real line. The compact sets in the real line $\mathbb{R}$ are exactly the closed and bounded sets (this is the Heine-Borel Theorem for $\mathbb{R}$). We can use this to see that the real line is a non-compact σ-compact locally compact Hausdorff topological group.

The rationals. The additive group $(\mathbb{Q}, +)$ endowed with the subspace topology inherited from the real line is a topological group. It is Hausdorff, but unlike the real line, it is not locally compact (and therefore also not compact). However, $\mathbb{Q}$ is σ-compact since it is countable.

The circle group. The set $\mathbb{S}^1 := \{z \in \mathbb{C} : |z| = 1\}$ under multiplication with the subspace topology inherited from the usual topology on $\mathbb{C}$ is a topological group, called the circle group; it is compact and Hausdorff.

The torus. Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, which is an abelian group under addition modulo 1. It is a quotient group of the real line, and hence endowed naturally with the quotient topology inherited from $\mathbb{R}$. Under this topology, $\mathbb{T}$ is a compact Hausdorff group. Note that $(\mathbb{T}, +)$ and $(\mathbb{S}^1, \cdot)$ are isomorphic as topological groups, and the map $x \mapsto \exp(2\pi i x)$ is a natural homeomorphic group isomorphism between them.

Infinite circle group. Let $\mathbb{S}^{\mathbb{N}} = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1 \times \ldots$ denote the product group consisting of countably many copies of $\mathbb{S}^1$. When endowed with the product topology, $\mathbb{S}^{\mathbb{N}}$ is a topological group. Since $\mathbb{S}^1$ is a compact Hausdorff group, by Tychonoff's theorem, $\mathbb{S}^{\mathbb{N}}$ is a compact Hausdorff group. We recall Tychonoff's theorem, which will be used again later in the course:

Theorem 84 (Tychonoff). *Let $(X_i)_{i \in I}$ be a collection (possibly uncountable) of topological spaces. If X_i is compact $\forall i \in I$, then, the product space $\prod_{i \in I} X_i$ is also compact.*

5.2.1. The Haar measure

Let G be a locally compact group, and let $\mathscr{B}_G$ denote Borel σ-algebra on G. A Borel measure μ on G is called *left translation invariant* if for every $A \in \mathscr{B}_G$ and $g \in G$, $\mu(g^{-1}A) = \mu(A)$, where $g^{-1}A = \{h \in G : gh \in A\}$. Analogously, μ is *right translation invariant* if $\mu(Ag^{-1}) = \mu(A)$ holds for all $A \in \mathscr{B}_G$ and $g \in G$.

Haar's Theorem. *Let G be a locally compact Hausdorff topological group. There is, up to a positive multiplicative constant, a unique Borel measure m_G on G satisfying the following properties:*

(i) The measure m_G is left translation invariant.
(ii) The measure m_G gives finite measure to any compact set.
(iii) The measure m_G gives positive measure to any open set.

The measure m_G guaranteed by Haar's Theorem is called the *left Haar measure* on G. In complete analogy, one can prove the existence and uniqueness of a *right Haar measure* on G.

The most familiar example of a Haar measure is the Lebesgue measure on the real line $\mathbb{R}$, which is both left and right translation invariant.

Remark 85. Here are some noteworthy remarks about the Haar measure.

- If G is compact then $m_G(G)$ is finite and positive, and hence, by normalizing m_G, we can assume that the Haar measure is a probability measure.
- If G is σ-compact then the Haar measure is a Radon measure (cf. the discussion in Section 3.4).
- If $\iota\colon G \to G$ denotes the map $\iota(x) = x^{-1}$ then m_G is a left Haar measure on G if and only if ιm_G, the push-forward of m_G under ι, is a right Haar measure on G. We see that the existence and uniqueness of a left Haar measure and a right Haar measure are equivalent.
- Topological groups for which the left Haar measure and the right Haar measure coincide are called *unimodular groups*. Examples of unimodular groups are abelian groups, compact groups, discrete groups (e.g., finite groups), semisimple Lie groups, and connected nilpotent Lie groups.

5.2.2. The Pontryagin dual

We now restrict our attention to abelian topological groups $(G,+)$.

Definition 86. Let G be a locally compact abelian topological group.

- A *continuous group character* of G is a continuous homomorphism from $(G,+)$ to the circle group $(\mathbb{S}^1,\cdot)$, i.e., a continuous map $\chi\colon G \to \mathbb{S}^1$ satisfying $\chi(a+b) = \chi(a)\chi(b)$ for all $a,b \in G$.
- The *Pontryagin dual* of G is the set of all continuous group characters of G and denoted by $\widehat{G}$. Since pointwise multiplication of continuous group characters is

a continuous group character, $\widehat{G}$ is a group. Moreover, when endowed with the topology given by uniform convergence on compact sets (that is, the topology induced by the compact-open topology on the space of all continuous functions from G to $\mathbb{S}^1$), $\widehat{G}$ is an abelian locally compact topological group.

Proposition 87. *We have $\widehat{\widehat{G}} \cong G$ and the canonical map $ev_G \colon G \to \widehat{\widehat{G}}$ given by $ev_G(g) \colon (\chi \mapsto \chi(g))$ is a homeomorphic isomorphism between $(G,+)$ and $(\widehat{\widehat{G}},\cdot)$.*

The proof of Proposition 87 is omitted.

Proposition 88. *Let G be a compact abelian topological group and let m_G denote the (normalized) Haar measure on G.*

(i) Finite linear combinations of continuous group characters form a uniformly dense subalgebra of $C(G)$.

(ii) The set $\{\chi : \chi$ is a continuous group character on $G\}$ forms an orthonormal basis of $L^2(G,\mathscr{B}_G,m_G)$.

Proof of part (i). Let $\mathscr{A} := \{\sum_{i=1}^n a_i\chi_i : a_i \in \mathbb{C}, \chi_i$ is a character of $G, \forall i = 1,\dots,n\}$, i.e, it is the set of finite linear combinations of continuous group characters. Note that $\mathscr{A}$ is clearly a subalgebra of $C(G)$ as it is closed under all the operations of the algebra $C(G)$. Moreover, $\mathscr{A}$ contains a non-zero constant function as it contains the trivial character defined by $\chi(g) = 1, \forall g \in G$. Finally, it is also closed under complex conjugation. Indeed, the conjugate of a character χ is defined by $\overline{\chi} : G \to \mathbb{S}^1$, $\overline{\chi}(g) = \overline{\chi(g)}$. One can see that this is well defined as for any $z \in \mathbb{C}$ such that $|z| = 1$, we have $\overline{z} = 1/z$, hence,

$$\begin{aligned}\overline{\chi}(g) &= \overline{\chi(g)}\\ &= \chi(g)^{-1}\\ &= \chi(g^{-1})\end{aligned}$$

since χ is an homomorphism. Moreover, since by Proposition 87 the map $\widehat{\widehat{G}} \cong G$ is injective, $\mathscr{A}$ separates points. In view of the Stone-Weierstrass Theorem (see Section 3.2), we conclude that $\mathscr{A}$ is dense in $C(G)$. □

5.3. Kronecker Systems

Definition 89. Given a measure preserving system $(X,\mathscr{A},\mu,T)$, a non-zero function $f \in L^2(X,\mathscr{A},\mu)$ is an *eigenfunction* if there exists a constant λ, called the *eigenvalue*, such that $f \circ T = \lambda f$ (where the equality is understood to hold μ-almost everywhere).

Observe that f is an eigenfunciton if and only if it is an eigenvector for the Koopman operator $U_T : L^2(X,\mathscr{A},\mu) \to L^2(X,\mathscr{A},\mu)$ and λ is the associated eigenvalue.

Since U_T is unitary, all its eigenvalues must have absolute value 1. The set of all eigenvalues of $(X,\mathscr{A},\mu,T)$ is called the *point-spectrum of* T and denoted by $\sigma(T)$.

Example 90. Let $(\mathbb{T},\mathscr{B}_{\mathbb{T}},m_{\mathbb{T}},T_\alpha)$ be a circle rotation, that is, $T_\alpha(x)=x+\alpha \bmod 1$ for some fixed $\alpha\in\mathbb{R}$. Then the function $f(x)=e(x)=e^{2\pi i x}$ is an example of an eigenfunction for this system, and its eigenvalue is $e(\alpha)$. As a matter of fact, all eigenfunctions for the transformation $Tx=x+\alpha \bmod 1$ are of the form $x\mapsto ce(nx)$ with $n\in\mathbb{Z}$ and $c\in\mathbb{C}$, and the corresponding eigenvalues are the numbers $e(n\alpha)$, $n\in\mathbb{Z}$. Hence, the point spectrum of T is given by

$$\sigma(T)=\{e(n\alpha):n\in\mathbb{Z}\}.$$

Note that $\sigma(T)$ is a subgroup of the circle group $(\mathbb{S}^1,\cdot)$, which is not a coincidence.

Lemma 91. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system.*

(i) The point-spectrum $\sigma(T)$ of T is a subgroup of $(\mathbb{S}^1,\cdot)$, where $\mathbb{S}^1$ denotes the unit circle in the complex plane.

(ii) If f and g are eigenfunctions with eigenvalues λ_f and λ_g, respectively, and $\lambda_f\neq\lambda_g$, then f and g are orthogonal.

(iii) If $(X,\mathscr{A},\mu,T)$ is ergodic then every eigenfunction f has constant modulus (i.e. $|f|$ is constant μ-almost everywhere) and every eigenvalue is simple (i.e. the eigenspace of every eigenvalue is 1-dimesnsional).

Proof. If f and g are eigenfunctions with eigenvalues λ_f and λ_g, respectively, then fg is an eigenfunction with eigenvalue $\lambda_f\lambda_g$ and $\overline{f}$ is an eigenfunction with eigenvalue $\overline{\lambda_f}=\lambda_f^{-1}$. This shows that $\sigma(T)$ is closed under products and taking inverses, proving that it is a subgroup of $(\mathbb{S}^1,\cdot)$. This proves part (i).

For part (ii), note that $\langle f,g\rangle=\langle Tf,Tg\rangle=\lambda_f\overline{\lambda_g}\langle f,g\rangle$, and hence $\langle f,g\rangle=0$.

For part (iii) observe $|f|\circ T=|\lambda_f||f|=|f|$. Since invariant functions in ergodic systems are almost everywhere constant, it follows that $|f|$ is almost everywhere constant. Finally, if f_1 and f_2 are two eigenfunctions with the same eigenvalue λ then f_1/f_2 is an invariant function and hence constant almost-everywhere by ergodicity. This shows that f_1 is a scalar-multiple of f_2, finishing the proof. □

Definition 92 (cf. page 21). A *(compact) group rotation* is a measure-preserving dynamical system $(G,\mathscr{B}_G,m_G,R)$ where $(G,+)$ is a compact abelian group, $R:G\to G$ is rotation by a fixed element $\alpha\in G$, that is, $R(x)=x+\alpha$ for all $x\in G$, $\mathscr{B}_G$ denotes the Borel σ-algebra on G, and m_G is the normalized Haar measure on G.

Proposition 93. *Let $(G,\mathscr{B}_G,m_G,R)$ be a group rotation. Then there exists an orthonormal basis for $L^2(G,\mathscr{B}_G,m_G)$ consisting of eigenfunctions.*

Proof. Recall that a continuous group character of a compact abelian group $(G,+)$ is a continuous homomorphism χ from G into the multiplicative group $\mathbb{S}^1\subseteq\mathbb{C}$. The collection of all continuous group characters of G, known as the Pontryagin dual of

G, forms an orthonormal basis for $L^2(G,\mathscr{B}_G,m_G)$ (see Proposition 88). Moreover, if χ is a continuous group character then $\chi(x+\alpha)=\chi(\alpha)\chi(x)$, so χ is an eigenfunction with eigenvalue $\chi(\alpha)$. □

The property described in Proposition 93 distinguishes an important class of measure-preserving systems.

Definition 94. A measure preserving system $(X,\mathscr{A},\mu,T)$ has *discrete spectrum* if there exists an orthonormal basis of $L^2(X,\mathscr{A},\mu)$ consisting of eigenfunctions. If $(X,\mathscr{A},\mu,T)$ is ergodic and has discrete spectrum then it is called a *Kronecker system*.

Theorem 95. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system, where $(X,\mathscr{A},\mu)$ is a standard probability space*[1]*. Then $(X,\mathscr{A},\mu,T)$ is a Kronecker system if and only if it is isomorphic to an ergodic group rotation.*

For the proof of Theorem 95 we need the following lemma.

Lemma 96. *Let $(G,+)$ be a compact abelian group and let $R\colon G\to G$ be rotation by a fixed element $\alpha\in G$. For any Borel probability measure ρ on G for which R is measure preserving and ergodic the resulting measure preserving system $(G,\mathscr{B}_G,\rho,R)$ is isomorphic to a group rotation.*

Proof. Let $H=\overline{\{n\alpha : n\in\mathbb{Z}\}}$, which is a closed subgroup of G, and consider the quotient group G/H. Let $\pi\colon G\to G/H$ denote the natural quotient map and let ρ^* be the push-forward of ρ under π. If ρ^* is not a point-mass then there exists a set $C\subseteq G/H$ with $0<\rho^*(C)<1$. The set $A=\pi^{-1}(C)$ is then a subset of G invariant under R and satisfying $0<\rho(A)<1$, which is impossible because $(G,\mathscr{B}_G,\rho,R)$ is ergodic. Thus ρ^* must be a point-mass, and let $H+u$ denote its support. It is now straightforward to show that the map $\psi\colon H\to G$ given by $\psi(x)=x+u$ is an isomorphism from $(H,\mathscr{B}_H,m_H,R)$ to $(G,\mathscr{B}_G,\rho,R)$, and we leave the details to the interested reader. Since $(H,\mathscr{B}_H,m_H,R)$ is a group rotation, the proof is complete. □

Proof of Theorem 95. That every group rotation has discrete spectrum is the content of Proposition 93. So it only remains to prove the converse under the additional assumption that we are dealing with a standard probability space and that T is ergodic.

Since $(X,\mathscr{A},\mu)$ is a standard probability space, we can assume without loss of generality that X is a compact metric space, $\mathscr{A}$ is the σ-algebra of Borel sets, and μ is a Borel probability measure. Our goal is to find a group rotation $(G,\mathscr{B}_G,m_G,R)$ such that $(X,\mathscr{A},\mu,T)$ and $(G,\mathscr{B}_G,m_G,R)$ are isomorphic. Let $\chi_1,\chi_2,\ldots$ be an orthonormal basis of $L^2(X,\mathscr{A},\mu)$ consisting of eigenfunctions and let $\lambda_1,\lambda_2,\lambda_3,\ldots$ denote the

[1] A *standard probability space* is any probability space that is measurably isomorphic to a probability space $(X,\mathscr{A},\mu)$ where X is a compact metric space, $\mathscr{A}$ is the σ-algebra of Borel sets, and μ is a Borel probability measure on X.

corresponding eigenvalues; note that this basis is countable because the Borel σ-algebra of a compact metric space is countably generated. In view of Lemma 91, part (iii), we can assume that χ_n takes values in $\mathbb{S}^1$. Let

$$\mathbb{S}^{\mathbb{N}} = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1 \times \ldots$$

and consider the map $\phi\colon X \to \mathbb{S}^{\mathbb{N}}$ given by

$$\phi(x) = (\chi_1(x), \chi_2(x), \chi_3(x), \ldots).$$

Also, let $\alpha = (\lambda_1, \lambda_2, \lambda_3, \ldots) \in \mathbb{S}^{\mathbb{N}}$ and consider the map $R\colon \mathbb{S}^{\mathbb{N}} \to \mathbb{S}^{\mathbb{N}}$ given by $R(y) = \alpha \cdot y$ for all $y \in \mathbb{S}^{\mathbb{N}}$. Note that

$$\phi \circ T(x) = R \circ \phi(x) \qquad \text{for } \mu\text{-a.e. } x \in X. \tag{5.3.1}$$

Let $\rho = \phi\mu$ denote the push-forward of μ under ϕ. We claim that $(X, \mathscr{A}, \mu, T)$ and $(\mathbb{S}^{\mathbb{N}}, \mathscr{B}_{\mathbb{S}^{\mathbb{N}}}, \rho, R)$ are isomorphic and that ϕ is an isomorphism between them.

To prove that ϕ is an isomorphism, it suffices to show that it is a factor map and that it is almost everywhere invertible. That ϕ is a factor map from $(X, \mathscr{A}, \mu, T)$ to $(\mathbb{S}^{\mathbb{N}}, \mathscr{B}_{\mathbb{S}^{\mathbb{N}}}, \rho, R)$ follows immediately from the definition of ρ and from (5.3.1). It remains to show that ϕ is almost everywhere invertible. Let $d\colon X \times X \to [0,\infty)$ be a metric on X. We claim for every $\varepsilon > 0$ there exists a set $\Omega_\varepsilon \subseteq X$ with $\mu(X \setminus \Omega_\varepsilon) \leqslant \varepsilon$ and the property that for any $x, y \in \Omega_\varepsilon$ we have

$$\phi(x) = \phi(y) \quad \Longrightarrow \quad d(x,y) \leqslant \varepsilon.$$

In order to find Ω_ε, let $B_1, \ldots, B_r$ be a finite collection of balls with diameter at most ε that cover the entire space X; such a cover exists because X is a compact metric space. Since $\chi_1, \chi_2, \ldots$ form a basis of $L^2(X, \mathscr{A}, \mu)$, there exists $N \in \mathbb{N}$ and $f_1, \ldots, f_r \in \operatorname{span}\{\chi_1, \ldots, \chi_N\}$ with $\|f_i - \mathbf{1}_{B_i}\|_{L^2} \leqslant \varepsilon/2r$. Thus, by Chebyshev's inequality, the set

$$\{x \in X : \max_{i=1,\ldots,r} |f_i(x) - \mathbf{1}_{B_i}(x)| \geqslant 1/2\}$$

has measure at most ε. Let Ω_ε be the complement of the above set. Suppose $x, y \in \Omega_\varepsilon$ and $\phi(x) = \phi(y)$. Since $\phi(x) = \phi(y)$, we have $\chi_n(x) = \chi_n(y)$ for all $n = 1, \ldots, N$ and therefore $f_i(x) = f_i(y)$ for all $i = 1, \ldots, r$. From $x, y \in \Omega_\varepsilon$ it follows that $|f_i(x) - \mathbf{1}_{B_i}(x)| < 1/2$ and $|f_i(y) - \mathbf{1}_{B_i}(y)| < 1/2$ for all $i = 1, \ldots, r$, and hence

$$\mathbf{1}_{B_i}(x) = \mathbf{1}_{B_i}(y) \ \text{ for all } i = 1, \ldots, r.$$

This can only happen when $d(x,y) \leqslant \varepsilon$, as claimed.

Next, let Ω denote the set of all $x \in X$ that belong to infinitely many $\Omega_{1/k}$ for $k \in \mathbb{N}$. By the Monotone Convergence Theorem, Ω has full measure. Moreover, for any $x, y \in \Omega$ we have

$$\phi(x) = \phi(y) \quad \Longrightarrow \quad x = y.$$

This completes the proof that ϕ is almost everywhere invertible.

Since we have proved that ϕ is an isomorphism, the systems $(X,\mathscr{A},\mu,T)$ and $(\mathbb{S}^{\mathbb{N}},\mathscr{B}_{\mathbb{S}^{\mathbb{N}}},\rho,R)$ are isomorphic. But in view of Lemma 96, the system $(\mathbb{S}^{\mathbb{N}},\mathscr{B}_{\mathbb{S}^{\mathbb{N}}},\rho,R)$ is isomorphic to a group rotation $(G,\mathscr{B}_G,m_G,R)$, which shows that $(X,\mathscr{A},\mu,T)$ is isomorphic to a group rotation $(G,\mathscr{B}_G,m_G,R)$. □

5.4. Weak Mixing Systems

In this section, we will explore what it means for a dynamical system to be "mixing". Heuristically, you can think of a mixing system as one in which the system's trajectories become increasingly "spread out" over time and gradually blend into one another. This process reflects how the system evolves in a way that makes its behavior appear more random as time progresses.

To help you visualize this concept, here are two videos that demonstrate mixing transformations:

- https://upload.wikimedia.org/wikipedia/commons/8/8c/Baker%27s_map_mixing.gif
- https://www.youtube.com/watch?v=hex1xhYFrr0

We begin by defining weak mixing, which is a key concept in understanding the broader idea of mixing.

Definition 97. A measure preserving system $(X,\mathscr{A},\mu,T)$ is *weak mixing* if the corresponding product system $(X\times X,\mathscr{A}\otimes\mathscr{A},\mu\otimes\mu,T\times T)$ is ergodic.

The following theorem states several equivalent properties to weak-mixing which explain the name.

Theorem 98. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system. Then the following are equivalent*

(i) $(X,\mathscr{A},\mu,T)$ is weak mixing.

(ii) For every ergodic measure preserving system $(Y,\mathscr{B},\nu,S)$, the product $(X\times Y,\mathscr{A}\otimes\mathscr{B},\mu\otimes\nu,T\times S)$ is ergodic.

(iii) For any two sets $A,B\in\mathscr{A}$ we have

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\left|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)\right|=0.$$

(iv) For any $f,g\in L^2(X,\mathscr{A},\mu)$ we have

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\left|\int_X T^n f\cdot g\,\mathrm{d}\mu-\int_X f\,\mathrm{d}\mu\int_X g\,\mathrm{d}\mu\right|=0.$$

(v) *For any $A,B \in \mathscr{A}$ there exists a subset $E \subseteq \mathbb{N}$ with upper density $\bar{d}(E) = 0$ such that*

$$\lim_{\substack{n\to\infty \\ n\notin E}} \mu(A \cap T^{-n}B) = \mu(A)\mu(B).$$

Condition (iii) makes it clear that every weak mixing system is ergodic. However, not every ergodic system is weak mixing. For example, any non-trivial system with discrete spectrum (such as a circle rotation) is not weak mixing.

Condition (ii) implies that if a system is weak mixing, then its product is also weak mixing. Therefore any number of self products yields a weak mixing system.

Proof of Theorem 98.

(i)⇒(iv) Replacing f with $f - \int_X f \, \mathrm{d}\mu$ we can assume that $\int_X f \, \mathrm{d}\mu = 0$. Using the Cauchy-Schwartz inequality we have

$$\begin{aligned}\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{n=1}^{N} \left| \int_X T^n f \cdot g \, \mathrm{d}\mu \right| \right)^2 &\leqslant \limsup_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \left| \int_X T^n f \cdot g \, \mathrm{d}\mu \right|^2 \\ &= \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \int_{X\times X} (f \otimes \overline{f}) \circ (T \times T)^n \cdot g \otimes \overline{g} \, \mathrm{d}(\mu \otimes \mu).\end{aligned}$$

Using the hypothesis that the product system $(X \times X, \mathscr{A} \otimes \mathscr{A}, \mu \otimes \mu, T \times T)$ is ergodic and applying von Neumann's Mean Ergodic Theorem to the function $f \otimes \overline{f} \in L^2(X \times X)$, we obtain

$$\begin{aligned}\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \int_{X\times X} (f \otimes \overline{f}) \circ (T \times T)^n \cdot g \otimes \overline{g} \, \mathrm{d}(\mu \otimes \mu) \\ = \int_{X\times X} f \otimes \overline{f} \, \mathrm{d}(\mu \otimes \mu) \int_{X\times X} g \otimes \overline{g} \, \mathrm{d}(\mu \otimes \mu).\end{aligned}$$

Observe that $\int_{X\times X} f \otimes \overline{f} \, \mathrm{d}(\mu \otimes \mu) = \left| \int_X f \, \mathrm{d}\mu \right|^2 = 0$, so the previous equation can be rewritten as

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \left| \int_X T^n f \cdot g \, \mathrm{d}\mu \right|^2 = 0,$$

finishing the proof.

(iv)⇒(iii) This is immediate by letting $f = \mathbf{1}_A$ and $g = \mathbf{1}_B$.

(iii)⇒(iv) Condition (iii) is the special case of (iv) when f and g are indicator functions. Since every L^2 function can be approximated by finite linear combinations of indicator functions, we deduce that (iv) holds for any $f, g \in L^2$.

(iii)⇒(v) Fix $m \in \mathbb{N}$ and set $A_m := \{n \in \mathbb{N} : |\mu(A \cap T^{-n}B) - \mu(A)\mu(B)| > 1/m\}$. Observe that

$$\frac{1}{N}\sum_{n=1}^{N}|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)|\geqslant\frac{|A_m\cap[1,N]|}{mN}$$

Taking the limit as $N\to\infty$ we conclude that $\bar{d}(A_m)=0$ for all $m\in\mathbb{N}$. For each $m\in\mathbb{N}$ let $N_m\in\mathbb{N}$ be such that for all $N>N_m$ we have $|A_m\cap[1,N]|\leqslant N/m$ and define

$$E=\bigcup_{m=1}^{\infty}(A_m\cap[N_m+1,N_{m+1}]).$$

Now observe that $A_m\subseteq A_{m+1}$ for all $m\in\mathbb{N}$, hence for each $N\in\mathbb{N}$, choosing m such that $N\in[N_m+1,N_{m+1}]$ we have $E\cap[1,N]\subseteq A_m\cap[1,N]$ and hence $|E\cap[1,N]|\leqslant N/m$. Taking $N\to\infty$ we conclude that $\bar{d}(E)=0$.
Finally, for each $m\in\mathbb{N}$, let $n>N_m$, then if $n\notin E$ we also have $n\notin A_m$ and so $|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)|<1/m$ concluding the proof.

(v)⇒(iii) Assuming (v), for every ε the set $\{n\in\mathbb{N}:\left|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)\right|>\varepsilon\}$ has density 0. On the other hand $|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)|\leqslant 1$ for every $n\in\mathbb{N}$, and hence

$$\limsup_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\left|\mu(A\cap T^{-n}B)-\mu(A)\mu(B)\right|\leqslant 2\varepsilon.$$

Since ε is arbitrary we conclude that (iii) holds.

(iv)⇒(ii) Let $(Y,\mathscr{B},\nu,S)$ be ergodic. In order to show that $(X\times Y,\mathscr{A}\otimes\mathscr{B},\mu\otimes\nu,T\times S)$ is ergodic, we will show that for any $F,G\in L^2(X\times Y)$,

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\int_{X\times Y}(T\times S)^nF\cdot G\,\mathrm{d}(\mu\otimes\nu)=\int_{X\times Y}F\,\mathrm{d}(\mu\otimes\nu)\int_{X\times Y}G\,\mathrm{d}(\mu\otimes\nu).\tag{5.4.1}$$

Since finite linear combinations of tensor functions of the form $(f_1\otimes f_2)(x,y)=f_1(x)f_2(y)$ form a dense subset of $L^2(X\times Y)$, it suffices to establish (5.4.1) when both f and g are tensor functions. Let $f(x,y)=f_1(x)f_2(y)\in L^2(X\times Y)$ and $g(x,y)=g_1(x)g_2(y)\in L^2(X\times Y)$ be arbitrary. Then (5.4.1) can be written as

$$\begin{aligned}\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N}\int_X T^nf_1\cdot g_1\,\mathrm{d}\mu\int_Y S^nf_2\cdot g_2\,\mathrm{d}\nu\\=\int_X f_1\,\mathrm{d}\mu\int_Y f_2\,\mathrm{d}\nu\int_X g_1\,\mathrm{d}\mu\int_Y g_2\,\mathrm{d}\nu.\end{aligned}\tag{5.4.2}$$

Since (5.4.2) is linear in f_1 and we can write $f_1=\int_X f_1\,\mathrm{d}\mu+\left(f_1-\int_X f_1\,\mathrm{d}\mu\right)$, we can separate the proof of (5.4.2) into two cases: when f_1 is a constant and when $\int_X f_1 \,\mathrm{d}\mu = 0$. For the first case, the left hand side of (5.4.2) is

$$\int_X f_1 \,\mathrm{d}\mu \int_X g_1 \,\mathrm{d}\mu \lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} \int_Y S^n f_2 \cdot g_2 \,\mathrm{d}\nu.$$

But now, using the Mean Ergodic Theorem for the ergodic system $(Y,\mathscr{B},\nu,S)$, it is clear that (5.4.2) holds in this case.

Next we establish (5.4.2) in the case that $\int_X f_1 \,\mathrm{d}\mu = 0$. Applying the triangle inequality and using Cauchy-Schwarz with f_2, g_2, we get

$$\begin{aligned}
&\left|\frac{1}{N}\sum_{n=1}^{N}\int_X T^n f_1\cdot g_1 \,\mathrm{d}\mu \int_Y S^n f_2 \cdot g_2 \,\mathrm{d}\nu\right| \\
&\qquad \leqslant \frac{1}{N}\sum_{n=1}^{N}\left|\int_X T^n f_1\cdot g_1 \,\mathrm{d}\mu \int_Y S^n f_2 \cdot g_2 \,\mathrm{d}\nu\right| \\
&\qquad \leqslant \|f_2\|_{L^2}\cdot\|g_2\|_{L^2}\cdot\left(\frac{1}{N}\sum_{n=1}^{N}\left|\int_X T^n f_1\cdot g_1 \,\mathrm{d}\mu\right|\right).
\end{aligned}$$

It follows from (iv) that this quantity converges to 0 as $N\to\infty$, establishing (5.4.2).

(ii)⇒(i) It suffices to show that if (ii) holds then $(X,\mathscr{A},\mu,T)$ is ergodic. To see this assume that $(X,\mathscr{A},\mu,T)$ is not ergodic and let $A\in\mathscr{A}$ be an invariant set such that $0<\mu(A)<1$. Let $(Y,\mathscr{B},\nu,S)$ be the (ergodic) one point system. Then $A\times Y$ is invariant for $T\times S$ and so $(X\times Y,\mathscr{A}\otimes\mathscr{B},\mu\otimes\nu,T\times S)$ would not be ergodic either.

□

Conditions (iii) and (iv) in Theorem 98 equally hold when the regular Cesàro averages are replaced by uniform Cesàro averages (which were introduced in Section 2.6), and the proofs presented work in that case as well. This yields two more equivalent characterizations of weak mixing.

Despite all the equivalent characterizations of weak mixing listed in Theorem 98, the most important one is still missing; it is given by the following theorem.

Theorem 99. *A measure preserving system $(X,\mathscr{A},\mu,T)$ is weak mixing if and only if it has no non-constant eigenfunctions.*

The proof Theorem 99 can be found on page 84 below. It relies an the Jacobs-de Leeuw-Glicksberg decomposition, which we state and prove in the next chapter.

5.5. Mixing Systems

Definition 100. A measure preserving system $(X,\mathscr{A},\mu,T)$ is *mixing* (or sometimes also referred to as *strong-mixing*) if for every $A, B \in \mathscr{A}$,

$$\lim_{n\to\infty} \mu(T^{-n}A \cap B) = \mu(A)\mu(B).$$

Proposition 101. *Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system. Then the following are equivalent.*

- *The system is mixing.*
- *For every $f, g \in L^2(X)$, $\lim_{n\to\infty} \int_X T^n f \cdot g \, \mathrm{d}\mu = \int_X f \, \mathrm{d}\mu \int_X g \, \mathrm{d}\mu$.*
- *For every $f \in L^2(X)$ with $\int_X f \, \mathrm{d}\mu = 0$, the orbit $T^n f$ converges to 0 in the weak topology.*

Proof. The equivalence between the first two follows from the fact that the set of finite linear combinations of indicator functions is dense in L^2. The equivalence between the last two is immediate, after replacing f with $\tilde{f} := f - \int_X f \, \mathrm{d}\mu$ and noticing that $\int_X \tilde{f} \, \mathrm{d}\mu = 0$. □

Every mixing system is weak mixing, and every weak mixing system if ergodic. The following corollary, a culmination of results proved up to this point, offers an illustrative juxtaposition of these central notions in ergodic theory.

Corollary 102. *Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system.*

- *T is ergodic if and only if for all $A, B \in \mathscr{A}$,*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \mu(T^{-n}A \cap B) = \mu(A)\mu(B).$$

- *T is weak mixing if and only if for all $A, B \in \mathscr{A}$,*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left|\mu(T^{-n}A \cap B) - \mu(A)\mu(B)\right| = 0.$$

- *T is mixing if and only if for all $A, B \in \mathscr{A}$,*

$$\lim_{N\to\infty} \mu(T^{-n}A \cap B) = \mu(A)\mu(B).$$

Proof. The first bullet point follows from Corollary 55. The second one follows from Theorem 98. The third is just the definition of (strong) mixing. □

There is also a notion of higher order mixing.

Definition 103. A measure preserving system $(X, \mathscr{A}, \mu, T)$ is *mixing of order k* if for every $A_1, \ldots, A_k \in \mathscr{A}$ and every $a_1, \ldots, a_k : \mathbb{N} \to \mathbb{N}$ with $\lim_{n\to\infty} a_i(n) - a_j(n) = \infty$ for every $1 \leqslant i < j \leqslant k$ one has

$$\lim_{n\to\infty} \mu\left(T^{-a_1(n)}A_1 \cap T^{-a_2(n)}A_2 \cap \cdots \cap T^{-a_k(n)}A_k\right) = \mu(A_1)\mu(A_2)\cdots\mu(A_k).$$

Notice that mixing of order 2 is the same a strong-mixing. It is clear that k-mixing implies $k-1$-mixing. It is a major open problem in ergodic theory whether the converse holds, even for $k = 3$.

5.6. Bernoulli Systems

Given a finite set Σ, called the *alphabet*, consider the set $\Sigma^{\mathbb{N}}$ of all infinite words in the alphabet Σ. The map $T\colon \Sigma^{\mathbb{N}} \to \Sigma^{\mathbb{N}}$ given by $T((x_n)_{n=1}^{\infty}) = (x_{n+1})_{n=1}^{\infty}$ is called the *left-shift*.

Definition 104. A measure preserving system $(X, \mathscr{A}, \mu, T)$ is called a *Bernoulli scheme* (or sometimes also *Bernoulli shift*) if $X = \Sigma^{\mathbb{N}}$, $\mathscr{A}$ is the σ-algebra of Borel sets on $\Sigma^{\mathbb{N}}$, T is the left shift and $\mu = \mu_0^{\mathbb{N}}$ is the product measure of some probability measure μ_0 on Σ. A measure-preserving system is called a *Bernoulli system* if it is isomorphic to a Bernoulli scheme.

Chapter 6

Spectral Theory of Measure Preserving Systems

We are now ready to describe one of the most profound phenomenon in ergodic theory, a dichotomy between eigenfunctions and weak mixing. Recall that a Kronecker system is spanned entirely by its eigenfunctions, whereas a weak mixing system has no non-trivial eigenfunctions at all. In general, a system exhibits a mixture of these two extremes. Nonetheless, it is still possible to separate these phenomena from one another. The result making this possible is called the Jacobs-de Leeuw-Glicksberg decomposition.

6.1. Herglotz's Theorem

As we have already done in pervious chapters, we identify the torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ with the half-open unit interval $[0,1)$, and we identify continuous functions g on $\mathbb{T}$ with continuous functions g on $[0,1)$ satisfying $\lim_{x\to 1^-} g(x) = g(0)$. For $x \in \mathbb{R}$, we write $e(x)$ to abbreviate $e^{2\pi i x}$.

Let μ be a finite Borel measure on $\mathbb{T}$. The *Fourier transform* of μ on $\mathbb{T}$ is the function $\hat{\mu}\colon \mathbb{Z} \to \mathbb{C}$ defined as

$$\hat{\mu}(m) = \int_{\mathbb{T}} e(mx)\, \mathrm{d}\mu(x), \qquad \forall m \in \mathbb{Z}.$$

Since μ is a finite measure, we have $|\hat{\mu}(m)| \leqslant \hat{\mu}(0) = \mu(\mathbb{T}) < \infty$ for all $m \in \mathbb{Z}$, and hence $\hat{\mu}$ is a bounded function. It is natural to ask what type of bounded functions on $\mathbb{Z}$ correspond to the Fourier transform of a finite Borel measure on $\mathbb{T}$. The answer to this question is provided by a classical result in Fourier analysis known as Herglotz's theorem.

Definition 105. A function $\varphi\colon \mathbb{Z} \to \mathbb{C}$ is called *non-negative definite* if for all $M \in \mathbb{N}$

and $\lambda_1,\ldots,\lambda_M \in \mathbb{C}$ one has

$$\sum_{i,j=1}^{M} \lambda_i \overline{\lambda_j} \varphi(i-j) \geqslant 0.$$

Equivalently, φ is *non-negative definite* if for any $M \in \mathbb{N}$ the matrix $A \in \mathbb{C}^{M\times M}$ defined by $A_{i,j} = \varphi(i-j)$ for all $1 \leqslant i,j \leqslant M$ is a non-negative definite matrix (sometimes also called positive semi-definite matrix).

Example 106. The indicator function of the even integers $\varphi(n) = \mathbf{1}_{2\mathbb{Z}}(n)$ is non-negative definite, whereas the indicator function of the odd integers $\varphi(n) = \mathbf{1}_{2\mathbb{Z}+1}(n)$ is not. Other examples of non-negative definite functions include $\varphi(n) = \mathbf{1}_{\{0\}}(n)$ and $\varphi(n) = e(nx)$ for any $x \in \mathbb{R}$.

Herglotz's Theorem. *A function $\varphi\colon \mathbb{Z} \to \mathbb{C}$ is non-negative definite if and only if there exists a finite Borel measure μ on $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ such that $\varphi = \hat{\mu}$, or in other words,*

$$\varphi(n) = \int_{\mathbb{T}} e(nx)\, \mathrm{d}\mu(x), \qquad \forall n \in \mathbb{Z}.$$

The measure μ provided by Herglotz's Theorem is called the *spectral measure* associated to φ.

Example 107. To help develop a better grasp on the notion of a spectral measure, we include here a short list of examples of non-negative functions and their corresponding spectral measures. Given a set $A \subseteq \mathbb{Z}$ let $\mathbf{1}_A$ denote the indicator function of A, and given $x \in \mathbb{T}$ let δ_x denote the Dirac point measure at x.

non-neg. def. fucn.	spectral meas.
1	δ_0
$\mathbf{1}_{2\mathbb{Z}}(n)$	$\frac{1}{2}(\delta_0 + \delta_{1/2})$
$e(nx)$	δ_x
$\mathbf{1}_{\{0\}}(n)$	Lebesgue measure

Proof of Herglotz's Theorem. First, let us show that if μ is a finite Borel measure on $\mathbb{T}$ then the function $\varphi(n) = \int_{\mathbb{T}} e(nx)\, \mathrm{d}\mu(x)$ is non-negative definite. Observe that for any $M \in \mathbb{N}$ and $\lambda_1,\ldots,\lambda_M \in \mathbb{C}$ we have

$$\begin{aligned}
\sum_{i,j=1}^{M} \lambda_i \overline{\lambda_j} \varphi(i-j) &= \sum_{i,j=1}^{M} \lambda_i \overline{\lambda_j} \int_{\mathbb{T}} e((i-j)x)\, \mathrm{d}\mu(x) \\
&= \int_{\mathbb{T}} \left(\sum_{i,j=1}^{M} \lambda_i e(ix) \overline{\lambda_j e(jx)} \right) \mathrm{d}\mu(x) \\
&= \int_{\mathbb{T}} \left| \sum_{i=1}^{M} \lambda_i e(ix) \right|^2 \mathrm{d}\mu(x) \\
&\geqslant 0,
\end{aligned}$$

proving that φ is non-negative definite.

It remains to show that if φ is non-negative definite then there exists a Borel measure μ such that $\varphi(n) = \int_{\mathbb{T}} e(nx)\,\mathrm{d}\mu(x)$ holds for all $n \in \mathbb{Z}$. Let $\Psi_N \colon \mathbb{T} \to \mathbb{C}$ be defined as

$$\Psi_N(x) = \frac{1}{N} \sum_{i,j=1}^{N} \varphi(i-j)\overline{e((i-j)x)}.$$

Since $\Psi_N(x)$ can be written as $\frac{1}{N}\sum_{i,j=1}^{N} e(-ix)\overline{e(-jx)}\varphi(i-j)$, the fact that φ is non-negative definite implies $\Psi_N(x) \geqslant 0$ for all $x \in \mathbb{T}$. Let μ_N denote the measure on $\mathbb{T}$ whose density function with respect to Lebesgue measure is Ψ_N. In other words, let μ_N be the measure on $\mathbb{T}$ uniquely determined by the property that

$$\int_{\mathbb{T}} g(x)\,\mathrm{d}\mu_N(x) = \int_0^1 g(x)\Psi_N(x)\,\mathrm{d}x$$

for all continuous functions $g \colon \mathbb{T} \to \mathbb{C}$. Note that for any $n \in \mathbb{Z}$,

$$\begin{aligned}
\int_{\mathbb{T}} e(nx)\,\mathrm{d}\mu_N(x) &= \frac{1}{N} \sum_{i,j=1}^{N} \varphi(i-j) \int_0^1 e((j-i+n)x)\,\mathrm{d}x \\
&= \frac{1}{N} \sum_{i,j=1}^{N} \varphi(i-j)\mathbf{1}_{\{n\}}(i-j) \\
&= \begin{cases} \left(1-\frac{|n|}{N}\right)\varphi(n), & \text{if } |n| \leqslant N, \\ 0, & \text{otherwise.} \end{cases}
\end{aligned}$$

In particular, this means that

$$\lim_{N\to\infty} \int_{\mathbb{T}} e(nx)\,\mathrm{d}\mu_N(x) = \varphi(n). \tag{6.1.1}$$

Since finite linear combinations of linear characters $x \mapsto e(nx)$, $n \in \mathbb{Z}$, are uniformly dense in $C(\mathbb{T})$ (due to the Stone-Weierstrass Theorem), we conclude that for all continuous $g \colon \mathbb{T} \to \mathbb{C}$ the limit

$$\lim_{N\to\infty} \int_{\mathbb{T}} g\mu_N \;\text{ exists.}$$

By the Riesz–Markov–Kakutani representation theorem (see Theorem 21), there exists a finite Borel measure μ on $\mathbb{T}$ such that for all continuous $g \colon \mathbb{T} \to \mathbb{C}$,

$$\lim_{N\to\infty} \int_{\mathbb{T}} g\mu_N = \int_{\mathbb{T}} g\mu.$$

In view of (6.1.1), we have $\int_{\mathbb{T}} e(nx)\mu(x) = \varphi(n)$, finishing the proof. $\square$

Corollary 108 (Existence of spectral measures for unitary operators). *Let $U \colon \mathscr{H} \to \mathscr{H}$ be a unitary operator on a Hilbert space $\mathscr{H}$. For each $f \in \mathscr{H}$ there is a unique* *finite Borel measure μ_f on $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ such that*

$$\langle U^n f, f\rangle = \int_{\mathbb{T}} e(nx)\, \mathrm{d}\mu_f(x), \qquad \forall n \in \mathbb{Z}. \tag{6.1.2}$$

Proof. Since U is unitary, it is straightforward to check that $\varphi(n) = \langle U^n f, f\rangle$ is a non-negative definite function, and hence the conclusion follows right away from Herglotz's Theorem. □

Definition 109. The measure μ_f defined through (6.1.2) is called the *spectral measure* of the function f with respect to a unitary operator U.

6.2. Wiener's Lemma

Wiener's Lemma. *Let μ be a finite Borel measure on $\mathbb{T}$. Then*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} |\hat{\mu}(n)|^2 = \sum_{x\in\mathbb{T}} |\mu(\{x\})|^2. \tag{6.2.1}$$

Proof. Consider the measure ν on $\mathbb{T}$ given by $\nu(B) = \int_{\mathbb{T}}\int_{\mathbb{T}} \mathbf{1}_B(x-y)\, \mathrm{d}\mu(x)\, \mathrm{d}\mu(y)$. Then we have

$$\begin{aligned}
& \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} |\hat{\mu}(n)|^2 = \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left| \int_{\mathbb{T}} e(nx)\, \mathrm{d}\mu(x) \right|^2 \\
\text{expand the square} \quad & = \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \int_{\mathbb{T}}\int_{\mathbb{T}} e(nx)e(-ny)\, \mathrm{d}\mu(x)\, \mathrm{d}\mu(y) \\
& = \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \int_{\mathbb{T}} e(nx)\, \mathrm{d}\nu(x) \\
\text{Dominated Convergence Theorem} \quad & = \int_{\mathbb{T}} \underbrace{\left(\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} e(nx) \right)}_{=\mathbf{1}_{\{0\}}(x)} \mathrm{d}\nu(x) \\
& = \int_{\mathbb{T}} \mathbf{1}_{\{0\}}(x)\, \mathrm{d}\nu(x) \\
& = \int_{\mathbb{T}}\int_{\mathbb{T}} \mathbf{1}_{\{0\}}(x-y)\, \mathrm{d}\mu(x)\, \mathrm{d}\mu(y) \\
& = \int_{\mathbb{T}} \mu(\{y\})\, \mathrm{d}\mu(y) \\
& = \sum_{y\in\mathbb{T}} |\mu(\{y\})|^2,
\end{aligned}$$

completing the proof. □

Wiener's Lemma was proved by Norbert Wiener in the late 1920s and early 1930s as part of his program to extend Fourier analysis to measures and generalized functions. The key idea is that Cesàro averages of $|\hat{\mu}(n)|^2$ detect the pure point part of a measure, while the continuous part contributes no mass to these averages.

6.3. Weak Mixing Functions

Definition 110. Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and let $f \in L^2(X,\mathscr{A},\mu)$. We say that f is a *weak mixing* function if for all $g \in L^2(X,\mathscr{A},\mu)$ one has

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\left|\langle U_T^n f, g\rangle\right| = 0. \tag{6.3.1}$$

Notice that a weak-mixing function f always satisfies $\int f \, \mathrm{d}\mu = 0$. Moreover, in view of Theorem 98, a system is weak mixing if and only if every function f with $\int f \, \mathrm{d}\mu = 0$ is a weak-mixing function.

Theorem 111. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and let $f \in L^2(X,\mathscr{A},\mu)$. The following are equivalent*

(i) The function f is weak mixing.

(ii) We have

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\left|\langle U_T^n f, f\rangle\right| = 0.$$

(iii) The spectral measure μ_f of f is continuous[1].

Proof of Theorem 111.

(ii)⇔(iii) An elementary argument shows that

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\left|\langle U_T^n f, f\rangle\right| = 0 \iff \lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\left|\langle U_T^n f, f\rangle\right|^2 = 0.$$

By the definition of the spectral measure, we have $\langle U_T^n f, f\rangle = \int_{\mathbb{T}} e(nx)\,\mathrm{d}\mu_f(x)$. Hence, f is weak mixing if and only if

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\left|\int_{\mathbb{T}} e(nx)\,\mathrm{d}\mu_f(x)\right|^2 = 0.$$

The claim now follows from Wiener's Lemma because μ_f is continuous if and only if it is non-atomic, which is equivalent to $\sum_{x\in\mathbb{T}} |\mu_f(\{x\})| = 0$.

[1]Recall that a Borel measure μ is called *continuous* if it is non-atomic, i.e., all singletons have zero measure.

(i)⇔(ii) The forward implication is immediate, so it only remains to prove the backwards direction. Fix $g \in L^2(X, \mathscr{A}, \mu)$. Our goal is to show (6.3.1). Consider the product system $(X \times X, \mathscr{A} \otimes \mathscr{A}, \mu \otimes \mu, T \times T)$ and define $F = f \otimes \overline{f}$ and $G = g \otimes \overline{g}$. Let F_{inv} denote the orthogonal projection of F onto the space of $T \times T$ invariant functions in $L^2(X \times X, \mathscr{A} \otimes \mathscr{A}, \mu \otimes \mu)$. We have

$$
\begin{aligned}
\lim_{N\to\infty} \left(\frac{1}{N} \sum_{n=0}^{N-1} \left| \langle U_T^n f, g \rangle \right| \right)^2 &\leqslant \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left| \langle U_T^n f, g \rangle \right|^2 \\
&= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left\langle U_{T\times T}^n F, G \right\rangle \\
&= \left\langle F_{\text{inv}}, G \right\rangle \\
&\leqslant \|G\|_{L^2} \|F_{\text{inv}}\|_{L^2} \\
&= \|G\|_{L^2} \left\langle F_{\text{inv}}, F_{\text{inv}} \right\rangle^{1/2} \\
&= \|G\|_{L^2} \left\langle F_{\text{inv}}, F \right\rangle^{1/2} \\
&= \|G\|_{L^2} \left(\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left\langle U_{T\times T}^n F, F \right\rangle \right)^{1/2} \\
&= \|G\|_{L^2} \left(\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left| \langle U_T^n f, f \rangle \right|^2 \right)^{1/2} \\
&\leqslant \|G\|_{L^2} \|f\|_{L^2} \left(\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \left| \langle U_T^n f, f \rangle \right| \right)^{1/2}.
\end{aligned}
$$

Jensen's/Cauchy-Schwarz Inequality

Mean Ergodic Theorem

Using $F = F_{\text{inv}} + F_{\text{erg}}$ and $F_{\text{inv}} \perp F_{\text{erg}}$

Since $\frac{1}{N} \sum_{n=0}^{N-1} |\langle U_T^n f, f \rangle|$ converges to 0, the proof is complete.

□

6.4. The Splitting $\mathscr{H}_{\text{c}} \oplus \mathscr{H}_{\text{wm}}$

Given a measure preserving system $(X, \mathscr{A}, \mu, T)$ define

$$\mathscr{H}_{\text{c}} = \overline{\text{span}\{f \in L^2(X, \mathscr{A}, \mu) : f \text{ is an eigenfunction}\}}$$

and

$$\mathscr{H}_{\text{wm}} = \{f \in L^2(X, \mathscr{A}, \mu) : f \text{ is a weak mixing function}\}.$$

The Jacobs-de Leeuw-Glicksberg decomposition. *We have $\mathscr{H}_{\text{c}} \perp \mathscr{H}_{\text{wm}}$ and $\mathscr{H}_{\text{c}} \oplus \mathscr{H}_{\text{wm}} = L^2(X, \mathscr{A}, \mu)$.*

Proof. If f is a weak mixing function and g is an eigenfunction then

$$|\langle f, g \rangle| = \frac{1}{N} \sum_{n=0}^{N-1} |\langle U_T^n f, U_T^n g \rangle| = \frac{1}{N} \sum_{n=0}^{N-1} |\langle U_T^n f, g \rangle| \xrightarrow{N\to\infty} 0.$$

This already shows that $\mathscr{H}_{\mathrm{wm}} \subseteq \mathscr{H}_{\mathrm{c}}^{\perp}$. To finish the proof, it remains to prove that if $f \notin \mathscr{H}_{\mathrm{wm}}$ then $f \notin \mathscr{H}_{\mathrm{c}}^{\perp}$, because then $\mathscr{H}_{\mathrm{c}} \perp \mathscr{H}_{\mathrm{wm}}$ and $\mathscr{H}_{\mathrm{c}} \oplus \mathscr{H}_{\mathrm{wm}} = L^2(X,\mathscr{A},\mu)$ as desired.

Suppose $f \notin \mathscr{H}_{\mathrm{wm}}$. Let μ_f be the spectral measure of f and observe that μ_f is not continuous due to Theorem 111. In particular, this means there exists some $\alpha \in \mathbb{T}$ such that $\mu_f(\{\alpha\}) > 0$. We claim that f correlates with an eigenfunction whose eigenvalue equals $e(\alpha)$. This will imply $f \notin \mathscr{H}_{\mathrm{c}}^{\perp}$ and finish the proof.

Write Y for $\mathbb{T}$, $\mathscr{B}$ for the Borel σ-algebra on $\mathbb{T}$, ν for the Haar measure on $\mathbb{T}$, and S for the map $x = x + \alpha$. Then $(Y,\mathscr{B},\nu,S)$ is a measure preserving system that we have already encountered numerous times and which we usually refer to as rotation by α. Let $g \in L^2(Y,\mathscr{B},\nu)$ denote the function $g(y) = e(y)$, $y \in Y$, and observe that g is an eigenfunction of $(Y,\mathscr{B},\nu,S)$ with eigenvalue $e(\alpha)$. Now consider the product system $(X \times Y, \mathscr{A} \otimes \mathscr{B}, \mu \otimes \nu, T \times S)$ and the function $F = f \otimes \overline{g}$ on this product. By the Mean Ergodic Theorem we have

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} U_{T\times S}^n F = F_{\mathrm{inv}}$$

where F_{inv} is the orthogonal projection of F onto the subspace of $(T \times S)$-invariant functions. Observe that

$$\begin{aligned}
U_{\mathrm{id}\times S^{-1}} F_{\mathrm{inv}} &= \lim_{N\to\infty} U_{\mathrm{id}\times S^{-1}} \Big(\frac{1}{N} \sum_{n=0}^{N-1} U_{T\times S}^n F \Big) \\
&= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} U_{T\times S}^n \big(U_{\mathrm{id}\times S^{-1}} F \big) \\
&= \lim_{N\to\infty} e(\alpha) \Big(\frac{1}{N} \sum_{n=0}^{N-1} U_{T\times S}^n F \Big) \\
&= e(\alpha) F_{\mathrm{inv}}.
\end{aligned}$$

This gives that $U_{T\times \mathrm{id}} F_{\mathrm{inv}} = U_{\mathrm{id}\times S^{-1}} U_{T\times S} F_{\mathrm{inv}} = U_{\mathrm{id}\times S^{-1}} F_{\mathrm{inv}}$, which proves

$$U_{T\times \mathrm{id}} F_{\mathrm{inv}} = e(\alpha) F_{\mathrm{inv}}. \tag{6.4.1}$$

Our next goal is to show $\langle F_{\mathrm{inv}}, F\rangle \neq 0$. We have

$$\begin{aligned}
\langle F_{\mathrm{inv}}, F\rangle &= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \langle U_{T\times S}^n F, F\rangle \\
&= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \langle U_T^n f, f\rangle \cdot \langle U_S^n \overline{g}, \overline{g}\rangle \\
&= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} e(-n\alpha) \langle U_T^n f, f\rangle \\
&= \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} e(-n\alpha) \int e(nx)\, \mathrm{d}\mu_f(x)
\end{aligned}$$

$$= \int \Big(\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} e(n(x-\alpha)) \Big)\, \mathrm{d}\mu_f(x)$$
$$= \mu_f(\{\alpha\}) > 0.$$

Define $h_y(x) = F_{\mathrm{inv}}(x,y)$. By Fubini's Theorem, we have

$$\langle F_{\mathrm{inv}}, F\rangle = \int_Y g(y) \Big(\int_X F_{\mathrm{inv}}(x,y)\overline{f}(x)\, \mathrm{d}\mu(x) \Big)\, \mathrm{d}\nu(y) = \int_Y g(y) \langle h_y, f\rangle\, \mathrm{d}\nu(y).$$

Since $\langle F_{\mathrm{inv}}, F\rangle \neq 0$, we conclude that there exists a positive measure set of $y \in Y$ for which $\langle h_y, f\rangle \neq 0$. By Fubini's Theorem, it follows from (6.4.1) that for ν-almost every y the function h_y is a eigenfunction for T with eigenvalue $e(\alpha)$. It follows that there exists some y such that both $\langle h_y, f\rangle \neq 0$ and h_y is an eigenfunction with eigenvalue $e(\alpha)$. Therefore $f \notin \mathscr{H}_c^{\perp}$ as was to be shown. □

Proof of Theorem 99. The Jacobs-de Leeuw-Glicksberg decomposition implies that a system $(X, \mathscr{A}, \mu, T)$ is weak mixing if and only if $\mathscr{H}_c = 1$. This is exactly the content of Theorem 99. □

Chapter 7

Entropy

Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system. Suppose we are given information about the approximate position of a point $x \in X$. What can we infer about the position of Tx? More generally, if I know the approximate position of $x, Tx, T^2x, \ldots, T^{n-1}x$ then how accurately can we predict the location of $T^n x$? The answers to these questions depend crucially on the properties of the transformation T. If the transformation T is *deterministic* in nature then the past trajectory of a point determines the imminent future trajectory and so information about $x, Tx, T^2x, \ldots, T^{n-1}x$ leads to a probable projection on the whereabouts of $T^n x$. Conversely, if the transformation is *chaotic* then the past orbit has little to no influence on the future course of the orbit, and knowing the approximate position of $x, T^x, \ldots, T^{n-1}x$ may not allow us to prognosticate the position of $T^n x$. The purpose of this chapter is to make these ideas mathematically precise using the notion of *entropy*.

7.1. Shannon Entropy

The concept of information entropy was introduced by Claude Shannon in 1948 and is also referred to as Shannon entropy. It has various interpretations. Sometimes it is interpreted as a measure of randomness, other times as a measure of information, and sometimes also as a measure of surprise. The ideas and concepts behind Shannon Entropy form the basis for the notion of entropy in ergodic theory.

Shannon entropy as the measure of information

When viewed in terms of information theory, entropy equals the average amount of information gained from an observation. But how can information be quantified? If grams is the measure of mass, seconds the measure of time, and meters the measure of distance, then what can be used to measure the amount of information?

To answer these questions, we need to conceptualize the idea of *information content*. Roughly speaking, information content is a function that associates a (non-negative real) number to the amount of information gained when apprehending new content, such as receiving a message, or observing the outcome of an experiment.

We can measure the quantity of information contained in a message based on the number of questions that must be asked in order to fully discover the content of the message. So, for example, the amount of information conveyed by a letter is the quantity of questions that we need to get answered in order to fully know the content of the letter. But for this to make sense, we first need to standardise the questions and the possible answers to ensure that asking for the number of questions is meaningful.

Here is how the questions are standardised so that their quantity is stable. First we agree upon the *search space*, the set of objects of our query. Only questions that concern (properties of) elements in the search space are allowed. Call a question *binary* if it only has two answers: "true" or "false". A binary question is called *perfect* if it splits the search space evenly into 2 equal sub-spaces. Then we define the information content of a message to be the number of perfect binary questions that this message answers.

To make this approach more mathematical, we need to start thinking in terms of "bits". Bits are either 0 or 1, which represent answers to binary questions. Then the number of questions that a message answers is the amount of bits that this message contains. So we have the answer to the question of what is used to measure the amount of information: it is *bits*!

For example, suppose our search space is given by the following collection of symbols:

$$\begin{array}{cccc} \color{red}{\blacksquare} & \color{red}{\square} & \color{red}{\blacktriangle} & \color{red}{\triangle} \\ \color{blue}{\blacksquare} & \color{blue}{\square} & \color{blue}{\blacktriangle} & \color{blue}{\triangle} \end{array}$$

This collection can be fully described using 3 perfect binary questions:

(a) *Is the object red?*

(b) *Does the object have an even number of corners?*

(c) *Is the object filled in?*

Now suppose we label the elements in $\{\color{red}{\blacksquare}, \color{red}{\square}\}$ as *good*, the elements in $\{\color{red}{\blacktriangle}, \color{red}{\triangle}, \color{blue}{\blacktriangle}, \color{blue}{\triangle}\}$ as *acceptable*, and the elements in $\{\color{blue}{\blacksquare}, \color{blue}{\square}\}$ as *bad*. Then if your message reads "You have a bad symbol" then you have received 2 bits of information, because you will be able to answer two questions, namely questions (a) and (b). On the other hand, if your message reads "You have an acceptable symbol" then you have received only one bit of information, because the message only allows for question (b) to be answered. The number of bits obtained from a message is its information content.

The *entropy* is the expected amount of information obtained from a question, or in other words, the expectation of the information content. Continuing with the above example, if every symbol has an equal probability to be chosen, then the probability of picking a good symbol is $p_1 = 0.25$, the probability of picking an acceptable symbol is $p_2 = 0.5$, and the probability of picking a bad symbol is $p_3 = 0.25$. So the entropy, or expected information content, of picking a symbol uniformly at random and being told whether it is good, acceptable, or bad equals

$$\begin{aligned} \text{entropy} &= -p_1 \log(p_1) - p_2 \log(p_2) - p_3 \log(p_3) \\ &= 0.25 * 2 + 0.5 * 1 + 0.25 * 2 \\ &= 1.5. \end{aligned}$$

Definition 112. Let $(X, \mathscr{A}, \mu)$ be a probability measure space and $\xi : X \to \{1, \ldots, r\}$ a discrete random variable. Let p_i denote the probability of the event $\{x \in X : f(x) = i\}$. The *Shannon entropy* of ξ is

$$H(\xi) = -\sum_{i=1}^{r} p_i \log_2(p_i).$$

Shannon entropy as the measure of surprise

How surprising is an event? Is it possible to measure surprise with a real number just as we measure other more graspable concepts in nature? The answers to these questions follow from Shannon's solution of the fundamental properties of information.

Let us try to find a function that describes the surprise of an event happening. Informally, if two events have the same probability then the occurrence of one event is just as surprising as the occurrence of the other. So the surprise should be a function of probability only. Let us write $I(p)$ to denote the surprise associated to an event with probability p.

If an event occurs with probability 1 then there is no uncertainty and hence the surprise should be zero. This yields the first condition for $I(p)$.

$$I(1) = 0. \tag{U_1}$$

Next, note that if an event A is more likely to happen than another event B, then it is less surprising when A occurs compared to when B occurs. So the more likely an event the less it should contribute to surprise. We conclude:

$$I(p) \text{ is a decreasing function in } p. \tag{U_2}$$

If two events have roughly the same probability of happening, then one should be approximately as surprising as the other, yielding yet another natural condition on $I(p)$.

$$I(p) \text{ is a continuous function in } p. \tag{U_3}$$

Finally, if event A has a certain amount of surprise, and event B has a certain amount of surprise, and you observe them together, and they are independent, it is reasonable that the amount of surprise adds,

$$I(pq) = I(p) + I(q). \tag{U_4}$$

Up to a scalar multiple, there is only one function that satisfies conditions (U_1)–(U_4), and that is

$$I(p) = -\log_2(p).$$

Now that we have found a mathematical description of surprise, we can calculate the expected surprise of a discrete random variable, which we call entropy. Indeed, the expected surprise should just be the expectation of $I(p)$.

Definition 113. Let $(X, \mathscr{A}, \mu)$ be a probability measure space and $\xi\colon X \to \{1, \ldots, r\}$ a discrete random variable. Let p_i denote the probability of the event $\{x \in X : f(x) = i\}$. The *Shannon entropy* of ξ is

$$H(\xi) = \sum_{i=1}^{r} p_i I(p_i) = -\sum_{i=1}^{r} p_i \log_2(p_i).$$

7.2. Entropy of a Partition

Recall that a partition of a probability space $(X, \mathscr{A}, \mu)$ is a finite or countably infinite collection of pairwise disjoint elements of $\mathscr{A}$ whose union equals X. We will use $\xi = \{A_1, \ldots, A_r\}$ and $\xi = \{A_1, A_2, \ldots\}$ to denote finite and countably infinite partitions, respectively. From a probabilistic point of view, a finite partition can be viewed as a discrete random variable $\xi\colon X \to \{1, \ldots, r\}$, and a countably infinite partition as a random variable $\xi\colon X \to \mathbb{N}$.

For any partition ξ we define $\sigma(\xi)$ to be the smallest σ-algebra containing the elements of ξ. We will call the elements of ξ the atoms of ξ. For every $x \in X$ we denote by $[x]_\xi$ the atom of ξ containing x. If ξ and η are partitions, then η is called a refinement of ξ, written $\xi \leqslant \eta$, if each atom of ξ is a union of atoms of η. The common refinement of ξ and η, denoted $\xi \vee \eta$, is the partition $\{A \cap B : A \in \xi,\ B \in \eta\}$. Notice that $\sigma(\xi \vee \eta) = \sigma(\xi) \vee \sigma(\eta)$.

A partition $\xi = \{A_1, A_2, \ldots\}$ of a probability space may be thought of as giving the possible outcomes $1, 2, \ldots$ of an experiment, with the probability of outcome i being $\mu(A_i)$. Inspired by Shannon entropy, we can associate a number $H(\xi)$ to ξ that describes the amount of uncertainty about the outcome of the experiment, or equivalently the amount of information gained by learning the outcome of the experiment.

Definition 114. Let $(X, \mathscr{A}, \mu)$ be a probability measure space and $\xi = \{A_1, A_2, \ldots\}$ a (finite or countably infinite) partition. The *entropy* of ξ is

$$H(\xi) = -\sum_{i\geqslant 1} \mu(A_i)\log_2(\mu(A_i)),$$

where $0\log 0$ is defined to be 0.

Recall that for $A,B \in \mathscr{A}$ the conditional measure of A with respect to B is

$$\mu(A|B) = \frac{\mu(A\cap B)}{\mu(B)}.$$

It describes the probability of the event A happening assuming that the event B has occurred. Given any partition ξ, we can then consider the partition $\xi \cap B = \{A\cap B : A \in \xi\}$ and the quantity

$$-\sum_{i\geqslant 1} \mu(A_i|B)\log_2\big(\mu(A_i|B)\big)$$

is the entropy of the partition $\xi \cap B$ with respect to the conditional measure $\mu(\cdot|B)$. This leads to the following definition.

Definition 115. If $\xi = \{A_1, A_2, \ldots\}$ and $\eta = \{B_1, B_2, \ldots\}$ are partitions, then the *conditional entropy of ξ given η* is defined to be

$$H(\xi|\eta) = \sum_{j\geqslant 1} \mu(B_j)\Big(-\sum_{i\geqslant 1} \mu(A_i|B_j)\log_2\big(\mu(A_i|B_j)\big)\Big). \tag{7.2.1}$$

Formula (7.2.1) should be viewed as a weighted average of entropies of the partition ξ conditioned on the individual atoms $B_j \in \eta$. It represents the average information gained from the outcome of ξ after being told the outcome of η.

Theorem 116. *Let $\xi = \{A_1, A_2, \ldots\}$, $\eta = \{B_1, B_2, \ldots\}$, and $\zeta = \{C_1, C_2, \ldots\}$ be (finite or countably-infinite) partitions of a probability space $(X, \mathscr{A}, \mu)$.*

(i) $H(\xi) \geqslant 0$, *with equality if and only if* $\mu(A) = 1$ *for some atom* $A \in \xi$.

"No information is gained by conducting an experiment where one of the outcomes is almost surely certain."

(ii) *If ξ is a finite partition with r atoms then $H(\xi) \leqslant \log r$, with equality if and only if $\mu(A) = 1/r$ for each atom $A \in \xi$.*

"Less bias implies more information."

(iii) $H(\xi \vee \eta) = H(\eta) + H(\xi|\eta)$.

"The amount of information gained by learning the outcome of ξ and η simultaneously equals the amount of information gained by first learning the outcome of η and then learning the outcome of ξ given that we already know the outcome of η."

(iv) $H(\xi) \geqslant H(\xi|\eta)$ *and* $H(\xi|\eta) \geqslant H(\xi|\eta \vee \zeta)$.

"The amount of information gained with preexisting knowledge is always smaller or equal than the amount gained without preexisting knowledge."

(v) $H(\xi|\xi) = 0$.

"If we already know the outcome of ξ then there is no information gained by observing the

outcome of ξ."

(vi) *Two partitions ξ and η are independent if and only if $H(\xi \vee \eta) = H(\xi) + H(\eta)$ if and only if $H(\xi|\eta) = H(\xi)$.*

"If two experiments are independent then knowing the outcome of one does not influence the amount of information gained by observing the outcome of the other."

(vii) *If $T : X \to X$ is a measure preserving transformation on $(X, \mathscr{A}, \mu)$ then $H(T^{-1}\xi) = H(\xi)$. Similarly, we have $H(T^{-1}\xi|T^{-1}\eta) = H(\xi|\eta)$.*

"Conducting and observing the outcome of an experiment now or 1 time-unit from now yields the same amount of information."

7.3. Connections to Entropy in Physics and Information Theory

Before learning about entropy in information theory and ergodic theory, the reader may have encountered the analogous notion in physics or information theory. Here is a link to an informative video that explains how one can view entropy in physics from a probabilistic angle, revealing the connection between classical entropy studied in thermodynamics, and the notion of entropy that we defined in the previous section: "What is entropy? – Jeff Phillips"

If you want to learn more about how to use entropy in applications, check out this video: "Solving Wordle using information theory"

7.4. Entropy of a Measure-Preserving Transformation

Fekete's Lemma. *Let $(a_n)_{n\in\mathbb{N}}$ be a sequence of real numbers satisfying*

$$a_{n+m} \leqslant a_n + a_m, \qquad \forall n, m \in \mathbb{N}. \tag{7.4.1}$$

Then the sequence $\frac{a_n}{n}$ converges and

$$\lim_{n\to\infty} \frac{a_n}{n} = \inf_{n\in\mathbb{N}} \frac{a_n}{n}.$$

Fekete's Lemma is used to define the measure-theoretic entropy of a measure preserving transformation. Let $(X, \mathscr{A}, \mu, T)$ be a measure preserving system and let ξ be a measurable partition of X. Consider the sequence $(a_n)_{n\in\mathbb{N}}$ defined by

$$a_n = H(\xi \vee T^{-1}\xi \vee \ldots \vee T^{-(n-1)}\xi) = H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big).$$

This sequence is sub-additive in the sense of (7.4.1). Indeed, combining parts (ii) and (iii) of Theorem 116 then gives

$$\begin{aligned} a_{n+m} &= H\Big(\bigvee_{i=0}^{n+m-1} T^{-i}\xi\Big) \\ &\leqslant H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) + H\Big(\bigvee_{i=n}^{n+m-1} T^{-i}\xi\Big) \end{aligned}$$

and using part (vii) of Theorem 116 now yields

$$H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) + H\Big(\bigvee_{i=n}^{n+m-1} T^{-i}\xi\Big) = H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) + H\Big(\bigvee_{i=0}^{m-1} T^{-i}\xi\Big) = a_n + a_m.$$

It follows from Fekete's Lemma that the limit

$$\lim_{n\to\infty} \frac{1}{n} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big)$$

always exists.

Definition 117. Let $(X, \mathscr{A}, \mu, T)$ be a measure-preserving system and ξ a (finite or countably infinite) partition of X with finite entropy, $H(\xi) < \infty$. The *entropy of T with respect to ξ* is

$$h(T,\xi) = \lim_{n\to\infty} \frac{1}{n} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big).$$

If $H(\xi) = \infty$ then we define $h(T,\xi) = \infty$.

Entropy quantifies the amount of uncertainty or randomness in a system. Heuristically, a system with low entropy is more deterministic than a system with high entropy. In particular, a system with zero entropy is deterministic in the sense that the past determines the future. The next theorem conceptualizes this heuristic.

Theorem 118 (Entropy conditioned on past). *Let $(X, \mathscr{A}, \mu, T)$ be a measure-preserving system and ξ a (finite or countably infinite) partition of X with finite entropy. Then*

$$h(T,\xi) = \lim_{n\to\infty} H\Big(T^{-n}\xi \Big| \bigvee_{i=0}^{n-1} T^{-i}\xi\Big).$$

Proof. For convenience, write $w(n) = H\big(T^{-n}\xi \big| \bigvee_{i=0}^{n-1} T^{-i}\xi\big)$. First, we observe that $w(n) \geqslant 0$ and, using (iv) and (vii) of Theorem 116,

$$\begin{aligned} w(n+1) &= H\Big(T^{-(n+1)}\xi \Big| \bigvee_{i=0}^{n} T^{-i}\xi\Big) \\ &\leqslant H\Big(T^{-(n+1)}\xi \Big| \bigvee_{i=1}^{n} T^{-i}\xi\Big) \end{aligned}$$

$$
\begin{aligned}
&= H\Big(T^{-n}\xi \Big| \bigvee_{i=0}^{n-1} T^{-i}\xi\Big) \\
&= w(n).
\end{aligned}
$$

So we conclude that $w(n)$ is a non-negative and non-increasing sequence, which means that $\lim_{n\to\infty} w(n)$ exists. Next, it follows from property (iii) of Theorem 116 that

$$
\begin{aligned}
H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) &= H\Big(\bigvee_{i=0}^{n-2} T^{-i}\xi\Big) + H\Big(T^{-(n-1)}\xi \Big| \bigvee_{i=0}^{n-2} T^{-i}\xi\Big) \\
&= H\Big(\bigvee_{i=0}^{n-2} T^{-i}\xi\Big) + w(n-1) \\
&= H\Big(\bigvee_{i=0}^{n-3} T^{-i}\xi\Big) + w(n-2) + w(n-1) \\
&\vdots \\
&= w(0) + w(1) + \ldots + w(n-1).
\end{aligned}
$$

It follows that

$$
h(T,\xi) = \lim_{n\to\infty} \frac{1}{n} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n} w(i) = \lim_{n\to\infty} w(n),
$$

where we have used the basic fact that the Cesàro average of a converging sequence converges to the limit of the sequence. □

Theorem 119 (Entropy conditioned on future)**.** *Let $(X,\mathscr{A},\mu,T)$ be a measure-preserving system and ξ a (finite or countably infinite) partition of X with finite entropy. Then*

$$
h(T,\xi) = \lim_{n\to\infty} H\Big(\xi \Big| \bigvee_{i=1}^{n} T^{-i}\xi\Big).
$$

Proof. Let us use $v(n)$ to abbreviate the expression $H\big(\xi \big| \bigvee_{i=1}^{n} T^{-i}\xi\big)$. An argument analogous to the one used at the beginning of the proof of Theorem 118 shows that $v(n)$ is a non-negative and non-increasing function in n and hence the limit

$$
\lim_{n\to\infty} v(n) = \lim_{n\to\infty} H\Big(\xi \Big| \bigvee_{i=1}^{n} T^{-i}\xi\Big)
$$

exists. Using Theorem 116 we see that

$$\begin{aligned} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) &= H\Big(\bigvee_{i=1}^{n-1} T^{-i}\xi\Big) + H\Big(\xi\Big|\bigvee_{i=1}^{n} T^{-i}\xi\Big) \\ &= H\Big(\bigvee_{i=0}^{n-2} T^{-i}\xi\Big) + v(n) \\ &\ \vdots \\ &= H\Big(\bigvee_{i=0}^{n-3} T^{-i}\xi\Big) + v(n-1) + v(n) \\ &\ \vdots \\ &= v(0) + v(1) + \ldots + v(n-1) + v(n). \end{aligned}$$

property **(iii)** of Theorem **116**

property **(vii)** of Theorem **116**

We conclude that

$$h(T,\xi) = \lim_{n\to\infty} \frac{1}{n} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n} v(i) = \lim_{n\to\infty} v(n),$$

where the last equality follows from the fact that the Cesàro average of a converging sequence equals the limit of that sequence. □

With Definition 117, and Theorems 118 and 119, we now have seen 3 different characterization of $h(T,\xi)$:

$$h(T,\xi) = \underbrace{\lim_{n\to\infty} \frac{1}{n} H\Big(\bigvee_{i=0}^{n-1} T^{-i}\xi\Big)}_{①} = \underbrace{\lim_{n\to\infty} H\Big(T^{-n}\xi\Big|\bigvee_{i=0}^{n-1} T^{-i}\xi\Big)}_{②} = \underbrace{\lim_{n\to\infty} H\Big(\xi\Big|\bigvee_{i=1}^{n} T^{-i}\xi\Big)}_{③}$$

Each of these formulas admits its own interpretation of entropy as a dynamical invariant, measuring the "randomness", "uncertainty" or "chaos" inside a dynamical system.

①: Entropy can be understood as the average information gain by observing an outcome over time. By interpreting a partition ξ as a random variable representing the outcomes of a random process (think of conducting an experiment that has finitely man outcomes, like rolling a 6-sided die or counting the number of passengers on the bus each morning during your commute to work), we can view $H\big(\bigvee_{i=0}^{n-1} T^{-i}\xi\big)$ as the expected information gained by observing the outcome of this random variable at n different points in time, from time 0 to time $(n-1)$. Since we have conducted n observations, it is natural to divide by n to normalize the expression, yielding the average information gain over time. According to this interpretation, zero entropy is indicative of a diminishing amount of new information obtained with each consecutive observation.

2: Entropy can also be viewed as a measurement of how much the past determines the future. The conditional entropy $H(\xi|\eta)$ measures the amount of uncertainty in observing ξ after having already observed η. In this sense, the expression $H\big(T^{-n}\xi\big|\bigvee_{i=0}^{n-1}T^{-i}\xi\big)$ measures the amount of uncertainty in the future state $T^n\xi$ subject to the knowledge of its past and present states $\bigvee_{i=0}^{n-1}T^{-i}\xi$. In other words, it quantifies how much additional information is needed to predict the outcome at time n given that we already know the outcomes at times $0,\dots,n-1$. According to this interpretation, a transformation has zero entropy if, as n gets larger and larger, there is less and less amount of new information obtained at time n once we have already learned what happened at times $0,\dots,n-1$.

3: Finally, we can interpret entropy as the level of reciprocality between cause and effect. It measures how much can be inferred about the initial state of the system knowing the future evolution.

Definition 120. The *entropy of T* is

$$h(T)=\sup_{\xi}h(T,\xi).$$

where the supremum is taken over all finite partitions of X.

Kolmogorov-Sinaĭ Theorem. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and suppose ξ is a generating partition of $(X,\mathscr{A},\mu,T)$, which means $\mathscr{A}=\sigma\big(\bigvee_{i=0}^{\infty}T^{-i}\xi\big)$. Then $h(T)=h(T,\xi)$.*

Lemma 121. *Let $(X,\mathscr{A},\mu,T)$ be a measure preserving system and suppose ξ is a generating partition of $(X,\mathscr{A},\mu,T)$. Then for any finite partition η of X with finite entropy we have*

$$\lim_{n\to\infty}H\Big(\eta\Big|\bigvee_{i=0}^{n-1}T^{-i}\xi\Big)=0.$$

Proof. Given $\delta>0$ and $n\in\mathbb{N}$ define

$$\mathscr{A}_\delta^{(n)}=\Big\{A\in\mathscr{A}:\exists B\in\sigma\Big(\bigvee_{i=0}^{n-1}T^{-i}\xi\Big)\text{ with }\mu(A\triangle B)\leqslant\delta\Big\}.$$

It is not hard to verify that $\bigcap_{\delta>0}\bigcup_{n\in\mathbb{N}}\mathscr{A}_\delta^{(n)}$ is a sub-sigma algebra of $\mathscr{A}$ containing $T^{-i}\xi$ for all $i\in\mathbb{N}\cup\{0\}$. Since ξ is a generating partition of X, it follows that

$$\mathscr{A}=\bigcap_{\delta>0}\bigcup_{n\in\mathbb{N}}\mathscr{A}_\delta^{(n)}.\tag{7.4.2}$$

Suppose $\eta=\{B_1,\dots,B_s\}$. From (7.4.2) it follows that for all $\delta>0$ there exists some $n\in\mathbb{N}$ and some sets

$$B_1',\dots,B_s'\in\sigma\Big(\bigvee_{i=0}^{n-1}T^{-i}\xi\Big)$$

such that $d(B_i \triangle B_i') < \delta$ holds for all $i = 1,\dots,s$. Define $\eta' = \{B_1',\dots,B_s'\}$. Since $\bigvee_{i=0}^{n-1} T^{-i}\xi$ is a refinement of η' (which means that every set in η' can be obtained by taking a union of atoms in $\bigvee_{i=0}^{n-1} T^{-i}\xi$), we have that

$$\eta' \vee \bigvee_{i=0}^{n-1} T^{-i}\xi = \bigvee_{i=0}^{n-1} T^{-i}\xi.$$

Using property (iv) of Theorem 116 we see that

$$H\Big(\eta\Big|\bigvee_{i=0}^{n-1} T^{-i}\xi\Big) = H\Big(\eta\Big|\eta' \vee \bigvee_{i=0}^{n-1} T^{-i}\xi\Big) \leqslant H(\eta|\eta').$$

So to finish the proof it suffices to show that $H(\eta|\eta')$ is small. Recall that

$$H(\eta|\eta') = \sum_{i=1}^{s} \mu(B_i')\Big(\sum_{j=1}^{s} \mu(B_i|B_j')\log_2\big(\mu(B_i|B_j')\big)\Big).$$

We have

$$\mu(B_i|B_j') = \frac{\mu(B_i \cap B_j')}{\mu(B_j')}.$$

Using that $B_j' \supseteq B_j \backslash (B_j \backslash B_j')$ we get

$$\mu(B_i \cap B_j') \geqslant \mu(B_i \cap (B_j \backslash (B_j \backslash B_j'))) \geqslant \mu(B_i \cap B_j) - \mu(B_j \backslash B_j') \geqslant \mu(B_i \cap B_j) - \delta.$$

On the other hand, since $B_j' \subseteq B_j \cup (B_j' \backslash B_j)$, we have

$$\mu(B_i \cap B_j') \geqslant \mu(B_i \cap B_j) + \mu(B_j' \backslash B_j) \leqslant \mu(B_i \cap B_j) + \delta.$$

It follows that

$$\mu(B_i|B_j') \begin{cases} \geqslant 1 - \frac{\delta}{\mu(B_i')}, & \text{if } i = j, \\ \leqslant \frac{\delta}{\mu(B_i')}, & \text{if } i \neq j. \end{cases}$$

It follows that

$$H(\eta|\eta') = \sum_{i=1}^{s} \mu(B_i')\Big(\sum_{j=1}^{s} \mu(B_i|B_j')\log_2\big(\mu(B_i|B_j')\big)\Big) \approx_\delta H(\eta|\eta) = 0.$$

□

Proof of the Kolmogorov-Sinaĭ Theorem. Suppose ξ is a generating partition of $(X, \mathscr{A}, \mu, T)$. Our goal is to show that for every finite partition η one has

$$h(T,\eta) \leqslant h(T,\xi),$$

because this will prove $h(T,\xi) = \sup_\eta h(T,\eta) = h(T)$ as desired. So let η be a finite partition of X. Using the definition $h(T,\eta)$ and then property (iii) of Theorem 116 twice, we get

$$
\begin{aligned}
h(T,\eta) &= \lim_{m\to\infty}\frac{1}{m}H\Big(\bigvee_{j=0}^{m-1}T^{-j}\eta\Big)\\
&\leqslant \lim_{m\to\infty}\frac{1}{m}H\Big(\Big(\bigvee_{j=0}^{m-1}T^{-j}\eta\Big)\vee\Big(\bigvee_{j=0}^{m-1}\bigvee_{i=0}^{n-1}T^{-j-i}\xi\Big)\Big)\\
&= \underbrace{\lim_{m\to\infty}\frac{1}{m}H\Big(\bigvee_{j=0}^{m-1}\bigvee_{i=0}^{n-1}T^{-j-i}\xi\Big)}_{[1]} + \underbrace{\lim_{m\to\infty}\frac{1}{m}H\Big(\bigvee_{j=0}^{m-1}T^{-j}\eta\Big|\bigvee_{j=0}^{m-1}\bigvee_{i=0}^{n-1}T^{-j-i}\xi\Big)}_{[2]}.
\end{aligned}
$$

On the one hand, we have

$$
\begin{aligned}
[1] &= \lim_{m\to\infty}\frac{1}{m}H\Big(\bigvee_{i=0}^{m+n-1}T^{-i}\xi\Big)\\
&= \lim_{m\to\infty}\frac{m+n}{m}\Big(\frac{1}{m+n}H\Big(\bigvee_{i=0}^{m+n-1}T^{-i}\xi\Big)\Big) = h(T,\xi).
\end{aligned}
$$

On the other hand,

$$
\begin{aligned}
[2] &\leqslant \lim_{m\to\infty}\frac{1}{m}\sum_{\ell=0}^{m-1}H\Big(T^{-\ell}\eta\Big|\bigvee_{j=0}^{m-1}\bigvee_{i=0}^{n-1}T^{-j-i}\xi\Big)\\
&\leqslant \lim_{m\to\infty}\frac{1}{m}\sum_{\ell=0}^{m-1}H\Big(T^{-\ell}\eta\Big|\bigvee_{i=0}^{n-1}T^{-\ell-i}\xi\Big)\\
&= H\Big(\eta\Big|\bigvee_{i=0}^{n-1}T^{-i}\xi\Big)\xrightarrow{n\to\infty}0.
\end{aligned}
$$

It follows that $h(T,\eta)\leqslant h(T,\xi)$, completing the proof. □

Corollary 122. *If the entropy of $(X,\mathscr{A},\mu,T)$ exceeds λ then every generating partition of $(X,\mathscr{A},\mu,T)$ must contain at least $\lfloor 2^{\lambda}\rfloor$ many atoms.*

Proof. Suppose ξ is a generating partition of $(X,\mathscr{A},\mu,T)$. Then

$$
\lambda < h(T) \underset{\text{Kolmogorov-Sinaĭ Theorem}}{=} h(T,\xi) = \lim_{n\in\mathbb{N}}\frac{1}{n}H\Big(\bigvee_{i=0}^{n-1}T^{-i}\xi\Big) \overset{\text{Fekete's Lemma}}{=} \inf_{n\in\mathbb{N}}\frac{1}{n}H\Big(\bigvee_{i=0}^{n-1}T^{-i}\xi\Big)\leqslant H(\xi) \underset{\text{Property (ii) of Theorem 116}}{\leqslant} \log(r),
$$

where r is the number of atoms in ξ. The claim follows. □

Proposition 123. *Let $(X,\mathscr{A},\mu,T)$ and $(Y,\mathscr{B},\nu,S)$ be measure preserving systems and consider the product system $(X\times Y,\mathscr{A}\otimes\mathscr{B},\mu\otimes\nu,T\times S)$. Then $h(T\times S)=h(T)+h(S)$.*

Theorem 124. *Let* $(X,\mathscr{A},\mu,T)$ *and* $(Y,\mathscr{B},\nu,S)$ *be measure preserving systems and suppose* $(Y,\mathscr{B},\nu,S)$ *is a factor of* $(X,\mathscr{A},\mu,T)$*. Then* $h(S)\leqslant h(T)$*.*

Corollary 125. *Let* $(X,\mathscr{A},\mu,T)$ *and* $(Y,\mathscr{B},\nu,S)$ *be measure preserving systems. If* $(X,\mathscr{A},\mu,T)$ *and* $(Y,\mathscr{B},\nu,S)$ *are isomorphic then* $h(T)=h(S)$*.*

Bibliography

[EW11] M. EINSIEDLER and T. WARD, *Ergodic theory with a view towards number theory*, *Graduate Texts in Mathematics* **259**, Springer-Verlag London, Ltd., London, 2011. MR 2723325 (2012d:37016). https://doi.org/10.1007/978-0-85729-021-2.

[Wal82] P. WALTERS, *An introduction to ergodic theory*, *Graduate Texts in Mathematics* **79**, Springer-Verlag, New York-Berlin, 1982. MR 648108 (84e:28017).