Topics in ergodic theory:

Notes 1

Tim Austin
timaustin@math.brown.edu
http://www.math.brown.edu/ timaustin

1 Introduction

In the first instance, a 'dynamical system' in mathematics is a mathematical object intended to capture the intuition of a 'system' in the real world that passes through some set of different possible 'states' as time passes. Although the origins of the subject lie in physical systems (such as the motion of the planets through different mutual positions, historically the first example to be considered), the mathematical study of these concepts encompasses a much broader variety of models arising in the applied sciences, ranging from the fluctuations in the values of different stocks in economics to the 'dynamics' of population sizes in ecology.

The mathematical objects used to describe such a situation consist of some set X of possible 'states' and a function $T : X \longrightarrow X$ specifying the (deterministic[1]) movement of the system from one state to the next. In order to make a useful model, one must endow the set X with at least a little additional structure and then assume that T respects that structure. Four popular choices, roughly in order of increasing specificity, are the following:

- X a measurable space and T measurable — this is the setting of *measurable dynamics*;

[1] There are also many models in probability theory of 'stochastic dynamics', such as Markov processes for which T is replaced by a probability kernel, which is something akin to a 'randomized function' – these lie outside the scope of this course.

- X a topological space and T continuous — *topological dynamics*;
- X a smooth manifold and T a smooth map — *smooth dynamics*;
- X a compact subset of $\mathbb{C}$ and T analytic — *complex dynamics*.

This course is basically concerned with the first of these possibilities, although we will also work with the second a few times. Smooth and complex dynamics are both highly-developed fields in their own right, but the extra structure that they have available makes them rather different from the measurable setting.

In fact the bare setting of measurable spaces and maps is still a little too broad: in order to develop a coherent theory we will make three additional assumptions:

- we will consider transformations of a measurable space (X, Σ) that preserve some probability measure μ on the space;
- we will assume our measurable spaces (X, Σ) are **countably generated**: that is, that there are $A_1, A_2, \ldots \in \Sigma$ which together generate the whole σ-algebra Σ;
- and we will work exclusively with transformations that are *invertible*.

Remarks 1. There is a theory of measure-preserving transformations on infinite-measure spaces, and even a theory for transformations which only leave some measure quasi-invariant (meaning that $T_{\#}\mu$ and μ are each absolutely continuous relative to the other, but they may not be equal). However, the assumption of an invariant probability measure gives us access to a huge range of phenomena and techniques that are not valid more generally, and as a consequence those other theories, although very deep in places, tend to give a much less precise understanding of a given dynamical system or question. On the other hand, the assumption of a finite invariant measure is also valid in sufficiently many cases that this theory is well worth pursuing in its own right.

2. The assumption that Σ is countably generated is to prevent our measure space from being 'too large', so that strange set-theoretic phenomena can impinge on our work. All measure spaces that arise in practice (such as from compact metric spaces) are countably generated.

3. The assumption that T be invertible is not quite so severe as it might seem. There is a standard construction, called the 'natural extension', whereby a single non-invertible μ-preserving map $T : X \longrightarrow X$ may be represented as a 'piece' of an invertible probability-preserving map on a larger space. This construction can be found, for instance, in Subsection 1.3.G of Petersen [Pet83], or on the first problem sheet. In view of this, we will not worry further about non-invertible transformations, and so settle on the following as our basic definition. $\lhd$

Definition 1.1 (Probability-preserving transformation)**.** *If (X, Σ, μ) is a probability space, then a **probability-preserving transformation (p.p.t.)** is a measurable function $T : X \longrightarrow X$ such that $\mu(T^{-1}(A)) = \mu(A)$ for all $A \in \Sigma$.*

However, having chosen to focus on this class of transformations, there is a different respect in which we can and will greatly generalize the work that we do: rather than considering a single p.p.t., many of our results will be formulated and proved for whole groups of p.p.t.s.

Definition 1.2 (Probability-preserving system)**.** *Given a countable group Γ, a **probability-preserving Γ-system (p.p. Γ-system**, or just **p.p.s.** if Γ is understood) is a probability space (X, Σ, μ) together with a collection $(T^\gamma)_{\gamma\in\Gamma}$ of p.p.t.s such that*

$$T^{\gamma\lambda} = T^\gamma \circ T^\lambda$$

for any $\gamma, \lambda \in \Gamma$. (Put another way: it is a homomorphism $T : \Gamma \longrightarrow \mathrm{Aut}(X, \Sigma, \mu)$ into the group of all *p.p.t.s on the space (X, Σ, μ).)*

While time has only one arrow, there are many applications in which a system exhibits many different kinds of symmetry, and a whole group of p.p.t.s can be used to describe this: for example, a probability measure describing the typical positions of the gas particles in a very cavernous room should be invariant under small shifts in any of the three spatial coordinate directions (at least approximately, and provided we ignore properties that are particular to the boundary of the room). It is therefore a happy occurrence that much of basic ergodic theory can be developed for more general classes of group than just $\mathbb{Z}$. In this course we will prove results for completely general countable groups Γ where possible, and in a few places will specialize to the case of $\mathbb{Z}^d$ (still better than just $\mathbb{Z}$) where some extra structure on the part of the group is needed.

Finally, a few more words are in order concerning the relation between ergodic theory and topological dynamics (the second setting introduced above). Although much of ergodic theory can be formulated for completely arbitrary probability spaces, certain technical arguments are made much easier by assuming that our probability space is actually a compact metric space with its Borel σ-algebra and a Borel probability measure. We will see later in the course that this is not really a restriction, in the sense that all p.p.s.s admit so-called ‘compact models’, but in the meantime we will simply assume this where it is helpful.

1.1 Examples and the classification problem

Ergodic theory is a subject underlain by a wealth of examples, and it is valuable to meet some of these before we begin to develop the general theory.

1. The identity transformation id on a probability space is obviously measurable and probability-preserving.

2. Other simple examples come from finite sets: if S is a finite set then a group Γ can act on it by permutations, i.e. through a homomorphism $\Gamma \longrightarrow \mathrm{Sym}(S)$, and this will obviously preserve the normalized counting measure on S.

3. **Compact group rotations** Suppose that Γ is our acting group, G is a compact metric group with Haar measure m_G (this measure is recalled below) and $\phi : \Gamma \longrightarrow G$ is a homomorphism. Then we define the resulting **rotation action** $R_\phi : \Gamma \curvearrowright (G, \mathcal{B}(G), m_G)$ by

$$R_\phi^\gamma(g) := \phi(\gamma) \cdot g.$$

   Note that the defining property of the Haar measure m_G is its invariance under all translations on G, so it is certainly R_ϕ-invariant.

   To give a more concrete example in case $\Gamma = \mathbb{Z}$, let $\alpha \in \mathbb{T}$ and define the resulting rotation $R_\alpha : \mathbb{T} \longrightarrow \mathbb{T}$ by $R_\alpha(t) = t + \alpha$. In this case $G = \mathbb{T}$ and the homomorphism ϕ above is given by $\phi(n) = n\alpha$. This more concrete example is referred to as a **circle rotation** and its obvious higher-dimensional analog on $\mathbb{T}^n$ as a **torus rotation**.

   Slightly more generally, if Γ, G and ϕ are as above and $H \leq G$ is a closed subgroup, then we can also define a rotation action R_ϕ on the compact homogeneous space G/H by

$$R_\phi^\gamma(gH) := \phi(g)gH.$$

For instance, if we consider the obvious action $\mathrm{O}(n) \curvearrowright \mathrm{S}^{n-1}$ by isometries and fix a distinguished point $u \in \mathrm{S}^{n-1}$, then we may naturally identify $\mathrm{Stab}(u) \subset \mathrm{O}(n)$ with a copy of $\mathrm{O}(n-1)$, and now since the action of $\mathrm{O}(n)$ is transitive on the sphere we obtain $\mathrm{O}(n)/\mathrm{O}(n-1) \cong \mathrm{S}^{n-1}$. Combining this with a homomorphism $\phi : \Gamma \longrightarrow \mathrm{O}(n)$ we obtain an action of Γ on the sphere by rigid motions, which obviously preserve the surface measure of S^{n-1}.

Such **rotation actions on compact homogeneous spaces** will play an important role later in the course.

4. **The adding machine and profinite actions** Various other special cases of compact group rotations are also of interest. One famous $\mathbb{Z}$-system is the **adding machine**. Consider the set $\{0,1\}^{\mathbb{N}}$ of infinite strings of 0s and 1s, identified as the group $\mathbb{Z}_2$ of dyadic integers: that is, addition is performed coordinate-wise modulo 2 but with carry to the right (this gives a different group law from $\{0,1\}^{\mathbb{N}}$ regarded simply as a direct product of copies of the two-element group). This becomes a compact group under addition, and we define $T \curvearrowright \mathbb{Z}_2$ to be the rotation by the element $\mathbf{1} := (1,0,0,\dots)$, so that

$$T(x_1,x_2,\dots) = \left\{ \begin{array}{ll} (x_1+1,x_2,\dots) & \text{if } x_1 = 0 \\ (0,0,\dots,0,x_{m+1}+1,x_{m+2},\dots) & \text{if } x_0 = \dots = x_m = 1 \text{ and } x_{m+1} = 0. \end{array} \right.$$

This is an example of a **profinite** action: given a group Γ, a profinite Γ-system (X,Σ,μ,T) is a system with

$$X = S_1 \times S_2 \times \cdots$$

for some sequence of finite sets S_1, S_2, ... with the property that if we let

$$X_m := S_1 \times S_2 \times \cdots \times S_m$$

and $\pi_m : X \longrightarrow X_m$ be the projection onto the first m coordinates, then $\pi_{m\#}\mu$ is the normalized counting measure on X_m and there are finite permutation actions $T_m : \Gamma \longrightarrow \mathrm{Sym}(X_m)$ such that $T_m \circ \pi_m = \pi_m \circ T$ for every m. Thus in quite a strong sense the system (X,Σ,μ,T) is generated by the sequence of systems $(X_m,\mathcal{P}(X_m),\pi_{m\#}\mu,T_m)$ acting on finite sets.

5. **Compact group automorphisms** Let G be a compact group again, but now suppose that $R : \Gamma \longrightarrow \mathrm{Aut}(G)$, where $\mathrm{Aut}(G)$ is the discrete group

of all continuous automorphisms of G: that is, R is an action of Γ such that R^γ is an automorphism of G for every γ. Then it can be proved that these R^γ also preserve m_G (see the first problem sheet), so $(G, \mathcal{B}(G), m_G, R)$ is another p.p.s. What similarities and differences does it have to the the first example?

Note that in case $\Gamma = \mathbb{Z}^d$, examples of this kind are often called **systems of algebraic origin**. Although we will treat them simply as interesting examples, their special algebraic structure makes possible a highly developed theory for them, which is well-treated in Schmidt [Sch95].

6. **Bernoulli shifts** A very different example arises from the basic data of probability theory: sequences of i.i.d. random variables.

   First suppose that $\Gamma = \mathbb{Z}$, let $[n] := \{1, 2, \dots, n\}$ and let $\mathbf{p} := (p_1, p_2, \dots, p_n)$ be a **stochastic vector**: that is, $p_i \geq 1$ and $\sum_{i=1}^n p_i = 1$. This $\mathbf{p}$ can naturally be interpreted as the values on singletons of a probability measure on $[n]$, and sometimes it will itself be referred to as a probability measure.

   Form $X := [n]^\mathbb{Z}$ with the product topology (which is compact and metrizable), let Σ be its Borel σ-algebra and let $\mu := \mathbf{p}^{\otimes \mathbb{Z}}$ be the product of copies of this measure $\mathbf{p}$. Now it is easy to check that the **right-shift** $T : X \longrightarrow X$ defined by

   $$T\big((x_n)_n\big) := (x_{n+1})_n$$

   preserves μ, and so defines a p.p.t. $(X, \mathcal{B}(X), \mu, T)$. The link with probability theory is seen by observing that if $\pi_i : X \longrightarrow [n]$ denotes the projection onto the coordinate indexed by $i \in \mathbb{Z}$, then on the probability space $(X, \mathcal{B}(X), \mu)$ these form an i.i.d. sequence of $[n]$-valued random variables. This system is called the **Bernoulli shift** over $\mathbf{p}$, and is often denoted simply by $\mathbf{B}(\mathbf{p})$.

   More generally, for any probability space (E, ν) and countable group Γ we could let $X := E^\Gamma$, $\mu := \nu^{\otimes \Gamma}$ and define $T : \Gamma \curvearrowright (X, \mathcal{B}(X), \mu)$ to be the coordinate right-shift given by

   $$T^\gamma\big((x_\lambda)_\lambda\big) = (x_{\lambda\gamma})_\lambda.$$

   This is sometimes referred to as the Γ**-Bernoulli shift** over (E, ν).

7. **Geometric examples** A wealth of examples can be obtained from smooth motion on a compact manifold as given by a differential equation or some

other natural geometric rule. Most classically, planetary motion can be described using a Hamiltonian system of differential equations, and the study of these systems has since taken on a life of its own. Other natural examples of such geometrical systems are geodesic flows, horocycle flows and billiard systems (whose dynamics is not technically 'smooth'). I won't describe these in detail here, but gentle introductions to some of these models can be found in Petersen [Pet83] and Glasner [Gla03].

Having introduced the above examples, it is high time to explain what ergodic theorists mean by two systems being 'essentially the same':

Definition 1.3. *Fix a countable discrete group* Γ. *Two* Γ*-systems* (X, Σ, μ, T) *and* (Y, Φ, ν, S) *are* ***isomorphic*** *if there are 'error sets'* $X_0 \in \Sigma$ *and* $Y_0 \in \Phi$ *that are respectively* T*-invariant and* S*-invariant and have* $\mu(X_0) = \nu(Y_0) = 0$*, and a measurable bijection* $\phi : X \setminus X_0 \longrightarrow Y \setminus Y_0$ *with measurable inverse such that the dynamics are intertwined:*

$$\phi \circ T^\gamma(x) = S^\gamma \circ \phi(x) \qquad \forall x \in X_0.$$

Remark The basic intuition here is that there should be an invertible map $X \longrightarrow Y$ which converts the group action on X to the group action on Y. The only wrinkle is that, as usual in the measure theoretic context, we prefer to overlook any possible 'bad' events that only have probability zero; since it's conceivable that one of our systems, say (X, Σ, μ, T), has some T-invariant subset X_0 on which the dynamics looks quite different from anything on Y, but which it happens doesn't support any of the measure μ, the above definition is framed to enable us to ignore such awkward but negligible events. $\lhd$

Which of our examples above are isomorphic?

First, it is worth recalling that any two atomless probability measures on compact metric spaces admit a huge infinity of different measure-preserving isomorphisms between them: this follows from theorems of Carathéodory and von Neumann which combined tell us that any such measure space is isomorphic to the unit interval $[0, 1]$ with Lebesgue measure, via a proof that gives a huge range of such isomorphisms. Since in practice our probability spaces will always arise as Borel measures on compact metric spaces, asking simply whether two such probability spaces are isomorphic is not very interesting.

However, the condition that an isomorphism should intertwine two different group actions is much more subtle, and leads to the above interesting question. It is known as the **classification problem**. While no good general answer to this question is known, we will see a couple of very striking partial results later: the Halmos-von Neumann Theorem treating different compact group rotations, and the Theorems of Kolmogorov-Sinai, Ornstein and Bowen which solve the isomorphism problem for different Bernoulli shifts over a wide range of different groups.

2 Some pre-requisites

This course will lean on some basic ideas from measure theory and functional analysis. You may struggle if you've not previously met measure spaces, Lebesgue integration or the basic properties of Banach spaces. At Brown these pre-requisites are all contained in MA221 and MA222, mostly in the former. I sketch what we will use later below, mostly without proofs; everything we need can also be found in most advanced introductions to analysis, such as Folland's popular text [Fol99].

2.1 Measure theory, probability and point-set topology

In this subsection I simply list some notions that we will need throughout the course:

- Measurable spaces as pairs (X, Σ) comprising a set equipped with a σ-algebra of its subsets. σ-subalgebras. Measures on a measurable space. The leading example of the Borel σ-algebra $\mathcal{B}(X)$ of a metric space X as that generated by the open sets.
- We've already mentioned the fact that any two atomless probability spaces built on the Borel σ-algebras of compact metric spaces (in fact, more generally, any separable complete metric spaces) are isomorphic, in the standard sense that there is a measure-respecting Borel bijection between them after we throw away a negligible set on either side.
- Basic integration theory: the definition of the Lebesgue integral, the monotone and dominated convergence theorems, and the definition of the Banach

spaces $L^1(\mu)$, $L^2(\mu)$ and $L^\infty(\mu)$ of a measure space (X, Σ, μ).

- Given a finite measure space (X, Σ, μ), another measurable space (Y, Φ) and a measurable map $\phi : X \longrightarrow Y$, the pushforward measure $\phi_\# \mu$ is defined by $\phi_\# \mu(A) := \mu(\phi^{-1}(A))$.

- The formation of the product of two measure spaces $(X \times Y, \Sigma \otimes \Phi, \mu \otimes \nu)$, and also the infinite product in the case of probability spaces.

- Conditional expectation: If (X, Σ, μ) is a probability space, $\Phi \leq \Sigma$ a σ-subalgebra and $f \in L^1(\mu)$ then there is a function $\mathsf{E}(f \mid \Phi) \in L^1(\mu|_\Phi)$ such that
$$\int_X \mathsf{E}(f \mid \Phi) g \, \mathrm{d}\mu = \int_X f g \, \mathrm{d}\mu \qquad \forall g \in L^1(\mu|_\Phi),$$
and the choice of $\mathsf{E}(f \mid \Phi)$ is unique up to almost-everywhere agreement.

- Basic definition of a compact metric space; formation of products of finitely or countably infinitely many such spaces, and the fact that they are still compact metric (Tychonoff's Theorem).

2.2 The space of finite Borel measures and the Riesz Representation Theorem

If (X, ρ) is a compact metric space, then the space $C(X)$ of continuous real-valued functions on X is a Banach space when equipped with the uniform norm $\|\cdot\|_\infty$, and moreover $C(X)$ is *separable*: indeed, since X is separable, if $(x_n)_n$ is a dense sequence in X, then the Stone-Weierstrass Theorem promises that rational-coefficient polynomial combinations of the distance functions

$$f_n : x \mapsto \rho(x_n, x)$$

comprise a countable uniformly dense subset of $C(X)$.

We will generally write $M(X)$ for the real vector space of all finite signed Borel measures on X, and $\mathrm{Pr}(X)$ for the subset of Borel probability measures.

Theorem 2.1 (Riesz Representation Theorem)**.** *Suppose that X is a compact metric space and let $C(X)$ be the Banach space of continuous real-valued functions* *on X equipped with the supremum norm. If μ is a signed Borel measure on X then we may define a linear functional $\psi_\mu \in C(X)^*$ by*

$$\psi_\mu(f) := \int_X f \, \mathrm{d}\mu.$$

This definition sets up a linear isometry between the space $M(X)$ of signed Borel measures on X with the total variation norm and the dual space $C(X)^*$ with its dual norm. Crucially, any bounded linear functional on $C(X)$ is represented by a signed measure in this way.

A measure μ is positive if and only if

$$\psi_\mu(1_X) = \|\psi_\mu\|$$

and is a probability measure if and only if in addition this common value is 1. □

In view of this theorem, we will generally abuse notation by regarding finite signed measures (such as probability measures) on X as being themselves members of $C(X)$.

Corollary 2.2. *If we endow $\mathrm{Pr}(X)$ with the topology inherited from the weak* topology on $C(X)^*$, then it is compact and metrizable in that topology.*

Sketch Proof This follows from the Riesz Representation Theorem and the Banach-Alaoglu Theorem, which implies that the closed unit ball of $C(X)^*$ is weak*-compact. Given this, it suffices to show that $\mathrm{Pr}(X)$ is a weak*-closed subset of that unit ball. However, the identification

$$\mathrm{Pr}(X) := \Big\{\mu \in C(X)^* : \ \textstyle\int 1_X \, \mathrm{d}\mu = 1\Big\} \cap \bigcap_{f \in C(X):\ \|f\| \le 1} \Big\{\mu \in C(X)^* : \ \Big|\textstyle\int f \, \mathrm{d}\mu\Big| \le 1\Big\}$$

expresses $\mathrm{Pr}(X)$ as an intersection of manifestly weak*-closed sets. Finally, note that since $C(X)$ is separable, if we let $(f_n)_n$ be a uniformly dense sequence in its unit ball then the weak* topology on $\mathrm{Pr}(X)$ agrees with that obtained from the metric

$$\rho(\mu, \mu') := \sum_{n=1}^{\infty} 2^{-n} \Big|\textstyle\int f_n \, \mathrm{d}\mu - \int f_n \, \mathrm{d}\mu'\Big|,$$

so this topology is also metrizable. □

2.3 Harmonic analysis

In this course our need for 'harmonic analysis' will focus on the analysis of continuous and measurable functions on *compact metric* groups (in particular, we will not venture into the more general setting of groups that are only *locally* compact). This means groups G that are equipped with a complete, compact metric ρ that is invariant under translation:

$$\rho(gh, gk) = \rho(hg, kg) = \rho(h, k) \qquad \forall g, h, k \in G$$

and for which the operations of inverse $G \longrightarrow G$ and multiplication $G \times G \longrightarrow G$ are continuous.

Theorem 2.3 (Haar measure)**.** *A group G such as above can be equipped with a unique Borel measure m_G that is left- and right-translation invariant,*

$$m_G(gE) = m_G(E) = m_G(Eg) \qquad \forall E \in \mathcal{B}(G),\ g \in G,$$

and is normalized to be a probability measure (i.e. $m_G(G) = 1$). □

The above theorem can be found in [Fol99] and also in texts on abstract harmonic analysis such as Hewitt and Ross [HR79].

Some examples of compact metric groups are the following, for each of which the unique Haar probability measure can easily be described directly.

1. Our most important example is the circle group

   $$\mathbb{T} := \mathbb{R}/\mathbb{Z}$$

   equipped with its quotient topology and with the group operation being addition modulo 1. This group is easily seen to be compact and metrizable. An isomorphic copy of this group is obtained as the set of unit vectors in $\mathbb{C}$

   $$\mathrm{S}^1 := \{z \in \mathbb{C} :\ |z| = 1\},$$

   where now the group operation is multiplication; clearly the map

   $$\mathbb{R} \longrightarrow \mathrm{S}^1 : \theta \mapsto \mathrm{e}^{2\pi\mathrm{i}\theta}$$

   is an onto continuous homomorphism with kernel $\mathbb{Z}$, and so defines the desired continuous isomorphism $\mathbb{T} \cong \mathrm{S}^1$ (the continuity of the inverse is an easy exercise).

2. More generally, the higher dimensional tori $\mathbb{T}^n$, $n \geq 2$, may be conceived as either Cartesian powers of $\mathbb{T}$ or as the higher-dimensional quotients $\mathbb{R}^n/\mathbb{Z}^n$.

3. Non-Abelian examples may be found in the form of the orthogonal groups $\mathrm{O}(n)$ (that is, the groups of linear isometries of $\mathbb{R}^n$, which are compact Lie groups) and their various closed subgroups.

4. We can always take countable Cartesian products of examples to form further examples, such as the infinite product $\mathbb{T}^{\mathbb{N}}$, which unlike the above is not a Lie group.

5. Finally, it is worth bearing in mind that any finite group is also an example.

In addition to the above examples, there are more complicated compact metric groups that are not 'finite-dimensional' after the fashion of Lie groups or finite groups. There is a theory telling us that any compact metric group can be generated as an inverse limit of these examples (this is a consequence of the Peter-Weyl Theorem), but we will not make any direct appeal to it in this course.

2.4 Spectral theorems

Suppose that $\mathfrak{H}$ is a complex Hilbert space. Then a bounded operator $T : \mathfrak{H} \longrightarrow \mathfrak{H}$ is **compact** if the image $T(B)$ of the unit ball $B \subset \mathfrak{H}$ is precompact. In case T is also self-adjoint, it is possible to give a very explicit picture of such an operator. The following theorem is treated more carefully and generally, for example, in Conway's book [Con90].

Theorem 2.4. *If T is a compact self-adjoint operator, then there are an orthonormal sequence $\xi_1, \xi_2, \ldots \in \mathfrak{H}$ and real numbers $\lambda_1, \lambda_2, \ldots$ such that $|\lambda_1| \geq |\lambda_2| \geq \ldots$, $|\lambda_i| \longrightarrow 0$ and*

$$T = \sum_{i \geq 1} \lambda_i P_i,$$

where P_i is the one-dimensional orthogonal projection onto $\mathbb{C} \cdot \xi_i$.

Proof **Step 1.** We will first show that we can find inductively an orthonormal sequence of eigenvectors $\xi_1, \xi_2, \ldots$ such that $T\xi_i = \lambda_i \xi_i$ for some $\lambda_i \in \mathbb{R}$.

Of course, if we have already found such ξ_1, ξ_2, ..., ξ_m, then T preserves the orthogonal complement $\operatorname{span}\{\xi_1, \xi_2, \dots, \xi_m\}^\perp$, because

$$\langle \zeta, \xi_i \rangle = 0 \quad \Rightarrow \quad \lambda_i \langle \zeta, \xi_i \rangle = \langle \zeta, T\xi_i \rangle = \langle T\zeta, \xi_i \rangle = 0,$$

using the fact that T is self-adjoint. So if we know how to find a single eigenvector for such a T, then applying that knowledge to $T|_{\operatorname{span}\{\xi_1,\xi_2,\dots,\xi_m\}^\perp}$ enables us to find ξ_{m+1}. Hence it suffices to show that any compact self-adjoint operator has an eigenvector.

Because $\overline{T(B)}$ is compact and the norm is a continuous function on $\mathfrak{H}$, it achieves a maximum on this set, say at $\xi \in \overline{T(B)}$. By definition it follows that $\|\xi\| = \|T\|_{\mathrm{op}}$.

Let $\zeta_n \in B$ be a sequence such that $T\zeta_n \longrightarrow \xi$ in norm, and now consider that

$$\begin{aligned} \|T\|_{\mathrm{op}}\|\xi\| = \|\xi\|^2 = \lim_n \langle \zeta_n, T^2\zeta_n \rangle \\ \leq \limsup_n \|\zeta_n\|\|T^2\zeta_n\| \leq \limsup_n \|T(T\zeta_n)\| = \|T\xi\|. \end{aligned}$$

Hence ξ also has the property that $\|T\xi\| = \|T\|_{\mathrm{op}}\|\xi\|$. Using the self-adjointness of T and the Cauchy-Schwartz inequality, this now implies that

$$\|T\|_{\mathrm{op}}\|\xi\|\|T\xi\| = \|T\xi\|^2 = \langle \xi, T^2\xi \rangle \leq \|\xi\|\|T^2\xi\|.$$

On the other hand, by the definition of the operator norm we have $\|T^2\xi\| \leq \|T\|_{\mathrm{op}}\|T\xi\|$, so in fact all the expressions above are equal and hence

$$\langle \xi, T^2\xi \rangle = \|\xi\|\|T^2\xi\|.$$

Since equality can occur in the Cauchy-Schwartz inequality only in the case of vectors differing by a positive real multiple, this implies that $T^2\xi = \lambda\xi$ for some $\lambda > 0$.

Therefore the subspace $\operatorname{span}\{\xi, T\xi\} \leq \mathfrak{H}$ is either one- or two-dimensional and invariant under T. If it is actually one-dimensional then we are done; if not, then the restriction $T|_{\operatorname{span}\{\xi,T\xi\}}$ is a self-adjoint operator on a two-dimensional Euclidean space, whose diagonalizability is well-known. Hence simply changing basis in $\operatorname{span}\{\xi, T\xi\}$ gives two orthogonal one-dimensional invariant subspaces, and once again the proof can proceed.

Step 2. Having found a sequence of eigenvectors ξ_1, ξ_2, ... as in Step 1 and normalized them all to be unit vectors, it only remains to observe that their corresponding eigenvalues must satisfy $\lambda_i \longrightarrow 0$, since otherwise T would not be compact: indeed, if there were a subsequence $(\lambda_{i_j})_j$ all of absolute value at least $\kappa > 0$, then the sequence of images

$$T\xi_{i_j} = \lambda_{i_j}\xi_{i_j} \quad j = 1, 2, \dots$$

lies in $T(B)$ and consists of vectors that are orthogonal and all have norm at least κ. This sequence would therefore have no further convergent subsequence, contradicting that $\overline{T(B)}$ is compact. On the other hand, since the procedure for producing the eigenvector sequence in Step 1 was greedy in the sense that

$$\|T\xi_{m+1}\| = \|T_{\mathrm{span}\{\xi_1,\dots,\xi_m\}^\perp}\|_{\mathrm{op}}\|\xi_{m+1}\| \qquad \forall m \geq 0,$$

it follows that $\|T_{\mathrm{span}\{\xi_1,\dots,\xi_m\}^\perp}\|_{\mathrm{op}} \longrightarrow 0$ as $m \longrightarrow \infty$ and hence that T must actually be zero on $\mathrm{span}\{\xi_1, \xi_2, \dots\}^\perp$, so we are left with the desired representation. □

References

[Con90] John B. Conway. *A course in functional analysis*, volume 96 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1990.

[Fol99] Gerald B. Folland. *Real analysis*. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York, second edition, 1999. Modern techniques and their applications, A Wiley-Interscience Publication.

[Gla03] Eli Glasner. *Ergodic Theory via Joinings*. American Mathematical Society, Providence, 2003.

[HR79] Edwin Hewitt and Kenneth A. Ross. *Abstract Harmonic Analysis, I (second ed.)*. Springer, 1979.

[Pet83] Karl E. Petersen. *Ergodic Theory*. Cambridge University Press, Cambridge, 1983.

[Sch95] Klaus Schmidt. *Dynamical systems of algebraic origin*, volume 128 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1995.

Ergodic Theory: Notes 2

1 The Ergodic Theorems

Suppose that $T \curvearrowright (X, \Sigma, \mu)$ is a p.-p.t. Pick $x \in X$, and consider its **orbit**: the set of points $\{T^n x : n \in \mathbb{Z}\}$. Within this one may consider the finite pieces $\{x, Tx, \dots, T^N x\}$ for larger and larger N. One basic intuition from the early days of ergodic theory is that, as N increases, these should 'fill up' some region of the state space, and should become more and more 'evenly distributed' within that region. This kind of phenomenon was first proposed as the 'ergodic hypothesis' by Boltzmann, in the case of a system describing the positions and momenta of all the particles in a gas. Birkhoff believed that for each initial state x, these orbit-pieces should become 'evenly distributed' among the set of all possible states having the same total energy as x (energy being a conserved quantity). Indeed, the word 'ergodic' was invented for the name of this hypothesis, taken from the Greek 'ergon' ('work') and 'odos' ('path').

In modern ergodic theory, we understand this idea in terms of the associated running averages of 'observables'. An 'observable' is just a function $f \in L^1(\mu)$, and we will consider the averages

$$S_N f(x) := \frac{1}{N} \sum_{n=1}^{N} f(T^n x).$$

The 'even distribution' of the orbits can be discussed in terms of the limiting behaviour of these averages. The intuition above essentially asserts that they should converge to the spatial average of f over some region of the state space X. We will shortly make this expectation precise.

This convergence question can be generalized to p.-p. Γ-systems for many other groups Γ, subject to certain abtract assumptions. In these notes we will state and prove the basic ergodic theorems for $\Gamma = \mathbb{Z}^d$, and will briefly mention later the much larger class of 'amenable' groups to which these results can also be extended.

For a $\mathbb{Z}^d$-action, the averages above are replaced by

$$S_N f(z) := \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} f(T^{\mathbf{n}}x),$$

where $[N] := \{1, 2, \ldots, N\}$.

Remark. Another way to think about ergodicity is that, given $x \in X$, the measure

$$\frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} \delta_{T^{\mathbf{n}}x}$$

should become 'smoothed' out in (X, Σ, μ) as $N \longrightarrow \infty$. This notion can be made precise if (X, T) is a topological dynamical system, as well as a p.-p. system. That will be the subject of the next set of notes. $\lhd$

1.1 Ergodicity and statement of the theorems

The obvious spatial average for an observable f is $\int_X f \, \mathrm{d}\mu$, but there are many system for which this need not be the correct limit.

Example. Suppose that $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ has the property that $X = X_1 \cup X_2$, a disjoint union of two measurable subsets, with $0 < \mu(X_1) < 1$, and such that these subsets are invariant under the dynamics: $T^{\mathbf{n}}X_i = X_i$ for $i = 1, 2$ and for all $\mathbf{n} \in \mathbb{Z}^d$. Now let $f := 1_{X_1}$. Then for all $N \geqslant 1$ we have

$$S_N f(x) = \begin{cases} 1 & \text{if } x \in X_1, \\ 0 & \text{if } x \in X_2. \end{cases}$$

$\lhd$

Definition 1. *A p.-p.* Γ*-system* (X, Σ, μ, T) *is* ***ergodic*** *if it has no decomposition into two* T*-invariant, positive-measure pieces as above: equivalently, if for every set* $A \in \Sigma$ *such that* $T^\gamma A = A$ *for all* A*, one has* $\mu(A) \in \{0, 1\}$.

Thus, the ergodic systems are the only ones for which our averages could possibly converge to $\int_X f \, \mathrm{d}\mu$ for all f, and we will see that this is the case.

In general, the spatial averages do always converge to something which we can describe quite simply, but we need to account for the failure of ergodicity. That can be done in terms of the **invariant σ-algebra** of a p.-p.s. (X, Σ, μ, T):

$$\Sigma^T := \{A \in \Sigma \,:\, T^\gamma A = A \; \forall \gamma \in \Gamma\}.$$

Before approaching the ergodic theorems, it is worth recording the following.

Proposition 2 (Characterizations of ergodicity)**.** *The following are equivalent:*

1. *$T : \Gamma \curvearrowright (X, \Sigma, \mu)$ is ergodic;*
2. *the σ-algebra Σ^T consists of only negligible and co-negligible sets;*
3. *if $A \in \Sigma$ is such that $\mu(T^\gamma A \triangle A) = 0$ for all γ, then $\mu(A) \in \{0, 1\}$;*
4. *if a measurable function $f : X \longrightarrow \mathbb{R}$ is such that $f(T^\gamma x) = f(x)$ for every $x \in X$ and $\gamma \in \Gamma$, then f is a.e. equal to a constant function;*
5. *if a measurable function $f : X \longrightarrow \mathbb{R}$ is such that $f(T^\gamma x) = f(x)$ for μ-a.e. x for every $\gamma \in \Gamma$, then f is a.e. equal to a constant function.*

Proof. *(1.$\Longleftrightarrow$ 2.)* Immediate from the definition of Σ^T.

(1. $\Longleftarrow$ 3.) Trivial.

(1. $\Longrightarrow$ 3.) If $\mu(T^\gamma A \triangle A) = 0$, then the set $B := \bigcap_{\gamma \in \Gamma} T^\gamma B$ is a countable intersection (because Γ is countable) of sets that all agree with A a.e., so $\mu(A \backslash B) = 0$. Also, B is strictly T-invariant by construction. Therefore (1) gives $\mu(A) = \mu(B) \in \{0, 1\}$.

(1. $\Longleftarrow$ 4.) If $A \in \Sigma$ is T-invariant, then its indicator function 1_A is a T-invariant. By assumption, this means 1_A is a.e. equal to a constant, and so A is either negligible (so $1_A = 0$ a.e.) or co-negligible (so $1_A = 1$ a.e.).

(1. $\Longrightarrow$ 4.) Now suppose that T is ergodic and that $f \in L^p(\mu)$ is T-invariant. Then each level set

$$A_q := \{x \in X : \ f(x) \leqslant q\} \quad \text{for } g \in \mathbb{Q}$$

is T-invariant, so by ergodicity, $\mu(A_q) \in \{0, 1\}$ for every $q \in \mathbb{Q}$. Since f takes finite values a.e., we must have $\mu(A_q) \longrightarrow 0$ as $q \longrightarrow -\infty$, and hence in fact $\mu(A_q) = 0$ for all sufficiently small q; similarly, $\mu(A_q) = 1$ for all sufficiently large q. Let

$$r := \sup\{q \in \mathbb{Q} : \ \mu(A_q) = 0\},$$

and let $A := \bigcup_{q \in \mathbb{Q} \cap (-\infty, r)} A_q$. This is measurable, as a countable union of measurable sets. On the other hand, $A_q \subseteq A_{q'}$ whenever $q \leqslant q'$, which implies easily that also $\mu(A_q) = 1$ whenever $q > r$. It follows that the sets

$$\{f < r\} = \bigcup_{q \in \mathbb{Q} \cap (-\infty, r)} A_q$$

and

$$\{f > r\} = \bigcup_{q \in \mathbb{Q} \cap (r,\infty)} (X \backslash A_q)$$

have measure zero, so $f = r$ a.s.

(4. $\Longleftarrow$ 5.) Trivial.

(4. $\Longrightarrow$ 5.) If $f(T^\gamma x) = f(x)$ for a.e. x for all γ, then the function $g(x) := \inf_{\gamma \in \Gamma} f(T^\gamma x)$ is measurable (as a countable infimum of measurable functions), Γ-invariant by construction, and agrees with f a.e. Therefore f, like g, must be equal to a constant a.e. □

As these proofs show, we can henceforth afford to be a little sloppy about the difference between 'almost invariant' and 'invariant'. Henceforth, the assertion that a function is 'invariant' will usually be taken to mean 'invariant a.e.'.

Remark. There is a technical sense in which an arbitrary system can be decomposed into a kind of 'union' of ergodic systems, but there may be a whole continuum of these ergodic systems involved, making this a rather complicated description to pursue; we will not treat it in these notes. However, we will need to discuss ergodicity (and also various strengthenings of it) repeatedly. ◁

We can now make a start on the main theorem of this set of notes.

Theorem 3 (Ergodic Theorems for $\mathbb{Z}^d$)**.** *Suppose that $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ and that $f \in L^1(\mu)$, and let*

$$S_N f := \frac{1}{N^d} \sum_{\mathbf{n} \in \{1,2,\dots,N\}^d} f \circ T^{\mathbf{n}}.$$

Then there is some function $\bar{f} \in L^1(\mu)$ that is invariant under the dynamics (i.e. $\bar{f} \circ T^{\mathbf{m}} = \bar{f}$ for all $\mathbf{m} \in \mathbb{Z}^d$) and such that

- *(Norm Ergodic Theorem[1]) $S_N f \longrightarrow \bar{f}$ in $\|\cdot\|_1$ as $N \longrightarrow \infty$;*
- *(Pointwise Ergodic Theorem[2]) there is some $X_0 \subseteq X$, depending on f, with $\mu(X_0) = 1$ and such that $S_N f(x) \longrightarrow \bar{f}(x)$ for all $x \in X_0$.*

The limit $\overline{f}$ is equal to $\mathsf{E}(f \mid \Sigma^T)$, the conditional expectation onto the σ-subalgebra of T-invariant sets. This is a.s. equal to $\int_X f \, \mathrm{d}\mu$ if T is ergodic.

[1]Due to von Neumann for $d = 1$, and also called the 'Mean Ergodic Theorem'

[2]Due to Birkhoff for $d = 1$, and also called the 'Individual Ergodic Theorem'

Corollary 4. *If $\Gamma = \mathbb{Z}^d$, then the following are equivalent:*

1. *$T : \Gamma \curvearrowright (X, \Sigma, \mu)$ is ergodic;*
2. *for any $f \in L^1(\mu)$ the averages $S_N(f)$ converge to a constant function $\bar{f}$ in $L^1(\mu)$;*
3. *for any $f \in L^1(\mu)$ the averages $S_N(f)$ converge pointwise to a constant function $\bar{f}$.* □

1.2 Proof of the Norm Ergodic Theorem

To the collection of p.-p.t.s $T^{\mathbf{n}} \curvearrowright (X, \Sigma, \mu)$, we can associate a collection of operators $U_T^{\mathbf{n}} : L^p(\mu) \longrightarrow L^p(\mu)$ (where p can be any given exponent in $[1, \infty]$) defined by

$$U_T^{\mathbf{n}} f(x) := f(T^{\mathbf{n}} x).$$

The Norm Ergodic Theorem is really just the assertion that the averages

$$\frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} U_T^{\mathbf{n}} f$$

converge in the function space $L^1(\mu)$ as $N \longrightarrow \infty$ when $p = 1$.

This proof can be presented in various ways; I have chosen to highlight certain features that will reappear in our treatment of multiple recurrence later. Other proofs of both the Norm and Pointwise Ergodic Theorems can be found in most books, at least for $\mathbb{Z}$-systems; Petersen [Pet83], in particular, covers several different aspects of these results.

We will first prove some results for functions and convergence in $L^2(\mu)$, and then convert these into $L^1(\mu)$ results. We do this in order to exploit the extra structure of the inner product on $L^2(\mu)$.

The heart of the proof is the following lemma.

Lemma 5. *If $f \in L^2(\mu)$ and*

$$\Big\| \frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} U_T^{\mathbf{n}} f \Big\|_2 \not\longrightarrow 0,$$

then f must correlate with a T-invariant L^2-function: that is,

$$\langle g, f \rangle = \int_X f \bar{g} \, \mathrm{d}\mu \neq 0 \qquad \textit{for some } g \in L^2(\mu) \textit{ such that } g \circ T = g.$$

Proof. By rescaling, we may assume without loss of generality that $\|f\|_2 = 1$. We will prove that there are $\delta > 0$ and a sequence $(g_i)_{i\geqslant 1}$ in $L^2(\mu)$ with

- $\|g_i\|_2 \leqslant 1$ for all i,
- $\langle g_i, f\rangle \geqslant \delta$ for all i, and
- $\|g_i \circ T^{\mathbf{p}} - g_i\|_2 \longrightarrow 0$ as $i \longrightarrow \infty$ for any fixed $\mathbf{p} \in \mathbb{Z}^d$.

Since the unit ball of $L^2(\mu)$ is compact in the weak topology (By the Banach-Alaoglu Theorem, since $L^2(\mu)$ is self-dual), we may always pass to a subsequence and so assume that $g_i \overset{\text{weak}}{\longrightarrow} g$ for some g in that unit ball. Now the second and third conditions ensure that $\langle f, g\rangle \geqslant \delta$ and $g \circ T^{\mathbf{p}} = g$ for all $\mathbf{p} \in \mathbb{Z}^d$, as required (since if $g \circ T^{\mathbf{p}} \neq g$ then there is some $h \in L^2(\mu)$ such that $\langle g \circ T^{\mathbf{p}} - g, h\rangle \neq 0$, but now by the definition of weak convergence this is the limit of $\langle g_i \circ T^{\mathbf{p}} - g_i, h\rangle \leqslant \|h\|_2\|g_i \circ T^{\mathbf{p}} - g_i\|_2 \longrightarrow 0$ as $i \longrightarrow \infty$, giving a contradiction).

The construction of the g_is rests on a careful re-arrangment of our initial assumption. Consider the squared norm of the averages of interest, and expand then out:

$$\Big\|\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} U_T^{\mathbf{n}} f\Big\|_2^2 = \frac{1}{N^{2d}}\sum_{\mathbf{n}\in[N]^d}\sum_{\mathbf{n}'\in[N]^d}\langle U_T^{\mathbf{n}} f, U_T^{\mathbf{n}'} f\rangle$$
$$= \frac{1}{N^{2d}}\sum_{\mathbf{n}\in[N]^d}\sum_{\mathbf{h}\in[N]^d-\mathbf{n}}\langle U_T^{\mathbf{n}} f, U_T^{\mathbf{n}+\mathbf{h}} f\rangle = \frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d}\Big\langle f, \frac{1}{N^d}\sum_{\mathbf{h}\in[N]^d-\mathbf{n}} U_T^{\mathbf{h}} f\Big\rangle,$$

where for the last equality we have used the simple fact that

$$\langle U_T^{\mathbf{n}} f, U_T^{\mathbf{n}} g\rangle = \int(\overline{f \circ T^{\mathbf{n}}}) \cdot (g \circ T^{\mathbf{n}})\,\mathrm{d}\mu = \int \overline{f} \cdot g\,\mathrm{d}\mu = \langle f, g\rangle,$$

which follows from the assumption that $T^{\mathbf{n}}$ is μ-preserving.

Thus, if this does not tend to zero as $N \longrightarrow \infty$, then there must be arbitrarily large N and selections of $\mathbf{n} \in [N]^d$ for which the value of

$$\Big\langle f, \frac{1}{N^d}\sum_{\mathbf{h}\in[N]^d-\mathbf{n}} U_T^{\mathbf{h}} f\Big\rangle$$

stays away from zero. Let N_i and $\mathbf{n}_i$ for $i \geqslant 1$ be a sequence of such selections, and let

$$g_i := \frac{1}{N_i^d}\sum_{\mathbf{h}\in[N_i]^d-\mathbf{n}_i} U_T^{\mathbf{h}} f.$$

These functions have the desired properties:

- g_i is an average of functions of norm at most 1 in $L^2(\mu)$, and so g_i itself has norm at most 1;
- the largeness of the inner product $\langle g_i, f\rangle$ was built into our choice of N_i and $\mathbf{n}_i$;
- and for any $\mathbf{p} \in \mathbb{Z}^d$ we have

$$\begin{aligned}
\|U_T^{\mathbf{p}} g_i - g_i\|_2 &= \Big\|\frac{1}{N_i^d}\sum_{\mathbf{h}\in[N_i]^d-\mathbf{n}_i} U_T^{\mathbf{h}+\mathbf{p}} f - \frac{1}{N_i^d}\sum_{\mathbf{h}\in[N_i]^d-\mathbf{n}_i} U_T^{\mathbf{h}} f\Big\|_2\\
&= \Big\|\frac{1}{N_i^d}\sum_{\mathbf{h}\in([N_i]^d-\mathbf{n}_i)\triangle([N_i]^d-\mathbf{n}_i+\mathbf{p})} U_T^{\mathbf{h}+\mathbf{p}} f\Big\|_2\\
&\leqslant \frac{|([N_i]^d-\mathbf{n}_i)\triangle([N_i]^d-\mathbf{n}_i+\mathbf{p})|}{N_i^d} = \mathrm{O}\Big(|\mathbf{p}|\frac{N_i^{d-1}}{N_i^d}\Big) \longrightarrow 0
\end{aligned}$$

as $i \longrightarrow \infty$, as required.

□

Proof of the Norm Ergodic Theorem. We first prove that if $f \in L^2(\mu)$ then the asserted convergence holds in $\|\cdot\|_2$. Let P be the orthogonal projection of $L^2(\mu)$ onto the closed subspace of T-invariant functions, so $Pf = \mathsf{E}(f \mid \Sigma^T)$, and decompose an arbitrary f as $(f - Pf) + Pf$. The first term of this decomposition, $f - Pf$, is orthogonal to all T-invariant functions, and hence by the above argument we must have

$$\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d}(f - Pf)\circ T^{\mathbf{n}} \longrightarrow 0 \qquad \text{in norm.}$$

On the other hand, Pf is T-invariant, so

$$\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d}(Pf)\circ T^{\mathbf{n}} = Pf \longrightarrow Pf \qquad \text{in norm.}$$

Putting these facts together shows that $\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f\circ T^{\mathbf{n}}$ converges in $\|\cdot\|_2$ to the T-invariant function Pf.

To turn this L^2 result into the desired L^1 result, simply recall that on the one hand $\|\cdot\|_1 \leqslant \|\cdot\|_2$ (for example, by the Cauchy-Schwartz inequality, since $\|f\|_1 = \langle 1_X, |f|\rangle$), and on the other that $L^2(\mu) \subset L^1(\mu)$ is a dense subspace for the $L^1(\mu)$ norm. In view of these facts, an arbitrary $f \in L^1(\mu)$ may be approximated in $L^1(\mu)$ by some $f' \in L^2(\mu)$, and now it follows that (a) $S_N f'$ converges to $\bar{f}'$ in

$\|\cdot\|_2$ and therefore certainly in $\|\cdot\|_1$, while (b) the difference $\|S_N f - S_N f'\|_1$ stays small as $N \longrightarrow \infty$ since each S_N is clearly a contraction on $L^1(\mu)$. Since the approximation by f' can be as good as we please, we must also have convergence of $S_N f$ in $L^1(\mu)$. $\square$

Remarks. *1.* A somewhat similar proof can be given that works directly in $L^1(\mu)$, where we don't have this inner-product trick to call on. However, this use of inner products will be more essential in the 'nonconventional' ergodic theorems we will discuss later, where it leads to a separate result by the name of the 'van der Corput estimate', and so it seems worthwhile to showcase it now.

2. Pushing the above argument only slightly further gives a norm ergodic theorem in L^p for every $p \in [0, \infty)$. More generally, however, it is an interesting and subtle question to decide for a given Banach space E and isometry $T : E \longrightarrow E$ whether the averages

$$\frac{1}{N}\sum_{n=1}^{N} T^n \xi$$

converge in norm for every $\xi \in E$. We will not explore this functional analytic question here, although one important case is covered on the second problem sheet. $\lhd$

1.3 Proof of the Pointwise Ergodic Theorem

This is rather more involved than the Norm Ergodic Theorem. The strategy is similar, but here it's easier to go straight to the $L^1(\mu)$ result.

Lemma 6. *Suppose that $f \in L^1(\mu)$ and that $A \in \Sigma$ is such that $\mu(A) > 0$ and*

$$S_N f(x) \not\longrightarrow 0$$

for all $x \in A$. Then there is a T-invariant $g \in L^\infty(\mu)$ such that $\langle f, g\rangle \neq 0$.

Obtaining the Pointwise Ergodic Theorem from this lemma is very similar to the Norm Ergodic Theorem.

Proof of Pointwise Ergodic Theorem from Lemma 6. Given $f \in L^1(\mu)$, decompose it as $(f - \mathsf{E}(f \mid \Sigma^T)) + \mathsf{E}(f \mid \Sigma^T)$; since conditional expectation is a contraction on $L^1(\mu)$, both of these summands are still in $L^1(\mu)$. Now, on the one hand, by the defining properties of conditional expectation we have $\int_X (f - \mathsf{E}(f \mid \Sigma^T)) g \, \mathrm{d}\mu = 0$ for every Σ^T-measurable $g \in L^\infty$, so that Lemma 6 implies

$$S_N(f - \mathsf{E}(f \mid \Sigma^T))(x) \longrightarrow 0 \qquad \text{for a.e. } x.$$

On the other, the function $\mathsf{E}(f \mid \Sigma^T)$ is itself T-invariant, so that $S_N(\mathsf{E}(f \mid \Sigma^T)) \equiv \mathsf{E}(f \mid \Sigma^T)$ for all N. □

It remains to prove Lemma 6. Our assumption tells us that either $\limsup_{N \longrightarrow \infty} S_N f(x) > 0$ or $\liminf_{N \longrightarrow \infty} S_N f(x) < 0$ on some positive-measure set; by replacing f with $-f$ if necessary we may assume the former. Since

$$\{x : \limsup_{N \longrightarrow \infty} S_N f(x) > 0\} = \bigcup_{m \geqslant 1} \{x : \limsup_{N \longrightarrow \infty} S_N f(x) > 1/m\}$$

and so $$\mu\{x : \limsup_{N \longrightarrow \infty} S_N f(x) > 0\} = \lim_{m \longrightarrow \infty} \mu\{x : \limsup_{N \longrightarrow \infty} S_N f(x) > 1/m\},$$

it follows that we can assume that $\limsup_{N \longrightarrow \infty} S_N f(x) > \eta$ for some $\eta > 0$ for all x in some positive-measure set A.

Similarly to Lemma 5, it will suffice to prove that there is a sequence $g_i \in L^\infty(\mu)$ with

- $\|g_i\|_\infty \leqslant 1$ for all i,
- $\langle g_i, f \rangle \geqslant \delta$ for all i for some fixed $\delta > 0$, and
- $\|g_i \circ T^{\mathbf{p}} - g_i\|_1 \longrightarrow 0$ as $i \longrightarrow \infty$ for any fixed $\mathbf{p} \in \mathbb{Z}^d$.

Given this, another weak*-compactness result will finish the proof: we have the standard identification $L^\infty(\mu) \cong L^1(\mu)^*$, and by the Banach-Alaoglu Theorem in the resulting weak* topology the unit ball of $L^\infty(\mu)$ is compact. Hence having found such a sequence $(g_i)_{i \geqslant 1}$, its boundedness ensures the existence of a weak* limit g along some subsequence, and now the second and third conditions ensure that $\langle f, g \rangle \geqslant \delta$ and $g \circ T^{\mathbf{p}} = g$ for all $\mathbf{p} \in \mathbb{Z}^d$, as required (note, in particular, that because we know the functions g_i are bounded in L^∞, the fact that they are approximately invariant only in L^1 is still enough to give a strictly invariant limit, since again if the limit were not strictly invariant then there would be some bounded function such that $\langle g_i \circ T^{\mathbf{p}} - g_i, h \rangle \neq 0$).

However, the construction of the g_i is a little more involved than for norm convergence. What we know is that for some positive-measure $A \in \Sigma$, for every $x \in A$

we can find arbitrarily large integers N_x such that

$$\frac{1}{N_x^d}\sum_{\mathbf{n}\in[N_x]^d} f(T^{\mathbf{n}}x) > \eta.$$

We would like to combine this scattered collection of large averages for the various individual points $x \in A$ into a single function g_i, which on the one hand gives a large inner product with f, and on the other is approximately invariant. The inner product with f should be estimated somehow in terms of the pointwise averages above. The approximate invariance will presumably arise because g is constructed from pieces that are constant on 'large patches' within each orbit, so that as for norm convergence we can ultimately exploit the estimate

$$\frac{|(\mathbf{p}+[M]^d)\triangle[M]^d|}{M^d} = \mathrm{O}\Big(|\mathbf{p}|\frac{M^{d-1}}{M}\Big) \longrightarrow 0 \text{ as } M \longrightarrow \infty.$$

To make something like this work in general we first need a simple (but clever) geometric result that will allow us to 'process' these large patches that we want to take from each orbit. It is a version of the well-known Vitali covering lemma.

Lemma 7 (Basic covering lemma)**.** *Suppose that $\mathcal{S}$ is a finite family of cubes in $\mathbb{Z}^d$. Then there is a subcollection $\mathcal{R} \subseteq \mathcal{S}$ such that*

- *it is pairwise disjoint:*

$$Q, Q' \in \mathcal{R} \text{ distinct} \quad \Longrightarrow \quad Q \cap Q' = \varnothing;$$

- *it covers a comparable volume to $\mathcal{S}$:*

$$\sum_{Q\in\mathcal{R}} |Q| = \Big|\bigcup \mathcal{R}\Big| \geqslant 3^{-d}\Big|\bigcup \mathcal{S}\Big|.$$

Proof. This follows by a greedy algorithm.

Begin by letting $Q_1 \in \mathcal{S}$ be cube of maximal size.

Now suppose we have already chosen $Q_1, Q_2, \dots, Q_m$ for some $m \geqslant 1$. If every other $Q' \in \mathcal{S}$ intersects Q_i for some $i \leqslant p$, then stop and let $\mathcal{R} := \{Q_1, \dots, Q_m\}$. Otherwise, let Q_{m+1} be a maximal-size member of

$$\{Q' \in \mathcal{S} :\ Q' \cap (Q_1 \cup \cdots \cup Q_m) = \varnothing\}.$$

Since $\mathcal{S}$ is finite, this algorithm must stop and give $\mathcal{R} = \{Q_1, \ldots, Q_m\}$ for some finite m. This family is pairwise disjoint by construction, so it remains to prove the volume estimate.

For any cube Q, let $3 \cdot Q$ denote the cube with the same centre as Q but three times the side length. Then for any $\mathbf{v} \in Q' \in \mathcal{S}$, either $Q' \in \mathcal{R}$, so that $\mathbf{v} \in \bigcup \mathcal{R}$; or $Q' \notin \mathcal{R}$ and so $Q' \cap Q_i \neq \varnothing$ for some $i \leqslant m$, since $\mathcal{R}$ is a maximal pairwise-disjoint family. Letting i be minimal with this property, it also follows that $|Q'| \leqslant |Q_i|$, for otherwise our algorithm would have chosen Q' instead of Q_i at step i. This now implies that $Q' \subseteq 3 \cdot Q_i$, and so overall we have shown that $\bigcup \mathcal{S} \subseteq \bigcup_{Q \in \mathcal{R}} 3 \cdot Q$. For the volumes, this gives

$$\Big|\bigcup \mathcal{S}\Big| \leqslant \Big|\bigcup_{Q \in \mathcal{R}} 3 \cdot Q\Big| \leqslant \sum_{Q \in \mathcal{R}} |3 \cdot Q| = 3^d \sum_{Q \in \mathcal{R}} = 3^d \Big|\bigcup \mathcal{R}\Big|,$$

where the last equality holds owing to the pairwise-disjointness of $\mathcal{R}$. □

Using this, we can complete the proof of Lemma 6.

Proof of Lemma 6. As already explained, it suffices to construct functions g_i that have the desired properties approximately.

If $y \in A$, then our assumption gives some $N_{y,i} > i$ such that $S_{N_{y,i}} f(y) > \eta$. The union of the sets $\{y \in A : \ N_{y,i} \leqslant M_i/2\}$ is equal to A, so now let M_i be so large that this set has measure at least $\mu(A)/2$. Now, for each $y \in X$, let its **forward box** be the subset $\{T^{\mathbf{m}} y : \ \mathbf{m} \in [M_i]^d\}$ of its orbit, and similarly let its **backward box** be $\{T^{-\mathbf{m}} y : \ \mathbf{m} \in [M_i]^d\}$.

Intuitively, if M_i is sufficiently large and $\mathbf{m} \in [M_i]^d$ is far from the boundary of this box, then we will also have

$$\mathbf{m} + [N_{T^{\mathbf{m}} y, i}]^d \subseteq [M_i]^d.$$

Let

$$K_i(y) := \big\{\mathbf{m} \in [M_i]^d : \ T^{\mathbf{m}} y \in A \text{ and } \mathbf{m} + [N_{T^{\mathbf{m}} y, i}]^d \subseteq [M_i]^d\big\} \subseteq [M_i]^d.$$

Given this, we apply the Basic Covering Lemma to the associated family of boxes

$$\mathcal{S}_{y,i} := \{\mathbf{m} + [N_{T^{\mathbf{m}} y, i}]^d : \ \mathbf{m} \in K_i(y)\}$$

to obtain a pairwise-disjoint subfamily $\mathcal{R}_i(y)$ such that

$$\Big|\bigcup \mathcal{R}_i(y)\Big| \geqslant 3^{-d} \Big|\bigcup \mathcal{S}_i(y)\Big|.$$

This, in turn, is at least $3^{-d}|K_i(y)|$, because $\bigcup \mathcal{S}_i(y)$ contains the set of points

$$\{\mathbf{m} + (1, 1, \ldots, 1) : \ \mathbf{m} \in K_i(y)\}.$$

(It is also easy to see that all these sets may be selected measurably in y; we won't worry about this here).

Finally, let

$$g_i(x) := \frac{1}{M_i^d} \sum_{\mathbf{n} \in [M_i]^d} 1_{\bigcup \mathcal{R}_i(T^{-\mathbf{n}} x)}(\mathbf{n}),$$

where $1_{\bigcup \mathcal{R}_i(T^{-\mathbf{n}} x)}$ denotes the indicator function of the union of all the rectangles $R \in \mathcal{R}_i(T^{-\mathbf{n}} x)$, so

$$1_{\bigcup \mathcal{R}_i(T^{-\mathbf{n}} x)}(\mathbf{k}) = \begin{cases} 1 & \text{if } \mathbf{k} \in \bigcup_{R \in \mathcal{R}_i(T^{-\mathbf{n}} x)} R \\ 0 & \text{else.} \end{cases}$$

This is an average of functions taking only the values 0 and 1, so certainly $0 \leqslant g_i \leqslant 1$. It remains to verify its nonzero inner product with f and its approximate invariance.

Inner product with f. This follows from some re-arrangement:

$$\begin{aligned} \langle g_i, f \rangle &= \frac{1}{M_i^d} \sum_{\mathbf{n} \in [M_i]^d} \int_X f(x) \cdot 1_{\bigcup \mathcal{R}_i(T^{-\mathbf{n}} x)}(\mathbf{n}) \, \mu(\mathrm{d}x) \\ &= \frac{1}{M_i^d} \sum_{\mathbf{n} \in [M_i]^d} \int_X f(T^{\mathbf{n}} y) \cdot 1_{\bigcup \mathcal{R}_i(y)}(\mathbf{n}) \, \mu(\mathrm{d}y) \quad \text{(change of vars)} \\ &= \int_X \frac{1}{M_i^d} \sum_{\mathbf{n} \in \bigcup \mathcal{R}_i(y)} f(T^{\mathbf{n}} y) \, \mu(\mathrm{d}y) \\ &\geqslant \int_X \frac{1}{M_i^d} \eta \Big| \bigcup \mathcal{R}_i(y) \Big| \, \mu(\mathrm{d}y), \end{aligned}$$

because by construction each of the pairwise-disjoint boxes in the family $\mathcal{R}_i(y)$ is such that the resulting average of f over that box is at least η. Crucially, the covering lemma has given us *disjoint* boxes, so summing over $\bigcup \mathcal{R}_i(y)$ is equivalent to summing over each individual box and then adding the results (there is no double-counting). Moreover, the covering lemma promises us that

$$\Big| \bigcup \mathcal{R}_i(y) \Big| = 3^{-d} |K_i(y)|,$$

so the above integral is at least

$$\int_X \frac{3^{-d}\eta|K_i(y)|}{M_i^d}\,\mu(\mathrm{d}y).$$

To turn this into a fixed lower bound, we need one more re-arrangement. Recalling the definition of $K_i(y)$, applying a crude lower bound and interchanging an integral and sum we obtain that the above equals

$$\begin{aligned}
&\int_X \frac{3^{-d}\eta}{M_i^d}\big|\big\{\mathbf{m}\in[M_i]^d :\ T^{\mathbf{m}}y\in A \text{ and } \mathbf{m}+[N_{T^{\mathbf{m}}y,i}]^d\subseteq[M_i]^d\big\}\big|\,\mu(\mathrm{d}y)\\
&\geqslant \int_X \frac{3^{-d}\eta}{M_i^d}\big|\big\{\mathbf{m}\in[M_i/2]^d :\ T^{\mathbf{m}}y\in A \text{ and } N_{T^{\mathbf{m}}y,i}\leqslant M_i/2\big\}\big|\,\mu(\mathrm{d}y)\\
&= \int_X \frac{3^{-d}\eta}{M_i^d}\sum_{\mathbf{m}\in[M_i/2]^d} 1_{\{z\in A:\ N_{z,i}\leqslant M_i/2\}}(T^{\mathbf{m}}y)\,\mu(\mathrm{d}y)\\
&= \frac{\eta 3^{-d}}{M_i^d}(M_i/2)^d\mu(\{y\in A:\ N_{y,i}\leqslant M_i/2\})\geqslant \eta 6^{-d}\mu(A)/2,
\end{aligned}$$

by our initial choice of M_i. This is positive and independent of i, as required.

Approximate invariance Finally, the approximate invariance of g_i follows by observing that

$$\begin{aligned}
&|g_i\circ T^{\mathbf{p}}(x)-g_i(x)|\\
&= \Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in[M_i]^d}\big(1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}+\mathbf{p}}x)}(\mathbf{n})-1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\big)\Big|\\
&= \Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in-\mathbf{p}+[M_i]^d}1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n}+\mathbf{p})-\frac{1}{M_i^d}\sum_{\mathbf{n}\in[M_i]^d}1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\big)\Big|\\
&\leqslant \Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in(-\mathbf{p}+[M_i]^d)\triangle[M_i]^d}1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\Big|\\
&\qquad+\Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in-\mathbf{p}+[M_i]^d}\big(1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n}+\mathbf{p})-1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\big)\Big|\\
&= \Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in(-\mathbf{p}+[M_i]^d)\triangle[M_i]^d}1_{\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\Big|\\
&\qquad+\Big|\frac{1}{M_i^d}\sum_{\mathbf{n}\in-\mathbf{p}+[M_i]^d}\big(1_{(\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)-\mathbf{p})\triangle\bigcup\mathcal{R}_i(T^{-\mathbf{n}}x)}(\mathbf{n})\big)\Big|:
\end{aligned}$$

the first of these terms tends to zero as $M_i \longrightarrow \infty$ uniformly in x because

$$\frac{|(-\mathbf{p}+[M_i]^d)\triangle[M_i]^d|}{M_i^d} = \mathrm{O}\Big(|\mathbf{p}|\frac{M_i^{d-1}}{M_i^d}\Big),$$

and the second has integral in x that tends to zero by taking the sum over $\mathbf{n}$ outside the integral, changing variables $T^{-\mathbf{n}}x \mapsto x$ for each $\mathbf{n}$ separately and then applying a similar estimate to each of the individual rectangles in the families $\mathcal{R}_i(x)$. □

1.4 Ergodic Theorems for amenable groups

Both the Norm and Pointwise Ergodic Theorems can be extended to cover the action of any group Γ that is **amenable**. Heuristically, this means that Γ contains 'approximately invariant' finite sets: to be precise, there is a sequence of finite subsets $F_1 \subseteq F_2 \subseteq \ldots \subseteq \Gamma$ such that $\bigcup_{N\geqslant 1} F_N = \Gamma$ and also such that for any fixed $\gamma \in \Gamma$ we have

$$\frac{|F_N \triangle F_N\gamma|}{|F_N|} \longrightarrow 0 \qquad \text{as } N \longrightarrow \infty.$$

This is called a **Følner sequence** of subsets of Γ. Given such a Følner sequence, a p.-p.s. $T : \Gamma \curvearrowright (X, \Sigma, \mu)$, and a function $f \in L^1(\mu)$, we can ask whether

$$S_N f := \frac{1}{|F_N|}\sum_{\gamma\in F_N} f \circ T^\gamma$$

converges in norm or pointwise. Norm convergence always holds, and can be proved similarly to our treatment of $\mathbb{Z}^d$. Pointwise convergence is again much harder, and requires some additional conditions on the Følner sequence. However, it is now known that for every amenable group and Følner sequence, there is some Følner subsequence along these the above averages convergen pointwise a.e.. This deep result was proved only recently by Linderstauss [Lin01].

All Abelian groups such as $\mathbb{Z}^d$ are amenable (for example, take $F_N := [N]^d$), and there are also many non-Abelian examples. However, there are also many non-amenable groups, such as free groups on two or more generators, and for these there are 'obvious' versions of the ergodic theorem that can fail (see the exercises).

References

[Lin01] Elon Lindenstrauss. Pointwise theorems for amenable groups. *Invent. Math.*, 146(2):259–295, 2001.

[Pet83] Karl E. Petersen. *Ergodic Theory*. Cambridge University Press, Cambridge, 1983.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: tim@cims.nyu.edu
URL: cims.nyu.edu/˜tim

Ergodic Theory

Notes 3: Topological systems and equidistribution

Having proved the basic ergodic theorems for $\mathbb{Z}^d$-systems, here we will look at some more concrete examples, and see how those theorems give new insights into their behaviour. As an application, we will show how Weyl's classical Polynomial Equidistribution Theorem can be deduced from an analysis of a special dynamical system.

1 Invariant measures for topological systems

Among dynamical systems that arise in science or in other parts of mathematics, it is often more natural to endow them with a topological structure than a purely measure-theoretic one. For example, many real-world dynamical systems are modeled by flows on manifolds, where the topology is obvious but the presence of an invariant measure may be less so.

Definition 1. *If Γ is a countable group, then a* ***topological*** Γ***-system*** *is a pair (X,T) in which X is a compact metric space and T is an action of Γ on X by homeomorphisms.*

Note that the compactness of X is a part of this definition.

Happily, for a topological $\mathbb{Z}^d$-system on a compact metric space — enough generality to cover a great many applications — one can always add some useful measure-theoretic structure. This is contained in the following simple but important theorem.

Theorem 2. *If (X,T) is a topological $\mathbb{Z}^d$-system, then X carries at least one T-invariant Borel probability measure.*

This was proved for single transformations (that is, $d = 1$) by Krylov and Bogolyubov in [KB37]. It actually holds for general amenable groups with essentially the same proof as theirs; once again, I have adopted the setting $\mathbb{Z}^d$-actions as a middle road.

Proof. Let x be any point of X, let δ_x the point-mass at x, and form the finite averages

$$\mu_N := \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} \delta_{T^{\mathbf{n}}x}.$$

Each μ_N is a Borel probability measure on X supported on (at most) N^d points. Moreover, these measures are approximately invariant as $N \longrightarrow \infty$. Indeed, for any fixed $\mathbf{k} \in \mathbb{Z}^d$ we have

$$\begin{aligned} T_*^{\mathbf{k}}\mu_N - \mu_N &= \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} T_*^{\mathbf{k}}\delta_{T^{\mathbf{n}}x} - \mu_N \\ &= \frac{1}{N^d} \sum_{\mathbf{n}\in\mathbf{k}+[N]^d} \delta_{T^{\mathbf{n}}x} - \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} \delta_{T^{\mathbf{n}}x} \\ &= \frac{1}{N^d} \sum_{\mathbf{n}\in(\mathbf{k}+[N]^d)\triangle[N]^d} \delta_{T^{\mathbf{n}}x}. \end{aligned}$$

After this cancellation, this is a sum of $|(\mathbf{k} + [N]^d)\triangle[N]^d| = \mathrm{O}(|\mathbf{k}|N^{d-1})$ point masses, and it is normalized by N^d. Therefore, for any $f \in C(X)$, we have

$$\Big|\int f\ \mathrm{d}(T_*^{\mathbf{k}}\mu_N - \mu_N)\Big| \leqslant \frac{\|f\|_\infty|\mathbf{k}|}{N} \longrightarrow 0$$

as $N \longrightarrow \infty$.

Since $\mathrm{Pr}(X)$ is sequentially compact in the vague topology, we may extract a vaguely convergent subsequence $(\mu_{N_m})_{m\geqslant 1}$. Letting μ be its limit, this is still a Borel probability measure on X, and it must now satisfy

$$\int f\ \mathrm{d}(T_*^{\mathbf{k}}\mu - \mu) = \lim_{m\longrightarrow\infty} \int f\ \mathrm{d}(T_*^{\mathbf{k}}\mu_{N_m} - \mu_{N_m}) = 0$$

for all $f \in C(X)$, so μ is T-invariant, as required. □

With a little more work, we can enhance the above theorem to obtain an *ergodic* invariant measure. We will not have an immediate use for this, and it does require another piece of background from Banach space theory, but it seems worth recording now.

They key is to consider the structure of the set of *all* invariant measures.

Lemma 3. *Suppose that* (X, T) *is a topological* $\mathbb{Z}^d$*-system and let* $\mathrm{Pr}^T(X) \subseteq \mathrm{Pr}(X)$ *denote the set of* T*-invariant Borel probability measures on* X*. Then* $\mathrm{Pr}^T(X)$ *is a nonempty convex set, and it is compact for the vague topology.*

Proof. We have just proved that $\mathrm{Pr}^T(X)$ is nonempty, and convexity follows easily: if $\mu_1, \mu_2 \in \mathrm{Pr}^T(X)$ and $0 \leqslant \alpha \leqslant 1$ then $\alpha\mu_1 + (1-\alpha)\mu_2$ satisfies:

$$\begin{aligned}&(\alpha\mu_1 + (1-\alpha)\mu_2)(T^{\mathbf{k}}A) = \alpha\mu_1(T^{\mathbf{k}}A) + (1-\alpha)\mu_2(T^{\mathbf{k}}A)\\ &= \alpha\mu_1(A) + (1-\alpha)\mu_2(A) = (\alpha\mu_1 + (1-\alpha)\mu_2)(A) \quad \forall A \in \mathcal{B}(X),\ \mathbf{k} \in \mathbb{Z}^d.\end{aligned}$$

To prove vague compactness, since $\mathrm{Pr}(X)$ is vaguely compact, it suffices to show that $\mathrm{Pr}^T(X)$ is vaguely closed in $\mathrm{Pr}(X)$. However, if $\mu_n \in \mathrm{Pr}^T(X)$ converges vaguely to $\mu \in \mathrm{Pr}(X)$, then for any $f \in C(X)$ and $\mathbf{k} \in \mathbb{Z}^d$ we have

$$\int f \circ T^{\mathbf{k}}\,\mathrm{d}\mu = \lim_{n\longrightarrow\infty} \int f \circ T^{\mathbf{k}}\,\mathrm{d}\mu_n = \lim_{n\longrightarrow\infty} \int f\,\mathrm{d}\mu_n = \int f\,\mathrm{d}\mu.$$

Hence $T^{\mathbf{k}}_*\mu = \mu$, and so $\mu \in \mathrm{Pr}^T(X)$, as required. □

This proposition puts us into position to apply another classical result about Banach spaces. Recall that, if E is a Banach space and $K \subset E$ is a closed convex subset, then $\xi \in K$ is an **extreme point** if it cannot be decomposed as

$$\alpha\xi_1 + (1-\alpha)\xi_2$$

with $0 < \alpha < 1$ and distinct $\xi_1, \xi_2 \in K$. Less formally, this asserts that 'ξ is not an internal point of any line segment lying wholly in K'. For instance, if K is a closed disk in $\mathbb{R}^2$, then its extreme points are simply its boundary points; if it is a closed square, then its extreme points are the four corners.

In this setting, a weak version of the Krein-Milman Theorem says that if, in addition, K is compact for a strong, weak or (in case E is a dual Banach space) weak* topology on E, then it must have some extreme points. This matters to us because of the following.

Lemma 4. *A probability measure $\mu \in \mathrm{Pr}^T(X)$ is an extreme point of $\mathrm{Pr}^T(X)$ if and only if it is ergodic.*

Proof. $(\Longrightarrow)$ If $\mu \in \mathrm{Pr}^T(X)$ is not ergodic, then there is some invariant set $A \in \mathcal{B}(X)$ with $0 < \mu(A) < 1$. Let $\alpha := \mu(A)$ and define

$$\mu_1(B) := \frac{1}{\mu(A)}\mu(B \cap A), \quad \mu_2(B) := \frac{1}{\mu(X\backslash A)}\mu(B \cap (X\backslash A)) \quad \forall B \in \mathcal{B}(X).$$

Each of these is a probability measure, and the invariance of A implies that each is T-invariant. They are distinct, since $\mu_1(A) = 1 \neq 0 = \mu_2(A)$, and for any $B \in \mathcal{B}(X)$ we have

$$\mu(B) = \mu(B \cap A) + \mu(B \cap (X\backslash A)) = \alpha\mu_1(A) + (1-\alpha)\mu_2(A).$$

So $\mu = \alpha\mu_1 + (1-\alpha)\mu_2$ and μ is not an extreme point.

($\Longleftarrow$) Suppose that $\mu \in \Pr^T(X)$ is ergodic and that it decomposes as $\mu = \alpha\mu_1 + (1-\alpha)\mu_2$; we will show that this can happen only if $\mu_1 = \mu_2 = \mu$. If $\mu(B) = 0$ for some $B \in \mathcal{B}(X)$, then we must also have $\mu_1(B) = \mu_2(B) = 0$, and so μ_1, μ_2 are both absolutely continuous with respect to μ. Therefore, by the Radon-Nikodým Theorem, there are $f_1, f_2 \in L^1(\mu)$ such that $f_i \geqslant 0$, $\int_X f_i \,\mathrm{d}\mu = 1$, and

$$\mu_i(B) = \int_B f_i \,\mathrm{d}\mu \quad \text{for } B \in \mathcal{B}(X).$$

Now the assumed T-invariance of each μ_i implies that

$$\int_B f_i \circ T^{\mathbf{k}} \,\mathrm{d}\mu = \int_{T^{-\mathbf{k}}B} f_i \,\mathrm{d}\mu = \int_B f_i \,\mathrm{d}\mu \quad \forall B \in \mathcal{B}(X),$$

and hence that $f_i \circ T^{\mathbf{k}} = f_i$. Since μ is ergodic, this implies that both f_is are μ-a.s. constant. Now the condition that $\int_X f_i \,\mathrm{d}\mu = 1$ implies that their constant values are 1, and hence that $\mu_i = \mu$ for $i = 1, 2$, as required. □

Corollary 5. *If (X, T) is a topological $\mathbb{Z}^d$-system, then X carries at least one* ergodic *T-invariant Borel probability measure.*

Proof. Theorem 2 gives that $\Pr^T(X)$ is nonempty, and Lemma 3 gives that it is convex and vaguely (= weak*) compact. The Krein-Milman Theorem therefore applies to give some extreme point $\mu \in \Pr^T(X)$, and now this is ergodic by the preceding lemma. □

Definition 6. *The set of T-ergodic (equivalently, extreme) measures in $\Pr^T(X)$ is denoted by $\mathcal{E}(\Pr^T(X))$.*

Remark. In fact, the above line of reasoning can be taken much further. After a slightly delicate argument to show that $\mathcal{E}(\Pr^T(X))$ is a G_δ-subset of $\Pr^T(X)$, one can apply Chôquet's Theorem to deduce that any $\mu \in \Pr^T(X)$ has a decomposition as an integral of ergodic invariant measures:

$$\mu = \int_\Omega \mu_\omega \,\nu(\mathrm{d}\omega),$$

where $(\Omega, \mathcal{F}, \nu)$ is some auxiliary probability space and the map $\Omega \longrightarrow \mathcal{E}(\Pr^T(X)) : \omega \mapsto \mu_\omega$ is measurable in a suitable sense. This is the first step towards proving the general 'ergodic decomposition' of an arbitrary p.p. system. We will return to this decomposition from a different point of view later in the course. ◁

2 Generic points and unique ergodicity

Having found an invariant measure, the ergodic theorems give us a relation with the ergodic averages of functions over orbits starting from individual points. In order to discuss these we make the following definition.

Definition 7 (Generic points)**.** *If (X,T) is a topological $\mathbb{Z}^d$-system and $\mu \in \mathrm{Pr}^T(X)$, then a point $x \in X$ is **generic** for μ if*

$$\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f(T^{\mathbf{n}}x) \longrightarrow \int f\,\mathrm{d}\mu \tag{1}$$

for every $f \in C(X)$. *(Thus, generic points are those that are 'good' for the Pointwise Ergodic Theorem for all continuous functions.)*

The Pointwise Ergodic Theorem tells us that for any one $f \in C(X)$, the convergence in (1) holds for μ-almost every $x \in X$. However, it seems to be asking much more that a point should give this convergence for *every* continuous function. Nevertheless, if μ is ergodic, then there are always plenty of generic points.

Lemma 8. *If (X,T) is a topological $\mathbb{Z}^d$-system and $\mu \in \mathcal{E}(\mathrm{Pr}^T(X))$, then the set of generic points for μ is measurable and has full μ-measure.*

Proof. First, recall that if X is compact with metric ρ, then the Banach space $C(X)$ is separable. This follows, for example, by first letting $(x_n)_{n\geqslant 1}$ be a countable dense sequence in X, and then applying the Stone-Weierstrass Theorem to the algebra of rational polynomial combinations of the functions

$$x \mapsto \rho(x, x_n).$$

Given this, let $(f_n)_{n\geqslant 1}$ be a countable dense sequence in $C(X)$, and, for each n, let

$$X_n := \Big\{x \in X : \frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f_n(T^{\mathbf{n}}x) \longrightarrow \int f_n\,\mathrm{d}\mu\Big\}.$$

An easy exercise shows that this is measurable, and the Pointwise Ergodic Theorem proves that it has full μ-measure. The countable intersection $Y := \bigcap_{n\geqslant 1} X_n$ still has full μ-measure, and any generic point lies in Y. To finish the proof, we will show that if $x \in Y$ then it is generic. Indeed, for any $f \in C(X)$ and $\varepsilon > 0$, there is some f_n with $\|f - f_n\|_\infty < \varepsilon/3$, and now, since $x \in X_n$, there is some N_0 such that

$$
\begin{aligned}
N \geqslant N_0 \quad &\Longrightarrow \quad \Big|\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f_n(T^{\mathbf{n}}x) - \int f_n\,\mathrm{d}\mu\Big| < \varepsilon/3 \\
&\Longrightarrow \quad \Big|\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f(T^{\mathbf{n}}x) - \int f\,\mathrm{d}\mu\Big| \\
&\qquad\qquad \leqslant \Big|\frac{1}{N^d}\sum_{\mathbf{n}\in[N]^d} f_n(T^{\mathbf{n}}x) - \int f_n\,\mathrm{d}\mu\Big| + 2\|f - f_n\|_\infty < \varepsilon.
\end{aligned}
$$

Since ε was arbitrary, this is the desired convergence. □

A large part of modern dynamics is given over to understanding how additional information about the topological system (X, T) constrains the possible invariant measures that it can carry. The simplest kind of constraint is that this measure should be unique.

Definition 9 (Unique ergodicity)**.** *A topological $\mathbb{Z}^d$-system (X, T) is **uniquely ergodic** if it has exactly one T-invariant measure.*

Before turning to examples and applications, we give the important reformulation of unique ergodicity in terms of generic points. In general, a topological system with invariant measure μ may still have a nonempty μ-negligible subset of points that are not generic for μ. However, it turns out that this always implies the existence of another invariant measure on the space, as expressed in the following result.

Theorem 10 (Generic points and unique ergodicity)**.** *For a topological $\mathbb{Z}^d$-system (X, T) with an invariant measure μ, the following are equivalent:*

- (X, T) *is uniquely ergodic;*
- strictly every $x \in X$ *is generic for* μ.

Proof. $(\Longrightarrow)$ This follows by re-applying the Krylov-Bogolyubov argument. If any $x \in X$ were not generic for μ, then there would be some $f \in C(X)$ and some integers $N_1 < N_2 < \dots$ such that

$$\frac{1}{N_m^d}\sum_{\mathbf{n}\in[N_m]^d} f(T^{\mathbf{n}}x) \longrightarrow \alpha \neq \int f\,\mathrm{d}\mu \quad \text{as } m \longrightarrow \infty.$$

By passing to a further subsequence if necessary, we could again assume that the sequence of measures

$$\mu_{N_m} := \frac{1}{N_m^d}\sum_{\mathbf{n}\in[N_m]^d} \delta_{T^{\mathbf{n}}x}$$

is vaguely convergent to some $\nu \in \Pr(X)$, which would therefore be T-invariant, by the same argument as before. On the other hand, it would also satisfy $\int f\,\mathrm{d}\nu = \alpha \neq \int f\,\mathrm{d}\mu$, and hence $\nu \neq \mu$. This would contradict the uniqueness of μ among T-invariant probability measures.

($\Longleftarrow$) If ν is any T-invariant probability measure, then we trivially have

$$\int f\,\mathrm{d}\nu = \lim_{N\longrightarrow\infty} \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} \Big(\int f(T^{\mathbf{n}}x)\,\nu(\mathrm{d}x)\Big)$$
$$= \int \Big(\lim_{N\longrightarrow\infty} \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} f(T^{\mathbf{n}}x)\Big)\,\nu(\mathrm{d}x).$$

However, if strictly every $x \in X$ is generic for μ, then the limit inside the last integral above converges to $\int f\,\mathrm{d}\mu$ for every x, and so, by Lebesgue's Dominated Convergence Theorem, we obtain $\int f\,\mathrm{d}\nu = \int f\,\mathrm{d}\mu$. Since f was arbitrary, this implies $\nu = \mu$. $\square$

3 Compact group rotations, generalizations and applications to equidistribution

We now turn to our application: Furstenberg's dynamical proof of the following famous number-theoretic result:

Theorem 11 (Weyl's Polynomial Equidistribution Theorem)**.** *Suppose that $p(X)$ is a real polynomial for which at least one coefficient other than the constant term is irrational. Then the sequence of fractional parts*

$$\{p(n)\}, \quad n = 1, 2, \ldots$$

is ***equidistributed modulo*** 1*: that is, for any $0 \leqslant a < b \leqslant 1$ we have*

$$\frac{|\{1 \leqslant n \leqslant N :\ a \leqslant \{p(n)\} < b\}|}{N} \longrightarrow b - a \qquad \text{as } N \longrightarrow \infty.$$

In fact, we will prove only the special case of Theorem 11 in which the leading coefficient of p is irrational, leaving the extension to the general case as an exercise.

The treatment below is essentially taken from Section 3.3 in [Fur81]. The crucial steps involve a study of unique ergodicity for a related topological system. The full system of interest will be revealed shortly, but first we establish unique ergodicity in a simpler example, which will then be used as a building block.

The acting group will simply be $\mathbb{Z}$ throughout this subsection. Also, if $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathbb{T}^d$ and $\mathbf{m} = (m_1, \ldots, m_d) \in \mathbb{Z}^d$, then we write

$$\mathbf{m} \bullet \alpha := m_1\alpha_1 + m_2\alpha_2 + \ldots + m_d\alpha_d,$$

an element of $\mathbb{T}$.

Proposition 12. *Suppose that $\alpha_1, \alpha_2, \ldots, \alpha_d \in \mathbb{T} = \mathbb{R}/\mathbb{Z}$ are **irrational and rationally independent**, meaning that for integers m_1, m_2, ..., m_d we have*

$$m_1\alpha_1 + m_2\alpha_2 + \ldots + m_d\alpha_d = 0 \mod 1 \quad \textit{only if} \quad m_1 = m_2 = \ldots = m_d = 0.$$

Let $T : \mathbb{T}^d \longrightarrow \mathbb{T}^d$ be the associated rotation transformation:

$$T(t_1, t_2, \ldots, t_d) := (t_1 + \alpha_1, t_2 + \alpha_2, \ldots, t_d + \alpha_d).$$

Let m be the Haar probability measure (that is, Lebesgue measure) on $\mathbb{T}^d$. Then m is T-invariant, and T is uniquely ergodic.

Proof. Let $\alpha = (\alpha_1, \ldots, \alpha_d)$, so that T is the rotation of $\mathbb{T}^d$ by α. For this system, it is easier to analyse the pointwise averages of a function $f \in C(\mathbb{T}^d)$ than to study some unknown invariant measure, so we will already see the usefulness of Theorem 10. By that theorem, it suffices to prove that for any $f \in C(\mathbb{T}^d)$ and $t = (t_1, \ldots, t_d) \in \mathbb{T}^d$ we have

$$\frac{1}{N}\sum_{n=1}^{N} f(t + n\alpha) \longrightarrow \int_{\mathbb{T}^d} f(s)\,\mathrm{d}s,$$

where $\int \cdot\,\mathrm{d}s$ denotes integration with respect to m.

By the Weierstrass Approximation Theorem, any $f \in C(\mathbb{T}^d)$ may be uniformly approximated by trigonometric polynomials of the form

$$g(t_1, \ldots, t_d) = \sum_{j=1}^{p} c_j \mathrm{e}^{2\pi\mathrm{i}(m_{j,1}t_1 + \cdots + m_{j,d}t_d)} = \sum_{j=1}^{p} c_j \mathrm{e}^{2\pi\mathrm{i}\mathbf{m}_j \bullet t}$$

for some coefficients $c_j \in \mathbb{C}$ and integers $m_{j,r}$, $1 \leqslant j \leqslant p$, $1 \leqslant r \leqslant d$. Clearly we may assume that the integer vectors $\mathbf{m}_j := (m_{j,1}, \ldots, m_{j,d})$ are all distinct. Also, by re-labeling and introducing a zero term if necessary, we may assume that $\mathbf{m}_1 = \mathbf{0}$. It therefore suffices to prove the above convergence for such functions g.

This now follows by elementary calculations. Substituting this formula for g into the pointwise averages starting from t and re-arranging gives

$$\frac{1}{N}\sum_{n=1}^{N} g(t+n\alpha) = \sum_{j=1}^{p} c_j\Big(\frac{1}{N}\sum_{n=1}^{N}\mathrm{e}^{2\pi\mathrm{i}(\mathbf{m}_j\bullet t+n(\mathbf{m}_j\bullet\alpha))}\Big)$$
$$= \sum_{j=1}^{p} c_j\mathrm{e}^{2\pi\mathrm{i}\mathbf{m}_j\bullet t}\Big(\frac{1}{N}\sum_{n=1}^{N}(\mathrm{e}^{2\pi\mathrm{i}\mathbf{m}_j\bullet\alpha})^n\Big).$$

These last inner averages may now be evaluated as geometric series. As $N\longrightarrow\infty$, they behave as $\mathrm{O}(\frac{1}{N}\frac{1}{|1-\exp(2\pi\mathrm{i}\mathbf{m}_j\bullet\alpha)|})$, and so tend to 0 unless $\mathbf{m}_j\bullet\alpha = 0$. By our assumptions on α, this is possible only if $\mathbf{m}_j = \mathbf{0}$, and therefore

$$\frac{1}{N}\sum_{n=1}^{N} g(t+n\alpha)\longrightarrow c_1 \quad \text{as } N\longrightarrow\infty$$

for strictly every $t\in\mathbb{T}^d$. On the other hand,

$$\int_{\mathbb{T}^d} g(s)\,\mathrm{d}s = \sum_{j=1}^{p} c_j\int_{\mathbb{T}^d}\mathrm{e}^{2\pi\mathrm{i}\mathbf{m}_j\bullet s}\,\mathrm{d}s = c_1,$$

since

$$\mathbf{m}_j\neq\mathbf{0} \quad\Longrightarrow\quad \int_{\mathbb{T}^d}\mathrm{e}^{2\pi\mathrm{i}\mathbf{m}_j\bullet s}\,\mathrm{d}s = 0.$$

□

This gives our first instance of unique ergodicity. To recover the full Polynomial Equidistribution Theorem, we will need a significant generalization of the above result. First we need another definition, which has a wider importance throughout ergodic theory.

Definition 13 (Skew-product transformation)**.** *Suppose that (Y,Φ) is a measurable space, that $T \curvearrowright Y$ is an invertible measurable transformation, that G is a compact metric group, and that $\sigma : Y\longrightarrow G$ is a measurable function. Then the* ***skew product of*** *T* ***by*** *σ is the transformation $T\ltimes\sigma \curvearrowright Y\times G$ defined by*

$$(T\ltimes\sigma)(y,g) := (Ty,\sigma(y)g).$$

The function σ is referred to as the ***cocycle*** *of this transformation (for reasons which relate to the extension of this definition to actions of more complicated groups).*

This new system is the ***topological skew product*** *if (Y,T) is a topological system and σ is a continuous function.*

A quick calculation now gives that the map $T \ltimes \sigma$ is a measurable transformation of $Y \times G$ with inverse given by

$$(T \ltimes \sigma)^{-1} = T^{-1} \ltimes \sigma^{(-1)},$$

where $\sigma^{(-1)}(y) := \sigma(T^{-1}y)^{-1}$, and that the higher iterates of $T \ltimes \sigma$ are given by

$$(T \ltimes \sigma)^n = T^n \ltimes \sigma^{(n)},$$

where

$$\sigma^{(n)}(y) := \sigma(T^{n-1}y) \cdot \sigma(T^{n-2}y) \cdot \cdots \cdot \sigma(y).$$

If T preserves a probability ν on Y, then $T \ltimes \sigma$ preserves the product measure $\nu \otimes m$, where m is the Haar probability measure on G. If $T \ltimes \sigma$ is a topological skew product, then it defines a homeomorphism of $Y \times G$.

A nice way to think about skew-products is as a generalization of group rotations so that they are 'relative' to another base system $(Y, \mathcal{B}(Y), \nu, T)$. The transformation $T \ltimes \sigma$ acts as a group rotation on each vertical fibre of $Y \times G$, but it also (i) moves fibres around, according to the movement of the points of Y under T, and (ii) rotates the fibres over different base points $y \in Y$ by different rotations, as given by the function $\sigma(y)$.

With this in mind, it is possibly not surprising that there is a 'relative' variant of Proposition 12, too.

Proposition 14. *Suppose that (Y, T) is a uniquely ergodic topological system with invariant probability ν, that $\sigma : Y \longrightarrow G$ is continuous, and that $\nu \otimes m$ is ergodic for the skew product $T \ltimes \sigma$. Then in fact the skew product is uniquely ergodic.*

Proof. *Step 1* If $(y, g) \in Y \times G$ is generic for $\nu \otimes m$, then so are all the 'fibrewise rotates' (y, gh). To see this, let $f \in C(X)$, and define

$$f_h(y, g) := f(y, gh).$$

This is still an element of $C(X)$, so the genericity of (y, g) for $\nu \otimes m$ gives

$$\begin{aligned}
\frac{1}{N}\sum_{n=1}^{N} f(T^n, \sigma^{(n)}(y)gh) &= \frac{1}{N}\sum_{n=1}^{N} f_h(T^n y, \sigma^{(n)}(y)g) \\
&\longrightarrow \int_{Y\times G} f_h(y', g')\,\nu(\mathrm{d}y')\,m(\mathrm{d}g') \\
&= \int_{Y\times G} f(y', g'h)\,\nu(\mathrm{d}y')\,m(\mathrm{d}g') \\
&= \int_{Y\times G} f(y', g')\,\nu(\mathrm{d}y')\,m(\mathrm{d}g'),
\end{aligned}$$

where the last equality holds because the measure m is rotation-invariant on G, and so integrating a function of $g'h$ over $m(\mathrm{d}g')$ is the same is simply intregrating the same function of g' over $m(\mathrm{d}g')$.

Step 2 The above shows that the set of points in $Y \times G$ that are generic for $\mu := \nu \otimes m$ does not depend on the second coordinate, so we may write this set as $A \times G$ for some measurable $A \subseteq Y$. On the other hand, since $\nu \otimes m$ is ergodic, almost every point is generic for it, and so we must have $\nu(A) = 1$.

Suppose that μ' is a $(T \ltimes \sigma)$-invariant probability measure; we must show that $\mu' = \mu$. The projection of μ' onto Y is certainly T-invariant, so since T is uniquely ergodic this projection of μ' must equal ν. This implies that $\mu'(A \times G) = \nu(A) = 1$, and hence that μ'-a.e. point is actually generic for μ. Now we applying the Pointwise Ergodic Theorem and the Lebesgue Dominated Convergence Theorem as in the proof of Proposition 12: for any $f \in C(Y \times G)$, we have

$$\begin{aligned}
\int_{Y\times G} f\,\mathrm{d}\mu' &= \int_{Y\times G} \Big(\lim_{N\longrightarrow\infty} \frac{1}{N}\sum_{n=1}^{N} f((T\ltimes\sigma)^n(y,g))\Big)\,\mu'(\mathrm{d}y,\mathrm{d}g)\\
&= \int_{Y\times G} \Big(\int_{Y\times G} f\,\mathrm{d}\mu\Big)\,\mu'(\mathrm{d}y,\mathrm{d}g) \quad \text{(by genericity for } \mu)\\
&= \int_{Y\times G} f\,\mathrm{d}\mu.
\end{aligned}$$

Hence that $\mu' = \mu$, as required. □

In view of this proposition, we can now obtain many more examples of uniquely ergodic topological systems by iterating the skew-product construction. Provided we start with a base system that is uniquely ergodic (such as one of the torus rotations of Proposition 12), at each subsequent step we need only show that the product measure is ergodic, and then the above proposition promises that it is uniquely ergodic. The following example is of this kind.

Corollary 15. *Suppose that α is irrational, and let $T \curvearrowright \mathbb{T}^d$ be given by*

$$T(t_1, t_2, \ldots, t_d) := (t_1 + \alpha, t_2 + t_1, \ldots, t_d + t_{d-1}).$$

Then the Haar measure m_d on $\mathbb{T}^d$ is invariant and uniquely ergodic for T.

Proof. The proof is by induction on d. For the base case $d = 1$ this is simply a one-dimensional special case of Proposition 12, so suppose now that $d > 1$ and that the result has already been proved for $d - 1$.

In this case we can write $T = S \ltimes \sigma$, where $S : \mathbb{T}^{d-1} \longrightarrow \mathbb{T}^{d-1}$ is the analogous transformation of one dimension less, and

$$\sigma(t_1, t_2, \ldots, t_{d-1}) := t_{d-1}.$$

By the inductive hypothesis, we know that m_{d-1} is uniquely ergodic for S, so by Proposition 14 we need only prove that $m_{d-1} \otimes m_1 = m_d$ is ergodic (rather than uniquely ergodic) for T.

To this end, suppose that $f = f \circ T$ for some $f \in L^2(m_d)$; we will show that f must be constant m_d-a.s. This can be done using basic Fourier analysis on the torus $\mathbb{T}^d$. Expanding f as a Fourier series gives

$$f(t) := \sum_{\mathbf{m} \in \mathbb{Z}^d} a_{\mathbf{m}} \mathrm{e}^{2\pi \mathrm{i} \mathbf{m} \bullet t},$$

and hence

$$f(Tt) := \sum_{\mathbf{m} \in \mathbb{Z}^d} a_{\mathbf{m}} \mathrm{e}^{2\pi \mathrm{i} \mathbf{m} \bullet Tt}.$$

Now a simple calculation gives

$$\begin{aligned} &\mathbf{m} \bullet Tt = m_1\alpha + m_1 t_1 + m_2(t_1 + t_2) + \cdots + m_d(t_{d-1} + t_d) \\ &= m_1\alpha + (m_1 + m_2)t_1 + (m_2 + m_3)t_2 + \cdots + (m_{d-1} + m_d)t_{d-1} + m_d t_d, \end{aligned}$$

and so equating the Fourier series for f and $f \circ T$ gives

$$\begin{aligned} &\sum_{\mathbf{m} \in \mathbb{Z}^d} a_{\mathbf{m}} \mathrm{e}^{2\pi \mathrm{i} \mathbf{m} \bullet t} \\ &\qquad = \sum_{\mathbf{m} \in \mathbb{Z}^d} \mathrm{e}^{2\pi \mathrm{i} m_1 \alpha} a_{\mathbf{m}} \mathrm{e}^{2\pi \mathrm{i}((m_1+m_2)t_1 + (m_2+m_3)t_2 + \cdots + (m_{d-1}+m_d)t_{d-1} + m_d t_d)}. \end{aligned}$$

Since the coefficients in these series are unique, this requires that

$$\begin{aligned} &a_{(m_1+m_2, \ldots, m_{d-1}+m_d, m_d)} = \mathrm{e}^{2\pi \mathrm{i} m_1 \alpha} a_{\mathbf{m}} \\ &\qquad\qquad \Longrightarrow \quad |a_{(m_1+m_2, \ldots, m_{d-1}+m_d, m_d)}| = |a_{\mathbf{m}}| \end{aligned}$$

for all $\mathbf{m} \in \mathbb{Z}^d$.

However, because f is square-integrable, by Parseval's Theorem its Fourier coefficients are square-summable: $\sum_{\mathbf{m} \in \mathbb{Z}^d} |a_{\mathbf{m}}|^2 < \infty$. In light of this, consider the orbits of the linear transformation

$$(m_1, m_2, \ldots, m_d) \mapsto (m_1 + m_2, m_2 + m_3, \ldots, m_d).$$

The penultimate terms of the resulting orbit of vectors progress as m_{d-1}, $m_{d-1} + m_d$, $m_{d-1} + 2m_d$, $\ldots$, and so are all distinct unless $m_d = 0$. It therefore follows that $a_{\mathbf{m}} = 0$ whenever $m_d \neq 0$, for otherwise the associated orbit would give an infinitude of summands within $\sum_{\mathbf{m}\in\mathbb{Z}^d} |a_{\mathbf{m}}|^2$ all of the same nonzero magnitude.

Hence,

$$f(t) = \sum_{(m_1,m_2,\ldots,m_{d-1})\in\mathbb{Z}^{d-1}} a_{(m_1,m_2,\ldots,m_{d-1},0)} \mathrm{e}^{2\pi\mathrm{i}(m_1t_1+\cdots+m_{d-1}t_{d-1})},$$

so f a.s. does not depend on the coordinate t_d. This means it is lifted from a function on the system $(\mathbb{T}^{d-1}, \mathcal{B}(\mathbb{T}^{d-1}), m_{d-1}, S)$. Since that system is ergodic by the inductive hypothesis, the invariance of f implies that it must be constant m-a.s., as required. □

Proof of Theorem 11. Let our polynomial be

$$p(X) = a_dX^d + a_{d-1}X^{d-1} + \ldots + a_0.$$

Assume that a_d is irrational, and hence so is $a_d/d!$. Now apply Corollary 15 to the transformation

$$T(t_1, t_2, \ldots, t_d) := (t_1 + a_d/d!, t_2 + t_1, \ldots, t_d + t_{d-1}).$$

This T is uniquely ergodic by that theorem, and hence by Theorem 10 every point of $\mathbb{T}^d$ is generic for T.

Now, a simple calculation shows that if we define

$$\begin{aligned}
p_d(X) &:= p(X),\\
p_{d-1}(X) &:= p_d(X+1) - p_d(X),\\
p_{d-2}(X) &:= p_{d-1}(X+1) - p_{d-1}(X),\\
&\ldots,\\
p_1(X) &:= p_2(X+1) - p_2(X),
\end{aligned}$$

then each p_i is a polynomial of degree i, $p_1(X) = (a_d/d!)X$, and we have

$$T^n(p_1(0), p_2(0), \ldots, p_d(0)) = (p_1(n), p_2(n), \ldots, p_d(n)) \mod 1.$$

Since the point $(p_1(0), \ldots, p_d(0))$ must be generic for Haar measure, it follows that for any continuous function of only the last coordinate, say $f(t_d)$, we have

$$\frac{1}{N}\sum_{n=1}^{N} f(p(n) \bmod 1) = \frac{1}{N}\sum_{n=1}^{N} f(T^n(p_1(0), \ldots, p_d(0)) \bmod 1) \longrightarrow \int_{\mathbb{T}} f(s)\,\mathrm{d}s.$$

This is the required equidistribution, except that f is a continuous function, whereas Weyl's result asks us to work instead with the indicator function $1_{[a,b)}$. To repair this we simply observe that for any $\varepsilon > 0$ there are $[0,1]$-valued continuous functions f_1, f_2 with $f_1 \leqslant 1_{[a,b)} \leqslant f_2$ and $\int |f_2 - f_1|\,\mathrm{d}m_{\mathbb{T}} < \varepsilon$, and so since we know equidistribution for these functions we obtain

$$\int_{\mathbb{T}} f_1(s)\,\mathrm{d}s \leqslant \liminf_{N\longrightarrow\infty} \frac{|\{1 \leqslant n \leqslant N:\ a \leqslant \{p(n)\} < b\}|}{N}$$
$$\leqslant \limsup_{N\longrightarrow\infty} \frac{|\{1 \leqslant n \leqslant N:\ a \leqslant \{p(n)\} < b\}|}{N} \leqslant \int_{\mathbb{T}} f_2(s)\,\mathrm{d}s,$$

where the left- and right-hand expressions here straddle the value $\int_{\mathbb{T}} 1_{[a,b)} = b - a$ and differ by at most ε. Since ε was arbitrary, the liminf and limsup must in fact both equal $b - a$, as required. □

For the general case of Theorem 11, if instead a_i is irrational but a_d, $a_{d-1}, \dots,$ a_{i+1} are not, then one may reduce to the case of irrational leading coefficient as follows. Let $r \geqslant 1$ be some integer such that each ra_j is also an integer for all $j = i+1, \dots, d$. Then the equidistribution of $\{p(n)\}$ clearly follows if we know it for each of the individual sequences $\{p(rn + b)\}$, $b = 0, 1, \dots, r-1$, and for each such b one may re-arrange this latter problem so that it does not involve the coefficients in degrees $i+1, \dots, d$. This last step is left as an exercise.

References

[Fur81] Hillel Furstenberg. *Recurrence in Ergodic Theory and Combinatorial Number Theory*. Princeton University Press, Princeton, 1981.

[KB37] Nicolas Kryloff and Nicolas Bogoliouboff. La théorie générale de la mesure dans son application à l'étude des systèmes dynamiques de la mécanique non linéaire. *Ann. of Math. (2)*, 38(1):65–113, 1937.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: `tim@cims.nyu.edu`
URL: `cims.nyu.edu/˜tim`

Ergodic Theory

Notes 4: The 'arithmetic' of systems, and notions of mixing

1 The 'arithmetic' of systems

1.1 Factors

It is high time we met the principal notion of a 'morphism' between p.-p.s.s. This can be explained in two different ways.

Definition 1 (Factors)**.** *If (X, Σ, μ, T) is a p.-p. Γ-system, then a **factor** of it is a σ-subaglebra $\Xi \leqslant \Sigma$ which is globally T-invariant, meaning that, for each $\gamma \in \Gamma$,*

$$B \in \Xi \quad \Longleftrightarrow \quad T^\gamma(B) \in \Xi.$$

Definition 2 (Factor maps)**.** *If (X, Σ, μ, T) and (Y, Φ, ν, S) are p.-p. Γ-systems, then a **factor map** from the first to the second is a measurable map $\pi : X \longrightarrow Y$ such that*

- *(measures agree) $\pi_* \mu = \nu$;*
- *(intertwining property) $\pi \circ T^\gamma(x) = S^\gamma \circ \pi(x)$ for μ-a.e. $x \in X$ for every $\gamma \in \Gamma$.*

If π is a factor map as above, then the pre-image σ-algebra

$$\pi^{-1}(\Phi) := \{\pi^{-1}(A) : A \in \Phi\} \leqslant \Sigma$$

is a factor of (X, Σ, μ, T). A factor map $\pi : (X, \Sigma, \mu, T) \longrightarrow (Y, \Phi, \nu, S)$ may be thought of as 'simulating' its target system using a 'part' of the domain system[1].

Moreover, it turns out that *all* factors actually come from factor maps modulo negligible sets, so in this sense the above notions are equivalent. We will now

[1] although beware that this may not correspond to the target being 'smaller' than the domain system in any rigorous sense, just as a proper surjection between infinite sets does not imply that they have different cardinality

use some new notation. If (X, Σ, μ) is a probability space and $\Xi, \Phi \leqslant \Sigma$ are σ-subalgebras, then Ξ and Φ **agree up to negligible sets**, written $\Xi = \Phi \mod \mu$, if for every $A \in \Xi$ there is some $B \in \Phi$ such that $\mu(A \triangle B) = 0$, and vice-versa. For instance:

- our definition of (X, Σ, μ) being **countably generated up to negligible sets** asserted that $\Sigma = \Sigma_0 \mod \mu$ for some countably generated $\Sigma_0 \leqslant \Sigma$;
- on the unit interval $[0, 1]$, the σ-algebra of all Lebesgue-measurable sets agrees with the Borel σ-algebra up to Lebesgue-negligible sets.

Proposition 3 (Existence of spatial models)**.** *If (X, Σ, μ, T) is a p.-p. Γ-system and $\Phi \leqslant \Sigma$ is a factor, then there are*

- *a topological Γ-system (Y, S),*
- *an S-invariant Borel probability measure ν on Y,*
- *and a factor map of p.-p. systems*

$$\pi : (X, \Sigma, \mu, T) \longrightarrow (Y, \mathcal{B}(Y), \nu, S)$$

such that $\Phi = \pi^{-1}(\mathcal{B}(Y)) \mod \mu$.

Proof. Step 1. First we will show that for any factor $\Phi \leqslant \Sigma$ is agrees with a countably generated factor up to negligible sets. Let $A_1, A_2, \ldots \in \Sigma$ be a sequence which is in dense the measure algebra of (X, Σ, μ), and for each i let $f_i \in L^1(\mu|_\Phi)$ be some choice of function representing the conditional expectation $\mathsf{E}(1_{A_i} \mid \Phi)$ (recall that conditional expectations are unique only up to agreement almost everywhere). Now let Φ_0 be the σ-algebra generated by all of the rational level sets of the functions $f_i \circ T^\gamma$ for $i \geqslant 1$ and $\gamma \in \Gamma$:

$$\Phi_0 := \sigma\text{-alg}\big(\{\{x : \ f_i(T^\gamma x) \leqslant q\} : \ i \geqslant 1, \ \gamma \in \Gamma, \ q \in \mathbb{Q}\}\big).$$

This Φ_0 is clearly globally T-invariant, since it is generated by a globally T-invariant collection of sets, and it is contained in Φ. On the other hand, for any $A \in \Phi$ and $\varepsilon > 0$ there is some $A_i \in \Sigma$ such that $\mu(A \triangle A_i) < \varepsilon$, and since $A \in \Phi$ this implies that $\|1_A - f_i\|_1 < \varepsilon$. Hence A must be within ε of the level set $\{f_i \leqslant 1/2\}$, and so any $A \in \Phi$ may be approximated in measure by elements of Φ_0. Letting B_i be a sequence of such approximants in Φ_0 with $\mu(A \triangle B_i) < 2^{-i}$, we see that $B := \bigcap_{j \geqslant 1} \bigcup_{i \geqslant j} B_i$ still lies in Φ_0 and has $\mu(A \triangle B) = 0$, as required.

Step 2. Suppose now that $A_1, A_2, \ldots$ is a dense sequence in (the measure algebra of) Φ.

Let $Y := \{0,1\}^{\Gamma\times\mathbb{N}}$ with its product topology and σ-algebra (bearing in mind that Γ is countable, so this product topology is metrizable). Define $\phi : X \longrightarrow Y$ by

$$\phi(x) := \big(1_{A_i}(T^\gamma(x))\big)_{(\gamma,i)\in\Gamma\times\mathbb{N}}.$$

We also equip Y with the right-coordinate-shift action of Γ:

$$S^\lambda\big((x_{\gamma,i})_{\gamma,i}\big) := (x_{\gamma\lambda,i})_{\gamma,i},$$

and simply define $\nu := \phi_*\mu$. It is clear that $\phi^{-1}(\mathcal{B}(Y))$ is the σ-algebra generated by the sets A_i and their shifts, and hence is exactly Φ. Therefore, to prove that ϕ is the desired factor map, it remains to show that it intertwines the actions. This follows by definition because

$$\phi(T^\lambda(x)) := \big(1_{A_i}(T^\gamma(T^\lambda(x)))\big)_{\gamma,i} = \big(1_{A_i}(T^{\gamma\lambda}(x))\big)_{\gamma,i} =: S^\lambda\Big(\big(1_{A_i}(T^\gamma(x))\big)_{\gamma,i}\Big),$$

as required. □

Definition 4 (Extensions)**.** *If $\pi : (X,\Sigma,\mu,T) \longrightarrow (Y,\Phi,\nu,S)$ is a factor map, then we also write that (X,Σ,μ,T) is an* ***extension of*** *(Y,Φ,ν,S)* ***through*** *π.*

In the special case $\Xi = \Sigma$, Proposition 3 gives something called a **compact model**:

Corollary 5. *For any p.-p. Γ-system (X,Σ,μ,T) there are a topological Γ-system (Y,S), a measure $\nu \in \mathrm{Pr}^S(Y)$ and a measurable map $\pi : (X,\Sigma) \longrightarrow (Y,\mathcal{B}(Y))$ which has $\pi_*\mu = \nu$, intertwines the actions T and S and is such that any $A \in \Sigma$ differs by a negligible set from some $\pi^{-1}(B)$.* □

The importance of this corollary is that if we have some question about p.-p.s. which involves only the properties of sets up to negligible sets, or of functions up to a.e. equality, then without loss of generality we can always assume that our system is defined by homeomorphisms on some compact metric space. We will occasionally make important use of this assumption in the sequel.

Corollary 6 seems close to asserting that π defines an isomorphism of p.-p. Γ-systems $(X,\Sigma,\mu) \longrightarrow (Y,\mathcal{B}(Y),\nu)$, in the sense explained in Lecture 1. The difference is that a true isomorphism should have a measurable inverse a.e., whereas Corollary 6 gives only the apparently-weaker conclusion that all measurable subsets of X are pull-backs of subsets of Y, up to negligible sets.

For a general countably generated measurable space (X,Σ), the weaker conclusion need not imply the stronger. However, for standard Borel spaces, they are equivalent. (This measure-theoretic fact will not be proved in the present notes.)

Corollary 6. *If (X, Σ, μ, T) is a p.-p. Γ-system defined on a standard Borel space, then it is isomorphic as a p.-p. system to to a topological Γ-system with an invariant measure.* $\square$

A simple use of Corollary 6 also gives a model for factor maps:

Proposition 7. *If $\phi : (X_1, \Sigma_1, \mu_1, T_1) \longrightarrow (X_0, \Sigma_0, \mu_0, T_0)$ is a factor map, and $\pi_0 : (X_0, \Sigma_0, \mu_0, T_0) \longrightarrow (Y_0, \mathcal{B}(Y_0), \nu_0, S_0)$ is a compact model, then we can find a compact model $\pi_1 : (X_1, \Sigma_1, \mu_1, T_1) \longrightarrow (Y_1, \mathcal{B}(Y_1), \nu_1, S_1)$ and a* continuous *factor map $\psi : Y_1 \longrightarrow Y_0$ such that the following diagram commutes:*

$$\begin{array}{ccc} X_1 & \xrightarrow{\pi_1} & Y_1 \\ {\scriptstyle\phi}\downarrow & & \downarrow{\scriptstyle\psi} \\ X_0 & \xrightarrow[\pi_0]{} & Y_0. \end{array}$$

Proof. Corollary 6 gives us a provisional compact model $\pi' : (X_1, \Sigma_1, \mu_1, T_1) \longrightarrow (Y', \mathcal{B}(Y'), \nu', S')$. We now enhance this to a larger model by defining $Y_1 := Y_0 \times Y'$,

$$\pi_1 : X_1 \longrightarrow Y_0 \times Y' : x \mapsto (\pi_0(\phi(x)), \pi'(x)),$$

and $\nu_1 := \pi_{1*}\mu$. Finally, let $\psi : Y_1 \longrightarrow Y_0$ be the projection onto the first coordinate. The required properties follow at once. $\square$

Remark. One of the important general problems of ergodic theory is the 'realization problem'. The above results show how a compact metric model can be found for a general p.-p.s., and the realization problem asks what more we can demand from these compact models. For instance, can we express a general $\mathbb{Z}$-system as

1. a compact model system $(Y, \mathcal{B}(Y), \nu, S)$ with S being *uniquely ergodic*, or

2. a system comprising a diffeomorphism of a smooth compact manifold M and an invariant measure that is given by a smooth density on M?

It is a deep theorem of Jewett and Krieger that the answer to question 1 is Yes; we will discuss some related constructions towards the end of the course. At time of writing, question 2 remains one of the most famous open problems of ergodic theory. $\lhd$

Let us also quickly mention a very important kind of factor which will reappear later. Suppose that $T : \Gamma \curvearrowright (X, \Sigma, \mu)$ and that $\Lambda \trianglelefteq \Gamma$ is a normal subgroup. If $A \in \Sigma$ is such that $T^\lambda(A) = A$ for all $\lambda \in \Lambda$, then for any $\gamma \in \Gamma$ we have

$$T^\lambda(T^\gamma(A)) = T^\gamma(T^{\gamma^{-1}\lambda\gamma}(A)) = A \quad \forall \lambda \in \Lambda,$$

since by assumption $\gamma^{-1}\Lambda\gamma = \Lambda$. Hence the σ-subalgebra

$$\Sigma^{T\upharpoonright\Lambda} := \{A \in \Sigma : \; T^{\lambda}(A) = A \;\forall \lambda \in \Lambda\}$$

is globally T-invariant.

Definition 8. *The factor $\Sigma^{T\upharpoonright\Lambda}$ is called the **Λ-partially invariant factor** of (X, Σ, μ, T).*

1.2 Joinings and disjointness

Factors give us a notion of 'parts' of a given p.-p.s.. It can also be important to study ways in which new, 'larger' systems can be built up from given ingredients.

Definition 9. *Given two p.-p. Γ-systems (X, Σ, μ, T), (Y, Φ, ν, S), a **joining** of them is a third system (Z, Ξ, θ, R) together with factor maps*

$$\begin{array}{ccc} & (Z, \Xi, \theta, R) & \\ \swarrow^{\pi} & & {}^{\psi}\searrow \\ (X, \Sigma, \mu, T) & & (Y, \Phi, \nu, S), \end{array}$$

such that the new σ-algebra Ξ is generated up to negligible sets by $\pi^{-1}(\Sigma)$ and $\psi^{-1}(\Phi)$ (so we cannot find a proper factor of (Z, Ξ, θ, R) with respect to which both factor maps π and ψ are still measurable).

Joinings of larger collections of systems are defined similarly.

Joinings were introduced into ergodic theory in Furstenberg's classic paper [Fur67], which still makes delightful reading. It can often be helpful to have the following more concrete picture of a joining in mind, whose proof requires only routine verifications.

Proposition 10. *Given a joining of two p.-p.s.s*

$$\begin{array}{ccc} & (Z, \Xi, \theta, R) & \\ \swarrow^{\pi} & & {}^{\psi}\searrow \\ (X, \Sigma, \mu, T) & & (Y, \Phi, \nu, S), \end{array}$$

the map $\phi : Z \longrightarrow X \times Y : z \mapsto (\pi(z), \psi(z))$ is measurable for the product σ-algebra $\Sigma \otimes \Phi$, it intertwines the actions R and $T \times S$, and the resulting measure $\lambda := \phi_*\theta$ is $(T \times S)$-invariant and has first and second marginals equal to μ and ν:

$$\lambda(A \times X) = \mu(A), \quad \lambda(X \times B) = \nu(B).$$

□

Thus, given the systems (X, Σ, μ, T) and (Y, Φ, ν, S), there is a canonical way to model a joining of them on the space $X \times Y$. Often a measure with marginals μ and ν and invariant under $T \times S$ is itself referred to as a joining of these two systems. This also makes contact with the notion of 'coupling' in probability theory. In ergodic-theoretic uses of joinings, the emphasis is crucially on the $(T \times S)$-invariance of λ.

This picture also makes it clear that any two systems always have at least one joining: the simple Cartesian product:

$$(X, \Sigma, \mu, T) \times (Y, \Phi, \nu, S) := \big(X \times Y, \Sigma \otimes \Phi, \mu \otimes \nu, (T^\gamma \times S^\gamma)_{\gamma \in \Gamma}\big).$$

In fact, with the use of factor maps we can also introduce a considerable generalization of this construction, calling on the following fact from measure theory:

Proposition 11 (Measure disintegration)**.** *Suppose that $\phi : X \longrightarrow Y$ is a Borel map between compact metric spaces, that μ is a Borel probability measure on X and that $\nu = \phi_*\mu$ is the pushforward measure of μ on Y. Then there is a function*

$$Y \ni y \mapsto \mu_y \in \Pr(X)$$

into the space of Borel probability measures on X with the property that for any $A \in \Pr(X)$ the function

$$Y \ni y \mapsto \mu_y(A) \in \mathbb{R}$$

is Borel measurable, we have

$$\mu(A) = \int_Y \mu_y(A)\, \phi_*\nu(\mathrm{d}y),$$

and μ_y is supported on $\phi^{-1}\{y\}$ for almost every $y \in Y$. The above integral expression for μ is referred to as its ***disintegration*** *over ϕ, and the individual fibre measures μ_y as its* ***disintegrands****. The disintegration is essentially unique, in that if $y \mapsto \mu'_y$ is another disintegration then $\mu_y = \mu'_y$ for ν-a.e. y.* □

Now, given a continuous factor map of compact models of systems

$$\pi : (X, \Sigma, \mu, T) \longrightarrow (Y, \Phi, \nu, S),$$

let $\mu = \int_Y \mu_y\, \nu(\mathrm{d}y)$ be the disintegration of μ over π. Owing to the essential uniqueness of the disintegration, we must have $T^\gamma_*\mu_y = \mu_{S^\gamma y}$ for ν-a.e. y. Now we form the new fibred product set $\tilde{X} := \{(x, x') \in X^2 : \pi(x) = \pi(x')\} \subseteq X^2$, which is easily seen to be closed, and is $(T \times T)$-invariant because π is a factor map.

We give $\tilde{X}$ its Borel σ-algebra, and on it we put the **relative product measure** defined by

$$\mu \otimes_\pi \mu(A) := \int_Y (\mu_y \otimes \mu_y)(A)\, \nu(\mathrm{d}y) \quad \text{for } A \in \mathcal{B}(\tilde{X}).$$

This is easily seen to be a Borel probability measure on $\tilde{X}$, and the property that $T^\gamma_* \mu_y = \mu_{S^\gamma y}$ implies that $\mu \otimes_\pi \mu$ is invariant under $T \times T$. Finally the two coordinate projections $\pi_i : \tilde{X} \longrightarrow X$, $i = 1, 2$, both push $\mu \otimes_\pi \mu$ back onto μ, so considered as factor maps these witness $(\tilde{X}, \mathcal{B}(\tilde{X}), \mu \otimes_\pi \mu, T \times T)$ as a joining of two copies of (X, Σ, μ, T), whose construction gives the important additional property that $\pi \circ \pi_1 = \pi \circ \pi_2$, and hence the commutative diagram

$$\begin{array}{ccccc} & & \tilde{X} & & \\ & \overset{\pi_1}{\swarrow} & & \overset{\pi_2}{\searrow} & \\ X & & & & X \\ & \underset{\pi}{\searrow} & & \underset{\pi}{\swarrow} & \\ & & Y. & & \end{array}$$

An equivalent description of the relative product measure $\mu \otimes_\pi \mu$ is that for any bounded measurable functions $f, g : X \longrightarrow \mathbb{R}$, one has

$$\int_{\tilde{X}} f(x)g(x')\, (\mu \otimes_\pi \mu)(\mathrm{d}(x, x')) = \int_X \mathsf{E}(f \mid \pi^{-1}(\Phi))\mathsf{E}(g \mid \pi^{-1}(\Phi))\, \mu(\mathrm{d}x).$$

1.3 Inverse limits

The final construction we record here is that of inverse limits. Suppose that we have a whole sequence of Γ-systems $(X_n, \Sigma_n, \mu_n, T_n)$ for $n \geqslant 0$, all of them compact models, together with **connecting maps**: continuous factor maps $\pi_n : (X_{n+1}, \Sigma_{n+1}, \mu_{n+1}, T_{n+1}) \longrightarrow (X_n, \Sigma_n, \mu_n, T_n)$, which we may visualize as a tower

$$\cdots \longrightarrow X_3 \xrightarrow{\pi_2} X_2 \xrightarrow{\pi_1} X_1 \xrightarrow{\pi_0} X_0.$$

(Note that given a tower of abstract systems and factor maps, an inductive application of Proposition 7 will give us compact models and continuous maps as here.)

Proposition 12. *In this situation, there are a compact model* Γ*-system* (X, Σ, μ, T) *together with a sequence of continuous factor maps* $\phi_n : X \longrightarrow X_n$ *such that*

- *(the factors are consistent)* $\pi_n \circ \phi_{n+1} = \phi_n$ *for all* $n \geqslant 0$*, and*

- *(the factors generate everything) Σ is generated up to negligible sets by its factors $\phi_n^{-1}(\Sigma_n)$.*

Moreover, this new system is essentially unique, in that given any other system (Y, Φ, ν, S) and factor maps $\psi_n : Y \longrightarrow X_n$ such that $\pi_n \circ \psi_{n+1} = \psi_n$ for all n, there is a factor map $\alpha : Y \longrightarrow X$ with $\psi_n = \phi_n \circ \alpha$ for all n. In particular, the σ-algebra Σ is generated by the sequence of σ-subalgebras $\phi_n^{-1}(\Sigma_n)$.

Proof. Let
$$X := \{(x_n)_{n \geqslant 1} : \ \pi_n(x_{n+1}) = x_n \ \forall n \geqslant 0\}.$$
Since we have chosen to work with compact models, this space is a decreasing intersection of compact subsets of $X_1 \times X_2 \times \cdots$, and so is itself a nonempty compact subset of this product.

For $m > n \geqslant 0$, let $\pi_{m,n} := \pi_n \circ \pi_{n+1} \circ \cdots \circ \pi_{m-1} : X_m \longrightarrow X_n$, so $\pi_n = \pi_{n+1,n}$ for each n. These are all still factor maps. Now, for each n, define $\nu_n \in \Pr(X_0 \times X_1 \times \cdots \times X_n)$ to be the pushforward of μ_n under the map
$$x \mapsto \big(\pi_{n,0}(x), \ \ldots, \ \pi_{n,n-2}(x), \ \pi_{n,n-1}(x), \ x\big),$$
so that this is a measure on the projection of X onto its first n coordinate factors. By taking a weak* limit we can extract from these a measure on X: for example, by picking a reference point $(y_0, y_1, \ldots) \in X_1 \times X_2 \times \cdots$, we may form the products $\nu_n \otimes \delta_{(y_{n+1}, y_{n+2}, \ldots)} \in \Pr(X_1 \times X_2 \times \cdots)$, and now these converge to a weak* limit $\mu \in \Pr(X)$ which is concentrated on X, because by the consistency among the μ_n these product measures eventually agree on any finite collection of coordinates in $X_1 \times X_2 \times \cdots$. Now if we define
$$\phi_n : X \longrightarrow X_n : (x_0, x_2, \ldots) \mapsto x_n$$
then $\phi_{n*}\mu = \mu_n$, and because of the definition of X we also have $\pi_n \circ \phi_{n+1} = \phi_n$. Also, if we define $T : \Gamma \curvearrowright (X, \mathcal{B}(X), \mu)$ by $T^\gamma = T_1^\gamma \times T_2^\gamma \times \cdots$, then each ϕ_n intertwines T with T_n.

Finally, suppose that (Y, Φ, ν, S) and $\psi_n : Y \longrightarrow X_n$ are as posited, and define $\alpha : Y \longrightarrow X$ by
$$\alpha(y) := (\psi_0(y), \psi_1(y), \ldots).$$
The assumption that $\pi_n \circ \psi_{n+1} = \psi_n$ implies that $\alpha(y) \in X$ for all y, and moreover it is clear by definition that α intertwines S with T (since this holds coordinate-wise in X) and that $\alpha_*\nu$ agrees with μ on any finite-dimensional projection of X, so that in fact $\alpha_*\nu = \mu$. □

Definition 13 (Inverse limit)**.** *The system (X, Σ, μ, T) together with the factor maps $(\phi_n)_{n \geqslant 0}$ is the **inverse limit** of the tower of systems $(X_n, \Sigma_n, \mu_n, T_n)$, $(\pi_n)_{n \geqslant 0}$.*

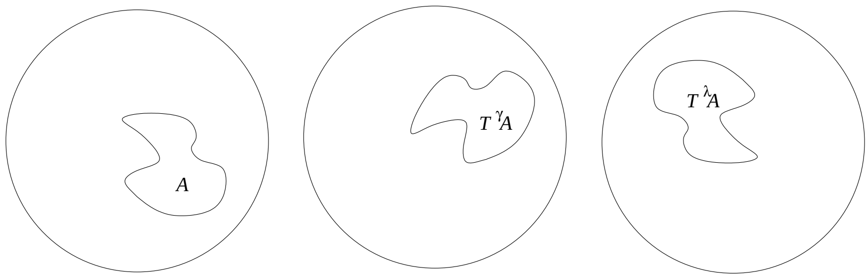

Figure 1: Sometimes, shape may be retained as well as size

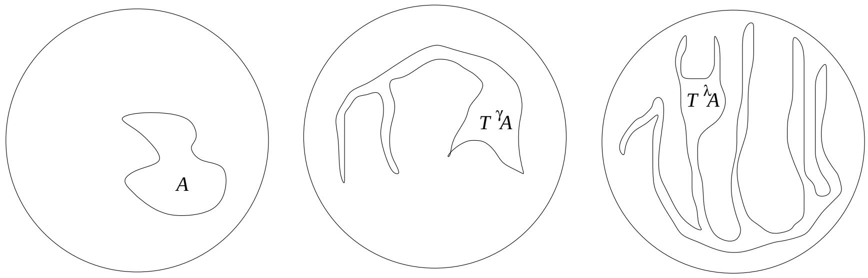

Figure 2: In other cases, shapes can become increasingly distorted

2 Weak and strong mixing

Two strengthenings of ergodicity, 'weak' and 'strong mixing', have emerged over time as important determiners of many different features of how a p.-p.s. behaves. We will now introduce these and begin to explore their basic properties.

In order to introduce the kinds of phenomenon that interest us, imagine a p.-p.s. (or single transformation, if you like) $T : \Gamma \curvearrowright (X, \Sigma, \mu)$, together with some subset $A \in \Sigma$ with $0 < \mu(A) < 1$, and suppose further that we have some natural way of visualizing the space X, say as a blob with A inside it as a smaller blob.

Now run the dynamics, and consider the images $T^{\gamma_i}(A)$ for some sequence $(\gamma_i)_i$ in Γ which tends to ∞: that is, which eventually leaves every finite subset of Γ. This is a sequence of sets of measure $\mu(A)$, because T preserves μ, but apart from that we know very little.

It could be that in some sense these images of A retain some features of their shape, as well as their size, as in Fig 2. On the other hand, it could be that, as $\gamma_i \longrightarrow \infty$, the 'shapes' of our image sets become more and more distorted, so that although their measures all equal $\mu(A)$, they spread that measure around the space X in an increasingly 'complicated' way, as in Fig 3.

Of course, we have not given rigorous meaning to the words 'shape' and 'complicated' here. Nevertheless, it turns out that one can give a precise form to these two intuitively-opposed possibilities. Perhaps the simplest way to do this for a given ergodic system $T : \Gamma \curvearrowright (X, \Sigma, \mu)$ is to ask about the behaviour of *pairs* of points $(x, y) \in X^2$. The idea is that in a system of the first kind above, where some sense of 'shape' is preserved, knowing the position of the T^γ-image of one point places some constraints on the T^γ-images of other points, whereas in a system of the second kind most pairs of distinct points will ultimately appear to be moving independently. The following definition gives a precise way to capture this intuition.

Definition 14 (Weak mixing)**.** *A p.-p.s.* $T : \Gamma \curvearrowright (X, \Sigma, \mu)$ *is* ***weakly mixing*** *if the product system* $T \times T : \Gamma \curvearrowright (X \times X, \Sigma \otimes \Sigma, \mu \otimes \mu)$ *is ergodic.*

The following is worth noting at once:

Lemma 15. *Weak mixing implies ergodicity.*

Proof. If $A \in \Sigma$ is T-invariant, then $A \times X \in \Sigma \otimes \Sigma$ is $(T \times T)$-invariant and has $\mu^{\otimes 2}(A \times X) = \mu(A)$. □

As our first intuitive situation above suggests, however, the reverse of this implication is false, as is witnessed by compact group rotations.

Example. Let $\alpha \in \mathbb{T}$ be irrational (that is, the image in $\mathbb{R}/\mathbb{Z}$ of an element of $\mathbb{R}\backslash\mathbb{Q}$), and let $R_\alpha : \mathbb{T} \longrightarrow \mathbb{T}$ be the resulting circle rotation. We have seen in Lecture 3 that R_α is uniquely ergodic, with Lebesgue measure m being the only invariant measure.

However, $(\mathbb{T}, \mathcal{B}(\mathbb{T}), m, R_\alpha)$ is not weakly mixing. Indeed, for any $A \in \mathcal{B}(\mathbb{T})$, the diagonal set

$$A' := \{(s, t) \in \mathbb{T}^2 : \ t - s \in A\}$$

is $(R_\alpha \times R_\alpha)$-invariant and has measure $m(A)$, which could be any value in $[0, 1]$.

A slight generalization of this shows that for any acting group Γ, nontrivial compact group rotations are never weakly mixing. More generally still, if G is a compact metric group, $\phi : \Gamma \longrightarrow G$ is a homomorphism and $H \leqslant G$ is a proper closed subgroup, then we can also define the rotation action R_ϕ on the compact homogeneous space G/H. In this case, since H is proper, we may pick two distinct left cosets $g_1 H \neq g_2 H$, and now since these are disjoint compact subsets of G there are open neighbourhoods U, V of the identity in G such that $V g_1 H U \cap V g_2 H U = \emptyset$. Hence the subset

$$\{(gH, g'H) \in (G/H)^2 : \ gHU \cap g'HU \neq \emptyset\}$$

is manifestly $(R_\phi \times R_\phi)$-invariant, is open and contains the diagonal $\{(gH, gH)\}$ and so has positive measure for $m_{G/H}^{\otimes 2}$, but also omits the open set $Vg_1 \times Vg_2$ and so has complement with positive measure. Therefore R_ϕ is not weakly mixing. $\triangleleft$

We will see an example of a weakly mixing transformation once we have a little more machinery. One important indication of the merit in Definition 14 is that it has many equivalent formulations.

Proposition 16. *The following are equivalent:*

1. (X, Σ, μ, T) *is weakly mixing;*
2. *for any other ergodic system* $S : \Gamma \curvearrowright (Y, \Phi, \nu)$ *the product* $T \times S : \Gamma \curvearrowright (X \times Y, \Sigma \otimes \Phi, \mu \otimes \nu)$ *is ergodic;*
3. *for any finite collection* $\{A_1, A_2, \ldots, A_k\} \subseteq \Sigma$ *and* $\varepsilon > 0$ *there is some* $\gamma \in \Gamma$ *such that* $|\mu(A_i \cap T^\gamma A_j) - \mu(A_i)\mu(A_j)| < \varepsilon$ *for all* $i, j \leqslant k$*;*
4. *for any finite collection*
$$\{f_1, f_2, \ldots, f_k\} \subseteq L_0^2(\mu) := \left\{f \in L^2_{\mathbb{C}}(\mu) : \textstyle\int_X f \, \mathrm{d}\mu = 0\right\}$$
and $\varepsilon > 0$ *there is some* $\gamma \in \Gamma$ *such that*
$$|\langle f_i, f_j \circ T^\gamma \rangle| = \left| \int_X \overline{f_i} \cdot (f_j \circ T^\gamma) \, \mathrm{d}\mu \right| < \varepsilon \quad \textit{for all } i, j \leqslant k;$$

and, in case $\Gamma = \mathbb{Z}^d$*,*

5. *for any* $f, g \in L_0^2(\mu)$ *there is a subset* $M \subset \mathbb{N}^d$ *of zero density, meaning that*
$$\frac{|M \cap \{1, 2, \ldots, N\}^d|}{N^d} \longrightarrow 0 \quad \textit{as } N \longrightarrow \infty,$$
such that $\langle f, g \circ T^{\mathbf{n}_i} \rangle \longrightarrow 0$ *for any sequence* $(\mathbf{n}_i)_{i \geqslant 1}$ *tending to* ∞ *in* $\mathbb{N}^d \backslash M$*;*
6. *for any* $A, B \in \Sigma$ *there is a subset* $M \subset \mathbb{N}^d$ *of zero density such that* $\mu(A \cap T^{\mathbf{n}_i} B) \longrightarrow \mu(A)\mu(B)$ *for any sequence* $(\mathbf{n}_i)_{i \geqslant 1}$ *tending to* ∞ *in* $\mathbb{N}^d \backslash M$*.*

Proof. (1. $\Longrightarrow$ 4.) Suppose that $\{f_1, f_2, \ldots, f_k\} \subseteq L_0^2(\mu)$ and $\varepsilon > 0$ are such that for every $\gamma \in \Gamma$ there are some $i, j \leqslant k$ for which $|\langle f_i, f_j \circ T^\gamma \rangle| \geqslant \varepsilon$. Consider the function
$$F(x, y) := \sum_{i=1}^k f_i(x) \overline{f_i(y)} \in L^2(\mu^{\otimes 2}).$$

Since $\int_X f_i \,\mathrm{d}\mu = 0$ for each i, it follows that

$$\int_{X^2} F \,\mathrm{d}\mu^{\otimes 2} = \sum_{i=1}^{k} \Big(\int_X f_i \,\mathrm{d}\mu \Big) \overline{\Big(\int_X f_i \,\mathrm{d}\mu \Big)} = 0, \quad \text{i.e. } F \in L^2_0(\mu^{\otimes 2}).$$

Let

$$K := \overline{\mathrm{conv}}\{F \circ (T^\gamma \times T^\gamma) : \gamma \in \Gamma\},$$

the closure of the set of convex combinations of the shifts $F \circ (T^\gamma \times T^\gamma)$ of the function F. Since $L^2_0(\mu^{\otimes 2}) \subseteq L^2(\mu^{\otimes 2})$ is $(T \times T)$-invariant, K lies inside this subspace.

We will next show that

$$\inf\{\|G\|_{L^2(\mu^{\otimes 2})} : G \in K\} \geqslant \varepsilon > 0.$$

Indeed, by continuity it suffices to check this when $G = \sum_{r=1}^s \alpha_r F \circ (T^{\gamma_r} \times T^{\gamma_r})$ is a finite convex combination, and in this case we can compute as follows:

$$\begin{aligned}
\|G\|_2^2 &= \sum_{r,r'=1}^{s} \alpha_r \alpha_{r'} \big\langle F \circ (T^{\gamma_r} \times T^{\gamma_r}), F \circ (T^{\gamma_{r'}} \times T^{\gamma_{r'}}) \big\rangle \\
&= \sum_{r,r'=1}^{s} \alpha_r \alpha_{r'} \sum_{i,i'=1}^{k} \big\langle (f_i \circ T^{\gamma_r}) \otimes \overline{(f_i \circ T^{\gamma_r})}, (f_{i'} \circ T^{\gamma_{r'}}) \otimes \overline{(f_{i'} \circ T^{\gamma_{r'}})} \big\rangle \\
&= \sum_{r,r'=1}^{s} \alpha_r \alpha_{r'} \sum_{i,i'=1}^{k} |\langle f_i \circ T^{\gamma_r}, f_{i'} \circ T^{\gamma_{r'}} \rangle|^2 \\
&= \sum_{r,r'=1}^{s} \alpha_r \alpha_{r'} \sum_{i,i'=1}^{k} |\langle f_i, f_{i'} \circ T^{\gamma_{r'}\gamma_r^{-1}} \rangle|^2 \\
&\geqslant \varepsilon^2 \sum_{r,r'=1}^{s} \alpha_r \alpha_{r'} = \varepsilon^2 \Big(\sum_{r=1}^{s} \alpha_r \Big)^2 = \varepsilon^2.
\end{aligned}$$

The point to this is that it gives us a new way (without averaging, since we do not have the Norm Ergodic Theorem for an arbitrary group Γ) of producing a non-trivial $(T \times T)$-invariant function. Indeed, since the Hilbert space norm $\|\cdot\|_{L^2_0(\mu^{\otimes 2})}$ is uniformly convex and has a positive lower bound on the closed convex set K, it attains a *unique minimum* on K (even though K is possibly non-compact). Suppose that minimum is attained at $G \in K$. By uniqueness, this G must be $(T \times T)$-invariant, so

$$\|G\|_2 \geqslant \varepsilon, \quad \int_{X \times X} G \,\mathrm{d}\mu^{\otimes 2} = 0 \quad \text{and} \quad G \circ (T^\gamma \times T^\gamma) = G \;\forall \gamma \in \Gamma.$$

Therefore G is a nonconstant invariant function in $L^2(\mu^{\otimes 2})$, and so $T \times T$ is not ergodic.

(4. $\Longrightarrow$ 3.) If $\{A_1, A_2, \ldots, A_k\} \subset \Sigma$ and $\varepsilon > 0$, then for each i we consider the **balanced function** $f_i := 1_{A_i} - \mu(A_i) \in L_0^2(\mu)$, and now by assumption 4. there is some $\gamma \in \Gamma$ such that for every $i, j \leqslant k$ we have

$$\begin{aligned}\varepsilon > |\langle f_i, f_j \circ T^\gamma \rangle| &= |\langle 1_{A_i} - \mu(A_i), 1_{T^{\gamma^{-1}}(A_j)} - \mu(A_j) \rangle| \\ &= |\mu(A_i \cap T^{\gamma^{-1}}(A_j)) - \mu(A_i)\mu(A_j)|.\end{aligned}$$

(3. $\Longrightarrow$ 2.) Suppose that property 3. holds, but that a Γ-system (Y, Φ, ν, S) has been found such that $T \times S$ is not ergodic. Let $E \in \Sigma \otimes \Phi$ be an invariant set with $\alpha := (\mu \otimes \nu)(E) \in (0, 1)$. We will prove that S cannot be ergodic.

First, using a standard property of the product measure, for any $\varepsilon > 0$ there is a finite collection of pairwise-disjoint rectangles $A_1 \times B_1, \ldots, A_k \times B_k$ with $A_i \in \Sigma$ and $B_i \in \Phi$ such that

$$\begin{aligned}&(\mu \otimes \nu)\Big(E \triangle \bigcup_{i \leqslant k} (A_i \times B_i)\Big) < \varepsilon \\ \Longrightarrow\ &(\mu \otimes \nu)\Big(\bigcup_{i \leqslant k} (A_i \times B_i)\Big) < \alpha + \varepsilon \text{ and } (\mu \otimes \nu)\Big(E \cap \bigcup_{i \leqslant k} (A_i \times B_i)\Big) > \alpha - \varepsilon.\end{aligned}$$

Since E and $(\mu \otimes \nu)$ are both $(T \times S)$-invariant, it follows that for any $\gamma \in \Gamma$ the above inqualities hold with each rectangle $A_i \times B_i$ moved to $T^\gamma A_i \times T^\gamma B_i$, and so combining these facts gives

$$\begin{aligned}(\mu \otimes \nu)\Big(\bigcup_{i \leqslant k} (A_i \times B_i)\Big) &= (\mu \otimes \nu)\Big(\bigcup_{i \leqslant k} (T^\gamma A_i \times S^\gamma B_i)\Big) < \alpha + \varepsilon \\ \text{and} \quad (\mu \otimes \nu)\Big(&\bigcup_{i \leqslant k} (A_i \times B_i) \cap \bigcup_{j \leqslant k} (T^\gamma A_j \times S^\gamma B_j)\Big) > \alpha - 2\varepsilon\end{aligned}$$

for all $\gamma \in \Gamma$.

Since the rectangles are pairwise disjoint, we have

$$\begin{aligned}(\mu \otimes \nu)\Big(\bigcup_{i \leqslant k} (A_i \times B_i) \cap \bigcup_{j \leqslant k} (T^\gamma A_j \times S^\gamma B_j)\Big) &= \sum_{i,j \leqslant k} (\mu \otimes \nu)((A_i \cap T^\gamma A_j) \times (B_i \cap S^\gamma B_j)) \\ &= \sum_{i,j \leqslant k} \mu(A_i \cap T^\gamma A_j) \cdot \nu(B_i \cap S^\gamma B_j)\end{aligned}$$

Now, using property 3, for any $\varepsilon > 0$ we can find $\gamma \in \Gamma$ such that $|\mu(A_i \cap T^\gamma A_j) - \mu(A_i)\mu(A_j)| < \varepsilon/k^2$ for all $i, j \leqslant k$. In this case the last expression above would lie within ε of

$$\sum_{i,j\leqslant k} \mu(A_i)\mu(A_j)\cdot\nu(B_i \cap S^\gamma B_j) = \Big\langle \sum_{i=1}^k \mu(A_i)1_{B_i}, \sum_{j=1}^k \mu(A_j)1_{B_j} \circ S^{\gamma^{-1}} \Big\rangle_{L^2(\nu)}.$$

Letting $g := \sum_{i=1}^k \mu(A_i)1_{B_i}$, we see that it is measurable, non-negative, and that

$$g(y) = \sum_{i=1}^k \mu(A_i)1_{B_i}(y) = \mu\Big\{x \in X : (x, y) \in \bigcup_{i=1}^k (A_i \times B_i)\Big\} \leqslant 1,$$

because the rectangles $A_i \times B_i$ are pairwise disjoint. So $0 \leqslant g \leqslant 1$ and, in addition, $\int_Y g\,\mathrm{d}\nu = \sum_{i=1}^k \mu(A_i)\nu(B_i) \leqslant \alpha + \varepsilon$.

However, the inequalities proved above give

$$\langle g, g \circ S^{\gamma^{-1}}\rangle \geqslant \alpha - 3\varepsilon \geqslant \int_Y g\,\mathrm{d}\nu - 4\varepsilon,$$

and the Cauchy-Schwartz inequality implies that

$$|\langle g, g \circ S^{\gamma^{-1}}\rangle| \leqslant \|g\|_2\|g \circ S^{\gamma^{-1}}\|_2 = \|g\|_2^2 = \int_Y g^2\,\mathrm{d}\nu,$$

so overall we have shown that $0 \leqslant g \leqslant 1$ and

$$\int g^2 \geqslant \int g - 4\varepsilon \quad \Longrightarrow \quad \int g(1-g) \leqslant 4\varepsilon.$$

By Chebychev's inequality this implies

$$\nu\{y \in Y : g(y)(1 - g(y)) \geqslant 4\sqrt{\varepsilon}\} < \sqrt{\varepsilon}.$$

Letting $C := \{y : g(y) \geqslant 1/2\}$, Fubini's Theorem now gives

$$\begin{aligned}(\mu \otimes \nu)\Big(\bigcup_{i=1}^k (A_i \times B_i) \triangle (X \times C)\Big) &= \int_{Y\setminus C} \mu\Big\{x : (x, y) \in \bigcup_{i=1}^k (A_i \times B_i)\Big\}\,\nu(\mathrm{d}y) \\ &\quad + \int_C \mu\Big\{x : (x, y) \notin \bigcup_{i=1}^k (A_i \times B_i)\Big\}\,\nu(\mathrm{d}y) \\ &= \int_{Y\setminus C} g + \int_C (1 - g) < 10\sqrt{\varepsilon}.\end{aligned}$$

Since our union of rectangles was an approximation to E, we obtain that $(\mu \otimes \nu)(E \triangle (X \times C)) < \varepsilon + 10\sqrt{\varepsilon}$. Since ε was arbitrary, it follows that E may be approximated arbitrarily well by sets depending only on the second coordinate in $X \times Y$, and so E must itself be equal to such a set $X \times D$ modulo a negligible set. This in turn requires that D be almost S-invariant of measure α, and hence that $\bigcap_{\gamma \in \Gamma} S^\gamma D$ has full measure in D and is strictly S-invariant, so S is not ergodic.

(2. $\Longrightarrow$ 1.) This takes two quick strides: first applying 2. with S the trivial action on a one-point space shows that T is ergodic, and now applying it again with $(Y, \Phi, \nu, S) := (X, \Sigma, \mu, T)$ proves property 1.

Now suppose that $\Gamma = \mathbb{Z}^d$.

(1. $\Longrightarrow$ 5.) This is similar to the implication (1. $\Longrightarrow$ 4.) but uses the Norm Ergodic Theorem. Suppose that $T \times T$ is ergodic and that $f, g \in L^2_0(\mu)$. The point is that the inner products $|\langle f, g \circ T^{\mathbf{n}} \rangle|$ are bounded by $\|f\|_2 \|g\|_2$ independently of $\mathbf{n}$, and therefore the assertion of convergence to zero away from a zero-density subset of $\mathbb{N}^d$ is equivalent to the assertion that

$$\frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} |\langle f, g \circ T^{\mathbf{n}} \rangle|^2 \longrightarrow 0 \quad \text{as } N \longrightarrow \infty.$$

However, the left-hand expression here is equal to

$$\frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} \big\langle f \otimes \overline{f}, (g \otimes \overline{g}) \circ (T \times T)^{\mathbf{n}} \big\rangle = \Big\langle f \otimes \overline{f}, \frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} (g \otimes \overline{g}) \circ (T \times T)^{\mathbf{n}} \Big\rangle.$$

Since $g \in L^2_0(\mu)$ it follows that $\int_{X \times X} g \otimes \overline{g} \, \mathrm{d}\mu^{\otimes 2} = 0$, and now since $T \times T$ is ergodic the Norm Ergodic Theorem gives that

$$\frac{1}{N^d} \sum_{\mathbf{n} \in [N]^d} (g \otimes \overline{g}) \circ (T \times T)^{\mathbf{n}} \longrightarrow 0 \quad \text{in } L^2(\mu^{\otimes 2}) \text{ as } N \longrightarrow \infty,$$

and hence the above inner products also tend to zero, as required.

(5. $\Longleftrightarrow$ 6.) Once again this follows in one direction by replacing each set A with its balanced function $f := 1_A - \mu(A)$, and in the other by approximating arbitrary functions by finite-valued simple functions and then considering their level sets individually.

(5. $\Longrightarrow$ 4.) If $\{f_1, f_2, \ldots, f_k\} \subset L^2_0(\mu)$ and $\varepsilon > 0$, then, by assumption, for each i, j there is some $M_{i,j} \subseteq \mathbb{N}^d$ having zero density such that $|\langle f_i, f_j \circ T^{\mathbf{n}} \rangle| < \varepsilon$ for all $\mathbf{n} \in \mathbb{N}^d \backslash M_{i,j}$. Since a finite union of zero-density sets still has zero density, we can find a point $\mathbf{n} \in \mathbb{N}^d \backslash \bigcup_{i,j \leqslant k} M_{i,j}$, and this now witnesses the property 4. $\square$

Remark. It is important to be aware that an action of a group Γ can be weakly mixing even though there is *some* sequence of group elements $\gamma_i \longrightarrow \infty$ such that the images $T^{\gamma_i}A$ do not become asymptotically independent, as in condition 3 above. Indeed, if any infinite subgroup $\Lambda \leqslant \Gamma$ is such that the subaction $(T^\lambda)_{\lambda\in\Lambda}$ is weakly mixing, then the whole action T is weakly mixing, even though some other subgroup $\Lambda' \leqslant \Gamma$ may have subaction which is not even ergodic. ◁

Now condition 2 above and a simple induction give the following.

Corollary 17. *If $T : \Gamma \curvearrowright (X, \Sigma, \mu)$ is weakly mixing then all of its Cartesian powers $T^{\times N} : \Gamma \curvearrowright (X^N, \Sigma^{\otimes N}, \mu^{\otimes N})$ are also weakly mixing.* □

In the next lecture we will explore a rather deeper necessary and sufficient condition for weak mixing, but that will have to wait until we have developed more tools.

The Bernoulli shifts offer examples of systems that are weakly mixing. A hint of this is offered by our intuitive discussion of the movement of pairs of points. Suppose that $\mu = \nu^{\otimes\Gamma}$ is a Bernoulli measure on the product space E^Γ and that $x, y \in E^\Gamma$ are chosen separately at random from μ. Then for any other $w, z \in E^\Gamma$ and any finite subset $F \subset \Gamma$, because F is finite we may choose a sequence of translates $F\gamma_1$, $F\gamma_2$, ...that are pairwise disjoint. For each of these translates there is some positive (although maybe very small) probability that both $x|_{F\gamma_i}$ is the γ_i-translate of $z|_F$ and $y|_{F\gamma_i}$ is the γ_i-translate of $w|_F$. Since these events are all independent when (x, y) are drawn independently from μ (because disjoint sets of coordinates are independent under μ), it follows that they must happen somewhere with probability 1, and hence that $(T^{\gamma_i}x, T^{\gamma_i}y) \approx (z, w)$ for some i. So it seems that almost every pair of points visits the vicinity of any other pair as we apply the diagonal action $T \times T$. However, to make rigorous contact with Definition 14 we need a little more care, and in particular we need the following further definition.

Definition 18 (Strong mixing)**.** *A p.-p.s. $T : \Gamma \curvearrowright (X, \Sigma, \mu)$ is* ***strongly mixing*** *if for any $A, B \in \Sigma$ we have*

$$\mu(A \cap T^\gamma B) \longrightarrow \mu(A)\mu(B) \quad \textit{as } \gamma \longrightarrow \infty \textit{ in } \Gamma.$$

Lemma 19. *Strong mixing implies weak mixing.*

Proof. This follows from condition 3 of Proposition 16: given $A_1, A_2, \dots, A_k \in \Sigma$ and ε, strong mixing implies that, for each pair $i, j \leqslant k$,

$$|\mu(A_i \cap T^\gamma A_j) - \mu(A_i)\mu(A_j)| < \varepsilon$$

for all γ outside some finite set $F_{i,j} \subset \Gamma$, so taking γ outside $\bigcup_{i,j} F_{i,j}$ witnesses that condition. □

Example. Let us now prove properly that a Bernoulli Γ-system is strongly mixing and hence weakly mixing. Recall first that this is the action

$$T^{\gamma} : (x_{\lambda})_{\lambda \in \Gamma} \mapsto (x_{\lambda\gamma})_{\lambda \in \Gamma}$$

defined on a product probability space $(E^{\Gamma}, \mathcal{B}(E^{\Gamma}), \nu^{\otimes \Gamma})$ for some auxiliary probability space (E, ν). For a finite subset $F \subset \Gamma$, let $\pi_F : E^{\Gamma} \longrightarrow E^F$ be the coordinate projection. If $A, B \subseteq E^{\Gamma}$ are measurable, then for any $\varepsilon > 0$ there are sets $A_0, B_0 \subseteq E^{\Gamma}$ which depend on only finitely many coordinates, say with $A_0 = \pi_F^{-1}(A_1)$ and $B_0 = \pi_H^{-1}(B_1)$, and such that $\mu(A \triangle A_0), \mu(B \triangle B_0) < \varepsilon/4$.

However, provided $\gamma \in \Gamma$ is outside the finite set $F \cdot H^{-1}$, it follows that F is disjoint from $H\gamma$, and hence that A_0 and $T^{\gamma^{-1}}(B_0)$ are strictly independent events under μ:

$$\mu(A_0 \cap T^{\gamma^{-1}}(B_0)) = \mu(A_0)\mu(B_0).$$

Combining this with the approximation of A, B by A_0, B_0 gives

$$|\mu(A \cap T^{\gamma^{-1}}(B)) - \mu(A)\mu(B)| < \varepsilon,$$

as required. $\lhd$

At this point, let us also note the obvious fact that weak and strong mixing are isomorphism invariants of systems, and so give us a way to distinguish systems up to isomorphism.

Corollary 20. *Compact homogeneous space rotations are never isomorphic to Bernoulli shifts, unless both are trivial.* $\square$

We have seen that

$$\text{Strong mixing} \implies \text{Weak mixing} \implies \text{Ergodicity},$$

where the second implication is strict (it cannot be reversed) for many groups (at least, any group that admits a nontrivial homomorphism to a compact group). For many groups, including every $\mathbb{Z}^d$, the first implication here is also strict, but this is harder to prove. It is not known whether *all* countably infinite groups have actions that are weakly mixing but not strongly mixing.

The examples we have seen of systems that are ergodic but not weakly mixing have been rotations on compact homogeneous spaces. In fact, it turns out that in a sense these are the *only* examples: any system which is not weakly mixing has a nontrivial factor which is isomorphic to one of these compact homogeneous space rotations. This fact will be the subject of the next lecture.

References

[Fur67] Harry Furstenberg. Disjointness in ergodic theory, minimal sets, and a problem in Diophantine approximation. *Math. Systems Theory*, 1:1–49, 1967.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: `tim@cims.nyu.edu`
URL: `cims.nyu.edu/~tim`

Ergodic Theory

Notes 5: More about mixing

1 Isometric actions

Having introduced weak and strong mixing in Notes 4, we continue to investigate weak mixing by proving a much more sophisticated equivalent characterization. The first step is a more detailed study of a special class of systems, generalizing the torus rotations.

Definition 1. *A p.-p. Γ-system is **isometric** if it is of the form $(X, \mathcal{B}(X), \mu, T)$, where (X, d) is a compact metric space and T^γ is an isometry of (X, d) for every γ. Often we denote this situation by $(X, d, \mathcal{B}(X), \mu, T)$*

*Slightly more generally, a p.-p. Γ-system is **Polish isometric** if it is as above, but where (X, d) is Polish but not necessarily compact.*

Despite their superficial similarity, actions by isometries form a *much* more restricted class than arbitrary topological systems.

Lemma 2. *The only weakly mixing isometric system is the trivial system.*

Proof. If $(X, d, \mathcal{B}(X), \mu, T)$ is isometric, then the metric $d : X \times X \longrightarrow [0, \infty)$ is $(T \times T)$-invariant. Therefore, if the system is weakly mixing, d must be $(\mu \otimes \mu)$-a.s. constant.

For any $\varepsilon > 0$, the compact space X may be covered by finitely many balls of radius ε, it follows that some balls of radius ε have positive measure, implying that $d(x, y) \leqslant 2\varepsilon$ for a positive-measure set of pairs $(x, y) \in X^2$. Since ε was arbitrary, this means that $d(x, y) = 0$ for $(\mu \otimes \mu)$-a.e. (x, y), and hence μ is equal to the Dirac mass δ_x for some (indeed, for μ-a.e.) x. $\square$

Lemma 3. *If (X, d) is a compact metric space, and the isometry group $\mathrm{Isom}(X, d)$ is given the topology of uniform convergence, then it is a compact metric topological group.*

Proof. The topology of uniform convergence on $\mathrm{Isom}(X, d)$ is generated by the metric

$$D(S,T) := \sup_{x \in X} d(Sx, Tx).$$

Since both S and T are d-preserving, one always has

$$\begin{aligned} D(S,T) = \sup_x d(x, S^{-1}(Tx)) &= \sup_x d(T^{-1}(Tx), S^{-1}(Tx)) \\ &= \sup_y d(T^{-1}y, S^{-1}y) = D(T^{-1}, S^{-1}), \end{aligned} \quad (1)$$

because both x and $y := Tz$ range over the whole of X in these suprema.

It is easily checked that composition and inversion define continuous operations

$$\mathrm{Isom}(X,d) \times \mathrm{Isom}(X,d) \longrightarrow \mathrm{Isom}(X,d) \quad \text{and} \quad \mathrm{Isom}(X,d) \longrightarrow \mathrm{Isom}(X,d)$$

for this metric, so it remains to check that it is complete and totally bounded.

Completeness If $(S_n)_n$ is a Cauchy sequence for D, then so is $(S_n^{-1})_n$, by (1). Both are uniformly convergent sequences of 1-Lipschitz functions $X \longrightarrow X$, and so both sequences converge uniformly to some limiting 1-Lipschitz functions. Having observed this, the fact that composition is continuous under the metric D implies that those limits functions are also inverse to each other, hence are isometries.

Cauchy subsequences Now let $(S_n)_{n \geqslant 1}$ be an arbitrary sequence in $\mathrm{Isom}(X, d)$. Since (X, d) is compact, it is separable, so we may choose a countable dense sequence $(x_m)_{m \geqslant 1}$ in X. Now, for each m, the sequence of images $(S_n(x_m))_{n \geqslant 1}$ lies in the compact space X. Therefore, a diagonal argument gives a subsequence $(S_{n_i})_{i \geqslant 1}$ such that $(S_{n_i}(x_m))_{i \geqslant 1}$ is Cauchy for each m, and also $(S_{n_i}(x_m))_{i \geqslant 1}$ is Cauchy for each m.

Now, for any $\varepsilon > 0$, we may choose some finite m_1 such that $X = \bigcup_{m \leqslant m_1} B_\varepsilon(x_m)$. Given this, we may choose some i_1 such that

$$\begin{aligned} i, j \geqslant i_1 \quad \Longrightarrow \quad & \max_{m \leqslant m_1} d(S_{n_i}(x_m), S_{n_j}(x_m)) < \varepsilon \\ & \qquad \Longrightarrow \quad \sup_{x \in X} d(S_{n_i}(x), S_{n_j}(x)) < 3\varepsilon. \end{aligned}$$

This proves that $(S_{n_i})_{i \geqslant 1}$ is Cauchy in the uniform topology. $\square$

For these systems, we now have the following generalization of the unique ergodicity of irrational circle rotations.

Lemma 4. *Let* (X, d) *be a compact metric space, let* $G := \mathrm{Isom}(X, d)$*, let* $\phi : \Gamma \longrightarrow \mathrm{Isom}(X, d)$ *be a homomorphism, and let* μ *be a* $\phi(\Gamma)$*-invariant and ergodic Borel probability measure on* X*. Let* $K := \overline{\phi(\Gamma)} \leqslant G$*. Then* μ *is supported on a single* Kx *for some* $x \in X$*, and is given by the formula*

$$\mu(A) = m\{k \in K : \ kx \in A\},$$

where m *is the Haar probability measure on* K*.*

Proof. Step 1. By assumption, any element of $\phi(\Gamma) = \{\phi(\gamma) : \ \gamma \in \Gamma\} \subseteq \mathrm{Isom}(X, d)$ preserves μ. We first show that this fact extends to any element of $K := \overline{\phi(\Gamma)}$. To see this, suppose that $f \in C(X)$ and $\phi(\gamma_n) \longrightarrow g \in K$. Then f is uniformly continuous, and this implies that $\|f \circ \phi(\gamma_n) - f \circ g\|_\infty \longrightarrow 0$, and hence

$$\int f \, \mathrm{d}\mu = \int f \, \mathrm{d}(\phi(\gamma_n)_* \mu) = \int f \circ \phi(\gamma_n) \, \mathrm{d}\mu \longrightarrow \int f \circ g \, \mathrm{d}\mu = \int f \, \mathrm{d}(g_* \mu).$$

Since f was arbitrary, this implies that $\mu = g_* \mu$. Since g was any element of K, it follows that K preserves μ.

Step 2. Since μ is a Borel probability measure on X, there is some $x \in X$ for which $\mu(B_\varepsilon(x)) > 0$ for all $\varepsilon > 0$ (otherwise, by compactness, we could cover X by finitely many open balls of measure 0). Also, μ-a.e. x has this property.

Fix such an x, and observe that its orbit Kx is a compact subset of X, because it is the image of the compact group K under the continuous maps $k \mapsto kx$. Now consider the function

$$f : Y \longrightarrow [0, \infty) : y \mapsto \mathrm{dist}(y, Kx).$$

For this function, we have $\{f < \varepsilon\} \supseteq B_\varepsilon(x)$, and hence

$$\mu\{f < \varepsilon\} \geqslant \mu(B_\varepsilon(x)) > 0 \quad \forall \varepsilon > 0.$$

On the other hand, it is continuous, hence measurable, and it is K-invariant by construction. Therefore, by ergodicity, it is μ-a.s. constant. Combining these facts, it must equal 0 almost surely, and hence

$$\mu\{f = 0\} = \mu(Kx) = 1.$$

That is, μ is supported on Kx. Since μ-a.e. point had the required property of x, it also follows that $Kx = Ky$ for μ-a.e. y.

Step 3. Finally, since μ is K-invariant, for any $A \in \mathcal{B}(X)$ we have $\mu(A) = k_*\mu(A)$ for all k. Averaging over k, this becomes

$$\mu(A) = \int_K k_*\mu(A)\, m(\mathrm{d}k) = \int_K \int_Y 1_A(ky)\, \mu(\mathrm{d}y) m(\mathrm{d}k) = \int_X \int_K 1_A(ky)\, m(\mathrm{d}k)\mu(\mathrm{d}y) = \int_X m\{k :\ ky \in A\}\, \mu(\mathrm{d}y).$$

However, if $Kx = Ky$, then $y = k_0x$ for some k_0, and now

$$m\{k :\ ky \in A\} = m\{k :\ kk_0x \in A\} = m\{k :\ kx \in A\},$$

because m is rotation-invariant. Therefore the above integrand is equal to $m\{k : kx \in A\}$ for μ-a.e. y, and so μ is given by the desired formula. □

Corollary 5. *In the setting above, let T denote the resulting p.-p. Γ-action on Kx: $T^\gamma y := \phi(\gamma)(y)$. Then there is an isomorphism of topological systems*

$$\overline{\psi} : (K/H, R_\phi) \xrightarrow{\cong} (Kx, T),$$

where

$$H := \mathrm{Stab}_K x = \{k \in K :\ kx = x\}$$

and R_ϕ is the compact homogeneous space rotation-action $K \curvearrowright K/H$, and such that $\mu = \overline{\psi}_* m_{K/H}$ for the Haar probability measure on $m_{K/H}$.

Proof. Consider the orbit-map $\psi : K \longrightarrow Kx$ defined by $\psi(k) = kx$. As remarked above, this is continuous map from one compact space to another, and it is surjective, by definition. Also, by a standard manipulation, if $h \in H := \mathrm{Stab}_K x$, then $\psi(kh) = khx = kx = \psi(k)$ for all k, so we obtain a well-defined map $\overline{\psi} : K/H \longrightarrow Kx$ by letting $\overline{\psi}(kH) := kx$. It is easily checked that this is still continuous, and it is now also injective because

$$\overline{\psi}(kH) = \overline{\psi}(k'H) \iff kx = k'x \iff (k^{-1}k')x = x \iff k^{-1}k' \in H \iff kH = k'H.$$

Therefore, $\overline{\psi}$ defines an isomorphism of topological systems

$$(K/H, R_\phi) \longrightarrow (Kx, T),$$

which therefore also respects the Borel σ-algebras. To finish, observe that $\mu = \overline{\psi}_* m_{K/H}$, by the formula for μ obtained in the lemma. □

This, ergodic isometric systems are simply the same (up to isomorphism) as the compact homogeneous-space rotations introduced previously.

We next prove that one can identify a canonical maximal isometric factor in an arbitrary ergodic system.

Lemma 6. *Let (X, Σ, μ, T) be a system in which Σ is countably-generated up to μ-negligible sets. Then this system has a unique largest factor (up to negligible sets) which is generated by a factor map to an isometric action. This is called its **Kronecker factor**.*

Proof. Let $\mathcal{C} \subseteq \Sigma$ be the collection of all sets that are pulled back under some isometric factor maps. We must show that there is a single isometric factor map that contains every element of $\mathcal{C}$ up to μ-negligible sets.

Since the measure-metric $(\mathfrak{A}_\mu, d_\mu)$ is separable, we may choose a sequence A_i in $\mathcal{C}$ such that $\{[A_i] : i \in \mathbb{N}\}$ is d_μ-dense in $\{[C] : C \in \mathcal{C}\}$. It now suffices to find a single isometric factor map that generates every A_i. For each i, let $\pi_i : (X, \Sigma, \mu, T) \longrightarrow (Z_i, \mathcal{B}(Z_i), \nu_i, R_i)$ be a sequence of isometric factor maps on compact metric spaces (Z_i, d_i) such that A_i lies in the μ-completion of $\pi_i^{-1}(\mathcal{B}(Z_i))$.

By normalizing, we may assume that $\operatorname{diam} d_i \leqslant 1$ for all i. Now the single map

$$\pi : x \mapsto (\pi_1(x), \pi_2(x), \ldots) \in \prod_{i \geqslant 1} Z_i$$

intertwines T with $R := R_1 \times R_2 \times \cdots$, which is still an isometric action for the metric

$$d((z_i)_i, (z'_i)_i) := \sum_{i \geqslant 1} 2^{-i} d_i(z_i, z'_i),$$

and π pushes μ to some R-invariant measure on Z. Since $A_i \in \pi^{-1}(\mathcal{B}(Z))$ for every i, this completes the proof. □

We will also need to know that under an ergodicity assumption, isometric and Polish isometric actions are essentially the same.

Lemma 7. *Suppose that $(\mathfrak{X}, d)$ is a Polish metric space, μ is a Borel probability measure on $\mathfrak{X}$, and $T : \Gamma \curvearrowright \mathfrak{X}$ is an μ-preserving action by isometries for which μ is ergodic. Then $\operatorname{spt} \mu$ is compact.*

Proof. Recall that $x \in \mathfrak{X}$ lies in $\operatorname{spt} \mu$ if and only if $\mu(B_r(x)) > 0$ for all $r > 0$. Since $\mu(\operatorname{spt} \mu) = 1$, this implies that for every $n \geqslant 1$ there is some $r_n > 0$ such that

$$\mu\{x : \mu(B_{1/n}(x)) > r_n\} > 1 - 2^{-n-1},$$

and hence, by the Borel-Cantelli Lemma, $\mu(Z) > 1/2$, where

$$Z := \bigcap_{n \geqslant 1} \{x : \ \mu(B_{1/n}(x)) > r_n\}.$$

On the other hand, since T preserves both d and μ, one has

$$\mu(B_{1/n}(x)) = \mu(T^\gamma B_{1/n}(x)) = \mu(B_{1/n}(T^\gamma x)) \quad \forall x \in X, \ \gamma \in \Gamma,$$

and so Z is T-invariant. Therefore $\mu(Z) = 1$, by ergodicity.

Finally, suppose that $n \geqslant 1$, and let $S \subseteq Z$ be a maximal set of points that are pairwise separated by distance at least $1/n$. Then, on the one hand, all the balls $B_{1/2n}(x)$ for $x \in S$ are disjoint and have measure at least ε_{2n}, so S must be finite. On the other hand, the family of balls $(B_{1/n}(x))_{x \in S}$ must cover the whole of Z, by the maximality of Z. Since n was arbitrarily large, this has shown that Z is totally bounded, and so $\mathrm{spt}\,\mu \subseteq \overline{Z}$ must be compact. □

2 Isometric actions are the obstructions to weak mixing

We can now state and prove a precise characterization of the failure of weak mixing.

Theorem 8. *If (X, Σ, μ, T) and (Y, Φ, ν, S) are ergodic p.-p. Γ-systems and $U \in \Sigma \otimes \Phi$ is $(T \times S)$-invariant, then U is measurable with respect to $\Xi \otimes \Lambda$, where Ξ and Λ are Kronecker factors of T and S, respectively. In particular, T is weakly mixing if and only its Kronecker factor is trivial.*

The proof of Theorem 8 can be presented in more than way. Our presentation will make use of the measure-metric spaces $(\mathfrak{A}_\mu, d_\mu)$ and $(\mathfrak{A}_\nu, d_\nu)$ associated to these p.-p.-systems: recall Lecture 1. As previously, we may assume that Σ and Φ are countably-generated up to negligible sets, so these measure-metric spaces are Polish.

Lemma 9. *Suppose that (X, Σ, μ) and (Y, Φ, ν) are complete probability spaces, that $\Lambda \leqslant \Phi$ is a ν-complete σ-subalgebra, and that U is in the $(\mu \otimes \nu)$-completion of $\Sigma \otimes \Phi$. For each $y \in Y$ let $U_y := \{x : \ (x, y) \in U\}$. Given the measurability of U, Fubini's Theorem implies that $U_y \in \Sigma$ for ν-almost every y.*

Assume further that the function

$$Y \longrightarrow \mathfrak{A}_\mu : y \mapsto [U_y]$$

is measurable with respect to Λ. Then U lies in the $(\mu \otimes \nu)$-completion of $\Sigma \otimes \Lambda$.

Proof. This proof has a very similar flavour to the implication ($3. \implies 2.$) in Proposition 16 of Notes 4.

First observe that for any $f \in L^1(\mu)$ and $g \in L^1(\nu)$, our assumption on U and Fubini's Theorem give that

$$\int 1_U \cdot (f \otimes g)\,\mathrm{d}(\mu \otimes \nu) = \int_Y g(y)\Big(\int_{U_y} f(x)\,\mu(\mathrm{d}x)\Big)\,\nu(\mathrm{d}y).$$

Since the expression $\int_{U_y} f(x)\,\mu(\mathrm{d}x)$ depends only on $[U_y]$, it is Λ-measurable as a function of y, and so the above is equal to

$$\int_Y \mathsf{E}(g\,|\,\Lambda)(y)\Big(\int_{U_y} f(x)\,\mu(\mathrm{d}x)\Big)\,\nu(\mathrm{d}y) = \int 1_U \cdot (f \otimes \mathsf{E}(g\,|\,\Lambda))\,\mathrm{d}(\mu \otimes \nu).$$

To use this fact, recall that, by the definition of $\Sigma \otimes \Phi$, for every $\varepsilon > 0$ there is a finite family of disjoint measurable rectangles $A_i \times B_i \in \Sigma \otimes \Phi$, $i = 1, 2, \dots, \ell$, such that

$$\Big\|1_U - \sum_i 1_{A_i} \otimes 1_{B_i}\Big\|_1 < \varepsilon.$$

The trick is to use this rectangle approximation to compare 1_U with itself:

$$\begin{aligned}
(\mu \otimes \nu)(U) &= \int 1_U \cdot 1_U\,\mathrm{d}(\mu \otimes \nu)\\
&\overset{\varepsilon}{\approx} \int 1_U \cdot \Big(\sum_i 1_{A_i} \otimes 1_{B_i}\Big)\,\mathrm{d}(\mu \otimes \nu)\\
&= \int 1_U \cdot \Big(\sum_i 1_{A_i} \otimes \mathsf{E}(1_{B_i}\,|\,\Lambda)\Big)\,\mathrm{d}(\mu \otimes \nu)\\
&= \int \mathsf{E}(1_U\,|\,\Sigma \otimes \Lambda) \cdot \Big(\sum_i 1_{A_i} \otimes \mathsf{E}(1_{B_i}\,|\,\Lambda)\Big)\,\mathrm{d}(\mu \otimes \nu)\\
&= \int \mathsf{E}(1_U\,|\,\Sigma \otimes \Lambda) \cdot \Big(\sum_i 1_{A_i} \otimes 1_{B_i}\Big)\,\mathrm{d}(\mu \otimes \nu)\\
&\overset{\varepsilon}{\approx} \int \mathsf{E}(1_U\,|\,\Sigma \otimes \Lambda) \cdot 1_U\,\mathrm{d}(\mu \otimes \nu).
\end{aligned}$$

Since ε was arbitrary, this implies that $f := \mathsf{E}(1_U\,|\,\Sigma \otimes \Lambda)$ is a function taking values in $[0, 1]$, and is such that

$$\int_X f\,\mathrm{d}(\mu \otimes \nu) = \int_U f\,\mathrm{d}(\mu \otimes \nu) = (\mu \otimes \nu)(U).$$

Therefore f must equal 1 on almost all of U and 0 on almost all of $(X \times Y)\backslash U$. Therefore U agrees up to a negligible set with $\{f = 1\}$, which lies in $\Sigma \otimes \Lambda$. $\square$

Proof of Theorem 8. Suppose that $U \subseteq X \times X$ is $(T \times S)$-invariant, and consider the function

$$\phi : Y \longrightarrow \mathfrak{A}_\mu : y \mapsto [U_y]$$

which is well-defined μ-a.e. Another easy exercise using the measurability of U shows that this function is also measurable.

On the Polish space $(\mathfrak{A}_\mu, d_\mu)$, define the Γ-action $\tilde{T}$ by $\tilde{T}^\gamma[A] := [T^\gamma A]$. This is a well-defined isometric action because T is μ-preserving.

Now, the $(T \times S)$-invariance of U implies that

$$U_{S^\gamma y} = T^\gamma U_y \quad \forall y \in X,\ \gamma \in \Gamma,$$

and so $\phi : (Y, S) \longrightarrow (\mathfrak{A}_\mu, \tilde{T})$ is equivariant. Let $\tilde{\mu} := \phi_* \nu$, a Borel probability measure on $\mathfrak{A}_\mu$. Then this gives that ϕ is a factor map from (Y, Φ, ν, S) to the Polish isometric system $(\mathfrak{A}_\mu, d_\mu, \mathcal{B}(\mathfrak{A}_\mu), \tilde{\mu}, \tilde{T})$. Since (Y, Φ, ν, S) is ergodic, so is the target system, and therefore Lemma 7 implies that $Z := \operatorname{spt} \tilde{\mu}$ is compact. Since ϕ a.s. takes values in Z, it still defines a factor map $Y \longrightarrow Z$, as required.

By construction, $y \mapsto [U_y]$ is measurable with respect to this factor, and therefore U is $(\Sigma \otimes \Lambda)$-measurable by Lemma 9. Applying the same reasoning with the rôles of X and Y reversed completes the proof. □

3 Representation-theoretic interpretation

This section will give another characterization of the failure of weak mixing, this time in terms of the unitary representation $U_T : \Gamma \curvearrowright L^2(\mu)$ associated to a p.-p. Γ-system (X, Σ, μ, T) by $U_T^\gamma f := f \circ T^{\gamma^{-1}}$. This is referred to as the **Koopman representation** associated to the action T.

The following is a counterpart to Theorem 8 which describes the failure of weak mixing in terms of this representation.

Theorem 10. *If (X, Σ, μ, T) is ergodic and $H \in L^2(\mu^{\otimes 2})$ is $(T \times T)$-invariant, then H is contained in the closed linear span of*

$$\{\phi \otimes \psi : \ \phi, \psi \text{ lie in finite-dimensional } U_T\text{-invariant subspaces}\}.$$

In case $\Gamma = \mathbb{Z}^d$, it is relatively easy to deduce this result from the description of the commuting family of unitary operators $U_T^{\mathbf{n}} f := f \circ T^{-\mathbf{n}}$ on $L^2(\mu)$ given by the spectral theorem. For a general group Γ, whose unitary representation theory may be very messy and complicated, this approach is unavailable, and even for $\mathbb{Z}^d$ it is instructive to take a more elementary approach. One tool that we will need is the spectral theorem for compact self-adjoint operators, recalled with the other preliminaries in Notes 1:

Theorem 11 (Spectral Theorem for compact self-adjoint operators)**.** *If M is a compact self-adjoint operator on $L^2(\mu)$ then there are an orthonormal sequence ξ_1, ξ_2, $\ldots \in L^2(\mu)$ and real numbers $\lambda_1, \lambda_2, \ldots$ such that $|\lambda_1| \geqslant |\lambda_2| \geqslant \ldots$, $|\lambda_i| \longrightarrow 0$ and*

$$M = \sum_{i \geqslant 1} \lambda_i P_{\xi_i},$$

where P_{ξ_i} is the one-dimensional orthogonal projection onto $\mathbb{C} \cdot \xi_i$. □

Our application of this result rests on the following important identification.

Lemma 12. *If (X, Σ, μ) is a probability space and $K \in L^2(\mu^{\otimes 2})$ then the operator $K\star : L^2(\mu) \longrightarrow L^2(\mu)$ defined by*

$$K \star f(x) := \int_X K(x,y) f(y)\, \mu(\mathrm{d}y)$$

is (bounded and) compact.

Proof. Let us first prove that $K\star$ is bounded, and more concretely that

$$\|K \star\|_{\mathrm{op}} \leqslant \|K\|_{L^2(\mu^{\otimes 2})}.$$

This follows from a simple appeal to the Cauchy-Schwartz inequality:

$$\begin{aligned} \|K \star f\|_2^2 &= \int_X \Big| \int_X K(x,y) f(y)\, \mu(\mathrm{d}y) \Big|^2 \mu(\mathrm{d}x) \\ &\leqslant \int_X \Big(\int_X |K(x,y)|^2\, \mu(\mathrm{d}y) \Big) \|f\|_2^2\, \mu(\mathrm{d}x) \\ &= \|K\|_{L^2(\mu^{\otimes 2})}^2 \|f\|_2^2. \end{aligned}$$

Having proved this, we now recall that as a square integrable function on $X \times X$, K may be approximated in $L^2(\mu^{\otimes 2})$ by finite sums of tensor product functions: that is, functions of the form

$$L(x,y) := \sum_{i=1}^{k} \phi_i(x) \psi_i(y) \quad \text{with } \phi_i, \psi_i \in L^2(\mu) \text{ for } i \leqslant k.$$

Let B denote the unit ball in $L^2(\mu)$. For any such L the associated operator $L\star$ has image contained in $\mathrm{span}\{\phi_1, \ldots, \phi_k\}$, so gives $L \star (B)$ contained in a finite-dimensional subspace and is therefore certainly compact. Since for any $\varepsilon > 0$ we may find such an L such that $\|K \star - L \star\|_{\mathrm{op}} \leqslant \|K - L\|_2 < \varepsilon/2$, and now the unit ball image $L \star (B)$ admits a finite covering by $(\varepsilon/2)$-balls, enlarging each of these to an ε-ball gives a finite covering of $K \star (B)$, so since ε was arbitrary this latter set is precompact. □

Proof of Theorem 10. First, let us write

$$H(x,y) := K(x,y) + \mathrm{i}L(x,y)$$

with $K(x,y) := \frac{1}{2}(H(x,y) + \overline{H(x,y)})$ and $L(x,y) := \frac{\mathrm{i}}{2}(\overline{H(x,y)} - H(x,y))$, so that K enjoys the skew-symmetry $K(y,x) = \overline{K(x,y)}$ and similarly for L. Clearly it suffices to prove the result for each of K and L separately. We give the proof for K, the case of L being identical.

By Lemma 12 the associated operator $K\star$ is compact, and in addition, the skew-symmetry of K amounts to the self-adjointness of the operator:

$$\begin{aligned}
\langle g, K\star f\rangle &= \int_X \overline{g(x)} \int_X K(x,y) f(y)\,\mu(\mathrm{d}y)\,\mu(\mathrm{d}x) \\
&= \int_X \int_X \overline{g(x)} K(x,y) f(y)\,\mu(\mathrm{d}y)\,\mu(\mathrm{d}x) \\
&= \int_X \int_X \overline{g(x) K(y,x)} f(y)\,\mu(\mathrm{d}x)\,\mu(\mathrm{d}y) \\
&= \langle K\star g, f\rangle.
\end{aligned}$$

Therefore Theorem 11 gives a decomposition

$$K\star = \sum_{i\geqslant 1} \lambda_i P_{\phi_i}$$

for some orthonormal sequence of eigenvectors ϕ_i with real eigenvalues $\lambda_i \longrightarrow 0$. Writing this out in terms of functions on $X \times X$, we have $P_{\phi_i} = (\phi_i \otimes \overline{\phi_i})\star$, and hence

$$K(x,y) = \sum_{i\geqslant 1} \lambda_i \phi_i(x) \overline{\phi_i(y)}$$

as a convergent series in $\|\cdot\|_{\mathrm{op}}$. To be careful, we should note at this point that because the ϕ_i are orthonormal, this convergence in operator norm implies

$$\int_{X\times X} K(\overline{\phi_j} \otimes \phi_j)\,\mathrm{d}\mu^{\otimes 2} = \langle \phi_j, K\star\phi_j\rangle = \lim_{m\longrightarrow\infty} \sum_{i=1}^{m} \lambda_i |\langle \phi_j, \phi_i\rangle|^2 = \lambda_j \|\phi_j\|_2^2 = \lambda_j,$$

and so since the functions $\overline{\phi_j}\otimes\phi_j$ are also orthonormal we must have $\sum_j |\lambda_j|^2 < \infty$ and therefore the above series also converges in $L^2(\mu^{\otimes 2})$.

Finally, the $(T \times T)$-invariance of K implies that $(K\star) \circ U_T^\gamma = U_T^\gamma \circ (K\star)$ for all $\gamma \in \Gamma$, and hence that all of the non-trivial eigenspaces of $K\star$ are U_T-invariant. Since these are finite-dimensional subspaces of $L^2(\mu)$ and each ϕ_i is contained in one of them, this fact and the above series expansion of K complete the proof. $\square$

4 The existence of weak mixing without strong mixing

To round out our discussion of weak and strong mixing, it seems worth proving that for $\Gamma = \mathbb{Z}$ they are not the same.

It is possible to describe concrete examples of p.-p.t.s that are weakly but not strongly mixing, but surprisingly tricky (two such are given in Section 4.5 of Petersen [Pet83]). However, even more surprisingly, it turns out that 'almost all' systems are examples (in the sense of Baire category), notwithstanding the difficulty of isolating any given one of them. This means that it is possible to prove their existence without describing one in detail (although, of course, some kind of description can always be recovered by unpicking the Baire category proof to produce a sequence that converges to an example).

Here we will take a slightly unusual approach to this Baire category argument because it is quicker, by working in a space of measures rather than of transformations.

Theorem 13. *Let S be the right-shift transformation on the Hilbert cube $X := [0,1]^{\mathbb{Z}}$, and let the convex set of invariant measures $\mathrm{Pr}^S(X)$ be given its vague topology, in which it is compact and metrizable. Then the subsets*

$$P := \{\mu \in \mathrm{Pr}^S(X) : \ (X, \mathcal{B}(X), \mu, S) \text{ is strongly mixing}\}$$

and

$$Q := \{\mu \in \mathrm{Pr}^S(X) : \ (X, \mathcal{B}(X), \mu, S) \text{ is weakly mixing}\}$$

are respectively meagre (that is, first category) and co-meagre (second category) in this topology. Consequently $Q \backslash P$ is nonempty.

Proof. Clearly there are two parts to this.

(1) In fact it suffices to treat the mixing properties of just one subset of X. Let $A := \{(x_i)_i \in X : \ x_0 \geqslant 1/2\}$. Clearly

$$\begin{aligned} P &\subseteq \{\mu \in \mathrm{Pr}^S(X) : \ \mu(A \cap S^n A) \longrightarrow \mu(A)^2 \text{ as } |n| \longrightarrow \infty\} \\ &\subseteq \Big\{\mu \in \mathrm{Pr}^S(X) : \ |\mu(A \cap S^n A) - \mu(A)^2| \leqslant \frac{1}{10}(\mu(A) - \mu(A)^2) \\ &\qquad\qquad\qquad\qquad\qquad\qquad \text{for } |n| \text{ sufficiently large}\Big\}, \end{aligned}$$

so it suffices to prove that this last set is meagre (so we can really prove our result with considerable room to spare). This, in turn, is contained in

$$\bigcup_{m \geqslant 1} \bigcap_{n \geqslant m} \Big\{\mu \in \mathrm{Pr}^S(X) : \ |\mu(A \cap S^n A) - \mu(A)^2| \leqslant \frac{1}{8}(\mu(A) - \mu(A)^2)\Big\},$$

which is a countable union of closed sets, so it remains to prove that each of these is nowhere-dense. Thus, given $m \geqslant 1$, $\mu \in \Pr^S(X)$ and some vague neighbourhood $U \ni \mu$, we must find $\mu' \in U$ such that

$$|\mu'(A \cap S^n A) - \mu'(A)^2| > \frac{1}{10}(\mu'(A) - \mu'(A)^2)$$

for some $n \geqslant m$.

To do this, first perturb μ very slightly (well within U) to assume that $0 < \mu(A) < 1$ (although it may be very close to 0 or 1). Now by shrinking the neighbourhood U if necessary, we may assume it is of the form

$$\Big\{\nu \in \Pr^S(X) : \Big|\int_X f_i \,\mathrm{d}\nu - \int_X f_i \,\mathrm{d}\mu\Big| < \varepsilon \;\forall i \leqslant k\Big\}$$

for some continuous functions $f_1, f_2, \ldots, f_k \in C(X)$ and $\varepsilon > 0$. Morever, since continuous functions that depend on only finitely many coordinates in X are uniformly dense in $C(X)$, by shrinking ε a little further we may replace each f_i with $f_i \circ \pi$ for the coordinate projection $\pi : X \longrightarrow [0,1]^F$ for some fixed finite $F \subset \mathbb{Z}$ and $f_i \in C([0,1]^F)$.

However, if we now let M be some integer that is much larger than both $\max\{|n| : n \in F\}$ and m, we may define a new probability measure to be the law of the random $[0,1]$-sequence produced as follows: draw $(x_i)_i$ from the measure μ, consider its $\{-M, -M+1, \ldots, M\}$-indexed block

$$x_{-M}, x_{-M+1}, \ldots, x_M,$$

and repeat this periodically. This defines a new (possibly non-invariant) measure $\theta \in \Pr([0,1]^{\mathbb{Z}})$ concentrated on strings which are periodic with period $2M+1$, and which clearly still satisfies $\theta(S^n A) = \mu(A)$ for all n. Now define μ' by re-randomizing over the $2M+1$ distinct translates of these strings:

$$\mu' := \frac{1}{2M+1}\sum_{n=-M}^{M} S^n_* \theta.$$

This recovers for us a shift-invariant measure, since the shifts of θ are $(2M+1)$-periodic. Among these shifts $S^n_*\theta$ for $-M \leqslant n \leqslant M$, at least $2M - 2\max\{|n| : n \in F\} - 2$ of them have the same marginal on $[0,1]^F$ as μ itself, so provided M is sufficiently large it follows that most of these translates are very close to μ and hence also that $\mu' \in U$. Morever we still have $\mu'(A) = \mu(A)$.

On the other hand, μ' is concentrated on strings that are $(2M+1)$-periodic, so if n is any non-zero multiple of $2M+1$ then $\mu'(A \triangle S^n A) = 0$ and therefore

$$\mu'(A \cap S^n A) = \mu'(A) > \mu'(A)^2 + \frac{1}{10}(\mu'(A) - \mu'(A)^2),$$

as required.

(2) This time we first recall the following equivalent characterization of the weak mixing of $(X, \mathcal{B}(X), \mu, T)$ proved in Notes 4: for all $A_1, A_2, \dots, A_k \in \mathcal{B}(X)$ and $\varepsilon > 0$ there is some n such that $|\mu(A_i \cap S^n A_j) - \mu(A_i)\mu(A_j)| < \varepsilon$ for all $i, j \leqslant k$. Moreover, it is clearly equivalent to test this condition on some countable μ-dense subfamily of $\mathcal{B}(X)$, and in the case of a probability measure on $[0,1]^{\mathbb{Z}}$ such a countable family is always offered by the collection of finite unions of products of rational intervals that depend on only finitely many coordinates. This is important, because it means we can choose a single sequence $A_1, A_2, \dots$ of Borel subsets of X, each depending on only finitely many coordinates, which is μ-dense in $\mathcal{B}(X)$ for every $\mu \in \mathrm{Pr}^S(X)$, and so we have

$$Q = \bigcap_{k \geqslant 1} \bigcap_{n \geqslant 1} \bigcup_{m \in \mathbb{Z}} \big\{\mu \in \mathrm{Pr}^S(X) : \, |\mu(A_i \cap S^m A_j) - \mu(A_i)\mu(A_j)| < 1/n \; \forall i, j \leqslant k\big\}.$$

Since this is a countable intersection, to prove it is co-meagre we must prove that each of the open sets

$$\bigcup_{m \in \mathbb{Z}} \big\{\mu \in \mathrm{Pr}^S(X) : \, |\mu(A_i \cap S^m A_j) - \mu(A_i)\mu(A_j)| < 1/n \; \forall i, j \leqslant k\big\}$$

is dense in $\mathrm{Pr}^S(X)$. Suppose that each A_i for $i \leqslant k$ depends only on coordinates in some finite set $E \subset \mathbb{Z}$, once again let U be an open subset of $\mathrm{Pr}^S(X)$, and arguing as before assume it is of the form

$$\Big\{\nu \in \mathrm{Pr}^S(X) : \, \Big|\int_X f_i \circ \pi \,\mathrm{d}\nu - \int_X f_i \circ \pi \,\mathrm{d}\mu\Big| < \varepsilon \; \forall i \leqslant k\Big\}$$

for some finite-dimensional projection $\pi : X \longrightarrow [0,1]^F$, continuous functions $f_1, f_2, \dots, f_k \in C([0,1]^F)$ and $\varepsilon > 0$. We will show that U must contain a measure μ' such that $|\mu'(A_i \cap S^m A_j) - \mu'(A_i)\mu'(A_j)| < 1/n$ for all $i, j \leqslant k$ for some $m \in \mathbb{Z}$.

To do this, again let M be an integer much larger than $\max\{|n| : \, n \in F\}$, and construct the law $\theta \in \mathrm{Pr}(X)$ of a new random string $(x_i)_i$ by now drawing each of the finite blocks

$$x_{kM}, x_{kM+1}, \dots, x_{(k+1)M-1}$$

independently from the corresponding finite-dimensional marginal of μ (of course, if we draw just one infinite string from μ, then these finite blocks of it will generally not be independent, so θ is usually 'more random' than μ). Once again, θ may not be S-invariant, but it is S^M-invariant and so the average

$$\mu' := \frac{1}{M} \sum_{i=1}^{M} S^i_* \theta$$

is S-invariant. If M is sufficiently large then this μ' lies in U, again because for most $n \in \{1, 2, \ldots, M\}$ the marginal of $S_*^n\theta$ onto $[0,1]^F$ agrees with that of μ. On the other hand, for any $m \geqslant M + \mathrm{diam}(E)$ the coordinate projections π_E and π_{E+m} are independent under μ', and hence

$$\mu'(A_i \cap S^m A_j) = \mu'(A_i)\mu'(A_j) \quad \forall i, j \leqslant k,$$

as required. □

Remark. The classical Baire category proof of the above result (parts (1) and (2) being due to Rokhlin and Halmos respectively, both in the 1940s) begins by fixing a compact metric space with a probability measure, such as $[0,1]$ with Lebesgue measure μ. One then defines the **coarse** (or **weak**) topology on the group $\mathrm{Aut}([0,1],\mu)$ of invertible Borel μ-preserving transformations of $[0,1]$ so that $S, T \in \mathrm{Aut}([0,1],\mu)$ are 'close' if for some finite list $A_1, A_2, \ldots, A_k \in \mathcal{B}([0,1])$ and some $\varepsilon > 0$ we have $\mu(TA_i \triangle SA_i) < \varepsilon$ for all $i \leqslant k$ (equivalently, if U_S and U_T are close in the strong operator topology of $\mathfrak{B}(L^2(\mu))$). It is routine to check that this defines a completely metrizable group topology on $\mathrm{Aut}([0,1],\mu)$, and now corresponding to our two steps above one can prove that the subsets of strongly and weakly mixing members of $\mathrm{Aut}([0,1],\mu)$ are meagre and co-meagre, respectively, so once again their difference (consisting of weakly-but-not-strongly mixing transformations) makes up most of the group.

This approach is a little more complicated than ours, because it is slightly more delicate to modify one transformation into another close to it that has some desired mixing property than it is to do so with measures in $\mathrm{Pr}^{\mathrm{shift}}([0,1]^{\mathbb{N}})$. However, working inside the group of transformations is in some sense more natural, since the arguments essentially don't depend on the choice of the underlying probability space, and also because they can be adapted to prove many other properties of 'generic' transformations which are not so easily accessed by looking at the space of shift-invariant measures. The approach via the space of transformations is nicely treated in Halmos' classic text [Hal60]. ◁

References

[Hal60] Paul R. Halmos. *Lectures on ergodic theory*. Chelsea Publishing Co., New York, 1960.

[Pet83] Karl E. Petersen. *Ergodic Theory*. Cambridge University Press, Cambridge, 1983.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: tim@cims.nyu.edu
URL: cims.nyu.edu/~tim

Ergodic Theory:

Notes 6: Introduction to Szemerédi's Theorem and Multiple Recurrence

The next few sets of notes will be concerned with 'multiple recurrence'. This ergodic-theoretic phenomenon corresponds to a famous result of arithmetic combinatorics, Szemerédi's Theorem, and some of its generalizations.

The classic introduction to this relation between combinatorics and ergodic theory is Furstenberg's book [Fur81]. Some of the treatment here will be closer to the survey [Aus10b].

1 Van der Waerden's Theorem, Szemerédi's Theorem and further generalizations

In 1927 van der Waerden gave a short, clever combinatorial proof of the following surprising fact:

Theorem 1 (Van der Waerden's Theorem [vdW27]). *For any fixed integers $c, k \geqslant 1$, if the elements of $\mathbb{Z}$ are coloured using c colours, then there is a nontrivial k-term arithmetic progression which is monochromatic: that is, there are some $a \in \mathbb{Z}$ and $n \geqslant 1$ such that*

$$a, a+n, \ldots, a+(k-1)n.$$

all have the same colour.

This is now one of the central results in the area of combinatorics called Ramsey Theory. The classic overview of this theory is the book [GRS90] of Graham, Rothschild and Spencer.

Note that it is *crucial* to allow both the start point a and the common difference $n \geqslant 1$ to be chosen freely. In some more difficult versions of this theorem it is possible, for example, to place certain restrictions on what choices we allow for n; but if we try to restrict ourselves a priori to any one value of n then the result is certainly false.

Prompted by van der Waerden's Theorem, in 1936 Erdős and Turán asked whether a deeper phenomenon might lie beneath it. Observe that for any c-colouring of $\mathbb{Z}$ and for any long interval in $\mathbb{Z}$, at least one of the colour-classes must occupy at least a fraction $1/c$ of the points in that interval. In [ET36] they asked whether this feature of a subset of $\mathbb{Z}$ already implies the presence of arithmetic progressions of any finite length k in that subset. It turns out that this is true, but its proof had to wait until Roth treated the case $k = 3$ in [Rot53], and then Szemerédi handled the general case in [Sze75], giving rise to what is now called Szemerédi's Theorem.

The formal statement requires the following definition.

Definition 2 (Upper Banach density)**.** *For $E \subseteq \mathbb{Z}^d$, its **upper Banach density** is the quantity*

$$\bar{\mathrm{d}}(E) := \limsup_{\min_{i\leqslant d}|N_i - M_i|\longrightarrow\infty} \frac{\left|E \cap \prod_{i\leqslant d}[M_i, N_i]\right|}{\prod_{i\leqslant d}(N_i - M_i)}.$$

That is, $\bar{\mathrm{d}}(E)$ is the supremum of those $\delta > 0$ such that one can find boxes in $\mathbb{Z}^d$ of arbitrary long side-lengths such that E contains at least a proportion δ of the lattice points in those box.

Theorem 3 (Szemerédi's Theorem)**.** *If $E \subseteq \mathbb{Z}$ has $\bar{\mathrm{d}}(E) > 0$, then for any $k \geqslant 1$ there are $a \in \mathbb{Z}$ and $n \geqslant 1$ such that*

$$\{a, a + n, \ldots, a + (k-1)n\} \subseteq E.$$

As already remarked, this directly implies van der Waerden's Theorem, because in any finite colouring of $\mathbb{Z}$ at least one of the colour-classes must have positive upper Banach density.

Szemerédi's proof of Theorem 3 is one of the virtuoso feats of modern combinatorics. It was also the first serious outing for several tools that have since become workhorses of that area of mathematics, particularly the Szemerédi Regularity Lemma in graph theory. However, shortly afterwards Furstenberg gave a new proof using ergodic theory. In [Fur77], he first showed the equivalence of Szemerédi's Theorem to an ergodic-theoretic phenomenon called 'multiple recurrence', and then proved that using newly-developed structural results in ergodic theory.

We will introduce multiple recurrence in the next subsection, but let us first bring the combinatorial side of the story closer to the present. Furstenberg and Katznelson quickly realized that Furstenberg's ergodic-theoretic proof could be considerably generalized, and in [FK78] they obtained a multidimensional version of Szemerédi's Theorem as a consequence:

Theorem 4 (Multidimensional Szemerédi Theorem)**.** *If $E \subseteq \mathbb{Z}^d$ has $\bar{\mathrm{d}}(E) > 0$, and $\mathbf{e}_1, \ldots, \mathbf{e}_d$, is the standard basis in $\mathbb{Z}^d$, then there are some $\mathbf{a} \in \mathbb{Z}^d$ and $n \geqslant 1$ such that*

$$\{\mathbf{a} + n\mathbf{e}_1, \ldots, \mathbf{a} + n\mathbf{e}_d\} \subseteq E$$

(so 'dense subsets contain the set of outer vertices of an upright right-angled simplex').

This easily implies Szemerédi's Theorem, because if $k \geqslant 1$, $E \subseteq \mathbb{Z}$ has $\bar{\mathrm{d}}(E) > 0$, and we define

$$\Pi : \mathbb{Z}^{k-1} \longrightarrow \mathbb{Z} : (a_1, a_2, \ldots, a_{k-1}) \mapsto a_1 + 2a_2 + \cdots + (k-1)a_{k-1},$$

then the pre-image $\Pi^{-1}(E)$ has $\bar{\mathrm{d}}(\Pi^{-1}(E)) > 0$, and an upright simplex found in $\Pi^{-1}(E)$ projects under Π to a nontrivial k-term progression in E. Similarly, by projecting from higher-dimensions to lower one can prove that Theorem 4 actually implies the following:

Corollary 5. *If $F \subset \mathbb{Z}^d$ is finite and $E \subseteq \mathbb{Z}^d$ has $\bar{\mathrm{d}}(E) > 0$, then there are some $\mathbf{a} \in \mathbb{Z}^d$ and $n \geqslant 1$ such that $\{\mathbf{a} + n\mathbf{b} : \mathbf{b} \in F\} \subseteq E$.* □

Furstenberg and Katznelson's proof of Theorem 4 rested on a formulation of multiple recurrence for actions of $\mathbb{Z}^d$ rather than $\mathbb{Z}$. Very surprisingly, Theorem 4 then went without a purely combinatorial proof for another twenty years, until a new approach using hypergraph theory was developed roughly in parallel by Gowers [Gow06] and Nagle, Rödl and Schacht [NRS06]. That hypergraph approach gave the first effective, finitary proof of this theorem: unlike the ergodic-theoretic approach, it also gives some bound in terms of $\bar{\mathrm{d}}(E)$ on how large a box in $\mathbb{Z}^d$ one must look in before being sure of finding a simplex contained in E. (In principle, one can extract such a bound from the Furstenberg-Katznelson proof, but it would be unimaginably poor.)

The success of Furstenberg and Katznelson's approach gave rise to a new subfield of ergodic theory often called Ergodic Ramsey Theory. It now contains several other results asserting that positive-density subsets of some kind of combinatorial structure must contain some further pattern of a special kind. Many of these have only be re-proven by purely combinatorial means (and with effective bounds) very recently. We will not state these in detail here, but only mention by name the IP Szemerédi Theorem of [FK85], the Density Hales-Jewett Theorem of [FK91] (finally given a purely combinatorial proof by the members of Gowers' 'Polymath 1' project in early 2009 [Pol09]), the Polynomial Szemerédi Theorem of Bergelson and Leibman [BL96] and the Nilpotent Szemerédi Theorem of Leibman [Lei98].

2 The phenomenon of multiple recurrence

In order to introduce Multiple Recurrence, it is helpful first to recall the probability-preserving version of Poincaré's classical Recurrence Theorem.

Theorem 6 (Poincaré Recurrence)**.** *If* (X, Σ, μ, T) *is a p.-p.t. and* $A \in \Sigma$ *has* $\mu(A) > 0$*, then there is some* $n \geqslant 1$ *such that* $\mu(A \cap T^{-n}A) > 0$.

Proof. The images $T^{-n}A$ are all subsets of the probability space X of equal positive measure, so some two of them must overlap. Once we have $\mu(T^{-n}A \cap T^{-m}A) > 0$ for $n \neq m$, the invariance of μ under T^n implies that also $\mu(A \cap T^{n-m}A) > 0$. □

Furstenberg's main result from [Fur77] strengthens this conclusion, by showing that in fact one may find several of the sets $T^{-n}A$, $n \in \mathbb{Z}$, that simultaneously overlap in a positive-measure set, where the relevant times n form an arithmetic progression.

Theorem 7 (Multiple Recurrence Theorem)**.** *If* (X, Σ, μ, T) *is a p.-p.t. and* $A \in \Sigma$ *has* $\mu(A) > 0$*, then for any* $k \geqslant 1$ *there is some* $n \geqslant 1$ *such that*

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

The Multidimensional Multiple Recurrence Theorem from [FK78] provides an analog of this for several commuting transformations.

Theorem 8 (Multidimensional Multiple Recurrence Theorem)**.** *If* $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ *is a p.-p.s. and* $A \in \Sigma$ *has* $\mu(A) > 0$ *then there is some* $n \geqslant 1$ *such that*

$$\mu(T^{-n\mathbf{e}_1}A \cap \cdots \cap T^{-n\mathbf{e}_d}A) > 0.$$

Note that for $d = 2$, simply applying the Poincaré Recurrence Theorem for the transformation $T^{\mathbf{e}_1 - \mathbf{e}_2}$ gives the desired conclusion.

Our goals in this course will be to prove Theorem 7 in the first case beyond Poincaré Recurrence, when $d = 3$. Two different ergodic-theoretic proofs of Theorem 8 can be found in [Fur81] and [Aus10b]. These are too long to be included in this course, but we will formulate and prove a related convergence result for some 'nonconventional' ergodic averages, which gives an introduction to some of the ideas.

First, let us prove the equivalence of Theorem 8 and Theorem 4. This equivalence (generalizing the case of subsets of $\mathbb{Z}$ and single transformations originally treated in [Fur77]) is often called the 'Furstenberg correspondence principle'. Although easy to prove, it has turned out to be a hugely fruitful insight into the relation between two different parts of mathematics. The version we give here essentially follows [Ber87].

Proposition 9 (Furstenberg correspondence principle). *If $E \subseteq \mathbb{Z}^d$, then there are a p.-p.s. $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ and a set $A \in \Sigma$ such that $\mu(A) = \bar{\mathrm{d}}(E)$, and for any $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k \in \mathbb{Z}^d$ one has*

$$\bar{\mathrm{d}}((E - \mathbf{v}_1) \cap (E - \mathbf{v}_2) \cap \cdots \cap (E - \mathbf{v}_k)) \geqslant \mu(T^{-\mathbf{v}_1} A \cap \cdots \cap T^{-\mathbf{v}_k} A).$$

In order to visualize this, observe that

$$(E - \mathbf{v}_1) \cap (E - \mathbf{v}_2) \cap \cdots \cap (E - \mathbf{v}_k)$$

is the set of those $\mathbf{n} \in \mathbb{Z}^d$ such that $\mathbf{n} + \mathbf{v}_i \in E$ for each $i \leqslant k$. Its density may be seen as the 'density of the set of translates of the pattern $\{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k\}$ that lie entirely inside E'. In these terms the above propositions shows that one can synthesize a p.-p.s. which provides a lower bound on this density for any given pattern in terms of the intersection of the corresponding shifts of the subset A.

Proof. First fix a sequence of boxes $R_j := \prod_{i \leqslant d}[M_{j,i}, N_{j,i}]$ such that $\min_{i \leqslant d}(N_{j,i} - M_{j,i}) \longrightarrow \infty$ and

$$\frac{|E \cap R_j|}{|R_j|} \longrightarrow \bar{\mathrm{d}}(E) \quad \text{as } j \longrightarrow \infty.$$

We can regard the set E as a point in the space $X := \mathcal{P}(\mathbb{Z}^d)$ of subsets of $\mathbb{Z}^d$, on which $\mathbb{Z}^d$ naturally acts by translation: $T^{-\mathbf{n}} B := B - \mathbf{n}$. This X can also be identified with the Cartesian product $\{0,1\}^{\mathbb{Z}^d}$ (by simply associating to each subset its indicator function), and this carries a compact metric product topology which makes (X, T) a topological system.

Now let

$$\nu_j := \frac{1}{|R_j|} \sum_{\mathbf{n} \in R_j} \delta_{T^{\mathbf{n}}(E)} \quad \text{for each } j,$$

the uniform measure on the patch of the T-orbit of E indexed by the large rectangle R_j. Because the side-lengths of the rectangles all tend to ∞, these measures are approximately invariant: that is, $\|T^{-\mathbf{n}}_{\#} \nu_j - \nu_j\|_{\mathrm{TV}} \longrightarrow 0$ as $j \longrightarrow \infty$ for any fixed $\mathbf{n} \in \mathbb{Z}^d$. Therefore, if we let $\mu \in \Pr(X)$ be any vague limit of them, then μ is strictly T-invariant. Suppose now that we have passed to a subsequence and may therefore write $\nu_j \longrightarrow \mu$ vaguely.

Finally, letting $A := \{H \in X : H \ni \mathbf{0}\}$ (this corresponds to the cylinder set $\{(\omega_{\mathbf{n}})_{\mathbf{n}} : \omega_{\mathbf{0}} = 1\}$ in $\{0,1\}^{\mathbb{Z}^d}$), we will show that this has the desired properties. By our initial choice of rectangles, we have

$$\mu(A) = \lim_{j \longrightarrow \infty} \frac{1}{|R_j|} \sum_{\mathbf{n} \in R_j} 1_{T^{\mathbf{n}}(E)}(\mathbf{0}) = \lim_{j \longrightarrow \infty} \frac{|E \cap R_j|}{|R_j|} = \bar{\mathrm{d}}(A),$$

where the first convergence holds because 1_A is a continuous function for the Cartesian product topology on X, and so vague convergence applies to it.

On the other hand, for any $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k \in \mathbb{Z}^d$, the indicator function $1_{T^{-\mathbf{v}_1}A\cap\cdots\cap T^{-\mathbf{v}_k}A}$ is also continuous on X, and so

$$\begin{aligned}\mu(T^{-\mathbf{v}_1}A \cap \cdots \cap T^{-\mathbf{v}_k}A) &= \lim_{j\longrightarrow\infty} \frac{1}{|R_j|} \sum_{\mathbf{n}\in R_j} 1_{T^{\mathbf{n}+\mathbf{v}_1}E\cap\cdots\cap T^{\mathbf{n}+\mathbf{v}_k}E}(\mathbf{0}) \\ &= \lim_{j\longrightarrow\infty} \frac{1}{|R_j|} \sum_{\mathbf{n}\in R_j} 1_{T^{\mathbf{v}_1}E\cap\cdots\cap T^{\mathbf{v}_k}E}(\mathbf{n}) \\ &\leqslant \bar{\mathrm{d}}(T^{\mathbf{v}_1}E \cap \cdots \cap T^{\mathbf{v}_k}E),\end{aligned}$$

since the upper Banach density is defined by a limsup over *all* rectangles of increasing side-lengths. □

Corollary 10. *Theorems 4 and 8 are equivalent.*

Proof. ($\Longrightarrow$) If (X, Σ, μ, T) is a p.-p. $\mathbb{Z}^d$-system and $A \in \Sigma$ with $\mu(A) > 0$, then by the Pointwise Ergodic Theorem the limit

$$f(x) := \lim_{N\longrightarrow\infty} \frac{1}{N^d} \sum_{\mathbf{n}\in[N]^d} 1_A(T^{\mathbf{n}}x)$$

exists a.s. and satisfies $\int_X f\,\mathrm{d}\mu = \mu(A) > 0$. Hence the set $\{x \in X : f(x) > 0\}$ has positive measure.

However, for each $x \in X$ we may apply Theorem 4 to the set of 'return times'

$$E_x := \{\mathbf{n} \in \mathbb{Z}^d : T^{\mathbf{n}}x \in A\}$$

(which has upper density at least $f(x)$) to find some $\mathbf{a}_x \in \mathbb{Z}^d$ and $n_x \geqslant 1$ such that

$$\{\mathbf{a}_x + n_x\mathbf{e}_1, \ldots, \mathbf{a}_x + n_x\mathbf{e}_d\} \subseteq E_x.$$

Moreover, it is easy to see that this choice of $\mathbf{a}_x$ and n_x can be made measurably in x. Since there are only countably many possible such choices, at least one of them must come up with positive probability inside the set $\{f > 0\}$, so there are $\mathbf{a} \in \mathbb{Z}^d$ and $n \geqslant 1$ such that

$$\mu\{x \in X : x \in T^{-\mathbf{a}-n\mathbf{e}_i}A\ \forall i \leqslant d\} > 0.$$

Applying the probability-preserving transformation $T^{\mathbf{a}}$ to this set, we see that it becomes $T^{-n\mathbf{e}_1}A \cap \cdots \cap T^{-n\mathbf{e}_d}A$, so this has positive measure, as required.

($\Longleftarrow$) This follows at once from Proposition 9: given E with $\bar{\mathrm{d}}(E) > 0$, that proposition produces a system and set with $\mu(A) > 0$ and such that the positivity of $\mu(T^{-n\mathbf{e}_1}A \cap \cdots \cap T^{-n\mathbf{e}_d}A)$ for some $n \geqslant 1$ implies that

$$(E - n\mathbf{e}_1) \cap (E - n\mathbf{e}_2) \cap \cdots \cap (E - n\mathbf{e}_d)$$

has positive upper density, and so is certainly nonempty. □

2.1 Nonconventional averages and the Furstenberg self-joining

A key aspect of Furstenberg and Katznelson's approach to Theorems 7 and 8 is that they give more than just the existence of one suitable time $n \geqslant 1$. In the multidimensional case, what they actually proved is the following.

Theorem 11 (Multidimensional Multiple Recurrence Theorem)**.** *If $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ is a p.-p.s. and $A \in \Sigma$ has $\mu(A) > 0$ then*

$$\liminf_{N \longrightarrow \infty} \frac{1}{N} \sum_{n=1}^{N} \mu(T^{-n\mathbf{e}_1}A \cap \cdots \cap T^{-n\mathbf{e}_d}A) > 0.$$

An analogous assertion for a single transformation lies behind Theorem 7. As in many other applications of ergodic theory, one finds that the averages in n behave more tamely than the term that appears for any particular value of n, and so are more amenable to analysis.

In this course, we will prove only one simple special case of Theorem 11 (referring to [Aus10a] or [Aus10b] for the general case). However, we will also answer a related and slightly easier question: whether the averages appearing in Theorem 11 actually converge, so that '$\liminf$' may be replaced with '$\lim$'.

Furstenberg and Katznelson's original proofs show that these non-negative sequences stay bounded away from zero, but do not show that their limits actually exist. Naturally, ergodic theorists quickly went in pursuit of this question in its own right. It was fully resolved only very recently, by Host and Kra [HK05] for the averages associated to a single transformation and then by Tao [Tao08] for those appearing in Theorem 11. Both are challenging proofs, and they are quite different from one another.

Host and Kra's argument builds on a long sequence of earlier work by several ergodic theorists, including [CL84, CL88a, CL88b, Rud95, Zha96, FW96, HK01] and several further references given there. An alternative proof of the main results of Host and Kra has been given by Ziegler in [Zie07], also resting on much of this earlier work. In addition to proving convergence, these efforts have led to a very detailed description of the limiting behaviour of these averages in terms of certain highly-structured factors of an arbitrary p.-p.s., in some sense generalizing our identification of the Kronecker factor as the ‘obstruction’ to weak mixing in the previous notes.

On the other hand, Tao’s proof of convergence for the multidimensional averages in Theorem 11 departs significantly from these earlier works, proceeding by first formulating a finitary, quantitative analog of the desired convergence assertion, and then making contact with the hypergraph-regularity theory developed for the new combinatorial proofs of Szemerédi’s Theorem in [NRS06, Gow06] (together with several other original insights in that finitary world).

In these notes, we will recount a more recent ergodic-theoretic proof of convergence, which avoids the transfer to a finitary combinatorial setting that underlies Tao’s proof, but also avoids the need for the much more detailed description that is implicit in the earlier approaches (such a description is still largely lacking for the multidimensional averages).

In fact, instead of the scalar averages appearing in the multiple recurrence theorems, it turns out to be more natural to study the functional averages

$$S_N(f_1,f_2,\dots,f_d) := \frac{1}{N}\sum_{n=1}^N (f_1\circ T^{n\mathbf{e}_1})\cdot(f_2\circ T^{n\mathbf{e}_2})\cdot\dots\cdot(f_d\circ T^{n\mathbf{e}_d})$$

for $f_1,f_2,\dots,f_d\in L^\infty(\mu)$. These are related by

$$\frac{1}{N}\sum_{n=1}^N \mu(T^{-n\mathbf{e}_1}A\cap\dots\cap T^{-n\mathbf{e}_d}A) = \int_X S_N(1_A,1_A,\dots,1_A)\,\mathrm{d}\mu,$$

so suitable convergence of the functional averages implies that of the scalar averages.

Theorem 12 (Convergence of nonconventional ergodic averages)**.** *If $T:\mathbb{Z}^d \curvearrowright (X,\Sigma,\mu)$ is a p.-p.s. and $f_1,f_2,\dots,f_d\in L^\infty(\mu)$, then the averages $S_N(f_1,f_2,\dots,f_d)$ converge in $\|\cdot\|_2$ as $N\longrightarrow\infty$.*

Remarks. 1. For $d=1$ this is just the usual Norm Ergodic Theorem (restricted to the case of bounded measurable functions), but for $d\geqslant 2$ the fact that for each n we multiply shifts of the functions f_i under different transformations means that it is not equivalent. An alternative way to view Theorem 12 is that we consider the functions

$$\frac{1}{N}\sum_{n=1}^N (f_1\otimes f_2\otimes\dots\otimes f_d)\circ(T^{\mathbf{e}_1}\times T^{\mathbf{e}_2}\times\dots\times T^{\mathbf{e}_d})^n$$

on the product space X^d, but instead of proving their convergence in $L^2(\mu^{\otimes d})$ (which would again follow from the usual NET) we prove their convergence in the L^2-space of the measure μ^Δ which is a copy of μ lifted to the diagonal set

$$\{(x, x, \dots, x) : x \in X\}.$$

Of course, this diagonal generally has measure zero under $\mu^{\otimes d}$, so this latter convergence does not follow from the NET or PET.

2. Using routine estimates from functional analysis one can strengthen Theorem 12 slightly by allowing functions $f_i \in L^{p_i}(\mu)$ such that $\frac{1}{p_1} + \cdots + \frac{1}{p_d} \leqslant \frac{1}{2}$, but we will not pursue this matter here.

3. Similarly to the basic ergodic theorems, it is also interesting to ask whether the above averages converge pointwise a.s. as $N \longrightarrow \infty$. This question is still open in almost all cases. Aside from $d = 1$ (corresponding to the Pointwise Ergodic Theorem), pointwise convergence is known only for the averages

$$\frac{1}{N}\sum_{n=1}^{N}(f_1 \circ T^n)(f_2 \circ T^{2n}) :$$

this is a very tricky result of Bourgain [Bou90]. $\lhd$

Aside from its natural interest, Theorem 12 is important because it enables a reformulation of Theorem 11 in terms of a convenient auxiliary construct. Clearly the fact of multiple recurrence is insensitive to our choosing a compact model. Having done so, we may consider the measures

$$\mu_N := \frac{1}{N}\sum_{n=1}^{N}\int_X \delta_{(T^{n\mathbf{e}_1}x, T^{n\mathbf{e}_2}x, \dots, T^{n\mathbf{e}_d}x)}\,\mu(\mathrm{d}x) \quad \text{on } X^d,$$

so that our N^{th} scalar average may also be written as $\mu_N(A \times A \times \cdots \times A)$, and more generally we have

$$\int_X S_N(f_1, f_2, \dots, f_d)\,\mathrm{d}\mu = \int_{X^d} f_1 \otimes f_2 \otimes \cdots \otimes f_d\,\mathrm{d}\mu_N.$$

The T-invariance of μ implies that each μ_N has all one-dimensional marginals equal to μ and is invariant under the diagonal transformations $T^{\mathbf{n}} \times T^{\mathbf{n}} \times \cdots \times T^{\mathbf{n}}$ for $\mathbf{n} \in \mathbb{Z}^d$. Therefore each μ_N is a joining of d copies of the system (X, Σ, μ, T) (a 'd-fold self-joining' of this system). Having chosen a compact model, if we pass to a vaguely convergent subsequence μ_{N_i}, then its limit μ^{F} also has these one-dimensional marginals and these invariances, and so is also a self-joining.

Next, for any product set $A_1 \times A_2 \times \cdots \times A_d$ the convergence of the averages $S_N(1_{A_1}, \ldots, 1_{A_d})$, implies that

$$\mu^{\mathrm{F}}(A_1 \times \cdots \times A_d) = \lim_{N \longrightarrow \infty} \int_X S_N(1_{A_1}, \ldots, 1_{A_d})\,\mathrm{d}\mu,$$

so in fact the limit measure μ^{F} is unique, and so the whole sequence μ_N must converge to it. Moreover, in addition to its invariance under each diagonal transformation $T^{\mathbf{n}} \times T^{\mathbf{n}} \times \cdots \times T^{\mathbf{n}}$, this limit μ^{F} is obtained as a limit of increasingly long averages under the 'off-diagonal transformation'

$$\vec{T} := T^{\mathbf{e}_1} \times T^{\mathbf{e}_2} \times \cdots \times T^{\mathbf{e}_d},$$

and so is invariant under this as well.

Definition 13. *The joining probability measure* μ^{F} *is the* ***Furstenberg self-joining*** *of the* $\mathbb{Z}^d$*-system* (X, Σ, μ, T).

We can now reformulate Theorem 11 as follows.

Theorem 14. *If* (X, Σ, μ, T) *is a p.-p.* $\mathbb{Z}^d$*-system and* $A \in \Sigma$ *then*

$$\mu^{\mathrm{F}}(A \times A \times \cdots \times A) = 0 \quad \textit{only if} \quad \mu(A) = 0.$$

Much of the more recent work in this area has gone into describing the Furstenberg self-joining in as much detail as possible. The Furstenberg self-joining will appear also in the proof of Theorem 12, in which we use its existence for $d-1$ transformations during the inductive proof of convergence for d transformations.

2.2 Some easier special cases

The special case of Theorem 7 with $k = 3$ is rather easier to handle than the general case, and will allow us to introduce some tools that we will need again later. The first is a basic estimate which underlies all of our subsequent work. It is a Hilbert-space version of a classical lemma of van der Corput concerning equidistribution of sequences (see, for instance, Kuipers and Niederreiter [KN74]).

Lemma 15 (Van der Corput estimate)**.** *Suppose that* $\mathfrak{H}$ *is a Hilbert space and* $(u_n)_{n \geqslant 1}$ *is a bounded sequence in* $\mathfrak{H}$*. Then*

$$\Big\| \frac{1}{N} \sum_{n=1}^{N} u_n \Big\| \longrightarrow 0$$

only if

$$\frac{1}{H}\sum_{h=1}^{H}\frac{1}{N}\sum_{n=1}^{N}\langle u_n, u_{n+h}\rangle \not\longrightarrow 0$$

as $N \longrightarrow \infty$ and then $H \longrightarrow \infty$ (that is, there are some $H_1 < H_2 < \dots$ and for each i some $N_{i,1} < N_{i,2} < \dots$ such that the above scalar averages at H_i and $N_{i,j}$ are uniformly bounded away from zero for all i and j).

This lemma is closely related to the proof of the 'inverse' result underlying our approach to the Norm Ergodic Theorem (Notes 2), according to which a function f whose ergodic averages do not vanish must have a nonzero inner product with an invariant function. This relationship goes quite deep: we will later use the van der Corput estimate to deduce other 'inverse theorems' for the nonconventional averages of Theorem 12.

Proof. Assume we have some $\varepsilon > 0$ and $N_1 < N_2 < \dots$ such that

$$\Big\|\frac{1}{N_i}\sum_{n=1}^{N_i} u_n\Big\| \geqslant \varepsilon \quad \forall i.$$

For any $H \geqslant 1$ we have

$$\Big\|\frac{1}{N}\sum_{n=1}^{N} u_n - \frac{1}{H}\sum_{h=1}^{H}\frac{1}{N}\sum_{n=1}^{N} u_{n+h}\Big\| = \mathrm{O}\Big(\frac{H}{N}\Big) \longrightarrow 0 \quad \text{as } N \longrightarrow \infty,$$

and so combining these facts gives

$$\liminf_{i\longrightarrow\infty}\Big\|\frac{1}{H}\sum_{h=1}^{H}\frac{1}{N_i}\sum_{n=1}^{N_i} u_{n+h}\Big\| \geqslant \varepsilon \quad \forall H \geqslant 1.$$

In this double average we can now exchange the order of averaging and then apply the triangle inequality to deduce that

$$\liminf_{i\longrightarrow\infty}\frac{1}{N_i}\sum_{n=1}^{N_i}\Big\|\frac{1}{H}\sum_{h=1}^{H} u_{n+h}\Big\| \geqslant \varepsilon \quad \forall H \geqslant 1.$$

Applying the Cauchy-Schwartz inequality to the average in n here gives

$$\liminf_{i\longrightarrow\infty}\frac{1}{N_i}\sum_{n=1}^{N_i}\Big\|\frac{1}{H}\sum_{h=1}^{H} u_{n+h}\Big\|^2$$

$$= \liminf_{i\longrightarrow\infty}\frac{1}{H^2}\sum_{h_1,h_2=1}^{H}\frac{1}{N_i}\sum_{n=1}^{N_i}\langle u_{n+h_1}, u_{n+h_2}\rangle \geqslant \varepsilon^2 \quad \forall H \geqslant 1.$$

Finally, by breaking the above averages into the two cases $\{h_1 \geqslant h_2\}$ and $\{h_1 < h_2\}$ and the re-parameterizing, say, $n + h_1$ as n and $h_2 - h_1$ as h, we see that the expression

$$\frac{1}{H^2}\sum_{h_1,h_2=1}^{H}\frac{1}{N_i}\sum_{n=1}^{N_i}\langle u_{n+h_1}, u_{n+h_2}\rangle$$

can be written as an average of the expressions

$$\frac{1}{H'}\sum_{h=1}^{H'}\frac{1}{N_i}\sum_{n=1}^{N_i}\langle u_n, u_{n+h}\rangle$$

over a range of large values of H' (say, $H^{1/2} \leqslant H' \leqslant H$), plus some error that vanishes as $N \longrightarrow \infty$ (correspondingly, at the rate $\mathrm{O}(H^{1/2}/N)$. Hence some of these latter averages in h and n must also stay large, as required. □

This simple estimate is surprisingly powerful. As a first application, one can deduce that Theorems 12 and 8 actually admit quite quick proofs for systems all of whose individual transformations are weakly mixing.

Theorem 16. *If $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ is such that every individual $T^{\mathbf{n}}$, $\mathbf{n} \in \mathbb{Z}^d$, is weakly mixing (that is, T is* ***totally weakly mixing****) then*

$$S_N(f_1, f_2, \ldots, f_d) \longrightarrow \Big(\int f_1\,\mathrm{d}\mu\Big)\Big(\int f_2\,\mathrm{d}\mu\Big)\cdots\Big(\int f_d\,\mathrm{d}\mu\Big)$$

in $\|\cdot\|_2$ as $N \longrightarrow \infty$. In particular, if $\mu(A) > 0$, then setting $f_1 = \ldots = f_k = 1_A$ gives

$$\frac{1}{N}\sum_{n=1}^{N}\mu(T^{-n\mathbf{e}_1}A \cap \cdots T^{-n\mathbf{e}_d}A) \longrightarrow \mu(A)^k > 0.$$

□

A proof can be obtained from Problem Sheet 4.

Matters are not so simple in the absence of weak mixing.

Example. Suppose that $R_\alpha \curvearrowright (\mathbb{T}, \mathcal{B}(\mathbb{T}), m_{\mathbb{T}})$ is an irrational, and hence ergodic, circle rotation, and let $f_1 : t \mapsto \mathrm{e}^{4\pi \mathrm{i}t} \in \mathbb{C}$ and $f_2 : t \mapsto \mathrm{e}^{-2\pi\mathrm{i}t}$. Then for any n we have

$$f_1(t + n\alpha)f_2(t + 2n\alpha) = \mathrm{e}^{4\pi\mathrm{i}t + 4n\pi\mathrm{i}\alpha}\mathrm{e}^{-2\pi\mathrm{i}t - 4n\pi\mathrm{i}\alpha} = \mathrm{e}^{2\pi\mathrm{i}t},$$

and so, even though T is ergodic, the averages

$$\frac{1}{N}\sum_{n=1}^{N}(f_1 \circ T^n)(f_2 \circ T^{2n})$$

are all equal to the non-constant function $\mathrm{e}^{2\pi \mathrm{i} t}$. Moreover, by letting A be a non-trivial level set of f_1 and using a little care, one can produce a set A such that $\frac{1}{N}\sum_{n=1}^{N} m_{\mathbb{T}}(A \cap T^{-n}A \cap T^{-2n}A)$ tends to some limit other than $m_{\mathbb{T}}(A)^3$. $\lhd$

By analogy with our previous structure theorem for weak mixing (Notes 5), our next hope might be that the above examples are the only possible 'bad guys'. In fact this is true for the case of T and T^2 (corresponding to the case $k = 3$ of the one-dimensional Szemerédi Theorem). For the more complicated cases something similar, albeit much more involved, takes place: that will not be covered in this course. In order to make this idea formal we need one more definition.

Definition 17 (Partially characteristic factor)**.** *If $T : \mathbb{Z}^d \curvearrowright (X, \Sigma, \mu)$ is a p.-p.s., then a factor $\Lambda \leqslant \Sigma$ is* ***partially characteristic in position*** *i for the averages S_N if*

$$\begin{aligned}\|S_N(f_1, f_2, \ldots, f_i, \ldots, f_d) - S_N(f_1, f_2, \ldots, \mathsf{E}(f_i \,|\, \Lambda), \ldots, f_d)\|_2 \\ = \|S_N(f_1, f_2, \ldots, f_i - \mathsf{E}(f_i \,|\, \Lambda), \ldots, f_d)\|_2 \longrightarrow 0\end{aligned}$$

as $N \longrightarrow \infty$ for any $f_1, f_2, \ldots, f_d \in L^\infty(\mu)$.

We will usually abbreviate 'partial characteristic in position 1' to just 'partially characteristic'.

The intuition here is that under the transformation $T^{\mathbf{e}_1}$, the fibres of the factor $\Lambda \leqslant \Sigma$ becomed 'twisted up' in such a complicated way that the movement of the functions f_2, …, f_d under the other transformations $T^{\mathbf{e}_i}$ for $i \geqslant 2$ keeps no track of it, and so asymptotically the behaviour of $f_1 \circ T^{\mathbf{e}_1}$ on these fibres becomes independent of the behaviours of these other functions. If one were lucky enough to find a highly-structured partially characteristic factor Λ, then in trying to prove convergence one could always replace f_1 by $\mathsf{E}(f_1 \,|\, \Lambda)$, and so simply reduce to the case when f_1 is Λ-measurable, which might be easier (or perhaps provide the seed for an argument by induction). Remarkably, this very optimistic idea turns out to be succesful in giving proofs of both convergence and multiple recurrence, although we are still some way from understanding which factors Λ will appear in general and how their structure will be useful for us.

However, in the special case of the averages

$$S_N(f_1, f_2) = \frac{1}{N}\sum_{n=1}^{N}(f_1 \circ T_1^n)(f_2 \circ T^{2n})$$

for a single ergodic p.-p.t. (X, Σ, μ, T) it turns out that the Kronecker factor is partially characteristic in both positions.

Lemma 18. *If* (X, Σ, μ, T) *is ergodic and* Λ *is its Kronecker factor, and if*

$$\|S_N(f_1, f_2)\|_2 \not\longrightarrow 0,$$

then $\mathsf{E}(f_i \,|\, \Lambda) \neq 0$ *for both* $i = 1, 2$ *(and so* Λ *is partially characteristic in both positions).*

Proof. We give the proof for f_1, the case of f_2 being similar. Applying Lemma 15 with

$$u_n := (f_1 \circ T^n)(f_2 \circ T^{2n}) \quad \text{for } n \geqslant 1,$$

we obtain that

$$\begin{aligned}
&\frac{1}{H}\sum_{h=1}^{H}\frac{1}{N}\sum_{n=1}^{N}\int_X (\overline{f_1} \circ T^n)(\overline{f}_2 \circ T^{2n})(f_1 \circ T^{n+h})(f_2 \circ T^{2n+2h})\,\mathrm{d}\mu \\
&\qquad = \frac{1}{H}\sum_{h=1}^{H}\frac{1}{N}\sum_{n=1}^{N}\int_X \overline{f_1}(\overline{f}_2 \circ T^n)(f_1 \circ T^h)(f_2 \circ T^{n+2h})\,\mathrm{d}\mu \\
&\qquad = \frac{1}{H}\sum_{h=1}^{H}\int_X \overline{f_1}(f_1 \circ T^h)\Big(\frac{1}{N}\sum_{n=1}^{N}(\overline{f_2}\cdot(f_2 \circ T^{2h})) \circ T^n\Big)\,\mathrm{d}\mu \not\longrightarrow 0
\end{aligned}$$

as $N \longrightarrow \infty$ and then $H \longrightarrow \infty$, where the first equality follows the T^n-invariance of μ, and the second by a re-arrangement.

However, the N-limit appearing inside the last integral above can now be evaluated by the Norm Ergodic Theorem. Using that T is ergodic, we obtain that

$$\frac{1}{H}\sum_{h=1}^{H}\Big(\int_X \overline{f_1}(f_1 \circ T^h)\,\mathrm{d}\mu\Big)\Big(\int_X \overline{f_2}(f_2 \circ T^{2h})\,\mathrm{d}\mu\Big) \not\longrightarrow 0$$

as $H \longrightarrow \infty$. By the Cauchy-Schwartz inequality, the non-vanishing of these implies that also

$$\begin{aligned}
&\frac{1}{H}\sum_{h=1}^{H}\Big|\int_X \overline{f_1}(f_1 \circ T^h)\,\mathrm{d}\mu\Big|^2 \\
&= \int_{X^2}(\overline{f_1} \otimes f_1)\cdot\frac{1}{H}\sum_{h=1}^{H}((f_1 \otimes \overline{f_1}) \circ (T \times T)^h)\,\mathrm{d}\mu^{\otimes 2} \\
&\longrightarrow \int_{X^2}(\overline{f_1} \otimes f_1)\cdot \mathsf{E}(f_1 \otimes \overline{f_1} \,|\, (\Sigma \otimes \Sigma)^{T \times T})\,\mathrm{d}\mu^{\otimes 2} \\
&\neq 0,
\end{aligned}$$

where we have now applied the Norm Ergodic Theorem for the system $(X^2, \Sigma^{\otimes 2}, \mu^{\otimes 2}, T^{\times 2})$ to evaluate the H-limit.

Finally, since $\mathsf{E}(f_1 \otimes \overline{f_1} \,|\, (\Sigma \otimes \Sigma)^{T\times T})$ is $(T \times T)$-invariant, Theorem 8 of Notes 5 gives that it is $(\Lambda \otimes \Lambda)$-measurable. Therefore the above non-vanishing implies also that

$$\mathsf{E}(f_1 \otimes \overline{f_1} \,|\, \Lambda \otimes \Lambda) = \mathsf{E}(f_1 \,|\, \Lambda) \otimes \overline{\mathsf{E}(f_1 \,|\, \Lambda)} \neq 0$$

and hence $\mathsf{E}(f_1 \,|\, \Lambda) \neq 0$, as required. □

Lemma 18 will imply our special case of Theorem 7 using the following lemma.

Lemma 19. *Suppose that G is a compact metric group, $U \subseteq G$ is a neighbourhood of the identity, and $g \in G$. Then there is an $L \in \mathbb{N}$ such that the set of integers*

$$\{n \in \mathbb{Z} : \; g^n \in U\}$$

has non-empty intersection with every interval in $\mathbb{Z}$ of length at least L.

Proof. Let $H := \overline{\{g^n : \; n \in \mathbb{Z}\}}$, a closed subgroup of G. Then $U \cap H$ is a neighbourhood of the identity in H, so it suffices to work within the subgroup H. Equivalently, this means we may assume that $\{g^n : \; n \in \mathbb{Z}\}$ is dense in G.

Having made this assumption, it follows that the collection of open sets

$$\{g^n U^{-1} : \; n \in \mathbb{Z}\}$$

covers G, because for any $h \in G$ our assumption of density gives some n such that $g^n \in hU$, hence $h \in g^n U^{-1}$. Therefore, by compactness, there is some $L \geqslant 1$ such that the subfamily $\{g^n U^{-1} : \; 1 \leqslant n \leqslant L\}$ covers G. However, for any $m \in \mathbb{Z}$ it now follows that

$$\bigcup_{m+1 \leqslant n \leqslant m+L} g^n U^{-1} = g^m \Big(\bigcup_{1 \leqslant n \leqslant L} g^n U^{-1} \Big) = g^m G = G.$$

Therefore, for any discrete interval $\{m+1, \dots, m+L\}$, there is some n in that interval with $g^n U^{-1} \ni e$, and hence $g^n \in U$. □

Proof of Theorem 7 in case $k = 3$. As discussed above, we will show that in this case

$$\liminf_{N \longrightarrow \infty} \int_X 1_A \cdot S_N(1_A, 1_A)\, \mathrm{d}\mu > 0,$$

where $S_N(f_1, f_2) := \frac{1}{N} \sum_{n=1}^{N} (f_1 \circ T^n)(f_2 \circ T^{2n})$.

First, by the theorem on compact models, we may assume that (X, T) is a topological dynamical system, and also that the factor Σ^T of T-invariant sets is generated by a continuous factor map $\xi : (X, \Sigma, \mu, T) \longrightarrow (Y, \Phi, \nu, S)$ to another compact model. Since ξ generates the T-invariant sets, we find that S must fix every Borel set of Y, and this implies that in fact $S = \mathrm{id}_Y$. Having passed to these compact models, the measure μ has a disintegration, say

$$\mu = \int_Y \mu_y \, \nu(\mathrm{d}y),$$

as given by Theorem 11 of Notes 4. By the essential uniqueness of this disintegration, and since μ is T-invariant, one obtains that μ_y is also T-invariant for a.e. y. On the other hand, $\mu(A) = \int_Y \mu_y(A) \, \nu(\mathrm{d}y)$, so it suffices to prove Theorem 7 for each of the systems (X, Σ, μ_y, T) for which $\mu_y(A) > 0$.

This is important, because one can now show that μ_y is be ergodic for T for a.e. y, so we have reduced our task to the ergodic case. This assertion requires a little thought, but we relegate it to Problem Sheet 4.

We therefore now assume that μ is ergodic for T. Let Λ be the Kronecker factor, so it is generated by a factor map

$$\pi : (X, \Sigma, \mu, T) \longrightarrow (Z, d, \mathcal{B}(Z), \nu, R)$$

whose target is a compact system. Let $f := \mathsf{E}(1_A \mid \Lambda)$. By Lemma 18 we have

$$\|S_N(1_A, 1_A) - S_N(f, f)\|_2 \longrightarrow 0.$$

Given two real or complex sequences $(a_n)_n$ and $(b_n)_n$, let $a_n \sim b_n$ stand for the assertion that $a_n - b_n \longrightarrow 0$. Then the above norm convergence gives

$$\begin{aligned} \frac{1}{N} \sum_{n=1}^{N} \mu(A \cap T^{-n} A \cap T^{-2n} A) &= \int_X 1_A \cdot S_N(1_A, 1_A) \, \mathrm{d}\mu \\ &\sim \int_X 1_A \cdot S_N(f, f) \, \mathrm{d}\mu \\ &= \int_X f \cdot S_N(f, f) \, \mathrm{d}\mu \\ &= \frac{1}{N} \sum_{n=1}^{N} \int_X f (f \circ T^n)(f \circ T^{2n}) \, \mathrm{d}\mu, \end{aligned}$$

where the equality of the third line follows from the definition of conditional expectation, because $S_N(f, f)$ is Λ-measurable. It therefore suffices to show that these last averages stay uniformly positive for $N \longrightarrow \infty$.

To do this, since f is Λ-measurable, it agrees a.s. with a function lifted from Z. Let us now write f for that function on Z instead. By basic properties of conditional expectation, we have $f \geqslant 0$ and $\int_Z f \, \mathrm{d}\nu = \mu(A) > 0$. This implies that one also has $\alpha := \int_Z f^3 \, \mathrm{d}\nu > 0$. It remains to show that these properties imply

$$\frac{1}{N} \sum_{n=1}^{N} \int_Z f(f \circ R^n)(f \circ R^{2n}) \, \mathrm{d}\nu > 0.$$

To prove this, first recall that R is an element of the compact group $G := \mathrm{Isom}(Z, d)$, metrized by the uniform metric

$$D(S, S') := \sup_{z \in Z} d(Sz, S'z).$$

Since f may be approximated in $\| \cdot \|_1$ by continuous functions on Z, an easy exercise shows that there is some $\delta > 0$ such that

$$S \in G, \ D(S, \mathrm{id}_Z) < \delta \quad \Longrightarrow \quad \|f - f \circ S\|_1 < \alpha/4.$$

Whenever n is such that $D(\mathrm{id}_Z, R^n) < \delta/2$, it follows that

$$D(\mathrm{id}_Z, R^{2n}) \leqslant D(\mathrm{id}_Z, R^n) + D(R^n, R^{2n}) < \delta.$$

Therefore for such n one has

$$\int_Z f(f \circ R^n)(f \circ R^{2n}) \, \mathrm{d}\nu > \int_Z f^3 \, \mathrm{d}\nu - 2\alpha/4 = \alpha/2.$$

Finally, Lemma 19 gives some $L \in \mathbb{N}$ such that the set of n for which this holds intersects every interval of at least L consecutive integers. These estimates together now imply that

$$\frac{1}{N} \sum_{n=1}^{N} \int_Z f(f \circ R^n)(f \circ R^{2n}) \, \mathrm{d}\nu \geqslant \frac{N - 2L}{N} \cdot \frac{1}{L} \cdot (\alpha/2) \quad \forall N,$$

and hence the limit infimum is at least $\alpha/2L > 0$, as required. □

The partially characteristic factor given by Lemma 18 can also be used to prove convergence in this case. This can be found on Problem Sheet 4. The next notes will prove convergence in the general case (Tao's Theorem).

References

[Aus10a] Tim Austin. Deducing the multidimensional Szemerédi theorem from an infinitary removal lemma. *J. Anal. Math.*, 111:131–150, 2010.

[Aus10b] Timothy Derek Austin. *Multiple recurrence and the structure of probability-preserving systems*. ProQuest LLC, Ann Arbor, MI, 2010. Thesis (Ph.D.)–University of California, Los Angeles.

[Ber87] V. Bergelson. Weakly mixing PET. *Ergodic Theory Dynam. Systems*, 7(3):337–349, 1987.

[BL96] V. Bergelson and A. Leibman. Polynomial extensions of van der Waerden's and Szemerédi's theorems. *J. Amer. Math. Soc.*, 9(3):725–753, 1996.

[Bou90] J. Bourgain. Double recurrence and almost sure convergence. *J. Reine Angew. Math.*, 404:140–161, 1990.

[CL84] Jean-Pierre Conze and Emmanuel Lesigne. Théorèmes ergodiques pour des mesures diagonales. *Bull. Soc. Math. France*, 112(2):143–175, 1984.

[CL88a] Jean-Pierre Conze and Emmanuel Lesigne. Sur un théorème ergodique pour des mesures diagonales. In *Probabilités*, volume 1987 of *Publ. Inst. Rech. Math. Rennes*, pages 1–31. Univ. Rennes I, Rennes, 1988.

[CL88b] Jean-Pierre Conze and Emmanuel Lesigne. Sur un théorème ergodique pour des mesures diagonales. *C. R. Acad. Sci. Paris Sér. I Math.*, 306(12):491–493, 1988.

[ET36] P. Erdős and P. Turán. On some sequences of integers. *J. London Math. Soc.*, 11:261–264, 1936.

[FK78] Hillel Furstenberg and Yitzhak Katznelson. An ergodic Szemerédi Theorem for commuting transformations. *J. d'Analyse Math.*, 34:275–291, 1978.

[FK85] Hillel Furstenberg and Yitzhak? Katznelson. An ergodic Szemerédi theorem for IP-systems and combinatorial theory. *J. d'Analyse Math.*, 45:117–168, 1985.

[FK91] Hillel Furstenberg and Yitzhak Katznelson. A Density Version of the Hales-Jewett Theorem. *J. d'Analyse Math.*, 57:64–119, 1991.

[Fur77] Hillel Furstenberg. Ergodic behaviour of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. d'Analyse Math.*, 31:204–256, 1977.

[Fur81] Hillel Furstenberg. *Recurrence in Ergodic Theory and Combinatorial Number Theory*. Princeton University Press, Princeton, 1981.

[FW96] Hillel Furstenberg and Benjamin Weiss. A mean ergodic theorem for $\frac{1}{N}\sum_{n=1}^{N} f(T^n x)g(T^{n^2}x)$. In Vitaly Bergleson, ABC March, and Joseph Rosenblatt, editors, *Convergence in Ergodic Theory and Probability*, pages 193–227. De Gruyter, Berlin, 1996.

[Gow06] W. T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combin. Probab. Comput.*, 15(1-2):143–184, 2006.

[GRS90] Robert L. Graham, Bruce L. Rothschild, and Joel H. Spencer. *Ramsey Theory*. John Wiley & Sons, New York, 1990.

[HK01] Bernard Host and Bryna Kra. Convergence of Conze-Lesigne averages. *Ergodic Theory Dynam. Systems*, 21(2):493–509, 2001.

[HK05] Bernard Host and Bryna Kra. Nonconventional ergodic averages and nilmanifolds. *Ann. Math.*, 161(1):397–488, 2005.

[KN74] L. Kuipers and H. Niederreiter. *Uniform distribution of sequences*. Wiley-Interscience [John Wiley & Sons], New York, 1974. Pure and Applied Mathematics.

[Lei98] A. Leibman. Multiple recurrence theorem for measure preserving actions of a nilpotent group. *Geom. Funct. Anal.*, 8(5):853–931, 1998.

[NRS06] Brendan Nagle, Vojtěch Rödl, and Mathias Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006.

[Pol09] D. H. J. Polymath. A new proof of the density Hales-Jewett theorem. Preprint, available online at `arXiv.org`: 0910.3926, 2009.

[Rot53] K. F. Roth. On certain sets of integers. *J. London Math. Soc.*, 28:104–109, 1953.

[Rud95] Daniel J. Rudolph. Eigenfunctions of $T \times S$ and the Conze-Lesigne algebra. In *Ergodic theory and its connections with harmonic analysis (Alexandria, 1993)*, volume 205 of *London Math. Soc. Lecture Note Ser.*, pages 369–432. Cambridge Univ. Press, Cambridge, 1995.

[Sze75] Endre Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975.

[Tao08] Terence Tao. Norm convergence of multiple ergodic averages for commuting transformations. *Ergodic Theory and Dynamical Systems*, 28:657–688, 2008.

[vdW27] B. L. van der Waerden. Beweis einer baudetschen vermutung. *Nieuw. Arch. Wisk.*, 15:212–216, 1927.

[Zha96] Qing Zhang. On convergence of the averages $(1/N)\sum_{n=1}^{N} f_1(R^n x) f_2(S^n x) f_3(T^n x)$. *Monatsh. Math.*, 122(3):275–300, 1996.

[Zie07] Tamar Ziegler. Universal characteristic factors and Furstenberg averages. *J. Amer. Math. Soc.*, 20(1):53–97 (electronic), 2007.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: `tim@cims.nyu.edu`
URL: `cims.nyu.edu/~tim`

Topics in Ergodic Theory
Notes 9: Introduction to entropy

1 Systems and processes

As we have seen previously, many natural examples of probability-preserving Γ-systems arise in the form

$$(A^\Gamma, \mathcal{B}(A^\Gamma), \nu, S)$$

for some compact (often finite) space A and some S-invariant $\nu \in \Pr A^\Gamma$, where S is the shift-action of Γ on A^Γ. Examples of this kind are called **shift systems**. We will now study shift systems more closely when $\Gamma = \mathbb{Z}^d$ and A is a finite set, referred to as the **alphabet**. Many of the results below can be extended to arbitrary discrete amenable groups Γ, but we will not explore that generality here.

The largest source of examples of shift systems for $\mathbb{Z}^d$, at least when $d \geqslant 2$, is statistical physics. For example, in some of the simplest models, A is $\{\pm 1\}$, and a configuration in $A^{\mathbb{Z}^d}$ represents the orientation, either 'up' or 'down', of the magnetic spins at all the lattice points in the microscopic structure of some magnetic solid. A shift-invariant measure ν on $A^{\mathbb{Z}^d}$ arises as a descriptor of the long-run average behaviour of such a system when it is in thermal equilibrium with its environment: the value $\nu(E)$ for $E \subseteq A^{\mathbb{Z}^d}$ represents the proportion of time that the microscopic state of the system spends in E.

On the other hand, the proof of Proposition 3 in Notes 4 showed that any p.-p. system is isomorphic to such an example when A is the Cantor space. Under the restriction $|A| < \infty$, this universality no longer holds (some examples will be given later), but one still expects that shift systems can be extremely general.

Entropy arises as one of the key parameters describing a shift system. It has an interpretation in statistical physics and probability, but here we will focus on what it can tell us about the general ergodic theory of shift systems. Good introductions to those other aspects of entropy can be found, for instance, in [Ell06, Var03].

First, let us set up some nomenclature for comparing abstract systems with shift systems. Suppose that (X, Σ, μ, T) is a $\mathbb{Z}^d$-system and that

$$\Phi : (X, \Sigma, \mu, T) \longrightarrow (A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S)$$

is a factor map to a shift system. Then the output of φ may be written in terms of separate coordinates,

$$\Phi(x) = (\varphi_n)_{n\in\mathbb{Z}^d}$$

for some maps $\varphi_n : X \longrightarrow A$, and now the equivariance of Φ requires that

$$\varphi_n = \varphi_0 \circ T^n \quad \forall n \in \mathbb{Z}^d.$$

Therefore Φ is completely determined by φ_0, as well as vice-versa. The map φ_0 will often be referred to as an **observable**, and it is **finite-valued** if $|A| < \infty$.

Definition 1. *An **observable** of (X, Σ, μ, T) is a Borel map $\varphi_0 : X \longrightarrow A$ to a compact metric space. It is **finite-valued** if A is a finite set, and it is **generating** if the σ-algebra generated by the collection of maps $\{\varphi \circ T^n : n \in \mathbb{Z}^d\}$ is the whole of Σ, up to negligible sets.*

*A $\mathbb{Z}^d$**-process** is a tuple $(X, \Sigma, \mu, T, \varphi)$ in which (X, Σ, μ, T) is a p.-p. $\mathbb{Z}^d$-system and $\varphi : X \longrightarrow A$ is a finite-valued observable.*

Two of the main questions motivating this last part of the course are the following:

- Given an abstract p.-p. system (X, Σ, μ, T), what are the processes that it can give rise to?
- Given two shift systems, when are they isomorphic as abstract p.-p. systems?

The second of these questions is essentially another version of the 'classification problem' mentioned in Notes 1, one of the main foundational problems in ergodic theory. Naïvely, it asks for effective criteria or invariants that classify p.-p.s.s up to isomorphism. By now it has become clear that among completely general systems, even for actions of $\mathbb{Z}$, this is asking too much; there are even abstract, rigorous reasons in descriptive set theory why this is so (see, for instance, Kechris [Kec10] for a general introduction).

However, some partial results are possible. We have already met the different kinds of mixing as some examples of natural isomorphism-invariants. Entropy will provide the next important isomorphism-invariant.

We will introduce entropy in three stages:

- first, we introduce Shannon's original notion of entropy for a probability distribution on a finite set, and prove some of its basic properties;
- then, we use Shannon entropy to define the entropy rate of a shift system;

- finally, we define the Kolmogorov-Sinai entropy of a p.-p. system by considering the entropy rates of all the processes that it gives rise to.

As our first application of this theory we will discuss its consequences for the following natural restricted instance of the classification problem:

> Given two finite probability spaces $([k], \nu)$ and $([\ell], \theta)$, when are the Bernoulli shifts $([k]^{\mathbb{Z}}, \mathcal{B}([k]^{\mathbb{Z}}), \nu^{\otimes \mathbb{Z}}, S)$ and $([\ell]^{\mathbb{Z}}, \mathcal{B}([\ell]^{\mathbb{Z}}), \theta^{\otimes \mathbb{Z}}, S)$ isomorphic?

Entropy provides a numerical invariant which can sometimes serve to distinguish two such Bernoulli shifts. In fact, it turns out that *among Bernoulli shifts*, entropy is a complete invariant: if two such shifts do have the same entropy, then they are isomorphic. This latter result is an extremely deep theorem of Ornstein. Importantly, the entropy of a Bernoulli shift can be written as a simple formula in terms of the one-dimensional marginal ν.

Theorem 2 (Ornstein's Isomorphism Theorem)**.** *The Bernoulli shifts above are isomorphic if and only if the measures ν and θ have the same entropy.*

Other nice treatments of this theory include Ornstein's original monograph [Orn74] and the more modern textbooks of Rudolph [Rud90] and Kalikow and McCutcheon [KM10].

2 Shannon entropy

2.1 Motivation and definition

Shannon entropy is a quantity associated to a probability distribution on a finite (or countable) set. Heuristically, it gives some indication of the 'effective size' of the distribution. It has the following intuitive properties:

- on a small set, all probability distributions have relatively small entropy;
- on a large set, the uniform distribution has high entropy, but for a distribution for which most (not necessarily all) of the mass is concentrated on a much smaller subset, the entropy is generally much smaller.

In general, of course, there are many parameters that conform to this intuition somehow. Crucially, the notion of entropy is also supposed to behave well under forming product measures, and this gives a clue as to how it should be defined.

Suppose that $\nu = (\nu_1, \nu_2, \ldots, \nu_k)$ is a probability distribution on $[k]$. Consider some very large integer $N \geqslant 1$, and form the product measure $\nu^{\otimes N}$ on $[k]^N$. Also fix some small $\varepsilon \in (0, 1)$, and consider the following question:

What is the minimal cardinality of a subset $T \subseteq [k]^N$ for which $\nu^{\otimes N}(T) > 1 - \varepsilon$ (that is, 'what is the effective size of the support of $\nu^{\otimes N}$, up to an error set of small measure?')?

Of course, determining that answer precisely is generally extremely delicate. However, intuitively the answer should grow exponentially in N (up to subexponential corrections), at least if we judge by the case in which ν is the uniform measure, for which T should simply be a subset of $[k]^N$ that contains at least $(1-\varepsilon)k^N$ points. So let us seek to evaluate the coefficient in this exponential.

To do this, first observe that by the Law of Large Numbers, if N is large enough and a string $\omega = (\omega_n)_{n \leqslant N}$ is drawn at random from $\nu^{\otimes N}$, then for each $i \leqslant k$ the frequency of the value i in the string ω should be roughly ν_i. Letting $f_i(\omega)$ be this frequency, it follows that one candidate set which should carry most of $\nu^{\otimes N}$ is the following:

$$\bigcup_{\substack{(n_1, n_2, \dots, n_k) \in \mathbb{N}^k : \\ \sum_i n_i = N \ \& \ (\nu_i - \delta)N \leqslant n_i \leqslant (\nu_i + \delta)N \ \forall i}} \{\omega : \ f_i(\omega) = n_i \ \forall i \leqslant k\}$$

for any $\delta > 0$ (where for given error-tolerance $\varepsilon > 0$, any $\delta > 0$ will work but N will need to be large enough in terms of δ). In information theory, this is called the **δ-typical set for** ν. Let us denote it $T_{\nu,\delta,N}$.

Now, the number of integer vectors $(n_1, n_2, \dots, n_k)$ such that $\sum_i n_i = N$ and $(\nu_i - \delta)N < n_i < (\nu_i + \delta)N$ for all i is at most N^k. Given such an integer sequence, the size of $\{\omega : \ f_i(\omega) = n_i \ \forall i\}$ is simply given by the multinomial coefficient

$$\frac{N!}{n_1! n_2! \cdots n_k!}.$$

We can now obtain an approximate answer to the above question as a consequence of Stirling's Approximation,

$$N! = (1 + \mathrm{o}(1)) \cdot \sqrt{2\pi}\sqrt{N}\mathrm{e}^{-N}N^N \quad \text{as } N \longrightarrow \infty.$$

Applying this to each factorial in the multinomial coefficient above gives

$$\frac{1}{(2\pi)^{(k-1)/2}} \cdot \frac{\sqrt{N}}{\sqrt{n_1 n_2 \cdots n_k}} \cdot \frac{N^N}{n_1^{n_1} n_2^{n_2} \cdots n_k^{n_k}}.$$

(where the factor e^{-N} cancels exactly with the factors e^{-n_i} because $\sum_i n_i = N$). Now, bearing in mind that we seek only an exponential growth rate in N, we may ignore the first two factors here. Defining $\nu'_i := n_i/N$, what remains can now be re-arranged to obtain

$$\begin{aligned}
&\frac{N!}{n_1!n_2!\cdots n_k!} \\
&= \frac{\exp(N\log N)}{\exp(\nu'_1 N\log(\nu'_i N) + \ldots + \nu'_k N\log(\nu'_k N))} \\
&= \frac{\exp(N\log N)}{\exp(\nu'_1 N\log N + \ldots + \nu'_k N\log N)\exp((\nu'_1\log\nu'_1 + \ldots + \nu'_k\log\nu'_k)N)},
\end{aligned}$$

using the identity $\log(ab) \equiv \log a + \log b$.

Finally, since $\nu'_1 + \ldots + \nu'_k = 1$, the numerator and the first exponential in the denominator here cancel to leave

$$\frac{N!}{n_1!n_2!\cdots n_k!} = \exp(-(\nu'_1\log\nu'_1 + \ldots + \nu'_2\log\nu'_2)N + \mathrm{o}(N)).$$

Since we have assumed that $|\nu_i - \nu'_i| < \delta$ for all i, this is equal to

$$\exp(-(\nu_1\log\nu_1 + \ldots + \nu_2\log\nu_2)N + \mathrm{O}(c(\delta)N))$$

for some $c(\delta)$ which tends to 0 as $\delta \longrightarrow 0$. Since $T_{\nu,\delta,N}$ is a union of fewer than N^k sets all of this size up to subexponential corrections, it follows this is also the approximate size of $T_{\nu,\delta,N}$.

On the other hand, all strings in the set $\{\omega : \ f_i(\omega) = n_i \ \forall i\}$ have the same probability under ν, equal to $\nu_1^{n_1}\nu_2^{n_2}\cdots\nu_k^{n_k}$. Any set T that carries at least $1-\varepsilon$ of the mass of $\nu^{\otimes N}$ must contain at least $(1-\varepsilon)$-proportion of that set $\{\omega : \ f_i(\omega) = n_i \ \forall i\}$ for which this probability is maximized, and that propertion will still have size approximately $\exp(-(\nu_1\log\nu_1 + \ldots + \nu_2\log\nu_2)N)$ up to subexponential corrections. Therefore this is the answer to our question up to exponential accuracy.

Definition 3. *The quantity*

$$\mathrm{H}(\nu) := -(\nu_1\log\nu_1 + \ldots + \nu_k\log\nu_k)$$

*is called the **Shannon entropy** of ν. Intuitively, it is the exponential growth rate of the 'effective number' of different strings in $[k]^N$ that one stands a good chance of seeing if one draws such a string at random from $\nu^{\otimes N}$.*

Thus, this definition is a natural consequence of the exponential behaviour of multinomial coefficients, as provided by Stirling's approximation.

2.2 First properties

We next record some simple properties of the entropy functional. The first can be proved by elementary calculus, and the second and third then follow from simple calculations and the cancavity proved in the first. They can also be proved by using the defining property of Shannon entropy and counting typical sets, but this is generally a more tedious approach.

Lemma 4. *The function*

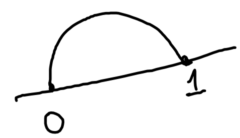

$$[0, \infty) \longrightarrow \mathbb{R} : t \mapsto \begin{cases} 0 & \text{if } t = 0 \\ -t \log t & \text{if } t > 0 \end{cases}$$

is strictly concave and crosses the X*-axis at* 0 *and* 1 *(so it is strictly positive on* $(0, 1)$*).* □

] how to prove this? Use lagrange multipliers.

Corollary 5. *For a fixed finite set* A*, the entropy functional*

$$\Pr A \longrightarrow [0, \infty) : \nu \mapsto \mathrm{H}(\nu)$$

is continuous, strictly concave, has image equal to $[0, \log |A|]$*, achieves* 0 *only on the point masses and achieves* $\log |A|$ *only on the uniform distribution.* □

Lemma 6. *If* ν_1*,* ν_2 *are probability measures on* A_1 *and* A_2*, and* $\mu \in \Pr(A_1 \times A_2)$ *is a coupling of* ν_1 *and* ν_2*, then*

$$\mathrm{H}(\mu) \leqslant \mathrm{H}(\nu_1) + \mathrm{H}(\nu_2),$$

with equality if and only if $\mu = \nu_1 \otimes \nu_2$*.* □

By the continuity of H on $\Pr A$ and the compactness of this space, the second part of this last lemma easily implies the following 'coercivity property':

Corollary 7. *For every finite sets* A_1, A_2 *and* $\varepsilon > 0$ *there is some* $\delta > 0$ *(depending on* ε*,* $|A_1|$ *and* $|A_2|$*) such that whenever* $\nu_i \in \Pr A_i$ *for* $i = 1, 2$*, if* μ *is a coupling of* ν_1 *and* ν_2 *for which*

$$\mathrm{H}(\mu) > \mathrm{H}(\nu_1) + \mathrm{H}(\nu_2) - \delta$$

$\delta(\varepsilon)$ Very close to uniform in TV.

then $\|\mu - \nu_1 \otimes \nu_2\|_{\mathrm{TV}} < \varepsilon$*.* □

Lemma 8 (Data-processing inequality)**.** *Let* A *and* B *be finite sets,* $\nu \in \Pr A$ *and* $\varphi : A \longrightarrow B$*. Then*

$$\mathrm{H}(\varphi_* \mu) \leqslant \mathrm{H}(\mu).$$

Proof. This provides a nice illustration of a proof in terms of typical sets. Let $\varepsilon > 0$, and suppose that $T \subseteq A^N$ is a set of minimal cardinality such that $\nu^{\otimes N}(T) > 1 - \varepsilon$. Then the set

$$S := \varphi^{\times N}(T) = \{(\varphi(a_1), \dots, \varphi(a_N)) :\ (a_1, \dots, a_N) \in T\} \subseteq B^N$$

satisfies

$$(\varphi_* \nu)^{\otimes N}(S) = \nu^{\otimes N}(T) > 1 - \varepsilon,$$

and has $|S| \leqslant |T|$, as an image of T. □

$H(\lambda) = \sum -\log \lambda_i$

Non colliding pts can be mapped to 1. If φ is invertible then we good.

2.3 Conditional entropy

Suppose now that A and B are finite sets, that $\lambda \in \Pr(A \times B)$, and that ν and μ are the marginals of λ on A and B. Let $\pi : A \times B \longrightarrow A$ be the coordinate projection. We may of course disintegrate λ over π:

$$\lambda = \int_A \delta_a \otimes \lambda_a \, \nu(\mathrm{d}a), \qquad (1)$$

∃ conditionale [illegible] $\lambda_a(c) = P(a, c)$

where $\lambda_a \in \Pr B$ is the conditional law of the B-coordinate given that the A-coordinate equals a (writing this as an integral, even though A is finite).

Definition 9. *In the above setting, the* ***conditional Shannon entropy of*** λ ***given*** π *is*

$$\mathrm{H}(\lambda \,|\, \pi) := \int_A \mathrm{H}(\lambda_a) \, \nu(\mathrm{d}a).$$

In the first place, conditional entropies are important because of their involvement in the calculation of the ordinary entropy of λ.

Lemma 10. *In the above setting, one has*

$$\mathrm{H}(\lambda) = \mathrm{H}(\nu) + \mathrm{H}(\lambda \,|\, \pi).$$

Proof. This is just a calculation using (1) and properties of log:

$$\begin{aligned}
\mathrm{H}(\lambda) &= -\sum_{a,b} \lambda\{(a,b)\} \log \lambda\{(a,b)\} = -\sum_{a,b} \nu\{a\}\lambda_a\{b\} \log(\nu\{a\}\lambda_a\{b\}) \\
&= -\sum_a (\nu\{a\} \log \nu\{a\}) \sum_b \lambda_a\{b\} - \sum_a \nu\{a\} \sum_b \lambda_a\{b\} \log \lambda_a\{b\} \\
&= \mathrm{H}(\nu) + \mathrm{H}(\lambda \,|\, \pi).
\end{aligned}$$

□

You get conditional entropy stuff because of
- disintegration of meas.
- properties of the logarithm

happens thanks to the log.

$$\lambda_a(b_2) := \begin{cases} \dfrac{\lambda(b_1 \times b_2)}{\nu(b_1)} & b_1 = a \\ 0 \end{cases}$$

$$H(\lambda_a) = \sum_\omega \lambda_a(\omega) \log \lambda_a(\omega)$$

Corollary 11. *If* $\lambda_1, \dots, \lambda_k \in \Pr A$ *and* $\lambda = \sum_{i=1}^{k} p_i \lambda_i$ *is a convex combination of them using some probability vector* $(p_1, \dots, p_k)$*, then*

$$\mathrm{H}(\lambda) \leqslant \mathrm{H}(p_1, \dots, p_k) + \sum_{i=1}^{k} p_i \mathrm{H}(\lambda_i).$$

$\varphi : \{1, \dots, k\} \times A \Rightarrow A$

Proof. Define the measure $\tilde{\lambda}$ on $[k] \times A$ by

$$\tilde{\lambda}(i, a) := p_i \lambda_i\{a\}.$$

This is a probability measure with marginals $(p_1, \dots, p_k)$ on $[k]$ and λ on A. Since λ is an image-measure of $\tilde{\lambda}$, the Data-Processing Inequality gives $\mathrm{H}(\lambda) \leqslant \mathrm{H}(\tilde{\lambda})$; on the other hand, the preceding lemma gives

$$\mathrm{H}(\tilde{\lambda}) = \mathrm{H}(p_1, \dots, p_k) + \sum_{i=1}^{k} p_i \mathrm{H}(\lambda_i).$$

□

3 The entropy rate of a shift system

Now suppose that $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S)$ is a shift $\mathbb{Z}^d$-system with a finite alphabet. For any subset $F \subseteq \mathbb{Z}^d$, let

$$\varphi^F : A^{\mathbb{Z}^d} \longrightarrow A^F$$

be the coordinate projection, and let $\nu^F := \varphi^F_* \mu$, the marginal distribution of the coordinates indexed by F.

Motivated by the results above for a sequence of i.i.d. A-valued random variables, we will now study the growth rate of the 'effective support' of the probability measures $\nu^F \in \Pr A^F$.

Having defined Shannon's entropy functional H in the previous section, this is most easily done by considering the entropy-values

$$\mathrm{H}(\nu^F)$$

as F increases through a sequence of rectangular boxes in $\mathbb{Z}^d$.

Lemma 12. *Given* ν*, the map* $F \mapsto \mathrm{H}(\nu^F)$*, defined on finite subsets of* $\mathbb{Z}^d$*, has the following properties:*

i) if $E \subseteq F \subseteq \mathbb{Z}^d$*, then* $\mathrm{H}(\nu^E) \leqslant \mathrm{H}(\nu^F)$*;*

ii) one has $\mathrm{H}(\nu^{E\cup F}) \leqslant \mathrm{H}(\nu^E) + \mathrm{H}(\nu^F)$ *for all* $E, F \subseteq \mathbb{Z}^d$.

Proof. If $E \subseteq F$, then $\nu^E = \varphi^{F,E}_* \nu^F$, where $\varphi^{F,E} : A^F \longrightarrow A^E$ is the coordinate-projection. Therefore part (i) follows from the data-processing inequality.

For part (ii), suppose first that $E \cap F = \varnothing$. Then $\nu^{E\cup F} \in \Pr(A^E \times A^F)$ is a coupling of ν^E and ν^F, so Lemma 6 gives

$$\mathrm{H}(\nu^{E\cup F}) \leqslant \mathrm{H}(\nu^E) + \mathrm{H}(\nu^F).$$

For general E and F, the desired result now follows by combining this special case with assertion (i), since

$$\mathrm{H}(\nu^{E\cup F}) = \mathrm{H}(\nu^{E\cup(F\setminus E)}) \leqslant \mathrm{H}(\nu^E) + \mathrm{H}(\nu^{F\setminus E}) \leqslant \mathrm{H}(\nu^E) + \mathrm{H}(\nu^F).$$

□

A subset of $\mathbb{Z}^d$ will be called a **rectangle** if it is the intersection of $\mathbb{Z}^d$ with a bounded rectangle in $\mathbb{R}^d$.

The subadditivity of the previous lemma implies that the values $\mathrm{H}(\nu^F)$ have a well-defined asymptotic behaviour which is exponential in $|F|$ as F increases through a sequence of rectangles. The following may be seen as a higher-dimensional analog of Fekete's Lemma (that is, the classical convergence result for subadditive sequences).

Lemma 13. *Given* ν *there is some* $h \geqslant 0$ *with the following property. For every* $\varepsilon > 0$ *there is an* $L > 0$ *such that whenever* $R \subseteq \mathbb{Z}^d$ *is a rectangle with all sides having length at least* L*, one has*

$$h|R| \leqslant \mathrm{H}(\nu^R) \leqslant (h+\varepsilon)|R|.$$

Proof. Let

$$h := \inf\left\{|C|^{-1}\mathrm{H}(\nu^C) : \ C \subseteq \mathbb{Z}^d \text{ a rectangle}\right\}.$$

We will show that this h has the asserted property. Given $\varepsilon > 0$, pick a rectangle C such that

$$|C|^{-1}\mathrm{H}(\nu^C) < h + \varepsilon.$$

Now choose $L \in \mathbb{N}$ large enough that the following holds:

> If $R \subset \mathbb{Z}^d$ is a rectangle with all side-lengths at least L, then there is a pairwise-disjoint family $\mathcal{C} = (C_1, C_2, \ldots, C_k)$ of translates of C, all contained in R, such that
>
> $$|C_1 \cup C_2 \cup \ldots \cup C_k| \geqslant (1-\varepsilon)|R|.$$

Having chosen this family $\mathcal{C}$, let $D := R\backslash(C_1 \cup C_2 \cup \cdots \cup C_k)$. Now a repeated application of Lemma 12 gives

$$\mathrm{H}(\nu^R) \leqslant \sum_{i=1}^{k} \mathrm{H}(\nu^{C_i}) + \mathrm{H}(\nu^D) \leqslant k\mathrm{H}(\nu^C) + |D|\mathrm{H}(\nu) < k|C|(h+\varepsilon) + \varepsilon|R|\mathrm{H}(\nu),$$

using that $\mathrm{H}(\nu^{C_i}) = \mathrm{H}(\nu^C)$ for all i, since ν is shift-invariant.

Since the sets C_i are disjoint and are all contained in R, the right-hand side above is not greater than

$$h|R| + \varepsilon(1 + \mathrm{H}(\nu))|R|,$$

so, since $\varepsilon > 0$ was arbitrary, this gives the desired upper bound.

On the other hand, one always has

$$\mathrm{H}(\nu^R) \geqslant h|R|$$

by the definition of h, so this completes the proof. $\square$

Definition 14. *The quantity h of the previous lemma is called the **entropy rate** (or **specific entropy**, or sometimes just **entropy**) of the shift system $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S)$. It is denoted $\mathrm{h}(\nu)$.*

*If $(X, \Sigma, \mu, T, \varphi)$ is a process, then its **entropy rate**, denoted $\mathrm{h}(\mu, T, \varphi)$, is defined to be the entropy rate of the resulting shift system obtained from the factor map*

$$(\varphi \circ T^n)_{n\in\mathbb{Z}^d} : (X, \Sigma, \mu, T) \longrightarrow (A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S).$$

3.1 The Shannon-McMillan Theorem

Having defined the entropy rate, we next turn to the first important theorem about it. It generalizes the simple fact that given $\nu \in \Pr A$, once N is large, the measure $\nu^{\otimes N}$ is mostly supported on the set $T_{\nu,\delta,N}$ of typical sequences, and gives all typical sequences roughly the same probability, to leading exponential order. This phenomenon generalizes to a shift system as follows.

Theorem 15 (Shannon-McMillan Theorem for $\mathbb{Z}^d$-systems)**.** *Let $|A| < \infty$, and let $\nu \in \Pr A^{\mathbb{Z}^d}$ be an* ergodic *shift-invariant measure. Let $h := \mathrm{h}(\nu)$. Then for every $\varepsilon > 0$ there is some $L \in \mathbb{N}$ for which the following holds:*

> *Whenever $R \subseteq \mathbb{Z}^d$ is a finite rectangle with all side-lengths at least L, there is a subset $Y \subseteq A^R$ such that*
>
> - *$\nu^R(Y) > 1 - \varepsilon$, and*

- *every* $a = (a_n)_{n\in R} \in Y$ *satisfies*

$$\mathrm{e}^{-(h+\varepsilon)|R|} < \nu^R(\{a\}) < \mathrm{e}^{-(h-\varepsilon)|R|}.$$

The ergodicity is crucial here. A first consequence of this result is that, as in the i.i.d. case, once a rectangle R is sufficiently large, the 'effective support' of the probability distribution $\nu^R \in \mathrm{Pr}(A^R)$ has size roughly $\mathrm{e}^{h|R|}$, and this effective size is realized by a sets of tuples in A^R that all have roughly the same probability under ν^R.

Theorem 15 is proved by showing that the quantities of interest are a kind of 'perturbation' of ergodic averages. Then, on the one hand, one applies the Mean Ergodic Theorem to those averages, and on the other one shows that the 'perturbation' becomes negligible as $N \longrightarrow \infty$.

The key to controlling the perturbation is the following basic inequality about entropy. Recall that if S is any set and $f : S \longrightarrow \mathbb{R}$, then $f^+ := \max\{f, 0\}$.

Lemma 16. *If* μ *and* ν *are two probability measures on a finite set* S*, then*

$$\int_S (-\log \mu\{s\} + \log \nu\{s\})^+ \,\mu(\mathrm{d}s) \leqslant \mathrm{e}^{-1}.$$

Proof. Let $T := \{s \in S : \ \mu\{s\} \leqslant \nu\{s\}\}$, so the integral of interest equals

$$\sum_{s\in T} \mu\{s\}\Big(-\log \frac{\mu\{s\}}{\nu\{s\}}\Big) = \sum_{s\in T} \nu\{s\}\Big(-\frac{\mu\{s\}}{\nu\{s\}} \log \frac{\mu\{s\}}{\nu\{s\}}\Big).$$

Since we restrict s to the set T on which $\frac{\mu\{s\}}{\nu\{s\}} \leqslant 1$, this sum is bounded by

$$\sum_{s\in T} \nu\{s\} \cdot \max_{0\leqslant t\leqslant 1}(-t\log t) \leqslant \max_{0\leqslant t\leqslant 1}(-t\log t) = \mathrm{e}^{-1},$$

where the last equality is a simple calculus exercise. □

Now consider a shift system $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S)$, and for any finite subset $F \subseteq \mathbb{Z}^d$ let

$$f_F : A^{\mathbb{Z}^d} \longrightarrow [0,\infty) : a \mapsto \begin{cases} -\log \nu^F\{\varphi^F(a)\} & \text{if } \nu^F\{\varphi^F(a)\} > 0 \\ 0 & \text{if } \nu^F\{\varphi^F(a)\} = 0. \end{cases}$$

Note that, by definition, $\nu^F\{\varphi^F(a)\} > 0$ for ν-a.e. a, so we are always in the first of these situations oustide of some ν-negligible event. We may also observe that

$$\nu^F\{\varphi^F(a)\} = \nu(\{a' : \ \varphi^F(a') = \varphi^F(a)\}),$$

so for each $a \in A^{\mathbb{Z}^d}$, we are effectively considering the event for another configuration $a' \in A^{\mathbb{Z}^d}$ that a' and a agree at the coordinates indexed by F, and recording (the negative logarithm of) the probability of this event.

This f_F is called the **information function** of the measure ν^F. Its first important property is the following simple calculation:

Lemma 17. *One has*

$$\int f_F \,\mathrm{d}\nu = \mathrm{H}(\nu^F).$$

□

In light of this calculation, the information function is often interpreted as a way of 'localizing' the total entropy $\mathrm{H}(\nu^F)$ to individual points of the space.

Corollary 18. *If $F = C_1 \cup C_2 \cup \cdots \cup C_m$ is a finite partition of a finite subset of $\mathbb{Z}^d$, then*

$$\int_{A^{\mathbb{Z}^d}} \Big(f_F - \sum_{i=1}^m f_{C_i}\Big)^+ \,\mathrm{d}\nu \leqslant \mathrm{e}^{-1}.$$

Proof. Let $S := A^F = A^{C_1} \times \cdots \times A^{C_m}$, and, for a tuple $a \in A^F$, let $a|_{C_i}$ be its projection onto the coordinates indexed by C_i. On A^F, consider the two measures

$$\nu^F \quad \text{and} \quad \nu^{C_1} \otimes \nu^{C_2} \otimes \cdots \otimes \nu^{C_m}.$$

In terms of these measures, the integral in question may be written as

$$\begin{aligned}
&\int_{A^F} \Big(-\log \nu^F\{a\} + \sum_{i=1}^m \log \nu^{C_i}\{a|_{C_i}\}\Big)^+ \,\nu^F(\mathrm{d}a)\\
&= \int_{A^F} \Big(-\log \nu^F\{a\} + \log(\nu^{C_1}\{a|_{C_1}\} \cdot \nu^{C_2}\{a|_{C_2}\} \cdots \nu^{C_m}\{a|_{C_m}\})\Big)^+ \,\nu^F(\mathrm{d}a)\\
&\qquad = \int_{A^F} \Big(-\log \nu^F\{a\} + \log(\nu^{C_1} \otimes \cdots \otimes \nu^{C_m})\{a\}\Big)^+ \,\nu^F(\mathrm{d}a).
\end{aligned}$$

This is of the form considered in Lemma 16, so it is bounded above by e^{-1}, as required. □

Proof of Theorem 15. First observe that the desired conclusion will follow if we show that

$$\||R|^{-1} f_R - h\|_1 \longrightarrow 0 \quad \text{as (min. side-length of } R) \longrightarrow \infty,$$

because

$$\||R|^{-1} f_R - h\|_1 = \int \frac{|-\log \nu^R(\{a\}) - h|R||}{|R|} \,\nu^R(\mathrm{d}a).$$

This already makes the result look very like the Norm Ergodic Theorem, except that the function $|R|^{-1}f_R$ is not a conventional ergodic average over the set of group-elements R. The key will be to make a comparison with a genuine ergodic average.

Step 1. Let $\varepsilon > 0$, and choose a rectangle C_0 so large that

$$\mathrm{H}(\nu^{C_0}) < (h+\varepsilon)|C_0|$$

(as is possible by Lemma 13). By translating and further enlarging C_0 if necessary, we may assume that it equals $[L]^d$ for some $L \in \mathbb{N}$.

Now, the function $g := |C_0|^{-1}f_{C_0} : A^{\mathbb{Z}^d} \longrightarrow [0,\infty)$ takes only finitely-many values, so is certainly in $L^1(\nu)$. We may therefore apply the Norm Ergodic Theorem for the action of the subgroup $L\cdot\mathbb{Z}^d \leqslant \mathbb{Z}^d$, to conclude that the averages

$$\frac{1}{M^d}\sum_{m\in[M]^d} g\circ S^{L\cdot m}$$

converge in $\|\cdot\|_1$ to some function $\overline{g} \in L^1(\mu)$ which is invariant under $S^{L\cdot m}$ for all $m \in \mathbb{Z}^d$.

By Lemma 17, $\int \overline{g}\,\mathrm{d}\nu = |C_0|^{-1}\mathrm{H}(\nu^{C_0}) \approx h$. If we knew that the subgroup $L\cdot\mathbb{Z}^d$ acted ergodically under S, then the Norm Ergodic Theorem would give that $\overline{g}$ is a.s. equal to this constant $|C_0|^{-1}\mathrm{H}(\nu^{C_0})$. However, since we do not know that this subgroup acts ergodically in general, we must be a little more careful.

Step 2. Now suppose $v \in C_0$, and let $C_1 := [ML]^d + v$ for some large M. One sees easily that $C_0 + L\cdot m \subseteq C_1$ for all $m \in [M-1]^d$, and that these translates of C_0 are pairwise-disjoint. Let

$$D := C_1 \backslash \bigcup_{m\in[M-1]^d}(C_0 + L\cdot m),$$

so $|D| = |C_1| - (M-1)^d L^d = \mathrm{O}(|C_1|/M)$.

Out next step is to control $\||C_1|^{-1}f_{C_1} - \overline{g}\|_1$. The key is to start by controlling only the positive part of this function. We have

$$\begin{aligned}\int(|C_1|^{-1}f_{C_1}-\overline{g})^+\,\mathrm{d}\nu &\leqslant |C_1|^{-1}\int\Big(f_{C_1}-\sum_{m\in[M-1]^d} f_{C_0}\circ S^{L\cdot m}\Big)^+\,\mathrm{d}\nu\\ &\quad+\int\Big(\frac{|C_0|}{|C_1|}\sum_{m\in[M-1]^d} g\circ S^{L\cdot m}-\overline{g}\Big)^+\,\mathrm{d}\nu\\ &\leqslant |C_1|^{-1}\int\Big(f_{C_1}-\sum_{m\in[M-1]^d} f_{C_0}\circ S^{L\cdot m}-f_D\Big)^+\,\mathrm{d}\nu\\ &\quad+|C_1|^{-1}\int f_D\,\mathrm{d}\nu\\ &\quad+\Big\|\frac{1}{(M-1)^d}\sum_{m\in[M-1]^d} g\circ S^{L\cdot m}-\overline{g}\Big\|_1,\end{aligned}$$

using that $|C_1|^{-1}f_{C_0}=M^{-d}g$. All three of these right-hand terms are now easily controlled:

- The first term is at most $|C_1|^{-1}\mathrm{e}^{-1}$, by Corollary 18, so this certainly tends to 0 as $M\longrightarrow\infty$.
- By Lemma 17 and Corollary 18, the second term is equal to

  $$|C_1|^{-1}\mathrm{H}(\nu^D)\leqslant\frac{|D|}{|C_1|}\mathrm{H}(\nu)=\mathrm{O}\Big(\frac{\mathrm{H}(\nu)}{M}\Big),$$

  which tends to 0 as $M\longrightarrow\infty$.
- The last term tends to zero as $M\longrightarrow\infty$ by the Norm Ergodic Theorem.

Therefore

$$\int(|C_1|^{-1}f_{C_1}-\overline{g})^+\,\mathrm{d}\nu\longrightarrow 0 \tag{2}$$

as $M\longrightarrow\infty$. On the other hand, Lemma 17 gives

$$\begin{aligned}\int|C_1|^{-1}f_{C_1}\,\mathrm{d}\nu=|C_1|^{-1}\mathrm{H}(\nu^{C_1})\geqslant h>|C_0|^{-1}\mathrm{H}(\nu^{C_0})-\varepsilon\\ =\int|C_0|^{-1}f_{C_0}\,\mathrm{d}\nu-\varepsilon=\int\overline{g}\,\mathrm{d}\nu-\varepsilon.\end{aligned}$$

Combining this with (2), one obtains that

$$\int(\overline{g}-|C_1|^{-1}f_{C_1})^+\,\mathrm{d}\nu\leqslant\int(\overline{g}-|C_1|^{-1}f_{C_1})\,\mathrm{d}\nu+\int(|C_1|^{-1}f_{C_1}-\overline{g})^+\,\mathrm{d}\nu<\varepsilon$$

for all sufficiently large M, and hence in fact

$$\|\overline{g} - |C_1|^{-1} f_{C_1}\|_1 < \varepsilon \tag{3}$$

for all sufficiently large M.

Step 3. Moreover, these last inequalities hold uniformly in our initial choice of 'shift vector' $v \in C_0$: that is, once M is sufficiently large, we have

$$\|(ML)^{-1} f_{[ML]^d} - \overline{g} \circ S^{-v}\|_1 = \|(ML)|^{-1} f_{[ML]^d + v} - \overline{g}\|_1 < \varepsilon \quad \forall v \in C_0,$$

and hence also

$$\Big\|(ML)^{-1} f_{[ML]^d} - \frac{1}{L^d} \sum_{v \in C_0} \overline{g} \circ S^{-v}\Big\|_1 < \varepsilon.$$

Now another look at the Norm Ergodic Theorem, together with the ergodicity of ν, gives that

$$\begin{aligned}\frac{1}{L^d} \sum_{v \in C_0} \overline{g} \circ S^{-v} &= \lim_{M \longrightarrow \infty} \frac{1}{(ML)^d} \sum_{v \in C_0} \sum_{m \in [M]^d} g \circ S^{Lm - v} \\ &= \lim_{M \longrightarrow \infty} \frac{1}{(ML)^d} \sum_{u \in [ML]^d} g \circ T^u = \int g \, \mathrm{d}\nu = |C_0|^{-1} \mathrm{H}(\nu^{C_0}).\end{aligned}$$

Therefore we have shown that

$$\||C_1|^{-1} f_{C_1} - h\|_1 < \||C_1|^{-1} f_{C_1} - |C_0|^{-1} \mathrm{H}(\nu^{C_0})\|_1 + \varepsilon < 2\varepsilon$$

provided C_1 is any rectangle with all side lengths sufficiently large.

Step 4. Finally, let M be large enough for the conclusion of Step 3, and let R be any rectangle all of whose sides are much longer than ML. Then, provided they are long enough, one can find a family of vectors $V \subseteq \mathbb{Z}^d$ such that the translates $[ML]^d + v$ for $v \in V$ are pairwise disjoint and all contained in R, and such that

$$E := R \backslash \bigcup_{v \in V} ([ML]^d + v)$$

has $|E| \leqslant \varepsilon |R| / \mathrm{H}(\nu)$. Having done so, another appeal to Corollary 18 in the same way as in Step 2 gives

$$\int \Big(|R|^{-1} f_R - |V|^{-1} \sum_{v \in V} (ML)^{-d} f_{[ML]^d + v}\Big)^+ \, \mathrm{d}\nu < \varepsilon$$

provided R is sufficiently large, and now, again as in Step 2, this implies that

$$\Big\| |R|^{-1} f_R - |V|^{-1} \sum_{v \in V} (ML)^{-d} f_{[ML]^d+v} \Big\|_1 < 2\varepsilon$$

provided R is large enough. Combining this with the conclusion of Step 3, it follows that

$$\| |R|^{-1} f_R - h \| < 4\varepsilon$$

for all sufficiently large rectangles R. Since $\varepsilon > 0$ was arbitrary, this completes the proof. □

The original Shannon-MacMillan Theorem was for $\mathbb{Z}$-systems, and is referred to in Information Theory as the Asymptotic Equipartition Property: see, for instance, [CT06, Chapter 3]. I believe the proof above (in the more general setting of amenable groups) is due to Moulin Ollagnier [MO83], building on earlier work of Kieffer [Kie75]. A nice presentation for general amenable groups can be found in [MO85, Section 4.4].

The theorem also has a strengthening to a pointwise-convergence result, due to Breiman, but the proof of that is rather trickier, and we will not use it in this course.

Theorem 19 (Shannon-McMillan-Breiman Theorem for $\mathbb{Z}^d$-actions)**.** *Given an ergodic $\nu \in \mathrm{Pr}^S A^{\mathbb{Z}^d}$ as before, one has*

$$N^{-d} f_{[N]^d}(a) \longrightarrow h \quad \text{for } \nu\text{-a.e. } a \in A^{\mathbb{Z}^d}.$$

□

The most immediate consequence of Theorem 15 is that we may reconnect with our initial intuition about the meaning of entropy.

Corollary 20. *In the setting of Theorem 15, for any $\alpha \in (0,1)$, one has*

$$\min\{|T| : T \subseteq A^R,\ \nu^R(T) \geqslant \alpha\} = \mathrm{e}^{h|R| + \mathrm{o}(|R|)}$$

as the minimal side-length of R tends to ∞.

Proof. Fix $\alpha \in (0,1)$ and also $\varepsilon > 0$, and let

$$T_R := \{a \in A^R :\ \mathrm{e}^{-(h+\varepsilon)|R|} < \nu^R(\{a\}) < \mathrm{e}^{-(h-\varepsilon)|R|}\}.$$

By Theorem 15, $\nu^R(T_R) \longrightarrow 1$ as R increases, and so for large enough R the required minimum is bounded above by $|T_R|$. Since every $a \in T_R$ satisfies $\nu^R(\{a\}) > \mathrm{e}^{-(h+\varepsilon)|R|}$, one has

$$1 \geqslant \nu^R(T_R) \geqslant |T_R| \mathrm{e}^{-(h+\varepsilon)|R|} \quad \Longrightarrow \quad |T_R| \leqslant \mathrm{e}^{(h+\varepsilon)|R|}.$$

On the other hand, if $\nu^R(T) \geqslant \alpha$ and R is so large that $\nu^R(T_R) > 1 - \alpha/2$, then also $\nu^R(T \cap T_R) < \alpha/2$. This requires that

$$\alpha/2 < \sum_{a \in T \cap T_R} \nu^R(\{a\}) \leqslant |T| \mathrm{e}^{-(h-\varepsilon)|R|} \quad \Longrightarrow \quad |T| \geqslant \mathrm{e}^{(h-\varepsilon)|R| + \log(\alpha/2)}.$$

Since these estimates both hold for large enough R for any $\varepsilon > 0$, this completes the proof. □

3.2 Affinity of the entropy rate

Theorem 15 will underly much of the later material in the course, but it applies only to ergodic measures. In some cases one cannot assume this, so we need to know what happens in the non-ergodic setting.

Happily, the situation is very simple. First recall that the set of invariant measures $\mathrm{Pr}^S A^{\mathbb{Z}^d}$ is a nonempty, compact, convex set for the vague topology. The entropy rate defines a function $\mathrm{h} : \mathrm{Pr}^S A^{\mathbb{Z}^d} \longrightarrow [0, \infty]$. It need not be continuous for the vague topology, but it is at least measurable, because

$$\mathrm{h}(\nu) = \lim_{L \longrightarrow \infty} L^{-d} \mathrm{H}(\nu^{[L]^d}),$$

and each function inside the right-hand limit is continuous, since it depends only on the finite marginal $\nu^{[L]^d} \in \mathrm{Pr}\, A^{[L]^d}$ and H is continuous on $\mathrm{Pr}\, A^{[L]^d}$.

Also, any $\nu \in \mathrm{Pr}\, A^{\mathbb{Z}^d}$ has an essentially unique ergodic decomposition: that it, a disintegration of the form

$$\nu = \int_Y \nu_y \, \theta(\mathrm{d}y),$$

where Y is an auxiliary compact metric space and the map $y \mapsto \mu_y : Y \longrightarrow \mathrm{Pr}^S A^{\mathbb{Z}^d}$ is measurable and takes values among the ergodic measures.

Proposition 21 (Affinity of the entropy rate)**.** *For any convex combination of invariant measures $\nu = \int_Y \nu_y \, \theta(\mathrm{d}y)$ on $A^{\mathbb{Z}^d}$ (not necessarily a disintegration), one has*

$$\mathrm{h}(\nu) = \int_Y \mathrm{h}(\nu_y) \, \theta(\mathrm{d}y).$$

Proof. Step 1. By definition,

$$\int_Y \mathrm{h}(\nu_y) \, \theta(\mathrm{d}y) = \int_Y \inf_{R \,\mathrm{rectangle}} |R|^{-1} \mathrm{H}(\nu_y^R) \, \theta(\mathrm{d}y),$$

and Lemma 13 shows that this is the same as

$$\int_Y \inf_{L \longrightarrow \infty} L^{-d} \mathrm{H}(\nu_y^{[L]^d}) \, \theta(\mathrm{d}y) = \int_Y \lim_{L \longrightarrow \infty} L^{-d} \mathrm{H}(\nu_y^{[L]^d}) \, \theta(\mathrm{d}y). \tag{4}$$

Now, all of the quantities $L^{-d}\mathrm{H}(\nu_y^{[L]^d})$ are non-negative, and also they are all bounded by

$$L^{-d} \cdot L^d \cdot \mathrm{H}(\nu_y) \leqslant \log |A|,$$

by Lemma 12 and then Corollary 5. Therefore the Dominated Convergence Theorem applies to the above right-hand integral in (4), equating it with

$$\lim_{L \longrightarrow \infty} L^{-d} \int_Y \mathrm{H}(\nu_y^{[L]^d})\, \theta(\mathrm{d}y).$$

This limit is greater than or equal to

$$\inf_{L \geqslant 1} L^{-d} \int_Y \mathrm{H}(\nu_y^{[L]^d})\, \theta(\mathrm{d}y),$$

but on the other hand, the left-hand integral in (4) is bounded above by this last infimum, so in fact all of these quantities are equal.

Step 2. Now fix an L. Suppose that $C \subseteq Y$ is measurable and $\theta(C) > 0$, and let $\nu_C := \theta(C)^{-1} \int_C \nu_y\, \theta(\mathrm{d}y)$. Since the function H is concave on $\Pr A^{[L]^d}$, Jensen's Inequality gives

$$\mathrm{H}(\nu_C^{[L]^d}) \geqslant \theta(C)^{-1} \int_C \mathrm{H}(\nu_y^{[L]^d})\, \theta(\mathrm{d}y).$$

If $\mathcal{P}$ is a partition of Y into positive-measure pieces, then we may average the above inequality over those pieces to obtain

$$\sum_{C \in \mathcal{P}} \mathrm{H}(\nu_C^{[L]^d}) \geqslant \int \theta(C) \mathrm{H}(\nu_y^{[L]^d})\, \theta(\mathrm{d}y). \tag{5}$$

On the other hand, the map $y \mapsto \nu_y^{[L]^d}$ is continuous and takes values in the compact set $\Pr A^{[L]^d}$, by Corollary 5. Therefore, if $\mathcal{P}$ is such that each of the sets $\{\nu_y^{[L]^d} \; : \; y \in C_i\}$ for $i = 1, 2, \ldots, k$ has sufficiently small diameter in $\Pr A^{[L]^d}$, then we may bring inequality (5) as close as we please to an equality. Combining this with the conclusion of Step 1, it follows that

$$\int_Y \mathrm{h}(\nu_y)\, \theta(\mathrm{d}y) = \inf_{L \geqslant 1} \inf_{\mathcal{P}} L^{-d} \sum_{C \in \mathcal{P}} \theta(C) \mathrm{H}(\nu_C^{[L]^d}) = \inf_{\mathcal{P}} \inf_{L \geqslant 1} L^{-d} \sum_{C \in \mathcal{P}} \theta(C) \mathrm{H}(\nu_C^{[L]^d}).$$

Step 3. It therefore suffices to show that

$$\lim_{L \longrightarrow \infty} L^{-d} \mathrm{H}(\nu^{[L]^d}) = \lim_{L \longrightarrow \infty} L^{-d} \sum_{C \in \mathcal{P}} \theta(C) \mathrm{H}(\nu_C^{[L]^d})$$

for any such partition $\mathcal{P}$. Since

$$\nu = \sum_{C \in \mathcal{P}} \theta(C)\nu_C,$$

another appeal to the concavity of H gives the inequality $\geqslant$, so we need only prove $\leqslant$. To this purpose, Corollary 11 gives

$$\mathrm{H}(\nu^F) \leqslant \mathrm{H}\big((\theta(C))_{C \in \mathcal{P}}\big) + \sum_{C \in \mathcal{P}} \theta(C)\mathrm{H}(\nu_C^{[L]^d}),$$

writing $(\theta(C))_{C \in \mathcal{P}}$ for the probability vector given by the ν-measures of the cells $C \in \mathcal{P}$. Since that is a fixed finite quantity (indeed, it is bounded by $\log|\mathcal{P}|$), we may divide by L^d and send $L \longrightarrow \infty$ to obtain

$$\lim_{L \longrightarrow \infty} L^{-d}\mathrm{H}(\nu^{[L]^d}) \leqslant \lim_{L \longrightarrow \infty} L^{-d} \sum_{C \in \mathcal{P}} \theta(C)\mathrm{H}(\nu_C^{[L]^d}),$$

as required. □

A more complete version of Proposition 21 can be obtained if one recalls that the ergodic decomposition may be realized as follows: if $\Phi \leqslant \mathcal{B}(A^{\mathbb{Z}^d})$ is the σ-algebra of invariant sets and $\nu \in \mathrm{Pr}^S A^{\mathbb{Z}^d}$, then the theorem on compact models gives a measurable, shift-invariant map $\pi : A^{\mathbb{Z}^d} \longrightarrow Y$ which generates Φ up to ν-negligible sets, and now $y \mapsto \nu_y$ may be defined on this space Y with the property that $\nu_y(\pi^{-1}\{y\}) = 1$ for ν-a.e. y.

Lemma 22. *Let $\nu = \int_Y \nu_y \,\theta(\mathrm{d}y)$ be as above, and let f_F denote the information function as in the proof of Theorem 15. Then*

$$\||R|^{-1} f_R - \mathrm{H}(\nu_{\pi(y)})\|_{L^1(\nu)} \longrightarrow 0$$

as the minimum side-length of R tends to ∞.

Proof. Theorem 15 shows that this holds for each of the norms $\|\cdot\|_{L^1(\nu_y)}$ for individual y, and the norm $\|\cdot\|_{L^1(\nu)}$ is just the average of these. □

Remark. Theorem 15 and Corollary 20 can both fail for non-ergodic measures; this is easily seen using Proposition 21 and some simple examples. ◁

References

[CT06] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition, 2006.

[Ell06] Richard S. Ellis. *Entropy, large deviations, and statistical mechanics*. Classics in Mathematics. Springer-Verlag, Berlin, 2006. Reprint of the 1985 original.

[Kec10] Alexander S. Kechris. *Global aspects of ergodic group actions*, volume 160 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2010.

[Kie75] J. C. Kieffer. A generalized Shannon-McMillan theorem for the action of an amenable group on a probability space. *Ann. Probability*, 3(6):1031–1037, 1975.

[KM10] Steven Kalikow and Randall McCutcheon. *An outline of ergodic theory*, volume 122 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.

[MO83] Jean Moulin Ollagnier. Théorème ergodique presque sous-additif et convergence en moyenne de l'information. *Ann. Inst. H. Poincaré Sect. B (N.S.)*, 19(3):257–266, 1983.

[MO85] Jean Moulin Ollagnier. *Ergodic theory and statistical mechanics*, volume 1115 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1985.

[Orn74] Donald S. Ornstein. *Ergodic theory, randomness, and dynamical systems*. Yale University Press, New Haven, Conn., 1974. James K. Whittemore Lectures in Mathematics given at Yale University, Yale Mathematical Monographs, No. 5.

[Rud90] Walter Rudin. *Fourier analysis on groups*. Wiley Classics Library. John Wiley & Sons Inc., New York, 1990. Reprint of the 1962 original, A Wiley-Interscience Publication.

[Var03] S. R. S. Varadhan. Large deviations and entropy. In *Entropy*, Princeton Ser. Appl. Math., pages 199–214. Princeton Univ. Press, Princeton, NJ, 2003.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: tim@cims.nyu.edu
URL: cims.nyu.edu/~tim

Topics in Ergodic Theory

Notes 10: Entropy of abstract systems

In the previous notes we introduced the Shannon entropy of a finitely-supported probability measure and the entropy rate of a shift system, and we proved the Shannon-McMillan Theorem, which asserted roughly that the marginal of a shift-invariant measure on a large enough rectangle is roughly uniformly supported on a certain collection of configurations.

Formally, the means of turning the entropy rate of shift systems into an isomorphism invariant of abstract systems is very simple: one simply supremizes over all shift-system factors.

Definition 1. *If (X, Σ, μ, T) is a p.-p. $\mathbb{Z}^d$-system, then its **Kolmogorov-Sinai (‘KS’) entropy** (or just **entropy**, or **metric entropy** in older texts) is*

$$\mathrm{h}(\mu, T) := \sup\left\{\mathrm{h}(\Phi_*\mu) : \ \Phi : X \longrightarrow A^{\mathbb{Z}^d},\ \Phi \circ T^m = S^m \circ \Phi\ \forall m \in \mathbb{Z}^d\right\},$$

where this supremum is allowed to take values in $[0, \infty]$.

However, making this definition is only half the work. In order to use it, one must find ways in which this supremum can be effectively computed or estimated. A priori, this could be challenging, because a given p.-p.s. admits a huge variety of factor maps to different shift systems.

The key fact is that the entropy rate of a shift system — which, a priori, depends on the product structure of that shift system — is actually invariant under abstract, ergodic-theoretic isomorphism. This will follow from a slightly stronger result: monotonicity under factor maps.

Theorem 2. *If A and B are finite alphabets, $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mu, S)$ is a shift system and $\Phi : A^{\mathbb{Z}^d} \longrightarrow B^{\mathbb{Z}^d}$ is measurable and shift-equivariant, then*

$$\mathrm{h}(\mu) \geqslant \mathrm{h}(\Phi_*\mu).$$

Corollary 3. *Isomorphic shift systems have equal entropy rate.* □

The proof of Theorem 2 will require some more preliminaries about measures on finite sets of the form A^R. The key new ingredient is that we will think of these, not just as finite sets, but as finite *metric spaces*.

1 Metric probability spaces

A **metric probability space** is a triple (X, d, μ) consisting of a metric space (X, d) and a probability measure $\mu \in \Pr X$.

First, if (X, d) is a metric space, $r \geqslant 0$ and $x \in X$, then we denote the **closed ball of radius** r **around** x by

$$B_r(x) := \{y \in X : \ d(x, y) \leqslant r\}.$$

More generally, for $S \subseteq X$, we set

$$B_r(S) := \bigcup_{x \in S} B_r(x).$$

Next, recall that for $r > 0$, then the r-covering number $\operatorname{cov}((X, d), r)$ is defined to be

$$\min\{|S| : \ B_r(S) = X\}.$$

Clearly $\operatorname{cov}((X, d), 0) = |X|$.

In the setting of metric probability spaces, we will need the following modification.

Definition 4. *If (X, d, μ) is a metric probability space, $\varepsilon > 0$ and $r \geqslant 0$, then the **ε-almost r-covering number** is defined by*

$$\operatorname{cov}_\varepsilon((X, d, \mu), r) := \min\{|S| : \ \mu(B_r(S)) > 1 - \varepsilon\}.$$

If the space (X, d) is clear from the context, then this will be abbreviated to $\operatorname{cov}_\varepsilon(\mu, r)$.

Lemma 5. *Suppose that $F : (X, d_X) \longrightarrow (Y, d_Y)$ is an L-Lipschitz map between compact metric spaces, and that $\mu \in \Pr X$. Then*

$$\operatorname{cov}_\varepsilon((Y, d_Y, F_*\mu), L\delta) \leqslant \operatorname{cov}_\varepsilon((X, d_X, \mu), \delta).$$

Proof. Since F is L-Lipschitz, for any $S \subseteq X$ one has $B_{L\delta}(F(S)) \supseteq F(B_\delta(S))$. Therefore, choosing S so that $\mu(B_\delta(S)) > 1 - \varepsilon$, we obtain

$$(F_*\mu)(B_{L\delta}(F(S))) \geqslant (F_*\mu)(F(B_\delta(S))) = \mu(F^{-1}(F(B_\delta(S)))) \geqslant \mu(B_\delta(S)) > 1 - \varepsilon,$$

and, of course, $|F(S)| \leqslant |S|$. □

1.1 Hamming metrics

Most of our concern will be with finite metric spaces that arise as follows.

Definition 6. *If A and S are finite sets, then the **Hamming metric** on A^S is defined by*

$$d_{\mathrm{H}}((a_s)_{s\in S},(a'_s)_{s\in S}) := |\{s\in S:\ a_s\neq a'_s\}|.$$

That is, it counts the number of coordinates in which $(a_s)_s$ and $(a'_s)_s$ differ.

The key in this lecture will be to think of the rectangle-marginals of a shift system $(A^{\mathbb{Z}^d},\mathcal{B}(A^{\mathbb{Z}^d}),\mu,S)$ as giving a family of metric probability spaces

$$(A^R,d_{\mathrm{H}},\mu^R),$$

and studying asymptotic features of the geometry of these spaces as R increases.

The first important tool to this purpose is an estimate for the cardinalities of balls in (A^N,d_{H}). We will see that it amounts to a re-use of some of the estimates in our initial motivation for Shannon entropy. It is most easily obtained as a special case of a different theorem: Cramér's Theorem about Large Deviations in probability. A more complete probabilistic account can be found, for instance, in [Kal02, Chapter 27].

Lemma 7. *Let X_1, X_2, ..., X_N be i.i.d. $\{0,1\}$-valued random variables on some background probability space $(\Omega,\mathcal{F},\mathsf{P})$ such that $\mathsf{P}(X_i=1)=p$. Then*

$$\mathsf{P}\Big\{\sum_{n=1}^N X_n\leqslant rN\Big\}\leqslant(\mathrm{e}^{\mathrm{H}(r,1-r)}p^r(1-p)^{(1-r)})^N\quad\forall r\in[0,p].$$

Proof. By Markov's Inequality, for any $s>0$ the probability in question may be estimated by

$$\begin{aligned}\mathsf{P}\Big\{\sum_{n=1}^N X_n\leqslant rN\Big\} &= \mathsf{P}\Big\{\prod_{n=1}^N \mathrm{e}^{sX_n}\leqslant \mathrm{e}^{srN}\Big\}\\ &\leqslant \mathrm{e}^{-srN}\mathsf{E}\prod_{n=1}^N \mathrm{e}^{sX_n}=(\mathsf{E}\mathrm{e}^{s(X_1-r)})^N\\ &= (\mathrm{e}^{-sr}(p\mathrm{e}^s+1-p))^N,\end{aligned}$$

using that these r.v.s are i.i.d. for the penultimate equality. This gives a bound for any choice of s. A little calculus now shows that the optimal choice is $s=$

$\log \frac{r(1-p)}{p(1-r)}$, and this gives

$$\begin{aligned} \mathrm{e}^{-sr}(p\mathrm{e}^s + 1 - p) &= \frac{p^r(1-r)^r}{r^r(1-p)^r}\Big(\frac{pr(1-p)}{p(1-r)} + 1 - p\Big) \\ &= \frac{p^r(1-p)^{1-r}}{r^r(1-r)^{1-r}} = \mathrm{e}^{\mathrm{H}(r,1-r)}p^r(1-p)^{1-r}, \end{aligned}$$

as required. □

Lemma 8. *For the balls in the Hamming metric space* (A^N, d_{H})*, one has that* $|B_{rN}(a)|$ *does not depend on* $a \in A^N$ *and satisfies*

$$|B_{rN}(a)| \leqslant \mathrm{e}^{\mathrm{H}(r,1-r)N}|A|^{rN} \quad \forall r \in [0, 1 - 1/|A|].$$

Proof. Fix $a \in A^N$, let ν be the uniform probability measure on A^N, and on the probability space (A^N, ν) consider the random variables

$$X_m(x_1, \dots, x_N) := 1_{\{x_n \neq a_n\}}.$$

Then these are i.i.d. $\{0,1\}$-valued random variables (for they all depend on independent coordinates) with common distribution given by

$$\nu(\{X_n = 1\}) = 1 - 1/|A|.$$

The relevance of these random variables is that

$$d_{\mathrm{H}}(a, x) = \sum_{n=1}^{N} X_n(x),$$

so now an appeal to Lemma 7 gives

$$\frac{|B_{rN}(a)|}{|A^N|} = \nu\Big\{\sum_{n=1}^{N} X_n \leqslant rN\Big\} \leqslant \Big(\mathrm{e}^{\mathrm{H}(r,1-r)}\frac{(1-1/|A|)^r}{|A|^{1-r}}\Big)^N \leqslant \mathrm{e}^{\mathrm{H}(r,1-r)N}\frac{|A|^{rN}}{|A|^N},$$

as required. □

1.2 The Wasserstein metric

Given a compact metric space (X, d), the space $\mathrm{Pr}\, X$ is compact in the topology of vague convergence. This latter topology is also metrizable, but in the sequel it will be important to work with certain natural metrics that generate it. The most important for us is defined in terms of couplings. Recall that if $\mu, \nu \in \mathrm{Pr}\, X$, then a (μ, ν)-coupling is some $\lambda \in \mathrm{Pr}\, X^2$ whose first and second marginals are μ and ν respectively. We now write $\mathrm{Cpl}(\mu, \nu)$ for the set of these. The product measure $\mu \otimes \nu$ is always an example.

Definition 9 (Wasserstein metric)**.** *If* $\lambda \in \mathrm{Cpl}(\mu, \nu)$*, then its* d***-cost*** *is the integral*

$$\int_{X^2} d(x, y)\, \lambda(\mathrm{d}x, \mathrm{d}y),$$

and the resulting ***Wasserstein*** *metric on* $\Pr X$ *is defined by*

$$\mathrm{W}_1^d(\mu, \nu) := \inf_{\lambda \in \mathrm{Cpl}(\mu,\nu)} \int_{X^2} d(x, y)\, \lambda(\mathrm{d}x, \mathrm{d}y).$$

In case d *is a Hamming metric, we abbreviate* $\mathrm{W}_1^{d_{\mathrm{H}}} =: \mathrm{W}_1^{\mathrm{H}}$

The subscript in 'W_1' refers to its location in a family of metrics W_p, $p \in [1, \infty]$, but we will not use the other members of this family.

Lemma 10. *The quantity* W_1^d *defines a metric on* $\Pr X$*, and convergence in* W_1 *implies convergence in the vague topology on* $\Pr X$*.*

Proof. First observe that, since $d : X^2 \longrightarrow [0, \infty)$ is continuous, the d-cost is a continuous functional on $\mathrm{Cpl}(\mu, \nu)$ for the vague topology. On the other hand, $\mathrm{Cpl}(\mu, \nu)$ is a vaguely closed subset of the compact set $\Pr X^2$, so the d-weight must attain its infimum over $\mathrm{Cpl}(\mu, \nu)$.

Clearly $\mathrm{W}_1(\mu, \nu) \geqslant 0$. If $\mathrm{W}_1(\mu, \nu) = 0$, then, by the above, there is some $\lambda \in \mathrm{Cpl}_d(\mu, \nu)$ such that $\int_{X^2} d\, \mathrm{d}\lambda = 0$. This implies that

$$\lambda\{d = 0\} = \lambda\{(x, x) : x \in X\} = 1,$$

and hence

$$\begin{aligned}\mu(A) = \lambda(A \times X) &= \lambda((A \times X) \cap \{(x, x) : x \in X\}) \\ &= \lambda((X \times A) \cap \{(x, x) : x \in X\}) = \lambda(X \times A) = \nu(A)\end{aligned}$$

for any Borel $A \subseteq X$, so $\mu = \nu$.

The symmetry of W_1 is obvious. Next, suppose that $\lambda_1 \in \mathrm{Cpl}(\mu, \nu)$ and $\lambda_2 \in \mathrm{Cpl}(\nu, \theta)$. Consider disintegrations of these over the second and first coordinates, respectively:

$$\lambda_1 = \int_X \lambda_{1,y} \otimes \delta_y\, \nu(\mathrm{d}y) \quad \text{and} \quad \lambda_2 = \int_X \delta_y \otimes \lambda_{2,y}\, \nu(\mathrm{d}y).$$

Now define $\lambda \in \Pr X^2$ by

$$\lambda = \int_X \lambda_{1,y} \otimes \lambda_{2,y}\, \nu(\mathrm{d}y).$$

On the one hand, this is an element of $\mathrm{Cpl}(\mu, \theta)$: for instance, for the first marginal one has

$$\lambda(A \times X) = \int_X (\lambda_{1,y} \otimes \lambda_{2,y})(A \times X)\, \nu(\mathrm{d}y)$$
$$= \int_X \lambda_{1,y}(A)\, \nu(\mathrm{d}y) = \lambda_1(A \times X) = \mu(A).$$

On the other hand, this λ satisfies

$$\begin{aligned}
\int d\, \mathrm{d}\lambda &= \int_X \int_{X^2} d(x,z)\, \lambda_{1,y}(\mathrm{d}x) \otimes \lambda_{2,y}(\mathrm{d}z)\, \nu(\mathrm{d}y) \\
&\leqslant \int_X \int_{X^2} (d(x,y) + d(y,z))\, \lambda_{1,y}(\mathrm{d}x) \otimes \lambda_{2,y}(\mathrm{d}z)\, \nu(\mathrm{d}y) \\
&= \int_X \int_X d(x,y)\, \lambda_{1,y}(\mathrm{d}x)\, \nu(\mathrm{d}y) + \int_X \int_X d(y,z)\, \lambda_{2,y}(\mathrm{d}z)\, \nu(\mathrm{d}y) \\
&= \int d\, \mathrm{d}\lambda_1 + \int d\, \mathrm{d}\lambda_2.
\end{aligned}$$

Taking infima over λ_1 and λ_2 gives

$$\mathrm{W}_1(\mu, \theta) \leqslant \mathrm{W}_1(\mu, \nu) + \mathrm{W}_1(\nu, \theta).$$

Finally, given $f_1, \dots, f_k \in C(X)$, let $M := \max_{i \leqslant k} \|f_i\|_\infty$. This max is finite, because X is compact, and for the same reason the functions f_i are uniformly continuous. Therefore, given also $\varepsilon > 0$, there is some $\delta > 0$ such that

$$d(x,y) < \delta \quad \Longrightarrow \quad |f_i(x,) - f_i(y)| < \varepsilon \quad \forall i \leqslant k.$$

Having chosen this δ, suppose that $\mu, \nu \in \Pr X$ and that $\int d\, \mathrm{d}\lambda < \varepsilon\delta$ for some $\lambda \in \mathrm{Cpl}(\mu, \nu)$. Then for any $i \leqslant k$ one has

$$\Big| \int f_i\, \mathrm{d}\mu - \int f_i\, \mathrm{d}\nu \Big| = \Big| \int (f_i(x) - f_i(y))\, \lambda(\mathrm{d}x, \mathrm{d}y) \Big| \leqslant \int |f_i(x) - f_i(y)|\, \lambda(\mathrm{d}x, \mathrm{d}y)$$
$$\leqslant \varepsilon \cdot \lambda\{d < \delta\} + 2M\lambda\{d \geqslant \delta\} \leqslant \varepsilon + 2M\delta^{-1} \int d\, \mathrm{d}\lambda \leqslant (1 + 2M)\varepsilon,$$

where the penultimate bound follows from Markov's Inequality. Since $\varepsilon > 0$ was arbitrary, this shows that the topology generated by W_1 is at least as fine as the vague topology. □

Remark. In fact, a little more work shows that W_1 precisely generates the vague topology, but we will not need that here. ◁

Lemma 11. *If* $r := \mathrm{W}_1(\mu, \nu)$*, then for any Borel* $A \subseteq X$ *one has*

$$\nu(B_\delta(A)) \geqslant \mu(A) - r/\delta \quad \forall \delta > 0.$$

Proof. Choose $\lambda \in \mathrm{Cpl}(\mu, \nu)$ minimizing $\int d \, \mathrm{d}\lambda$. Let $B := X \backslash B_\delta(A)$. Then $A \times B \subseteq \{d > \delta\}$, and so

$$\delta\lambda(A \times B) = \int \delta 1_{A\times B} \, \mathrm{d}\lambda \leqslant \int d \, \mathrm{d}\lambda = r.$$

On the other hand,

$$\lambda(A \times B) = \lambda(A \times X) - \lambda(A \times B_\delta(A)) \geqslant \mu(A) - \nu(B_\delta(A)).$$

Combining these inequalities completes the proof. □

Corollary 12. *If* $r := \mathrm{W}_1(\mu, \nu)$*, then*

$$\mathrm{cov}_{\varepsilon + r/\delta}(\mu, \delta + \delta') \leqslant \mathrm{cov}_\varepsilon(\nu, \delta') \quad \forall \varepsilon, \delta, \delta' > 0.$$

Proof. Observe that for any $S \subseteq X$ one has $B_{\delta+\delta'}(S) \supseteq B_\delta(B_{\delta'}(S))$. Therefore, if $S \subseteq X$ is such that $\nu(B_{\delta'}(S)) > 1 - \varepsilon$, then the preceding lemma gives

$$\mu(B_{\delta+\delta'}(S)) \geqslant \mu(B_\delta(B_{\delta'}(S))) \geqslant \nu(B_{\delta'}(S)) - r/\delta > 1 - (\varepsilon + r/\delta).$$

□

1.3 The $\overline{\mathrm{d}}$-metric

In the setting of a shift action $S : \mathbb{Z}^d \curvearrowright A^{\mathbb{Z}^d}$, one can formulate a metric on $\mathrm{Pr}^S A^{\mathbb{Z}^d}$ analogous to the Wasserstein metric on $\mathrm{Pr}\, A^N$, normalized by N. Intuitively, we should like to mimic the infimum defining the Wasserstein metric, but with an integrand that captures the 'fraction of coordinates' where two configurations in $A^{\mathbb{Z}^d}$ differ. This 'fraction' does not always make sense, but we can get around this in view of the fact that we work exclusively with shift-invariant measures.

Definition 13. *The metric* $\overline{\mathrm{d}}$ *on* $\mathrm{Pr}^S A^{\mathbb{Z}^d}$ *is defined by*

$$\overline{\mathrm{d}}(\mu, \nu) := \inf \big\{ \lambda(\{(a, a') : \ a_0 \neq a'_0\}) : \ \lambda \in \mathrm{J}(\mu, \nu) \big\},$$

where $\mathrm{J}(\mu, \nu) \subseteq \mathrm{Pr}(A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d})$ *is the set of all joinings of* μ *and* ν*.*

The connection with Wasserstein metrics is made more formal by the following observation:

Lemma 14. *For any $\lambda \in \mathrm{J}(\mu,\nu)$ and any $R \subseteq \mathbb{Z}^d$, one has*

$$\int_{A^{\mathbb{Z}^d}\times A^{\mathbb{Z}^d}} d_{\mathrm{H}}(a|_R, a'|_R)\,\lambda(\mathrm{d}a,\mathrm{d}a') = |R|\lambda\{a_0 \neq a'_0\}.$$

If λ is ergodic for the shift-action on $(A\times A)^{\mathbb{Z}^d}$, then one also has

$$\lambda\{a_0 \neq a'_0\} = \lim_{L\longrightarrow\infty} L^{-d} d_{\mathrm{H}}(a|_{[L]^d}, a'|_{[L]^d}) \quad \textit{for } \lambda\textit{-a.e. } (a,a').$$

Proof. The first equality follows because

$$\int d_{\mathrm{H}}(a|_R, a'|_R)\,\lambda(\mathrm{d}a,\mathrm{d}a') = \sum_{m\in R}\int 1_{\{(S^m a)_0 \neq (S^m a')_0\}}\,\lambda(\mathrm{d}a,\mathrm{d}a'),$$

and all of the integrands on the right are equal to $\lambda\{a_0 \neq a'_0\}$ by the shift-invariance of λ. In case λ is ergodic, the second equality now follows from the same formula and the Pointwise Ergodic Theorem. □

Note, however, that $\overline{\mathrm{d}}$ does not itself arise in the same way as a Wasserstein metric, because a priori the limit

$$\lim_{L\longrightarrow\infty} L^{-d} d_{\mathrm{H}}(a|_R, a'|_R)$$

does not make sense for all pairs (a,a').

Also, since the functional $\lambda \mapsto \lambda\{a_0 \neq a'_0\}$ is continuous for the vague topology on $\Pr A^{\mathbb{Z}^d}$, one also has that the infimum in Definition 13 is always attained at some λ.

2 Isomorphism-invariance of the entropy rate

The first step towards Theorem 2 is to obtain another expression for the entropy rate of an ergodic shift-invariant meausre, giving is some more 'robustness'.

Proposition 15. *If $\nu \in \Pr A^{\mathbb{Z}^d}$ is ergodic and shift-invariant, then for every $\varepsilon,\delta \in (0,1)$ and $\eta > 0$ one has*

$$\mathrm{cov}_\varepsilon\big((A^R, d_{\mathrm{H}}, \nu^R), \delta|R|\big) \geqslant \mathrm{e}^{(\mathrm{h}(\nu) - \mathrm{H}(\delta,1-\delta) - \delta\log|A| - \eta)|R|} \tag{1}$$

for all sufficiently large rectangles R. Consequently, for any $\varepsilon \in (0,1)$ one has

$$\mathrm{h}(\nu) = \sup_{\delta>0}\liminf_{R\uparrow}\frac{1}{|R|}\log\mathrm{cov}_\varepsilon\big((A^N, d_{\mathrm{H}}, \nu^R), \delta|R|\big),$$

where the $\liminf$ *is taken as all side lengths of R tends to ∞.*

Proof. If R is sufficiently large and $S \subseteq A^R$ satisfies $\nu^R(B_{\delta|R|}(S)) > 1-\varepsilon$, then the support-counting corollary of the Shannon-McMillan Theorem (Corollary 20 of Notes 9) gives

$$|B_{\delta|R|}(S)| \geqslant 2\mathrm{e}^{(h-\eta)|R|}.$$

On the other hand,

$$|B_{\delta|R|}(S)| \leqslant \sum_{x\in S} |B_{\delta|R|}(x)| \leqslant |S|\mathrm{e}^{(\mathrm{H}(\delta,1-\delta)+\delta\log|A|)|R|},$$

by Lemma 8. Combining these inequalities gives

$$|S| \geqslant 2\mathrm{e}^{(h-\eta-(\mathrm{H}(\delta,1-\delta)+\delta\log|A|))|R|}.$$

The second conclusion follows by letting $\eta \longrightarrow 0$ and $\delta \longrightarrow 0$, since then also $\mathrm{H}(\delta,1-\delta)+\delta\log|A| \longrightarrow 0$. ☐

The importance of this proposition is that the covering numbers in question for $\delta > 0$ behave more 'smoothly' under factor maps than the actual count of the minimal number of configurations needed to support most of ν^R. The kind of robustness we gain is that two different, but small, values of δ will both asymptotically give good approximations to $\mathrm{h}(\nu)$. This will be used in conjunction with the following.

Lemma 16. *Let $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mu, S)$ be a shift system and $\Phi : A^{\mathbb{Z}^d} \longrightarrow B^{\mathbb{Z}^d}$ a shift-equivariant map. Then for every $\delta > 0$ there are $L, N \geqslant 1$ such that for any rectangle R with all sides longer than N, there is an L-Lipschitz map*

$$\Psi : A^R \longrightarrow B^R$$

such that

$$\int d_{\mathrm{H}}\big(\Phi(a)|_R, \Psi(a|_R)\big)\,\mu(\mathrm{d}a) < \delta|R|.$$

Therefore

$$\mathrm{W}_1^{\mathrm{H}}(\nu^R, \Psi_*\mu^R) < \delta|R|.$$

Proof. Since $\varphi_0 : A^{\mathbb{Z}^d} \longrightarrow B$ is measurable, it may be approximated in measure by a map depending on only finite many coordinates: that is, there are $K \geqslant 1$ and a map $\psi : A^{[-K;K]^d} \longrightarrow B$ such that

$$\mu(\{a : \varphi_0(a) \neq \psi(a|_{[-K;K]^d})\}) < \delta. \tag{2}$$

Now suppose that R is a rectangle with all sides longer than some $N \in \mathbb{Z}$, and partition it as $R_1 \cup D$, where

$$R_1 := \{m \in R : m + [-K;K]^d \subseteq R\}.$$

Easy estimates give $|D| \leqslant 2dK|R|/N$, so for fixed K this is $\mathrm{O}(|R|/N)$. Pick arbitrarily some letters $b_m \in B$ for $m \in D$, and now define a map $\Psi : A^R \longrightarrow B^R$ as follows:

$$\Psi(a) := (\psi_m(a))_{a\in R} := \begin{cases} \psi(a|_{m+[-K;K]^d}) & \text{if } m \in R_1 \\ b_m & \text{if } m \in D. \end{cases}$$

Let $L := (2K+1)^d$.

Each output coordinate $\psi_m(a)$ depends on at most $(2K+1)^d = L$ input coordinate (if $m \in D$, then it does not depend on a at all), This implies that if a and a' differ in at most d coordinates, then their Ψ-images can differ in at most Ld coordinates: that is, that Ψ is L-Lipschitz.

On the other hand, one has

$$\begin{aligned} \int d_{\mathrm{H}}\big(\Phi(a)|_R, \Psi(a|_R)\big)\,\mu(\mathrm{d}a) &= \int \sum_{m\in R} 1_{\{\varphi_m(a)\neq\psi_m(a|_R)\}}\,\mu(\mathrm{d}a) \\ &= |D| + \sum_{m\in R_1} \mu(\{a:\ \varphi_m(a) \neq \psi_m(a|R)\}) \\ &= |D| + \sum_{m\in R_1} \mu(\{a:\ \varphi_0(T^m a) \neq \psi(T^m a|_{[-K;K]^d})\}) \\ &= |D| + |R_1|\mu(\{a:\ \varphi_0(a) \neq \psi(a|_{[-K;K]^d})\}). \end{aligned}$$

By (2) and the fact that $|D| = \mathrm{O}(|R|/N)$, this will still be $< \delta|R|$ provided we chose N large enough.

The last conclusion follows because the measure

$$\lambda := \int \delta_{(\Phi(a)|_R,\Psi(a|_R))}\,\mu(\mathrm{d}a)$$

is a coupling of ν^R (since $\Phi_*\mu = \nu$) and $\Psi_*\mu^R$, so

$$\mathrm{W}_1^{\mathrm{H}}(\nu^R, \Psi_*\mu^R) \leqslant \int d_{\mathrm{H}}(b,c)\,\lambda(\mathrm{d}b,\mathrm{d}c) = \int d_{\mathrm{H}}\big(\Phi(a)|_R, \Psi(a|_R)\big)\,\mu(\mathrm{d}a).$$

□

Proof of Theorem 2. Step 1. In case μ is ergodic, by Proposition 15, it suffices to prove the following:

> For every $\delta, \varepsilon > 0$ there are some $\delta', \varepsilon' > 0$ such that
>
> $$\mathrm{cov}_\varepsilon(\nu^R, \delta|R|) \leqslant \mathrm{cov}_{\varepsilon'}(\mu^R, \delta'|R|)$$
>
> for all sufficiently large R.

Choose $r > 0$ so small that $\varepsilon - r/\delta > 0$. The preceding lemma gives $L > 0$ and, for all sufficiently large R, an L-Lipschitz map $\Psi : A^R \longrightarrow B^R$ such that $\mathrm{W}_1^{\mathrm{H}}(\nu^R, \Psi_*\mu^R) < r|R|$. By Corollary 12 and then Lemma 5, we obtain

$$\mathrm{cov}_\varepsilon(\nu^R, \delta|R|) \leqslant \mathrm{cov}_{\varepsilon - r/\delta}(\Psi_*\mu^R, \delta|R|/2) \leqslant \mathrm{cov}_{\varepsilon - r/\delta}(\mu^R, \delta|R|/2L).$$

Letting $\delta' := \delta/L$ and $\varepsilon' := \varepsilon - r/\delta$, this completes the proof.

Step 2. If μ is not ergodic, let $\mu = \int_Y \mu_y\, \theta(\mathrm{d}y)$ be an ergodic decomposition of it. Then applying Φ_* gives $\Phi_*\mu = \int_Y \Phi_*\mu_y\, \theta(\mathrm{d}y)$, and each image measure $\Phi_*\mu_y$ is still ergodic. Therefore, applying Step 1 to the measures μ_y and recalling the affinity of h (Proposition 21 in Notes 9), we obtain

$$\mathrm{h}(\mu) = \int_Y \mathrm{h}(\mu_y)\, \theta(\mathrm{d}y) \geqslant \int_Y \mathrm{h}(\Phi_*\mu_y)\, \theta(\mathrm{d}y) = \mathrm{h}(\Phi_*\mu).$$

□

The importance of Theorem 2 for the KS entropy is the following.

Corollary 17. *If $\varphi : X \longrightarrow A$ is a generating observable for a p.-p.s. (X, Σ, μ, T), and $\Phi : X \longrightarrow A^{\mathbb{Z}^d}$ is the resulting factor map, then*

$$\mathrm{h}(\mu, T) = \mathrm{h}(\Phi_*\mu).$$

Proof. The inequality $\mathrm{h}(\mu, T) \geqslant \mathrm{h}(\Phi_*\mu)$ is immediate. In the reverse direction, suppose that $\Psi = (\psi_m)_m : X \longrightarrow B^{\mathbb{Z}^d}$ is any other factor map. Since φ is generating, we have

$$\Phi^{-1}(\mathcal{B}(A^{\mathbb{Z}^d})) = \sigma\text{-alg}((\varphi_m)_{m \in \mathbb{Z}^d}) = \Sigma \geqslant \Psi^{-1}(\mathcal{B}(B^{\mathbb{Z}^d})) \qquad \mathrm{mod}\ \mu.$$

In particular, the level-set partition $(\psi_0^{-1}\{b\})_{b \in B}$ is measurable with respect to $\Phi^{-1}(\mathcal{B}(A^{\mathbb{Z}^d}))$ up to μ-negligible sets. We may adjust ψ_0 on a μ-negligible set without changing $\Psi_*\mu$, and so we may actually assume that ψ_0 is measurable with respect to $\Phi^{-1}(\mathcal{B}(A^{\mathbb{Z}^d}))$, and therefore $\psi_0 = \xi_0 \circ \Phi$ for some $\xi_0 : A^{\mathbb{Z}^d} \longrightarrow B$. Letting

$$\Xi := (\xi_0 \circ T^m)_{m \in \mathbb{Z}^d} : A^{\mathbb{Z}^d} \longrightarrow B^{\mathbb{Z}^d},$$

this has now created a commutative diagram of factor maps

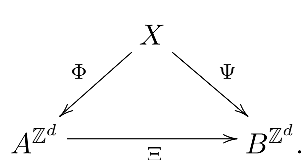

Therefore Theorem 2 gives

$$\mathrm{h}(\Psi_*\mu) = \mathrm{h}(\Xi_*\Phi_*\mu) \leqslant \mathrm{h}(\Phi_*\mu),$$

completing the proof. $\square$

This corollary is a version for $\mathbb{Z}^d$-systems of the 'Kolmogorov-Sinai Theorem'. It has the following reassuring further corollary.

Corollary 18. *For any shift system $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \nu, S)$, its abstract KS entropy is the same as its entropy rate.* $\square$

3 First examples and applications

3.1 Bernoulli shifts

It is high time we completed our initially-promised application of entropy: distinguishing Bernoulli shifts.

Lemma 19. *For a Bernoulli shift $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mathbf{p}^{\otimes\mathbb{Z}^d}, S)$ over some finite probability space $(A, \mathbf{p})$, the KS entropy is $\mathrm{H}(\mathbf{p})$.*

Proof. By Corollary 17, it suffices to check one generating observable. One such is given by $\pi_0 : (a_m)_m \mapsto a_0$, since $(\pi_0 \circ S^m)_m = \mathrm{id}_{A^{\mathbb{Z}^d}}$, and for this observable the marginal distribution on a rectangle R is just $\mathbf{p}^{\otimes R}$, so the entropy rate is

$$\lim_R |R|^{-1}\mathrm{H}(\mathbf{p}^{\otimes R}) = \lim_R |R|^{-1} \cdot |R| \cdot \mathrm{H}(\mathbf{p}) = \mathrm{H}(\mathbf{p}).$$

$\square$

Corollary 20. *If $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mathbf{p}^{\otimes\mathbb{Z}^d}, S) \cong (B^{\mathbb{Z}^d}, \mathcal{B}(B^{\mathbb{Z}^d}), \mathbf{q}^{\otimes\mathbb{Z}^d}, S)$ for some finite probability spaces $(A, \mathbf{p})$ and $(B, \mathbf{q})$, then $\mathrm{H}(\mathbf{p}) = \mathrm{H}(\mathbf{q})$.* $\square$

The reverse of this implication is Ornstein's Theorem, which will be discussed more later in the course.

3.2 Torus rotations

Recall that a $\mathbb{Z}^d$-action by **torus rotations** is a system of the form $(\mathbb{T}^D, \mathcal{B}(\mathbb{T}^D), m, R)$, where m is the Haar measure on $\mathbb{T}^D$ and R is given by

$$R^{(m_1,\dots,m_d)}(t) := t + m_1\alpha_1 + \cdots + m_d\alpha_d$$

for some fixed elements $\alpha_1, \dots, \alpha_d \in \mathbb{T}^D$. We have already seen that these systems are not weak mixing. Insofar as 'positive entropy' is an indicator of 'very random', the following comes as no surprise.

Proposition 21. *All ergodic torus rotations have KS entropy equal to zero.*

This extends to general, possibly non-ergodic isometric systems with a little extra work, but we omit the details.

Proof. Let $0 \in \mathbb{T}^D$ be the point $(0, 0, \ldots, 0)$, and consider $\mathbb{T}^d$ endowed with the metric $|\cdot|$ arising from the Euclidean metric on $\mathbb{R}^D$.

Step 1. The key fact for this proof is the following:

For any $\delta \in (0, 1/2)$, the observable

$$1_{B_\delta(0)} : \mathbb{T}^D \longrightarrow \{0, 1\}$$

is generating.

To see this, suppose that $s, t \in \mathbb{T}^d$ are distinct. Because R is assumed ergodic, the subgroup $\mathbb{Z}\alpha_1 + \cdots + \mathbb{Z}\alpha_d$ must be dense in $\mathbb{T}^D$. Since s and t are distinct, it follows that there is some $m \in \mathbb{Z}^d$ such that $r := m_1\alpha_1 + \cdots + m_d\alpha_d$ has the following property:

$$R^m s = s + r \in \operatorname{int} B_\delta(0) \quad \text{and} \quad R^m t = t + r \in \mathbb{T}^D \backslash B_\delta(0),$$

and hence

$$1_{B_\delta(0)}(R^m s) = 1 \neq 0 = 1_{B_\delta(0)}(R^m t).$$

Moreover, there is a fixed bound $M(\delta, |s-t|)$ (which may blow up as either $\delta \longrightarrow 0$ or $|s-t| \longrightarrow 0$) such that this m may be found in $[M]^d$. This implies that, for any $\delta > 0$ and $r > 0$ there is an $M \in \mathbb{N}$ such that the map

$$\mathbb{T}^D \longrightarrow \{0, 1\}^{[M]^d} : s \mapsto (1_{B_\delta(0)}(R^m s))_{m \in [M]^d}$$

distinguishes any two points that are distance at least r apart: that is, the level-sets of this map all have diameter less than r. Letting $M \longrightarrow \infty$, it follows that any open subset of $\mathbb{T}^d$ may be approximated in measure by unions of level-sets of these maps, and hence, overall, $1_{B_\delta(0)}$ is a generating obserable.

Step 2. The result of Step 1 shows that

$$\mathrm{h}(m, R) = \mathrm{h}(\Phi_{\delta *} m) \quad \forall \delta > 0.$$

On the other hand, the subadditivity of entropy gives

$$\begin{aligned} \mathrm{h}(\Phi_{\delta *} m) &\leqslant \mathrm{H}(m(B_\delta(0)), 1 - m(B_\delta(0))) \\ &= -m(B_\delta(0)) \log m(B_\delta(0)) - (1 - m(B_\delta(0))) \log(1 - m(B_\delta(0))). \end{aligned}$$

As $\delta \longrightarrow 0$ we have $m(B_\delta(0)) \longrightarrow 0$, and then this last bound also tends to $-0 \log 0 - 1 \log 1 = 0$, so $\mathrm{h}(m, R)$ must be zero. $\square$

4 Some properties of KS entropy

4.1 Continuity under $\overline{\mathrm{d}}$

We next give a continuity result for the entropy rate, and then for the KS entropy. The appropriate notion of approximation for shift-invariant measures is given by the $\overline{\mathrm{d}}$-metric.

Lemma 22. *Fix a finite alphabet A. Then for every $\varepsilon > 0$ there is a $\delta > 0$, depending on ε and $|A|$, such that*

$$\text{for } \mu, \nu \in \mathrm{Pr}^S A^{\mathbb{Z}^d}, \quad \overline{\mathrm{d}}(\mu,\nu) < \delta \quad \Longrightarrow \quad |\mathrm{h}(\mu) - \mathrm{h}(\nu)| < \varepsilon.$$

Proof. Step 1. First assume that μ and ν are ergodic, and let $\lambda \in \mathrm{J}(\mu,\nu)$ be a joining for which $\lambda\{a_0 \neq a_0'\} = \overline{\mathrm{d}}(\mu,\nu)$. Choose $\delta < \varepsilon^2$. Then for any $R \subseteq \mathbb{Z}^d$ one has $\lambda^R \in \mathrm{Cpl}(\mu^R, \nu^R)$ and

$$\int d_{\mathrm{H}}(a,a')\, \lambda^R(\mathrm{d}a, \mathrm{d}a') = |R|\lambda\{a_0 \neq a_0'\} < \delta|R|,$$

so $\mathrm{W}_1^{\mathrm{H}}(\mu^R, \nu^R) < \delta$. For any $\eta > 0$, applying inequality (1) and Corollary 12 as in the proof of Theorem 2 gives

$$\begin{aligned}
\mathrm{h}(\mu) - \eta &\leqslant \frac{1}{|R|}\log \mathrm{cov}_{\varepsilon}(\mu^R, \sqrt{\delta}|R|) + \mathrm{H}(\sqrt{\delta}, 1-\sqrt{\delta}) + \sqrt{\delta}\log|A| \\
&\leqslant \frac{1}{|R|}\log \mathrm{cov}_{\varepsilon-\sqrt{\delta}}(\nu^R, \sqrt{\delta}|R|/2) + \mathrm{H}(\sqrt{\delta}, 1-\sqrt{\delta}) + \sqrt{\delta}\log|A| \\
&\leqslant \frac{1}{|R|}\log \mathrm{cov}_{\varepsilon-\sqrt{\delta}}(\nu^R, 0) + \mathrm{H}(\sqrt{\delta}, 1-\sqrt{\delta}) + \sqrt{\delta}\log|A|,
\end{aligned}$$

provided all sides of R are long enough. As the minimum side-length of R increases to ∞, Proposition 15 turns this into

$$\mathrm{h}(\mu) \leqslant \mathrm{h}(\nu) + \mathrm{H}(\sqrt{\delta}, 1-\sqrt{\delta}) + \sqrt{\delta}\log|A|,$$

where the error term on the right tends to 0 as $\delta \longrightarrow 0$.

Step 2. Now consider general μ and ν. Let δ be chosen as in the ergodic case, and now assume that $\lambda \in \mathrm{J}(\mu,\nu)$ with $\lambda\{a_0 \neq a_0'\} \leqslant \delta^2$. Let $\lambda = \int_Y \lambda_y\, \theta(\mathrm{d}y)$ be an ergodic decomposition of λ, let $\pi_i : A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d} \longrightarrow A^{\mathbb{Z}^d}$ for $i = 1,2$ be the first and second coordinate projections, respectively, and let $\mu_y := \pi_{1*}\lambda_y$ and $\nu_y := \pi_{2*}\lambda_y$. Both μ_y and ν_y are ergodic, as images of the ergodic measure λ_y. On the one hand,

$$\delta^2 > \lambda\{a_0 \neq a_0'\} = \int_Y \lambda_y\{a_0 \neq a_0'\}\, \theta(\mathrm{d}y).$$

Therefore, letting $U := \{y \in Y : \lambda_y\{a_0 \neq a'_0\} < \delta\}$, Markov's Inequality gives $\theta(U) \geqslant 1 - \delta$.

On the other hand, the affinity of the entropy rate gives

$$|\mathrm{h}(\mu) - \mathrm{h}(\nu)| = \Big|\int_Y \mathrm{h}(\mu_y) - \mathrm{h}(\nu_y)\,\theta(\mathrm{d}y)\Big| \leqslant \int_Y |\mathrm{h}(\mu_y) - \mathrm{h}(\nu_y)|\,\theta(\mathrm{d}y)$$

$$\leqslant \int_U |\mathrm{h}(\mu_y) - \mathrm{h}(\nu_y)|\,\theta(\mathrm{d}y) + \int_{Y\setminus U} |\mathrm{h}(\mu_y) - \mathrm{h}(\nu_y)|\,\theta(\mathrm{d}y) < \varepsilon(1-\delta) + \delta\log|A|,$$

because all $\mathrm{h}(\mu_y)$ and $\mathrm{h}(\nu_y)$ must lie in $[0, \log|A|]$. Since we may choose δ smaller if necessary, this completes the proof. □

Definition 23. *Let $\varphi : X \longrightarrow A$ and $\psi : X \longrightarrow B$ be finite-valued observables on a probability space (X, μ), and let $\delta > 0$. Then φ is δ-**almost** ψ-**measurable** if there is a map $\xi : B \longrightarrow A$ such that*

$$\mu\{\varphi \neq \xi \circ \psi\} < \delta.$$

*It is ψ-**measurable** if this holds with $\delta = 0$.*

Corollary 24. *Let $\varepsilon > 0$ and A and B be finite alphabets. Then there is a $\delta > 0$, depending on ε and $|A|$ but not $|B|$, for which the following holds. Suppose that $\varphi : X \longrightarrow A$ and $\psi : X \longrightarrow B$ are finite-valued observables on a p.-p. system (X, Σ, μ, T) such that φ is δ-almost ψ-measurable. Let $\Phi : X \longrightarrow A^{\mathbb{Z}^d}$ and $\Psi : X \longrightarrow B^{\mathbb{Z}^d}$ be the corresponding equivariant maps. Then*

$$\mathrm{h}(\Psi_*\mu) \geqslant \mathrm{h}(\Phi_*\mu) - \varepsilon.$$

Proof. Let δ depend on ε and $|A|$ as in Lemma 22, let ξ be as in Definition 23, and let $\Xi : A^{\mathbb{Z}^d} \longrightarrow B^{\mathbb{Z}^d}$ be the factor map resulting from ξ. Then $\overline{d}(\Phi_*\mu, \Xi_*\Psi_*\nu) < \delta$, so Lemma 22 gives

$$\mathrm{h}(\Psi_*\mu) \geqslant \mathrm{h}(\Xi_*\Psi_*\mu) \geqslant \mathrm{h}(\Phi_*\mu) - \varepsilon.$$

□

Corollary 25. *Suppose that $\mathcal{P}_1 \leqslant \mathcal{P}_2 \leqslant \dots$ is an increasing sequence of finite measurable partitions of (X, Σ) which together generate Σ up to μ-negligible sets. For each i, let $\psi_i : X \longrightarrow [|\mathcal{P}_i|]$ be an observable whose level-sets are $\mathcal{P}_i$. Then*

$$\mathrm{h}(\mu, T) = \lim_{i\longrightarrow\infty} \mathrm{h}(\Phi_{i*}\mu),$$

where $\Phi_i : X \longrightarrow [|\mathcal{P}_i|]^{\mathbb{Z}^d}$ is the equivariant map resulting from φ_i, and this limit is non-decreasing. □

4.2 Entropy and joinings

Lemma 26. *If (X, Σ, μ, T) and (Y, Φ, ν, S) are p.-p.s.s and $\lambda \in \mathrm{J}(\mu, \nu)$, then*

$$\mathrm{h}(\lambda, T \times S) \leqslant \mathrm{h}(\mu, T) + \mathrm{h}(\nu, S).$$

This is an equality in case $\lambda = \mu \otimes \nu$.

Proof. Let $\mathcal{P}_1 \leqslant \mathcal{P}_2 \leqslant \cdots \leqslant \Sigma$ and $\mathcal{Q}_1 \leqslant \mathcal{Q}_2 \leqslant \cdots \leqslant \Phi$ be finite measurable partitions of X and Y generating Σ and Φ up to negligible sets, respectively. Let $\varphi_i : X \longrightarrow [|\mathcal{P}_i|]$ and $\psi_i : Y \longrightarrow [|\mathcal{Q}_i|]$ be observables generating $\mathcal{P}_i$ and $\mathcal{Q}_i$, and let Φ_i and Ψ_i be the resulting equivariant maps. Then $\mathcal{P}_i \otimes \mathcal{Q}_i$ is a sequence of measurable partitions of $X \times Y$ generating $\Sigma \otimes \Phi$ up to λ-negligible sets, and each $\mathcal{P}_i \otimes \mathcal{Q}_i$ is generated by the finite-valued observable $\Phi_i \times \Psi_i : X \times Y \longrightarrow [|\mathcal{P}_i|] \times [|\mathcal{Q}_i|]$. Therefore Corollary 25 gives

$$\begin{aligned} \mathrm{h}(\lambda, T \times S) &= \lim_{i \longrightarrow \infty} \mathrm{h}((\Phi_i \times \Psi_i)_*\lambda) = \lim_{i \longrightarrow \infty} \lim_{L \longrightarrow \infty} L^{-d}\mathrm{H}(((\Phi_i \times \Psi_i)_*\lambda)^{[L]^d}) \\ &\leqslant \lim_{i \longrightarrow \infty} \lim_{L \longrightarrow \infty} \big(L^{-d}\mathrm{H}((\Phi_{i*}\mu)^{[L]^d}) + L^{-d}\mathrm{H}((\Psi_{i*}\nu)^{[L]^d})\big) \\ &= \mathrm{h}(\mu, T) + \mathrm{h}(\nu, S). \end{aligned}$$

The inequality here results from the subadditivity of Shannon entropy. In case $\lambda = \mu \otimes \nu$, this step becomes the equality

$$\mathrm{H}(((\Phi_i \times \Psi_i)_*\lambda)^{[L]^d}) = \mathrm{H}((\Phi_{i*}\mu)^{[L]^d}) + \mathrm{H}((\Psi_{i*}\nu)^{[L]^d}),$$

leading to $\mathrm{h}(\mu \otimes \nu, T \times S) = \mathrm{h}(\mu, T) + \mathrm{h}(\nu, S)$. □

This is an obvious analog of Lemma 6 in Notes 9 for Shannon entropy. Note, however, the 'coercivity' part of that lemma can fail here: for instance, if (X, Σ, μ, T) is any non-trivial system with entropy zero, then any joining of two copies of this system has entropy $\leqslant 2\mathrm{h}(\mu, T) = 0$, but there are such examples which admit many different self-joinings (for instance, among the torus rotations).

4.3 Entropy and ergodic decomposition

We can also generalize the affinity of the entropy rate function to KS entropy. Now suppose that (X, Σ, μ, T) is a compact model of a p.-p. system, let $\Phi \leqslant \Sigma$ be its factor of invariant sets, and let $\pi : X \longrightarrow Y$ be a factor map to another compact metric space which generates Φ up to μ-negligible sets. In this case, because the underlying spaces are compact, one may disintegrate μ over the factor map π to obtain an ergodic decomposition

$$\mu = \int_Y \mu_y \, \nu(\mathrm{d}y), \quad \text{where } \nu := \pi_*\mu.$$

Proposition 27 (Integral formula for KS entropy)**.** *In the setting above, one has*

$$\mathrm{h}(\mu, T) = \int_Y \mathrm{h}(\mu_y, T)\,\nu(\mathrm{d}y).$$

Proof. This requires a preliminary consideration of Σ. Since Σ is the Borel σ-algebra of X, it is countably generated, for example by the collection of all rational-radius balls centred at the points of a countable dense subset of X. We may therefore choose an increasingly fine sequence $\mathcal{P}_1 \leqslant \mathcal{P}_2 \leqslant \cdots$ of finite Borel partitions of X which together generate Σ. For each i, we may now also choose a finite-valued observable $\varphi_i : X \longrightarrow \{1, 2, \dots, |\mathcal{P}_i|\}$ whose level sets are precisely the cells of $\mathcal{P}_i$. Because the $\mathcal{P}_i$ are increasing and together generate Σ, for any Borel probability measure on X, any Borel set may be approximated in measure by $\mathcal{P}_i$-measurable sets as $i \longrightarrow \infty$. Given a finite-value observable ψ on X and $\delta > 0$, we may apply this fact to the level-sets of ψ to conclude that ψ is is δ-almost φ_i-measurable for all sufficiently large i. Also, since $\mathcal{P}_i \leqslant \mathcal{P}_{i+1}$, the observable φ_i is measurable with respect to φ_{i+1} for each i.

Therefore, Corollary 25 may be applied to both μ and all of the μ_ys, giving

$$\mathrm{h}(\mu, T) = \lim_{i\longrightarrow\infty} \mathrm{h}(\Phi_{i*}\mu) = \lim_{i\longrightarrow\infty} \mathrm{h}\Big(\int_Y \Phi_{i*}\mu_y\,\nu(\mathrm{d}y)\Big) = \lim_{i\longrightarrow\infty} \int_Y \mathrm{h}(\Phi_{i*}\mu_y)\,\nu(\mathrm{d}y),$$

by the affinity of the entropy rate for a fixed space of shift-invariant measures. However, we also have that the map $y \mapsto \mathrm{h}(\Phi_{i*}\mu_y)$ is non-decreasing in i, because each $\Phi_{i*}\mu_y$ is an image of $\Phi_{(i+1)*}\mu_y$. Therefore the Monotone Convergence Theorem and another appeal to Corollary 25 give

$$\lim_{i\longrightarrow\infty} \int_Y \mathrm{h}(\Phi_{i*}\mu_y)\,\nu(\mathrm{d}y) = \int_Y \lim_{i\longrightarrow\infty} \mathrm{h}(\Phi_{i*}\mu_y)\,\nu(\mathrm{d}y) = \int_Y \mathrm{h}(\mu_y, T)\,\nu(\mathrm{d}y).$$

□

References

[Kal02] Olav Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: tim@cims.nyu.edu
URL: cims.nyu.edu/˜tim

Ergodic Theory

Notes 11: Entropy, couplings and joinings

Having defined the abstract Kolmogorov-Sinai entropy and established its basic properties, we will now begin to study some more delicate ergodic-theoretic features of shift systems. Our main goal is the following theorem of Sinai.

Theorem 1 (Sinai's Theorem)**.** *If (X, Σ, μ, T) is a p.-p. $\mathbb{Z}^d$-system of entropy h, and $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mathbf{p}^{\otimes \mathbb{Z}^d}, S)$ is a Bernoulli shift with $\mathrm{H}(\mathbf{p}) \leqslant h$, then there is a factor map*

$$(X, \Sigma, \mu, T) \longrightarrow (A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mathbf{p}^{\otimes \mathbb{Z}^d}, S).$$

The constraint $\mathrm{H}(\mathbf{p}) \leqslant h$ is clearly necessary, because $\mathrm{H}(\mathbf{p})$ equals the KS entropy of the Bernoulli shift, and KS entropy is monotone under factor maps. Choosing some $\mathbf{p}$ for which $\mathrm{H}(\mathbf{p}) = h$, this result asserts, in a sense, that all of the KS entropy in an abstract system can be accounted by a Bernoulli-system factor of it (although not at all uniquely).

We will assume $h < \infty$ in our treatment of Sinai's Theorem, but this assumption can be removed with some extra technicalities.

Sinai's Theorem immediately implies the following.

Corollary 2 (Weak Ornstein Theorem)**.** *If $(A, \mathbf{p})$ and $(B, \mathbf{q})$ are finite probability spaces with $\mathrm{H}(\mathbf{p}) = \mathrm{H}(\mathbf{q})$, then each of*

$$(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mathbf{p}^{\otimes \mathbb{Z}^d}, S)$$

and

$$(B^{\mathbb{Z}^d}, \mathcal{B}(B^{\mathbb{Z}^d}), \mathbf{q}^{\otimes \mathbb{Z}^d}, S)$$

admits a factor map onto the other. □

The full Ornstein Theorem gives that these two Bernoulli shifts are actually isomorphic. Its proof is a significant enhancement of the proof of Sinai's Theorem, and will not be covered in this course.

1 Entropy-transport inequalities

One of the phenomena lying behind Sinai's Theorem is the ability to construct W_1-small couplings between a product measure and another measure which resembles it in some weaker sense. The following is an example of an 'entropy-transport' inequality (where 'transport' refers to the Wasserstein metric).

Theorem 3. *Given a finite probability space $(A, \mathbf{p})$ and $\varepsilon > 0$, there is $\delta > 0$ for which the following holds: If $N \geqslant 1$, and $\mu \in \Pr A^N$ satisfies*

- *(entropy approximation) $\mathrm{H}(\mu) \geqslant N(\mathrm{H}(\mathbf{p}) - \delta)$, and*
- *(marginal approximation)*

$$\sum_{n=1}^{N} \|\pi_{n*}\mu - \mathbf{p}\|_{\mathrm{TV}} < \delta N,$$

where $\pi_n : A^N \longrightarrow A$ is the N^{th} coordinate projection, then

$$\mathrm{W}_1^{\mathrm{H}}(\mu, \mathbf{p}^{\otimes N}) < \varepsilon N.$$

Thus, if an unknown probability measure μ on A^N has roughly the same entropy and, on average, roughly the same marginals as $\mathbf{p}^{\otimes N}$, this μ must be close to $\mathbf{p}^{\otimes N}$ in the Wasserstein metric.

There are many ways to prove Theorem 3. We will prove it via the following more precise result, which in turn uses an adaptation of an argument of Marton.

Theorem 4. *Given a finite alphabet A, there is a $C < \infty$, depending only on $|A|$, for which the following holds: If $N \geqslant 1$, and $\mu \in \Pr A^N$ has one-dimensional marginals $\theta_1, \ldots, \theta_N$, then*

$$\mathrm{W}_1\Big(\mu, \bigotimes_{n \leqslant N} \theta_n\Big) \leqslant C\sqrt{N\Big(\sum_{n \leqslant N} \mathrm{H}(\theta_n) - \mathrm{H}(\mu)\Big)}.$$

Observe that, by the strong subadditivity of H (Lemma 6 of Notes 9) and induction on N, one always has

$$\mathrm{H}(\mu) \leqslant \sum_{n \leqslant N} \mathrm{H}(\theta_n),$$

with equality if and only if $\mu = \bigotimes_{n \leqslant N} \theta_n$. Theorem 4 makes this quantitative in terms of the Wasserstein metric.

We begin with a sequence of elementary lemmas; then show how these combine to give Theorem 4; and then use Theorem 4 to deduce the slightly more flexible Theorem 3.

Lemma 5. *For any compact metric space* (X, d) *and* $\mu, \nu \in \operatorname{Pr} X$*, one has*

$$\mathrm{W}_1(\mu, \nu) \leqslant \|\mu - \nu\|_{\mathrm{TV}} \cdot \operatorname{diam}(X, d).$$

Proof. Let $\alpha := \|\mu - \nu\|_{\mathrm{TV}}$, and recall that $\alpha = 1 - (\mu \wedge \nu)(X)$, where $\mu \wedge \nu$ is the unique largest Borel measure on X that is less than or equal to both μ and ν. If $\alpha = 0$, then $\mu = \nu$ so we may use the diagonal coupling. If $\alpha > 0$, define $\lambda \in \operatorname{Pr} X^2$ by

$$\lambda = \int_X \delta_{(x,x)}\, (\mu \wedge \nu)(\mathrm{d}x) + \frac{1}{\alpha}(\mu - \mu \wedge \nu) \otimes (\nu - \mu \wedge \nu).$$

Then

$$\begin{aligned}\lambda(A \times X) = (\mu \wedge \nu)(\{x :\ (x, x) \in A \times X\}) + \frac{1}{\alpha}(\mu - \mu \wedge \nu)(A) \cdot (\nu - \mu \wedge \nu)(X) \\ = (\mu \wedge \nu)(A) + (\mu - \mu \wedge \nu)(A) = \mu(A),\end{aligned}$$

using that $(\nu - \mu \wedge \nu)(X) = 1 - (\mu \wedge \nu)(X) = \alpha$. Similarly, $\lambda(X \times B) = \nu(B)$, and finally, one has

$$\begin{aligned}\int_{X^2} d\, \mathrm{d}\lambda = \int_X d(x, x)\, (\mu \wedge \nu)(\mathrm{d}x) + \frac{1}{\alpha} \int_{X^2} d(x, y)\, (\mu - \mu \wedge \nu)(\mathrm{d}x) \cdot (\nu - \mu \wedge \nu)(\mathrm{d}y) \\ \leqslant \frac{1}{\alpha} \cdot \operatorname{diam}(X, d) \cdot (\mu - \mu \wedge \nu)(X) \cdot (\nu - \mu \wedge \nu)(X) \leqslant \alpha \cdot \operatorname{diam}(X, d).\end{aligned}$$

□

Lemma 6. *Let* (X, d_X) *and* (Y, d_Y) *be compact metric spaces, and endow* $X \times Y$ *with the metric*

$$d((x, y), (x', y')) := d_X(x, x') + d_Y(y, y').$$

Let $\theta \in \operatorname{Pr} X$*, suppose that* $\mu_1, \mu_2 \in \operatorname{Pr}(X \times Y)$ *both have first marginal equal to* θ*, and let*

$$\mu_i = \int_X \delta_x \otimes \nu_{i,x}\, \theta(\mathrm{d}x)$$

for $i = 1, 2$ *be the disintegrations over the first coordinate. Then*

$$\mathrm{W}_1^d(\mu_1, \mu_2) \leqslant \int_X \mathrm{W}_1^{d_Y}(\nu_{1,x}, \nu_{2,x})\, \theta(\mathrm{d}x).$$

Proof. Let $\lambda_x \in \Pr Y^2$ be a $(\nu_{1,x}, \nu_{2,x})$-coupling realizing $\mathrm{W}_1^{d_Y}(\nu_{1,x}, \nu_{2,x})$ for each x. (Technically, one should make sure that λ_x can be chosen measurably in x, but we will apply this lemma only in the case of finite sets X and Y, so ignore this point.) Now define $\lambda \in \Pr((X \times Y)^2)$ by

$$\lambda(\mathrm{d}x_1, \mathrm{d}y_1, \mathrm{d}x_2, \mathrm{d}y_2) := \int_X \delta_{(x,x)}(\mathrm{d}x_1, \mathrm{d}x_2) \otimes \lambda_x(\mathrm{d}y_1, \mathrm{d}y_2)\, \theta(\mathrm{d}x).$$

The marginal of this on the first copy of $X \times Y$ is

$$\int_X \delta_x \otimes \nu_{1,x}\, \theta(\mathrm{d}x) = \mu_1,$$

and similarly it is μ_2 on the second copy of $X \times Y$. Now an easy computation gives

$$\int_{(X\times Y)^2} d\,\mathrm{d}\lambda = \int_X \Big(d_X(x,x) + \int_{Y^2} d_Y(y_1,y_2)\, \lambda_x(\mathrm{d}y_1, \mathrm{d}y_2)\Big)\, \theta(\mathrm{d}x)$$

$$= \int_X \mathrm{W}_1^{d_Y}(\nu_{1,x}, \nu_{2,x})\, \theta(\mathrm{d}x),$$

so this is an upper bound for $\mathrm{W}_1^d(\mu_1, \mu_2)$. □

Lemma 7. *Let (X, d_X) and (Y, d_Y) be as above, and now let $\theta_i \in \Pr X$ and $\nu_i \in \Pr Y$ for $i = 1, 2$. Then*

$$\mathrm{W}_1^d(\theta_1 \otimes \nu_1, \theta_2 \otimes \nu_2) \leqslant \mathrm{W}_1^{d_X}(\theta_1, \theta_2) + \mathrm{W}_1^{d_Y}(\nu_1, \nu_2).$$

Proof. Let $\lambda_X \in \mathrm{Cpl}(\theta_1, \theta_2)$ and $\lambda_Y \in \mathrm{Cpl}(\nu_1, \nu_2)$ be couplings achieving the Wasserstein distances, and let λ be the product measure

$$\lambda(\mathrm{d}x_1, \mathrm{d}y_1, \mathrm{d}x_2, \mathrm{d}y_2) := \lambda_X(\mathrm{d}x_1, \mathrm{d}x_2) \otimes \lambda_Y(\mathrm{d}y_1, \mathrm{d}y_2).$$

A simple check shows that $\lambda \in \mathrm{Cpl}(\theta_1 \otimes \nu_1, \theta_2 \otimes \nu_2)$ and

$$\int_{(X\times Y)^2} d\,\mathrm{d}\lambda = \int_{X^2} d_X\, \mathrm{d}\lambda_X + \int_{Y^2} d_Y\, \mathrm{d}\lambda_Y.$$

□

Corollary 8. *Let (X, d_X) and (Y, d_Y) be as above, let $\mu \in \Pr(X \times Y)$ with disintegration*

$$\mu = \int_X \delta_x \otimes \nu_x\, \theta(\mathrm{d}x)$$

over the first coordinate, and let $\theta' \in \Pr X$ and $\nu' \in \Pr Y$. Then

$$\mathrm{W}_1^d(\mu, \theta' \otimes \nu') \leqslant \int_X \mathrm{W}_1^{d_Y}(\nu_x, \nu')\, \theta(\mathrm{d}x) + \mathrm{W}_1^{d_X}(\theta, \theta').$$

Proof. The triangle inequality for W_1 gives

$$\mathrm{W}_1^d(\mu, \theta' \otimes \nu') \leqslant \mathrm{W}_1^d(\mu, \theta \otimes \nu') + \mathrm{W}_1^d(\theta \otimes \nu', \theta' \otimes \nu'),$$

and now the two terms on the left may be bounded using Lemmas 6 and 7 respectively. □

Lemma 9. *Suppose that (X, μ) is a probability space, that A is a finite set, that $x \mapsto \nu_x$ is a measurable map $X \longrightarrow \Pr A$, and that $\nu := \int_X \nu_x \, \mu(\mathrm{d}x)$. Then there is some $C > 0$, depending only on $|A|$, such that*

$$\int_X \|\nu - \nu_x\|_{\mathrm{TV}} \, \mu(\mathrm{d}x) \leqslant C \sqrt{\mathrm{H}(\nu) - \int_X \mathrm{H}(\nu_x) \, \mu(\mathrm{d}x)}.$$

Proof. The function H is smooth and strictly concave on $\Pr A$, with Hessian matrix that is negative-definite everywhere. Since this Hessian is also continuous, and $\Pr A$ is compact, it follows that there is some $c > 0$ such that

$$\mathrm{H}(\nu') \leqslant \mathrm{H}(\nu) + \nabla \mathrm{H}(\nu) \bullet (\nu' - \nu) - c\|\nu' - \nu\|_{\mathrm{TV}}^2 \quad \forall \nu' \in \Pr A.$$

Letting $\nu' := \nu_x$ and integrating over $\mu(\mathrm{d}x)$, the linear term vanishes because $\nu = \int \nu_x \, \mu(\mathrm{d}x)$, leaving

$$\begin{aligned} \int_X \mathrm{H}(\nu_x) \, \mu(\mathrm{d}x) &\leqslant \mathrm{H}(\nu) - c \int_X \|\nu_x - \nu'\|_{\mathrm{TV}}^2 \, \mu(\mathrm{d}x) \\ \Longrightarrow \quad \int_X \|\nu_x - \nu'\|_{\mathrm{TV}} \, \mu(\mathrm{d}x) &\leqslant \sqrt{\int_X \|\nu_x - \nu'\|_{\mathrm{TV}}^2 \, \mu(\mathrm{d}x)} \\ &\leqslant \sqrt{c^{-1}} \sqrt{\mathrm{H}(\nu) - \int_X \mathrm{H}(\nu_x) \, \mu(\mathrm{d}x)}. \end{aligned}$$

□

Proof of Theorem 4. The case $N = 1$ is trivial, since then $\mu = \theta_1$. So suppose $N \geqslant 2$, let C be as in Lemma 9, and suppose that the result is known with this choice of C for product spaces of length at most $N - 1$. Let $\mu \in \Pr A^N$ have one-dimensional marginals $\theta_1, \dots, \theta_N$. Let $\pi^- : A^N \longrightarrow A^{N-1}$ be the projection onto the first $N - 1$ coordinates, let $\mu^- := \pi^-_* \mu$, and let

$$\mu = \int_{A^{N-1}} \delta_a \otimes \nu_a \, \mu^-(\mathrm{d}a)$$

be the disintegration over π^-. Writing A^N as the product $A^{N-1} \times A$, Corollary 8 gives

$$\mathrm{W}_1^{\mathrm{H}}\Big(\mu, \bigotimes_{n\leqslant N} \theta_n\Big) \leqslant \int_{A^{N-1}} \mathrm{W}_1(\nu_a, \theta_N)\, \mu^-(\mathrm{d}a) + \mathrm{W}_1^{\mathrm{H}}\Big(\mu^-, \bigotimes_{n\leqslant N-1} \theta_n\Big).$$

For the first term on the right, Lemmas 9 and 5 give

$$\begin{aligned}\int_{A^{N-1}} \mathrm{W}_1(\nu_a, \theta_N)\, \mu^-(\mathrm{d}a) &\leqslant \int_{A^{N-1}} \|\nu_a - \theta_N\|_{\mathrm{TV}}\, \mu^-(\mathrm{d}a) \\ &\leqslant C\sqrt{\mathrm{H}(\theta_N) - \int_{A^{N-1}} \mathrm{H}(\nu_a)\, \mu^-(\mathrm{d}a)}.\end{aligned}$$

For the second, the inductive hypothesis gives

$$\mathrm{W}_1^{\mathrm{H}}\Big(\mu^-, \bigotimes_{n\leqslant N-1} \theta_n\Big) \leqslant C\sqrt{(N-1)\Big(\sum_{n\leqslant N-1} \mathrm{H}(\theta_n) - \mathrm{H}(\mu^-)\Big)}.$$

Adding these estimates, and using the elementary inequality

$$\sqrt{(N-1)a} + \sqrt{b} \leqslant \sqrt{N(a+b)} \quad \forall a, b \in [0, \infty),$$

we obtain the upper bound

$$C\sqrt{N}\sqrt{\sum_{n\leqslant N} \mathrm{H}(\theta_n) - \mathrm{H}(\mu^-) - \mathrm{H}(\mu \,|\, \pi^-)},$$

Recalling that $\mathrm{H}(\mu) = \mathrm{H}(\mu^-) + \mathrm{H}(\mu \,|\, \pi^-)$, this is equal to

$$C\sqrt{N}\sqrt{\sum_{n\leqslant N} \mathrm{H}(\theta_n) - \mathrm{H}(\mu)}.$$

□

Proof of Theorem 3. By a simple induction on N, Lemmas 5 and 7 give

$$\mathrm{W}_1^{\mathrm{H}}\Big(\bigotimes_{n\leqslant N} \theta_n, \mathbf{p}^{\otimes N}\Big) \leqslant \sum_{n=1}^{N} \|\theta_n - \mathbf{p}\|_{\mathrm{TV}}.$$

Combined with Theorem 4 and the triangle inequality for W_1, this gives

$$\mathrm{W}_1^{\mathrm{H}}(\mu, \mathbf{p}^{\otimes N}) \leqslant C\sqrt{N}\sqrt{\sum_{n\leqslant N} \mathrm{H}(\theta_n) - \mathrm{H}(\mu)} + \sum_{n=1}^{N} \|\theta_n - \mathbf{p}\|_{\mathrm{TV}}.$$

Now suppose $\varepsilon > 0$. Since H is continuous on the compact set $\mathrm{Pr}\, A$, there is some $\delta > 0$ such that

$$\|\theta - \theta'\|_{\mathrm{TV}} < \delta \quad \Longrightarrow \quad |\mathrm{H}(\theta) - \mathrm{H}(\theta')| < \varepsilon.$$

On the other hand, H is bounded by $\log |A|$, so, making δ even smaller if necessary, this can be turned into

$$\sum_{n \leqslant N} \|\theta_n - \mathbf{p}\|_{\mathrm{TV}} < \delta N$$
$$\Longrightarrow \quad \Big| \sum_{n \leqslant N} \mathrm{H}(\theta_n) - N\mathrm{H}(\mathbf{p}) \Big| \leqslant \sum_{n \leqslant N} |\mathrm{H}(\theta_n) - \mathrm{H}(\mathbf{p})| < \varepsilon N \quad (1)$$

(since the left-hand side here implies that $\|\theta_n - \mathbf{p}\|_{\mathrm{TV}}$ is very small for most values of n, and the remaining terms are bounded by $\log |A|$, hence cannot contribute too much to the sum).

Therefore, if $\mathrm{H}(\mu) \geqslant N(\mathrm{H}(\mathbf{p}) - \delta)$ and also $\sum_{n \leqslant N} \|\theta_n - \mathbf{p}\|_{\mathrm{TV}} < \delta N$, then inequality (1) gives

$$\mathrm{W}_1^{\mathrm{H}}(\mu, \mathbf{p}^{\otimes N}) < C\sqrt{N}\sqrt{\varepsilon N + \delta N} + \delta N = (C(\varepsilon + \delta) + \delta)N.$$

Since ε and δ may be chosen arbitrarily small, this completes the proof. □

2 Estimates on the $\bar{\mathrm{d}}$-metric

Now consider the shift action on some $A^{\mathbb{Z}^d}$. The following is a dynamical-systems analog of Theorem 3.

Theorem 10. *For every finite probability space $(A, \mathbf{p})$ and every $\varepsilon > 0$, there is a $\delta > 0$ for which the following holds: If $\mu \in \mathrm{Pr}^S\, A^{\mathbb{Z}^d}$ satisfies*

- $\|\mu^{\{0\}} - \mathbf{p}\|_{\mathrm{TV}} < \delta$, *and*
- $\mathrm{h}(\mu) > \mathrm{H}(\mathbf{p}) - \delta$,

then

$$\bar{\mathrm{d}}(\mu, \mathbf{p}^{\otimes \mathbb{Z}^d}) \leqslant \varepsilon.$$

Proof. Given $|A|$ and ε, let δ be chosen according to Theorem 3.

For each $L \in \mathbb{N}$, the measure $\mu^{[L]^d} \in \mathrm{Pr}\, A^{[L]^d}$ has all one-dimensional marginals equal to $\mu^{\{0\}}$, hence close to $\mathbf{p}$. Also, it satisfies

$$\mathrm{H}(\mu^{[L]^d}) \geqslant L^d \mathrm{h}(\mu) > L^d(\mathrm{H}(\mathbf{p}) - \delta),$$

by the definition of h as an infimum over rectangles. Therefore, Theorem 3 gives some $\lambda_L^0 \in \mathrm{Cpl}(\mu^{[L]^d}, \mathbf{p}^{\otimes [L]^d})$ such that

$$\int_{A^{[L]^d} \times A^{[L]^d}} d_{\mathrm{H}}\, \mathrm{d}\lambda_L^0 < \varepsilon L^d.$$

Using this λ_L^0, we define can $\lambda_L^1 \in \Pr(A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d})$ as a product measure over side-L boxes. First, we identify $\mathbb{Z}^d$ with $[L]^d \times (L \cdot \mathbb{Z})^d$ by writing it as a union of all the boxes obtained from $[L]^d$ by translating by elements of $(L \cdot \mathbb{Z})^d$. This then leads to the identification

$$A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d} = A^{[L]^d \times (L\cdot\mathbb{Z})^d} \times A^{[L]^d \times (L\cdot\mathbb{Z})^d} = \big(A^{[L]^d} \times A^{[L]^d}\big)^{(L\cdot\mathbb{Z})^d},$$

and in terms of this right-hand product space we let

$$\lambda_L^1 := (\lambda_L^0)^{\otimes (L\cdot\mathbb{Z})^d}.$$

This measure on $A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d}$ need not be invariant under the shift action on coordinate, but it is invariant under shifting by any element of $L \cdot \mathbb{Z}^d$, so we may produce a fully shift-invariant measure by setting

$$\lambda_L := L^{-d} \sum_{v \in [L]^d} S_*^v \lambda_L^1.$$

We now establish two estimates concerning these measures.

- Let $K \in \mathbb{N}$, and let $\pi_i^{[-K,K]^d} : A^{\mathbb{Z}^d} \times A^{\mathbb{Z}^d} \longrightarrow A^{[-K,K]^d}$ for $i = 1$ (resp. $i = 2$) be the coordinate projection via the first (resp. second) copy of $A^{\mathbb{Z}^d}$. Then

$$\pi_{i*}^{[-K,K]^d} \lambda_L = L^{-d}\Big(\sum_{v \in [L]^d \text{ s.t. } v + [L]^d \supseteq [-K,K]^d} \pi_{i*}^{[-K,K]^d}(S_*^v \lambda_L^1) + \sum_{v \in [L]^d \text{ s.t. } v + [L]^d \not\supseteq [-K,K]^d} \pi_{i*}^{[-K,K]^d}(S_*^v \lambda_L^1) \Big).$$

  On the one hand, if $[-K,K]^d \subseteq v + [L]^d$, then the marginal of $S_*^v \lambda_L^1$ on the first set of coordinates in $[-K,K]^d$ is just the marginal of λ_L^0 on some copy of $[-K,K]^d$ through the first set of coordinates, and this is precisely $\mu^{[-K,K]^d}$, because the first marginal of λ_L^0 is $\mu^{[L]^d}$. On the other hand, the second term inside the right-hand side above consists of at most $4dKL^{d-1}$

terms (coming from the 'boundary' of $v + [L]^d$). Putting these facts together gives that

$$\pi_{1*}^{[-K,K]^d}\lambda_L = \mu^{[L]^d} + \mathrm{O}\Big(\frac{dK}{L}\Big),$$

and similarly

$$\pi_{2*}^{[-K,K]^d}\lambda_L = \mathbf{p}^{\otimes[L]^d} + \mathrm{O}\Big(\frac{dK}{L}\Big).$$

Therefore

$$\pi_{1*}\lambda_L \longrightarrow \mu \quad \text{and} \quad \pi_{2*}\lambda_L \longrightarrow \mathbf{p}^{\otimes\mathbb{Z}^d}$$

in the vague topology as $L \longrightarrow \infty$.

- Secondly, observe that

$$\lambda_L\{a_0 \neq a_0'\} = L^{-d}\sum_{v\in[L]^d}\lambda_L^0\{(a_m, a_m')_{m\in[L]^d} : a_v \neq a_v'\}$$

$$= L^{-d}\int_{A^{[L]^d}\times A^{[L]^d}} d_{\mathrm{H}}\,\mathrm{d}\lambda_L < \varepsilon.$$

Therefore, letting $\lambda \in \mathrm{Pr}^S(A^{\mathbb{Z}^d}\times A^{\mathbb{Z}^d})$ be any subsequential limit of $(\lambda_L)_{L\geqslant 1}$ in the vague topology, the first of the above facts implies that $\lambda \in \mathrm{J}(\mu, \mathbf{p}^{\otimes\mathbb{Z}^d})$, and the second implies that

$$\lambda\{a_0 \neq a_0'\} \leqslant \limsup_{L\longrightarrow\infty}\lambda_L\{a_0 \neq a_0'\} \leqslant \varepsilon,$$

as required. □

The property given by the preceding theorem, slightly generalized, is enough to warrant its own name.

Definition 11. *A shift-system* $(A^{\mathbb{Z}^d}, \mathcal{B}(A^{\mathbb{Z}^d}), \mu, S)$ *is **finitely determined** if for every* $\varepsilon > 0$ *there are* $K \in \mathbb{N}$ *and* $\delta > 0$ *for which the following holds: If* $\mu \in \mathrm{Pr}^S A^{\mathbb{Z}^d}$ *satisfies*

$$\|\nu^{[-K;K]^d} - \mu^{[-K;K]^d}\|_{\mathrm{TV}} < \delta$$

and

$$\mathrm{h}(\mu) > \mathrm{h}(\nu) - \delta,$$

then also

$$\overline{\mathrm{d}}(\mu,\nu) < \varepsilon.$$

In terms of this property, we can now formulate the full strength of Ornstein's Theorem.

Theorem 12 (Ornstein's Theorem)**.** *For any $h \in [0, \infty)$, let $\mathbf{B}$ be a choice of Bernoulli shift of entropy rate h. Then for any other shift-system $\mathbf{X}$ with entropy rate h, the following are equivalent:*

i) $\mathbf{X}$ *is finitely determined.*

ii) $\mathbf{X} \cong \mathbf{B}$.

Theorem 10 is the first step towards proving that (ii) $\Longrightarrow$ (i). We will not complete the proof of either direction in Ornstein's Theorem, but will use the property of being finitely determined in the proof of Sinai's Theorem.

TIM AUSTIN
COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER ST, NEW YORK NY 10012, USA

Email: tim@cims.nyu.edu
URL: cims.nyu.edu/˜tim