

Ergodic Theory

LECTURE NOTES

Contents

0	Preliminaries from Measure Theory	3
0.1	Algebras and σ -Algebras	3
0.2	Measures and Measure spaces	5
0.3	Measurable Functions and Integrals	10
0.4	L^p Spaces	12
0.5	Convergence Theorems	13
1	Measure Preserving Systems	15
1.1	Definition and Examples	15
1.2	Recurrence	21
1.3	Ergodicity	22
2	Von Neumann's Mean Ergodic Theorem	27
2.1	Koopman Operator	27
2.2	The Splitting $\mathcal{H}_{\text{inv}} \oplus \mathcal{H}_{\text{erg}}$	28
2.3	The Mean Ergodic Theorem	29
2.4	Uniform Mean Ergodic Theorem	30
2.5	Consequences of the Mean Ergodic Theorem	31

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 $\mathcal{P}(X)$ to be the power set of X , that is, the set formed by all subsets of X .

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

- (i) (Contains empty set) We have $\emptyset \in \mathcal{A}$;
- (ii) (Closed under complements) If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$;
- (iii) (Closed under finite unions) If $A_1, \dots, A_k \in \mathcal{A}$, then $A_1 \cup \dots \cup A_k \in \mathcal{A}$.

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

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

If $\mathcal{A}$ on X is an algebra (or a σ -algebra) then a subset of X is called *$\mathcal{A}$ -measurable* if it is an element of the σ -algebra $\mathcal{A}$.

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

Example 2 (Examples of algebras).

- The collection of subsets of X which are either finite or co-finite (meaning that their complement is finite) 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 $\mathcal{A}$ and a countable family of sets $\{A_n\}_{n \in I} \subseteq \mathcal{A}$, $I \subseteq \mathbb{N}$, we have that $\bigcap_{n \in I} A_n \in \mathcal{A}$. This follows from De Morgan's law $(\bigcap_{n \in I} A_n)^c = \bigcup_{n \in I} A_n^c \in \mathcal{A}$, and property (iv). Here are some basic examples of σ -algebras.

Example 3 (Examples of σ -algebras).

- $\mathcal{P}(X)$ and $\{\emptyset, X\}$ are σ -algebras.
- For any subset $A \subseteq X$, $\mathcal{A} = \{\emptyset, A, X \setminus A, X\}$ 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 $\mathcal{A}_1, \mathcal{A}_2 \subseteq \mathcal{P}(X)$, we have that $\mathcal{A}_1 \cap \mathcal{A}_2$ is also a σ -algebra. More generally, for any (possibly uncountable) family of σ -algebras $\mathcal{A}_i \subseteq \mathcal{P}(X)$, $i \in I$, the intersection $\bigcap_{i \in I} \mathcal{A}_i$ is a σ -algebra.

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

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

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 (Examples of σ -algebras generated by collection of subsets).

- The σ -algebra on X consisting of all countable and co-countable subsets of X is the σ -algebra generated by the collection of all singletons of X .
- Let X_1, X_2 be two sets, and $\mathcal{A}_1, \mathcal{A}_2$ be σ -algebras on X_1 and X_2 respectively. We define $\mathcal{A}_1 \otimes \mathcal{A}_2$ 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 $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$. Note that $\mathcal{A}_1 \otimes \mathcal{A}_2$ is called the *product σ -algebra* generated by $\mathcal{A}_1$ and $\mathcal{A}_2$.

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). 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 $\mathcal{B}_X$, or simply $\mathcal{B}$. 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 (Examples of Borel σ -algebras).

- 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 referred to as the *letters* of Σ . The infinite product $\Sigma^{\mathbb{N}}$ can then be identified with the set of all infinite strings in these letters. Observe that the natural topology on Σ is the discrete topology, whose basis consists of singletons, i.e., sets consisting of 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 $\mathcal{M} \subseteq \mathcal{P}(X)$ is a collection of subsets of X having the following properties:

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

Note that both $\mathcal{P}(X)$ 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 $\mathcal{A} \subseteq \mathcal{P}(X)$ be an algebra and let $\mathcal{S}$ be the smallest monotone class containing $\mathcal{A}$. Then we have $\sigma(\mathcal{A}) = \mathcal{S}$.

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 modelling 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 $(X, \mathcal{A})$, where X is a set and $\mathcal{A} \subseteq \mathcal{P}(X)$ is a σ -algebra, is called a *measurable space*, and any set $A \in \mathcal{A}$ is called a *measurable set*.

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

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

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

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

If $(X, \mathcal{A})$ measurable space and μ is a measure on it then the triple $(X, \mathcal{A}, \mu)$ 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 $(X, \mathcal{A}, \mu)$ is said to be a finite measure space if μ satisfies $\mu(X) < \infty$, and if in addition $\mu(X) = 1$, $(X, \mathcal{A}, \mu)$ is called a probability space.

$(X, \mathcal{A}, \mu)$ is called a σ -finite measure space if X is a countable union of elements of $\mathcal{A}$ of finite measure.

We now state different useful properties about measures.

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

- (i) (Finite unions) For any positive integer n and disjoint sets $A_1, A_2, \dots, A_n \in \mathcal{A}$, 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 $A, B \in \mathcal{A}$ and $A \subseteq B$, then $\mu(A) \leq \mu(B)$.
 (iii) (Countable subadditivity) For any countable family of sets $\{A_n\}_{n \in I} \subseteq \mathcal{A}$, $I \subseteq \mathbb{N}$, not necessarily disjoint, we have

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

- (iv) (Continuity) If $A_1 \subseteq A_2 \subseteq \dots \in \mathcal{A}$, then

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

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

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

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 $(X, \mathcal{A}, \mu)$ be a measure space. We say that a property holds μ -almost everywhere on X (sometimes abbreviated as μ -a.e.) if 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 $\mathcal{A}$ be a σ -algebra on X and define $\mu(A) = 0, \forall A \in \mathcal{A}$. Then $(X, \mathcal{A}, \mu)$ is a measure space and μ is referred to as the *null measure* on $(X, \mathcal{A})$.

Counting measure. Let X be a set, and for any $A \in \mathcal{P}(X)$ define $\mu(A) = |A|$, where $|A|$ denotes the cardinality of A . Then $(X, \mathcal{P}(X), \mu)$ 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 $(X, \mathcal{A})$ be a measurable space, and $x \in X$. Then we define the *Dirac measure* δ_x by

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

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

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

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

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

The set function $B \mapsto \mu(B|A)$ is a measure on $\mathcal{A}$ 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 $(X_1, \mathcal{A}_1, \mu_1)$ and $(X_2, \mathcal{A}_2, \mu_2)$ be two measure spaces. Let $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$ 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 $(X, \mathcal{A})$ which satisfies $\mu(A_1 \times A_2) = \mu_1(A_1)\mu_2(A_2)$ for every $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$.

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 , and let $\mathcal{A}$ the Borel σ -algebra on $\Sigma^\mathbb{N}$ as previously introduced in Example 7. 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 coordinates of I .

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

Definition 15 (Hausdorff topological space). 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 (locally compact topological space). 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 $\mathcal{B}$ the Borel σ -algebra defined on X . Then, any measure μ defined on $\mathcal{B}$ 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, $\mathcal{B}$ the Borel σ -algebra defined on X , and μ a finite Borel measure on $\mathcal{B}$. 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 $\mu(K) \geq \mu(X) - \varepsilon$), μ is called a Radon measure. These conditions guarantee that the measure is compatible in some sense with the topology of the space. An useful property of the Radon measure is that it makes the map $f \mapsto \int f d\mu$, 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 $\mathcal{B}_\mathbb{R}$ 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 $\mathcal{B}_{[0,1]}$ 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 $(\mathbb{R}^n, \mathcal{B}_{\mathbb{R}^n})$, the Lebesgue measure μ on $\mathcal{B}_{\mathbb{R}^n}$ 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 $\mathcal{B}_{\mathbb{R}^n}$:

- (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 17 (Atom). Given a measure space $(X, \mathcal{A}, \mu)$, a set $A \in \mathcal{A}$ 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 $(X, \mathcal{A})$ 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 $A \in \mathcal{A}$ 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 18 (Pre-measure). Let $\mathcal{A}$ be an algebra on a set X . A set function $\mu_0: \mathcal{A} \rightarrow [0, \infty]$ is called a *pre-measure* on $(X, \mathcal{A})$ if $\mu_0(\emptyset) = 0$ and, for every countable (or finite) sequence $A_1, A_2, \dots \in \mathcal{A}$ of pairwise disjoint sets whose union lies in $\mathcal{A}$, we have

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

Theorem 19 (Carathéodory). Let $\mathcal{A}$ be an algebra on a set X . Any pre-measure μ_0 on $\mathcal{A}$ extends to a measure μ on the σ -algebra $\sigma(\mathcal{A})$ generated by $\mathcal{A}$. Moreover, if μ_0 is σ -finite then this extension is unique and σ -finite.

Theorem 20 (Riesz-Markov-Kakutani representation theorem). Let X be a locally compact Hausdorff space. Let $\mathcal{C}(X)$ be the Banach space of continuous functions on X with the norm $\|f\|_{\mathcal{C}(X)} = \sup_{x \in X} |f(x)|$. If $l: \mathcal{C}(X) \rightarrow \mathbb{C}$ is a positive linear functional on $\mathcal{C}(X)$, then there exists a unique Radon measure μ on X such that

$$l(f) = \int f \, d\mu,$$

for all $f \in \mathcal{C}(X)$.

0.3. Measurable Functions and Integrals

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

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

Definition 21 (Simple function). A function $f: X \rightarrow \mathbb{R}$ is called *simple* if

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

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

defined to be

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

Definition 22 (Measurable function). A function $g: X \rightarrow \mathbb{R}$ is called measurable if $g^{-1}(A) \in \mathcal{A}$ for any (Borel) measurable set $A \in \mathcal{B}_{\mathbb{R}}$.

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

Proposition 23. *Let $f, g: X \rightarrow \mathbb{R}$ 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 $\max\{f, g\}$

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 24. *Let $g: X \rightarrow \mathbb{R}_{\geq 0}$ 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 $f_n(x) \leq f_{n+1}(x)$ for all $x \in X$ and $n \in \mathbb{N}$) such that $\lim_{n \rightarrow \infty} f_n(x) = g(x)$ for each $x \in X$.*

Definition 25 (Integral of non-negative measurable function). Let $g: X \rightarrow \mathbb{R}_{\geq 0}$ 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 24. Then the *integral* of g is defined to be

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

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

Observe that the expression $\int g \, d\mu$ defined above is guaranteed to exist since $f_n(x) \leq f_{n+1}(x)$ 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 26 (Integral of general measurable function). Given a measurable function $g: X \rightarrow \mathbb{R}$, g has in general a unique decomposition $g = g_+ - g_-$, where $g_+(x) = \max\{g(x), 0\}$ and $g_-(x) = \max\{-g(x), 0\}$ 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

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

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

Proposition 27. *Let $f, g: X \rightarrow \mathbb{R}$. If f is integrable and g is measurable with $|g| \leq f$, 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 28 (Integral over measurable set). Let $f: X \rightarrow \mathbb{R}$ be an integrable function, and A be a measurable set. The *integral of f over A* is defined as

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

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 29 ($\mathcal{L}^p$ space). Let $(X, \mathcal{A}, \mu)$ be a measure space. For $1 \leq p < \infty$, we define the set $\mathcal{L}^p(X, \mathcal{A}, \mu)$ (sometimes also denoted $\mathcal{L}^p(\mu)$) to be the set of all measurable functions $f: X \rightarrow \mathbb{R}$ such that $\int |f|^p \, d\mu < \infty$.

Definition 30 (L^p space). We define an equivalence relation on $\mathcal{L}_\mu^p$ by $f \sim g$ if $\int |f - g|^p \, d\mu = 0$ and we write $L_\mu^p = \mathcal{L}_\mu^p / \sim$ for the space of equivalence classes. Elements of L_μ^p 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 $\|\cdot\|_p$ by:

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

We now give the definition of the L_μ^p in the case $p = \infty$.

Definition 31 (Essential supremum). The *essential supremum* is the generalization to measurable functions of the supremum of a continuous function, and is defined

by:

$$\|f\|_\infty = \inf\{\alpha \in \mathbb{R}_{\geq 0} : \mu(\{x \in X : f(x) > \alpha\}) = 0\}.$$

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

Proposition 32. *Under the norm $\|\cdot\|_p$ previously defined, L_μ^p is a Banach space, for every $1 \leq p \leq \infty$.*

Proposition 33. *For $1 \leq p < q \leq \infty$ we have $L_\mu^p \supsetneq L_\mu^q$ for any finite measure space, with strict inclusion except in some degenerate cases.*

Finally we turn to integration of functions of several variables.

Theorem 34 (Fubini–Tonelli). *Let $(X, \mathcal{A}, \mu)$ and $(Y, \mathcal{B}, \nu)$ be two measure spaces and let f be a non-negative integrable function on the product space $(X \times Y, \mathcal{A} \otimes \mathcal{B}, \mu \otimes \nu)$. Then, for μ -almost every $x \in X$ the function $y \mapsto f(x, y)$ is integrable, and for ν -almost every $y \in Y$ the function $x \mapsto f(x, y)$ is integrable, and we have*

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

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

$$\int f d\mu = \lim_{n \rightarrow \infty} \int f_n d\mu.$$

In particular, if $\lim_{n \rightarrow \infty} \int f_n d\mu < \infty$, 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 behaviour of the integral under pointwise (or almost everywhere) convergence: The integral is lower semi-continuous under almost everywhere convergence.

Theorem 36 (Fatou's lemma). *Let $(f_n)_{n \in \mathbb{N}}$ be a sequence of non-negative real-valued measurable functions on the measure space $(X, \mathcal{A}, \mu)$. Then, $f = \liminf_{n \rightarrow \infty} f_n$ is measurable and*

$$\liminf_{n \rightarrow \infty} \int f_n \, d\mu \geq \int f \, d\mu = \int \liminf_{n \rightarrow \infty} f_n \, d\mu$$

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

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

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

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

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

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

is a probability measure on $\mathcal{A}$ 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 $A \in \mathcal{A}$, we have $\mu(T^{-1}(A)) = \mu(A)$. Thus, saying that T preserving the measure μ (as defined in Definition 38) is equivalent to stating that μ is invariant under T ; the two terms express the same property and we will use them interchangeably throughout these lecture notes.

Example 39. Imagine a computer program with the capability to generate uniformly at random and without bias a real number x in the interval $[0, 1) = \{x \in \mathbb{R} : 0 \leq x < 1\}$. 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 $T: [0, 1) \rightarrow [0, 1)$ 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 40 (Measure preserving system). A measure preserving system is a quadruple $(X, \mathcal{A}, \mu, T)$ where $(X, \mathcal{A}, \mu)$ is a probability space and $T: X \rightarrow X$ is a measure preserving transformation.

Examples

The following examples illustrate the above definitions and serve as a guide for the concepts and results presented in later sections.

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

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

Rotation on m points. Given an integer $m \geq 2$, let $X = \{0, 1, \dots, m-1\}$, which we can identify with the finite cyclic group of order m . Let $\mathcal{A}$ be the power set of $\{0, 1, \dots, m-1\}$ and let $T: \{0, 1, \dots, m-1\} \rightarrow \{0, 1, \dots, m-1\}$ 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 $(X, \mathcal{A}, \mu, T)$ is called *rotation on m points*.

Circle rotations. Let $X = [0, 1)$, endowed with the Borel σ -algebra $\mathcal{A}$ and the Lebesgue measure μ . Given $\alpha \in \mathbb{R}$ we consider the map $T = T_\alpha: X \rightarrow X$ 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 $\mathcal{B}_G$ and the (normalized) Haar measure m_G , then for any fixed $\alpha \in G$, the map $R : x \mapsto x + \alpha$ preserves m_G and hence $(G, \mathcal{B}_G, m_G, R)$ 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 39 above. Take $(X, \mathcal{A}, \mu)$ to be the unit interval $[0, 1)$ equipped with its Borel σ -algebra and Lebesgue measure. Let $T : X \rightarrow X$ 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:

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

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 behaviour 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 $(X, \mathcal{A}, \mu, T)$, where $X = [0, 1)^n$, $\mathcal{A}$ 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, the 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 $x_0, y_0 \in \{0, 1, \dots, N-1\}$. 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

topology¹. In view of Tychonoff's theorem, X is compact. Let $\mathcal{A}$ be the Borel σ -algebra on X generated by the cylinder sets. 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 $T: X \rightarrow X$ 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 $(X, \mathcal{A}, \mu, T)$ 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 $(X, \mathcal{A}, \mu, T)$ where $X = \Sigma^{\mathbb{N}}$, $\mathcal{A}$ 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 $(X, \mathcal{A}, \mu)$, where $X = [0, 1]^2$ is the unit square, $\mathcal{A}$ is the Borel σ -algebra on X and μ is the two-dimensional Lebesgue measure. We define the Baker's map $T: [0, 1]^2 \rightarrow [0, 1]^2$ by :

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

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

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

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

¹This 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 $i \leq n$, then $y \in U$.

the dough, and finally, we compress the result in order to get back the unit square. This process is a chaotic map from the unit square into itself and it has this very nice 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 $(X, \mathcal{A}, \mu, T)$ and $(Y, \mathcal{B}, \nu, S)$, we define their *product* to be the measure preserving system $(Z, \mathcal{C}, \lambda, R)$, where $(Z, \mathcal{C}, \lambda) = (X \times Y, \mathcal{A} \otimes \mathcal{B}, \mu \times \nu)$ is the product of the probability spaces $(X, \mathcal{A}, \mu)$ and $(Y, \mathcal{B}, \nu)$, and $R: Z \rightarrow Z$ is defined as $R(x, y) = (Tx, Sy)$.

Skew-products. Let $X = [0, 1]^2$, let $\mathcal{A}$ be the Borel σ -algebra and let μ be the Lebesgue measure. Fix $\alpha \in \mathbb{R}$ and let $T: X \rightarrow X$ be the map $T(x, y) = (x + \alpha \bmod 1, y + x \bmod 1)$. Then $(X, \mathcal{A}, \mu, T)$ 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 $f, g \in C([0, 1])$

$$\int_0^1 \int_0^1 f(x + \alpha \bmod 1) g(x + y \bmod 1) dy dx = \int_0^1 \int_0^1 f(x) g(y) dy dx,$$

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 $(X, \mathcal{A}, \mu, T)$ be a measure preserving system and let $A \in \mathcal{A}$ with $\mu(A) > 0$. Then for some $n \in \mathbb{N}$ we have*

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

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 $A, T^{-1}A, T^{-2}A, \dots$ 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

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

□

Corollary 41. *Let $(X, \mathcal{A}, \mu, T)$ be a measure preserving system and let $A \in \mathcal{A}$. 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 $x, Tx, T^2x, \dots$ 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 system² 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

²As 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 $x, Tx, T^2x, T^3x, \dots$ 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 ??) 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 42 (Ergodicity). A measure preserving system $(X, \mathcal{A}, \mu, T)$ is *ergodic* if for every set $A \in \mathcal{A}$,

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

Henceforth, let us call a set $A \in \mathcal{A}$ *strictly invariant* if $T^{-1}A = A$ and *almost everywhere invariant* if $\mu(A \Delta T^{-1}A) = 0$. Similarly, we call a measurable function $f: X \rightarrow \mathbb{C}$ *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 43. Let $(X, \mathcal{A}, \mu, T)$ be a measure preserving system. The following are equivalent:

- (i) $(X, \mathcal{A}, \mu, T)$ is ergodic;
- (ii) If $A \in \mathcal{A}$ is almost everywhere invariant then either $\mu(A) = 0$ or $\mu(A) = 1$;
- (iii) If $f: X \rightarrow \mathbb{C}$ is measurable and strictly invariant then f is equal to a constant almost everywhere.
- (iv) If $f: X \rightarrow \mathbb{C}$ 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 $A \in \mathcal{A}$ 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 $f: X \rightarrow \mathbb{C}$ 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) $\implies$ (iii). Suppose $f: X \rightarrow \mathbb{C}$ is a measurable and strictly invariant function. Recall that the essential supremum of a measurable function is defined as

$$\text{esssup } f = \inf\{\alpha \in \mathbb{R} : \mu(\{x \in X : f(x) > \alpha\}) = 0\}.$$

For any $\alpha < \text{esssup } f$ 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 $\alpha < \text{esssup } f$ then this implies that f is almost everywhere equal to $\text{esssup } f$, finishing the proof. $\square$

Examples

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

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

Then $(X, \mathcal{A}, \mu)$ is a finite probability space. A map $T: X \rightarrow X$ 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 $(1 \ 2 \ \dots \ n)$ constitutes an ergodic transformation on $\{1, \dots, n\}$, since the invariant subsets are $\emptyset$ and X . On the other hand, the permutation $(1 \ 2)(3 \ \dots \ n)$ is not ergodic since the sets $\{1, 2\}$ and $\{3, \dots, n\}$ are invariant subsets which have measure $\frac{2}{n}$ and $\frac{n-2}{n}$ respectively.

Circle rotations. Consider the probability space $(X, \mathcal{A}, \mu)$ where $X = [0, 1)$, $\mathcal{A}$ is the Borel σ -algebra on X , and μ is the Lebesgue measure. Given $\alpha \in \mathbb{R}$, consider the rotation by alpha $T: X \rightarrow X$ defined by $Tx = x + \alpha \pmod{1}$. 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 $(X, \mathcal{A}, \mu)$ where $X = [0, 1)$, $\mathcal{A}$ is the Borel σ -algebra on X , and μ is the Lebesgue measure. This time, consider the doubling map $Tx = 2x \pmod{1}$. 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 $[0, 1]^2, \mathcal{A} \otimes \mathcal{A}, \mu \otimes \mu$, and the map $T \times T(x, y) = (px \bmod 1, py \bmod 1)$.

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 $T_A : \mathbb{T}^n \rightarrow \mathbb{T}^n$ 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 $(X, \mathcal{A}, \mu)$ be a probability space and $T : X \rightarrow X$ a measure preserving transformation on it. Fix some $A \in \mathcal{A}$ 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 $\mathcal{A}|_A$, which consists of the restriction of $\mathcal{A}$ on A , i.e. $\mathcal{A}|_A := \{B \cap A : B \in \mathcal{A}\}$. We now define $\mu|_A$ to be the probability measure on A defined by :

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

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

Chapter 2

Von Neumann's Mean Ergodic Theorem

2.1. Koopman Operator

Definition 44 (Koopman operator). Given a measure preserving transformation $T: X \rightarrow X$ on a probability space $(X, \mathcal{A}, \mu)$, we call the linear operator $U_T: L^2(X, \mathcal{A}, \mu) \rightarrow L^2(X, \mathcal{A}, \mu)$ 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 45. The operator U_T is an isometry, which means $\langle U_T f, U_T g \rangle = \langle f, g \rangle$ for all $f, g \in L^2(X, \mathcal{A}, \mu)$. In particular, $\|U_T f\|_{L^2} = \|f\|_{L^2}$ for all $f \in L^2(X, \mathcal{A}, \mu)$.

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

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

which proves that U_T is isometric. □

2.2. The Splitting $\mathcal{H}_{\text{inv}} \oplus \mathcal{H}_{\text{erg}}$

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

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

In view of Proposition 43, the system $(X, \mathcal{A}, \mu, T)$ is ergodic if and only if $\mathcal{H}_{\text{inv}}$ consists only of almost everywhere constant functions.

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

$$f = g - g \circ T \tag{2.2.1}$$

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

$$\mathcal{H}_{\text{erg}} = \overline{\{f \in L^2(X, \mathcal{A}, \mu) : f \text{ is a coboundary}\}}. \tag{2.2.2}$$

Note that $\mathcal{H}_{\text{inv}}$ and $\mathcal{H}_{\text{erg}}$ are both invariant subspaces of $L^2(X, \mathcal{A}, \mu)$ under U_T , by which we mean that $U_T \mathcal{H}_{\text{inv}} \subseteq \mathcal{H}_{\text{inv}}$ and $U_T \mathcal{H}_{\text{erg}} \subseteq \mathcal{H}_{\text{erg}}$. 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 $\mathcal{H}_{\text{erg}}$ is the orthocomplement of $\mathcal{H}_{\text{inv}}$.

Theorem 46. We have $\mathcal{H}_{\text{inv}} \perp \mathcal{H}_{\text{erg}}$ and $\mathcal{H}_{\text{inv}} \oplus \mathcal{H}_{\text{erg}} = L^2(X, \mathcal{A}, \mu)$.

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

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

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

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

The following is an immediate corollary of Theorem 46.

Corollary 47. *For every $f \in L^2(X, \mathcal{A}, \mu)$ there exist unique $f_{\text{inv}} \in \mathcal{H}_{\text{inv}}$ and $f_{\text{erg}} \in \mathcal{H}_{\text{erg}}$ such that*

$$f = f_{\text{inv}} + f_{\text{erg}}. \quad (2.2.3)$$

Note that f_{inv} in (2.2.3) is exactly the orthogonal projection of f onto the space $\mathcal{H}_{\text{inv}}$ and, likewise, f_{erg} is the orthogonal projection of f onto the space $\mathcal{H}_{\text{erg}}$.

2.3. The Mean Ergodic Theorem

Here is Von Neumann's Mean Ergodic Theorem.

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

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} U_T^n f = f_{\text{inv}} \quad \text{in } L^2\text{-norm}, \quad (2.3.1)$$

where f_{inv} is the orthogonal projection of f onto $\mathcal{H}_{\text{inv}}$ as guaranteed by (2.2.3).

Proof. According to (2.2.3) we can write $f = f_{\text{inv}} + f_{\text{erg}}$. Hence

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

Clearly, we have $\frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\text{inv}} = f_{\text{inv}}$, because f_{inv} is invariant under U_T . Thus, to finish the proof of (2.3.1), it suffices to show

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} U_T^n f_{\text{erg}} = 0 \quad \text{in } L^2\text{-norm} \quad (2.3.2)$$

for all $f_{\text{erg}} \in \mathcal{H}_{\text{erg}}$. In view of (2.2.2), we can assume that f_{erg} is a coboundary, i.e., there exists $g \in L^2(X, \mathcal{A}, \mu)$ such that $f_{\text{erg}} = g - U_T g$. But if $f_{\text{erg}} = g - U_T g$ then the sum in (2.3.2) is telescoping, yielding

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

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

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

and (2.3.2) follows. $\square$

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

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} U_T^n f = \int f \, d\mu \quad (2.3.3)$$

in L^2 -norm.

Proof. In light of (2.3.1), it suffices to show that if the system $(X, \mathcal{A}, \mu, T)$ is ergodic then $f_{\text{inv}} = \int f \, d\mu$. So assume $(X, \mathcal{A}, \mu, T)$ is ergodic and let $f \in L^2(X, \mathcal{A}, \mu)$ be arbitrary. According to Corollary 47, there exist unique $f_{\text{inv}} \in \mathcal{H}_{\text{inv}}$ and $f_{\text{erg}} \in \mathcal{H}_{\text{erg}}$ such that $f = f_{\text{inv}} + f_{\text{erg}}$. By definition, f_{inv} is an almost everywhere invariant function. Therefore, by part (iv) of Proposition 43, f_{inv} is almost everywhere equal to a constant, which we denote by c . To finish the proof of (2.3.3), it only remains to show that $\int f \, d\mu = c$. Let $\mathbf{1}$ denote the function that is constant equal to 1 everywhere. Then

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

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

2.4. Uniform Mean Ergodic Theorem

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

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

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

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

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

$$\lim_{N-M \rightarrow \infty} \frac{1}{N-M} \sum_{n=M}^{N-1} U_T^n f = f_{\text{inv}} \quad \text{in } L^2\text{-norm,} \quad (2.4.1)$$

where f_{inv} is the orthogonal projection of f onto $\mathcal{H}_{\text{inv}}$ as guaranteed by (2.2.3).

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. $\square$

2.5. Consequences of the Mean Ergodic Theorem

Corollary 48. *A measure preserving system $(X, \mathcal{A}, \mu, T)$ is ergodic if and only if for every $A, B \in \mathcal{A}$,*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} \mu(T^{-n}A \cap B) = \mu(A)\mu(B). \quad (2.5.1)$$

Proof. If the system is not ergodic, then by definition there exists a strictly invariant set $A \in \mathcal{B}$ with $0 < \mu(A) < 1$. Taking $B = X \setminus A$, we see that $\mu(A)\mu(B) > 0$ but $T^{-n}A \cap B = \emptyset$ for every n , contradicting (2.5.1).

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

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

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

$$\lim_{N \rightarrow \infty} \int \left(\frac{1}{N} \sum_{n=0}^{N-1} U_T^n 1_A \right) \cdot 1_B \, d\mu = \int \mu(A) \cdot 1_B \, d\mu = \mu(A)\mu(B),$$

completing the proof. $\square$

Setting $A = B$ in Corollary 48 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 49. 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 .

Khinchine's recurrence theorem. *Let $(X, \mathcal{B}, \mu, T)$ be a measure preserving system, let $A \in \mathcal{B}$ 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*

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

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 $N_k - M_k \rightarrow \infty$ as $k \rightarrow \infty$ and such that $\mu(A \cap T^{-n}A) \leq \mu^2(A) - \varepsilon$ for all $n \in [M_k, N_k)$. Applying Uniform Mean Ergodic Theorem to the indicator function 1_A of A we have

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

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

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

It follows that

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

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

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)**.