Lecture notes in ergodic theory

Mikko Stenlund

March 30, 2015

(Please send feedback to mikko.stenlund@helsinki.fi)

Contents

Preface 5

Chapter 1. Introduction 7

Chapter 2. Continuous and measurable transformations 9
1. Ergodic theorem for continuous transformations 9
2. Pushforward and invariant measures 12
3. Existence of invariant measures 15
4. Extreme points of the convex set $\mathcal{P}_T$ for continuous transformations 20
5. Ergodicity 22
6. Ergodic decomposition: a (very) brief overview 28

Chapter 3. Recurrence and ergodicity 31
1. Recurrence theorems of Poincaré and Kac 31
2. Preliminaries: Conditional expectation 34
3. Von Neumann's mean ergodic theorem 38
4. Birkhoff's almost sure ergodic theorem 41
5. Maximal ergodic theorem 43

Chapter 4. Mixing 47
1. Weak and strong mixing 47
2. Preliminaries: Convergence of conditional expectations along nested sequences of sub-sigma-algebras 52
3. Kolmogorov automorphisms and exact endomorphisms 53
4. Bernoulli and Markov shifts 55

Chapter 5. Introduction to information and entropy 61
1. Concepts 61
2. Basic identities for information and entropy 64
3. Shannon–McMillan–Breiman theorem 66
4. Applications of the S–M–B theorem 70

Appendix A. Some measure and integration theory 71
1. Positive, signed and complex measures 71
2. Monotone class theorem and π–λ lemma 72
3. Monotone and dominated convergence 73
4. Radon–Nikodym theorem 74

Appendix B. Some functional analysis 75
1. Projection operators 75
2. Extension of continuous linear functionals 76
3. Identifying subspaces using continuous linear functionals 76
4. Compact convex subsets of locally convex spaces 77

5. Representation of functionals as measures 79
6. Weak topology and convergence of measures 80

Appendix. Bibliography 85

Appendix. Index 87

Preface

These are the lecture notes for a graduate course in ergodic theory taught by the author at the University of Helsinki in the spring semester of 2015. None of the material presented is the original work of the author. Rather, the notes attempt to give a concise introduction to the subject, assuming no prior exposure to it, by collecting topics of the author's choice from various sources. Consequently, the outcome is a cross section of ergodic theory intended to help the reader quickly embark on more comprehensive studies in the field with the aid of research literature and the mentioned references.

There are many books containing excellent accounts of ergodic theory of varying extent. We particularly mention the classics by Walters [17], Petersen [10], Sinai [16] and Parry [9]. Brin and Stuck [2] includes a concise and clear introduction to ergodic theory, in addition to different topics in the theory of dynamical systems. Pollicott and Yuri [12] is more geared for the undergraduate student, but contains interesting material and many applications not discussed in these notes. In addition to the published books, the lecture notes by Sarig [15], Hochman [7] and Bakhtin [1] are highly recommended. This set of notes has been influenced by all the preceding references, with special emphasis on Parry's book.

It is assumed that the reader is familiar with the basics of measure and integration theory, as well as elementary topology, and is comfortable with the notions of Banach and Hilbert spaces (in particular L^p spaces) and bounded linear operators on them. We will take advantage of several results in real and functional analysis, which, however, are all recalled in the appendices; general references in this regard include Rudin [13, 14], and Dunford and Schwartz [3], while Phelps [11] treats the Choquet theory of compact convex sets. It will help the reader to have some knowledge of probability theory, especially the concept of conditional expectation given a sigma-algebra, but it is not required. The few facts we need from probability theory are explained in the text, using Durrett [4] and Williams [18] as references.

CHAPTER 1

Introduction

Deterministic time evolution of a "system" is abstractly modeled by a family of maps

$$T_t : X \to X \ , \quad t \in \mathbb{R}_+ \text{ or } \mathbb{N} \ ,$$

depending on whether time is continuous or discrete. Here the set X is called the state space and its elements are called states of the system. If $x \in X$ is the state of the system at time 0, then its state at any time $t \geq 0$ is $T_t(x) \in X$. In particular, $T_0 = \mathrm{id}_X$. By determinism, it is evident that

$$T_t \circ T_s = T_{t+s}$$

for all $s, t \geq 0$. The identity above simply means that evolving the initial state first by s time units and then by another t time units results in the same final state as evolving the initial state directly by $t + s$ time units. In particular,

$$T_n = \underbrace{T_1 \circ \cdots \circ T_1}_{n\text{-fold composition}} = T_1^n \ , \quad n \in \mathbb{N} \ .$$

In discrete time, deterministic time evolution is consequently specified by the compositions $T^n = T \circ \cdots \circ T$ — or iterates — of a single map

$$T : X \to X \ .$$

Example 1.1. *Suppose $X \subset \mathbb{R}^d$ and $F : X \to \mathbb{R}^d$ are such that the initial value problem*

$$y' = F(y) \ , \quad y(0) = x \tag{1.1}$$

has a unique solution $y : \mathbb{R}_+ \to X$ for all $x \in X$. In other words, given $x \in X$, there exists a unique function satisfying $y'(t) = F(y(t))$, for all $t \in \mathbb{R}_+$, together with the initial condition $y(0) = x$. Let us write $T_t(x) = y(t)$ to emphasize the initial condition. The function $\tilde{y}(t) = y(t + s) = T_{t+s}(x)$ ($s \geq 0$ fixed) satisfies

$$\tilde{y}'(t) = y'(t + s) = F(y(t + s)) = F(\tilde{y}(t)) \ , \quad \tilde{y}(0) = y(s) \ .$$

By uniqueness, we must have $\tilde{y}(t) = T_t(y(s))$, or $T_{t+s}(x) = T_t(T_s(x))$.

Broadly speaking, ergodic theory can be viewed as the study of the statistical behavior of the trajectories $(T^t(x))_{t\geq 0}$ of a deterministic system. The topic has its origin in physics, more precisely in statistical mechanics and thermodynamics. Namely, suppose equation (1.1) describes the motion of a large number of gas molecules in a container. A physical quantity, say the temperature of the gas, is represented by a function

$$f : X \to \mathbb{R} \ ,$$

because it is determined by the state of the system. Then $f(T^t(x))$ is the temperature of the gas at time t, assuming the initial state of the system was $x \in X$. Physicists were routinely assuming that the average value $\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$ of repeated measurements of the temperature would in the long run simply be computable from the average value of f over all possible states:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) = \int_X f(x)\,\mathrm{d}x \ .$$

There was subsequently a need to justify such assumptions mathematically.

Although ergodic theory is still relevant to physics, it has found applications in many other fields of science and areas of mathematics. These include ecology, biology, information theory, probability theory and number theory, to name some. For this reason it was left intentionally vague what the word "system" above means; it depends on the situation. Likewise, "deterministic time evolution" does not exclude random behavior. To get a feel of the various applications ergodic theory might have, here are two examples of results that can be proved with ergodic-theoretic means:

Example 1.2 (Law of large numbers). *Suppose $Y_1, Y_2, \dots$ is a sequence of real-valued, independent and identically distributed random variables such that the expected value $\mathrm{E}(|Y_1|) < \infty$. Then, almost surely,*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} Y_i = \mathrm{E}(Y_1) \ .$$

For instance, if one tosses a fair coin infinitely many times, then the fraction of head obtained tends to 50%.

Example 1.3 (Szemerédi). *If $N \subset \mathbb{N}$ is an arbitrary set of natural numbers having positive upper density, then N contains an arithmetic progression*

$$m, m+k, m+2k, \dots, m+(n-1)k$$

of an arbitrary length n.

Here $m, k \in \mathbb{N}$ generally depend on n. N is said to have positive upper density if there exist sequences $a_n, b_n \in \mathbb{N}$, $n \geq 1$, and $\delta > 0$ such that $\lim_{n\to\infty} b_n - a_n = \infty$ and

$$\frac{|N \cap [a_n, b_n]|}{|[a_n, b_n]|} > \delta$$

for all $n \geq 1$. ($[a_n, b_n] = \{a_n, a_n + 1 \dots, b_n\}$ and $|A|$ denotes the number of elements in a set $A \in \mathbb{N}$.)

In these notes we concentrate on discrete time, in two settings:

(1) X is a compact metric space and T is continuous.
(2) $(X, \mathscr{B})$ is a measurable space and T is measurable.

Of course, a topological space X can always be endowed with the Borel sigma-algebra, which renders every continuous map measurable. While the second setting is more general, it will be beneficial to start from the first one.

CHAPTER 2

Continuous and measurable transformations

We begin our journey into ergodic theory by considering continuous transformations $T : X \to X$ of a compact metric space X. This means that the preimage $T^{-1}U \subset X$ is open for all open sets $U \subset X$. Soon after, we will consider measurable transformations $T : X \to X$ of a measurable space $(X, \mathscr{B})$. This means that X is a set, $\mathscr{B}$ is a sigma-algebra on X, and the preimage $T^{-1}A \subset X$ is measurable for all measurable sets $A \subset X$ (that is, $T^{-1}A \in \mathscr{B}$ for all $A \in \mathscr{B}$). When X is a compact metric space, we endow it with the Borel sigma-algebra (the smallest sigma-algebra containing all open sets), after which any continuous transformation is measurable. In other words, the latter setting is more general than the former.

In these notes the words "map" and "transformation" are used interchangeably. Likewise, "a transformation of a space" and "a transformation on a space" mean the same thing.

1. Ergodic theorem for continuous transformations

In this section X is a compact metric space and $T : X \to X$ is a continuous transformation. It is informative to begin by introducing certain results for this special — albeit interesting — case, because measure theory will not play a role in the discussion. We will nevertheless encounter ideas that will be useful in the study of general measurable transformations of measurable spaces. The Banach space of continuous functions $f : X \to \mathbb{C}$ equipped with the uniform norm $\|f\|_\infty = \sup_{x\in X} |f(x)|$ is denoted by $C(X)$. Some standard results from functional analysis will be used, which are recalled in Appendix B.

By ergodic theorems we mean results concerning the convergence of the **time averages**[1]

$$\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$$

of a function $f : X \to \mathbb{C}$ along the **trajectory** $(T^i(x))_{i=0}^{\infty}$ of a point $x \in X$. If the limit exists *for a given function* f *and a given point* x, we denote it by $f_+(x)$:

$$f_+(x) = \lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) \ . \tag{2.1}$$

In this section we always assume that $f \in C(X)$. Nevertheless, the above limit is to be understood one point at a time, as the limit of a sequence of complex numbers.

We begin by defining two special classes of functions.

[1] Such averages are also called ergodic averages or Birkhoff averages.

Definition 2.1. *A function $f \in C(X)$ is called* **invariant**

$$f = f \circ T \ .$$

Thus,

$$I = \{f \in C(X) \, : \, f = f \circ T\}$$

is the set of all invariant functions.

A function $f \in C(X)$ is called a **coboundary** *if there exists $g \in C(X)$ such that*

$$f = g - g \circ T \ .$$

Thus,

$$B = \{g - g \circ T \, : \, g \in C(X)\}$$

is the set of all coboundaries.

Both I and B are clearly linear subspaces of $C(X)$. Moreover, I is closed but B is generally *not* closed. Note that $f \in \bar{B}$ if and only if there exists a sequence $f_n \in B$, $n \geq 1$, which converges to f in $C(X)$: $\lim_{n\to\infty} \|f_n - f\|_\infty = 0$.

Let us immediately make the following observations:

Lemma 2.2. *Let $f \in C(X)$.*

(1) *If $f_+(x)$ exists for some $x \in X$, then $(f \circ T)_+(x)$ and $f_+(T(x))$ exist and*
$$(f \circ T)_+(x) = f_+(T(x)) = f_+(x) \ .$$
(2) *If $f_+(x)$ exists for all $x \in X$ and if $f_+ \in C(X)$, then $f_+ = f_+ \circ T$. That is,*
$$f_+ \in I \ .$$

Proof. The first claim follows from $\frac{1}{n}\sum_{i=0}^{n-1}(f \circ T) \circ T^i = \frac{1}{n}\sum_{i=0}^{n-1}(f \circ T^i) \circ T = \frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i + \frac{1}{n}(f \circ T^n - f)$, because f is bounded. The second claim is a corollary of the first one. □

Invariant functions and coboundaries play an important role in ergodic theory. For example, we have the following result:

Lemma 2.3. *For invariant functions,*

$$f_+ = f \ , \quad f \in I \ ,$$

where the convergence is uniform. For (limit points of the set of) coboundaries,

$$f_+ = 0 \ , \quad f \in \bar{B} \ ,$$

where the convergence is uniform. In particular,

$$I \cap \bar{B} = \{0\} \ .$$

Proof. (1) If $f \in I$, then $\frac{1}{n}\sum_{i=0}^{n-1} f \circ (T^i(x)) = \frac{1}{n}\sum_{i=0}^{n-1} f(x) = f(x)$ for all $x \in X$.

(2) On the other hand, if $f \in B$, there exists $g \in C(X)$ such that $f = g - g \circ T$. Then $\frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i = \frac{1}{n}g - \frac{1}{n}g \circ T^n$, which converges to 0 uniformly by the boundedness of g. The result now extends to all $f \in \bar{B}$ by approximation: given any $\varepsilon > 0$, there exist $f_\varepsilon \in B$ such that $\|f - f_\varepsilon\|_\infty \leq \frac{\varepsilon}{2}$ and N_ε such that $\left\|\frac{1}{n}\sum_{i=0}^{n-1} f_\varepsilon \circ T^i\right\|_\infty \leq \frac{1}{2}\varepsilon$ for all $n \geq N_\varepsilon$. Since $\left\|\frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i - \frac{1}{n}\sum_{i=0}^{n-1} f_\varepsilon \circ T^i\right\|_\infty \leq \frac{1}{2}\varepsilon$, we see that $\left\|\frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i\right\|_\infty \leq \varepsilon$ if $n \geq N_\varepsilon$. Because ε was arbitrary, $\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i = 0$ uniformly.

(3) If $f \in I \cap \bar{B}$ then $f = f_+ = 0$ by the preceding parts of the lemma. □

Clearly, the subspaces I and $\bar{B}$ are special for the convergence of time averages. We have seen that $f_+(x)$ exists for all $x \in X$ and that $f_+ \in C(X)$, provided $f \in I \oplus \bar{B}$.[2] In fact, the next theorem shows that this is true if and only if $f \in I \oplus \bar{B}$. The theorem we present is a slight generalization of a similar theorem appearing in William Parry's textbook [9].

Theorem 2.4 (Ergodic theorem for continuous transformations)**.** *Let X be a compact metric space and $T : X \to X$ a continuous transformation. Let*

$$A_p = \{f \in C(X) \,:\, f_+(x) \text{ exists for all } x \in X \text{ and } f_+ \in C(X)\}$$

and

$$A_u = \left\{ f \in C(X) \,:\, \frac{1}{n}\sum_{i=0}^{n-1} f \circ T^i \text{ converges uniformly to } f_+ \in C(X) \right\}$$

be the subspaces of $C(X)$ on which the time averages converge pointwise and uniformly, respectively, to a continuous limit. Then

$$A_p = A_u = I \oplus \bar{B} \ .$$

A few remarks on the theorem are in order. First, the only distinction of A_p and A_u is that if f belongs to the former, then the convergence to f_+ is pointwise (for every point of X), whereas in the case of the latter it is uniform. In the letter case $f_+ \in C(X)$ is automatic, but in the former this has to be imposed separately. Thus, it has turned out that *pointwise convergence of the time averages to a continuous limit is equivalent to uniform convergence of the time averages.* Moreover, this happens if and only if f can be written as $f = g + h$, where g is invariant and h is a coboundary or, more generally, the limit of a uniformly converging sequence of coboundaries. The second remark is that the theorem says absolutely nothing about the size of the subspaces A_p, A_u and $I \oplus \bar{B}$ — just that they coincide. The theorem in Parry's book is the special case that $A_p = C(X) \Leftrightarrow A_u = C(X) \Leftrightarrow I \oplus \bar{B} = C(X)$. Another ergodic theorem (Theorem 2.24) for continuous transformations, which identifies a sufficient (but not necessary) condition for $A_p = A_u = I \oplus \bar{B} = C(X)$, will be proved shortly, once enough machinery has been developed.

Of course, we are assuming here that X is a compact metric space and T is continuous, and we are working on the space of continuous functions $C(X)$. Nevertheless, this theorem is a good place to embark on a journey toward more general results. In particular, it outlines a strategy to prove the much harder ergodic theorem of Birkhoff for measurable transformations later on.

Proof of Theorem 2.4. We follow [9]. It is clear that A_p and A_u are linear normed subspaces of $C(X)$. It is equally clear that $I \oplus \bar{B} \subset A_u \subset A_p$, by Lemma 2.3. It remains to show that $A_p \subset I \oplus \bar{B}$. To this end, we define a continuous linear projection operator on the normed space A_p by assigning to each f its limit time average f_+:

$$P : A_p \to I : Pf = f_+ \ .$$

[2]We write $I \oplus \bar{B}$ for $I + \bar{B}$ to remind the reader that I and $\bar{B}$ are closed subspaces and $I \cap \bar{B} = \{0\}$.

This is indeed defined for every $f \in A_p$. Since f_+ is an invariant function and $(f_+)_+ = f_+$, we have $\operatorname{im} P = I \subset A_p$ and $P^2 = P$, so P is indeed a projection operator on A_p. (The reader should check that P is linear and continuous.) By Lemma B.1,

$$A_p = \operatorname{im} P + \ker P = I + \ker P \ , \quad I \cap \ker P = \{0\} \ ,$$

where $\ker P$ is a subspace of A_p, closed in the subspace topology.

Since $\bar{B} \subset \ker P$, the remaining task is to show that $\bar{B} = \ker P$. To this end, we employ Lemma B.3: it suffices to check that if $L : A_p \to \mathbb{C}$ is an arbitrary continuous linear functional which vanishes on $\bar{B}$ (in particular on B), then it also vanishes on $\ker P$. Let L be such a functional. Since $f \circ T^i - f \circ T^{i+1} \in B$ for all $i \geq 0$ and all $f \in A_p$, we have $L(f \circ T^i - f \circ T^{i+1}) = 0$, which results in $L(f) = L(f \circ T) = L(f \circ T^2) = \cdots$.[3] On the other hand, if $f \in \ker P$, then $f_+ = Pf = 0$, which means that

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i = 0 \quad \underline{\text{pointwise}}. \tag{2.2}$$

If we can exchange the order of the limit and the functional L in

$$0 = L(0) = L\left(\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i\right) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} L(f \circ T^i) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} Lf = Lf \ ,$$

then L indeed vanishes on $\ker P$ and we are done. The subtlety here is that the convergence in (2.2) is not uniform. (If it were uniform, the desired property would follow from the continuity of L.) Instead, since the sequence $\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i$, $n \geq 1$, is uniformly bounded, we can appeal to the dominated convergence theorem for linear functionals (Theorem B.9) and reach the desired conclusion.[4] □

2. Pushforward and invariant measures

Ergodic theory concerns the study of the trajectories $(T^i(x))_{i=0}^{\infty}$ of the initial points $x \in X$ for a map $T : X \to X$. Given x, its trajectory is completely determined by the iterates of T, i.e., it is deterministic. Nevertheless, it is highly interesting to *pick the initial point x randomly*, according to some probability measure m. This means that x is a random variable, with values in X, such that

$$\text{Probability}(x \in A) = \mathrm{m}(A) \tag{2.3}$$

for any measurable set $A \subset X$. (We say that x is distributed according to m, or that m is the distribution of x.) For each realized value of x its trajectory is still given deterministically by the iterates of T — but since x is random, also the trajectory $(T^i(x))_{i=0}^{\infty}$ is random; it is actually a sequence of random variables $T^i(x)$, $i \geq 0$. (Sequences of random variables are called stochastic processes). For a random initial point it makes sense to ask how the trajectory behaves with a high probability, and so on. An obvious question arises: if the distribution of x is m, what is the distribution of $T(x)$ (and of the higher iterates

[3] Note that each $f \circ T^j$ belongs to the domain A_p of L since $(f \circ T^j)_+ = f_+$ by Lemma 2.2.

[4] Strictly speaking, the dominated convergence theorem applies to continuous linear functionals defined on the entire space $C(X)$, while the functional L is only defined on the subspace A_p. The way around this is to use the Hahn–Banach extension theorem (Theorem B.2): the functional L admits an extension to a functional $\ell : C(X) \to \mathbb{C}$ which is continuous, with the same norm as L, and coincides with L on the subspace A_p. To complete the proof one simply writes ℓ in place of L below (2.2) in the proof.

$T^i(x)$, $i \geq 2$)? More precisely, given a measurable set $A \subset X$, what is the probability of the event $Tx \in A$? Since $Tx \in A \Leftrightarrow x \in T^{-i}A$, in view of (2.3) we have

$$\text{Probability}(T(x) \in A) = \mathrm{m}(T^{-1}A) \ , \tag{2.4}$$

provided that $T^{-1}A$ is measurable. The right side specifies the distribution of $T(x)$, which leads to the following definition:

Definition 2.5. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. If m is a measure on $(X, \mathscr{B})$, its* **pushforward** *is the measure $T_*\mathrm{m}$ on $(X, \mathscr{B})$ defined by*

$$T_*\mathrm{m}(A) = \mathrm{m}(T^{-1}A) \ , \quad A \in \mathscr{B} \ .$$

This also defines the map

$$T_* : \mathrm{m} \mapsto T_*\mathrm{m} \ ,$$

from the set of all measures on $(X, \mathscr{B})$ into itself.

Exercise 2.6. *Check that $T_*\mathrm{m}$ is indeed a measure on $(X, \mathscr{B})$, and that it is a probability measure if m is. If T_*^n denotes the n-fold composition of T_*, check that*

$$T_*^n = (T^n)_* \ , \quad n \geq 1 \ .$$

Exercise 2.7. *A family of probability measures we often encounter are the measures δ_x, $x \in X$, defined by*

$$\delta_x(A) = 1_A(x) \ , \quad A \in \mathscr{B} \ .$$

In other words, δ_x assigns measure 1 to any set containing x (including the singleton $\{x\}$ if it is measurable) and measure 0 otherwise. It is called the **point mass** *at x. Prove that*

$$T_*\delta_x = \delta_{T(x)} \ .$$

To recapitulate, (2.4) sates that if m is the distribution of x, then the pushforward $T_*\mathrm{m}$ is the distribution of $T(x)$. Note also that $(T^2)_*\mathrm{m} = T_*(T_*\mathrm{m})$ is the distribution of $T^2(x)$ and so on. Of course, in general the distributions of x and $T(x)$ are different, $\mathrm{m} \neq T_*\mathrm{m}$. For a special measure m, which suitably reflects the properties of the map T, $\mathrm{m} = T_*\mathrm{m}$ is possible. This motivates the following definition:

Definition 2.8. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. A measure m on $(X, \mathscr{B})$ is called an* **invariant measure of** T, *if*

$$T_*\mathrm{m} = \mathrm{m} \qquad \textit{meaning} \qquad \mathrm{m}(T^{-1}A) = \mathrm{m}(A) \ , \quad A \in \mathscr{B} \ .$$

In this case $(X, \mathscr{B}, \mathrm{m}, T)$ — or just T — is called a **measure-preserving transformation**, *briefly* **mpt**. *If m is a probability measure, it is called a* **probability-preserving transformation**, *briefly* **ppt**.

Thus, invariance of the initial measure m means precisely that x and $T(x)$ are identically distributed. It follows that all $T^i(x)$, $i \geq 0$, are identically distributed. (In fact, the stochastic process $(T^i(x))_{i\geq 0}$ is then stationary.) Beware, however, that the random variables $T^i(x)$ are far from being independent: if $T^i(x)$ is known, then $T^{i+j}(x)$ is completely determined for all $j \geq 1$.

Convention. In these notes a measure is positive ($\mathrm{m}(A) \geq 0$ for all $A \in \mathscr{B}$) and finite ($\mathrm{m}(X) < \infty$), unless otherwise stated; see Section 1 in Appendix A. Such a measure can be normalized to a probability measure (m is positive and $\mathrm{m}(X) = 1$) excluding the trivial case in which the measure is identically zero. Thus, the study of mpts reduces to the study of ppts, which is what we will focus on.

To avoid any confusion, let us record this as a definition; see Section 1 of Appendix A for the terminology used.

Definition 2.9. *Let $(X, \mathscr{B})$ be a measurable space. We write $\mathcal{M}(X, \mathscr{B})$ for the set of finite (positive) measures, $\mathcal{M}_s(X, \mathscr{B})$ for the set of signed measures, and $\mathcal{P}(X, \mathscr{B})$ for the set of probability measures. If T is a measurable transformation of $(X, \mathscr{B})$, we write $\mathcal{P}_T(X, \mathscr{B})$ for the set of invariant probability measures. When the underlying measurable space is clear, we simply write $\mathcal{M}$, $\mathcal{M}_s$, $\mathcal{P}$ and $\mathcal{P}_T$. Clearly $\mathcal{P}_T \subset \mathcal{P} \subset \mathcal{M} \subset \mathcal{M}_s$.*

As mentioned, for a measure m to be invariant it has to be special and take into account the properties of the map T. Here are two examples of ppts we will return to several times later on:

Exercise 2.10 (Rotations of the circle). *Let X be the circle $\mathbb{S}^1$ obtained from the unit interval $[0, 1]$ by identifying the end points, and let $T : X \to X : x \mapsto x + \alpha \pmod 1$, where $\alpha \in (0, 1)$ is a fixed number. Note that T describes a rotation of the circle with circumference 1 by the angle $2\pi\alpha$. Show that the Lebesgue measure is invariant. What does this mean from the probabilistic point of view above?*

Exercise 2.11 (Angle doubling map). *Let X be the circle $\mathbb{S}^1$ as above, and let T be the angle doubling map $T : X \to X : x \mapsto 2x \pmod 1$. Show that the Lebesgue measure is invariant. What does this mean from the probabilistic point of view above?*

At this point, it is not clear whether an invariant measure exists for a given map and whether there could be more than one. We can, however, already make a useful elementary observation:

Lemma 2.12. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. Then $\mathcal{P}$ and $\mathcal{P}_T$ are convex subsets of $\mathcal{M}$.*

Exercise 2.13. *Prove Lemma 2.12.*

In particular, Lemma 2.12 sheds some light on the uniqueness issue: if there are two invariant measures $\mathrm{m}_1 \neq \mathrm{m}_2$, then there are *uncountably many* of them, as $(1-t)\mathrm{m}_1 + t\mathrm{m}_2$ is invariant for each $t \in [0, 1]$. We will return to such considerations later, and for this reason the following trivial warmup example is in order:

Exercise 2.14. *Let $X = [0, 1]$ and let T be the identity map id_X. Show that any $\mathrm{m} \in \mathcal{P}$ is invariant. In particular, δ_x is an invariant measure for any $x \in X$.*

The preceding example demonstrates that there exist maps with uncountably many invariant probability measures — corresponding to the extreme points of the convex set $\mathcal{P}_T$ — which cannot be expressed as convex combinations of any other invariant probability measures. It turns out that this feature is quite typical. For instance, the angle doubling map and rotations of the circle by a rational angle have it.

Let us also present a useful characterization of pushforward and invariant measures:

Lemma 2.15. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map.*

(1) Given two measures $\mathrm{m}, \mathrm{m}' \in \mathcal{P}$, we have $\mathrm{m}' = T_*\mathrm{m}$ if and only if

$$\int_X f \, \mathrm{dm}' = \int_X f \circ T \, \mathrm{dm} \; , \qquad f \in L^1(X, \mathscr{B}, \mathrm{m}') \; .$$

(2) A measure $\mathrm{m} \in \mathcal{P}$ *is invariant if and only if*

$$\int_X f \, \mathrm{dm} = \int_X f \circ T \, \mathrm{dm} \ , \qquad f \in L^1(X, \mathscr{B}, \mathrm{m}) \ .$$

Equivalently, the classes $L^1(X, \mathcal{B}, \mathrm{m}')$ and $L^1(X, \mathcal{B}, \mathrm{m})$ can be replaced by bounded measurable functions, bounded nonnegative measurable functions, simple functions, or indicator functions of measurable sets as is evident from the proof below. The formulation above was chosen because if, say, the invariance of m has been established, then the integral identity is immediately at our disposal for all absolutely integrable functions. But to reiterate, in the opposite direction it is sufficient to check the integral identity for indicator functions in order to infer that m is invariant.

Proof of Lemma 2.15. (1) If $\mathrm{m}' = T_*\mathrm{m}$, the integral identity holds for indicator functions 1_A, $A \in \mathscr{B}$:

$$\int_X 1_A \, \mathrm{dm}' = \mathrm{m}'(A) = \mathrm{m}(T^{-1}A) = \int_X 1_{T^{-1}A} \, \mathrm{dm} = \int_X 1_A \circ T \, \mathrm{dm} \ .$$

The property extends first to all nonnegative $\mathscr{B}$-measurable functions f by the monotone convergence theorem (Theorem A.3), because there exists an increasing sequence of simple functions $s_n \uparrow f$. Then it extens to all $f \in L^1(X, \mathscr{B}, \mathrm{m}')$ by the decomposition of $f = f_+ - f_-$ into its positive and negative part. (In particular, we have shown that $f \in L^1(X, \mathscr{B}, T_*\mathrm{m}) \Rightarrow f \circ T \in L^1(X, \mathscr{B}, \mathrm{m})$.) In the opposite direction, let $A \in \mathscr{B}$ and set $f = 1_A$. As above, the integral identity yields $\mathrm{m}'(A) = \mathrm{m}(T^{-1}A)$, so that $\mathrm{m}' = T_*\mathrm{m}$. (2) is an immediate consequence of (1) by setting $\mathrm{m}' = \mathrm{m}$. □

We finish with an exercise introducing a concept needed later.

Exercise 2.16. *Let* $(X_i, \mathscr{B}_i, \mathrm{m}_i, T_i)$, $i \in \{1, 2\}$, *be probability-preserving transformations. The* **product transformation** $(X_1 \times X_2, \mathscr{B}_1 \times \mathscr{B}_2, \mathrm{m}_1 \times \mathrm{m}_2, T_1 \times T_2)$ *is defined by*

$$(T_1 \times T_2)(x, y) = (T_1(x), T_2(y)) \ .$$

Show that $T_1 \times T_2$ *is a probability-preserving transformation.*

[Notation: $\mathscr{B}_1 \times \mathscr{B}_2$ *is the sigma-algebra generated by the measurable rectangles* $A_1 \times A_2$ *with* $A_1 \in \mathscr{B}_1$ *and* $A_2 \in \mathscr{B}_2$*. The product measure* $\mathrm{m}_1 \times \mathrm{m}_2$ *is the unique measure on* $\mathscr{B}_1 \times \mathscr{B}_2$ *satisfying* $\mathrm{m}(A_1 \times A_2) = \mathrm{m}_1(A_1)\,\mathrm{m}_2(A_2)$*.]*

3. Existence of invariant measures

Earlier, we introduced the concept of an invariant measure, a measure m satisfying the identity $\mathrm{m} = T_*\mathrm{m}$. However, it is not clear at all whether such measures exist. In this section we obtain a partial answer to this question:

Theorem 2.17. *Let* X *be a compact metric space and* $T : X \to X$ *a continuous map. Then there exists at least one invariant Borel probabillity measure* m.

In general there may be many invariant measures, also for continuous maps on compact metric spaces, as we observed in Exercise 2.14.

We will actually give two alternative proofs of Theorem 2.17, because each is informative in its own right. Both proofs rely on the fact that, when T is a continuous transformation of a compact metric space, $\mathcal{P}$ is compact (as we will see), and the first one also uses the convexity of $\mathcal{P}$.

Before we can speak about compactness, the topology must be specified, of course. We endow $\mathcal{M}_s \supset \mathcal{P} \supset \mathcal{P}_T$, the set of signed measures, with the so-called **weak topology of measures**. It is described in Section 6 of Appendix B in some detail; we only present the two facts about it that are relevant to us.

Fact 1: A sequence $\mathrm{m}_n \in \mathcal{M}_s$, $n \geq 1$, converges to $\mathrm{m} \in \mathcal{M}_s$ — written $\mathrm{m}_n \Rightarrow \mathrm{m}$ — if and only if

$$\lim_{n\to\infty} \int_X f \,\mathrm{dm}_n = \int_X f \,\mathrm{dm}\ , \qquad f \in C(X)\ .$$

Fact 2: The weak topology is metrizable on the subset $\mathcal{P}$. In particular, a map $F : \mathcal{P} \to \mathcal{P}$ is continuous if and only if

$$\mathrm{m}_n \Rightarrow \mathrm{m} \quad \text{implies} \quad F(\mathrm{m}_n) \Rightarrow F(\mathrm{m})\ .$$

A related observation is that in the current setting it is natural to work with continuous functions $f \in C(X)$ instead of more general measurable functions $f : X \to \mathbb{C}$, and this usually amounts to little loss of generality. In particular, the following version of Lemma 2.15 will be needed soon:

Lemma 2.18. *Let X be a compact metric space and $T : X \to X$ a continuous map.*

(1) Given two measures $\mathrm{m}, \mathrm{m}' \in \mathcal{P}$, we have $\mathrm{m}' = T_*\mathrm{m}$ if and only if

$$\int_X f \,\mathrm{dm}' = \int_X f\circ T \,\mathrm{dm}\ , \qquad f \in C(X)\ .$$

(2) A measure $\mathrm{m} \in \mathcal{P}$ is invariant if and only if

$$\int_X f \,\mathrm{dm} = \int_X f\circ T \,\mathrm{dm}\ , \qquad f \in C(X)\ .$$

Exercise 2.19. *Prove Lemma 2.18.*

[Hint: Since the Borel sigma-algebra is generated by the open sets, m' is the pushforward of m if and only if $\mathrm{m}'(U) = \mathrm{m}(T^{-1}U)$ for all open sets $U \subset X$. The indicator function 1_U of any open set U can be approximated by a bounded sequence of continuous functions.]

Returning to the question whether an invariant probability measure exists, observe that the identity

$$\mathrm{m} = T_*\mathrm{m} \tag{2.5}$$

can be viewed as a **fixed-point problem** for the map $T_* : \mathcal{P} \to \mathcal{P}$.

3.1. The first proof: finding an invariant measure as a fixed point. We will prove that there indeed exists a fixed point, borrowing a classical idea from the topology of Euclidean spaces: recall that if $K \subset E$ is a compact convex subset of a Euclidean space E and $g : K \to K$ is a continuous function, then Brouwer's fixed point theorem guarantees the existence of a fixed point $g(x) = x \in K$. The Schauder–Tychonoff fixed point theorem (Theorem B.4) generalizes this result to infinite dimensions, which is required for our purposes. Before we can apply it, we need to establish the desired structure of $\mathcal{P}$:

Lemma 2.20. *Let X be a compact metric space. Then the set $\mathcal{P}$ of all Borel probability measures on X is convex and compact in the weak topology of measures.*

PROOF. We have already proved convexity. Since the weak topology is metrizable on $\mathcal{P}$, compactness is equivalent to every sequence $\mathrm{m}_n \in \mathcal{P}$, $n \geq 1$, having a subsequence m_{n_k}, $k \geq 1$, which converges weakly to a limit $\mathrm{m} \in \mathcal{P}$. But note that a sequence of probability measures is bounded ($\mathrm{m}_n(X) = 1$), so Lemma B.12 applies and there exists a subsequence $\mathrm{m}_{n_k} \Rightarrow \mathrm{m} \in \mathcal{M}$. This means that $\lim_{n\to\infty} \int_X f \, \mathrm{dm}_n = \int_X f \, \mathrm{dm}$ for all $f \in C(X)$. Setting $f = 1$, we see that m is a probability measure, so $\mathcal{P}$ is compact. □

We also need to establish the continuity of T_* on $\mathcal{P}$:

Lemma 2.21. *Let X be a compact metric space and $T : X \to X$ a continuous map. Then the map $T_* : \mathcal{P} \to \mathcal{P}$ is continuous in the weak topology of measures.*

PROOF. It suffices to show that $\mathrm{m}_n \Rightarrow \mathrm{m}$ implies $T_*\mathrm{m}_n \Rightarrow T_*\mathrm{m}$. The second condition means that $\int_X f \, \mathrm{d}(T_*\mathrm{m}_n) \to \int_X f \, \mathrm{d}(T_*\mathrm{m})$ for all $f \in C(X)$. By Lemma 2.18, this is equivalent to $\int_X f \circ T \, \mathrm{dm}_n \to \int_X f \circ T \, \mathrm{dm}$. Since $f \circ T$ is continuous, the latter is true by the assumption $\mathrm{m}_n \Rightarrow \mathrm{m}$. □

We need one last fact from Appendix B, which the reader is advised to take for granted:

Fact 3: The set $\mathcal{M}_s$ of all signed Borel measures is a (real) vector space, and equipped with the weak topology of measures it is a locally convex space.

PROOF OF THEOREM 2.17. We now know that $\mathcal{P}$ is a compact convex subset of the locally convex space $\mathcal{M}_s$ and that $T_* : \mathcal{P} \to \mathcal{P}$ is a continuous map. Hence, the Schauder–Tychonoff fixed point theorem yields the existence of $\mathrm{m} \in \mathcal{P}$ such that (2.5) holds, meaning that m is an invariant Borel probability measure. □

3.2. The second proof: finding an invariant measure as the limit of approximately invariant measures. Recall from analysis that fixed points $x = g(x)$ are often sought by iteration: one first picks a point x_0, constructs a sequence x_i, $i \geq 0$, with $x_i = g(x_{i-1}) = g^{\circ i}(x_0)$, and then tries to show that $\lim_{i\to\infty} x_i = x$ exists. If g is continuous, then the limit x is a fixed point.

We are now tempted to pick a measure $\mu \in \mathcal{P}$ and to show that $T_*^i\mu$ converges weakly to a fixed point m of T_* — an invariant measure. This generally fails. However, it is a useful observation to make that if $T_*^i\mu$ converges to m, then so does the averaged sequence

$$\mathrm{m}_n = \frac{1}{n} \sum_{i=0}^{n-1} T_*^i \mu \ , \quad n \geq 1 \ ,$$

and the limit is m. But what if the former does not converge? Since $\mathcal{P}$ is compact, both $T_*^i\mu$ and m_n do have converging *subsequences*, no matter which measure μ is. Our preference for the averaged sequence is revealed by the following exercises, which the reader is encouraged to solve:

Exercise 2.22. *Show that if there is a weakly converging subsequence m_{n_k}, $k \geq 1$, then its limit m is an invariant measure.*

[Hint: Compare m_{n_k} and $T_*\mathrm{m}_{n_k}$.]

A similar statement is generally *not* true for a converging *subsequence* of $T_*^i\mu$, $i \geq 0$:

Exercise 2.23. *Let X be the set of two points $\{0,1\}$ endowed with the discrete metric: $d(x,y)=0$ if $x=y$ and $d(x,y)=1$ otherwise. Let $T:X\to X$ be the map that exhanges the points: $T(0)=1$ and $T(1)=0$. Show that T has a unique invariant probability measure, and that any other probability measure μ satisfies*

$$\mu = T_*^2\mu = T_*^4\mu = \cdots \quad \text{and} \quad \mu \neq T_*\mu = T_*^3\mu = T_*^5\mu = \cdots .$$

Proof of Theorem 2.17. Let $\mu \in \mathcal{P}$ be arbitrary and define the averaged sequence m_n, $n \geq 1$, as above. Since m_n is a convex combination of probability measures, $\mathrm{m}_n \in \mathcal{P}$. Since $\mathcal{P}$ is compact (Lemma 2.20) in the weak topology of measures, there exists a weakly converging subsequence. By Exercise 2.22 the limit is an invariant measure. □

3.3. Uniqueness of an invariant measure. As we saw in Theorem 2.17, a continuous map on a compact metric space always has an invariant measure. On the other hand, Exercise 2.14 revealed that a given map may have (even uncountably) many of them. In this section we characterize those continuous maps that have precisely one invariant measure, following Parry's book [9]. Uniqueness turns out to be intimately related to the convergence of the time averages $\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$ *and to whether the limit f_+ in (2.1) is a constant function.* For brevity, we write $g \in \mathbb{C}$ if $g: X \to \mathbb{C}$ is a constant function.

Theorem 2.24 (Another ergodic theorem for continuous transformations)**.** *Let X be a compact metric space and $T:X\to X$ a continuous map. Then T has a unique invariant Borel probability measure m if and only if one of the following equivalent conditions is satisfied:*

(1) Given any $f \in C(X)$, the limit $f_+(x)$ exists for all $x \in X$ and $f_+ \in \mathbb{C}$.

(2) Given any $f \in C(X)$, $\frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i$ converges uniformly to $f_+ \in \mathbb{C}$.

(3) $C(X) = \mathbb{C} \oplus \bar{B}$.

If any one of these is satisfied, then the constant functions are the only invariant ones:

$$I = \mathbb{C} .$$

Moreover, the constant f_+ is determined by the unique invariant measure:

$$f_+ = \int_X f \, \mathrm{dm} . \tag{2.6}$$

Proof. (i) We first prove that (1)–(3) are equivalent and imply $I=\mathbb{C}$. By Theorem 2.4, pointwise convergence of the time averages to a continuous limit (such as a constant function) is equivalent to uniform convergence, so (1) ⇔ (2). Since $f_+=0$ for $f\in\bar{B}$ and $f_+=f$ for $f\in\mathbb{C}$, we obtain (3) ⇒ (1). By Theorem 2.4 we have (1) ⇒ $C(X) = I \oplus \bar{B}$. Introducing the linear operator $P: C(X) \to I : Pf = f_+$, we note that $I = \operatorname{im} P$. But (1) states that $\operatorname{im} P = \mathbb{C}$, so $I=\mathbb{C}$ and (1) ⇒ (3).

(ii) Suppose there are two different invariant measures m_1 and m_2. Then there exists a continuous function $f\in C(X)$ such that $\int_X f\,\mathrm{dm}_1 \neq \int_X f\,\mathrm{dm}_2$. Note that by invariance $\int_X f\,\mathrm{dm}_k = \int_X f\circ T^i\,\mathrm{dm}_k$ for each $i\geq 1$ and $k=1,2$. If (1) holds, then

$$\int_X f\,\mathrm{dm}_k = \int_X \frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i\,\mathrm{dm}_k \to \int_X f_+\,\mathrm{dm}_k = f_+$$

as $i \to \infty$ for each $k = 1, 2$, which is contradicts $\mathrm{m}_1 \neq \mathrm{m}_2$. Here we used the dominated (or bounded) convergence theorem. Hence (1) $\Rightarrow$ the invariant measure m is unique and (2.6) holds.

Finally, suppose (1) does not hold. Then there exists $f \in \mathbb{C}$ such that either (a) $f_+(x)$ fails to exists for some $x \in X$ or (b) f_+ exists on all of X but is not constant. In either case, we next construct two distinct invariant measures using a method similar to the one in the (second) proof of Theorem 2.17 in Section 3.2, which completes the proof.

Case (a): There exist two subsequences $\mathcal{N}_1, \mathcal{N}_2 \subset \mathbb{N}$ along which the limits disagree, that is, $\lim_{n\to\infty,\, n\in\mathcal{N}_k} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) = c_k \in \mathbb{C}$ $(k = 1, 2)$ where $c_1 \neq c_2$. Note first that

$$f(T^i(x)) = \int_X f \circ T^i \, \mathrm{d}\delta_x = \int_X f \, \mathrm{d}(T^i_*\delta_x)$$

by Lemma 2.18. Thus, the measures

$$\mu_n = \frac{1}{n}\sum_{i=0}^{n-1} T^i_*\delta_x \in \mathcal{P} \ , \quad n \geq 1 \ ,$$

satisfy

$$\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) = \int_X f \, \mathrm{d}\mu_n \ .$$

But $\mathcal{P}$ is compact in the weak topology of measures (Lemma 2.20), so both sequences $(\mu_n)_{n\in\mathcal{N}_k}$ must have weakly convergent subsequences: there exist $\mathcal{N}_k' \subset \mathcal{N}_k$ and $\mathrm{m}_k \in \mathcal{P}$ $(k = 1, 2)$ such that $\mu_n \Rightarrow \mathrm{m}_k$ as $n \to \infty$ with $n \in \mathcal{N}_k'$. We observed earlier (Exercise 2.22) that such limits m_k are always invariant measures. Since $c_1 \neq c_2$, we must have $\mathrm{m}_1 \neq \mathrm{m}_2$.

Case (b): There exist two points $x_1, x_2 \in X$ for which $f_+(x_1) \neq f_+(x_2)$. A construction similar to the one in case (a) again yields two distinct invariant measures. We leave the details to the reader. □

The property of a map having a unique invariant measure has been given a special name:

Definition 2.25. *Suppose $(X, \mathscr{B})$ is a measurable space and $T : X \to X$ is a measurable map. If T has a unique invariant probability measure, we say that T is* **uniquely ergodic***.*

In other words, Theorem 2.24 characterizes uniquely ergodic continuous maps on a compact metric space. We finish our discussion with an example of unique ergodicity:

Exercise 2.26 (Irrational rotations are uniquely ergodic, the angle doubling map is not)**.** *A sequence of numbers $(x_n)_{n\geq 0} \subset [0, 1]$ is called* **uniformly distributed** *if*

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} 1_I(x_i) = |I| \tag{2.7}$$

holds for any interval $I \subset [0, 1]$. Here $|I|$ is the length of the interval. It is not hard to see that (2.7) is equivalent to the property that

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f(x_i) = \int_0^1 f(x) \, \mathrm{d}x$$

holds for any continuous function $f : [0,1] \to \mathbb{C}$ satisfying $f(0) = f(1)$. Moreover, Hermann Weyl has proved that (2.7) is also equivalent to

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} e^{2\pi i k x_i} = 0 \ , \quad k \in \mathbb{Z} \setminus \{0\} \ .$$

Now, let T be a rotation of the circle as in Exercise 2.10.

(a) Show that T is uniquely ergodic if α is irrational.

(b) Show that T has uncountably many invariant probability measures if α is rational. [Hint: Show that there exists $q \geq 0$ such that any point $x \in X$ is q-periodic: $T^q(x) = x$. Then consider the pushforward measures $T_*^i \delta_x$, $i \geq 0$.]

Finally, let T be the angle doubling map in Exercise 2.11.

(c) Show that T is not uniquely ergodic.

4. Extreme points of the convex set $\mathcal{P}_T$ for continuous transformations

As we observed in Theorem 2.24, the property of having a unique invariant measure m has to do with the convergence of the time averages $\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$ to a constant $f_+ \in \mathbb{C}$, and that constant is given by the invariant measure, namely $f_+ = \int_X f \, \mathrm{dm}$; see (2.6). We will extend such observations to the more general case of having a whole family of invariant measures. For this reason, we will study the structure of the set $\mathcal{P}_T$ in this section and the next in some detail.

Let us summarize what we know about the set of invariant measures thus far:

Theorem 2.27. *Let X be a compact metric space and $T : X \to X$ a continuous map. Then the set $\mathcal{P}_T$ of all invariant Borel probability measures is nonempty, convex and compact in the weak topology of measures.*

Proof. Convexity of $\mathcal{P}_T$ was proved in Lemma 2.12 and $\mathcal{P}_T \neq \emptyset$ in Theorem 2.17. Since $\mathcal{P}_T \subset \mathcal{P}$ and $\mathcal{P}$ is compact (Lemma 2.20), it suffices to check that $\mathcal{P}_T$ is closed. To that end, suppose $\mathrm{m}_n \in \mathcal{P}_T$ converges weakly to $\mathrm{m} \in \mathcal{P}$. Since T_* is continuous (Lemma 2.21), we have $\mathrm{m}_n = T_*\mathrm{m}_n$ converges weakly to $T_*\mathrm{m}$. This implies $\mathrm{m} = T_*\mathrm{m}$, or $\mathrm{m} \in \mathcal{P}_T$. □

Since convex combinations of invariant measures are invariant, it is reasonable to suspect that there might exists a special family of invariant measures such that all the others can be obtained by convex combinations of them. Namely, recall that if K is a compact convex set of a Euclidean space and E is the set of its **extreme points**, then any $x \in K$ can be expressed as a convex combination $x = \lambda_1 e_1 + \cdots + \lambda_n e_n$, for some weights $\lambda_1, \ldots, \lambda_n \in (0,1)$ with $\lambda_1 + \cdots + \lambda_n = 1$ and extreme points $e_1, \ldots, e_n \in E$ [5]. In our infinite dimensional setting the situation is far from obvious, but owing to Theorem 2.27 we can appeal to the Choquet theorem (Theorem B.6) and obtain a similar representation of an

[5] A point $x \in K$ of a convex set K is called an extreme point of K if it is not an interior point of a chord connecting two distinct points of K. In other words, if $x = (1-t)u + tv$ for some $u, v \in K$ and $t \in (0,1)$, then $x = u = v$.

arbitrary invariant measure in $\mathcal{P}_T$ as a **barycenter**, a generalized convex combination, of the extreme points of $\mathcal{P}_T$:

Theorem 2.28. *Let X be a compact metric space and $T : X \to X$ a continuous map. Denote by $\mathcal{E}_T \subset \mathcal{P}_T$ the set of extreme points of the compact convex set $\mathcal{P}_T$. Given any $\mathrm{m} \in \mathcal{P}_T$, there exists a Borel probability measure λ supported on $\mathcal{E}_T$ such that*

$$\mathrm{m} = \int_{\mathcal{E}_T} \mu \, \mathrm{d}\lambda(\mu) \ . \tag{2.8}$$

In particular, $\mathcal{E}_T$ is nonempty.

Before giving the proof, let us make some remarks and observations. The representation (2.8) means that

$$\mathrm{m}(A) = \int_{\mathcal{E}_T} \mu(A) \, \mathrm{d}\lambda(\mu) \ , \quad A \in \mathscr{B} \ ,$$

or, equivalently, that

$$\int_X f \, \mathrm{dm} = \int_{\mathcal{E}_T} \left(\int_X f \, \mathrm{d}\mu \right) \mathrm{d}\lambda(\mu) \ , \quad f \in C(X) \ .$$

Thus, given any invariant measure m, the measure $\mathrm{m}(A)$ of a Borel set is determined by taking a suitable average of the extreme measures $\mu(A)$, where the weights are determined by λ. Since the weights are generally not unique even in finite dimensional convex combinations, the proof given below leaves the question of uniqueness of the measure λ open.[6] Note that λ is a measure on the set of measures $\mathcal{E}_T$, which is perhaps a bit abstract, but there is no problem at all from the measure-theoretic point of view because $\mathcal{P}_T$ has a topology — the weak topology of measures — and $\mathcal{E}_T$ is in fact a Borel subset with respect to it. By definition, the maps $\mu \mapsto \int_X f \, \mathrm{d}\mu$, $f \in C(X)$, are continuous with respect to the said topology. That λ is supported on $\mathcal{E}_T$ means that λ is a measure on $\mathcal{P}_T$ and $\lambda(\mathcal{P}_T \setminus \mathcal{E}_T) = 0$.

PROOF OF THEOREM 2.28. Recall that the set $\mathcal{M}_s$ of all signed Borel measures equipped with the weak topology of measures is a locally convex space and that this topology is metrizable on $\mathcal{P}$. Since $\mathcal{P}_T \subset \mathcal{P}$ is compact and convex, we may appeal to the Choquet Theorem (Theorem B.6), which implies the representation part. On the other hand, since $\mathcal{P}_T$ is nonempty, it now follows that $\mathcal{E}_T$ is nonempty. □

We now illustrate Theorem 2.28 with the simple example encountered in Exercise 2.14.

Exercise 2.29. *Let T be the identity map on $X = [0, 1]$ and recall that every measure is invariant, $\mathcal{P}_T = \mathcal{P}$. Show that the extreme points are precisely the point masses, i.e., $\mathcal{E}_T = \{\delta_x : x \in X\}$. Thus, Theorem 2.28 implies that any $\mathrm{m} \in \mathcal{P}$ is the barycenter of point masses:*

$$\mathrm{m} = \int_{\{\delta_x \, : \, x \in X\}} \mu \, \mathrm{d}\lambda(\mu)$$

for some Borel probability measure supported on $\{\delta_x : x \in X\}$. This is not very surprising: check directly that

$$\mathrm{m} = \int_X \delta_x \, \mathrm{dm}(x) \ ,$$

[6] In our ergodic-theoretic setting λ is essentially unique; the set $\mathcal{P}_T$ can be thought of as an infinite dimensional simplex, not just a convex set.

which means that $\mathrm{m}(A) = \int_X \delta_x(A)\,\mathrm{dm}(x)$ *for all* $A \in \mathscr{B}$. *Note the distinction between the two representations, however, that the measure* λ *is defined on the set of measures* $\mathcal{P}$.

As a slightly less trivial example, we will soon see how the invariant measures of rational rotations of the circle can be decomposed into generalized convex combinations of its extreme measures.

5. Ergodicity

We continue our study of the convex set $\mathcal{P}_T$ and the set $\mathcal{E}_T$ of its extreme points with the eventual goal of characterizing the convergence of the time averages $\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$, in the general setting of measurable transformations. This calls for the introduction of the concepts of invariant sets and ergodic measures and their basic properties.

Definition 2.30. *Let* $(X, \mathscr{B})$ *be a measurable space and* $T : X \to X$ *a measurable map.*

(1) A set $A \in \mathscr{B}$ *is* **(strictly) invariant** *if* $T^{-1}A = A$.
(2) Given $\mathrm{m} \in \mathcal{P}_T$, *a set* $A \in \mathscr{B}$ *is* **almost invariant** *if* $\mathrm{m}(T^{-1}A \,\Delta\, A) = 0$.[7]

The invariance condition $T^{-1}A = A$ has a dynamical characterization:

$x \in A \Rightarrow T(x) \in A$, i.e., the trajectory of a point in A cannot escape from A

and

$T(x) \in A \Rightarrow x \in A$, i.e., the trajectory of a point in A^c cannot enter A.

Almost invariance of a set with respect to an invariant measure means that similar properties hold, except on a set of measure zero. A measurable set need not be invariant even if it is made of trajectories:

Example 2.31. *Let* $X = \{0, 1\}$, $\mathscr{B}$ *the sigma-algebra of all subsets of* X, *and* $T \equiv 0$. *The set* $\{0\}$ *consists of a trajectory, but it is not invariant since* $T^{-1}\{0\} = X$.

It is useful to make the following general observation about invariant measures:

Exercise 2.32. *If* $\mathrm{m} \in \mathcal{P}_T$, *then*

$$\mathrm{m}(A \setminus T^{-1}A) = \mathrm{m}(T^{-1}A \setminus A)\ , \quad A \in \mathscr{B}\ . \tag{2.9}$$

Hence, a set $A \in \mathscr{B}$ is almost invariant if and only if either $\mathrm{m}(A \setminus T^{-1}A) = 0$ or $\mathrm{m}(T^{-1}A \setminus A) = 0$.

Exercise 2.33 (A dynamical characterization of almost invariant sets)**.** *Prove that* $A \in \mathscr{B}$ *is almost invariant if and only if the trajectory of almost every* $x \in A$ *is contained in* A.

Note that the almost invariance of a set depends on the measure, while the invariance of a set is a strictly set-theoretic concept. This distinction is often not very significant for practical purposes. Of course, every invariant set is almost invariant with respect to any measure, but there turns out to be a partial converse:

Lemma 2.34. *Suppose* $(X, \mathscr{B}, \mathrm{m}, T)$ *is a probability-preserving transformation and that* $A \in \mathscr{B}$ *is an almost invariant set. There exists a strictly invariant set* $A_0 = T^{-1}A_0 \in \mathscr{B}$ *such that* $\mathrm{m}(A_0 \,\Delta\, A) = 0$.

[7]Here $A \,\Delta\, B = (A \setminus B) \cup (B \setminus A)$ denotes the symmetric difference of the sets $A \subset X$ and $B \subset X$.

Exercise 2.35. *Prove Lemma 2.34.*

[Hint: Define A_0 as the set of all those points whose trajectories visit A infinitely often. That is, $A_0 = \{x \in X : T^i(x) \in A$ for infinitely many $i \geq 1\} = \cap_{n=1}^{\infty} \cup_{i=n}^{\infty} T^{-i}A$. Show that, for all $i \geq 1$, $\mathrm{m}(T^{-i}A \,\Delta\, A) = 0$ holds, and that this implies $\mathrm{m}(A_0 \,\Delta\, A) = 0$.]

Invariant probability measures m and (almost) invariant sets A are closely related. The next exercise shows that if $0 < \mathrm{m}(A) < 1$, m can be decomposed into the convex combination

$$\mathrm{m} = \mathrm{m}(A)\mathrm{m}(\,\cdot \mid A) + (1 - \mathrm{m}(A))\mathrm{m}(\,\cdot \mid A^c)$$

of its conditional parts on A and A^c. In particular, m cannot be an extreme point of $\mathcal{P}_T$.

Exercise 2.36. *Suppose $(X, \mathscr{B}, \mathrm{m}, T)$ is a probability-preserving transformation and that $A \in \mathscr{B}$ is an almost invariant set with $0 < \mathrm{m}(A) < 1$. Observe that A^c has the same properties, and show that the conditional measures $\mathrm{m}(B \mid A) = \frac{1}{\mathrm{m}(A)}\mathrm{m}(B \cap A)$ and $\mathrm{m}(B \mid A^c) = \frac{1}{\mathrm{m}(A^c)}\mathrm{m}(B \cap A^c)$, $B \in \mathscr{B}$, are both invariant.*

If A is moreover (strictly) invariant, show that the restrictions $(A, \mathscr{B} \cap A, \mathrm{m}(\,\cdot \mid A), T|_A)$ and $(A^c, \mathscr{B} \cap A^c, \mathrm{m}(\,\cdot \mid A^c), T|_{A^c})$ define probability-preserving transformations.

(Notation: Given $E \in \mathscr{B}$, $\mathscr{B} \cap E = \{B \cap E : B \in \mathscr{B}\}$ and $T|_E : E \to X : T|_E(x) = T(x)$.)

Interpretation: T separates into the two components $T|_A$ and $T|_{A^c}$ which "do not interact".

If, say, A contained a smaller invariant set A_1, one could further decompose m_A into its conditional parts on A_1 and A_1^c, and so on, which would lead to convex combinations $\mathrm{m} = \sum_{i=1}^{n} \mathrm{m}(E_i)\mathrm{m}(\,\cdot \mid E_i)$ of invariant measures on smaller and smaller invariant sets E_i with $\cup_{i=1}^{n} E_i = A \cup A^c = X$. Let us be heuristic for the moment. It benefits intuition to think that extreme measures (which by definition cannot be decomposed as above) "live on minimal invariant sets" which cannot be split into smaller invariant sets. Then, repeating the above procedure of conditioning the measure as many times as possible, one should arrive at a representation of m as a (possibly infinite) convex combination of extreme measures, like in (2.8), without the continuity and compactness assumptions of Theorem 2.28. We will return to this topic in Section 6.

We already defined invariant functions in the context of a continuous map $T : X \to X$ on a compact metric space X as continuous functions satisfying the relation $f = f \circ T$. In the general measure-theoretic setting this definition is not sufficient, even if X is a topological space.

Definition 2.37. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map.*

(1) A measurable function $f : X \to \mathbb{C}$ is **(strictly) invariant** *if $f = f \circ T$ on X.*
(2) Given an invariant measure m, a measurable function $f : X \to \mathbb{C}$ is **almost invariant** *if $f = f \circ T$ almost everywhere with respect to m.*

There is no loss of generality in assuming that an invariant function is real valued.

Invariant sets and invariant functions are closely related:

Lemma 2.38. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map.*

(1) $A \in \mathscr{B}$ is an invariant set if and only if 1_A is an invariant function.
(2) $A \in \mathscr{B}$ is an almost invariant set if and only if 1_A is an almost invariant function.

(3) *A function $f : X \to \mathbb{R}$ is invariant if and only if its level sets $\{f > t\} = \{x \in X : f(x) > t\}$, $t \in \mathbb{R}$, are invariant.*
(4) *A function is almost invariant if and only if its level sets are almost invariant.*

Equivalently, one can use the level sets $\{f \geq t\}$, $\{f < t\}$ or $\{f \leq t\}$ in (3) and (4).

Exercise 2.39. *Prove Lemma 2.38.*

Of course, an invariant function is almost invariant with respect to any invariant measure. But we also have the following counterpart of Lemma 2.34:

Lemma 2.40. *If f is an almost invariant function with respect to* m*, then there exists a strictly invariant function f_0 such that $f = f_0$ almost everywhere with respect to* m*.*

Proof. We follow the book of Brin and Stuck [**2**]. Given f, let us define the measurable functions $\phi_i : X \to \mathbb{C} : x \mapsto f(T^i(x)) - f(x)$, $i \geq 1$. Then each $A_i = \phi_i^{-1}\{0\}$ is a measurable set, and $\mathrm{m}(A_i) = 1$ because f is almost invariant. Note that the set

$$A = \bigcap_{i \geq 1} A_i = \{x \in X : f(T^i(x)) \text{ is constant for all } i \geq 0\}$$

is measurable and $\mathrm{m}(A) = 1$. The intuition is that A is a good set, on which the invariance property of the function f holds everywhere, even under repeated iterations of T. Using A, we construct a modification f_0 of f.

The first candidate for f_0 that springs to mind is $g = 1_A f$, which is measurable and agrees with f on the set A of full measure. However, this candidate is not necessarily invariant, because the set A is not necessarily invariant: If $x \in A$ then $T(x) \in A$, meaning $A \subset T^{-1}A$, but there is no reason for the opposite inclusion to hold. To see that g is not invariant, suppose that $x \in A^c$ and $T^n(x) \in A$ for some $n \geq 1$. Then $g(x) = 0$ and $g(T^n(x)) = f(T^n(x))$, but there are no guarantees whatsoever that $f(T^n(x)) = 0$, so g fails to be invariant.

To overcome the issue above, we consider the set B consisting of all those points whose trajectories visit A at least once (and therefore never leaves A again). It can be expressed as

$$B = \bigcup_{n \geq 0} T^{-n}A = \{x \in X : \exists\, n \geq 0 \text{ such that } f(T^i(x)) \text{ is constant } \underline{\text{for all } i \geq n}\} .$$

This is a measurable set, and $\mathrm{m}(B) = 1$ because $A \subset B$. The set is also invariant, but f is not invariant on $B \setminus A$, so also the candidate $1_B f$ for f_0 fails. Fortunately, a minor correction works: If $x \in B$, then there exists a minimal integer $n(x) \geq 0$ such that $f(T^i(x)) = f(T^{n(x)}(x))$ for all $i \geq n(x)$, so we define

$$f_0(x) = \begin{cases} f(T^{n(x)}(x)) & \text{if } x \in B, \\ 0 & \text{if } x \in B^c. \end{cases}$$

This functions agrees with f on A, because $n(x) = 0$ for $x \in A$. We also have $n(x) = k$ precisely for $x \in T^{-k}A \setminus \cup_{n=0}^{k-1} T^{-n}A$, so f_0 is measurable. As the reader may check, f_0 is invariant. □

The following concept has a key role in ergodic theory, as we will soon learn.

Definition 2.41. *An invariant measure* $\mathrm{m} \in \mathcal{P}_T$ *is* **ergodic***, if each invariant set* $A = T^{-1}A \in \mathcal{B}$ *satisfies* $\mathrm{m}(A) \in \{0,1\}$.

In words, a measure is ergodic if all invariant sets are trivial.

Exercise 2.42. *Let* $T : \mathbb{S}^1 \to \mathbb{S}^1 : T(x) = x + \alpha \pmod 1$ *be a rotation of the circle with* *rational* $\alpha \in (0,1)$. *Then every* $x \in X$ *is periodic with the same period* q*; see Exercise 2.26. (1) Show that the measures* $\mathrm{m}_x = \frac{1}{q}(\delta_x + \delta_{T(x)} + \cdots + \delta_{T^{q-1}(x)})$ *are ergodic. (2) Show that these are in fact the only ergodic probability measures.*

[Hint: In (2), write α *in irreducible form to determine the minimal* q*. Then show that an ergodic probability measure must assign measure* 1 *to a set* $I \cup T^i I \cup \cdots \cup T^{q-1} I$*, where* $I \subset \mathbb{S}^1$ *is a closed interval which can be made arbitrarily short.]*

Since the ongoing discussion is somewhat lengthy and the topic may seem to have drifted away from the convergence of time averages, some reassuring remarks are in order. First, recall Definition 2.25, the definition of unique ergodicity. The terminology is explained by the next exercise:

Exercise 2.43. *Suppose* $(X, \mathscr{B})$ *is a measurable space and that* $T : X \to X$ *is uniquely ergodic. Prove that the (unique) invariant measure* $\mathrm{m} \in \mathcal{P}_T$ *is ergodic.*

Next, recall that if T is a uniquely ergodic continuous map on a compact metric space X, then the time averages $\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) = \int_X f \,\mathrm{dm}$ for all $x \in X$ and all $f \in C(X)$, where m is the unique (thus ergodic) invariant measure; see Theorem 2.24. In the general case of a measurable map $T : X \to X$ on a measurable space $(X, \mathscr{B})$, we will see that

$$\text{m is ergodic} \iff \lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) = \int_X f \,\mathrm{dm}\ , \ \text{m-a.e. } x \text{ and all } f \in L^1(X, \mathscr{B}, \mathrm{m})\ . \tag{2.10}$$

(This result is known as the Birkhoff ergodic theorem.) Therefore, the study of ergodic measures is at the heart of understanding limits of time averages, and a little patience is warranted.

There are several equivalent characterizations of ergodicity, including the following:

Theorem 2.44. *A measure* $\mathrm{m} \in \mathcal{P}_T$ *is ergodic if and only if one of the following conditions is satisfied:*

(1) *If* $A \in \mathscr{B}$ *is invariant, then* $\mathrm{m}(A) \in \{0,1\}$.
(2) *If* $A \in \mathscr{B}$ *is almost invariant, then* $\mathrm{m}(A) \in \{0,1\}$.
(3) *If* $f : X \to \mathbb{R}$ *is invariant, then* f *is constant* *almost* *everywhere.*
(4) *If* $f : X \to \mathbb{R}$ *is almost invariant, then* f *is constant almost everywhere.*
(5) *For any* $A \in \mathscr{B}$ *with* $\mathrm{m}(A) > 0$*, we have* $\mathrm{m}(\cup_{i=1}^{\infty} T^{-i}A) = 1$.
(6) *For any* $A, B \in \mathscr{B}$ *with* $\mathrm{m}(A) > 0$ *and* $\mathrm{m}(B) > 0$*, there exists* $i \geq 1$ *such that* $\mathrm{m}(T^{-i}A \cap B) > 0$.

Before proving Theorem 2.44, let us make some remarks.

Equivalently, the field $\mathbb{R}$ can be replaced with $\mathbb{C}$. Moreover, the class of measurable functions $f : X \to \mathbb{R}$ can be replaced by $L^p(X, \mathscr{B}, \mathrm{m})$ for any $p \in [1, \infty]$, in particular by bounded measurable functions.

Warning concerning (3): Earlier, we studied the ergodic theory of continuous maps and called a *continuous* function $f \in C(X)$ invariant if $f \circ T = f$ on X. For that discussion it was sufficient to restrict to $f \in C(X)$. However, even if $T : X \to X$ is a continuous map on a compact metric space, ergodicity *does not* follow even if every *continuous* invariant function turns out to be constant.

Condition (5) has an important dynamical interpretation: If A is *any* set of positive measure, then the trajectory of *almost every point in* X *will eventually visit the set* A. In fact, it does so infinitely many times!

Exercise 2.45. *Show that if* m *is ergodic and* $A \in \mathscr{B}$ *with* $\mathrm{m}(A) > 0$*, then the trajectory of almost every* $x \in X$ *visits* A *infinitely often.*

[Hint: Consider the sets $\cup_{i=n}^{\infty} T^{-i}A$*,* $n \geq 1$*.]*

Thus, we may informally say that m is ergodic, if the trajectory of almost every point in X will explore every region of positive measure in X — no matter how small. The following exercise further illuminates this point.

Exercise 2.46. *Suppose* X *is a compact metric space and that* $T : X \to X$ *is continuous and uniquely ergodic. Assume also that the invariant measure* m *satisfies* $\mathrm{m}(U) > 0$ *for every nonempty open set* $U \subset X$*. Show that the trajectory of* every *point* $x \in X$ *is dense.*

[Hint: Use Theorem 2.24.]

Given the preceding discussion, it is not entirely unreasonable that in the ergodic case the time averages $\sum_{i=0}^{n-1} f(T^i(x))$ along the trajectory of a typical point will converge to the space average $\int_X f \, \mathrm{dm}$, as claimed (still without proof) in (2.10).

Proof of Theorem 2.44. (1) is the defining property of ergodicity, so it remains to prove that conditions (1)–(6) are equivalent. Of course, (2) $\Rightarrow$ (1) and (4) $\Rightarrow$ (3).

(1) $\Rightarrow$ (2): If (1) holds and $A \in \mathscr{B}$ is almost invariant, then Lemma 2.34 guarantees that there exists an invariant set $A_0 \in \mathscr{B}$ with $\mathrm{m}(A \,\Delta\, A_0) = 0$. Since $\mathrm{m}(A_0) = 0$, also $\mathrm{m}(A) = 0$.

(3) $\Rightarrow$ (4): This is a similar application of Lemma 2.34 as above.

(4) $\Rightarrow$ (2): If $A \in \mathscr{B}$ is almost invariant, then 1_A is almost invariant by Lemma 2.38 and by (4) constant almost everywhere. Hence $\mathrm{m}(A) = \int_X 1_A \, \mathrm{m} \in \{0, 1\}$.

(2) $\Rightarrow$ (3): Let f be invariant. Every level set $\{f > t\}$ is invariant by Lemma 2.38. Hence, $\mathrm{m}(\{f > t\}) \in \{0, 1\}$ for all $t \in \mathbb{R}$. This implies (exercise!) that $\mathrm{m}(\{f = c\}) = 1$ for some $c \in \mathbb{R}$.

Having proved that (1)–(4) are equivalent, we will now complete the proof by showing that (2) $\Rightarrow$ (5) $\Rightarrow$ (6) $\Rightarrow$ (1).

(2) $\Rightarrow$ (5): Let $A \in \mathscr{B}$ and $\mathrm{m}(A) > 0$, and denote $B = \cup_{i=1}^{\infty} T^{-i}A$. This set is almost invariant: $T^{-1}B \subset B$ and $\mathrm{m}(T^{-1}B) = \mathrm{m}(B)$ imply $\mathrm{m}(T^{-1}B \,\Delta\, B) = 0$. Thus, $\mathrm{m}(B) \in \{0, 1\}$. But $\mathrm{m}(B) \geq \mathrm{m}(T^{-1}A) = \mathrm{m}(A) > 0$, so $\mathrm{m}(B) = 1$.

(5) $\Rightarrow$ (6): Let $A, B \in \mathscr{B}$ be sets of positive measure. By assumption $\mathrm{m}(\cup_{i=1}^{\infty} T^{-i}A) = 1$, so $\mathrm{m}(\cup_{i=1}^{\infty} B \cap T^{-i}A) = \mathrm{m}(B \cap \cup_{i=1}^{\infty} T^{-i}A) = \mathrm{m}(B) > 0$. Thus, $\mathrm{m}(B \cap T^{-i}A) > 0$ at least for one $i \geq 1$.

(6) $\Rightarrow$ (1): Assume there exists an invariant set $A \in \mathscr{B}$ with $0 < \mathrm{m}(A) < 1$. Then $0 = \mathrm{m}(A \cap A^c) = \mathrm{m}(T^{-i}A \cap A^c)$ for all $i \geq 1$, which contradicts (6). $\square$

Exercise 2.47. *Recall the angle doubling map $T : x \mapsto 2x \pmod 1$ in Exercise 2.11. Show that the Lebesgue measure is ergodic.*

[Hint: Any L^2 function f of period 1 has a unique representation as the Fourier series $f(x) = \sum_{k\in\mathbb{Z}} \hat{f}(k)e^{2\pi ikx}$, where $\hat{f}(k) = \int_0^1 e^{-2\pi ikx} f(x)\,\mathrm{d}x$ satisfies $\lim_{|k|\to\infty} \hat{f}(k) = 0$.]

The next theorem reveals the connection of ergodicity to the structure of the set $\mathcal{P}_T$ of all invariant measures.

Theorem 2.48. *The set of ergodic measures is exactly the set $\mathcal{E}_T$ of extreme points of $\mathcal{P}_T$.*

For the proof of Theorem 2.48 we will need the next lemma. Recall that a measure μ is absolutely continuous with respect to another measure m, if $\mathrm{m}(A) = 0$ implies $\mu(A) = 0$. In that case μ has a density $f = \frac{d\mu}{\mathrm{dm}} \in L^1(X, \mathscr{B}, \mathrm{m})$ with respect to m, meaning that

$$\mu(A) = \int_A \mathrm{d}\mu = \int_A f\,\mathrm{dm}\;, \quad A \in \mathscr{B}\;.$$

This is the Radon–Nikodym Theorem (Theorem A.6).

Lemma 2.49. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. If μ and m are invariant probability measures and μ is absolutely continuous with respect to m, then the Radon–Nikodym density $f = \frac{d\mu}{\mathrm{dm}}$ is almost invariant ($f = f \circ T$ a.e.) with respect to m.*

Proof. We follow Michael Hochman's lecture notes [7]. By Lemma 2.38 it suffices to show that any level set $A = \{f > t\}$, $t \in \mathbb{R}$, is almost invariant. First, we recall (2.9) implies

$$\mathrm{m}(A \setminus T^{-1}A) = \mathrm{m}(T^{-1}A \setminus A)\;,$$

because m is invariant. But also μ is invariant, so

$$\int_{A\setminus T^{-1}A} f\,\mathrm{dm} = \mu(A \setminus T^{-1}A) = \mu(T^{-1}A \setminus A) = \int_{T^{-1}A\setminus A} f\,\mathrm{dm}\;.$$

On $A \setminus T^{-1}A$, $f > t$, while on $T^{-1}A \setminus A$, $f \leq t$. If $\mathrm{m}(A \setminus T^{-1}A) = \mathrm{m}(T^{-1}A \setminus A) > 0$, we run into the contradiction $t < t$. Hence $\mathrm{m}(A \setminus T^{-1}A) = \mathrm{m}(T^{-1}A \setminus A) = 0$, and A is almost invariant. $\square$

Proof of Theorem 2.48. If m is an extreme point of $\mathcal{P}_T$, then it is ergodic by Exercise 2.36. In the other direction, let m be ergodic. Suppose there exist two measures $\mu, \nu \in \mathcal{P}_T$ and a number $t \in (0,1)$ such that $\mathrm{m} = (1-t)\mu + t\nu$. Then μ (and ν) is absolutely continuous with respect to m, because $\mu(A) = 0$ if $\mathrm{m}(A) = 0$. It follows that it has a Radon-Nikodym density $f = \frac{d\mu}{dm} \in L^1(X, \mathscr{B}, \mathrm{m})$, which by Lemma 2.49 is almost invariant with respect to m. Since m is ergodic, f is constant almost everywhere by Theorem 2.44. Since $f \geq 0$ and $\int_X f\,\mathrm{dm} = \nu(X) = 1$, this constant must be 1. But this means that $\mu = \mathrm{m}$, and further that m is an extreme point of $\mathcal{P}_T$. $\square$

The next corollary further establishes the role of ergodic measures as the building blocks of other invariant measures. Informally, it states that two distinct ergodic measures "live on disjoint subsets".

Corollary 2.50. *Let* $\mathrm{m}_1, \mathrm{m}_2 \in \mathcal{E}_T$ *be two ergodic probability measures. Then they either coincide or are mutually singular. In other words, either* $\mathrm{m}_1 = \mathrm{m}_2$ *or there exists* $A \in \mathscr{B}$ *such that* $\mathrm{m}_1(A) = 1$ *and* $\mathrm{m}_2(A) = 0$.

Exercise 2.51. *Prove Corollary 2.50 with the aid of Theorem 2.48.*

[Hint: Consider the measure $\mathrm{m} = \frac{1}{2}\mathrm{m}_1 + \frac{1}{2}\mathrm{m}_2$.*]*

6. Ergodic decomposition: a (very) brief overview

In the case of continuous transformations on a compact metric space, we were able to conclude that any invariant measure $\mathrm{m} \in \mathcal{P}_T$ can be expressed as a barycenter

$$\mathrm{m} = \int_{\mathcal{E}_T} \mu \, \mathrm{d}\lambda(\mu) \ . \tag{2.11}$$

of the extreme points of $\mathcal{P}_T$, i.e., the ergodic measures; see Theorem 2.28. In particular, since $\mathcal{P}_T$ is nonempty, it was concluded that $\mathcal{E}_T$ is nonempty. The reasoning was based on the observations that $T_* : \mathcal{P} \to \mathcal{P}$ is a continuous map, and that $\mathcal{P}_T$ is a convex and compact subset of $\mathcal{P}$ in the weak topology of Borel measures.

For a general measurable transformation on a measurable space we also know that $\mathcal{P}_T$ is convex — but perhaps empty — and its extreme points are precisely the ergodic measures. As discussed in Section 5, it is tempting to think that a representation of an invariant measure similar to (2.11) exists. However, we can no longer argue that T_* is continuous, nor that $\mathcal{P}_T$ is compact. In order to deal with this, we have to impose a mild assumption, namely that the measurable space $(X, \mathscr{B})$ be a standard Borel space:

Definition 2.52. *A topological space* X *is called a* **Polish space***, if it is separable and metrizable with a complete metric. The measurable space* $(X, \mathscr{B})$*, where* $\mathscr{B}$ *is the Borel sigma-algebra, is then called a* **standard Borel space***. If, moreover,* m *is a probability measure on* $\mathscr{B}$*, the probability space* $(X, \mathscr{B}, \mathrm{m})$ *is called a* **standard probability space***.*

Many ergodic theory books choose to work with Lebesgue spaces instead. These are isomorphic to completions of standard probability spaces, so there is not a big difference between the two. The nuisance with a Lebesgue space is that the completion of a sigma-algebra depends on the measure, while the Borel sigma-algebra does not. Only pathological probability spaces fail to be standard (or completions thereof). Nevertheless, every standard probability space with no atoms[8] is isomorphic to the unit interval equipped with the Borel sigma-algebra and the Lebesgue measure!

The result sought for is the one below. We will not provide the proof, which can be found, for instance, in the lecture notes [7, 15].

Theorem 2.53 (Ergodic decomposition on a standard Borel space)**.** *Let* $(X, \mathscr{B})$ *be a standard Borel space and* $T : X \to X$ *a measurable transformation. Given any* $\mathrm{m} \in \mathcal{P}_T$*, there exists a family* $\{\mathrm{m}_x \in \mathcal{E}_T : x \in X\}$ *of ergodic Borel probability measures such that the map* $x \mapsto \mathrm{m}_x(A)$ *is Borel measurable for all* $A \in \mathscr{B}$ *and*

$$\mathrm{m} = \int_X \mathrm{m}_x \, \mathrm{dm}(x) \ .$$

[8] A set $A \in \mathscr{B}$ with $\mathrm{m}(A) > 0$ is called an atom if there exists no measurable subset $B \subset A$ such that $0 < \mathrm{m}(B) < \mathrm{m}(A)$. In a standard probability space, an atom consists of a single point of positive measure and a set of measure zero. In any probability space, there can be at most countably many atoms.

The displayed formula, called the **ergodic decomposition** of m, is just shorthand notation for $\mathrm{m}(A) = \int_X \mathrm{m}_x(A)\,\mathrm{dm}(x)$, $A \in \mathscr{A}$. Equivalently, $\int_X f\,\mathrm{dm} = \int_X \left(\int_X f\,\mathrm{dm}_x\right)\mathrm{dm}(x)$ for all $f \in L^1(X, \mathscr{B}, \mathrm{m})$.

The intuition regarding the ergodic decompositions is that the set X admits a partition ε into "minimal invariant sets" — the ergodic components — which is in general uncountable. If ε_x denotes the partition element containing $x \in X$, then m_x is the conditional measure of m on ε_x. The latter is a Borel probability measure on X which is supported on ε_x, $\mathrm{m}_x(\varepsilon_x) = 1$, although in general $\mathrm{m}(\varepsilon_x) = 0$. The existence of such a system of conditional measures calls for the assumption that the probability space be standard.

Example 2.54. *In Exercise 2.29 the trivial example of the identity map on the interval* $[0,1]$ *was considered. It was determined that an arbitrary measure* m*, which is of course invariant, has the representation* $\mathrm{m} = \int_{[0,1]} \delta_x\,\mathrm{dm}(x)$*, where the point masses* δ_x *are precisely the ergodic measures. This is an example of an ergodic decomposition with* $\mathrm{m}_x = \delta_x$*; the ergodic components are just the singletons* $\{x\}$*,* $x \in [0,1]$*.*

From Exercise 2.42 we get a similar ergodic decomposition for rational rotations of the circle, in which case m_x *is the measure* $\frac{1}{q}(\delta_x + \delta_{T(x)} + \cdots + \delta_{T^{q-1}}(x))$*; the ergodic components are the periodic trajectories* $\{x, T(x), \ldots, T^{q-1}(x)\}$*,* $x \in \mathbb{S}^1$*.*

CHAPTER 3

Recurrence and ergodicity

In this chapter we continue the study of ergodicity and time averges, with the final goal of proving (2.10).

Recall that ergodicity can be stated as follows, in terms of visits to sets: Given any set $A \in \mathscr{B}$ of positive measure, $\mathrm{m}(A) > 0$, the trajectory of almost every point $x \in X$ will eventually visit A (infinitely often). If m is not ergodic, then all bets are off. In particular, there exists an invariant set $B \in \mathscr{B}$ with $0 < \mathrm{m}(B) < 1$, so the the trajectory of any $x \in B^c$ will never visit B: $T^i(x) \in B^c$ for all $i \geq 0$. One may nevertheless wonder whether the trajectory of a point x in a given set A will eventually *return to the same set*. Such points are called **recurrent with respect to the set** A. The question of recurrence is addressed next.

1. Recurrence theorems of Poincaré and Kac

In this section we study the typicality of recurrent points: If $(X, \mathscr{B}, \mathrm{m}, T)$ is a probability-preserving transformation and $A \in \mathscr{B}$ is a given set, how typical is it that the trajectory of a randomly chosen point $x \in A$ will eventually *return* to A, meaning that $T^i(x) \in A$ for some $i \geq 1$? The answer is astounding: it happens almost surely, infinitely many times, *without any additional assumptions!*

Theorem 3.1 (Poincaré recurrence theorem)**.** *Suppose* $(X, \mathscr{B}, \mathrm{m}, T)$ *is a ppt. Given* $A \in \mathscr{B}$*, let* $B = \{x \in A : T^i(x) \in A$ *for infinitely many* $i \geq 1\}$*. Then* $B \in \mathscr{B}$ *and*

$$\mathrm{m}(B) = \mathrm{m}(A) \ .$$

The reader is invited to compare the result with Exercise 2.45, and to note that ergodicity is *not* required in the Poincaré recurrence theorem.

Proof of Theorem 3.1. First note that we can write $B = A \cap \cap_{n=0}^{\infty} \cup_{i=n}^{\infty} T^{-i}A$, so B is a measurable set. Define $A_n = \cup_{i=n}^{\infty} T^{-i}A \in \mathscr{B}$, $n \geq 0$. Then $A_0 \supset A_1 \supset \cdots$ and

$$A = A \cap A_0 \supset A \cap A_1 \supset \cdots \supset \cap_{n=0}^{\infty} A \cap A_n = B \ .$$

In particular, we obtain $\mathrm{m}(A \cap A_n) \downarrow \mathrm{m}(B)$ as $n \to \infty$. To complete the proof, it suffices to show that $\mathrm{m}(A \cap A_n) = \mathrm{m}(A \cap A_0)$ for all n. Since $A_n = T^{-n}A_0$ and m is invariant, $\mathrm{m}(A_n) = \mathrm{m}(A_0)$. This implies $\mathrm{m}(A \cap A_n) = \mathrm{m}(A \cap A_0)$, because $A_n \subset A_0$. □

Despite the short proof, the Poincaré's theorem is powerful and very useful, as it applies to all probability-preserving transformation.

Example 3.2. *Let* $T : \mathbb{S}^1 \to \mathbb{S}^1 : x \mapsto x + \alpha \pmod 1$ *be a rotation of the circle and* $A \subset \mathbb{S}^1$ *a Cantor set of positive Lebesgue measure. Since the Lebesgue measure is invariant, the Poincaré recurrence theorem guarantees that, for almost every* $x \in A$*, the rotated point* $T^i(x)$ *is in the Cantor set* A *for infinitely many* $i \geq 1$*. Without any assumption on* α

and the structure of A, this is not immediately obvious; although for an irrational α the trajectories are dense, whereby they visit any interval, a Cantor set does not contain any intervals.

Given a set $A \in \mathscr{B}$ with $\mathrm{m}(A) > 0$, define the (first positive) hitting time

$$n_A(x) = \inf\{n \geq 1 \,:\, T^n(x) \in A\}, \quad x \in X,$$

with the convention that $\inf \emptyset = \infty$.

The Poincaré recurrence theorem guarantees that $n_A(x) < \infty$ for almost every $x \in A$. On the other hand, if T is ergodic, then $\mathrm{m}(\cup_{i\geq 1} T^{-i}A) = 1$ by Theorem 2.44, so $n_A(x) < \infty$ for almost every $x \in X$. (Note that $\cup_{i\geq 1} T^{-i}A$ is the set of those points whose trajectories visit A at least once: $x \in \cup_{i\geq 1} T^{-i}A \Leftrightarrow x \in T^{-i}A$ for some $i \geq 1$, which means $T^i(x) \in A$).

Poincaré's recurrence theorem tells us only that the trajectory of almost every point $x \in A$ will eventually return to the set A — that is, $T^n(x) \in A$ — and it actually does so infinitely often, but it does not yield any information about the frequency of these returns. Provided that the map T is ergodic, such information can be obtained from the following result:

Theorem 3.3 (Kac recurrence theorem)**.** *Suppose $(X, \mathscr{B}, \mathrm{m}, T)$ is an ergodic ppt. For any $A \in \mathscr{B}$ with $\mathrm{m}(A) > 0$ and any $f \in L^1(X, \mathscr{B}, \mathrm{m})$ we have*

$$\int_X f \,\mathrm{dm} = \int_A \sum_{i=0}^{n_A - 1} f \circ T^i \,\mathrm{dm} \ .$$

We emphasize that n_A is a function, so the number of terms in the sum is not constant. Choosing the constant function $f = 1$, Theorem 3.3 leads to $\int_A n_A \,\mathrm{dm} = 1$, which can be restated as a quantitative version of the Poincaré recurrence theorem:

Corollary 3.4. *Suppose $(X, \mathscr{B}, \mathrm{m}, T)$ is an ergodic ppt and $A \in \mathscr{B}$ with $\mathrm{m}(A) > 0$. Then the expected return time is*

$$\int_A n_A \,\mathrm{dm}_A = \frac{1}{\mathrm{m}(A)} \ ,$$

where m_A denotes the conditional probability measure on A, that is, $\mathrm{m}_A(B) = \frac{1}{\mathrm{m}(A)}\mathrm{m}(B \cap A)$. In particular, the expected return time is finite.

The finiteness of $\int_A n_A \,\mathrm{dm}_A$ is by itself a stronger result than the Poincaré recurrence theorem. (Of course ergodicity is assumed here unlike in Poincaré's theorem.) For example, if the set A occupies 1% of the state space X, then on average it takes 100 iterations of T before the trajectory of a typical point $x \in A$ returns to A.

Example 3.5. *To continue Example 3.2, assume that α is irrational, so the Lebesgue measure m is ergodic by Exercise 2.26. By the Kac recurrence theorem, one may expect that it takes $\frac{1}{\mathrm{m}(A)}$ rotations before the trajectory of a point returns to the Cantor set.*

There are various ways of proving Kac's theorem, of which we now present only one, following Omri Sarig's lecture notes **[15]**.

Proof of Theorem 3.3. We will first assume that f is *bounded* and *nonnegative*, as is common in measure-theoretic proofs, and finally extend to the general case.

Let us make the following observation: Denoting

$$1(n_A > n) = 1_{\{x\in X\,:\,n_A(x)>n\}} = 1_{T^{-n}A^c\cap\cdots\cap T^{-1}A^c}$$

for convenience, we have for all $n \geq 0$ that

$$\int_X 1(n_A > n)\, f\circ T^n \,\mathrm{dm} = \int_A 1(n_A > n)\, f\circ T^n \,\mathrm{dm} + \int_{A^c} 1(n_A > n)\, f\circ T^n \,\mathrm{dm}\ .$$

The nuisance is that the second integral on the right side is over A^c; the idea of the proof is to recursively express it as a sum of integrals over A. By basic identities such as $1_B\circ T = 1_{T^{-1}B}$ and by the invariance of m, we first see that

$$\begin{aligned}\int_{A^c} 1(n_A > n)\, f\circ T^n \,\mathrm{dm} &= \int_X 1_{A^c}1_{T^{-n}A^c\cap\cdots\cap T^{-1}A^c}\, f\circ T^n \,\mathrm{dm}\\ &= \int_X 1_{T^{-n}A^c\cap\cdots\cap T^{-1}A^c\cap A^c}\, f\circ T^n \,\mathrm{dm}\\ &= \int_X (1_{T^{-n}A^c\cap\cdots\cap T^{-1}A^c\cap A^c}\, f\circ T^n)\circ T \,\mathrm{dm}\\ &= \int_X 1_{T^{-n-1}A^c\cap\cdots\cap T^{-1}A^c}\, f\circ T^{n+1} \,\mathrm{dm}\\ &= \int_X 1(n_A > n+1)\, f\circ T^{n+1} \,\mathrm{dm}\ .\end{aligned}$$

This leads to the recursion formula

$$\int_X 1(n_A > n)\, f\circ T^n \,\mathrm{dm} = \int_A 1(n_A > n)\, f\circ T^n \,\mathrm{dm} + \int_X 1(n_A > n+1)\, f\circ T^{n+1} \,\mathrm{dm}$$

for all $n \geq 0$. Since $1(n_A > 0) = 1$ and $T^0 = \mathrm{id}_X$, the recursion formula yields

$$\int_X f\,\mathrm{dm} = \int_A \sum_{n=0}^{N-1} 1(n_A > n)\, f\circ T^n \,\mathrm{dm} + \int_X 1(n_A > N)\, f\circ T^N \,\mathrm{dm}\ , \quad N\geq 0\ . \tag{3.1}$$

Now, let us take the limit $N\to\infty$ in (3.1). To that end, note that

$$\left|\int_X f\circ 1(n_A > N)\, T^N \,\mathrm{dm}\right| \leq \|f\|_\infty\, \mathrm{m}(n_A > N) \to 0\ ,$$

because $n_A < \infty$ for almost every $x\in X$ by ergodicity (as pointed out above), which implies[1] $\lim_{N\to\infty}\mathrm{m}(n_A > N) = 0$. On the other hand, $\sum_{n=0}^{N-1} f\circ T^n\, 1(n_A > n)$ is increasing in N, so the monotone convergence theorem and some rearranging of nonnegative series yield

$$\begin{aligned}\int_X f\,\mathrm{dm} &= \int_A \sum_{n=0}^{\infty} 1(n_A > n)\, f\circ T^n \,\mathrm{dm} = \int_A \sum_{n=0}^{\infty}\sum_{j=n+1}^{\infty} 1(n_A = j)\, f\circ T^n \,\mathrm{dm}\\ &= \int_A \sum_{j=1}^{\infty}\sum_{n=0}^{j-1} 1(n_A = j)\, f\circ T^n \,\mathrm{dm} = \int_A \sum_{j=1}^{\infty} 1(n_A = j) \sum_{n=0}^{n_A-1} f\circ T^n \,\mathrm{dm}\ .\end{aligned}$$

Here $\sum_{j=1}^{\infty} 1(n_A = j) = 1(n_A > 0) = 1$, so the proof is complete for bounded nonnegative functions f. A nonnegative integrable function is the limit of an increasing sequence

[1]Note that the sequence is decreasing, $\{n_A > N+1\}\subset\{n_A > N\}$, so that $\lim_{N\to\infty}\mathrm{m}(n_A > N) = \mathrm{m}(\cap_{N\geq 0}\{n_A > N\}) = \mathrm{m}(n_A = \infty) = 0$.

of bounded nonnegative functions, so the result extends to that case by the monotone convergence theorem. The general case of an integrable function is proved by splitting the function into its positive and negative part. □

2. Preliminaries: Conditional expectation

Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space. A **sub-sigma-algebra** is a family $\mathscr{A} \subset \mathscr{B}$ which itself is a sigma-algebra. For example, the **trivial sigma-algebra**

$$\mathscr{N} = \{\emptyset, X\}$$

is a sub-sigma-algebra of every sigma-algebra on X. Recall from basic measure theory that, given any $f \in L^1(X, \mathscr{B}, \mathrm{m})$, the set function defined by

$$\mu_f(A) = \int_A f \, \mathrm{dm} \ , \quad \underline{A \in \mathscr{A}} \ , \tag{3.2}$$

is a finite complex measure on $(X, \mathscr{A})$. Let it be emphasized that the domain of definition is intentionally restricted to the sub-sigma-algebra $\mathscr{A}$. What is more, μ_f is absolutely continuous with respect to m, which can also be viewed as a measure on $(X, \mathscr{A})$. These simple observations have a powerful conclusion. Namely, by the Radon–Nikodym theorem (Theorem A.6), μ_f has a density $h_f \in L^1(X, \mathscr{A}, \mathrm{m})$ with respect to m, meaning that

$$\mu_f(A) = \int_A h_f \, \mathrm{dm} \ , \quad A \in \mathscr{A} \ . \tag{3.3}$$

Moreover, h_f is unique in the sense that any other candidate must agree with it almost everywhere. This is a nontrivial conclusion; there is the important distinction between (3.2) and (3.3) that the function f in general is *not* $\mathscr{A}$-measurable, so $f \neq h_f$.

Definition 3.6. *Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space, $\mathscr{A}$ a sub-sigma-algebra and $f \in L^1(X, \mathscr{B}, \mathrm{m})$. A function h_f is called the* **conditional expectation of f given $\mathscr{A}$** *if*

(1) h_f is $\mathscr{A}$-measurable, and
(2) for all $A \in \mathscr{A}$,

$$\int_A f \, \mathrm{dm} = \int_A h_f \, \mathrm{dm} \ . \tag{3.4}$$

We then denote

$$h_f = \mathrm{E}(f|\mathscr{A}) \ .$$

It can be checked that the preceding conditions imply $h_f \in L^1(X, \mathscr{A}, \mathrm{m})$. By the Radon–Nikodym discussion above, h_f exists and is unique in the L^1 sense. Let us also point out that the second condition holds if and only if

$$\int_X f g \, \mathrm{dm} = \int_X \mathrm{E}(f|\mathscr{A}) \, g \, \mathrm{dm}$$

for all bounded $\mathscr{A}$-measurable functions g. Note that the functions f and $\mathrm{E}(f|\mathscr{A})$ have the same average (or expectation):

$$\int_X f \, \mathrm{dm} = \int_X \mathrm{E}(f|\mathscr{A}) \, \mathrm{dm} \ .$$

It should be emphasized that $\mathrm{E}(f|\mathscr{A})$ is a function. It also depends on the measure m, but the dependence is left implicit when there is no danger of confusion about the underlying measure.

Before discussing the meaning of the construction, let us present some standard facts about conditional expectation, similar to those of ordinary expectation. The proofs, which are straightforward but which we omit, can be found in probability theory textbooks such as [4, 18].

Theorem 3.7. *Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space and $\mathscr{A}$ be a sub-sigma-algebra.*

(1) *(Linearity) The map $L^1(X, \mathscr{B}, \mathrm{m}) \to L^1(X, \mathscr{A}, \mathrm{m}) : f \mapsto \mathrm{E}(f|\mathscr{A})$ is linear.*
(2) *(Positivity) If $0 \leq f \in L^1(X, \mathscr{B}, \mathrm{m})$, then*
$$\mathrm{E}(f|\mathscr{A}) \geq 0 \ .$$
(3) *(Monotone convergence) Let $f_n, f \in L^1(X, \mathscr{B}, \mathrm{m})$ be such that $0 \leq f_n \uparrow f$. Then*
$$E(f_n|\mathscr{A}) \uparrow E(f|\mathscr{A}) \quad \textit{almost everywhere.}$$
(4) *(Jensen inequality) Let $f \in L^1(X, \mathscr{B}, \mathrm{m})$ be real valued and $\phi : \mathbb{R} \to \mathbb{R}$ be a convex Borel measurable function such that $\phi \circ f \in L^1(X, \mathscr{B}, \mathrm{m})$. Then*
$$\phi \circ \mathrm{E}(f|\mathscr{A}) \leq \mathrm{E}(\phi \circ f|\mathscr{A})$$

The Jensen inequality has a particularly important application: If $p \in [1, \infty)$ and $f \in L^p(X, \mathscr{B}, \mathrm{m})$, then $\mathrm{E}(f|\mathscr{A}) \in L^p(X, \mathscr{A}, \mathrm{m})$ because
$$\|\mathrm{E}(f|\mathscr{A})\|_p^p = \int_X |\mathrm{E}(f|\mathscr{A})|^p \,\mathrm{dm} \leq \int_X \mathrm{E}(|f|^p|\mathscr{A}) \,\mathrm{dm} = \int_X |f|^p \,\mathrm{dm} = \|f\|_p^p \ .$$
Likewise, for $f \in L^\infty(X, \mathscr{B}, \mathrm{m})$,
$$\|\mathrm{E}(f|\mathscr{A})\|_\infty \leq \mathrm{E}(\|f\|_\infty|\mathscr{A}) = \|f\|_\infty.$$
Briefly, the map
$$L^p(X, \mathscr{B}, \mathrm{m}) \to L^p(X, \mathscr{A}, \mathrm{m}) : f \mapsto \mathrm{E}(f|\mathscr{A})$$
is a linear positive contraction for any $p \in [1, \infty]$. As we will see shortly, it is also a projection, meaning that $\mathrm{E}(\mathrm{E}(f|\mathscr{A})|\mathscr{A}) = \mathrm{E}(f|\mathscr{A})$.

In elementary probability theory, the conditional expectation of $f \in L^1(X, \mathscr{B}, \mathrm{m})$ given a set $A \in \mathscr{B}$ of positive measure is defined to be $\mathrm{E}(f|A) = \frac{1}{\mathrm{m}(A)} \int_A f \,\mathrm{dm}$. The notion of conditional expectation given a sub-sigma-algebra just defined is a broad generalization of this elementary notion:

Exercise 3.8. *Let $\{A_1, \ldots, A_n\}$, $n \in \mathbb{N} \cup \{\infty\}$, be a finite or countable partition of X into measurable sets A_i with $\mathrm{m}(A_i) > 0$. That is, $\cup_{i=1}^n A_i = X$ and $A_i \cap A_j = \emptyset$ if $i \neq j$. Let $\mathscr{A} = \sigma(\{A_1, \ldots, A_n\})$. Show that, for almost every x,*
$$\mathrm{E}(f|\mathscr{A})(x) = \frac{1}{\mathrm{m}(A_i)} \int_{A_i} f \,\mathrm{dm}$$
where A_i is the element of the partition containing x. Note that the conclusion does not change if the definition of a partition is relaxed to $\mathrm{m}(X \setminus \cup_{i=1}^n A_i) = 0$ and $\mathrm{m}(A_i \cap A_j) = 0$ for $i \neq j$.

Example 3.9. *Consider the interval $[0, 1)$. For a fixed $n \geq 1$, the partition $\{[2^{-n}(i - 1), 2^{-n}i) : 1 \leq i \leq 2^n\}$ generates a sub-sigma-algebra $\mathscr{A}_n \subset \mathscr{B}$ of the Borel sigma-algebra. We have*
$$\mathrm{E}(f|\mathscr{A}_n)(x) = \frac{1}{2^{-n}} \int_{2^{-n}(i-1)}^{2^{-n}i} f(y) \,\mathrm{d}y$$

for $2^{-n}(i-1) \le x < 2^{-n}i$. *Then* $\mathrm{E}(f|\mathscr{A}_n)$ *is constant on each partition element, with the constant being the average of* f *on that element. In a sense,* $\mathrm{E}(f|\mathscr{A}_n)$ *is the function that best approximates* f *while being constant on partition elements.*

In general, it is very useful to think of $\mathrm{E}(f|\mathscr{A})$ as the best $\mathscr{A}$-measurable approximation of a given $f \in L^1(X, \mathscr{B}, \mathrm{m})$; the bigger the sub-sigma-algebra $\mathscr{A}$ the better the approximation. The following facts are also standard and consistent with the approximation idea. The proofs can be found in the textbooks [4,18], but the reader is encouraged to give them a try.

Theorem 3.10. *Let* $(X, \mathscr{B}, \mathrm{m})$ *be a probability space.*

(1) If $\mathscr{A}_1 \subset \mathscr{A}_2$ *are two sub-sigma-algebras and* $f \in L^1(X, \mathscr{B}, \mathrm{m})$*, then*
$$\mathrm{E}(\mathrm{E}(f|\mathscr{A}_1)|\mathscr{A}_2) = \mathrm{E}(f|\mathscr{A}_1) \quad \textit{and} \quad \mathrm{E}(\mathrm{E}(f|\mathscr{A}_2)|\mathscr{A}_1) = \mathrm{E}(f|\mathscr{A}_1)$$
(2) If $\mathscr{A}$ *is a sub-sigma-algebra,* f *is* $\mathscr{A}$*-measurable and* $g, fg \in L^1(X, \mathscr{B}, \mathrm{m})$*, then*
$$\mathrm{E}(fg|\mathscr{A}) = f\,\mathrm{E}(g|\mathscr{A}) \ .$$
(3) If $f \in L^1(X, \mathscr{B}, \mathrm{m})$*, then*
$$\mathrm{E}(f|\mathscr{N}) = \int_X f\,\mathrm{dm} \quad \textit{and} \quad \mathrm{E}(f|\mathscr{B}) = f \ ,$$
where $\mathscr{N} = \{\emptyset, X\}$ *is the trivial sigma-algebra.*

(4) In L^2 *the conditional expectation has a special role: If* $\mathscr{A}$ *is a sub-sigma-algebra, then*
$$L^2(X, \mathscr{B}, \mathrm{m}) \to L^2(X, \mathscr{A}, \mathrm{m}) : f \mapsto \mathrm{E}(f|\mathscr{A})$$
is the orthogonal projection onto the subspace $L^2(X, \mathscr{A}, \mathrm{m})$ *of* $L^2(X, \mathscr{B}, \mathrm{m})$*.*

The first fact implies as a special case the earlier claim that the conditional expectation is a projection. The proof of the last fact is based on the splitting $f = f - \mathrm{E}(f|\mathscr{A}) + \mathrm{E}(f|\mathscr{A})$, where $f - \mathrm{E}(f|\mathscr{A})$ and $\mathrm{E}(f|\mathscr{A})$ are orthogonal by the second fact. In particular, given any $f \in L^2(X, \mathscr{B}, \mathrm{m})$, the conditional expectation $\mathrm{E}(f|\mathscr{A})$ is the unique $g \in L^2(X, \mathscr{A}, \mathrm{m})$ which minimizes the distance $\|f - g\|_2$ — it is the best $\mathscr{A}$-measurable approximation of f.

Before leaving the general topic of sub-sigma-algebras and conditional expectation, we make a few simple observations that will serve us in the future.

If $T : X \to X$ is a measurable map and $\mathscr{A}$ is a sub-sigma-algebra of $\mathscr{B}$, we define
$$T^{-n}\mathscr{A} = \{T^{-n}A \,:\, A \in \mathscr{A}\}$$
for each $n \ge 1$.

Exercise 3.11. *Show that each* $T^{-n}\mathscr{A}$ *is a sub-sigma-algebra of* $\mathscr{B}$*.*

Given two sub-sigma-algebras $\mathscr{A}_1, \mathscr{A}_2 \subset \mathscr{B}$, the meaning of an expression such as $\mathscr{A}_1 \subset \mathscr{A}_2$ is the usual set-theoretic one. On a probability space $(X, \mathscr{B}, \mathrm{m})$ sets of measure zero do not usually matter. For this reason, we write
$$\mathscr{A}_1 \subset \mathscr{A}_2 \pmod{\mathrm{m}}$$
if for every $A_1 \in \mathscr{A}_1$ there exists $A_2 \in \mathscr{A}_2$ such that $\mathrm{m}(A_1 \,\Delta\, A_2) = 0$. We also write
$$\mathscr{A}_1 = \mathscr{A}_2 \pmod{\mathrm{m}}$$

if $\mathscr{A}_1 \subset \mathscr{A}_2 \pmod{\mathrm{m}}$ and $\mathscr{A}_2 \subset \mathscr{A}_1 \pmod{\mathrm{m}}$. For example, $\mathscr{A} = \mathscr{N} \pmod{\mathrm{m}}$ means that the sub-sigma-algebra $\mathscr{A} \subset \mathscr{B}$ contains only sets of measure 0 or 1. The next lemma shows that differences (mod m) do not affect the conditional expectations. We present the proof for completeness, but it can be skipped.

Lemma 3.12. *Let $\mathscr{A}_1 = \mathscr{A}_2 \pmod{\mathrm{m}}$ be two sub-sigma-algebras. For any $f \in L^1(X, \mathscr{B}, \mathrm{m})$,*

$$\mathrm{E}(f|\mathscr{A}_1) = \mathrm{E}(f|\mathscr{A}_2) \quad \textit{almost everywhere.}$$

Proof. Denote $g_i = \mathrm{E}(f|\mathscr{A}_i)$. This is the unique element (equivalence class of functions) in $L^1(X, \mathscr{A}_i, \mathrm{m}) \subset L^1(X, \mathscr{B}, \mathrm{m})$ such that

$$\int_X f\, 1_A \,\mathrm{dm} = \int_X g_i\, 1_A \,\mathrm{dm} \ , \quad A \in \mathscr{A}_i \ .$$

We claim that $g_1 = g_2$ a.e. Because $\mathscr{A}_1 = \mathscr{A}_2 \pmod{\mathrm{m}}$, we obtain

$$\int_X f\, 1_A \,\mathrm{dm} = \int_X g_i\, 1_A \,\mathrm{dm} \ , \quad A \in \mathcal{A}(\mathscr{A}_1 \cup \mathscr{A}_2) \ ,$$

where $\mathcal{A}(\mathscr{A}_1 \cup \mathscr{A}_2)$ is the algebra generated by $\mathscr{A}_1 \cup \mathscr{A}_2$. Denote by $\mathscr{M}_i$ the collection of all $A \in \mathscr{B}$ such that

$$\int_X f\, 1_A \,\mathrm{dm} = \int_X g_i\, 1_A \,\mathrm{dm} \ .$$

This is a monotone class, and we have just checked that $\mathcal{A}(\mathscr{A}_1 \cup \mathscr{A}_2) \subset \mathscr{M}_i$. By the monotone class theorem, $\sigma(\mathscr{A}_1 \cup \mathscr{A}_2) \subset \mathscr{M}_i$, meaning

$$\int_X f\, 1_A \,\mathrm{dm} = \int_X g_i\, 1_A \,\mathrm{dm} \ , \quad A \in \sigma(\mathscr{A}_1 \cup \mathscr{A}_2) \ .$$

Because $g_i \in L^1(X, \sigma(\mathscr{A}_1 \cup \mathscr{A}_2), \mathrm{m})$, we have $g_i = \mathrm{E}(f|\sigma(\mathscr{A}_1 \cup \mathscr{A}_2))$, $i = 1, 2$. □

Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation. We denote by

$$\mathscr{I} = \{A \in \mathscr{B} \,:\, \mathrm{m}(A \,\Delta\, T^{-1}A) = 0\}$$

the family of almost invariant sets and by

$$\mathscr{I}_0 = \{A \in \mathscr{B} \,:\, A = T^{-1}A\}$$

the family of invariant sets.

Exercise 3.13. *Prove the following:*

(1) *$\mathscr{I}$ and $\mathscr{I}_0$ are sub-sigma-algebras of $\mathscr{B}$.*
(2) *$\mathscr{I}$ and $\mathscr{I}_0$ satisfy*

$$\mathscr{I}_0 = T^{-1}\mathscr{I}_0$$

and

$$\mathscr{I} = T^{-1}\mathscr{I} = \mathscr{I}_0 \pmod{\mathrm{m}} \ .$$

(3) m *is ergodic if and only if $\mathscr{I} = \mathscr{I}_0 = \mathscr{N} \pmod{\mathrm{m}}$.*
(4) *A function is $\mathscr{I}_0$-measurable or $\mathscr{I}$-measurable if and only if it is invariant or almost invariant, respectively.*

Exercise 3.14. *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation. Show that if $\mathscr{A}$ is a sub-sigma-algebra and $f \in L^1(X, \mathscr{B}, \mathrm{m})$, then*

$$\mathrm{E}(f|\mathscr{A}) \circ T = \mathrm{E}(f \circ T|T^{-1}\mathscr{A}) \ . \tag{3.5}$$

(This identity is not true if m *is not an invariant measure.)*

The identity in (3.5) may look slightly strange, but it is easy to remember: We may think of $T^{-1}\mathscr{A}$ as the future sub-sigma-algebra (it consists of the sets $T^{-1}A$, $A \in \mathscr{A}$, involving the future event $T(x) \in A$) and of $g(T(x))$ as the future value of a function g. Thus, (3.5) states that the future value of the conditional expectation $\mathrm{E}(f|\mathscr{A})$ coincides with the conditional expectation of the future value of f given the future sub-sigma-algebra.

An application of the observations above is that

$$\mathrm{E}(f \circ T|\mathscr{I}) = \mathrm{E}(f|\mathscr{I}) = \mathrm{E}(f|\mathscr{I}) \circ T$$

holds almost everywhere, for all $f \in L^1(X, \mathscr{B}, \mathrm{m})$.

3. Von Neumann's mean ergodic theorem

We return to the convergence of time averages. The next result is a historical landmark in ergodic theory. In fact, von Neumann considered the theorem one of his three greatest achievements. His original proof was based on the spectral theory on Hilbert spaces developed by Stone and him. The much simpler proof given here follows the presentation of Parry's book [9], which gives credit for the simplification to Hopf. It avoids spectral theory, but linear isometries on a Hilbert space will be studied. Before we formulate von Neumann's theorem, let us elucidate how linear operators on Hilbert spaces arise in the study of time averages.

Definition 3.15. *Let $(X, \mathscr{B})$ be a measurable space and T a measurable map. The linear operator*

$$Uf = f \circ T$$

acting on the vector space of measurable functions is called the **Koopman operator***.*

Definition 3.16. *A linear operator $L : V \to V$ on a normed space V is an* **isometry** *if*

$$\|Lv\| = \|v\| \ , \quad v \in V \ .$$

Exercise 3.17. *Prove that the Koopman operator satisfies $\|Uf\|_p = \|f\|_p$ for all $f \in L^p(X, \mathscr{B}, \mathrm{m})$ and all $p \in [1, \infty)$. What can be said about the case of $L^\infty(X, \mathscr{B}, \mathrm{m})$?*

Theorem 3.18 (Von Neumann ergodic theorem)**.** *Let U be an isometry on a Hilbert space H. Denote by $I = \{x \in H : Ux = x\}$ the subspace of U-invariant vectors, and let $P : H \to I$ be the orthogonal projection onto I. Then*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} U^i x = Px \ , \quad x \in H \ .$$

In particular, if $(X, \mathscr{B}, \mathrm{m}, T)$ is a probability-preserving transformation and $\mathscr{I} = \{A \in \mathscr{B} : \mathrm{m}(T^{-1}A \,\Delta\, A) = 0\}$ is the sub-sigma-algebra of almost invariant sets, then

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i = \mathrm{E}(f|\mathscr{I}) \ , \quad f \in L^2(X, \mathscr{B}, \mathrm{m}) \ ,$$

where convergence takes place in $L^2(X, \mathscr{B}, \mathrm{m})$. If m is ergodic, then

$$\mathrm{E}(f|\mathscr{I}) = \int_X f \, \mathrm{dm}$$

almost everywhere.

We remark that any $f \in L^1(X, \mathscr{B}, \mathrm{m})$ has the same average as $\mathrm{E}(f|\mathscr{I})$:

$$\int_X \mathrm{E}(f|\mathscr{I})\,\mathrm{dm} = \int_X f\,\mathrm{dm} \tag{3.6}$$

according to the properties of conditional expectation. In the ergodic case $\mathrm{E}(f|\mathscr{I}) = \int_X f\,\mathrm{dm}$ almost everywhere, because the function $\mathrm{E}(f|\mathscr{I})$ is almost invariant and, as such, constant almost everywhere.

Note that the convergence in the second part of the theorem takes place in the norm $\|\cdot\|_2 = \|\cdot\|_{L^2(X,\mathscr{B},\mathrm{m})}$. In other words, $\lim_{n\to\infty} \|\frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i - \mathrm{E}(f|\mathscr{I})\|_2 = 0$. This is both satisfactory and unsatisfactory; on the one hand, convergence in norm is a strong notion of convergence, but on the other hand, in ergodic theory one would also like to have almost sure convergence. In the next section we will prove Birkhoff's almost sure ergodic theorem, which guarantees almost sure convergence for all f in the larger space $L^1(X, \mathscr{B}, \mathrm{m})$, as well as convergence in the weaker L^1 norm. There is a version of von Neumann's theorem for measure-preserving transformations applicable to any of the Banach spaces $L^p(X, \mathscr{B}, \mathrm{m})$ with $p \in [1, \infty)$, which guarantees that the time averages of an L^p function converge in the corresponding norm $\|\cdot\|_p$; see, e.g., [**17**]. Such a theorem can be proved as a corollary of the L^2 version or of Birkhoff's theorem. Namely, the following equivalence among the L^p spaces holds:

Exercise 3.19. *Suppose there exists $p \in [1, \infty)$ such that*

$$\lim_{n\to\infty} \left\| \frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i - \mathrm{E}(f|\mathscr{I}) \right\|_p = 0 \ , \quad f \in L^p(X, \mathscr{B}, \mathrm{m}) \ .$$

Show that the same statement then holds in fact for all $p \in [1, \infty)$.

[Hint: A truncation argument works; given any $M > 0$, split $f = f1_{\{f\le M\}} + f1_{\{f>M\}}$, establish the corresponding result first for the bounded part, and then show that the unbounded part does not contribute in the limit $M \to \infty$.]

Recall that, according to Theorem 2.4, the time averages of $f \in C(x)$ converge to a continuous function $f_+ \in C(X)$ if and only if $f \in I \oplus \bar{B}$. There I was the subspace of continuous invariant functions and B was the subspace of continuous coboundaries. The proof of the von Neumann ergodic theorem we present is based on a similar splitting $H = I \oplus \bar{B}$ of the Hilbert space with the following definitions:

Definition 3.20. *Let H be a Hilbert space and $U : H \to H$ an isometry. An element $x \in H$ is called* **invariant** *if $x = Ux$. Thus*

$$I = \{x \in H \,:\, x = Ux\}$$

is the (closed) subspace of all invariant elements.

An element $x \in H$ is called a **coboundary** *if there exists $y \in H$ such that $x = y - Uy$. Thus*

$$B = \{y - Uy \,:\, y \in H\}$$

is the (generally not closed) subspace of all coboundaries.

A remark is in order. In the case where $H = L^2(X, \mathscr{B}, \mathrm{m})$ and U is the Koopman operator of a probability-preserving transformation, $f \in I$ means that $f = Uf = f\circ T$ holds in the almost sure sense of L^2; that is, f is an almost invariant L^2 function. Likewise

$f \in B$ means that there exists $f \in L^2(X, \mathscr{B}, \mathrm{m})$ such that $f = g - g \circ T$ almost everywhere; such functions are also called L^2 coboundaries.

The following observation is at the heart of von Neumann's theorem. The reader is invited to compare it with Lemma 2.3, which was equally important for Theorem 2.4.

Lemma 3.21. *Let H be a Hilbert space and $U : H \to H$ an isometry. Then*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} U^i x = \begin{cases} x & \text{if } x \in I, \\ 0 & \text{if } x \in \bar{B}, \end{cases}$$

where the convergence takes place in the norm of the Hilbert space.

Exercise 3.22. *Prove Lemma 3.21.*

Thus, the proof of Theorem 3.18 amounts to showing that $H = I \oplus^{\perp} \bar{B}$, where the symbol $\oplus^{\perp}$ means that I and $\bar{B}$ are orthogonal complements, $I = \bar{B}^{\perp}$. We are ready to enter the proof after recalling some elementary facts about isometries. Below, U^* is the adjoint of U, which is the continuous linear operator $H \to H$ satisfying

$$\langle Ux, y\rangle = \langle x, U^* y\rangle \ , \quad x, y \in H \ .$$

Lemma 3.23. *Let H be a Hilbert space.*

(1) An isometry $U : H \to H$ is one-to-one and continuous.

(2) A continuous linear operator $U : H \to H$ is an isometry if and only if

$$\langle x, y\rangle = \langle Ux, Uy\rangle \ , \quad x, y \in H \ .$$

(3) A continuous linear operator $U : H \to H$ is an isometry if and only if

$$U^* U = I \ .$$

(4) If $U : H \to H$ is an isometry, then

$$Ux = x \quad \Longleftrightarrow \quad U^* x = x \ .$$

Exercise 3.24. *Prove Lemma 3.23.*

[Hint: In (4) express $\|x - Ux\|^2$ in term of U^*.]

Proof of Theorem 3.18. Suppose first that $H = I \oplus^{\perp} \bar{B}$ and that $P : H \to I$ is the orthogonal projection onto I. Then $Px = x$ for $x \in I$ and $Px = 0$ for $x \in \bar{B}$. Lemma 3.21 then implies that $\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} U^i x = Px$ for all $x \in H$. It remains to prove that $H = I \oplus^{\perp} \bar{B}$.

$\bar{B}^{\perp} \subset I$: If $x \in \bar{B}^{\perp}$, then $\langle x, b\rangle = 0$ for any $b \in \bar{B}$. In particular, $\langle x, Uy - y\rangle = 0$ for any $y \in H$, which implies $\langle U^* x - x, y\rangle = 0$ or $U^* x = x$. Equivalently, $Ux = x$, meaning $x \in I$.

$I \subset \bar{B}^{\perp}$: If $x \in I$, then $U^* x = x$, which implies $\langle x, Uy - y\rangle = \langle U^* x - x, y\rangle = 0$ for any $y \in H$. Thus, $\langle x, b\rangle$ for any $b \in B$. This actually holds for any $b \in \bar{B}$: if $b_n \in B$ converges in H to $b \in \bar{B}$, then $0 = \langle x, b_n\rangle \to \langle x, b\rangle$. Hence, $x \perp \bar{B}$.

In the case of the probability-preserving transformation, the Hilbert space is $H = L^2(X, \mathscr{B}, \mathrm{m})$ and the isometry is $U : H \to H : Uf = f \circ T$. We only need a few remarks. Firstly, the subspace $I \subset L^2(X, \mathscr{B}, \mathrm{m})$ of invariant elements is precisely the subspace of almost invariant L^2 functions. Moreover, f is almost invariant if and only if f is $\mathscr{I}$-measurable, so $I = L^2(X, \mathscr{I}, \mathrm{m})$. The orthogonal projection P onto I is thus $\mathrm{E}(\,\cdot\,|\mathscr{I})$ by Theorem 3.10. In the ergodic case, $\mathrm{E}(f|\mathscr{I}) = \int_X f\,\mathrm{dm}$ as explained below the theorem. □

4. Birkhoff's almost sure ergodic theorem

In this section we introduce the most fundamental result in ergodic theory, due to George Birkhoff, which can be viewed as an almost sure version of von Neumann's ergodic theorem. The result has several names in the literature, such as the pointwise ergodic theorem, almost sure ergodic theorem and individual ergodic theorem. We first state the theorem,

Theorem 3.25 (Birkhoff ergodic theorem)**.** *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation and $\mathscr{I} = \{A \in \mathscr{B} : \mathrm{m}(T^{-1}A \,\Delta\, A) = 0\}$ the sub-sigma-algebra of almost invariant sets. Then, for all $f \in L^1(X, \mathscr{B}, \mathrm{m})$,*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i = \mathrm{E}(f|\mathscr{I}) \ ,$$

where the convergence takes place both <u>almost surely</u> and in the L^1 norm. If m is ergodic, then

$$\mathrm{E}(f|\mathscr{I}) = \int_X f\,\mathrm{dm}$$

almost everywhere.

Recall from (3.6) that the limit $\mathrm{E}(f|\mathscr{I})$ satisfies $\int_X \mathrm{E}(f|\mathscr{I})\,\mathrm{dm} = \int_X f\,\mathrm{dm}$, and that in the ergodic case $\mathrm{E}(f|\mathscr{I}) = \int_X f\,\mathrm{dm}$ almost everywhere.

From Birkhoff's theorem, we get immediately two additional characterizations of ergodicity, the first of which was promised in (2.10).

Theorem 3.26. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. A measure $\mathrm{m} \in \mathcal{P}_T$ is ergodic if and only if one of the following equivalent conditions is satisfied:*

(1) For all $f \in L^1(X, \mathscr{B}, \mathrm{m})$ and almost every $x \in X$,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) = \int_X f\,\mathrm{dm} \ .$$

(2) For all $A \in \mathscr{B}$ and almost every $x \in X$,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} 1_A(T^i(x)) = \mathrm{m}(A) \ .$$

In other words, ergodicity means precisely that the time averages converge to the space average almost surely. Condition (2) has an interesting dynamical interpretation as well: it states that the trajectory of a typical point visits every region of the state space more or less regularly, in such a way that the asymptotic frequency of visits is proportional to the measure of the region. Recall that we already knew that there are infinitely many visits by Exercise 2.45, but had no quantitative information about the times of visits. Thus, much like the Kac recurrence theorem (or Corollary 3.4) quantifies the time of first return to a set (starting from a point in the same set), the Birkhoff ergodic theorem quantifies the times of the infinitely many visits to a set (starting from a point in X).

PROOF OF THEOREM 3.26. Ergodicity implies (1), which implies (2). If $A \in \mathscr{B}$ is an invariant set, then 1_A is an invariant function, $1_A = 1_A \circ T^i$ for all $i \geq 1$, and (2) implies that $1_A(x) = \mathrm{m}(A)$ for almost every $x \in X$. This implies $\mathrm{m}(A) \in \{0, 1\}$, so (2) implies ergodicity. □

The proof of Theorem 3.25 is somewhat more involved than that of the von Neumann ergodic theorem: Birkhoff's theorem applies to the larger class $L^1(X, \mathscr{B}, \mathrm{m})$ of functions f, and also gives information about the convergence of the time averages $\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))$ for individual points $x \in X$. Nevertheless, the particular proof we present — which follows Garsia [5, 6] as in Parry [9] — proceeds along very similar lines to the above proof of von Neumann's theorem. It is based on the splitting $L^1(X, \mathscr{B}, \mathrm{m}) = I \oplus \bar{B}$ with the following definitions:

Definition 3.27. *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation. Then*

$$I = \{f \in L^1(X, \mathscr{B}, \mathrm{m}) \, : \, f = f \circ T\}$$

is the (closed) subspace of $L^1(X, \mathscr{B}, \mathrm{m})$ of all almost invariant functions.

A function $f \in L^1(X, \mathscr{B}, \mathrm{m})$ is called an L^1 **coboundary** *if there exists $g \in L^1(X, \mathscr{B}, \mathrm{m})$ such that $f = g - g \circ T$ almost everywhere. Thus*

$$B = \{g - g \circ T \, : \, g \in L^1(X, \mathscr{B}, \mathrm{m})\}$$

is the (generally not closed) subspace of all L^1 coboundaries.

Analogously to Lemma 3.21, we have the following Lemma:

Lemma 3.28. *In the above setting,*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i = \begin{cases} f & \text{if } f \in I, \\ 0 & \text{if } f \in \bar{B}, \end{cases}$$

where the convergence takes place <u>in the L^1 norm</u>, and the equality holds almost everywhere (that is, in the sense of L^1). In the case of I, the convergence is also almost sure.

Exercise 3.29. *Prove Lemma 3.28.*

It will indeed turn out $L^1(X, \mathscr{B}, \mathrm{m}) = I \oplus \bar{B}$, which is an important part of the proof. Then the preceding lemma implies a convergence result for all functions in this space. But we can already see a missing element: we need *almost sure* convergence, also on $\bar{B}$. The problem is that a general $f = g - g \circ T \in B$, the function g is *not a bounded function*. Let us give a name to the special case when g is bounded:

Definition 3.30. *A measurable function f is called an L^∞* **coboundary** *if there exists $g \in L^\infty(X, \mathscr{B}, \mathrm{m})$ such that $f = g - g \circ T$ almost everywhere. We write*

$$B_\infty = \{g - g \circ T \, : \, g \in L^\infty(X, \mathscr{B}, \mathrm{m})\} \, .$$

It is obvious that $B_\infty \subset B$ is a dense subspace of $\bar{B}$ in the topology of $L^1(X, \mathscr{B}, \mathrm{m})$:

$$\bar{B}_\infty = \bar{B} \, . \tag{3.7}$$

Exercise 3.31. *Show that*

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i = 0 \quad \underline{\textit{almost everywhere}} \quad \textit{if } f\in B_\infty \ . \tag{3.8}$$

Since $L^1(X,\mathscr{B},\mathrm{m}) = I\oplus \bar{B} = I\oplus \bar{B}_\infty$, it would seem that we are almost done; we just need to extend (3.8) to all of $\bar{B}_\infty$ by closing the subspace B_∞. However, the closure in $\bar{B}_\infty$ here is in the sense of L^1, not L^∞. To accomplish such an extension of the almost sure convergence to all of $\bar{B}_\infty$ (hence to all of $L^1(X,\mathscr{B},\mathrm{m})$), we need a powerful tool — the maximal ergodic theorem. We introduce it next, and finish the proof of the Birkhoff ergodic theorem after that.

Here is a fun application of Birkhoff's theorem:

Exercise 3.32. *Let $(X,\mathscr{B},\mathrm{m},T)$ be a probability-preserving transformation. Show that, for almost every $x\in X$ and all $f\in L^1(X,\mathscr{B},\mathrm{m})$, $\lim_{n\to\infty}\frac{1}{n}f(T^n(x))=0$.*

5. Maximal ergodic theorem

Theorem 3.33. [2] *Let $f\in L^1_{\mathbb{R}}(X,\mathscr{B},\mathrm{m})$. Denote $S_k=\sum_{i=0}^{k-1} f\circ T^i$, $k\geq 0$, with the convention that $S_0=0$. Also denote $M_n=\max_{0\leq k\leq n} S_k$, $n\geq 0$. Then*

$$\int_{\{M_n>0\}} f\,\mathrm{dm}\geq 0\ ,\quad n\geq 0\ ,$$

where $\{M_n>0\}=\{x\in X\ :\ \max_{0\leq k\leq n} S_k(x)>0\}$.

Observe that $M_n\geq 0$ because $S_0=0$. Accordingly, $X=\{M_n>0\}\cup\{M_n=0\}$. Here $\{M_n>0\}$ is the set of those points $x\in X$ for which the time average $\frac{1}{k}\sum_{i=0}^{k-1} f(T^i(x))$ is greater than 0 at least for one $k\in[1,n]$. By the maximal ergodic theorem, the conditional mean of f on this set is positive. On the complement $\{M_n=0\}$, $f\leq 0$ because $S_1=f$. Hence, conditioned on the event that the time average is at most 0 for all $k\in[1,n]$, the mean of f is nonpositive:

$$\int_{\{M_n=0\}} f\,\mathrm{dm}\leq 0\ ,\quad n\geq 0\ .$$

Thus the maximal ergodic theorem tells something rather nontrivial about the connection between time and space averages.

Proof of Theorem 3.33. (Following Garsia [5].) We prove the result with the aid of the Koopman operator $Uf=f\circ T$. Note that it can be viewed as a linear operator $L^1_{\mathbb{R}}(X,\mathscr{B},\mathrm{m})\to L^1_{\mathbb{R}}(X,\mathscr{B},\mathrm{m})$, and that $S_k=\sum_{i=0}^{k-1} U^i f$. The proof is based on three fundamental observations, the first of which is the trivial identity

$$f+US_k=S_{k+1}\ , \tag{3.9}$$

while the second one is the bound

$$U\left(\max_{0\leq k\leq n} S_k\right)\geq \max_{0\leq k\leq n} US_k\ . \tag{3.10}$$

[2]In the proof the only property about the map $(X,\mathscr{B},\mathrm{m},T)$ used is that the Koopman operator $Uf=f\circ T$ is a positive contraction. Hence, a similar result is true for any positive contraction U on $L^1(X,\mathscr{B},\mathrm{m})$ with the definition of S_k taken to be $S_k=\sum_{i=0}^{k-1} U^i f$.

We save the proof of this bound for last and take it now for granted. The third fundamental observation is that U is a positive contraction:

$$h \geq 0 \Rightarrow Uh \geq 0 \quad \text{and} \quad \|Uh\|_1 \leq \|h\|_1 \ .$$

It follows from (3.9)–(3.10) that

$$f + UM_n \geq f + \max_{0\leq k\leq n} US_k = \max_{0\leq k\leq n}(f + US_k) = \max_{0\leq k\leq n} S_{k+1} = \max_{1\leq k\leq n+1} S_k \geq \max_{1\leq k\leq n} S_k \ .$$

If $x \in \{M_n > 0\}$, then $\max_{1\leq k\leq n} S_k(x) = \max_{0\leq k\leq n} S_k(x) = M_n(x)$, because $S_0 = 0$. Therefore, we have shown that

$$f(x) + UM_n(x) \geq M_n(x) \ , \quad x \in \{M_n > 0\} \ .$$

This implies

$$\int_{\{M_n>0\}} f \,\mathrm{dm} \geq \int_{\{M_n>0\}} M_n - UM_n \,\mathrm{dm} \ .$$

The right side is nonnegative:

$$\int_{\{M_n>0\}} M_n - UM_n \,\mathrm{dm} = \int_X M_n - UM_n \,\mathrm{dm} \geq 0 \ .$$

The first equality follows from the fact that $M_n \geq 0$ (recall $S_0 = 0$), so either $M_n > 0$ or $M_n = 0$. In the second inequality we also used the fact that U is a positive contraction:

$$0 \leq \|M_n\|_1 - \|UM_n\|_1 = \int_X |M_n| - |UM_n| \,\mathrm{dm} = \int_X M_n - UM_n \,\mathrm{dm} \ .$$

It remains to prove (3.10). Given $f, g \in L^1_{\mathbb{R}}(X, \mathscr{B}, \mathrm{m})$, we have the decomposition $\max(f, g) = (f - g)_+ + g$. This yields $U(\max(f, g)) = U((f - g)_+) + Ug \geq Ug$, because U is positive. Switching the roles of f and g, we also have $U(\max(f, g)) \geq Uf$, so $U(\max(f, g)) \geq \max(Uf, Ug)$. The general case is proved by induction. □

We now present a straightforward corollary, which makes it more obvious how the maximal ergodic theorem is relevant to the limit behavior of time averages.

Corollary 3.34. *For all $f \in L^1(X, \mathscr{B}, \mathrm{m})$ and $\lambda > 0$,*

$$\mathrm{m}\left\{x \in X \ : \ \limsup_{n\to\infty} \left|\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))\right| > \lambda\|f\|_1\right\}$$
$$\leq \mathrm{m}\left\{x \in X \ : \ \sup_{n\geq 1} \left|\frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x))\right| > \lambda\|f\|_1\right\} \leq \frac{1}{\lambda} \ .$$

Proof. Only the last inequality requires a proof. If we can prove it for $|f|$ in place of f, then the result follows also for f, so it suffices to assume that $f \geq 0$. This has the advantage that all the absolute values disappear from the formulas. Let us now define the sets

$$A_n = \left\{x \in X \ : \ \max_{1\leq k\leq n} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) > \lambda\|f\|_1\right\}$$

for each $n \geq 1$. Then

$$\left\{x \in X \ : \ \sup_{n\geq 1} \frac{1}{n}\sum_{i=0}^{n-1} f(T^i(x)) > \lambda\|f\|_1\right\} = \bigcup_{n\geq 1} A_n \ .$$

Defining the function $g = f - \lambda\|f\|_1$, we have

$$A_n = \left\{ x \in X \,:\, \max_{1\le k\le n} \frac{1}{n}\sum_{i=0}^{n-1}(f(T^i(x)) - \lambda\|f\|_1) > 0 \right\}$$
$$= \left\{ x \in X \,:\, \max_{1\le k\le n} \frac{1}{n}\sum_{i=0}^{n-1} g(T^i(x)) > 0 \right\} = \left\{ x \in X \,:\, \max_{0\le k\le n} \sum_{i=0}^{n-1} g(T^i(x)) > 0 \right\} .$$

By the maximal ergodic theorem,

$$\int_{A_n} g \,\mathrm{dm} \ge 0 \quad \text{so that} \quad \mathrm{m}(A_n)\,\lambda\|f\|_1 \le \int_{A_n} f \,\mathrm{dm} \le \|f\|_1 \; ,$$

which implies

$$\mathrm{m}(A_n) \le \frac{1}{\lambda} \; , \quad n \ge 1 \; .$$

Because $A_1 \subset A_2 \subset \cdots$, we get $\mathrm{m}(\cup_{n\ge1} A_n) = \sup_{n\ge1} \mathrm{m}(A_n) \le \frac{1}{\lambda}$, and the proof is complete. □

We finish the section with the proof of Birkhoff's theorem.

Proof of Theorem 3.25. Let us first show that $L^1(X, \mathscr{B}, \mathrm{m}) = I \oplus \bar{B}$. To accomplish that — very similarly to the proofs of Theorems 3.18 and 2.4 — we define the linear projection operator $P : L^1(X, \mathscr{B}, \mathrm{m}) \to I : Pf = \mathrm{E}(f|\mathscr{I})$. Then $L^1(X, \mathscr{B}, \mathrm{m}) = \operatorname{im} P \oplus \ker P = I \oplus \ker P$, and we need to establish $\ker P = \bar{B}$. Since $\|\mathrm{E}(f|\mathscr{I})\|_1 = \|f\|_1$, the operator P is continuous, so $\ker P$ is closed. Moreover, $\mathrm{E}(g - g\circ T|\mathscr{I}) = \mathrm{E}(g|\mathscr{I}) - \mathrm{E}(g\circ T|\mathscr{I}) = \mathrm{E}(g|\mathscr{I}) - \mathrm{E}(g|T^{-1}\mathscr{I}) = 0$, so $B \subset \ker P$, which implies $\bar{B} \subset \ker P$.

We proceed to show that $\ker P \subset \bar{B}$ with the aid of Lemma B.3, just as in the proof of Theorem 2.4.[3] Namely, it suffices to check that if $L : L^1(X, \mathscr{B}, \mathrm{m}) \to \mathbb{C}$ is an arbitrary continuous linear functional which vanishes on $\bar{B}$ (in particular on B), then it also vanishes on $\ker P$.

Given a functional $L : L^1(X, \mathscr{B}, \mathrm{m}) \to \mathbb{C}$, by the representation theorem of continuous linear functionals on L^1 (Theorem B.10), there exists $h \in L^\infty(X, \mathscr{B}, \mathrm{m})$ such that $Lf = \int_X fh \,\mathrm{dm}$. If L vanishes on B, we have $L(g - g\circ T) = 0$ for all $g \in L^1(X, \mathscr{B}, \mathrm{m})$. This means that $\int_X g\circ T h \,\mathrm{dm} = \int_X gh \,\mathrm{dm}$ for all $g \in L^1(X, \mathscr{B}, \mathrm{m})$. In particular, $\int_X h\circ T h \,\mathrm{dm} = \int_X h^2 \,\mathrm{dm} = \int_X (h\circ T)^2 \,\mathrm{dm}$, so $\int_X (h\circ T - h)^2 \,\mathrm{dm} = 0$, which means that h is almost invariant. Thus, h is $\mathscr{I}$-measurable. Using the last fact, we have $Lf = \int_X fh \,\mathrm{dm} = \int_X \mathrm{E}(f|\mathscr{I}) h \,\mathrm{dm}$. If $f \in \ker P$, then $\mathrm{E}(f|\mathscr{I}) = 0$, so $Lf = 0$. Hence, $\ker P \subset \bar{B}$ and $L^1(X, \mathscr{B}, \mathrm{m}) = I \oplus \bar{B}$.

In view of Lemma 3.28, equation (3.7) and Exercise 3.31, it remains to prove that (3.8) holds on $\bar{B}_\infty$, not just on B_∞. To this end, we apply Corollary 3.34. Let us denote

$$R(x; f) = \limsup_{n\to\infty} \left| \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x)) \right| \; .$$

[3]In the proof of von Neumann's theorem, showing $H = I \oplus \bar{B}$ was done in a slightly more direct fashion, because it was possible to take advantage of the inner product structure of H.

Since $R(x; f - f_\infty) = R(x; f)$ for all $f_\infty \in B_\infty$, Corollary 3.34 guarantees that, for all $\lambda > 0$, $f \in \bar{B}_\infty$ and $f_\infty \in B_\infty$, we have

$$\mathrm{m}\{x \in X \,:\, R(x; f) \geq \lambda \|f - f_\infty\|_1\} \leq \frac{1}{\lambda} \ .$$

Now, given such f and λ, pick $f_\infty \in B_\infty$ in such a way that $\|f - f_\infty\|_1 \leq \frac{1}{\lambda^2}$. Then $\{x \in X \,:\, R(x; f) \geq \frac{1}{\lambda}\} \subset \{x \in X \,:\, R(x; f) \geq \lambda \|f - f_\infty\|_1\}$, so that

$$\mathrm{m}\left\{x \in X \,:\, R(x; f) \geq \frac{1}{\lambda}\right\} \leq \frac{1}{\lambda} \ .$$

It is obvious that the set on the left side increases with λ and $\{x \in X \,:\, R(x; f) > 0\} = \cup_{\lambda > 0} \left\{x \in X \,:\, R(x; f) \geq \frac{1}{\lambda}\right\}$. Hence,

$$\mathrm{m}\left\{x \in X \,:\, R(x; f) > 0\right\} = 0 \ , \quad f \in \bar{B}_\infty = \bar{B} \ ,$$

which is what was to be shown. □

CHAPTER 4

Mixing

In this chapter we will introduce several so-called mixing properties of measure-preserving transformations, all of which are stronger than ergodicity.

1. Weak and strong mixing

To motivate the subject, it is very illuminating to present two more characterizations of ergodicity.

Theorem 4.1. *Let $(X, \mathscr{B})$ be a measurable space and $T : X \to X$ a measurable map. A measure $\mathrm{m} \in \mathcal{P}_T$ is ergodic if and only if one of the following equivalent conditions is satisfied:*

(1) For all $f, g \in L^2(X, \mathscr{B}, \mathrm{m})$,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \int_X f \circ T^i \, \bar{g} \, \mathrm{dm} = \int_X f \, \mathrm{dm} \int_X \bar{g} \, \mathrm{dm} \ . \tag{4.1}$$

(2) For all $A, B \in \mathscr{B}$,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \mathrm{m}(T^{-i}A \cap B) = \mathrm{m}(A)\, \mathrm{m}(B) \ ,$$

that is,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \big(\mathrm{m}(T^{-i}A \cap B) - \mathrm{m}(A)\, \mathrm{m}(B)\big) = 0 \ . \tag{4.2}$$

Given $A, B \in \mathscr{B}$ with $\mathrm{m}(B) > 0$, we have $\frac{1}{\mathrm{m}(B)} \mathrm{m}(T^{-i}A \cap B) = \mathrm{m}(T^{-i}A|B)$. Since $\mathrm{m}(A) = \mathrm{m}(T^{-n}A)$, condition (4.2) thus has the following dynamical interpretation: given that x is picked randomly in B according to the distribution m, the conditional probability that $T^n(x) \in A$ converges *in arithmetic mean* to the measure of A — whatever the sets A and B are.

Proof of Theorem 4.1. If (1) holds, we take $f = 1_A$ and $g = 1_B$. Since $1_A \circ T = 1_{T^{-1}A}$ and $1_{T^{-1}A} 1_B = 1_{T^{-1}A \cap B}$, (2) follows. Assuming (2), let A be an invariant set and $B = A^c$. Then $0 = \mathrm{m}(A \cap A^c) = \mathrm{m}(A)\,\mathrm{m}(A^c)$, which implies $\mathrm{m}(A) \in \{0, 1\}$, so m is ergodic. Finally, if m is ergodic and $f \in L^2(X, \mathscr{B}, \mathrm{m})$, the von Neumann ergodic theorem implies that $\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i = \int_X f \, \mathrm{dm}$ in the L^2 norm. This implies $\lim_{n\to\infty} \langle \frac{1}{n} \sum_{i=0}^{n-1} f \circ T^i, g \rangle = \langle \int_X f \, \mathrm{dm}, g \rangle = \int_X f \, \mathrm{dm} \int_X \bar{g} \, \mathrm{dm}$ for any $g \in L^2(X, \mathscr{B}, \mathrm{m})$. □

It is natural to ask whether under some conditions the series in (4.2) might converge absolutely, or whether the summand might even converge to zero without taking the arithmetic mean. This leads to the following definitions:

Definition 4.2. *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation. The measure* m *is called* **weakly mixing** *if*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \left|\mathrm{m}(T^{-i}A \cap B) - \mathrm{m}(A)\,\mathrm{m}(B)\right| = 0 \tag{4.3}$$

and **mixing** *(or* **strongly mixing***) if*

$$\lim_{n\to\infty} \mathrm{m}(T^{-n}A \cap B) = \mathrm{m}(A)\,\mathrm{m}(B) \tag{4.4}$$

for all $A, B \in \mathscr{B}$.

Condition (4.4) is equivalent to $\lim_{n\to\infty}(\mathrm{m}(T^{-n}A\cap B) - \mathrm{m}(T^{-n}A)\mathrm{m}(B)) = 0$, because $\mathrm{m}(A) = \mathrm{m}(T^{-n}A)$. In probability theory, two events $A, B \in \mathscr{B}$ are called independent, if $\mathrm{m}(A \cap B) = \mathrm{m}(A)\mathrm{m}(B)$. When $\mathrm{m}(B) > 0$, this is equivalent to $\mathrm{m}(A|B) = \mathrm{m}(A)$. From this point of view, mixing means that the events $T^{-n}A$ and B become asymptotically independent, which has the following dynamical interpretation: given that x is picked randomly in B according to the distribution m, the conditional probability that $T^n(x) \in A$ converges to the measure of A. There are two points two be made. First, the choice of the sets A and B does not matter for the conclusion. Second, the convergence is genuine, in that taking an arithmetic mean first is not required unlike in the case with ergodicity. While ergodicity means that trajectories starting from a given region B eventually (and frequently) visit any other region A of the state space X, mixing means that a fraction $\approx \mathrm{m}(B)$ of the trajectories are actually visiting B at any given instant after a sufficiently long time has passed.

Example 4.3. *Suppose we add some milk in a cup of coffee; say 1 part of milk in 4 parts of coffee, bringing the milk concentration in the cup to 20%. Our experience tells us that stirring the mixture with a spoon will eventually mix the milk in the coffee uniformly. That is, the concentration of milk in any region in the mixture approaches 20%. We may model the situation as follows: let X represent the liquid in the cup, $B \subset X$ the region initially holding all the milk, $T : X \to X$ the action of stirring the cup once with a spoon, and* m *the natural measure of volume (which is invariant, because the liquid is incompressible). Now, consider an arbitrary small region $A \subset X$ in the liquid. The milk molecules that will be found in the region A after stirring n times correspond to the initial region $T^{-n}A\cap B$, so the milk concentration in the region A at that time will be $\frac{\mathrm{m}(T^{-n}A\cap B)}{\mathrm{m}(A)}$, and our experience is that this number approaches $\mathrm{m}(B) = 20\%$, which is consistent with the mathematical definition of mixing.*

Let us already present characterizations of weak mixing and mixing in terms of functions analogous to the characterization (4.1) of ergodicity.

Theorem 4.4. *The measure* m *is weakly mixing if and only if*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \left| \int_X f \circ T^i\, \bar{g}\,\mathrm{dm} - \int_X f\,\mathrm{dm} \int_X \bar{g}\,\mathrm{dm} \right| = 0 \ , \tag{4.5}$$

and mixing if and only if

$$\lim_{n\to\infty} \int_X f \circ T^n\, \bar{g}\,\mathrm{dm} = \int_X f\,\mathrm{dm} \int_X \bar{g}\,\mathrm{dm} \ , \tag{4.6}$$

for all $f, g \in L^2(X, \mathscr{B}, \mathrm{m})$.

Exercise 4.5. *Prove Theorem 4.4.*

[Hint: Simple functions. The two parts can be proved simultaneously with the aid of Lemma 4.9 below.]

Exercise 4.6. *Recall from Exercise 2.47 that the angle doubling map* $T : x \mapsto 2x \pmod 1$ *is ergodic. Show that it is mixing.*

[Hint: Fourier. Consider $f \in L^2$ *and the functions* $g(x) = e^{2\pi i m x}$, $m \in \mathbb{Z}$, *which span* L^2*.]*

Given the discussion above, the next result is obvious.

Lemma 4.7. *Invariant measures have the following hierarchy:*

$$\text{(strong) mixing} \quad \Longrightarrow \quad \text{weak mixing} \quad \Longrightarrow \quad \text{ergodicity.}$$

The hierarchy in Lemma 4.7 is strict, in that none of the implications can be reversed. For example, ergodicity does not imply weak mixing:

Exercise 4.8. *Recall from Exercise 2.26 that a rotation of the circle by an irrational angle is ergodic. Show that it is not weakly mixing.*

[Hint: $f(x) = g(x) = e^{2\pi i x}$*.]*

Weak mixing, though, is not a whole lot weaker than mixing. The following well-known lemma will help convince the reader of their near equivalence.

A set $\mathcal{N} \subset \mathbb{N}$ is said to have **zero density** if

$$\lim_{n\to\infty} \frac{1}{n} |\mathcal{N} \cap \{1, \ldots, n\}| = 0 ,$$

where $|\cdot|$ denotes the cardinality of a set.

Lemma 4.9 (Koopman–von Neumann)**.** *A bounded sequence of numbers* $a_n \in \mathbb{C}$, $n \geq 0$, *satisfies*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} |a_i| = 0$$

if and only if there exists an index set $\mathcal{N} \subset \mathbb{N}$ *of zero density such that*

$$\lim_{n\to\infty,\, n\notin\mathcal{N}} a_n = 0 .$$

The proof can be found, for instance, in Walters [17]. As an immediate application we observe that a measure-preserving transformation is weakly mixing if and only if the mixing condition (4.4) holds after possibly excluding a subsequence of zero density. This result is so useful for proofs that we formulate it as a theorem.

Theorem 4.10 (Weak mixing is nearly mixing)**.** *The measure* m *is weakly mixing if and only if there exists an index set* $\mathcal{N} \subset \mathbb{N}$ *of zero density such that*

$$\lim_{n\to\infty,\, n\notin\mathcal{N}} \mathrm{m}(T^{-n}A \cap B) = \mathrm{m}(A)\,\mathrm{m}(B)$$

for all $A, B \in \mathscr{B}$*.*

In fact, it is not straightforward to give examples of weakly mixing transformations that are not mixing, but they do exist; see, e.g., Parry [9] for an example due to Kakutani and von Neumann, or Sarig [15] for an example due to Chacon.

Ergodicity and weak mixing are related through product transformations. Recall from Exercise 2.16 that if $(X, \mathscr{B}, \mathrm{m}, T)$ is an mpt, then the product mpt $(X \times X, \mathscr{B} \times \mathscr{B}, \mathrm{m} \times \mathrm{m}, T \times T)$ is determined by $(T \times T)(x, y) = (T(x), T(y))$.

Theorem 4.11. *The following conditions are equivalent:*

(1) T is weakly mixing.
(2) $T \times T$ is ergodic.
(3) $T \times T$ is weakly mixing.

The reader should beware of the corollary

$$T \text{ is ergodic} \quad \not\Longrightarrow \quad T \times T \text{ is ergodic} \quad \Longrightarrow \quad T \text{ is ergodic.}$$

Exercise 4.12. *It was shown in Exercise 2.26 that the Lebesgue measure* m *is ergodic for the map $T : \mathbb{S}^1 \to \mathbb{S}^1 : x \mapsto x + \alpha \pmod 1$ if α is irrational. Prove that the Lebesgue measure $\mathrm{m} \times \mathrm{m}$ on the two dimensional torus $\mathbb{T}^2 = \mathbb{S}^1 \times \mathbb{S}^1$ is not ergodic for $T \times T : \mathbb{T}^2 \to \mathbb{T}^2 : (x, y) \mapsto (x + \alpha, y + \alpha) \pmod 1$.*

[Remark: It is not part of the exercise, but it can be shown similarly to Exercise 2.26 that $(x, y) \mapsto (x + \alpha, y + \beta) \pmod 1$ is uniquely ergodic if and only if $\{\alpha, \beta, 1\}$ are rationally independent, which means that $k_1\alpha + k_2\beta + k_3 = 0$ for $k_i \in \mathbb{Q}$ only when $k_1 = k_2 = k_3 = 0$.]

On the other hand, strong mixing of T and $T \times T$ are equivalent:

Exercise 4.13. *Show that*

$$T \text{ is mixing} \quad \Longleftrightarrow \quad T \times T \text{ is mixing.}$$

[Hint: Mimic the proof of Theorem 4.11 below.]

Example 4.14. *The map $T \times T : \mathbb{T}^2 \to \mathbb{T}^2 : (x, y) \mapsto (2x, 2y) \pmod 1$ is mixing, because $T : \mathbb{S}^1 \to \mathbb{S}^1 : (x, y) \mapsto 2x \pmod 1$ is mixing by Exercise 4.6.*

Proof of Theorem 4.11. We follow Walters [17]. The basic ideas used in the proof are that (a) before considering general sets in $\mathscr{B} \times \mathscr{B}$, we consider the rectangles $A \times B \in \mathscr{B} \times \mathscr{B}$ which generate the product sigma-algebra, and that (b) we view convergence in arithmetic mean as ordinary convergence excluding a subsequence of zero density as in Lemma 4.9.

Since (3) $\Rightarrow$ (2), it suffices to prove the implications (2) $\Rightarrow$ (1) $\Rightarrow$ (3).

(1) $\Rightarrow$ (3): Let $A_i, B_i \in \mathscr{B}$ $(i = 1, 2)$. By Theorem 4.10, there exist index sets $\mathcal{N}_i \subset \mathbb{N}$ $(i = 1, 2)$ of zero density such that

$$\lim_{n \to \infty, n \notin \mathcal{N}_i} \mathrm{m}(T^{-n} A_i \cap B_i) = \mathrm{m}(A_i)\, \mathrm{m}(B_i) \ .$$

Since $(T \times T)^{-n}(A_1 \times A_2) \cap (B_1 \times B_2) = (T^{-n} A_1 \cap B_1) \times (T^{-n} A_2 \cap B_2)$, we see that

$$\mathrm{m} \times \mathrm{m}((T \times T)^{-n}(A_1 \times A_2) \cap (B_1 \times B_2)) = \mathrm{m}(T^{-n} A_1 \cap B_1)\, \mathrm{m}(T^{-n} A_2 \cap B_2) \ .$$

The right side converges to $\mathrm{m}(A_1)\mathrm{m}(B_1)\mathrm{m}(A_2)\mathrm{m}(B_2) = \mathrm{m}\times\mathrm{m}(A_1\times A_2)\,\mathrm{m}\times\mathrm{m}(B_1\times B_2)$, as $n\to\infty$ with the constraint $n\notin\mathcal{N}_1\cup\mathcal{N}_2$. Since $\mathcal{N}_1\cup\mathcal{N}_2$ has zero density, Lemma 4.9 yields

$$\lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\big|\mathrm{m}\times\mathrm{m}((T\times T)^{-i}(A_1\times A_2)\cap(B_1\times B_2)) - \mathrm{m}\times\mathrm{m}(A_1\times A_2)\,\mathrm{m}\times\mathrm{m}(B_1\times B_2)\big| = 0\ .$$

Fact: It follows that the same result holds for two arbitrary sets $\mathbf{A},\mathbf{B}\in\mathscr{B}\times\mathscr{B}$ in place of the rectangles $A_1\times A_2$ and $B_1\times B_2$.[1]

(2) ⇒ (1): Let $A,B\in\mathscr{B}$. The ergodicity of $T\times T$ implies

$$\begin{aligned}\lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\mathrm{m}(T^{-i}A\cap B) &= \lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\mathrm{m}\times\mathrm{m}((T\times T)^{-i}(A\times X)\cap(B\times X))\\ &= \mathrm{m}\times\mathrm{m}(A\times X)\,\mathrm{m}\times\mathrm{m}(B\times X) = \mathrm{m}(A)\,\mathrm{m}(B)\end{aligned}$$

as well as

$$\begin{aligned}\lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}(\mathrm{m}(T^{-i}A\cap B))^2 &= \lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\mathrm{m}\times\mathrm{m}((T\times T)^{-i}(A\times A)\cap(B\times B))\\ &= \mathrm{m}\times\mathrm{m}(A\times A)\,\mathrm{m}\times\mathrm{m}(B\times B) = (\mathrm{m}(A)\,\mathrm{m}(B))^2\ .\end{aligned}$$

(In particular, we just showed that T is ergodic.) Combining the two, we arrive at

$$\lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}(\mathrm{m}(T^{-i}A\cap B)-\mathrm{m}(A)\,\mathrm{m}(B))^2 = 0\ .$$

By Lemma 4.9, there exists an index set $\mathcal{N}\subset\mathbb{N}$ of zero density such that

$$\lim_{n\to\infty,\,n\notin\mathcal{N}}(\mathrm{m}(T^{-n}A\cap B)-\mathrm{m}(A)\,\mathrm{m}(B))^2 = 0\ ,$$

or, equivalently,

$$\lim_{n\to\infty,\,n\notin\mathcal{N}}\mathrm{m}(T^{-n}A\cap B) = \mathrm{m}(A)\,\mathrm{m}(B)\ ,$$

which shows that T is weakly mixing. □

We will soon encounter even stronger notions of mixing, namely the Kolmogorov property and exactness, followed by the Bernoulli and Markov properties. Before introducing these concepts is possible, it is necessary to dwell a little on sub-sigma-algebras and conditional expectations, which is what we do next.

[1]This can be proved, for instance, using the π–λ lemma: Let $\mathscr{R} = \{E_1\times E_2 : E_1,E_2\in\mathscr{B}\}$ be the collection of rectangles. Then $\mathscr{R}$ is a π-system, which generates the product sigma-algebra, i.e., $\sigma(\mathscr{R}) = \mathscr{B}\times\mathscr{B}$. Here $\sigma(\mathscr{R})$ denotes the smallest sigma-algebra containing $\mathscr{R}$. Given $A_1\times A_2\in\mathscr{R}$, denote by $\mathscr{L}$ the collection of all those $\mathbf{B}\in\mathscr{B}\times\mathscr{B}$ for which $\lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\big|\mathrm{m}\times\mathrm{m}((T\times T)^{-i}(A_1\times A_2)\cap\mathbf{B}) - \mathrm{m}\times\mathrm{m}(A_1\times A_2)\,\mathrm{m}\times\mathrm{m}(\mathbf{B})\big| = 0$. Then $\mathscr{L}$ is a λ-system, and we have shown that $\mathscr{R}\subset\mathscr{L}$. By the π–λ lemma, $\mathscr{B}\times\mathscr{B} = \sigma(\mathscr{R})\subset\mathscr{L}$, so convergence actually takes place for any rectangle $A_1\times A_2$ and any $\mathbf{B}\in\mathscr{B}\times\mathscr{B}$. Similarly, $A_1\times A_2$ can then be replaced with any $\mathbf{A}\in\mathscr{B}\times\mathscr{B}$. We skip the details.

2. Preliminaries: Convergence of conditional expectations along nested sequences of sub-sigma-algebras

2.1. Sequences of sub-sigma-algebras. Recall that the intersection of an arbitrary family of sigma-algebras is a sigma-algebra. Given a decreasing sequence of sub-sigma-algebras $\mathscr{B} \supset \mathscr{A}_1 \supset \mathscr{A}_2 \supset \cdots$, the notation

$$\mathscr{A}_n \downarrow \mathscr{A}$$

means that $\mathscr{A} = \cap_{n\geq 1}\mathscr{A}_n$. Of course, $\mathscr{A} \subset \mathscr{A}_n \subset \mathscr{B}$. Likewise,

$$\mathscr{A}_n \downarrow \mathscr{A} \pmod{\mathrm{m}}$$

means that $\mathscr{A}$ is a sub-sigma-algebra satisfying $\mathscr{A} = \cap_{n\geq 1}\mathscr{A}_n \pmod{\mathrm{m}}$.

Given an increasing sequence of sub-sigma-algebras $\mathscr{A}_1 \subset \mathscr{A}_2 \subset \cdots \subset \mathscr{B}$, the union $\cup_{n\geq 1}\mathscr{A}_n$ is not necessarily an algebra, let alone a sigma-algebra; the notation

$$\mathscr{A}_n \uparrow \mathscr{A}$$

means that $\mathscr{A}$ is the sigma-algebra $\sigma(\cup_{n\geq 1}\mathscr{A}_n)$ generated by the union, i.e., the smallest sigma-algebra containing each $\mathscr{A}_n$. Obviously $\mathscr{A}_n \subset \mathscr{A} \subset \mathscr{B}$. Similarly,

$$\mathscr{A}_n \uparrow \mathscr{A} \pmod{\mathrm{m}}$$

means that $\mathscr{A}$ is a sub-sigma-algebra satisfying $\mathscr{A} = \sigma(\cup_{n\geq 1}\mathscr{A}_n) \pmod{\mathrm{m}}$.

Example 4.15. *Recall Example 3.9, where we constructed a sub-sequence of sigma-algebras $\mathscr{A}_n$, $n \geq 1$, with the aid of the partitions of the interval $[0,1)$ into subintervals $[2^{-n}(i-1), 2^{-n}i)$, $1 \leq i \leq 2^n$. Since increasing n leads to a finer partition, we have $\mathscr{A}_1 \subset \mathscr{A}_2 \subset \cdots$. It is tempting to conclude that the piecewise constant approximation $\mathrm{E}(f|\mathscr{A}_n)$ approaches $f \in L^1(X)$ as $n \to \infty$. In fact, it does: It is not hard to check that $\mathscr{A}_n \uparrow \mathscr{B} \pmod{\mathrm{m}}$, from which $\lim_{n\to\infty} \mathrm{E}(f|\mathscr{A}_n) = \mathrm{E}(f|\mathscr{B}) = f$ follows by the theorem below.*

Theorem 4.16 (Lévy–Doob upward/downward convergence theorem)**.** *Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space and $\mathscr{A}, \mathscr{A}_n$, $n \geq 1$, sub-sigma-algebras. Suppose that either $\mathscr{A}_n \uparrow \mathscr{A}$ or $\mathscr{A}_n \downarrow \mathscr{A}$. Then*

$$\lim_{n\to\infty} \mathrm{E}(f|\mathscr{A}_n) = \mathrm{E}(f|\mathscr{A}) \quad \text{almost surely and in } L^p,$$

for all $f \in L^p(X, \mathscr{B}, \mathrm{m})$ and for all $p \in [1, \infty)$.

The result is also known as a martingale convergence theorem, because the sequence $\mathrm{E}(f|\mathscr{A}_n)$ is a martingale if $\mathscr{A}_n \uparrow \mathscr{A}$ and a reverse (or backward) martingale if $\mathscr{A}_n \downarrow \mathscr{A}$. Many books, such as [**9**, **4**, **18**], present just the L^1 version of the Lévy–Doob upward/downward theorem theorem. The L^p version above, together with its proof, can be found, e.g., in James Norris' lecture notes [**8**]. We omit the proof, but it is worthwhile to point out that — analogously to Exercise 3.19 — the L^1 version implies the L^p version:

Exercise 4.17. *Let either $\mathscr{A}_n \uparrow \mathscr{A}$ or $\mathscr{A}_n \downarrow \mathscr{A}$. Suppose there exists $p \in [1, \infty)$ such that*

$$\lim_{n\to\infty} \|\mathrm{E}(f|\mathscr{A}_n) - \mathrm{E}(f|\mathscr{A})\|_p = 0 \ , \quad f \in L^p(X, \mathscr{B}, \mathrm{m}) \ .$$

Show that the same statement then holds in fact for all $p \in [1, \infty)$.

[Hint: Truncate f as in Exercise 3.19.]

For an increasing sequence $\mathscr{A}_1 \subset \mathscr{A}_2 \subset \cdots$ of sub-sigma-algebras it is obviously true that $L^p(X, \mathscr{A}_1, \mathrm{m}) \subset L^p(X, \mathscr{A}_2, \mathrm{m}) \subset \cdots$. As an immediate application of Theorem 4.16, we also obtain

Corollary 4.18. *Suppose $\mathscr{A}_n \uparrow \mathscr{A} \subset \mathscr{B}$. Then $\cup_{n=1}^{\infty} L^p(X, \mathscr{A}_n, \mathrm{m})$ is a dense subspace of $L^p(X, \mathscr{A}, \mathrm{m})$.*

The sub-sigma-algebras $T^{-n}\mathscr{A} = \{T^{-n}A : A \in \mathscr{A}\}$ introduced in Section 2 arise naturally in ergodic theory. For example, if $f : X \to \mathbb{C}$ is a $\mathscr{A}$-measurable function, then $f \circ T^n$ is $T^{-n}\mathscr{A}$-measurable. In general, there is no inclusion relation between $T^{-n}\mathscr{A}$ and $T^{-m}\mathscr{A}$ for $n \neq m$. If $T : X \to X$ has a measurable inverse, the sub-sigma-algebras $T^n\mathscr{A}$, $n \geq 1$, are defined similarly: $T^n\mathscr{A} = \{T^nA : A \in \mathscr{B}\}$. In the following sections we study situations in which the sub-sigma-algebras $T^{-n}\mathscr{A}$ or $T^n\mathscr{A}$ form nested sequences.

3. Kolmogorov automorphisms and exact endomorphisms

We are now ready to introduce two closely related mixing properties, one concerning non-invertible and the other invertible transformations.

A measure-preserving transformation $(X, \mathscr{B}, \mathrm{m}, T)$ is said to be invertible, if $T : X \to X$ is an invertible map and the inverse T^{-1} is measurable, i.e., $(T^{-1})^{-1}A = TA \in \mathscr{B}$ for all $A \in \mathscr{B}$. Then $(X, \mathscr{B}, \mathrm{m}, T^{-1})$ is an mpt. Measure-preserving transformations are also known as **endomorphisms**[2] and invertible mpts as **automorphisms**.

Definition 4.19. *A probability-preserving transformation $(X, \mathscr{B}, \mathrm{m}, T)$ is* **exact**, *or an* **exact endomorphism**, *if*

$$T^{-n}\mathscr{B} \downarrow \mathscr{N} \pmod{\mathrm{m}} .$$

Definition 4.20. *An invertible probability-preserving transformation $(X, \mathscr{B}, \mathrm{m}, T)$ is a* **Kolmogorov automorphism**, *or* **K-automorphism**, *if there exists a sub-sigma-algebra $\mathscr{A} \subset \mathscr{B}$ such that*

$$T^{-1}\mathscr{A} \subset \mathscr{A} \; ; \quad T^n\mathscr{A} \uparrow \mathscr{B} \pmod{\mathrm{m}} \quad \text{and} \quad T^{-n}\mathscr{A} \downarrow \mathscr{N} \pmod{\mathrm{m}} .$$

Note that, in the case of the full sigma-algebra, $T^{-1}\mathscr{B} \subset \mathscr{B}$ is automatic, so it is not explicitly stated as a condition for exactness. If T is invertible, then $T^{-1}\mathscr{B} = \mathscr{B}$, so an invertible transformation in never exact, save for the trivial case $\mathscr{B} = \mathscr{N} \pmod{\mathrm{m}}$. In some sense the exactness condition $T^{-n}\mathscr{B} \downarrow \mathscr{N} \pmod{\mathrm{m}}$ means that T is far from being invertible. Exactness and the Kolmogorov property are analogous concepts for non-invertible and invertible transformations.

Perhaps at first it seems mysterious from the definitions why the properties introduced above are regarded as mixing properties.

Theorem 4.21. *Exact endomorphisms and Kolmogorov automorphisms are mixing.*

In standard ergodic theory books the theorem is proved using spectral theory. We present a conceptually simple alternative proof, because it makes it crystal clear how the conditions on the sigma-algebras imply mixing in each case.

[2]Some authors reserve the term endomorphism to the case where T is onto.

PROOF OF THEOREM 4.21. We prove mixing by showing (4.6) for all $f, g \in L^2(X, \mathscr{B}, \mathrm{m})$. By subtracting the right side of (4.6) from the left one, it is clearly sufficient to prove that

$$\lim_{n\to\infty} \int_X f \circ T^n \, g \, \mathrm{dm} = 0$$

holds for all $f, g \in L^2(X, \mathscr{B}, \mathrm{m})$ such that $\int_X g \, \mathrm{dm} = 0$.

Exact case: Let $f, g \in L^2(X, \mathscr{B}, \mathrm{m})$ with $\int_X g \, \mathrm{dm} = 0$. Then $f \circ T^n$ is $T^{-n}\mathscr{B}$-measurable. By the rules of conditional expectation,

$$\int_X f \circ T^n \, g \, \mathrm{dm} = \int_X f \circ T^n \, \mathrm{E}(g|T^{-n}\mathscr{B}) \, \mathrm{dm} \; .$$

Thus, the Cauchy–Schwarz inequality implies

$$\left| \int_X f \circ T^n \, g \, \mathrm{dm} \right| \le \|f \circ T^n\|_2 \|\mathrm{E}(g|T^{-n}\mathscr{B})\|_2 = \|f\|_2 \|\mathrm{E}(g|T^{-n}\mathscr{B})\|_2 .$$

Since $T^{-n}\mathscr{B} \downarrow \mathscr{N}$ and $g \in L^2(X, \mathscr{B}, \mathrm{m})$, the Lévy–Doob upward/downward convergence theorem implies that, in the L^2 norm,

$$\mathrm{E}(g|T^{-n}\mathscr{B}) \to \mathrm{E}(g|\mathscr{N}) = \int_X g \, \mathrm{dm} = 0 \; .$$

This completes the proof in the exact case.

Kolmogorov case: Let $k \ge 0$, $f \in L^2(X, T^k\mathscr{A}, \mathrm{m})$ and $g \in L^2(X, \mathscr{B}, \mathrm{m})$ with $\int_X g \, \mathrm{dm} = 0$. Then $f \circ T^n$ is $T^{k-n}\mathscr{A}$-measurable, so the assumption $T^{-n}\mathscr{A} \downarrow \mathscr{N}$ implies

$$\int_X f \circ T^n \, g \, \mathrm{dm} = \int_X f \circ T^n \, \mathrm{E}(g|T^{k-n}\mathscr{A}) \, \mathrm{dm} \to 0$$

precisely as above. Note the restriction $f \in L^2(X, T^k\mathscr{A}, \mathrm{m})$, however, which is not sufficient for mixing. To prove that we have the corresponding result for any $f \in L^2(X, \mathscr{B}, \mathrm{m})$, we use the condition $T^k\mathscr{A} \uparrow \mathscr{B}$. If $f \in L^2(X, \mathscr{B}, \mathrm{m})$, then $\mathrm{E}(f|T^k\mathscr{A}) \in L^2(X, T^k\mathscr{A}, \mathrm{m})$. Hence,

$$\lim_{n\to\infty} \int_X \mathrm{E}(f|T^k\mathscr{A}) \circ T^n \, g \, \mathrm{dm} = 0$$

by what we have proved so far. Since $T^k\mathscr{A} \uparrow \mathscr{B}$, the Lévy–Doob upward/downward convergence theorem implies that, in the L^2 norm,

$$\mathrm{E}(f|T^k\mathscr{A}) \to \mathrm{E}(f|\mathscr{B}) = f \; .$$

By taking k sufficiently large, we see that

$$\begin{aligned} \limsup_{n\to\infty} \left| \int_X f \circ T^n \, g \, \mathrm{dm} \right| &\le \limsup_{n\to\infty} \left| \int_X (f - \mathrm{E}(f|T^k\mathscr{A})) \circ T^n \, g \, \mathrm{dm} \right| \\ &\le \|f - \mathrm{E}(f|T^k\mathscr{A})\|_2 \|g\|_2 \end{aligned}$$

is arbitrarily small. The proof is now complete.[3] □

[3]We ignored the minor detail that the sequences of sigma-algebras converge to $\mathscr{N}$ and $\mathscr{B}$ only (mod m). For example, $T^{-n}\mathscr{B} \downarrow \mathscr{N}$ (mod m) means that $T^{-n}\mathscr{B} \downarrow \cap_{n\ge1} T^{-n}\mathscr{B}$ where $\cap_{n\ge1} T^{-n}\mathscr{B} = \mathscr{N}$ (mod m). This is no cause of worry, because if two sub-sigma-algebras satisfy $\mathscr{A}_1 = \mathscr{A}_2$ (mod m), then the conditional expectations satisfy $E(f|\mathscr{A}_1) = E(f|\mathscr{A}_2)$ almost everywhere; see Lemma 3.12.

We learn from the proof that the conditions for exactness and the Kolmogorov property are actually tailor-made for mixing.

We will see examples of exact endomorphisms and Kolmogorov automorphisms next.

4. Bernoulli and Markov shifts

In this section we introduce the so-called Markov and Bernoulli shifts, which form important classes of probability-preserving transformations.

For each integer $k \geq 2$, we define $S_k = \{1, \dots, k\}$. We call S_k **the set (or space) of k symbols**. Thus, the cartesian products

$$S_k^{\mathbb{N}} = \prod_{i \in \mathbb{N}} S_k \quad \text{and} \quad S_k^{\mathbb{Z}} = \prod_{i \in \mathbb{Z}} S_k$$

consist, respectively, of one-sided sequences $x = (x_i)_{i \in \mathbb{N}}$ and two-sided sequences $x = (x_i)_{i \in \mathbb{Z}}$ of symbols $x_i \in S_k$. In explicit expressions involving two-sided sequences, it is necessary to single out the zeroth symbol to avoid confusion; we single it out with a star, which clarifies the distinction between the two-sided sequences $(\dots, x_{-1}, \overset{*}{x}_0, x_1, \dots)$ and $(\dots, x_0, \overset{*}{x}_1, x_2, \dots)$, etc.

A natural map acting on sequences of each type is the **shift map** defined by

$$(\tau x)_i = x_{i+1}$$

for each $i \in \mathbb{N}$ or $\mathbb{Z}$. More precisely, we define the one-sided shift (of k symbols)

$$\tau : S_k^{\mathbb{N}} \to S_k^{\mathbb{N}} : (x_0, x_1, x_2 \dots) \mapsto (x_1, x_2, \dots)$$

and the two-sided shift (of k symbols)

$$\tilde{\tau} : S_k^{\mathbb{Z}} \to S_k^{\mathbb{Z}} : (\dots, x_{-1}, \overset{*}{x}_0, x_1, \dots) \mapsto (\dots, x_0, \overset{*}{x}_1, x_2, \dots) \ .$$

Thus, τ and $\tilde{\tau}$ simply shift each symbol one place to the left, save for τ which also annihilates the zeroth symbol of the original sequence. The two-sided shift $\tilde{\tau}$ is clearly invertible and the one-sided shift τ is onto but not one-to-one.

Let us construct a probability-preserving transformations out of τ and $\tilde{\tau}$. This calls for a sigma-algebra and an invariant probability measure. To that end, let Σ_k denote the sigma-algebra on S_k which comprises all subset of Σ_k. Then we can endow $S_k^{\mathbb{N}}$ and $S_k^{\mathbb{Z}}$, respectively, with the product sigma-algebras

$$\Sigma_k^{\mathbb{N}} \quad \text{and} \quad \Sigma_k^{\mathbb{Z}} \ .$$

Let us recall what these are. Given any finite set of symbols $s_0, \dots, s_m \in S_k$, $m \geq 0$, the set

$$[s_0, \dots, s_m] = \{x \in S_k^{\mathbb{N}} \ : \ x_i = s_i \ , \ 0 \leq i \leq m\}$$

is called a **one-sided cylinder set**. The product sigma algebra $\Sigma_k^{\mathbb{N}}$ is the sigma-algebra generated by all one-sided cylinder sets. Likewise, given any $s_l, \dots, s_m \in S_k$, $l \leq 0 \leq m$, the set

$$[s_l, \dots, \overset{*}{s}_0, \dots, s_m] = \{x \in S_k^{\mathbb{Z}} \ : \ x_i = s_i \ , \ l \leq i \leq m\}$$

is called a **two-sided cylinder set**, and the sigma-algebra $\Sigma_k^{\mathbb{Z}}$ is the one generated by all two-sided cylinder sets. Finally, note that a probability measure λ on the finite measurable space (S_k, Σ_k) is simply a probability vector $(\lambda_1, \dots, \lambda_k)$ where $\lambda_i = \lambda(i)$. Given such a measure, we can consider the product measures

$$\lambda^{\mathbb{N}} \text{ on } (S_k^{\mathbb{N}}, \Sigma_k^{\mathbb{N}}) \quad \text{and} \quad \lambda^{\mathbb{Z}} \text{ on } (S_k^{\mathbb{Z}}, \Sigma_k^{\mathbb{Z}}).$$

Here $\lambda^{\mathbb{N}}$ is the unique probability measure satisfying

$$\lambda^{\mathbb{N}}([s_0, \dots, s_m]) = \lambda_{s_0} \cdots \lambda_{s_m}$$

for all cylinder sets. The case of $\lambda^{\mathbb{Z}}$ is similar.

Exercise 4.22. *Show that $\lambda^{\mathbb{N}}$ and $\lambda^{\mathbb{Z}}$ are invariant measures with respect to τ and $\tilde{\tau}$, respectively, regardless of the choice of λ.*

[Hint: Work with cylinder sets.]

Definition 4.23. *Let $k \geq 2$ and let $\lambda = (\lambda_1, \dots, \lambda_k)$ be a probability measure with $\lambda_i > 0$, $1 \leq i \leq k$.*

(1) The ppt $(S_k^{\mathbb{N}}, \Sigma_k^{\mathbb{N}}, \lambda^{\mathbb{N}}, \tau)$ is called a **one-sided Bernoulli shift** *of k symbols, or the one-sided $(\lambda_1, \dots, \lambda_k)$-shift.*

(2) The ppt $(S_k^{\mathbb{Z}}, \Sigma_k^{\mathbb{Z}}, \lambda^{\mathbb{Z}}, \tilde{\tau})$ is called a **two-sided Bernoulli shift** *of k symbols, or the two-sided $(\lambda_1, \dots, \lambda_k)$-shift.*

The terms Bernoulli endomorphism and Bernoulli automorphism are also used.

The case $k = 1$ is excluded as trivial. The strict positivity of λ is a natural assumption; if, say, $\lambda_k = 0$, the set S_k could be reduced to S_{k-1} without any loss of generality.

Before proceeding, let us pause to give the Bernoulli shifts a probabilistic interpretation, which also explains the name. Suppose we have a dice with k possible outcomes, namely $1, \dots, k$, such that the number i always comes out with probability λ_i, completely independently of any previous throws. We throw the dice infinitely many times, and let ξ_i denote the random outcome of the ith throw. Then $(\xi_i)_{i \geq 0}$ is a sequence of random variables, which are independent and identically distributed (i.i.d. for short), each having the distribution λ. Such sequences of repeated independent trials are called **Bernoulli processes** in probability theory. In a realized experiment, every ξ_i assumes a concrete value in S_k; by independence of the throws, we have

$$\text{Probability}(\xi_n = s_0, \dots, \xi_{m+n} = s_m) = \lambda_{s_0} \cdots \lambda_{s_m} = \lambda^{\mathbb{N}}([s_0, \dots, s_m])$$

for any $m, n \geq 0$ and any symbols $s_0, \dots, s_m \in S_k$. The random sequence $\xi = (\xi_i)_{i \geq 0}$ is called **stationary**, because the right side is independent of n. Another way of writing this is

$$\text{Probability}(\tau^n \xi \in [s_0, \dots, s_m]) = \text{Probability}(\xi \in [s_0, \dots, s_m]) , \tag{4.7}$$

from which it is seen that the stationarity of ξ is equivalent to the invariance of the measure $\lambda^{\mathbb{N}}$ with respect to the shift τ. The independence of the throws is equivalent to the invariant measure being the *product* measure $\lambda^{\mathbb{N}}$. There is a one-to-one correspondence of one-sided Bernoulli shifts in ergodic theory and sequences of i.i.d. random variables in probability theory. In particular, the one-sided $(\frac{1}{2}, \frac{1}{2})$-shift corresponds to tossing a fair coin infinitely many times.

Independence is a rather restrictive assumption. We therefore discuss the more general setting of Markov chains with finite state space. These correspond to so-called Markov shifts in ergodic theory, of which Bernoulli shifts are special cases. Again, we generate a random sequence $(\xi_i)_{i\geq 0}$ of symbols in Σ_k, but this time the dice has a memory of the previous throw: the outcome ξ_i of the ith throw, $i \geq 1$, depends on the preceding outcome ξ_{i-1} but no outcome $\xi_0, \ldots, \xi_{i-2}$ before that. More precisely, we first assume that the random outcome ξ_0 has some distribution $\lambda = (\lambda_1, \ldots, \lambda_k)$, that is, Probability$(\xi_0 = r) = \lambda_r$; for the remaining throws we assume that

$$\text{Probability}(\xi_i = s_i \,|\, \xi_0 = s_0, \ldots, \xi_{i-1} = s_{i-1}) = \text{Probability}(\xi_i = s_i \,|\, \xi_{i-1} = s_{i-1}) = p_{s_{i-1},s_i}$$

holds for all $i \geq 1$ and all symbols. Here $p_{r,s}$ is the probability that the outcome $r \in \Sigma_k$ is followed by the outcome $s \in \Sigma_k$. In particular,

$$p_{r,s} \geq 0 \quad \text{and} \quad \sum_{1\leq s\leq k} p_{r,s} = 1 \ .$$

Thus, the probabilities $p_{r,s}$ can be viewed as the entries of the matrix

$$P = (p_{r,s})_{1\leq r,s\leq k}$$

where each row $(p_{r,1}, \ldots, p_{r,k})$ is a probability vector. Such a random sequence $(\xi_i)_{i\geq 0}$ is called a **Markov chain** with state space Σ_k, transition matrix P and initial distribution λ. We note that the Bernoulli process with i.i.d. outcomes is recovered as the special case in which every row of P equals λ and λ is strictly positive (i.e., $p_{r,s} = \lambda_s > 0$). We will next see how the Markov case also yields a probability-preserving transformation.

If the distribution of ξ_0 is $\lambda = (\lambda_1, \ldots, \lambda_k)$, what is the distribution of ξ_1? Since ξ_0 and ξ_1 are generally not i.i.d., the probability that $\xi_1 = s$ is

$$\sum_{1\leq r\leq k} \lambda_r \, \text{Probability}(\xi_1 = s | \xi_0 = r) = \sum_{1\leq r\leq k} \lambda_r \, p_{r,s} \ .$$

Hence, the distribution of ξ_1 is given by the probability vector λP. Similarly, it can be checked that, if the distribution of ξ_0 is λ, then the distribution of ξ_i is λP^i. Thus, the random variables ξ_i are identically distributed (but not necessarily independent) if and only if

$$\lambda P = \lambda \ . \tag{4.8}$$

It is said that λ is a **stationary distribution**, if (4.8) holds.

Exercise 4.24. *Show that the Markov chain $(\xi_i)_{i\geq 0}$ is stationary if and only if the initial distribution λ is stationary. In other words,* (4.7) *holds for all n and all cylinder sets $[s_0, \ldots, s_m]$ if and only if* (4.8) *holds.*

There exists a well-defined probability measure m on $(S_k^{\mathbb{N}}, \Sigma_k^{\mathbb{N}})$ induced by λ and P, called the distribution of the Markov chain $\xi = (\xi_i)_{i\geq 0}$, such that

$$\text{Probability}(\xi \in A) = \mathrm{m}(A) \ , \quad A \in \Sigma_k^{\mathbb{N}} \ .$$

This measure is uniquely determined by its value on the cylinder sets $A = [s_0, \ldots, s_m]$. Thus, m is the unique measure which satisfies

$$\mathrm{m}([s_0, \ldots, s_m]) = \lambda_{s_0} \, p_{s_0,s_1} \cdots p_{s_{m-1},s_m}$$

Exercise 4.25. *Show that the Markov chain $(\xi_i)_{i\geq 0}$ is stationary if and only if the measure* m *induced by λ and P is invariant with respect to the one-sided shift map $\tau : S_k^{\mathbb{N}} \to S_k^{\mathbb{N}}$.*

The measure m can be extended to a probability measure $\widetilde{\mathrm{m}}$ on $(S_k^{\mathbb{Z}}, \Sigma_k^{\mathbb{Z}})$, which is invariant with respect to the two-sided shift $\tilde{\tau} : S_k^{\mathbb{Z}} \to S_k^{\mathbb{Z}}$. This is done by demanding that

$$\widetilde{\mathrm{m}}([s_l, \ldots, \overset{*}{s}_0, \ldots, s_m]) = \widetilde{\mathrm{m}}([\overset{*}{s}_l, \ldots, s_0, \ldots, s_m]) = \mathrm{m}([s_l, \ldots, s_0, \ldots, s_m])$$

for all $l \leq 0 \leq m$ and all symbols.

Definition 4.26. *Let $\lambda = (\lambda_1, \ldots, \lambda_k)$ be a stationary distribution for the transition matrix P.*

(1) The ppt $(S_k^{\mathbb{N}}, \Sigma_k^{\mathbb{N}}, \mathrm{m}, \tau)$ is called a **one-sided Markov shift** *of k symbols, or the one-sided (P, λ)-shift.*

(2) The ppt $(S_k^{\mathbb{Z}}, \Sigma_k^{\mathbb{Z}}, \widetilde{\mathrm{m}}, \tilde{\tau})$ is called a **two-sided Markov shift** *of k symbols, or the two-sided (P, λ)-shift.*

The terms Markov endomorphism and Markov automorphism are also used.

It should be emphasized that the only distinction between Markov and Bernoulli shifts is the measure on the product space $(S_k^{\mathbb{N}}, \Sigma_k^{\mathbb{N}})$ or $(S_k^{\mathbb{Z}}, \Sigma_k^{\mathbb{Z}})$. In the Bernoulli case it is the product measure, corresponding to i.i.d. Bernoulli process, while in the Markov case the measure is more generally induced by a transition matrix and a stationary distribution. Of course, Bernoulli shifts are special cases of Markov shifts, as discussed earlier.

The case that concerns us the most is when there exists a power $N > 0$ such that P^N has strictly positive entries:

$$(P^N)_{r,s} > 0 \ , \quad 1 \leq r, s \leq k \ .$$

Such a transition matrix P is called **irreducible and aperiodic**. We borrow without proof the following result from the theory of Markov chains:

Lemma 4.27. *If the $k \times k$ transition matrix P is irreducible and aperiodic, there exists a unique stationary distribution, i.e., a probability vector $\lambda = (\lambda_1, \cdots, \lambda_k)$ satisfying (4.8). Moreover, λ is strictly positive ($\lambda_s > 0$, $1 \leq s \leq k$) and $\lim_{n\to\infty}(P^n)_{r,s} = \lambda_s$.*

The last statement means that each row of P^n converges to the probability vector λ.

Definition 4.28. *A Markov shift is called irreducible and aperiodic if the corresponding transition matrix P is irreducible and aperiodic.*

Every Bernoulli shift is an irreducible and aperiodic Markov shift, since $p_{r,s} = \lambda_s > 0$.

Having introduced the concepts, we are now ready for the main result concerning the mixing properties of Markov and Bernoulli shifts.

Theorem 4.29. *An irreducible and aperiodic one-sided Markov shift is an exact endomorphism. An irreducible and aperiodic two-sided Markov shift is a Kolmogorov automorphism.*

PROOF. The proof follows Parry [9] and Walters [17]. It is not complicated, but requires a bit of attention.

One-sided case: Let $\mathscr{B} = \Sigma_k^{\mathbb{N}}$, which is the sigma-algebra generated by all cylinder sets $[s_0, \ldots, s_m]$, $m \geq 0$. The objective is to show that $\tau^{-n}\mathscr{B} \downarrow \mathscr{N}$ (mod m), which means that the sub-sigma-algebra $\mathscr{B}_\infty = \cap_{n\geq 1}\tau^{-n}\mathscr{B}$ satisfies the condition $\mathrm{m}(A) \in \{0,1\}$ for all $A \in \mathscr{B}_\infty$. The strategy is to prove that

$$\mathrm{m}(A \cap [s_0,\dots,s_m]) = \mathrm{m}(A)\,\mathrm{m}([s_0,\dots,s_m]) \tag{4.9}$$

holds for any $A \in \mathscr{B}_\infty$ and any cylinder set. Namely, this implies (exercise!) that the sigma-algebras $\mathscr{B}_\infty$ and $\mathscr{B}$ are independent, meaning

$$\mathrm{m}(A \cap B) = \mathrm{m}(A)\,\mathrm{m}(B)\ , \quad A \in \mathscr{B}_\infty\ ,\ B \in \mathscr{B}\ .$$

In particular, for any $A \in \mathscr{B}_\infty \subset \mathscr{B}$, we then have $\mathrm{m}(A) = \mathrm{m}(A)^2$, or $\mathrm{m}(A) \in \{0,1\}$ as desired.

It remains to prove (4.9). To this end, we first note that the sub-sigma-algebra $\tau^{-n}\mathscr{B}$ is generated by sets of the form $\tau^{-n}[r_0,\dots,r_l] = \{x \in S_k^{\mathbb{N}} : x_n = r_0,\dots,x_{n+l} = r_l\}$. Next, let $B = [s_0,\dots,s_m]$ and $C = [r_0,\dots,r_l]$ be two cylinder sets. If $n \geq m+2$, we have

$$\begin{aligned}
\mathrm{m}(B \cap \tau^{-n}C) &= \sum_{1\leq t_1,\dots,t_{n-m-1}\leq k} \mathrm{m}([s_0,\dots,s_m,t_1,\dots,t_{n-m-1},r_0,\dots,r_l]) \\
&= \sum_{1\leq t_1,\dots,t_{n-m-1}\leq k} \lambda_{s_0}\, p_{s_0,s_1}\cdots p_{r_{l-1},r_l} \\
&= \lambda_{s_0}\, p_{s_0,s_1}\cdots p_{s_{m-1},s_m}(P^{n-m})_{s_m,r_0}\, p_{r_0,r_1}\cdots p_{r_{l-1},r_l} \\
&= (P^{n-m})_{s_m,r_0}\lambda_{r_0}^{-1}\cdot \mathrm{m}(B)\,\mathrm{m}(C) \\
&= (P^{n-m})_{s_m,r_0}\lambda_{r_0}^{-1}\cdot \mathrm{m}(B)\,\mathrm{m}(\tau^{-n}C)\ .
\end{aligned}$$

Since P is irreducible and aperiodic, Lemma 4.27 applies, so $\lim_{n\to\infty}(P^{n-m})_{s_m,r_0}\lambda_{r_0}^{-1} = 1$. Thus, for any $m \geq 0$ and $\varepsilon \in (0,1)$, there exists $n_0 = n_0(m,\varepsilon)$ such that

$$(1-\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(\tau^{-n}C) \leq \mathrm{m}(B \cap \tau^{-n}C) \leq (1+\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(\tau^{-n}C)$$

holds for all $n \geq n_0$, and all $B = [s_0,\dots,s_m]$ and $C = [r_0,\dots,r_l]$, $l \geq 0$. This implies (exercise!)

$$(1-\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(A) \leq \mathrm{m}(B \cap A) \leq (1+\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(A)\ , \quad A \in \tau^{-n}\mathscr{B}\ ,$$

for all $n \geq n_0$ and all cylinder sets $B = [s_0,\dots,s_m]$, where m has the value fixed earlier. In particular, if A is an arbitrary set in the intersection $\mathscr{B}_\infty = \cap_{n\geq 1}\tau^{-n}\mathscr{B}$, we see that

$$(1-\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(A) \leq \mathrm{m}(B \cap A) \leq (1+\varepsilon)\,\mathrm{m}(B)\,\mathrm{m}(A)$$

holds for all and all cylinder sets $B = [s_0,\dots,s_m]$, with an arbitrary $m \geq 0$. Since ε was arbitrary, (4.9) follows.

Two-sided case: Let $\mathscr{B} = \Sigma_k^{\mathbb{Z}}$ and let $\mathscr{A}$ be the sub-sigma-algebra generated by the "one-sided" cylinder sets $[\overset{*}{s}_0,\dots,s_m]$, $m \geq 0$. It follows from $\tilde{\tau}^{-1}[\overset{*}{s}_0,\dots,s_m] = \cup_{1\leq s\leq k}[\overset{*}{s}, s_0\dots,s_m]$ that $\tilde{\tau}^{-1}\mathscr{A} \subset \mathscr{A}$. Next, we show that $\tilde{\tau}^n\mathscr{A} \uparrow \mathscr{B}$. For each $n \geq 1$, $\tilde{\tau}^n\mathscr{A}$ is a sub-sigma-algebra of $\mathscr{B}$ containing all cylinder sets of the form $[s_l,\dots,\overset{*}{s}_0,\dots,s_m]$, where $-n \leq l \leq 0 \leq m$. Thus, $\cup_{n\geq 1}\tilde{\tau}^n\mathscr{A} \subset \mathscr{B}$ contains all cylinder sets $[s_l,\dots,\overset{*}{s}_0,\dots,s_m]$, $l \leq 0 \leq m$, which generate $\mathscr{B}$, so $\sigma(\cup_{n\geq 1}\tilde{\tau}^n\mathscr{A}) = \mathscr{B}$. In other words, $\tilde{\tau}^n\mathscr{A} \uparrow \mathscr{B}$. Finally, $\tilde{\tau}^{-n}\mathscr{A} \downarrow \mathscr{N}$ (mod $\widetilde{\mathrm{m}}$) is proved similarly to the one-sided case. □

Exercise 4.30. *Bridge the gaps in the proof concerning the independence of $\mathscr{B}_\infty$ and $\mathscr{B}$. [Hint: π–λ lemma.]*

To conclude, we have established the following hierarchy of mixing properties:

<u>non-invertible ppt</u>		<u>invertible ppt</u>
one-sided Bernoulli shift		two-sided Bernoulli shift
$\Downarrow$		$\Downarrow$
one-sided irreducible aperiodic Markov shift		two-sided irreducible aperiodic Markov shift
$\Downarrow$		$\Downarrow$
exact endomorphism		Kolmogorov automorphism
$\Searrow$		$\Swarrow$
	mixing	
	$\Downarrow$	
	weakly mixing	
	$\Downarrow$	
	ergodic	

The reader should now be equipped to solve the following informative exercise, by an approach similar to the one in the exactness proof of a one-sided Markov shift above.

Exercise 4.31. *Recall that the angle doubling map* $T : \mathbb{S}^1 \to \mathbb{S}^1 : (x, y) \mapsto 2x \pmod 1$ *is mixing by Exercise 4.6. Prove that it is exact.*

[Hint: The dyadic intervals $I_{m,i} = [2^{-m}(i-1), 2^{-m}i]$, $1 \le i \le 2^m$, $m \ge 1$, *generate the Borel sigma-algebra on* $\mathbb{S}^1$. *The dyadic intervals, with* m *fixed, can be parametrized in a clever way: Denote* $[s_0, \dots, s_{m-1}] = \{x \in [0,1] : \sum_{i=0}^{m-1} 2^{-i-1}s_i \le x \le \sum_{i=0}^{m-1} 2^{-i-1}s_i + 2^{-m}\}$ *for all* $s_i \in \{0, 1\}$, $0 \le i \le m-1$. *Then* $I_{m,i} = [s_0, \dots, s_m - 1]$ *for suitable indices. Writing* $J_0 = [0, \frac{1}{2}]$ *and* $J_1 = [\frac{1}{2}, 1]$, *we have* $x \in [s_0, \dots, s_{m-1}]$ *if and only if* $T^i(x) \in J_{s_i} \pmod 1$, $0 \le i \le m-1$.*]*

CHAPTER 5

Introduction to information and entropy

The rest of these lecture notes focuses on the connection between information theory and ergodic theory, and is based mostly on Parry [9]. On the one hand, we will prove a fundamental result in information theory with ergodic-theoretic means. On the other hand, we will see how the information-theoretic notion of entropy can be be exploited in ergodic theory.

1. Concepts

In this section we introduce and motivate the concepts of information and entropy.

Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space. Suppose $x \in X$ is picked at random according to m, and that our task is to try to guess which point x is. From the probabilistic point of view, the task would be perfectly successful if we could identify a set $C \in \mathscr{B}$ of measure zero which contains x. Of course, without any information given on x, we can only say that Probability$(x \in C) = \mathrm{m}(C)$ for each $C \in \mathscr{B}$. Now, suppose someone reveals to us that $x \in A$ (where $A \in \mathscr{B}$ is some particular set). We then gain information on x; the smaller $\mathrm{m}(A)$, the more information. Motivated by this, we associate to each set A an amount of gained information, $\varphi(\mathrm{m}(A)) \geq 0$, which only depends on the measure $\mathrm{m}(A)$ of the set.

To rephrase, we think of A as the actual information gained, and of $\varphi(\mathrm{m}(A))$ as the amount of information gained. The latter is a numerical quantifier of the former.

Above, $\varphi : (0, 1] \to [0, \infty)$ is a function still to be determined. (For reasons soon to be discovered, we exclude 0 from the domain.) It turns out that the choice of φ is essentially unique, after imposing a natural condition. Namely, suppose we receive two *independent* pieces of information, which do not affect each other in any way: $x \in A$ and $x \in B$, where the sets $A, B \in \mathscr{B}$ satisfy $\mathrm{m}(A \cap B) = \mathrm{m}(A)\,\mathrm{m}(B)$. Then the amount of gained information associated to $A \cap B$ should be the sum of the amounts associated to A and B separately, or $\varphi(\mathrm{m}(A)\,\mathrm{m}(B)) = \varphi(\mathrm{m}(A)) + \varphi(\mathrm{m}(B))$. Thus, we demand that the function φ satisfy

$$\varphi(st) = \varphi(s) + \varphi(t) \ , \quad 0 < s, t \leq s + t \leq 1 \ .$$

Finally, since $x \in X$ represents worthless information and $x \in A$, $\mathrm{m}(A) \approx 1$, nearly worthless information, we might set $\varphi(1) = 0$ and demand $\varphi(t) \approx 0$ if $t \approx 1$. But let us just assume that φ is bounded in a neighborhood $(1 - \delta, 1]$ of 1 for some $\delta > 0$. By the exercise below, we are now forced to take

$$\varphi(t) = -\log t \ , \quad 0 < t \leq 1 \ ,$$

up to a constant multiple. Observe that $\lim_{t \to 0+} \varphi(t) = \infty$ is compatible with the idea that $x \in A$ with $\mathrm{m}(A) = 0$ corresponds to perfect information on x.

Exercise 5.1. *Suppose* $\psi : \mathbb{R} \to \mathbb{R}$ *satisfies* $\psi(s+t) = \psi(s) + \psi(t)$ *and is bounded in a neighborhood of* 0. *Show that* $\psi(t) = t\psi(1)$ *for all* $t \in \mathbb{R}$. *As an application, show that* $\varphi(t) = C \log t$ *for some constant* C.

[Hint: Start by proving the identity $\psi(t) = t\psi(1)$ *for rational numbers* t. *Assuming* ψ *is continuous at* 0, *extend the identity to all of* $\mathbb{R}$. *Then show that continuity at* 0 *follows from boundedness in a neighborhood of* 0. *As for the application, note that* φ *admits a unique extension from* $(0,1]$ *to* $(0,\infty)$ *satisfying* $\varphi(st) = \varphi(s) + \varphi(t)$.*]*

Above we learned that the number

$$I(A) = -\log \mathrm{m}(A)$$

is a reasonable quantifier of the amount of information gained from a set $A \in \mathscr{B}$. We are now ready to introduce the information function and the entropy. They are defined with respect to multiple sets $A_1, A_2, \ldots$ forming a partition of X.

Definition 5.2. *Let* $\alpha = \{A_1, A_2, \ldots\} \subset \mathscr{B}$ *be a countable or finite collection of sets. We say that* α *is a* **countable partition** *if* $\mathrm{m}(X \setminus \cup_{i\geq 1} A_i) = 0$ *and* $\mathrm{m}(A_i \cap A_j) = 0$ *for* $i \neq j$.

Definition 5.3. *Let* α *be a countable partition. The function*

$$I(\alpha) = -\sum_{A\in\alpha} 1_A \, \log \mathrm{m}(A)$$

from X *into* $[0,\infty)$ *is the* **information function of** α. *Its expected value*

$$H(\alpha) = \int_X I(\alpha)\,\mathrm{dm} = -\sum_{A\in\alpha} \mathrm{m}(A)\log \mathrm{m}(A)$$

is the **entropy of** α.[1]

The number $H(\alpha)$ is called entropy because of its similarity to entropy in statistical mechanics. The latter has the expression $-k_B \sum_i p_i \log p_i$, where $k_B > 0$ is the Boltzmann constant and p_i is the probability that the state i of the statistical mechanical system at issue is occupied. We will not discuss this further.

If α is a countable partition, almost every $x \in X$ belongs to a unique $A \in \alpha$. It is convenient to denote this unique partition element by $\alpha(x)$. We then have simply

$$I(\alpha)(x) = -\log \mathrm{m}(\alpha(x))$$

for almost every x, which connects in an obvious way the information function of a partition to the information gained from individual sets. The information function tells us the amount of information $I(\alpha)(x)$ gained from the partition α on each point x. It is (almost) constant on partition elements. Observe that the entropy $H(\alpha)$ can be viewed as the weighted average of $-\log \mathrm{m}(A)$ with weight $\mathrm{m}(A)$ over all $A \in \alpha$. Alternatively, we have

$$H(\alpha) = -\int_X \log \mathrm{m}(\alpha(x))\,\mathrm{dm}\ ,$$

which is the expected value of $I(\alpha(x))$.

[1] In the definition of $I(\alpha)$ we exclude any terms with $\mathrm{m}(A) = 0$ from the sum. In the definition of $H(\alpha)$ this is automatic by assuming $0 \log 0 = 0$. Similar conventions will be used hereafter. There is deliberate overloading of the symbol I, which should not cause confusion.

Then most interesting question is how much *additional* information is gained from a set or a countable partition when some prior information has already been given. This leads to the notions of conditional information and entropy, which we next introduce. Suppose we have already been given information in the form of a set $A \in \mathscr{B}$, after which we are given additional information in the form of a set $B \in \mathscr{B}$. The most obvious quantifier of the amount of additional information is the number

$$I(B|A) = \varphi(\mathrm{m}(B|A)) = -\log \mathrm{m}(B|A)$$

where $\mathrm{m}(B|A) = \frac{\mathrm{m}(A\cap B)}{\mathrm{m}(A)}$ stands for the conditional probability of B given A. This is the right choice, because it results in the right amount of total information gained: altogether, we end up receiving the informations A and B, or $A \cap B$, and

$$-\log \mathrm{m}(A \cap B) = -\log(\mathrm{m}(B|A)\mathrm{m}(A)) = -\log \mathrm{m}(A) - \log \mathrm{m}(B|A) \ .$$

The left side is the total amount of information gained, which on the right side is expressed as the sum of two amounts — the information gained from A and the additional information gained from B given A. This immediately leads us to define

$$I(\beta|A) = -\sum_{B\in\beta} 1_B \ \log \mathrm{m}(B|A)$$

as the information function of a countable partition β given a set A. As natural as the definition is, it is not sufficient for our purposes; we need to be able to receive information in the form of a countable partition, or even a sub-sigma-algebra. This is accomplished in a straightforward manner. Namely, the elementary conditional probability $\mathrm{m}(B|A)$ given a set A is replaced by the general conditional probability $\mathrm{m}(B|\mathscr{A}) = \mathrm{E}(1_B|\mathscr{A})$ given a sub-sigma-algebra $\mathscr{A}$.

Definition 5.4. *Let β be a countable partition and $\mathscr{A}$ a sub-sigma-algebra. The function*

$$I(\beta|\mathscr{A}) = -\sum_{B\in\beta} 1_B \ \log \mathrm{m}(B|\mathscr{A})$$

is the **conditional information (function) of** *β* **given** *$\mathscr{A}$. Its expected value*

$$H(\beta|\mathscr{A}) = \int_X I(\beta|\mathscr{A})\,\mathrm{dm} = -\int_X \mathrm{m}(B|\mathscr{A}) \log \mathrm{m}(B|\mathscr{A})\,\mathrm{dm}$$

is the **conditional entropy of** *β* **given** *$\mathscr{A}$. If α is a countable partition, we often write $I(\beta|\alpha)$ for $I(\beta|\sigma(\alpha))$ and $H(\beta|\alpha)$ for $H(\beta|\sigma(\alpha))$.*

Note that $\mathrm{m}(B|\mathscr{A})$ is a function but $H(\beta|\mathscr{A})$ is a constant. The expression for $H(\beta|\mathscr{A})$ requires a small argument.[2] It can very well happen that $H(\beta|\mathscr{A}) = \infty$; we will return to this later.

Being given no prior information at all corresponds to conditioning on the trivial sigma-algebra $\mathscr{N} = \{\emptyset, X\}$. Since $\mathrm{m}(B|\mathscr{N}) = \mathrm{E}(1_B|\mathscr{N}) = \int_X 1_B\,\mathrm{dm} = \mathrm{m}(B)$, we obtain

$$I(\beta|\mathscr{N}) = I(\beta) \quad \text{and} \quad H(\beta|\mathscr{N}) = H(\beta)$$

[2]Because $1_B \ \log \mathrm{m}(B|\mathscr{A})$ is non-negative, Tonelli's theorem guarantees that $\int_X I(\beta|\mathscr{A})\,\mathrm{dm} = -\sum_{B\in\beta}\int_X 1_B \ \log \mathrm{m}(B|\mathscr{A})\,\mathrm{dm}$. Using the monotone convergence theorem $-\int_X 1_B \ \log \mathrm{m}(B|\mathscr{A})\,\mathrm{dm} = \lim_{M\to\infty}\int_X 1_B \ \min(M, -\log \mathrm{m}(B|\mathscr{A}))\,\mathrm{dm}$. Because the function $\min(M, -\log \mathrm{m}(B|\mathscr{A}))$ is bounded and $\mathscr{A}$-measurable, the properties of conditional expectation imply that the previous expression equals $\lim_{M\to\infty}\int_X \mathrm{m}(B|\mathscr{A}) \min(M, -\log \mathrm{m}(B|\mathscr{A}))\,\mathrm{dm}$, which equals $-\int_X \mathrm{m}(B|\mathscr{A}) \log \mathrm{m}(B|\mathscr{A})\,\mathrm{dm}$, again by the monotone convergence theorem.

as intuition suggests. If we have already been given more information in $\mathscr{A}$ than is contained in β, no additional information is gained:

Exercise 5.5. *Show that*

$$I(\beta|\mathscr{A}) = 0 \iff H(\beta|\mathscr{A}) = 0 \impliedby \beta \subset \mathscr{A} \pmod{\mathrm{m}}$$

(The last implication can be reversed.)

In practice, we will always condition on a sigma-algebra generated by a partition. The next exercise focuses on this case.

Exercise 5.6. *Let α and β be a countable partitions. Show that, for almost every x,*

$$I(\beta|\alpha)(x) = I(\beta|\alpha(x))(x) = I(\beta(x)|\alpha(x)) \ .$$

Thus, the definition of $I(\beta|\alpha)$ generalizes the definition of $I(\beta|A)$, which generalizes the definition of $I(B|A)$. Also, give a verbal interpretation to the consequence

$$H(\beta|\alpha) = \int_X I(\beta(x)|\alpha(x))\,\mathrm{dm}(x) = -\int_X \log \mathrm{m}(\beta(x)|\alpha(x))\,\mathrm{dm}(x) \ .$$

[Hint: Exercise 3.8.]

2. Basic identities for information and entropy

Let α and β be countable partitions. We write $\alpha \leq \beta$ of $\beta \geq \alpha$ if every element of α is a union of elements in β. We then say alternatively that α is coarser than β, β is finer than α, or β is a refinement of α. For example $\{[0,\frac{1}{2}],[\frac{1}{2},1]\} \leq \{[0,\frac{1}{4}],[\frac{1}{4},\frac{2}{4}],[\frac{2}{4},\frac{3}{4}],[\frac{3}{4},1]\}$. Intuitively, $\alpha \leq \beta$ means that β contains at least the same information as α. A similar concept can be defined $\pmod{\mathrm{m}}$. We also define the countable partition

$$\alpha \vee \beta = \{A \cap B \ : \ A \in \alpha, B \in \beta\} \ ,$$

called the **join** of α and β. Clearly $\alpha \vee \beta$ is the coarsest partition which is finer than α and β. The join $\alpha \vee \beta$ represents the total information in α and β. The join of any finite number of countable partitions is defined and interpreted similarly.

Lemma 5.7 (Basic identities)**.** *Let α, β and γ be countable partitions. Then*

$$I(\alpha \vee \beta|\gamma) = I(\alpha|\gamma) + I(\beta|\alpha \vee \gamma)$$

and

$$H(\alpha \vee \beta|\gamma) = H(\alpha|\gamma) + H(\beta|\alpha \vee \gamma) \ .$$

Since the left side is symmetric in α and β, of course

$$I(\alpha \vee \beta|\gamma) = I(\beta|\gamma) + I(\alpha|\beta \vee \gamma)$$

also holds. The lemma is so natural it could have been guessed: Suppose we have been given the information γ. Then the total amount of additional information gained from α *and* β will be the same whether we receive α and β simultaneously or one after the other.

Proof of Lemma 5.7. We check that the left side equals the right side. To that end, we appeal to the general fact that $I(\alpha|\gamma)(x) = I(\alpha(x)|\gamma(x)) = \log \mathrm{m}(\gamma(x)) - \log \mathrm{m}(\alpha(x) \cap \gamma(x))$, for almost every x, found in Exercise 5.6. To simplify the notation, suppose x is fixed and denote $A = \alpha(x)$, $B = \beta(x)$ and $C = \gamma(x)$. (This makes sense for any x in the complement of a set of measure zero, which is sufficient for our purposes). In particular, we note that $(\alpha \vee \beta)(x) = A \cap B$ and $(\alpha \vee \gamma)(x) = A \cap C$. The claim then becomes

$$\begin{aligned} & \log \mathrm{m}(C) - \log \mathrm{m}(A \cap B \cap C) \\ = \; & \log \mathrm{m}(C) - \log \mathrm{m}(A \cap C) + \log \mathrm{m}(A \cap C) - \log \mathrm{m}(A \cap B \cap C) \ , \end{aligned}$$

which is a tautology. □

The basic identities have the following elementary, but useful, consequences:

Corollary 5.8. *Let α, β and γ be countable partitions. Then*

$$\beta \leq \gamma \quad \Longrightarrow \quad I(\alpha \vee \beta|\gamma) = I(\alpha|\gamma) \quad \textit{and} \quad H(\alpha \vee \beta|\gamma) = H(\alpha|\gamma)$$

as well as

$$\alpha \leq \beta \quad \Longrightarrow \quad I(\alpha|\gamma) \leq I(\beta|\gamma) \quad \textit{and} \quad H(\alpha|\gamma) \leq H(\beta|\gamma) \ .$$

All of the implications above are perfectly natural from the point of view of gained information. The reader is invited to interpret the claims in that way.

Exercise 5.9. *Prove Corollary 5.8.*

Although not a consequence of the basic identities, the following lemma complements Corollary 5.8 nicely:

Lemma 5.10. *Let α, β and γ be countable partitions. Then*

$$\beta \leq \gamma \quad \Longrightarrow \quad H(\alpha|\beta) \geq H(\alpha|\gamma) \ .$$

More generally, if $\mathscr{A}_1 \subset \mathscr{A}_2$ are sub-sigma-algebras, then $H(\alpha|\mathscr{A}_1) \geq H(\alpha|\mathscr{A}_2)$.

Again, the statement is quite intuitive: the less prior information we have been given, the more information we expect to gain when we receive the new information α. What is counterintuitive, however, is that $\beta \leq \gamma$ does <u>not</u> by itself imply $I(\alpha|\beta) \geq I(\alpha|\gamma)$; the inequality only holds for the expected values.

Proof of Lemma 5.10. Because $\beta \leq \gamma$ implies $\sigma(\beta) \subset \sigma(\gamma)$, it suffices to prove that, for any sub-sigma-algebras $\mathscr{A}_1 \subset \mathscr{A}_2$ and any $A \in \alpha$,

$$\int_X \mathrm{m}(A|\mathscr{A}_1) \log \mathrm{m}(A|\mathscr{A}_1) \, \mathrm{dm} \leq \int_X \mathrm{m}(A|\mathscr{A}_2) \log \mathrm{m}(A|\mathscr{A}_2) \, \mathrm{dm} \ ,$$

which can be restated as

$$\int_X \phi \circ \mathrm{E}(1_A|\mathscr{A}_1) \, \mathrm{dm} \leq \int_X \phi \circ \mathrm{E}(1_A|\mathscr{A}_2) \, \mathrm{dm} \ ,$$

where $\phi(x) = x \log x$, $0 \leq x \leq 1$ $(0 \log 0 = 0)$. The key observation is that ϕ is convex and bounded. The conditional Jensen inequality (Theorem 3.7) thus applies, so

$$\phi \circ \mathrm{E}(f|\mathscr{A}_1) \leq \mathrm{E}(\phi \circ f|\mathscr{A}_1)$$

holds for any integrable function $f : X \to [0, 1]$. We choose $f = \mathrm{E}(1_A|\mathscr{A}_2)$. Because $\mathscr{A}_1 \subset \mathscr{A}_2$, $\mathrm{E}(f|\mathscr{A}_1) = \mathrm{E}(\mathrm{E}(1_A|\mathscr{A}_2)|\mathscr{A}_1) = \mathrm{E}(1_A|\mathscr{A}_1)$. This yields

$$\phi \circ \mathrm{E}(1_A|\mathscr{A}_1) \leq \mathrm{E}(\phi \circ \mathrm{E}(1_A|\mathscr{A}_2)|\mathscr{A}_1)$$

Integrating both sides finishes the proof because of the identity $\int_X \mathrm{E}(g|\mathscr{A}) \, \mathrm{dm} = \int_X g \, \mathrm{dm}$. □

3. Shannon–McMillan–Breiman theorem

In this section we formulate and prove a fundamental result in information theory, known as the Shannon–McMillan–Breiman theorem. It resembles the Birkhoff ergodic theorem, and in fact the proof relies on it. The result is thus also known as the ergodic theorem of information theory.

A straightforward connection between information theory and ergodic theory arises as follows. Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation. Recall that if $\mathscr{A}$ is a sub-sigma-algebra of $\mathscr{B}$, then $T^{-1}\mathscr{A} = \{T^{-1}A : A \in \mathscr{A}\}$ is another sub-sigma-algebra. Similarly, if α is a countable partition, then

$$T^{-1}\alpha = \{T^{-1}A : A \in \alpha\}$$

is another countable partition. Clearly $T^{-1}\mathscr{A}$ and $T^{-1}\alpha$ comprise information about the future: $x \in T^{-1}A \Leftrightarrow T(x) \in A$. Recall from Exercise 3.14 the identity

$$\mathrm{E}(f|\mathscr{A}) \circ T = \mathrm{E}(f \circ T | T^{-1}\mathscr{A}) \ ,$$

which has the intuitive interpretation that the future value of $\mathrm{E}(f|\mathscr{A})$ equals the conditional expectation of the future value of f given the future sub-sigma-algebra $\mathscr{A}$. This leads to a corresponding identity for the conditional information function, by observing that $f = 1_A$ gives

$$\mathrm{m}(A|\mathscr{A}) \circ T = \mathrm{m}(T^{-1}A | T^{-1}\mathscr{A}) \ .$$

Lemma 5.11. *Let $(X, \mathscr{B}, \mathrm{m})$ be a probability space. If T is a measure-preserving transformation, $\mathscr{A}$ a sub-sigma-algebra and β a countable partition, then*

$$I(\beta|\mathscr{A}) \circ T = I(T^{-1}\beta | T^{-1}\mathscr{A}) \quad \textit{and} \quad H(\beta|\mathscr{A}) = H(T^{-1}\beta|T^{-1}\mathscr{A}) \ .$$

In particular, infer that a countable partition β has the same entropy as $T^{-1}\beta$:

$$H(\beta) = H(T^{-1}\beta) \ .$$

Exercise 5.12. *Prove Lemma 5.11.*

Let us introduce the following notations. If α_n and $\mathscr{A}_n$, $n \geq 1$, are a sequence of countable partitions and sub-sigma-algebras, respectively, then

$$\bigvee_{i=1}^{n} \alpha_i = \alpha_1 \vee \cdots \vee \alpha_n = \{A_1 \cap \cdots \cap A_n : A_i \in \alpha_i \, , \, 1 \leq i \leq n\}$$

and

$$\bigvee_{i=1}^{n} \mathscr{A}_i = \mathscr{A}_1 \vee \cdots \vee \mathscr{A}_n = \sigma\left(\bigcup_{i=1}^{n} \mathscr{A}_i\right) .$$

These are countable partitions and sub-sigma-algebras, respectively. In the latter case we can also set $n = \infty$. It is obvious that

$$\sigma\left(\bigvee_{i=1}^{n} \alpha_i\right) = \bigvee_{i=1}^{n} \sigma(\alpha_i) \ .$$

When conditioning on a sigma-algebra generated by a countable partition, we use the shorthand notation

$$I\left(\beta \middle| \bigvee_{i=1}^{n} \alpha_i\right) = I\left(\beta \middle| \sigma\left(\bigvee_{i=1}^{n} \alpha_i\right)\right) .$$

This also makes sense for $n = \infty$.

Definition 5.13. *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation and α a countable partition with $H(\alpha) < \infty$. The number*

$$h(T,\alpha) = H\left(\alpha \,\middle|\, \bigvee_{i=1}^{\infty} T^{-i}\alpha\right)$$

is called the **entropy of T with respect to** α.

The next result is the promised ergodic theorem of information theory:

Theorem 5.14 (Shannon–McMillan–Breiman theorem)**.** *Let $(X, \mathscr{B}, \mathrm{m}, T)$ be a probability-preserving transformation and α a countable partition of X satisfying $H(\alpha) < \infty$. Denote*

$$f = I\left(\alpha \,\middle|\, \bigvee_{i=1}^{\infty} T^{-i}\alpha\right)$$

and by $\mathscr{I}$ the sub-sigma-algebra of almost invariant sets. Then

$$\lim_{n\to\infty} \frac{1}{n} I\left(\bigvee_{i=0}^{n-1} T^{-i}\alpha\right) = \mathrm{E}(f|\mathscr{I}) \ ,$$

where the convergence takes place both almost surely and in the L^1 norm. Moreover,

$$\lim_{n\to\infty} \frac{1}{n} H\left(\bigvee_{i=0}^{n-1} T^{-i}\alpha\right) = h(T,\alpha) \ .$$

In the ergodic case,

$$\mathrm{E}(f|\mathscr{I}) = h(T,\alpha) \ .$$

Let us pause to discuss the significance of the theorem to ergodic theory. In the ergodic case — which is of the greatest importance — it states that, for a typical point $x \in X$, the value $I(\vee_{i=0}^{n-1} T^{-i}\alpha)(x) = -\log \mathrm{m}((\vee_{i=0}^{n-1} T^{-i}\alpha)(x))$ is arbitrarily close to $n\,h(T,\alpha)$, provided n is large enough. Alternatively, the partition element $A = (\vee_{i=0}^{n-1} T^{-i}\alpha)(x)$ containing x has measure

$$\mathrm{m}(A) \approx e^{-n\,h(T,\alpha)} \ .$$

We can interpret this more dynamically, in terms of the trajectory of x: If $A_i = \alpha(T^i(x))$, $i \geq 0$, denotes the partition element of α which contains the point $T^i(x)$, then $A = A_0 \cap T^{-1}A_1 \cap \cdots \cap T^{-n-1}A_{n-1}$ yields

$$\mathrm{m}(A_0 \cap T^{-1}A_1 \cap \cdots \cap T^{-n-1}A_{n-1}) \approx e^{-n\,h(T,\alpha)} \ .$$

Hence, the set of all of those points $y \in X$ whose trajectories $y, T(y), \ldots, T^{n-1}(y)$ up to time $n-1$ pass through the *exact same* sets $A_0, \ldots, A_{n-1}$ has measure $\approx e^{-n\,h(T,\alpha)}$. The measure thus decreases exponentially in n. We are led to conclude that, typically, two points starting off in the same element of the partition α end up in different partition elements exponentially quickly under iterating the map T. How quickly this happens is characterized by the entropy of T with respect to α, which of course depends on both the partition and the map. Thus, we have a dynamical interpretation of $h(T,\alpha)$: the bigger the entropy, the more efficiently the partition can distinguish between the trajectories of distinct points.

[▶ Exercises on ppts here or later?]

Concerning the condition $H(\alpha) < \infty$, recall that in the Birkhoff ergodic theorem the function $f : X \to \mathbb{R}$ has to be integrable, i.e., $f \in L^1(X, \mathscr{B}, \mathrm{m})$. Analogously, in the Shannon–McMillan–Breiman theorem, $H(\alpha) < \infty$ is an integrability condition; it states precisely that $I(\alpha) \in L^1(X, \mathscr{B}, \mathrm{m})$. It also implies the integrability of $f = I(\alpha | \vee_{i=1}^{\infty} T^{-i}\alpha)$, because $H(\alpha|\mathscr{A}) \leq H(\alpha|\mathscr{N}) = H(\alpha)$ for any sub-sigma-algebra $\mathscr{A}$.

The next two results are needed for the proof of Theorem 5.14. The proof of the first one, which we omit, can be found in Parry [**9**]. The second one is a corollary of the first.

Lemma 5.15 (Chung's lemma)**.** *Suppose that β is a countable partition satisfying $H(\beta) < \infty$ and that $\mathscr{A}_1 \subset \mathscr{A}_2 \subset \cdots$ is an increasing sequence of sub-sigma-algebras. Then*

$$\int_X \sup_{n \geq 1} I(\beta|\mathscr{A}_n)\, \mathrm{dm} \leq H(\beta) + 1\ .$$

Theorem 5.16. *Suppose that β is a countable partition satisfying $H(\beta) < \infty$ and that $\mathscr{A}_1 \subset \mathscr{A}_2 \subset \cdots$ is an increasing sequence of sub-sigma-algebras with $\mathscr{A}_n \uparrow \mathscr{A}$. Then*

$$\lim_{n \to \infty} I(\beta|\mathscr{A}_n) = I(\beta|\mathscr{A})$$

almost everywhere and in the L^1 norm. Moreover,

$$H(\beta|\mathscr{A}_n) \downarrow H(\beta|\mathscr{A})\ .$$

Exercise 5.17. *Prove Theorem 5.16.*

We are now equipped to prove the Shannon–McMillan–Breiman theorem. The proof may appear lengthy, but the ideas used are at this point very simple.

Proof of Theorem 5.14. Since $H(\alpha) < \infty$, Chung's lemma implies that f is integrable. Hence, $\int_X \mathrm{E}(f|\mathscr{I})\, \mathrm{dm} = \int_X f\, \mathrm{dm} = H(\alpha|\vee_{i=1}^{\infty} T^{-i}\alpha) = h(T, \alpha)$. The claim about entropy follows from the one about the information function, because convergence in the L^1 norm implies convergence of integrals. If m is ergodic, $\mathscr{I} = \mathscr{N} \pmod{\mathrm{m}}$, which yields $\mathrm{E}(f|\mathscr{I}) = \int_X f\, \mathrm{dm} = h(T, \alpha)$.

It remains to check that $\lim_{n\to\infty} \frac{1}{n} I(\vee_{i=0}^{n-1} T^{-i}\alpha) = \mathrm{E}(f|\mathscr{I})$. Observe that, for arbitrary countable partitions, $I(\alpha \vee \beta) = I(\alpha) + I(\beta|\alpha)$ and $I(\beta) \circ T = I(T^{-1}\beta)$; the first one follows from the basic identity with $\gamma = \mathscr{N}$ and the second from Lemma 5.11 with $\mathscr{A} = \mathscr{N}$. Using these identities, we obtain

$$I\left(\bigvee_{i=0}^{n-1} T^{-i}\alpha\right) = I\left(\alpha \vee \bigvee_{i=1}^{n-1} T^{-i}\alpha\right) = I\left(\alpha \,\middle|\, \bigvee_{i=1}^{n-1} T^{-i}\alpha\right) + I\left(\bigvee_{i=0}^{n-2} T^{-i}\alpha\right) \circ T\ ,$$

because $\vee_{i=1}^{n-1} T^{-i}\alpha = T^{-1}(\vee_{i=0}^{n-2} T^{-i}\alpha)$. Let us now write $f_n = I(\alpha|\vee_{i=1}^{n} T^{-i}\alpha)$, $n \geq 0$ (with the understanding that $f_0 = I(\alpha)$). In the new notation, the identity above reads

$$I\left(\bigvee_{i=0}^{n-1} T^{-i}\alpha\right) = f_{n-1} + I\left(\bigvee_{i=0}^{n-2} T^{-i}\alpha\right) \circ T\ .$$

This is a recursion relation resulting in the expansion

$$I\left(\bigvee_{i=0}^{n-1} T^{-i}\alpha\right) = f_{n-1} + f_{n-2} \circ T + \cdots + f_0 \circ T^{n-1} = \sum_{i=0}^{n-1} f_{n-1-i} \circ T^i\ .$$

The right side resembles a Birkhoff sum to which we would like to apply the Birkhoff ergodic theorem. The problem is that there are several functions f_i appearing in the sum. The strategy is to show that $\frac{1}{n}\sum_{i=0}^{n-1} f_{n-1-i}\circ T^i$ has the same limit as $\frac{1}{n}\sum_{i=0}^{n-1} f\circ T^i$, to which Birkhoff's theorem does apply. More precisely, since

$$\left|\frac{1}{n}I\left(\bigvee_{i=0}^{n-1}T^{-i}\alpha\right)-\mathrm{E}(f|\mathscr{I})\right|\le\left|\frac{1}{n}\sum_{i=0}^{n-1}f\circ T^i-\mathrm{E}(f|\mathscr{I})\right|+\frac{1}{n}\sum_{i=0}^{n-1}|f_{n-1-i}-f|\circ T^i\ ,$$

where the first term on the right converges to zero almost everywhere and in the L^1 norm, it suffices to show that

$$\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}g_{n-1-i}\circ T^i=0 \tag{5.1}$$

almost everywhere and in the L^1 norm, where, to simplify the notation, we have denoted $g_n=|f_n-f|$, $n\ge 0$.

Since $\sigma(\vee_{i=1}^{n}T^{-i}\alpha)\uparrow\sigma(\vee_{i=1}^{\infty}T^{-i}\alpha)$, Theorem 5.16 implies that

$$\lim_{n\to\infty}g_n=0$$

almost everywhere and in the L^1 norm. In particular,

$$\begin{aligned}\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\int_X g_{n-1-i}\circ T^i\,\mathrm{dm}&=\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\int_X g_{n-1-i}\,\mathrm{dm}\\&=\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}\int_X g_i\,\mathrm{dm}=0\ .\end{aligned}$$

Hence, (5.1) holds in L^1, and it remains to prove that it also holds almost everywhere.

To prove almost sure convergence, observe that, for any $m\in\{0,1,\dots,n-1\}$,

$$\begin{aligned}\frac{1}{n}\sum_{i=0}^{n-1}g_{n-1-i}\circ T^i&=\frac{1}{n}\sum_{i=0}^{n-m-1}g_{n-1-i}\circ T^i+\frac{1}{n}\sum_{i=n-m}^{n-1}g_{n-1-i}\circ T^i\\&\le\frac{1}{n}\sum_{i=0}^{n-m-1}G_m\circ T^i+\frac{1}{n}\sum_{i=n-m}^{n-1}G_0\circ T^i\ ,\end{aligned}\tag{5.2}$$

where

$$G_m=\sup_{n\ge m}g_n\le\sup_{n\ge m}I\left(\alpha\middle|\bigvee_{i=1}^{n}T^{-i}\alpha\right)+I\left(\alpha\middle|\bigvee_{i=1}^{\infty}T^{-i}\alpha\right)\ .$$

Obviously $G_0\ge G_1\ge\cdots$. Since $\lim_{n\to\infty}g_n=0$, we conclude that $G_m\downarrow 0$ almost everywhere, as m increases to ∞. What is more, G_0 (and G_m, $m\ge 1$) is integrable by Lemma 5.15. In particular, $\sum_{i=0}^{m-1}G_0\circ T^i$ is integrable, so the Birkhoff ergodic theorem implies that

$$\frac{1}{n}\sum_{i=n-m}^{n-1}G_0\circ T^i=\frac{n-m}{n}\cdot\frac{1}{n-m}\left(\sum_{i=0}^{m-1}G_0\circ T^i\right)\circ T^{n-m}$$

converges to zero almost everywhere as $n \to \infty$; see Exercise 3.32. By (5.2), we are led to conclude that

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} g_{n-1-i} \circ T^i \leq \limsup_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-m-1} G_m \circ T^i = \mathrm{E}(G_m|\mathscr{I}) \ ,$$

where Birkhoff's theorem was used again. But $m \geq 0$ is arbitrary; taking $m \to \infty$ proves that the left side is zero because $\lim_{m\to\infty} \mathrm{E}(G_m|\mathscr{I}) = 0$ by the monotone convergence part in Theorem 3.7. □

4. Applications of the S–M–B theorem

APPENDIX A

Some measure and integration theory

For convenience, we recall routinely used terminology and facts from measure theory and Lebesgue's integration theory.

1. Positive, signed and complex measures

Let X be a set. A **sigma-algebra** $\mathscr{B}$ on X is a collection of subsets of X satisfying the following properties: (1) $\emptyset, X \in \mathscr{B}$; (2) if $A \in \mathscr{B}$, then $A^c = X \setminus A \in \mathscr{B}$; and (3) if $A_i \in \mathscr{B}$, $i \in \mathbb{N}$, then $\cup_{i\in\mathbb{N}} A_i \in \mathscr{B}$. Thus, a sigma-algebra is closed under taking complements and countable unions of its elements. Of course, it is then also closed under countable intersections. The elements $A \in \mathscr{B}$ are called **measurable sets** and $(X, \mathscr{B})$ is called a **measurable space**.

A **measure** m on a sigma-algebra $\mathscr{B}$ (or on measurable space $(X, \mathscr{B})$) is a function $\mathscr{B} \to [0, \infty]$ satisfying the following properties: (1) $\mathrm{m}(\emptyset) = 0$; (2) If $A_i \in \mathscr{B}$, $i \in \mathbb{N}$, are disjoint ($A_i \cap A_j = \emptyset$ for $i \neq j$), then $\mathrm{m}(\cup_{i\in\mathbb{N}} A_i) = \sum_{i\in\mathbb{N}} \mathrm{m}(A_i)$. The last property is called countable additivity. Then $(X, \mathscr{B}, \mathrm{m})$ is called a **measure space**. Since $\mathrm{m}(A) \geq 0$ for each $A \in \mathscr{B}$, measures defined this way are also called **positive measures**. Note that $\mathrm{m}(A) = \infty$ is allowed for a measure. We say that a measure is **finite**, if $\mathrm{m}(X) < \infty$ (and hence $\mathrm{m}(A) < \infty$ for all $A \in \mathscr{B}$), and **sigma-finite** if X can be written as a countable union $\cup_{i\in\mathbb{N}} X_i$ of disjoint sets $X_i \in \mathscr{B}$ such that $\mathrm{m}(X_i) < \infty$ for each i.

Complex measures $\mathrm{m} : \mathscr{B} \to \mathbb{C}$ and real measures $\mathrm{m} : \mathscr{B} \to \mathbb{R}$ satisfy similar properties as measures, but always take finite values by definition. A real measure is a complex measure, of course. Real measures are also called **signed measures**, which is our choice. A positive measure is a **probability measure** if $\mathrm{m}(X) = 1$, in which case $(X, \mathscr{B}, \mathrm{m})$ is a **probability space**. We use the following notations:

$$\begin{aligned}
\mathcal{M}(X, \mathscr{B}) &= \text{set of (positive) finite measures on } (X, \mathscr{B}) \\
\mathcal{P}(X, \mathscr{B}) &= \text{set of probability measures on } (X, \mathscr{B}) \\
\mathcal{M}_s(X, \mathscr{B}) &= \text{set of signed measures on } (X, \mathscr{B}) \\
\mathcal{M}_c(X, \mathscr{B}) &= \text{set of complex measures on } (X, \mathscr{B})
\end{aligned}$$

Often we write $\mathcal{M}$ instead of $\mathcal{M}(X, \mathscr{B})$, and so on, when there is no danger of confusion.

When X is a compact metric space (more generally, a topological space), $\mathscr{B}$ is always the **Borel sigma-algebra**, the smallest sigma-algebra containing the open sets. A measure — positive, signed or complex — on $\mathscr{B}$ is then called a **Borel measure**. Thus, $\mathcal{M}(X, \mathscr{B})$ is the set of positive finite Borel measures, and so on. The elements $A \in \mathscr{B}$ are called **Borel sets**.

Observe that $\mathcal{M}_s$ is a (real) vector space. Every signed measure $\mathrm{m} \in \mathcal{M}_s$ admits the splitting

$$\mathrm{m} = \mathrm{m}_+ - \mathrm{m}_-$$

into its positive part $\mathrm{m}_+ \in \mathcal{M}$ and negative part $\mathrm{m}_- \in \mathcal{M}$, both of which are finite. This allows to define a norm, the **total variation norm**

$$\|\mathrm{m}\|_{\mathrm{TV}} = \mathrm{m}_+(X) + \mathrm{m}_-(X) \ , \tag{A.1}$$

which turns $\mathcal{M}_s$ into a normed space. Likewise, a total variation norm can be defined on the vector space $\mathcal{M}_c$. It is a relevant fact for us that $\mathcal{P}$ is a convex subset (but not a linear subspace) of $\mathcal{M}_s$ and $\mathcal{M}_c$.

2. Monotone class theorem and π–λ lemma

Many properties in measure theory are difficult to check directly for every element of a sigma-algebra, but considerably easier to check for a subfamily of the sigma-algebra. If this subfamily generates the sigma-algebra, then the property in question can often be extended to the full sigma-algebra with the aid of the monotone class theorem or the π–λ lemma. The two are closely related, and choosing one over the other in practice is a matter of convenience.

Let X be a set. A collection $\mathscr{M}$ of its subsets is called a **monotone class**, if it satisfies the following properties: (1) $X \in \mathscr{M}$; (2) $\mathscr{M}$ is closed under countable unions of increasing sets: if $A_1 \subset A_2 \subset \cdots$ where $A_i \in \mathscr{M}$, then $\cup_{i\geq 1} A_i \in \mathscr{M}$; and (3) $\mathscr{M}$ is closed under countable intersections of decreasing sets: if $A_1 \supset A_2 \supset \cdots$ where $A_i \in \mathscr{M}$, then $\cap_{i\geq 1} A_i \in \mathscr{M}$.

The intersection of an arbitrary collection of monotone classes is a monotone class and the family of all subsets of X is a monotone class. Therefore, given any collection $\mathscr{A}$ of subsets, the smallest monotone class containing $\mathscr{A}$ is well defined. It is called the monotone class generated by $\mathscr{A}$, and we denote it by $\mathscr{M}(\mathscr{A})$.

A collection $\mathscr{A}$ of subsets is called an **algebra**, if it satisfies the following properties: (1) $\emptyset, X \in \mathscr{A}$; (2) If $A \in \mathscr{A}$, then $A^c = X \setminus A \in \mathscr{A}$; (3) If $A, B \in \mathscr{A}$, then $A \cup B \in \mathscr{A}$. Thus, an algebra is closed under taking complements and finite unions (and intersections) of its elements.

Theorem A.1 (Monotone class theorem)**.** *Let $\mathscr{A}$ be an algebra of subsets of X. Then $\mathscr{M}(\mathscr{A}) = \sigma(\mathscr{A})$.*

In other words, the monotone class generated by an algebra coincides with the sigma-algebra generated by the algebra. In applications one wishes to demonstrate that some property P holds for all elements of a sigma-algebra $\mathscr{B}$. One approach proceeds through the following steps:

(1) Identify a generating algebra $\mathscr{A} \subset \mathscr{B} = \sigma(\mathscr{A})$ such that property P holds for all $A \in \mathscr{A}$.
(2) Show that $\mathscr{M} = \{A \subset X \,:\, A \text{ has property } P\}$ is a monotone class.
(3) Note that $\mathscr{A} \subset \mathscr{M}$. This implies $\mathscr{M}(\mathscr{A}) \subset \mathscr{M}$.
(4) By the monotone class theorem, $\mathscr{B} = \sigma(\mathscr{A}) \subset \mathscr{M}$, so property P holds for all $A \in \mathscr{B}$.

A collection of subsets $\mathscr{P}$ is called a **π-system**, if (1) $\mathscr{P} \neq \emptyset$ and (2) if $A, B \in \mathscr{P}$, then $A \cap B \in \mathscr{P}$. (The letter π refers to the word product.)

A collection of subset $\mathscr{L}$ is called a **λ-system**, if it satisfies the following properties: (1) $X \in \mathscr{L}$; (2) $A^c \in \mathscr{L}$ for all $A \in \mathscr{L}$; and (3) $\mathscr{L}$ is closed under countable unions of disjoint sets: if $A_i \in \mathscr{L}$, $i \in \mathbb{N}$, are disjoint, then $\cup_{i \in \mathbb{N}} A_i \in \mathscr{L}$. (The letter λ refers to the word limit.)

Theorem A.2 (π–λ lemma)**.** *Let $\mathscr{P}$ be a π-system and $\mathscr{L}$ a λ-system with $\mathscr{P} \subset \mathscr{L}$. Then $\sigma(\mathscr{P}) \subset \mathscr{L}$.*

In other words, if a π-system is contained in a λ-system, then also the sigma-algebra generated by the π-system is contained in the λ-system. Another approach to prove that every element of a sigma-algebra $\mathscr{B}$ has some property P is the following:

(1) Identify a generating π-system $\mathscr{P} \subset \mathscr{B} = \sigma(\mathscr{P})$ such that property P holds for all $A \in \mathscr{P}$.
(2) Show that $\mathscr{L} = \{A \subset X \,:\, A \text{ has property } P\}$ is a λ-system.
(3) Note that $\mathscr{P} \subset \mathscr{L}$.
(4) By the π–λ lemma, $\mathscr{B} = \sigma(\mathscr{P}) \subset \mathscr{L}$, so property P holds for all $A \in \mathscr{B}$.

3. Monotone and dominated convergence

The following theorems are needed for interchanging the order of integration and taking a limit, in order to conclude that $\lim_{n\to\infty} \int_X f_n \,\mathrm{dm} = \int_X \lim_{n\to\infty} f_n \,\mathrm{dm}$ under suitable conditions.

Theorem A.3 (Monotone convergence theorem)**.** *Let $f_n : X \to [0, \infty]$, $n \geq 1$, be a <u>non-decreasing</u> sequence of measurable functions. Then the pointwise limit $f : X \to [0, \infty] : f(x) = \lim_{n\to\infty} f_n(x)$ is a measurable function and*

$$\lim_{n\to\infty} \int_X f_n \,\mathrm{dm} = \int_X f \,\mathrm{dm} \ .$$

The monotone convergence theorem is often used in conjunction with the fact that a measurable function $f \geq 0$ is the pointwise limit of a non-decreasing sequence of simple functions, finite linear combinations $s_n = \sum_{j=1}^{m(n)} \alpha_{n,j} 1_{E_{n,j}}$, where each $\alpha_{n,j}$ is a nonnegative constant and $E_{n,j}$ is a measurable set.[1] Thus it suffices to prove a statement about integrals for all indicator functions: it extends to simple functions by linearity and to nonnegative measurable functions by Theorem A.3. After this, one may again appeal to linearity and extend the result to functions $f \in L^1(X, \mathscr{B}, \mathrm{m})$ by splitting $f = f_+ - f_-$ into its positive and negative part.

When the sequence f_n is not monotone but converges pointwise, one can often dominate it by an absolutely integrable function $g \geq |f_n|$, which again permits interchanging a limit and an integral:

Theorem A.4 (Dominated convergence theorem)**.** *Let $f_n : X \to \mathbb{C}$, $n \geq 1$, be a sequence of measurable functions having a pointwise limit f almost everywhere. If there exists a*

[1] In fact, the Lebesgue integral of a measurable function $f \geq 0$ can be *defined* as $\int f \,\mathrm{dm} = \sup_{0 \leq s \leq f} \int s \,\mathrm{dm}$, where the supremum is taken over all simple functions satisfying $0 \leq s \leq f$.

function $g \in L^1(X, \mathscr{B}, \mathrm{m})$ *with* $\sup_{n\geq 1}|f_n| \leq g$ *almost everywhere, then* f *is integrable and*

$$\lim_{n\to\infty}\int_X f_n\,\mathrm{dm} = \int_X f\,\mathrm{dm}\ .$$

A dominating *integrable* function excludes sequences "escaping to infinity". For example, if $f_n = 1_{[n,n+1]}$, then $\lim_{n\to\infty} f_n = 0$ but $\int_X f_n\,\mathrm{dm} = 1$ for all $n \geq 1$. It also excludes the following "blowup": if $f_n = n1_{[0,\frac{1}{n}]}$, then $\lim_{n\to\infty} f_n = 0$ except at 0 but again $\int_X f_n\,\mathrm{dm} = 1$ for all $n \geq 1$.

The dominated convergence theorem has a useful corollary for finite measures:

Theorem A.5 (Bounded convergence theorem)**.** *Let the measure* m *be finite:* $\mathrm{m}(X) < \infty$*. Let* $f_n : X \to \mathbb{C}$*,* $n \geq 1$*, be a sequence of measurable functions having a pointwise limit* f *almost everywhere. If there exists* $M \in [0,\infty)$ *such that* $\sup_{n\geq 1}|f_n| \leq M$ *almost everywhere, then* f *is measurable and* $|f| \leq M$ *almost everywhere (hence integrable) and*

$$\lim_{n\to\infty}\int_X f_n\,\mathrm{dm} = \int_X f\,\mathrm{dm}\ .$$

4. Radon–Nikodym theorem

Theorem A.6 (Radon–Nikodym theorem)**.** *Let* m *be a sigma-finite measure and* μ *a finite complex measure, which is absolutely continuous with respect to* m*. Then there exists a unique*[2] *function* $h \in L^1(X, \mathscr{B}, \mathrm{m})$ *such that*

$$\mu(A) = \int_A h\,\mathrm{dm}\ , \quad A \in \mathscr{B}\ .$$

The function h *is called the Radon-Nikodym derivative (or density) of* μ *with respect to* m*, and is denoted by* $\frac{d\mu}{\mathrm{dm}}$*.*

[2]Elements of L^p spaces are equivalence classes of functions, so uniqueness means uniqueness of the equivalence class. Any two representatives of the equivalence class agree almost everywhere.

APPENDIX B

Some functional analysis

We recall some facts from functional analysis, which will be used throughout the lectures. Many of the results could be generalized and are stated without proofs. Both the generalizations and their proofs can be found in standard textbooks such as [13,14,3].

Let V and V' be normed spaces over the field $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. A linear operator $L : V \to V'$ is continuous if and only if it is bounded, meaning that the operator norm

$$\|L\| = \sup_{v \in V \setminus \{0\}} \frac{\|Lv\|}{\|v\|}$$

is finite. (Here the norms are the ones of the appropriate spaces.) In the special case $V' = \mathbb{K}$ we say that L is a **linear functional**. The **dual space** V^* of V is the normed space consisting of all continuous linear functionals $L : V \to \mathbb{K}$, equipped with the operator norm.

The **kernel** and **image** of a linear operator $L : V \to V'$ are the subspaces

$$\ker L = \{v \in V \ : \ Lv = 0\}$$

and

$$\operatorname{im} L = \{v' \ : \ Lv = v' \text{ for some } v' \in V\} \ ,$$

respectively. If L is continuous, $\ker L = L^{-1}\{0\}$ is closed.

Let V be a vector space and E and F two subspaces. The **sum of** E **and** F is the subspace $E + F = \{e + f \ : \ e \in E, \ f \in F\}$. Each element $v \in E + F$ has a unique representation $v = e + f$ with $e \in E$ and $f \in F$ if and only if $E \cap F = \{0\}$. In this case $E + F$ is called a **direct sum**. If V is a normed space, a direct sum is denoted by

$$E \oplus F$$

if the subspaces are *closed*. If V is an inner product space and the *closed* subspaces are *orthogonal*, $E \perp F$, we write

$$E \oplus^{\perp} F$$

for the direct sum.

1. Projection operators

A linear projection operator $P : V \to V$ is one that satisfies $P^2 = P$. On some occasions we will benefit from the following result:

Lemma B.1. *If V is a vector space and $P : V \to V$ a linear projection operator, then*

$$V = \ker P + \operatorname{im} P \quad \textit{and} \quad \ker P \cap \operatorname{im} P = \{0\} \ .$$

In addition, $I - P$ is a linear projection operator, and the relations

$$\operatorname{im} P = \ker (I - P) \quad \textit{and} \quad \ker P = \operatorname{im} (I - P)$$

hold.

If V is a normed space and $P : V \to V$ is a continuous linear projection, then $\ker P$ *and* $\operatorname{im} P$ *are moreover closed subspaces:*

$$V = \ker P \oplus \operatorname{im} P \ .$$

If H is a Hilbert space *and $P : H \to H$ is a* self-adjoint *(meaning $P = P^*$) continuous linear projection, then* $\ker P$ *and* $\operatorname{im} P$ *are moreover orthogonal complements:*

$$H = \ker P \oplus^{\perp} \operatorname{im} P \ .$$

Thus, P and $I - P$ are the orthogonal projections onto the closed subspaces $\operatorname{im} P$ *and* $\ker P$*, respectively.*

Proof. Trivially $\ker P \subset \operatorname{im}(I - P)$, because $v \in \ker P$ implies $v = (I - P)v \in \operatorname{im}(I - P)$. The identity $P^2 = P$ is equivalent to $P(I - P) = 0$, which means that $\operatorname{im}(I - P) \subset \ker P$. Thus $\ker P = \operatorname{im}(I - P)$; reversing the roles of P and $I - P$ we also get $\operatorname{im} P = \ker(I - P)$. Take an arbitrary $v \in V$. Then $v = (I - P)v + Pv$ where $(I - P)v \in \ker P$ and $Pv \in \operatorname{im} P$. Thus $V = \ker P + \operatorname{im} P$. If $v \in \ker P \cap \operatorname{im} P$, then $v = Pw$ for some $w \in V$ with $0 = Pv = P^2 w = Pw = v$. Hence, $\ker P \cap \operatorname{im} P = \{0\}$.

If P is continuous, then $\ker P = P^{-1}\{0\}$ is closed. Since also $I - P$ is continuous, $\operatorname{im} P = \ker(I - P)$ is closed as well. Thus, $V = \ker P \oplus \operatorname{im} P$.

If P is self-adjoint, the claim is that $\ker P$ and $\operatorname{im} P$ are orthogonal complements. Let $x \in \ker P$ and $y \in \operatorname{im} P$, so that $Px = 0$ and $y = Pz$ for some $z \in H$. Then $\langle x, y\rangle = \langle x, Px\rangle = \langle Px, z\rangle = 0$. Hence, $\ker P \perp \operatorname{im} P$. Since $H = \ker P \oplus \operatorname{im} P$, the proof is complete. □

2. Extension of continuous linear functionals

It is occasionally necessary to extend a continuous linear functional $L : E \to \mathbb{K}$ from a subspace E of a normed space V to a continuous linear functional $\ell : V \to \mathbb{K}$ on the entire space. By an extension it is meant that L and ℓ coincide on the subspace E: $Le = \ell e$ for all $e \in E$. The Hahn–Banach extension theorem guarantees that such an extension is always possible, in such a way that L and ℓ have the same operator norm.

Theorem B.2 (Hahn–Banach extension theorem)**.** *Let V be a normed space and $E \subset V$ an arbitrary subspace. If $L : E \to \mathbb{K}$ is a countinuous linear functional, there exists an extension $\ell : V \to \mathbb{K}$ such that the operator norms $\|\ell\|_{V \to \mathbb{K}} = \|L\|_{E \to \mathbb{K}}$.*

3. Identifying subspaces using continuous linear functionals

On a few occasions we need to determine whether two closed subspaces $E \subset F \subset V$ actually satisfy $E = F$. To this end, the following lemma will be useful:

Lemma B.3. *Suppose $E \subset F$ are subspaces of a normed space V. Let $L : V \to \mathbb{K}$ be an* arbitrary *continuous linear functional which vanishes on E ($E \subset \ker L \Leftrightarrow Le = 0 \;\forall e \in E$). If L vanishes also on F, then $\bar{E} = \bar{F}$.*

Note that in the lemma it is necessary to consider *all* functionals vanishing on E. Secondly, the reason the closures of E and F appear is that, in a sense, continuous linear functionals cannot distinguish between closed and non-closed subspaces: if L vanishes on a subspace E, then it also vanishes on $\bar{E}$ by continuity.

PROOF OF LEMMA B.3. Trivially $\bar{E} \subset \bar{F}$, so we are left with proving the opposite inclusion. If $\bar{F} \not\subset \bar{E}$, then there exists a continuous linear functional L which vanishes on $\bar{E}$ but does not vanish on all of $\bar{F}$. (We will prove this in the paragraph below.) Equivalently, L vanishes on E but not on all of F. If such a functional does not exists, then $\bar{F} \subset \bar{E}$. This is what the lemma states.

Let us finally prove the existence of L as claimed above. If $\bar{F}$ is strictly larger than $\bar{E}$, fix $f_0 \in \bar{F} \setminus \bar{E}$. For each $v \in \operatorname{span}\{f_0\} \oplus \bar{E}$, there exist unique $t(v) \in \mathbb{K}$ and $e(v) \in \bar{E}$ such that $v = t(v)f_0 + e(v)$. It is clear that $t : \operatorname{span}\{f_0\} \oplus \bar{E} \to \mathbb{K}$ is a linear functional. Obviously t vanishes on $\bar{E}$. Moreover, t is continuous, because it is continuous at 0: If $v_n \to 0$, then $t(v_n)f_0 + e(v_n) \to 0$, which is only possible if $t(v_n) \to 0$ (and $e(v_n) \to 0$). By the Hahn–Banach extension theorem (Theorem B.2), t extends to a continuous linear functional $L : V \to \mathbb{K}$ on the entire normed space V in such a way that L and t agree on $\operatorname{span}\{f_0\} \oplus \bar{E}$. □

4. Compact convex subsets of locally convex spaces

In this section we recall facts concerning compact convex sets of locally convex spaces. We also recall all necessary definitions. For further background on locally convex spaces, we refer the reader to [**14**]; the results on convex sets below are part of what is known as Choquet theory, for which [**11**] is a good reference.

Let V be a vector space. A set $K \subset V$ is **convex**, if the convex combination $tu + (1-t)v \in K$ for all $t \in [0,1]$ and all $u, v \in K$. In other words, K contains the entire chord connecting two arbitrary points of K. The **convex hull** of a set $E \subset V$ is the smallest convex set $K \subset V$ containting E. (This exists and is unique, because V is convex and the intersection of convex sets is convex.) The **closed convex hull** is the closure of the convex hull. A point $e \in K$ of a convex set K is called an **extreme point**, if it cannot be represented as the convex combination of two distinct points of K. For example, the extreme points of a closed interval are its end points, the extreme points of a triangle are its vertices and the extreme points of a disk are its boundary points.

A **topological vector space** V is a vector space equipped with a topology which makes (1) every point a closed set and (2) the operations of addition ($V \times V \to V : (v_1, v_2) \to v_1 + v_2$) and scalar multiplication ($\mathbb{K} \times V \to V : (\alpha, v) \mapsto \alpha v$) continuous. A **locally convex space** is a topological vector space in which any neighborhood U of any point $v \in V$ contains a convex neighborhood K: $v \in K \subset U$. Normed spaces are obvious examples of locally convex spaces, the open balls being the convex neighborhoods.

In the text we will need to find a solution to an equation of the form $S(v) = v$ where S is a continuous map mapping a compact convex subset of a locally convex space into itself. Such a solution is called a fixed point of S. We will implement the following generalization of the classical Brouwer fixed point theorem from Euclidean spaces to arbitrary locally convex spaces.

Theorem B.4 (Schauder–Tychonoff fixed point theorem)**.** *Let V be a locally convex space. If $K \subset V$ is a nonempty compact convex set and $S : K \to K$ is a continuous map, then S has a fixed point in K. That is, there exists $v \in K$ such that $S(v) = v$.*

Note that S need not be linear in the theorem.

The next theorem establishes a relationship between a compact convex set and its extreme points.

Theorem B.5 (Krein–Milman theorem)**.** *Let V be a locally convex space. If $K \subset V$ is a compact convex set and E the set of its extreme points, then K is exactly the closed convex hull of E.*

A trivial example is the closed unit interval $K = [0, 1]$ with $E = \{0, 1\}$; the closed convex hull of $E = K$. The theorem guarantees that there the set of extreme points is always sufficiently large so that its closed convex hull coincides with the compact convex set. In particular, the set of extreme points cannot be empty (unless K is empty).

In fact, the Krein–Milman theorem can be restated as follows: under the same assumptions, given an arbitrary point $v \in K$, there exists a Borel probability measure λ supported on the closure $\bar{E}$ (which means that $\lambda(K \setminus \bar{E}) = 0$) such that v is the **barycenter** (a generalized convex combination)

$$v = \int_{\bar{E}} e \, \mathrm{d}\lambda(e) \ .$$

Therefore, the following theorem under the additional assumption that K be metrizable is sharper, as it states that an arbitrary point of a compact convex set is the barycenter of its extreme points — not just of points in the closure $\bar{E}$:

Theorem B.6 (Choquet theorem)**.** *Let V be a locally convex space. If $K \subset V$ is a <u>metrizable</u> compact convex set, E the set of its extreme points and v an arbitrary point in K, then E is a Borel set* [1] *and there exists a Borel probability measure λ supported on E (meaning $\lambda(K \setminus E) = 0$) such that*

$$v = \int_{E} e \, \mathrm{d}\lambda(e) \ .$$

To shed light on the statement, consider the compact convex set $[0, 1]^2 \subset \mathbb{R}^2$. The set of its extreme points is $E = \{e_1 = (0, 0), e_2 = (0, 1), e_3 = (1, 0), e_4 = (1, 1)\}$. Say, for the center $v = (\frac{1}{2}, \frac{1}{2})$ we have $v = \frac{1}{2}e_1 + \frac{1}{2}e_4 = \int_E e \, \mathrm{d}\lambda(e)$ where $\lambda = \frac{1}{2}(\delta_{e_1} + \delta_{e_4})$. Likewise, we also have $v = \frac{1}{2}e_2 + \frac{1}{2}e_3 = \int_E e \, \mathrm{d}\lambda'(e)$ where $\lambda' = \frac{1}{2}(\delta_{e_2} + \delta_{e_3})$. This shows that the measure λ does not have to be unique, because the convex combination required to represent v is not unique, even in finite dimensions. (It is unique when K is a finite dimensional simplex, such as a triangle). In finite dimensions, $d < \infty$, the measure-theoretic formulation is superfluous, because a finite convex combination of $\leq d + 1$ extreme points is always sufficient (Carathéodory's theorem) — although the combination obviously depends on v and although E may be uncountable (as in the case of a disk). In infinite dimensions, which is the case of interest to us, the formulation given in terms of barycenters and measures is necessary.

[1] Without metrizability, E is not generally a Borel set in the infinite dimensional case. The Choquet–Bishop–de Leeuw theorem is a generalization of the Choquet theorem to the non-metrizable setting.

5. Representation of functionals as measures

Note that if $(X, \mathscr{B}, \mathrm{m})$ is a measure space, then the relation

$$Lf = \int_X f \, \mathrm{dm} \; , \quad f \in L^1(X, \mathscr{B}, \mathrm{m}) \; ,$$

defines a linear functional $L : L^1(X, \mathscr{B}, \mathrm{m}) \to \mathbb{C}$. Positivity of the measure ($\mathrm{m}(A) \geq 0$ for all $A \in \mathscr{B}$) is equivalent to positivity of the functional ($f \geq 0 \Rightarrow Lf \geq 0$). In the positive case, m is a probability measure ($\mathrm{m}(X) = 1$) if and only if $L1 = 1$. The following theorem is a highly useful converse of our observation for functionals of continuous functions.

Theorem B.7 (Riesz–Markov–Kakutani representation theorem)**.** *Let X be a compact metric space and $\mathscr{B}$ its Borel sigma-algebra. If $L : C(X) \to \mathbb{C}$ is a positive linear functional, there exists a unique positive Borel measure* m *on $(X, \mathscr{B})$ such that*

$$Lf = \int_X f \, \mathrm{dm} \; , \quad f \in C(X) \; .$$

Moreover, the following are equivalent:

(1) $L1 < \infty$,
(2) m *is a finite measure,*
(3) L *is continuous.*

The last part of the theorem follows immediately from

$$|Lf| \leq \|f\|_\infty \, \mathrm{m}(X) = \|f\|_\infty \, L1 \leq \|f\|_\infty \, \|L\| \; .$$

Exercise B.8. *Fix $x \in X$ and define the functional $L : C(X) \to \mathbb{C} : Lf = f(x)$. Observe that L is positive and linear. Identify the Borel measure* m *representing L in Theorem B.7.*

[Hint: Start with indicator functions.]

The Riesz representation theorem yields a **one-to-one correspondence** (bijection) between continuous positive linear functionals $C(X) \to \mathbb{C}$ and finite positive Borel measures on X. In fact, the correspondence is an isometric isomorphism when the linear functionals are equipped with the operator norm $\|L\|$ and the measures with the total variation norm $\|\mathrm{m}\|_{\mathrm{TV}}$ in (A.1). There are alternate versions of the theorem yielding similarly the one-to-one correspondences

$$\mathcal{M}_s \leftrightarrow C_{\mathbb{R}}(X)^* \quad \text{and} \quad \mathcal{M}_c \leftrightarrow C(X)^* \; . \tag{B.1}$$

Here $C_{\mathbb{R}}(X)$ is the space of real-valued continuous functions and $C_{\mathbb{R}}(X)^*$ is the space of continuous linear functionals $C_{\mathbb{R}}(X) \to \mathbb{R}$. This is immediate, if one accepts the fact that any $L \in C(X)^*$ can be decomposed into $L = L_1 - L_2 + \mathrm{i}L_3 - \mathrm{i}L_4$, where each $L_i \in C(X)^*$ is positive.

In Section 3 of Appendix A, we discussed conditions allowing to take a limit inside of an integral, so that $\lim_{n\to\infty} \int_X f_n \, \mathrm{dm} = \int_X \lim_{n\to\infty} f_n \, \mathrm{dm}$. In view of the representation theorem above, one is immediately led to ask under what conditions

$$\lim_{n\to\infty} Lf_n = L\Big(\lim_{n\to\infty} f_n\Big) \tag{B.2}$$

holds true. There is a subtle point to be made here. If f_n converges to f in $C(X)$ — that is, uniformly — then (B.2) holds true if L is continuous. However, we will encounter situations where a similar conclusion is required when the sequence converges pointwise. Fortunately, we can now combine the dominated convergence theorem (Theorem A.4 or A.5) with Theorem B.7, and obtain an analogous convergence result for functionals:

Theorem B.9 (Dominated convergence theorem for linear functionals)**.** *Let X be a compact Hausdorff space and $f_n \in C(X)$, $n \geq 1$, a sequence of functions converging pointwise to $f \in C(X)$. If $\sup_{n\geq 1} \|f_n\|_\infty < \infty$, then (B.2) holds for all continuous linear functionals $L : C(X) \to \mathbb{C}$.*

Proof. Assuming L is positive, there exists a Borel measure m on $(X, \mathscr{B})$ such that $Lh = \int_X h \,\mathrm{dm}$ for all $h \in C(X)$. Since $\sup_{n\geq 1} |f_n| < \infty$, the dominated (or bounded) convergence theorem for integrals implies $\lim_{n\to\infty} Lf_n = \lim_{n\to\infty} \int_X f_n \,\mathrm{dm} = \int_X f \,\mathrm{dm} = L(f)$. A general continuous linear functional $L : C(X) \to \mathbb{C}$ can be written as $L_1 - L\_2\mathrm{i}L_3 - \mathrm{i}L_4$, where each of the linear functionals on the right side is continuous and positive. The preceding argument applies to each of them separately, so the proof is complete. □

There is also a representation theorem for linear functionals acting on L^p spaces:

Theorem B.10 (Representation of continuous linear functionals on L^p.)**.** *Let $(X, \mathscr{B}, \mathrm{m})$ be a measure space, where m is sigma-finite and positive. Let $p \in [1, \infty)$ and $q \in (1, \infty]$ be such that $\frac{1}{p} + \frac{1}{q} = 1$. For any continuous linear functional $L : L^p(X, \mathscr{B}, \mathrm{m}) \to \mathbb{C}$, there exists a unique function $h \in L^q(X, \mathscr{B}, \mathrm{m})$ such that*

$$Lf = \int_X f\, h \,\mathrm{dm}\ , \quad f \in L^p(X, \mathscr{B}, \mathrm{m})\ .$$

Moreover, $\|L\| = \|h\|_q$.

6. Weak topology and convergence of measures

In this section X is a compact metric space. The dual space $C(X)^*$ of $C(X)$ is the space of all continuous linear functional $L : C(X) \to \mathbb{C}$ equipped with the operator norm

$$\|L\| = \sup_{f\in C(X)\setminus\{0\}} \frac{|Lf|}{\|f\|_\infty}\ .$$

We will begin by introducing a weak notion of convergence for sequences of functionals $L_n \in C(X)^*$, because this will immediately yield a useful notion of convergence for sequences of finite measures m_n with the aid of the Riesz–Markov–Kakutani representation theorem (Theorem B.7).

In the operator norm topology, a sequence of continuous linear functionals L_n, $n \geq 1$, converges if there exists a continuous linear functional L such that $\lim_{n\to 0} \|L_n - L\|^* = 0$. In other words, given any $\varepsilon > 0$, there exists $N \geq 1$ such that $|L_n f - Lf| \leq \varepsilon\|f\|$ holds for all $f \in C(X)$ if $n \geq N$. For many purposes this notion of convergence is too strong. To understand this, consider the example where $X = [0,1]$ and $L_n \in C([0,1])^*$ is the evaluation functional $L_n f = f(\frac{1}{n})$. One is tempted to think that L_n converges to $L \in C([0,1])^* : Lf = f(0)$ because $L_n f \to f(0)$ for any fixed $f \in C([0,1])$. But this is *not* the case: if $f_n \in C([0,1])$ satisfies $f_n(0) = 0$, $f_n(\frac{1}{n}) = 1$, $f_n \geq 0$ and $\|f_n\|_\infty = \sup_{x\in X} |f_n(x)| = 1$, then $|L_n f_n - Lf_n| = 1 = \|f_n\|_\infty$. Hence $\|L_n - L\| \geq 1$ for all $n \geq 1$.

A remedy to the preceding ailment would be to use a weaker topology on $C(X)^*$, in the sense of which $L_n \to L$ provided that $L_n f \to Lf$ for all $f \in C(X)$. This is accomplished by endowing $C(X)^*$ with the weakest topology such that the map $E_f : C(X)^* \to \mathbb{C} : L \mapsto Lf$ is continuous for every $f \in C(X)$. This is called the topology of pointwise convergence, or the **weak-* topology of linear functionals** (pronounced "weak star topology"). If a sequence converges in this topology, which is referred to as **weak-* convergence of linear functionals**, we write

$$L_n \overset{w^*}{\to} L \quad \text{or} \quad w^*\text{-}\lim_{n\to\infty} L_n = L \ .$$

Since each E_f is continuous, we see that $L_n \overset{w^*}{\to} L$ implies $L_n f \to Lf$ for all $f \in C(X)$. The implication can be reversed:

Fact: The weak-* topology is metrizable on any ball $B_r = \{L \in C(X)^* : \|L\| \le r\}$. A metric is given by $d(L_1, L_2) = \sum_{i\ge 1} 2^{-i}|L_1 f_i - L_2 f_i|$, where $\{f_i\}_{i\ge 1}$ is a countable bounded subset of $C(X)$.[2]

Suppose $L_n f \to Lf$ for all $f \in C(X)$. Then L and L_n, $n \ge 1$, are in B_r for some $r > 0$.[3] On B_r, the weak-* topology is metrizable by d. But $d(L_n, L) = \sum_{i\ge 1} 2^{-i}|L_n f_i - Lf_i| \to 0$, so $L_n \overset{w^*}{\to} L$. In conclusion,

$$L_n \overset{w^*}{\to} L \quad \Longleftrightarrow \quad L_n f \to Lf \ , \quad f \in C(X) \ .$$

For instance, in the earlier example $L_n(f) = f(\frac{1}{n}) \to f(0) = L(f)$ for all $f \in C([0,1])$; in other words, convergence does take place in the weak-* topology, $L_n \overset{w^*}{\to} L$.

We now turn to the convergence of measures. Recall that the Riesz–Markov–Kakutani representation theorem (Theorem B.7) establishes a one-to-one correspondence between functionals and measures: the map $\Phi : \mathcal{M}_c \to C(X)^*$, where $\Phi(\mathrm{m})$ is the functional $\Phi(\mathrm{m})(f) = \int_X f\, \mathrm{dm}$, is a bijection. The weak-* topology on $C(X)^*$ induces a topology on $\mathcal{M}_c$, which turns Φ into a homeomorphism.[4] Recalling the definition of the weak-* topology on $C(X)^*$, the induced topology on $\mathcal{M}_c$ — called the **weak topology of measures** — is thus the weakest topology such that the map $\mathcal{M}_c \mapsto \mathbb{C} : \mathrm{m} \mapsto \int_X f\, \mathrm{dm}$ is continuous for all $f \in C(X)$, and it is metrizable on subsets bounded in the total variation norm (see remark below (A.1)). What is of most importance to us is the notion of convergence it yields for sequences: the measures m_n **converge weakly** to m, written $\mathrm{m}_n \Rightarrow \mathrm{m}$, if and only if the corresponding functionals converge in the weak-* sense.[5] In other words,

$$\mathrm{m}_n \Rightarrow \mathrm{m} \quad \Longleftrightarrow \quad \int_X f\, \mathrm{dm}_n \to \int_X f\, \mathrm{dm} \quad \forall f \in C(X) \ .$$

---

[2] This is so, because the normed space $C(X)$ is separable: Let $\{f_i\}_{i\ge 1}$ be a countable dense subset of the unit ball $\{f \in C(X) : \|f\|_\infty \le 1\}$ of $C(X)$. Then $d(L_1, L_2) \le \|L_1 - L_2\| < \infty$, and it can be easily checked that d is a metric on the *set* $C(X)^*$. In fact, the topology induced by this metric coincides with the weak-* topology on any ball B_r (but not on the entire $C(X)^*$).

[3] This follows from the uniform boundedness principle, since $C(X)$ is a Banach space.

[4] If $\Phi : X \to Y$ is a map and Y is a topological space, the topology on X induced by Φ consists of the preimages $\Phi^{-1}U$ of open sets of Y. Then Φ is continuous. In case Φ is a bijection, it is also open, and therefore a homeomorphism. Indeed, $\Phi(V) = U$ is open in Y for any open set $V = \Phi^{-1}U$ in X.

[5] Analysis oriented mathematicians use the term "weak-* convergence" also in the setting of Borel measures, while the term "weak convergence" is favored by probabilists.

Weak convergence is typically the satisfactory notion of convergence for measures. For instance, coming back to the earlier example, we have $L_n f = \int_{[0,1]} f \, \mathrm{d}\delta_{\frac{1}{n}}$ and $Lf = \int_{[0,1]} f \, \mathrm{d}\delta_0$. Here δ_x denotes the Borel probability measure with its entire mass at the point $x \in [0,1]$, that is, $\delta_x(A) = 1_A(x)$. The above discussion thus amounts to the statement that $\delta_{\frac{1}{n}} \Rightarrow \delta_0$ as $n \to \infty$; the point mass at $\frac{1}{n}$ converges weakly to the point mass at 0.

Let us briefly mention that requiring $\mathrm{m}_n(A) \to \mathrm{m}(A)$ for all measurable sets A results in a stronger notion of convergence, which is usually too strong. For example, δ_n does not converge to δ_0 in this strong sense, which can be seen by taking $A = \{0\}$. The failure to converge in this sense comes from the boundary set ∂A having a nonzero measure with respect to the limit δ_0, and this is a general phenomenon. In fact, the weak convergence $\mathrm{m}_n \Rightarrow \mathrm{m}$ is equivalent to the statement that $\mathrm{m}_n(A) \to \mathrm{m}(A)$ for all measurable sets such that $\mathrm{m}(\partial A) = 0$. Such sets are called "continuity sets of m". Finally, convergence in the total variation norm, $\|\mathrm{m}_n - \mathrm{m}\|_{\mathrm{TV}} \to 0$, is an even stronger notion, and corresponds to convergence of linear functionals in the operator norm.

We chose to work with complex and signed measures instead of just finite positive measures because $\mathcal{M}_c$ and $\mathcal{M}_s$ are vector spaces. What is more, the following holds:

Lemma B.11. *Equipped with the weak topology of measures, $\mathcal{M}_c$ and $\mathcal{M}_s$ are locally convex spaces.*

Sketch of proof. Since $C(X)$ is a normed space, its dual $C(X)^*$ is a locally convex space *in the weak-* topology.* The latter induces a locally convex topology on $\mathcal{M}_c$, because the linear homeomorphism Φ maps neighborhoods to neighborhoods and convex sets to convex sets. The case of signed measures is similar. □

Next, we present a lemma which states that norm-bounded sets of measures are (relatively sequentially) compact in the weak topology of measures. In our application the measures will be positive, actually probability measures, so we restrict the formulation to that case. Note that $\|\mathrm{m}\|_{\mathrm{TV}} = \mathrm{m}(X)$ for positive measures.

Lemma B.12. *Let X be a compact metric space and $\mathrm{m}_n \in \mathcal{M}$, $n \geq 1$, a sequence of positive Borel measures on $(X, \mathscr{B})$ which is bounded: $\sup_{n\geq 1} \mathrm{m}_n(X) < \infty$. Then there is a subsequence m_{n_k}, $k \geq 1$, converging weakly to a limit $\mathrm{m} \in \mathcal{M}$.*

Alternatively, the lemma can be stated as follows: if $L_n : X \to \mathbb{C}$, $n \geq 1$, is a bounded sequence ($\sup_{n\geq 1} \|L_n\| < \infty$) of continuous positive linear functionals, then it has a weak-* convergent subsequence L_{n_k}, $k \geq 1$. This is a consequence of the Banach–Alaoglu theorem stating that a norm-closed ball of the dual space of a normed space is compact in the weak-* topology, but we have decided to give a direct proof.

Proof of Lemma B.12. We need to show that there exist a subsequence m_{n_k}, $k \geq 1$, and a Borel measure m such that $\lim_{k\to\infty} \int_X f \, \mathrm{dm}_{n_k} = \int_X f \, \mathrm{dm}$ for all $f \in C(X)$. The proof relies on the fact that, for a compact metric space X, the Banach space $C(X)$ is separable, meaning that it has a countable dense set $f_k \in C(X)$, $k \geq 1$, and proceeds by a familiar diagonalization argument.

Since the sequence $\int_X f_1 \, \mathrm{dm}_n \in \mathbb{C}$, $n \geq 1$, is bounded, it has a convergent subsequence with a limit $c(f_1) \in \mathbb{C}$. We can express this as follows: there exists $\mathcal{N}_1 \subset \mathbb{N}$ such that $\lim_{n\to\infty,\, n\in\mathcal{N}_1} \int_X f_1 \, \mathrm{dm}_n = c(f_1)$. We repeat the procedure for f_2, obtaining that there exists $\mathcal{N}_2 \subset \mathcal{N}_1$ such that $\lim_{n\to\infty,\, n\in\mathcal{N}_2} \int_X f_2 \,\mathrm{dm}_n = c(f_2) \in \mathbb{C}$. Repeating this procedure over and over, we produce a nested sequence of index sets $\mathbb{N} \supset \mathcal{N}_1 \supset \mathcal{N}_2 \supset \cdots$ such that $\lim_{n\to\infty,\, n\in\mathcal{N}_k} \int_X f_k \,\mathrm{dm}_n = c(f_k) \in \mathbb{C}$ for each $k \geq 1$. Finally, we form one more index set $\mathcal{N}$ consisting of $n_1 < n_2 < \cdots$ in such a way that $n_k \in \mathcal{N}_k$ for each $k \geq 1$. Since the sets $\mathcal{N}_k$ are nested, we have actually shown that $\lim_{n\to\infty, n\in\mathcal{N}} \int_X f_k \,\mathrm{dm}_n = c(f_k)$ for all $k \geq 1$. Since the functions f_k form a dense set in $C(K)$, it follows that $\lim_{n\to\infty,\, n\in\mathcal{N}} \int_X f \,\mathrm{dm}_n = c(f) \in \mathbb{C}$ for an arbitrary $f \in C(K)$. We leave it to the reader to check that $c : C(K) \to \mathbb{C}$ is a continuous positive linear functional. By Theorem B.7, there exists a finite positive Borel measure m such that $c(f) = \int_X f \,\mathrm{dm}$ for all $f \in C(X)$. □

Bibliography

[1] Yuri Bakhtin. *Some topics in ergodic theory*. Lecture notes. Available from: http://www.cims.nyu.edu/~bakhtin/notes.pdf.

[2] Michael Brin and Garrett Stuck. *Introduction to dynamical systems*. Cambridge University Press, Cambridge, 2002. Available from: http://dx.doi.org/10.1017/CBO9780511755316, doi:10.1017/CBO9780511755316.

[3] Nelson Dunford and Jacob T. Schwartz. *Linear operators. Part I*. Wiley Classics Library. John Wiley & Sons, Inc., New York, 1988. General theory, With the assistance of William G. Bade and Robert G. Bartle, Reprint of the 1958 original, A Wiley-Interscience Publication.

[4] Rick Durrett. *Probability: theory and examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, fourth edition, 2010. Available from: http://dx.doi.org/10.1017/CBO9780511779398, doi:10.1017/CBO9780511779398.

[5] Adriano M. Garsia. A simple proof of E. Hopf's maximal ergodic theorem. *J. Math. Mech.*, 14:381–382, 1965.

[6] Adriano M. Garsia. *Topics in almost everywhere convergence*, volume 4 of *Lectures in Advanced Mathematics*. Markham Publishing Co., Chicago, Ill., 1970.

[7] Michael Hochman. *Notes on ergodic theory*. 2013. Lecture notes. Available from: http://math.huji.ac.il/~mhochman/courses/ergodic-theory-2012/notes.final.pdf.

[8] James Norris. *Advanced Probability*. Lecture notes. Available from: http://www.statslab.cam.ac.uk/~james/Lectures/ap.pdf.

[9] William Parry. *Topics in ergodic theory*, volume 75 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2004. Reprint of the 1981 original.

[10] Karl Petersen. *Ergodic theory*, volume 2 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989. Corrected reprint of the 1983 original.

[11] Robert R. Phelps. *Lectures on Choquet's theorem*, volume 1757 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, second edition, 2001. Available from: http://dx.doi.org/10.1007/b76887, doi:10.1007/b76887.

[12] Mark Pollicott and Michiko Yuri. *Dynamical systems and ergodic theory*, volume 40 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1998.

[13] Walter Rudin. *Real and complex analysis*. McGraw-Hill Book Co., New York, third edition, 1987.

[14] Walter Rudin. *Functional analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, Inc., New York, second edition, 1991.

[15] Omri Sarig. *Lecture Notes on Ergodic Theory*. 2009. Lecture notes. Available from: http://www.wisdom.weizmann.ac.il/~sarigo/506/ErgodicNotes.pdf.

[16] Yakov Sinai. *Topics in ergodic theory*, volume 44 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1994.

[17] Peter Walters. *An introduction to ergodic theory*, volume 79 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1982.

[18] David Williams. *Probability with martingales*. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge, 1991. Available from: http://dx.doi.org/10.1017/CBO9780511813658, doi:10.1017/CBO9780511813658.

Index

λ-system, 73
π–λ lemma, 73
π-system, 73

Algebra, 72
Almost invariant function, 23
Almost invariant set, 22
Angle doubling map, 14
Automorphism, 53

Banach–Alaoglu theorem, 82
Barycenter, 21, 78
Bernoulli process, 56
Bernoulli shift, 56
Birkhoff ergodic theorem, 41
Borel measure, 71
Borel set, 71
Borel sigma-algebra, 71
Bounded convergence theorem, 74

Choquet theorem, 78
Coboundary, 10, 42
Coboundary of a linear operator, 39
Conditional entropy, 63
Conditional expectation, 34
Conditional information function, 63
Convex hull, 77
Convex set, 77
Countable partition, 62
Cylinder set, 55

Dominated convergence theorem, 73
Dominated convergence theorem for linear functionals, 80
Dual space, 75

Endomorphism, 53
Entropy, 62
Entropy of a transformation with respect to a countable partition, 67
Ergodic decomposition, 29
Ergodic measure, 25
Exact endomorphism, 53
Extreme point, 20, 77

Hahn–Banach extension theorem, 76

Image of a linear operator, 75
Information function, 62
Invariant element of a linear operator, 39
Invariant function, 10, 23
Invariant measure, 13
Invariant set, 22
Irreducible and aperiodic transition matrix, 58
Isometry, 38

Join of partitions, 64

Kac recurrence theorem, 32
Kernel of a linear operator, 75
Kolmogorov automorphism (or K-automorphism), 53
Koopman operator, 38
Krein–Milman theorem, 78

Lévy–Doob upward/downward convergence theorem, 52
Linear functional, 75
Locally convex space, 77

Markov chain, 57
Markov shift, 58
Maximal ergodic theorem, 43
Measurable sets, 71
Measurable space, 71
Measure space, 71
Measure-preserving transformation, mpt, 13
Mixing, 48
Monotone class, 72
Monotone class theorem, 72
Monotone convergence theorem, 73

Poincare recurrence theorem, 31
Point mass, 13
Probability measure, 71
Probability space, 71
Probability-preserving transformation, ppt, 13
Product transformation, 15
Pushforward measure, 13

Radon–Nikodym theorem, 74
Recurrent point, 31

Riesz–Markov–Kakutani representation theorem, 79
Rotation of the circle, 14

Schauder–Tychonoff fixed point theorem, 77
Shift, 55
Sigma-algebra, 71
Stationary distribution, 57
Stationary random sequence, 56
Strong mixing, 48
Sub-sigma-algebra, 34

Time average, 9
Topological vector space, 77
Trajectory, 9
Trivial sigma-algebra, 34

Uniformly distributed sequence, 19
Uniquely ergodic transformation, 19

Von Neumann ergodic theorem, 38

Weak convergence of measures, 81
Weak mixing, 48
Weak topology of measures, 81
Weak-* convergence of linear functionals, 81
Weak-* topology of linear functionals, 81

Zero density of a set of natural numbers, 49